Global Information Assurance Certification Paper

An Exploration on Creating a Secure Enclave within a Lower Security Environment.

GSEC Practical Assignment
Version 1.4b Option 1
Leon Buker

Abstract

Traditionally classified networks have been created in isolation from unclassified networks. Whilst this air gapped solution can be highly secure, it is also potentially expensive and cumbersome (from both the user and administrators points of view).

This paper attempts to explore the options around presenting information classified as protected, on a network only classified to in-confidence. The risks involved are explored, and a number of possible solutions are put forward.

Some Assumptions

This paper has been written around a number of assumptions that stem from a desire for it to be applicable in a commercial setting. There are two main things that fall out from this. First, an eye needs to be kept on the cost at all times. In a commercial setting, it is common that if it is going to cost more to secure *x* than it will cost if *x* is compromised, then the risk of *x* being compromised will be accepted instead of being mitigated. The second main assumption is that the commercial world revolves around the Microsoft product suite. All desktops are assumed to be running Windows XP, and that Microsoft Office is a requirement. The server environment is assumed to be Windows 2000.

There is a final assumption. Because this paper is exploring adding a more secure domain over a less secure one, it assumes that the less secure one already exists. Basic physical security elements, policies, procedures and processes are already assumed to be in place for the basic network, which is assumed to be of an in-confidence level.

It must also be pointed out that this paper has been written from an Australian perspective. Security terms such as in-confidence and protected are being taken from the Australian Commonwealth Protective Security Manual (PSM), and may differ from other international definitions.

Why Combine Multiple Security Domains?

One of the main reasons for combining multiple security domains is to reduce the potential cost of maintaining completely isolated operating environments. This cost is borne out directly in terms of the extra hardware that is required to provide dual environments. But there is also the hidden cost of lost

productivity. Cleared users do not have immediate access to all information in an integrated manner when isolated systems are used.

What Do We Need to Protect

Before discussing a way of creating a secure enclave, there needs to be an investigation of what functional requirements need to be met by this additional environment.

There needs to be a level of separation between the low security in-confidence network, and the new high security protected one. This needs to be done on a technical level to provide a separation of data, but also to allow users to be aware of which environment they are working in at any stage. One of the big concerns when access to multiple security domains is combined, is the risk that data could be accidentally declassified is greatly magnified.

The three 'A's of authentication, authorisation and accounting will need to be followed for all access to and within the protected network. This may mean more stringent requirements will be enforced against the in-confidence network, as the protected network will rely upon it to some extent.

Three main services will be required from the protected network: a secure way of storing and accessing files, a secure way of printing, and secure application provision. To achieve this, a secure file server, a secure printing solution, a method of securing the workstations which will be accessing the data, and the network which ties all of these components together will also need to be created and secured.

The above points all have a very technical focus. Over and above this, there needs to be a concerted education programme. Particularly as in a combined security domain environment there is a lot of responsibility placed on the users of the system to help enforce protection and separation of data between the two environments. Security policies and procedures will also need to be updated. It would be recommended that security reviews be conducted on a regular basis, particularly after the initial implementation, to help increase user awareness of the need for these new policies and procedures to be followed.

Creating a Secure Enclave

What follows is one possible solution that can be used to create a secure enclave. The principle being followed is one of defence in depth. The basic layout of the network can be seen in Figure 1. To keep the diagram somewhat simple, all of the extra hardware required to create a fully redundant clustered environment has been omitted. But for an actual implementation the firewalls, VPN concentrators and key servers would need to be duplicated and available in a fail-over configuration in order to provide a high level of availability.

Notes on VPN Configuration

In this particular design the client workstations only access servers in the protected zone through a virtual private network (VPN) tunnel. An IPSec VPN

tunnel is used to provide encryption and authentication of all network traffic between the in-confidence and protected networks.

