



UNIVERSITY OF TARTU

INSTITUTE OF COMPUTER SCIENCE

Cloud Computing – Lecture 12

Hybrid Clouds

27 April 2020

Shivananda Poojara

Pelle Jakovits

Satish Srirama

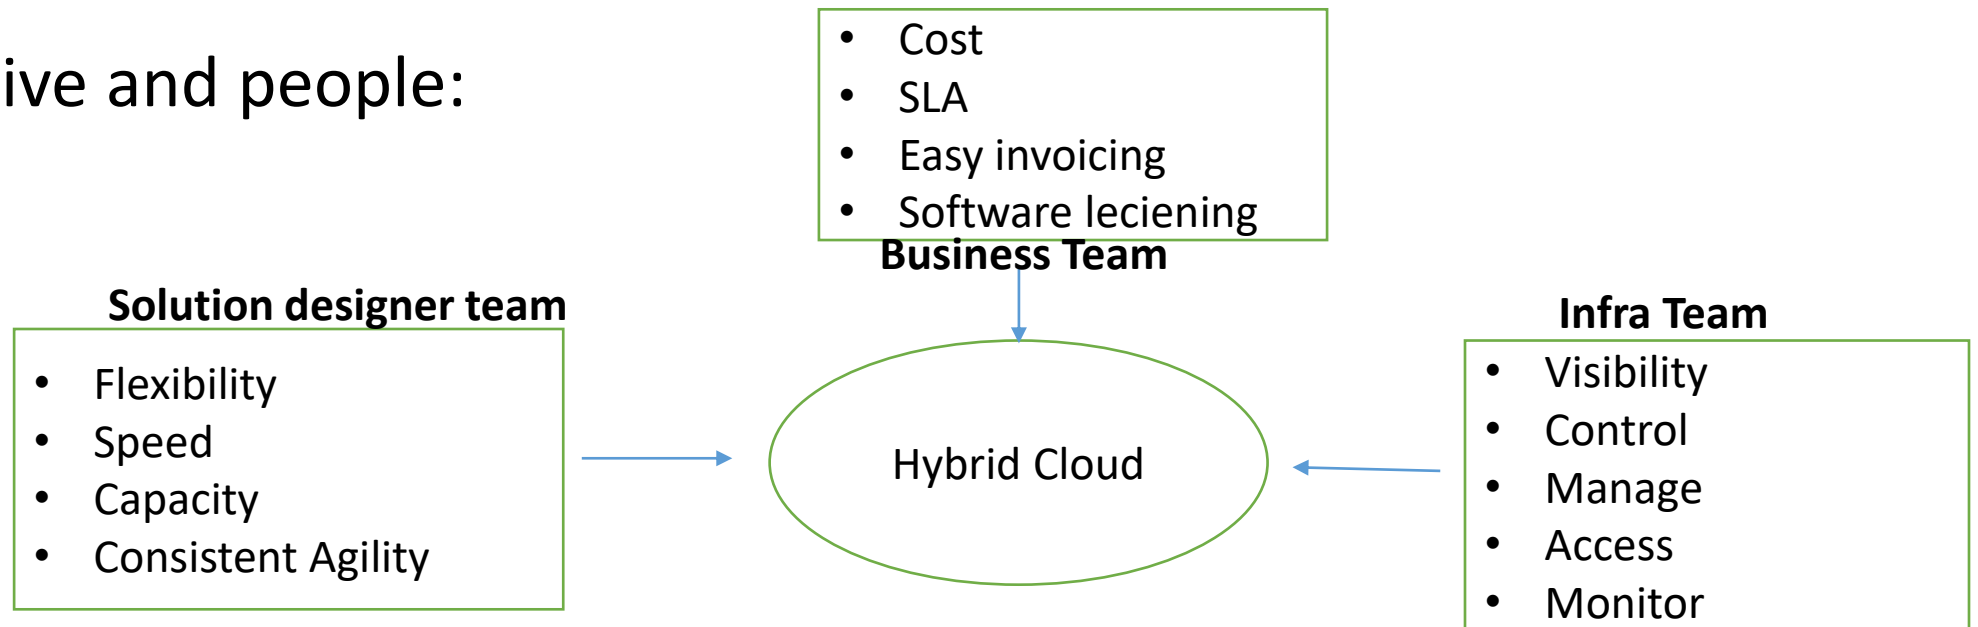


Outline

- Hybrid clouds definition and architecture
- Necessity, key strategies and workloads
- Hybrid cloud deployment options -Azure stack, AWS hybrid clouds, VMware solutions
- Cloud management platform – ManageIQ
- Aneka Hybrid cloud
- Advantages and disadvantages of hybrid clouds
- Insights to multi clouds

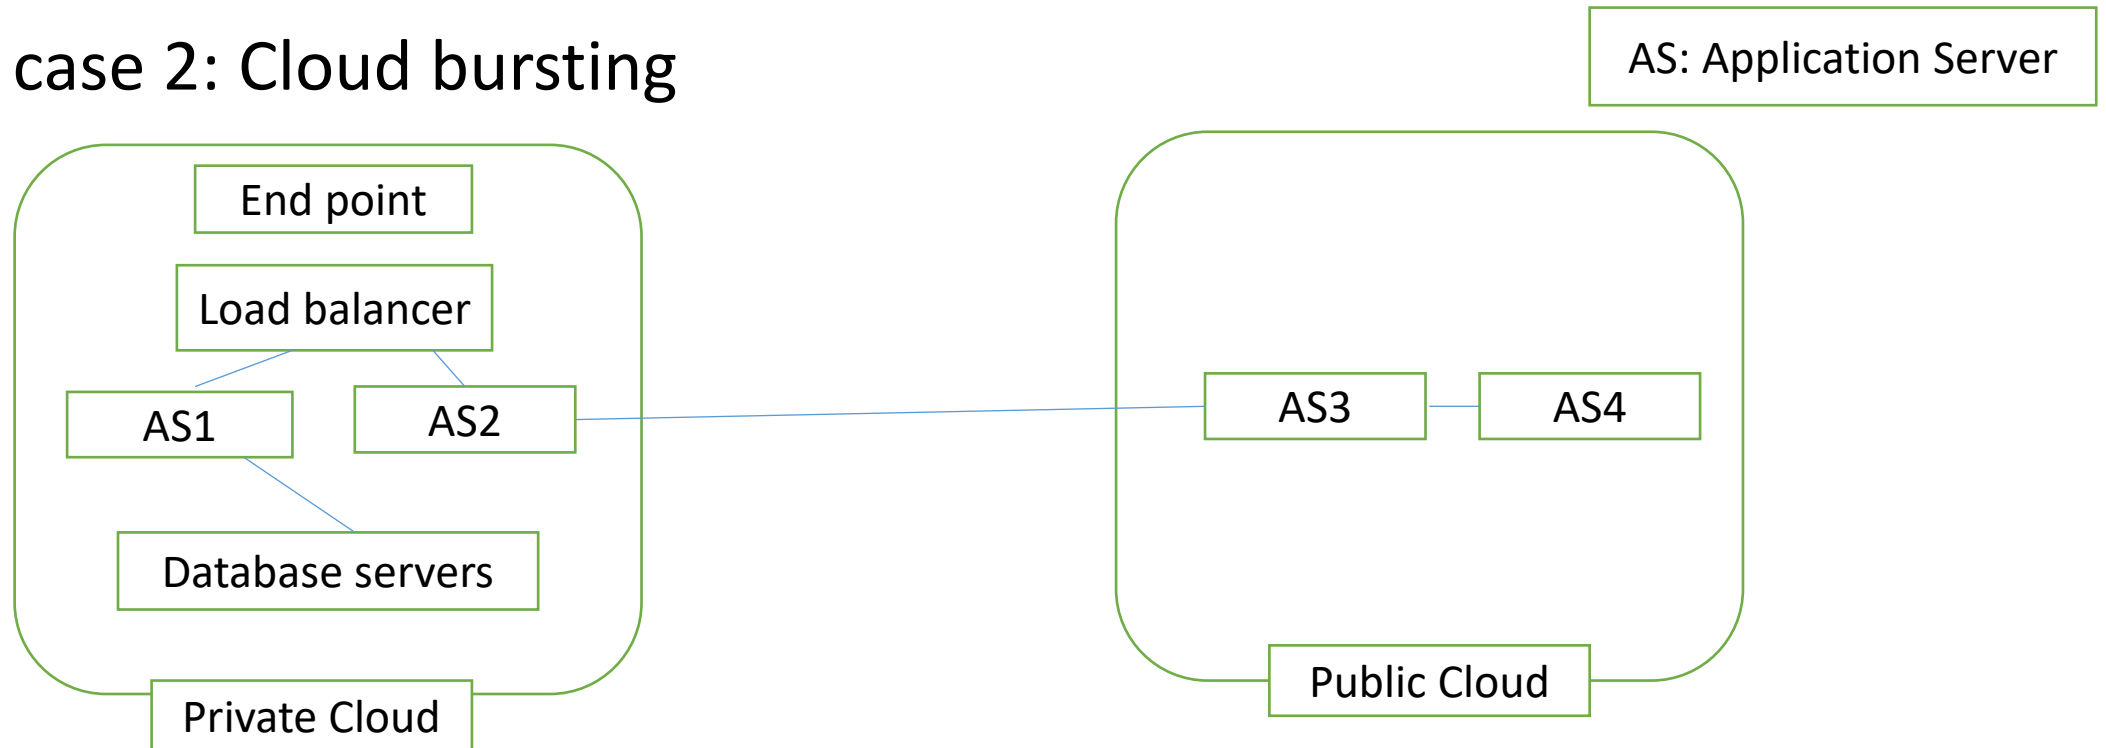
Hybrid clouds

- Hybrid clouds is combination of **on premise** infrastructure with **off premise resource** or public clouds such as AWS, IBM,GCP or Microsoft Azure.
- **Cloud bursting**, Shrink by provisioning external resources when required and unleased when not needed.
- Perspective and people:



