Communications Security Establishment
Centre de la sécurité des télécommunications

CANADIAN CENTRE FOR CYBER SECURITY

CYBER CENTRE DATA CENTRE VIRTUALIZATION REPORT: BEST PRACTICES FOR DATA CENTRE VIRTUALIZATION

(ITSP.70.010)

March 2020

PRACTITIONER

TLP:WHITE

Canada

FOREWORD

ITSP.70.010 Cyber Centre Data Centre Virtualization Report: Best Practices for Data Centre Virtualization is an UNCLASSIFIED publication, issued under the authority of the Head, Canadian Centre for Cyber Security (Cyber Centre). For more information, email or phone our Contact Centre:

Cyber Centre Contact Centre
contact@cyber.gc.ca
613-949-7048 or 1-833-CYBER-88

EFFECTIVE DATE

This publication takes effect on (03/27/2020).

REVISION HISTORY

Revision	Amendments	Date
1	First release.	March 27, 2020

OVERVIEW

To improve efficiency and streamline operations, many organizations are virtualizing their data centres. However, some are doing so without a proper understanding of the security implications of their decisions on the overall security posture of the data centre. By increasing their consolidation ratios (i.e. the number of virtual machines that can operate on each physical host machine), many organizations may unintentionally compromise their network security architecture. Moreover, any attempt to replace many of the physical isolation mechanisms with logical equivalents may introduce unacceptable risk into their network architecture.

Not only are virtualized data centres (VDCs) susceptible to most of the vulnerabilities inherent in traditional data centres, but they introduce new vulnerabilities specific to this environment. While complex virtualization threats (e.g. System Management Mode (SMM) rootkits, side-channel attacks, and hyperjacking) must be addressed, it is the more practical threats that are of greater risk to the VDC. These practical threats include things such as misconfiguration, compromise of the management interface, and conventional attacks on both the hypervisor and virtual machines (VMs).

A VDC must address the threats, inherent vulnerabilities, and characteristic risks to data centres, as well as those specific to complex virtualized environments. Fortunately, VDCs can be made secure by using several safeguards and best practices. These safeguards and best practices involve addressing each layer of the virtual environment as well as addressing interactions between the various layers. This report provides an overview of a VDC, presents the inherent vulnerabilities in a VDC, and provides recommendations on how to design a secure VDC.

TABLE OF CONTENTS

LIST OF FIGURES

1 INTRODUCTION

Many organizations are consolidating traditional data centres into VDCs[1] to save money through improved consolidation ratios and economies of scale. Technological developments, such as multi-core processor and hyper-threading technologies, have dramatically increased potential consolidation ratios, making virtualization an operational imperative.

However, many organizations proceed without a proper understanding of the security implications of virtualization, in general, and the VDC, specifically. Not only are VDCs filled with many of the vulnerabilities inherent in traditional data centres, but they introduce new vulnerabilities. Similarly, by increasing consolidation ratios, many organizations may unintentionally compromise their network security architecture.

Organizations must have a proper understanding of the security implications of their decisions on the overall security posture of the VDC. Also, they must consider the security implications of the decisions made in every step of the design and the deployment of the VDC. Only then can organizations be confident that the resulting infrastructure has been implemented to mitigate identified vulnerabilities against known threats.

1.1 POLICY DRIVERS

There are several Government of Canada (GC) policies that address IT security requirements. GC departments must ensure that all IT security policies and procedures align with the following Treasury Board of Canada Secretariat (TBS) policies:

- *Policy on Government Security* (PGS) [1][2];
- *Directive on Security Management* [2];
- *Directive on Privacy Practices* [3];
- Financial Administration Act [4];
- *Guideline on Acceptable Network and Device Use* [5]; and
- *Policy on Management of Materiel* [6].

Non-GC organizations can use these policies as reference materials when developing their own IT security policy frameworks.

[1] A VDC refers to a data centre that uses virtualization (data centre virtualization), to a considerable degree, to abstract compute, network, storage, security, and manage resources from the underlying hardware.

[2] Numbers in square brackets refer to references cited in the Supporting Content section of this document.

1.2 APPLICABLE ENVIRONMENTS

This document provides guidance only for unclassified IT systems that may hold sensitive information or assets that, if compromised, could reasonably be expected to cause injury to an individual interest, such as a person or an organization (i.e. personal information[3] and business information[4]). Within the GC context, this guidance can be applied to IT systems that hold Protected A and/or Protected B information.

This document does not provide guidance for IT systems that hold **highly sensitive information or assets of an individual interest** (i.e. Protected C information in the GC context) and **sensitive information or assets of a national interest** (i.e. classified information[5]). IT systems that hold this type of information require additional design considerations that are not within the scope of this document.[6]

[3] As defined in the *Privacy Act* [7] and the *Personal Information Protection and Electronic Documents Act* [8], personal information is "information about an identifiable individual that is recorded in any form".

[4] Business information in this context refers to information that may reasonably be expected to cause injury to an organization, as defined in subsection 20(1) of the *Access to Information Act* [9].

[5] Within the GC context, classified information is any information or assets that, if compromised, could reasonably be expected to cause injury to the national interest, defence, and maintenance of the social, political and economic stability of Canada. Information is classified at the Confidential, Secret, and Top Secret levels depending on the type of information and the potential injury.

[6] Contact the Cyber Centre Contact Centre for guidance regarding cryptographic solutions in PROTECTED C or Classified environments.

1.3 RELATIONSHIP TO THE IT RISK MANAGEMENT PROCESS

ITSG-33 *IT Security Risk Management: A Lifecycle Approach* [10] outlines two levels of IT security risk management activities: organizational-level activities and information system-level activities. These two levels of activities are outlined in Figure 1.

