



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Richard Wanner
GIAC Security Essentials Certification (GSEC)
Version 1.4b

NetTop for Data Privacy through Secure Desktops

Abstract

This paper takes a building block approach to describing how the NSA's NetTop solution may be used in workstation architectures to create a secure computing platform for use in environments where sensitive or classified data must be processed in conjunction with non-sensitive data. The paper begins with a discussion of the challenges of secure computing in environments with data of varying degrees of sensitivity. Then the building blocks of the NetTop solution, SELinux and VMWare are highlighted. Finally a couple of potential architectures are described using this solution to provide a multi-workstation platform and when combined with personal firewalls and VPN technologies it can provide a secure platform for multi-sensitivity computing environments.

Table of Contents

Abstract.....	1
Introduction	3
Solution Components.....	5
Security-Enhanced Linux (SELinux)	6
Domains	6
Mandatory Access Control.....	6
VMWare	7
Putting it Together – NetTop	8
Architectures	8
Conclusions	13
Cited References	14
Other Material	14

Table of Figures

Figure 1: Typical Workspace in an Air-Gap Processing Environment [1]	3
Figure 2: Multiple Networks, Multiple Desktops	4
Figure 3: NetTop Workstation Architecture	5
Figure 4: Multiple Networks, Single Desktop	9
Figure 5: Two NIC NetTop Workstation	10
Figure 6: Single Network, Single Desktop.....	11
Figure 7: Multiple Processing Environments, one Network.....	12

© SANS Institute 2004, All rights reserved. Author retains full rights.

Introduction

Certain processing environments require special care in maintaining confidentiality of data. In certain heterogeneous processing environments, such as, government, military, and intelligence organizations some data is of sufficient sensitivity to require clear separation of processing environments. This has usually resulted in the creation of multiple distinct processing environments with dedicated workstations and dedicated networks. In the past the solution employed in this situation is an air-gap strategy; multiple independent processing environments with dedicated networks and dedicated workstations. Figure 1 shows a typical workspace where an air-gap solution has been deployed. A typical user workspace will require one workstation and corresponding network to access the Internet, yet another workstation and corresponding network for non-sensitive intranet access for accessing corporate applications, and a third workstation and network for access to sensitive or classified applications.



Figure 1: Typical Workspace in an Air-Gap Processing Environment [1]

Figure 2 portrays a simplified network diagram for an air-gap processing environment. Note the presence of multiple distinct networks, and workstations for each processing environment.




Figure 2: Multiple Networks, Multiple Desktops

The primary advantage of this type of deployment is that the physical separation of the processing environments prevents leakage of data between the networks. The only way data can cross between networks is by conscious physical effort by a person.

The downsides of an air-gap architecture are numerous.

- Increased capital or lease costs for workstations and networking equipment.
- Increased deployment, support, and maintenance costs.
- Users requiring access to multiple processing domains need to have multiple workstations in their work area. This leads to cluttered workspaces, and increased facilities costs due to the requirement for larger work areas.
- Physical security of sensitive networks is complicated due to the extension of the sensitive network into employee office areas.
- Areas with restricted space, or limited networking facilities, will not be able to meet all the processing needs of its users. A typical environment where this sort of data separation would be required is the military. Unfortunately, the military has unique IT problems in that some environments such as shipboard, or mobile facilities do not have the space or IT facilities to deploy an air-gap solution.

However, modern technology has provided a way to provide multiple processing domains using one workstation and potentially one network, while still providing a logical gap between the domains. The National Security Agency (NSA) has been promoting an architecture called NetTop. NetTop employs Security-Enhanced Linux to provide a secure operating system platform and VMWare to provide multiple virtual workstations on the same physical hardware. Hewlett-Packard (HP) has licensed this technology and is selling a commercial version of this architecture. The addition of VPN technology can provide logically separated networks, which permits the processing of data of different sensitivities over one network. The remainder of this document describes the components of the NetTop solution and some sample architectures which can be deployed using this solution.

Solution Components

The NetTop solution consists of two primary components; Security-Enhanced Linux (SELinux), and VMWare. Figure 3 shows a typical high-level NetTop workstation architecture.