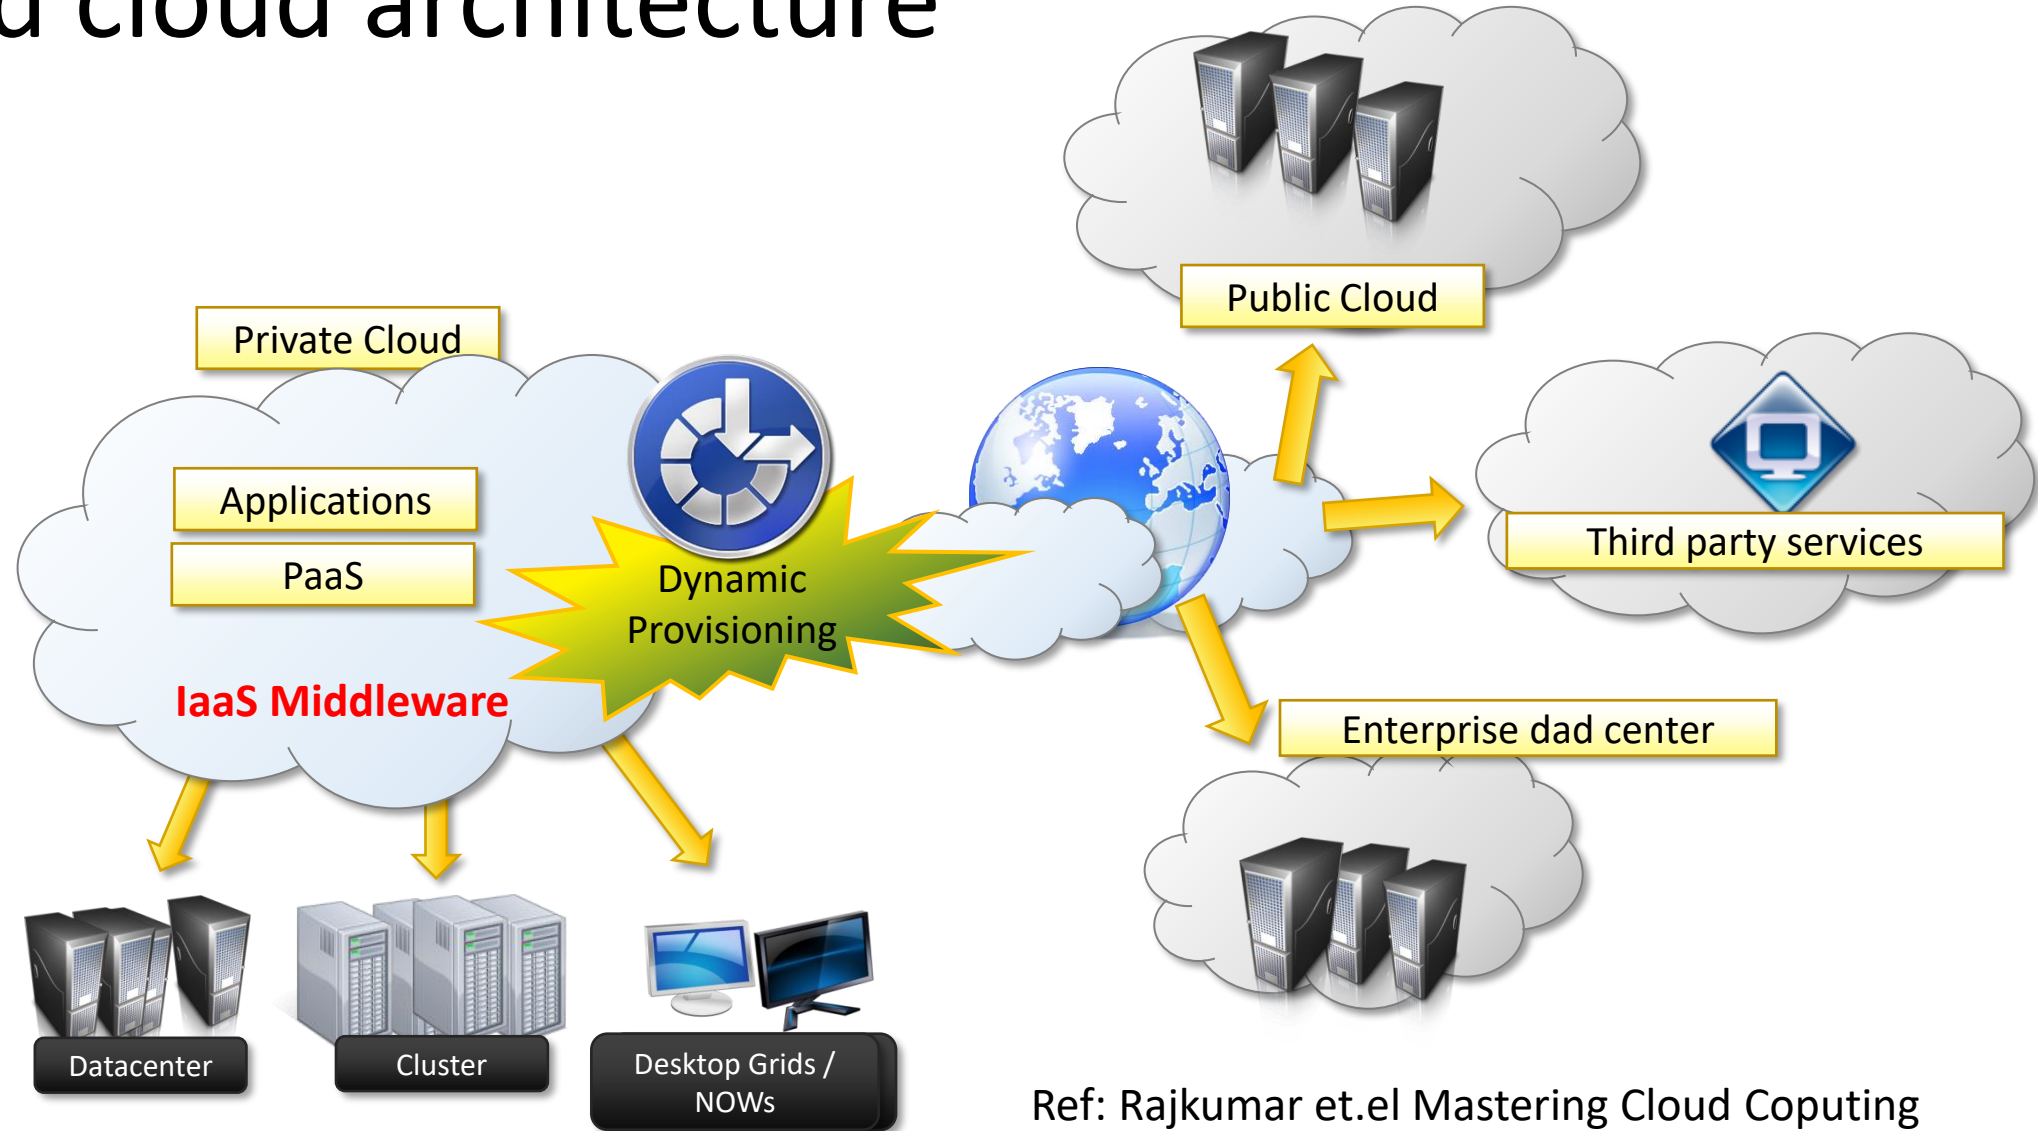
Necessity of hybrid cloud

- Use Case 1: Untested workloads with uncertain demand.
- Use case 2: Cloud bursting



- Use 3: Disaster recovery and high availability

Hybrid cloud architecture



Ref: Rajkumar et.al Mastering Cloud Computing

Key Strategies for Hybrid Cloud Success

The cloud operating system that gives you the ability to **monitor**, **manage**, and **orchestrate** across all environments using a single set of tools.

- Choosing a framework **a cloud Operating systems** like opens tack, cloud stack that will allow to manage workloads on-premises and in the cloud.
- **Modernize** your on-premises environments in accordance with that framework.
- Choose only **public clouds** and CSPs **compatible** with that framework.

Source <https://cdw-prod.adobecqms.net/content/dam/cdw/on-domain-cdw/brands/nutanix/o-reilly-ebook-designing-and-building-a-hybrid-cloud.pdf>

Capabilities required -hybrid cloud

- **Integration:** Integration of data and processes needs to be seamless across the deployments,
- **Data localization:** Security and compliance
- **Operational visibility and management:** Unified management
- **Security services:** Policies and business rules are needed to understand and monitor information about workloads
- **DevOps:** Single unified way to manage and synchronize applications
- **Integration services:** Approaches that allows application code to be easily moved from one environment to another.

Possible Workloads in hybrid cloud

- Workload- is an independent service or collection of code that can be executed.
- Workloads need to be executed based on: Criticality, Level of security, interaction with other environments.
- Types of workloads
 - **Batch** - customer bills and account statements
 - **Analytic** – high volume, complexity, and importance
 - **Transactional**- automation of business processes
 - **Database**- high amount of Input/Out (I/O) cycles

Assessing the needs of hybrid cloud

1. Assessing the current state of operations
 - Identify the current IT infra locations.
 - Identify the status of infra like utilization, cost, physical space.
 - Identify the current cloud provider impact on your business.
 - Identify workload characteristics of your business.
2. Assessing the future workloads:
 - Identify new applications, new products planning in a year.
 - Planning any updates for business applications.
 - Prioritize the workloads and assess type of application, resource requirements, data locations.
3. Create decision matrix
 - Decision matrix can help to determine where to locate each workload going forward and how to assess anticipated costs.

