Defence Research and Development Canada
Recherche et développement pour la défense Canada

Applying Virtual Machine Technology to Achieve Multi-Level Security

A Conceptual Technical Overview

Glen Henderson and Larry Tremblay

Defence R&D Canada – Ottawa
CONTRACT REPORT
DRDC Ottawa CR 2006-087
March 2006

Canada

Applying Virtual Machine Technology to Achieve Multi-Level Security

A Conceptual Technical Overview

Glen Henderson
Cinnabar Networks, a Division of Bell Security Solutions Inc.

Larry Tremblay
Cinnabar Networks, a Division of Bell Security Solutions Inc.

Prepared by:

Cinnabar Networks, A Division of Bell Security Solutions, Inc.
265 Carling Ave., Suite 200
Ottawa, Ontario
K1S 2E1

Project Manager: Glen Henderson (613-296-3716)
Contract number: W7714-5-3171
Contract Scientific Authority: Dr. Steve Zeber (613-991-1388)

Defence R&D Canada – Ottawa

Contract Report
DRDC Ottawa CR 2006-087
March 2006

Abstract

This document presents a conceptual technical overview for utilizing virtual machine (VM) technology to achieve a Multi-Level Secure (MLS) solution. The goal of this effort is to define a solution architecture that can leverage the savings in space and infrastructure that can be realized through the use of virtual system images while still adhering to the security principles that define an MLS environment. As part of this effort, an emphasis has been made to illustrate the isolation that exists between virtual machines and the hosting environment in the areas of information processing, information storage and information transmission. This document provides an architectural approach that utilizes VM images that are distributed according to a flexible, yet secure, policy. This policy defines the conditions under which potentially sensitive system images can be distributed and accessed. It is the position of this paper that VM technology can be leveraged to provide a more effective MLS solution while still maintaining the needed separation to ensure sensitive data assets are protected.

This page intentionally left blank.

Executive summary

Applying Virtual Machine Technology to Achieve Multi-Level Security

Henderson, G., Tremblay, L.; DRDC Ottawa CR 2006-087; Defence R&D Canada – Ottawa; March 2006.

Introduction

This paper presents the result of research into the use of virtual machine (VM) technology to reduce the hardware costs, maintenance effort, logistical issues and physical space needed to host an MLS solution. Such a solution is shown to reduce these financial and opportunity costs while in no way compromising the principles that define MLS processing. This paper provides a comprehensive view as to how and where VM (and supplementary technologies) can be combined to multi-level and/or multi-caveat separation through a single user interface. Presented as a conceptual technical overview, this paper describes an end-to-end secure processing model leveraging current technologies to meet or improve upon existing IT practices while maintaining the current level of isolation for sensitive information assets.

Results

The required level of isolation between guest images and the host infrastructure can be met through the use of VM technologies. When a secure workstation is combined with a trusted mechanism for distributing VM images, the processing environment can leverage this isolation to ensure that there is a high degree of confidence in the separation of sensitive network data needed to achieve an MLS solution. The distribution mechanism presented in this paper is a centrally configured policy server that can make distribution decisions based on a variety of criteria beyond the simple matching of classification level of the network to the clearance level of the user. These advanced criteria include: where is this user located, what is the defence condition of the environment, is the user the current duty officer, etc. An additional feature of the proposed solution is the ability to store information on a portable storage device and have this information integrated with the distributed VM image. As a user obtains and loads an image, the personalized information, properly protected via a cryptographic solution, is merged into the processing space, tailoring the environment to each user's specification. In this way, VM images do not need to carry any personal information and can be used generically by any user on any workstation.

Significance

The solution architecture presented in this paper provides benefits in terms of reduced costs for operating an MLS solution in terms of hardware, maintenance, logistical and physical space. This solution also provides new capabilities to the manner by which control over information access is gated.

Future plans

A detailed solution architecture will refine the solution proposed in this paper to reflect the actual solution configuration. This architecture should be written to closely meet the requirements taken from this document and the target community to ensure that the solution will be acceptable to the users. This activity will also define the system, network and security configuration for components in the solution space.

In conjunction with the previous activity, an effort should be made to research an appropriate policy server component. It has been recognized that the policy server will likely require an element of software develop to create a component that provides the required policy enforcement and policy decision points. The extent to which this will be extension of existing software as opposed to a complete custom solution will be a primary decision point for this effort.

It is recommended that an initial prototype be developed to prove the viability of the proposed solution. The prototype can also serve to identify deployment constraints and obtain solution metrics. The prototype can also serve as a demonstrator to illustrate the usefulness of the solution and foster interest in pursuing a full solution deployment.

Table of contents

List of figures

List of tables

1. Introduction

1.1 Background

The need for Multilevel Security (MLS) has been recognized since the beginning of the mainstream use of computing technology. This need is seen most clearly in the defence community. The ability to apply security labels to information (classification levels) and limit access to users that meet the needed level of trust to access that information (clearance level) forms a base requirement for the processing of sensitive information. This broad statement encompasses two significant points, namely:

1. It should not be possible for a user to access an information resource that is classified beyond the user's clearance level (the no-read-up property); and
2. There should be mechanisms in place to prevent the leakage of information that is classified at a higher level to be made accessible to users at a lower clearance level (the no-write-down property).

Among the most successful approaches to achieving an MLS solution the Bell-LaPadula[1] security model has been most readily adopted for military uses. This model successfully adapts itself to a military context, while maintaining true to the principles of the MLS requirements. The Bell-LaPadula model was innovative in that the focus was the transfer information (e.g. messages, files, device usage) thus presenting MLS in a dynamic context that is more suitable for modern computing environments.

It was quickly recognized that a hierarchical set of classification levels did not provide the needed granularity over sensitive information that pertains to a particular community of interest. The ability to affect access controls over a subset of information resources at a single classification level has been instituted in terms of compartments or caveats. To access information that has been labelled with a caveat, a user's security level must also include the caveat designation.

The attainment of an MLS solution still remains a difficult problem to solve given the high level confidence that must be provided to assure the isolation of the information resources at each security level. Software based MLS solutions are expensive to develop and maintain and still remain potentially vulnerable to information leakage (both read-up and write-down) through system design flaws, viruses and application vulnerabilities. As a result a common method for achieving an MLS solution is to provide additional hardware to have several systems and networks in parallel, with each environment running at a separate classification level and no physical linkage between environments. This has been traditionally seen as the only method for assuring the isolation of information resources at varying classification levels.

While this solution incurs additional hardware and maintenance costs, there is an added drawback to the parallel hardware MLS solution in operational environments where space is at a premium, for example, on board naval vessels. What is needed is a solution that provides the isolation gained through separate hardware environments without the need to deploy hardware infrastructure. It is the intent of this paper to show that such a solution is possible through the use of Virtual Machine (VM) technology.

1.2 Project Purpose

This paper presents the result of research into the use of VM technology to reduce the hardware costs, maintenance effort, logistical issues and physical space needed to host an MLS solution. Such a solution must be shown to reduce these financial and opportunity costs while in no way compromising the principles that define MLS processing. This paper provides a comprehensive view as to how and where VM (and supplementary technologies) can be combined to multi-level and/or multi-caveat separation through a single user interface. Presented as a conceptual technical overview, this paper describes an end-to-end secure processing model leveraging current technologies to meet or improve upon existing IT practices while maintaining the current level of isolation for sensitive information assets.

1.3 Project Scope

The development of this conceptual technical overview is limited to the following elements.

1. This investigation should focus on what can be achieved through the application of current technologies and specifically reference commercial-off-the-shelf (COTS) solutions that are appropriate for the uses identified within this document.
2. This investigation must show how VM technology will exist with and leverage other elements of the deployment environment such as:

    a. Networking devices;

    b. Networks; and

    c. Encryption technologies.

3. This investigation must identify technical limitations and operational constraints that will impact the proposed architecture.

1.4 Critical Success Factors

Within the context of the project purpose and scope, a discussion with project stakeholders resulted in the identification of the following critical success factors for a VM-based MLS solution. These factors will guide the development of the proposed solution architecture.

1. The solution must meet or exceed the existing level of isolation for information resources as achieved by the physical separation of processing environments. Where this degree isolation cannot be met, the reasons for the degradation in isolation should be described, an indication of the severity of the security concern should be provided and methods to mediate this concern should be suggested.

2. The solution should provide benefits in terms of reduced costs (hardware, maintenance, logistical and space) and optionally bringing new capabilities to the manner by which control over information access is gated.

3. The solution must not incur undue negative performance in terms of system boot time, system response time or network performance. Preliminary hardware recommendations should be provided to ensure this level of performance should be provided where possible.

4. The solution must be sufficiently robust as to be acceptable for an operational military environment.

5. The solution must be able to scale to meet any anticipated demands (i.e. user, network, data, or application scalability).

1.5 Document Overview

This conceptual technical overview presents information in the following sections.

Operational Requirements: This section sets the stage for the examination of applicable VM technologies by stating the perceived requirements to which any proposed solution architecture must comply.

Technology Overview: This section will provide a review of technologies and products that will have a role in defining the proposed solution architecture. Specific details, such as how VM technologies achieve data separation, will be presented. The goal of this section is to develop an understanding of the capabilities of the technologies in play and build upon this understanding to propose a valid solution architecture.

Solution Architecture: This section presents the complete solution architecture within the context of the stated goals for this investigation. This section is divided into 3 layers, each of which shows a progression in the use of VM technology to develop a secure and robust solution.

Comprehensive View of Components: The section provides more details regarding the nature and interaction between elements in the proposed solution architecture. Specifically, this section relates elements in the solution architecture back to the stated requirements and illustrates how the solution meets the project goals.

End to End Security model: This section examines the solution at all points along the data processing path to illustrate how information isolation and protected is achieved.

This paper concludes with a discussion of the following:

1. Technical limitations of the proposed architecture;

2. Operational constraints that would apply to an implementation of this architecture; and

3. Recommended future research to continue to develop this model.

1.6 Solution Approach

In discussion with project stakeholders, it was recognized that an iterative approach to presenting this conceptual technical overview would be appropriate. It is the position of the authors that VM technology can be used not only to meet the project goals, but also that the application of this technology can be leveraged to achieve further improvements to the security, maintenance and operation of an MLS solution. However, in order to clearly demark the application of VM as opposed to extending a VM-based solution, the solution architecture is presented in 3 layers. Each subsequent layer builds upon the infrastructure and capabilities of the previous layers. It is it believed that this manner of presentation will assist in expressing the capabilities of the application of various technologies as well as identifying the future benefits such capabilities will bring to an IT community. Further, the presentation of a layered approach to MLS solution design will aid in the definition of future steps, allowing some each layer to be further defined (or developed) in sequence with each iteration bringing more functionality into the hands of the community being served by the MLS technology. A high-level description of each layer is given below:

1. The use of existing VM technology to achieve processing environment isolation;

2. Extending the solution to include the concept of downloadable VM images and portable devices; and

3. Using the VM/MLS solution in the context of policy-based image deployment architecture.

When viewed as an aggregated solution, all layers provide a comprehensive solution to achieving an MLS solution using virtual machine technology.

2. Operational Requirements

This section establishes a context for the examination of applicable virtual machine technologies by stating the perceived requirements to which any proposed solution architecture must comply in order to achieve the project objectives.

2.1 Multi-Level Security Principles

This section details the requirements dictated by the need for a multi-level security (MLS) solution.

2.1.1 Data Classification Separation

The ability to share information, most particularly with a military context, must be tempered with sufficient controls to prevent the leakage of sensitive information. This can be achieved by adequately isolating computing environments that house information at different sensitivity levels, gating access to these environments and establishing rigid controls for the exchange of information between environments. The traditional approach to enforcing multiple security levels has been for organizations to operate a separate computing infrastructure for each environment.

Multi-Level Security (MLS) solutions allow for isolated computing environments within a single discrete network of servers, storage devices and networking equipment. MLS has two primary principles:

1. Controls must exist to prevent users from accessing information at a higher classification than their authorization permits; and

2. Controls must prevent unauthorized users from declassifying information.

Effectively implemented, MLS systems ensure that data can be consolidated onto a single infrastructure, while maintaining the highest levels of assurance that it is only accessible by authorized users. Effectiveness, in this context, means that the following characteristics have been addressed in the design, development and deployment of the solution:

1. The MLS environment must control access to sensitive resources. As such, an MLS environment must operate according to a set of rules and assurances that limit data access to authorized users. This can be achieved through the implementation of the following functions.

   Mandatory access control (MAC) – With mandatory access control, access is restricted based on the sensitivity of the information and the authorization of the user. These controls cannot be bypassed or altered by anyone other than an authorized security administrator.

   Discretionary access control (DAC) – In addition to mandatory controls, systems can also allow for some degree of discretionary access control. This is accomplished through the use of discretionarily assigned access control lists that identify the users that can access a given resource and their level of authority (e.g. read, update, delete) with regard to that resource.

Both the resource owner and the security administrator can determine who can access the resource and with what authority.

All resources and devices within an MLS environment receive a security label. Bit-map checks are performed against requests to use the device. For example, if a print request by an authorized user is made to print a top secret document on a particular printer, a bit map check will compare the device's clearance level with the document's. If they match, permission will be granted; if not, it will be denied.

2. The system cannot allow the reuse of a given resource (e.g. system component / device) until it is purged of residual data.

3. The system must enforce accountability by requiring each user to be identified and by creating audit records that associate security-related events with the users that initiate them. Each MLS system user is assigned an identity that corresponds to that user's security label. This identity must be established through a trusted Identification and authentication process. A supporting auditing function must associate security-related events (such as file access) with the user that caused the event. The audit record uses the security label to show when the data was accessed, the level of authority that was required and the actions that were taken.

4. The system must label all hardcopy and electronic data with relevant security information through a security labelling process which includes electronic labels for information resources.

5. The system must be able to hide the names of data sets, files and directories from users who do not have the "need-to-know" to access them. The names of files, data sets and directories are only displayed to users with access authority. Users without a "need-to-know" will not see the file or object listed or displayed.

   An MLS system will prevent a user from declassify data by "writing down" to a lower classification level than the classification at which the data was originally created. The user cannot 'write down' the data by labelling it secret or sensitive, in order to grant access to users with less than a top-secret designation. However, a user can create a document with a lower classification level than their current clearance level. If a subject is to be simultaneously granted write access to an object and read access to a second object, the classification of the first object must dominate the classification of the second object. The "no write down" rule is not in relation to subject classification, but rather in relation to the classifications of all of the various objects at the particular simultaneous point in time.

These functions must be present and effective in order to achieve a trusted MLS solution.

2.1.2 Multi-Caveat Separation (MCS)

MCS is a complementary extension to MLS in that information resources can be compartmentalized according to communities of interest. These compartments, within sensitivity levels, enforce the need to know security principle. Hierarchical security designations generally do not map well outside of military and similar environments, which are rigidly defined and controlled.

There are a few major differences between MCS and MLS:

1. All information within a given category exists at a specific classification level.

2. Some MLS concepts such as the No Write Down principle do not apply to MCS categories, since the former are designed to prevent leakage from high security levels to low security levels. However, controls must be in place to ensure that there is sufficient isolation between compartments.

3. MCS provides more discretionary control over information resources.

Any proposed MLS solution must include the ability to compartmentalize information.

