

Virtualization & Cloud Computing

Corso di: Big Data Architectures

Prof. Paolo Nesi

Department of Systems and Informatics, University of Florence


Via S. Marta 3, 50139, Firenze, Italy

tel: +39-055-2758515, fax: +39-055-2758570

DISIT Lab, Sistemi Distribuiti e Tecnologie Internet

<http://www.disit.dinfo.unifi.it/>

paolo.nesi@unifi.it
<http://www.disit.dinfo.unifi.it/nesi>



Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access



Datacenter, Definitions

- A computer factory/farm in which servers/computers are called HOSTS and are hosted and organized:
 - power, net maintenance, etc.
 - As: industrial computers, blades
 - 19 inc by xU (high)
- They can be exploited for private purposes, grid, cloud computing, renting/hosting, etc.



Inside a rack

- Including:
 - Computer,
 - SAN, NAS,
 - FW,
 - Switches,
 - etc.




Server, Host



H100


Infrastructure, definition

- A set of datacenters and clusters
 - Cluster: group of servers plus storage (if any)
- A storage as a set of
 - NAS: Network Area Storage
 - SAN: Storage Area Network
- Etc..



SAN vs NAS

- In most cases they are wrongly considered the same staff
 - It is true that: NAS/SAN typically are HD in Raid connected/shared to/by Servers
- **NAS: Network Area Storage**
 - Network HD sharing content at level of File via TCP/IP, HTTP, NFS, CIFS, etc., protocols
 - Multiple servers may access to them, the HD can be easily mounted, and see on the OS
 - Reduced performances with respect to SAN
- **SAN: Storage Area Network**
 - HD system segmented in LUN, mounted by the Server(host)/VM OS at level of disk block
 - The OS accessing to them has to format with its own file system
 - Multiple server may access to them
 - Possible protocols: low level IP, iSCSI, Fiber, etc.



DAS vs NAS vs SAN Storage Technologies

Direct Attached | Network Attached | Storage Area Network

DAS


Direct Attached Storage



- Direct Connection
- Single Server
- Low Cost

NAS

Network Attached Storage



- File Sharing
- Multiple Users
- Scalable & Affordable

SAN

Storage Area Network



- High Performance
- High Availability
- Enterprise Storage

	DAS	NAS	SAN
Connection	Direct	Network	Dedicated Network
Access	Single Server	Multiple Users	Multiple Servers
Cost	Low	Medium	High
Performance	High	Moderate	Very High

Consideration

- The fact that we have servers and storage is not a key for cloud
- Cloud is not a synonymous of data center
- Since 60' John McCarthy supposed to have a **number of servers providing services for all**
- Servers have been typically used
- Close to the end of the 90', the first Software as a Service, SaaS, started.
—Since.....



Housing

- Moving HW, Moving servers from local premise to an infrastructure:
- Buying HW and finding a safe location to placing,
 - Keeping HW on maintenance, H24
 - Making updates of SW if needed
 - Keeping HW on Power and UPS, Continuity, H24
 - Keeping Connection active, replicated, ...
 - Keeping HW on constant temperature, H24
 - Keeping HW on safe, H24, vigilance
 - Performing backups



Rack




Server stack

Motivations for Cloud computing and Virtualization

- **Computer Based Physical Systems (servers) are**
 - **not flexible and expensive**
 - The HW/SW has to be dimensioned for the workload worst case and not for the real/typical one,
 - The needed computational and network capabilities change over time
 - HW/SW can be hardly reused for different purposes since specific allocation is performed
 - **subjected to HW/SW failure**, and thus:
 - Software has to be installed again when changing hardware
 - Software has to be updated by skilled personnel
 - OS has to be reinstalled in most cases

Cloud computing and Virtualization

- **Computer Based Physical Systems (servers)** are
 - Subjected to high costs of setup, installation and use:
 - Power, conditioning, network connection, etc.
 - not scalable, to scale up frequently implies hardware change/re-engineering, (it was the gold period of the Moore's Law, of performance duplication every two years)
 - Data duplication, migration,
 - sharing databases, etc.
 - Distribution of activities among HW
 - Balancing or workload,
 - Add more HW,




Final Motivations

- **Reduction of costs for HW and operating system SW maintenance.**
 - High costs to guarantee **high availability**, HA: 99,999% of up time
 - High costs to guarantee high reliability
 - High costs to follow the HW/SW technological trends for performances, computational needs
 - Critical mass is needed to justify HW costs
- **Sharing resources**
 - Among for multiple applications and solutions
- **Needs of High flexibility** in terms of features
 - Most of the SW is becoming a services (licensing per year/month or users) and not anymore a product.
 - Many vendors provide complex SW systems in terms of services, on the basis of their consumption via network connection

Hosting, ..., definitions

- **HPC: high performance computing**
 - Solution based on cloud/grid for parallel execution of algorithms and tools.
- **Hosting a machine (computer/VM) into a datacenter**
 - According to some SLA (service Level Agreement), contract...
 - Renting a Computer/VM/Cluster into a data center
- **Hosting web portal into a datacenter/cluster**
 - Renting a web space via some SLA, contract, monthly rate to publishing web pages: service httpd
 - Additional services: mysql, php, asp, ftp, ssh, https, etc.
 - Features:
 - Space on disk, networking
 - Domain space, etc.



- Parties and purpose
- Scope of services
- Service availability / uptime
- Performance commitments
- Support service levels
- Incident management
- Maintenance and planned downtime
- Backup and disaster recovery
- Security obligations
- Data protection and data ownership
- Monitoring and reporting
- Exclusions
- Service credits and remedies
- Customer responsibilities
- Escalation process
- Change management
- Termination and persistent SLA failure
- Review and governance

Service Level Agreement

Severity	Example	Initial response	Target restoration
Critical / P1	Hosting service unavailable	15 min	2 hours
High / P2	Major degradation	30 min	4 hours
Medium / P3	Limited functionality	4 hours	1 business day
Low / P4	General request	1 business day	Best effort

Monthly availability | Service credit

- $\geq 99.9\%$ | 0%
- 99.0%–99.89% | 10%
- 98.0%–98.99% | 20%
- $< 98.0\%$ | 30% \

Availability


Availability % ▲	Downtime per year ^[2] ◆	Downtime per month ◆	Downtime per week ◆	Downtime per day ◆
55.555555% ("nine fives")	162.33 days	13.53 days	74.92 hours	10.67 hours
90% ("one nine")	36.53 days	73.05 hours	16.80 hours	2.40 hours
95% ("one and a half nines")	18.26 days	36.53 hours	8.40 hours	1.20 hours
97%	10.96 days	21.92 hours	5.04 hours	43.20 minutes
98%	7.31 days	14.61 hours	3.36 hours	28.80 minutes
99% ("two nines")	3.65 days	7.31 hours	1.68 hours	14.40 minutes
99.5% ("two and a half nines")	1.83 days	3.65 hours	50.40 minutes	7.20 minutes
99.8%	17.53 hours	87.66 minutes	20.16 minutes	2.88 minutes
99.9% ("three nines")	8.77 hours	43.83 minutes	10.08 minutes	1.44 minutes
99.95% ("three and a half nines")	4.38 hours	21.92 minutes	5.04 minutes	43.20 seconds
99.99% ("four nines")	52.60 minutes	4.38 minutes	1.01 minutes	8.64 seconds
99.995% ("four and a half nines")	26.30 minutes	2.19 minutes	30.24 seconds	4.32 seconds
99.999% ("five nines")	5.26 minutes	26.30 seconds	6.05 seconds	864.00 milliseconds
99.9999% ("six nines")	31.56 seconds	2.63 seconds	604.80 milliseconds	86.40 milliseconds
99.99999% ("seven nines")	3.16 seconds	262.98 milliseconds	60.48 milliseconds	8.64 milliseconds
99.999999% ("eight nines")	315.58 milliseconds	26.30 milliseconds	6.05 milliseconds	864.00 microseconds
99.9999999% ("nine nines")	31.56 milliseconds	2.63 milliseconds	604.80 microseconds	86.40 microseconds

HA

			99,8	99,98	99,998	99,999	
			0,2	0,02	0,002	0,001	
	60						
ore	24		17,52	1,752	0,1752	0,0876	ore anno
	365						
minuti	525600		1051,2	105,12	10,512	5,256	minuti
	100	0,2					
	525600	X					
			35,04	3,504			


Infrastructure as a Service, IaaS

- **the provision of infrastructure** with a relevant processing power, storage, network and other basic resources independently of application services of any type.
 - Example: a server in rent, a VM for rent
- Thus, it is possible to use an infrastructure of a provider to run your application, against a payment based on the renting (consumption) of the infrastructure, leaving the responsibility of the HW, OS, and SW (libraries, middleware, backup, storage, network, etc.) as well as of the application itself to the provider.
 - Amazon EC2 is an example of IaaS service.




