IaaS, Private Cloud și Virtualizare

UBB, 2023.03.15

Tudor Damian

- **IT Consultant** & **Trainer**
  - Cloud Strategy & Governance
  - Cybersecurity
  - IT Risk Management
  - Business Process Optimization
  - Digital Transformation
- **Co-founder** @ ITCamp & ITCamp Community
- **Contact**: tudy.ro

ITCAMP

C|EH™ Certified Ethical Hacker

MVP Microsoft Most Valuable Professional

Agenda

- Private Cloud și IaaS - introducere
- Arhitecturi de virtualizare
- Soluții existente
  - Type 1 (bare-metal)
    - Full / paravirtualized (VMWare, Hyper-V, Xen)
  - Type 2 (hosted)
    - OS-assisted (KVM, VirtualBox, Virtuozzo/OpenVZ)
  - Bonus: containers
- Studiu de caz: Hyper-V & System Center
- Azure Stack
- PowerShell

Private Cloud și IaaS

Sursă imagine: wordle.net

dacă la început lumea
era ~~destul de~~ reticentă...

… de câțiva ani încoace, ~~toți~~ mulți vor să se mute…

...dar știu toți la ce să se aștepte?

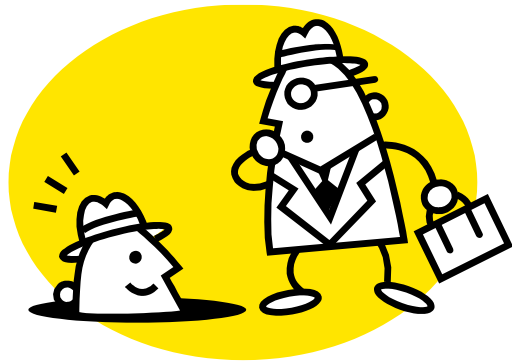

Varianta tradițională, ineficientă

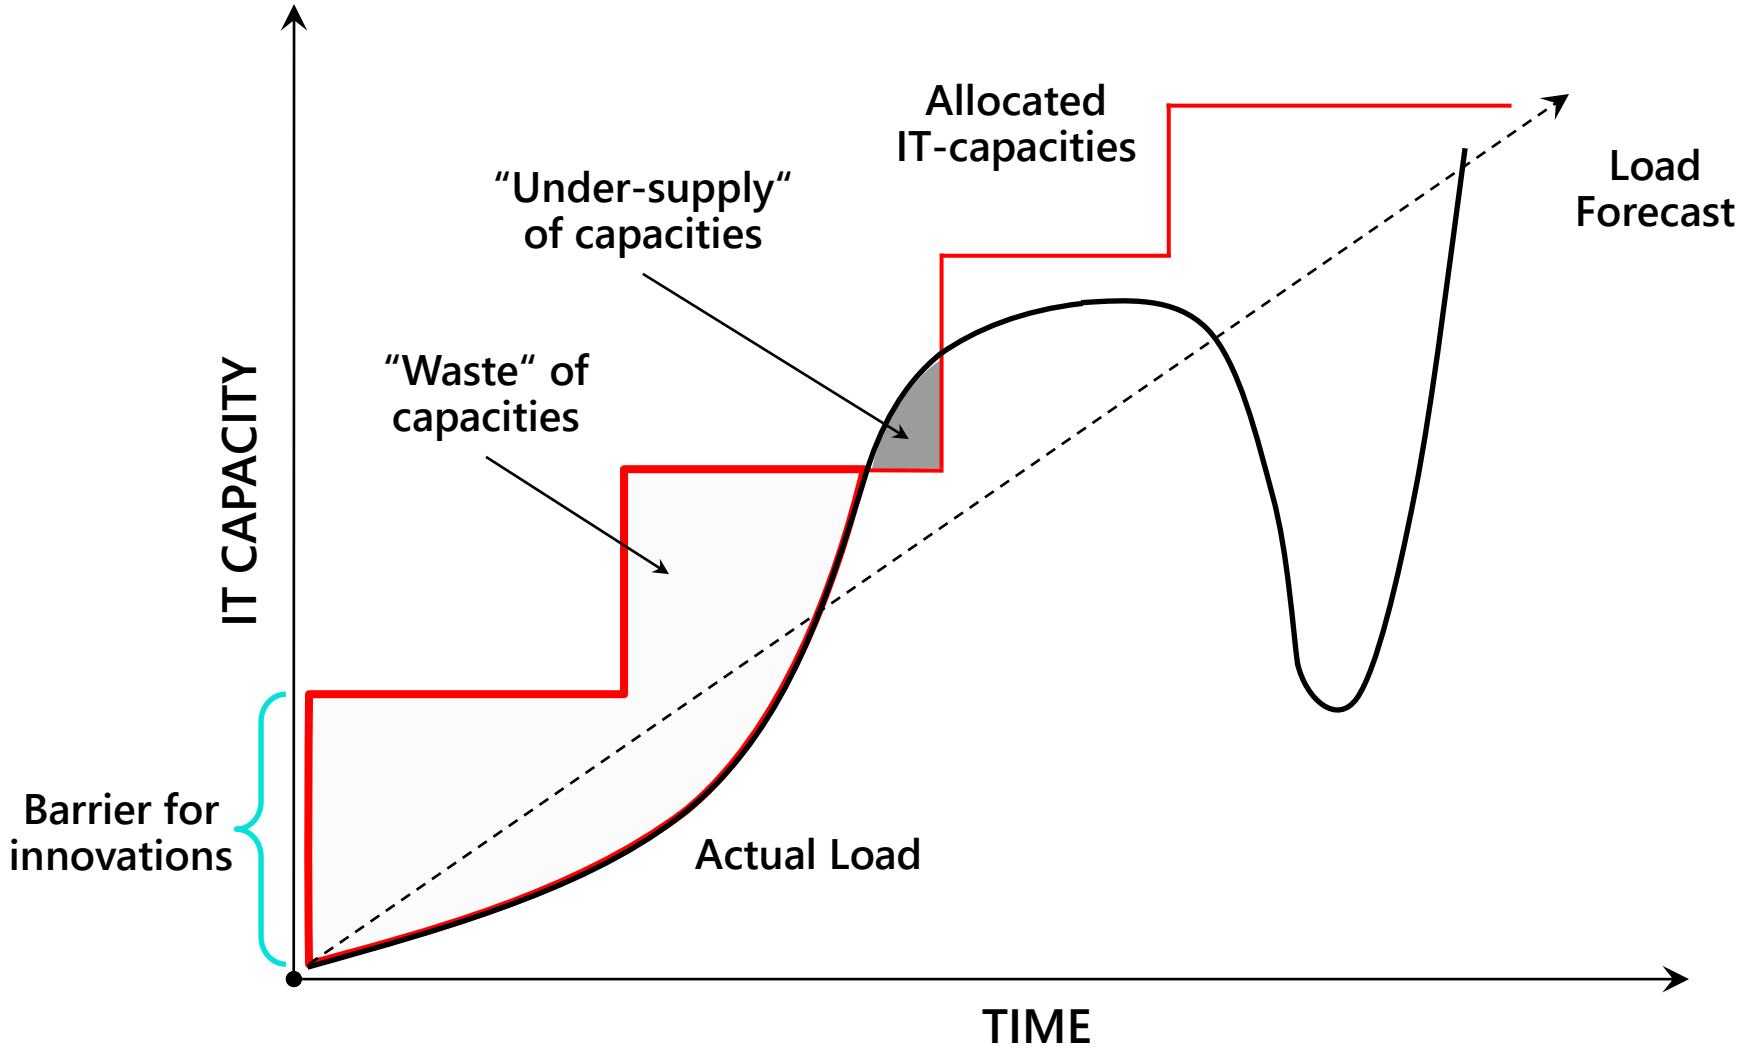

Source: Microsoft Cloud Continuum

Într-un Cloud, totul e mult mai dinamic

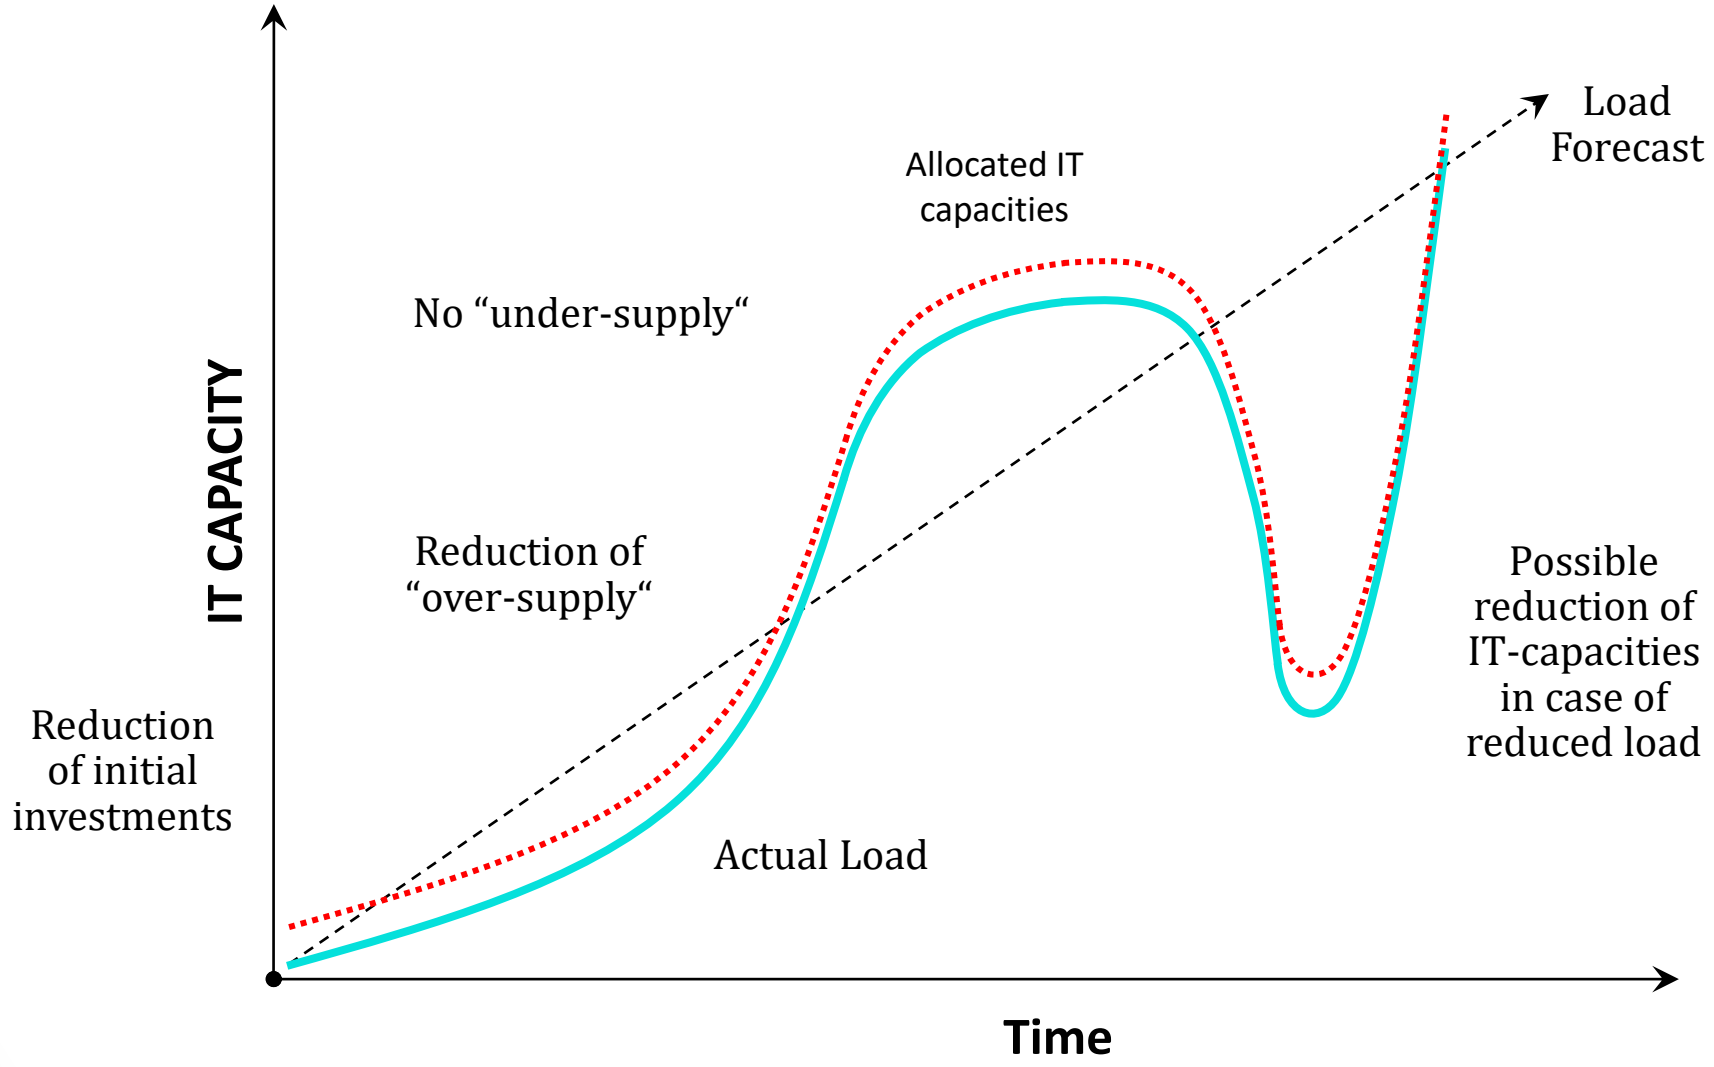

Source: Microsoft Cloud Continuum

Workload patterns în Cloud

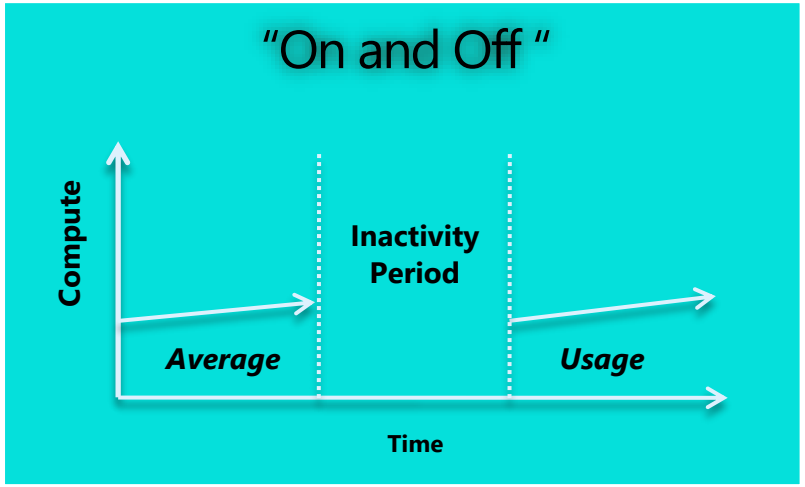

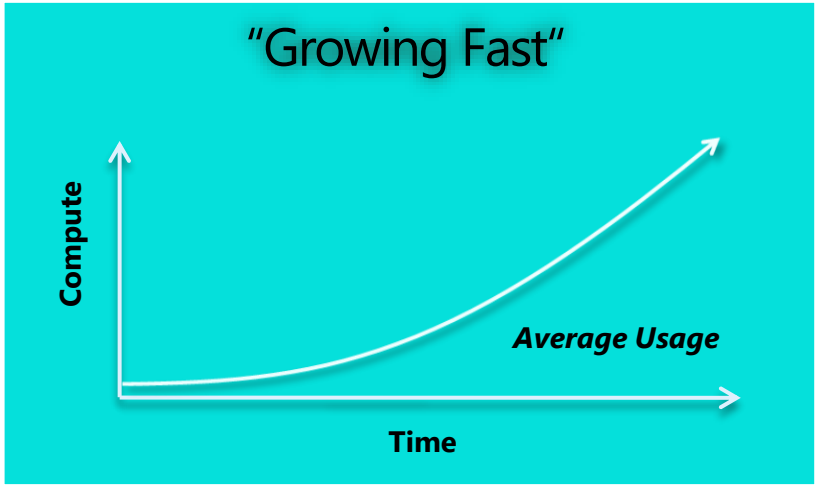

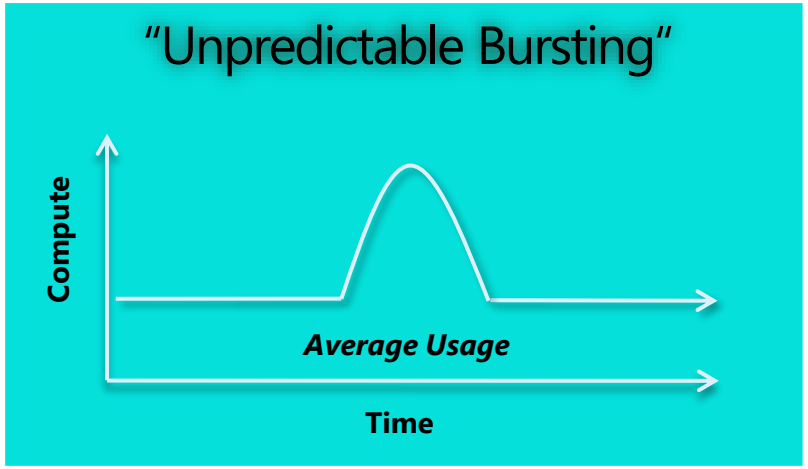

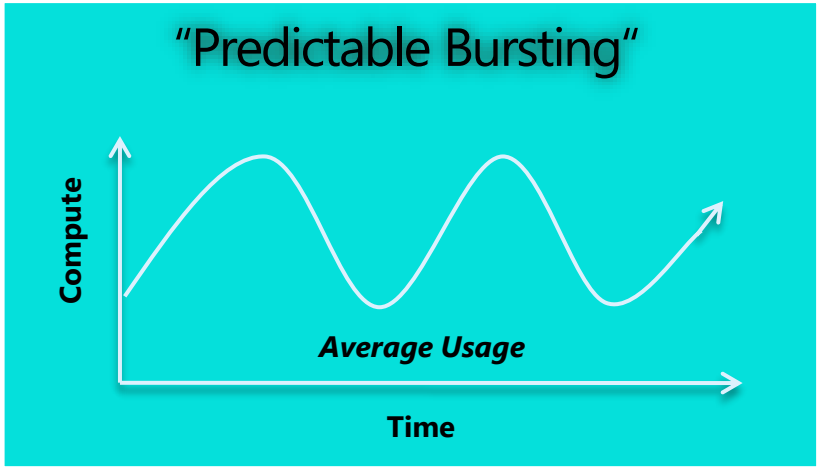

Source: Microsoft Cloud Continuum

Tipuri de servicii

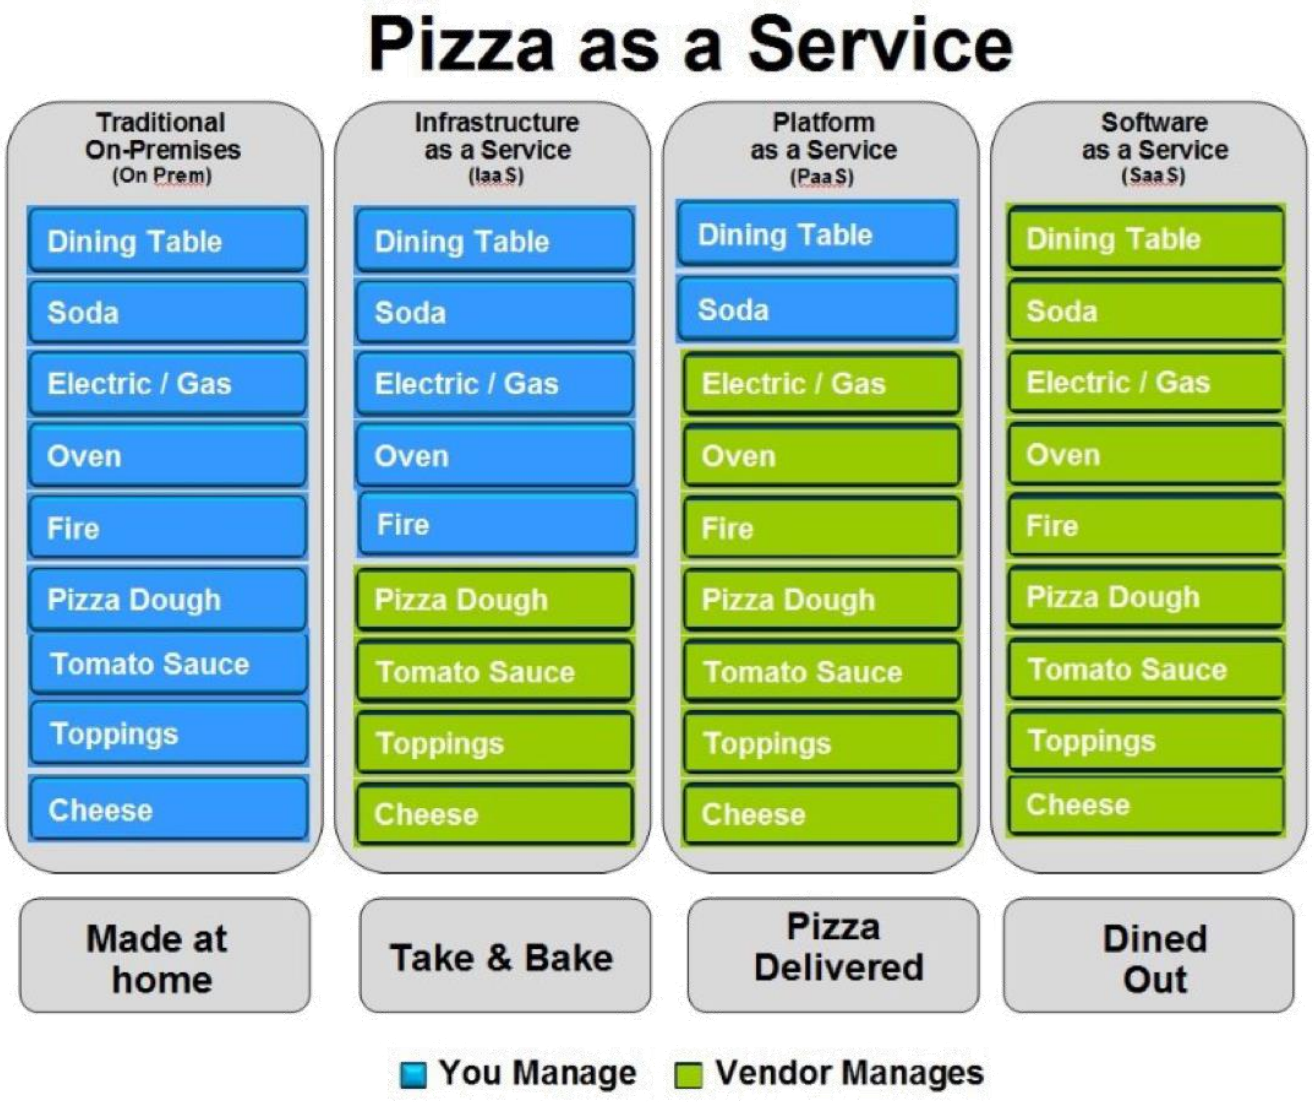

Tipuri de servicii ”Cloud”

