

Microsoft Windows Common Criteria Evaluation

Microsoft Windows Server 2008 R2 Hyper-V

Security Target

Document Information	
Version Number	2.6
Updated On	Thursday, January 12, 2012

This is a preliminary document and may be changed substantially prior to final commercial release of the software described herein.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. This work is licensed under the Creative Commons Attribution-NoDerivs-NonCommercial License (which allows redistribution of the work). To view a copy of this license, visit <http://creativecommons.org/licenses/by-nd-nc/1.0/> or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

The example companies, organizations, products, people and events depicted herein are fictitious. No association with any real company, organization, product, person or event is intended or should be inferred.

© 2012 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Visual Basic, Visual Studio, Windows, the Windows logo, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

TABLE OF CONTENTS

<u>1</u>	<u>ST INTRODUCTION</u>	<u>6</u>
1.1	SECURITY TARGET (ST) REFERENCE	6
1.2	TOE REFERENCE	6
1.3	TOE OVERVIEW.....	6
1.3.1	SUMMARY OF THE TOE SECURITY FUNCTIONS.....	6
1.3.2	TOE USAGE.....	7
1.3.3	DEFINITIONS.....	7
1.4	TOE DESCRIPTION	9
1.4.1	INTENDED METHOD OF USE	12
1.4.2	SUMMARY OF SECURITY FUNCTIONALITY	12
1.5	HARDWARE REQUIREMENTS AND SUPPORTED GUEST OPERATING SYSTEMS	13
1.5.1	MEMORY.....	14
1.5.2	PROCESSORS	14
1.5.3	NETWORKING.....	14
1.5.4	STORAGE	15
1.5.5	OTHER HARDWARE COMPONENTS.....	16
1.5.6	SUPPORTED GUEST OPERATING SYSTEMS.....	16
1.5.7	INTEGRATION SERVICES (ENLIGHTENMENTS).....	18
1.6	TOE GUIDANCE	20
1.7	REFERENCES.....	20
<u>2</u>	<u>CONFORMANCE CLAIM</u>	<u>21</u>
<u>3</u>	<u>SECURITY PROBLEM DEFINITION</u>	<u>22</u>
3.1	INTRODUCTION.....	22
3.2	THREATS	22
3.3	SECURITY POLICIES	23
3.4	ASSUMPTIONS.....	24
<u>4</u>	<u>SECURITY OBJECTIVES</u>	<u>27</u>
4.1	SECURITY OBJECTIVES FOR THE TOE	27
4.2	SECURITY OBJECTIVES FOR THE TOE ENVIRONMENT	28
4.3	SECURITY OBJECTIVES RATIONALE	31
4.3.1	COMPLETE COVERAGE: THREATS AND ORGANIZATIONAL SECURITY POLICIES.....	31

5	<u>EXTENDED COMPONENTS DEFINITION.....</u>	<u>38</u>
6	<u>SECURITY REQUIREMENTS.....</u>	<u>40</u>
6.1	SECURITY FUNCTIONAL REQUIREMENTS	40
6.1.1	SECURITY FUNCTIONAL REQUIREMENTS IMPLEMENTED BY HYPER-V R2	43
6.1.1.1	Security Audit (Class FAU)	43
6.1.1.2	User Data Protection (Class FDP).....	44
6.1.1.3	Identification and Authentication (Class FIA)	46
6.1.1.4	Security Management (Class FMT).....	47
6.1.1.5	Protection of the TSF (Class FPT).....	49
6.1.1.6	Resource Utilisation (Class FRU).....	49
6.2	SECURITY ASSURANCE REQUIREMENTS	50
6.3	SECURITY REQUIREMENTS RATIONALE	51
6.3.1	INTERNAL CONSISTENCY OF REQUIREMENTS	51
6.3.2	COMPLETE COVERAGE: SECURITY OBJECTIVES	52
6.3.3	SECURITY REQUIREMENTS COVERAGE	54
6.3.4	SECURITY REQUIREMENTS DEPENDENCY ANALYSIS.....	55
6.3.5	OPERATIONS PERFORMED ON THE SFRS.....	57
6.4	TOE SUMMARY SPECIFICATION RATIONALE.....	65
6.4.1	SECURITY FUNCTIONS JUSTIFICATION	65
6.4.2	ASSURANCE MEASURES JUSTIFICATION.....	65
7	<u>TOE SUMMARY SPECIFICATION</u>	<u>66</u>
7.1	ARCHITECTURE OF THE TOE.....	66
7.1.1	WINDOWS HYPERVISOR	68
7.1.2	VMBUS.....	72
7.1.3	OPERATING SYSTEM ENLIGHTENMENTS (OUTSIDE OF THE TSF).....	72
7.1.4	VIRTUAL SERVICE CLIENTS (OUTSIDE OF THE TSF).....	73
7.1.5	VIRTUALIZATION SERVICE PROVIDER.....	73
7.1.6	VM WORKER PROCESSES	74
7.1.7	WMI PROVIDER	74
7.1.8	AZMAN AS THE FRAMEWORK FOR ACCESS CONTROL TO MANAGEMENT OPERATIONS OF HYPER-V R2	77
7.2	SECURITY FUNCTIONALITY PROVIDED BY HYPER-V R2.....	77
7.2.1	AUDIT FUNCTIONALITY	78
7.2.2	USER AND TSF DATA PROTECTION FUNCTIONALITY	78
7.2.3	IDENTIFICATION FUNCTIONALITY	79
7.2.4	SECURITY MANAGEMENT FUNCTIONALITY	79
7.2.5	TSF PROTECTION FUNCTIONALITY.....	81
7.2.6	TSF PROTECTION DURING LIVE MIGRATION.....	81

7.2.7 RESOURCE UTILIZATION FUNCTIONALITY..... 82

1 ST Introduction

This is version 2.6 of the Security Target for Microsoft Windows Server 2008 R2 Hyper-V RTM (called “Hyper-V” in this document). This document is based on the Security Target used for the previous evaluation of Microsoft Windows Server 2008 Hyper-V RTM and has been modified to represent the changes and enhancements made for Release 2.

1.1 Security Target (ST) reference

Title: Microsoft Windows Server Core 2008 R2: Hyper-V Server Role, Hyper-V Security Target

Version: 2.6

Keywords: Virtualization, hypervisor

This document is the Security Target for the Common Criteria evaluation of the Microsoft Server 2008 R2 Hyper-V RTM virtualization product. It is conformant to the Common Criteria for Information Technology Security Evaluation Version 3.1 Release 3[CC].

1.2 TOE Reference

The TOE is the Microsoft Server 2008 R2 Hyper-V (referenced as Hyper-V R2 in this document). This product is a hypervisor or virtual machine monitor for Intel processors with Intel’s VT-x support and for AMD processors with AMD VT support.

1.3 TOE Overview

The Target of Evaluation (TOE) is the Microsoft Hyper-V R2 virtualization part of the Server 2008 R2 product. Hyper-V R2 allows the definition of partitions that have separate address spaces where they can load an operating system and applications operating on top of this operating system. The TOE consists of the Hyper-V specific components of Microsoft Server 2008 R2 Server Core executing in the root partition with the Hyper-V R2 hypervisor running in a hypervisor configuration.

An operating system within such a partition has access to virtualized peripheral devices where access to those devices is controlled by Hyper-V R2. An operating system may either access devices using the same I/O related instructions as on a real system or it may use a specific interface offered by Hyper-V R2 (called the VMBus) to communicate with Hyper-V R2 for access to peripheral devices. In the first case the operating system can only access the devices virtualized by Hyper-V R2. When using the VMBus defined interface, an operating system in a guest partition needs to install “enlightenments” that set up the VMBus communication and use the “synthetic” devices accessible via VMBus. Note that the “enlightenments” within a guest operating system is part of the TOE, but not part of the TSF.

1.3.1 Summary of the TOE security functions

The TOE offers separation of partitions, controls access of partitions to resources like virtual hard disks or virtual network adapter, controls the management of the TOE and enforces a privilege-based management policy, allows auditing of security critical events, controls access of administrative users in the root partition to TOE management functions, and enforces quota for CPU time for partitions. The TOE also offers live migration of partitions from one instance of the TOE to another instance, provided

the underlying platforms of both instances of the TOE satisfy the compatibility criteria required for live migration. Live migration will neither allow a partition to increase its privileges nor undermine the security of any of the two instances of the TOE involved in the migration.

1.3.2 TOE usage

The TOE can be used to consolidate several physical servers based on Intel's x86 architecture onto one machine. The TOE allows the definition of so called partitions. Each instantiation of the TOE has one dedicated partition, called the root partition, and a variable number of so called "guest partitions". Resources accesses by guest partitions are virtualized by the TOE, i. e. the TOE performs a "translation" of the virtual resource accesses by a guest partition onto the real resources available to the TOE. Such resources include virtual CPUs, main memory, virtual hard disks, virtual network adapters, virtual CD/DVD drives or floppy disk drives as well as virtual video adapter and virtual mouse and keyboard. The root partition is used for support of the resource virtualization and for TOE management activities.

Each guest partition can take over the tasks of one physical server. Each guest will have its own operating system installed and is restricted by the TOE to the use of the resources that are assigned to the partition. The assignment of resources to partitions is performed by administrative roles defined in the root partition. The TOE allows separating each guest partition from others with a comparable degree as if they were executing on separate physical servers.

See the TOE description section in this chapter for more detailed information about the TOE structure and the security functionality implemented by the TOE.

1.3.3 Definitions

This section defines some of the TOE specific terms used throughout this document.

Hypervisor and Partitions

A hypervisor is a layer of software that sits just above the hardware and beneath one or more operating systems. Its primary job is to provide isolated execution environments called partitions. Each partition is provided with its own set of (physical or virtual) hardware resources (memory, devices, CPU cycles). The hypervisor is responsible for controlling and arbitrating access to the underlying hardware where necessary.

Guests

Software running within a non-root partition is referred to as a guest. A guest might consist of a full-featured operating system like Windows Vista or a small, special-purpose kernel. The hypervisor is "guest-agnostic".

Specialized Partitions

In most respects, all partitions are equal. However, some partitions may be granted special privileges or assigned specialized functions.

In the evaluated version of Hyper-V R2 there is only partition that has special privileges, which is called the root partition. The root partition acts as the default owner of all hardware resources. It is also typically in charge of power management, plug and play, and hardware failure events. The root partition is also responsible for creating and managing other partitions and assigning hardware resources. In some respects, it acts like the service processor on a machine with hardware partitioning facilities. To start Hyper-V R2 first the Server 2008 R2 instance that is supposed to run in the root partition is started. In the evaluated configuration, this is the Server 2008 R2 Server Core system. This then starts the Hypervisor, which creates the root partition and “moves” the running instance of Server 2008 R2 into the root partition and assigns all devices to this partition.

Virtualization and Emulation

The hypervisor provides support for hardware virtualization. Virtualization provides multiple logical instances of CPUs and other hardware resources. These logical instances are mapped onto physical hardware resources using a variety of techniques. One such technique is emulation, the simulation of a processor or device using software. While the hypervisor facilitates processor and device emulation, the architecture attempts to provide or facilitate alternatives to emulation for performance reasons.

Legacy Guests, Worker Processes and Enlightenment

The hypervisor provides support for legacy guests. A legacy guest is an operating system that has no knowledge of the fact that it is running within a virtualized environment. Legacy guests require substantial infrastructure including a system BIOS and a wide variety of emulated devices. This infrastructure is not provided directly by the hypervisor. A separate piece of code called a worker process provides this infrastructure to legacy guests. Worker processes are part of the root partition and receive specific intercepts (i.e. notifications that specific events have occurred within a guest). This support for guest-mode intercept handlers provides added flexibility and reduced complexity within the hypervisor. Note that a worker process exists for each guest partition.

Device emulation provides broad compatibility, but it results in poor performance. Other operational assumptions made by legacy guests also add to virtualization overhead. This virtualization performance overhead can be mitigated by enlightening a legacy operating system. An enlightened guest has knowledge of the fact that it is running within a virtualized environment and changes its behavior accordingly. Various degrees of enlightenment are possible. For example, a guest might use a specialized block device driver to talk to an idealized virtual disk using a fast communication channel between itself and the root partition. More extensive enlightenment involves modifying or extending a guest hardware abstraction layer (HAL) so it talks to an idealized “synthetic” interrupt controller or access model specific register (MSR) that do not exist in real hardware.

Snapshots

A snapshot is a collection of data about a partition and its current state that allows restarting the partition in this state. A Hyper-V R2 snapshot therefore includes all of the information and data that is required to roll back the status of a partition to the state when the snapshot was taken. Information that is collected when taken a snapshot include:

- Partition configuration settings (the contents of the .vmc file)
- Virtual network settings
- The current state of all virtual hard disks (VHDs) that are attached to the partition
- State information for the partition

Live Migration

The live migration function allows moving a running partition from one instance of the TOE (the “source” to another instance of the TOE (the “destination”). The TOE will verify that both instances of the TOE satisfy defined compatibility criteria which ensure that software executed in the partition see the same operational environment from the underlying platform, which consists of the hardware visible to the partition and the Hyper-V R2 interfaces and functions. Live migration requires all virtual hard disks used by the partition to be migrated to reside on cluster shared volumes. Live migration is not possible for partitions that use virtual hard disks on media local to the “source” instance of the TOE.

Live migration ensures that the security properties of a partition are maintained throughout the migration process and on the destination system regardless of the software executed in the partition being migrated.

1.4 TOE description

Hyper-V R2 consists of a small hypervisor layer that uses the virtualization support functions of the Intel (Intel-VT) and AMD (AMD-V) to define and separate guest partitions, the Server 2008 R2 Server Core installation in the root partition and the “enlightenments” that can be used in guest partitions to use the virtualization functions more efficiently. For more information about the hardware support for virtualization provided by Intel and AMD, please read [INTEL-VT] and [AMD-V]. Note that the “enlightenments” are not part of the TSF portion of the TOE. The TOE also includes the guidance documentation required to securely install, configure and operate Hyper-V R2 with the Server 2008 R2 core in the root partition.

A partition is defined by assigning a memory region (called “guest physical memory” GPA), a number of virtual processors and a set of virtualized devices and resources. Hyper-V R2 differentiates between the single “root” partition and multiple “guest” partitions. The root partition is part of the trusted components of the TOE since it is responsible for the management of the partitions as well as the virtualization of devices and resources other than physical memory pages and virtual processors. All real devices except memory pages and processors are allocated to the root partition and all the mapping

between the virtualized devices accessible by a guest partition and the real device that backs the virtual device is done by the root partition. The root partition runs the Server Core configuration of Server 2008 R2 with the configuration defined in [Hyper-V R2_ECG].

When the administrator starts a partition, first the worker process for that partition is started. This worker process reads the configuration data of the partition and based on this data it defines the events within the partition that it wants to intercept. Additional events to be intercepted are defined by the root partition when the hypervisor layer is instructed by the root partition to start the new partition. The hypervisor itself defines additional events it wants to intercept and handle. Whenever such an event happens, a trap into the hypervisor is generated by the processor hardware and the hypervisor either handles the event itself or passes it on to the root partition to be handled there. For events where the hypervisor needs the support of the root partition, the hypervisor generates a message to the root partition, which then performs the required actions like mapping the access to a virtualized device to the real device that backs the virtual device. When this has been done, the root partition sends a message back to the hypervisor defining how to respond to the partition. The hypervisor then reflects this response back to the partition.

The hypervisor also provides a set of functions to the partitions they can use. One of those functions can be used by a partition to determine it is running on Hyper-V R2 and to determine the version of Hyper-V R2 that is installed. The software within the partition (usually a server operating system) can then install software that makes use of specific hypervisor functions that are not available if the operating system would execute directly on the system's hardware. This software, called "enlightenments" can then optimize the performance of the operating system and use the flexibility provided by the hypervisor.

Beside the hypervisor calls that allow the use of specific functions of the hypervisor layer, this also includes a direct and fast communication mechanism between individual guest partitions and the root partition. This communication mechanism, called VMBus, allows a guest partition to "shortcut" requests for access to a virtualized device allocated to the partition and sends this request directly to the root partition without involving the hypervisor layer and without undergoing the burden of virtualizing every hardware interface of this device. Using VMBus the root partition can also return any response to the request directly to the partition. The software within a partition that implements the "shortcuts" for accessing virtualized devices within a guest partition is called "Virtualization Service Client" or VSC. The corresponding software component within the root partition that handles the requests transferred via VMBus is called a "Virtualization Service Provider" or VSP.

Hyper-V R2 ensures that the physical memory allocated to guest partitions do not overlap with memory allocated to another guest partition, providing each guest partition with its own, isolated guest physical address space. Real processors within a system are assigned by the hypervisor to guest partitions on a time-slice basis. There is no guarantee that the same virtual processor is always backed by the same real processor and real processors are not bound to a specific guest partition.

Hyper-V R2 just ensures that when a real processor is assigned to a virtual processor, which then is executing on behalf of code within a guest partition, no information is passed via the processor's state or register from its previous assignment to a virtual processor of another guest partition.

Hyper-V R2 therefore enforces the separation of guest partitions as well as an access control policy between partition and virtualized devices. In addition Hyper-V R2 allows assigning maximum quota for CPU time and memory for guest partitions to avoid that a single partition is able to use all of those resources.

The management of Hyper-V R2 and its configuration is performed within the root partition. The Server 2008 R2 components within the root partition provide the generic Server 2008 R2 functionality for user authentication, auditing and audit management, file access control and process isolation and management. In the evaluated configuration the root partition is solely used for the management of Hyper-V R2, the virtualization of devices and passing remote users through to guest partitions using the RDP protocol. No untrusted program is allowed to be installed on the Server 2008 R2 instance within the root partition. An organization that wants to have generic applications operating on Server 2008 R2 needs to install Server 2008 R2 in one of the guest partitions and install those programs on this instantiation of Server 2008 R2. The generic security functionality of Windows Server 2008 R2 mentioned above has been evaluated separately in the evaluation of Windows Server 2008 R2. Readers interested in the security functionality and security properties provided by Windows Server 2008 R2 are referred to the Windows Server 2008 R2 Security Target.