Applications on the Cloud

- **Single-tier applications**, such as small web portals, do not demand a complex system, and
 - do not justify the acquisition of a host server !!!???
- **Multi-tier applications have multiple servers**, such as:
 - High performance web server, high number of users
 - SN: Social networks
 - CMS: content management systems
 - CRM: customer relationship management
 - CDN: Content delivering network
 - ERP: enterprise resource planning
 - P2P torrent tracker: see Piratebay
- *All typically parallel applications that may need grid, computation, storage, etc. may run on cloud, etc.*
-





Structure

- Elements of cloud Computing
- Virtualization 
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access

Concept of Virtual Server

- **Hyp:**
 - you can make the image (in a file) of the HD of a Physical Server (with its Operating System and files) and
 - put in execution the OS from the HD (file) by a program of "emulation"
- **Thus we have:**
 - a computer with an operating system OS
 - that hosts a program of "emulation",
 - which in turn, it is capable to run in separate processes the images of the corresponding OS VM



Typical Features for renting a Computer or a Virtual Machine into a Datacenter (1/2)



- **Requesting** to have one or more Computers/Hosts and/or VMs
- **Hardware, they have a counterpart in the SLA:**
 - CPU: 32/64 bit, number of cores/CPU, frequency of work, intel/amd/arm, etc.
 - Eventual vCOREs
 - RAM memory: size and frequency of work
 - HD: space, speed (7.2, 15Kgiri); SSD ?
 - security level/RAID, type SAS/SATA, SCSI
 - NAS/SAN, Network area storage/SAN, fiber/internet: size, RAID, etc.
 - Network features:
 - Number of connections/cards, NICs
 - number of IP addresses, static/DHCP
 - Transfer rate: minimum guaranteed, maximum possible, down/upload
 - Maximum transferred bytes: per day, per month, etc.

Typical Features for renting a Computer or a Virtual Machine into a Datacenter (2/2)

- **Software:**
 - Operating systems: versions, with maintenance, etc.
 - Software preinstalled into the Computer/VM, see in the following
 - Eventual SDK
 - Eventual connection with Cloud Services: Back, emails, fax, recovery, etc.
- **Services:**
 - Periodic back up: details on HD space
 - Access to VM/Computer: remote desktop or KVM tool
 - Reboot or not of the Server, for example via Plesk.


Virtual Machine, Virtualization

- **Virtual Machine:**
 - An image of an operating system that can be put in execution into a real host/computer creating a virtual computer that exploits a part of all of the host resources
 - E.g.: Host may be Linux-like while the VM may be Window, Mac, Linux, ...
- **Virtualization**
 - Transforming a physical computer into a VM, virtual machine, hosted on some Host computer
- **Hypervisor** (VM Monitor) to manage the several VMs on the host





The Virtualization

- Process by which it is possible to create a virtual machine, VM,
- For example by:
 - An installation DVD
 - Cloning a physical machine (P2V)
 - Clone from VM
 - Template from..







The Hypervisor

- It is the essential OS, that can be installed in a Host and can put in execution and manage one or more VM
- The Host is a server HW that host the Hypervisor
- ESX of VMware is an Hypervisor.



Performance Monitoring of VMs on a Host



Cloud computing and Virtualization

- **Most of the Host/datacenter capabilities** are not exploited at 100% in every time instant
 - They are typically present in large industries/institutions or in large services: google, amazon, tiscali, Dada, ibm, cnr, etc
 - Their size is typically defined on the basis of the workload worst case
- ***If they are big***: the exploitation of the remaining resources for cloud/hosting is a solution to recover money, since a non working machine (like a caw) is a costs without any return.
- ***If they are small***: it could be a solution to host machines on a professional infrastructure to reduce annual costs for HW/SW
 - Delegating to cluster owners the costs for *maintenance, renovating hardware, renovating software, network costs, back up costs, power supply, etc.*

Benefits of Virtual Machine

- **Main benefits:**

- Separation of OS+SW with respect to the needed HW
- Exploiting legacy solutions which can be wrapped into a VM and protected with a physical firewall without reinstalling and recompiling old applications.
- The simple upgrade can be obtained by giving more CPU instead of changing HW and reinstalling all SW elements.

- **For example:**

- An old Linux Server hosting several web portals with a old versions of: MySQL, PHP, etc.. and many configuration aspects: users, mailing lists, etc. very time consuming to port on a new server
- An old Cobol application running only on an old Windows 2000 Server, which cannot be recompiled into a new Windows Server 2008 at 64 bits without spending months of work.
- A Cobol application running on machine based on an IBM system 36....
- Etc.

VM on PC workstations


- On a single Computer it is possible to put in execution a VM by using a standalone Hypervisor
- **VMware Workstation** can play on Windows:
 - VMware Virtual Machines with: MS Win, Linux, Mac, ChromeOS, etc.
- **VMware Player** can play on Windows or Linux
 - Free of charge
 - VMware Virtual Machines with: MS Win, Linux, Mac, ChromeOS, etc
- **VMware Fusion** can play on MAC:
 - VMware Virtual Machines with: MS Win, Win 7, ChromeOS, etc.
- **Microsoft Virtual PC** on Windows 7:
 - Free of charge
 - Create VM with Win XP, Linux Ubuntu,

Software/Components preinstalled into rented Computers/VMs

- **Several kinds of software components/tools** that can be accessible on rented VM and/or Computers.
- **Availability of HW/SW:** for example at 98% or 99.999%
- **Typical components** that could be requested:
 - DB: MySQL
 - FTP: server and client
 - Web Server: Apache, IIS
 - Add-on: PHP, Perl, Python, cache tools on several levels
 - SMTP address, antispam
 - Web Application Server: TomCat,
 - And: Antivirus, backup, email, Drupal, Joomla, etc...
- **A full server can be customized, so that any other tool can be installed as well from the user**

Cloud Computing with VMs


Several Hypervisors on a Clusters



Clusters, generic definition


- **For Cluster, it is typically intended:**

- Several kinds of computers and/or VMs used to compose a group of resources, that can be managed together, and may be sharing the same domain or not.
- a set of computers/VMs into a datacenter which are dedicated to a unique problem, for example:
 - A microgrid, an Hadoop cluster,
 - A social network.....
 - A web portal with its multi-tier servers: front-end portals, balancer, database computers, backoffice microgrid nodes, etc.



From Clustering VM to Cloud

- The first step has been the exploitation of unused resources to provide them as a service to third party.



Approaches for Virtualization

- **Hardware Emulation**



- Full emulation of Hardware devices and features
- It is possible to use an original OS without changes, may be with some drivers installed, but **not kernel** changes
- Higher isolation among VMs, strong robustness, limited efficiency.
- **Used by** VMware core
- Typically 10% of overhead

- **Paravirtualization (total or only on some devices)**

- The execution is performed via specific API and the hosted OS has to be modified to use them instead of the original HW
- Lower isolation, higher efficiency
- Lower robustness: VM crashes may crash the whole system
- **Used by:** HP-VM, Xen (both of them which can also go in emulation mode)
- **Used by VMware:** VMXNET (100 Gbps net), PVSCSI in vSphere4
- Typically 2% of overhead


Para-virtualization

- It is intended as a solution to cut out a part of the I/O stack, for example for the HD access or network.
- The following example is related to HP VM v. 3.5



“Speeds and Feeds”

Optimization for the Highest Consolidation Ratios



VM Scale Up

- 8-way vSMP and 255 GB of RAM per VM

Hardware Scale Up

- 64 cores and 512GB of physical RAM

Hardware Assist Purpose Built Scheduler

- Lowest CPU overhead

Hardware Assist Page Sharing Ballooning

- Maximum memory efficiency

VMXNET3 VMDirectPath I/O

- Wirespeed network access

Storage stack optimization VMDirectPath I/O


- Greater than 360k iops per second
Lower than 20 microsecond latency

■ Current ■ NEW
DISIT lab, Big Data Arc. 2026

L1
50 Cores
X
2 TRAM

G NIC

ESX 4.0 Performance with SQL Server 2008




- ESX achieves 90% of native performance on 4.0 vCPU VM
- Workload transaction latency unchanged between ESX 4.0 and Native

Host: risorse primarie


Host Profile

- Memory Reservation
- Storage
- Networking
- Date and Time
- Firewall
- Security
- Services
- Users and User Groups
- ~~Security~~



Limiting VM Resources

- **VM Resources** (CPU, Mem, HD, net.) consist also in providing support for:






—

Dynamically providing resources over the reserved values that can be negotiation into the SLA/ contract.


- **Controlling and limiting** access/cost and the exploitation of HW resources:
 - A limit on the number of CPUs
 - A limit on the number of Clocks, over of a reserved number of clocks
 - A limit on the maximum size of the RAM, over of the reserved number of Mbytes
 - A limit on the size of the HD size (SAN/NAS access)
 - A limit on the number of network cards, number of Mbps, etc.