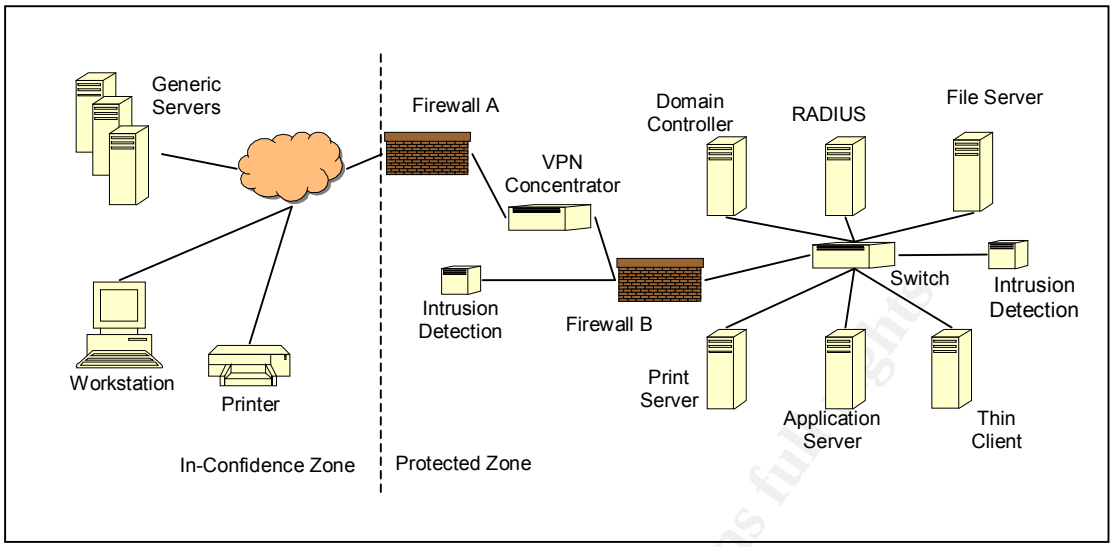

Figure 1 High Level Network Diagram

During the process of establishing the VPN tunnel the user will be authenticated a number of times. In this design the Cisco 3000 series VPN concentrators have been selected (primarily as they are an industry standard). The VPN client on the workstation initially authenticates against the concentrator using a group name and password. A set of properties for the VPN tunnel can be associated with the group name, and these are forced on the client during this initial authentication stage. Some of these properties include hours of access, number of simultaneous connections, encryption and authentication protocols and configuration of the client's built-in personal firewall.

The personal firewall plays an important role in this environment as a technique called split tunnelling will be used. What this means is that the user will be able to simultaneously access data on the in-confidence and protected networks. This is one of the design goals. There is also a significant risk involved as the user's workstation becomes a bridge between these two environments. The personal firewall software is configured to only allow outbound connections on the in-confidence network. This will protect their workstation to some extent from an attack. This will not work if a trojan application has been installed on the workstation, and initiates an outbound connection to an attacker's machine once the VPN tunnel has been established. This can be mitigated to some extent by configuring the personal firewall to only allow out bound connections on known ports. It will also be important to run virus scanning software which can also pick up on the presence of malicious applications.

It is possible to limit the range of destination IP addresses that the client can access from the VPN concentrator. By using VLANs and filters on the concentrator, set per authentication group, it is possible to hide or expose

servers in the protected zone depending on which group the user authenticated against.

It must be pointed out that VLANs are not considered secure on their own, although there is documentary evidence to support their use as an effective security mechanism (Pollino & Schiffman). The risk associated with VLANs in this environment is low, as all users authorised to access this environment (i.e. that are able to establish a VPN tunnel in the first place) have the required security clearance. VLANs are being used to help enforce the principle of need to know.

At this stage the user is required to enter their username and password. The VPN concentrator passes these details to a RADIUS server. (At least two RADIUS servers are required to be configured in a network load balanced configuration in order to provide high availability). In this implementation Microsoft's Internet Authentication Server (IAS) was selected to allow for easy Active Directory integration. The RADIUS server is able to authenticate the user by confirming the username password combination in Active Directory, and then is also able to evaluate a set of policies (more on this later) to confirm or deny that the user has authorisation to establish this VPN tunnel. RADIUS accounting can be used to help develop a detailed audit trail. It is also possible to use the extensible authentication protocol (EAP) to support the use of smart cards or biometrics as an additional authentication measure at this point.

The domain controllers in the protected zone are part of the same Active Directory forest and domain as the in-confidence network. This means that the same user account can be used to access both environments, making it easier and providing a more integrated experience for the user. There is obviously an element of risk associated with making such a choice, but in this case usability has been selected over security.

