GIAC Security Essentials Certification (GSEC) Practical Assignment
Version 1.4b – Option 2

Growing Pains

Andrew Edwards

Contents

Introduction

This document describes part of my work as a security consultant on a recent project; working for a client going through an explosive period of growth in personnel and scope. The client originally comprised of a small group of people developing an idea. The idea was eventually backed financially, warranting a move into a brand new building with a tenfold increase in staff. During the 18 month period of the project there was a 400% growth in staff on top of that. The scope of the groups work in budgetary terms increased 1000%.

I haven't worked on any reasonably sized project that hasn't changed in some way as it went along, but this project changed more than most. I am however limiting the scope of this document to the security infrastructure of the client and the main types of traffic passing through them as well as some elements of the underlying network.

It would not be helpful to you, as the reader of this document (and extremely difficult for me as the author), if I attempted to relay the events in chronological order; let alone include all of the changes in direction that went on, so I have split the document into three sections:-

1. The Original Design.
2. The Improved Design.
3. Conclusions.

My role changed throughout the course of the project; initially I was just asked to implement the firewalls and proxy servers. When I pointed out some of the technical and security issues with the original design it was clear that although security elements had been included, a deeper understanding of security hadn't been applied. I ended up as part of the design team responsible for security with a team of five people at one stage. I have identified risks where applicable with regards to Confidentiality, Availability and Integrity.

The original design was never implemented in its entirety, as things started moving towards the improved design even as it was being implemented

The changes to the security infrastructure due to technical necessity, new requirements and risk mitigation are covered in the second section.

The final section shows that although security has improved technically, the company itself was growing so fast the processes could not be put in place to support it even if the company had wanted to.

The Original Design

The following diagram shows a simplified version of the network that was due to be implemented at RCN (Random Company Name) Ltd.

