MANAGEMENT OF TASKS IN VIRTUALIZED ENVIRONMENTS

Asil ALKAYA
Adnan Menderes University / Faculty of Business Administration
Title Asst. Prof. Dr.
E-mail: asil.alkaya@adu.edu.tr

─Abstract ─

Virtualization has opened a wide new range of opportunities for managing computer resources since it allows isolating different tasks in a single node without any interference with others. Moreover, a virtual machine is easier to monitor and offers a powerful but simple interface for managing its allocated resources which brings new chances for resource management. In addition to this, machine maintenance acquires a new level thanks to virtualization since it allows managing machines as software making it a perfect environment for executing tasks and complex services.

Nevertheless, dealing with some virtualization capabilities, such as the creation, implies an effort for the user in order to take benefit from them. In order to avoid this problem, task management should be achieved which allows executing tasks and providing virtualized environments to the users without any extra effort in an efficient way. This is a virtualized environment management which aims to provide virtual machines that fulfills with the user requirements in terms of software and system capabilities.

Key Words: *Virtualization, task management, network*
JEL Classification: C890

1. INTRODUCTION

Task management is a moving target. The task is made more complex if the network is distributed amongst in-house, remote and 'in the cloud' infrastructure. And then, on top of that, virtualized servers and desktops that don't increase the hardware footprint but increase the number of available services. But while the network management tasks become more demanding, more organizations are looking to cut costs and rationalize resources.

Configuring and monitoring the virtualized environment offers extensive capabilities for configuring each virtual server and its hosted virtual machines. In order to bring the servers into the overall operations; applications, networking and storage, management tools must be able to interact with the physical and virtual components seamlessly. This is where the 'virtual management tools' come in.

2. VIRTUALIZATION

Abstraction of the hardware resources into multiple execution environments comes from need to make more effective use of hardware. Two main approaches exist:

- Full virtualization: Hardware is *completely* emulated by the virtual machine
- Paravirtualization: The virtual machine provides an application programming interface (API) and the guest operating system is modified to run on the virtual machine

2.1. Benefits of virtualization

Virtualization is not just used only for development and testing purposes but also used for consolidating physical servers. It's a change in information technology infrastructure so that:

- Creates hardware independence and mobility
- Isolation from conflicts and service availability

- Manage downtime and disaster recovery

2.2. History

International Business Machines (IBM) developed Control Program-40 (CP-40) which emulated the System/360 architecture for multiple users in 1964 and then released VM/370 for the System/370 which included virtual machine support, real device support and greater hardware exploitation. IBM also developed versions of MVS, UNIX, DOS/VSE and PC/DOS to run under VM in 1972.Through 1970's, virtualization is eclipsed by microcomputers. In 1981, IBM announced Extended Architecture (XA) which, among other things, had specialized input-output (I/O) processors that were part of the hardware. VMWare Workstation is released in 1999 and VMware ESX Server is released in 2001. Up to 2003, all the virtual machines prepared for commercial use only, but by announcing the first public release of XEN server, an open source platform is intended to be used for academic and educational purposes also. Sun Microsystems announced the Sun xVM in 2007 and VirtualBox in 2008.

2.1. Hypervisor technology

A hypervisor, also called a virtual machine manager, is a program that allows multiple operating systems to share a single hardware host. The hypervisor controls the host processor and resources and ensures that each 'virtual machine' get enough resources to do what they do.

A popular method of virtualization is paravirtualization using a hypervisor to manage the guest operating system (OS) also called Virtual Machine Monitor (VMM) The term hypervisor comes from the hyper call made by the guest OS to the virtual machine which is similar to a supervisor call made by an operating system to the kernel.

The hypervisor manages the operation levels of the guest OS by creating a virtual kernel mode and virtual user mode. Privileged instructions are paravirtualized and are validated and executed by the hypervisor on either the hardware or the host OS.