To limit administrative control of who has access to the protected network requires the careful delegation of permissions across Active Directory. All user accounts will be required to have their remote access dial-in permissions set to "control access through remote access policy". This is only available if the domain is a Windows 2000 native mode domain. It will force all authorisations to occur on the IAS RADIUS server against the policies defined there. It is possible to control access to this setting through ACLs in Active Directory; hence the ability to change this setting can be limited to a small number of authorised administrators.

The IAS RADIUS server's policies will primarily be configured to check that a user is a member of a particular group before authorising the user. The easiest way to control administrative access to these groups is to place them in a separate OU in Active Directory. ACLs can then be adjusted on this OU to once again limit access to only authorised administrators. A number of other properties can be tested for through IAS RADIUS server policies. Refer to the document Internet Authentication Service for Windows 2000 for further information on this.

Notes on Firewall Configuration

The configuration of firewall 'A' is quite simple. All traffic is to be denied, except for VPN tunnels and Active Directory replication traffic. The number of ports that are required to be opened for Active Directory traffic are numerous and provide some security concern. The protected network needs to be defined as being in a separate Active Directory site. This means that the replication partners, known as inter-site links, can be defined and controlled. This means that those ports opened in the firewall will only be allowed to connect to certain IP addresses.

The biggest problem with Active Directory replication is that it uses RPC over random high ports. This would require opening the firewall to an unacceptable level. It is possible to reduce the range of ports that RPC uses, but there are functional side effects (Riley). Another solution is to use the L2TP IPSec functionality that is built in to Windows 2000. The details for this approach can be read about in Riley. The downside is that when encryption is used like this, both firewalls and network intrusion detection devices will be unable to intelligently inspect the contents of these packets as they pass through the network.

Firewall 'B' is able to use stateful packet inspection on the unencrypted traffic leaving the VPN concentrator. Packets from users will have a different IP address allocated to packets that originate from the concentrators themselves. This means it is possible to hide administrative parts of the network from users, such as the IAS RADIUS server, domain controllers and network intrusion detection devices without affecting the functionality of the concentrators and the network.

Notes on Intrusion Detection

Network intrusion detection devices have been placed in two locations on the network: directly after the VPN concentrator and on the internal segment connecting all of the servers. There is no point in placing detection devices in front of the concentrator, as all this traffic will be encrypted.

Host intrusion detection can be installed on the servers, and optionally even on the workstations. Due to the high level of tuning that would be required for host intrusion detection to be considered useful in such a diverse environment, it would most likely only be deployed for the most paranoid of applications.

Some Other Configuration Notes

There are two other minor but important considerations that need to be addressed. All servers and network devices will use the network time protocol (NTP) to synchronise their clocks. This will make the reading of logs infinitely easier. The time service built in to Windows 2000 servers uses the simple network time protocol (SNTP). Although already available, this protocol has not been selected as within an Active Directory site it is only accurate to within two seconds and between sites can be as far out as 20 seconds. NTP does not display this limitation. (Brandolini & Green)

All logs will be collected and copied to a centralised logging server. This is important for two reasons. Firstly, if the local logs are compromised on a server there is some chance that the exported copy will not have been compromised as it is on a different server. Secondly, the Windows 2000 Active Directory is a multi-threaded multi-master application. This means that changes can be made on any domain controller and will be replicated through out the environment; but the only place this change will be audited is on the domain controller where the initial change was made. To get a complete historical picture of the Active Directory, the logs of all domain controllers need to be amalgamated. This amalgamation will occur on the central logging server.

Securing Servers

A lot has already been written on securing Windows 2000 as a file server, and hence not a lot of detail will be covered here. In 2002 Windows was certified to be EAL4 compliant and passing the common security criteria. Guidelines outlining how to implement servers using these guidelines are available from Microsoft (see References).

While it will be important to keep these servers patched against any security issues, it must be noted that each patch will need to be checked if it has any effect on the EAL4 rating of the server in question.

All servers will need to be racked in a C-class cabinet (tamper evident, see PSM) in a secured facility. All network cabling to these cabinets must be in clear conduit. It must also be possible to inspect all sides of these cabinets (although they can be placed in a single group).

A separate backup mechanism from that used in the in-confidence network will need to be used to ensure the appropriate treatment and separation of data is followed.