Figure 3: NetTop Workstation Architecture

The underlying operating system is Fedora Linux or Redhat Linux with SELinux extensions. This provides a stable, secure operating system base. Virtual workstations are created inside of VMWare virtual machines (VMs), thus creating a logical air-gap. The following sections break down these components in more detail.

Security-Enhanced Linux (SELinux)

One of the National Security Agency's (NSA) mandates is to protect the information assets of the United States [2]. Because of this mandate they take a leading role in producing security standards for the US federal government agencies and military [3].

One of the projects the NSA has been driving is Security-Enhanced Linux (SELinux). SELinux adds "enhanced" security features such as domains and Mandatory Access Control (MAC) to Linux. Essentially, it adds a security server and support utilities to a standard Linux kernel. This security server provides strong typing for objects, multiple processing domains in the same physical host, and an object and domain access policy to SELinux. These features provide the ability to restrict user and process access to system resources such as files, other processes and memory [4].

In the 2.4 and early 2.6 Linux kernels, SELinux is configured using kernel patches and utility RPMs. By the time the 2.6 kernel is released in Fedora Core 2 and eventually in Redhat Enterprise Linux, SELinux support will be an integral part of the 2.6 Linux kernel [5].

It's ok. You can shout, "Stop! My brain is full!" I have introduced a lot of complicated sounding concepts in last few paragraphs. So let's go back and look at some of those concepts so you can clearly understand SELinux's capabilities and so you will be better equipped to understand how SELinux supports the NetTop solution.

Domains

SELinux provides support for domains and types. In SELinux every process including user shells are placed into separate domains, and every object on the system has a distinct type. Each domain can be thought of as a separate security domain. SELinux policy defines what access each domain has to each type as well as limiting interaction between domains. Users or processes in one domain cannot access objects in a different domain without the appropriate domain transitions or access rights defined in the SELinux policy [6] This combined with Mandatory Access Control permits a very fine-grained ability to limit access to objects.

Mandatory Access Control

Typical access control in Linux is based on a Discretionary Access Control (DAC) model [7]. DAC is identity based, and owner controlled. In a DAC based system each user on a Linux server is assigned a userid. This userid permits the user to identify himself or herself to the server. The user is allowed to create, delete, or modify objects, such as files, within his scope of control, and through permissions can restrict which users can access these objects. In other words who can

access an object is at the discretion of the owner of the object. This model is quite flexible. However this flexibility is also its biggest pitfall. The granularity of control is too broad, and cannot be restricted sufficiently to maintain the confidentiality of the object. Users, and by extension programs and processes have full control over the access given to files they create. For example once the owner of the object provides read access to an object, there is nothing stopping others from making a copy of the object and passing it on to others that the owner did not intend to have access.

What SELinux provides is a Mandatory Access Control (MAC) model. MAC mandates fine-grained access to all objects on the system [8]. In SELinux objects are assigned labels or attributes. These attributes are used to provide fine-grained control of access to object. In SELinux the fine-grained control and adherence to security domains is managed by the SELinux policy.

A Role-Based Access Control (RBAC) component assigns each user to a role. Under SELinux, each role has a list of domains that the role can operate in. The user cannot access domains which are not accessible from her role.

From a NetTop point of view this provides a mechanism to logically separate computing environments into different domains, thus reducing the likelihood access to sensitive processing domains even if the underlying SELinux machine is compromised.

Further detail on SELinux is beyond the scope of this paper. More information can be found on the SELinux home page at the NSA at <http://www.nsa.gov/selinux/index.html>.

VMWare

VMWare is a commercial application which permits the running of multiple virtual machines (VM) on the same physical workstation [9]. VMWare also permits the creation of virtual switches which can permit each VM to have logically separate networking from other virtual machines on the same physical hardware. VMWare runs on either Windows or Linux host operating system. In this case we are looking to use an SELinux host. But even though the host operating system is Linux the guest operating systems, the operating systems running in the VMs, can be any operating system supported by VMWare. This provides a great deal of flexibility from a deployment point of view.