Model for Memory

- Memory
 - Total host memory
 - For each VM
 - Max memory for the OS
 - Reserved Memory
 - Used Memory (time)
 - Limited memory
- Be careful to overcommitted resources on the Host respect to the total CPU and Mem of the hosted VMs





VM: Memory Definitions




- **Reserved:** not available for the other VMs on the same host!!
- **Active:** memory actually used
- **Ballon:** Memory requested by the VM to the host
 - This memory is shared among several hosts
- **Shared:** memory available for sharing
- **Swapped:** amount of memory that has been swapped into the HD by the VMKernel of the host
 - This parameter has to be low as much as possible
 - Increase the reserved to make it low or reduce the number of VM on the host

Model for CPU

- CPUs
 - Total host Clock
 - For each VM
 - Max #cks/s host
 - Reserved #cks/s
 - Used #cks/s
 - Limited #cks/s
- Ck/s $\rightarrow$ Hz
- 1 GHz $\rightarrow$ 1000 MHz
- # of CPUs, vCPU




■ Usage ■ Ready ■ Ready for 0 ■ Ready for 1 ■ Ready for 2 ■ Ready for 3



Model for CPU

- A host may have multiple CPUs with multiple Cores
- A host with 4 CPUs at 3Ghz and 6 core each has
 - 24 cores at 3Ghz
 - 72.000 MHz
- All VMs on the same host share this amount
- CPU cores vs vCPUs



Virtual Resources, 1/2

- **The idea of Virtual Resources** (CPU, Mem, HD, net.) consists in providing a number of resources larger than those that physically available and manage them virtually and/or dynamically
- **For example, one Host may have 2 VM and HW resources:**

- 2 cores at 1300 Mhz
- 1.8 Gbyte RAM
- 2 network cards
- Best case:
 - CPU=400+800 Mhz
 - RAM=400+400 Mbyte
 - 2 Network cards shared
- Worst case → no resources enough:
 - CPU=3 cores at 1300 Mhz
 - RAM=2 Gbyte
 - 2 Network cards shared



VM1:

- 1 CPU 1300 Mhz, 400 Mhz reserved
- 1 Gbyte RAM max, 400 Mbyte reserved
- 2 network cards, 2 IPs

VM2:

- 2 CPU 1300 Mhz, 800 Mhz reserved
- 1 Gbyte RAM max, 400 Mbyte reserved
- 2 network cards, 2 IPs

Virtual Resources, 2/2

- The server hosting the VMs
 - Min Mhz:
 - 400Mhz + 800Mhz
 - **Max Mhz:**
 - **1300Mhz + 2*1300Mhz**
 - Min RAM:
 - 400Mbyte + 400Mbyte
 - **Max RAM:**
 - **1000Mbyte + 1000Mbyte**
 - Network:
 - No limits on the number of virtual IP addresses/cards

VM1:



1 CPU 1300 Mhz, 400 Mhz reserved
1 Gbyte RAM max, 400 Mbyte reserved
2 network cards, 2 IPs

VM2:

2 CPU 1300 Mhz, 800 Mhz reserved
1 Gbyte RAM max, 400 Mbyte reserved
2 network cards, 2 IPs


Model for network

- An host may have multiple network cards/connectors with one physical IP each or some of them bridged
- Each single VM may have:
 - access to a number of networks (virtual toward physical)
 - Installed multiple network cards
 - Each card can have one IP and may be connected to a network
- This means that an host may be connected to multiple networks with multiple IP
- Network cards may be connected to:
 - Physical switch
 - Virtual switch




Example of Cloud Networking

- Example of allocation of a three tiers solution
- 3 different networks (N1, N2, N3)
- 2 Host: H1, H2
- 3 VM: VM1, VM2, VM3
- → SW1 is convenient to be software (full virtual). If it is HW, the VM1 can be moved from H1 to H2, dynamically
- → SW2 could be fully virtual if VM3 is moved on H1




Example of Cloud Networking

- Example of allocation of a three tiers solution
- 3 different networks (N1, N2, N3)
- 2 Host: H1, H2
- 3 VM: VM1, VM2, VM3
- → SW1 is convenient to be software (full virtual). If it is HW, the VM1 can be moved from H1 to H2, dynamically
- → SW2 could be fully virtual if VM3 is moved on H1



Example of Cloud Networking

- Example of allocation of a three tiers solution
- 3 different networks (N1, N2, N3)
- 2 Host: H1, H2
- 3 VM: VM1, VM2, VM3
- → SW1 is convenient to be software (full virtual). If it is HW, the VM1 can be moved from H1 to H2, dynamically
- → SW2 could be fully virtual if VM3 is moved on H1




Example of Cloud Networking

- Tutti gli host che stanno nello stesso cluster devono avere la stessa configurazione di rete in termini di NIC e Switch altrimenti le VM non possono essere spostate da un host all'altro in modo safe.
- È compito del gestore evitare che possano accadere cose di questo tipo. Che una VM vada a finire in un Host che non ha le risorse necessarie anche in termini di configurazione di rete.

Virtual Switch: vSwitch0

[Remove...](#) [Properties...](#)




Virtual Switch: vSwitch1

[Remove...](#) [Properties...](#)



Virtual Switch: vSwitch2

[Remove...](#) [Properties...](#)



Network and Virtual Networks

- The same VM with access to 2 different network via real network adapters
- A virtual network


Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access

Model for HD, Storage


- Storage:
 - HD contained into VM
 - NAS: External storage to VM and mounted for their usage
 - SAN for storing VMs
 - In any case connected with some Operating Systems (may be dedicated)

- HD with fault tolerant solution:
 - Redundancy: replications into the Storage
 - Distributed/federate/replicated: replications into multiple storage




SDN

Internal replications: HD RAID



RAID 0: HD0+HD1

RAID 1: HD0 RT cloned on HD1




RAID 5: gestione del guasto di un disco
valido fino a 14

In caso di rottura di uno, il quarto
viene ricostruito a caldo

RAID 6: per avere la doppia parità


Internal replications: RAID 50 to scale



- RAID 50: A1,A2,A3,A4,A5,A6 and three Parities
 - Min 9HD, can support failures of three HDs
 - improves upon the performance of RAID 5
- In this case:
 - Size: 6 over 9 HD, you lose a HD per R5
 - Read/Write speed up: $n(m - 1) \times \text{HD}$
 - $\rightarrow n=3, m=3 \rightarrow 6x$ in this case



Schema RAID 60 con Hot Spare



Legenda



Dati
(contiene i dati utente)



Parità
(informazioni di parità RAID 6)



Hot Spare
(riserva per sostituzione automatica)


Caratteristiche

- ✓ RAID 60 = striping su gruppi RAID 6
- ✓ Tollera fino a 2 guasti per set RAID 6
- ✓ Hot spare per sostituzione automatica
- ✓ Buon compromesso tra capacità, prestazioni e resilienza

Simple configuration

- **VM stored on a storage and executed on a remote hypervisor host**

- At the start is uploaded on ESX memory and RUN
- Each write is passed back to the storage, SAN
- All via network connection




- **VM stored and executed on the same hypervisor host**

- At the start is uploaded on ESX memory and RUN
- Each write is passed back to the storage
- All via internal communication
- → higher performances
- → lower flexibility




Structure


- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access




Cluster of Hosts in DRS: network config




Cluster of Hosts in DRS: all needed config iSCSI





Schema SAN – RAID 60 – Doppio Storage – Doppio Alimentatore – Doppio RAID



Legenda collegamenti

- Fabric 1 (SAN A)
- Fabric 2 (SAN B)
- Replica tra storage
- Alimentazione

Caratteristiche principali

- RAID 60 (striping + doppia parità) per alte prestazioni e resilienza
- Doppio storage (primario e secondario)
- Doppio controller RAID per ogni storage
- Doppio alimentatore per ogni storage
- Doppia fabric SAN (A e B) con switch ridondati
- Replica tra storage (sincrona o asincrona)
- Connessioni ridondate da ogni server a entrambe le fabric
- Alta disponibilità e tolleranza ai guasti


Vantaggi

- Continuità operativa (nessun single point of failure)
- Elevata tolleranza ai guasti
- Alte prestazioni e scalabilità
- Protezione dei dati (RAID 60 + replica)
- Manutenzione senza interruzioni
- Adatto a ambienti mission critical

Elettrogeno, Double UPS for Servers


Schema di alimentazione di un datacenter

Gruppo elettrogeno, gruppi di continuità (UPS) e server c/er con alimentatori destri e sinistri




Alta disponibilità: in caso di mancanza rete, il gruppo elettrogeno alimenta gli UPS che mantengono in servizio i server. Ogni server è dotato di due alimentatori (sinistro e destro) collegati a linee indipendenti (A e B).

UPS, gruppi di continuità



Structure


- Elements of cloud Computing
- Virtualization
- Storage
- Clusters
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access



Model for storage consumption



- HDD Typical:
 - 100 Mbyte per sec. (high/max value)
 - Seek time 5-7 Msec.
 - 5000-7500-10000-15000 rpm (revolution per minute), spindle
 - Thus latency: (5.56, 4.17, 3.00, 2.00) respectively
 - Format Block size: 2Mbyte, 4, ..8
- Cache, no cache;
different kind:
- SSD:
 - 30% faster





Storage Consumption


- Each host may have
 - internal Storage one or more HD
 - External storage may be mounted (SAN or NAS)
- External storage may be shared or reserved for a specific Host
- Each VM has its permanent memory consumption:
 - Image for the VM
 - + space for snapshots (HD incremental changes + memory when the snapshot is performed on running)

Examples


Host

Space Utilization for 145-Storage-A-4

By File Type




By Virtual Machines (Top 5)




Virtual Machine

Space Utilization By File Type



Space Utilization By Datastores (Top 5)





Virtual Machine: Snapshots

- **Working on VM Snapshots**

- Creating a Snapshot:

- A point from which it is possible to reboot, restart
- Consuming HD space
- Making back up since the core image of the VM is not changed, changes are confined in the files representing the last status “you are here” and not the previous conditions

- Restarting from a past snapshot

- Losing current point: “you are here”, to avoid this do another snapshot!!