Hyper-V R2 is distributed as part of the Windows Server software editions (Windows Server 2008 R2 Datacenter, Windows Server 2008 R2 Enterprise, and Windows Server 2008 R2 Standard). Those include the Hyper-V R2 role, with the x64 version of the remote management tools, and integration services for the supported versions of the Windows operating system. Integration services for the supported versions of Linux distributions are distributed through the Microsoft Connect Web site and are identified as Linux Integration Components for Microsoft Hyper-V R2.

The Hyper-V R2 management tools are available separately to allow remote management of a server running Hyper-V R2. Packages are available to install the tools Windows Server 2008 R2.

The guidance documentation provided for Hyper-V R2 consists of the interface descriptions for the Hypervisor Calls, the description of the management interfaces (WMI interfaces for virtualization) and the description of the AzMan interfaces. In addition there is guidance for the secure installation, configuration and start-up of the TOE. Specific guidance for non-administrative users is not required, since the only non-administrative 'users' are the child partitions and remote users that are passed through to a guest partition they are allowed to access. Neither of those does require specific guidance other than the description of the programming interfaces that a partition can use resp. the list of VMs they may attempt to connect to.

Administrative users or users that are allowed to connect to a partition are identified and authenticated by the TOE environment (the Server 2008 R2 instance in the root partition) using either locally defined and managed credentials or through the use of an Active Directory service. The protection and correct

administration of the server operating the Active Directory service as well as the protection of the communication link to this server has to be ensured by the TOE environment.

1.4.1 Intended method of use

Hyper-V R2 is used in conjunction with Windows Server 2008 R2 (as the operating system in the root partition) to provide a computing environment that allows creating partitions within a single computer system where each partition can load and operate an operating system and its applications very much like the operating system and its application would execute directly on real hardware. Each guest partition gets virtualized hardware resources (memory, processors, storage, network devices, interrupt controller etc.) assigned where the virtualization of those resources is performed by the TOE. This allows the operator of a real server system to have multiple virtual servers installed on one server hardware system and assign each virtual server the resources it needs. The management functions of Hyper-V R2 allow an authorized administrator to modify this allocation of resources to virtualized servers in order to optimize the use of resources and react to the different needs for resources by the virtualized servers.

Virtualized servers in the partitions are separated from each other such that they can only communicate using virtualized network functions or shared storage devices. This is the same way they would also communicate with each other if they were installed on separate hardware systems and the intention of the virtualization provided by Hyper-V R2 is to provide a comparable degree of separation between virtual servers as there is between real servers on different hardware while being able to optimize the use of resources and ease the management of multiple servers.

For security reasons the evaluated configuration does not allow the installation of untrusted software in the root partition.

1.4.2 Summary of security functionality

Hyper-V R2 provides the following primary security functionality:

- Access control between partitions and virtualized resources
- Auditing of security critical events detected by Hyper-V R2 (where the audit records are then submitted to the Windows Server 2008 R2 audit subsystem in the root partition and stored together with audit record created by this instance of Windows Server 2008 R2)
- Object reuse for all resources managed by Hyper-V R2
- Management of the Hyper-V R2 configuration including the configuration of the partitions
- Maximum quota for defined resources assigned to partitions (CPU time, memory, disk storage)
- Live migration of partitions from one instance of the TOE to another instance of the TOE ensuring the integrity and consistency of the partition being migrated

In addition Hyper-V R2 provides the following architectural properties:

- TSF protection against tampering from guest partitions and network devices

- Separation between the guest partitions
- Reference mediation for access of guest partitions to protected resources (including virtualized devices)
- Non-bypassability of the reference mediation
- Maintaining the separation mechanism provided by the underlying hardware when virtualizing resources and devices or responding to hypervisor calls for a guest partition.

The Windows Server 2008 R2 instance in the root partition provides the following security functionality which is used by Hyper-V:

- Identification and authentication of administrative users and users that request to be passed through to a guest partition
- Management and protection of the audit trail
- Access control of administrative users to Windows Server 2008 R2 management objects
- Access control to files and devices used
- Management of users and access rights (to Windows Server 2008 R2 objects)

Those security functions of Server 2008 R2 have been evaluated in a separate evaluation of this product. Therefore in the evaluation of Hyper-V R2 those security functions provided by Server 2008 R2 are handled as security functions provided by another trusted IT product.

1.5 Hardware Requirements and supported guest operating systems

Note: the requirements defined in this section apply for the version of Hyper-V R2 subject to this evaluation. Future versions of Hyper-V R2 may have different hardware requirements and may support other guest operating systems.

To install and use the Hyper-V R2 role, the following is required:

- A x64-based processor. Hyper-V R2 is available in 64-bit editions of Windows Server 2008 R2 — specifically, the 64-bit editions of Windows Server 2008 R2 Standard, Windows Server 2008 R2 Enterprise, and Windows Server 2008 R2 Datacenter.

Hyper-V R2 is not available for 32-bit (x86) editions or Windows Server 2008 R2 for Itanium-Based Systems. However, the Hyper-V R2 management tools are available for 32-bit editions.

- Hardware-assisted virtualization. This is available in processors that include a virtualization option—specifically processors with Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) technology.

- Hardware-enforced Data Execution Prevention (DEP) must be available and enabled. Specifically, you must enable Intel XD bit (execute disable bit) or AMD NX bit (no execute bit).

Hardware support for second-level address translation is not required, but will be used if the underlying platform provides such support.

1.5.1 Memory

The maximum amount of memory that can be used is determined by the operating system, as follows:

- For Windows Server 2008 R2 Enterprise and Windows Server 2008 R2 Datacenter, the physical computer can be configured with up to 1 TB of physical memory, and partitions that run either of those editions can be configured with up to 1 TB of memory per partition.
- For Windows Server 2008 R2 Standard, the physical computer can be configured with up to 32 GB of physical memory, and partitions that run either of those editions can be configured with up to 31 GB of memory per partition.

1.5.2 Processors

Hyper-V R2 is supported on physical computers with up to 64 logical processors. A logical processor can be a core processor or a processor using hyper-threading technology. One can configure up to 4 virtual processors on a partition. However, the number of processors supported by a guest operating system might be lower.

The following are some examples of supported systems and the number of logical processors they provide:

- A single-processor/dual-core system provides 2 logical processors.
- A single-processor/quad-core system provides 4 logical processors.
- A dual-processor/dual-core system provides 4 logical processors.
- A dual-processor/quad-core system provides 8 logical processors.
- A quad-processor/dual-core system provides 8 logical processors.
- A quad-processor/dual-core, hyper-threaded system provides 16 logical processors.
- A quad-processor/quad-core system provides 16 logical processors.

1.5.3 Networking

Hyper-V R2 provides the following networking support:

- Each partition can be configured with up to 12 virtual network adapters — 8 can be the “network adapter” type and 4 can be the “legacy network adapter” type. The network adapter type provides better performance and requires a dedicated driver that is included in the integration services packages.

- Each virtual network adapter can be configured with either a static or dynamic MAC address.
- Each virtual network adapter offers integrated virtual local area network (VLAN) support and can be assigned a unique VLAN channel.
- Hyper-V R2 supports an unlimited number of virtual networks with an unlimited number of partitions per virtual network.

Note: Hyper-V R2 cannot connect a virtual network to a wireless network adapter. As a result, wireless networking capabilities can not be provided to partitions.

1.5.4 Storage

Hyper-V R2 supports a variety of physical storage options (using the device drivers of the Windows Server 2008 R2 instance in the root partition):

- Direct-attached storage: Serial Advanced Technology Attachment (SATA), external Serial Advanced Technology Attachment (eSATA), Parallel Advanced Technology Attachment (PATA), Serial Attached SCSI (SAS), SCSI, USB, and Firewire.
- Storage area networks (SANs): Internet SCSI (iSCSI), Fibre Channel, and SAS technologies can be used.
- Network-attached storage
- Storage in cluster-shared volumes (which is required for all storage used by partitions that are subject to live migration)

Hyper-V R2 supports configuring a partition to use the following types of virtual storage.

- Virtual hard disks of up to 2040 GB. Hyper-V R2 can use fixed virtual hard disks, dynamically expanding virtual hard disks, and differencing disks.
- Virtual IDE devices. Each partition supports up to 4 IDE devices. The startup disk (sometimes referred to as the boot disk) must be attached to one of the IDE devices. The startup disk can be either a virtual hard disk or a physical disk.
- Virtual SCSI devices. Each partition supports up to 4 virtual SCSI controllers, and each controller supports up to 64 disks. This means that each partition can be configured with as many as 256 virtual SCSI disks.
- Physical disks. Physical disks attached directly to a partition (sometimes referred to as pass-through disks) have no size limitation other than what is supported by the guest operating system.
- Partition storage capacity. Using virtual hard disks, each partition supports up to 512 TB of storage. Using physical disks, this number is even greater depending on what is supported by the guest operating system.

- Partition snapshots. Hyper-V R2 supports up to 50 snapshots per partition.

1.5.5 Other hardware components

The following is information about the other types of virtual hardware components that can be used with Hyper-V R2.

Virtual DVD drive	<p>A partition has 1 virtual DVD drive by default when you create the partition. Partitions can be configured with up to 3 DVD drives, connected to an IDE controller. (Partitions support up to 4 IDE devices, but one device must be the startup disk.)</p> <p>A virtual DVD drive can access CDs and DVDs, either .iso files or physical media. However, only one partition can be configured to access a physical CD/DVD drive at a time.</p>
Virtual COM port	<p>Each partition is configured with 2 virtual serial (COM) ports that can be attached to a named pipe to communicate with a local or remote physical computer.</p> <p>Note: No access to a physical COM port is available from a partition.</p>
Virtual floppy drive	<p>Each partition is configured with 1 virtual floppy drive, which can access virtual floppy disk (.vfd) files.</p> <p>Note: No access to a physical floppy drive is available from a partition.</p>

1.5.6 Supported guest operating systems

The following operating systems are supported for use in a partition as a guest operating system. 32-bit and 64-bit guest operating systems can be executed at the same time on one server running Hyper-V R2. Note that the TSF does not make any security-related assumptions based on the type of guest operating system used – see Section 1.3.1. Therefore, other operating systems not explicitly listed below, such as future versions of the listed ones, may also be used as guests.

- The following editions of Windows Server 2008 (32-bit and 64-bit) and Server 2008 R2 (64 bit) can be used as a supported guest operating system in a partition configured with 1, 2, or 4 virtual processors:
 - Windows Server 2008 R2 Standard (without Hyper-V R2)
 - Windows Server 2008 R2 Enterprise (without Hyper-V R2)
 - Windows Server 2008 R2 Datacenter (without Hyper-V R2)
 - Windows Server 2008 Standard without Hyper-V

- Windows Server 2008 Enterprise without Hyper-V
- Windows Server 2008 Datacenter without Hyper-V
- Windows Web Server 2008
- Windows Server 2008 HPC Edition
- The following editions of Windows Server 2003 can be used as a supported guest operating system in a partition configured with 1 or 2 virtual processors:
 - Windows Server 2003 R2 Standard Edition with Service Pack 2
 - Windows Server 2003 R2 Enterprise Edition with Service Pack 2
 - Windows Server 2003 R2 Datacenter Edition with Service Pack 2
 - Windows Server 2003 Standard Edition with Service Pack 2
 - Windows Server 2003 Enterprise Edition with Service Pack 2
 - Windows Server 2003 Datacenter Edition with Service Pack 2
 - Windows Server 2003 Web Edition with Service Pack 2
 - Windows Server 2003 R2 Standard x64 Edition with Service Pack 2
 - Windows Server 2003 R2 Enterprise x64 Edition with Service Pack 2
 - Windows Server 2003 R2 Datacenter x64 Edition with Service Pack 2
 - Windows Server 2003 Standard x64 Edition with Service Pack 2
 - Windows Server 2003 Enterprise x64 Edition with Service Pack 2
 - Windows Server 2003 Datacenter x64 Edition with Service Pack 2
- The following versions of Windows 2000 can be executed in a partition configured with 1 virtual processor:
 - Windows 2000 Server with Service Pack 4
 - Windows 2000 Advanced Server with Service Pack 4
- The following Linux distributions can be executed in a partition configured with 1 virtual processor:
 - Suse Linux Enterprise Server 10 with Service Pack 2 (x86 edition)
 - Suse Linux Enterprise Server 10 with Service Pack 2 (x64 edition)

- Suse Linux Enterprise Server 10 with Service Pack 1 (x86 edition)
- Suse Linux Enterprise Server 10 with Service Pack 1 (x64 edition)
- The following 32-bit and 64-bit versions of Windows 7 and Windows Vista can be executed in a partition:
 - Windows 7 x86 (configured with 1, 2 or 4 virtual processors)
 - Windows 7 x64 (configured with 1, 2 or 4 virtual processors)
 - Windows Vista Business with Service Pack 1 (configured with 1 or 2 virtual processors)
 - Windows Vista Enterprise with Service Pack 1 (configured with 1 or 2 virtual processors)
 - Windows Vista Ultimate with Service Pack 1 (configured with 1 or 2 virtual processors)
- The following versions of Windows XP can be executed in a partition:
 - Windows XP Professional with Service Pack 3 (configured with 1 or 2 virtual processors)
 - Windows XP Professional with Service Pack 2 (configured with 1 virtual processor)
 - Windows XP Professional x64 Edition with Service Pack 2 (configured with 1 or 2 virtual processors)

It is worth to note that Windows Server 2008 R2 has been evaluated as executing in a guest partition of Hyper-V as one the platforms.

1.5.7 Integration services (Enlightenments)

Integration services are available for supported guest operating systems as described in the following table. While those integration services are part of the Hyper-V TOE, they are not part of the Hyper-V TSF since they operate as part of a guest partition and can therefore not be protected by Hyper-V from tampering and bypassing. The evaluation of Hyper-V ensures that the Hyper-V security functions are enforced even when the integration services within a guest partition have been tampered with or are bypassed.

Guest operating system	Device and service support
Windows Server 2008 R2	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows Server 2008 (64-bit editions)	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization,

Guest operating system	Device and service support
	data exchange, heartbeat, and online backup
Windows Server 2008 (x86 editions)	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows Server 2003 (x64 editions) with Service Pack 2	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows Server 2003 (x86 editions) with Service Pack 2	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows 2000 Server with Service Pack 4	Drivers: IDE, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, and heartbeat
Windows 2000 Advanced Server with Service Pack 4	Drivers: IDE, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, and heartbeat
Suse Linux Enterprise Server 10 (x64 edition) with Service Pack 1 or 2	Drivers only: IDE, SCSI, networking, and mouse
Suse Linux Enterprise Server 10 (x86 edition) with Service Pack 1 or 2	Drivers only: IDE, SCSI, networking, and mouse
Windows Vista (64-bit editions) with Service Pack 1	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows Vista (x86 editions) with Service Pack 1	Drivers: IDE, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, heartbeat, and online backup

Guest operating system	Device and service support
Windows XP Professional (x86 editions) with Service Pack 2 or 3	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, and heartbeat
Windows XP Professional x64 Edition with Service Pack 2	Drivers: IDE, SCSI, networking, video, and mouse Services: operating system shutdown, time synchronization, data exchange, and heartbeat

1.6 TOE Guidance

The following guidance documentation exists for the startup and configuration as well as the operation of the TOE:

- Windows Server Core 2008 R2: Hyper-V Server Role, Hyper-V R2 Evaluated Configuration Guide [Hyper-V R2_ECG]

This document contains the entry point for installing, configuring and operating the TOE in its evaluated configuration. This document contains pointers to other documents that are also considered guidance, where [Hyper-V R2_ECG] has precedence whenever those documents contradict [Hyper-V R2_ECG].

1.7 References

[AMD-V]	AMD64 Technology AMD64 Architecture Programmer's Manual Volume 2: System Programming, Revision 3.17, June 2010
[CC]	Common Criteria for Information Technology Security Evaluation, Version 3.1 Revision 3, July 2009, Part 1 to 3, CCMB-2009-07-001 to CCMB-2009-07-003
[CEM]	Common Methodology for Information Technology Security Evaluation, Version 3.1 Revision 3, July 2009, CCMB-2008-07-004
[INTEL-VT]	Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2, Order Number: 253669-037US, January 2011
Hyper-V R2_ECG	Microsoft Server Core 2008 R2: Hyper-V Server Role, Hyper-V R2 Evaluated Configuration Guide, Version 2.0, tbd

2 Conformance Claim

This Security Target is conformant to the Common Criteria for Information Technology Security Evaluation Version 3.1 Revision 3 [CC]. It is CC Part 2 extended and Part 3 conformant, with a claimed Evaluation Assurance Level of EAL4, augmented by ALC_FLR.3.

This Security Target does not claim conformance with any Protection Profile.

3 Security Problem Definition

3.1 Introduction

The intended purpose of the TOE is to allow the operation of different guest partitions on a single hardware server where each guest partition will operate like its own (virtualized) server system. Like in a data center that hosts a number of server systems used for different purposes and operate potentially even for different organizations, the separation of the guest partitions within the TOE is the main task to be achieved by the TSF. In addition, a guest partition should not be able to access TSF controlled resources other than those the TSF has assigned to the guest partition.

The configuration of the partitions is assumed to be performed by trusted administrator that are well trained and perform their duties in accordance with the policies defined by the organization that operates the TOE.