Use cases

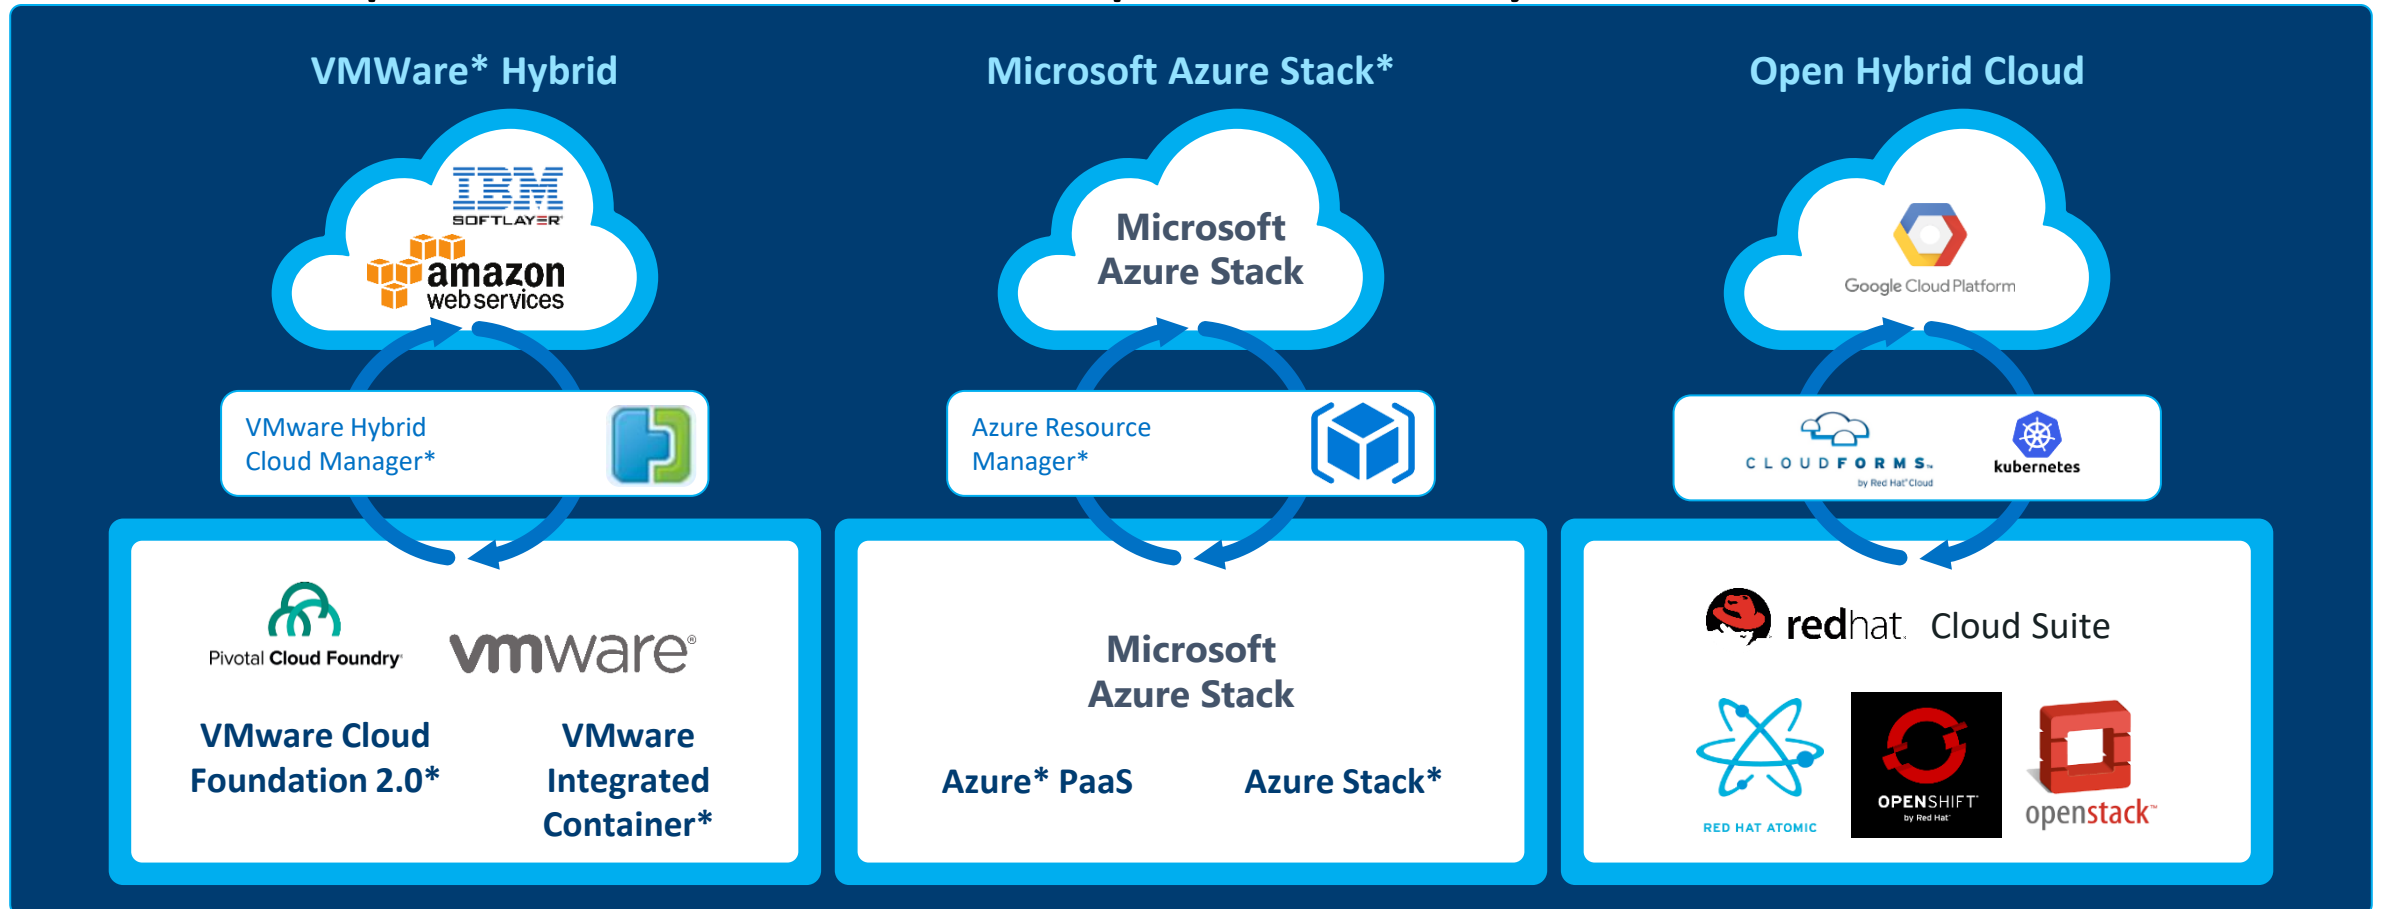
- **Health care sector** : Security and compliance in storing and movement of health data.
- **Finance sector**: Flexibility, power, scale, and seamless connectivity.
- **Government**: Use cloud email and collaboration tools and a mixture of public and private cloud infrastructure.
- **Retail**: Always in (99.999%) uptime support in ecommerce operations.
- **ecommerce**: Enormous web traffic and scalability.

Source: <https://www.harbourit.com.au/blog/everything-you-need-to-know-about-hybrid-cloud/>

Options for hybrid cloud deployment

- Vendor specific hybrid clouds
 - Microsoft azure stack, VMware Hybrid Cloud
- Customizable hybrid clouds
 - Based on your choices

Example for Vendor specific hybrid clouds



Microsoft Azure stack



- Extended Azure services and capabilities to environment of choice— from the datacenter to edge locations and remote offices—with Azure Stack.
- Build, deploy and run hybrid and edge computing apps consistently across IT ecosystem, with flexibility for diverse workloads.
- It has three solutions:
 - Azure Stack Edge- Machine learning at edge, IoT and data transfer from edge to cloud
 - Azure Stack HCI- Remote offices, high performance workloads
 - Azure Stack Hub- Connected and disconnected scenarios, data sovereignty