- Deleting a past snapshot

- Recovering HD space, removing a past restarting point

- Removing all snapshots

- **Defragmenting images of the HDs into the VM**


VM Snapshots

- **VM Snapshots can be at VM Off or ON**
 - Snapshots of running VM have implications...
- **Removing Snapshots**
 - **Defragmenting images of the HDs** into the VM
 - Consolidating the changes in the previous version
- VMware WS has an automated Snapshot model to plan the periodic snapshotting of the VM, for example:
 - every hour, day, week, ...
- A way to make back up
 - A different way can be to clone the VM on different host or NAS. In most cases, the cloning implies the lost of performed snapshots

Snapshot process


• Original version

- 1 VM
- 12 Gb ram
- 300Gbyte hd
- 72+102+12+8 Gbyte snapshots



• After T1 Removal


- 1 VM
- 12 Gb ram
- 300Gbyte hd
- 72+122+12+8 Gbyte snapshots



Snapshot process


• Original version

- 1 VM
- 12 Gb ram
- 300Gbyte hd
- 72+102+12+8 Gbyte snapshots




• After T1 Rimoval

- 1 VM
- 12 Gb ram
- 300Gbyte hd
- 72+122+12+8 Gbyte snapshots




Virtual Machine: Snapshots, 1/3




When a number of snapshots are in place:

the access to HD is deteriorating since the data are not contiguous, the data access implies multiple seek for multiple accesses at different segments of the same data in different file representing the different progressive snapshots!


Virtual Machine: Snapshots, 2/3



Virtual Machine: Snapshots, 3/3




VM evolution, from virtualization to cloud tech.



Structure


- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access





Cloud Computing with VMs

Several Hypervisors on a Clusters



Cloud Computing

- **A Cloud:** a type of parallel and distributed system consisting of a collection of interconnected and virtualized computers, VMs
 - They are dynamically provisioned and presented as one or more unified computing resources
 - based on service-level agreements, SLA, established through negotiation between the service provider and consumers
- Subset of grid computing where the allocated process are virtual computers
- **Implies**
 - An alternative way to have local servers or GRIDs
 - Outsourcing: HW and SW tools, they may be grid elements
 - Outsourcing: network, CPU, memory, HD, etc.
 - Definition of some service agreement, monthly rate, minimum network capability, kind of HW, mem space, minimum level of CPU/mem, etc.

Cloud vs GRID computing, HPCN

- **Cloud computing** is an evolution of GRID computing
- **Renting a GRID** service means:
 - Parallelize the algorithm, demand the execution, wait for the results, etc.
 - To do not know where the processes are executed. In most cases batch processing such as on Globus
- **Renting a VM/Computer** into a Cloud means:
 - Simple contracts, remote access to servers
 - Get access to virtual or physical resources at 100%
 - Privacy of the allocated processes in your VM/Host
 - Processes running on your preferred OS and do not have to be recompiled as in most Grid solutions.
 - Scalability in terms of number of CPU, Mem, network performance, computers, etc. OR you have to change your OWN architecture
 - Simple creation and reconfiguration of multitier solutions
 - Creating fault tolerant solutions
 - Balancing load of CPU, Memory, network, etc.

Cloud vs SuperComputers

- **Cloud Computing** is based on a set of Hosts for hosting VM or providing access to single/multiple Servers fully at disposal of the customers accessing to the cloud via remote access (see later).
 - They are MIMD computers
 - Hosts may contains hypervisors to host VM of diff. OSs.
 - Hosts may be single computers with Linux, Windows,...
- **Super Computers**, such as Blue Gene/P:
 - Processor PowerPC 450 with 4 cores, 850 Mhz
 - Each board: 32 CPU processors
 - Each Rack: 32 boards
 - [Jülich Research Centre](#) in Germania has a Blue Gene with 65536 processors → 167 teraFlops, the strongest computer in the world!



System Property	<u>Amazon</u> Elastic Compute Cloud (EC2)	<u>Google</u> App Engine	<u>Microsoft</u> Live Mesh	<u>Sun</u> Network.com (Sun Grid)	<u>GRIDS Lab</u> Aneka
Focus	Infrastructure	Platform	Infrastructure	Infrastructure	Software Platform for enterprise Clouds
Service Type	Compute, Storage (Amazon S3)	Web application	Storage	Compute	Compute
Virtualisation	OS Level running on a Xen hypervisor	Application container	OS level	Job management system (Sun Grid Engine)	Resource Manager and Scheduler
Dynamic Negotiation of QoS Parameters	None	None	None	None	SLA-based Resource Reservation on Aneka side.
User Access Interface	Amazon EC2 Command-line Tools	Web-based Administration Console	Web-based Live Desktop and any devices with Live Mesh installed	Job submission scripts, Sun Grid Web portal	Workbench, Web-based portal
Web APIs	Yes	Yes	Unknown	Yes	Yes
Value-added Service Providers	Yes	No	No	Yes	No
Programming Framework	Customizable Linux-based Amazon Machine Image (AMI)	Python	Not applicable	Solaris OS, Java, C, C++, FORTRAN	APIs supporting different programming models in C# and other .Net supported languages

Comparing Computing Services

- In the previous table, different solution for computing services in the network are compared.
- They are mainly: cloud computing, grid services, application services,
- The solutions taken have been selected as representative of their category
 - Please see the slides on GRID for more details and a wider comparison on the grid solutions
 - Please see the slides regarding the general distributed systems for multitier applications

Advantages of Cloud Computing

- **Grids are difficult to use and to maintain:**
 - GRID customers have too many different needs that make the creation of fully open grid very difficult.
- **Cloud computing**
 - HW/SW is hosting virtual computers that can be moved to other solutions with low costs
 - Lower costs since the HW/SW is seen as a service
 - No maintenance, centralized services such as back up, scaling, etc.
 - Lower costs for Small Business:
 - Reduction of costs since the admortment/depreciation can be performed by who is exploiting the computer
 - Scalability similar to grid:
 - Horizontal scaling
 - Parallelization as MIMD

Goals of Cloud Computing


- ***Scalability.***
 - scaling with workload demands so that performance and compliance with service levels remain on target
- ***Availability.***
 - users of Internet applications expect them to be up and running every minute of every day, i.e.: h24, 24/7
- ***Reliability***
 - physical system components rarely fail, but it happen. So that, they can be replaced without disruption.
 - Today, reliability means that applications do not fail and most importantly they do not lose data, and the service is not stopped.
- ***Security & privacy.***
 - Applications need to provide access only to authorized, authenticated users, that need to be able to trust that their data is secure.

Goals of Cloud Computing



- ***Flexibility and agility:***
 - Adapt rapidly to changes of business conditions by increasing the velocity at which applications are delivered into customer hands. E.g.: more CPU, more clock, more memory, more network cards, etc.
- ***Serviceability:***
 - In the past this meant using servers that could be repaired without, or with minimal, downtime.
 - Today it means that an application's underlying infrastructure components can be updated or even replaced without disrupting its characteristics including availability and security.
- ***Efficiency:***
 - differentiates the cloud computing. The process allocation and costs have to be very effective with respect to the investment.

Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access




- **Classificazione NIST**
 - Software as a Service (SaaS)
 - Platform as a Service (PaaS)
 - Infrastructure as a Service (IaaS)
- **Business Process as a Service (BPaaS)**
 - Aggiunto in seguito
- **Middleware as a Service (MaaS)**
- **Mobility as a Service (MaaS)**
- **Everything as a Service (XaaS)**



Modello Generale


NIST



CSP

Definizioni

- **Private Cloud.** abilitata per operare soltanto per un'organizzazione. può essere gestita dalla stessa organizzazione o da parte di terzi.
- **Community Cloud.** condivisa da più organizzazioni a supporto di una singola community che ha interessi e obiettivi comuni. Questa può essere gestita dalle stesse organizzazioni o da terzi in modalità on-premise e off-premise.
- **Public Cloud.** resa disponibile in maniera pubblica ed è di proprietà di un organizzazione che vi gestisce la vendita di servizi cloud.
- **Hybrid Cloud.** Infrastruttura composizione di due o più cloud (siano essi private, community o pubblici), rimangono entità separate, ma comunque accomunate da standard o tecnologie proprietarie che abilitano un certo livello di portabilità di dati e/o applicazioni di migrazione e/o bursting.




Infrastructure as a Service (IaaS)

- erogazione di servizi infrastrutturali relativi a capacità elaborativa, storage, rete e altri elementi di base assolutamente indipendenti da servizi applicativi di qualunque tipo.
- Si utilizza quindi l'infrastruttura messa a disposizione dal cloud service provider, CSP, per eseguire la propria applicazione,
 - pagamento in base al consumo dell'infrastruttura
 - lasciando sotto la responsabilità dell'utente la gestione del sistema operativo, dell'eventuale middleware e della parte di runtime, oltre che dell'applicazione stessa.
- Amazon EC2 è un esempio di servizio IaaS.

CSP

Host
Host
VM



Infrastructure as a Service (IaaS)

- erogazione di servizi infrastrutturali relativi a:
 - CPU come clock, storage as HD, rete, mem
 - altri elementi di base assolutamente indipendenti da servizi applicativi di qualunque tipo.
- infrastruttura messa a disposizione dal provider per eseguire la propria applicazione:
 - Affitto di Server e/o macchine virtuali
 - pagamento in base al consumo dell'infrastruttura
 - lascia sotto la responsabilità dell'utente la gestione del sistema operativo, dell'eventuale middleware e della parte di runtime, oltre che dell'applicazione stessa.
- Amazon EC2 è un esempio di servizio IaaS.