2.1.3 MLS Guards

In many MLS/MCS deployments, there remains a need to transfer information between sensitive networks. The process that is responsible for managing this process in a controlled manner is known as an MLS Guard. Guards reside at the interconnection point between networks, that is, the security boundaries, and act as a trusted bridge between these environments. Information flow is often restricted to a one-way transfer (usually from low-to-high). However, bidirectional data flow is often appropriate for MCS environments. In either case, controls must be applied to the information flow to ensure that the data transfer occurs according to stated guidelines and security principles. As defined by these guidelines, guards will coordinate the information transfer by applying of the following controls:

1. Verification the security label associated with information to be transferred from the 'high' security network is appropriate for the classification level of the "lower" security network;

2. Application of content controls (e.g. a sensitive word check);

3. Prevention of the transfer of malicious code;

4. Recognition of the authorizing agent requesting the information transfer;

5. General network protection mechanisms (e.g. DoS attacks);

6. Accountability and Auditing Functions; and

7. Flag inappropriate or questionable data transfer requests to the network security officer.

The controls over which the data transfer request must comply will depend on the nature of the transfer. The progression of low-to-high, inter-caveat and high-to-low data transfers represents an increasing level of risk for data disclosure. The controls used by MLS guards should always be appropriate to mitigate this risk.

2.2 Solution Specific Principles

This intent of this research effort is to examine the use of virtual machine technology as a means for achieving an MLS solution. As part of the context for this effort, the following capabilities were identified as required for any proposed MLS solution documented during this investigation.

1. The proposed solution must allow simultaneous access to all networks through a single interface.

2. The proposed solution must be able to meet security policy and certification and accreditation (C&A) requirements. Section *8*:*Certification and Accreditation* provides additional detail regarding the VM/MLS solution in the context of a C&A process.

3. Technology Overview

The following section proves a description of some of the currently available technologies and solutions that are expected to have a role in the proposed MLS solution. The purpose of this section is to lay the foundation for the expression of the proposed architecture by describing these technologies in terms of capabilities, innovations and contributions to a secure computing environment. Where appropriate, specific examples of implemented and available products will be provided, however, at this stage of the MLS solution definition, no specific product is stated as a required solution dependency.

3.1 Virtual Machines

The original meaning of virtual machine (VM) is the creation of a number of different identical execution environments on a single computer, each of which exactly emulates the host computer for a user. Each user's virtual environment presents the user with what appears to be their own, private, computing platform for their exclusive use. Virtual machine technology today is generally oriented towards providing a number of virtual environments for a single user. Today's VM software typically allows the user to host numerous virtual machines within the virtual machine environment (VM environment), which may or may not execute the host system's operating system (OS). As an example using current technology, it is not uncommon for a user to host the VM environment on a Linux host, with Windows executing in a virtual environment (or vice-versa). This allows the user to use one physical computer to perform work in two different operating systems. Further, each virtual machine can connect to separate networks, if so desired.

In the case of MLS and MCS systems, we can use this technology to maintain the required separation between the different levels and/or caveats. One VM could host a network connection to a SECRET network, while another hosts a connection to an UNCLASSIFIED network. Yet a third VM could host a connection to a network carrying Canadian Eyes Only information, and so on. In this fashion, a single physical computer can host a number of environments at differing levels of classification and caveat. This approach requires less space for computer equipment since a single system provides all the processing functions that would normally take several systems.

The major barrier to using VM technology for MLS and MCS systems is the actual data separation. In most VM environments, there is no restriction in place to prevent the movement of data from one VM to another. Clearly, it is not acceptable to allow users to arbitrarily copy information from a TOP SECRET environment down to an UNCLASSIFIED environment.

In the United States, the National Security Agency (NSA) has recognised the benefits of using VM technology, as well as the problem of data separation inherent in the VM environment. In response, they launched an initiative to help create a VM solution that eliminates the data separation issues. The result is a Hewlett Packard product offering called NetTop .

The NSA provides the following high-level description of NetTop.

> *NetTop® is the result of a significant and continuing research effort based on the technical challenge referenced above. Broadly, NetTop is a building block-based architecture and associated prototypes, predominately based on COTS, and designed to address a series of different information assurance requirements. As a by-product, the NetTop architecture reduces the physical and environmental footprint issues typically encountered in high-level information assurance solutions.*
>
> *NetTop incorporates typical COTS user hardware and software found in most offices, schools, and homes. This technology is then combined with an underlying host operating system, virtual machine monitor, virtual network hubs, network encryptions, and a filtering router that allows multiple machine environments to run simultaneously and to access multiple networks all from the same physical platform.*
>
> *Additional research has been performed to address issues associated with the use of "thin clients," methods of providing increased assurance levels, and techniques that can provide failure detection.*
>
> *The benefit of the NetTop architecture is that it removes security functionality from the control of the end-user OS and applications. Important security functions such as communications encryption can be placed in a separate protected environment that cannot be influenced by user software. Similarly, an isolated filtering router function is used to provide protection from rudimentary network attacks. The modularity of the NetTop architecture and the use of standard TCP/IP networking to connect virtual machines facilitates simple replacement or upgrade of individual components.*

At this time, there are NetTop implementations available from Hewlett-Packard and Trusted Computer Systems. In both cases, the host system is the NSA's Security-Enhanced Linux (SELinux), with VMWare providing the VMs[5]. The VMWare software is modified to meet the NetTop standard. This results in a VM solution that is compliant with US DCID6/3 Protection Level 4.

A single physical computer hosting a NetTop system can provide VMs that enforce required MLS separation. When viewed as a hierarchy of progressively more sensitive computing environments, restrictions are placed on how information can be transferred between virtual machines:

1. Information transfer is blocked from higher level (more sensitive) VMs to lower level VMs; and

2. Information transfer is permitted (with some restrictions) from lower level VMs to higher level VMs.

The leveraging of the NetTop approach to isolation of sensitive networks will play a significant role in the development of a VM-based solution to achieving Multi-Level Security (the VM/MLS solution).

3.2 Portable Memory Storage Options

A thumb drive is one term used to refer to a portable memory storage technology that provides NAND-type flash memory accessible through a Universal Serial Bus (USB) interface. Most recent vendor thumb drive implementations utilize the USB 2.0 standard and allow for storage capacity of up to 4 Gig of memory with 16 Gig storage option to be available in the near future. While the general design of the thumb drive (IC memory, USB controller and interface) remains similar across most implementations, the applications to which these devices have been applied indicate the true flexibility of this technology. For example, recent PC systems have the ability to boot from flash drives, allowing network administrators to house an operating system with troubleshooting tools on a portable device.

Thumb drives have been used as an element of a security infrastructure. In addition to providing encrypted file systems to protect the information on the drive and biometric readers integrated in the USB thumb drive architecture, thumb drives have been used as physical authentication tokens allowing a system to be used only if the thumb drive is present.

As part of the read-only memory (ROM) profile, which exists for any USB enabled device, thumb drives contain a unique "vendor ID" and a "product ID" which allows the device to be differentiated from any other USB device (including different USB thumb drives models from the same manufacturer). Certain host systems, such as the Linux hotplug technology, automated the process of detecting USB hardware connection, detecting the device type and loading the appropriate driver. As part of the hotplug design, it is possible to specify a whitelist of recognized and allowable devices. Any device that is not specifically listed will not be assigned USB bus resources and will not, in effect, be attached to the system. Note that the USB 2.0 specification does not include the ability to distinguish between two USB devices that have the same "vendor ID" and "product ID".

In short, USB thumb drives can be used to provide the following key capabilities:

1. Securely store a large amount of system, application data (including system preferences and state information);

2. Act as an authentication token to unlock the system; and

3. Uniquely identify the user.

These capabilities will be leveraged to enhance the proposed VM/MLS solution.

3.3 Hardware Cryptographic Modules

The files that hold the VM images (or any other component of the proposed VM/MLS solution) can be protected via encryption. The security achieved through encrypting files comes at a price, namely, processing time. In the case of encrypting VM image files, the file must be decrypted prior to being loaded by the VM software, and when the VM is shut down, it must be encrypted again prior to storage. From the user's point of view, long encryption and decryption times may foster a perception of poor performance. In a defence condition, long delays for decryption and encryption may be unacceptable as it may directly impact the ability of a user to respond to an

immediate threat. A worst-case scenario where a user has to shut down a running VM in a damaged area and quickly restart it in an undamaged area underscores the drawback of slower encryption.

Given the size of image files, the use of hardware cryptographic modules (HCM) is likely to be very useful in alleviating long wait times. The addition of a HCM to the host system would allow for much faster decryption of the image files prior to being started by the VM software, as well as much faster encryption of the image after the VM is shut down. Certainly, the system will function well enough without a HCM, but user and performance constraints may dictate their use.

Transport security is achieved in the HP NetTop through the use of VPN technology. Within each VM image, a VPN is used to establish a secure channel to a VPN concentrator that can redirect the communication channel to the appropriate target network. In this way, isolation is achieved for communication channels that use a single Network Interface Card (NIC) on the workstation. This is discussed in more detail in section *4.1.1.1*:*Guest/Guest Isolation*. The NetTop solution does not dictate any specific VPN solution; however, the choice of other VM/MLS components may restrict the list of compatible VPN solutions.

3.4 Secure Operating Systems

The host operating system underlying the VM software must be a secure operating system. Just as an executing VM must be isolated from other VM's and the host operating system, there must be protections in place to prevent attackers from compromising the host operating system to use as a platform to attack an executing VM.

An example operating system is SELinux (Security Enhanced Linux)[3], which is a version of the Linux operating system with security extensions developed by the NSA. As with any Unix-type operating system, Linux can be secured against external attackers using standard well-known hardening methods. SELinux extends the hardening process by adding policy enforcement and role-based access controls to enforce access control between processes, allowing true multilevel security.

This type of access control can be used to specify processes that even super user access cannot touch. The VM software and executing VM's can be designated as untouchable using this mechanism. Thus, if an attacker manages to compromise the system and gain this level of access, which usually ensures unrestricted access to the system, the access controls will prevent the attacker from accessing the VM's, protecting the sensitive data being processed within them.

3.5 Policy-Based Networking

Within a policy-based network, access to network based resources is controlled and tracked based on characteristics of the user, the resource being accessed and the environment in which these elements interact. The rules that gate how access is granted are expressed in terms of a policy. The nature of the policy to be implemented can vary greatly, depending in the nature and operation environment in question.

At the core of a policy-based network is the policy server that acts as the arbiter between the submission of a request for information (e.g. a file) and the execution of that request. The policy engine, the logic core within the policy server, determines the manner by which the policy is expressed and evaluated. Many security products include policy engines for specific tasks such as network routing, email virus protection and workflow scheduling. LDAP solutions, including Microsoft Active Directory, have the ability to provide access control in a policy context for specific applications. Recent trends, however, have indicated the need to be able to express policy and policy queries to meet the need of general access control challenges.

Several initiatives have been put forward which meet this need for a general approach to security policy management. The extensible Access Control Markup Language (XACML)[4],[7],[8] defines the processing environment and protocol for exercising policy decisions in a generalized manner. Components in an XACML-based policy solution include the following:

- The Policy Enforcement Point (PEP) Any entity attempting to access a protected resource must go through a PEP which exists as part of the service that is delivering the information resource (e.g. a file server or web server). The PEP will format an XACML-based request using the requester's attributes, the resource in question, the action, and other information pertaining to the request and forward the request on to the Policy Decision Point. Based on the response from the PDP, the access to the information resource will be allowed or denied.
- The Policy Decision Point (PDP) matches an access control request (in XACML format) with the stated policy and provides a response to the PEP to allow or prevent access.

An XACML solution benefits from the fact that the technology is standard-based yet generic, allowing the development of a customized policy definition for the target usage. Additionally, the manner by which policy decisions are made allows for flexible response conditions as it pertains to the policy. For example, given a XACML-based policy, a request to access a specific resource may result in three different conditions for different users:

- User A is granted access to the resource;
- User B is denied access to the resource; and
- User C is granted access to the resource under conditions of increased auditing.

There are several XACML / Policy Server implementations in various stages of development. The most complete implementation, as of this writing, is provided by Sun Microsystems . Additionally, it has been recognized that limitations on the XACML standard can be addressed by using complementary protocols, such as the Security Assertion Markup Language (SAML). For example, XACML messages do not inherently carry trust information, being more focused on the expression of authorization requests. SAML allows the creation of an interface through which authentication and authorization requests can be transmitted. Combining XACML policy expression with SAML security assertion management leads to a powerful and complete solution. Given that these two standards operate in a similar domain, it is widely expected that these two standards will be merged into a single standard.

The technologies identified in this section are expected to play a role in the design of a VM/MLS solution, specifically addressing areas of security, performance and flexibility.

4. Solution Architecture

This section provides an overview of the proposed VM/MLS solution. In keeping with the objective of this engagement, the solution is presented in three layers:

1. The use of existing VM technology to achieve processing environment isolation;

2. Extending the solution to include the concept of downloadable VM images and portable devices; and

3. Using the VM/MLS solution in the context of policy-based image deployment architecture.

The first layer is based on existing technologies and available COTS solutions. Each subsequent layer introduces new functionality that leverages the solution architecture from previous layers.

4.1 Virtual Machines as Protected Processing Environments

This section addresses the general use of VM software and the SELinux/NetTop environment specifically, as it is the only solution presently available (as of this writing) that can enforce all of the requirements for MLS/MCS usage. Where appropriate, security enforcement that is imposed by NetTop to maintain MLS compliance is identified.

4.1.1 Environment Isolation

Isolation of the VMs is the keystone of a VM solution for MLS/MCS usage. The highest priority is ensuring isolation of the VMs hosting guest operating systems (Guest/Guest). VMs should also be incapable of interacting directly with the host operating system (Host/Guest).

VM isolation begins with the VM environment running on top of the host operating system. The VM environment performs virtualization of the hardware on the host system (known as the hardware abstraction layer or HAL) for each hosted VM. A Virtual Machine Monitor (VMM) acts as the arbitrator between VMs and the host system. The VMM enumerates the services and devices present in the host system (processor, memory, disk, I/O ports, network interfaces, etc.), and does not generally allow VMs to directly access the physical hardware for a given device. The VMM is also responsible for handling processor activities and memory functions for VMs.

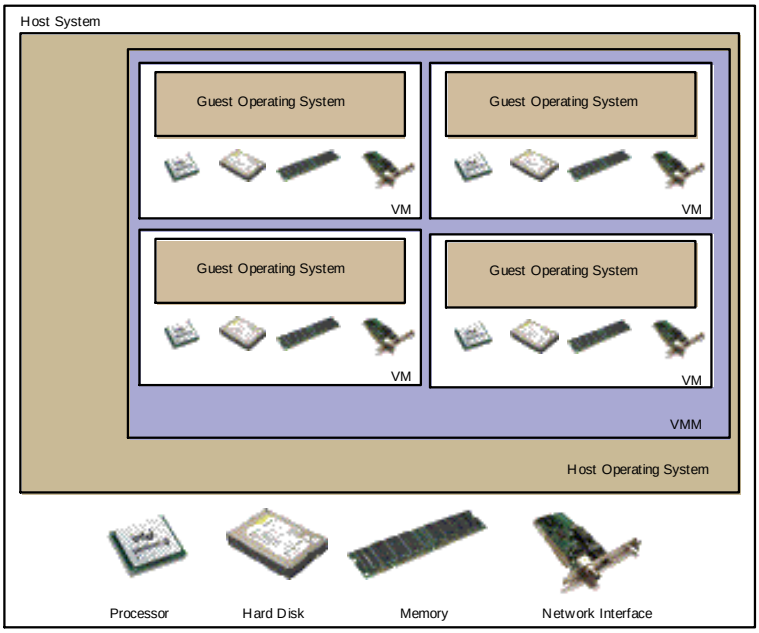

Figure 1: Typical VM Environment

As shown in Figure 1, the VMM presents each VM with a virtual copy of the host system. Each VM in the VM environment believes it is interacting directly with the host system. In actuality, access to the system and its devices is handled by the VMM. From the VM's point of view, it uses the same calls and procedures as if it had true physical access to a device, rendering the use of the HAL to perform device function transparent to the VM and the user. Instructions from VMs go to the VMM, which applies any access controls or other restrictions as needed, then executes them on the physical processor or performs operations in the physical memory. This is the basis of isolation between each VM and the host system.