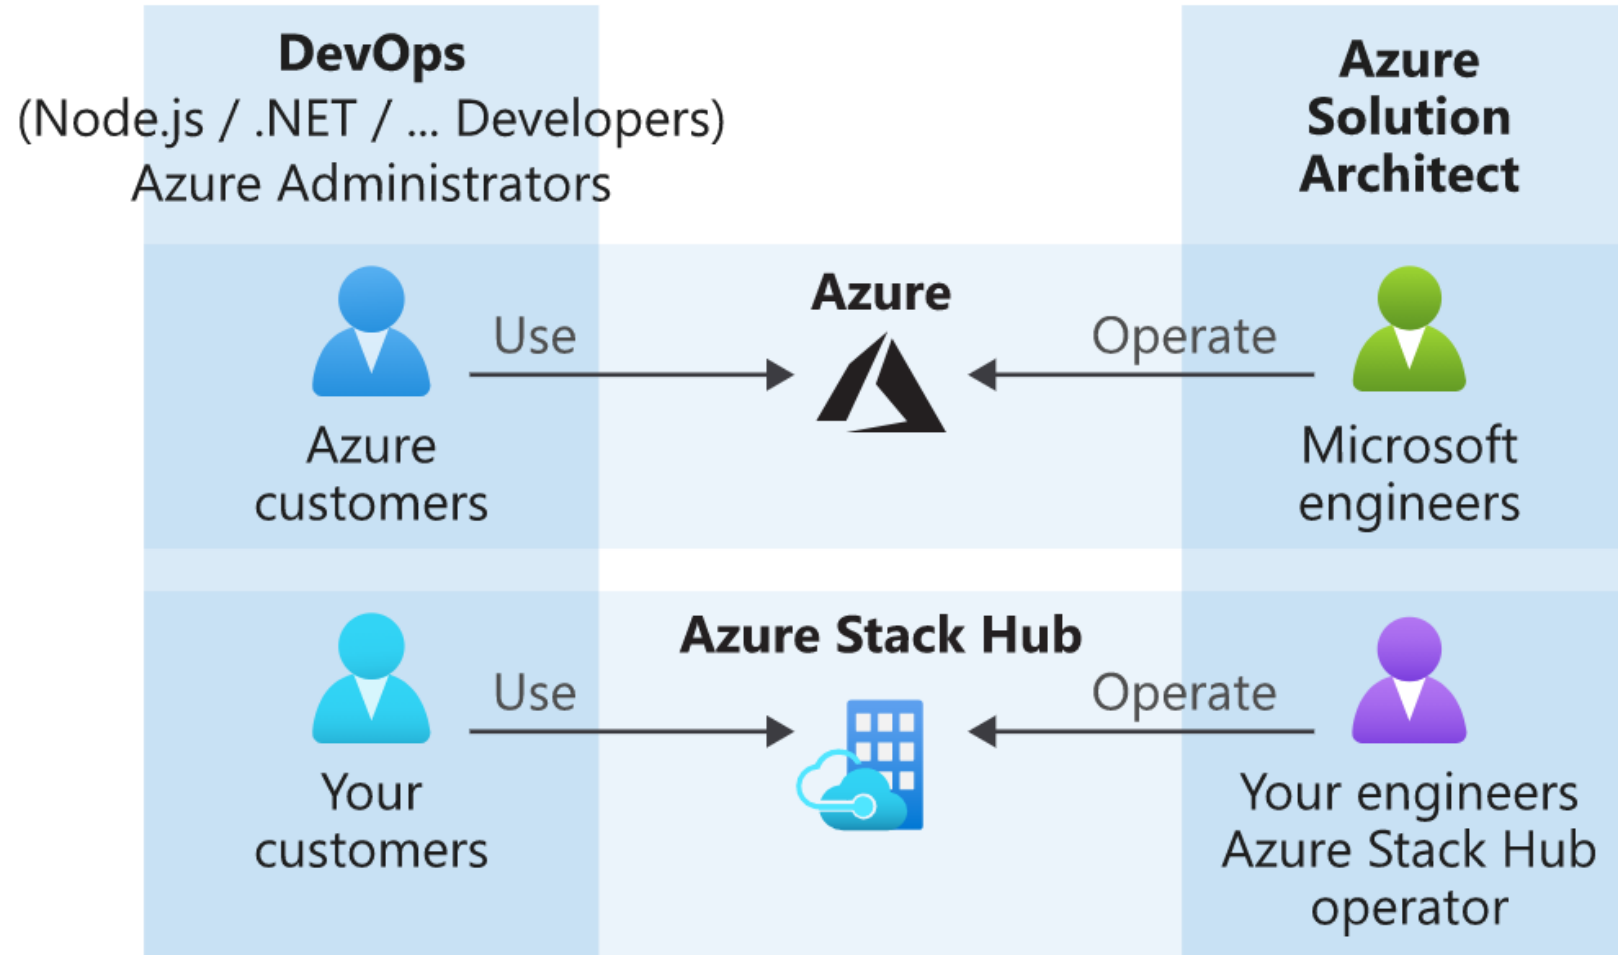
Azure Stack Hub

- Azure stack hub is extension of Azure that provides on premise environment to run the applications and deliver azure services locally on the datacentre.
- Its main focus on data residency - Azure Stack Hub is an on-premises appliance. Customers fully own and control the appliance, access to the appliance, and any data stored on the appliance.
- A customer can alternatively elect to connect an Azure Stack Hub appliance to global Azure or to the Internet in a hybrid workload scenario

Azure stack hub architecture

- It consists of integrated systems in racks of 4-16 servers built by trusted hardware vendors and delivered in straightforward to customer datacentre.
- Solution provider will assist in deploying and configuring the servers and setup ready to go azure solutions on-premise.
- Users can apply consistent DevOps processes whether they are connected to Azure or not.
- Integrated systems are manufactured by trusted hardware vendors in connection with Microsoft.

How is Azure Stack Hub managed?



Azure Stack Edge

- Used to run the workloads and get quick actionable insights right at where data is created.
- Its basically hardware as a service with azure stack edge built with all azure services.
- These devices can be managed from cloud using azure management tools.
- With Azure IoT Edge, can manage and deploy containers from IoT hub and integrate with Azure IoT solutions at the edge with rugged options using Kubernetes with multi-node and virtual machine support.
- Azure stack edge devices are enabled with NVIDIA T4 Tensor Core GPU and Intel VPU to run ML models.

Azure Stack Edge Pro and Mini Series devices

Pro

- 1U rack-mountable appliance, optimised for conditions inside a datacentre or branch location.
- Available with 1-2 NVIDIA T4 GPUs



Pro R

- Ruggedised datacentre-grade power with a built-in NVIDIA T4 GPU, in a transportable case for remote locations.
- Available options: With or without Uninterruptable Power Supply (UPS)



Mini R

- Ruggedized, battery-operated device-small enough to fit into a backpack-designed for harsh environments and disconnected scenarios. Includes a built-in Intel VPU for edge processing.



AWS Hybrid Cloud

- AWS hybrid cloud solutions and services can help to run and manage applications wherever they may need to reside.
- With AWS, can use the same infrastructure, services, APIs, and tools wherever need it – from the cloud, to on premises, and to the edge.
- This is for some need to remain on premises or in a specific location due to low latency, local data processing, or data residency requirements.
- Hybrid Infrastructure solutions- AWS outposts, AWS wavelength, AWS local zones
- Rugged & Disconnected Edge – AWS Snowball, AWS Snowcone

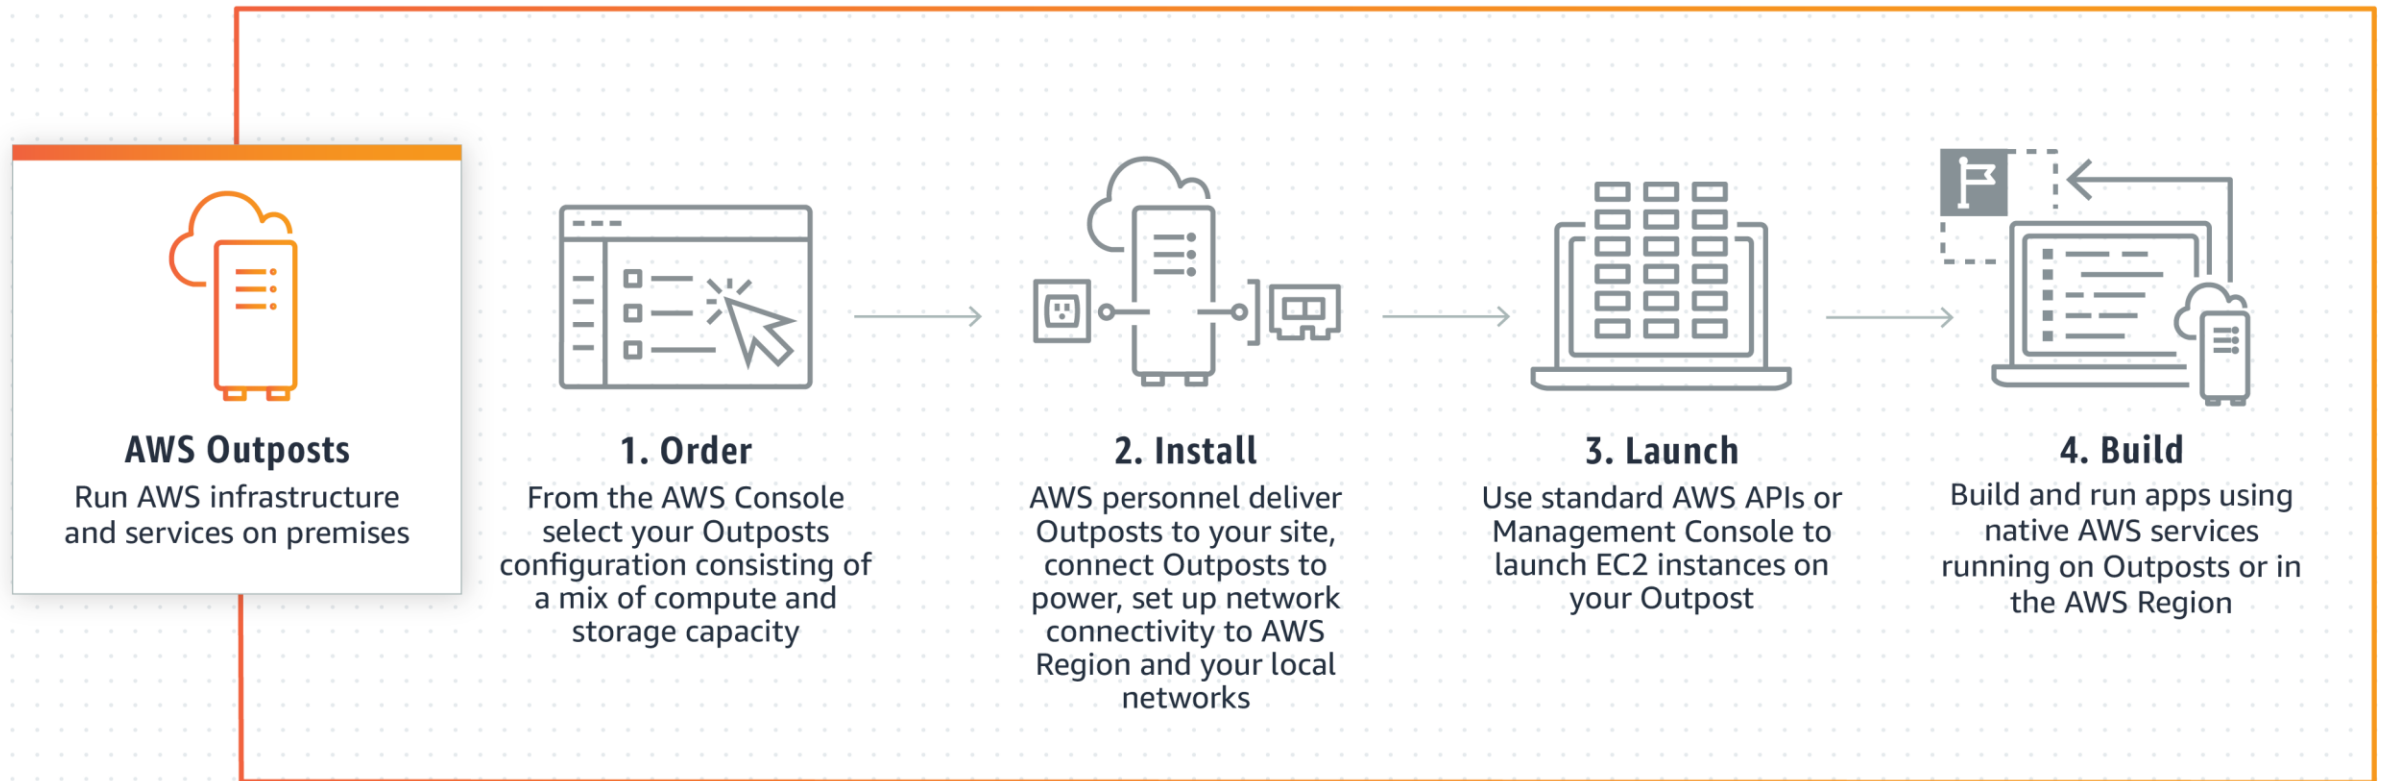
AWS Outposts



AWS Outposts

- Fully managed service offers the same AWS infrastructure, AWS services, APIs, and tools.
- AWS compute, storage, database, and other services run locally on Outposts.
- Outposts is available as a 42U rack that can scale from 1 rack to 96 racks to create pools of compute and storage capacity
- Outposts will be available in two smaller form factors, 1U and 2U rack-mountable servers for locations with limited space or capacity requirements

How it works?