2.2. Hardware virtualization assistance

The Popek and Goldberg Formal Requirements for Virtualizable Third Generation Architectures are a set of requirements for sufficient hardware virtualization:

- **Equivalence –** A program running under VMM should exhibit the same behaviour if run on the machine directly
- **Resource Control –** The VMM should be in complete control of the virtualized resources
- **Efficiency** – Major of machine instructions should be allowed to execute with VMM intervention

3. VIRTUAL ENVIRONMENT

A virtual environment mainly composed by three different layers: the data infrastructure, the node management (VRMM), and the global Scheduler. The data infrastructure offers a distributed storage for supporting virtualization capabilities such as migration and checkpoint support. VRMM is in charge of creating and maintaining the whole virtual machine lifecycle. Finally, the scheduling layer is in charge of distributing tasks and VMs among the physical nodes.

3.1. Features

Features of a virtual environment include:

- Create customized VM in order to execute tasks
- Tasks can run in a full customized environment
  - Pre-created tasks such as: Tomcat server. mySQL server
- Use of different schedulers with different features
- Abstract from virtualization difficulties
  - Creation
  - Migration
  - Checkpointing
- Efficiency
- Easy to deploy

- Modular and easy to extent

3.2. Virtualization layer

Virtual machines are developed for mainframe systems to allow multiple users to safely share those expensive machines. As computers became cheaper, the motivation behind virtualization faded and processor architectures like the Intel x86 were developed without some of the features needed to support virtualization.

Figure-1: Comparison of layers with and without virtulization

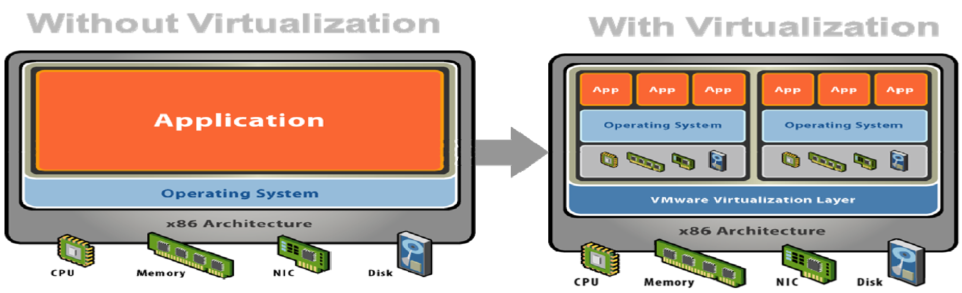

When problems like server proliferation and the need to run multiple applications in dedicated operating systems started becoming serious issues for IT managers and software developers, .the virtual machine concept has been taken into consideration. Virtualization technology provides each virtual machine with a true representation of an x86 and amd64 computers, complete with processor, memory, networking interfaces and storage devices by the virtualization layer that gives the virtual machines direct access to the underlying x86 hardware – an important distinction from the much slower emulation technology that must process all virtual machine operations in software.

3.2. Virtualization server

After consolidating servers into virtual machines, a huge improvement in server utilization can be measured by average CPU usage. IT administrators generally dedicate a single x86 server per application because of issues like application compatibility problems when running on a single OS image, software vendor requirements or internal billing constraints. The result of that “one application, one server” practice is a datacenter filled to capacity with servers running at 5-15% utilization. With multiple servers consolidated on a virtualization Server host, servers can be operated reliably at 60-80% CPU utilization – a huge improvement in the return on investment in their server hardware. This higher CPU utilization is possible only with the intelligent CPU resource controls made possible by the bare-metal architecture of virtualization server. Normally, running a server at high CPU utilization creates the risk of a critical service being starved for CPU if others experience a load peak. Virtualization server eliminates that risk with its CPU resource controls that assign CPU capacity to virtual machines on a “fair share” basis and it also allows an absolute minimum level of CPU capacity to be allocated to critical virtual machines

3.3 Key properties of virtual machines

The hypervisor partitions a server into multiple virtual machines. VMs are basically just files and can be treated like files. This gives VMs amazing properties that then enable a much broader set of capabilities beyond just running many VMs on one server, not possible in the physical world. These properties are:

1) Partitioning: With hypervisor, it is able to run many different OS-es on the same machine
2) Isolation: These VMs running on a physical machine are independent and unaware of each other. Infecting one doesn’t mean affecting all the others.
3) Encapsulation: Since the entire OS and the applications are packaged up into a set of files – manipulation of VMs becomes very much simpler than their physical counterparts
4) Hardware independence: Virtual machines can be created on any x86 and amd64 hardware and moved to any other x86 and amd64 hardware; in many cases live

Advantages of deploying multiple virtual machines on a single physical server include: performance, stability, scalability and cross-platform support.