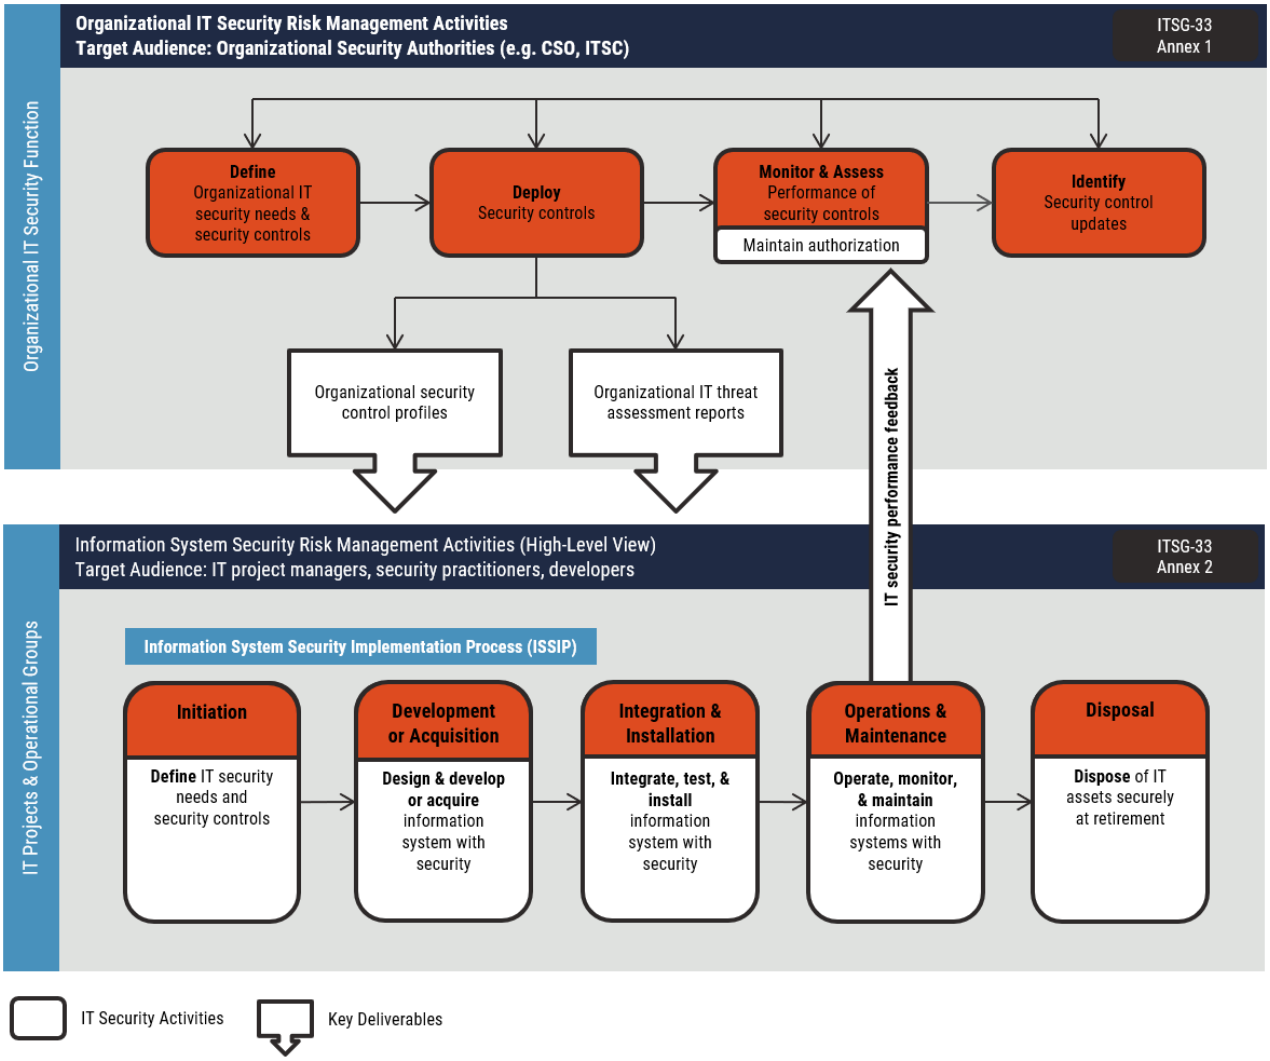

Figure 1 IT Security Risk Management Process

1.3.1 ORGANIZATIONAL-LEVEL ACTIVITIES

Organizational-level activities should be integrated into the organization's security program to plan, manage, assess and improve the management of IT security-related risks faced by the organization. This document will need to be considered during the Define, Deploy, and Monitor and Assess phase. These activities are described in detail in Annex 1 of ITSG-33 [10].

1.3.2 INFORMATION SYSTEM-LEVEL ACTIVITIES

Information system-level activities should be integrated into an information system lifecycle to ensure the IT security needs of supported business activities are met, the appropriate security controls are implemented and operating as intended, and the continued performance of the implemented security controls is assessed, reported back, and acted upon to address any issues. Organizations should consider the guidance in this document during the following phases of the information system security implementation process (ISSIP):

- Initiation;
- Development/Acquisition;
- Integration and Installation; and,
- Operations and Maintenance.

These activities are described in detail in Annex 2 of ITSG-33 [10].

1.4 APPROACH

The best practices presented in this report were developed using several sources, including industry sources (both vendor and analyst), US government sources (e.g. the National Institute of Standards and Technology [NIST]), and in-house expertise.

This report provides best practices for securing a VDC. To secure a VDC, the organization is responsible for implementing these best practices in a way that supports their business processes and provides the level of security assurance required. Security assurance is discussed in considerable detail in ITSG-33 [10], specifically in Section 8 of Annex 2. According to ITSG-33 [9], the security assurance level (SAL) consists of a pre-selected set of security assurance requirements that gives a degree of confidence in the adequacy of the security engineering and documentation work performed by the project team. Ultimately, the implemented security controls should perform as intended and satisfy the organization's business needs for security.

It is believed that an organization that implements the best practices outlined in this document, with a SAL of 3, could have a reasonable amount of confidence that its VDC is protected from virtualization-specific threats executed by highly sophisticated threat actors (i.e. categorized at the Threat deliberate 5 [Td5]) level and below). In this report, the term **secure** refers to this state. It is not intended to imply an absolute state of security.

1.5 SCOPE

This report addresses data centre virtualization, which encompasses a variety of virtualization activities aimed at creating a VDC. One of these virtualization activities is server virtualization. Server virtualization uses full virtualization[7] and, most often, bare-metal virtualization[8]. Desktop virtualization, which uses hosted virtualization[9], is out of scope due to key architectural and security differences. Application virtualization and operating system virtualization[10] are also outside of the scope of this report.

The best practices provided in this report are applicable to all VDCs. Consequently, this guidance is vendor-agnostic. We have removed vendor-specific security guidance from this report.

This report focuses primarily on technical and design issues related to the VDC. While many security issues in a VDC are operational, these issues are largely out of scope. For a comprehensive look at operational security controls, refer to ITSG-33 [10].

1.6 ASSUMPTIONS

This report assumes that you have a basic understanding of virtualization concepts and terms.

[7] Full virtualization was the first type of x86 virtualization to be developed, and it remains the most prominent type to date. It is referred to as full virtualization because the guest operating system and associated applications run on a VM that is fully abstracted from the underlying hardware by the virtualization layer.

[8] In bare-metal virtualization, the hypervisor runs directly on the system hardware.

[9] In hosted virtualization, which is another type of full virtualization, the hypervisor runs on top of a standard operating system and relies on it for all interaction with the underlying hardware.

[10] Operating system virtualization, sometimes referred to as shared kernel virtualization, differs from full virtualization. Applications are hosted in virtual environments on a common operating system rather than in VMs with a separately installed operating system.

2 VIRTUALIZED DATA CENTRE OVERVIEW

2.1 OVERVIEW

Virtualization introduces a layer of abstraction between the underlying physical resource and the service requesting the resource. In a VDC, virtualization is used to abstract many of the physical resources including servers, networks, and storage. By abstracting these resources, the VDC can achieve high consolidation ratios[11] and considerable flexibility. However, the challenge is to achieve these objectives without compromising security (see note below). This section examines the seven integral layers in VDCs. These layers, illustrated in Figure 2, provide the basic structure for discussion:

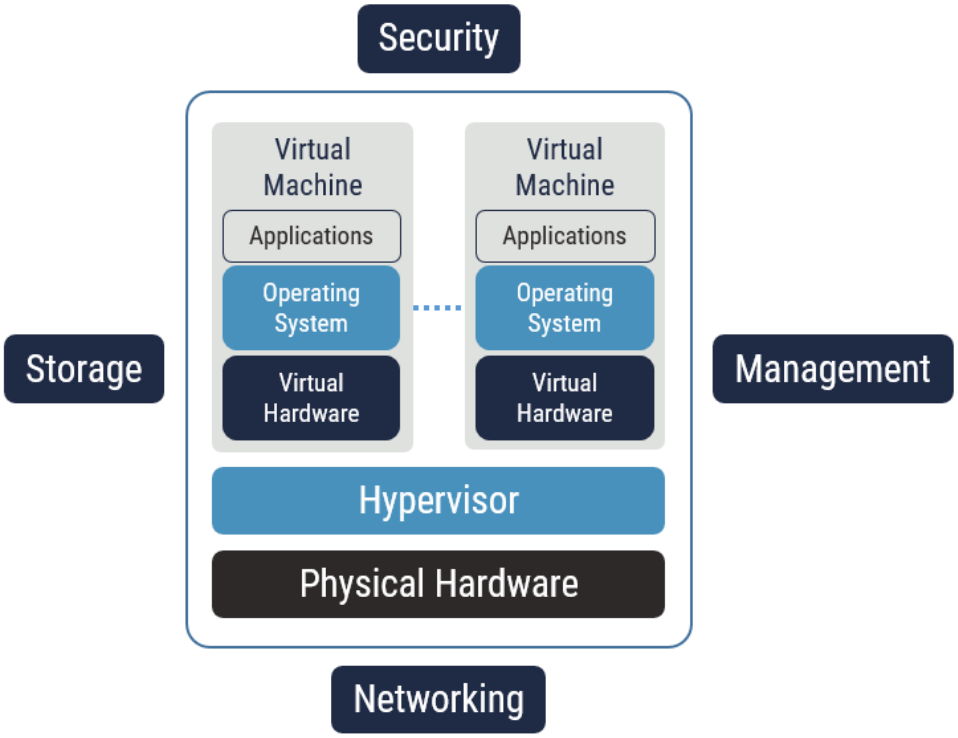

Figure 2 VDC Layers

Security/Flexibility Trade-off: Organizations implementing a VDC are primarily interested in virtualization to improve consolidation ratios and to obtain greater flexibility. Improved consolidation ratios translate to cost savings. By abstracting physical resources, virtualization allows organizations to make changes easily in response to internal requirements or external stimuli. Also, virtualization greatly improves the ease with which resources can be managed and provisioned.

However, this flexibility can come at the expense of security. The trade-off between security and flexibility is a recurring theme throughout this document (see VDC reference architecture best practices #1 in Section 3).

[11] Consolidation ratio: the number of virtual devices that can run on each physical device.

2.2 NETWORKING

Networking within a traditional data centre is relatively straightforward, but networking in a VDC is more complex. This section examines the networks required in a VDC. It also examines using network security zones in a VDC.