3.2 Threats

The **threat agents** in such a system are therefore the following:

- Subjects within a guest partition that attempt to interfere with the configuration or operation of other guest partitions, that attempt to escalate their privileges within the TOE, attempt to access or use resources not assigned to the partition, or attempt to escalate their privileges within a partition using a way that they would not have when the software in the partition would execute on a real server system.
- External entities that attempt to attack the TOE via an external interface in order to tamper with the TSF or the TSF data.
- Administrators that may misconfigure the TOE such that it does not protect assets in accordance with an organizations policy.

The first two threat agents are assumed to have a basic attack potential (compatible with the target assurance level). Administrators are assumed to be trustworthy and not assumed to actively attack the security functionality of the TOE. Especially they are not assumed to deliberately attempt to extend their privileges. As a result the administrator is only seen as a threat agent to the respect that the guidance documentation may be misleading in a way that the administrator configures the TOE to an insecure state where the guidance does not provide sufficient information that this state is insecure.

The **assets** that need to be protected are the resources assigned to a partition, the virtualized and synthetic devices managed by the root partition and the TSF data, which includes the configuration data of the partitions and the audit data generated by the TOE.

The TOE is designed to counter the following threats:

T.CONFIGURATION_CHANGE	The lack of TSF-enforced constraints on the ability of an authorized subject to invoke or dictate how the TOE is reconfigured may result in the TOE transitioning to an insecure (unknown, inconsistent, etc) state.
T.DENIAL_OF_SERVICE	A malicious subject may block others from specific system resources (system memory, persistent storage, and processing time) via a resource exhaustion attack.
T.UNAUTHORIZED_ACCESS	A subject may gain access to resources or TOE security management functions for which it is not authorized according to the TOE security policy.
T.PARTITION_COMPROMISE	The TSF of a correctly configured TOE may allow software within a guest partition to escalate its processor or memory related privileges in a way that would not be possible when the software is executed on the real hardware. (Note: The software within a guest partition is not subject to this evaluation and therefore any flaw within a guest partition's software that allows such a privilege escalation is beyond the scope of this evaluation. Also the software a hypervisor-aware operating system installs when executing on the TOE is not part of this evaluation).

Table 2: List of Threats

3.3 Security Policies

An organizational security policy is a set of rules or procedures imposed by an organization upon its operations to protect its sensitive data. The following operational security policies are supported by the TOE:

P.ACCOUNTABILITY	The TOE shall provide the capability to make available information regarding the occurrence of security-relevant events.
P.CONFIGURATION	The TOE shall provide functions that allow authorized administrators to perform the setup and initialization of the TOE in accordance with the security policies for configuration and administration issued by the organization responsible for the operation of the TOE.

Table 3: List of Organisational Security Policies

3.4 Assumptions

This section describes the security aspects of the environment in which the TOE is intended to be used. It includes information about the physical, personnel, procedural, and connectivity aspects of the environment.

The TOE is assured to provide effective security measures in a cooperative non-hostile environment only if it is installed, managed, and used correctly. The operational environment must be managed in accordance with user/administrator guidance documentation. The following specific conditions are assumed to exist in an environment where the TOE is employed:

A.ADMIN	Authorized administrators of the TOE are assumed to be knowledgeable and trustworthy to follow the guidance and not misuse their privileges. This applies also to domain administrators that manage the AD-DS used by the TOE.
A.PHYSICAL	It is assumed that the non-IT environment provides the TOE with appropriate physical security commensurate with the value of the IT assets protected by the TOE.
A. SUBJECT_ALLOCATION	It is assumed that properly trained trusted administrators will create and manage the configuration data of partitions.
A.USER_SUBJECT_BINDING	It is assumed that users are authenticated and bound to subjects in the root partition with the correct security attributes assigned to the subjects, allowing the TOE to base access decisions upon such security attributes.
A.DEFINED_INSTALL	It is assumed that the administrator installs and configures the TOE in accordance with the guidance provided for the installation and configuration of the TOE.
A.REMOTE_ADMIN	It is assumed that remote administration is performed only using properly protected communication links.
A.REMOTE_IT_PRODUCTS	It is assumed that any other IT product that may be used to support the authentication of administrators, used to protect communication links, or used to assist administrators in their administrative tasks is trusted to perform its security related functions correctly and does not include side effects that may allow unauthorized persons to perform administrative functions on the TOE

	<p>or perform administrative functions other than those explicitly initiated by the trusted administrator. Any communication to such a trusted IT product is assumed to be protected against unauthorized interception or modification of the network traffic. This applies especially to the communication with the Active Directory Service (AD-DS) used by the TOE environment (Server 2008 R2 in the root partition).</p> <p>Note: The Windows Server 2008 R2 in the root partition can operate either in a stand-alone or in domain member configuration as covered by the evaluation of Windows Server 2008 R2. Operating the Windows Server 2008 R2 instance in the root partition as a domain controller would contradict the assumption A.CLEAN_ROOT below and is not covered by this evaluation.</p>
A.CLEAN_ROOT	It is assumed that no additional software than the one specified in the configuration guidance is installed in the root partition.
A.CORRECT_HARDWARE	It is assumed that the underlying hardware of the TOE operates correctly as described in the hardware manuals and does not expose undocumented critical side effects
A.PARTITION_CONNECT	It is assumed that users that are allowed to connect to a particular guest partition via the Remote Desktop Protocol (RDP) have all the same right to access information within this partition.
A.MEMORY_MANAGEMENT	The Windows Server 2008 R2 instance that is running in the root partition provides memory management services to other components running in the server instance or the root partition with kernel-mode privileges. It exposes a dedicated and functionally-complete kernel memory management API to these components.
A.LIVE_MIGRATION	The communication link used in the communication for live migration of partitions is assumed to be protected by the operational environment from unauthorized access and undetected modification of the data

	transferred as part of the live migration process.
--	--

Table 4: List of Assumptions

4 Security Objectives

This section defines the security objectives of the TSF and its supporting environment. Security objectives, categorized as either IT security objectives or non-IT security objectives, reflect the stated intent to counter identified threats, comply with any organizational security policies identified, or both. All of the identified threats and organizational policies are addressed under one of the following categories.

4.1 Security Objectives for the TOE

O.PARTITION_ACCESS	The TOE will ensure that subjects within a partition gain only authorized access to exported resources assigned to the partition.
O.AUDIT_GENERATION	The TOE will provide the capability to generate audit records for security relevant auditable events.
O.AUTHORIZED_SUBJECT	The TOE will ensure that only authorized subjects are allowed to access TSF data. Note: this applies for access to TSF data where the access control is performed by the TOE. Access fully controlled by the TOE environment is covered by the evaluation of the access control functions in the IT environment of the TOE.
O.INIT_SECURE_STATE	The TOE will provide functions that allow administrators to setup and configure the TOE such that it is started in a secure state where all the other security objectives are enforced.
O.MANAGE	The TOE will provide all the functions necessary to support the administrative users and authorized subjects in their management of the TOE security functions and configuration data, and restrict these functions from use by unauthorized subjects.
O.RESIDUAL_INFORMATION	The TOE will ensure that any information contained in a protected resource is not released to subjects when the resource is reallocated.
O.RESOURCE_ALLOCATION	The TOE will provide mechanisms that enforce constraints on the allocation of TOE resources

	assigned to a partition.
O.PARTITION_ISOLATION	The TOE will provide mechanisms to protect each guest partition from unauthorized interference by other guest partitions.
O.PRESERVE_PART_PRIV	The TOE will preserve the hardware separation functions within a partition such that software within the partition is able to implement its own policy for separation in the same way as it would be when executing directly on the underlying hardware
O.LIVE_MIGRATION	The TOE will preserve the integrity of a partitions data, the integrity of a partitions privileges and security attributes, and the overall context the partition is operated in when the running partition is migrated from one instance of the TOE to another one. The TOE will prohibit live migration if the source and destination instance of the TOE are not sufficiently compatible to provide the identical operational environment to the partition on both the source and the destination instance of the TOE.

Table 5: List of Security Objectives for the TOE

4.2 Security Objectives for the TOE environment

The TOE is assumed to be complete and self-contained and, as such, not dependent upon any other products to perform properly. However, certain objectives with respect to the general operating environment must be met. The following are the non-IT security objectives:

OE.PHYSICAL	Physical security will be provided for the TOE by the non-IT environment commensurate with the value of the IT assets protected by the TOE.
OE.HW_MECHANISMS	The underlying processor implements the privileges, ring separation, memory protection, and virtualization support as described in the hardware manuals. The underlying hardware implements interrupt handling, bus configuration, device controller configuration, timers and time, and other I/O related aspects as described in the

	related hardware manuals.
OE.HW_SIDE_EFFECTS	The underlying hardware has no undocumented side effects that may interfere with the security functions implemented by the TOE.
OE.USER_SUBJECT_BINDING	The TOE environment (in this case the Windows Server 2008 R2 instance in the root partition) will ensure that users are correctly authenticated and processes acting on behalf of such users are bound to the user indicating correctly the security attributes the process has. The authentication methods and the management of user accounts allowed in the evaluation of Windows Server 2008 R2 for a stand-alone or domain connected configuration can be used.
OE.SUBJECT_ALLOCATION	A properly trained trusted individual will create configuration vectors such that, for those partitions to which subjects are allocated, each partition is allocated one or more subjects (i.e., subjects with homogeneous access requirements, or subjects with heterogeneous access requirements) that are appropriate for the policy abstraction supported by the TOE.
OE.TRUSTED_INDIVIDUAL	Any individual allowed to perform procedures upon which the security of the TOE may depend must be trusted with assurance commensurate with the value of the IT assets. This applies also for individuals managing functions within the TOE environment that the TOE depends on.
OE.REMOTE_ADMIN	Any communication link used for remote administration or communication with another trusted IT product must be protected from unauthorized access or interference by unauthorized persons or systems. Note: The Security Target for Server 2008 R2 included security functionality allowing the protection of communication links. It is up to the persons responsible for the setup and

	configuration of a system using the TOE to decide if the functionality of Server 2008 R2 is used for the protection of communication links or if they want to use other functions or mechanisms.
OE.REMOTE_IT_PRODUCT	<p>Any remote IT product used to assist in the authentication of administrative users, in the administration activities or in the storage of TSF data must protect all security related data against unauthorized access and must ensure that it performs the function the TOE expects from this product correctly and without any side effects that could undermine the security of the TOE.</p> <p>Note: It is common to use a Domain Controller to centrally manage services and accounts. This aspect of management has been addressed in the security evaluation of Server 2008 R2 and this Security Target bases this objective for the TOE environment on the assessment performed in the Server 2008 R2 evaluation as far as Active Directory is concerned.</p>
OE.NETWORK	The physical networks the TOE is connected to, have controls that protect against attacks on the physical layer (e. g. high voltage) and against attacks on the data link layer (layer 2).
OE.TRUSTED_ROOT	The root partition is installed as defined for the evaluated configuration and has no software applications installed that are not required for the virtualization support or the management of Hyper-V R2. Especially there is no untrusted application installed in the root partition.
OE.PARTITION_USER	Procedures in the TOE environment exist to ensure that users that are allowed to connect to a specific partition have the right to access all information processed by this partition.
OE.LIVE_MIGRATION	Measures within the TOE environment will ensure the confidentiality and integrity of the data transferred as part of the migration of a running

	partition between two instances of the TOE.
--	---

Table 6: List of Security Objectives for the TOE environment

4.3 Security objectives rationale

This section provides a rationale for the existence of each threat, policy statement, security objective, and component that comprise the security target.

4.3.1 Complete coverage: threats and organizational security policies

This section provides evidence demonstrating coverage of the threats and Organizational Security Policies (OSPs) by both the IT and non-IT security objectives. The following table shows this objective to threat and policy mapping, and the table is followed by a discussion of the coverage for each threat and OSP.

T.CONFIGURATION_CHANGE	<p>O.MANAGE</p> <p>O.AUTHORIZED_SUBJECT</p> <p>OE.USER_SUBJECT_BINDING</p> <p>OE.TRUSTED_INDIVIDUAL</p> <p>OE.REMOTE_ADMIN</p> <p>OE.REMOTE_IT_PRODUCT</p>
T.DENIAL_OF_SERVICE	O.RESOURCE_ALLOCATION
T.UNAUTHORIZED_ACCESS	<p>O.PARTITION_ACCESS</p> <p>O.AUTHORIZED_SUBJECT</p> <p>O.RESIDUAL_INFORMATION</p> <p>O.PARTITION_ISOLATION</p> <p>O.LIVE_MIGRATION</p> <p>OE.PHYSICAL</p> <p>OE.USER_SUBJECT_BINDING</p> <p>OE.NETWORK</p> <p>OE.REMOTE_IT_PRODUCT</p> <p>OE.TRUSTED_ROOT</p>

	OE.LIVE_MIGRATION
T.PARTITION_COMPROMISE	O.PRESERVE_PART_PRIV O.LIVE_MIGRATION OE.HW_MECHANISMS OE.HW_SIDE_EFFECTS OE.LIVE_MIGRATION

Table 7: Mapping Threats to Security Objectives

P.ACCOUNTABILITY	O.AUDIT_GENERATION OE.USER_SUBJECT_BINDING
P.CONFIGURATION	O.INIT_SECURE_STATE OE.SUBJECT_ALLOCATION OE.TRUSTED_INDIVIDUAL OE.REMOTE_IT_PRODUCT OE.TRUSTED_ROOT

Table 8: Mapping Organisational Security Policies to Security Objectives

Threats

T.CONFIGURATION_CHANGE

The lack of TSF-enforced constraints on the ability of an authorized subject to invoke or dictate how the TOE is reconfigured may result in the TOE transitioning to an insecure (unknown, inconsistent, etc) state.

This threat is addressed by the security objective O.MANAGE that requires the existence of appropriate management functions together with the security objective O.AUTHORIZED_SUBJECT, which requires that only authorized administrators can perform those management functions which in combination with OE.USER_SUBJECT_BINDING ensures that configuration changes performed by the TOE can only be performed by user authorized for such operations. OE.TRUSTED_INDIVIDUALS supports countering this threat by requiring that those administrators are trusted to perform their job correctly and not misuse their privileges. OE.REMOTE_ADMIN also supports countering this threat by requiring that remote administration facilities are protected against unauthorized subjects attempting to access the communication link. OE.REMOTE_IT_PRODUCT addresses the issue that a remote product used for the

management of the TOE configuration may be used by an attacker or an unauthorized person to impersonate as an authorized administrator and modify the configuration of the TOE.

T.DENIAL_OF_SERVICE

A malicious subject may block others from system resources (e.g., system memory, persistent storage, and processing time) via a resource exhaustion attack.

This threat is countered by security objective O.RESOURCE_ALLOCATION requesting that constraints exist on the allocation of resources to partitions, which in turn prohibits a denial of service attack caused by resource exhaustion.

T.UNAUTHORIZED_ACCESS

A subject may gain access to resources or TOE security management functions for which it is not authorized according to the TOE security policy.

This threat is countered by the security objective O.PARTITION_ACCESS, which requests an access control mechanism for resources exported to partitions. Concerning access to partition configuration data and other TSF data, this threat is addressed by the security objective O.AUTHORIZED_SUBJECT, which requires restricting access to this data to authorized administrators (the authentication of administrative users is performed by Server 2008 R2 in the root partition). The aspect of subjects getting access to information by residuals left in resources assigned to the subject is addressed by O.RESIDUAL_INFORMATION. The aspect of a guest partition getting access to data from other partitions by sharing data with another guest partition is addressed by the security objective O.PARTITION_ISOLATION. The aspect of a subject gaining unauthorized access during or after being migrated to a different instance of the TOE is addressed by O.LIVE_MIGRATION.

OE.PHYSICAL supports countering the threat by prohibiting unauthorized persons to access TOE resources by physically manipulating the TOE hardware.

OE.USER_SUBJECT_BINDING supports countering the threat by assuring that subjects that request access have correctly been bound to a user and that this user has been successfully authenticated.

OE.NETWORK supports countering the threat by prohibiting attempts to use irregular network signals to attack the TOE and potentially damage the network adapter in way that may cause changes in the TOE configuration.

OE.REMOTE_IT_PRODUCT supports countering the threat by prohibiting unauthorized access to the TOE using a remote product for accessing the TOE that can not be trusted to perform its functions correctly.

OE.TRUSTED_ROOT supports countering the threat by prohibiting that any untrusted application is installed in the root partition. This prohibits that any untrusted program may attempt to use the Server 2008 R2 system call interface in order to explore a way to bypass the access control policies enforced. Only the defined applications required for virtualization support and management are installed and those are part of the TSF.

OE.LIVE_MIGRATION supports countering the threat by prohibiting an external entity manipulating data transferred during live migration in a way that would give the migrated subject different access to resources than what the subject had on the system it was migrated from.

T.PARTITION_COMPROMISE

The TSF of a correctly configured TOE may allow software within a guest partition to escalate its processor or memory related privileges in a way that would not be possible when the software is executed on the real hardware. (Note: The software within a guest partition is not subject to this evaluation and therefore any flaw within a guest partition's software that allows such a privilege escalation is beyond the scope of this evaluation. Also the software a hypervisor aware operating system installs when executing on the TOE is not part of this evaluation).

The threat of software within a partition escalating its hardware privileges within the partition by using functions provided by the TSF is addressed by security objective O.PRESERVE_PART_PRIV, which requires that this is not possible. The security objective O.LIVE_MIGRATION addresses this threat by ensuring that a partitions privileges as well as the partitions view of the underlying platform are maintained.

OE.HW_MECHANISMS supports countering the threat by requiring the hardware to work correctly in accordance with its specification. The TOE relies on the documented functions of the underlying hardware to correctly enforce the same separation and privilege policy that is enforced when the software runs on a real system.

OE.HW_SIDE_EFFECTS supports countering the threat by requiring the hardware to not expose undocumented side effects which, although not contradicting the specification, may undermine the separation within a partition.

OE.LIVE_MIGRATION supports countering the threat by requiring integrity protection for the data transferred as part of a live migration operation.