Secure Printing

There are four main issues to be considered in achieving secure printing. The primary one relates to print jobs being left on the printer and accessed by other users of a lower security clearance. If there are printers with different security levels available, it is desirable that a process is in place to help ensure that data goes to the appropriate printer. The third issue is that access to spool locations be restricted. And finally that printer network traffic is encrypted.

One possible solution is to use Capella Technologies' SecureJet product. SecureJet is a DIMM based solution designed to secure and track printing. It is simply placed in a free DIMM slot within an HP printer. (HP printers are assumed, as they currently supply approximately 80% of the market).

Print jobs are encrypted by the printer driver as they leave the workstation and hence before being sent to a designated print server and then printer, where it

is held on the printer's hard disk. To release the print job, the recipient enters a PIN code using the printer's front-panel keypad and display. Once the PIN is validated, the print job is decrypted and printed.

This certainly addresses the first and last requirements. Access to spool locations (third requirement) can be controlled through the careful use of NTFS ACLs on the spool directories on the print server. Because SecureJet encrypts the whole transmission of the print job, there is some extra protection afforded to the spool area. Auditing can also be done across the spool directories. But the SecureJet solution also creates another location where the print job is spooled: on the hard drive incorporated into the printer. It must be noted that disposal of this printer will need to meet the regulations surrounding the disposal of digital media containing protected information as outlined in the Protective Security Manual.

This solution does not address the concern of users sending print jobs to printers that are not SecureJet enabled. There are two simple solutions to this problem. One is to only have SecureJet enabled printers on the network and hence require all users to use a secure printing model. Another solution involves the use of thin client tools. Thin client tools can be used to present the user with a very specific experience, including which printers are available from particular applications. This is explored more in the section "Accessing the Secure Enclave."

One of the downsides of this solution from an administration point of view is that it does not integrate into the current Active Directory structure for authentication at the printer. A more integrated solution would make use of smart card technology that would require two factor authentication of the user against Active Directory.

A careful investigation of the encryption technique used by SecureJet will need to be carried out. From the documentation available it appears that DES is used to encrypt the print jobs. DES is no longer considered a secure encryption technique, being trivial to break using modern technology.

Accessing the Secure Enclave

There are two main ways that can be used to access the protected network through the infrastructure discussed above. It is possible to make a direct connection (using the VPN), accessing resources in the same way as they are accessed on the in-confidence network. Alternatively, a thin client application server can be used (over the VPN).

Directly accessing the secure enclave, via the VPN infrastructure, allows for a more integrated experience for the user. Adding a thin client solution will add another layer of complexity to both the user and the network administrators. An enterprise level thin client solution will also add more to the final cost of the solution.

Direct access also places a lot of responsibility on the user to correctly follow procedures and policies to ensure the correct handling of protected data.

The use of Citrix MetaFrame as a thin client solution will provide a relatively integrated solution compared to Microsoft's Terminal Server thin client solution. MetaFrame allows for the seamless delivery of applications to the desktop, making applications appear to be running locally and not requiring the user to understand the concept of having two desktops.

A thin client solution allows for a level of control over the user's environment that cannot be easily gained over a normal desktop. It is possible to restrict the ability of the clipboard to stop users doing a cut and paste between applications running on their desktop (i.e. on the in-confidence network) and those running on the terminal server (i.e. on the protected network).

Access to local printers can also be blocked with thin client, restricting protected applications to only print to those printers already mapped on the thin client server. These printers would of course be the ones fitted out with an appropriate secure printing solution.

Secure file server access could also be changed if only thin client access is used. One of the main concerns with direct access is that users are trusted to only copy things into the protected environment, and not to place copies on the in-confidence network. This can be enforced through the thin client solution, by not allowing local drive mappings or network drive mappings to be made. Of course another solution will need to be found on how data can initially enter the protected network.

The decision on whether to use thin client technology or not will depend on the availability of initial capital, whether the applications in use are compatible with thin client technology, and the future strategy for the IT infrastructure. If there is a high demand for the protected network, so that the majority of users are using it, it would most probably be more cost effective to look at raising the overall security rating of the network from in-confidence to protected, rather than providing a thin client solution on such a large scale where fat clients have already been deployed successfully.

Something that is discussed in Finnegan, Neumann and Lipner is the use of virtual machines. Whilst it doesn't provide the level of integration that is required in this theoretical solution, it is worth mentioning.