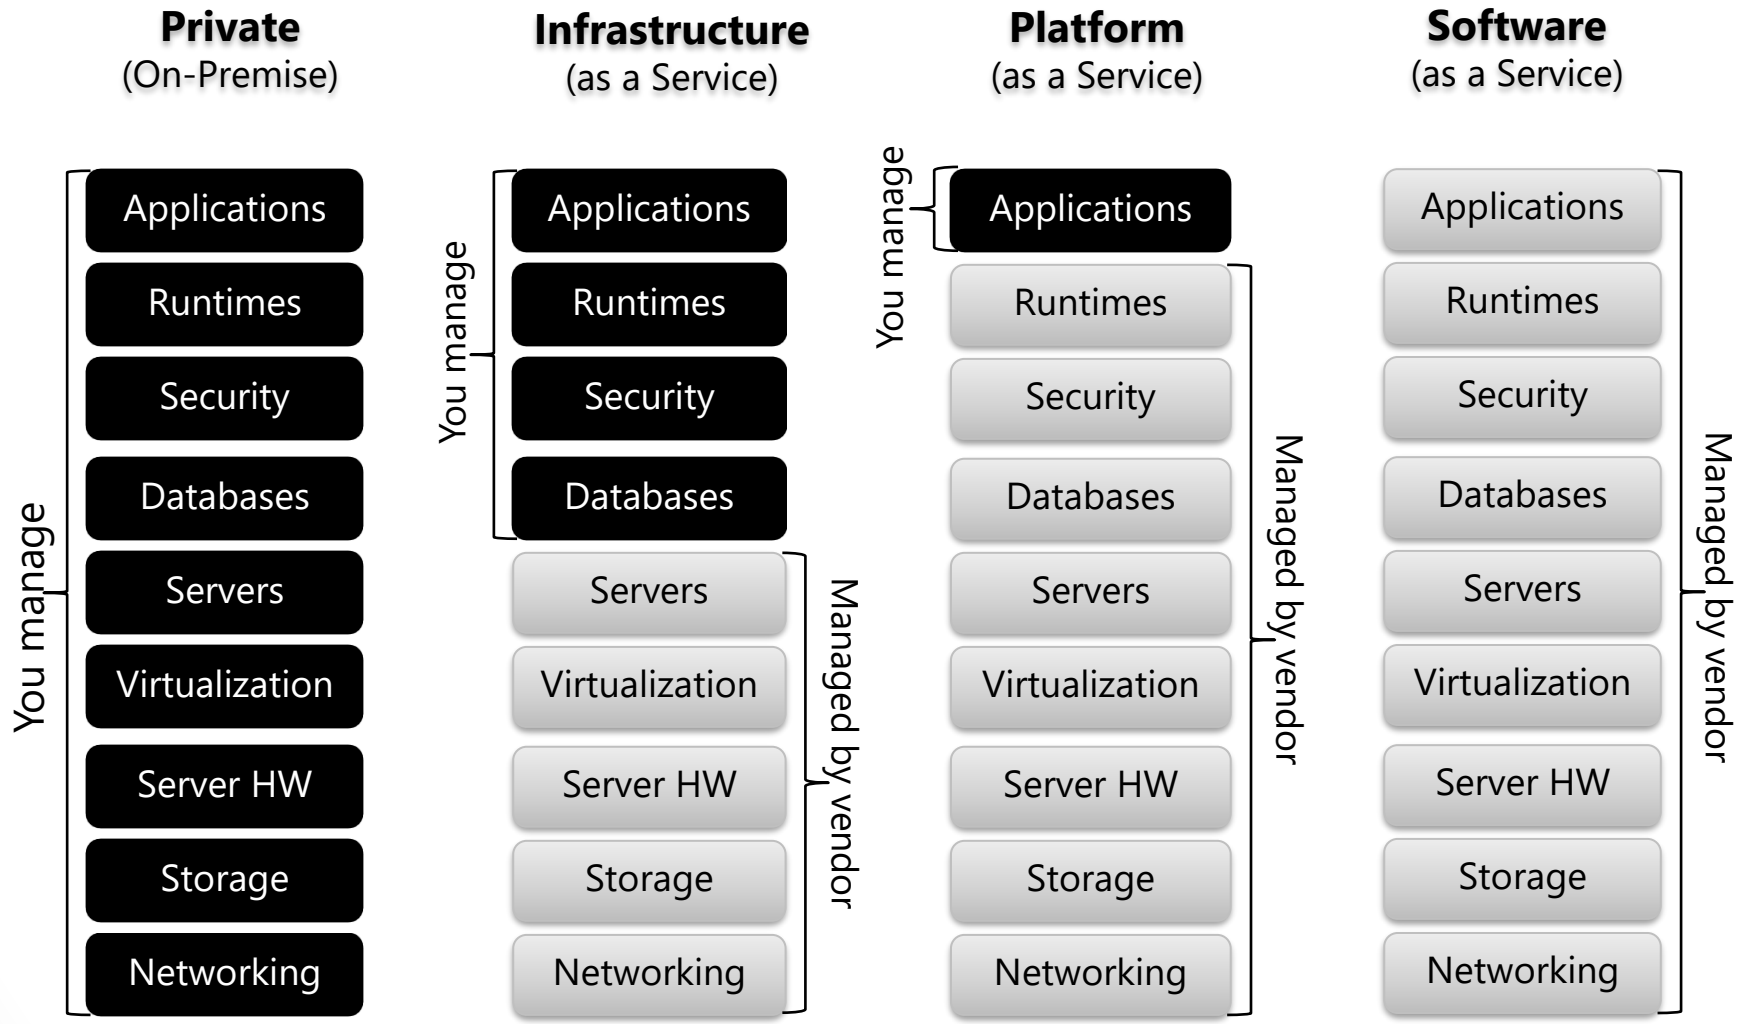

Source: Microsoft Cloud Continuum

Shared Responsibility in the Cloud

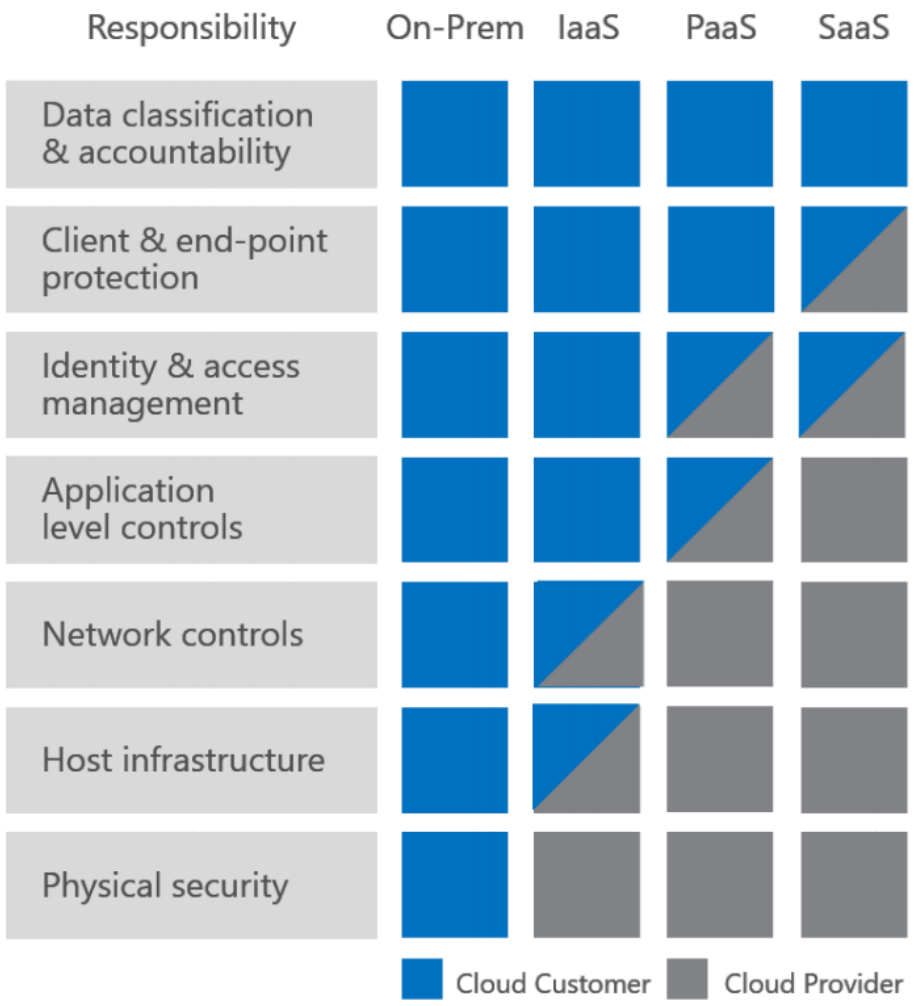

Ce spune industria IT?

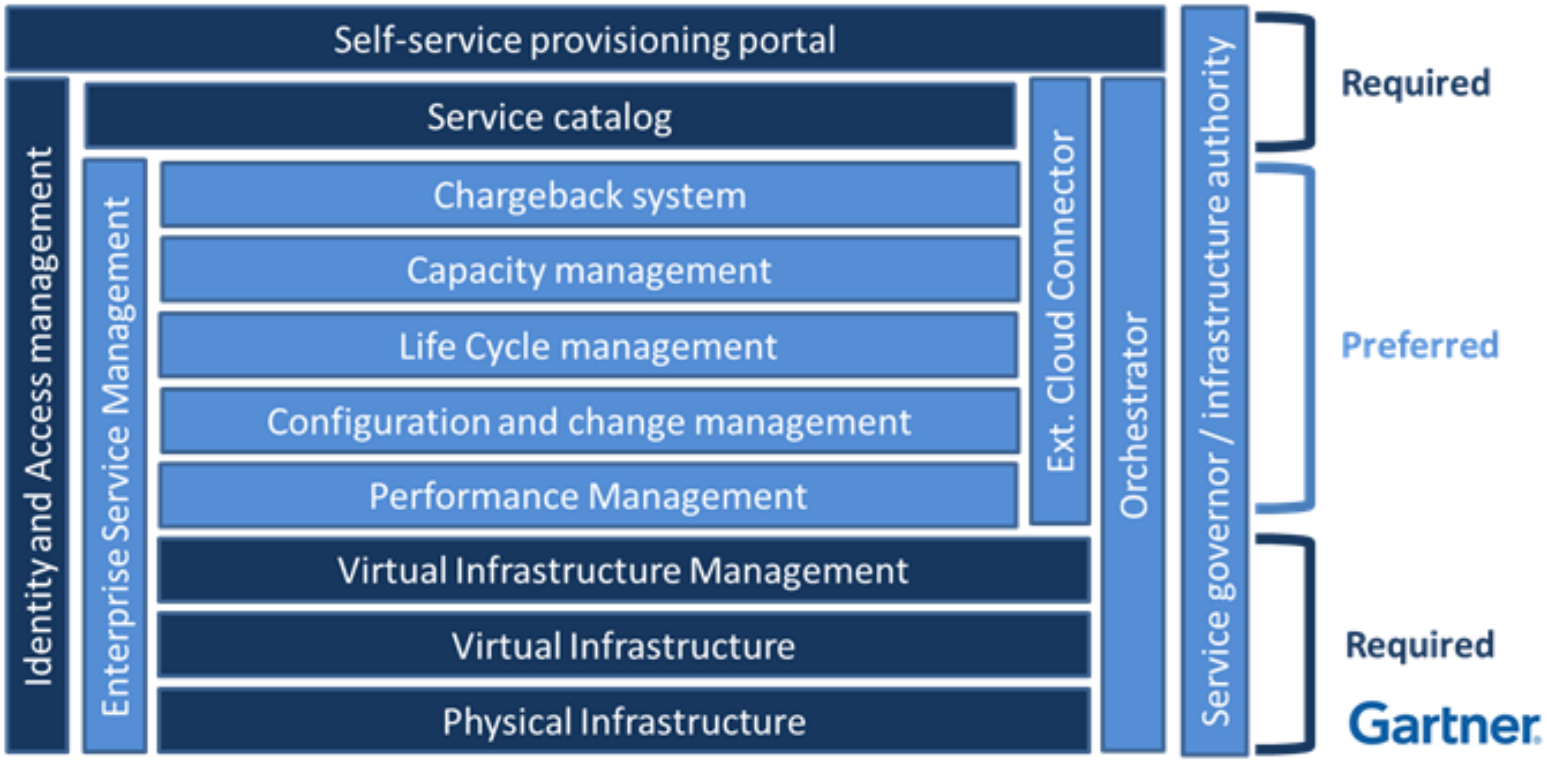

Componentele unui "Private Cloud"

3rd party extensions

Management

Virtualization

Compute Network Storage

- Pooled Resources
- Virtualization
- Elasticity
- Scalability
- Continuous Availability
- Predictability
- Usage-Based
- Multi-Tenancy
- Security
- Automation
- Service management

Source: Microsoft Cloud Continuum

VIRTUAL? PRIVATE? CLOUD?

a. de ce vreau să fac asta?

b. pentru cine fac asta**?**

c. ce vreau să ofer?

SLA, compliance (PCI, ISO)

la ce ne uităm când construim un

?

1. buget

cost per kWh, preț per U, costuri legate de bandă/conectivitate, personalul tehnic, etc.

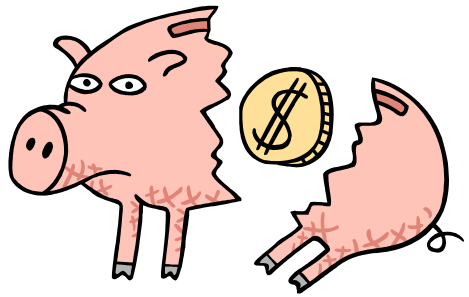

2. arhitectură

structură, scalabilitate, fiabilitate, redundanță, securitate, flexibilitate

3. hardware

server (MIPS/MOPS), storage (I/O-OPS),
network (1Gbps, 10Gbps, infiniband, fiber)

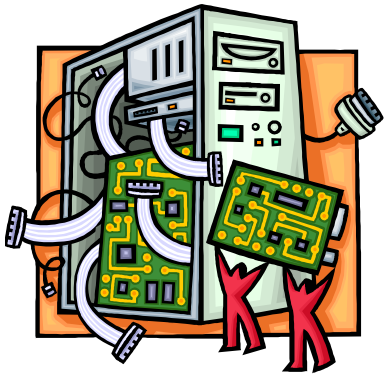

4. virtualization layer

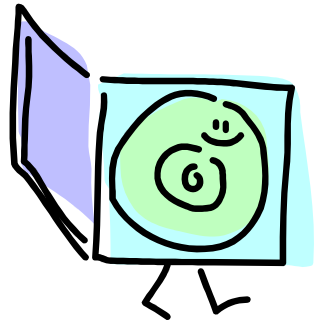

5. management / monitorizare

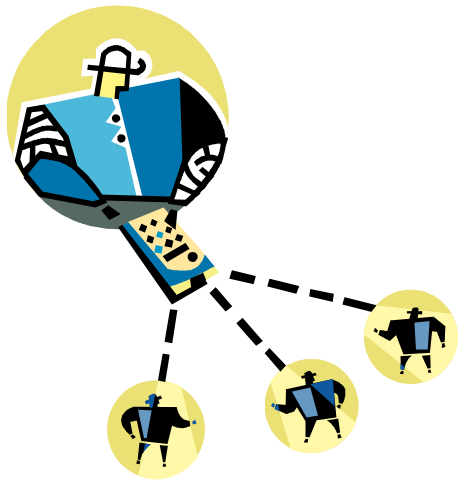

6. procese interne

politici de securitate/incident-response, ușurința de a detecta/repara problemele, disaster recovery, high availability și timpul necesar aducerii unui nod online

7. oameni

echipa tehnică, mentenanță post-implementare

Recapitulare – Private Cloud

1. Buget
2. Arhitectură
3. Hardware
4. Virtualizare
5. Management și monitorizare
6. Procese interne
7. Oameni

Recapitulare – Private Cloud

1. Buget
2. Arhitectură
3. Hardware
4. Virtualizare
5. Management și monitorizare
6. Procese interne
7. Oameni

ARHITECTURI DE VIRTUALIZARE

când lucrăm cu virtualizarea,
ajungem să auzim câțiva termeni
uzuali...

virtualizare

mașină virtuală

hypervisor

paravirtualizare

microkernelized hypervisor

monolithic hypervisor

synthetic device drivers

parent partition

binary translation

un pic de istorie

virtualizarea e veche

primul val, IBM CP-40

CP-40 a intrat în producție în ianuarie 1967

atunci, ca și acum, atracția o constituia reducerea costurilor

era bazat pe sistemul S/360

S/360 a introdus facilități de memorie virtuală și adresare pe 32-bit

CP-40 a extins S/360, oferind un mediu complet virtualizat

astfel, CP-40 avea suport pentru
14 mașini virtuale

a urmat CP-67, bazat pe S/360-67

iar în 1972, IBM a finalizat
VM Facility 370, sau VM/370

cu ocazia asta vin și termenii...

Virtual Machine (VM)

Control Program (hypervisor)

a apărut și o nouă mașină, S/370

iar VM/370 a reușit să virtualizeze eficient SO destinate S/360 și S/370

...și erau vreo câteva :)

OS/360, DOS/360, OS/370, DOS/370, MVS, CMS, CMS/370

așa că virtualizarea și-a atins scopul: **reducerea costurilor**

...în final s-a ajuns la z/VM care rulează acum pe sistemele mainframe IBM z10 și IBM zEnterprise (z196/z114)

chiar și cu VM/370,
costurile erau încă ridicate

al doilea val, microprocesoarele

1977, Macintosh

Apple II

1988, Connectix

Virtual PC (pentru Macintosh)

Virtual PC putea rula Windows, OS/2, Linux pe hardware **Macintosh**

în 2003, Microsoft cumpără Connectix

Virtual PC e modificat, și astfel userii Windows pot să ruleze și alte SO

MS-DOS 6.22, Windows 95, 98, NT 4.0, Me, OS/2, 2000, 2003, XP, Vista, 2008

al treilea val, VMware

1999, VMware produce

VMWare Workstation

2001, VMware (GSX) Server (virtualizare server-level)

trecem și la partea tehnică

chiar dacă la bază, conceptele VMware și Virtual PC erau similare cu CP-40

totuși, atât din Virtual PC, cât și din VMware, lipsea hypervisor-ul

ambele se bazau pe
existența unui SO ”gazdă” (host)

host-ul era astfel intermediar (ineficient) între VM și hardware

ulterior, au apărut îmbunătățiri

posibilitatea de a rula VM pe hardware nespecializat

migrare de la arhitecturi hosted la arhitecturi bazate pe hypervisor

VMware ESX, Hyper-V, Xen

folosesc o forma de hypervisor

reintroducerea suportului hardware:

Intel VT, AMD-V

tendințe

chiar dacă micro-virtualizarea tinde să egaleze soluția IBM de acum 50+ ani

au apărut și elemente inovatoare, majoritatea din partea VMware

live server migration
(VMotion)

dynamic load balancing
(Distributed Resource Scheduling)

real-time failover
(HA Clustering)

sau din partea Intel

Virtualization for Directed I/O (VT-d)
pentru procesarea input-output

și în același timp, virtualizarea a adus cu ea probleme de securitate

la început, avantajele virtualizării, gen izolarea aplicațiilor, au fost afectate

procesoarele cu virtual-assist
sunt un bun exemplu

un guest putea să acceseze direct alt guest,
ignorând politicile de securitate

sau mai rău, exploit-ul Blue Pill

tipuri de atacuri

jailbreak attacks (escapes)

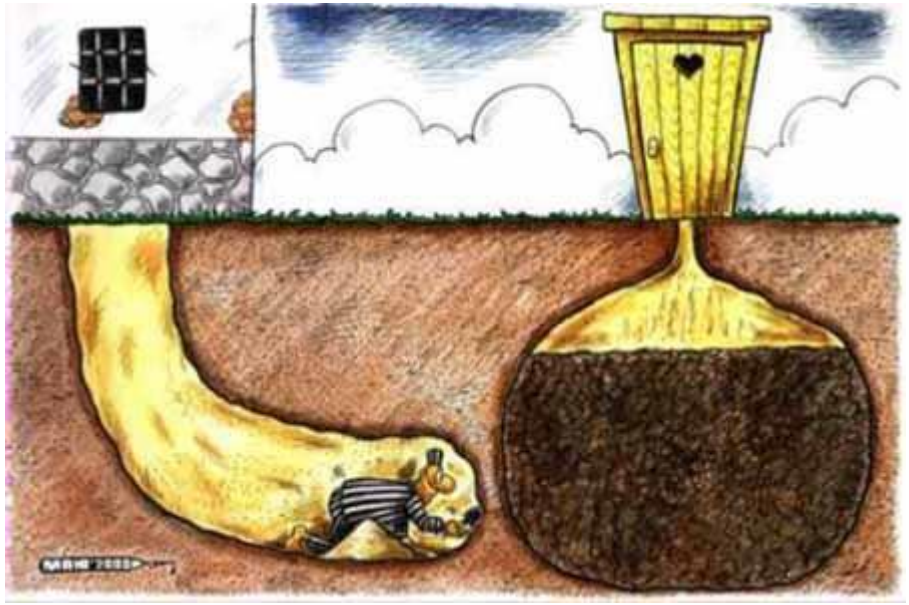

migration attacks

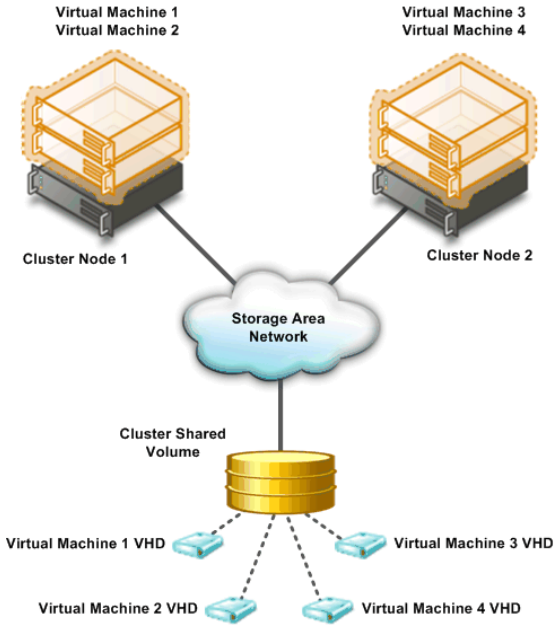

virtual / physical
network service attacks

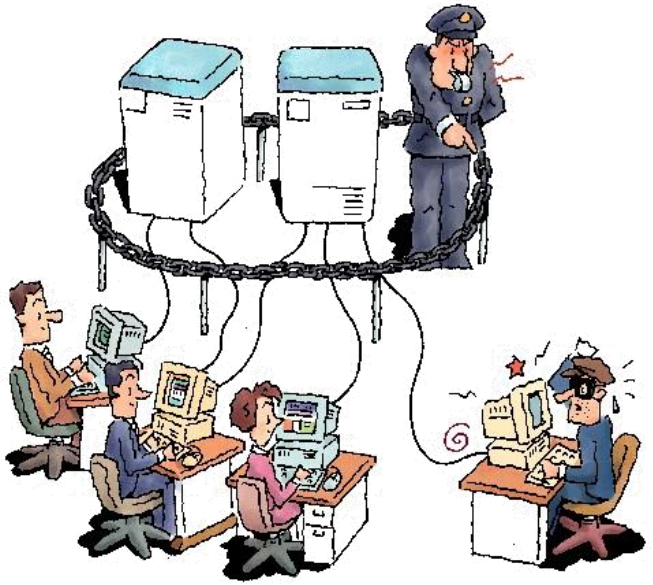

encryption attacks

soluția? facilități de securitate adăugate la noul hardware

chiar și așa, virtualizarea va rămâne, atât timp cât își îndeplinește scopul

începem cu câteva explicații

de la soluțiile de **virtualizare**,
toată lumea așteaptă automat la:

izolare eficientă
securitate
performanță
ușurință în administrare

tehnic, virtualizarea se poate realiza în mai multe moduri

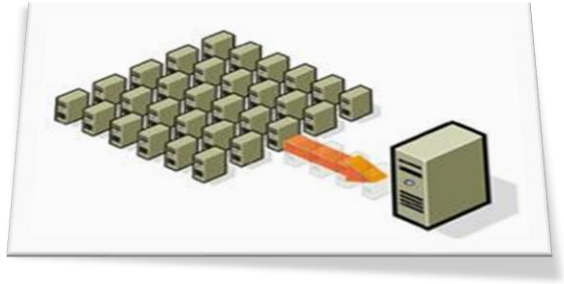

full virtualization (Type 1)

guest nemodificat, migrare ușoară spre medii virtualizate

paravirtualizare

guest modificat pentru a elimina nevoia de ”binary translation”

oferă avantaje de performanță în anumite circumstanțe, însă e nevoie de o versiune modificată de kernel instalată pe guest

virtualizare asistată hardware

AMD-V, Intel VT

prima generație a inclus doar virtualizare CPU, generațiile următoare vin cu suport de virtualizare pe memorie și I/O

hosted virtualization (Type 2)