2.2.1 REQUIRED NETWORKS

Four networks are typically required in a VDC. These networks can be either physically separated networks, logically separated networks, or a combination of both[12]. Physically separated networks are helpful in that they provide complete separation but at the expense of cost and convenience. In contrast, logically separated networks are convenient and inexpensive to implement but may not provide the level of separation necessary. The four required networks are the data network, the management network, the storage network, and the live migration network. These networks are described below.

2.2.1.1 DATA NETWORK

The data network is used for all communications to, from, and between production servers in the VDC. It is the same as the core network that exists in all physical network environments. The data network should be isolated from the other networks for confidentiality, integrity, and availability.

2.2.1.2 MANAGEMENT NETWORK

In a VDC, the management network is used for all communications between the virtualization management interface and the hypervisor. The management network separates management traffic used for administration or maintenance (e.g. monitoring, logging, software updates) from the actual data processed and stored within the VDC. As a result, the management network should be isolated and secured to prevent a threat actor from gaining privileged access and compromising the virtual infrastructure and to ensure availability of the management network during an event on the data network.

2.2.1.3 STORAGE NETWORK

A Storage Area Network (SAN)[13] is used for all communications between the hypervisor and the storage array. It is used primarily to transfer VMs from the storage array, on which they are stored, to the server on which they are running. The SAN should be isolated to prevent a threat actor from gaining access to VMs while they are in transit. Some storage technologies (e.g. fibre channel) need to use a separate physical network for technological reasons.

---

[12] Theoretically, all four types of traffic could use the same physical network without any additional controls. This network consolidation, without using logical separation, should be strongly discouraged as it does not afford adequate protection to traffic.

[13] While a SAN is typically used in a data centre, nothing precludes using Network-Attached Storage (NAS). NAS, as the name implies, is basically a storage subsystem that is connected to the network. It runs a stripped-down operating system capable of providing both a file system and storage, in much the same way as a network file server. Although the term SAN will be used throughout the report, much of the guidance is applicable to NAS.

It should be noted that the storage network discussed in this section refers to the storage network required for virtualization. Specifically, this storage network is used to transfer VMs from the disk arrays, on which they are stored, to the physical servers on which they are hosted. In addition to virtualization storage, there is a requirement for enterprise data storage. Enterprise data storage should use a separate storage infrastructure.

2.2.1.4 LIVE MIGRATION NETWORK

The live migration network is used to transfer VMs between servers located in the same network security zone for reasons of performance or maintenance. The live migration network should be isolated for confidentiality and integrity to prevent a threat actor from gaining access to VMs, and specifically VM memory, while the VMs are in transit. Isolation also prevents availability due to conflict between VM migration data network operations.

2.2.2 NETWORK SECURITY ZONES

Network security zones allow for the logical partitioning of the data network according to a common security policy and common security requirements. Services with common security requirements are grouped into the same network security zone and protected according to a common security policy comparable with the sensitivity of the service provided and/or data processed. Also, network security zones support a defence-in-depth strategy. That is, more sensitive resources are protected by successive safeguards. *ITSP.80.022 Baseline Security Requirements for Network Security Zones* [11] and *ITSG-38 Network Security Zoning – Design Considerations for Placement of Services within Zones* [12] are two publications that provide guidance on implementing network security zones within the GC.

Network security zones are applicable in VDCs just as they are in their physical counterparts. However, in VDCs, extra care needs to be taken to ensure that using virtualization does not unintentionally compromise the integrity of network security zones nor compromise the boundary separating network security zones (see VDC reference architecture best practices #2 in Section 3).

The network security zones required in a VDC are illustrated in Figure 3:

- Public Zone (PZ);
- Public Access Zone (PAZ);
- Operations Zone (OZ);
- Restricted Zone (RZ); and
- Management Zone (MZ).

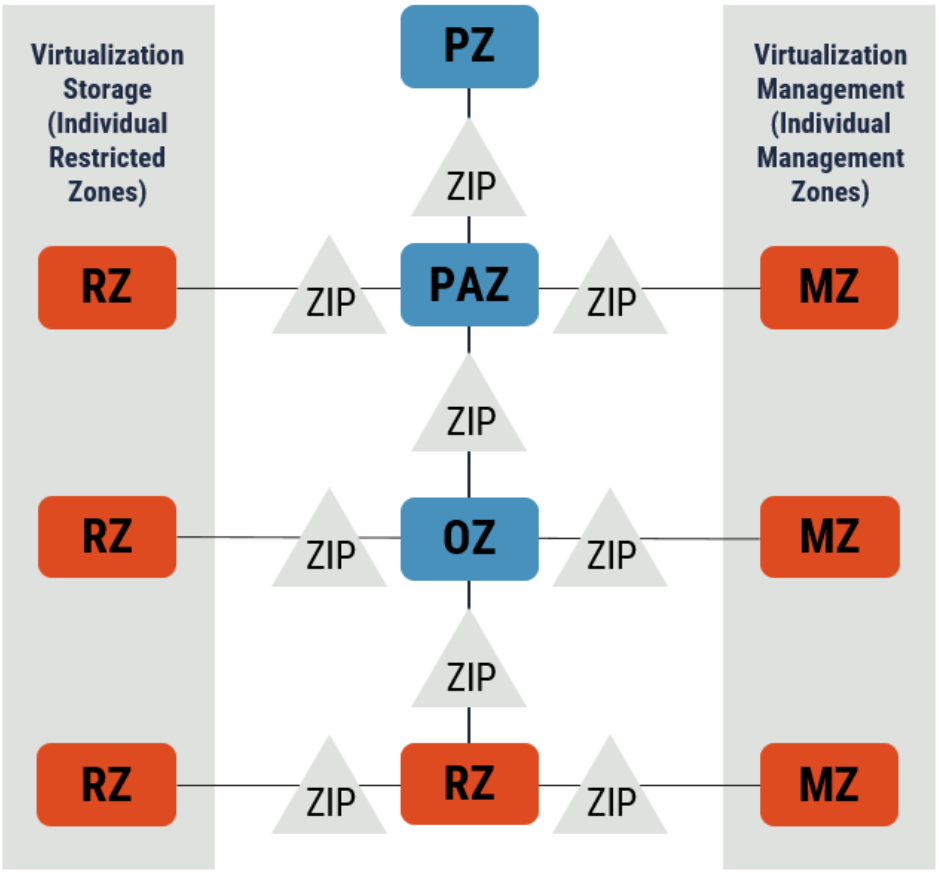

Figure 3 VDC Network Security Zones

Boundary/Zone Interface Points (ZIPs): Security boundaries, or ZIPs, are located between each network security zone. They typically contain firewalls, or filtering routers, that limit the permitted communications in and out of the zone. In addition, they often include Intrusion Detection Systems (IDS) that monitor traffic flow to detect abnormal behaviour and Intrusion Prevention Systems (IPS) that respond to any abnormal behaviour detected. Virtualization of security appliances in the boundary is discussed in Section 2.8.

2.2.2.1 PUBLIC ZONE

A Public Zone (PZ) is a network that is outside the VDC's control.

2.2.2.2 PUBLIC ACCESS ZONE

The Public Access Zone (PAZ) mediates access between the Operations Zone (OZ) and the PZ. The PAZ serves to terminate protocol sessions between the OZ and PZ in both directions and to initiate new protocol sessions to the desired destination, which is typically accomplished using proxy servers. Hosts with critical applications or data should not be hosted in the PAZ. Access to other networks from the VDC should only occur through the PAZ.

2.2.2.3 OPERATIONS ZONE

The Operations Zone (OZ) is the standard environment for routine operations. In a VDC, it will contain most servers, which are accessible from external end users through the PAZ proxy services. VDCs could also have end-user systems (i.e. workstations) which are also located in the OZ. However, we recommend that a perimeter exist between OZ end-user systems and OZ servers.

2.2.2.4 RESTRICTED ZONE

The Restricted Zone (RZ) contains sensitive data and services, including business critical services and repositories of sensitive information. In a VDC, the Storage Network, (defined in Section 2.2.1.3), is a specific implementation of the RZ for virtual image storage.

2.2.2.5 MANAGEMENT ZONE

The Management Zone (MZ) is the location where the administrative systems which manage systems hosted in the PAZ, OZ, and RZ are situated. Care needs to be taken to ensure that the MZ does not compromise the separation provided by using the other network security zones (see VDC reference architecture best practices #3 and #4 in Section 3).

2.2.3 NETWORK FUNCTION VIRTUALIZATION

In a VDC, consolidation and virtualization can be implemented on the networking components. Using Software Defined Networking (SDN) and Network Function Virtualization (NFV) can reduce networking components to a commodity. Using these technologies can greatly assist with the optimization and consolidation of a VDC, but some care must be taken to plan and implement the security features available with SDN, as any traffic virtualized in this method becomes invisible to standard security appliances and firewalls that lie outside the virtualized environment.