NB: Server names and network addresses have been modified from the original

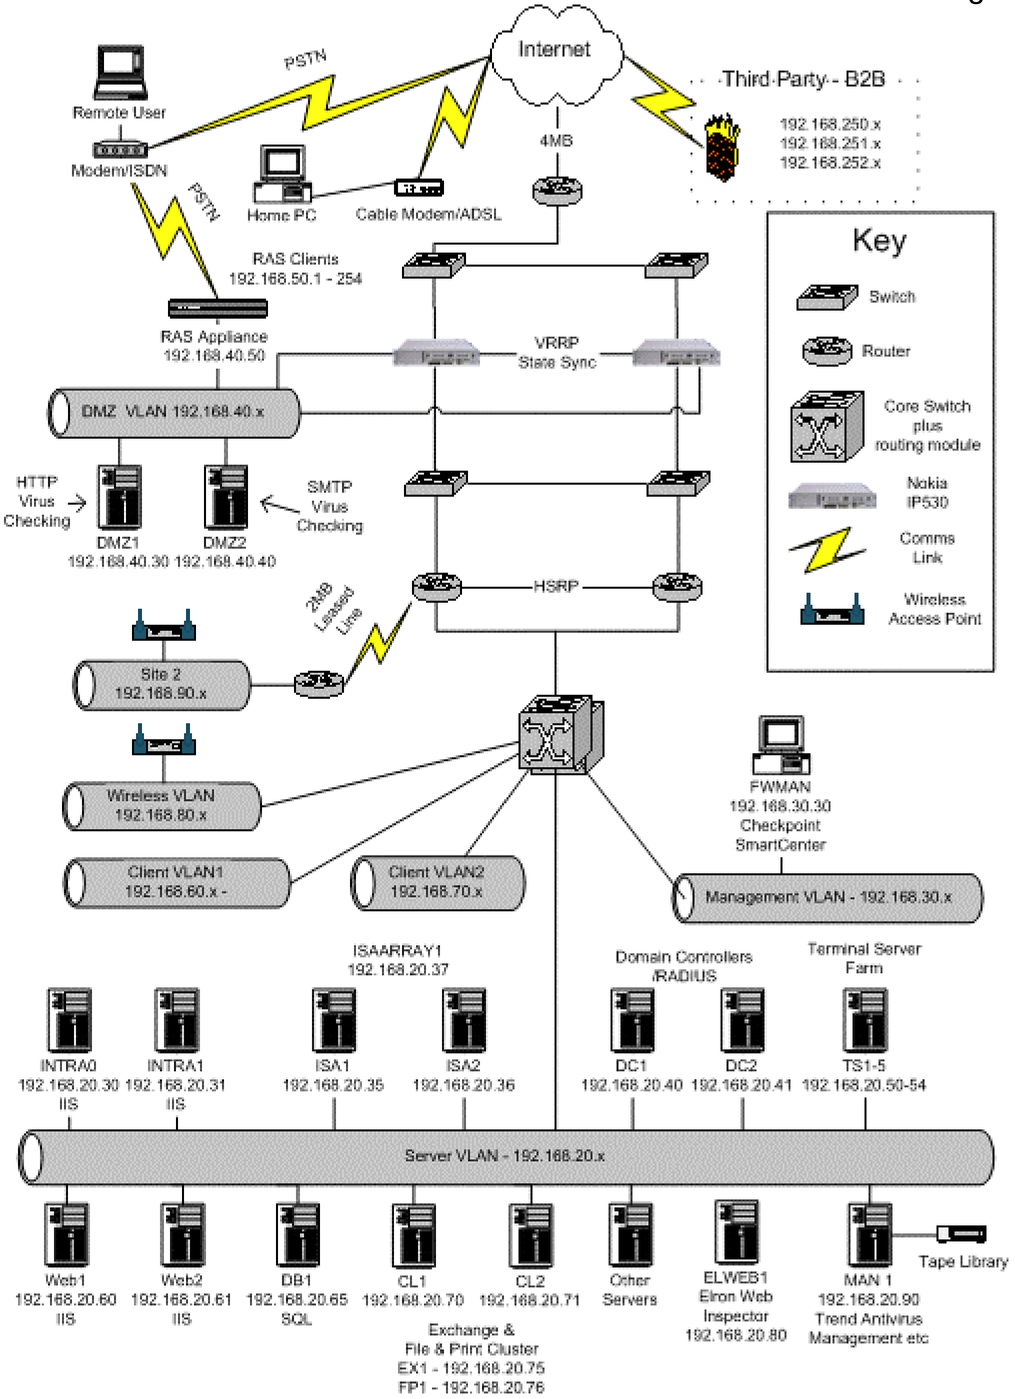

Figure 1 - Original Logical Network Design

Network

The Internet router is owned and maintained by BT and only supports one connection to RCN Ltd's network due to the level of service purchased from BT. This router connects to one 3500 series Cisco switch; another 3500 series switch is connected to the first switch to make a resilient pair. Each switch has one connection to one of the two Nokia IP530 appliances. The Nokias connect to each other for state table synchronisation. Below the Nokias is another pair of interconnected 3500 series switches. Each Nokia connects to one of that pair. The Nokias also have a connection to the DMZ (Demilitarised Zone) VLAN (Virtual Local Area Network). This VLAN exists as a number of ports on the core switches.

The Internal 3500's have one connection each, to one of a pair of interconnected routers that use HSRP (Hot Standby Routing Protocol) for resilience. These routers have a connection to one of the 6500 series core switches. The rest of the internal network is made up of a number of VLANs that have ports spread across both core switches. The VLANs create logical partitions in the network for Servers, clients, wireless access etc.

There is a link between a router at the main site and one at the remote site over a 2Mb leased line. The remote site has a mixture of wired and wireless users.

Assorted network facts:-

- The wireless sections of the network use the 802.11b standard.
- WEP (Wired Equivalent Privacy) is used as the encryption protocol with 128 bit keys.
- SSID's set to default and are broadcasted.
- The internal network uses a private address range.
- All network devices have unused ports disabled.
- Rooms where network equipment is installed require key card access.
- All VLAN's on the core switches can access each other.
- All cabling and patch panels are CAT6 rated.

Risks

Confidentiality

Traffic on the internal network is generally unencrypted.

The weak points of 802.11b are well known and well documented; there are many articles available when searching the Internet. The fact that WEP hadn't even been enabled at the second site means that anyone with a wireless card in range of an access point could access devices at both sites. "…the majority of access points are being deployed without WEP even being enabled. That's the equivalent of leaving your front door virtually unlocked and standing wide open[1]." (Ellison, paragraph 6)

The DMZ is a VLAN on the core switches that specifically accepts traffic from the Internet ('dirty'). If an external Denial of Service attack were aimed at the DMZ there is a potential that the core switches may be affected, resulting in a severe network outage. The Remote Access device is connected to the DMZ VLAN, and as a consequence, a large number of ports have been opened on the firewall to let legitimate traffic through from a range of source addresses that are supplied to the RAS user when they authenticate successfully. Access from the DMZ servers to the internal LAN is limited. An attacker could use the DMZ server as a platform to attack the internal network using a spoofed address from the range a RAS user is given. This would give them a large number of ports to infiltrate the internal network.

No intrusion detection capability.

There are no ACL's configured on the network devices. This means all internal VLANs including the Management VLAN are accessible from every other VLAN.

Availability

There is a single point of failure with the Internet connection. There is only one router and one ISP (Internet Service Provider).

The availability of the network could be compromised by a Denial of Service (DOS) attack capable of making its way onto the DMZ VLAN via the firewall affecting the core network.

Firewalls

I implemented and configured a pair of Nokia IP530's running Checkpoint 4.1 and IPSO-3.4.1-FCS11. The IPSO operating system is based on FreeBSD and has been specifically hardened for its role on the appliance.

VRRP (Virtual Router Redundancy Protocol) was configured so that if any of the configured interfaces of the primary firewall should fail all traffic would pass through the hot standby firewall that is configured to use the same ruleset. As Checkpoint's state synchronisation was being used, all existing sessions are maintained.

All traffic is subject to NAT (Network Address Translation) at the Nokias. All traffic is implicitly denied; rules must be created to allow traffic through the firewall. Anti-spoofing is enabled.

A site-to-site VPN exists between RCN Ltd and a third party who provide web development services. A pre shared secret is used for authentication; 3DES and MD5 are used for encryption and integrity respectively. No NAT is required as there is no address overlap with the third party. The third party has limited access to specific machines using specific protocols. They only have local accounts on those machines. Traffic through the firewall from the 3rd party is logged.

Home/remote users can connect to RCN Ltd using the SecuRemote VPN client from Checkpoint. Users connecting via this method had a second account created in the Win2k active directory. When they attempt to connect they are prompted for a username and password. They supply the requested information and the firewall checks the information via RADIUS (Remote Access Dial-In User Service). This account has no domain rights whatsoever. When the user authenticates successfully they can access the server VLAN using specific protocols. The majority of users access the terminal server farm so they can have their usual desktop at home; other users have laptops and require the same level of access as they do on the LAN. When a user accesses an internal resource they must authenticate again using a second username and password.

The management of the firewalls wasn't thought of in the initial design. Originally only firewall licences had been ordered. I pointed out that a separate machine to the firewalls running the Checkpoint SmartCenter software would be required with the appropriate licence. A Checkpoint NG SmartCenter licence was purchased and I installed the software on a desktop machine that I called FWMAN. As stated previously, the firewall management was an afterthought so a server had not been bought for this purpose and a desktop was all that was available. This machine was situated in the room used by most of the staff involved in the project. Cisco Works was then added as this seemed to have been forgotten about too. Network-wise, this desktop existed in the Management VLAN.