Microsoft Virtual Server / Virtual PC
VMWare Workstation
VirtualBox
Linux KVM

Tipuri de virtualizare

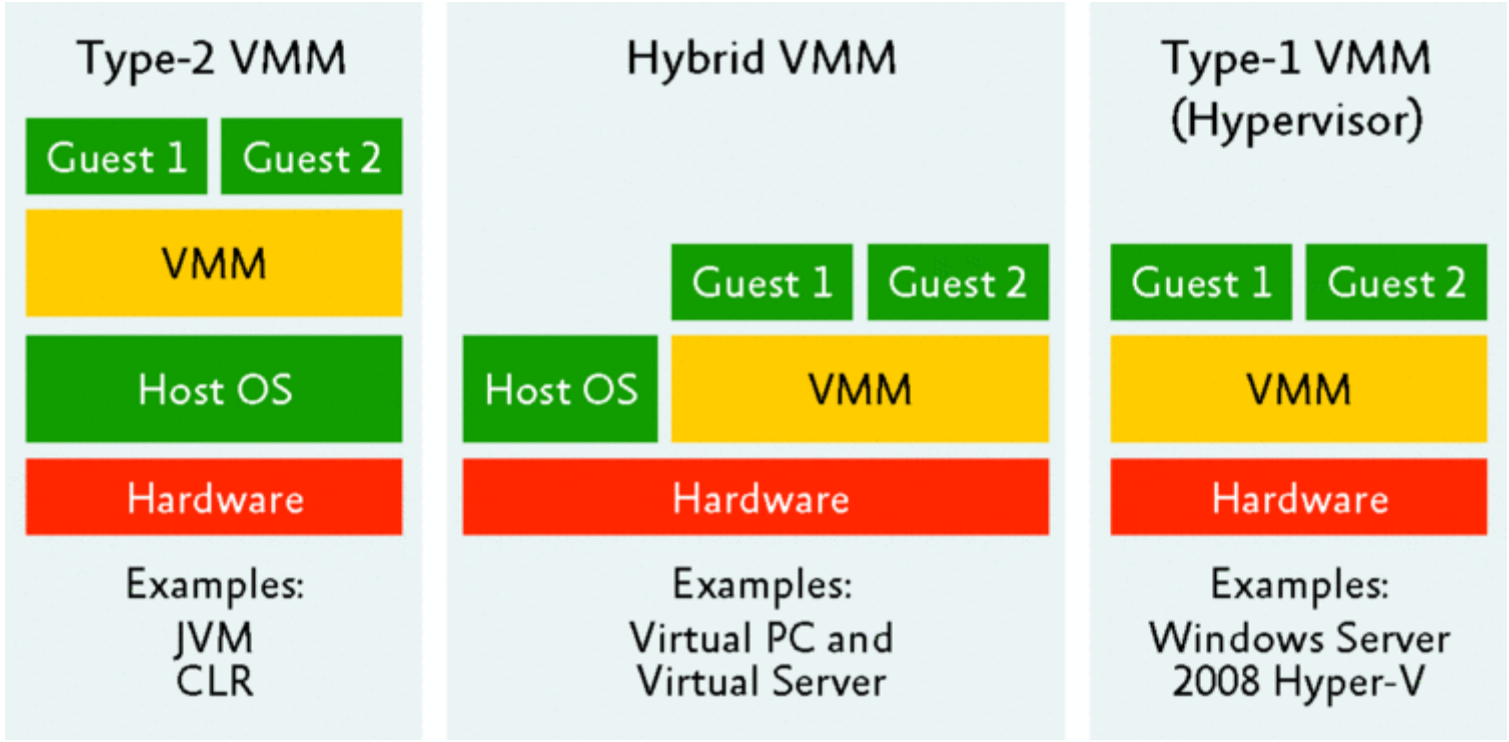

Type I Virtualization Overview

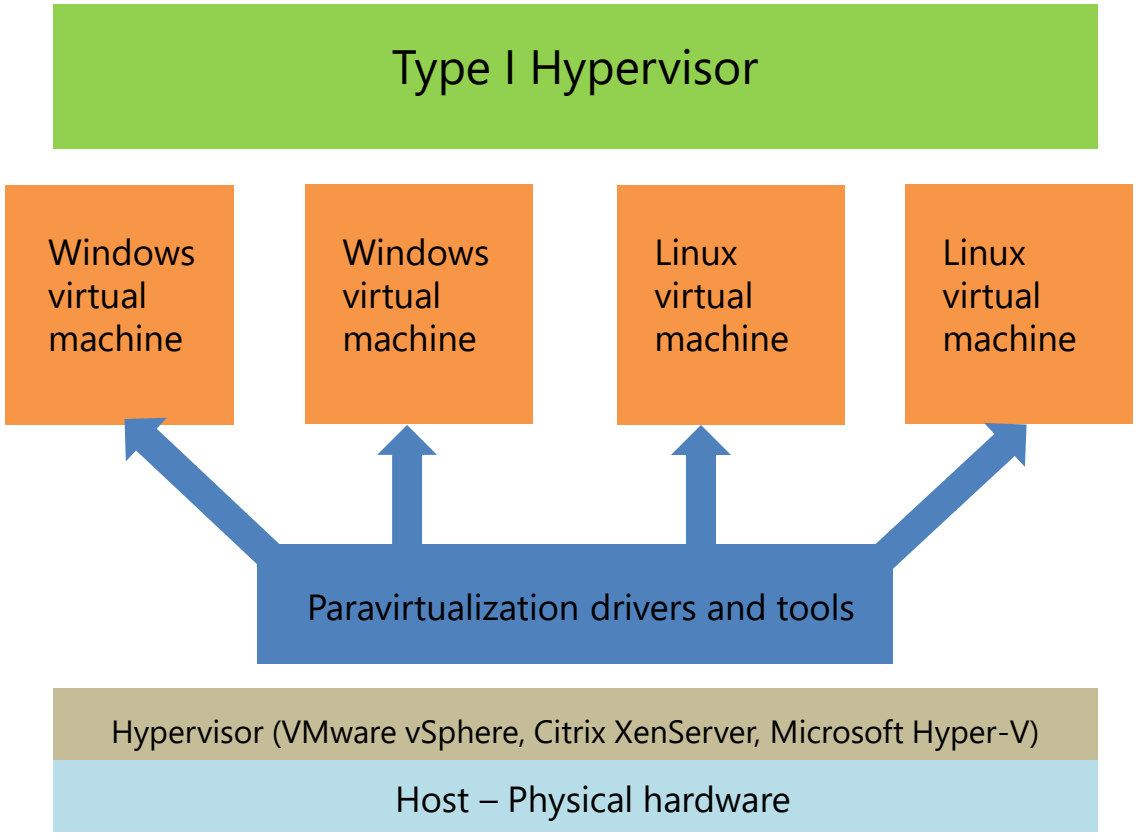

Type II Virtualization Overview

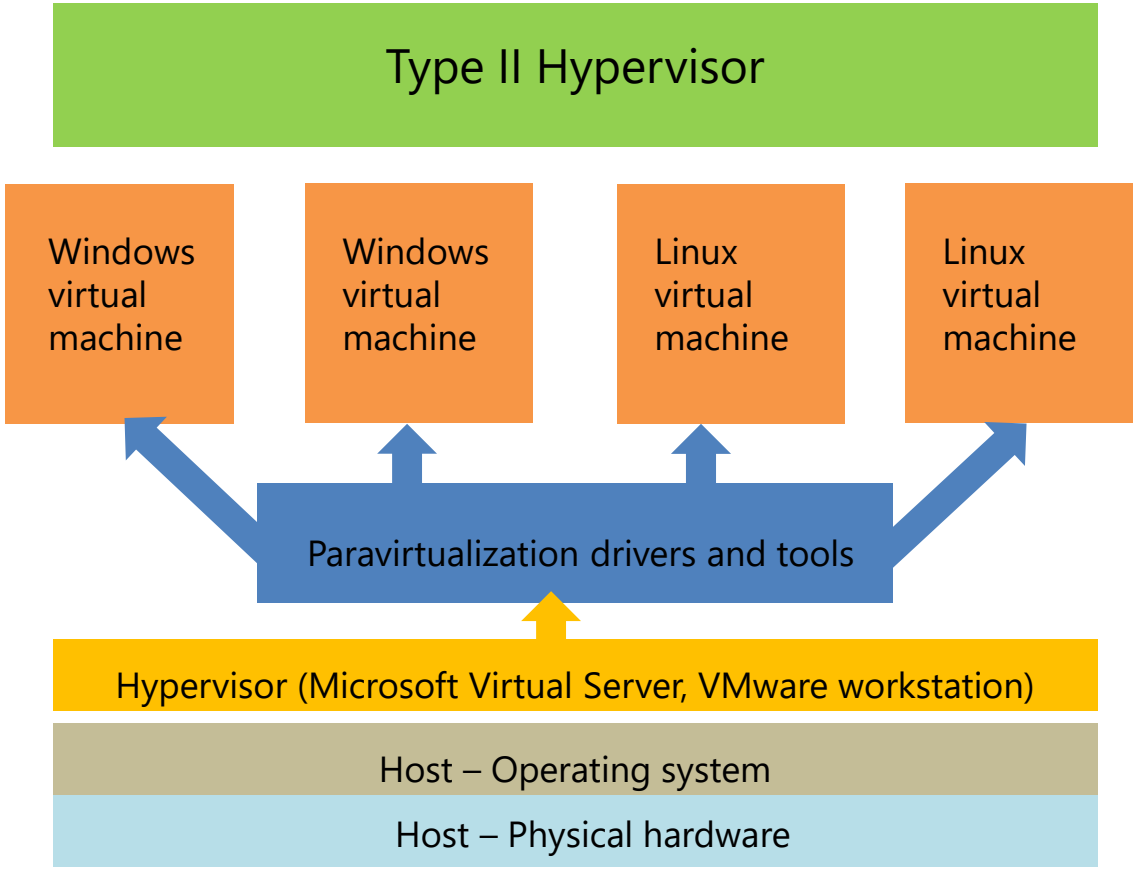

OS virtualization

Virtuozzo / OpenVZ

Windows Containers

WHAT IS A CONTAINER?

