

Enabling the New Era of Cloud Computing: Data Security, Transfer, and Management

Yushi Shen
Microsoft Corporation, USA

Yale Li
Microsoft Corporation, USA

Ling Wu
EMC² Corporation, USA

Shaofeng Liu
Microsoft Corporation, USA

Qian Wen
Endronic Corp, USA

A volume in the Advances in Systems
Analysis, Software Engineering, and High
Performance Computing (ASASEHPC)
Book Series



Managing Director:	Lindsay Johnston
Editorial Director:	Myla Merkel
Production Manager:	Jennifer Yoder
Publishing Systems Analyst:	Adrienne Freeland
Development Editor:	Austin DeMarco
Acquisitions Editor:	Kayla Wolfe
Typesetter:	Lisandro Gonzalez
Cover Design:	Jason Mull

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: <http://www.igi-global.com>

Copyright © 2014 by IGI Global. All rights reserved. No part of this publication may be reproduced, stored or distributed in any form or by any means, electronic or mechanical, including photocopying, without written permission from the publisher. Product or company names used in this set are for identification purposes only. Inclusion of the names of the products or companies does not indicate a claim of ownership by Yushi Shen, Yale Li, Ling Wu, Shaofeng Liu, Qian Wen, and IGI Global of the trademark or registered trademark.

Library of Congress Cataloging-in-Publication Data

Shen, Yushi, 1978-
Enabling the new era of cloud computing : data security, transfer, and management / by Yushi Shen, Yale Li, Ling Wu, Shaofeng Liu, and Qian Wen.
pages cm
Includes bibliographical references and index.
Summary: "This book discusses cloud computing as an emerging technology and its critical role in the IT industry upgrade and economic development in the future"-- Provided by publisher.
ISBN 978-1-4666-4801-2 (hardcover) -- ISBN 978-1-4666-4802-9 (ebook) -- ISBN 978-1-4666-4803-6 (print & perpetual access) 1. Cloud computing. I. Title.

QA76.585.S54 2014
004.67'82--dc23

2013027879

This book is published in the IGI Global book series Advances in Systems Analysis, Software Engineering, and High Performance Computing (ASASEHPC) (ISSN: 2327-3453; eISSN: 2327-3461)

British Cataloguing in Publication Data

A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

For electronic access to this publication, please contact: eresources@igi-global.com.

machine, that even if the physical machine that runs multiple virtual machines has failed all together, the virtual machines can be failed over to other physical machines immediately. (EMC² Corporation, Cloud computing foundations)

Hypervisor

The hypervisor is a software that does server virtualization. It enables multiple operating systems to run concurrently on a physical host computer, and to interact directly with the physical resources of the host computer. Hypervisor provides the attributes for the physical server that lies underneath the virtualized machines, running different operating systems. Hypervisor is the primary component of virtualization that enables computer system to partition hardware resources, such as CPU and memory, into virtualized resources.


Hypervisor has two components: the kernel and the virtual machine manager. The kernel works as the operating system, handling such tasks as process creation, file system management,

resources scheduling, IO stack etc. The virtual machine monitor (VMM), which resides below the operating system layer, is responsible for handling and sending the virtual machines' requests, also executing commands. When a virtual machine is created, resources such CPU, memory and I/O devices are assigned to the virtual machine. To execute processes, these resources need to be managed according to a time schedule on the physical machine. The VMM handles these requests and communications from the virtual level down to the physical level. The VMM's job also includes allocating and managing the system processor, memory, IO devices and other hardware resources that correspond to each individual virtual machine. When a virtual machine starts running, the controls are transferred to the VMM.

There are chiefly two kinds of Hypervisor: the bare-metal hypervisor and the hosted hypervisor.

- For the bare-Metal hypervisor, the hypervisor runs directly on the hardware. The Hypervisor itself functions as an operat-

*Figure 1. Bare-Metal Hypervisor
(EMC² Corporation - Virtualized data center and cloud infrastructure)*



ing system, resides on ring0 processor, and executes commands against the hardware. This type of hypervisor requires certified hardware, so that appropriate drivers are available to communicate with the hardware. Since the bare-metal hypervisor is directly installed on the X86 based hardware, it could access the hardware resources more efficiently, and is scalable. When databases or ERP applications are being deployed in a production environment, the bare-metal hypervisor is most likely to be used, because it has much less overhead, and more hardware resources can be delicate to the application that runs on the virtual machine. The bare-metal hypervisor is the most predominant hypervisor, being used in the virtualized data centers. It is also the direction of the cloud virtualization.

- The hosted hypervisor is a hypervisor that runs inside the operating system. It is installed and run as an application on top of an operating system. Since it is running on top of an operating system, it supports a broader range of hardware configurations. One may have the Windows OS or Linux installed on the host machine, then VMware workstation or Microsoft Hyper-V can be installed and run as an application within the operating system environment. Instead of the hypervisor being at the operating system level, it is another application, and other applications can be running within the hypervisor application.

The hosted hypervisor focuses on the development process. For a developer using a windows OS machine, but needs to have the Linux environment to develop an application, Linux can be installed in the virtual machine and development done on the same laptop, while other applications continue to run in the Windows environment.

Types of Computer Virtualization


The X86 CPU architecture offers four levels of privilege known as ring0, 1, 2 and 3. In the traditional X86 architecture, operating system kernels expect direct CPU access running in Ring 0, which is the most privileged level. With virtualization, the virtual machine monitor can sit on Ring 0, and the guest operating systems sit on top of the VMM, so that the VMM can interact with physical resources and the guest operating systems.