From a NetTop point of view VMWare provides a mechanism to run virtual workstations and virtual networks of different sensitivities on the same physical workstation.

Putting it Together – NetTop

So, what have we got? NetTop is a Fedora Linux base, with SELinux extensions. This provides us with a secure computing base that is resistant to attack from the network, and permits separation of processing into separate domains. VMWare provides the ability to create multiple virtual workstations on the same computing platform. It also provides, through virtual networks, the ability to segregate the network inside the computer, so data of different sensitivities does not co-exist on the same network.

This in a nutshell is the NSA's NetTop architecture [10]. HP also sells a commercial solution based on this architecture [1], although it does not at this time take advantage of SELinux integration into the 2.6 Linux kernel, nor does it support VMWare version 4. However HP does have plans to update the Linux kernel and VMWare version in future releases.

Architectures

This section describes some architectures which can be used in conjunction with the NetTop architecture. This is by no means an exhaustive list of all possible architectures, but rather is designed to show a sampling of how the NetTop architecture could be used in a typical multi-sensitivity processing environment.

Before we can proceed, we need to add another component to round out the architecture. NetTop provides us with a secure computing platform that supports multiple virtual workstations, however it does not protect against network based attacks against the virtual workstations. To provide a level of security against these type of attacks, a personal firewall should be used. I recommend a centrally managed personal firewall such as InfoExpress's CyberArmor [11] or other similar products available. By using a centrally managed firewall in all VMs, separate policies can be applied to each VM thus providing a more conservative policy to sensitive processing domains, and a more liberal one where appropriate in less sensitive processing domains.

This brings us to the first architecture. In the legacy situation where multiple physical networks of varying sensitivities are present in the facility, NetTop provides us with the ability to eliminate the extra desktops and only use one desktop for connectivity to multiple networks. Figure 4 shows a high-level network view of what this topology looks like. Please note that each workstation has two network interface cards, one on the non-sensitive network, and the other on the sensitive network.