VM vs Container

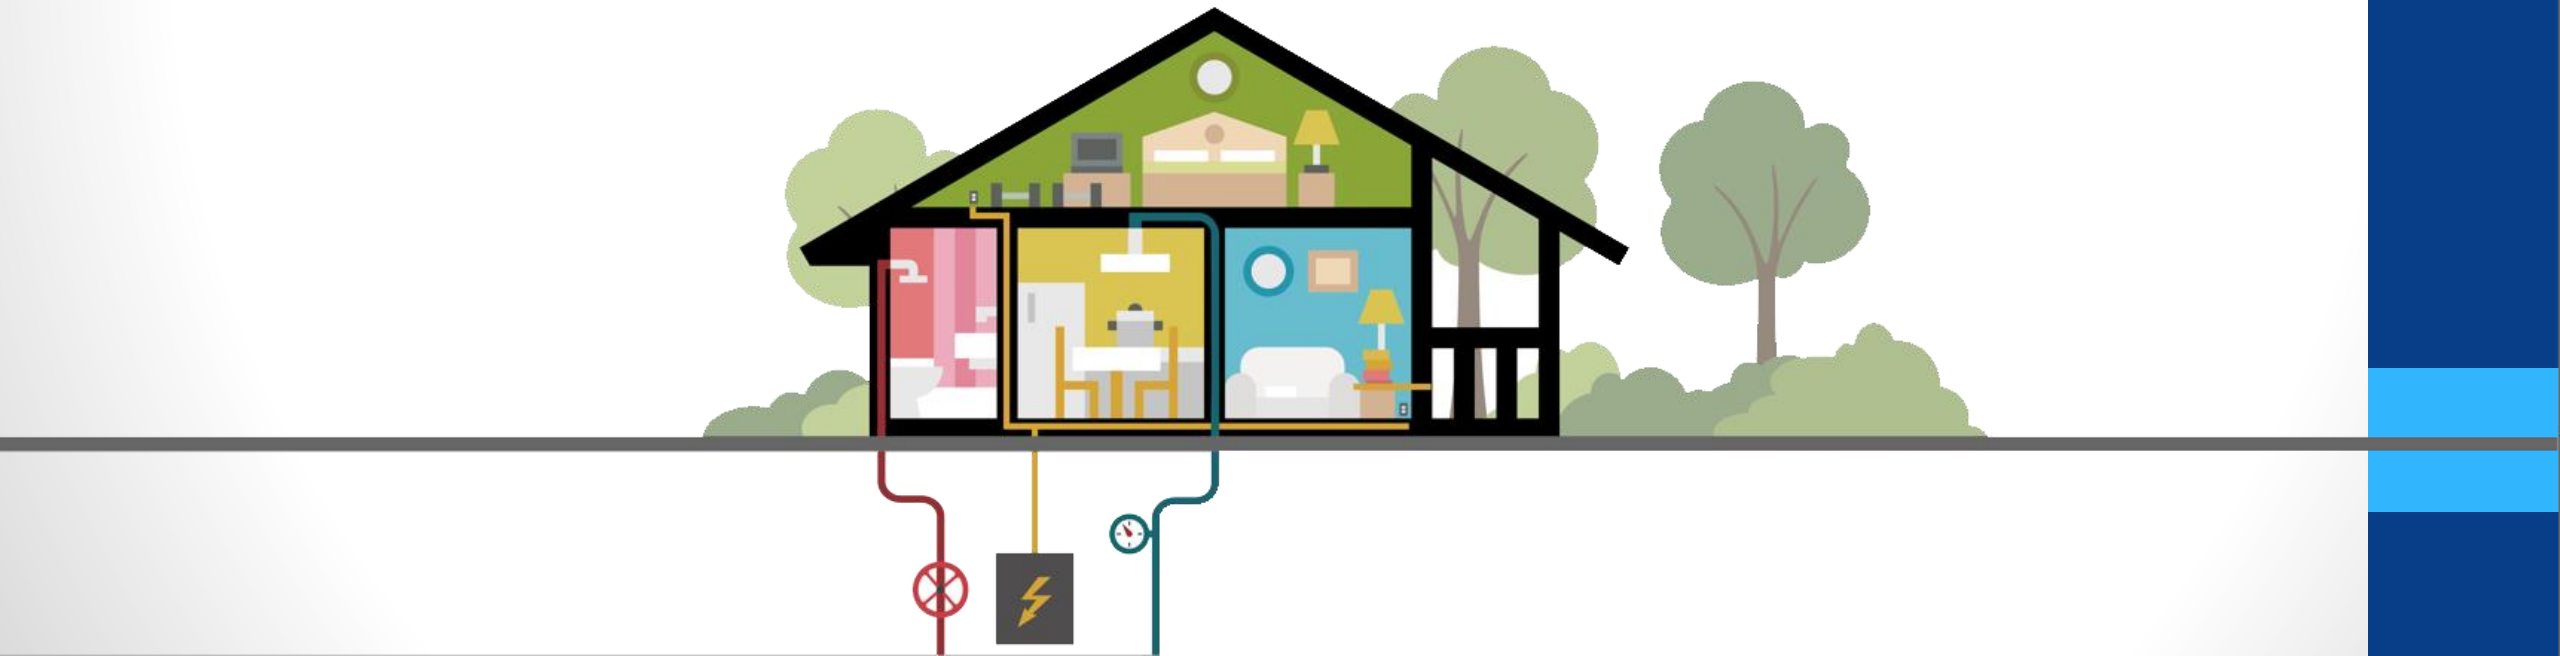

VM vs Container

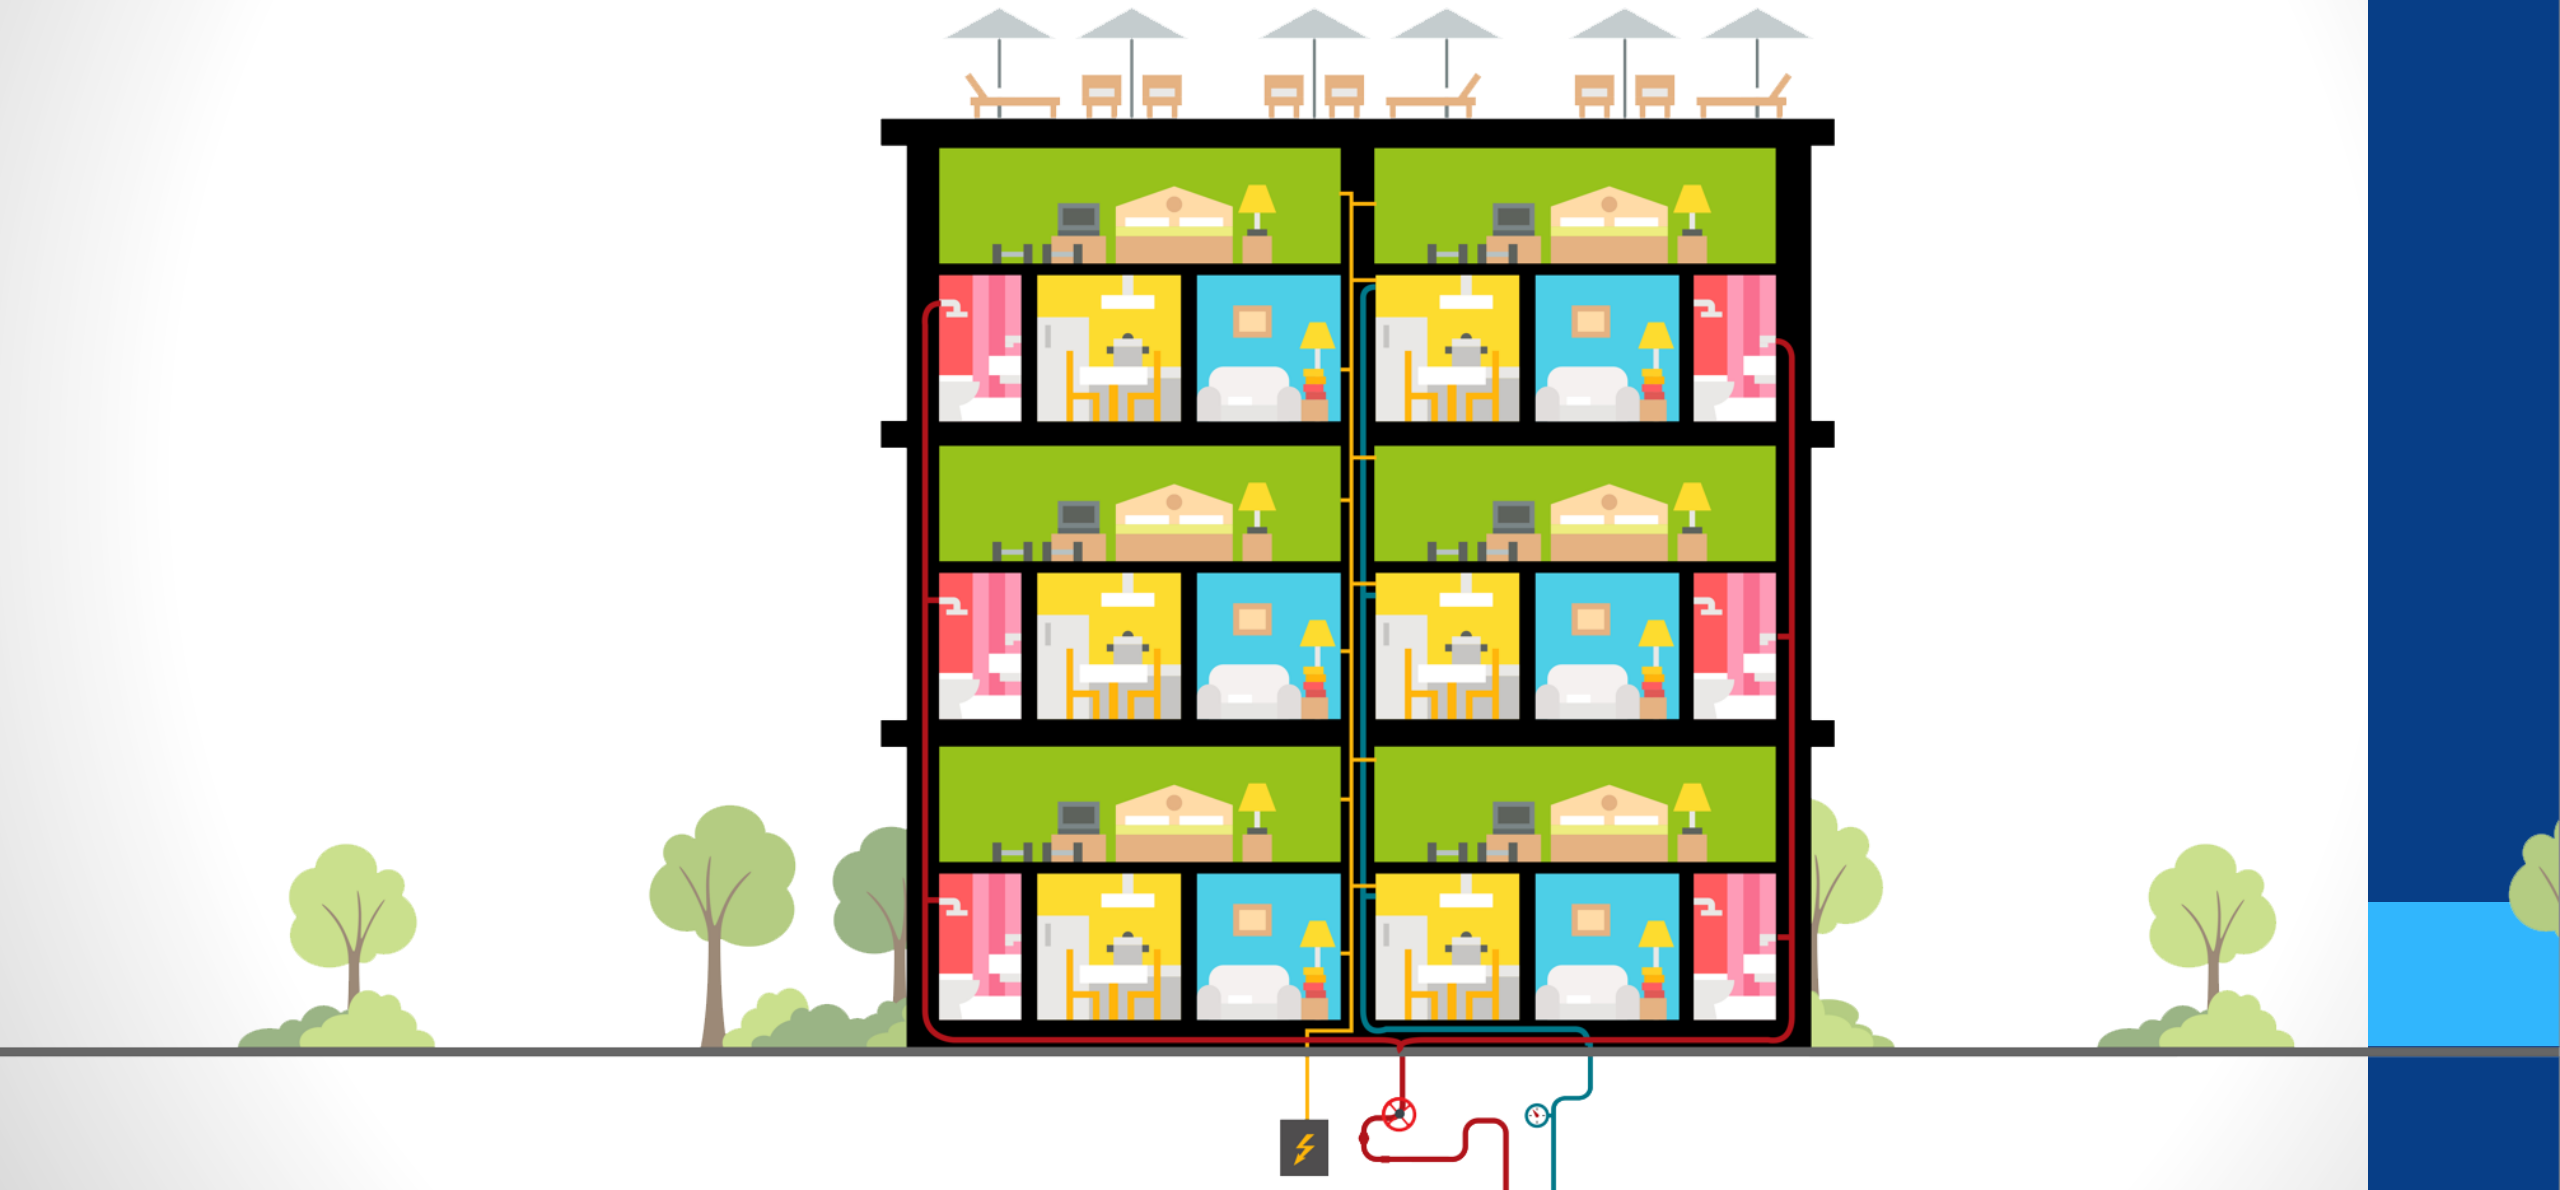

VM vs Container

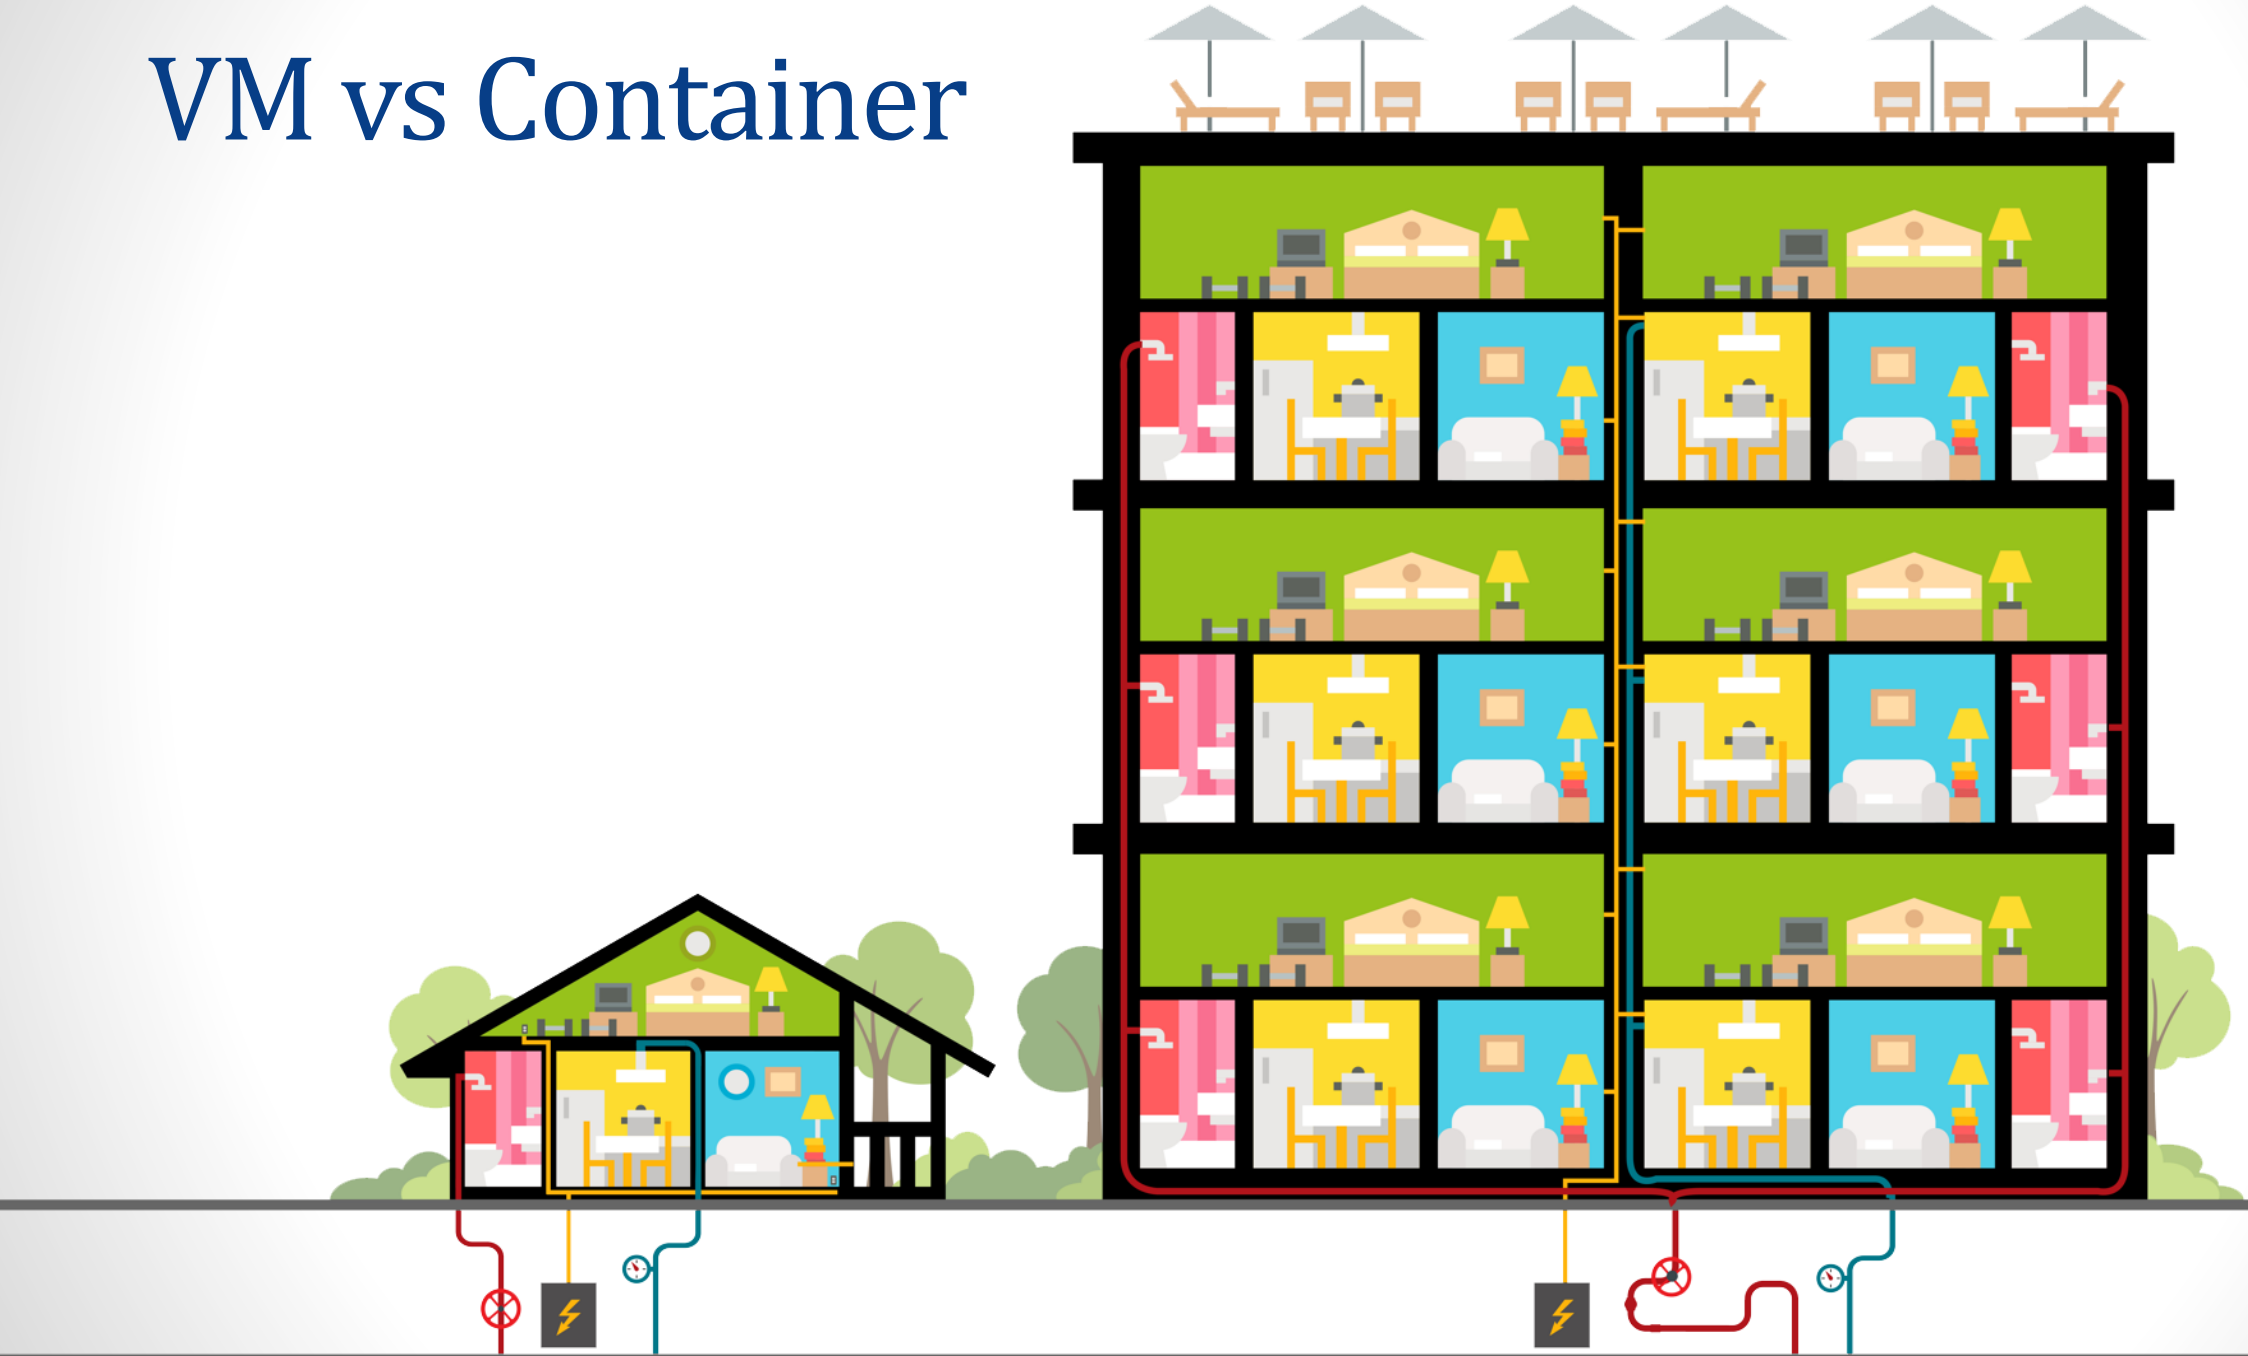

Ce e un container?