Infrastructure as a Service (IaaS)


- Una serie di Host, VM e NAS/SAN
- Un gestore di macchine virtuali
- Un sistema per il deploy basato su templates
- Un sistema di monitoring and alarm su risorse fisiche e parametri di hosts
- Un sistema per l'accounting e il deploy automatico
- Alcuni servizi per l'accesso diretto a Host e VM

Area IaaS	AWS	Microsoft Azure	Google Cloud	Oracle Cloud (OCI)	IBM Cloud
Compute / VM	Amazon EC2	Azure Virtual Machines	Compute Engine	OCI Compute – Virtual Machines	Virtual Servers for VPC
Autoscaling	EC2 Auto Scaling	Virtual Machine Scale Sets	Managed Instance Groups / Autoscaling	Instance Pools / Autoscaling	Instance Groups / Auto Scale for VPC
Block Storage	Amazon EBS	Azure Managed Disks	Hyperdisk / Persistent Disk	OCI Block Volumes	Block Storage for VPC
Object Storage	Amazon S3	Azure Blob Storage	Cloud Storage	OCI Object Storage	IBM Cloud Object Storage
File Storage	Amazon EFS / FSx	Azure Files	Filestore	OCI File Storage	File Storage for VPC
Rete virtuale	Amazon VPC	Azure Virtual Network (VNet)	Virtual Private Cloud (VPC)	Virtual Cloud Network (VCN)	IBM Cloud VPC
Subnet / IP privati	VPC Subnets	VNet Subnets	VPC Subnets	VCN Subnets	VPC Subnets
Firewall infrastrutturale	Security Groups / Network ACL	Network Security Groups	VPC Firewall Rules / Firewall Policies	Network Security Groups / Security Lists	Security Groups / Network ACL
Load Balancing	Elastic Load Balancing: ALB, NLB, GWLB	Azure Load Balancer, Application Gateway	Cloud Load Balancing	OCI Load Balancer / Network Load Balancer	Load Balancer for VPC
VPN	AWS Site-to-Site VPN / Client VPN	Azure VPN Gateway	Cloud VPN	OCI Site-to-Site VPN	VPC VPN
Connessione privata on-prem	AWS Direct Connect	Azure ExpressRoute	Cloud Interconnect	OCI FastConnect	IBM Cloud Direct Link
Bare Metal	EC2 Bare Metal Instances	Azure BareMetal Infrastructure*	Bare Metal Solution*	OCI Bare Metal Instances	Bare Metal Servers for VPC
GPU / acceleratori	EC2 GPU Instances	N/NC/ND-series VM	GPU VM / accelerator-optimized instances	GPU VM e Bare Metal	GPU / AI Optimized instances
Storage locale NVMe	Instance Store	Local / Temporary SSD a seconda VM	Local SSD	Dense I/O / local NVMe	Instance Storage / NVMe su profili supportati
Modello di pagamento base	Pay-as-you-go + Savings Plans / Reserved / Spot	Pay-as-you-go + Reservations / Savings Plan / Spot	Pay-as-you-go + Committed Use Discounts / Spot	Pay-as-you-go + commitment options	Hourly/monthly + reservations
Infrastructure as Code	CloudFormation / Terraform	ARM / Bicep / Terraform	Infrastructure Manager / Terraform	Resource Manager / Terraform	Schematics / Terraform

Platform as a Service (PaaS)


- **Per i developers:**
 - erogazione di servizi applicativi di base come sistemi operativi, middleware, linguaggi, tecnologie di base dati e l'ambiente runtime necessari per eseguire l'applicazione,
 - L'applicazione rimane l'unica cosa sotto la responsabilità dell'utente, oltre alla definizione del modello (e.g., numero e dimensione dei server, datacenter, caratteristiche del networking) da utilizzare per l'esecuzione dell'applicazione.
- Google AppEngine è un esempio di Platform as a Service.
- A livello PaaS viene anche collocato l'insieme dei servizi MaaS, Middleware as a Service.

Platform as a Service (PaaS)



PaaS

- **Middleware per accesso ai servizi di base**
 - Middleware as a Service, MaaS
 - API per accesso a questi servizi
- **Sistema di sviluppo** per applicazioni che possono usare le API
- **Servizi e risorse base**
- **Business Process Modeling:**
 - permette definire come allocare e configurare le macchine virtuali ed i servizi sul cloud.
 - A livello SaaS: Crea procedure, script, etc. per il deploy e la configurazione di VM, Applicazioni, servizi, etc. anche riconfigurazioni e trasformazioni
- **Processo di lavoro**



Area PaaS	AWS	Microsoft Azure	Google Cloud	Oracle Cloud (OCI)	IBM Cloud
Application Platform	Elastic Beanstalk, App Runner	Azure App Service	App Engine, Cloud Run	OCI Application Container / servizi applicativi OCI	IBM Cloud Code Engine
Container PaaS / Serverless containers	App Runner, ECS Fargate	Azure Container Apps	Cloud Run	OCI Container Instances	Code Engine
Functions / FaaS	AWS Lambda	Azure Functions	Cloud Run functions	OCI Functions	Functions tramite Code Engine
Managed Kubernetes	Amazon EKS	Azure Kubernetes Service (AKS)	Google Kubernetes Engine (GKE)	Oracle Kubernetes Engine (OKE)	IBM Cloud Kubernetes Service
Database SQL gestito	Amazon RDS / Aurora	Azure SQL Database, Azure Database for PostgreSQL/MySQL	Cloud SQL, AlloyDB	Autonomous AI Database, MySQL HeatWave	Db2 on Cloud, Databases for PostgreSQL
Database NoSQL gestito	DynamoDB	Cosmos DB	Firestore / Bigtable	NoSQL Database	Cloudant
In-memory / Cache	ElastiCache	Azure Managed Redis	Memorystore	OCI Cache	Databases for Redis
Messaging / Queue	SQS, SNS	Service Bus, Event Grid	Pub/Sub	Queue, Streaming	Event Streams, Messages for RabbitMQ
Event-driven platform	EventBridge	Event Grid	Eventarc	Events	Event Streams
API Management	Amazon API Gateway	Azure API Management	Apigee, API Gateway	OCI API Gateway	IBM API Connect
Integration PaaS	Step Functions, AppFlow	Logic Apps	Workflows, Application Integration	Oracle Integration	App Connect
Data warehouse / analytics PaaS	Redshift	Synapse Analytics / Fabric	BigQuery	Autonomous Data Warehouse	Db2 Warehouse
ETL / Data Integration	AWS Glue	Azure Data Factory	Dataflow / Data Fusion	OCI Data Integration	DataStage
Streaming analytics	Kinesis	Event Hubs / Stream Analytics	Dataflow + Pub/Sub	OCI Streaming	Event Streams
AI/ML Platform	SageMaker	Azure Machine Learning / Azure AI Foundry	Vertex AI	OCI Data Science / Generative AI	watsonx
DevOps / CI-CD	CodePipeline, CodeBuild	Azure DevOps / GitHub Actions	Cloud Build, Cloud Deploy	OCI DevOps	Continuous Delivery / Tekton
Monitoring applicativo	CloudWatch / X-Ray	Azure Monitor / Application Insights	Cloud Monitoring / Trace	OCI Monitoring / APM	Instana / IBM Cloud Monitoring

Software as a Service (SaaS)

- **Per gli utenti finali**
 - erogazione di servizi applicativi di qualunque tipo, accessibili indipendentemente dalla collocazione e dal tipo di device utilizzato.
 - Non è eseguita un'applicazione proprietaria del cliente, ma il cliente stesso paga il diritto (mediante licenza o canone di affitto) di utilizzo di un'applicazione messa a disposizione dal provider, senza preoccuparsi di come essa venga realizzata e gestita nel cloud.
 - L'unica preoccupazione del cliente in questo caso, oltre ovviamente alla scelta della corretta applicazione che soddisfi le sue necessità, è gestire il numero di licenze richieste in funzione del numero di utenti.
- Salesforce.com Customer Relationship Management (CRM) è un esempio di soluzione in cui il software è venduto in modalità as a service.