Figure-2: Provisioning activities of physical and virtual environments

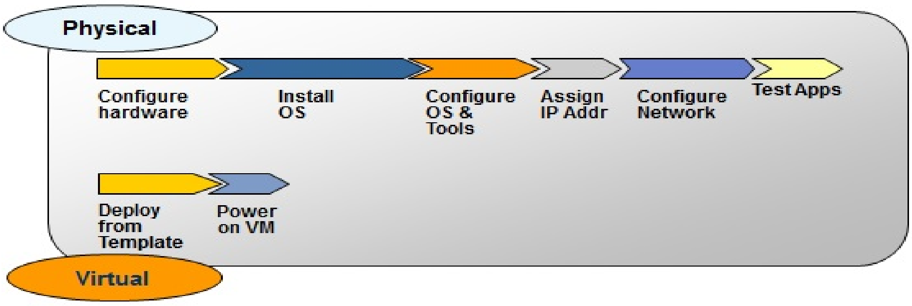

Table 1: Performance metrics for a virtual environment

Host	Service	Status	Duration	Attempt	Last Check	Status Information
localhost	Root Partition	Ok	2h 40m 55s	1/4	2013-04-13 05:07:18	DISK OK - free space: / 5552 MB (76% inode=89%):
	HTTP	Ok	2h 42m 10s	1/4	2013-04-13 05:06:03	HTTP OK HTTP/1.1 200 OK - 2732 bytes in 0.002 seconds
	Current Load	Ok	2h 43m 25s	1/4	2013-04-13 05:04:48	OK - load average: 0.40, 0.46, 0.33
	PING	Ok	2h 41m 32s	1/4	2013-04-13 05:06:41	PING OK - Packet loss = 0%, RTA = 0.07 ms
	Total Processes	Ok	2h 39m 2s	1/4	2013-04-13 05:04:12	PROCS OK: 80 processes with STATE = RSZDT
	SSH	Ok	2h 40m 17s	1/4	2013-04-13 05:07:56	SSH OK - OpenSSH_5.3 (protocol 2.0)
	Swap Usage	Ok	2h 39m 40s	1/4	2013-04-13 05:03:36	SWAP OK - 100% free (2015 MB out of 2015 MB)

Host	Service	Status	Duration	Attempt	Last Check	Status Information
	Current Users	Ok	2h 42m 47s	1/4	2013-04-13 05:05:26	USERS OK - 0 users currently logged in