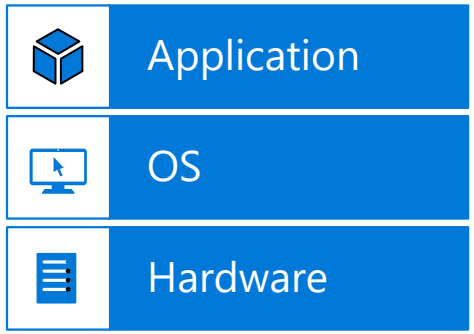

Traditional virtual machines = hardware virtualization

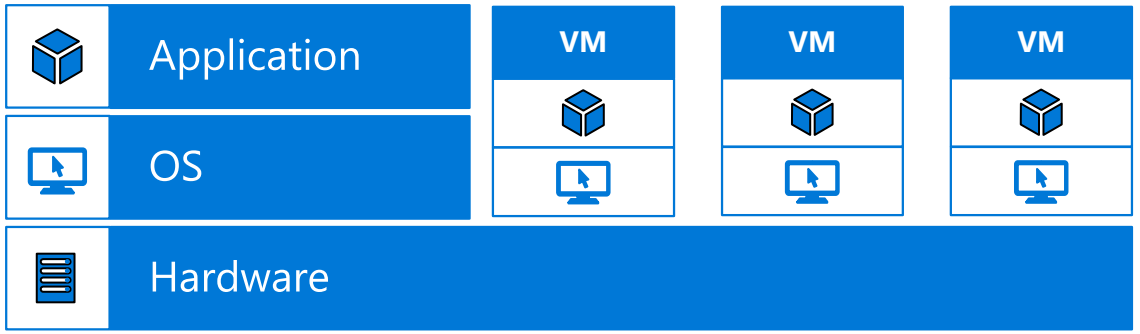

Traditional virtual machines = hardware virtualization

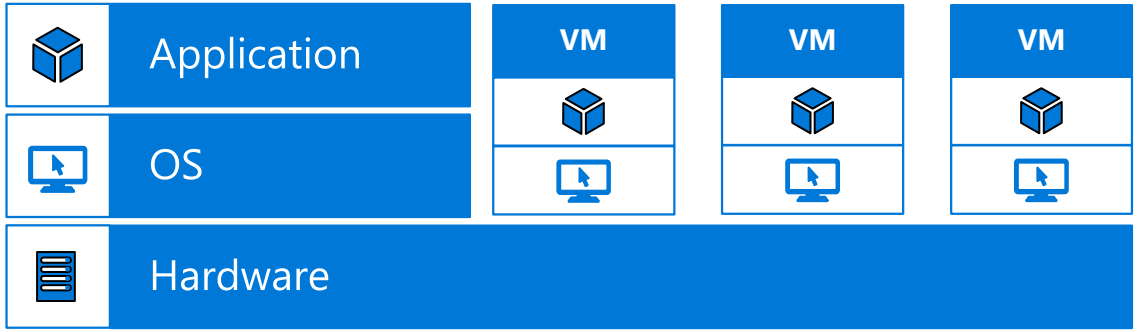

Traditional virtual machines = hardware virtualization

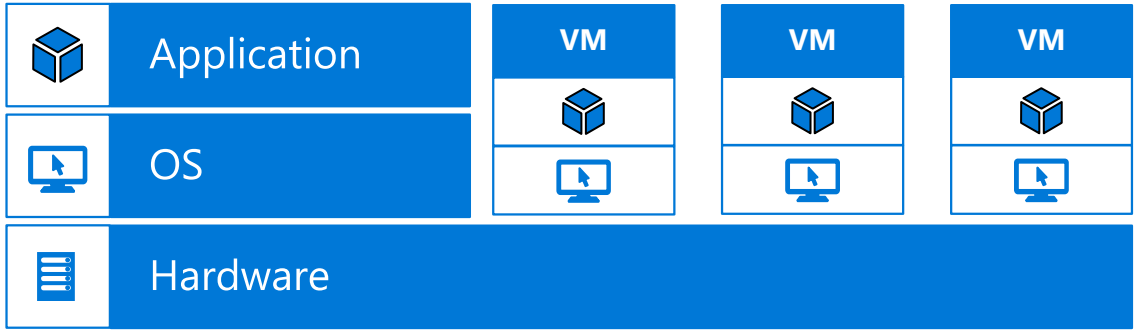

Containers = Operating system virtualization

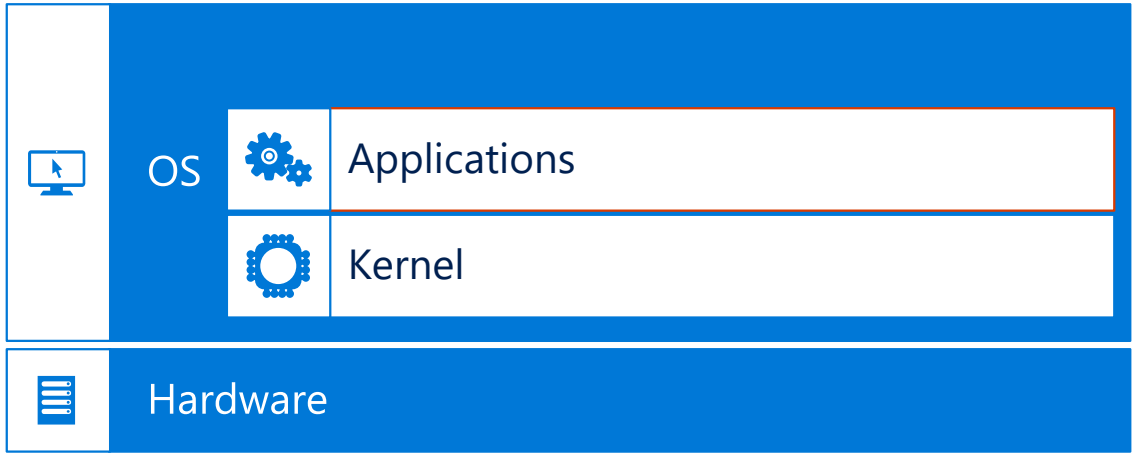

Traditional virtual machines = hardware virtualization

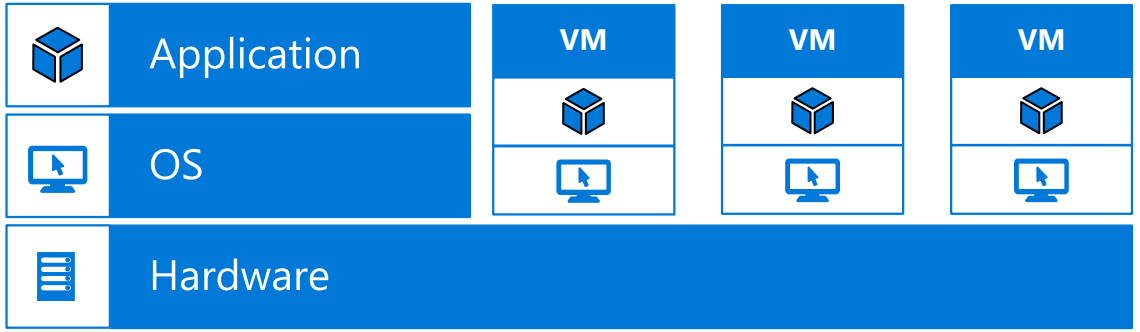

Containers = Operating system virtualization

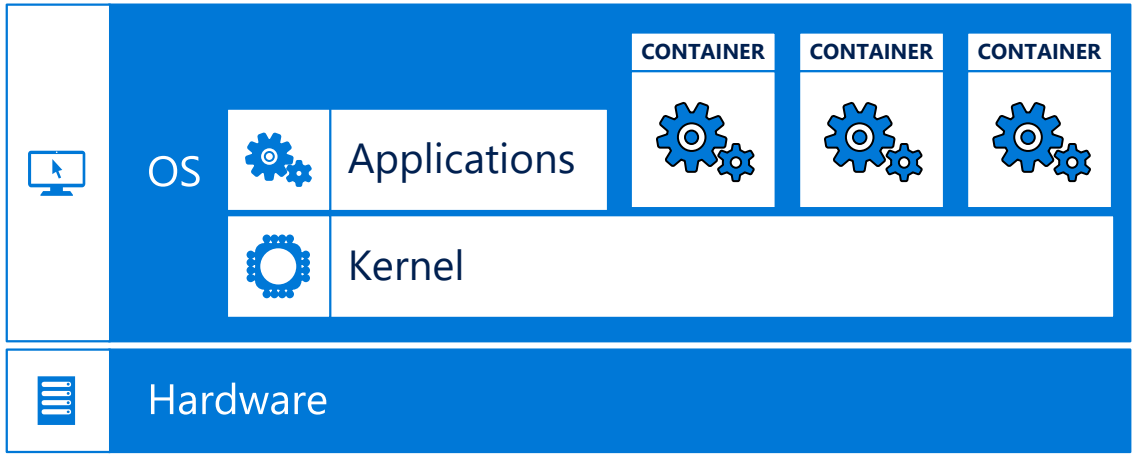

Traditional virtual machines = hardware virtualization

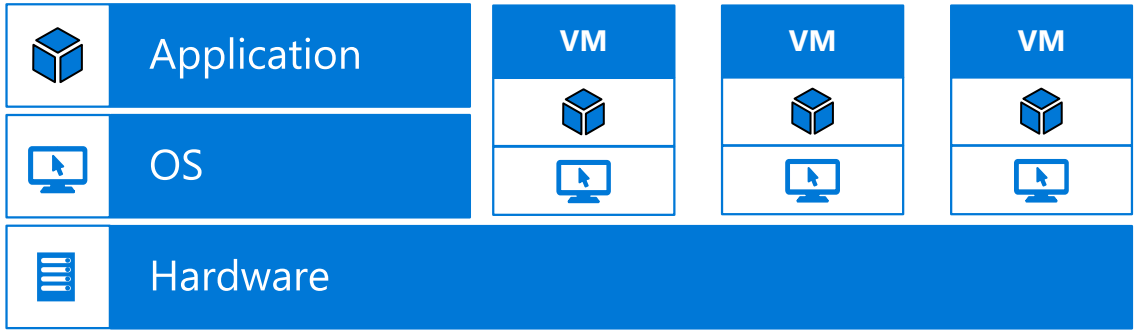

Containers = Operating system virtualization

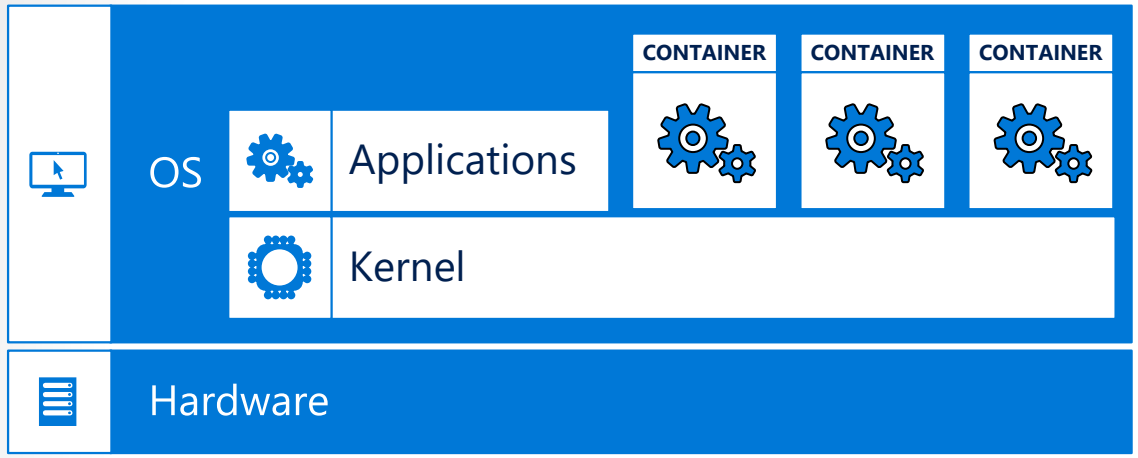

Traditional virtual machines = hardware virtualization

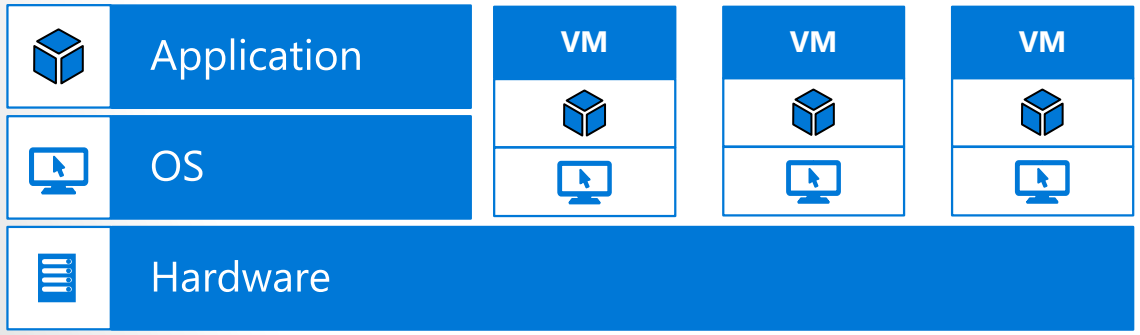

Containers = Operating system virtualization

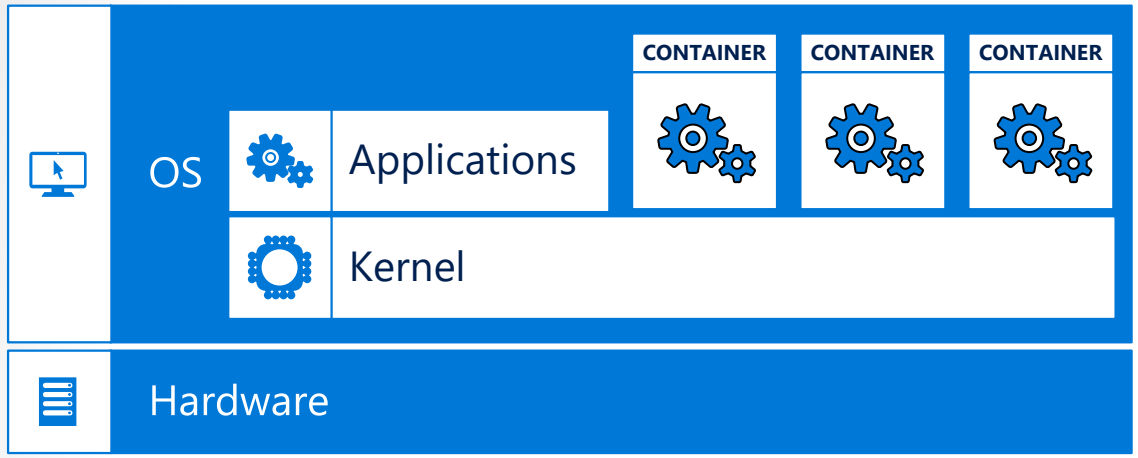

Traditional virtual machines = hardware virtualization

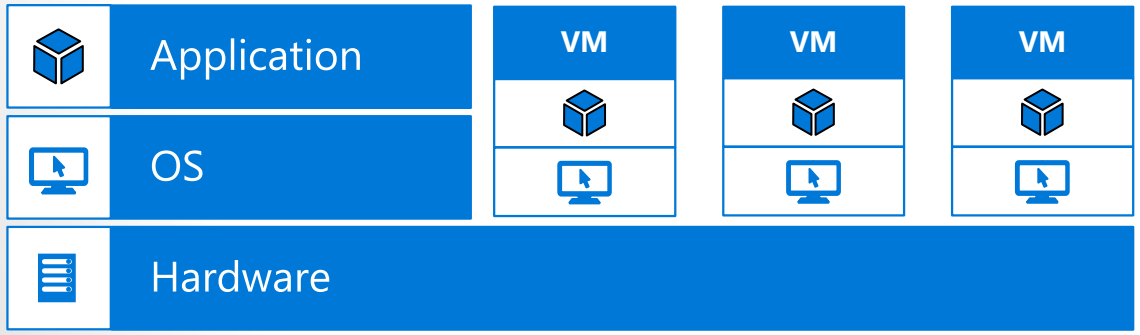

Windows Server Containers

Maximum speed and density

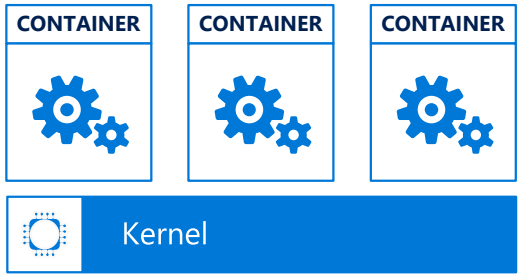

Containers = Operating system virtualization

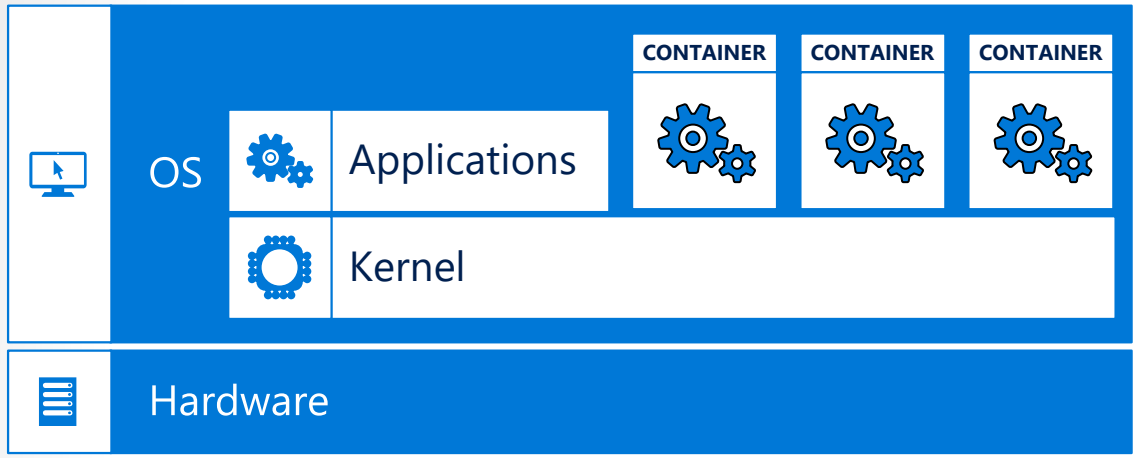

Traditional virtual machines = hardware virtualization

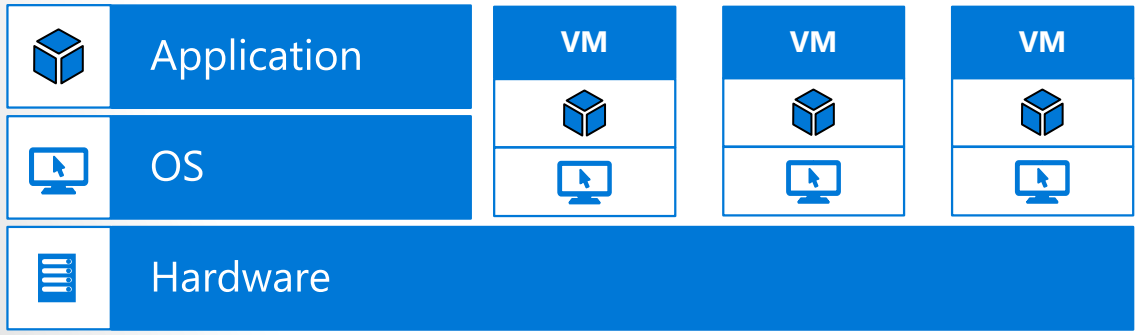

Windows Server Containers

Maximum speed and density

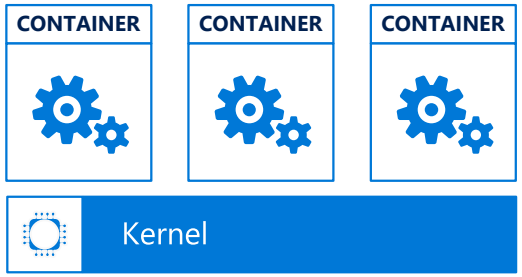

Hyper-V Containers

Isolation plus performance

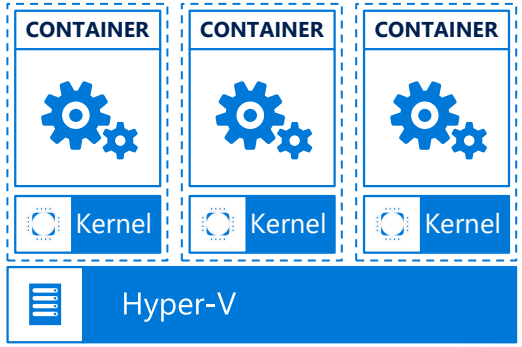

să discutăm puțin despre

Hyper-V / Xen / VMWare
KVM / Virtuozzo / OpenVZ

Arhitectura Hyper-V

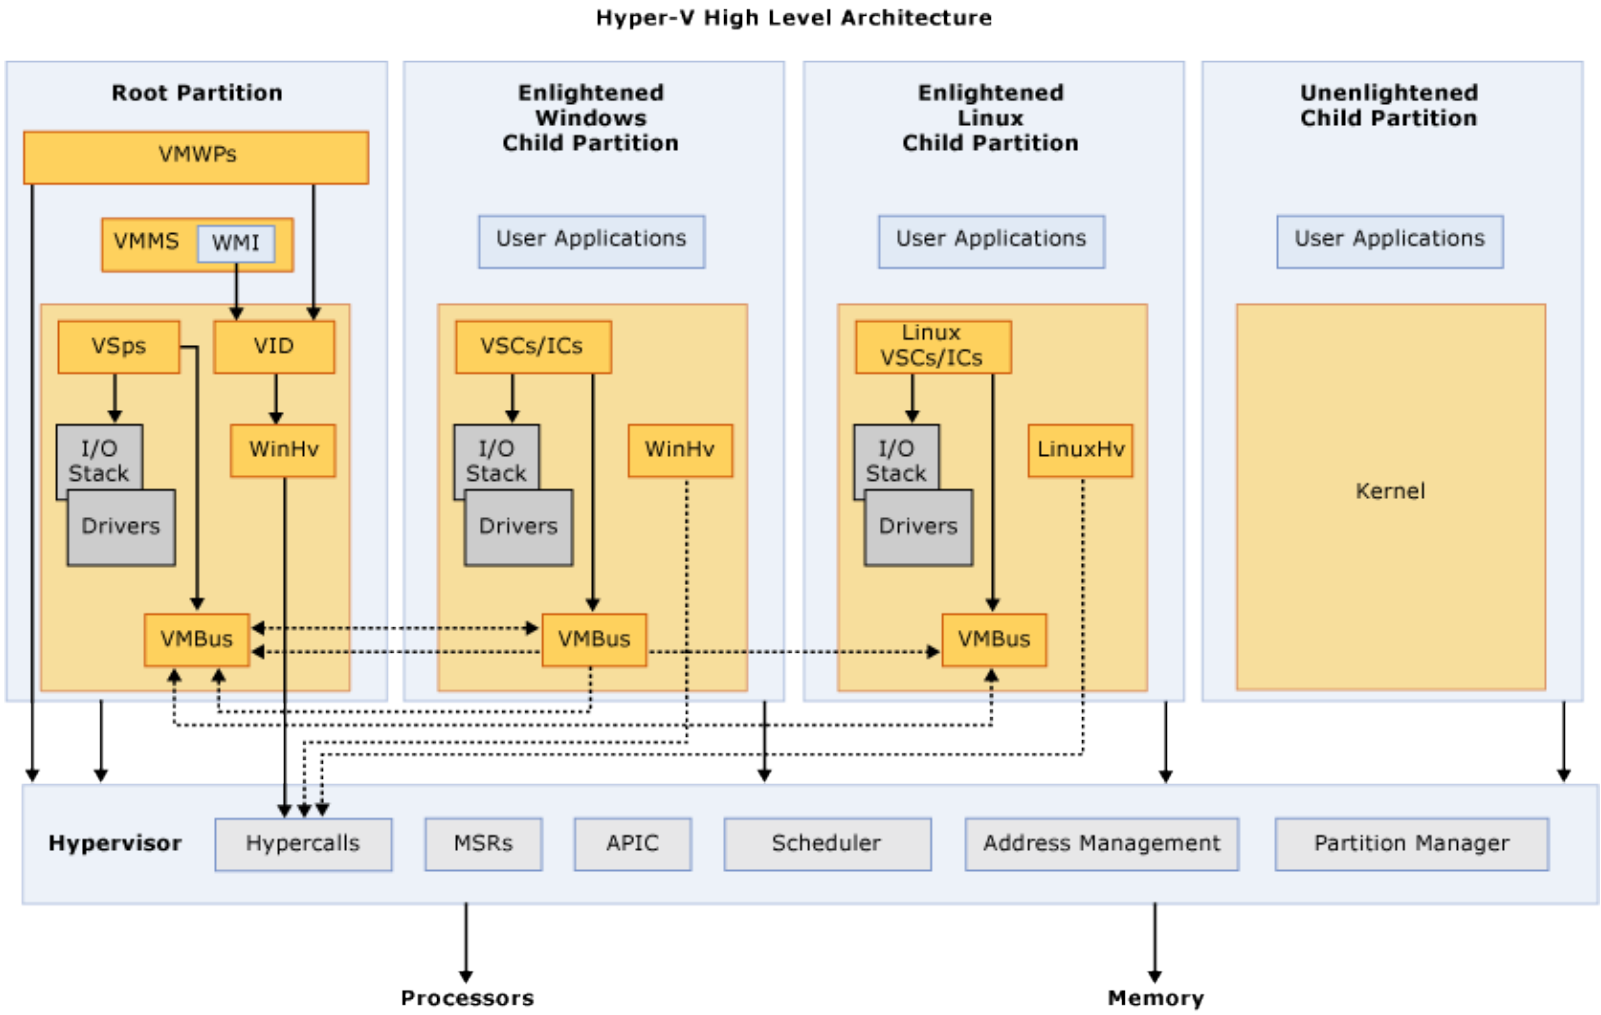

iar cerințele nu sunt exagerate...

x64

DEP

Intel VT / AMD-V

SLAT (W8 Client)

Xen

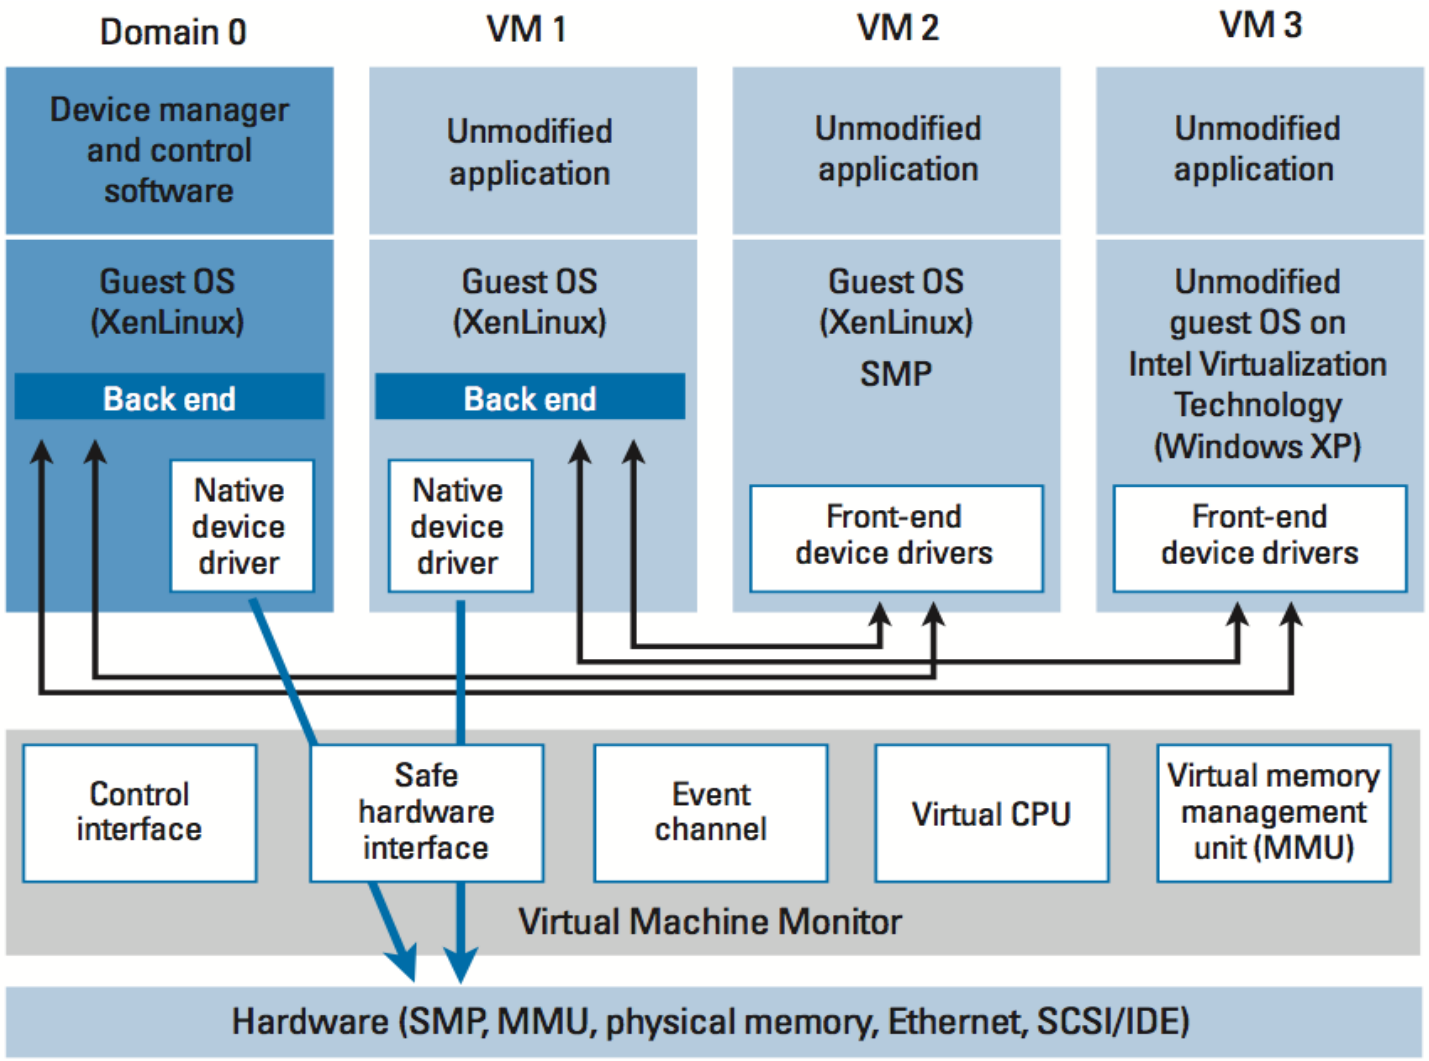
Domain 0
VM 1
VM 2
VM 3
Device manager and control software
Unmodified application
Unmodified application
Unmodified application
Guest OS (XenLinux)
Guest OS (XenLinux)
Guest OS (XenLinux) SMP
Unmodified guest OS on Intel Virtualization Technology (Windows XP)
Back end
Back end
Native device driver
Native device driver
Front-end device drivers
Front-end device drivers
Control interface
Safe hardware interface
Event channel
Virtual CPU
Virtual memory management unit (MMU)
Virtual Machine Monitor
Hardware (SMP, MMU, physical memory, Ethernet, SCSI/IDE)