2.2.3.1 MICRO-SEGMENTATION

Network security zones group assets together by sensitivity and function, providing boundary controls for traffic, broadly protecting the network from compromise external to the boundary, commonly referred to as North-South traffic. Micro-segmentation is a network security technique for virtualized networks, further logically dividing the network into distinct security segments down to the endpoint by defining security controls and services for each unique segment. This technique occurs within the broader security zoning. Micro-segmentation can be used to protect every VM in an enterprise network with policy-driven, application-level security controls at the workload level. This defends the interior of the network, commonly referred to as East-West traffic, from potential compromise.

2.3 PHYSICAL HARDWARE

While the physical hardware layer is the only layer that is largely unchanged from a traditional computing environment, there are a few important differences. Hardware compatibility, processors, network adapters, and switches are four important aspects of the physical hardware requirements for VDCs.

2.3.1 HARDWARE COMPATIBILITY

The majority of VDCs utilize bare-metal hypervisors that run directly on the physical hardware. However, while hypervisor support for hardware is extensive, not all servers are supported. As a result, hypervisor vendors typically maintain a hardware-compatibility list of supported systems.

2.3.2 PROCESSORS

Multi-core processors, as the name implies, have two or more independent Central Processing Unit (CPU) cores. Hyper-threading technology (HTT) basically simulates an additional processor for each CPU core. For example, a dual-core CPU with hyper-threading is seen by a traditional operating system as if it were a quad-core CPU. These two technologies, along with additional Random-Access Memory (RAM), increase the number of VMs that can be supported on a single system.

In addition, since late 2005, most x86 processors used in VDCs provide hardware-assisted virtualization. Although the details differ slightly between implementations, hardware-assisted virtualization is basically an extension to the x86 architecture that optimizes performance and improves security for virtualized environments. Hardware-assisted virtualization is discussed in more detail in Section 2.4.1.

2.3.3 NETWORK ADAPTERS

In a traditional data centre, most servers have a limited number of network adapters. In a VDC, servers typically have a greater number of network adapters to support the various networks required and to provide redundancy.

2.3.4 NETWORKING DEVICES

Using SDN, network functions such as switching, routing, and firewalls can be implemented in software using the underlying fabric of physical switches. Multiple networks can be overlaid upon the same infrastructure logically separated via Virtual Local Area Networks (VLANs). Not all switches are compatible with the various implementations of SDN and care should go into planning whether to utilize SDN and in which technology to invest.

2.4 HYPERVISOR

The hypervisor, also called a Virtual Machine Monitor (VMM), has direct access to the physical hardware of the system. The hypervisor, (illustrated in Figure 2), is responsible for partitioning and scheduling access to the physical server resources amongst the VMs running on it, while ensuring the appropriate level of isolation for the VMs. Bare-metal hypervisor is almost exclusively used in VDCs (see VDC reference architecture best practices #5 in Section 3). This section will examine three aspects of the hypervisor: protection rings, resource partitioning, and virtual networking.

2.4.1 PROTECTION RINGS

Protection rings are hierarchical privilege levels which were implemented in x86 processors to provide fault tolerance and malicious code protection. An understanding of protection rings is important for understanding several the vulnerabilities discussed in Section 3.3. Hardware-assisted virtualization (illustrated in Figure 4) is most commonly implemented in VDCs and consists of the following protection rings:

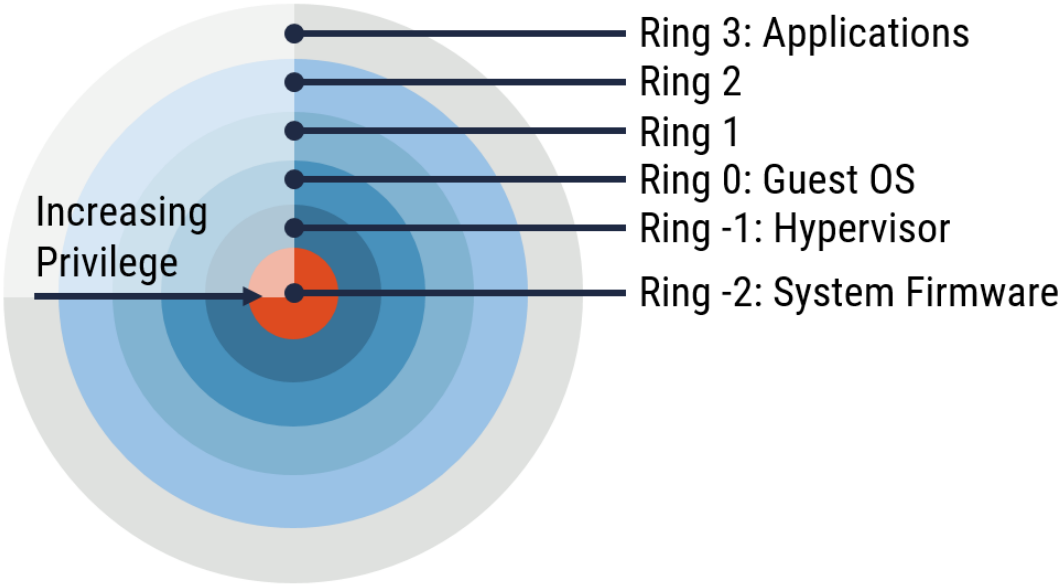

Figure 4 Protection Rings for x86 Hardware-assisted Virtualization

- **Ring 3**: Applications always run in Ring 3 regardless of whether they are traditional, fully virtualized, or hardware-assisted virtualization architecture.
- **Ring 2**: Ring 2 was intended to be used for device drivers in a traditional x86 architecture. It is not used for either full virtualization or hardware-assisted virtualization.
- **Ring 1**: Ring 1 was intended to be used for device drivers in a traditional x86 architecture. In full virtualization, the guest operating systems run in Ring 1. Since operating systems normally run in Ring 0, they are tricked into believing that they are running in Ring 0. The hypervisor accomplishes this by providing binary translation of a privileged operating system request. In hardware-assisted virtualization, Ring 1 is not used.
- **Ring 0**: Ring 0 is the level in which the operating system interacts with system hardware in a traditional x86 architecture. In full virtualization, the hypervisor runs in the same protection ring as the kernel and the device drivers. However, in hardware-assisted virtualization, Ring 0 is reserved for the guest operating systems. Since the guest operating systems are expecting to run in Ring 0, binary translation is not necessary.
- **Ring -1**: Ring -1 is a super-privileged protection ring that is introduced in hardware-assisted virtualization. It is reserved for the hypervisor.

- **Ring -2**: The System Management Module (SMM) code runs in a specially protected part of system memory and is accessible only by system firmware. This mode, which is sometimes referred to as Ring -2, is the most privileged CPU operation mode on x86 processors.

2.4.2 RESOURCE PARTITIONING

The hypervisor partitions hardware resources amongst the VMs that it hosts. This partitioning is typically accomplished logically rather than physically. In logical partitioning, the hypervisor shares the hardware resources amongst the VMs. The hypervisor can allocate these resources in such a way that any one VM cannot consume more than its share of the resources. These resource consumption limits prevent a VM, malicious or not, from consuming an excessive amount of resources and thereby depriving the other hosted VMs of resources. In physical partitioning, the hypervisor assigns distinct physical resources (e.g., disk partitions, disk drives, and network adapters) to each VM. While this approach limits the resources that can be consumed by any one VM, it also requires separate hardware resources, thus defeating the main purpose of virtualization. In a VDC employing bare-metal hypervisors, VMs can only interact with one another through permitted network communications. There should be no other communication path available. In contrast, hosted virtualization on a workstation typically supports cutting and pasting, as well as file transfer between VMs, which removes many of the separation controls offered by bare-metal virtualization.

2.4.3 VIRTUAL NETWORKING

VMs hosted on the same physical server, referred to as co-located VMs, can use the hypervisor for network communications with one another rather than a physical network. Using SDN, this can be extended throughout the data centre to any physical server connected to the switching fabric. Virtual networking can be a security concern as this network traffic is largely concealed from the view of traditional security infrastructure. Virtual networks consist of the following components:

- **Virtual Network Adapters**: Virtual network adapters, or Virtual Network Interface Cards (vNICs), are configured for each VM to dictate how the VM will connect to the physical and virtual networks.
- **Virtual Switches (vSwitches)**: vSwitches are the virtual equivalent of physical switches. They are used within virtual networking to interconnect VMs over the internal network, and to connect virtual network adapters with physical network adapters.
- **Portgroups**: A portgroup is an element of the vSwitch that enables it to be partitioned in much the same manner as ports on a physical switch are assigned to VLANs. vNICs are assigned to portgroups on a virtual switch. vNICs can communicate with other vNICs assigned to the same portgroup but are unable to communicate with vNICs assigned to other portgroups on the same virtual switch.