Area SaaS	Microsoft	Google	Salesforce	Oracle	SAP	ServiceNow
Produttività / Office	Microsoft 365: Word, Excel, PowerPoint, Outlook	Google Workspace: Docs, Sheets, Slides, Gmail	—	—	—	—
Email aziendale	Exchange Online / Outlook	Gmail	—	—	—	—
File sharing / Documenti	OneDrive, SharePoint	Google Drive	Files / Salesforce Platform	Oracle Content Management	SAP Document Management	Document Management
Collaboration / Meeting	Microsoft Teams	Google Meet / Chat	Slack	Oracle Collaboration / Fusion tools	SAP Work Zone	Employee Center / collaboration workflow
CRM Vendite	Dynamics 365 Sales	—	Sales Cloud	Oracle Fusion Cloud Sales	SAP Sales Cloud	—
Customer Service / Contact Center	Dynamics 365 Customer Service	—	Service Cloud	Oracle Fusion Service	SAP Service Cloud	Customer Service Management
Marketing	Dynamics 365 Customer Insights	—	Marketing Cloud	Oracle Fusion Marketing	SAP Emarsys / Marketing solutions	—
ERP	Dynamics 365 Finance / Business Central	—	—	Oracle Fusion Cloud ERP	SAP S/4HANA Cloud	—
Finance / Accounting	Dynamics 365 Finance	—	—	Oracle Fusion Financials	SAP S/4HANA Cloud Finance	Finance & Supply Chain Workflows
Supply Chain	Dynamics 365 Supply Chain Management	—	—	Oracle Fusion Cloud SCM	SAP S/4HANA Cloud / SAP SCM solutions	Supply Chain Workflows
HR / HCM	Dynamics 365 Human Resources / Microsoft HR integrations	—	—	Oracle Fusion Cloud HCM	SAP SuccessFactors	HR Service Delivery
Procurement	Dynamics 365 / Power Platform integrations	—	—	Oracle Fusion Procurement	SAP Ariba	Procurement Service Management
Analytics / BI	Power BI	Looker / Looker Studio	Tableau	Oracle Analytics Cloud	SAP Analytics Cloud	Performance Analytics
IT Service Management	—	—	—	—	—	ServiceNow ITSM
IT Operations Management	—	—	—	—	—	ServiceNow ITOM
Security Operations	Microsoft Defender suite	Google Workspace Security	Salesforce Security	Oracle Fusion security functions	SAP security solutions	ServiceNow Security Operations
Workflow / Low-code SaaS	Power Apps / Power Automate	AppSheet	Salesforce Platform / Flow	Oracle Fusion / Visual Builder	SAP Build	ServiceNow AI Platform
AI integrata	Microsoft 365 Copilot / Dynamics 365 Copilot	Gemini for Workspace	Agentforce / Einstein	Oracle AI for Fusion Applications	SAP Joule	ServiceNow AI / AI Agents




Confronto IaaS vs PaaS vs SaaS

Responsabilità	IaaS	PaaS	SaaS
Applicazioni	Cliente	Cliente	Provider
Dati	Cliente	Cliente	Cliente/provider secondo il servizio
Runtime	Cliente	Provider	Provider
Middleware	Cliente	Provider	Provider
Sistema operativo	Cliente	Provider	Provider
Virtualizzazione	Provider	Provider	Provider
Server	Provider	Provider	Provider
Storage fisico	Provider	Provider	Provider
Networking fisico	Provider	Provider	Provider
Gestione tecnica richiesta al cliente	Alta	Media	Bassa
Personalizzazione	Alta	Medio-alta	limitata
Esempio	EC2, Azure VM	App Service, Cloud Run	Microsoft 365, Salesforce, SAP S/4HANA Cloud


Business Process as a Service (BPaaS):

- erogazione di servizi non esclusivamente riferiti ad ambiti applicativi ma direttamente alle funzionalità di business o di processo, potenzialmente trasversali rispetto alle piattaforme applicative.
- un processo di business mappato interamente nel cloud (composto da servizi, applicazioni web, applicazioni legacy, servizi di integrazione, etc.).
- Il processo di business è un pattern di servizi ed include problemi di sicurezza, costi, scalabilità connessione fra local e cloud bidirezionale, il cloud può essere un burst per l'azienda, e può sgravare i costi nel momento del bisogno.



Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access



High Availability: how to reach it

- The high availability has to be guaranteed only by the integration of features of:
 - ***High Reliability***
 - ***High Serviceability***
 - ***Fault tolerance***
 - ***Migration of VM to different HW***
 - ***Disaster recovering***

High Availability, how to achieve it

- May be achieved by redundancy of :
 - Servers and Services in hot spare or balancing
 - VM based architectures
 - HW/SW: power supply, network connections, etc.

CPU / STORAGE

① – **Load Balancing/Balancer** according to server traffic/requests

Rnd Ext

- Server Clustering, multitier solution

① – **Hot Spare:**


Rnd Int

- Server: cloned server to be used when the main is not functioning: heartbeat to detect the server availability and thus failover. (heartbeat signal to communicate the correct running of a process/CPU)
- HD: Raid based on SAN/NAS, LUN, in host


– **Mixed: balancing and hot spare**


High Availability: Hot spare

- **Fully cloned** servers to be alternatively used when the running one fails
- **Internal Network**
 - *Heartbeat* to detect the server availability and thus failover.
 - *[To keep servers aligned on context and data]*
- **Shared data** storage is a simplification and optional.
 - A different solution may be to have a cloned storage to keep aligned among the two servers



One master and 4 hotspares in heartbeat






Network and Virtual Networks

- The same VM with access to 2 different network via real network adapters
- A virtual network

- Partitioned Application Data
-
- The diagram shows three MySQL database instances, each labeled 'DATABASE' and 'MySQL' with the 'opensolaris' logo. The first two instances are active, with red circles around their bottom sections containing the number '2'. The third instance is crossed out with a large red 'X'. Dashed blue arrows originate from a stack of three cylinders labeled 'Partitioned Application Data' and point to the bottom sections of the first two active MySQL instances. A caption at the bottom right reads 'Application Data Distributed Across Horizontally Scaled Database Instances'.
- Application Data Distributed Across Horizontally Scaled Database Instances


High Availability: Load Balancing

- Service distributed/cloned on more servers
 - According to different policies
 - Round robin
 - Network traffic
- Single transfer rate capability on the front-end of the load balancer
- Sensing the availability of the server on balancer +
 - heartbeat solutions to understand if the servers are alive or not
- Common SAN
- That is a Cluster of servers on the cloud !




High Availability

- **High Availability**, available 99.999 % (called “Five Nines”) percent of the time.
- **Five Nines** is the term for saying a service or system will be up almost 100 percent of the time.
- In case of failure:
 - the path changes to guarantee the service




Balancing multi-tier, 3-tier



in Arc
vite

Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access



vSphere infrastructure of VMware

V65

- **High level features:**


- HA: high availability
- DRS: Distributed Resource Scheduling
- Creating Fault Tolerance architectures
- DPM: datacenter power management, based on VMotion
- Converting VM into VM for infrastructure, from physical to VM
- vApp: are Virtual Application Services
- Cloning and Moving VM
- Making Templates for VM
- Backup VM and thus virtual servers

FT

- **Low level features:**


- VMotion: VM moving among Hosts
- Dynamic increment of: CPU ck, MEM, net...

Summary of VMware vSphere 4.0



VM Cloning

Create a new VM from an existing one (on the same or another host/datastore)



- 1 Select the source VM (VM1).
- 2 Choose clone type (full or linked).
- 3 Copy VM configuration and virtual disks.
- 4 Assign a new name and settings (CPU, RAM, network, etc.).
- 5 Start the cloned VM (VM2) and update unique identifiers (e.g. hostname, IP, machine ID).

Summary


- **Cloning:** creates a new VM from an existing one (copy of configuration and disks).
- **Migration:** moves an existing VM to another host, either its storage, its execution, or both.

VM Migration (between two hosts)

Move a VM from one host to another.

Storage Migration (only storage)


Move the VM's virtual disks to a datastore on the other host. The VM remains running on the original host.



- 1 Select the VM (running or powered off).
- 2 Move/copy the virtual disks to the target datastore (e.g. Storage vMotion).
- 3 The VM continues to run on Host A, now using the datastore on Host B.

Execution Migration (only execution)

Move the running VM from one host to another, keeping the same storage. The VM keeps running during the migration.



- 1 Initiate live migration of the running VM (e.g. vMotion).
- 2 The VM's execution (CPU and memory state) is moved to Host B.
- 3 Storage remains on the same datastore (no disk movement).

Migration types



Storage migration
Move the virtual disks
(independent of execution).



Execution migration
Move the running VM
(CPU and memory state).

Combined migration (storage + execution)





Possible to move both storage and execution (e.g. enhanced vMotion), useful when datastores are not shared.

- **Results of cloning**
 - Two VMs with different MAC addresses on NICs
- **Results of Migration**
 - One VM in different location
 - The same MAC addresses on the NICs


Simple VMotion of VMware vSphere (migration)

- **you need:**
 - VM Off
 - Compatibility among host network
 - virtual network present
- Steps performed by supervisor:
 1. Copy HD image(s)
 2. Completed ? clean the old HD copy



Simple VMotion of VMware vSphere

- **you need:**
 - VM (on or Off)
 - Compatibility among host CPU and network
 - virtual network present
- **Steps performed by supervisor:**
 1. Copy HD image(s)
 2. aligning OS and CPU status (RUN)
 3. off-the old and then on-new (one clock)
 4. Clean the old HD copy



VMware DRS, Distrib. Res. Scheduling


- DRS is used to balance the workload among Hosts

- Moving VMs is a tools for balancing the workload on Hosts

..