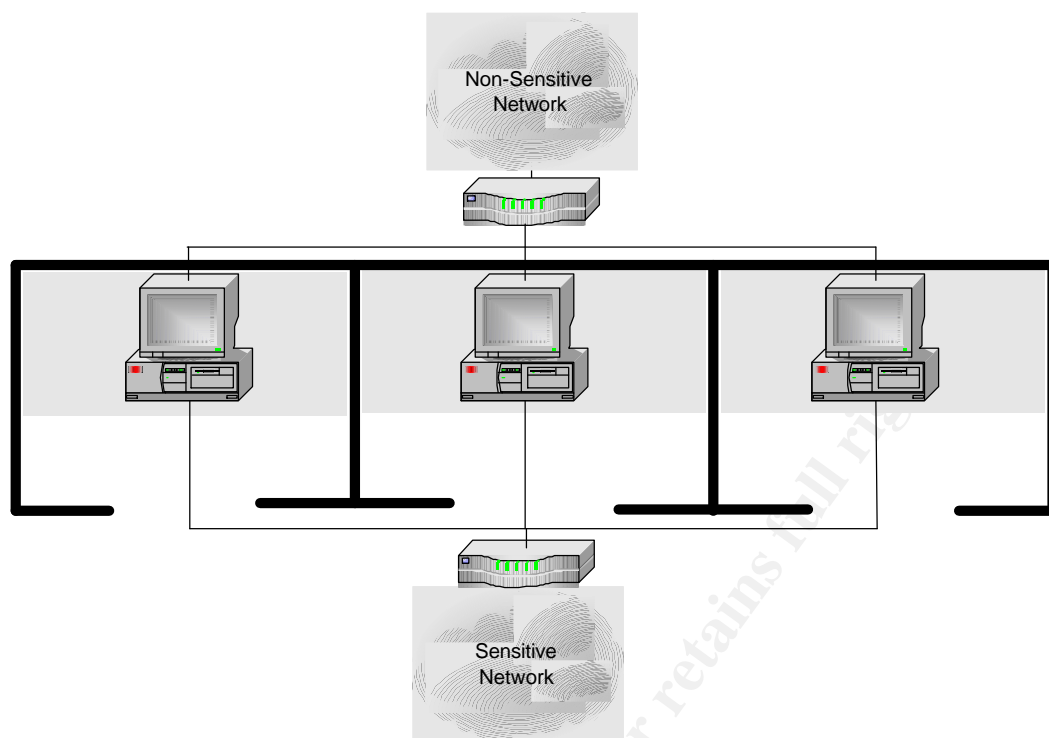


Figure 4: Multiple Networks, Single Desktop

Figure 5 shows a high-level view of the workstation architecture for this scenario. In this case, VMWare not only provides multiple processing environments by providing one VM for sensitive processing, and one for non-sensitive processing, but VMWare also provides multiple virtual networks to provide logical separation of the network data within the machine. Each of these virtual networks can be associated with a physical network interface in the host to permit the appropriate VM to access only the network which it should have access to.

In this scenario since the SELinux host operating system only exists to provide services to the VMWare application, the SELinux host operating system does not need, and should not be assigned an IP address. This reduces the possibility of a network based attack against the host operating system being used to compromise the VMWare based processing environments.

Each processing environment can be configured to require authentication, thus each processing environment can have different credentials to authenticate the user.

A personal firewall is installed in each VM to provide increased security to those processing environments. By using a centrally managed personal firewall, different policies can be applied to each processing environment depending on the sensitivity of the environment and the requirement to access network resources. For example the non-sensitive processing environment most

probably contains the majority of the corporate services, so it will require a more liberal firewall policy to permit access to file shares, printers, databases, etc. The sensitive environment probably contains very few services, and the firewall policy can be more strictly limited.




Figure 5: Two NIC NetTop Workstation

Our next architecture, as shown in Figure 6, shows an architecture where multiple processing environments are required, but only one physical network is present. This network will need to be used to carry both non-sensitive and sensitive data on the same medium. Fortunately, encryption technology provides us with the ability to secure the sensitive network traffic using Virtual Private Networks (VPNs). The VPN provides secure key exchange using IKE/ISAKMP and 128 bit or greater symmetric encryption for the actual VPN tunnel. All data

bound for the sensitive network will be encapsulated inside of the encrypted tunnel. Thus, using a VPN we can create virtual networks for encapsulating sensitive network traffic while still using a non-secured network as transport.




Figure 6: Single Network, Single Desktop

While either a gateway-to-gateway VPN, or a client to gateway could be used in this architecture, a client-to-gateway VPN provides several advantages. Firstly, gateway-to-gateway VPNs generally use a pre-shared secret key to establish the VPN tunnel. If this pre-shared secret is compromised it is possible to compromise the VPN. With a client-to-gateway VPN the key is negotiated using IKE/ISAKMP, thus making the compromise of the tunnel substantially more difficult. Secondly, client-to-gateway VPNs generally require a user to authenticate to establish a VPN tunnel with the gateway. This authentication can be used to provide a secondary mechanism to validate the identity of the user in the sensitive processing environment.

Figure 7 shows a high-level view of the workstation architecture for this scenario. In this case, VMWare once again provides multiple processing environments by providing one VM for sensitive processing, and one for non-sensitive processing. But in this scenario, because a VPN is used to segregate the networks, VMWare only provides one virtual network within the host to handle both the sensitive and non-sensitive network traffic.

In this scenario as well, the SELinux host operating system only exists to provide services to the VMWare application, so once again the SELinux host operating system does not need, and should not be assigned an IP address. This reduces the possibility of a network based attack against the host operating system being used to compromise the VMWare based processing environments.