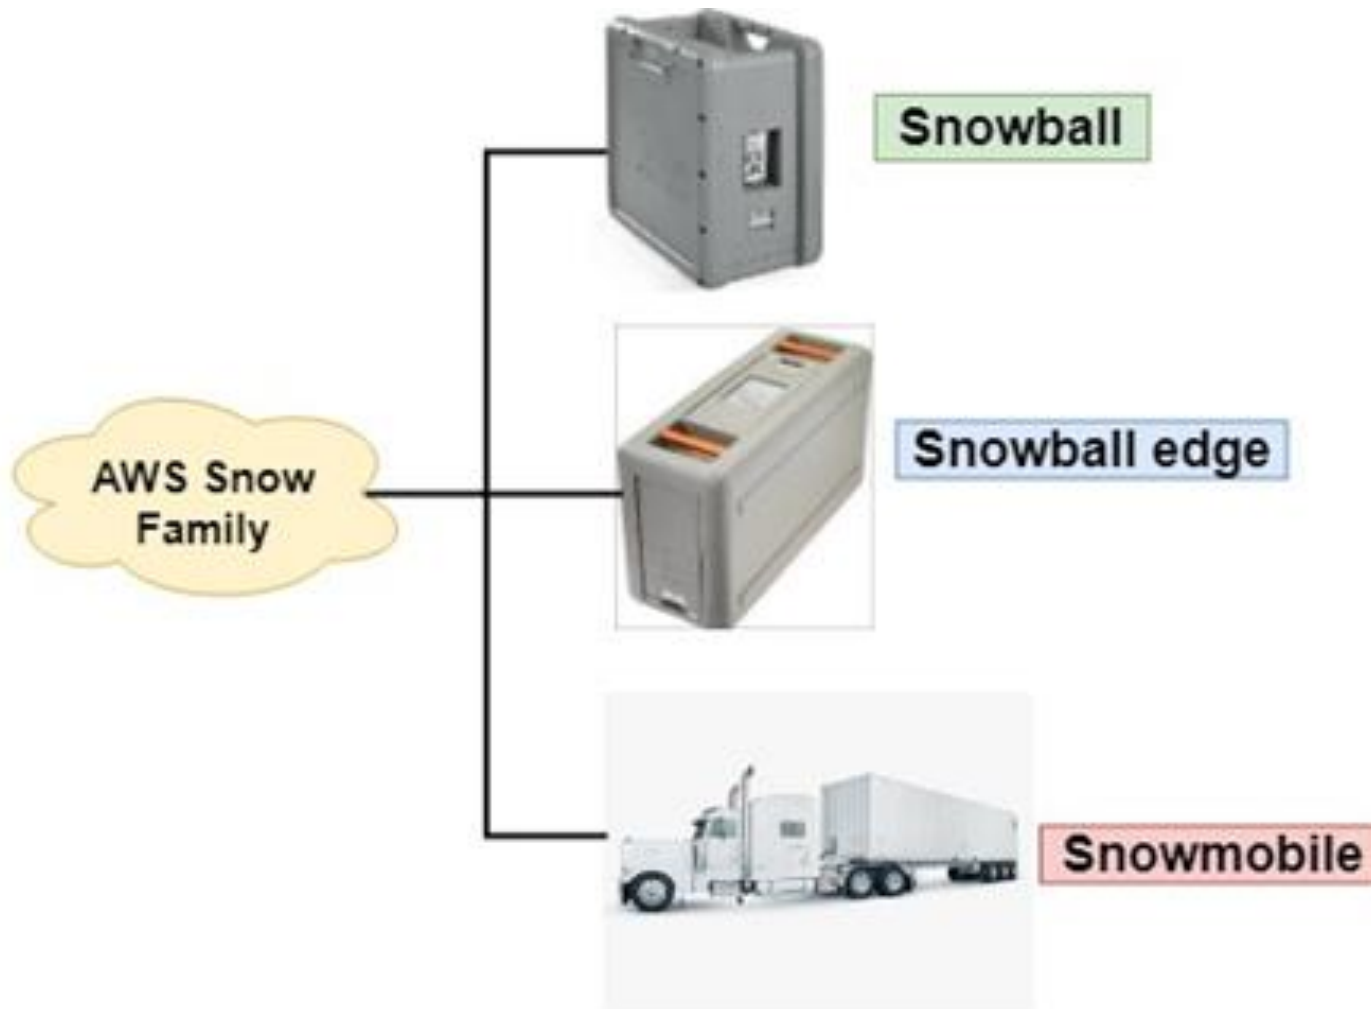


Rugged & Disconnected Edge – AWS Snowball



- Used for edge computing, data migration and edge storage.
- It has two options
 - Snowball Edge Storage Optimized devices - block storage and Amazon S3-compatible object storage, and 40 vCPUs.
 - Snowball Edge Compute Optimized - 52 vCPUs, block and object storage, and an optional GPU.
- These devices may also be rack mounted and clustered together to build larger temporary installations.
- Snowball supports specific Amazon EC2 instance types and AWS Lambda functions
- Develop and test in the AWS Cloud, then deploy applications on devices in remote locations to collect, pre-process, and ship the data to AWS


AWS Snow family



AWS Snowball

- Snowball is a petabyte-scale data transport solution
- It is a streamline bringing the data into aws and bypassing an internet.
- high network costs, long transfer time, and a security issue have been resolved by using Snowball

How Snowball moves data into and out of AWS



AWS Snowball Edge

- Snowball Edge is a 100 TB data transfer device with on-board storage and compute capabilities.
- Snowball Edge is like an AWS data center that you can bring on-premises.
- Run Lambda functions from Snowball edge
- Faster Data transfer: It transfers the data with a speed of up to 100 GB/second.
- Clustering: You can cluster Snowball edges together to form a local storage tier and process your data on-premises to achieve 99.99% data durability across 5-10 devices

Example: Aircraft engine manufacturer

AWS Snowball mobile

- A Snowmobile is an exabyte-scale data transfer service.
- It can transfer large amounts of data in and out of AWS.
- You can transfer 100 PB per Snowmobile, a 45-foot long ruggedized shipping container, pulled by a semi-trailer truck
- Transferring data with Snowmobile is secure, fast and cost-effective.

VMware Hybrid Cloud

- VMware hybrid cloud is backed by portfolio of services providing the software defined building blocks for virtualized infrastructure to build and run the applications.
- VMware has its vendor specific solutions to build hybrid cloud
- It primarily consists of Software defined data center components:
 - **vSphere & vCenter** for compute virtualization and management
 - **vSAN** for storage virtualization
 - **NSX** for network virtualization