la VMWare,
lucrurile stau puțin diferit

VMWare ESX Server

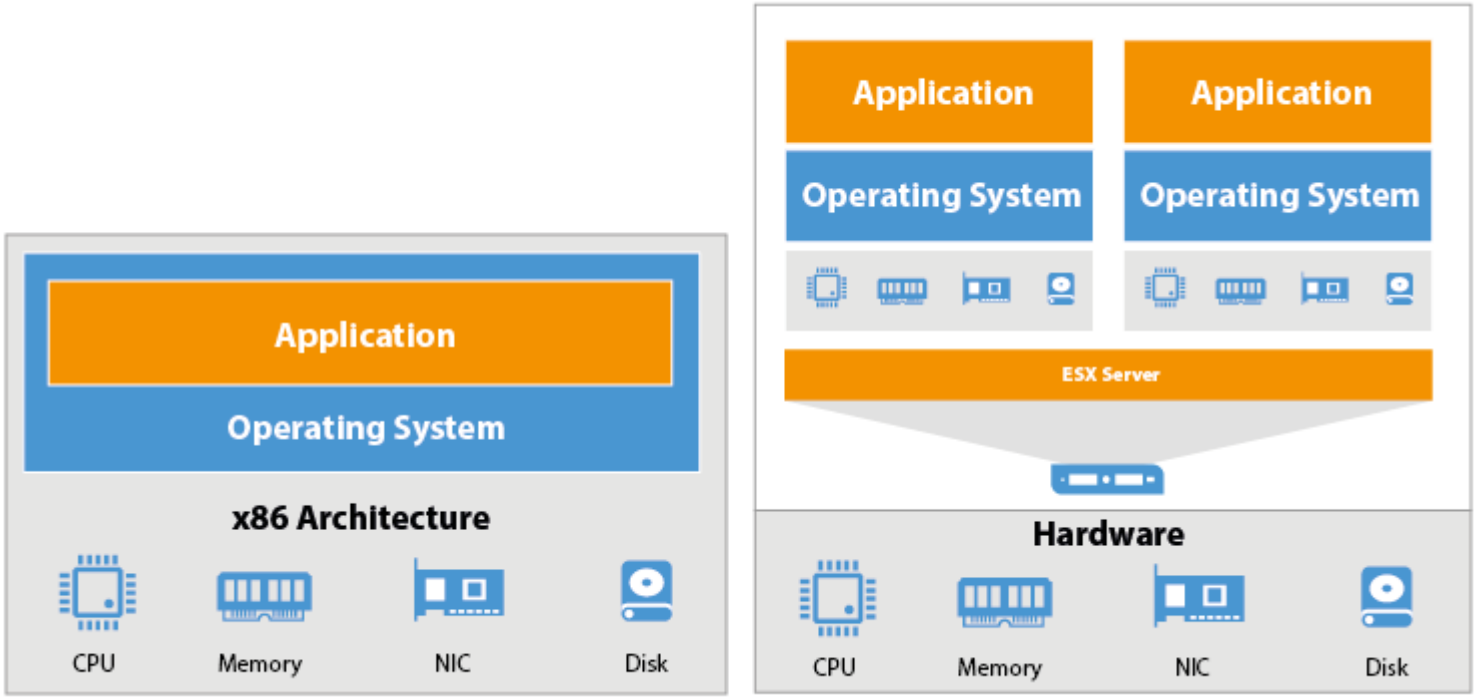

Hyper-V – microkernelized hypervisor

VMWare – monolithic hypervisor

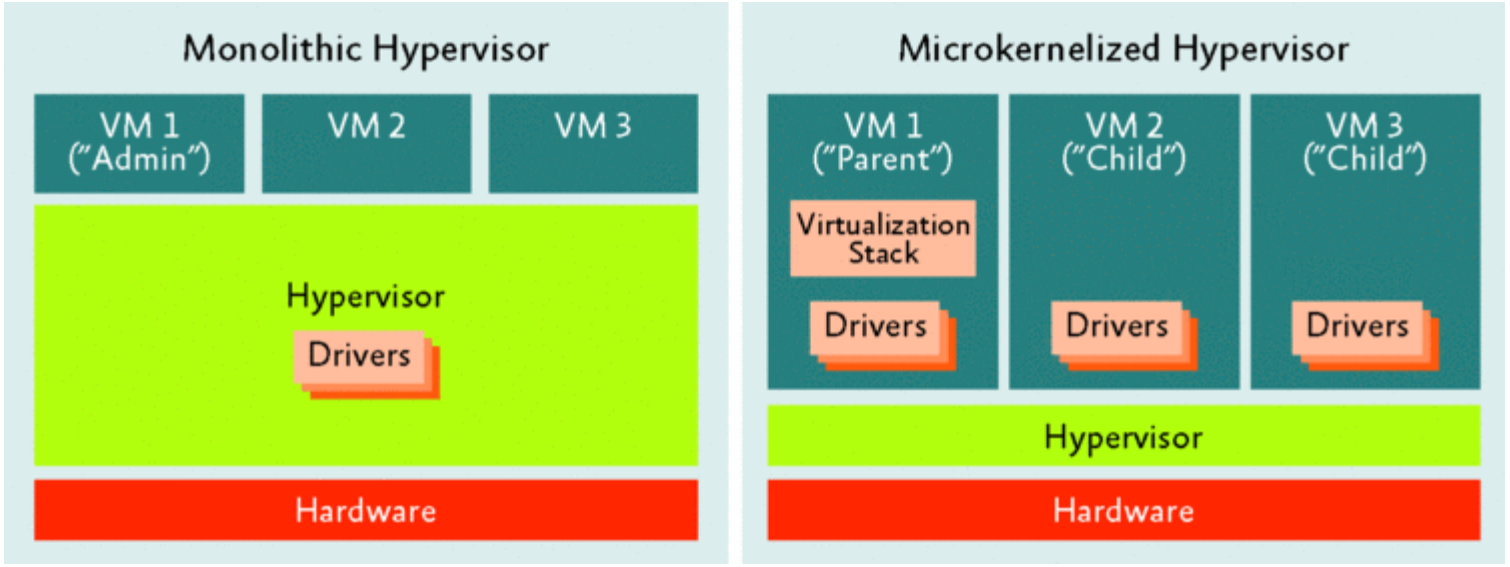

Hyper-V hypervisor (Windows 10)

hvax64.exe (AMD) – 1-2MB

hvix64.exe (Intel) – 1-2MB

VMWare hypervisor – ~150-200MB

Hyper-V – ”synthetic” drivers

VMWare – emulated drivers

KVM / Qemu

KVM – no emulation, user-space program

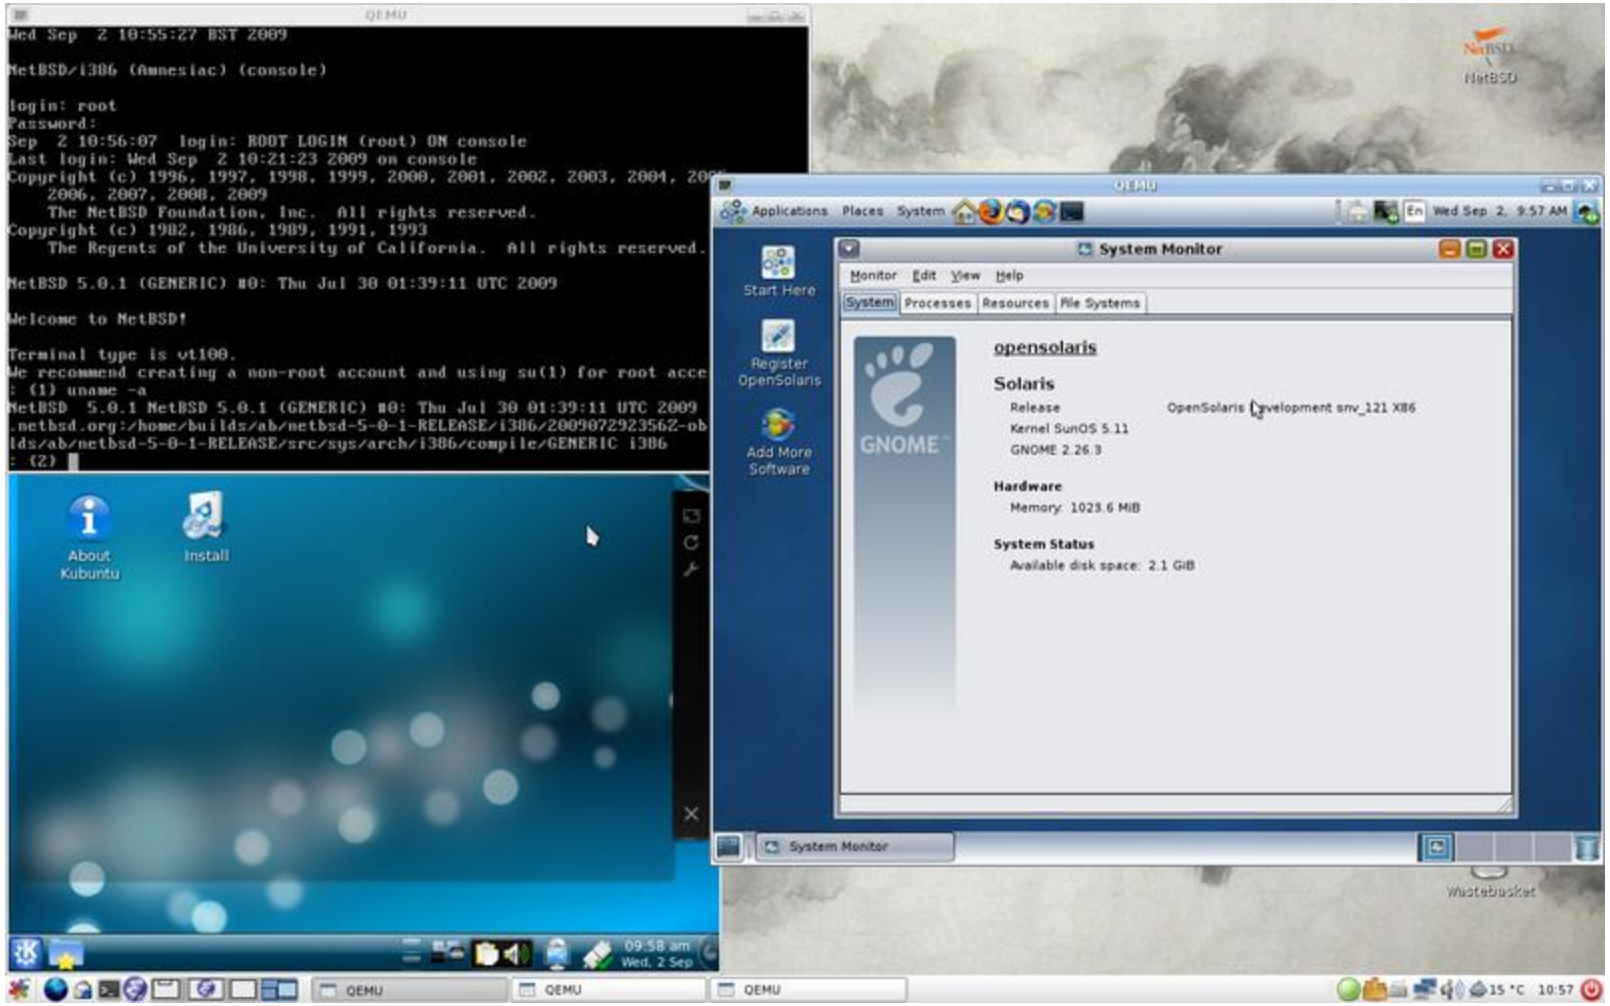

VMM architecture: KVM

- Hypervisor integrat în Linux (code based)

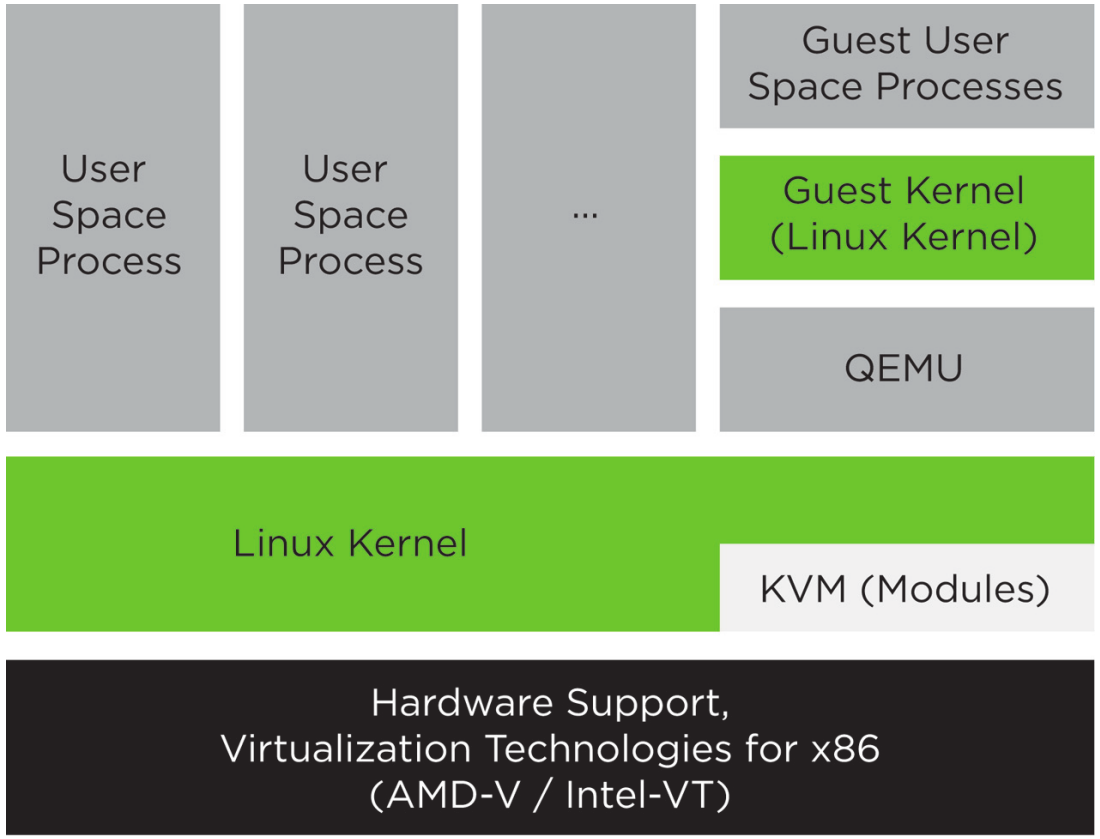

- Hypervisor: Kernel module
- Guest OS: User space process (QEMU for x86 emulation)
- Are nevoie de HW virtualization extensions

Virtuozzo / OpenVZ Containers

VZ / OpenVZ oferă container-based virtualization

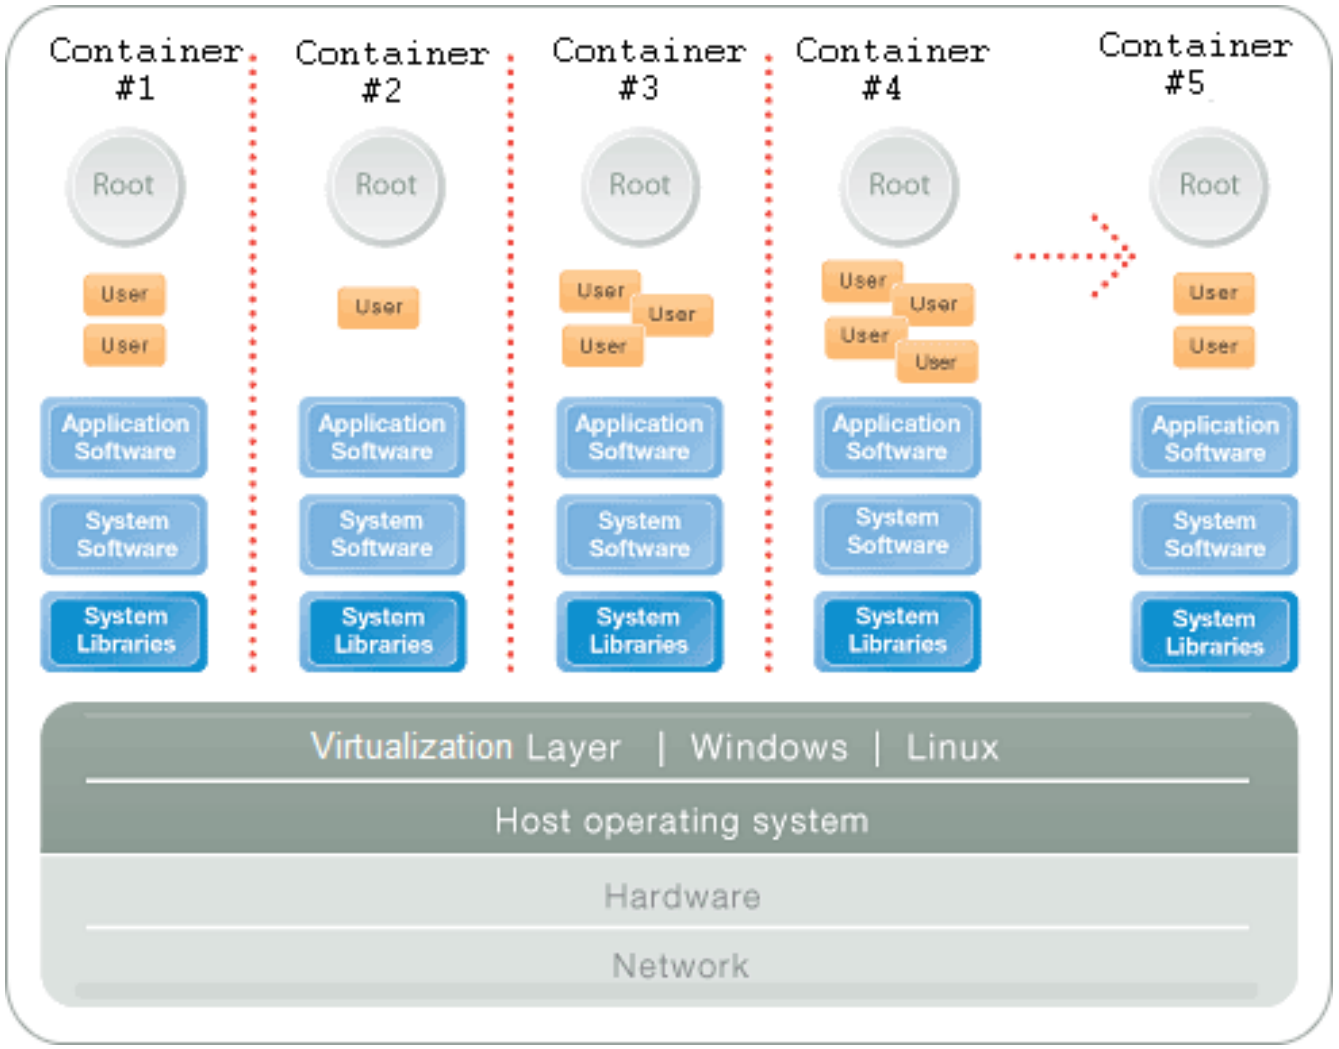

soluții bazate pe virtualizare

consolidare hardware

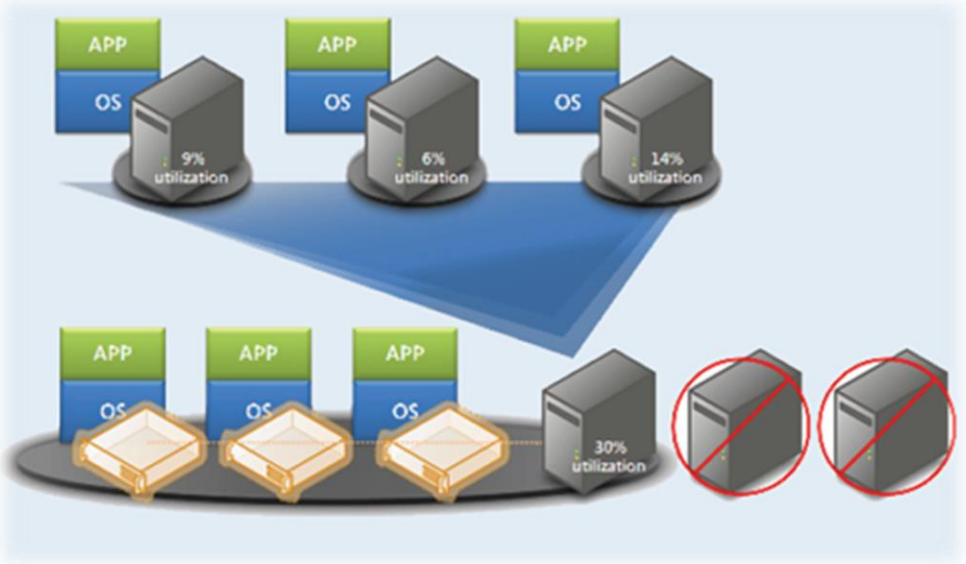

high availability

testing

disaster recovery

virtual private clouds

sandboxing

http://en.wikipedia.org/wiki/Sandbox_(computer_security)

forensic analysis

honeypots / honeynets

http://en.wikipedia.org/wiki/Honeypot_(computing)

HYPER-V OVERVIEW

Hyper-V architecture

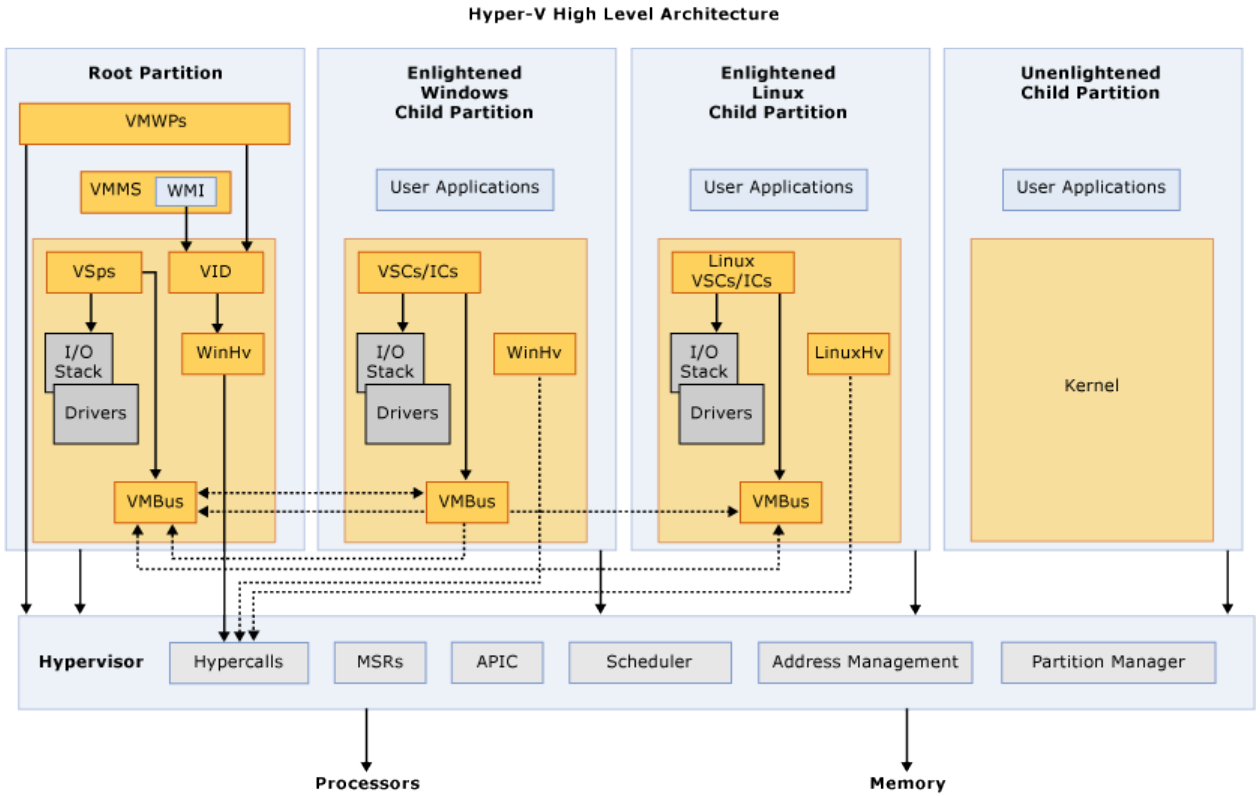

How does it work?