SAN

DRS and vMotion in a Single vSphere Cluster




VM Configuration	
VM Hardware Version	Compatible (e.g. vHW 19)
VM Storage	On shared SAN datastore
Network	Connected to port group (vDS or vSS)
vMotion Compatibility	EVC enabled (same CPU compatibility across hosts)
DRS Settings	VM is in DRS-enabled cluster (with appropriate rules/affinity)



Host Configuration	
Hypervisor	VMware ESXi (same version)
CPU Compatibility	Enable EVC
Networking	vMotion network configured (dedicated, e.g. 10/25GbE)
Storage	Connected to shared SAN (Fibre Channel or iSCSI)
Cluster	All hosts added to DRS-enabled cluster
DRS	Set to Manual / Partially Automated / Fully Automated

SAN Configuration	
Type	Enterprise SAN (Fibre Channel or iSCSI)
Datastores	VMFS datastores accessible by all hosts in the cluster
Multipathing	Enabled (e.g. via VMware PSP/MPIO)
LUN Zoning	Present same LUNs to all hosts
Capacity	Plan for growth and performance
High Availability	Redundant controllers, paths, and links

vSphere High Availability (Full HA) - Architecture



Reyle




- How Full HA Works**
- 1 HA agents on each host monitor the health of other hosts.
 - 2 If a host fails, the cluster detects the failure (via heartbeats).
 - 3 HA automatically restarts the affected VMs on surviving hosts in the cluster.
 - 4 VMs are restarted from the shared SAN datastores.
 - 5 No manual intervention is required (for VM restart).

Reyle

VMware DPM, Power Management


- DPM consolidates workloads to reduce power consumption
 - Cuts power and cooling costs
 - Automates management of energy efficiency
 - optimizing host resources
- Based on VMotion

75%




VMware Fault Tolerance, FT

- ❑ Single identical VMs running in lockstep on separate hosts
- ❑ Zero downtime, zero data loss failover for all virtual machines in case of hardware failures
- ❑ Single common mechanism for all applications and Operating systems
- ❑ Need to have a storage shared by the same VM
- ❑ The VM itself can be stored in the SAN




VM in Fault Tolerance (FT) – Architecture




HA: High Availability of vSphere, FT

- **When a host fails**, the running VM on the host may be turned ON on another hosts
 - Just the time to turn on again the host
- **HOT Spare** solution:
 - It is also possible to keep aligned 2 distinct hosts to make a faster switch OFF→ON of the VM on the faulty host
 - implies to have duplicated resources: Host, CPU etc.



Incremental Backup of a VM



Key Takeaways

- 1 First backup = full copy of the VM
- 2 Next backups = only changed data blocks
- 3 Faster backups and lower storage usage
- 4 Restore = full backup + all required incrementals


Restore chain

To restore a VM, the full backup and all required incremental backups are combined in order.



DRS – vMotion Approach

Automatic VM load balancing and seamless migration using shared SAN storage




VM Configuration	
VM Hardware Version	Compatible with vSphere (eg. vHW 19)
VM Storage	On shared SAN datastore
Network	Connected to distributed port group (vDS)
vMotion Compatibility	Same CPU compatibility mode (EVC) across hosts
DRS Settings	VM is in a DRS-enabled cluster (with appropriate rules/affinity)

Host Configuration	
Hypervisor	VMware ESXi (same version across cluster)
CPU Compatibility	Enable EVC (Enhanced vMotion Compatibility)
Networking	vMotion VMkernel (dedicated, 10/25GbE)
Management	Connected to vCenter Server
Storage	Connected to shared SAN (Fibre Channel or iSCSI)
Cluster	Add hosts to DRS-enabled cluster
DRS	Set automation level (Manual / Partially Automated / Fully Automated)

SAN Configuration	
Type	Enterprise SAN (Fibre Channel or iSCSI)
Datastores	VMFS datastores accessible by all hosts in the cluster(s)
Replication	Configured between Site A and Site B (Synchronous or Asynchronous)
Multipathing	Enable redundant paths (e.g. via VMware PSP/MPIO)
LUN Zoning	Present same LUNs to all hosts in cluster
Capacity	Plan for growth and DR requirements
High Availability	Redundant controllers, paths, and links


VM Converter, the migration

- **Conversion possibilities, migration possibilities**
- **P2V, from Physical → Virtual machine**
 - Reusing legacy servers into stronger and new HW machines
 - From ISO CD of an OS → VM
- **V2V, from Virtual → to Virtual**
 - Import/export a VM from/to different standards
 - From VM Workstation → Infrastructure VM
 - From Infrastructure VM → template for VM with some parameters
 - From Infrastructure VM → VM Workstation
 - From Infrastructure VM → Infrastructure VM changing parameters
 -



Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access




Security on the cloud

- Protecting VMs from external access
- Protecting VMs each other in the cloud
- Technologies:
 - Accessing to other VM via dedicated virtual networks, using Virtual Networking, Virtual Switch
 - Avoiding shared disk, at least using authenticated connections
 - Using Firewall
 - Communicating with other VMs via protected connections: protected WS, HTTPs, SSL, SFTP, etc.


Virtual Networking

- Per isolare meglio dei guest OS da altri che non sono nello stesso trust domain
- isolare i virtual switch dei singoli trust domain.
- Solo i guest che condividono lo stesso domain hanno schede di rete virtuali sullo stesso virtual switch.
- virtual switch su porte logiche del sistema host che non hanno indirizzo ip configurato.

Figure 1 - Isolating virtual machines to separate network trust domains




VMware Vsphere VMsafe



- Creates a new, stronger layer of defense – fundamentally changes protection profile for VMs running on VMware Infrastructure
- Protect the VM by inspection of virtual components (CPU, Memory, Network and Storage)
- Complete integration and awareness of VMotion, Storage VMotion, HA, etc.
- Provides an unprecedented level of security for the application and the data inside the VM


Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access




vCenter vSphere

- The so called Inventory
- Datacenter
 - Cluster00
 - Host....
 - Host....
 - Application:
 - VM1
 - VM2
 - VM3
 - VM4
 - Cluster01
 - H2..
 - ...





Cluster vs Resource Pool


- The so called Inventory
- Datacenter
 - Cluster00
 - Host....
 - Host....
 - Application:
 - VM1
 - VM2
 - VM3
 - VM4
 - Cluster01
 - H2..
 - ...




Some Snapshots of vSphere



Storage View



VM	Multipathing Status	Space Used	Snapshot Space	Disks
Train01_WindowsXP_01	Partial/No Redundancy	1.94 GB	0.00 B	1
VM2-XP	Partial/No Redundancy	3.38 GB	0.00 B	1
VM4-RHEL5	Partial/No Redundancy	8.10 MB	0.00 B	1
VM2-W2K3	Partial/No Redundancy	4.50 GB	256.07 MB	1
W2K3 Template	Partial/No Redundancy	4.00 GB	0.00 B	1
XP Template	Partial/No Redundancy	3.00 GB	0.00 B	1



Overview

Map Relationships:

Custom Map ▾


Host Options

- ☒ Host to VM
- ☒ Host to Network
- ☐ Host to Datastore

VM Options

- ☐ Fault Tolerance relationships
- ☐ VM to Network
- ☐ VM to Datastore
- ☐ Show only powered on VMs

[Apply Relationships](#)



Monitoring the Solution

- Monitoring and assessing performance at level of:
 - Datacenter, Cluster, Host
 - Virtual Machine from outside
 - Virtual Machine inside:
 - this has to be performed by using tools inside the VM operating system
 - Windows:
 - System monitoring hosts, detailed performances
 - Linux:
 - Top or other tools

Main Monitoring methods

- Agent and Agentless
 - Problems to maintain the certification and assistance level
- Protocols:
 - SSH
 - API of hosts hypervisors
 - Storage protocols
 - Network monitoring protocols
- Call from applications and services to the Collector:
 - Call REST
 - Call to WebService
 - Write on database

Performance Analysis, Cluster/DC

'eolo - Client per terminal server'

EOLO - vSphere Client

File Edit View Inventory Administration Plug-ins Help

Home Inventory Hosts and Clusters Search Inventory

prometeo1.regione.toscana.it VMware ESX, 4.0.0, 208167

Summary Virtual Machines Performance Configuration Tasks & Events Alarms Permissions Maps Storage Views Hardware Status


Name, State or Guest OS contains: Clear

Name	Status	VMware Tools Status	VM Version	State	Provisioned Sp...	Used Space	Host CPU - MHz	Host Mem - MB
crio8	✓	Normal	OK	7	Powered On	75,00 GB	75,00 GB	198
saturnotest	✓	Normal	OK	7	Powered On	35,15 GB	35,15 GB	0
crio2	✓	Normal	OK	7	Powered On	79,00 GB	79,00 GB	1127
carontex	✓	Normal	OK	7	Powered On	12,00 GB	12,00 GB	265
azan	✓	Normal	OK	7	Powered On	15,50 GB	15,50 GB	88
crio4test	✓	Normal	OK	7	Powered On	48,00 GB	48,00 GB	884
elio6test	✓	Normal	OK	7	Powered On	43,00 GB	43,00 GB	176
elio6	✓	Normal	OK	7	Powered On	43,00 GB	43,00 GB	154
crio7	✓	Normal	OK	7	Powered On	48,00 GB	48,00 GB	66
ganimede	✓	Normal	Out of date	4	Powered On	99,66 GB	99,66 GB	22
giapeto	✓	Normal	Out of date	4	Powered On	53,02 GB	53,02 GB	44


Recent Tasks

DISIT lab, Big Data Arc. 2026

Datacenter space consumed again time



Performance Analysis, Single VM



Performance Chart Legend





Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
	naa.600300570...	Write rate	Average	KBps	1001	1727	626	1058,693
	naa.600300570...	Read rate	Average	KBps	0	820	0	18,045
	192.168.1.10	Usage	Average	KBps	3281	6992	1564	2594,006
	192.168.1.10	Write rate	Average	KBps	1110	6019	716	1218,615
	192.168.1.10	Read rate	Average	KBps	2170	2279	446	1374,872
	naa.600300570...	Read rate	Average	KBps	2170	2279	446	1356,693
	naa.600300570...	Write rate	Average	KBps	1001	1727	626	1058,693