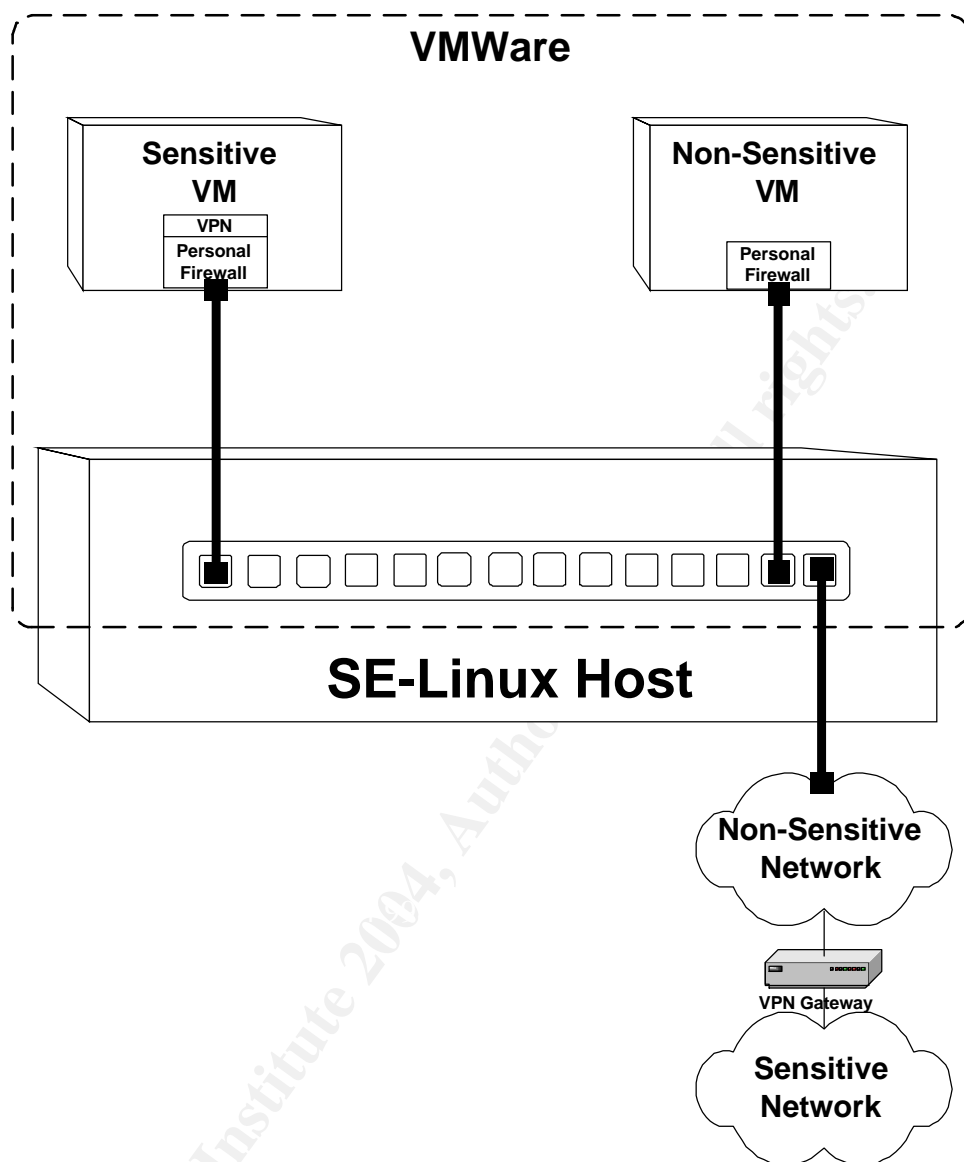


Figure 7: Multiple Processing Environments, one Network

Each processing environment can be configured to require authentication, thus each processing environment can have different credentials to authenticate the user.

Once again, a personal firewall is installed in each VM to provide increased security to those processing environments from network based attacks.

The client-to-gateway VPN provides a secondary authentication mechanism for access to the network, as well as providing an encrypted transport mechanism for the sensitive network data over the non-sensitive network. Once the VPN terminates at the VPN gateway, the sensitive network traffic will be able to traverse the sensitive network the same as any directly connected host.