Risks

Confidentiality

The remote users generally have Windows 2000 (Win2k) as their desktop software. I only made the Win2k and Windows XP SecuRemote clients available as self installing packages and these were pre installed before the users took possession of their laptops. The operating system had not been hardened and the users do not have a desktop firewall. "…most home computers are insecurely configured. If an attacker discovers the home computer and takes it over, they may be able to use their access to the computer to leverage access to the corporate network over the employee's VPN connection.[2]" (Cole, Fossen, Northcutt and Pomeranz, Volume 2, p.979).

Users may write their passwords down as they have a specific username and password that is different to their domain username and password. An attacker could use this username and password to get themselves onto the network and discover more about the internal network.

Availability

There is only a single firewall manager and it is installed on a desktop that has no RAID controller. If the single hard disk should fail the firewall manager would have to be restored from backup. This desktop had not been added to the backup schedule!

Outlook Web Access

Outlook Web Access (OWA) is a method whereby a user can access their mail via an Internet browser and do most of the things a user can do from a normal mail client e.g. Outlook. OWA had been configured in such a way that an Internet user would type in a URL i.e. http://owa.rcn.co.uk/exchange in their browser. This name would resolve to a valid Internet address that was actually a virtual address configured on the firewall. This destination address was then modified using (NAT) to the virtual address of a cluster, the active node of which was running Internet Information Server (IIS) and Exchange. The user would be prompted for a username and password which would travel across the Internet in clear text.

Risks

Any HTTP based attack that the IIS server was vulnerable to that wasn't stopped by the firewall would directly affect the active node of an exchange cluster that ran the entire organisations e-mail. This risk could compromise Confidentiality, Availability or Integrity depending on the attack and what the attacker's goals were.

Confidentiality

Anyone able to intercept the traffic involved in the logon process could read the username and password of the user using OWA.

Proxy Servers

Two servers running Microsoft Internet Security and Acceleration Server 2000 (ISA Server) were configured as an array in cache only mode meaning the firewall functionality was not enabled. Being in array means there is a single point of configuration and that configuration is stored in Active Directory. Each machine had a single interface. Network Load Balancing (NLB) was configured to add some resilience to the array. Incoming requests would be load balanced across both of the machines.

ISA Server was configured to only allow members of the domain Internet access. Group Policy was modified so that user's proxy settings for their browser would be configured to use the ISA array. As the access was restricted to domain accounts, the ISA Server required authentication. To anyone logged onto the domain the authentication process was seamless. Due to the authentication each user's requests were logged against their username enabling tracking of the users downloading the most content for example.

Risks

Availability

A 'heartbeat' is a message to another array member containing information about the state of the cluster; 'heartbeats' occur every second. If a failure occurs in one of the nodes making it unable to send five heartbeat messages in a row the remaining node assumes responsibility for all traffic destined for the shared virtual address. NLB works at the network level. If the ISA Server services were stopped, NLB would continue to run quite merrily. In a two node cluster, 50% of outbound requests would fail.

In a switched environment a switch learns the MAC (Media Access Control) addresses of the network cards attached to it. Using multicast mode, the array members retain their own MAC address for their individual IP addresses and share a multicast MAC address for the virtual address. Switches don't automatically register multicast MAC addresses so the switch would send the traffic to all switch ports in the hope that it gets to the right place. This can result in switch flooding, thereby affecting the availability of the network.

Anti-virus

Trend Micro NeatSuite had been purchased to handle the Anti-virus needs of RCN Ltd. NeatSuite is made up of several components including; ScanMail for Exchange, ServerProtect for Windows 2000, OfficeScan for desktops and InterScan VirusWall for the Internet gateway.

ScanMail was installed on the Exchange cluster, ServerProtect on all servers and OfficeScan on all desktops. All of the products were configured to check the Trend Management component on MAN1 for updates daily. The laptops also check whenever they connect to the LAN.

The component that I had particular involvement in was InterScan VirusWall. The product is capable of scanning SMTP, HTTP and FTP. The product can be installed so that these functions are split. Only 2 machines had been purchased and the SMTP and HTTP scanning components were given their own machine. It was decided that FTP would not be allowed in the environment and this was blocked by the firewall.

The design included the use of CVP (Content Vector Protocol) functionality of the Checkpoint firewall software. CVP enables the Checkpoint firewall to pass traffic to a compatible content checking program on another machine. That program checks the traffic against certain criteria depending on its purpose. If the traffic passes the criteria, the content checking software informs the Checkpoint firewall and the traffic is allowed on to its destination. If the criteria test is failed the traffic is blocked. In this case for example, the firewall will pass all inbound and outbound messages to the SMTP VirusWall server to be scanned for viruses. Inbound messages are also checked to make sure they are actually destined for the correct internal domain i.e. recipient@RCN.co.uk protecting against being used as a mail relay. The SMTP VirusWall also adds a disclaimer in the form of an attachment to each outbound message and checks the content of the message for offensive words. HTTP content that has been requested is checked for malicious code by the HTTP VirusWall before it is returned to the user.

Risks

Confidentiality

E-mails are not encrypted. E-mails are placed in a queue before they are scanned. They are stored in a plaintext readable format with a file extension of .tmp. These files can be read in this state. If I can read inbound and outbound messages in this manner it is highly likely the administrators at the recipients end will have the same capability. If the external recipients company has an SMTP gateway in their DMZ that has been compromised it may be possible for others to read that e-mail also.

Integrity

Messages could be modified in transit by removing a message from the queue, modifying it and then placing the message back in the queue for distribution. The organisation provides no digital signature functionality.

Availability

As there is only one server of each content checking type configured, there is a single point of failure. If either the scanning application or network connectivity for example is lost the firewall will not let HTTP or SMTP traffic through, depending on which content checking server fails, as the messages or HTTP content cannot be verified.

URL Checking

Elron Web Inspector as it was called at the time (now Zixcorp Web Inspector) was deployed to monitor web requests. The product used various configurable methods to deny access to content deemed objectionable.