4.1.1.1 Guest/Guest Isolation

Isolation between guest operating systems, each running in a separate virtual machine, is described here in detail by examining each of the significant areas where isolation may be compromised.

Processor Isolation: The VMM provides each VM with its own virtual processor that is a true representation of the host system's processor. As far as a VM is concerned, it is executing all instructions directly on the host system's processor. Instructions fall into two modes – privileged and unprivileged. On the target host platform for NetTop, instructions executed by a VM in an unprivileged mode are in fact directly executed on the host system processor as they cannot influence the operation of the host operating system or other executing VMs. Instructions executed in a privileged mode (e.g. the VM's operating system kernel) are communicated from the VM to the VMM, or more precisely, a binary translator within the VMM. The binary translator makes subtle translations in the instructions to remove instructions that may prove problematic for security or system stability, and then executes them in proxy for the VM on the physical processor. For example, if the VM is executing a privileged instruction to disable

interrupts, it is not desirable to disable interrupts on the host system, as this would likely interfere with the operation of the host system and other executing VMs. The VMM can replace the instructions to actually disable interrupts on the host system with null instructions, execute the modified instructions in proxy for the VM, and return the result to the calling VM. Then, to simulate the disabled interrupt state for the requesting VM, the VMM will not pass interrupt messages to the VM until it issues instructions to re-enable interrupts. In this fashion, guest/guest isolation is achieved.

Memory Isolation: The VMM maintains its own table of pages in the host system's physical memory associated with VMs, known as shadow pages. When a VM is configured, it is allocated a defined amount of memory, which any system information tests executed within the VM will report as the total system memory. For example, if the host system has 16 gigabytes of physical memory, and a VM is allocated 2GB, any query from within the VM will show the allocated 2GB and no more. All memory operations within the VM, including direct memory accesses, appear as normal to applications within the VM, and the memory allocated to the VM appears as a contiguous block regardless of how many disjoint pages of the host system's memory the VM may be using. When a memory operation passes from the VM to the VMM, the VMM translates the memory location as seen by the VM to the memory location in the host system's physical memory actually being used. The VM has no knowledge at all of the host system's physical memory, and cannot directly access it in any way. The VMM handles all translations and accesses, and maintains isolation of each VM's memory spaces by not allowing any VM to access host system memory allocated to another VM.

Hard Drive Isolation: When a VM is created, a very large file is created on the host system's mass storage device – called a VM image, or just image. This file must be large enough to fit the guest operating system to be hosted in the VM, data and applications to be hosted in the VM, and any incidental storage needed (virtual memory, etc.). The guest operating system sees this file as its primary mass storage device. For example, if a 20 gigabyte image is created and Windows XP installed on it, when a user is using that Windows XP environment, it will report the system disk size as 20 gigabytes. This is repeated for all images hosted by the VM environment on a given host system. As with memory, the VMM prevents a VM from directly accessing the host system's hard disk, and from accessing hard disk space allocated to any other VM.

Most VM environments allow access to the host system's mass storage as an option, but this is easily disabled using a configuration setting. Similarly, VMs may be configured to map another VM image to allow accessibility between VMs, but this is also disabled using a configuration setting. In the NetTop environment, these settings are disabled by default.

System Device Isolation: As noted previously, each VM believes it has direct access to the host system's devices. In reality, the VMM controls all access to devices. When a VM uses a device, it believes it has sole control over it. This means that for storage and data devices, such as USB devices, CD-ROM drives, communication ports, and so on, any VM may generally read or write to the device without hindrance. This creates a problem if, for example, a CD containing Top Secret data is placed in a CD-ROM drive where a VM with lower clearance can access it.

An important concept is that of device locking. A VM may request exclusive access to a device (e.g. a CD-ROM), causing the VMM to lock the device. This means that the VM using the device has exclusive access until it purposely releases the device. While a device is locked, no

other VM may access it, and this is enforced by the VMM. By this mechanism, isolation can be enforced for storage and data devices. In the NetTop environment, access to floppy drives, CD-ROM drives, sound cards, etc. can only be accessed when locked by the VM.

For network interfaces, each VM may each be assigned an individual physical network interface card (NIC) present in the host system that connects to the appropriate network, see Figure 2. Standard network protections at the host operating system level (firewalls, access control lists, authentication, etc.) then enforce the Guest/Guest isolation by not permitting unauthorised VMs to connect and request addresses. Similarly when a VM is using a particular NIC in NetTop, it is locked by the VMM so that no other VM may use it.

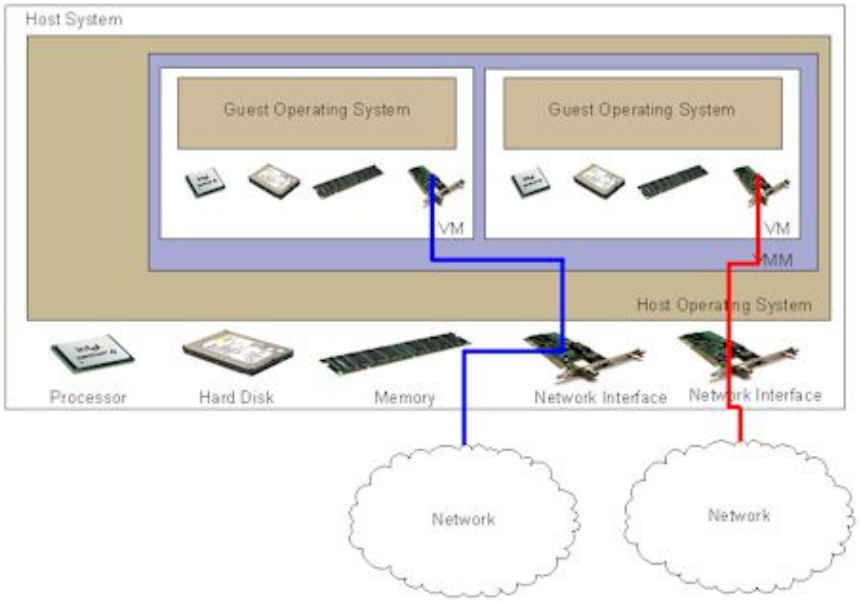

Figure 2: Multiple Networks with Multiple NICs

A single NIC can also be shared by several VMs to connect to multiple networks of varying level, see Figure 3. Because network traffic is routed from all active VMs to a single NIC, the use of virtual private networks (VPN) is required. The base assumption is that any operations within the VMM are controlled and secure. VPNs ensure that all data sent over the NIC is encrypted before it leaves the boundary of the VMM, and so security is established before leaving the secured area.

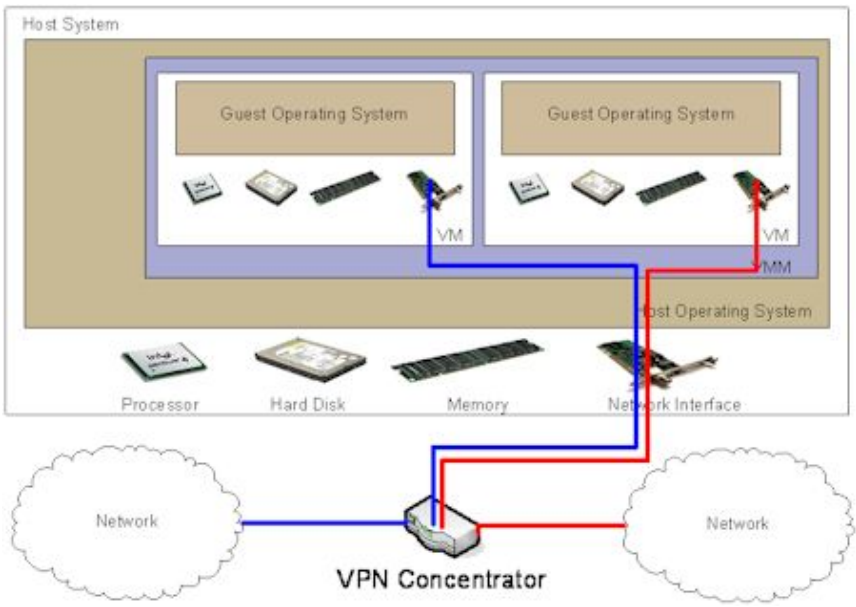

Figure 3: Multiple Network Connections over one NIC

Data Isolation: In the standard VMM, users are permitted to copy and paste data between VMs. The NetTop environment enforces separation of data between VMs by disabling the copy and paste functionality. Provision exists within NetTop to allow transfer from lower level VMs to higher level VMs, but not vice-versa. This is accomplished via a "data pump" incorporated into the VMM. The data pump works in conjunction with the VMM to ensure that data may be passed from low to high, but not vice-versa.

4.1.1.2 Host/Guest Isolation

The Host/Guest isolation arises from the VMM and the SELinux kernel extensions. The SELinux kernel extensions strictly enforce the security policy. This prevents users from opening terminal windows, and only allows them to interact with the NetTop desktop application that contains the VM environment.

The VMM prevents direct access to the host system by VMs running in the VM environment. As noted in Section 4.1.1.1 (Guest/Guest Isolation), the VMM presents the VM with a virtual copy of the host system, but does not allow it direct access to any of the host system's resources. Most of the mechanisms given there for preventing Guest/Guest interaction also effect Host/Guest isolation.

4.1.2 Limitations

VM images themselves can be a point of attack in the system. An attacker with access to the VM image can extract data from it. In most VM environments, images are stored on the host system's local hard drive. While the host operating system is a secure one, it is also possible that an attacker may simply remove the hard disk physically to be worked on in privacy elsewhere.

If a user's image resides on a specific host system, there is obvious difficulty in using a different system. In an environment such as a naval ship, there might be several locations where a given user might want to start a VM to access a network. Storing images on the host system means that users now need to have several copies of their image for all the locations they may wish to use. If a user changes settings or stores a file on one image, it will not exist in any other copy of their image, causing repetition of effort to keep all images current. The situation may arise where sensitive information is stored on an insecure medium for ease of copying between images.

4.2 Portable Guest Images

It is the position of this paper that existing VM/MLS solution approaches (specifically, the approach used by NetTop) can be extended to improve the flexibility and security of the chosen model. The main focus of the enhancements detailed in this section is to propose a thin-client architecture for loading and running virtual machines.

The traditional definition of a thin client system is a host system that performs all the application processing, but stores nothing locally, and often has no hard drive. A thin-client system downloads needed applications from the server, runs it locally and returns any updated data to the server. The next time the program is to be run, it must be downloaded again. The advantages to using a thin-client approach to loading guest operating systems onto a secure workstation are described in this section.

4.2.1 Secure Hosting Platform

If VM images are obtained form some external source (e.g. a portable device or a file server location), then the host system needs only to have the host operating system and VM environment present on it. Attackers who gain access to the host system will find they have accessed a security-hardened system that contains no actual (i.e. sensitive) data. Assuming access to these images is tightly controlled, for example, not left unattended at the host system, gaining access to the host system will not result in information compromise or disclosure.

Portable storage devices for VM images offer several advantages over storing the images on the host system. An image, complete with the user's settings, data, and preferences can be carried along to any host system connected to the network(s) they need to connect to. Users can protect their image and data without having to leave it on a host system that may be vulnerable to theft or attack. It also allows the host system to be configured as a true thin client.

It is significant to note however, that such a secure platform will be used for specific tasks in hosting the VM/MLS solution and, therefore, will have specific hardware requirements to achieve these tasks. For the proposed solution, it is expected that this secure platform will provide the following functionality:

1. Host a hardened, secure operating system[6]; (e.g. SELinux)

2. Host application software for hosting virtual machine environments; (e.g. VMware)

3. Allow the loading of memory-based virtual machines;

4. Allow the interaction of USB-based portable storage devices (e.g. thumb drives); and

5. Download VM images from the local network.

As such, it is anticipated that the following specific hardware requirements will be needed when defining the secure hosting platform:

- A large amount of system RAM;
- A high speed network and interface 1 GHz network adapters, cabling and infrastructure; and
- One additional USB port, in addition to the ports that are already using for the existing infrastructure, to support the use of a thumb drive.

It is important to consider that the virtual machine solution itself will impose specific hardware/software requirements for the hosting platform. For example, to support the use of VMware, the hosting platform must have an X Windows compatible video card and the software modules to support the parallel port must be available and enabled.

It is also noteworthy that the use of the local hard disk can be completely eliminated from the configuration of the hosting platform. Using a software distribution method known as LiveLinux , it becomes possible to launch the operating system and hosting software from a bootable CD. In this way the hosting platform contains no locally stored information. This approach it not recommended for the VM/MLS solution at this time. It suffices to establish an environment where there is no sensitive information stored on local systems and sensitive information that is loaded to the secure platform is zeroized on system shutdown. This leaves a choice in how to configure the host system. Two potential approaches are provided below:

1. One viable configuration alternative is a true thin client configuration. The host system is a complete computer, but has no hard drive. It is also equipped with a very large amount of RAM, multiple gigabytes worth. When the host system is started, it retrieves an image of the host operating system and VM software from a server, and executes it exclusively from memory. As VMs are loaded, they are also loaded into and executed from memory entirely.

2. A second configuration is a hybrid thin client configuration. The host system contains a hard drive, which contains only the host operating system and VM software. When started, the host operating system and VM software are automatically started as per any normal computer. As VMs are loaded, they are loaded into and executed from memory entirely. As with the true thin client configuration, a very large amount of RAM is needed, but it is reduced somewhat since memory for the host operating system and VM software is not needed.

The main point to be made here is that the amount of system memory that is available will dictate the number of simultaneous virtual machines that can be launched simultaneously. In addition to the operating system and the virtual machine hosting software, the following information will use system memory:

Storing local VM images: When a classified environment is to be launched (e.g. TS workstation) the VM image for access to this environment must be downloaded to the secure workstation.

Loading VM images: System memory will be used as part of the act of activating the VM environment. This is particularly significant since there will be no swap space available to the guest operating system since disk access is forbidden.

In summary, the secure platform must be able to hold a substantial amount of memory. The actual amount that is needed will have to be determined through experimentation since it is heavily dependent on:

1. The size of the individual VM images;

2. The number of simultaneous images that are needed; and

3. The level of performance that is needed from the VM images.

4.2.2 Memory Based Images

Added security can be obtained through the use of random access memory (RAM) drives. If the host system is equipped with enough memory to create a mass storage device in memory, VM images can be loaded onto it and executed directly from memory. The approach for this process is to create a RamDisk, essentially an allocation of system memory that has been set aside to act as a representation of a traditional hard disk based file system. RamDisks are primarily known for the following two properties:

1. They are faster (on the order of 1000 times faster) than their hard disk equivalents; and

2. The storage is volatile (if power to the system is terminated, all information stored in this drive is lost).

What is proposed is that virtual machine images should be loaded into a RamDisk file system and processed at this location. With VMs executing in memory, any loss of power to the host system will cause all trace of a VM to disappear completely. In the case of an attack in progress, simply turning off the host system will immediately deny an attacker access to the VM and any sensitive information contained within the guest operating system hosted in that image. This maps directly to the reuse requirement in an MLS solution. After a user session is completed, the secure workstation can be restarted and all sensitive information, which has been loading through the VM images, will be irrevocably lost. As indicated in the following diagram, the portions of the solution that are in memory and therefore active are obtained from solution components that are loaded from the hard disk and RamDisk file systems.