Organisational Security Policies

P.ACCOUNTABILITY

The TOE shall provide the capability to make available information regarding the occurrence of security relevant events.

This organizational security policy is addressed by the security objective O.AUDIT_GENERATION, which requires the TOE to provide an audit function for security relevant events that allows making subjects responsible for their actions.

This policy is supported by OE.USER_SUBJECT_BINDING, which ensures that subjects are correctly bound to users allowing the ensured identification of the user responsible for an action.

P.CONFIGURATION

The TOE shall provide functions that allow authorized administrators to perform the setup and initialization of the TOE in accordance with the security policies for configuration and administration issued by the organization responsible for the operation of the TOE.

This security policy is addressed by the security objective O.INIT_SECURE_STATE which requires the TOE to provide the administrative functions that allow an authorized administrator to setup and initialize the TOE such that it starts in a secure state.

OE.SUBJECT_ALLOCATION supports this organisational security policy such that authorized administrators do not define configurations that contradict the separation policy an organization wants to enforce.

OE.TRUSTED_INDIVIDUAL supports this organisational security policy by requiring authorized administrators to not misuse his privileges.

OE.REMOTE_IT_PRODUCT supports this organisational security policy by requiring that remote IT products used for managing or accessing the TOE are trusted to correctly pass their user's actions to the TOE and not to send requests to the TOE that have not been initiated by the user of the product.

OE.TRUSTED_ROOT supports this organisational security policy by requiring that no untrusted application is installed in the root partition that may attempt to tamper with the configuration as defined by the trusted administrator.

Mapping security objectives for the operational environment to the assumption backing them

The following table maps the security objectives for the operational environment to the assumptions that back those objectives.

OE.PHYSICAL	This security objective for the operational environment is backed by assumption A.PHYSICAL.
OE.HW_MECHANISMS	This security objective for the operational environment is backed by assumption A.CORRECT_HARDWARE.
OE.HW_SIDE_EFFECTS	This security objective for the operational environment is backed by assumption A.CORRECT_HARDWARE.
OE.USER_SUBJECT_BINDING	This security objective for the operational environment is backed by assumption A.USER_SUBJECT_BINDING.
OE.SUBJECT_ALLOCATION	This security objective for the operational environment is backed by assumption A.SUBJECT_ALLOCATION
OE.TRUSTED_INDIVIDUAL	This security objective for the operational environment is backed by assumption A.ADMIN.
OE.REMOTE_ADMIN	This security objective for the operational environment is backed by assumption A.REMOTE_ADMIN.
OE.REMOTE_IT_PRODUCT	This security objective for the operational environment is backed by assumption A.REMOTE_IT_PRODUCTS.
OE.NETWORK	This security objective for the operational environment is backed by assumption A.PHYSICAL.
OE.TRUSTED_ROOT	This security objective for the operational environment is backed by assumptions A.Defined_INSTALL and A.CLEAN_ROOT.
OE.PARTITION_USER	This security objective for the operational environment is backed by assumption A.PARTITION_CONNECT

OE.LIVE_MIGRATION	This security objective for the operational environment is backed by assumption A.LIVE_MIGRATION
-------------------	---

Table 9: Mapping security objectives for the operational environment to assumptions

5 Extended Components Definition

This Security Target includes one extended components. This extended component defines a subset of the component FAU_GEN.1 as defined in part 2 of the CC.

This extended component needed to be defined since Hyper-V uses the audit trail interfaces provided by Windows Server 2008 R2 for trusted components that want to store their audit records in the common audit trail provided by Windows Server 2008 R2. The existing SFR of the CC do not offer the possibility of auditing without providing timestamps, which is the intended security behavior of the TOE. While Hyper-V collects all the information for its own audit records and formats the audit records, it does not generate an audit record for the start and the stop of the audit system (this is done by Windows Server 2008 R2, which controls the audit trail). Hyper-V also does not store the time and date into the audit record. This is also a function that is performed by Windows Server 2008 R2 when a trusted component submits an audit record for inclusion into the Windows Server 2008 R2 audit trail.

This way of handling audit records is common for a large number of products that on the one hand want to produce audit records and on the other hand do not want to implement the functions to manage the audit trail, protect the audit trail, or process audit records for evaluation. Many applications use audit trails provided either by the underlying operating system or by a dedicated audit component in the IT environment (e. g. a log host). In order to have a single time source most of those dedicated audit systems that provide an interface to other product to submit audit records for storing them in the audit trail will put the time stamp with the time and date into the audit records themselves rather than relying on the other systems to synchronize their time sources.

The extended requirement FAU_GEN_SUB.1 has been derived as a subset from the SFR FAU_GEN.1 as listed in part 2 of the CC. The requirement for auditing the start and stop of the audit system has been dropped compared to FAU_GEN.1.1 and in FAU_GEN.1.2 the requirement to include the time and date has been dropped. As a consequence of dropping the inclusion of the time and date, the extended component no longer has a dependency on FPT_STM.1

FAU_GEN_SUB.1 Audit data generation

Hierarchical to: No other components.

Dependencies: none

FAU_GEN_SUB.1.1 The TSF shall be able to generate an audit record of the following auditable events:

- a) All auditable events for the [selection, choose one of: *minimum, basic, detailed, not specified*] level of audit; and
- b) [assignment: other specifically defined auditable events].

FAU_GEN_SUB.1.2 The TSF shall record within each audit record at least the following information:

- a) type of event, subject identity (if applicable), and the outcome (success or failure) of the event; and
- b) For each audit event type, based on the auditable event definitions of the functional components included in the PP/ST, [assignment: *other audit relevant information*].

6 Security Requirements

This section defines the security functional and security assurance requirements that apply for the evaluation of Microsoft Hyper-V R2.

6.1 Security Functional Requirements

As mentioned in the TOE introduction and explained in more detail in the TOE summary specification, the TSF for this evaluation consist of the hypervisor layer and the Hyper-V R2 related software within the root partition. The root partition also includes a Server 2008 R2 Server Core installation which is augmented by the components specific to Hyper-V R2.

The security functionality described in this Security Target is partly implemented by the hypervisor layer and partly by the Hyper-V R2 specific components in the root partition. The implementation of some security functions of Hyper-V use supporting security services of the Windows Server 2008 R2 instance in the root partition. In cases where the implementation of a security function of Hyper-V uses security functionality of the Windows Server 2008 R2 instance in the root partition, this is described in application notes.

In addition, assignments or selections performed on the components are marked in **bold** and refinements performed are marked in **bold and underlined**.

The security functional requirements have been derived using the following paradigms for Hyper-V R2:

- Hyper-V R2 has two types of “users”:
 - Human users that act as administrators to configure and manage the TOE.
 - Partitions as entities outside of the TOE that “use” the functions of the TSF.

Note that the human users need to be identified and authenticated while the partitions are created and operate under the complete control of the TOE and therefore only need to be identified.
- Hyper-V R2 has two types of access control policies:
 - A “Partition Management Access Control Policy” which controls the actions of administrative users. This policy allows controlling access of administrative users to Hyper-V management operations. The functionality allows grouping of operations to “tasks” and the assignment of tasks and/or individual operations to roles. Roles can then be assigned to administrative users. Since this grouping is dynamic and can be done in an installation defined way, no fixed “roles” are defined in this Security Target.
 - A “Partition Access Control Policy” which controls the access of partitions to virtualized devices
- The security attributes used in the enforcement of the policies are the following:

- Human users (administrators):
 - Identity of the user (which is provided by Windows Server 2008 R2)

Note that user role within Hyper-V may be dynamically defined within Hyper-V using the Hyper-V management functions. Roles are just sets of Hyper-V related management privileges that may be assigned to a user.

- Partitions:
 - Identity of the Partition
 - Partition privileges
- Hyper-V R2 configuration data:
 - Access control lists associated with the data

Note that for the partition management access control policy the TOE relies on the generic access control functionality of Server 2008 R2 very much like any application relies on an operating system's access control policy to protect data it stores in operating system provided persistent data container like files..

- Virtualized devices:
 - Access control list associated with the virtualized device
- Residual information protection
 - Hyper-V R2 removes all information from resources before a resource is assigned to a partition
- Audit Policy:
 - Hyper-V R2 allows to audit the following Hyper-V R2 specific events:
 - Failure to start the hypervisor
 - Partition creation
 - Partition deletion
 - Critical hypervisor error

For the management and protection of the audit trail as well the evaluation of the audit records by authorized administrators the generic functions of Windows Server 2008 R2 are used. Those include:

- Access control for the audit trail

- Selection of events that are audited
- Tools to review the audit trail
- Actions performed to ensure that audit records are not lost
- Management Policy:
 - Hyper-V R2 allows authorized administrators to manage the following Hyper-V R2 specific aspects:
 - Creation and deletion of partitions
 - Assignment of virtualized resources to partitions
 - Definition of maximum quota of resources (CPU time, memory) for partitions
- Protection of TSF data:
 - Hyper-V establishes a separate domain for the hypervisor as well as for the root partition. The VMBus functionality allows the root partition to share dedicated memory with a guest partition. This memory is only used for fast data transfer between a guest partition and the root partition and the root partition does not store any TSF data in memory shared with a guest partition.
 - For live migration Hyper-V ensures that the target system of the migration is compatible to the source system such that the guest system that is migrated sees exactly the same (virtualized) platform on both the source and the target system.
 - For live migration Hyper-V ensures that a migrated guest partition will only be started on the target system when the complete system state has been correctly transferred. If this can not be verified, the guest partition is resumed on the source system and the (incompletely) migrated partition on the target system is deleted and never activated.

This security functionality relies on the following security functions implemented in Windows Server 2008 R2:

- Identification and authentication of administrative users
While Windows Server 2008 R2 performs the authentication of users, the management of the Hyper-V specific user security attributes that define the management actions the user can perform on Hyper-V TSF data, objects and functions is controlled by Hyper-V itself.
- Management of the Windows Server 2008 R2 security attributes of administrative users
- Access control policy for Server 2008 R2 files and objects
- Management of the Server 2008 R2 access control policy

- Generation of audit events specific for Server 2008 R2
- Management and protection of the audit trail
- Review of audit records
- Management and provision of date and time

6.1.1 Security Functional Requirements implemented by Hyper-V R2

6.1.1.1 Security Audit (Class FAU)

FAU_GEN_SUB.1 (H) Audit data generation

FAU_GEN_SUB.1.1 The TSF shall be able to generate an audit record of the following auditable events:

1. All auditable events for the **not specified** level of audit; and
2. **The following hypervisor specific events:**
 - **Failure to start the hypervisor**
 - **Creation of a partition (by the hypervisor)**
 - **Deletion of a partition (by the hypervisor)**
 - **Failure condition detected within the hypervisor**
3. **The following partition management specific events:**
 - **Access checks performed by AzMan on Hyper-V R2 management operations**
 - **Reconfigure partition (Virtual Machine)**

Application Note: Hyper-V R2 uses the Windows Server 2008 R2 audit system. Startup and shutdown of this audit system are recorded by Server 2008 R2 in the root partition without involvement of the Hyper-V R2 specific functionality. The audit functions of Windows Server 2008 R2 are also responsible to write the time and date into the audit record.

Application Note: Modifications to the Hyper-V R2 AzMan policy are audited by the IT environment.

FAU_GEN_SUB.1.2 The TSF shall record within each audit record at least the following information:

- a) type of event, subject identity (if applicable), and the outcome (success or failure) of the event; and

- b) For each audit event type, based on the auditable event definitions of the functional components included in the PP/ST, **and no other security relevant information.**

Application Note: Date and Time are inserted by the Server 2008 R2 audit function in the root partition..

FAU_GEN.2 (H) User identity association

FAU_GEN.2.1 For audit events resulting from actions of identified users, the TSF shall be able to associate each auditable event with the identity of the user (**partition or the administrator**) that caused the event.

Application Note: Hyper-V R2 supports two types of users: partitions and administrators. The identity of the “user” that causes the event is recorded for both types of users. The identity of administrators is obtained from Windows Server 2008 R2.

6.1.1.2 User Data Protection (Class FDP)

FDP_ACC.1 (H-PM) Subset access control (Hyper-V R2 partition management)

FDP_ACC.1.1 The TSF shall enforce the **Partition Management Access Control Policy** on **subjects acting on behalf of an administrative user, individual management operations as objects and performing the operation as functions.**

Application Note: The subjects are the processes executing on behalf of the Windows Server 2008 R2 defined users..

FDP_ACC.1 (H-DA) Subset access control (Hyper-V R2 device access)

FDP_ACC.1.1 The TSF shall enforce the **Partition Device Access Control Policy** on **partitions as subjects, virtualized and synthesized devices as objects and device access as function.**

FDP_ACF.1 (H-PM) Security attribute based access control (Hyper-V R2 partition management)

FDP_ACF.1.1 The TSF shall enforce the **Partition Management Access Control Policy** to objects based on the following:

- a) **The identity of the administrative user, his roles, the tasks and operations assigned to the roles and the additional BizRules assigned to the tasks.**

Application Note: The identity of the administrative user is established and authenticated by the Server 2008 R2 functions in the root partition. Roles in Hyper-V are defined as a collection of tasks and therefore it is up to the installation if roles are defined and what those roles are.

FDP_ACF.1.2 The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed: **access (i. e performing an operation) is allowed if either the operation itself or a task that contains the operation is assigned to a role that has been assigned to the administrative user and the BizRules assigned to one of the tasks involved in the access decision don't disallows access.**

FDP_ACF.1.3 The TSF shall explicitly authorise access of subjects to objects based on the following additional rules: **none**

FDP_ACF.1.4 The TSF shall explicitly deny access of subjects to objects based on the following additional rules: **no additional rules.**

FDP_ACF.1 (H-DA) Security attribute based access control (Hyper-V R2 device access)

FDP_ACF.1.1 The TSF shall enforce the **Partition Device Access Control Policy** to objects based on the following:

- a) **the type of device**
- b) **the identity of the partition**
- c) **the configuration data for the partition.**

FDP_ACF.1.2 The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed: **a partition is allowed to access a virtualized or synthesized device if the partition configuration has this device assigned to the partition.**

FDP_ACF.1.3 The TSF shall explicitly authorise access of subjects to objects based on the following additional rules: **the device is a virtualized S3 Trio Video Card or a virtualized Intel 440 BX chipset.**

FDP_ACF.1.4 The TSF shall explicitly deny access of subjects to objects based on the following additional rules: **no additional rule.**

FDP_RIP.1 (H) Subset residual information protection

FDP_RIP.1.1 The TSF shall ensure that any previous information content of a resource is made unavailable upon the **allocation of the resource** to the following objects: **virtual**

memory allocated to a partition, virtual and synthesized devices allocated to a partition, virtual processors allocated to a partition.

6.1.1.3 *Identification and Authentication (Class FIA)*

FIA_ATD.1 (H) User attribute definition (Hyper-V R2 partitions)

FIA_ATD.1.1 The TSF shall maintain the following list of security attributes belonging to individual users:

- a) **For partitions:**
 - a. **Partition identifier (number)**
 - b. **Partition configuration data**
 - c. **Partition privileges**
- b) **For administrative users:**
 - a. **none**

Application Note: A human user becomes an administrative user by assigning him the ability to perform management operations for Hyper-V controlled objects and functions. Being an administrator is therefore not a user security attribute but is implicit with the assignment of the ability to perform management functions.

FIA_UID.2 (H) User identification before any action (Hyper-V R2 partitions)

FIA_UID.2.1 The TSF shall require each user (**partition**) to be successfully identified before allowing any other TSF-mediated actions on behalf of that user.

FIA_USB.1 (H) User-subject binding

FIA_USB.1.1 The TSF shall associate the following user security attributes with subjects acting on the behalf of that user:

- a) **For subjects acting on behalf of partitions:**
 - a. **The identity of the partition**
 - b. **The configuration data of the partition**
 - c. **The privilege vector of the partition**

FIA_USB.1.2 The TSF shall enforce the following rules on the initial association of user security attributes with subjects acting on the behalf of users:

- **Every subject acting on behalf of a partition will be assigned the security attributes associated with the partition on whose behalf the subject will act.**

FIA_USB.1.3 The TSF shall enforce the following rules governing changes to the user security attributes associated with subjects acting on the behalf of users: **subjects acting on behalf of partitions can not add additional security attributes beyond those initially assigned.**

6.1.1.4 Security Management (Class FMT)

FMT_MSA.1 (H) Management of security attributes

FMT_MSA.1.1 The TSF shall enforce the **Partition Management Access Control Policy** to restrict the ability to **query and modify** the security attributes **partition configuration data** to **authorized administrators that have the authority for those operations assigned in the AzMan policy for Hyper-V R2.**

Application Note: The AzMan authorization store can be maintained in an Active Directory database or a local XML file. In case of those authorizations stored in an Active Directory database, the AD Domain controller needs to protect the data and the functions to manage the data in accordance with the assumption A.REMOTE_IT_PRODUCTS. This applies to all management functions where the management of security attributes or policies can either be done locally (and then is covered by the local security policy) or remotely in an Active Directory database (in which case Active Directory needs to ensure the data is protected from unauthorized access). The AzMan functions that control those authorities are part of Hyper-V

FMT_MSA.2 (H) Secure security attributes

FMT_MSA.2.1 The TSF shall ensure that only secure values are accepted for **partition configuration data.**

FMT_MSA.3 (H-PM) Static attribute initialization

FMT_MSA.3.1 The TSF shall enforce the **Partition Management Access Control Policy** to provide **restrictive** default values for security attributes that are used to enforce the SFP.

FMT_MSA.3.2 The TSF shall allow the **authorized administrator with access to the AzMan policy** to specify alternative initial values to override the default values when an object or information is created.

Application Note: When Hyper-V is installed at least one user has to be given the permission to access the AzMan policy in order to be able to manage the administrative functions.