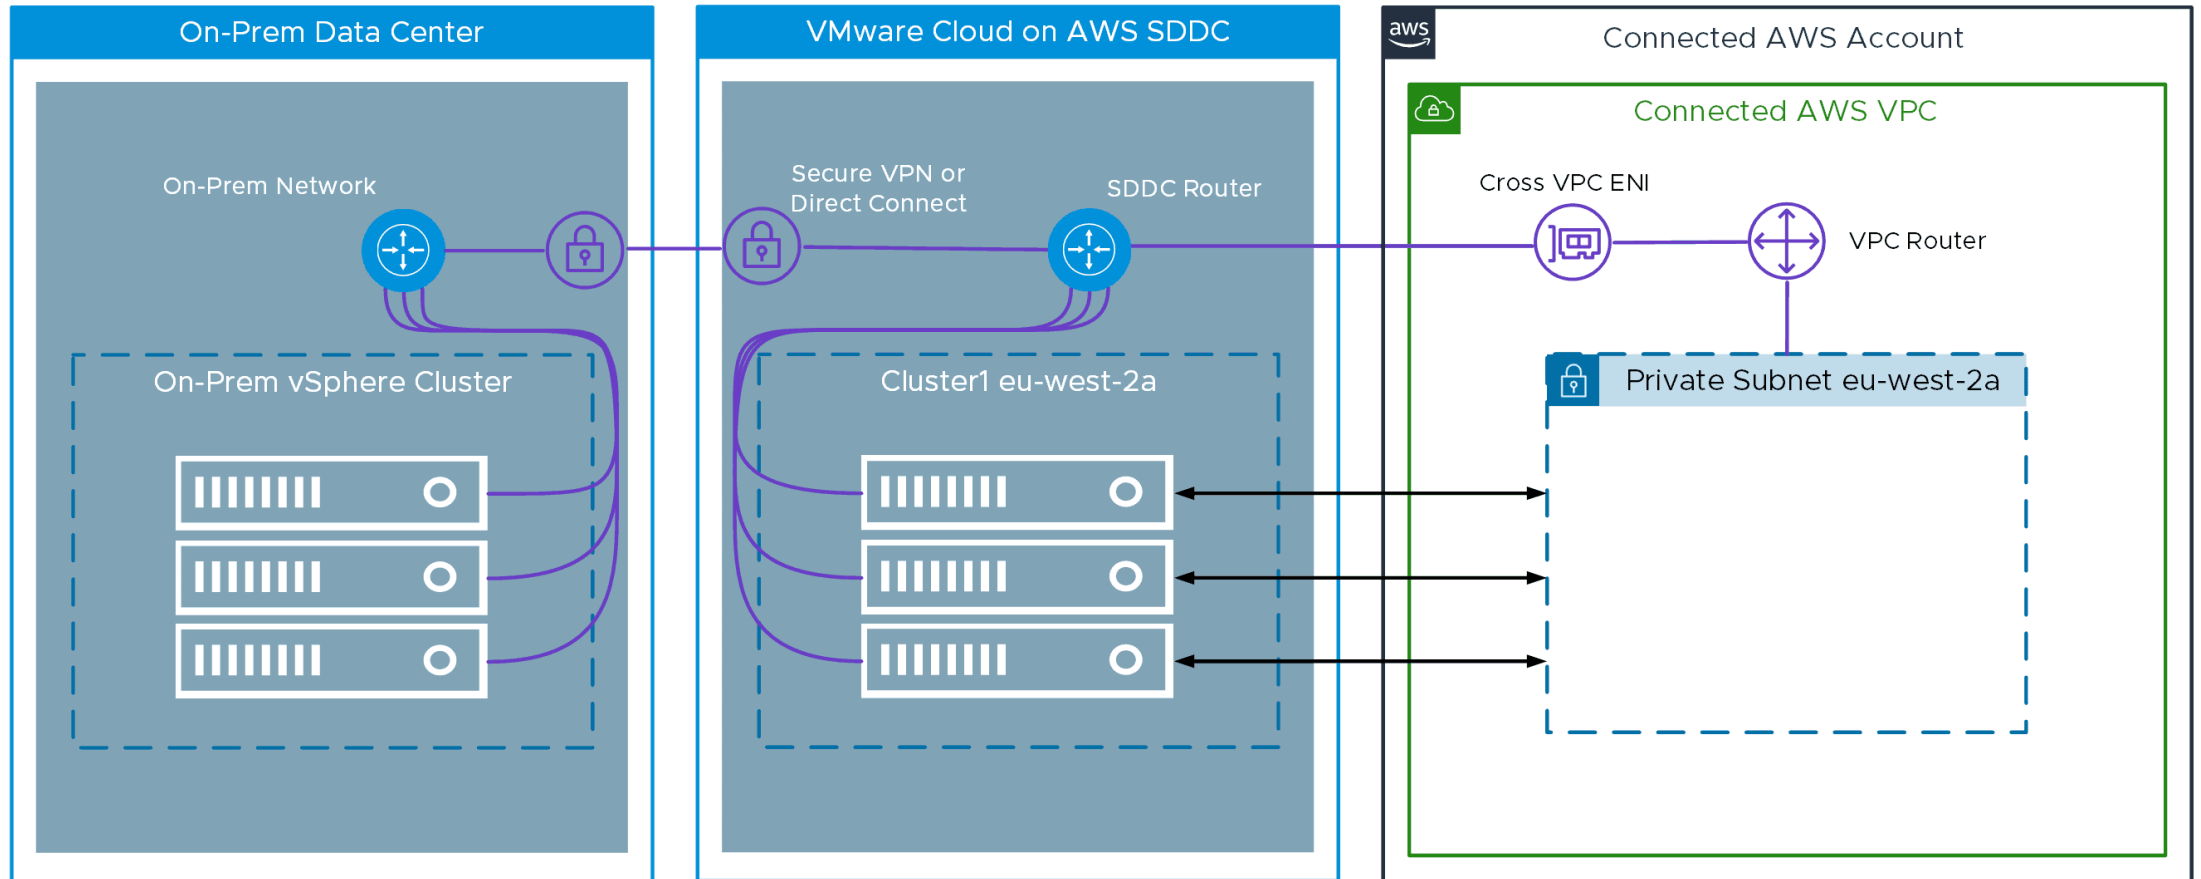
VMware options to design a hybrid cloud setup

- **VMware Cloud Foundation** with embedded Kubernetes capabilities
- **VMware Cloud Foundation on Dell EMC** physically installed in your facilities and maintained as a fully-managed service
- **Hyperconverged Infrastructure (HCI) powered by VMware vSAN** available on Dell EMC VxRail and vSAN ready nodes from multiple hardware vendors
- **VMware Cloud on AWS jointly engineered by VMware and AWS** and powered by the VMware Cloud Foundation SDDC stack running on AWS infrastructure

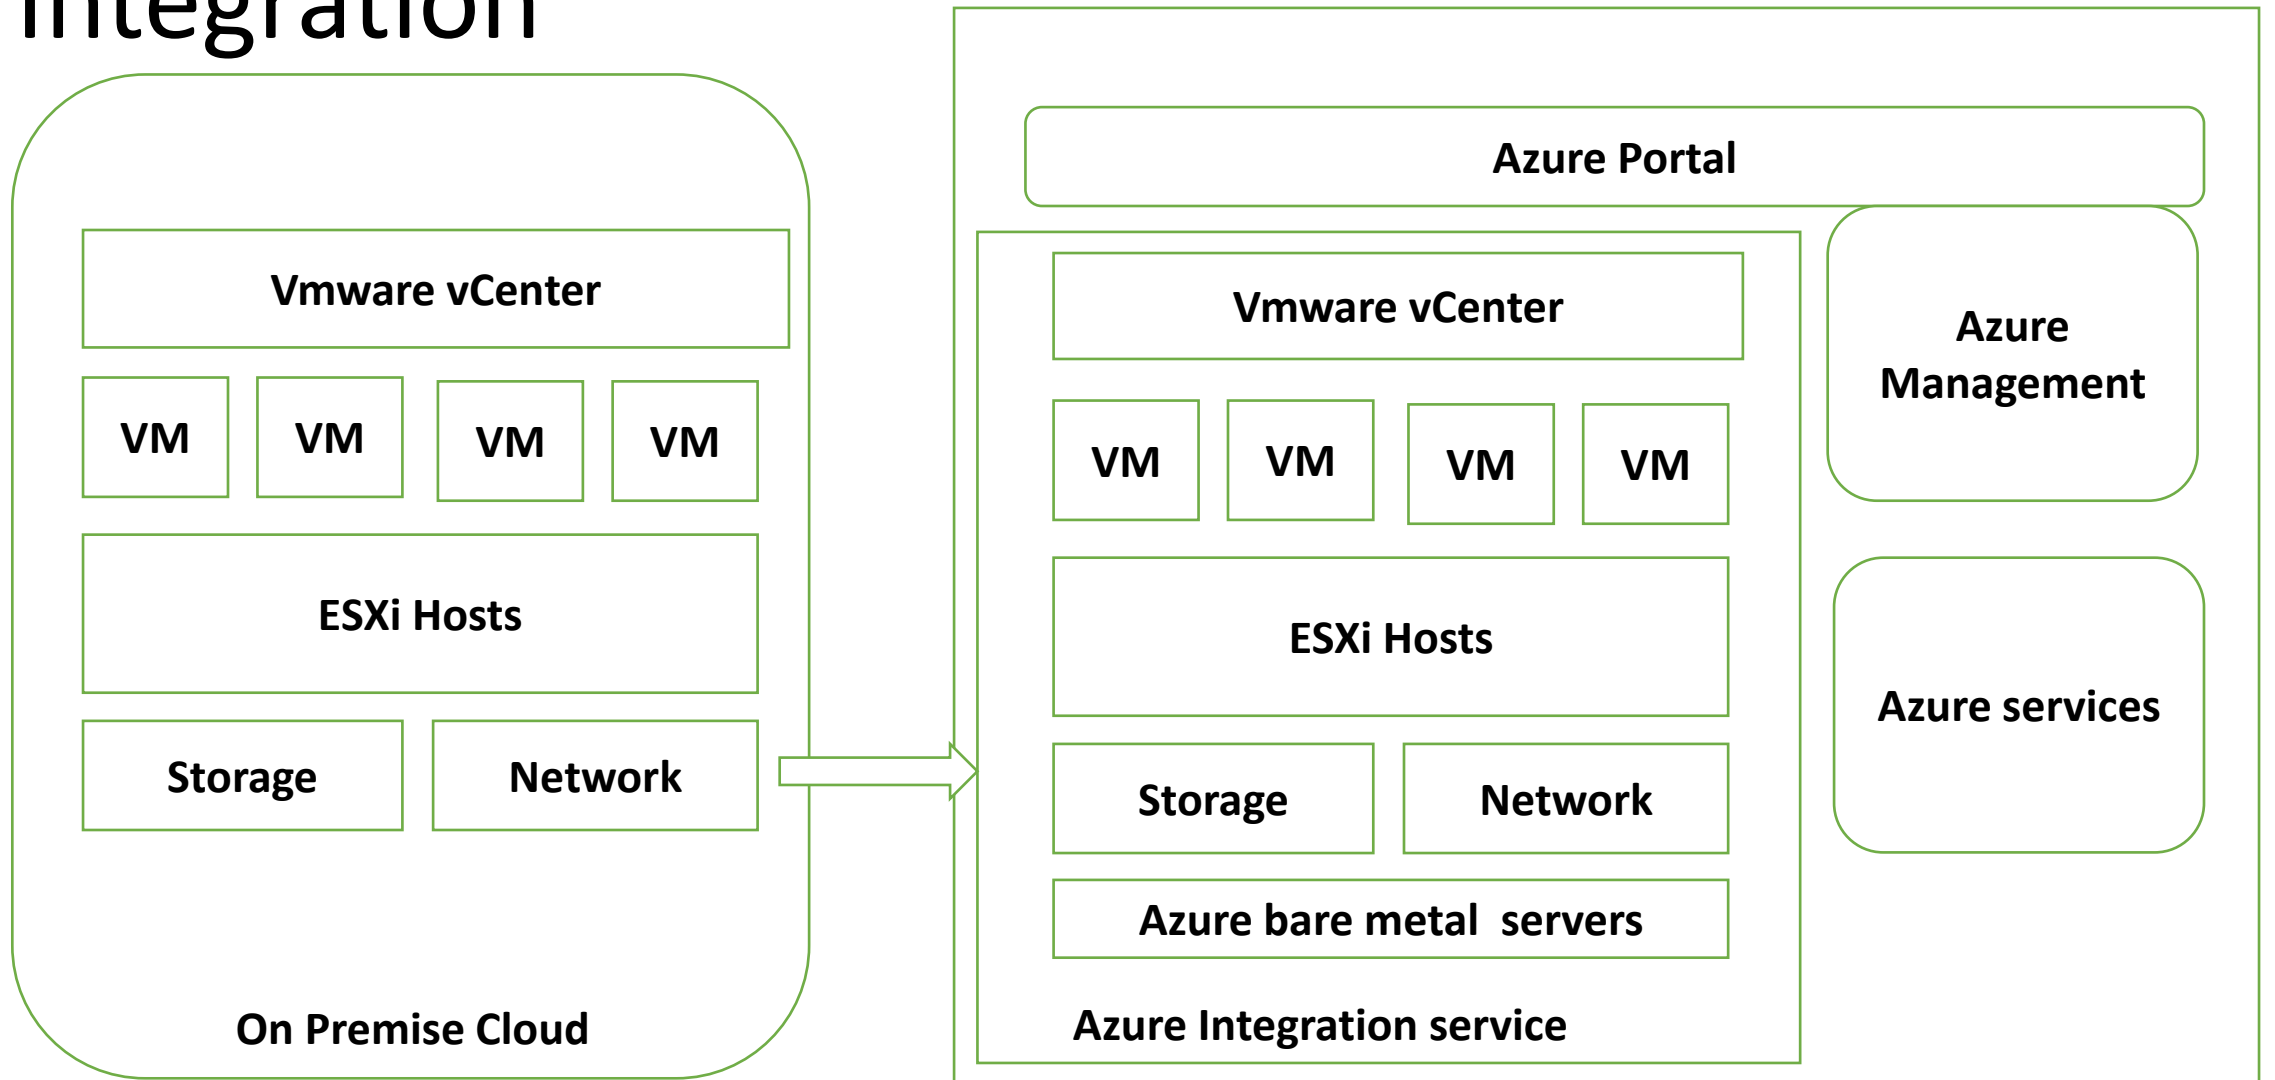
VMware options to design a hybrid cloud setup