It is worth re-iterating that this proposed model would have VM images being transferred to the RamDisk from their originating source location (e.g. a central distribution point or portable drive) and then loaded into the virtual machine environment form this location. There are no “writes” back to the original location so any changes to the system configuration would not be saved. However, as will be documented in section *4.2.3.2:Personalized Settings*, the user’s custom settings and data may be written back to the user’s portable drive allowing the retention of the user’s working environment.

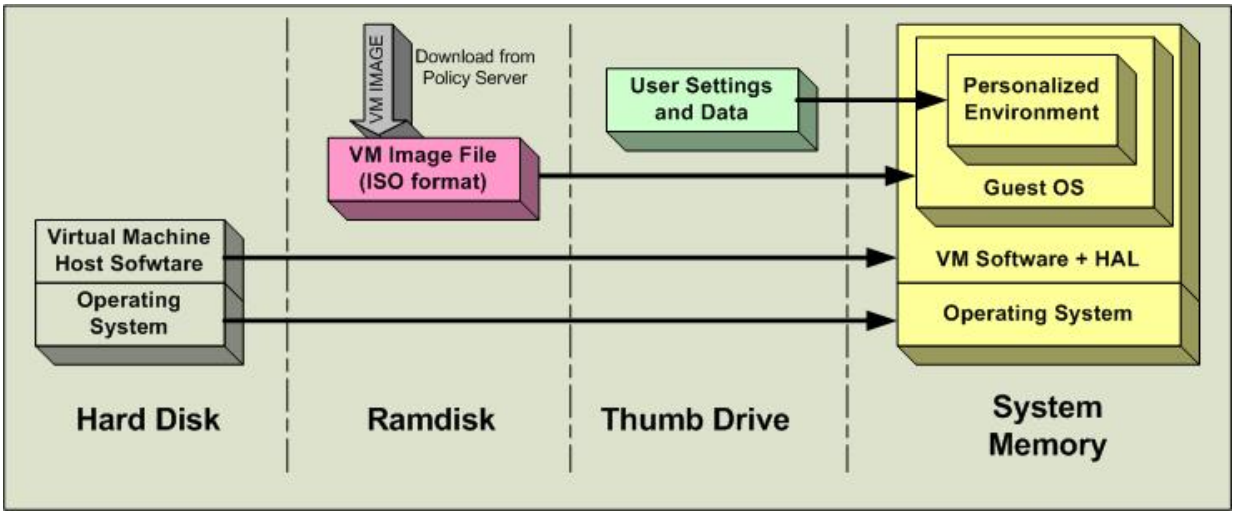

Figure 4:Hard Disk and Ram Disk Information Sources

When the workstation is powered down or reset, only the hard disk hosted information will remain available for access. Similarly, it is a simpler and surer process to wipe memory locations of sensitive material rather than disk-based locations. Traditionally, wiping hard disks has been a technologically challenging process when many protections are in place to protect the integrity of hard disk stored data, for example, journaling file systems. To truly wipe hard disks requires a low level format or GoC approved software designed for this purpose. Both methods take a substantial amount of time to ensure that the information has truly been removed from the disk.

By contrast, an equivalent zeroization function that clears the memory location previously used to host a VM image can be used with a high degree of confidence that the image cannot be recovered. As previously stated, if zeroization must be guaranteed, the system can be powered down, eliminating all data in volatile memory.

4.2.3 Thumbdrives

USB flash drives (thumb drives) have become very common recently. With advances in memory technology, the density of a very small (1cm x 3cm) USB flash drive can be 4GB today, and it is fully expected that larger storage devices will become available over time. Used as a portable storage device for VM images, several thumb drives can be unobtrusively held in a pocket, or on a lanyard around a user's neck. Within the context of the proposed VM/MLS solution, it is anticipated that thumb drives could be leveraged to provide the capabilities and services documented in the following sections.

4.2.3.1 Token Based Authentication

As a USB device, thumb drives provide certain identification information that is common to all devices that use the USB for connectivity. This information is obtained through the standardized USB communication protocol that provides the design and manufacturing requirements for any USB enabled device. Among the information that is expected as part of the connectivity negotiation with USB devices are the following pieces of information:

Vendor ID: a unique identifier that is associated with a specific manufacturer

Product ID: a unique identifier, within the context of a specific vendor, that is associated with a specific product design and/or revision.

With these two pieces of information it is possible to uniquely identify a specific USB device model. This links in closely with access control methodologies that can allow or deny specific types of hardware from connecting to host systems. For example, the Linux hotplug technology is responsible for detecting and connecting USB devices to the host system. The hotplug solution includes the ability to define whitelists and blacklists to allow or deny, respectively, specific hardware from connecting to the host system. Within the context of a USB-based token identification solution, it is possible to restrict the connection of USB thumb drives to a specific product or set of products. This would require an attacker to obtain one of the approved thumb drive models to launch an attack against the VM/MLS solution.

However, as there is no concept of a serial number in the USB device ROM architecture, it is not possible to differentiate between two identical instances of the same USB product. Again using the Linux hotplug example, it is not possible to allow or deny specific thumb drive from accessing the solution. It falls, therefore, to security providers to implement a method for uniquely identifying the token and the user in a 2-factor authentication scheme. There are many examples of this type of solution that can be leveraged to combine a user's credentials with credentials stored on the thumb drive. One such example is based on electronic card access solutions.

1. A unique identifier for each user is randomly generated and stored at the authentication server.

2. The user's credentials (username/password) are used to hash the unique identifier.

3. This hash value is stored on the token.

4. When a user authenticates, the username/password and hash are sent to the authentication server.

5. The authentication server will allow access if:

    a. The username/password are correct; and

    b. The calculated hash, using the locally stored unique identifier matches the submitted hash value.

More complex authentication schemes can build upon this solution; for example, a challenge/response solution can be put in place to eliminate the need to submit user credentials.

In addition, all USB devices have a small area where unique string data can be stored, which may be useful for key material, possibly removing the need for a separate token for image decryption, or acting as an identifier for a public-key system to retrieve the private key needed for decryption from a secured repository. Access to these string data memory areas requires the need for a specialized reader; however, access to this memory area is not restricted per se. In summary,

there are existing hardware and software based mechanisms that can be employed to tightly bind the user to a specific thumb drive and achieve a 2-factor authentication solution.

4.2.3.2 Personalized Settings

One of the prime motivators for having portable images is the capability for users to maintain their personalised settings within their own guest operating system environment. From simple things such as desktop arrangement, to application settings and data files stored in the VM, the biggest advantage of using portable images is that the user has this custom environment no matter where they start their VM. The question then arises, is it necessary that the entire VM image be made portable?

An alternate possibility is to create a standard VM image for all users, configured so that the user's personalised settings are the only data elements transferred to portable media between sessions. Among the items to be saved to portable media would be information such as:

- User interface layout (desktop icons, menus, etc.);
- Application configuration data specific to the user;
- OS configuration data specific to the user;
- Networking configuration data; and
- User files.

The VM image may reside on the host system's hard disk or may be requisitioned from a central server at start-up (more detail about this in 4.2.3.3). The image would contain the operating system and applications only. VM operation would follow these steps:

1. User starts the VM using the standard image.
2. User's personalised settings are copied or made available from portable media to the running VM.
3. User conducts business operations using the guest operating system hosted within the VM image.
4. User calls for the VM to shut down.
5. User's personalised settings are copied or saved back to portable media.
6. VM shuts down without saving the user's settings as part of the VM image.

Depending on the operating system that the VM contains, exporting of the user's personalized settings will require varying levels of effort. In a Unix-type environment, the user's "/home" directory could reside on portable media, and the operating system would mount the portable media directly, linking it to the operating system so that all accesses for user settings and files directly access the portable media. In a Windows environment, most user files and some settings are stored in a unique directory for each user, so a similar technique may be used. The task becomes somewhat more complex because of the existence of the Windows registry, though, used

to store many application settings in a centralised location (though there are ways to deal with this).

Windows Vista, the next version of Windows forthcoming, uses the idea of virtualised registry. This presents each user with their own registry, which may be used to advantage here for storing a user's settings on portable media without needing special handling to make it work. Further research into this would be required as the details of how this technology operates are not fully known yet.

Given the sensitive nature of information that may be stored on a user's portable storage device, it is essential to address the question of how best to protect this information. Encrypting the user's personalised settings on the portable media is the obvious answer, but it does introduce additional solution complexity. If the information is encrypted on the media, it cannot be directly mounted or accessed by the operating system. Some possibilities for this include:

RamDisks: The personalised settings can be decrypted to a RamDisk, which is then mounted or accessed normally by the operating system, instead of mounting the portable media directly. At shutdown, the contents of the RamDisk are encrypted and copied back to the portable media

Encrypted Disk: Third-party software such as TrueCrypt allow encrypted disks. On-the-fly encryption and decryption of data to and from the encrypted disk is performed as part of the normal operation of the system. User personalised settings could be stored on portable media with this software installed as part of the VM image, allowing the encrypted disk to be mounted and accessed normally, with encryption happening without user involvement.

The advantage of encrypting only the user's personalised settings is a huge reduction in the amount of data requiring encryption and decryption. This leads to a significant performance gain during VM start-up and shutdown, possibly avoiding the need to use hardware cryptographic modules to increase performance.

Care must be taken with the contents of the VM image, though. If the images reside on the host system, they should not include any information that an attacker might extract to aid their attacks against the network to which the VM connects. An attacker must not be able to simply start the VM and use it as a platform to launch an attack from. Encryption of the images may still be desirable, if they are not served to the thin client from a protected server. Of course, encrypting the images may lead back to the necessity of hardware crypto for performance.

Further research is warranted to determine the performance of completely portable VM images versus portable user settings only with VM images resident on the host system or images served remotely.

4.2.3.3 Remote Storage of Images

As will be further discussed in section *4.3:Policy-Based Access Control,* the intent is to have users download VM images from a centrally managed and access-controlled location to ensure currency and integrity of the environments hosted within these images. However, there may be cases where there is a requirement to use the applications hosted within a specific shipboard VM

image outside of the ship's operations zone. To support such a scenario, it is proposed that the secure workstation provide the user the ability to download the image to the thumb drive.

There is a set of requirements that must be in place to support this capability:

- The thumb drive in question must have sufficient capacity to hold the VM image;
- There must be an equivalent secure workstation available at the target destination where the VM image will be used; and
- As with the personalized information on the thumb drive, the VM image should be protected via encryption.

While not defined at this point, it is proposed that metadata about the VM image be sent along with the image and stored on the device. This metadata can be used to:

1. Ensure the integrity of the image;
2. Define characteristics of the VM image such as classification level; and
3. Define additional conditions under which this image can be used (e.g. specific workstation ID).

The VM image metadata can be used in conjunction with a policy server (see section *4.3 Policy-Based Access Control* At this stage in the solution architecture, the ability for downloadable images to thumb drives is seen as an extension to the primary centrally managed VM image deployment approach. It suffices to say that the ability to download image to thumb drives is possible, given the current state of the technology, and that these images can be protected in a manner which complies with certification and accreditation requirements.

4.2.4 Summary of Approach

To summarize, the proposed architecture includes the following components.

Secure workstations that store no sensitive information but allow access to secure environments through the hosting of guest operating systems. The HP NetTop approach using SELinux and VMware is currently the expected system and application software configuration that will meet the VM/MLS solution requirements.

Centrally deployed VM images that are downloaded on an as needed basis and loaded onto the secure workstations. Note that these VM images contain a guest operating system for a specific classification level that includes the applications, network connectivity configuration, and access to shared services to connect to a suitably protected processing environment.

Personalized Portable Storage Device thumbdrives which allow the user to:

1. Access the infrastructure via strong authentication;
2. Experience a customized environment by loading user specific data and settings from the thumb drive; and

3. Potentially transport the environment to another processing facility in a secure manner

These elements, used in conjunction with a policy-based solution architecture server (see section *4.3 Policy-Based Access Control*) provide what is perceived to be the necessary components to achieve a VM/MLS solution architecture.

4.2.4.1 Cryptographically Protected Information

As a brief side note, a discussion is provided relating to the mechanisms that should be put in place to protect data at all stages of processing. For the purpose of this discussion there are 3 types of information that must be examined.

System and application data: This software, which resides on the secure workstation, does not in and of itself store sensitive information. It provides the environment in which sensitive guest operating systems can be hosted. As host to guest isolation has been provided through the approach taken by NetTop, protections are in place that prevent the leakage of sensitive data to the workstation.

Guest operating systems: The VM images themselves, which store the guest operating systems, exist at the deployment server for download, in a RamDisk on the secure workstation and, potentially, on the user's thumb drive for transport to another processing facility. A risk assessment should be made to determine the level of protection needed to safeguard these images from loss of confidentiality and integrity. However, is should be noted that the inherent solution architecture provides a degree of protection that may mitigate the need for encryption of these images. Most significantly:

1. The VM images can be hosted on a deployment server that is placed in a protected environment, away from access by casual users.

2. On the workstation, the VM images exist in volatile memory which is erased when a user's session terminates

These images provide the means to access sensitive information, but system, application and data access can still be controlled through string authentication and access control measures. No sensitive information, other than system and network configuration, should exist in the VM images.

If encryption is deemed necessary to protect some VM images, performance of the cryptographic transformations can be increased to meet response time requirements through the use of hardware cryptographic modules. NetTop supports encryption of VM images, and will decrypt an image before loading it for added security. However, it is unknown at this time if hardware cryptographic modules can be used in conjunction with NetTop's built-in encryption. Another option is encryption of images outside of NetTop. With the addition of a dedicated module (that would be evaluated by SELinux authorities) to the SELinux base, a hardware cryptographic module could be used to decrypt images very quickly to reduce the time needed to load a VM from its image.

Personalized setting and user data: This information is expected to reside on a user's personalized thumb drive and it is anticipated that encryption will have a play here. There are many solutions that provide volume-based encryption that can be leveraged for protecting user information. More appropriately, solutions such as TrueCrypt allow for device based encryption and is specifically geared towards protection memory based devices such as thumb drives. Again, a risk assessment must be undertaken to determine the sensitivity level of the information that may be stored on the thumb drive and the necessary protections that must be put in pace to safeguard this information. It suffices to say at this time that thumb drive encryption is a necessary component to the VM/MLS solution and must be included in the architectural design.

At this time, it is not clear if the VM images themselves will contain sensitive information. A string case could be made that, in spite of the lack of any inherent sensitive data in the VM image, the configuration information and security policy information should be treated as sensitive information. In this case, the transmission between the policy distribution point and the workstation should be protected against theft or modification. Deploying a separate channel security solution is not needed since the VPN concentrator solution can provide the channel security for this communication. However, the selection of VPN solution will then need to support the Host based operating system, that is, SELinux.

4.3 Policy-Based Access Control

This section provides a high level description of how granular access to information resources can be achieved by combining the previously documented technology solution with a policy based network solution. Specifically, this section proposes to enhance the isolation, compartmentalization and access control of a multi-level security solution through the enforcement of a defined policy. Building on existing work sponsored by DRDC [Bacic2003] policy, the position of this paper is to define policy as business rules codified into a logic that permits the mediation of requests for access to valuable resources or assets. Specific conditions, presented to a policy engine results in predefined actions being invoked that uphold the security principles that have been coded into the rules enforced by the policy server.