- VSP (virtualization service provider)
  - Hyper-V component in the”parent partition”
  - Communicates with the hardware drivers
  - Gives access to the host resources
- VSC (virtualization service client)
  - Drivers for”synthetic devices” installed in the enlightened guest OS
  - Exposes every virtual device and translates I/O requests
  - There’s always a VSP/VSC pair
- VMBus (virtual machine bus)
  - A high-speed point-to-point in-memory bus
  - Allows the communication between VSPs and VSCs through Hyper-V
- For Linux, every VSC has a DIM (Driver Interface Mapper)
  - DIMs interact with the Linux Kernel like any other driver
  - There’s also a “VSC core” based on each existing VSP

Emulated vs. enlightened

- Emulated drivers
  - Drivers are “emulated”
  - All requests targeting the hardware (HDD, network, etc) are not direct
  - “Translated” in both directions (VM-hw, hw-VM) by the hypervisor
  - They bring in a performance overhead
  - The emulated drivers are pretty similar to what we had in Virtual Server, with some enhancements:
    - Video = S3 Trio64+ SVGA (VESA)
    - Network = Intel/DEC ”Tulip” 21x4x
    - IDE = Intel 440BX chipset MB
- ”Enlightened” drivers
  - Also known as “synthetic drivers”
  - These make the VM ”hypervisor aware”
  - Written especially for virtualized environments => paravirtualization
  - They’re basically just pointers to the drivers in the ”parent partition”
  - Huge performance boost!

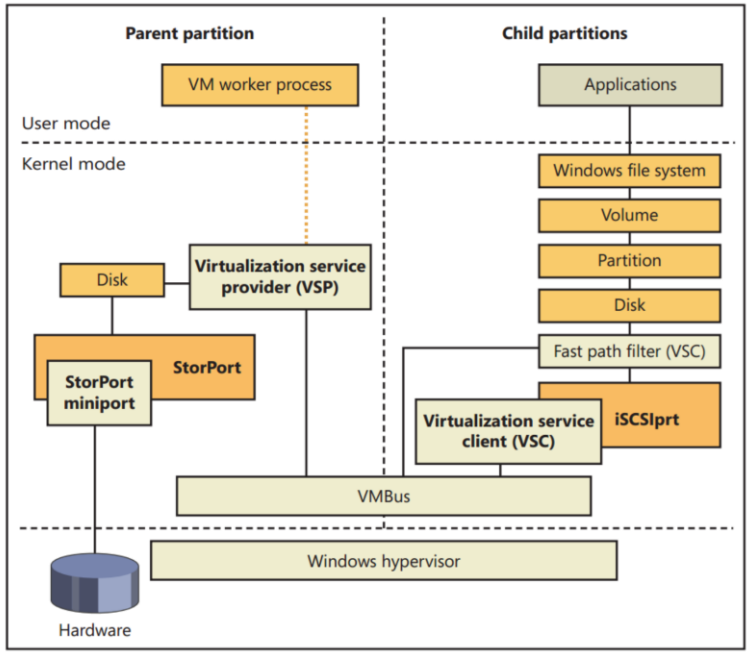

IS on Windows

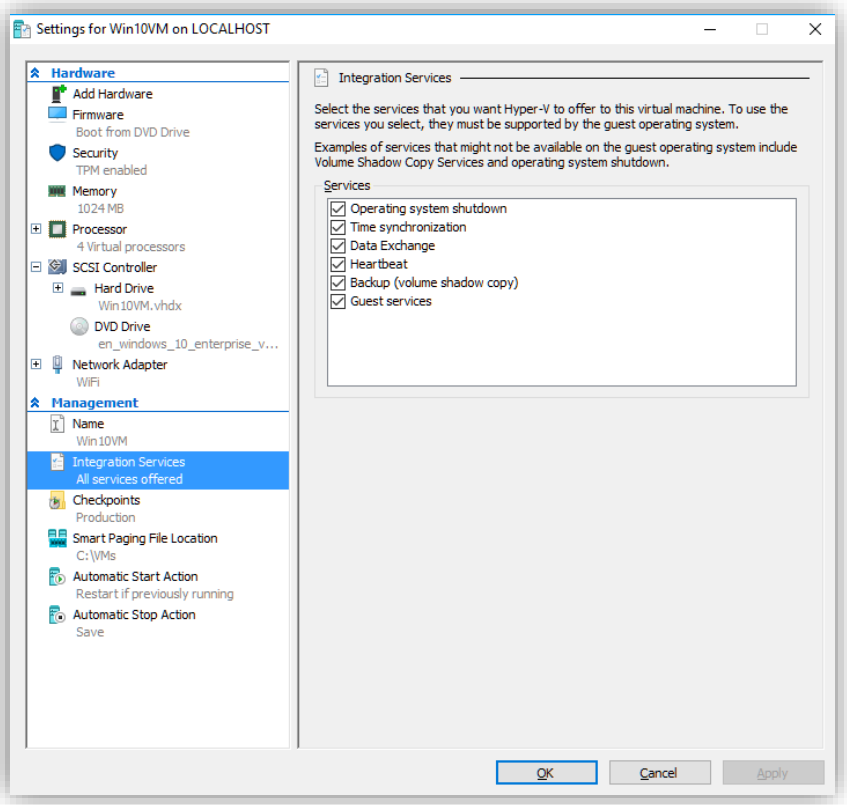

IS on Windows

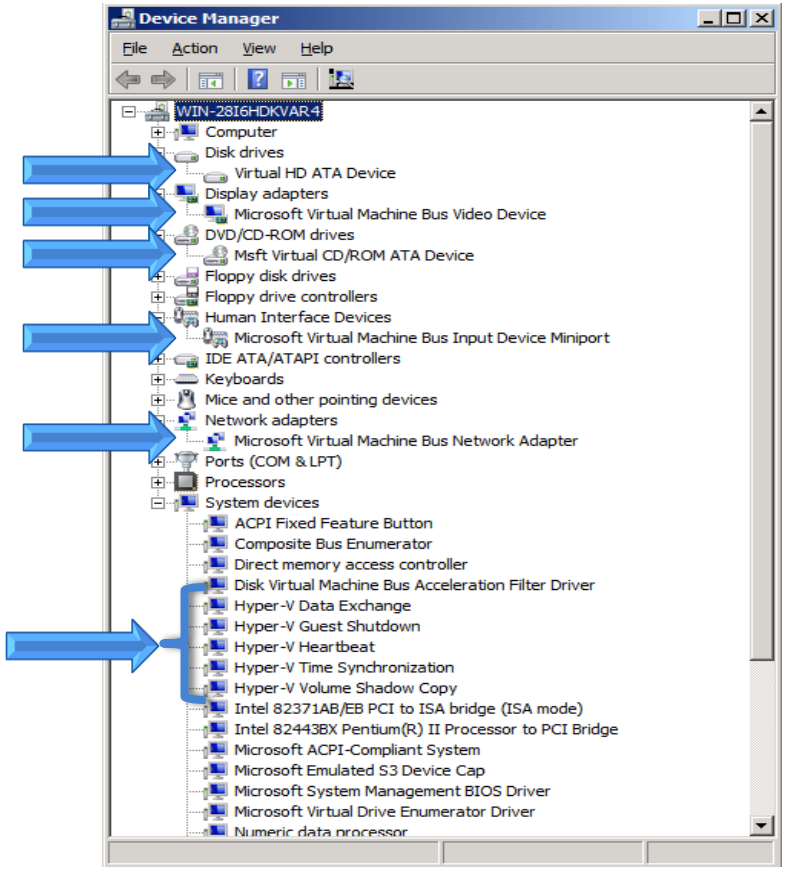

Hyper-V before WS 2012

- Hyper-V history
  - Introduced with Windows Server 2008
  - New version in Windows Server 2008 R2
  - Update with Windows Server 2008 R2 SP1
- Two manifestations of Hyper-V
  - Hypervisor-based virtualization feature of Windows Server 2008 R2
  - Microsoft Hyper-V Server (free download) containing only:
    - Windows Hypervisor
    - Windows Server driver model
    - Virtualization components
    - Failover Clustering
- Windows Server 2008 R2
  - VM mobility with live migration
  - VM high-availability improvements
  - Improved management of datacenters
  - Simplified method for physical and virtual computer deployments
  - Easier adding and removing VM storage
  - Improved virtual networking performance
  - Hyper-V processor compatibility mode
  - Improved memory management

Hyper-V in WS 2012/R2

- Performance & scalability improvements
- Disaster recovery
  - Hyper-V Replica
- Storage improvements
  - VHDX, Virtual Fiber Channel in the Guest (MPIO), NFS Storage with RDMA support, ODX, native deduplication, 4K native disk, iSCSI Target, NTFS online scan/repair
- New migration scenarios
  - Live Migration with no Clustering (via SMB share), LM with no shared storage, Live Storage Migration
- ~150 new PowerShell cmdlets
  - In-box metering
- Snapshots, backup and control
  - Online snapshot merge, Windows Server Backup Support, DM improvements, VM Priority
- Networking improvements
  - NIC Teaming
  - Extensible Virtual Switch
  - Bandwidth QoS
  - Dynamic Virtual Machine Queue (DVMQ)
  - Single Root I/O Virtualization (SR-IOV)
  - Receive Side Scaling (RSS)
  - Receive Side Coalescing (RSC)
  - IPSec Task Offload
  - Address Virtualization
    - Generic Routed Encapsulation (GRE)
    - Address Rewrite

WS 2016 - Hyper-V & Clustering

- Connected Standby support
- Discrete Device Assignment
- Host resource protection
- Hot add and remove for network adapters and memory
- Hyper-V Manager improvements
- Integration Services through WU
- Linux Secure Boot
- Nested virtualization
- Production Checkpoints
- Rolling Hyper-V Cluster Upgrade
- Clustered VM start order priority
- Storage quality of service (QoS)
- Shielded virtual machines
- VM configuration file format
- VM configuration version
- Windows Containers
- Windows PowerShell Direct
- Node Fairness
- Site Awareness
- VM Storage Resiliency
- Cluster Node Quarantine
- Workgroup & multi-domain clusters
- Cluster Cloud Witness

https://technet.microsoft.com/en-us/library/dn765471.aspx

WS 2016 - Storage

- Storage Spaces Direct
- Storage Replica
- Storage Quality of Service
- Data Deduplication improvements
- SMB hardening improvements
- Work Folders
- ReFS

https://technet.microsoft.com/en-us/library/dn765475.aspx

WS 2016 - Network

- Software Defined Networking (SDN) infrastructure
- Network Controller
- Network Function Virtualization (NFV)
- Datacenter Firewall
- RAS Gateway
- Software Load Balancer (SLB) and Network Address Translation (NAT)
- Standardized protocols (REST, JSON, WSMAN/OMI, SNMP)
- Flexible encapsulation technologies (VxLAN & NVGRE)

https://technet.microsoft.com/en-us/library/mt126106.aspx

Instrumente de management

SUITA SYSTEM CENTER

System Center - overview

Hyper-V & System Center

- **Virtual Machine Manager** – Manage physical and virtualized infrastructure
- **Operations Manager** – Monitor physical and virtualized infrastructure
- **Data Protection Manager** – Backup VMs and hosts
- **Orchestrator** – Automate VMs in a workflow
- **Configuration Manager** – Secure and manage VM operating systems
- **Endpoint Protection** – Protect VMs from virus and malware
- **Service Manager** – Provide service offerings with VMs
- **Operations Management Suite** (formerly Azure Operational Insights) – Looks for best practices for virtual machines and hosts
- **Windows Azure Pack** – Offer an Azure-like experience to employees and customers, built on top of your own private cloud

LINUX ON HYPER-V

How does it work?

- VSP (virtualization service provider)
  - Hyper-V component in the”parent partition”
  - Communicates with the hardware drivers
  - Gives access to the host resources
- VSC (virtualization service client)
  - Drivers for”synthetic devices” installed in the enlightened guest OS
  - Exposes every virtual device and translates I/O requests
  - There’s always a VSP/VSC pair
- VMBus (virtual machine bus)
  - A high-speed point-to-point in-memory bus
  - Allows the communication between VSPs and VSCs through Hyper-V
- For Linux, every VSC has a DIM (Driver Interface Mapper)
  - DIMs interact with the Linux Kernel like any other driver
  - There’s also a “VSC core” based on each existing VSP

Emulated vs. enlightened

- Emulated drivers
  - Drivers are “emulated”
  - All requests targeting the hardware (HDD, network, etc) are not direct
  - “Translated” in both directions (VM-hw, hw-VM) by the hypervisor
  - They bring in a performance overhead
  - The emulated drivers are pretty similar to what we had in Virtual Server:
    - Video = S3 Trio64+ SVGA (VESA)
    - Network = Intel/DEC ”Tulip” 21x4x
    - IDE = Intel 440BX chipset MB

- ”Enlightened” drivers
  - Also known as “synthetic drivers”
  - These make the VM ”hypervisor aware”
  - Written especially for virtualized environments => paravirtualization
  - They’re basically just pointers to the drivers in the ”parent partition”
  - **Huge performance boost!**

The story so far

- What's been done so far?
  - July 2009 – Microsoft contributes with over 20.000 lines of code in the Linux kernel
  - December 2009 – The drivers (in staging) are included in the 2.6.32 Kernel
  - July 1st 2010 – Microsoft presents at the Red Hat Summit
  - Official support list:

http://technet.microsoft.com/en-us/library/cc794868(WS.10).aspx


Pigs do fly: Microsoft unleashes 20,000 lines of Linux code

By Mary Jo Foley | July 20, 2009, 8:58am PDT

Summary

Nope, that's not a typo in the headline. Microsoft is releasing three Microsoft-developed

Microsoft is releasing three Microsoft-developed Linux drivers to the Linux community for possible inclusion in the Linux source tree.

This is the first time Microsoft has made Microsoft-developed code available directly to the Linux community. The Redmondians have released various pieces of code under different open-source licenses

- Linux Integration Services 2.1 (July 2010)
  - SUSE Linux Enterprise Server (10 SP1/SP2/SP3, 11)
  - Red Hat Enterprise Linux (5.2, 5.3, 5.4, 5.5)
- Linux Integration Services 3.4 (September 2012)
  - Red Hat Enterprise Linux 5.5-5.9 & 6.0-6.3 x86 and x64 (Up to 4 vCPU)
  - CentOS 5.5-5.8 & 6.0-6.3 x86 and x64 (Up to 4 vCPU)
  - SUSE Linux Enterprise Server 10 SP4 & 11 SP1/SP2
- Linux Integration Services 3.5 (December 2013)
  - Red Hat Enterprise Linux (RHEL) 5.5-5.8, 6.0-6.3 x86 and x64
  - CentOS 5.5-5.8, 6.0-6.3 x86 and x64

The road to “enlightment”

- Driver support for synthetic devices (v1 – original distro code, created by Citrix)
  - Synthetic network controller & Synthetic storage controller (IDE/SCSI)
- Fastpath Boot Support for Hyper-V (v2.0 – December 2009)
  - Block VSC – increased boot performance
- Timesync (v2.1 – July 2010)
  - The clock inside the virtual machine will remain synchronized with the clock on the host
- Integrated Shutdown (v2.1 – July 2010)
  - VMs can be cleanly shut down from Hyper-V
- Symmetric Multi-Processing (SMP) Support (v2.1 – July 2010)
  - Supported Linux distributions can use up to 4 virtual processors (VP) per virtual machine
- Heartbeat (v2.1 – July 2010)
  - Allows the host to detect whether the guest is running and responsive
- Pluggable Time Source (v2.1 – July 2010)
  - A pluggable clock source module is included to provide a more accurate time source to the guest.
- KVP (Key Value Pair) Exchange (v3.1 – July 2011)
  - Information about the running Linux VM can be obtained by using the Key Value Pair exchange functionality on the host (FQDN, Linux IS version, IP addresses, OS version/distro/kernel, CPU architecture x86/x64)
- Integrated Mouse Support (v3.2 – January 2012)
  - The cursor is no longer bound to the VMConnect window when used with the Linux Graphical User Interface
- SMP 32 vCPU support (v3.3 – June 2012)
  - Support for up to 32 vCPUs on certain distros
- Live Migration (v3.4 – September 2012)
  - Linux virtual machines can undergo live migration for load balancing purposes
- Jumbo Frames (v3.4 – September 2012)
  - Linux virtual machines can be configured to use Ethernet frames with more than 1500 bytes of payload

Linux Integration Services Development

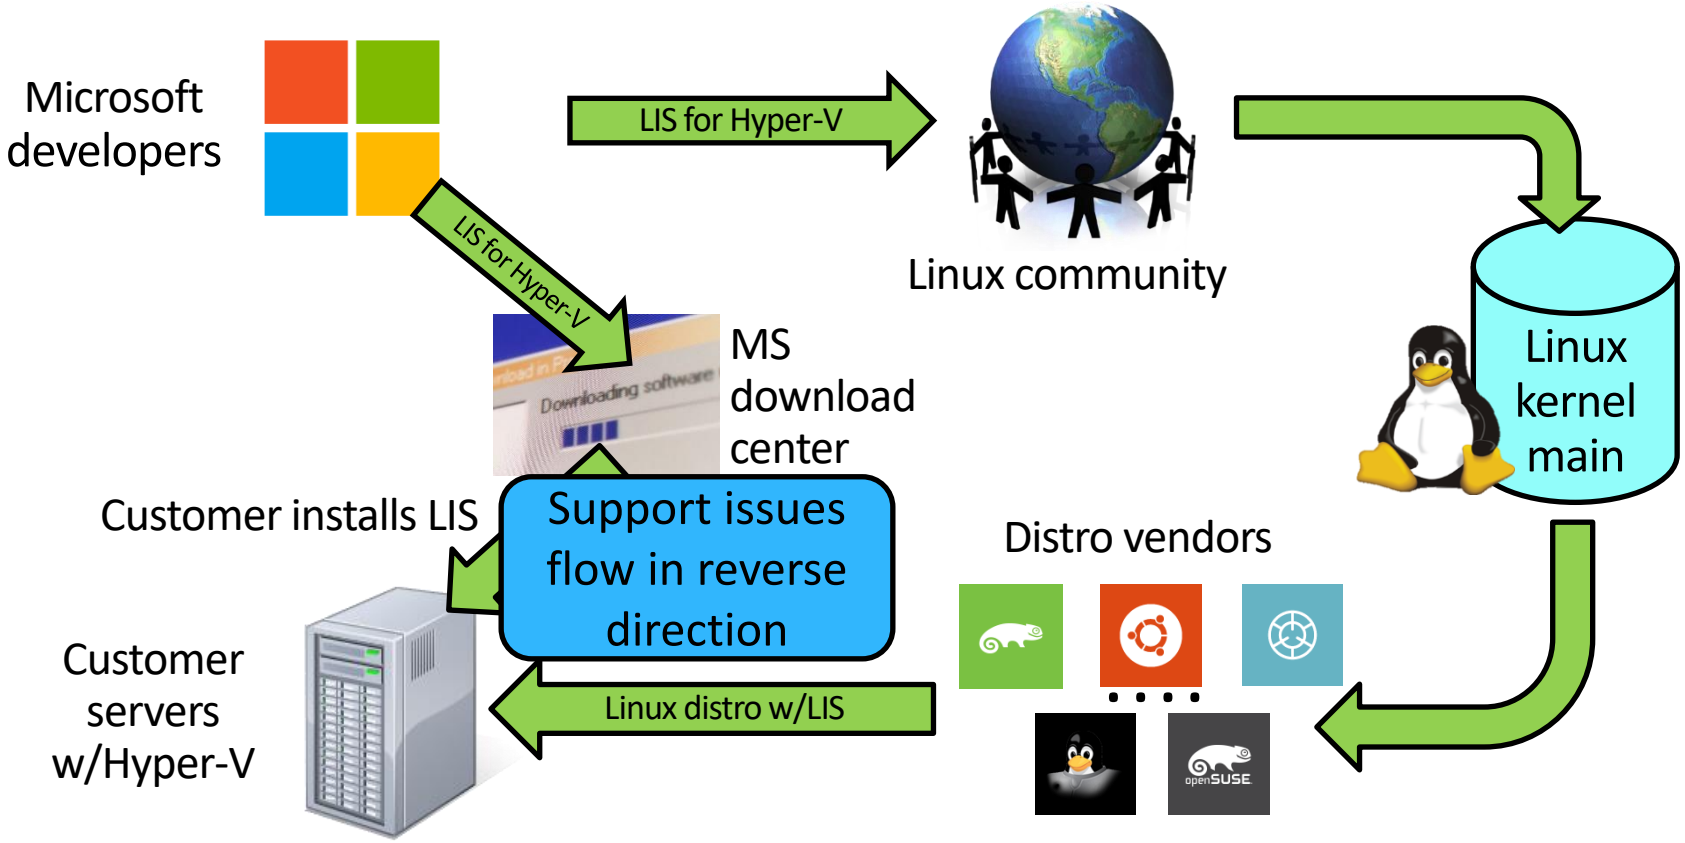

Microsoft & Linux/Unix

	Linux						UNIX		
	Red Hat	SUSE	CentOS	Ubuntu	Debian	Oracle	AIX	HP-UX	Solaris
Operations Manager	✓	✓	✓	✓	✓	✓	✓	✓	✓
Configuration Manager	✓	✓	✓	✓	✓	✓	✓	✓	✓
Endpoint Protection	✓	✓	✓	✓	✓	✓	*No Plans*		
Data Protection Manager	✓	✓	✓	✓	✓	✓			
Virtual Machine Manager	✓	✓	✓	✓	✓	✓			
Hyper-V	✓	✓	✓	✓	✓	✓			
Azure IaaS	✓	✓	✓	✓	✓	✓			

Linux IS components

- hv_vmbus – communication with the host
- hv_storvsc – storage VSC
- hv_netvsc – network VSC
- hv_timesource – pluggable time source
- hv_mouse – enlightened mouse
- hv_utils – graceful shutdown, timesync, heartbeat
- hv_balloon – dynamic memory
- hv_vss_daemon – live backups