Figure-3: Root partitioning

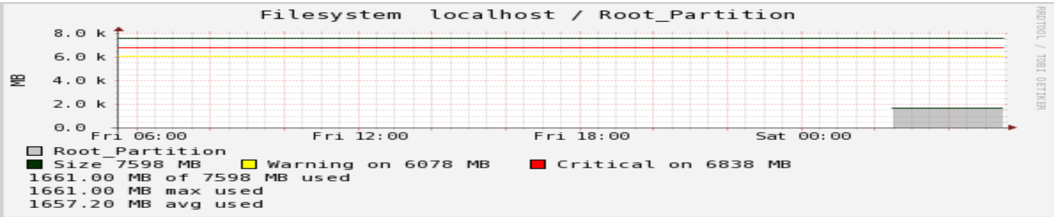

Figure-4: CPU load

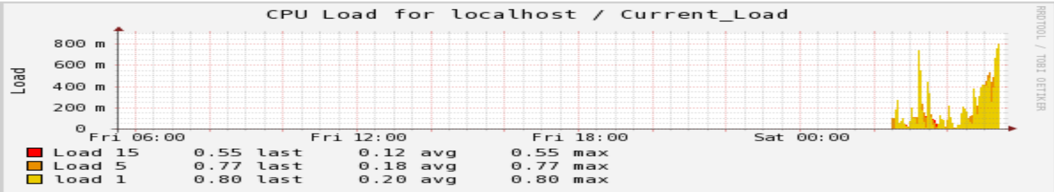

Figure-5: Ping time between virtual environments

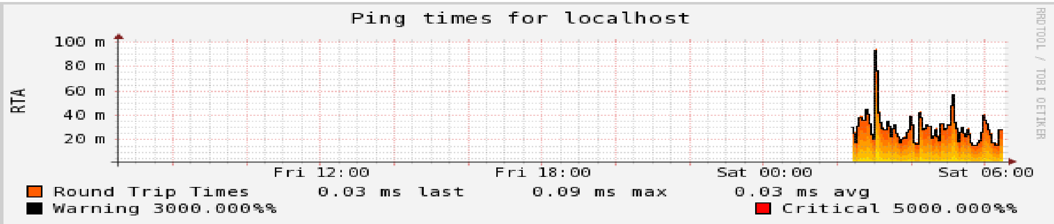

Figure-6: Swap usage

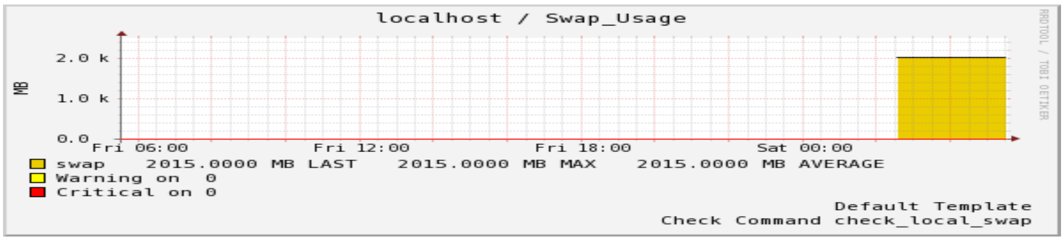

4. CONCLUSION

Operations center includes hosts and services. Tactical overview is about network health (connectivity and up/down hosts and services), event handlers and notification services.In virtual environments, detection of open hosts and services have been taken into consideration. Services that are expected to be inspected should include root partitioning, http web server, current load on host(s), ping (communication within the virtual environment), ssh server, swap usage control for the localhost (main operating system) and user control.Performance metrics can be shown such as disk usage, cpu usage, memory usage, load and swap space utilization. Performance graphs are can be specialized on hosts (localhost) and graph explorer can be built for visual experience.

There are pitfalls implementing the virtual environment such as hardware failure, over or under use of resources, related operational processors need to be executed and skills shortage (configuration, tuning and troubleshooting). So, firms should prepare a management plan and make use of information technology infrastructure library which establishes management complexity to install virtual environments to the physical servers.

BIBLIOGRAPHY

Gerald J. Popek and Robert P. Goldberg (1974). "Formal Requirements for Virtualizable Third Generation Architectures". Communications of the ACM 17 (7): 412 –421.

N.M.M.K. Chowdhury, R. Boutaba, Network virtualization: state of the art and research challenges, IEEE Communications Magazine 47 (7) (2009) 20–26

J.D. Touch, Y.-S. Wang, L. Eggert, G. Finn, A Virtual Internet Architecture, Tech. Rep. TR-570, USC/Information Sciences Institute (2003).

N. Fujita, J.D. Touch, V. Pingali, Y.-S. Wang, A dynamic topology and routing management strategy for virtual IP networks, IEICE Transactions on Communications E89-B (9) (2006) 2375–2384.

A. Sundararaj, P. Dinda, Towards virtual networks for virtual machine grid computing, in: Proceedings of the Third USENIX Virtual Machine Research and Technology Symposium (VM’04), 2004, pp. 177–190.

A. Jun, A. Leon-Garcia, A virtual network approach to network resources management, in: Proceedings of the Canadian Conference on Broadband Research (CCBR’98), 1998..

A. Sundararaj, P. Dinda, Towards virtual networks for virtual machine grid computing, in: Proceedings of the Third USENIX Virtual Machine Research and Technology Symposium (VM’04), 2004, pp. 177–190.

M. Feridan, M. Moser, A. Tanner, Building an abstraction layer for management systems integration, in: Proceedings of the First IEEE/ IFIP International Workshop on End-to-End Virtualization and Grid Management (EVGM’2007), 2007, pp. 57–60