The previous section has identified an architecture by which virtual machine images are loaded onto a secured workstation on an as needed basis. These images are loaded into volatile memory, ensuring that no leakage of information occurs after the VM image, and it associated processing space, has been released. The question remains as to where these images are acquired from and who can gain access to and load these images. This paper proposes that these functions are controlled via a policy based decision process. A high-level view of a policy solution requires the definition of the following components.

4.3.1 Policy Elements

This section details the characteristic elements, that is, active and passive entities and attributes of these entities, which will be present in a policy-based virtual machine deployment strategy for the VM/MLS solution.

4.3.1.1 Primary Entities

In order to define what transpires within a system we must utilize a generic term to define the objects upon which a security system will operate. We define an entity as the unified security object against which all security related functions are performed. The simplest policy model would include a subject and an object. In the VM/MLS solution, the subject would be a user and the object would be the VM image with the default operation to access the image. In such a model, we would expect access to be granted if and only if the user's clearance level meets or exceeds the classification level associated with the VM image.

Both entities, the active entity of the user and the passive entity of the VM image, must be validated prior to being part of the policy decision. In the case of the user (the subject) an authentication process must take place to verify the identity of the user. This process can be enhanced through a two-factor authentication process using the thumb drive as a hardware token. Similarly, the VM image can be verified through a digital signature check on the file to ensure that it has not been replaced or altered.

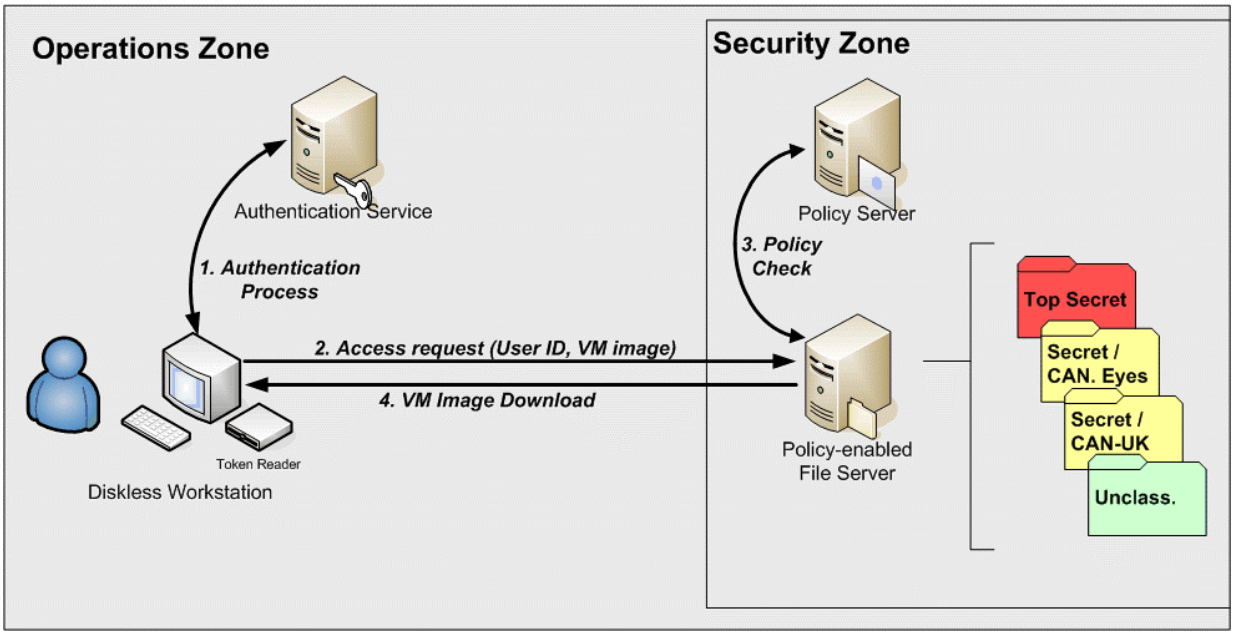

Figure 5:A simple policy-based VM access model

The above diagram illustrates the main processes in a simple policy based decision to allow a user to acquire and load a VM image on the local workstation. The user presents credentials to the workstation, which are then validated by an authentication service. The user provides the authenticated user identity with an indication of the desired VM image to the policy enforcement point (PEP). The PEP, in this case, is a file server that makes a call to the policy engine to validate the request to ensure it is in compliance with the stated policy. If the policy allows the user to download the chosen image, an "allow" response is returned to the file server which then allows the image to be provided to and loaded on the diskless workstation.

4.3.1.2 Location Based Policy Decisions

With the simple policy model in place, extensions to this model are proposed to add to the flexibility of the solution. It is expected that the target environment, namely on board Canadian Navy vessels, there will be segregated areas which allow processing of information at a higher level of sensitivity. For example, workstations on the bridge may allow the processing of Top Secret information whereas crew quarters may only allow unclassified data processing. Workstations can therefore be added to the model as a separate entity. A workstation takes on the clearance level for the processing area in which it resides. If the workstation can be uniquely identified and it's location within the ship (i.e. the processing zone in which it resides) is defined on the policy, this information can be added to the access control decision. In this case, a user with access to Top Secret information will not be able to access the VM image outside a zone that has not been cleared to process information at this level of sensitivity.

In the following diagram, three zones are identified and a policy is in place that allows only certain VM images to be downloaded to and loaded on workstations in specific zones. The Operations zone only allows unclassified VM images. The Security Zone allows Unclassified and Secret VM images. The High Security Zone allows only Secret and Top Secret VM images, but not unclassified images.

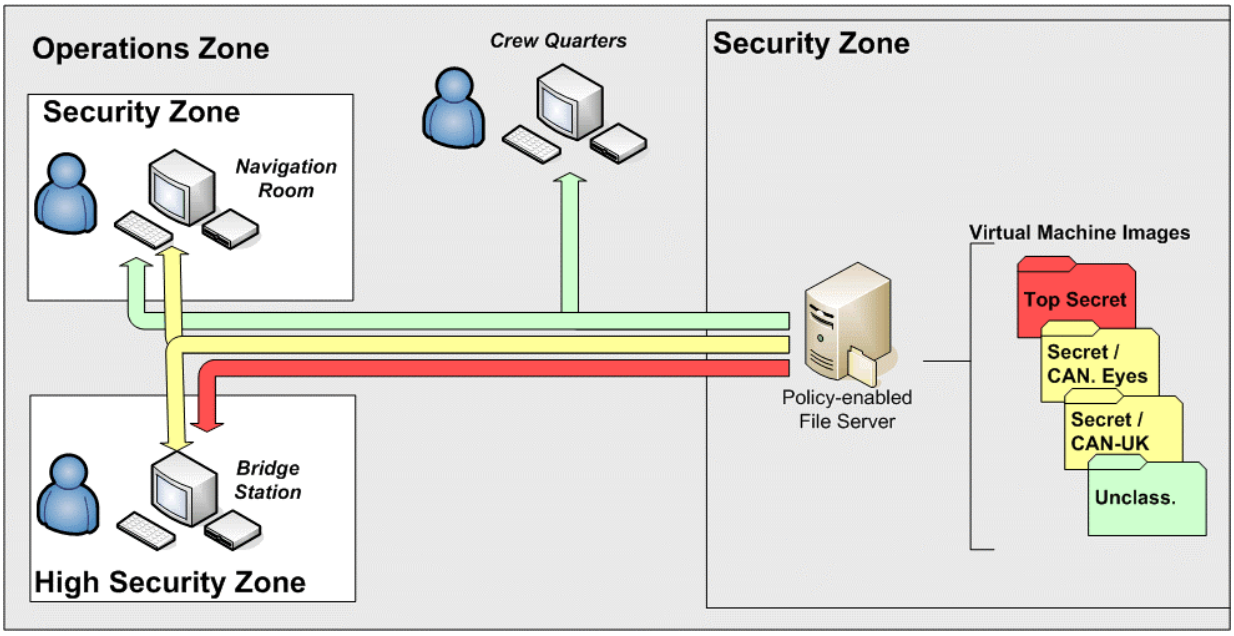

Figure 6:Location Aware Policy Decisions

The workstation identity can be derived through unique attributes of the system (e.g. MAC address, processor identifier). The presence of an unrecognized system in the environment or the movement of an existing system to a new zone in the target environment, based on the submission of a policy request form a previously unknown host, can be flagged as a security event. The policy administrator can handle this scenario accordingly:

1. for a valid addition/movement of a system to a security zone, this policy change can be accepted; or

2. for an invalid addition/movement of a system to a security zone security procedures should be followed to respond to a potential security breach.

This policy model can also potentially leverage integrity tools. The integrity of the system, as determined by 3rd party integrity tools such as Tripwire, can be included as an attribute of the workstation entity. The policy engine can dictate specific actions in response to integrity violation as determined by these tools.

4.3.1.3 Additional Policy Elements

The following elements are proposed to illustrate how a customized policy can add to the flexibility and security of the MLS solution. These elements are attributes associated with the subject (e.g. a user's role) or the access request (e.g. the time of day) and can be sent as part of the policy operation check or determined by the policy engine as part of the processing of that request. In some cases, external information apart from the policy engine and the requestor will have to be obtained (e.g. duty schedules). The policy elements presented below are provided to exemplify the capabilities of a policy-based network. A complete list of attributes should be determined as part of a requirements gathering process for the development of an appropriate security policy.

Duty Officer State: Certain information resources should only be available to the duty officer. Duty officer is a role that is assigned to a group of users at the same level of clearance. However, there can only be a single duty officer active at a given time. A user that is cleared to access sensitive information mat be denied access if they are not the active duty officer. This is a case where a simple allow/deny decision is not sufficient since the off-duty officer may have a valid reason to access these resources. A decision may be returned from the policy server to allow access with conditions. Such conditions may include a heightened level of auditing or the requirement for duty officer authorization to allow off-duty access to sensitive information.

Hours of Operation: Certain environments may only be accessible within specific hours. Again, the policy may dictate that the system can be accessed with conditions. Alternatively, specific user groups may be granted after hours use whereas a wider community can access the environment only during the approved window of availability.

Ship Location: Multiple policies can be written and approved for use under different conditions. For example, while a ship is in port, it is more vulnerable to physical and electronic attack. Therefore, a more restrictive policy should be engaged to strictly limit the users' accessibility to and capabilities on the shipboard information network. Once outside the zone of concern a less restrictive policy, geared to "at sea" operations, can be instituted allowing normal IT activities to resume.

Defence Condition: Under heightened defence conditions, a ship may opt to relax policy restrictions in the interest of responding to an imminent threat. Removing certain restrictions may allow for faster command decisions and execution of orders. In essence the risk of not responding in time to specific threats outweighs the threat posed by policy restrictions. For example, during combat operations, access to highly sensitive systems can be made available to all senior officers anywhere in the ship. This would allow the establishment of an alternate command center should the primary command center be damaged/destroyed.

4.3.2 Operations / Responses

Within the policy engine context, the subject/object interaction is defined in terms of operations. Given the user as subject and VM image as object, the following operations are seen as providing useful functions for a policy-based solution for accessing and loading VM images. In most cases, these operations will be used by functions within the loading application software on the secure workstation.

List: Obtain a list of all VM images available through the workstation interface

Access: For a given VM image, does the current authenticated user's rights allow them to access the VM image? This decision will be mitigated by other elements dictated above (workstation location, time of day, etc…)

Load: Obtain and load the selected VM image on the workstation

Save: Allow the VM image to be saved to a local resource

Personalize: Allow the use of personalized settings for the given user.

The policy should be able to dictate a specific response from a set of actions that matches the security policy and upholds the intent of the security practices for the information processing environment. Some examples of appropriate responses are given below.

Deny: This represents a denial of access, given the submitted request. Depending on the nature of the request, this response may generate a security event.

Allow / Allow with conditions: This represents a scenario where access is granted. Default security mechanisms can be set for this response or specific conditions can be attached to the acceptance of this request. For example, enhanced auditing may be specified for this session or secondary authorization may be needed to allow this session to continue.

Security Alert: A policy violation may trigger a security incident requiring immediate response (e.g. alerting security guards).

Honeypot: A policy violation may result in the presentation of a fictitious environment, designed to gather information about the person that has initiated the session for forensic activity and possible criminal prosecution.

The exact nature of and selection of the various responses should be determined through a requirements gathering process to fully define and express the appropriate policy.

4.3.3 Authentication / Identification

Mechanisms must be put in place that provide for unique identification of every individual as the access the secure workstation. The policy-based solution can leverage existing identification and authentication mechanisms, including the use of token-based authentication for highly sensitive

environments. As stated earlier, the thumb drive technology may be leveraged as a hardware token for some environments.

4.3.4 Accountability / Auditing

Mechanisms must be in place to uniquely identify and authenticate users and to dynamically track their actions within an application or system. Security controls, regardless of type, are not foolproof and to ensure a system of recourse after a security breach a non-circumventable, unalterable, continuous audit mechanism must be in place and operational.

It is expected that each isolated environment hosted via the VM/MLS solution will have its own internal requirements for and implementation of accountability. The policy based front end for accessing these isolated environments must work in conjunction with these practices to accurately track what images has been accessed by which user as well as the conditions under which this access took place (where, when, etc..). It should be possible for the policy to dictate to the hosted environment additional conditions of operation, such as the need for heightened accountability for a particular session.

While it is assumed that each isolated environment has the ability to generate, collect store, filter and protect information that related to security events that are triggered within that environment, the same degree of protection must be applied to audit events that are triggered through the policy-based interface to these virtualized environment.

4.3.5 Sample Usage

Given the previous description of policy definition and enforcement, the proposed policy-based VM/MLS solution can be described as in the following diagram. The process for gaining access to a virtual machine based segregated network connection is defined by following this process.

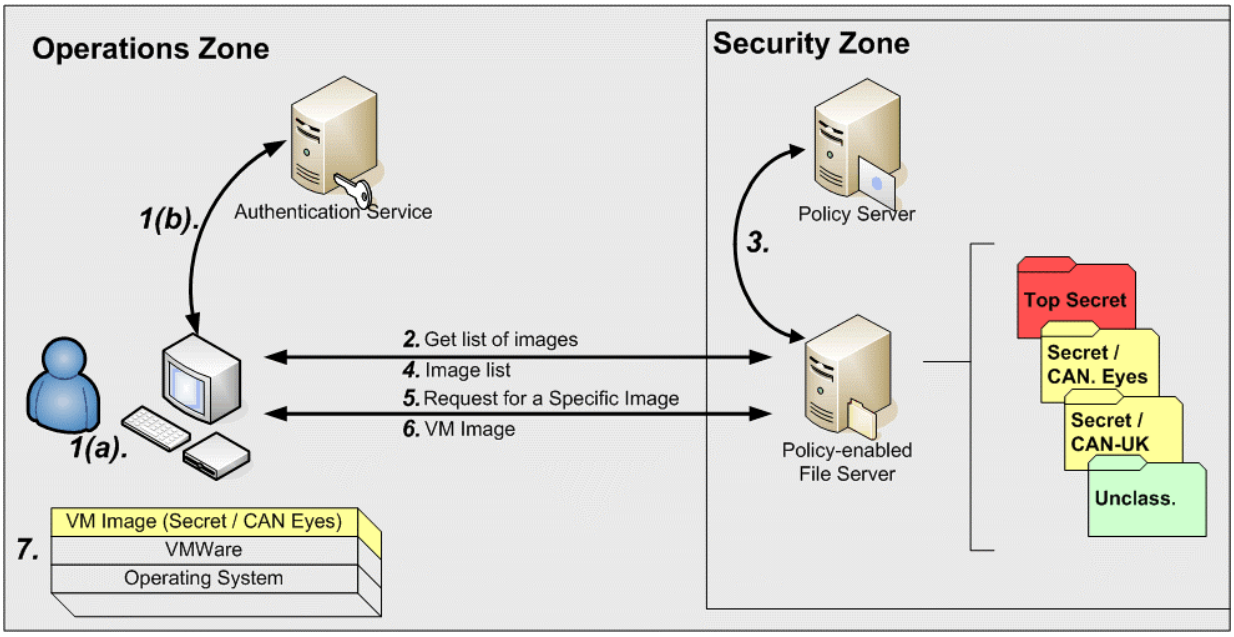

Figure 7:Sample User Session

Step 1. A user inserts their thumb drive to the diskless workstation (1a) and authenticates to the hosting infrastructure (1b). The thumb drive may be leveraged to act as a hardware token for 2-factor authentication.