A port was configured to receive a copy of all traffic bound for the ISA Array (port mirroring). If the request wasn't approved Elron Web Inspector would send a packet that closed the connection on the client, meaning the client would not receive the content. The user who made the request was also logged.

Risks

Availability

There was only one Elron Web Inspector server. If the server failed objectionable content could be accessed.

The Improved Design

On top of the issues highlighted in the first section, the client added a few new requirements.

1. They were going to host a website internally that must only be accessible using SSL (Secure Sockets Layer). It transpired it would in fact be two websites, one being a test website.
2. The downtime for this service must be kept to an absolute minimum.

   ”I would rather buy another six servers than have half an hour's downtime.”
   - Corporate Services Director

3. An Intrusion Detection System (IDS) should be added to the network.
4. Traffic must pass through two firewalls of different types.

The following diagram shows a simplified version of my improved network design that was proposed to and approved by the client.

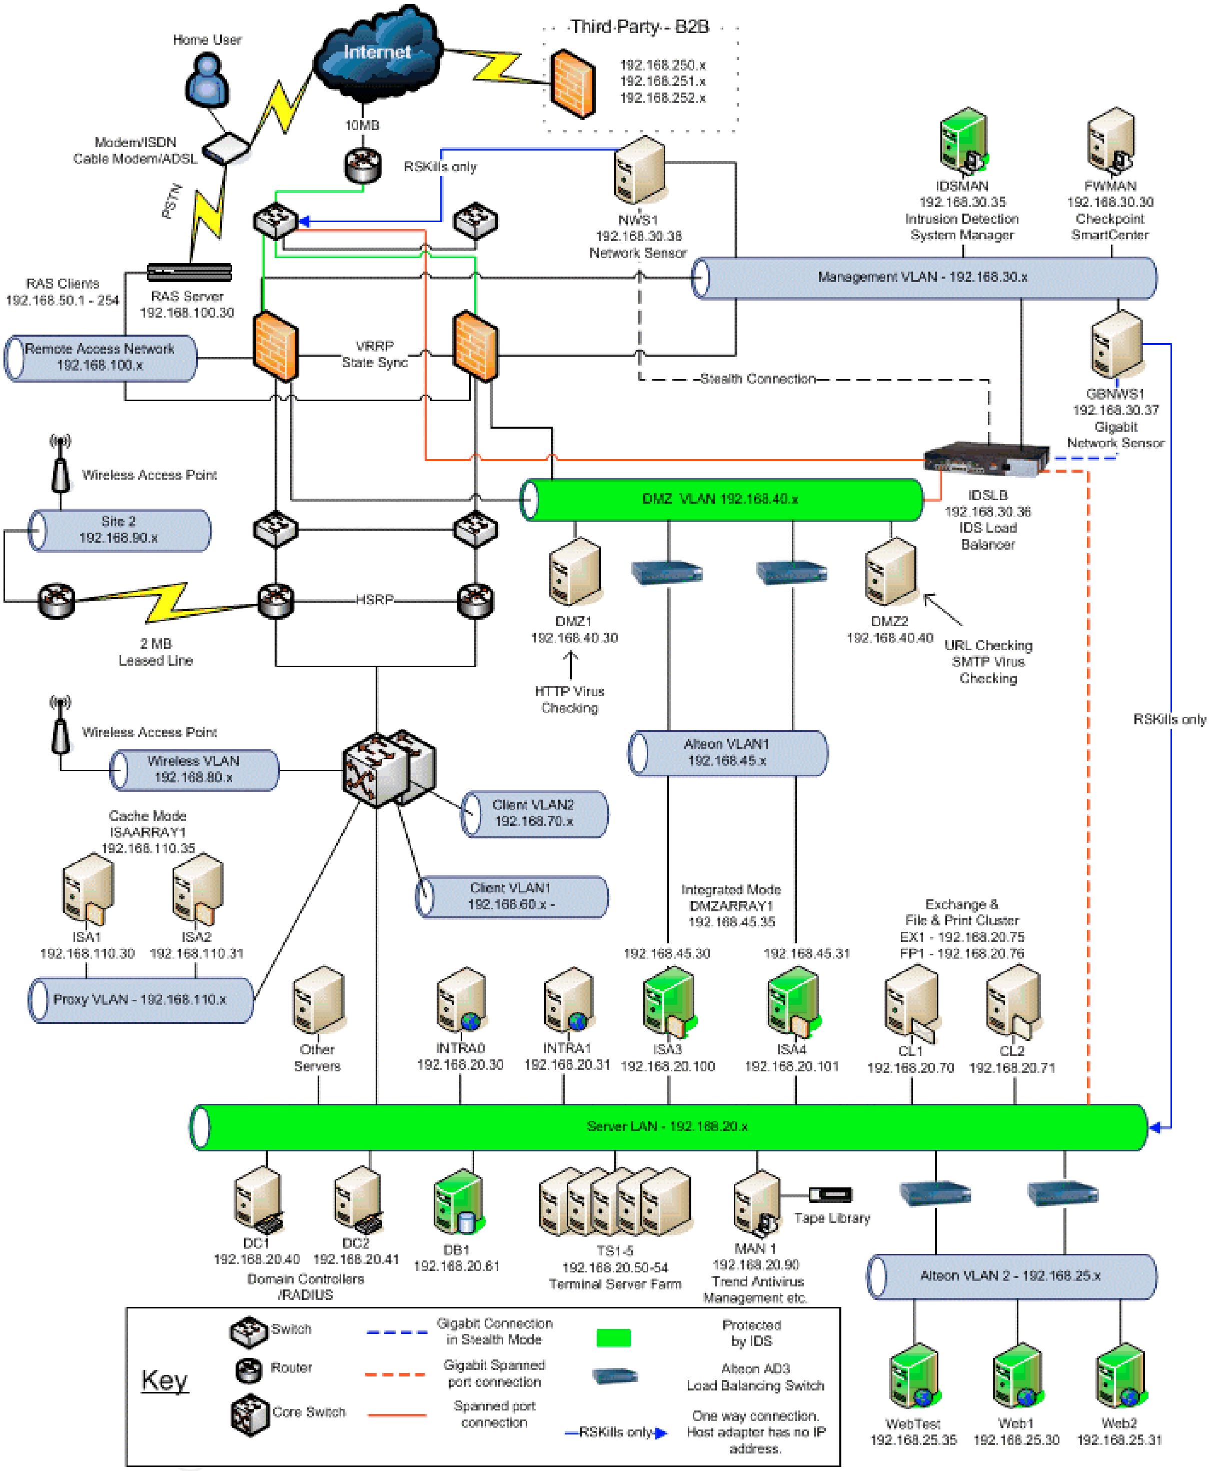

Network

The improved design introduces 2 additional 48 port switches so that 3 VLANs can be created separate to the core network.

1. Remote Access
2. DMZ
3. Management

The Remote Access VLAN will only contain the Remote Access device. This means the large number of ports that are open on the firewall are only accessible by users who have been authenticated. Traffic must pass through the Nokia firewalls either from or to this network.

The DMZ VLAN accepts traffic from external sources and has a limited access to the internal network using known services from known machines. Traffic must pass through either the Nokia firewalls or the ISA Server firewalls either from or to this network.

The Management VLAN is only accessible through the firewall so access is restricted. Traffic must pass through the Nokia firewall either from or to this network. The switches used for the creation of these VLANS are a resilient pair and both have the same VLAN configuration. Servers have connections to both switches using adapter teaming so that failure of an individual switch does not impact service. The ISA Servers are a resilient pair and have their own capability to deal with a switch failure; each ISA Server has one connection to the DMZ VLAN on different switches. As the switches have 48 ports each, there is room to grow any of these VLANs without additional change to the network infrastructure being required.

SPAN Ports are required so that traffic for a particular VLAN can be copied and passed to an Intrusion Detection device. A limitation of the 48 port switches is that they only support the creation of two SPAN ports. Only the DMZ VLAN has had a SPAN port created on both switches.

Depending on system load a second SPAN port could be created on the Remote Access LAN as there is a potential that the Remote Access device could be found by an attacker using a technique called war dialling. This technique involves dialling numbers using a modem until a corresponding device is found. The attacker would then attempt to gain access to the network.

The Management VLAN is not monitored by an Intrusion Detection device. External connections are not allowed into this VLAN. Internal connections are limited by port and host. Remote access via VPN can be given for a small number of staff for administration.

There was no need for routing changes internally as when the DMZ and Management VLANs were moved the traffic destined for those VLANs was passed to the firewalls as the default gateway; the firewalls then routed the traffic appropriately.

The core switches had their port capacity increased by 48 ports per switch so that additional servers required by the project could be connected and leave room for growth.