FMT_MSA.3 (H-DA) Static attribute initialization

- FMT_MSA.3.1 The TSF shall enforce the **partition device access control policy** to provide **restrictive** default values for security attributes that are used to enforce the SFP.
- FMT_MSA.3.2 The TSF shall allow the **administrator authorized by the Hyper-V R2 AzMan policy** to **modify partition configuration data** to specify alternative initial values to override the default values when an object or information is created except for those security attributes that nobody is allowed to modify.

FMT_MTD.1 (H-1) Management of TSF data

- FMT_MTD.1.1 The TSF shall restrict the ability to **change_default and modify** the **partition's access to hypervisor calls related to inter-partition communication and the creation of guest partitions to no user or partition**.

Application Note: This requirement has been included to describe the requirement that a partition's access to hypervisor calls is static and can not be changed dynamically.

FMT_MTD.1(H-2) Management of TSF data

- FMT_MTD.1.1 The TSF shall restrict the ability to **change_default and modify** the **resource quota assigned to partition to administrators authorized by the Hyper-V R2 AzMan policy to perform those actions**.

FMT_MTD.1(H-3) Management of TSF data

- FMT_MTD.1.1 The TSF shall restrict the ability to **modify** the **Hyper-V R2 AzMan policy to administrators authorized to access the Hyper-V R2 authorization store**.

FMT_REV.1 (H) Revocation

- FMT_REV.1.1 The TSF shall restrict the ability to revoke **the assignment of devices** associated with the **partitions** under the control of the TSF to **administrators**.
- FMT_REV.1.2 The TSF shall enforce the rules **that the administrator needs to be authorized by the Hyper-V R2 AzMan policy to perform those actions**.

Application Note: Revocation is not bound to one specific role but to a privilege that can be assigned to installation defined management roles. The specification of FMT_REV.1 within part 2 of the CC is too narrow and therefore FMT_REV.1.2 has been used to correctly describe the way revocation can be performed in the TOE.

FMT_SMF.1 (H) Specification of Management Functions

FMT_SMF.1.1 The TSF shall be capable of performing the following management functions:

- Management of partition configuration data
- Management of virtual switches
- Defining, deleting, starting and stopping partitions
- Management of the Hyper-V R2 AzMan policy.

6.1.1.5 Protection of the TSF (Class FPT)

FPT_TDC.1 (H) Inter-TSF TSF data consistency

FPT_TDC.1.1 The TSF shall provide the capability to consistently interpret **the complete state of a partition** when shared between the TSF and another trusted IT product **(which needs to be another instance of the TOE) when transferred as part of live migration.**

FPT_TDC.1.2 The TSF shall use **the semantics of the underlying hardware platform and the semantics of the virtualized resources** when interpreting the TSF data from another trusted IT product **that it received as part of a live migration operation.**

FPT_TEE.1 (H) Testing of external entities

FPT_TEE.1.1 The TSF shall run a suite of tests **when live migration of a partition is requested** to check the fulfillment of **the compatibility criteria required for live migration.**

FPT_TEE.1.2 If the test fails, the TSF shall **prohibit live migration of a partition.**

6.1.1.6 Resource Utilisation (Class FRU)

FRU_RSA.1 (H) Maximum quotas

FRU_RSA.1.1 The TSF shall enforce maximum quotas of the following resources: **maximum CPU time per virtual CPU, maximum amount of the partition (guest) physical memory that partitions can use over a specified period of time.**

6.2 Security Assurance Requirements

The SARs for the TOE are the EAL 4 components augmented with ALC_FLR.3 as specified in Part 3 of the CC. No operations are applied to the assurance components.

Assurance Class	Assurance components
ADV: Development	ADV_ARC.1 Security architecture description
	ADV_FSP.4 Complete functional specification
	ADV_IMP.1 Implementation representation of the TSF
	ADV_TDS.3 Basic modular design
AGD: Guidance documents	AGD_OPE.1 Operational user guidance
	AGD_PRE.1 Preparative procedures
ALC: Life-cycle support	ALC_CMC.4 Production support, acceptance procedures and automation
	ALC_CMS.4 Problem tracking CM coverage
	ALC_DEL.1 Delivery procedures
	ALC_DVS.1 Identification of security measures
	ALC_LCD.1 Developer defined life-cycle model
	ALC_TAT.1 Well defined development tools
	ALC_FLR.3 Systematic flaw remediation
ASE: Security Target evaluation	ASE_CCL.1 Conformance claims
	ASE_ECD.1 Extended components definition
	ASE_INT.1 ST introduction
	ASE_OBJ.2 Security objectives
	ASE_REQ.2 Derived security requirements
	ASE_SPD.1 Security problem definition
	ASE_TSS.1 TOE summary specification
ATE: Tests	ATE_COV.2 Analysis of coverage
	ATE_DPT.1 Testing: security enforcing modules ¹

¹ The previous Security Target listed ATE_DPT.2 instead of ATE_DPT.1. Including ATE_DPT.2 as part of EAL4 was a mistake in CC Version 3.1 prior to Release 3. CC Version 3.1 Release 3 now has corrected this flaw and therefore it has also been corrected in this Security Target.

Assurance Class	Assurance components
	ATE_FUN.1 Functional testing
	ATE_IND.2 Independent testing – sample
AVA: Vulnerability assessment	AVA_VAN.3 Focused vulnerability analysis

Table 10: Security assurance components (EAL 4 augmented by ALC_FLR.3)

6.3 Security requirements rationale

This section provides the rationale for the internal consistency and completeness of the security functional requirements defined in this Security Target.

6.3.1 Internal consistency of requirements

This section shows that the security functional requirements are internally consistent and mutually supportive:

FAU_GEN_SUB.1 (H) and FAU_GEN.2 (H) require that the TOE provides the capability to audit specific events related to Hyper-V R2 functions. The events are general failure conditions that either do not allow a correct start of the hypervisor or failure conditions detected by the hypervisor. In addition the creation and deletion of partitions shall generate an audit event.

This functionality is supported by the identification of the partitions (defined in FIA_ATD.1 (H) which defines the partition identifier as a security attribute and FIA_UAU.2 (H) which requires the identification of a partition). The audit trail management, protection and audit record evaluation are not performed by Hyper-V and therefore not claimed in this Security Target. Those functions are performed by the instance of Windows Server 2008 R2 executing in the root partition. Hyper-V uses the interfaces Windows Server 2008 R2 provides for placing audit records into the Windows Server 2008 R2 audit trail. As part of the functionality provided by this interface, Windows Server 2008 R2 places the time and date into the audit record while all other information in the Hyper-V related audit records is provided by Hyper-V.

The functionality related to the Partition Device Access Control Policy is defined in FDP_ACC.1 (H-DA) and FDP_ACF.1 (H-DA) which define the subjects, objects, types of access and access control rules. In addition FMT_MSA.3 (H) defines that restrictive default values (no access) are assigned for the security attributes (access control entries). As a result, for a partition to have access to a device that is controlled by the device access control policy, the right to access must be assigned by an administrator when creating or managing the partition configuration data. FMT_MSA.1 (H) defines that only an authorized administrator is allowed to query and modify the partition configuration data and FMT_REV.1 defines that the revocation of a partitions access to a device is restricted to an administrator. This restricts all management activities to authorized administrators. On the other hand, the authorizations of individual administrators themselves with respect to access to specific partition configuration data and management functions can be controlled by the partition management access control policy defined by FDP_ACC.1 (H-PM) and FDP_ACF.1 (H-PM). To be able to start this system of management restrictions

for authorized administrators, the TOE needs to start with an administrator that is not restricted and which can then define other administrative users with more restricted rights. This is defined in FMT_MSA.3 (H-PM), which requires permissive default values for the initial administrator.

The specific aspects of live migration have been addressed by the two requirements FPT_TDC.1 (H) and FPT_TEE.1 (H). FPT_TDC.1 (H) addresses the requirement that both the source and the destination instance of the TOE need to ensure that all interactions of a partition with the underlying hardware and the TOE is handled consistently. This requires that the real and virtualized environments a partition sees have to be fully compatible on both the source and the target. The only differences visible to a partition between the source and the destination instance of the TOE can be those that clearly have no security implication (e. g. timing differences). FPT_TEE.1 (H) ensures that a TOE tests for the required hardware compatibility before allowing a live migration operation to be started, thereby ensuring that there are no differences in the operation of the underlying hardware that would result in a potential inconsistency when running the subject to be migrated on the destination hardware.

6.3.2 Complete coverage: security objectives

This section demonstrates that the functional components selected for this profile provide complete coverage of the defined security objectives. The mapping of components to security objectives is depicted in the following table.

Security Objective	Related Security Functional Requirement
O.PARTITION_ACCESS	FDP_ACC.1 (H-DA) FDP_ACF.1 (H-DA) FIA_ATD.1 (H) FIA_UID.2 (H) FIA_USB.1 (H) FMT_MSA.3 (H-DA)
O.AUDIT_GENERATION	FAU_GEN_SUB.1 (H) FAU_GEN.2 (H)
O.AUTHORIZED_SUBJECT	FMT_MSA.1(H) FMT_MSA.3(H-PM) FMT_MSA.3(H-DA) FMT_MTD.1(H-2) FMT_MTD.1(H-3) FMT_REV.1 (H) FMT_SMF.1 (H)

Security Objective	Related Security Functional Requirement
O.INIT_SECURE_STATE	FMT_MSA.2 (H) FMT_MTD.1 (H-1)
O.MANAGE	FDP_ACC.1 (H-PM) FDP_ACF.1 (H-PM) FMT_MSA.1 (H) FMT_MSA.3 (H-PM) FMT_MTD.1 (H-2) FMT_MTD.1 (H-3) FMT_REV.1 (H) FMT_SMF.1 (H)
O.RESIDUAL_INFORMATION	FDP_RIP.1 (H)
O.RESOURCE_ALLOCATION	FRU_RSA.1 (H)
O.LIVE_MIGRATION	FPT_TDC.1 (H) FPT_TEE.1 (H)
O.PARTITION_ISOLATION	ADV_ARC ¹
O.PRESERVE_PART_PRIV	ADV_ARC ¹

Table 11: Mapping Security Objectives to Security Functional Requirements

¹ Following the CEM, security objectives would need to be mapped to SFRs contributing to satisfy the security objective. With objectives that define properties like O.PARTITION_ISOLATION and O.PRESERVE_PART_PRIV, this is not possible in the same way as for functional requirements. Properties need to be enforced by the TSF as whole and can not be simply mapped to a set of individual functions. CC Version 2.3 and earlier included the SFR families FPT_SEP and FPT_RVM to address the two properties of separation and complete reference mediation. SFRs from those families could be used to map the related objectives to. In CC Version 3 those families of SFRs have been replaced by the new assurance family of ADV_ARC. It seems therefore natural to now map the security objectives related to properties to the assurance class ADV_ARC.

O.PARTITION_ACCESS is fully addressed by the access control policy defined by FDP_ACC.1 (H-DA) and FDP_ACF.1 (H-DA), the identification of the partitions as defined by FIA_ATD.1 (H), FIA_UID.2 (H) and FIA_USB.1 (H), and the management of the access control policy defined in FMT_MSA.3 (H-DA).

O.AUDIT_GENERATION is fully addressed by the audit generation of the hypervisor as defined in FAU_GEN_SUB.1 and FAU_GEN.2.

O.AUTHORIZED_SUBJECT is addressed by the management functions of Hyper-V that require a check to validate the subject performing the management activity is actually allowed to do so. This includes the SFRs FMT_MSA.1(H), FMT_MSA.3(H-PM), FMT_MSA.3(H-DA), FMT_MTD.1(H-2), FMT_MTD.1(H-3), FMT_REV.1 (H), and FMT_SMF.1 (H).

O.INIT_SECURE_STATE is addressed by FMT_MSA.2 (H) together with FMT_MTD.1 (H-1).

O.MANAGE is addressed by the various SFRs of the Security Management class as well as the management access control policy (FDP_ACC.1 (H-PM) and FDP_ACF.1 (H-PM)) which defines the access rights and privileges an administrative user needs to have to perform management operations.

O.RESIDUAL_INFORMATION is addressed by FDP_RIP.1 (H).

O.RESOURCE_ALLOCATION is addressed by FRU_RSA.1 (H).

O.LIVE_MIGRATION is fully addressed by the combination of FPT_TDC.1 (H) and FPT_TEE.1 (H).

The two security objectives O.PARTITION_ISOLATION and O.PRESERVE_PART_PRIV define architectural properties and are therefore not mapped to SFRs but to the TOE architecture (ADV_ARC.1).

6.3.3 Security requirements coverage

The following table shows that each security functional requirement addresses at least one objective.

SFR	Security Objective addressed by the SFR
FAU_GEN_SUB.1 (H)	O.AUDIT_GENERATION
FAU_GEN.2 (H)	O.AUDIT_GENERATION
FDP_ACC.1 (H-DA)	O.PARTITION_ACCESS
FDP_ACC.1 (H-PM)	O.MANAGE
FDP_ACF.1 (H-DA)	O.PARTITION_ACCESS
FDP_ACF.1 (H-PM)	O.MANAGE
FDP_RIP.1 (H)	O.RESIDUAL_INFORMATION
FIA_ATD.1 (H)	O.PARTITION_ACCESS
FIA_UID.2 (H)	O.PARTITION_ACCESS
FIA_USB.1 (H)	O.PARTITION_ACCESS
FMT_MSA.1 (H)	O.MANAGE, O.AUTHORIZED_SUBJECT

SFR	Security Objective addressed by the SFR
FMT_MSA.2 (H)	O.INIT_SECURE_STATE
FMT_MSA.3 (H-PM)	O.MANAGE, O.AUTHORIZED_SUBJECT
FMT_MSA.3 (H-DA)	O.PARTITION_ACCESS, O.AUTHORIZED_SUBJECT
FMT_MTD.1 (H-1)	O.INIT_SECURE_STATE
FMT_MTD.1 (H-2)	O.MANAGE, O.AUTHORIZED_SUBJECT
FMT_MTD.1 (H-3)	O.MANAGE, O.AUTHORIZED_SUBJECT
FMT_REV.1 (H)	O.MANAGE, O.AUTHORIZED_SUBJECT
FMT_SMF.1 (H)	O.MANAGE, O.AUTHORIZED_SUBJECT
FPT_TDC.1 (H)	O.LIVE_MIGRATION
FPT_TEE.1 (H)	O.LIVE_MIGRATION
FRU_RSA.1 (H)	O.RESOURCE_ALLOCATION

Table 12: Mapping Security Functional Requirements to Security Objectives

This table shows that each security functional requirement contributes to satisfy at least one security objective.

6.3.4 Security requirements dependency analysis

The following table shows the dependencies which exist. In this table the SFRs are listed with their dependencies. In the case of multiple instantiations of SFRs it also shows, which of those SFRs actually resolves the dependency.

SFR	Dependency	Resolved?
FAU_GEN_SUB.1 (H)	none	Yes
FAU_GEN.2 (H)	FAU_GEN.1 FIA_UID.1	Yes (by FAU_GEN_SUB.1 (H). FAU_GEN_SUB.1 includes the user/subject identity in the same way as FAU_GEN.1 does, so the dependency is satisfied). Yes (by FIA_UID.2 (H) for

SFR	Dependency	Resolved?
		partitions) No (for human users. See the note) ¹
FDP_ACC.1 (H-PM)	FDP_ACF.1 (H-PM)	Yes
FDP_ACC.1 (H-DA)	FDP_ACF.1 (H-DA)	Yes
FDP_ACF.1 (H-PM)	FDP_ACC.1 (H-PM) FMT_MSA.3 (H-PM)	yes
FDP_ACF.1 (H-DA)	FDP_ACC.1 (H-DA) FMT_MSA.3 (H-DA)	yes
FDP_RIP.1 (H)	none	Yes
FIA_ATD.1 (H)	none	Yes
FIA_UID.2 (H)	none	Yes
FIA_USB.1 (H)	FIA_ATD.1 (H)	Yes
FMT_MSA.1 (H)	FDP_ACC.1 or FDP_IFC.1 FMT_SMR.1 FMT_SMF.1	Yes No ² Yes
FMT_MSA.2 (H)	FDP_ACC.1 or FDP_IFC.1 FMT_MSA.1 FMT_SMR.1	Yes Yes No ²
FMT_MSA.3 (H-PM)	FMT_MSA.1 FMT_SMR.1	Yes No ²
FMT_MSA.3 (H-DA)	FMT_MSA.1 FMT_SMR.1	Yes No ²
FMT_MTD.1 (H-1)	FMT_SMR.1 FMT_SMF.1	No ³ No ³
FMT_MTD.1 (H-2)	FMT_SMR.1 FMT_SMF.1	No ² Yes

SFR	Dependency	Resolved?
FMT_MTD.1 (H-3)	FMT_SMR.1 FMT_SMF.1	No ⁴ No ⁴
FMT_REV.1 (H)	FMT_SMR.1	No ²
FMT_SMF.1 (H)	none	Yes
FPT_TDC.1 (H)	none	Yes
FPT_TEE.1 (H)	none	Yes
FRU_RSA.1	none	Yes

Table 13: Dependencies between Security Functional Requirements

- 1 Human users that want to perform Hyper-V management operations are identified by the Windows Server 2008 R2 operating system operating in the root partition. Hyper-V obtains the identity of those users from Windows Server 2008 R2, includes those identities in audit records that require the identity of the administrator and also bases its management decisions on the identity obtained from Windows Server 2008 R2. So, while formally the dependency is not satisfied within the TOE itself, the dependency is satisfied by the IT environment.
- 2 As described the TOE supports roles that can be dynamically created by an installation. Those roles are groups of privileges for individual management operations. Since those roles are not predefined, it is neither possible nor useful to use FPT_SMR.1 which is designed to describe predefined roles that always exist for a TOE. So, while formally the dependency is not satisfied, the system still supports roles and allows those roles to be assigned to users.
- 3 The SFR is stated to describe that the specific aspect can not be managed (modified) during operation regardless of the privileges an administrator has. Therefore as the SFR is stated there is no dependency on any roles or privileges.
- 4 While the TOE provides functionality to manage the AzMan policy store, the protection of this store is enforced by the access control policy of Windows Server 2008 R2. This is the same approach as many application use to protect their TSF data stored in operating system protected storage objects.