Step 2. Upon successful presentation of the user's credentials, a request is sent to the VM image Policy Enforcement Point (PEP) to acquire a list of available VM images, based on known attributes of the user and the environment from where the request originated. The PEP provides a file sharing mechanism (e.g. Samba shares) to gate access to the VM images.

Step 3. The PEP sends a properly formatted policy request to the Policy Decision Point (PDP) for each available VM image using the credentials and attributes of the user, the hosting workstation, the image and the operational environment to determine if each VM image should (as per the policy be made available to the user.

Step 4: The list of available images is returned to the image loading interface and is presented to the user.

Step 5. The user selects a specific image from the list for loading onto the diskless workstation. This results in another policy request to determine of there are additional conditions that must be associated with the user's session.

Step 6. The PEP allows the VM image to be downloaded to the workstation

Step 7. This image is loaded into volatile memory for execution.

4.3.6 Advantages

In addition to the ability to enforce a flexible policy for access images, a policy-based provides the following advantages to a VM/MLS solution.

Easier Software Maintenance: As with any thin client solution, centrally managed operating system images are easier to manage and maintain. Updates to the image to include patches, configuration changes or new applications are immediately assessable to workstations. Operating system backups are not necessary so long as the images themselves are backed up.

Easier hardware Maintenance: As each workstation is identical, it is a simple matter to replace faulty hardware with standby systems. The same set of standby systems can be used in any of the security zones since no sensitive information is stored on the workstations.

Portable Images: If the workstation solution component is replicated across environments (e.g. different ships, in port, etc) images can be transported between these environments.

4.3.7 VM Image Deployment Approach

The mechanism by which the VM images are deployed (downloaded) to secure workstations does not need to be specified at this time. It suffices to say that there are many methods by which this can be achieved (e.g. file shares, LDAP). However, some generalized statements that can be made about the characteristics and capabilities that this mechanism should have.

1. The solution must be flexible enough to allow for close linkage with the policy server. A solution that works through references (e.g. LDAP) would have some configuration and revision control advantages over other solution options.

2. The solution must either provide or work with the solution mandated authentication mechanism.

3. Similarly, the solution must work with any data protection mechanisms that are in use within the problem space (e.g. encryption, signing).

4. The solution must be able to scale to allow for use under conditions of increased user demand without significant performance impact.

A detailed architectural design may impose other requirements on the selection of the deployment solution.

4.3.8 VM/MLS Without Policy-Based Access Control

Policy-based VM image distribution, using a PDP/PEP approach, provides a great deal of flexibility for gating access to isolated networking environments under a variety of conditions and constraints. However, it should be noted that including this capability as a solution requirement will incur additional time and effort to design, develop, test and certify the VM/MLS solution deployment. Therefore, a brief discussion of alternate methods for distributing images is provided.

One method would be to avoid the dynamic VM image distribution issue entirely by having images stored on the workstations themselves with each VM associated with a specific caveat separated network. Workstations would only hold those images which connect to networks that meet the environmental conditions in which the workstation resides. That is, a workstation in the SECRET zone would not store images that can connect to the TOP SECRET network. The operating system could be used to enforce access control, based on strongly authenticated user identity, to the VM images that are appropriate for the current user. It is noteworthy that such a solution will eliminate the need to download the images on demand and will, therefore, reduce network traffic and image start-up time.

While there are no specific security concerns associated with this approach, there is loss of functionality and flexibility in the following areas:

1. When a VM image must be updated, the effort to deploy images across workstations is significant. Instead of a centrally accessible copy of the VM image, the image is deployed to many systems, each of which must be updated when the master image is updated.

2. Access control decisions are made at each workstation across the network; therefore, the effort to distribute and maintain this policy would be significant. For example, if the operating system is used to enforce access control restrictions on VM images, each system must be independently configured and verified to ensure that the access control mechanism is correct.

3. If access control decisions are being made at many locations across the network, the effort to collect audit access information to perform forensic analysis is increased. It is notable, however, that a VM/MLS level audit will always be secondary to the auditing that occurs within the isolated network environments themselves. However, if a malicious attacker is able to circumvent the operating system for unauthorized access to VM images, it is likely that the host audit trail will be compromised as well, making it difficult to trace where the unauthorized access took place.

An alternate approach to VM image distribution would be to store the images at a central location and require users to download the images prior to use. This scenario mitigates many of the concerns of the workstation stored image approach. Images that are stored centrally (e.g. on a network accessible file share or FTP server) can be easily updated if the need arises and the new image would then be used in subsequent network access. Access control occurs at a central location which could be closely linked with the user management infrastructure to allow the access control policy to be changed quickly should the need arise. Similarly, VM/MLS level auditing can be tracked at a central location simplifying the process for any forensic effort.

It should be noted that this compromise approach has much in common with the policy-based networking approach but removes the requirement to develop the policy server solution components. Some of the features of the policy-based approach would not be available in a simplified central repository solution. If, however, the VM/MLS solution were designed in a modular fashion, it should be possible to migrate to the full policy-based solution without significant disruption to the initial, simplified central image repository approach. A modular design requires that interface points within the solution space be independent and well structured. This will allow the replacement of specific components with improved versions of these components that provide additional features. The image distribution / policy enforcement point

should be designed to be modular in this fashion so as to allow the VM/MLS solution to be able to deliver improved service over time. In actual fact, modularity in design should be a requirement for the entire VM/MLS solution, allowing any solution components to be easily replaced with improved versions that will allow the solution to meet future demands.

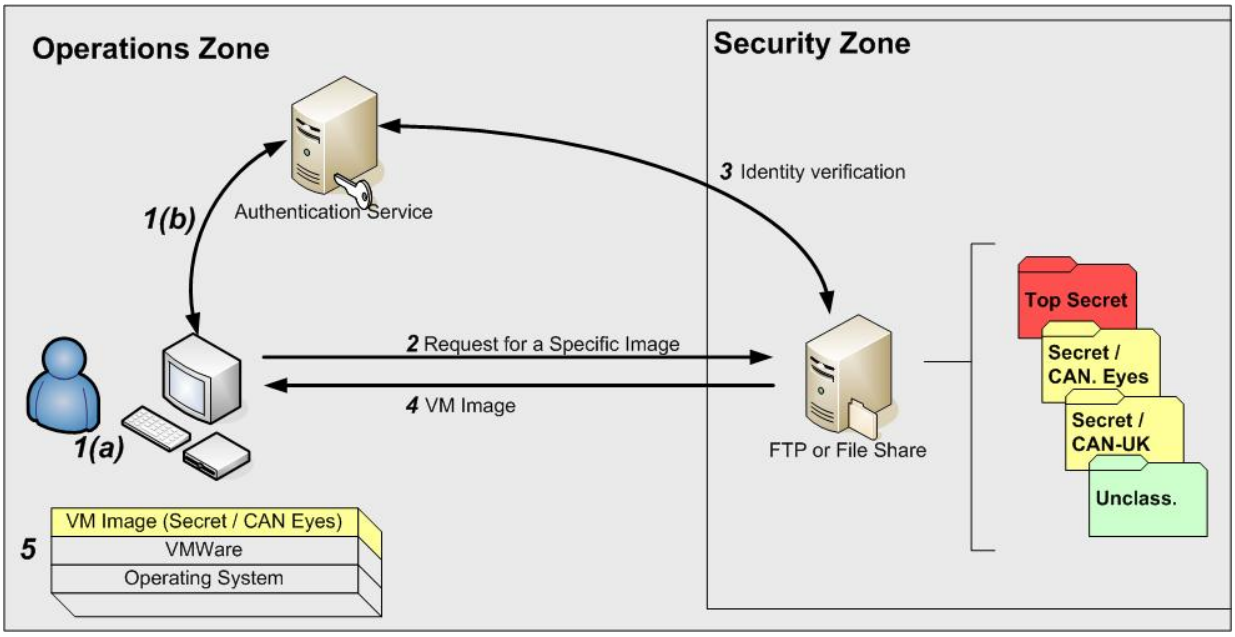

Figure 8:Alternate Distribution Method

The above diagram illustrates this alternate distribution method. The user workstation and authentication process remains unchanged from the policy-based model. To access a VM image a user would connect to a centralized distribution service. This service could then verify the identity of the user through existing identity verification routines that would in conjunction with the service (e.g. a pluggable authentication module). The service would then provide the user with the VM images for which the user has the authority to access. Note that this model does not include the ability to tailor the list of available images based on the location of the user.

A third VM image distribution method would be to store VM images locally, but arrange to have these images updated periodically via a batch process. In this way, master VM images would be distributed to the intended workstations as they are created. A master list of images could be maintained which would indicate to which workstations each image should be deployed. If an image is updated (e.g. virus updates, new software patches), this image could be transmitted to each destination workstation through a batch process. A summary of each approach is provided in the following table:

Table 1: Summary of Alternate VM Distribution Methods

	Method 1 VMs stored on the workstation	**Method 2** VMs downloaded from a central location	**Method 3** VMs pushed to workstations as needed
New VM image deployment	Manual process. Coordination of which images go to which workstation is a complex manual problem.	Simple to deploy new images. A new image is placed at the distribution point and all subsequent VM access requests will receive this image	Simple to deploy new images. A new image is pushed to the workstations after they have been created.
Auditing issues	VM image access would be logged locally.	Auditing can occur at a central point using authenticated credentials to track image use.	VM image access would be logged locally.
Policy enforcement	No specific policy. Access to the system would grant access to the VM images.	A simple policy can be put in place with a logical interface point for introducing a more complex policy server solution	A simple policy would push images to only those workstations that are authorized to use them. Access control list may be distributed as well.
Scalability	Challenging to scale as each workstation must be configured separately.	Additional distribution points can be added to serve specific communities.	Solution can scale to meet increased demand in terms of images and systems.
Network usage	No impact to the network	Substantial network use for the acquisition of VM images as they are downloaded prior to use.	Network impact can be scheduled to off-peak hours
Time to boot	Fast boot, no download needed	Performance impact since images must be downloaded prior to use. A fast network is needed to ensure the solution can operate within acceptable time limits	Only the local image is used so the time to boot does not incur a network download cost.
Flexibility	Inflexible. A workstation can only use those images that are stored locally.	Allows some flexibility in that a new user/workstation can quickly be enabled to acquire and use a VM image	Somewhat flexible, but not easy to respond to critical events.
Availability	VM image access is not reliant on the network or distribution points	Design must avoid the creation of a single point of failure at the distribution point.	Local images used, therefore images are always available.
Level of Effort	Simple to establish, more difficult to maintain.	More difficult to establish, simpler to maintain	Most challenging approach since the push-technology solution must be investigated.

5. Comprehensive View of Components

This section provides a comprehensive view of the solution components in the context of the previously identified MLS/MCS and project requirements details in section *2:Operational Requirements*. It is important to recognize that the proposed VM/MLS solution attempts to maintain existing isolation between security environments that is currently achieved through physical separation. The virtual machine technology is being used to simulate physical separation through memory, process and network isolation. Once this isolation is achieved, the existing MLS/MCS physical protections can be inherited by the VM/MLS solution. For example, the physical separation of environments dictates that there are separate user communities in each environment. This would continue to be true in the VM/MLS environment where each VM connect to separate infrastructure and, therefore, contains separate user communities.

The following table details each MLS/MCS requirement and responds to the implication and safeguard as it pertains to the VM/MLS solution.

Table 2: MLS/MCS requirements and the VM/MLS solution

Solution Requirement	***MLS/VM Implication***	***Solution Safeguard***
Mandatory Access Controls	As each VM image simulates the existing physical separation approach to MLS/MCS, the user experience within each VM will inherit the domain segregation (system high) properties of each isolated environment.	Inherited protection based on leveraging the existing configuration where there is physical separation between networks.
	Where VM images operate on shared resources (e.g. workstations, network), isolation between VM environments must be assured.	The HP NetTop solution provides assurances that data leakage between VM environments is not possible. This solution has been reviewed and approved by the NSA.
	Access to VM images must be controlled to ensure that users have access only to those images appropriate for their level of clearance.	A central VM image distribution center, in conjunction with a policy server that enforces access control linking classification levels of images with the clearance level of the user, can provide assurance that access to the VM images that connect to sensitive networks is restricted.

Solution Requirement	*MLS/VM Implication*	*Solution Safeguard*
Discretionary Access Controls	As each VM image simulates the existing physical separation approach to MLS/MCS, the user experience within each VM will inherit the discretionary access controls to information resources within each network	There is no role for DAC outside these simulated physically segregated networks. That is, the right to access, obtain and load VM images is defined by the central policy server and cannot be discretionally assigned except through the policy server.
Bitmap checks	Any devices that are accessed within this solution must be associated with a specific security label and access to those resources must be restricted based on the user's clearance level.	Within the network environment, there is inherited protection based on leveraging the existing configuration where there is physical separation between networks. Devices attached to the local workstation will be cleared to the level that is deemed appropriate for the processing environment. The policy server will not allow the loading of VM images in operational zones that do not have the necessary level of clearance. Devices attached to the workstation are connected to only one virtual machine at a time, thus, there is no sharing of resources between virtual machines.
Reuse Issues	All resources used must be properly sanitized before being re-used.	As all virtual machine operations occur within the RamDisk, when the VM is released the memory location is cleared and the RamDisk is destroyed.
Identity and Authentication	User identity must be verified prior to accessing resources.	The authentication server will determine the user's identity, based on 2-factor authentication using the USB thumb drive as a hardware token. This identity will be established and utilized within the policy server to properly grant access and rights to the VM images.
Auditing	Security events must be audited.	The authentication server, distribution mechanism, workstation and policy server can all be used in conjunction with a consolidated and trusted audit server to track security events and generate an audit log which can be used to trace user activities. Within each VM environment, the existing audit trail will be generated based on established practices and applied to the events that occurred within the segregated networks.