Three further VLANs were created. Two for the devices to be load balanced by the Alteon Load Balancing Switches, Alteon VLAN1 would have the external interfaces of the array of ISA Servers in integration mode. Alteon VLAN2 would house a pair of web servers for the live service and a test web server for UAT (User Acceptance Testing). The Proxy VLAN was just for the outbound proxy servers.

Wireless

WEP was configured at Site 2. In addition to that the Wireless Access Points were configured to disable broadcasting of SSID's (Service Set Identifier). It was recommended that another project was initiated that would introduce the up and coming (at the time) 802.11i standard using either LEAP (Lightweight Extensible Authentication Protocol) or PEAP (Protected Extensible Authentication Protocol). "The new 802.11i wireless standard uses two approaches to provide better security and protection than WEP." (Harris, Shon, p.28)[3]

The actual reconfiguration of the network was performed by CCIE qualified network engineers hired by RCN.

Firewalls

Original Firewalls

I realised when I was trying to download the latest version of the Checkpoint software that a software subscription hadn't been ordered. This means even hotfixes cannot be applied to the firewall. It annoys me intensely that you have to pay to get hotfixes that fix flaws in the product the company has paid a lot of money for. To be fair though this subscription fee also includes full product upgrades as they become available. In the long run I suppose it works out. Software subscription licences were duly purchased.

I migrated the Checkpoint management software to a more robust server platform. The server platform had a fresh installation of Checkpoint NG Feature Pack 3, Hotfix 1, management components. This server was added to the backup schedule of MAN1. Backups occurred through the firewall and the appropriate ports were opened between MAN1 and the firewall manager to achieve this.

An additional management tool was installed called Nokia Horizon Manager V1.2 and was used to automatically backup the two Nokia IP530 firewalls on a daily basis. It could also be used to apply future operating system updates.

If a firewall needed to be restored the configuration and application files would be restored from within the Nokia Horizon Manager software and the firewall policy reapplied.

If the management server were to fail the firewalls would carry on working although no updates to the firewall policy could be made. Logs would be stored locally on the firewall during this time. When the management server is restored from tape the logs would be uploaded from the firewall and policy changes can be made.
Log files would be stored for 3 months locally on the firewall manager and older log files would be available on tape for at least 6 months.

Four additional ports were added to the Nokia IP530's to support the two additional subnets that would protected by them.

Additional Firewalls

To meet the requirement that traffic destined for the internally hosted website should pass through two types of firewall, two ISA servers running Microsoft ISA Server Enterprise Edition Service Pack 1 Feature Pack 1) were configured as members of an enterprise array in integrated mode. SSL connections would be terminated at the ISA Servers and a new HTTP connection will be made from the ISA Servers to the internal Web server to complete the connection. A 128-bit SSL certificate is stored on the ISA Servers. This authenticates RCN's site to the user and sets up an encrypted tunnel. The certificate was obtained from VeriSign.

An AEP SureWare Runner 2000v3 SSL Acceleration card that supports 2000 (1024 bit) SSL transactions/sec was installed in each ISA Server to ensure response times to the site are kept to a minimum. These SSL accelerators were chosen purely on price/performance.