6.3.5 Operations performed on the SFRs

The following table maps the SFRs as stated in this Security Target to the definitions of the components in part 2 of the CC to show how the operations on the components have been performed.

--	--

<p>FAU_GEN_SUB.1.1</p> <p>The TSF shall be able to generate an audit record of the following auditable events:</p> <ol style="list-style-type: none"> 1. All auditable events for the not specified level of audit; and 2. The following hypervisor specific events: <ul style="list-style-type: none"> • Failure to start the hypervisor • Creation of a partition (by the hypervisor) • Deletion of a partition (by the hypervisor) • Failure condition detected within the hypervisor 3. The following partition management specific events: <ul style="list-style-type: none"> • Access checks performed by AzMan on Hyper-V R2 management operations • Reconfigure partition (Virtual Machine) 	<p>FAU_GEN_SUB.1.1</p> <p>The TSF shall be able to generate an audit record of the following auditable events:</p> <ol style="list-style-type: none"> 1. All auditable events for the [selection, choose one of: <i>minimum, basic, detailed, not specified</i>] level of audit; and 2. [assignment: <i>other specifically defined auditable events</i>].
<p>FAU_GEN_SUB.1.2</p> <p>The TSF shall record within each audit record at least the following information:</p> <ol style="list-style-type: none"> a) type of event, subject identity (if applicable), and the outcome (success or failure) of the event; and b) For each audit event type, based on the auditable event definitions of the functional components included in the PP/ST, and no other security relevant information. 	<p>FAU_GEN_SUB.1.2</p> <p>The TSF shall record within each audit record at least the following information:</p> <ol style="list-style-type: none"> a) type of event, subject identity (if applicable), and the outcome (success or failure) of the event; and b) For each audit event type, based on the auditable event definitions of the functional components included in the PP/ST, [assignment: <i>other audit relevant information</i>].

<p>FAU_GEN.2.1</p> <p>For audit events resulting from actions of identified users, the TSF shall be able to associate each auditable event with the identity of the user (<u>partition or the administrator</u>) that caused the event.</p>	<p>FAU_GEN.2.1</p> <p>For audit events resulting from actions of identified users, the TSF shall be able to associate each auditable event with the identity of the user that caused the event.</p>
<p>FDP_ACC.1.1</p> <p>The TSF shall enforce the Partition Management Access Control Policy on Server 2008 R2 subjects acting on behalf of an administrative user, individual management operations as objects and performing the operation as functions.</p>	<p>FDP_ACC.1.1</p> <p>The TSF shall enforce the [assignment: <i>access control SFP</i>] on [assignment: <i>list of subjects, objects, and operations among subjects and objects covered by the SFP</i>].</p>
<p>FDP_ACC.1.1</p> <p>The TSF shall enforce the Partition Device Access Control Policy on partitions as subjects, virtualized and synthesized devices as objects and device access as function.</p>	<p>FDP_ACC.1.1</p> <p>The TSF shall enforce the [assignment: <i>access control SFP</i>] on [assignment: <i>list of subjects, objects, and operations among subjects and objects covered by the SFP</i>].</p>
<p>FDP_ACF.1.1</p> <p>The TSF shall enforce the Partition Management Access Control Policy to objects based on the following:</p> <p>The identity of the administrative user, his roles, the tasks and operations assigned to the roles and the additional BizRules assigned to the tasks.</p>	<p>FDP_ACF.1.1</p> <p>The TSF shall enforce the [assignment: <i>access control SFP</i>] to objects based on the following: [assignment: <i>list of subjects and objects controlled under the indicated SFP, and for each, the SFP-relevant security attributes, or named groups of SFP-relevant security attributes</i>].</p>
<p>FDP_ACF.1.2</p> <p>The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed: access (i. e performing an operation) is allowed if either the operation itself or a task that contains the operation is assigned to a role that has been assigned to the administrative user and the BizRules assigned</p>	<p>FDP_ACF.1.2</p> <p>The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed: [assignment: <i>rules governing access among controlled subjects and controlled objects using controlled operations on controlled objects</i>].</p>

<p>to one of tasks involved in the access decision don't disallows access.</p>	
<p>FDP_ACF.1.3</p> <p>The TSF shall explicitly authorize access of subjects to objects based on the following additional rules: none</p>	<p>FDP_ACF.1.3</p> <p>The TSF shall explicitly authorize access of subjects to objects based on the following additional rules: [assignment: <i>rules, based on security attributes, that explicitly _authorize access of subjects to objects</i>].</p>
<p>FDP_ACF.1.4</p> <p>The TSF shall explicitly deny access of subjects to objects based on the following additional rules: . no additional rules</p>	<p>FDP_ACF.1.4</p> <p>The TSF shall explicitly deny access of subjects to objects based on the following additional rules: [assignment: <i>rules, based on security attributes, that explicitly deny access of subjects to objects</i>].</p>
<p>FDP_ACF.1.1</p> <p>The TSF shall enforce the Partition Device Access Control Policy to objects based on the following:</p> <ul style="list-style-type: none"> a) the type of device b) the identity of the partition c) the configuration data for the partition. 	<p>FDP_ACF.1.1</p> <p>The TSF shall enforce the [assignment: <i>access control SFP</i>] to objects based on the following: [assignment: <i>list of subjects and objects controlled under the indicated SFP, and for each, the SFP-relevant security attributes, or named groups of SFP-relevant security attributes</i>].</p>
<p>FDP_ACF.1.2</p> <p>The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed: a partition is allowed to access a virtualized or synthesized device if the partition configuration has this device assigned to the partition.</p>	<p>FDP_ACF.1.2</p> <p>The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed: [assignment: <i>rules governing access among controlled subjects and controlled objects using controlled operations on controlled objects</i>].</p>
<p>FDP_ACF.1.3</p> <p>The TSF shall explicitly authorize access of subjects to objects based on the following additional rules: the device is a virtualized S3 Trio Video Card or a virtualized Intel 440 BX</p>	<p>FDP_ACF.1.3</p> <p>The TSF shall explicitly authorize access of subjects to objects based on the following additional rules: [assignment: <i>rules, based on security attributes, that explicitly authorize</i></p>

<p>chipset.</p>	<p><i>access of subjects to objects</i>].</p>
<p>FDP_ACF.1.4</p> <p>The TSF shall explicitly deny access of subjects to objects based on the following additional rules: no additional rule.</p>	<p>FDP_ACF.1.4</p> <p>The TSF shall explicitly deny access of subjects to objects based on the following additional rules: <i>[assignment: rules, based on security attributes, that explicitly deny access of subjects to objects]</i>.</p>
<p>FDP_RIP.1.1</p> <p>The TSF shall ensure that any previous information content of a resource is made unavailable upon the allocation of the resource to the following objects: virtual memory allocated to a partition, virtual and synthesized devices allocated to a partition, virtual processors allocated to a partition.</p>	<p>FDP_RIP.1.1</p> <p>The TSF shall ensure that any previous information content of a resource is made unavailable upon the <i>[selection: allocation of the resource to, deallocation of the resource from]</i> the following objects: <i>[assignment: list of objects]</i>.</p>
<p>FIA_ATD.1.1</p> <p>The TSF shall maintain the following list of security attributes belonging to individual users:</p> <p>a) For partitions:</p> <ul style="list-style-type: none"> a. Partition identifier (number) b. Partition configuration data c. Partition privileges <p>b) For administrative users:</p> <ul style="list-style-type: none"> a. none 	<p>FIA_ATD.1.1</p> <p>The TSF shall maintain the following list of security attributes belonging to individual users: <i>[assignment: list of security attributes]</i>.</p>
<p>FIA_UID.2.1</p> <p>The TSF shall require each user (partition) to be successfully identified before allowing any other TSF-mediated actions on behalf of that user.</p>	<p>FIA_UID.2.1</p> <p>The TSF shall require each user to be successfully identified before allowing any other TSF-mediated actions on behalf of that user.</p>
<p>FIA_USB.1.1</p> <p>The TSF shall associate the following user</p>	<p>FIA_USB.1.1</p> <p>The TSF shall associate the following user</p>

<p>security attributes with subjects acting on the behalf of that user:</p> <p>a) For subjects acting on behalf of partitions:</p> <ul style="list-style-type: none"> a. The identity of the partition b. The configuration data of the partition c. The privilege vector of the partition 	<p>security attributes with subjects acting on the behalf of that user: [assignment: <i>list of user security attributes</i>].</p>
<p>FIA_USB.1.2</p> <p>The TSF shall enforce the following rules on the initial association of user security attributes with subjects acting on the behalf of users:</p> <p>Every subject acting on behalf of a partition will be assigned the security attributes associated with the partition on whose behalf the subject will act.</p>	<p>FIA_USB.1.2</p> <p>The TSF shall enforce the following rules on the initial association of user security attributes with subjects acting on the behalf of users: [assignment: <i>rules for the initial association of attributes</i>].</p>
<p>FIA_USB.1.3</p> <p>The TSF shall enforce the following rules governing changes to the user security attributes associated with subjects acting on the behalf of users: subjects acting on behalf of partitions can not add additional security attributes beyond those initially assigned.</p>	<p>FIA_USB.1.3</p> <p>The TSF shall enforce the following rules governing changes to the user security attributes associated with subjects acting on the behalf of users: [assignment: <i>rules for the changing of attributes</i>].</p>
<p>FMT_MSA.1.1</p> <p>The TSF shall enforce the Partition Management Access Control Policy to restrict the ability to query and modify the security attributes partition configuration data to authorized administrators that have the authority for those operations assigned in the AzMan policy for Hyper-V R2.</p>	<p>FMT_MSA.1.1</p> <p>The TSF shall enforce the [assignment: <i>access control SFP(s), information flow control SFP(s)</i>] to restrict the ability to [selection: <i>change_default, query, modify, delete, [assignment: other operations]</i>] the security attributes [assignment: <i>list of security attributes</i>] to [assignment: <i>the _authorized identified roles</i>].</p>
<p>FMT_MSA.2.1</p> <p>The TSF shall ensure that only secure values</p>	<p>FMT_MSA.2.1</p> <p>The TSF shall ensure that only secure values</p>

<p>are accepted for partition configuration data.</p>	<p>are accepted for [assignment: <i>list of security attributes</i>].</p>
<p>FMT_MSA.3.1</p> <p>The TSF shall enforce the Partition Management Access Control Policy to provide restrictive default values for security attributes that are used to enforce the SFP.</p>	<p>FMT_MSA.3.1</p> <p>The TSF shall enforce the [assignment: <i>access control SFP, information flow control SFP</i>] to provide [selection, choose one of: <i>restrictive, permissive, [assignment: other property]</i>] default values for security attributes that are used to enforce the SFP.</p>
<p>FMT_MSA.3.2</p> <p>The TSF shall allow the administrator authorized by the Hyper-V R2 AzMan policy to modify partition configuration data to specify alternative initial values to override the default values when an object or information is created except for those security attributes that nobody is allowed to modify.</p>	<p>FMT_MSA.3.2</p> <p>The TSF shall allow the [assignment: <i>the _authorized identified roles</i>] to specify alternative initial values to override the default values when an object or information is created.</p>
<p>FMT_MTD.1.1</p> <p>The TSF shall restrict the ability to change_default and modify the partition's access to hypervisor calls related to inter-partition communication and the creation of guest partitions to nobody.</p>	<p>FMT_MTD.1.1</p> <p>The TSF shall restrict the ability to [selection: <i>change_default, query, modify, delete, clear, [assignment: other operations]</i>] the [assignment: <i>list of TSF data</i>] to [assignment: <i>the _authorized identified roles</i>].</p>
<p>FMT_MTD.1.1</p> <p>The TSF shall restrict the ability to change_default and modify the resource quota assigned to partition to administrators authorized by the Hyper-V R2 AzMan policy to perform those actions.</p>	<p>FMT_MTD.1.1</p> <p>The TSF shall restrict the ability to [selection: <i>change_default, query, modify, delete, clear, [assignment: other operations]</i>] the [assignment: <i>list of TSF data</i>] to [assignment: <i>the _authorized identified roles</i>].</p>
<p>FMT_MTD.1.1</p> <p>The TSF shall restrict the ability to modify the Hyper-V R2 AzMan policy to administrators authorized to access the Hyper-V R2 authorization store.</p>	<p>FMT_MTD.1.1</p> <p>The TSF shall restrict the ability to [selection: <i>change_default, query, modify, delete, clear, [assignment: other operations]</i>] the [assignment: <i>list of TSF data</i>] to [assignment:</p>

	<i>the_ uthorized identified roles].</i>
<p>FMT_REV.1.1</p> <p>The TSF shall restrict the ability to revoke the assignment of devices associated with the partitions under the control of the TSF to administrators.</p>	<p>FMT_REV.1.1</p> <p>The TSF shall restrict the ability to revoke [assignment: <i>list of security attributes</i>] associated with the [selection: <i>users, subjects, objects, [assignment: other additional resources]</i>] under the control of the TSF to [assignment: <i>the_ uthorized identified roles</i>].</p>
<p>FMT_REV.1.2</p> <p>The TSF shall enforce the rules that the administrator needs to be authorized by the Hyper-V R2 AzMan policy to perform those actions.</p>	<p>FMT_REV.1.2</p> <p>The TSF shall enforce the rules [assignment: <i>specification of revocation rules</i>].</p>
<p>FMT_SMF.1.1</p> <p>The TSF shall be capable of performing the following management functions:</p> <ul style="list-style-type: none"> • Management of partition configuration data • Management of virtual switches • Defining, deleting, starting and stopping partitions • Management of the Hyper-V R2 AzMan policy. 	<p>FMT_SMF.1.1</p> <p>The TSF shall be capable of performing the following management functions: [assignment: <i>list of management functions to be provided by the TSF</i>].</p>
<p>FPT_TDC.1.1</p> <p>The TSF shall provide the capability to consistently interpret the complete state of a partition when shared between the TSF and another trusted IT product <u>(which needs to be another instance of the TOE) when transferred as part of live migration.</u></p>	<p>FPT_TDC.1.1</p> <p>The TSF shall provide the capability to consistently interpret [assignment: <i>list of TSF data types</i>] when shared between the TSF and another trusted IT product.</p>
<p>FPT_TDC.1.2</p> <p>The TSF shall use the semantics of the underlying hardware platform and the semantics of the virtualized resources when</p>	<p>FPT_TDC.1.2</p> <p>The TSF shall use [assignment: <i>list of interpretation rules to be applied by the TSF</i>] when interpreting the TSF data from another</p>

interpreting the TSF data from another trusted IT product <u>that it received as part of a live migration operation.</u>	trusted IT product.
FPT_TEE.1.1 The TSF shall run a suite of tests when live migration of a partition is requested to check the fulfillment of the compatibility criteria required for live migration.	FPT_TEE.1.1 The TSF shall run a suite of tests [selection: <i>during initial start-up, periodically during normal operation, at the request of an _authorized user, [assignment: other conditions]</i>] to check the fulfillment of [assignment: <i>list of properties of the external entities</i>].
FPT_TEE.1.2 If the test fails, the TSF shall prohibit live migration of a partition.	PT_TEE.1.2 If the test fails, the TSF shall [assignment: <i>action(s)</i>].
FRU_RSA.1.1 The TSF shall enforce maximum quotas of the following resources: maximum CPU time per virtual CPU, maximum amount of the partition (guest) physical memory that partitions can use over a specified period of time.	FRU_RSA.1.1 The TSF shall enforce maximum quotas of the following resources: [assignment: <i>controlled resources</i>] that [selection: <i>individual user, defined group of users, subjects</i>] can use [selection: <i>simultaneously, over a specified period of time</i>].

Table 14: Mapping SFR components of this ST to the component definition in part 2

This mapping shows that all the operations have been performed correctly and marked accordingly.

6.4 TOE Summary Specification Rationale

6.4.1 Security functions justification

The TOE summary specification contains for each section that describes security functionality a part that maps the security functionality to the security functional requirements. This provides the link between the security functional requirements and the security functionality described in the TOE summary specification.

6.4.2 Assurance measures justification

The assurance measures are those defined by the EAL 4 assurance level, augmented by ALC_FLR.3. The selection of the EAL 4 assurance level is consistent with the assumed attack potential of the threat agents and is an assurance level used in the evaluation of products of the same kind.

7 TOE Summary Specification

This section describes the architecture and the security functions of Hyper-V R2. It starts with a description of architecture and the components followed by a description of the security functions.

7.1 Architecture of the TOE

The TOE consists of the following components:

- A small hypervisor responsible for the separation of the guest partitions and the management of virtual storage and virtual processors.
- A root partition (running the Server 2008 Server Core with some Hyper-V R2 specific additions) that performs management activities and device virtualization
- Optional TOE software loaded into individual partitions (Virtualization Service Clients, VMBus driver and Enlightenments). This software is executed within a partition and is therefore not part of the TSF.

In addition the TOE requires server hardware that satisfies the hardware requirements for Hyper-V R2 as laid out in chapter 1.