The concept is that a Windows XP workstation is setup with two local accounts. Each account is locked down to the extent that it can do nothing other than run a particular virtual machine. One virtual machine is used to access the in-confidence network, the other the protected network. Separate network cards can be installed in the workstation, allowing for a complete separation of infrastructure from the workstation back.

The user is able to switch between virtual machines quickly and easily by the using the fast-user switching mechanism that is part of Windows XP. But, the user is not able to run both virtual machines concurrently. This particular

solution could be expanded to allow access to multiple secure networks whilst still only requiring a single workstation per desk.

Some Other Technologies & Future Directions

There are a number of sophisticated technologies available to provide a solution along these lines. Most of this technology is being used in military environments, where budgetary and compatibility concerns are not always paramount.

Some of these solutions have involved the creation of wholly new operating systems, such as Trusted Solaris and Trusted HP-UX. This can provide a high level of mandatory security, by integrating controls into files, and other objects as well as users and processes that may be carried out. By placing this control on an operating system level, the chances are reduced of an authorised user unwittingly running a malicious application and declassifying data.

Whilst this might deliver an optimal solution from a security point of view, it does impose severe restrictions on available applications; in particular there is no version of Microsoft Office currently available for these platforms.

Microsoft is working on what it currently calls the Next-Generation Secure Computing Base (NGSCB); you may have heard it called Palladium previously. This may provide a fairly sophisticated solution as it will be "employing a unique hardware and software architecture … [to] create a protected computing environment inside of a Windows PC…" ("Next Generation Secure Computing Base"). Unfortunately it is not currently available, and hence is rather submerged in marketing hype.

Conclusion

There is still a lot more work to be done in exploring all of the possible security issues involved in collapsing two security domains to this extent. At this stage there is little technology available that can be used to help enforce the separation of security domains once they have been collapsed. The primary responsibility still rests with the user. For some organisations this level of risk will be unacceptable. For others, the cost savings will justify taking this risk. One possible application for collapsing these domains is to reap initial cost savings, whilst progressing towards upgrading an in-confidence network to a protected level.

References

Administering Your HP-UX Trusted System. August 1996.
URL: http://docs.hp.com/hpux/pdf/B2355-90121.pdf
(25 July 2003).

Australian Attorney-General's Department. Commonwealth Protective Security Manual. Canberra: Attorney-General's Department, 2000.

Brandolini, Shala and Green, Darin. The Windows Time Service. April 2001.
URL: http://www.microsoft.com/technet/prodtechnol/windows2000serv/maintain/operate/wintime.asp
(24 December 2002)

Cisco VPN 3000 Series Concentrators – Cisco Systems.
URL: http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/index.html
(15 July 2003).

Finnegan, Sean, Neumann, Bill, and Lipner, Steve. "The Challenge of Providing Access to Multiple Security Domains."
URL: http://www.microsoft.com/usa/presentations/FinneganLipnerNeumann_SecuritySummitWest03.ppt
(18 July 2003).

Internet Authentication Service for Windows 2000. 3 May 2000.
URL: http://www.microsoft.com/windows2000/docs/IAS.doc
(24 July 2003).

Next Generation Secure Computing Base.
URL: http://www.microsoft.com/resources/ngscb/default.mspx
(24 July 2003).

Riley, Steve. Active Directory Replication Over Firewalls. Top IT Tasks. March 2001.
URL: https://www.microsoft.com/technet/ittasks/tasks/adrepfir.asp
(21 July 2003).

Pollino, David and Schiffman, Mike. Secure Use of VLANs: An @stake Security Assessment. 7 August 2002.
URL: http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/stake_wp.pdf
(20 July 2003).

SecureJet 4.0 Secure & Control Printing, Copying, Scanning and Fax Usage. Fall 2002.
URL: http://www.capellatech.com/pdfs/securejet_whitepaper.pdf
(5 August 2003)

Trusted Solaris Operating System.
URL: http://wwws.sun.com/software/solaris/trustedsolaris/index.html
(25 July 2003).

Windows 2000 Security Configuration Guide. Version 1.0. 4 October 2002.
URL: http://download.microsoft.com/download/8/c/c/8cc94365-13d6-4975-bf69-9d4cd16a01a7/W2kCCSCG.pdf
(5 July 2003)