- **Azure VMware Solution**
- **Google Cloud VMware Engine**
- **IBM Cloud for VMware Solutions**
- **Oracle Cloud VMware Solution**

VMware Hybrid cloud solution using AWS




Hybrid Cloud: VmWare and Microsoft Azure Integration

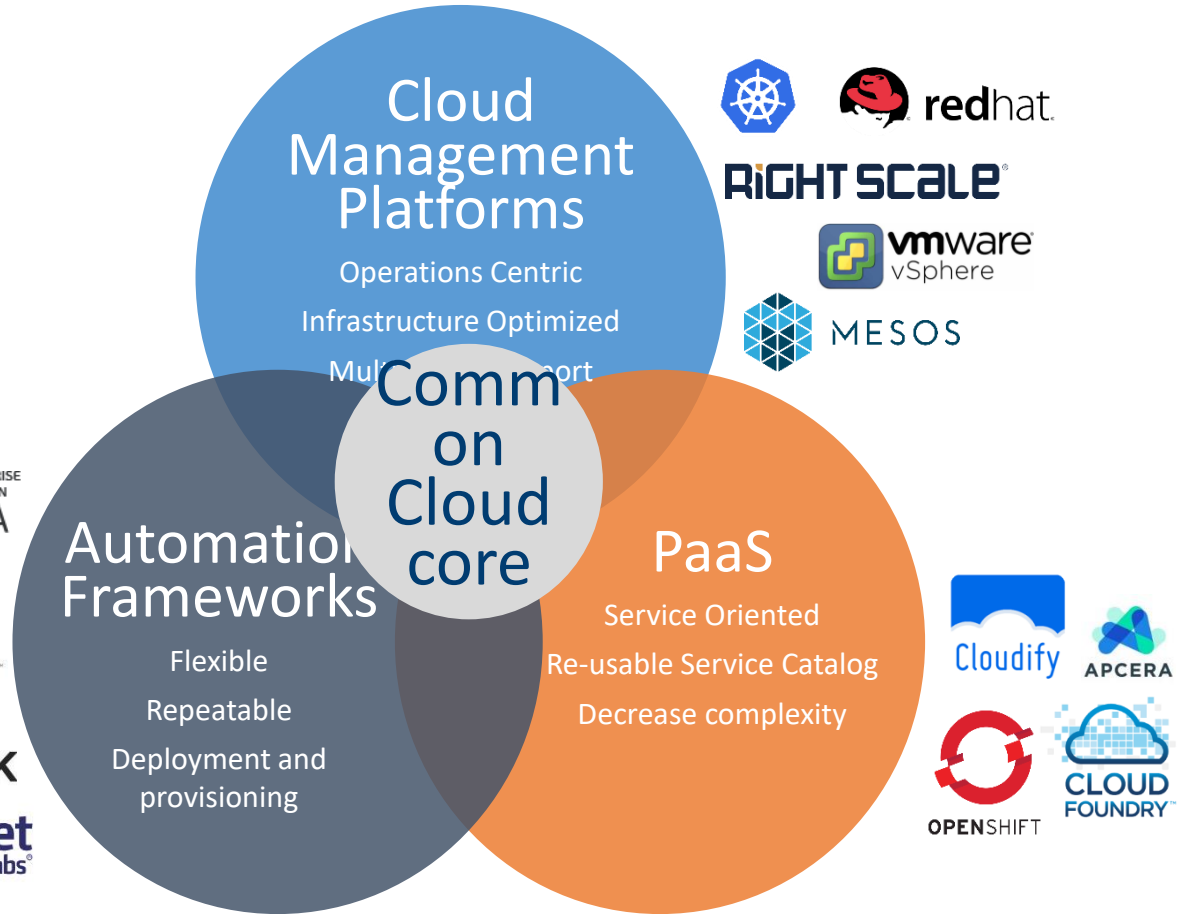


Customizable hybrid clouds

Private cloud of your choice



Public cloud of your choice

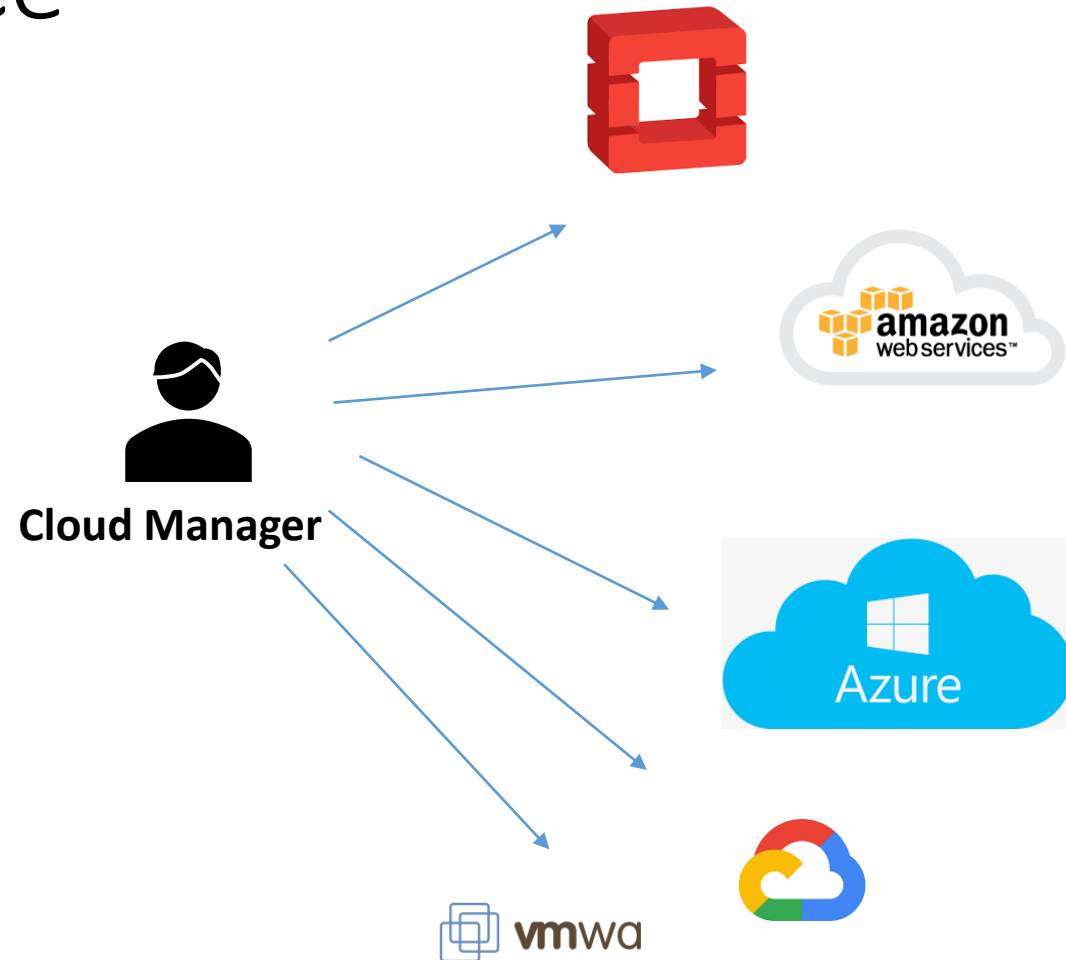


Hybrid cloud management platforms

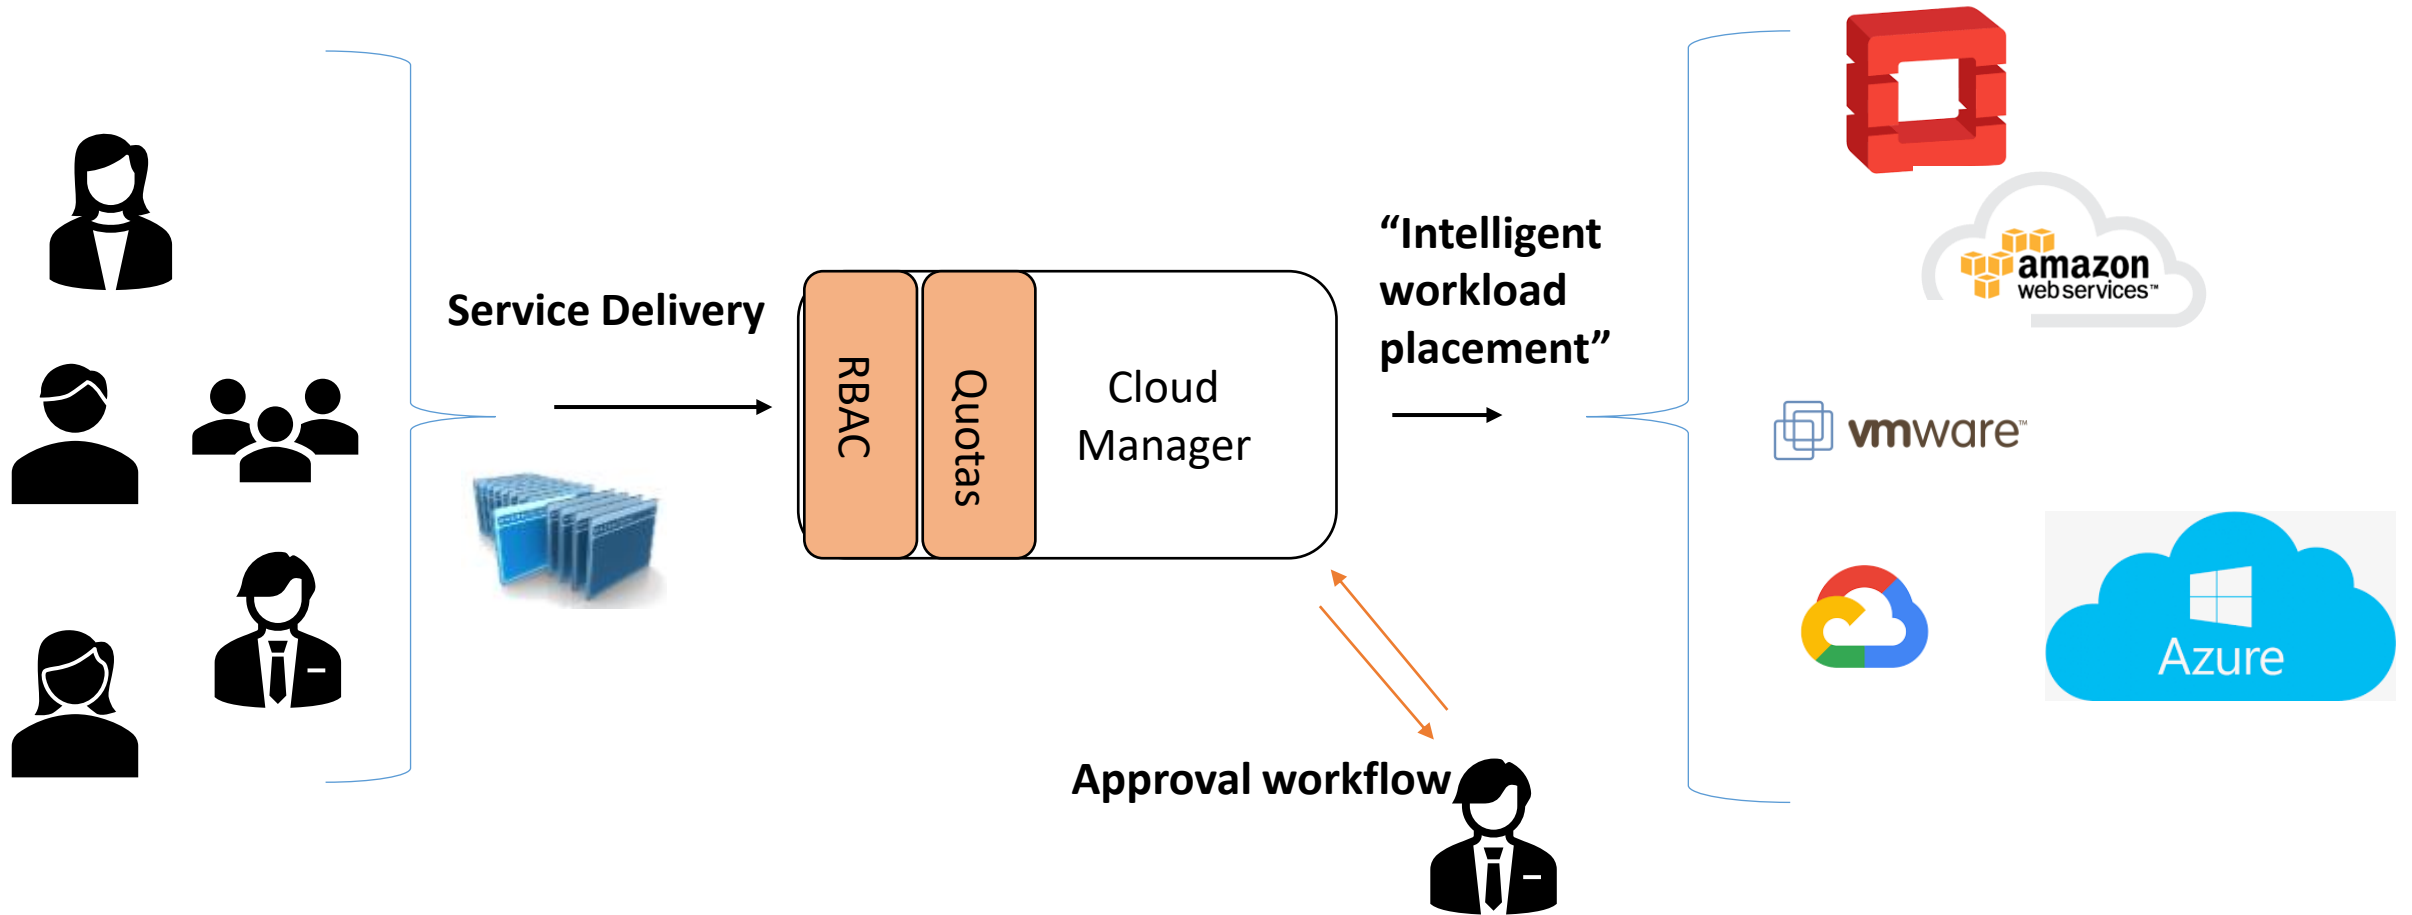