The following figure shows the general architecture of the TOE:

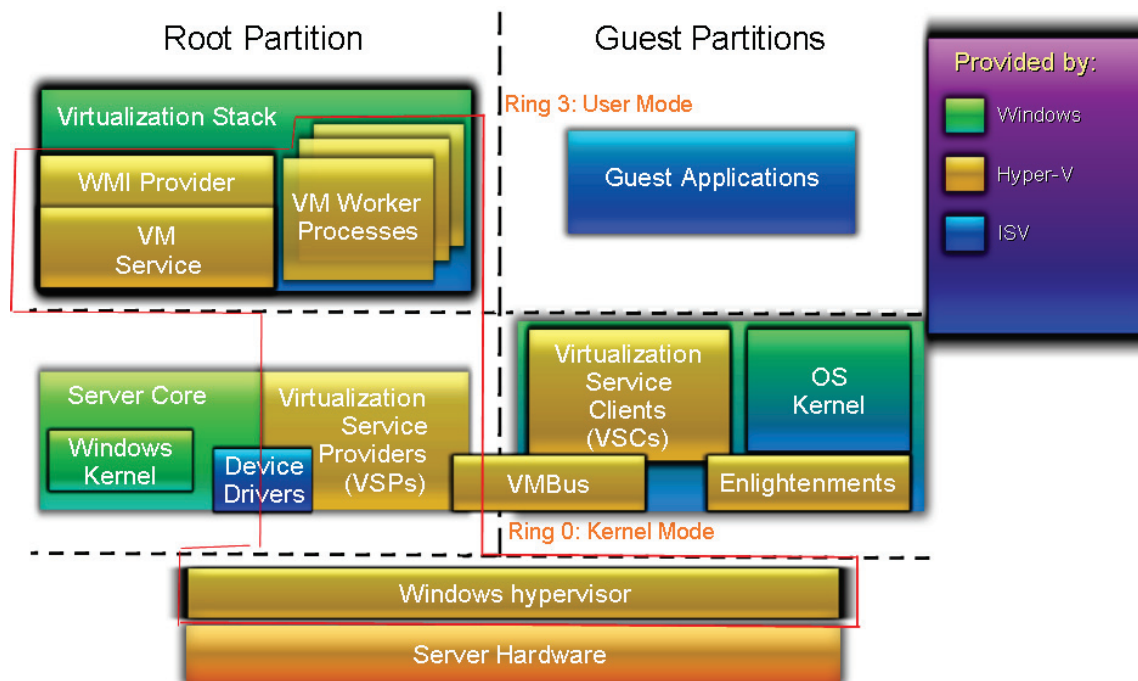


Figure 1: Hyper-V R2 architecture components

The figure shows the architectural components of the TOE with the red line indicating the TSF boundary. For simplicity the figure shows only one guest partition, but of course Hyper-V R2 supports multiple concurrent guest partitions. The example guest partition shown in the figure is for a partition that is “hypervisor aware” and has installed “Virtualizations Service Clients” (VSCs) and “Enlightenments”.

The colors in the figure indicate Hyper-V R2 specific software (gold), Server Core 2008 software (green), and (potentially) third party software (blue). Note that no part of the TSF contains third party software.

To be more precise the following is a list of the components of the picture above listing the components that are part of the TOE:

1. Windows hypervisor
this is the component that provides the hypervisor services, i. e. all the functions that execute not in the “virtualized” mode provided by the underlying processors. This component is part of the TOE and also part of the TSF.
2. Virtualization Service Providers (VSPs)
this component provides the virtualization services of Hyper-V that are integrated into the root partition. They are extension to the Windows Server 2008 R2 kernel executing in the root partition. They provide the interfaces of the root partition to the Windows hypervisor, the device virtualization functions for virtual hard disks/storage devices and virtual switches as well as support functions for the management of the partitions. This component is part of the TOE and part of the TSF. Note that the emulation of keyboard, mouse, video and other devices on a motherboard is performed within the VM worker process.
3. VMBus
this component provides a communication channel a guest partition can use to communicate with the root partition without going through the hypervisor. It is used to send requests from a guest partition to the root partition and exchange data between the guest partition and the root partition. VMBus is divided into a subcomponent executing as part of the root partition and a subcomponent executing in the guest partition. VMBus is part of the TOE but only the subcomponent executing in the root partition is part of the TSF.
4. VM Worker Processes
there is one instance of a VM worker process per guest partition. VM worker processes operate as applications on top of the operating system in the guest partition. The main task of the VM worker process is the emulation of the virtual motherboard with its (virtualized) devices and to perform live migration. VM worker processes are part of the TOE and also part of the TSF.
5. WMI Provider
WMI is the Windows Management Interface and a subset of this interface is dedicated to Hyper-V management. This Hyper-V WMI Provider is part of the TOE and also part of the TSF.

6. VM Service

this component implements the Hyper-V specific management functions including the access control to management functionality. This component is part of the TOE and part of the TSF.

7. Virtualization Service Clients (VSCs)

this component includes functions that use VMBus for communicating with root partition for efficient use of virtualized devices. This component is part of the TOE but since it executes as part of an untrusted guest partition it is not part of the TSF.

8. Enlightenments

this component is loaded by the kernel of Hyper-V “aware” operating system to use hypervisor calls where those provide more efficient operation than full platform emulation. The dedicated use of hypervisor calls can drastically reduce the number of traps into the hypervisor and thereby significantly enhance the performance of an operating system in a guest partition. This component is part of the TOE but since it executes as part of an untrusted guest partition it is not part of the TSF.

As part of its boot process an operating system for a X86-platform needs to identify the capabilities provided by the underlying hardware platform in order to load the correct device drivers and other platform dependent code. As a part of this processing the operating system will issue the CPUID instruction to identify the CPU family, the exact CPU type and the specific capabilities supported by the CPU. Hyper-V virtualizes this instruction and indicates in the values returned that the operating system is operating on Hyper-V and even indicates the version of Hyper-V used. This information allows a Hyper-V aware operating system to load the set of enlightenments for the version of Hyper-V it is executing on.

9. Device drivers

Those are not part of Hyper-V but part of Windows Server 2008 R2. The only exemption is the device driver for networking devices which has one specific function (VMQ) for the support of virtualization. VMQ support is part of Hyper-V and has been included in this evaluation as part of the Hyper-V TOE and is also part of the TSF.

All other parts of Windows Server 2008 R2 are part of the TOE environment.

The following section explains the purpose of the individual components shown in the figure:

7.1.1 Windows Hypervisor

By design, the hypervisor is a small and relatively simple piece of code. It runs on x64 systems and supports both 32-bit and 64-bit guests

The hypervisor runs in its own context and runs at a privilege level higher than that of guests. From this perspective, it makes sense to view the hypervisor as running at “ring -1” (where rings 0 through 3 are used for guests).

The hypervisor is designed with concurrency in mind, and it is a design goal to provide scalability to large numbers of processors.

In an effort to minimize complexity, the hypervisor is not preemptable. External interrupts and inter-processor interrupts are disabled while code is running within the hypervisor. (SMIs and NMIs will remain enabled.)

The hypervisor is logically divided into two distinct strata. The lower layer is a simple microkernel that supports memory allocation, threads, signaling, and a hardware abstraction mechanism. The second layer runs on top of the microkernel and provides virtualization services (partitions, virtual processors, address translation, hypercalls, etc.). This layer is necessarily platform-specific because it virtualizes the underlying processor and interrupt controller hardware.

Partitions

A *partition* is the basic unit of isolation supported by the hypervisor. A partition is made up of a physical address space and one or more virtual processors. Furthermore, a partition can be assigned specific hardware resources (memory, devices and CPU cycles) and certain permissions and access rights. Each partition has a unique partition identity (unique per re-boot of the hypervisor, not globally unique), which is represented as a 64-bit number.

Partitions also have a set of attributes assigned to them by the hypervisor. The security relevant attributes of a partition are:

- Amount of per virtual processor CPU time reserved for the partition
- Maximum amount of per virtual processor CPU time the partition is allowed to use
- Ability to create partitions (always disallowed for guest partitions in the evaluated configuration. This default value can not be changed)
- Access to memory pool related hypervisor calls
- Ability to initiate inter-process communication with other guest partitions (always disallowed for guest partitions in the evaluated configuration. This default value can not be changed)
- Access to debugging related hypervisor calls
- Access to CPU power management related hypervisor calls

Address Translation

The hypervisor provides a physical memory virtualization facility. This allows each partition to have a zero-based contiguous physical address space. Virtual processors support all x86 paging features and memory access modes (including large pages, global pages, write protection control, PAE, four-level page tables, paging disabled, etc.).

The hypervisor implements this behavior by means of two levels of address translation. The first level is controlled by the guest, allowing it to define a virtual address (VA) to guest physical address (GPA) translation. This is done via standard page tables maintained by the guest.

Second-level address translation is provided by the hypervisor without knowledge of the guest. This allows the hypervisor to virtualize physical memory, mapping guest physical addresses (GPAs) to system physical addresses (SPAs). The layout of a guest's physical address space is defined at partition creation time and can be modified at runtime. When hardware support for second-level address translation is available, the TOE will use it to optimize performance.

Virtual Processors

Each partition has one or more *virtual processors* associated with it. A virtual processor is a virtualized instance of an x86 processor complete with user-level and system-level registers.

A virtual processor is scheduled on a physical processor (more specifically, on a hardware thread or core). The hypervisor does not use hard affinities, so a virtual processor may move from one physical processor to another. The hypervisor schedules virtual processors according to specified scheduling policies and constraints.

Note that while virtual processors are similar to threads within a traditional kernel, they differ from threads in the sense that they don't have an assigned address space. Rather, the guest operating system is able to switch the address space of a virtual processor by modifying its page table base (i.e. reload CR3 on x86 processors).

The hypervisor virtualizes all modes of an x86 processor including real mode, 16-bit and 32-bit protected mode, v86, long mode, etc. If virtualization support is missing for one or more processor modes (e.g. real mode), the hypervisor will provide the necessary emulation layer.

The hypervisor does not support dynamic addition and removal of logical processors. All potential logical processors must be declared at boot time.

Invoking the hypervisor

If guest code is actively executing, there are three ways code can enter the hypervisor:

1. **Interrupts:** Asynchronous events including external interrupts, IPIs, NMIs, SMIs, SIPIs, etc. (see section below on Interrupt Routing and Delivery)
2. **Intercepts:** Synchronous events the hypervisor has requested to intercept within the guest
3. **Hypercalls:** A programmatic interface to the hypervisor (analogous to "kernel calls")

Interrupt Routing and Delivery

The hypervisor is responsible for routing interrupts to individual guests. All real hardware interrupts are delivered to the hypervisor first (regardless of whether interrupts are currently masked by the guest). The hypervisor can then decide whether to handle the interrupt itself or reflect it to a partition.

Interrupts are used to signal a variety of asynchronous events from assigned hardware devices, VSPs, TSPs, other virtual processors, etc. While some interrupts are targeted at specific virtual processors (e.g.

the virtual equivalent of IPIs), most interrupts are targeted at a partition. The hypervisor is responsible for routing the interrupt to an appropriate virtual processor.

From the perspective of a virtual processor, all interrupts delivered by the hypervisor are dispatched in the same way as traditional interrupts. On x86 processors, the behavior is defined by the guest's interrupt descriptor table (IDT). Most interrupts delivered by the hypervisor are maskable, so the guest can choose to hold off interrupts.

Intercept Handling, Redirection and Reflection

When the hypervisor is invoked due to some action performed by a guest (e.g. an access to an unmapped page, execution of a HLT instruction, etc.), the hypervisor can choose to do one of three things:

1. **Silently handle the intercept and return to the guest.** For example, the hypervisor may receive a page fault intercept because the page table entry in its shadow page table has not yet been populated despite the fact that the guest has already mapped the page in its page tables. In this case, the hypervisor would fill in the corresponding shadow page table entry and return back to the guest in such a way that the guest was unaware that the intercept occurred.
2. **Redirect the intercept to the partition related worker process in the root partition.** The worker process is a piece of code running in the root partition that is notified when an intercept occurs. This allows, for example, emulation of legacy devices. The hypervisor defines the events it would like to intercept. Examples on the x86 include: I/O port accesses, MSR accesses, CPUID, HLT, specific faults and exceptions, accesses to specified guest physical page ranges, and triple faults. Some intercepts (e. g. those related to address translation) are handled directly by the hypervisor. All other intercepts are forwarded to the worker processes to be handled there.
3. **Reflect the intercept back to the guest.** The hypervisor can simulate an exception within the guest. For example, if the hypervisor receives a page fault intercept and determines that the accessed page is unmapped within the guest's page tables, it can reflect the page fault back to the guest. The hypervisor also reflects the response it receives back from the worker process on intercepts

Device Ownership and Assignment

The hypervisor effectively "owns" the following hardware facilities: CPUs, memory, and APICs. All other hardware facilities and devices are assigned to the root partition.

Memory Intercepts

Intercepts can be requested for specific GPAs e. g. to simulate memory mapped I/O device register. The hypervisor supports read, write and execute intercepts. When a memory intercept is detected by the hypervisor, it sends an intercept message to the worker process related to the partition. The worker process typically responds by removing the intercept (e.g. paging in memory to back the GPA) and allowing the intercepted instruction to be re-executed. Alternatively, the external monitor can choose to

complete the intercepted instruction in software. This is required for certain emulation scenarios (e.g. MMIO, VGA frame buffer accesses, etc.).

7.1.2 VMBus

VMBus is a component that allows inter-partition communication between a guest partition and the root partition. It is implemented as a ring buffer shared between the guest and root partition but the VMBus protocol also allows defining a broader range of shared memory between the two partitions used to transfer large amount of data especially in response to I/O request from a guest partition. VMBus is implemented by a driver in both the root and the guest partition.

The VMBus driver in a guest partition is optional and requires the operating system to be “hypervisor aware” and have such a driver available. Since such a VMBus driver in a guest partition operates as a part of the host operating system in the guest partition, it is considered to be outside of the TSF.

Communication via VMBus is used to access “synthetic devices” for

- Storage
- Networking
- Video
- USB

Synthetic devices are implemented using specific communication protocols over VMBus. Those are:

- SCSI and iSCSI (RFC 3720)
- RNDIS (Remote Network Driver Interface Specification)
- RDP (Remote Desktop Protocol)

The “Virtual Service Provider” in the root partition intercepts the protocols and the commands sent via those protocols by a guest partition, interprets them and passes them to specific handler within the root partition to perform the real I/O operations.

7.1.3 Operating System Enlightenments (outside of the TSF)

An operating system can detect that it operates on Hyper-V R2 using the CPUID processor instruction. This instruction will be intercepted by the Hypervisor and returns a value indicating that the operating system is operating on a specific version of Hyper-V R2. The operating system can then use this information to install “Virtual Service Clients” that use the VMBus interface to communicate with the root partition for I/O operations and can install “Enlightenments” that make use of hypervisor calls to optimize the operation of the operating system on top of Hyper-V R2 (e. g. memory management).

Enlightenments operate as part of the guest operating system and are considered to be outside of the TSF.

7.1.4 Virtual Service Clients (outside of the TSF)

Virtual Service Clients are an optional part of a guest partition and provide an interface between the guest operating system’s device access functions and the VMBus driver. Requests within the guest partition to access a virtual device can be directed to a virtual service client for the device, which then sends the request to the root partition via the VMBus communication interface. The response received from the root partition is then passed back to the virtual service client, which passes this response (eventually with some transformation) back to the requestor.

Note that virtual service clients operate within a guest partition as part of the guest partition’s operating system and are therefore considered to be outside of the TSF.

7.1.5 Virtualization Service Provider

The virtualization service provider is part of the root partition and takes requests for virtual device access passed by a guest partition via VMBus (using the protocols supported by VMBus) and redirects those requests to the functions that interpret the commands and translate them into the I/O operations on the real devices attached to the TOE. The following figure shows the flow of an example for a request to a virtual hard disk.

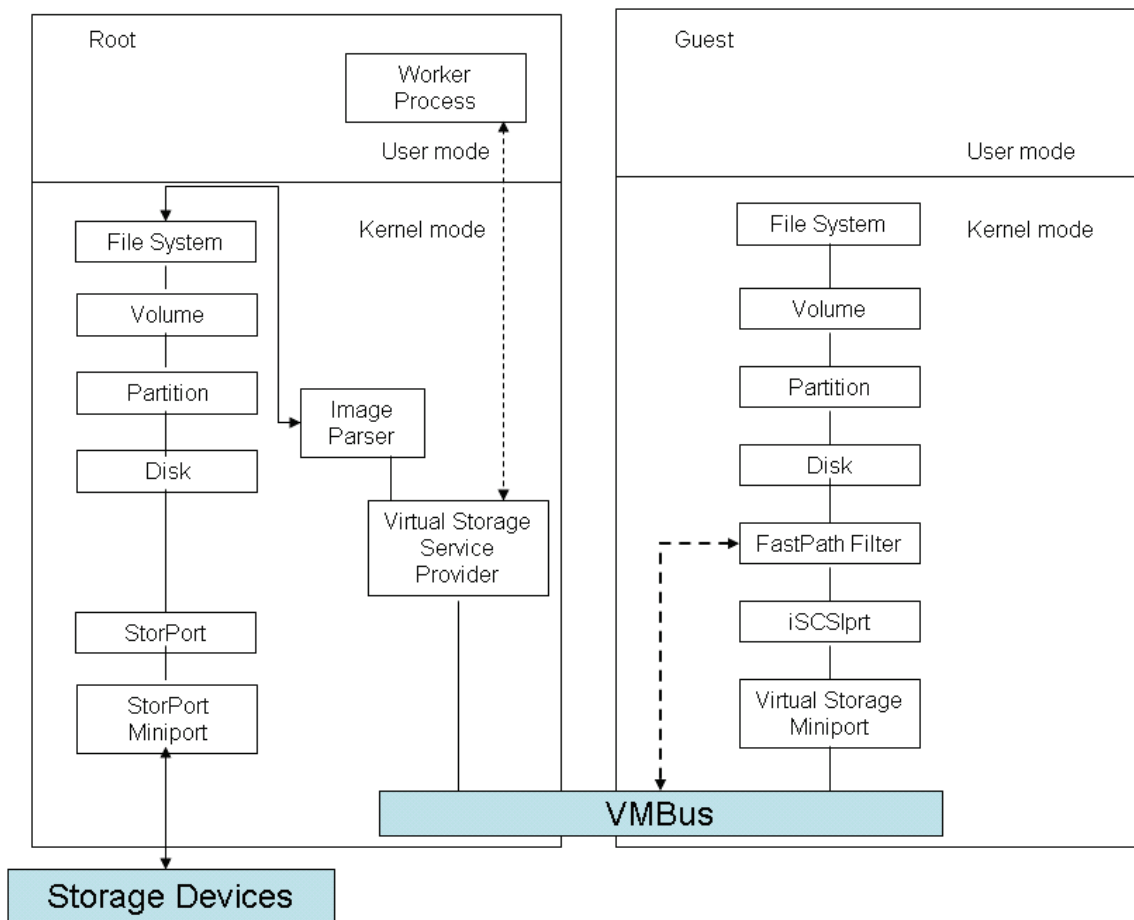


Figure 2: Example of a flow for access to a virtual hard disk (VHD)

In this example the request of an application within a guest partition is translated by the operating system within that guest partition into a request sent via VMBus to the root partition. Within the root partition this request is directed to the Virtual Storage Server which is the entry point of the Virtualization Service Provider for storage devices. Within the VSP and the Image Parser this request is translated to a request for access to a file that represents the virtual hard disk on the real storage media. In the case of a read request the data is read from this file and transferred back via the Virtual Storage Server and the VMBus to the guest partition, where the operating system passes this data back to the application.