2.5 VIRTUAL MACHINE (VM)

A VM encapsulates an information system, including hardware, operating system, and applications, into a file. The operating system and applications, which are fully abstracted from the underlying hardware by the hypervisor, interact with the hypervisor to access the physical hardware of the system. This section will examine the three components of the VM: the virtual hardware, the operating system, and the applications.

2.5.1 VIRTUAL HARDWARE

The virtual hardware is the virtual representation of the physical hardware. The hypervisor virtualizes the hardware for each VM so that the guest operating system behaves as if it were running on a physical system. Specifically, each guest appears to have its own hardware, including the CPU, memory, storage (hard disk, CD/DVD), storage controllers, network interface cards (typically Ethernet), display and sound drivers, keyboard and mouse. However, VMs have no direct access to hardware; they only have visibility to virtual devices.

In the case of memory, each VM has access to RAM that it believes exists solely for its purpose. The physical memory is partitioned amongst VMs. When the memory is allocated to a VM, it is zeroed out. Virtualization software uses several methods to allocate memory in excess of the physical memory available, thus reducing the cost of memory.

2.5.2 OPERATING SYSTEM

A VM is capable of hosting most traditional enterprise operating systems. The operating system is identical in nearly all respects to an operating system running on a physical server. Also, the operating system is often unaware that it is running in a virtualized environment.

2.5.3 APPLICATIONS

A VM is capable of hosting almost any type of application. The application is identical in nearly all respects to an application running in a traditional environment. Also, the applications are typically unaware that they are running in a virtualized environment.

2.6 MANAGEMENT

Separation of duties refers to dividing roles and responsibilities so that a single individual cannot subvert a critical process. Separation of duties requires that, for sets of transactions, no single individual be allowed to execute all transactions within the set. The most commonly used examples are the separate transactions needed to initiate a payment and to authorize a payment. No single individual should be capable of executing both transactions. The principle of separation of duties is extremely important in a VDC.

To comply with the principle of separation of duties, administrative roles in a VDC should be partitioned. There is a tendency in a VDC to want to combine administrative roles to streamline administration. However, this temptation should be resisted as it runs contrary to the principle of separation of duties. Privileged access should only be provided to a limited number of

trusted individuals within the organization. Also, rather than provide an excessive amount of privilege, an organization might want to consider implementing a break-glass policy.[14]

Management roles that can be found in a VDC, as illustrated in Figure 5, typically include the following parts:

- **Application Management:** Applications in VDC are managed as in traditional computing environments.
- **Operating System Management:** Operating system management is like the same role in traditional computing environments. The only difference is that in a VDC many operating systems can be hosted on one physical server.
- **Virtualization Management:** Virtualization management is a new management layer specific to the VDC. Through virtualization management an attacker can manipulate VMs, including starting them up, shutting them down, migrating them to different physical servers, taking snapshots and reverting to previous snapshots. An attacker can also use virtualization management to deploy rogue VMs or even make changes to virtual switches. To mitigate these threats, we recommend that a segregated role be assigned for virtualization management. Also, we recommend that a virtualization management role should be assigned for each network security zone to limit an administrator's span of control. This approach would necessitate multiple virtualization management roles collaborating to make changes spanning multiple network security zones (see VDC reference architecture best practices #6 in Section 3).
- **Infrastructure Management:** Infrastructure management within the VDC includes both physical systems and the networking infrastructure (e.g. routers and switches).
- **Security Management:** Security management encompasses configuring and maintaining boundary security controls and system security controls. This role is slightly different in a VDC. The role has all the same responsibilities as in a traditional data centre. However, it also encompasses the hypervisor layer, as well as virtualization-specific vulnerabilities.

[14] A break-glass policy is an access control policy that can be overridden in times of need. For example, in a data centre the break-glass policy governing least privilege may need to be overridden if a number of administrators were simultaneously afflicted by a virulent strain of flu and unable to perform their duties. However, a break-glass policy is always implemented with compensating audit controls. Consequently, any actions taken under this policy would be heavily audited to ensure that no unauthorized actions take place.

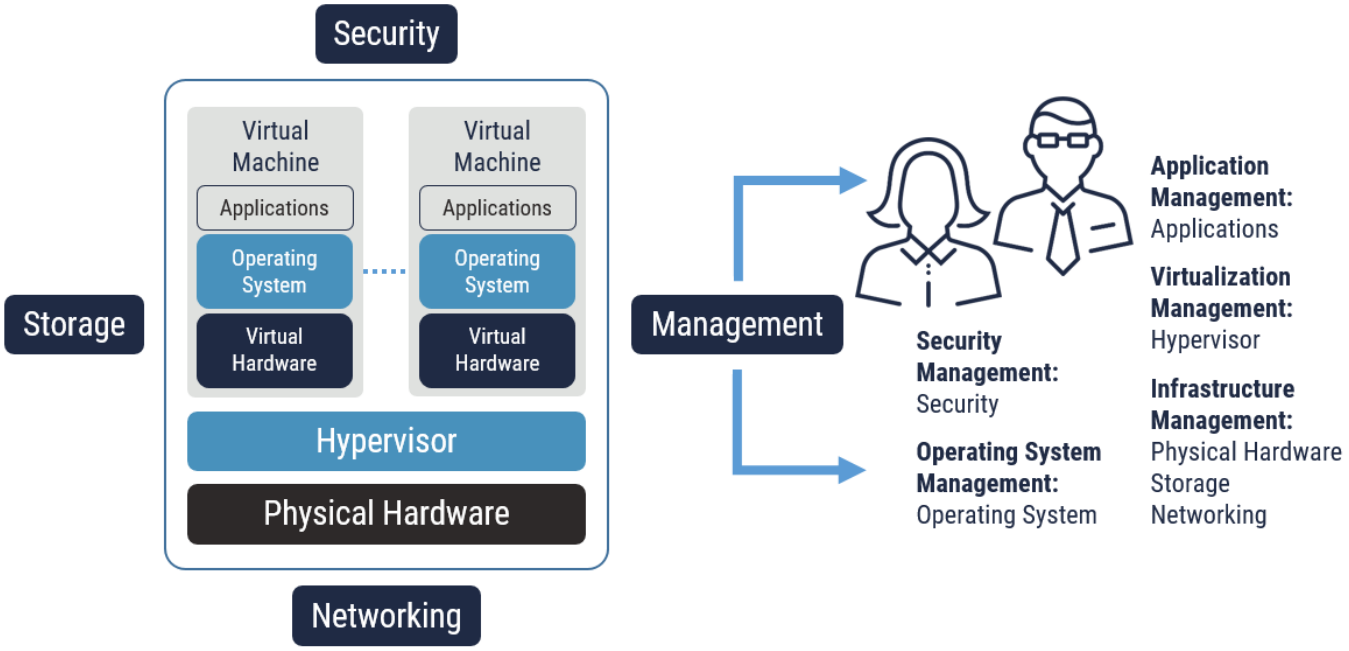

Figure 5 Management Roles in a VDC

2.7 STORAGE

Most SANs in VDCs are either fibre channel (FC) or Internet Small Computer System Interface (iSCSI).

FC, the general name for a set of standards being developed by the American National Standards Institute (ANSI), is a serial link that supports the fibre channel protocol (FCP), as well as higher level protocols such as SCSI. FC requires using a separate FC network. Benefits of FC include increased transfer speed and current market leadership position. A variation, Fibre Channel over Ethernet (FCoE), layers FC over Ethernet. FC SANs support three different types of authentication mechanisms:

- Diffie Hellman Challenge-Handshake Authentication Protocol (DH-CHAP)[15];
- Fibre Channel Authentication Protocol (FCAP)[16]; and
- Fibre Channel Password Authentication Protocol (FCPAP)[17].

iSCSI, an Internet Engineering Task Force (IETF) standard (Request for Comments (RFC) 3270, RFC 3783), allows using the SCSI protocol over a Transmission Control Protocol/Internet Protocol (TCP/IP) network. This is beneficial as it allows an iSCSI SAN to use standard Ethernet technology. Benefits of iSCSI include simplicity and cost. iSCSI SAN supports using the Challenge-Handshake Authentication Protocol (CHAP) and IP Security (IPsec) framework.

[15] DH-CHAP uses a pre-shared secret for authentication.

[16] FCAP uses digital certificates for authentication.

[17] FCPAP uses passwords for authentication.