ISA Servers have some in-built intrusion detection capabilities licensed through ISS (Internet Security Systems) these capabilities will be enhanced by the addition of ISS RealSecure Server Sensor software (try saying that after a drinking game). By installing the Server Sensor software, the ISA Server can be turned into an in-line intrusion prevention device. As the traffic has to pass through the ISA Server all of the hosts behind the ISA Server are protected because the ISA Server can block those attacks at the external interface. The ISA Servers own security is strengthened as the server sensor monitors the event logs and triggers alerts if suspicious activity is detected.

Resiliency

Windows 2000's NLB function can only be configured on one interface per machine. With an ISA Server in integrated mode there must be minimum of two interfaces. Which side gets load balanced? The all important website had web servers that needed to be load balanced also.

It had also been stated that downtime was to be avoided (English understatement). As I described earlier in the document NLB is not application aware. If the ISA Server services were stopped NLB would still accept traffic causing a 50% loss of service in a two-node array. I had to find a load balancing solution that could balance traffic across both sides of the ISA array and would respond to a fault at the application layer (layer 7 of the OSI model) for the website traffic.

I chose some Alteon switches from Nortel Networks. They perform a layer 7 health check via the ISA servers configured as a load balanced group. If this health check fails the server is removed from the load balanced group until a successful health check occurs, thus providing protection from a service failure.

The diagram shows four Alteon switches, as it is a logical diagram. There are actually only two physical Alteons and one of these is a hot standby. VRRP is used again in the same manner as it is used for the firewalls.

I had been on a site where a colleague of mine installed Alteons for this kind of purpose and had good results. He was the one who actually implemented them for this client too.

Outlook Web Access

A certificate was purchased for the OWA URL. This was installed on the new ISA Array and the ISA Server configured to publish the OWA site and perform SSL to HTTP bridging. The users would now type https://owa.rcn.co.uk/exchange to access their webmail. The traffic was then subject to NAT and was directed the virtual address of the Alteons in the DMZ to be load balanced across the ISA array. The URL was checked to make sure the path was allowed and the SSL tunnel was terminated at this point. Then it was checked by the Server Sensor. A new connection was then initiated to the active node of the Exchange cluster to retrieve the content. This meant that the domain usernames and passwords would no longer be passed over the Internet in clear text and the Exchange server was further protected as the traffic was also checked by the IDS Server Sensor software as it passes through the ISA Servers.

Proxy Servers

The only thing that changed in the improved design was the fact that the ISA array was put in its own VLAN. This meant switch flooding shouldn't be a problem. I wasn't sure if this would work at the time but it was implemented and the network guys haven't complained about switch flooding.

I've since found another option, "Just make sure you configure any recalcitrant Cisco devices with static ARP entries for the NLB array MAC address and you'll do fine." (Shinder)[4]

URL Checking

Elron Web Inspector was removed from the design entirely. If the single server failed web traffic would have continued unchecked. Also, colleagues were surprised while testing, that they could access pages they had assumed would be blocked. I personally didn't configure it and wasn't given a chance to check the configuration and test. Added to that, there were certain 'thick client' users who had been given permission to bypass the proxy servers. This would mean the request wouldn't have even been seen by the Elron server where it had initially been situated.

It became apparent to me that the logical place to place this control was at the gateway. I did some research and decided on using Websense in conjunction with the CVP functionality of the Checkpoint firewalls. As the outbound mail isn't particularly time dependent and the utilisation of resources on the server checking inbound and outbound mail for viruses (DMZ2) was low I decided to install the software on DMZ2. This also freed up the server used for Elron to be used as one of the other additional servers required by the improved design.

Websense was configured to block certain categories of website and proved to block all attempted sites throughout testing.

Resilience was still an issue for me and was pointed out as a risk.

Anti-virus

The SMTP solution was changed from the original as it just didn't work as intended. E-mails were scanned for viruses ok; however, the CVP method relied on the Checkpoint software on doing the MX record lookup for mail delivery. If the first entry on the list failed other entries on the list were not tried at all. Because of this flaw in the SMTP Security Server of the Checkpoint firewall software the TrendMicro InterScan VirusWall for SMTP component was installed in standalone mode making it act as a smart host. This meant all outbound messages were forwarded by the Exchange server to DMZ2. DMZ2 did the necessary lookups before forwarding the message.

I had all sorts of problems with the HTTP virus scanning. It did work, technically. What I mean by that is that malicious code could be detected and stopped. The problem was that file downloads of any appreciable size did not complete successfully and the general web browsing experience was not acceptable to the client. When the HTTP scanning was bypassed, download and browsing performance were fine.

I worked with the product vendors support channels to fix both the downloading issue and also the performance issue. The downloading problem was fixed by upgrading to NG Feature Pack 3. Time was marching on with other areas of the project, making it harder and harder to test potential solutions on the live firewalls. Eventually the client ran out of patience and decided they would live with the potential risk rather than the actual performance problem. Since then several hotfixes have been brought out that mention the HTTP Security Server in the release notes that may well solve the problems.

Intrusion Detection

The Intrusion Detection System itself underwent a few design iterations and ended up like it did for various reasons.

1. It had been indicated the management and monitoring of the security infrastructure may be outsourced.
2. I had a finite budget for this part of the design.
3. The client wanted as wide a coverage as possible.
4. The client wanted evidence that the investment they were making was protecting them from actual intrusion attempts.

The IDS comprises of one Gigabit Network Sensor (RealSecure Gigabit Network Sensor 7.0 for Windows 2000) that monitors the DMZ and Server VLAN. One Network Sensor (RealSecure Network Sensor 7.0 for Windows 2000) that monitors traffic external to the firewalls. Seven Server Sensors (RealSecure Server Sensor 6.5 for Windows 2000) that monitor specific hosts. An IDS management server (SiteProtector 2.0 for Windows 2000) and an IDS load balancing switch (Toplayer AS3532). RealSecure Network Sensors inspect network packets and look for signatures that could indicate an attack against the network.