7.1.6 VM Worker Processes

The VM worker processes are used for the virtualization of emulated devices. There is one separate VM worker process for each guest partition and all those VM worker processes operate as regular “user” processes on top of the Server 2008 R2 operating system in the root partition. Each VM worker process operates under a different SID, which represents the guest partition within the root partition.

When the hypervisor intercepts an instruction related to device emulation (like access to an I/O port or access to memory mapped I/O), it signals this instruction and the guest partition identity (and virtual processor) to the root partition. The root partition signals this to the worker process that corresponds to the guest partition. The worker process performs the emulation and sends the response back to the Server 2008 R2 kernel of the root partition (using specific ioctl interfaces), which in turn signals this response back to the hypervisor. The hypervisor takes the appropriate action to signal the result of the emulated instruction back to the guest partition.

A virtualized S3 Trio Video Card and a virtualized Intel 440 BX chipset are always available to a partition, providing the minimal virtualized hardware features that allow a partition to boot. The Intel 440 BX chipset provides the programming interfaces to the PCI bus, IDE controller, DMA controller, Interrupt controller, UART, floppy disk controller, keyboard, mouse, timer, real time clock, and power management. Note that even if the programming interfaces for all those devices exist, it depends on the configuration of the partition if all those programming interfaces are backed by virtualized devices (very much like in a real machine).

The VM worker process provides also a virtualized BIOS that is mapped into the guest partition.

7.1.7 WMI Provider

Windows Management Instrumentation (WMI) is the infrastructure for management data and operations on Windows-based operating systems. The WMI provider for virtualization and the hypervisor API enable developers, to quickly build custom tools, utilities, and enhancements for the virtualization platform. The WMI interfaces can manage defined aspects of the virtualization services.

Hyper-V R2 exposes a rich interface which permits the administrator to monitor and control the partition environment. Most of Hyper-V R2 can be controlled via the WMI provider, which allows easy,

yet powerful customization of the partitions. For further details on what can be controlled with the WMI provider, see the following topics:

- Virtual System Management Service
- Networking
- Resource Management

Virtual System Management Service

The Virtual System Profile describes the objects which make up a partition. These objects include the base system, the devices that make up the system, the settings for the system and its devices, and the management service that performs operations on the system.

The physical computer and its hosted partitions are each represented by the ComputerSystem class. The ComputerSystem instance representing the physical computer is associated, via the HostedDependency association, to the ComputerSystem instances representing the partitions on the physical computer.

Each partition is associated with a Virtual System Global Setting Data (VSGSD) instances and one or more Virtual System Setting Data (VSSD) instances. For each VSSD there is a series of Resource Allocation Setting Data (RASD) objects. The RASD objects describe the settings for each device in a partition. Together, the VSGSD, VSSDs and RASDs describe the configuration of the partition.

The VSGSD represents the global settings for a partition. These settings are independent of those described in the VSSD, and thus do not change if a snapshot is applied to the partition. The VSSD is used to describe the virtualization-specific settings of a partition. An instance of the VSSD may describe either the current settings for the partition, or the settings of the partition at the time of a snapshot.

None of the objects in the model support direct creation, modification or deletion through intrinsic Put or Delete operations. Instead, the Virtual System Management Service (VSMS) contains methods to manipulate the objects.

Networking

The networking profile describes the objects used for configuring the system to allow partitions to communicate over the network. The global networking objects, used to configure the network switch in the root partition, include the Msvm_VirtualSwitchManagementService, Msvm_VirtualSwitch, and Msvm_SwitchPort classes. The partition networking objects, used to configure the network interface card (NIC) in the partition, include the Msvm_EmulatedEthernetPort, Msvm_ResourceAllocationSettingData, Msvm_VmLANEndpoint and Msvm_SwitchLANEndpoint classes.

The root of the global networking profile is the Msvm_VirtualSwitch class. This class represents a virtual switch device in the root partition. Msvm_VirtualSwitch is associated with instances of the Msvm_SwitchPort class, which represents the ports on the virtual switch. Instances of the

Msvm_VirtualSwitch and Msvm_SwitchPort classes are created, deleted, and connected via the Msvm_VirtualSwitchManagementService class (not shown in the diagram above).

Virtual Switch Management Service (VSMS) represents the networking service present on a single Hyper-V R2 host and contains methods for Msvm_VirtualSwitchManagementService used to control the definition, modification, and destruction of global networking resources such as virtual switches, switch ports and internal Ethernet ports.

The representation of the ethernet NIC device in the partition looks very similar to that of any other device, as described in the Virtual System Management Service. The Msvm_EmulatedEthernetPort and Msvm_SyntheticEthernetPort classes represent the virtual NIC device, and are configured via an associated Resource Allocation Setting Data (RASD) instance. The only unusual characteristic of this representation is that, when the partition is instantiated and in turn creates the Msvm_EmulatedEthernetPort and Msvm_SyntheticEthernetPort devices, it also creates an associated Msvm_VmLANEndpoint instance for the virtual NIC. Similarly, when the partition is saved or turned off and the Msvm_EmulatedEthernetPort and Msvm_SyntheticEthernetPort instances are destroyed, the associated Msvm_VmLANEndpoint instance is also destroyed. The purpose of the Msvm_VmLANEndpoint is to serve as a bridge for connecting two networking ports to each other. In this case, it is used to connect a virtual NIC to a port on the virtual switch device. In other words, it connects the Msvm_EmulatedEthernetPort and Msvm_SyntheticEthernetPort instances on the partition to a particular Msvm_SwitchPort instance on the virtual switch. To connect a switch to the outside, you must bind the physical ethernet port to the Msvm_VirtualSwitch through BindExternalEthernetPort. Adversely, when connecting a switch to the host networking stack, or internal NIC, use ConnectInternal to have a partition talk to the host and not the outside world. Msvm_ActiveConnection connects a switch port to the Msvm_SwitchLanEndpoint to which the port is connected inside of Hyper-V R2. The existence of this object means that the switch port and the Msvm_SwitchLanEndpoint are actively connected and the ethernet port associated with Msvm_VMLanEndPoint can communicate with the network through the switch port.

Resource Management

The Resource Virtualization Profile provides the means by which a client can discover the virtual resources supported by the virtualization system. It also describes the capacity – or number of allocations – that is supported for each type of virtual resource.

Two different classes of virtual resources are defined by the Resource Virtualization Profile:

- Shared Resource: Represents the resources managed by Hyper-V R2 that are, or are capable of being shared among multiple partitions. Msvm_Processor is an example of a shared resource.
- Synthetic Resource: Represents the virtual resources that have no corresponding real resource. Msvm_SyntheticEthernetPort is an example of a synthetic resource.

The resource pool is used to collect a class of Hyper-V R2 managed resources so that it can be easily discovered while its capabilities and settings can be described in a central location.

From the resource pool, one can access the associated Allocation Capabilities (AC). This class describes the capabilities of the resource described by this resource pool. For instance, it may indicate whether the `Msvm_EmulatedEthernetPort` represented by this resource pool supports virtual LANs (VLANs) or filters.

The AC Profile defines the means by which one can discover the valid range of and default settings for a given virtual resource. An AC object is associated with each resource pool. Four Resource Allocation Setting Data (RASD) objects are associated with the AC object to describe the minimum, maximum, default and incremental values for the given resource's allocation. Together, these classes describe the overall range of supported capabilities. The `Msvm_AllocationCapabilities` instance provides an anchor point for the set of `Msvm_ResourceAllocationSettingData` instances that specify the default and valid range of settings for a virtual resource. The `Msvm_SettingsDefineCapabilities` association class provides the link between the AC instance and the minimum, maximum, incremental and default settings for a resource supported by the virtualization platform.

7.1.8 AzMan as the framework for access control to management operations of Hyper-V R2

Authorization Manager (AzMan) is a set of COM-based runtime interfaces that allow applications to easily manage and verify a user's request to access application defined resources (usually operations defined by an application).

AzMan allows to group operations to "tasks", define roles as a collection of permissions to tasks and/or operations, and assign roles to users. In addition one can define "BizRules", which are scripts that are evaluated when the access check is performed. Those allow evaluating specific conditions at runtime like the time of day. BizRules can be associated with tasks.

AzMan uses an "authorization store" to store the policy including the tasks, roles and assignments of roles to users. The authorization store in the evaluated configuration is either a XML file stored on the local disk or stored in an Active Directory database to which the Server 2008 R2 instance in the root partition is connected.

AzMan allows auditing modifications to the policy store as well as all access checks performed.

AzMan also allows defining a "scope" as a collection of resources to which a policy applies.

7.2 Security functionality provided by Hyper-V R2

This section presents the security functionality of Hyper-V R2 specific components and the mapping to the SFRs as defined the previous chapter of this Security Target.

In general Hyper-V R2 implements the following security functionality (using the headings from part 2 of the CC):

- Generation of Hyper-V R2 specific audit records
- User and TSF data protection - protection of partition configuration data, partition resources and devices allocated to a partition

- Identification of partitions
- Management of partition configuration
- Resource utilization quota management and enforcement

7.2.1 Audit functionality

The hypervisor layer as well as the Hyper-V R2 specific components in the root partition are able to generate auditable events. Those events are stored and managed by the Server 2008 R2 audit function. Security relevant events generated by the hypervisor are signaled to the root partition and the root partition generates and records the audit record for those events. The security relevant events signaled by the hypervisor are:

- The creation of a partition
- The deletion of a partition
- A failure condition detected internal to the hypervisor component

In addition to those events the Hyper-V R2 specific components within the root partition are capable to generate audit records for the following events:

- Access checks performed by AzMan

Each audit records is sent to the Event Logger in the Server 2008 R2 Server Core in the root partition, which inserts the time and data and stores the record in the event log.

Note that modifications to the AzMan policy for Hyper-V R2 are audited by the IT environment.

Mapping to SFRs

This implements the security functional requirements FAU_GEN_SUB.1 (H) and FAU_GEN.2 (H).

7.2.2 User and TSF data protection functionality

The Hyper-V R2 user data protection function performs

- Discretionary access control of administrators to administration objects
- Discretionary access control of partitions to virtualized and synthetic devices
- Residual data protection

Since some Hyper-V R2 TSF data is stored in storage objects those are also protected using the Server 2008 R2 data protection functions and not functions of the Hyper-V R2 specific parts. In the case of direct access to the AzMan policy store this is the only protection (very much like any application that stores its configuration data in files provided by the operating system).

Mapping to SFRs

This implements the security functional requirements FDP_ACC.1 (H-PM), FDP_ACC.1 (H_DA), FDP_ACF.1 (H_PM), FDP_ACF.1 (H_DA) and FDP_RIP.1 (H).

7.2.3 Identification functionality

The TOE requires each partition to be identified and this is ensured by the TSF which assigns an identifier to a partition. This identifier is used to identify the partition when it traps into the hypervisor and when a communication channel to the root partition (via VMBus) is established. Since guest partitions are created and fully controlled by the TSF, authentication of the partitions is not necessary. When a partition is started, the privileges and the virtualized devices are assigned to the partition in accordance with the partition's configuration data. Some privileges like the privilege to create partitions or the privilege to create ports that might be used for direct communication between guest partitions are not assigned to any guest partition and this default value can not be changed.

Mapping to SFRs

This implements the security functional requirements FIA_ATD.1 (H), FIA_UID.2 (H) and FIA_USB.1 (H)

7.2.4 Security Management Functionality

Hyper-V R2 can be managed either directly using the management functions within the root partition, or remotely using a client that connects to the WMI and performs the management activities using this interface. In both cases the administrative user that wants to perform management activities is authenticated as a regular user by the Server 2008 R2 authentication function. The Hyper-V related management activities are then enforced by Hyper-V with respect to the Hyper-V privileges assigned to that user.

Hyper-V R2 specific configuration data (like the AzMan policy store) is stored in objects that are protected by the standard Server 2008 R2 access control functions. Direct access to those objects is therefore managed by the access control management functions of Server 2008 R2.

Hyper-V R2 uses the "Authorization Manager" (AzMan) to define a role-based access control model for managing Hyper-V R2. The Hyper-V specific components in the root partition protects a set of Hyper-V R2 management operations. A default scope is provided as well as the set of operations that can be controlled. Tasks and roles can be defined by an application that uses the AzMan API for the management of the policy store. The command-line scripts HVRemote.wsf and HVRoles.wsf can be used to manage the authorization store of the AzMan policy for Hyper-V R2 locally in the root partition.

The authorization store for Hyper-V R2 can be found in the file "C:\ProgramData\Microsoft\Windows\Hyper-V R2\Initialstore.xml" in the root partition. This file needs to be selected when defining or editing role definitions, defining or editing task definitions, or defining or editing scopes. Roles can then be assigned to users, defining the management activities they are allowed to perform.

The authorization policy for Hyper-V R2 has a set of operations which can be used to group to tasks and/or be assigned to roles. Those are:

- Operations on Hyper-V R2 services
 - Read Service Configuration
 - Reconfigure Service
 - View Virtual Switch Management Service
- Operations on Virtual Machines
 - Create Virtual Machine
 - Delete Virtual Machine
 - Change Virtual Machine Authorization Scope
 - Start Virtual Machine
 - Stop Virtual Machine
 - Pause and Restart Virtual Machine
 - Reconfigure Virtual Machine
 - View Virtual Machine Configuration
 - Allow Input to Virtual Machine
 - Allow Output from Virtual Machine
- Operations on Virtual Switches
 - Create Virtual Switch
 - Delete Virtual Switch
 - Modify Switch Settings
 - View Switches
 - Create Virtual Switch Port
 - Delete Virtual Switch Port
 - Connect Virtual Switch Port
 - Disconnect Virtual Switch Port
 - Modify Switch Port Settings

- View Switch Ports
- View LAN Endpoints
- Change VLAN Configuration Port
- View VLAN Settings
- Create Internal Ethernet Port
- Delete Internal Ethernet Port
- View Internal Ethernet Port
- Bind External Ethernet Port
- Unbind External Ethernet Port
- View External Ethernet Ports

Whenever a user attempts to perform an administrative operation on Hyper-V R2, the AzMan interfaces for checking access are called to validate the user's rights to perform the management function.

In addition AzMan may use an authorization store in an Active Directory database the Server 2008 R2 instance in the root partition is connected to.

Mapping to SFRs

This implements the security functional requirements FMT_MSA.1 (H), FMT_MSA.2 (H), FMT_MSA.3 (H-PM), FMT_MSA.3 (H-DA), FMT_MTD.1 (H-1), FMT_MTD.1 (H-2), FMT_MTD.1 (H-3), FMT_REV.1 (H), FMT_SMF.1 (H)

7.2.5 TSF Protection Functionality

Hyper-V R2 protects its TSF code and data by maintaining separate address spaces for those that are protected from access or unmediated interference by untrusted subjects or subjects outside of the TOE that may communicate with the TOE via its network interfaces. So the main TSF protection functionality of the TOE is the maintenance of this separate address spaces and the mediation of all references to TSF data and protected user data.

Mapping to SFRs

TSF Protection Functionality is implemented by the TOE architecture.

7.2.6 TSF Protection during Live Migration

Hyper-V R2 allows under for live migration of a running partition from one system running Hyper-V R2 to another system running Hyper-V R2. The partition does not need to be stopped for this operation and the transition itself will be transparent to the partition being migrated. The migrated partition will have

the same security attributes and the migrated partition will not see a different interface or a different behavior of the underlying platform (hardware and Hyper-V R2) from the source system.

In order to allow for live migration, Hyper-V R2 checks for a number of prerequisites that need to be satisfied to successfully migrate a partition. In summary, those include a check of the underlying hardware platform for compatibility as well as a check that all virtual hard disks of the partition to be migrated are implemented on cluster shared volumes which allow both the source and the destination system to provide the same access to those virtual hard disks.

Mapping to SFRs

TSF Protection during Live Migration Functionality is implemented by the security functional requirements FPT_TDC.1 (H) and FPT_TEE.1 (H).

7.2.7 Resource utilization functionality

Hyper-V R2 provides functions that allow an authorized administrator to define the maximum amount of memory and CPU time a guest partition is allowed to use. Hyper-V R2 controls the use of those resources and ensures that a given guest partition does not use more of those resources than allocated to the partition. This prohibits that a single guest partition consumes all resources of a specific type and thereby cause a denial of service for other guest partitions.

Mapping to SFRs

This implements the security functional requirement FRU_RSA.1 (H)