Performance Analysis, Single VM


Getting Started Summary Resource Allocation Performance Tasks & Events Alarms Console Permissions Maps Storage Views

Overview Advanced

Realtime Summary for opnet.axmedis.org




Performance Analysis, Single VM




Performance Analysis of VM on the Host

CPU Usage (Top 10)




testinterop.eclap.eu-running test.eclap.eu-running test.eclap.eu-old
bpnetmardaniele-2-running

Memory Usage (Top 10)




testinterop.eclap.eu-running test.eclap.eu-running bpnetmardaniele-2-running
test.eclap.eu-old

Network (Mbps) (Top 10)



test.eclap.eu-running bpnetmardaniele-2-running testinterop.eclap.eu-running
test.eclap.eu-old

Disk (KBps) (Top 10)



testinterop.eclap.eu-running test.eclap.eu-running bpnetmardaniele-2-running
test.eclap.eu-old

Problematiche del Monitoraggio

- Parametri da monitorare: IaaS, PaaS, e SaaS:
 - Condizioni e parametri di stato degli HOST
 - Parametri di consumo degli HOST
 - Spazio storage, snapshots, etc.
 - Condizioni e parametri di stato delle VM e del sistema operativo
 - Parametri di consumo delle VM
 - Spazio disco, memoria, connessione
 - Condizioni e parametri di stato delle applicazioni e dei servizi dentro le VM
 - Parametri di consumo delle applicazioni e dei servizi dentro le VM:
 - E.g.: Numero di fatture, numero di utenti


Monitorare perche' → HA, business consum.

- Actions:
 - Allarmi gialli/rossi (75%/90%) se si superano soglie definite
 - Invio di: email, sms, etc.
 - Attivazione di riconfigurazioni, cambi di configurazione, estensioni di risorse
 - Attivazione di moving
 - Shutdown di emergenza
 - Network off/on
 - Cambio di billing in base alla SLA
 - Etc. etc.

Aspetto	IaaS – Infrastructure as a Service	PaaS – Platform as a Service	SaaS – Software as a Service
Oggetto principale del monitoraggio	VM, host, rete, storage, sistema operativo	Applicazioni, runtime, container, middleware, database gestiti	Applicazione e servizio percepito dall'utente
CPU / RAM	✓ Monitoraggio dettagliato di VM e istanze	✓ Principalmente a livello applicativo/container	⚠ Normalmente non visibile al cliente
Storage	✓ Capacità, IOPS, latenza, throughput	✓ Metriche del servizio storage/database	⚠ Gestito principalmente dal provider
Rete	✓ Traffico, latenza, packet loss, interfacce	✓ Connessioni e traffico applicativo	⚠ Principalmente disponibilità del servizio
Sistema operativo	✓ Responsabilità importante del cliente	⚠ Generalmente gestito dal provider	✗ Non gestito dal cliente
Middleware / Runtime	✓ / ⚠ Dipende dall'architettura	✓ Elemento centrale del monitoraggio	✗ Normalmente nascosto
Applicazione	✓ Tramite strumenti APM/log	✓ Monitoraggio centrale	✓ Disponibilità, errori e funzionalità
Log	✓ OS, VM, rete e applicazioni	✓ Applicazione, runtime e servizi	✓ Audit log e log applicativi resi disponibili
Metriche tipiche	CPU, RAM, disk usage, IOPS, bandwidth, uptime	response time, throughput, error rate, CPU/container, DB latency	availability, response time, error rate, login failure
APM – Application Performance Monitoring	Opzionale	Molto importante	Principalmente gestito dal provider
Disponibilità / Uptime	Monitoraggio delle singole risorse	Monitoraggio piattaforma + applicazione	Parametro fondamentale
SLA	Principalmente infrastrutturale	Piattaforma e servizi	Disponibilità del servizio finale
Responsabilità cliente	Alta	Media	Bassa
Responsabilità provider	Infrastruttura fisica e cloud	Infrastruttura + piattaforma/runtime	Quasi tutto lo stack
Esempi di strumenti	AWS CloudWatch, Azure Monitor, Google Cloud Monitoring, Zabbix, Prometheus	Azure Monitor / Application Insights, Google Cloud Operations, Datadog, New Relic, Prometheus	Dashboard del provider, audit/logging SaaS, synthetic monitoring esterno


Tool	Tipologia	Ideale per	Metriche infrastruttura	Log	APM / Tracing	Kubernetes	Multi-cloud	Complessità	Costo indicativo
Prometheus + Grafana	Open source	Kubernetes, container, microservizi	✓ Ottimo	⚠ Con Loki/strumenti esterni	⚠ Con OpenTelemetry/Jaeger	✓ Eccellente	✓	Medio-alta	Basso, ma richiede gestione
Zabbix	Open source / enterprise	Server, VM, rete, datacenter	✓ Eccellente	✓	⚠ Limitato rispetto agli APM puri	✓	✓	Media	Basso
Datadog	SaaS observability	Cloud, microservizi, applicazioni	✓	✓	✓ Eccellente	✓	✓ Eccellente	Bassa-media	Medio-alto
Dynatrace	Enterprise observability	Grandi infrastrutture e applicazioni complesse	✓	✓	✓ Eccellente	✓	✓	Media	Alto
New Relic	SaaS observability /APM	Applicazioni e DevOps	✓	✓	✓ Eccellente	✓	✓	Bassa-media	Medio-alto
Azure Monitor	Cloud-native	Ambienti Microsoft Azure	✓	✓	✓ Application Insights	✓ AKS	⚠ Sì, ma Azure-centric	Bassa in Azure	A consumo
AWS CloudWatch	Cloud-native	Ambienti AWS	✓	✓	✓ tracing/integration	✓ EKS	⚠ AWS-centric	Bassa in AWS	A consumo
Google Cloud Monitoring	Cloud-native	Ambienti Google Cloud	✓	✓	✓	✓ GKE	⚠ GCP-centric	Bassa in GCP	A consumo

Architettura ICARO



Structure

- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access



Cloud SW

Vendor / piattaforma	Focus principale	Virtualizzazione VM	Kubernete s / container	Private cloud	Hybrid / Multi- cloud	Automa zione / IaC	Punti distintivi
Broadcom / VMware Cloud Foundation	Private cloud enterprise	✓ Molto forte	✓	✓ Eccellente	✓	✓	Stack integrato per compute, storage, networking e Kubernetes; VCF 9.1 è orientato anche a workload AI. Broadcom
Nutanix Cloud Platform	HCI e hybrid multicloud	✓ Molto forte	✓	✓	✓ Eccellente	✓	Integra hypervisor AHV, storage, networking, DR e gestione cloud in una piattaforma unificata. Nutanix
Red Hat OpenShift	PaaS / Kubernetes enterprise	⚠ Secondario	✓ Eccellente	✓	✓ Eccellente	✓	Molto orientato ad applicazioni cloud-native, container, DevOps e Kubernetes enterprise.
Microsoft	Hybrid cloud + Azure	✓	✓	✓	✓ Eccellente	✓	Forte integrazione tra infrastruttura locale e Azure, particolarmente adatta a ecosistemi Microsoft.
OpenStack (OS)	Cloud IaaS open source	✓ Eccellente	✓	✓ Eccellente	✓	✓	Gestisce pool di compute, storage e networking tramite API; molto flessibile ma richiede competenze operative elevate. openstack.org
Canonical	OpenStack / Kubernetes / Linux cloud	✓	✓	✓	✓	✓	Offre distribuzioni e supporto enterprise per OpenStack e infrastrutture cloud open source. Canonical
SUSE / Rancher (OS)	Kubernetes management	⚠	✓ Eccellente	✓	✓ Eccellente	✓	Particolarmente adatto alla gestione centralizzata di cluster Kubernetes distribuiti
HashiCorp	Infrastructure as Code e cloud automation	✗	⚠	⚠	✓ Eccellente	✓ Eccellente	Terraform consente di gestire infrastrutture cloud e on-premise con un workflow dichiarativo comune e oltre 1.000 provider. Sentinel HashiCorp Developer


Structure


- Elements of cloud Computing
- Virtualization
- Storage
- Hosts vs SAN
- VM SnapShots
- Cloud Computing
- IAAS, PAAS, SAAS
- High Availability
- vSphere Infrastructure
- Security on the Cloud
- An example and monitoring
- Cloud Vendors
- remote Access



How to **WORK** with VMs in the Cloud

- **KVM solutions**
 - Local access via local KVM
 - Local server access via HTTP
- **Windowing Terminal**
 - MS Windows Remote Desktop
 - X Terminal to linux
- **Remote Solutions:**
 - VNC, Radmin, G Remote Desktop, etc.
 - VNC: Also possible via HTTP port
- **Telnet, char based consol, SSH**
 - VT100 terminal for example







- kvm




RJ45 KVM




Windows Remote Desktop



Remote machine
Windows 2003





Local machine
Windows Vista



VNC remote connection

Remote machine
Windows XP



Local machine
Windows Vista