2.8 SECURITY

Virtualization security is concerned with securing the VDC and is discussed throughout this report. As such, there is no need to address it specifically in this section. However, security virtualization, whereby physical security appliances are virtualized, is being increasingly employed in VDCs. The following section will examine varying approaches to security virtualization in a VDC.

2.8.1 SECURITY VIRTUALIZATION

Security appliances (e.g. firewalls, IDS) can be virtualized in much the same way as traditional servers. By doing so, you are left with what is known as a virtual security appliance (VSA). A VSA is a virtual appliance that ideally consists of a hardened operating system and a single security application. VSAs typically consist of a single security application to be consistent with the principle of isolating security functions from one another (see VDC reference architecture best practices #7 in Section 3). The benefits of security virtualization include cost savings, increased flexibility, and connectivity to the virtualized network infrastructure.

Virtual machine introspection (VMI) is one technique for externally monitoring the runtime state of a system-level virtual machine. For virtual machine introspection, the runtime state can be defined broadly to include processor registers, memory, disk, network, and any other hardware-level events.

VMI is a way to protect a security application from attack by malicious software. The reason for this is that the software interface between a virtual machine and a hypervisor is relatively small. This makes it easier to implement and verify than with the relatively larger interface between an operating system and its applications. Other uses for VMI include software debugging and systems management.[18]

Figure 6 illustrates three diverse approaches to security virtualization.

[18] https://link.springer.com/referenceworkentry/10.1007/978-1-4419-5906-5_647

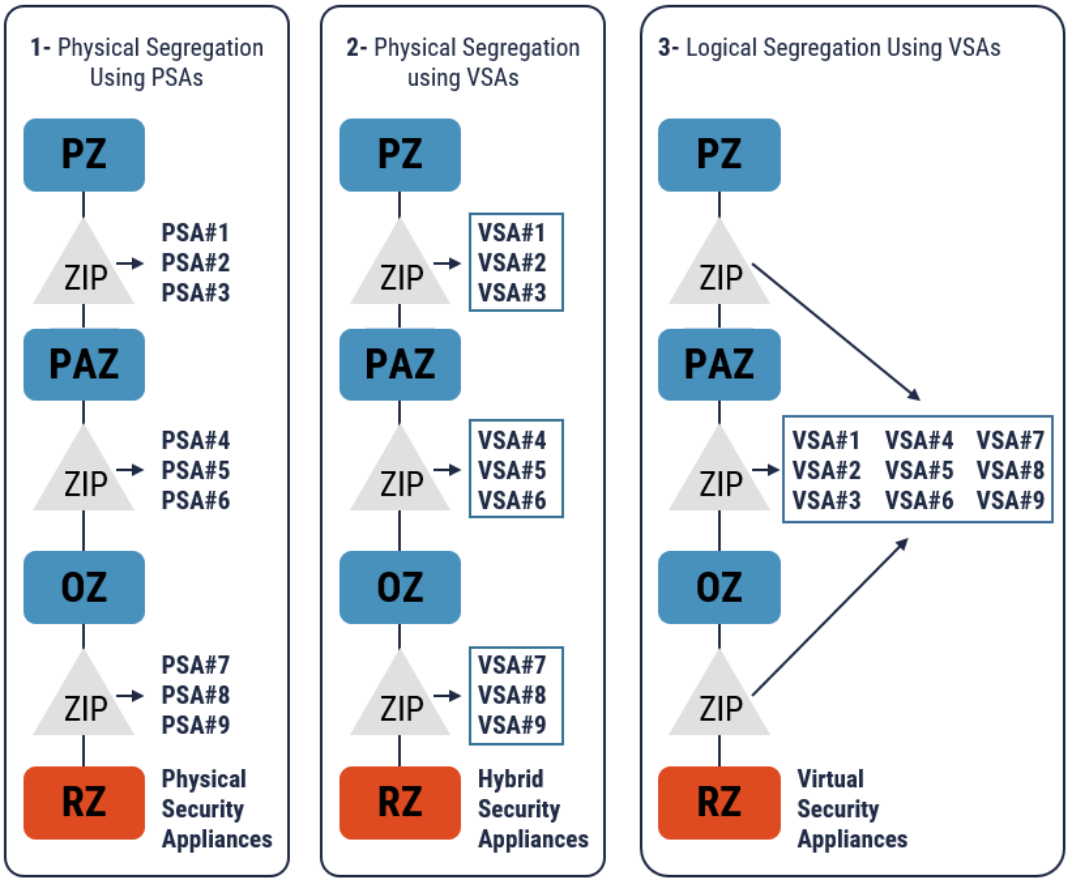

Figure 6 Security Virtualization Options

1. **Segregated boundary with physical security appliances (physical)**. This physical approach is used in most traditional data centres. Not only is each boundary a distinct physical implementation (see VDC reference architecture best practice #8 in Section 3), but each is composed of physical security appliances (PSAs). The security benefit to this approach is that all network traffic traveling between network security zones is forced to travel through a purely physical implementation of the boundary. The disadvantage is that using and managing a relatively large number of PSAs is less flexible.

2. **Segregated boundary with virtual security appliances (hybrid)**. This hybrid approach combines the distinct physical implementation of the boundary (consistent with VDC referenced architecture best practice #8 in Section 3) using VSAs. The advantage of this approach, in terms of security, is that all network traffic traveling between network security zones is forced to travel through a physically segregated implementation of the boundary. The advantage, in terms of flexibility, is that a common boundary device with VSAs can be implemented for use at each boundary.

3. **Consolidated boundary with virtual security appliances (logical)**. This logical approach fully virtualizes the boundary. Not only are the boundaries hosted on the same physical device, but VSAs are leveraged to provide the requisite security services. The security concern is that by mixing network security zones on the same server you dramatically increase the likelihood of human error leading to a misconfiguration that compromises your VDC network security architecture. The advantage of this approach, in terms of flexibility, is that a single physical device with a common set of VSAs can be implemented for all boundaries.

3 VDC REFERENCE ARCHITECTURE BEST PRACTICES

This section provides a list of VDC reference architecture best practices. The best practices, including the nine presented above, have been consolidated below to facilitate review.

1. **Only consolidate/virtualize after due consideration.** With virtualization, the tendency is to want to consolidate/virtualize everything in the data centre. Consolidating/virtualizing, without due consideration for the security implications, is likely to adversely affect the overall security posture of the VDC.
2. **Do not span network security zones with virtualization.** We strongly recommend that VMs from multiple network security zones not be co-located[19] on the same physical hardware in a VDC. In other words, servers should only be used to host VMs belonging to the same network security zone.
3. **Implement a separate MZ for each network security zone**. We strongly recommend that physically separate MZs be implemented for each network security zone in a VDC. There should not be any direct connectivity between MZs.
4. **Implement a separate RZ for each network security zone's virtualization storage**. We strongly recommend that physically separate RZs be implemented for each network security zone's virtualization storage network in a VDC. There should not be any direct connectivity between these RZs.
5. **Use a bare-metal hypervisor**. We strongly recommend that a bare-metal hypervisor be used in a VDC. A bare-metal hypervisor runs directly on the physical hardware. In contrast, a hosted hypervisor runs as an application on top of a host operating system. Hosted hypervisors are inherently less secure as they expose the virtualization layer (hypervisor) and virtualized systems (VMs) to the vulnerabilities in the host operating system. In contrast, a bare-metal hypervisor is purpose built, having a smaller code base, and fewer layers of components, and will inherently have fewer vulnerabilities to exploit.
6. **Segregate management functions**. Within a VDC, management functions should be segregated. Traditional data centre roles should be supplemented with a virtualization management role. This separate virtualization management role should be assigned for each network security zone to prevent changes spanning multiple network security zones.
7. **Isolate security functions**. Security functions need to be isolated from both the systems that they are intended to protect and from other security functions. In a VDC, boundary security functions could be hosted on either a physical or virtual security appliance. Also, PSAs and VSAs should only provide a single security function to effectively isolate that security function from other security functions provided by other security appliances.
8. **Use a segregated implementation of the boundary**. In a VDC, we strongly recommend that a segregated implementation of the boundary be used. This approach will prevent misconfiguration or a security flaw from effectively compromising the VDC security architecture. In a segregated implementation of the boundary, either PSAs or VSAs could be used.

[19] In a VDC, co-location refers to the practice of hosting VMs on the same physical server.