MICROSOFT AZURE STACK

Azure Stack

The power of Azure with the control of the datacenter

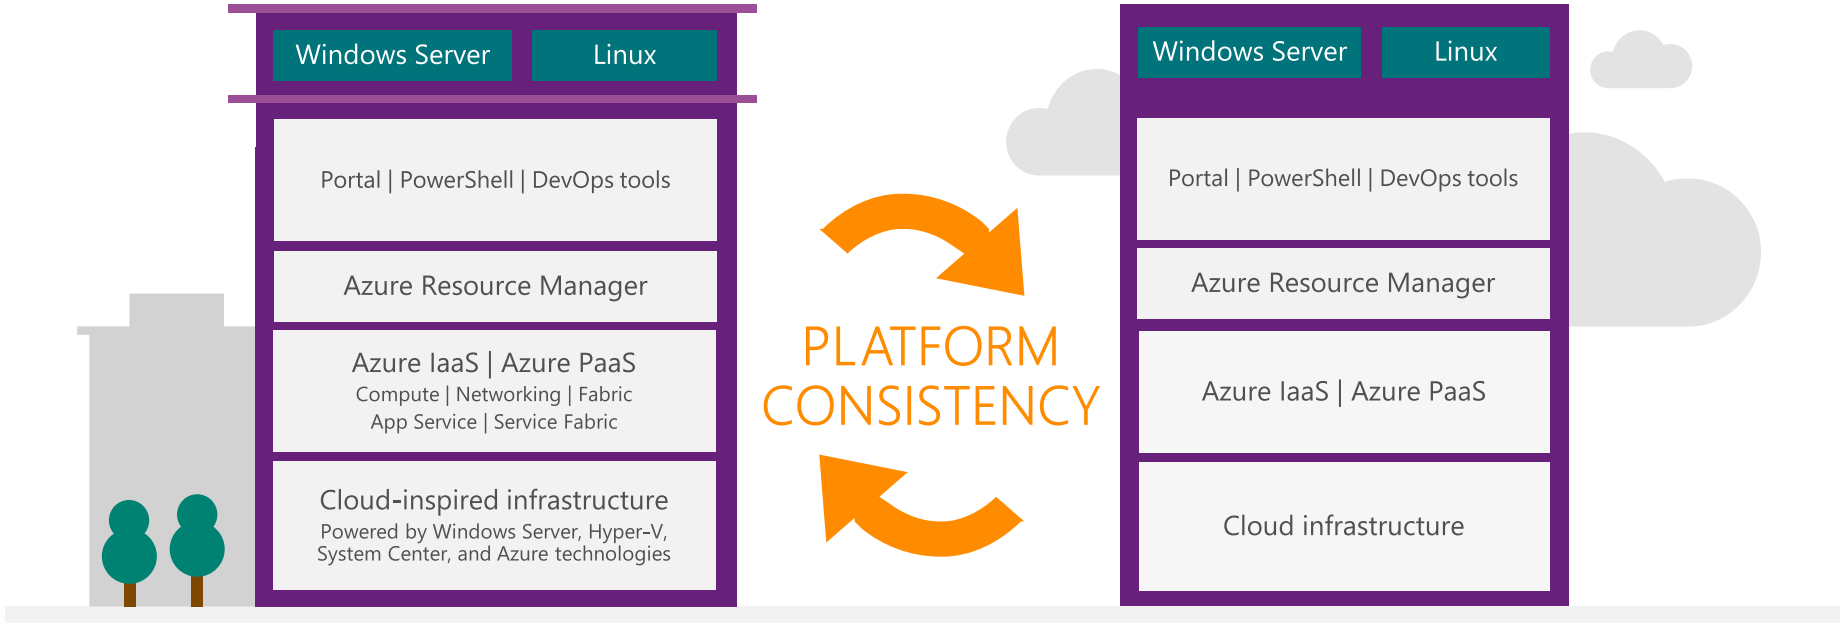

Microsoft Azure Stack
Private | Hosted

Microsoft Azure
Public

Azure Stack - Overview

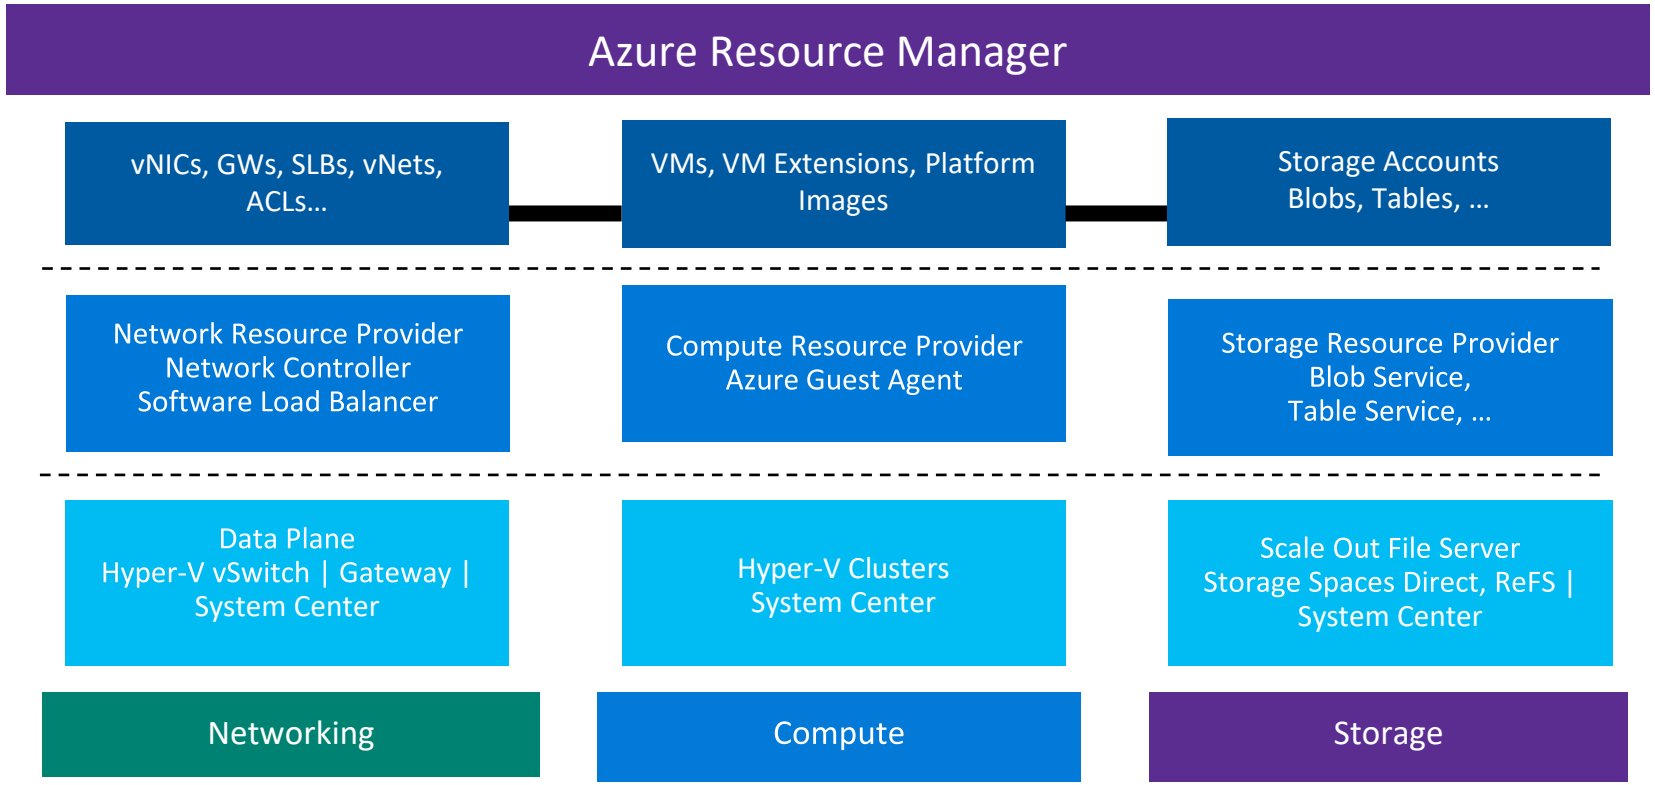

Cloud-inspired infrastructure

Cloud-optimized application platform | Cloud-consistent service delivery

Compute cluster	*Network controller*	*Storage cluster*	*Host Guardian Service*
Virtual machines	*Virtual networks*	*File servers*	*Shielded VMs*
Standard APIs	*Standard APIs*	*Standard APIs*	*Standard APIs*
Compute Industry-standard **servers**	Networking Physical **network**	Storage Industry-standard **disks**	Security TPM-enabled **hardware**




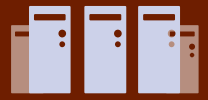

The deployment model is evolving

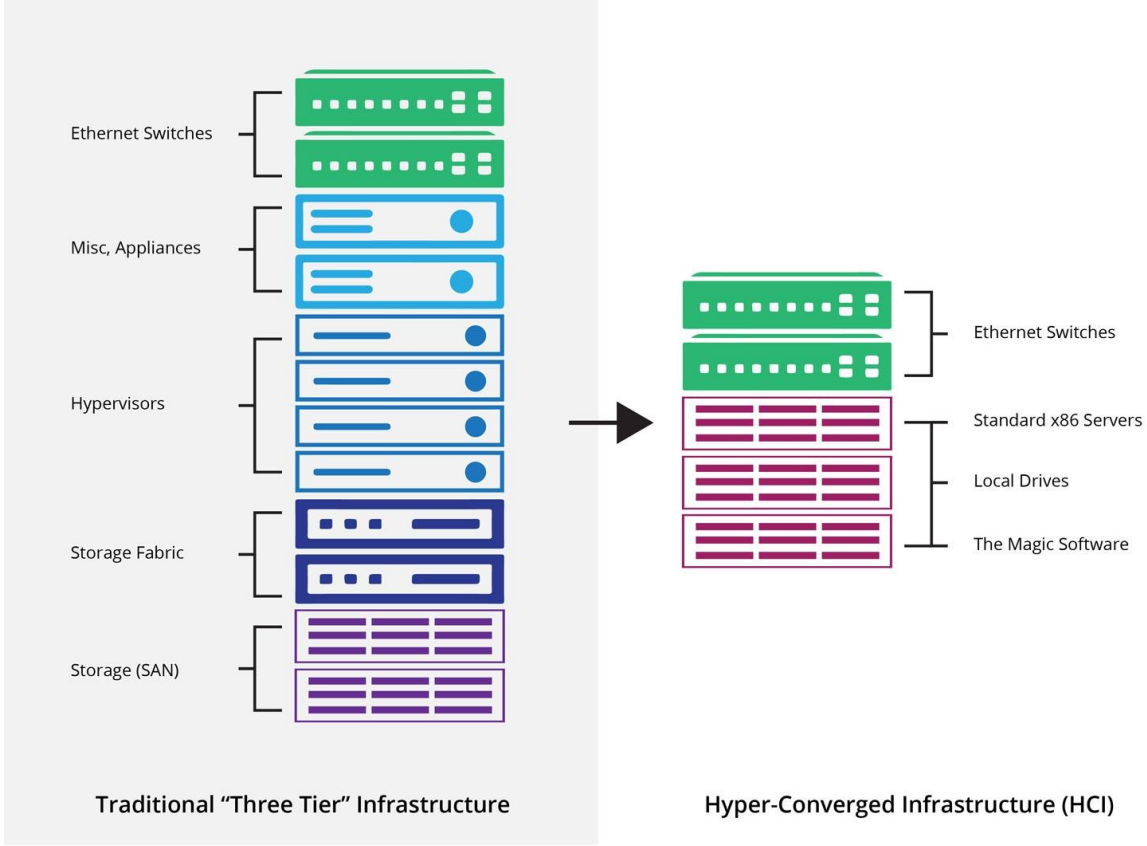

Azure Stack HCI (Hyperconverged)

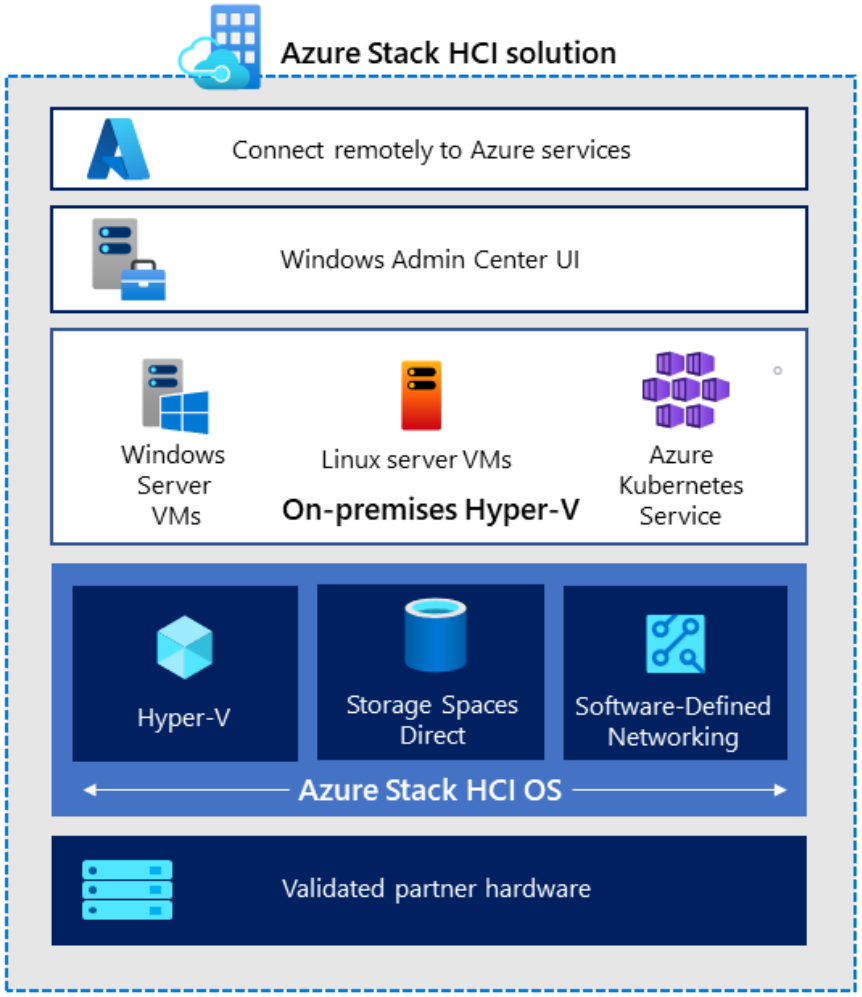

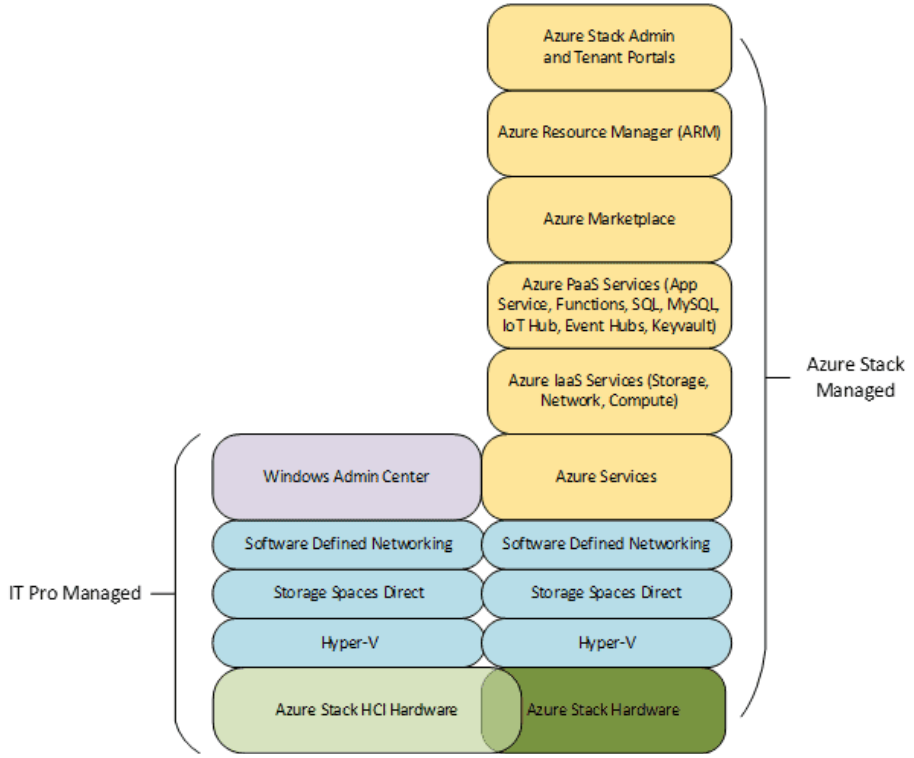

Soluții de scripting

POWERSHELL

Powershell v1.0

- Command Line Interpreter – CLI
- Script Execution Engine
- Help online vast
- Comenzi build-in (cmdlets), extensibile
- Suportă WMI, COM, CMD, etc.
- Construit peste .NET Framework 2.0

- Suport pentru:
  - Windows Server, Exchange, SQL, SCOM, SCVMM, SCDPM, Compute Cluster, OpenXML, IBM WebSphere MQ, Active Directory, Lotus Domino, VMWare Infrastructure, Windows 7, WDK, NetApp Data ONTAP, macOS, Linux, etc.

PowerShell – before & after

	INAINTE	ACUM
GUI	MMC	GUI-uri bazate pe PowerShell
Shell interactiv	CMD	PowerShell
Scripting	BAT în CMD	PowerShell
COM	WMI (VBScript)	PowerShell

Securitate în PowerShell

- Secure by design & by default
- Fisierele script sunt asociate cu Notepad
  - Scripturile nu vor rula
- CLI cere permisiunea de executare pentru scripturi neverificate
- Execuția PS se face:
  - Prin 'cmdlets', (programe .NET, scrise de un developer, compilate într-un DLL și încărcate de un script PowerShell)
  - Script-uri PowerShell ('.ps1')
  - Funcții PowerShell
  - Programe executabile

PowerShell v2.0

- Control Remote (PowerShell Remoting)
- Îmbunătățiri ale engine-ului (cmdlets, operators, debugging)
- Try-Catch-Finally
- GUI pentru PowerShell (PowerShell ISE)
- Hosting
  - - Run-space restricționat
  - - Delegare drepturi
- Background Jobs
- Operații tranzacționale
- Eventing
- Network File Transfer
- API-uri noi

PowerShell v3.0

- New commandlets
  - $Pv3 = Get-Command *
  - $Pv3.count
- PowerShell ISE
  - Single Command Pane
  - IntelliSense
  - Snippets
  - Collapsible regions
- Updatable Help
  - Update-Help
- Windows PowerShell Workflows
  - Parallel, Sequence & InlineScript keywords
- Remote Get-Module
  - Get-Module implicit remoting

PowerShell v4.0 and v5.0

- PowerShell Desired State Configuration (DSC)
  - Enables deployment and management of configuration data for software services and their environments
- Improvements in job scheduling
- Development using classes
  - Using formal syntax and semantics similar to OOP languages (Class, Enum and other keywords have been added)
- ConvertFrom-String cmdlet
- Archiving (zip) native support
- Script tracing feature
  - Enables detailed tracking and analysis of PowerShell scripting usage on a system
  - All script blocks are logged to ETW (Event Tracing for Windows) event log, under Microsoft-Windows-PowerShell/Operational
- And a lot more:
  - https://technet.microsoft.com/en-us/library/hh857339.aspx

Powershell v6.0 and v7.0

- Support for Windows, macOS and Linux
- The -Parallel switch for the ForEach-Object cmdlet to help handle parallel processing
- Near parity with Windows PowerShell in terms of compatibility with built-in Windows modules
- A new error view
- The Get-Error cmdlet
- Pipeline chaining operators (&& and ||) that allow conditional execution of the next cmdlet in the pipeline
  - The ?: operator for ternary operation
  - The ??= operator that only assigns a value to a variable when the variable's existing value is null
  - The ?? operator for null coalescing
- Cross-platform Invoke-DscResource
- Return of the Out-GridView cmdlet
- Return of the -ShowWindow switch for the Get-Help

PowerShell DSC

https://msdn.microsoft.com/en-us/powershell/dsc/overview

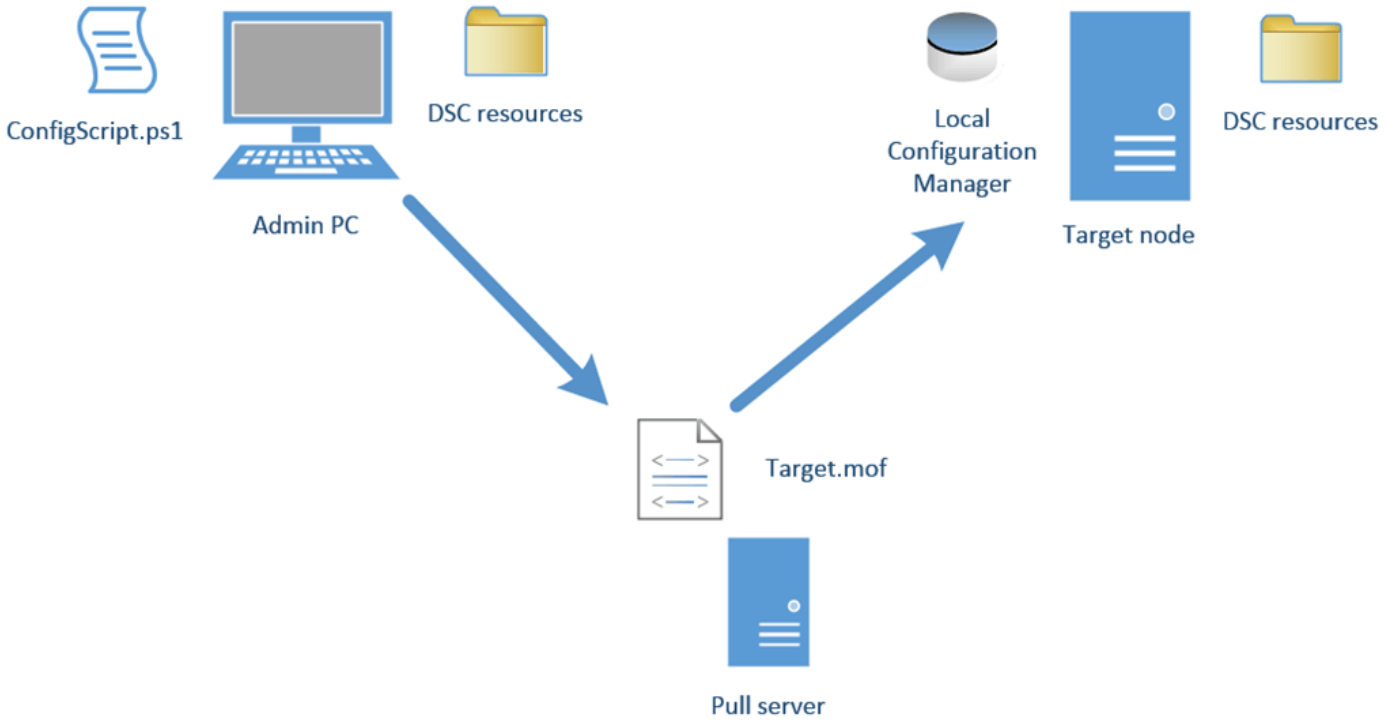

THANKS ☺