In Brief, virtualization acts as an operating system. The operating system sits on the highest CPU level, which is ring0. Applications typically do not interact with hardware directly; they usually interact with the operating system for recourse and command executions. The user applications typically run in ring 3 with less privilege. So the challenge for virtualization is that the hypervisor needs to control the lower levels of privilege. The virtualization technique enables the hypervisor to sit on the lowest level of the processor, in order to interact with the physical hardware, and mask the operating system from having to see itself.

In full virtualization, the VMM sits below the operating system in Ring 0, emulates the underlying physical resources, and presents them to the guest operating system. The guest operating system is expected to sit in ring 0, the virtualization technique makes it believe that it is actually sitting in the higher ring with less privileges to the processor architecture. The guest operating system on the virtual machine is unaware that it is being virtualized. The host operating system might think that it is sitting on the lowest Ring 0 level of the processor architecture, but in reality it is actually sitting on the top of the hypervisor. The hypervisor can completely decouple the guest operating system from the underlying hardware.

All the commands are executed at the hypervisor level. The kernel is doing the interaction with the physical hardware, while the VMM is passing

Figure 2. Hosted Hypervisor
(EMC² Corporation - Virtualized data center and cloud infrastructure)



the guest operating system, doing the binary translation of the commands through hypervisor down to the physical hardware that lies underneath. All the commands, such as handling, timer controls, IOs, are executed at the hypervisor level, and the virtual machine is communicating through the virtual machine manager.

In full virtualization, if the console is opened up before powering up the virtual machine, the virtual machine BIOS setting is to come up. VMware ESX, ESXI and Microsoft Hyper-V that runs in the server core environment are some examples. Please be aware that the Microsoft Hyper-V can be run as an application within the windows environment. In a special Windows Server Core installation, which installs the most basic components, the Hyper-V server role can be installed, which distinguishes the operating system to be a

virtual machine itself, and layers the hypervisor underneath it. This installation makes Windows Hyper-V similar to the infrastructure layer as ESX in VMware. VMware and Microsoft are market leaders in the full virtualization technologies.

Para-virtualization is also called the OS assisted virtualization. In Para-virtualization, the operating system is aware of itself being virtualized. The guest operating system sits in Ring 0 with the Hypervisor beneath it. Rather than the hypervisor sitting on that level and doing all the translation for the virtual machine monitor, the Para-virtualization guest operating system sits there and interact directly with the hypervisor. Para-virtualization product examples are the open source Xen hypervisor and VMware Linux.

Hardware assisted virtualization introduces virtualization in the X86 processor architecture,

and uses hypervisor-aware CPU to provide assistance to the hypervisor. With hardware assisted virtualization, the operating system could directly access the physical resources. It allows a fully virtual operating system to run in Ring 0, which gives the operating system direct access to the system resources without the control from the Virtual Machine Monitor. In 2006, Intel and AMD have released the first generation hardware assisted features, and have announced future development roadmaps, which includes hardware support for memory virtualization, as well as for I/O devices. The Intel Virtualization Technology (VT) and AMD virtualization (AMD V) both target privileged instructions with a new CPU execution mode feature, allowing the Virtual Machine Monitor to run in a new root mode below ring 0. They are building the chip set that increases CPU overhead, and allows the virtualization of the X86 instruction sets, while decreasing the hypervisor overhead. (EMC² Corporation - Cloud infrastructure and services.)

The graph below describes the Ring level for all three types of virtualizations:

Virtual Machine

Virtual Machine Files

A virtual machine is made of a set of files that reside in the underlying hypervisor file system. A virtual machine could be looked at from two different perspectives. From the user's perspective, a virtual machine is a generic set of hardware that runs an operating system and user applications. Like a physical machine, it has the same component as the physical machine, such as the CPU, memory, hard disk, network interface cards and other IO devices. From a hypervisor perspective, which manages and runs a virtual machine, a virtual machine is a set of files that include configuration files, virtual disk files, virtual BIOS files, virtual machine swap file and a log file.

The virtual machine configuration file stores the details of virtual machine configuration information, such as the virtual machine name, guest operating system, the number and types of virtual disks, the number of CPUs and the size of memory, the type of network adaptors and the associated MAC addresses, SCSI controller types and disk types.

The virtual disk file stores the contents in the disk drive belonging the virtual machine, and sits on the hypervisor file system. To the virtual machine, the virtual disk file appears and works as a physical disk drive. A virtual machine could have multiple virtual disk files, each representing a single disk.

Virtual BIOS files store the virtual machine's BIOS information. The virtual machine swap file is the paging file for the virtual machine, which backs up the virtual machine's RAM contents. This file is present only when the virtual machine is running. If we have allocated certain amount of memory and reserved the memory, the swap file is to show this difference.

The log file records virtual machine activities, such as the time the virtual machine got started, its activities and etc. It helps in troubleshooting when there is an application failure or some server problems.

File Systems in the Virtual Environment

The hypervisor has a file system. The virtual machine file system (VMFS) is a clustered file system that stores virtual machine files. Cluster SANS, internal direct attached storage, and external storage are presented to the hypervisor as virtual disks, formatted and attached to the server. The virtual disks are stored as files on a VMFS. VMFS allows multiple virtual machines to concurrently read and write data from the same storage device. The virtual machine file system is deployed on the FC and iSCSI storage, apart from local storages. Because we have the underlining cluster