9. **Secure every layer of the VDC**. Since attackers will invariably concentrate their efforts on the least secure point in the VDC, every effort must be made to secure every layer of the VDC.
10. **Use physically separate networks for data, management, storage and live migration networks.** In a VDC, we recommend that physically separate networks be implemented for data, management, storage, and live migration networks rather than VLANs/VSANs for security and performance reasons. Using physically separate networks will ensure that privileged networks are effectively isolated from all other traffic.
11. **Use logical controls to provide fine-grained separation to supplement physically separate network security zones.** It should be noted that while network security zones are used to provide coarse-grained separation of information systems, they do not preclude using more fine-grained separation within the network security zones themselves. The concern is that applications sharing a network security zone can communicate with one another without any of the communication having to transit a security boundary. There may be a requirement for this network communication to transit a security boundary to undergo the appropriate level of scrutiny. This can be accomplished by using logical controls such as VLANs or micro-segmentation available through SDN.
12. **Use Layer 2 security options to mitigate Layer 2 attacks**. Although Layer 2 attacks have been around for years, they are still effective against poorly configured or misconfigured switches. To address this, care needs to be taken to develop secure switch configurations and ensure that they get deployed throughout the VDC.
13. **Use authentication and encryption to protect sensitive network communications.** In addition to physically isolating sensitive networks, authentication and encryption can be used to further protect sensitive network communications. However, it should be noted that using encryption will preclude using network monitoring. For this reason, it may be necessary to terminate encryption at ZIPs to inspect traffic before re-encrypting for further transit.
14. **Employ configuration management and system security to secure the hypervisor**. The hypervisor should be secured in the same way that traditional kernels are secured. Specifically, it should be secured by reducing the attack surface (i.e., disabling/removing unnecessary services), hardening the remaining services (using appropriate hardening guidance), and implementing the most recent patches. Vulnerability scans and compliance checks should be undertaken periodically to ensure the effectiveness of these steps.
15. **Use resource management capabilities to mitigate the vulnerability of resource contention**. Resource management controls, such as reservations and resource limits, should be used to ensure that VMs have sufficient resources and that they cannot consume excessive resources in the event of compromise.
16. **Use introspection, and related Application Programming Interfaces (APIs), to protect VMs.** Introspection, and related APIs, provided by the hypervisor enable the monitoring of VMs, including their memory, processes and networking, without having to install security agents in the VM. Although this approach also has its disadvantages, we recommend using this approach in the VDC to lessen conventional vulnerabilities of VMs and the lack of visibility of virtual networks.
17. **Implement virtualized networking security options provided by SDN solutions, which can place micro-segmentation security policies and firewalls on each independent endpoint of the network.** Carefully configuring these options can result in securing the network in question to a higher level than was previously available with traditional systems.

18. **Do not blend virtualized information systems with different security requirements.** ITSG-33 [10] describes a detailed approach for determining the security control profile, and ultimately the individual security controls that are applicable for a given information system. Part of this process involves assessing the confidentiality, integrity and availability requirements of the information system. We strongly recommend that information systems encapsulated in VMs be hosted on a server with VMs having similar security requirements.
19. **Co-locate virtualized information systems belonging to different organizations on the same physical hardware only after careful consideration.** Co-location can facilitate the compromise of VMs. Also, an attacker can purposely attempt to co-locate a VM on a server to conduct an attack. Consequently, VMs belonging to different organizations should only be co-located on the same physical hardware after careful consideration.
20. **Physically isolate cryptographic services**. Centralized cryptographic services, including Certificate Authorities (CAs), should not be hosted in VMs due to the vulnerability to side-channel attacks. Ideally, cryptographic services should be physically isolated to prevent these types of attacks.
21. **Secure each VM as if it were a physical system.** VMs should be secured in the same way that physical systems are secured using a combination of configuration management and system security. Specifically, they should be secured by reducing the attack surface (e.g. disabling and removing unnecessary services), hardening the operating system and applications (using appropriate hardening guidance), and implementing the most recent patches. Periodically, carry out vulnerability scans and compliance checks to ensure the effectiveness of these steps. Templates should be used to build secure baselines. VMs can then be built from the templates.
22. **Secure access to management interfaces.** Access to the management interface should be secured as much as possible by following the four steps below:
    - Require strong, centrally authenticated, two-factor authentication;
    - Use a secure, dedicated management console;
    - Ensure that management communications are secure; and
    - Disable unused management interfaces and harden any interfaces used.
23. **Implement effective image management, including using templates.** Image management should be implemented to mitigate VM sprawl, including rogue and dormant VMs. This includes limiting the creation of VMs to specific administrative roles, using templates, and the judicious storage and use of snapshots.
24. **Use access controls and/or encryption to mitigate VM theft**. VMs are susceptible to theft at multiple locations within the VDC. A combination of physical access controls, logical access controls, and/or encryption can be used to mitigate this vulnerability. However, these controls must be applied at multiple layers to mitigate this vulnerability throughout the VDC.

4 SUMMARY

To improve consolidation ratios and streamline operations, many organizations are choosing to virtualize their data centres. Doing so requires a proper understanding of the security implications on the overall security posture of the data centre.

Fortunately, VDCs can be made more secure by using several safeguards and best practices. These safeguards and best practices involve addressing each layer of the virtual environment as well as addressing interactions between the various layers. This report provided an overview of a VDC, presented the inherent vulnerabilities in a VDC, and provided recommendations on how to design a secure VDC.

4.1 CONTACTS AND ASSISTANCE

If your organization would like more information on best practices for data centre virtualization, please contact:

Cyber Centre Contact Centre
contact@cyber.gc.ca
613-949-7048 or 1-833-CYBER-88

5 SUPPORTING CONTENT

5.1 LIST OF ABBREVIATIONS

Term	Definition
ANSI	American National Standards Institute
API	Application Programming Interface
BP	Best Practice
CCCS	Canadian Centre for Cyber Security (Cyber Centre)
CHAP	Challenge-Handshake Authentication Protocol
CPU	Central Processing Unit
DH-CHAP	Diffie-Hellman - Challenge Handshake Authentication Protocol
DHCP	Dynamic Host Configuration Protocol
FC	Fibre Channel
FCAP	Fibre Channel Authentication Protocol
FCPAP	Fibre Channel Password Authentication Protocol
FCoE	Fibre Channel over Ethernet
GC	Government of Canada
HTT	Hyper-Threading Technology
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
IPS	Intrusion Prevention Systems
IPsec	Internet Protocol Security
iSCSI	Internet Small Computer System Interface
IT	Information Technology
ITSG	Information Technology Security Guidance
MZ	Management Zone
NAS	Network Attached Storage
NFV	Network Function Virtualization
NIST	National Institute of Standards Technology
OZ	Operations Zone
PAZ	Public Access Zone
PSA	Physical Security Appliance
PZ	Public Zone

Term	Definition
RAM	Random Access Memory
RFC	Request for Comments
RZ	Restricted Zone
SAL	Security Assurance Level
SAN	Storage Area Network
SDN	Software Defined Networking
SMM	System Management Module
SSH	Secure Shell
TBS	Treasury Board of Canada Secretariat
Td5	Threat Deliberate 5
TCP/IP	Transmission Control Protocol/Internet Protocol
TPM	Trusted Platform Module
VDC	Virtualized Data Centre
VLAN	Virtual Local Area Network
VMM	Virtual Machine Monitor
vNIC	Virtual Network Interface Card
VPN	Virtual Private Network
VSA	Virtual Security Appliance
VSAN	Virtual Storage Area Network
ZIP	Zone Interface Points