This sort of architecture also provides some major advantages over the air-gap architecture. The first advantage is that the physical part of the sensitive network does not need to be deployed all the way to the desktop. In fact the physical workstation can be deployed at a substantial distance from the actual sensitive network, in a separate building or facility, a different country, or even on a different continent. Also, if this workstation architecture is deployed universally when access to the sensitive network is required, the sensitive network can be completely secured in a secured room, or facility with its only exposure to other networks being the external interface of the VPN gateway. This is substantially more secure and substantially cheaper from a deployment and support point of view than deploying the sensitive network all the way to the desktop.

Conclusions

This document has attempted to show that SELinux when combined with VMWare can be used to create a secure desktop architecture which when combined with personal firewalls, and VPN technology can be used provide an alternative to multiple desktops, and multiple networks, in environments where sensitive data needs to be processed such as military and government facilities.

This sort of architecture can be used to reduce the complexity, deployment costs, and maintenance costs of deploying multiple desktops and multiple networks, and can even extend the processing of sensitive data to environments where it would not previously have been feasible such as aboard ship, or mobile installations.

Cited References

- [1] HP Inc., HP NetTop Product Brochure
- [2] S Rajnic, An Introduction to the NSA's Security-Enhanced Linux, <http://www.sans.org/rr/papers/32/232.pdf>
- [3] DiBona, Chris. "Security Enhanced (SE) Linux, The OS that Came in from the Cold"; Linux Magazine, September 2001, http://www.linux-mag.com/2001-09/se_linux_01.html
- [4] National Security Agency, Security Enhanced Linux Frequently Asked Questions (FAQ), <http://www.nsa.gov/selinux/faq.html>
- [5] J Lettice, Red Hat Fedora plans Linux kernel 2.6 for April, <http://www.theregister.co.uk/content/4/34614.html>
- [6] K. Thompson et al., The UnOfficial SELinux FAQ, <http://www.crypt.gen.nz/selinux/faq.html>
- [7] M. Curphey et al., A Guide to Building Secure Web Applications, Chapter 8. Access Control and Authorization, Discretionary Access Control, <http://www.cgisecurity.com/owasp/html/ch08.html#id2859716>
- [8] M. Curphey et al., A Guide to Building Secure Web Applications, Chapter 8. Access Control and Authorization, Mandatory Access Control, <http://www.cgisecurity.com/owasp/html/ch08s02.html>
- [9] VMWare Inc, VMWare 4 User Manual, http://vmwaresvca.www.conxion.com/software/ws40_manual.pdf
- [10] R Meushaw, D Simard, A Network on a Desktop, NSA Tech Trend Notes, p. 3, <http://www.vmware.com/pdf/TechTrendNotes.pdf>
- [11] InfoExpress Corporation, CyberArmor – Enterprise Class Personal Firewall, http://www.infoexpress.com/security_products/firewall_overview.php

Other Material

- P Loscocco, S Smalley, Integrating Flexible Support for Security Policies into the Linux Operating System <http://www.nsa.gov/selinux/freenix01-abs.html>
- P Loscocco, S Smalley, Meeting Critical Security Objectives with Security-Enhanced Linux <http://www.nsa.gov/selinux/ottawa01-abs.html>

S Smalley, Configuring the SELinux Policy,
<http://www.nsa.gov/selinux/doc/policy2.pdf>

D. McCullagh, R Zarate, Super-Secure Linux, Inch by Inch,
<http://www.wired.com/news/linux/0,1411,53004,00.html>

G. Gross, SELinux Aims For Security Certification Among Cautious IT Professionals, <http://www.landfield.com/isn/mail-archive/2002/Apr/0002.html>

R Meushaw, D Simard, Commercial Technology in High Assurance Applications, NSA Tech Trend Notes, p. 2, <http://www.vmware.com/pdf/TechTrendNotes.pdf>

W Jackson, NSA Hones Secure Desktop To Run Multiple Oses,
http://www.gcn.com/22_28/tech-report/23568-1.html

E Leuning, NSA Attempting To Design Hack-Proof Computer,
<http://zdnet.com.com/2100-11-527810.html>

R Lemos, NSA Looks to Linux for Virtual Security, <http://news.com.com/2100-1001-251927.html?legacy=cnet>

© SANS Institute 2004, Author retains full rights.