Monitored Networks

- 192.168.20.x (Server VLAN)
- 192.168.40.x (DMZ)
- Traffic between the Internet router and the firewalls

The RealSecure Server Sensors inspect network traffic directed at a specific host at multiple points in the network stack, allowing it to monitor and act on traffic that may have been encrypted.

Server Sensors

- IDS Manager (IDSMAN)
- ISA3
- ISA4
- Webtest
- Web1
- Web2
- DB1

Network

The Server VLAN has a Gigabit port configured as a span port to cope with all of the traffic on the Server VLAN. In reality a Gigabit Sensor can only cope with up to 600Mb. Traffic utilisation is so low that the entire server VLAN does not create this level of throughput. The DMZ VLAN has a single 100Mb port configured as a span port. Traffic in the DMZ is to major extent limited by the Internet pipe and this was now 10Mb.

The SPAN ports from both networks are connected to an IDS Load Balancer that consolidates all of the information from the separate VLAN's to one gigabit port that has a Gigabit Network Sensor running in Stealth Mode. Therefore one sensor was protecting two areas of the network. If the volume of traffic increases multiple span ports can be created on the switches and the load balancer allows for adding sensors easily too and as such adds flexibility to the design. I included this device in the design as flexibility was extremely important for a client growing at such a rate.

The Network Sensor used to monitor traffic external to the firewalls has a connection to the external switch in stealth mode. Both Network Sensors have a connection into the Management VLAN that is used to send event information and alerts to IDSMAN.

Stealth mode means that that interface has no IP address on the monitored segment. Network Sensors cannot be detected by an attacker and are immune to IP attacks on this interface. No reporting traffic appears on the production segment.

Network Sensors

External Network

The NWS in front of the firewall can detect attacks specifically directed at the firewalls including denial of service attacks. This is particularly important as the firewall acts as the control point for data flowing into the internal network. This sensor can also detect attacks attempting to pass through legitimately opened ports on the firewall.

Certain types of attack are prevented by configuring the NWS to actively terminate sessions. An example of something that is terminated at this level would be someone trying to telnet over port 25 (SMTP). This port is open on the firewall to allow inbound e-mail. Use of telnet over this port is evidence of an attack attempt. The network sensor recognises this attack and sends a reset packet to the source address of the attack and terminates the connection (RSKill). The reset packet has a spoofed source address which is the address of the machine the attack is aimed at. Attacks however are expected at the perimeter and the policy applied to this sensor reflects that. A port scan for example is an everyday occurrence and is treated as a low priority event. If this were to happen internally the priority would be higher. Reports can be run that identify if attacks are consistently being performed from specific IP addresses. Further action can then be taken.

Internal Networks

A Gigabit Network Sensor monitors the Server VLAN (192.168.20.x) and the DMZ (192.168.40.x). This sensor detects attacks that have passed through the firewall or have originated internally. Attacks directed at any server in those networks will be picked up by the sensor. This can also be used to check that VPN connections and SSL tunnels through the firewall have not been used to launch an attack. The Gigabit Sensor is configured to actively respond to attacks. As the firewall is configured to stop spoofing, reset packets for connections to the DMZ that are intrusion attempts, will not make it to their intended destination. The attempt is still logged however, with a high priority. The Internal Network is more critical for this function.

Server Sensors

Server Sensors have the ability to block attacks on a per host basis. A Server Sensor is installed on IDSMAN as this server controls what attacks are checked on all sensors, the responses to those attacks and all the alerts. If this machine was compromised evidence of further attack could be lost.

Sensors are installed on the ISA Servers and will block inbound attacks destined for the web servers or the Exchange cluster running OWA (Outlook Web Access) services.

Sensors are also installed on the web servers and DB1. These sensors will give protection from internal as well as external users.

Events are sent to IDSMAN via the firewall.

Management

IDSMAN runs SiteProtector which is used as the central controlling point for the Network and Server Sensors.

Policies are created on this management station and applied to the sensors. These policies will determine what attacks are monitored and what actions are taken should an attack be detected.

The Toplayer AS3532 IDS load balancing switch is managed from IDSMAN via a web interface.

Installation of updates for new attack signatures and service packs etc. for all sensors are performed from IDSMAN.

Backup

Both Network Sensors had a Ghost image taken when they were fully implemented. In the event of failure the Network Sensor would be recovered by fixing or replacing the hardware applying the Ghosted image and then reapplying the sensor policy. IDSMAN is backed up by MAN1. All hosts having Server Sensors installed are backed up as part of the normal backup routine.

Conclusion

Lots of time effort and money have been spent in trying to improve security at RCN Ltd. It can be argued that indeed security has improved. For example, inbound HTTP traffic for OWA has a very different journey through the network.

NOTE: The actual paths and sequences are much more complicated than described below; I describe a simplified view of this path that adequately demonstrates the point.

Before the changes:-

1. HTTP request is translated into a virtual address and routed to the external interface of the active firewall.
2. Traffic arrives at the firewall.
3. The firewall modifies the destination address to be the active node of the exchange cluster.
4. Traffic arrives at the exchange cluster.
5. User is prompted to authenticate.
6. Password travels over the Internet in clear text.
7. Authenticated user accesses content.