5.2 GLOSSARY

Term	Definition
Break-Glass Policy	An access control policy that allows for the overriding of standard practice in the time of need or emergency. For example, in a data centre, the policy governing least privilege may need to be overridden using a break-glass provision, if several administrators are simultaneously afflicted by a virulent strain of flu and unable to perform their duties. A break-glass policy is always implemented with compensating audit controls; any actions taken under this policy would be heavily audited to ensure that no unauthorized actions take place.
Commingle	In the context of virtualization, refers to placing VMs with different security requirements and/or VMs from multiple clients on the same hypervisor in order maximize efficiency and consolidation ratios.
Consolidation Ratio	The number of virtual servers operating on each physical host machine.
Diffie Hellman Challenge-Handshake Authentication Protocol (DH-CHAP)	An Internet standard for the authentication of devices connecting to a Fibre Channel switch, using a secure key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication.
Digital Signature Standard (DSS)	The digital signature algorithm (DSA) developed by the US National Security Agency (NSA) to generate a digital signature for the authentication of electronic documents. DSS was put forth by NIST in 1994 and has become the US government standard for authentication of electronic documents. DSS is specified in Federal Information Processing Standard (FIPS) 186.
Fibre Channel Authentication Protocol (FCAP)	An optional authentication mechanism employed between any two devices or entities on a Fibre Channel (FC) network using certificates or optional keys.
Fibre Channel Password Authentication Protocol (FCPAP)	An optional password-based authentication and key exchange protocol which is utilized in Fibre Channel networks. It is used to mutually authenticate Fibre Channel ports to one another other.
Full Virtualization	The first type of x86 virtualization to be developed and the most prominent type to date. It is referred to as full virtualization because the guest operating system and associated applications run in a VM which is fully abstracted from the underlying hardware by the virtualization layer.
Hosted Virtualization	A type of full virtualization where the hypervisor runs on top of a standard operating system and relies on it for all interaction with the underlying hardware.
Internet Small Computer System Interface (iSCSI)	A transport-layer protocol that describes how Small Computer System Interface (SCSI) packets should be transported over a TCP/IP network. iSCSI works on top of the Transport Control Protocol (TCP) and allows the SCSI command to be sent end-to-end over local-area networks (LANs), wide-area networks (WANs) or the Internet.
Introspection	A technique for externally monitoring the runtime state of a virtual machine. without having to install security agents in the virtual machine
Operating System Virtualization	A type of virtualization, sometimes referred to as shared-kernel virtualization, which differs from full virtualization in that applications are hosted in virtual environments on a common operating system rather than in VMs with a separately installed operating system.
System Management Mode (SMM)	An operating mode of x86 central processor units (CPUs) in which all normal execution, including the operating system, is suspended. An alternate software system, which usually resides in the computer's firmware, or a hardware-assisted debugger, is then executed with high privileges.

Term	Definition
Virtualized Data Centre (VDC)	A data centre employing virtualization to a considerable degree in order to abstract compute, network, storage, security, and management resources from the underlying hardware.

5.3 REFERENCES

Number	Reference
1	Treasury Board of Canada Secretariat. *Policy on Government Security*. 1 July 2009.
2	Treasury Board of Canada Secretariat. *Directive on Security Management.* 1 July 2019.
3	Treasury Board of Canada Secretariat. *Directive on Privacy Practices*. 6 May 2014.
4	Department of Justice. *Financial Administration Act.* 12 March 2018.
5	Treasury Board of Canada Secretariat. *Guideline on Acceptable Network and Device Use.* 30 May 2016.
6	Treasury Board of Canada Secretariat. *Policy on Management of Materiel.* 26 June 2006.
7	Department of Justice Canada. *Privacy Act.* (R.S.C., 1985, c. P-21).
8	Department of Justice Canada. *Personal Information Protection and Electronic Documents Act.* (S.C. 2000, c.5).
9	Department of Justice Canada. *Access to Information Act.* (R.S.C., 1985, c. A-1).
10	Canadian Centre for Cyber Security. *ITSG-33 IT Security Risk Management: A Lifecycle Approach*, December 2014.
11	Canadian Centre for Cyber Security. *ITSP.80.022 Baseline Security Requirements for Network Security Zones*.
12	Canadian Centre for Cyber Security. *ITSG-38 Network Security Zoning – Design Considerations for Placement of Services within Zones.* 1 May 2009.

5.4 BIBLIOGRAPHY

5.4.1 CYBER CENTRE PUBLICATIONS

ITSG-33 – IT Security Risk Management: A Lifecycle Approach. November 2012.
ITSP.30.031 V3 User Authentication Guidance for Information Technology Systems. April 2018.
ITSP.40.062 Guidance on Securely Configuring Network Protocols. August 2007
ITSP.40.111 Cryptographic Algorithms for Unclassified, Protected A, and Protected B Information. August 2016.
ITSG-22 Baseline Security Requirements for Network Security Zones, June 2007
TRA-1 - Harmonized Threat & Risk Assessment (TRA) Methodology. October 2007.
TSCG-01\E: – Technology Supply Chain Guidelines – Contracting Clauses for Telecommunications Equipment and Services. October 2010

5.4.2 SSC PUBLICATIONS

Isolation in a Software Defined Data Centre, Shared Services Canada, January 2015.

5.4.3 OTHER PUBLICATIONS

A. Desnos, E. Filiol and I. Lefou. Detecting (and creating!) a HVM rootkit (aka BluePill-like), September 2009.
A. Sawani, SubVirt: Implementing Malware with Virtual Machines, 2006.
A. Tereshkin & R. Wojtczuk, Introducing Ring -3 Rootkits, Invisible Things Lab, July 2009.
B. Danev, et al. Enabling Secure VM-vTPM Migration in Private Clouds, Department of Computer Science, ETH Zurich, Switzerland, 2011.
B. Dolan-Gavitt., et al. Virtuoso: Narrowing the Semantic Gap in Virtual Machine Introspection, School of Computer Science, Georgia Institute of Technology, 2011.
D. Dai Zovi, Hardware Virtualization Rootkits, Matasano, 2006.
I. Kyte, P. Zavarsky, D. Lindskog & R. Ruhl. Detection of Hardware Virtualization Based Rootkits by Performance Benchmarking, Concordia University College of Alberta, 2012.
J. Rutkowska & A. Tereshkin, IsGameOver() Anyone?, Version 1.01, Invisible Things Lab, August 2007.
J. Rutkowska, Introducing Blue Pill, Invisible Things Lab, 22 June 2006.
J. Stewart, Practical Considerations in Virtual Machine Covert Channels, Thesis, East Stroudsburg University of Pennsylvania, 6 May 2011.
J. Xiao, Z. Xu, H. Huang & H. Wang, POSTER: A Covert Channel Construction in a Virtualized Environment, 2012.
K. Figueroa, M. Figueroa and A. Williams, VLAN Layer 2 Attacks: Their Relevance and Their Kryptonite, Defcon 16, 2008.
K. Kortchinsky, CLOUDBURST, A VMware Guest to Host Escape Story, Immunity Inc., 2009.
K. Nance, B. Hay & M. Bishop, Virtual Machine Introspection – Observation or Interference? IEEE Security & Privacy, 2008.
K. Suzaki, K. Iijima, T. Yagi and C. Artho, Memory Deduplication as a Threat to the Guest OS, National Institute of Advanced Industrial Science and Technology, 2011.

M. Green, Attack of the week: Cross-VM side-channel attacks, 26 October 2012.
M. Myers & S. Youndt, An Introduction to Hardware-assisted Virtual Machine Rootkits, Crucial Security, 7 August 2007.
Ptacek, Goldsmith & Lawson. Don't Tell Joanna, the Virtualized Rootkit is Dead, Matasano Security, 2007.
R. Wojtczuk and J. Rutkowska, Attacking Intel Trusted Execution Technology, Invisible Things Lab, February 2009.
R. Wojtczuk and J. Rutkowska, Attacking Intel TXT via SINIT Code Execution Hijacking, Invisible Things Lab, November 2011.
R. Wojtczuk, Subverting the Xen Hypervisor, Invisible Things Lab, 2008.
S. Bahram, et al., DKSM: Subverting Virtual Machine Introspection for Fun and Profit, 2010.
S. Convery, Hacking Layer 2: Fun with Ethernet Switches, Cisco Systems, 2002.
S. Embleton, S. Sparks and C. Zou, SMM Rootkits: A New Breed of OS Independent Malware, University of Central Florida, 22 September 2008.
Separation of Duties, NIST, 1995
Special Publication 800-125 - Guide to Security for Full Virtualization Technologies, NIST, January 2010.
SYSRET 64-bit Operating System Privilege Escalation Vulnerability on Intel CPU Hardware, CERT, 12 June 2012.
T. Garfinkel & M. Rosenblum, A Virtual Machine Introspection Based Architecture for Intrusion Detection, Computer Science Department, Stanford University, 2003.
T. Garfinkel & M. Rosenblum. When Virtual is Harder than Real: Security Challenges in Virtual Machine Based Computing Environments, Stanford University Department of Computer Science, 2005.
T. Ristenpart, et al. Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds, 2011.
Trusted Platform Module Summary, Trusted Computing Group.
V. Scarlata, et al., TPM Virtualization: Building a General Framework, Intel Corporation, 2008.
VMware vSphere 6.5: Install, Configure, Manage, VMware Education Services, 2017.
vTPM: Virtualizing the Trusted Platform Module, IBM Research Report, IBM, 14 February 2006.
Y. Bulygin & D. Samyde, Chipset Based Detection and Removal of Virtualization Malware a.k.a. DeepWatch, Intel Corporation, 2008.
Y. Xu, et al., An Exploration of L2 Cache Covert Channels in Virtualized Environments, 2011.
Y. Zhang, et al., Cross-VM Side Channels and Their Use to Extract Private Keys, 2012.
Z. Wu, Z. Xu & H. Wang, Whispers in the Hyper-space: High-speed Covert Channel Attacks in the Cloud, The College of William and Mary, 2012.