Solution Requirement	***MLS/VM Implication***	***Solution Safeguard***
Labelling	Security labels must be applied to all sensitive resources	The VM images will have electronic labels associated with them to ensure that the policy server makes the proper association between clearance level and VM image classification
Data Hiding	Users should not see information for which they have no access.	The policy server will provide a list of accessible VM images based on the user's access rights, their location, and other attributes. VM images that are not accessible will not be presented to the user.
Compartmentalization	It should be possible to segregate communities access based on a need-to-know	The policy server can grant access rights based on classification and/or caveats. The policy server is flexible to gate access based on many criteria.
Guards	The transition of information between networks must be done in a controlled and rigid manner to reduce information leakage.	The SELinux/VMware/NetTop solution provides for a *secure message pump* mechanism to transfer information between VM environments. All other methods for information sharing (e.g. cut and paste) can be disabled, preventing information leakage.
Single Workstation	It should be possible to access any image from a single workstation, based on the user's access rights and other session attributes.	A properly configured workstation can load and use any properly configured and deployed VM image.
Simultaneous Multiple Domain Access	It should be possible to have multiple domains active on a single workstation at one time.	A user can have as many VM environments active at one time as needed, subject to the hardware configuration (especially the amount of available system RAM) of the workstation.

6. End-to-End Security

This section summarizes the VM/MLS solution approach within the context of end-to-end security. Specifically, this section examines the state of sensitive data at all points along the solution architecture to identify where safeguards are in place to protect the data and ensure the integrity of the MLS architecture.

6.1 Zone Isolation

Currently, classified environments are kept separate using physical separation through replicated infrastructure (workstations, network devices, cabling). Within the proposed architecture there is an attempt to eliminate the need for replicated infrastructure. That is, multiple sensitive networks are made accessible through: a single workstation and a single network (cabling) infrastructure. In operational zones, infrastructure replication may still be used but only for the purpose of eliminating single points of failure and improve the robustness of the solution.

It is recognized that network accessible resources, such as file servers, databases, and application servers, for each classified network will still be physically separated. Therefore, the existing physical zone separation will still be used to isolate these networks for centrally stored network resources (e.g. in a server room). Network traffic will be routed from the shared network infrastructure to the appropriate zone for access to that secure network's information resources. The following diagram illustrates this concept

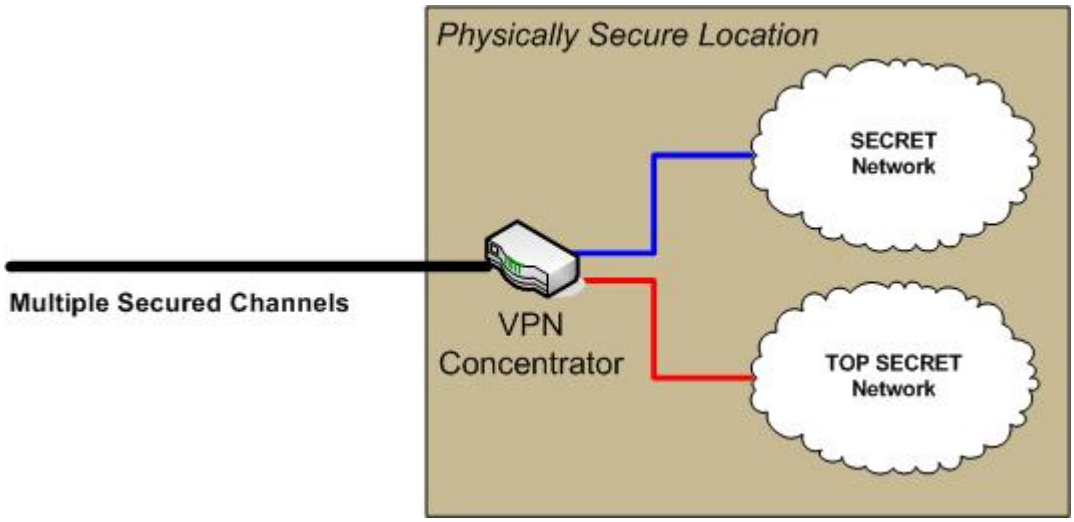

Figure 9:Zone Isolation through Physical Separation

The core solution component in this model is the VPN concentrator that established the private network tunnel to access the appropriate zone.

6.2 Network Isolation

At the network level, isolation of environments is achieved in part by encapsulation all network traffic between a specific VM image and its target sensitive network into a virtual private network. Each VM image will establish a separate VPN connection that is routed through a VPN concentrator and directed to the appropriate zone for processing. In this way, all VM images can share the same network resources (e.g. cabling, NIC cards) while the transmitted information remains isolated from the other networks.

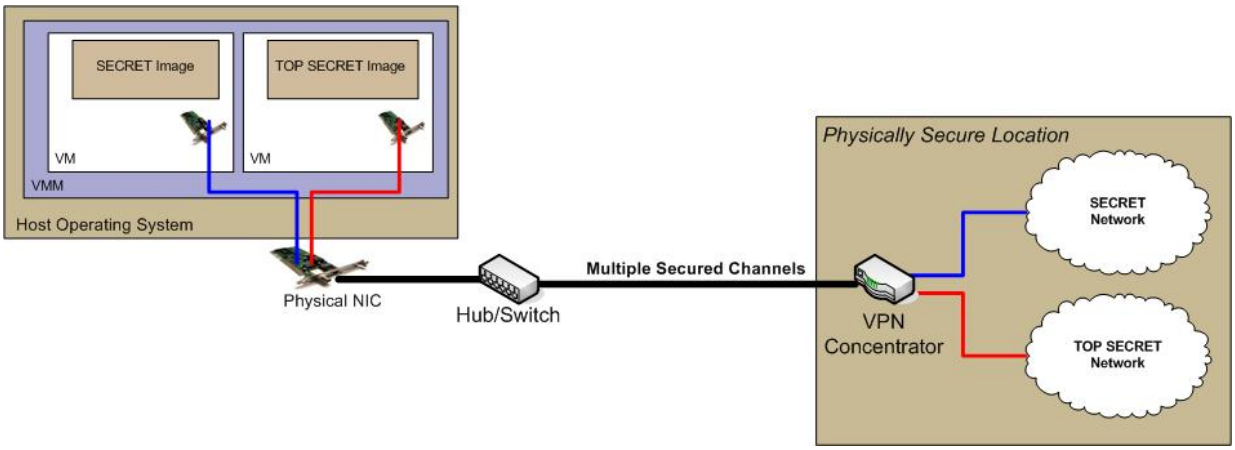

Figure 10:Network Isolation through VPN Channels

Without the VM image and appropriately configured VPN session, it is not possible access the target environment. Acquisition of the VM images is controlled through a policy server allowing only certain VM images to be used by select individuals under specified conditions.

6.3 Guest OS / Workstation Isolation

VM isolation is achieved through virtualization of the hardware on the host system for each hosted VM. If the policy server allows a VM image to be acquired and loaded, it exists with a self-contained virtual environment, where use of system memory, processing, device usage and messaging is tightly controlled by the Virtual Machine Monitor. With the VMM acting as arbitrator between VMs and the host system, the hosted guest operating systems cannot directly access the physical hardware for a given device, including the processor, the memory resources and interface-based devices. Devices must be bound to a specific virtual machine prior to being used and then release that device before it can be used again.

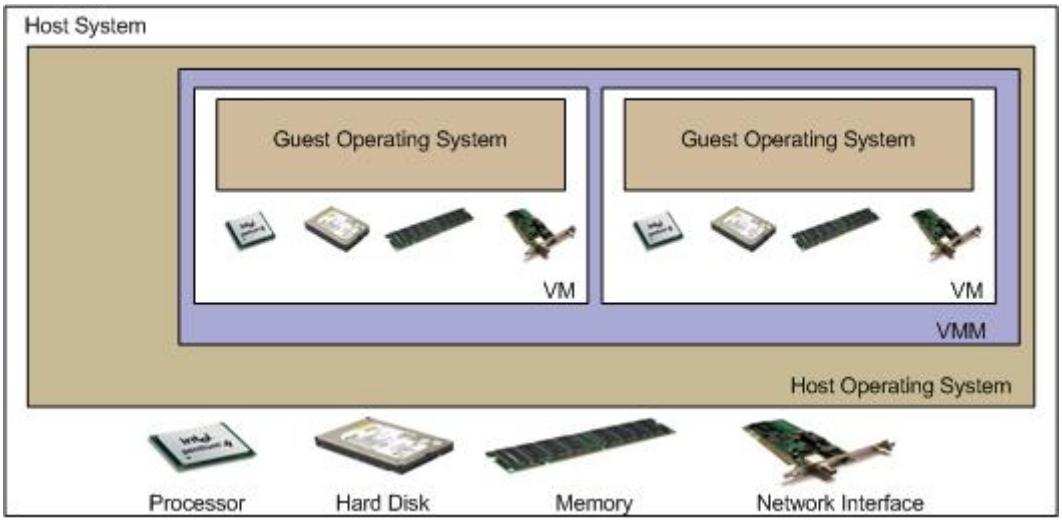

Figure 11:VM Isolation through Hardware Abstraction

All virtual machine functioning takes place within system memory, eliminating the leakage of information through swap files, temporary file, and core dumps. The act of requesting a VM image results in the creation of a RamDisk where the VM image is stored locally and loaded into the virtual machine environment. On shutdown, the RamDisk is cleared and destroyed, eliminating any residual information that may be present within that memory construct.

6.4 Control Over VM Images

Images are obtained from a central repository that only allows the distribution of a specific image if it meets specific criteria:

1. A match between the classification/caveat of the image and the user's clearance level

2. A match between the classification level of the image and the clearance level of the zone in which the image is to be used

3. Adherence to policy as it applies to environmental conditions such as duty shift schedule, defence condition, and ship status

4. Adherence to policy as it applies to security constraints such as sign-off authorization from a superior officer or the availability of enhanced auditing.

Additionally, the central housing of VM images with the policy server allows tight control over images from a maintenance perspective. Items such as patches, virus scanning signatures, new applications, new security tool can be quickly added to the guest operating system deployed on the image with the assurance this these new images will then be placed into production and made available to the user community. The policy server can match the revision level of the image to be loaded with the latest known image to ensure that outdated VM images are not used. Additionally, integrity checking can be incorporated into the loading function to ensure that the VM images have not been maliciously modified prior to use.

7. Technical Limitations / Operational Constraints

As this document provide a high-level technical overview of a proposed VM/MLS solution there still remains many areas of technical investigation to prove the technology is appropriate for the intended usage. It is expected that a detailed technical architecture will identify additional and more specific technical limitations and, hopefully, provide recommendations on methods to mitigate the impact of these limitations.

At this time, the following technical limitations are known to impact the proposed VM/MLS solution.

1. The heavy reliance on system memory to achieve the isolation goals of the solution may limit the size and number of simultaneous virtual machines on a single workstation. This limitation can be mitigated through the use of specialized hardware with a large amount of RAM. A detailed architecture should investigate this issue to provide specific metrics for memory usage and suggest a typical workstation hardware configuration, including system memory needs.

2. VM images will be downloaded at the beginning of each user session (e.g. each duty shift). It is expected that these images will be large, on the order of 1 Gig in size. As a result, the method by which these images are downloaded will dictate how quickly a guest operating system will be obtained and loaded.

3. It is expected that hardware cryptographic modules will play a role in improving the performance of encryption functions. It is not known to what extent solution components, such as NetTop, can leverage HCMs (although NetTop does support encryption of VM images). The result is that loading encrypted images may have to be a two-stage process, namely, decryption and loading. It is also notable that hardware cryptographic modules may also be used to improve the performance of transport layer data protection. There selection of appropriate and compatible products for this purpose must be determined.

4. Care must be applied in the detailed architecture to identify and mitigate sign points of failure for solution components such as the VPN concentrator.

5. It is not clear, at this time, how a single thumb drive can be used to store state information for multiple active VM images. In a worst-case scenario, separate tokens will be needed for each VM environment to which the user has access.

Similarly, the following issues are perceived to be operational constraints on the proposed VM/MLS solution. The focus in this section is a list of current capabilities, under the existing MLS solution that will no longer apply under the new solution, or new tasks, functions or responsibilities that must be undertaken to support the new solution.

1. The role of VM administrator by which VM images are created, maintained and made available for deployment must be defined and added to a system administrator's responsibilities. While this added responsibility will result in additional work, it is expected that the overall solution provides a cost savings in terms of administrative and maintenance overhead since a significant amount of hardware replication can be eliminated.

2. As system state information is saved to the user's thumb drive only on session termination, users must follow shut down procedures to ensure that information is properly and safely stored on their portable device.

3. Control over thumb drive usage, specifically as it applies to its usage within the policy-based network, will be an added task to the administration of this solution. This include binding a user to a portable device, integrating this information in the policy server and applying lifecycle management to these devices (e.g. revocation if lost).

8. Certification and Accreditation

It is recognized that a critical factor to the success of the VM/MLS technology deployment is the ability for this solution to achieve approval through a formal Certification and Accreditation (C&A) process. There are several existing C&A methodologies which could be applied to this solution; however, given the nature of the deployment environment and the nature of the processes that will operate within this environment, it has been determined that the Canadian Department of National Defence Information System Certification and Accreditation Guideline[2] (C&A Guide) is the most appropriate methodology in this instance.

This section will provide a discussion of the VM/MLS solution components in the context of the methodology provided in the C&A Guide. This section will map the salient points from the C&A Guide to the significant components from the proposed solution. The resulting discussion will examine each component (including the complete end-to-end solution) in the context of the major C&A processes. Specifically, this section will address the following C&A implications of the solution components:

- What elements of the C&A process have been achieved to date and what is likely achievable in the near term;
- What challenges/obstacles will be encountered during the C&A process;
- What is the suggested route for proceeding through the C&A process; and
- What is the expected timeframe (level of effort) for the C&A effort.

The ultimate goal of this effort is to receive accreditation, or formal acceptance, of the adequacy of a system's security and functionality. Acceptance of the information gained through the certification activities implies that there is a fundamental understanding of:

1. The level of protection that the system will provide in its stated configuration; and

2. The risks associated with the deployment and use of such a system.

To achieve accreditation, certification information must be gained about the system under evaluation. This is a formal evaluation of the security components that comprise the system's security safeguards. This evaluation will also examine the system as a whole to assess the degree to which these safeguard work together to provide a holistic security posture.

Within the context of the VM/MLS solution each of the main solution components, as defined in Figure 7: Sample User Session, will factor in the certification process:

- Virtual Machines;
- Secure operating systems;
- Authentication Server;
- Hardware cryptographic modules;
- Policy server /supported services; and

- The Aggregated VM/MLS Solution.

A discussion of each component is provided in the context of the C&A implications are provided in subsequent sections.

Table 3: Existing / Recent Certification Activities

Component	Description
Virtual Machines	As described in Section 3.1 Virtual Machines, the HP NetTop solution, being a combination of VMware virtual machine technology in combination with SELinux, has achieved 1.8.50 Certification from NSA (August 2003). This certification includes the ability to domain separation up to Top Secret.
Host Operating System	Within the NetTop solution, SELinux has been certified to operate within the NetTop solution architecture. For the development of a virtual machine-based solution, there will generally be a limited number of operating systems that will be supported by the VM software. For example, VMware 4.5 supports specific Linux distributions, including Red Hat Enterprise Server, SUSE and Mandrake. There is commitment to pursue certification for the Red Hat Enterprise Server product. Red Hat Enterprise Server 4 is currently EAL 3 certified and Novell SUSE Linux Enterprise Server 9, has been EAL 4 certified according to the Common Criteria rating. Red Hat Enterprise Linux 5 is currently in the process for certification to achieve EAL 4. Successful EAL 4 certification will indicates that a product meets government security standards for assured information sharing within and across government agencies.
Authentication Server	It is expected that the deployment environment will have pre-existing authentication infrastructure. It is further anticipated that the Entrust suite of products will be used for PKI applications. Entrust has achieved certification for many of its security products, including the Entrust Authority Security Manager 7.0 which has received an EAL 4+ evaluation.
Hardware Cryptographic Modules	There are existing hardware cryptographic accelerators that can be used in conjunction with the previously stated solution components to improve the performance of the encryption/signing activities. The SafeNet/Chrysalis line of security devices not only provides the needed functionality, but also have been successfully achieved FIPS 140-2 validation. Validation of the HCM component will aid in the meeting of the project C&A goals.
Policy-Based Networking	Currently, there are no known certification activities involving policy-based networking components. As such, it is anticipated that this component of the proposed VM/MLS solution will present the most significant challenge to successfully achieving C&A. At this time, it has not been determine if the policy server requirements can be met through existing solutions, can be met by extending existing solutions or must be developed as a new product. Extending existing solutions will not necessarily translate to a faster route to certification if these solutions have not been evaluated in a formal context. In actuality, a custom solution may be the simplest route to certification since the solution can be specifically designed to solve the VM/MLS problem space, thus reducing the overall scope of the solution evaluation process.