After the changes:-

1. HTTPS request is translated into a virtual address and routed to the external interface of the active firewall.

2. The traffic is mirrored to the SPAN port on the external switch and checked by a Network Sensor. This sensor can send a reset packet if an intrusion attempt is detected. Although SSL is being used, protocol anomalies can still be detected.
3. Simultaneously the traffic arrives at the firewall.
4. The firewall modifies the destination address to be the virtual address which is active on the active Alteon Load Balancing switch.
5. As the traffic is passed to the Alteon the Gigabit Network sensor sees the traffic and scans it. Any alerts are sent to a console. The traffic is still encrypted at this stage.
6. The traffics destination is modified once more to be the external interface of one of the ISA Servers.
7. The URL is checked for validity. If not the connection is dropped.
8. The Server Sensor checks for lower OSI layer intrusion signatures and has an opportunity to block the connection.
9. The SSL traffic is decrypted.
10. The Server Sensor has access to the unencrypted packets and checks them again.
11. The ISA Server creates a new connection to the active Exchange node.
12. The traffic is scanned by the GNS. A reset packet could be sent at this stage should an intrusion attempt be detected.
13. Traffic arrives at the exchange cluster.
14. User is prompted for authentication.
15. Encrypted password travels over the Internet.
16. User accesses content.

The website users follow roughly the same kind of path except they are also subject to a server sensor on the destination web server via more load balancing switches. This Server Sensor again, has the opportunity to block the connection.

Before the modifications RCN Ltd had no way of knowing if an intrusion had been made apart from some basic firewall functionality or if there were obvious results like a website defacement.

The IDS system also added some protection and detection of internal users or others who had gained access to the internal network physically or via the wireless connection.

The wireless access point at the 2nd site had some extra security with the WEP keys being configured and the SSIDs not being broadcast anymore.

I could go on citing examples but I hope they are clear enough within the Improved Design section. The important conclusion I have made is that a lot of time and effort has been wasted.

I have been a techie for years working on various technologies and for different kinds of company. Technically speaking I can stand up and defend my design although it may not be perfect and there are many aspects that have not been covered in this document as I had to limit the scope somehow. I may well use different products if I had the chance to go back in time to improve the technical solution that has been described. The biggest problem with the work I have done is that it is just a technical solution or more accurately part of one.

I had not really considered the lack of a test environment a security risk. During the whole security infrastructure implementation there were several other projects happening simultaneously. The company's aversion to downtime and aggressive pace of change became increasingly risky to the business. Because there was no test environment, changes were planned in theory and far too many changes were scheduled for far too small a window. The computer room was fully re-patched four times because it got so messy. It was allowed to get that way because the staff had to work so quickly trying to get those changes in, so other project timescales weren't affected. If a cable had to be stretched across cabinets or there was only a 5m cable to hand for a job that required a 1m cable, so be it.

The training I have received from SANS and through studying for my CISSP has opened my eyes to a much bigger picture and taught me some technical things too.

Something I was ignorant of was VLAN security issues. Having done a little research I'm surprised I hadn't come across it before but I come from an operating system architecture background. That's my excuse and I'm sticking to it. An excellent paper by Steve A Rouiller[5] concludes "attacking VLANs is tough but it's possible" (Rouiller, p.10) but there are also things that can be done to mitigate those risks. I since contacted the network guy at RCN Ltd and he informed me the switches that contain the RAN, DMZ and Management VLANs do not use any form of trunk, only VLAN access ports and that they are secure as they can be in his opinion. The core switches were another matter however but the risk was viewed as being low.

RCN Ltd has not got a security policy. They invented requirements on a per project basis. When I suggested a security awareness program for instance, the Corporate Services Director refused. He was happy to have his users left in ignorance. He seemed to think the users would be confused by it all. They have since been hit by blaster and welchia. There are two reasons for this. The users were not given enough information to update their software and virus signatures themselves. Investment was not made in systems that could do it automatically for them.

After the SANS conference in Hammersmith, London, I spoke to the IT Manager about assets so that future work could be targeted at the most important assets first. He informed me that not only did he not know but that nobody would be able to answer my questions. They were not even prepared to consider the ideas seriously as they were focused on other things entirely.

The company were happy to have all of the well known things a company should have as far as security is concerned. They had all the components described throughout this document after all such as firewalls and IDS. They had not given time to understanding how they applied to their business or understood the consequences of having them even though I amongst others had tried to explain it to them.

I suppose this mirrors childhood and adolescence though. The company was growing at an alarming rate; it was confused and far too caught up in itself to look forwards. It didn’t fully understand where it was going and wouldn’t listen to or couldn’t fully comprehend the advice being passed on to it. It would copy the popular things everyone was doing. The parent organisation that was funding everything was just demanding results and doing tick box checks.

Thankfully, the company had some budget issues and had to suspend progress for a while and get used to working in a stable state. A bit like higher education; a distinct lack of funds and making do with what they’d got for a while.

Their experiences in this situation have added weight to the recommendations made that were initially ignored or thought of as low priority. Early next year the finances are due to come back online and I am confident the growth of the company will be managed better. I hope they’ll begin to listen and understand the recommendations are business enablers with real life examples from their own experiences rather than draconian restrictions. Maybe I’ll be able to explain it better too.

References

[1] Ellison, Craig. "Exploiting and Protecting 802.11b Wireless Networks." 24 Sep. 2001.
URL: http://www.extremetech.com/article2/0,3973,11388,00.asp (24 Nov. 2003).

[2] Cole, Eric, Fossen, Jason, Northcutt, Stephen, and Pomeranz, Hal. SANS Security Essentials with CISSP CBK Version 2.1, Volume 2. SANS Press, 2003.

[3] Harris, Shon. "Greater WLAN Security with 802.11i." Windows & .NET Magazine. August 2003 (2003): 28 – 32.

[4] Shinder M.D., W, Thomas. "Using NLB with ISA Server Part 2: Layer 2 Fun with Unicast and Multicast Modes." 6 Feb. 2003.
URL: http://www.isaserver.org/articles/basicnlbpart2.html (28 Nov 2003).

[5] Rouiller, A, Steve. "Virtual LAN Security: weaknesses and countermeasures." 2003.
URL: http://www.giac.org/practical/GSEC/Steve_A_Rouiller_GSEC.pdf (4 Dec 2003).