It is significant to note that the certification and accreditation process is related to implementations, rather than general architectures or products. While the use of products that have been approved through an evaluation process (e.g. common criteria or FIPS) will be

beneficial, the certification process is more concerned with implementation details such as configuration and processes. As a result, it is difficult to apply previous certification activities to the current certification process since there are many variables that change from implementation to implementation, such as local departmental policies and data sensitivity. For this reason, there is no formal reference for certified GoC solution architectures since no two systems are exactly alike.

As a final point on this topic, it is important to take the certification process in the spirit in which it is intended. Given that an implementation is able to show that the following risk management actions are in place:

- Certified/approved cryptographic solution within the system space;
- Good security practices for the system;
- Good change control management; and
- Execution of a Threat and Risk Assessment with implementation of the recommended changes to ensure that the residual risk is reduced to an acceptable level.

Then the C&A process can be successfully completed, no matter what has previously been done in terms of similar solution architectures. It is the position of this paper than C&A for the proposed solution is achievable.

8.1 Information System Security Components Mapped to C&A Deliverables

This section presents solution specific issues that will impact the C&A process for the VM/MLS solution. Information is presented in the context of the C&A deliverables that are expected to be provided to the evaluation team as part of a comprehensive C&A effort. The C&A implications listed below pertain the VM/MLS solution in its entirety, however, where specific components will impact the process, these components are identified specifically.

8.1.1 Concept of Operation (CONOP)

In defining the CONOP document, the purpose of the system and the criticality of the problem that is being solved by the VM/MLS solution must be understood by the key players that will use and maintain the system. From a system perspective, the CONOP must demonstrate the basic operational rules implemented by the system and presents a broad description of the anticipated system in support to the organization. Within this context, the policy rules, and policy server that enforces them, will require an in depth description to indicate how they control the distribution of sensitive VM images. This description may be best demonstrated through the development of scenarios; however, a formalized statement of how the VM distribution policy enforces access control protection will be needed for the C&A process.

Additionally, the CONOP document will require a description of the operational security practices, including identification of the classes or categories of users that will be involved in the use, maintenance or security functioning of the solution components. Note that the

inclusion of each solution component will incur specific tasks, the responsibility for which must reside with an existing or new user role. For example, the installation, definition, operation, modification, and auditing of the policy server are system component tasks that must be delegated to specific staff. Similarly, existing roles (e.g. security officer) should be examined in the context of the VM/MLS solution to identify new responsibilities that are inherent in the deployment of the new system

8.1.2 Threat Risk Assessment (TRA)

A TRA constitutes the core of the C&A process since many data sources, environmental concerns and threat scenarios are examined in the context of achieving an acceptable level of risk. The initial phase of the TRA process will derive the Statement of Sensitivity (SoS) for the classes of information that will be evaluated in the TRA process. For the VM/MLS solution, this will be a simple process, as the sensitivity level of the various compartmentalized environments that are accessed through the VM images will be known.

The sensitivity of additional information that is incurred as part of the VM/MLS solution, such as the policy data and system configuration information will have to be assessed during the TRA process.

One significant SoS concern is the availability requirement of operational data and services. If MLS environments are only accessible through the VM images that are, in turn, gained through the policy server and deployment mechanism, the VM/MLS infrastructure may be a single point of failure. Appropriate redundancy must be put in place to ensure that access to environments is not restricted by failures at the VM/MLS infrastructure.

8.1.2.1 Physical Security

VM/MLS solution components must be physically protected. The policy server, distribution point and VPN infrastructure must be located in a zone where physical access controls are in place to restrict unauthorized access. Workstations must be protected in such a manner than the level of physical access is appropriate for the sensitivity level of the information that will be processed on the system (as defined by the VM image distribution policy).

The solution components associated with the authentication service may also require physical protection. This need will be a function of the assurance level that can be attributed to the authentication solution that is chosen for deployment.

All systems that host solution components should be hardened to minimize the attack surface presented to potential attackers.

Since the workstation ID will be used to identify the location of the system, protections must be put in place to prevent the system from being moved to less secure environment.

8.1.2.2 Personnel Security

Responsibility for management of the VM/MLS solution components must be restricted to personnel that have been properly screened to a level appropriate for the sensitivity of the data.

8.1.2.3 Procedural Security

While it is expected that the individual compartmentalized environments will have pre-existing security policies, the VM/MLS solution will introduce an additional policy enforcement point in terms of the distribution of VM images to persons that meet the appropriate criteria (identification, location, role, etc…).

Management activities involving the operation of the VM/MLS solution must be assigned, monitored and auditing as appropriate for ensuring the correct security practices are being soloed. These activities include:

a. System (software/hardware) maintenance on VM/MLS components;

b. Infrastructure maintenance;

c. Definition and maintenance of VM images;

d. VM/MLS distribution policy definition and review; and

e. Monitoring of system state and auditing of system usage.

Procedures must be in place for each identified activity to ensure that the associated tasks are carried out in a manner that does not compromise the security of the solution and takes place with full knowledge and approval of the data owners.

8.1.2.4 IT Security

The following section identifies technical safeguards issues that will be examined during the TRA process.

a. Emissions Security: Given the shipboard nature of the VM/MLS solution, it is expected that it could operate in two modes, depending on the proximity of the ship to potential eavesdropping. If the ship is in port/close to another vessel, Top Secret VM images could be restricted to emissions protected zones/devices. Once at sea, the restrictions on Top Secret processing could be relaxed in the face of a reduced eavesdropping threat. Solution components could be deployed on EMSEC-approved hardware, should the TRA identify eavesdropping as a significant risk.

b. Cryptographic Security: Information that is written to portable storage devices may require protection via cryptographic transformation, depending on the sensitivity level of the compartmentalized environment to which the data pertains. The selected cryptographic hardware/software must meet approved standards for this use.

Hardware and software components that perform the authentication process must similarly be approved prior to use.

c. Computer Security: Safeguards must be in place that ensure that the operating system remains intact, retains its installed configuration and continues to operate as intended. This applies not only to the VM/MLS infrastructure components and workstations, but also to the guest operating systems that will be hosted at these workstations. The use of a certified operating system will aid in this process, however, additional safeguards may be required, such as integrity checking software and consolidated auditing routines.

d. Network Security: While not specifically stated in this document, it is expected that network security devices will be in place to maintaining integrity of the network configuration to prevent unauthorized re-routing, modification, destruction or disclosure of information or data. Both active and passive devices, such as firewall and IDS solutions, respectively, will have a role in achieving Network Security for the solution space. A more detailed network architecture will identify the security components that are needed and the TRA process will validate the selection, deployment and configuration of these devices.

8.1.3 Contingency Planning

Given the nature of the environment to which the VM/MLS solution will be deployed, it is recognized that the solution will potentially be placed in an operational warfare scenario. As such, there must be a development of continuity plans for both normal and combat operations. In both scenarios, the goal is to respond to a catastrophe and return to an operational state as quickly as possible so as not to jeopardize the mission. The difference between the scenarios is the speed at which the response must happen and the level of functionality that are needed to be deemed operational.

For example, in combat, the need to respond quickly to service failure is high, however, not all functions are immediately needed for resumption of the ship's mission. After a catastrophic event, the need to access weapons and navigation tools has priority over lesser function such as those that are available within the Unclassified network. As stated earlier in this document, the VM/MLS solution can provide assistance in more quickly regaining normal operational state in an emergency by relaxing the policy to allow access to sensitive network from alternate locations within the ship.

The contingency plan that is prepared as part of the C&A process should be mindful of the features of the VM/MLS solution as the following questions are addressed to establish the plan's content:

- What are the critical operational issues to overcome in an emergency?
- What are the critical human resource issues to overcome?
- What are the specific or unique contingency measures?
- What are the critical physical location issues to overcome?

- What are the critical computer availability issues?

These questions must be addressed in the context of the TRA process that will drive the contingency planning requirements.

8.1.4 Change management Plan

It is assumed that there are existing change management practices in place for each of the compartmentalized environments to which the VM/MLS solution will enable access. The solution itself must track, in a formalized manner, changes to the system solution configuration as it pertains to:

- New VM images that are released for use;
- Changes to the VM image distribution policy;
- Changes to the VM/MLS solution component software;
- Changes to the host operating system configuration;
- Change to the hardware deployments, including the location of hardware within each security processing zone;
- Changes to the configuration of infrastructure components; and
- Changes to the tool used to monitor, maintain and audit solution components.

The change management practices must be established to ensure that any system changes are made in such a manner that they do not compromise system security and are made with full knowledge of the system and data owners.

8.2 Level of Effort

The recommended practice for effective attainment of C&A is to interleave the C&A process with the system development activities. As such, the development of the CONOP, TRA, contingency and change management plans should occur with the development of the system. This will substantially reduce the time needed to meet the C&A project objectives. This interleaving of activities will necessitate the allocation of additional time for the various tasks in the system design and development activities; however, the overall C&A effort will be substantially reduced. By interleaving these activities, information gathering, threat identification, safeguard selection and risk management issues are identified as they are encountered and can be immediately captured by the C&A team

An initial step of C&A planning is required at the forefront of this effort to ensure that the correct personnel have been included in the C&A team, that the C&A goals are fully understood among the team, that solution assumptions are properly documented, that a reporting structure is in place for the sharing of information and that a reasonable schedule has been set. The DND IS C&A guide provides advice on how to best establish the C&A plan and how to interleave C&A activities during a system design / development cycle.

The use of certified and approved solution components will similarly reduce the effort to achieve the C&A objective. Approval can come from many areas of security evaluation depending on the nature of the component; including FIPS 140-2 validation; Common Criteria evaluation and RCMP approved equipment selection.

It is expected that the policy server component will be the most challenging element to attaining C&A for the VM/MLS solution. The lack of certified policy server solutions, the potential need for customized software development and the need to establish and test the policy server functions will result in substantial effort to prove the robustness and security of this component. While it is not known at this time if the effort to create a policy server will require the creation of a new solution or the extension of existing tools, it is anticipated that some level of software development will be needed. The DND IS C&A Guidelines provides a Security Checklist for Software Development activities as part of a discussion of Information Technology Security (ITSEC). This checklist gives guidance on the use of a software development life cycle, the enforcement of software development controls and quality assurance. Linking ITSEC guidelines to the C&A process will greatly assist in achieving the C&A objectives for any software development activities.

9. Future Research

The following is a list of logical next steps in the development of the proposed MLS/VM solution.

Requirements analysis: The conceptual technical overview has been written with a high level understanding of the solution requirements. To progress the solution to the point where a prototype solution can be achieved, a more detailed understanding of the operational needs and constraints on the eventual solution is needed. It is expected that this effort will require a series of interviews with key personnel that will have a role in the deployment and use of the proposed solution.

Detailed solution architecture: A detailed solution architecture will refine the solution proposed in this paper to reflect the actual solution configuration. This architecture will be written to closely meet the requirement taken from the previous activity to ensure that the solution will be acceptable to the target user community. This activity will also define the system, network and security configuration for components in the solution space.

Policy Server development: In conjunction with the previous activity, this activity will focus on the creation of the policy server component. It has been recognized that the policy server will likely require an element of software develop to create a component that provides the required policy enforcement and policy decision points. The extent to which this will be extension of existing software as opposed to a complete custom solution will be a primary decision point for this effort.

Prototype / Proof-of-Concept: It is recommended that an initial prototype be developed to prove the viability of the proposed solution. The prototype can also serve to identify deployment constraints and obtain solution metrics. The prototype can also serve as a demonstrator to illustrate the usefulness of the solution and foster interest in pursuing a full solution deployment.

As stated in this document, the C&A process can be most effectively achieved when the C&A activities are integrated into the system design, development and validation phases. As a result, it is recommended that any work activities include a C&A component to the assigned deliverables.

References

[1] Bell, D.D., L.J. La Padula, Secure Computer System: Unified Exposition and Multics Interpretation, ESD-TR-75-306. Bedford, MA: ESD/AFSC, Hanscom AFB.

[2] Government of Canada, Canadian Department of National Defence Information System Certification and Accreditation Guideline

[3] Peter Loscocco, Stephen Smalley, Meeting Critical Security Objectives with Security-Enhanced Linux, Proceedings of the 2001 Ottawa Linux Symposium, July 2001.

[4] John Merrells, XACML: XML Access Control, Parthenon Computing

[5] Tony Musgrave, Hewlett-Packard, HP NetTop: High Assurance Computing using SELinux and Virtual Machines, 2004

[6] Stephen Smalley, Timothy Fraser, A Security Policy Configuration for the Security-Enhanced Linux, NAI Labs Technical Report, February 2001.

[7] Sun Microsystems, A Brief Introduction to XACML, March 14, 2003

[8] Manish Verma, XML Security: Control information access with XACML, Second Foundation, Oct 2004

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Cinnabar Networks, a division of Bell Security Solutions Inc.
Suite 200, 265 Carling Avenue
Ottawa ON K1S 2E1

2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)

UNCLASSIFIED

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C, R or U) in parentheses after the title.)

Applying Virtual Machine Technology to Achieve Multi-Level Security (U)

4. AUTHORS (last name, followed by initials – ranks, titles, etc. not to be used)

Henderson, G., Tremblay, L.

5. DATE OF PUBLICATION (Month and year of publication of document.)

March 2006

6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)

56

6b. NO. OF REFS (Total cited in document.)

8

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Contract Report

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

DRDC Ottawa, NIO Section
3701 Carling Avenue, Ottawa ON K1A 0Z4

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)

15BQ03

9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)

W7714-5-3171

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DRD-6-041-1

10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

DRDC Ottawa CR 2006-087

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

(X) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

This document presents a conceptual technical overview for utilizing virtual machine (VM) technology to achieve a Multi-Level Secure (MLS) solution. The goal of this effort is to define a solution architecture that can leverage the savings in space and infrastructure that can be realized through the use of virtual system images while still adhering to the security principles that define an MLS environment. As part of this effort, an emphasis has been made to illustrate the isolation that exists between virtual machines and the hosting environment in the areas of information processing, information storage and information transmission. This document provides an architectural approach that utilizes VM images that are distributed according to a flexible, yet secure, policy. This policy defines the conditions under which potentially sensitive system images can be distributed and accessed. It is the position of this paper that VM technology can be leveraged to provide a more effective MLS solution while still maintaining the needed separation to ensure sensitive data assets are protected.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

virtualization multi-level policy VM MLS

Defence R&D Canada

Canada's leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

www.drdc-rddc.gc.ca