- Hybrid cloud management tool helps to automate, orchestrate the public and private clouds by enforcing the policies by obeying the quotas with automated workflow management.
- Features
 - Self service provisioning
 - Cloud workload management
 - Show back/Chargeback
 - Bursting workloads
 - Capacity planning and management
 - Leveraging Existing infrastructure

Seamless self service

- Role based delegation
- Self service portals
- Self service catalogs
- Automated provisioning
- Quotas and Chargeback



Deployment of Self-Service resources



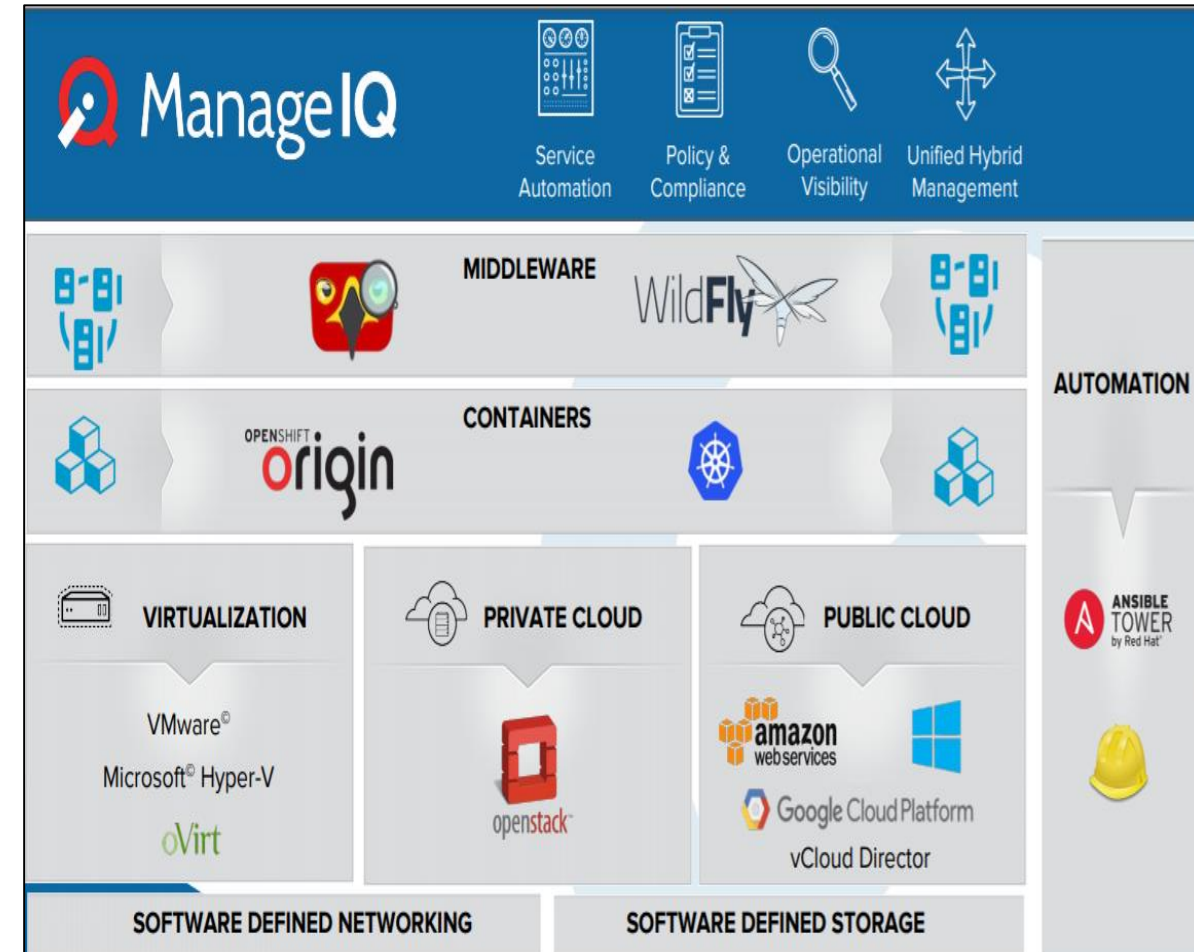
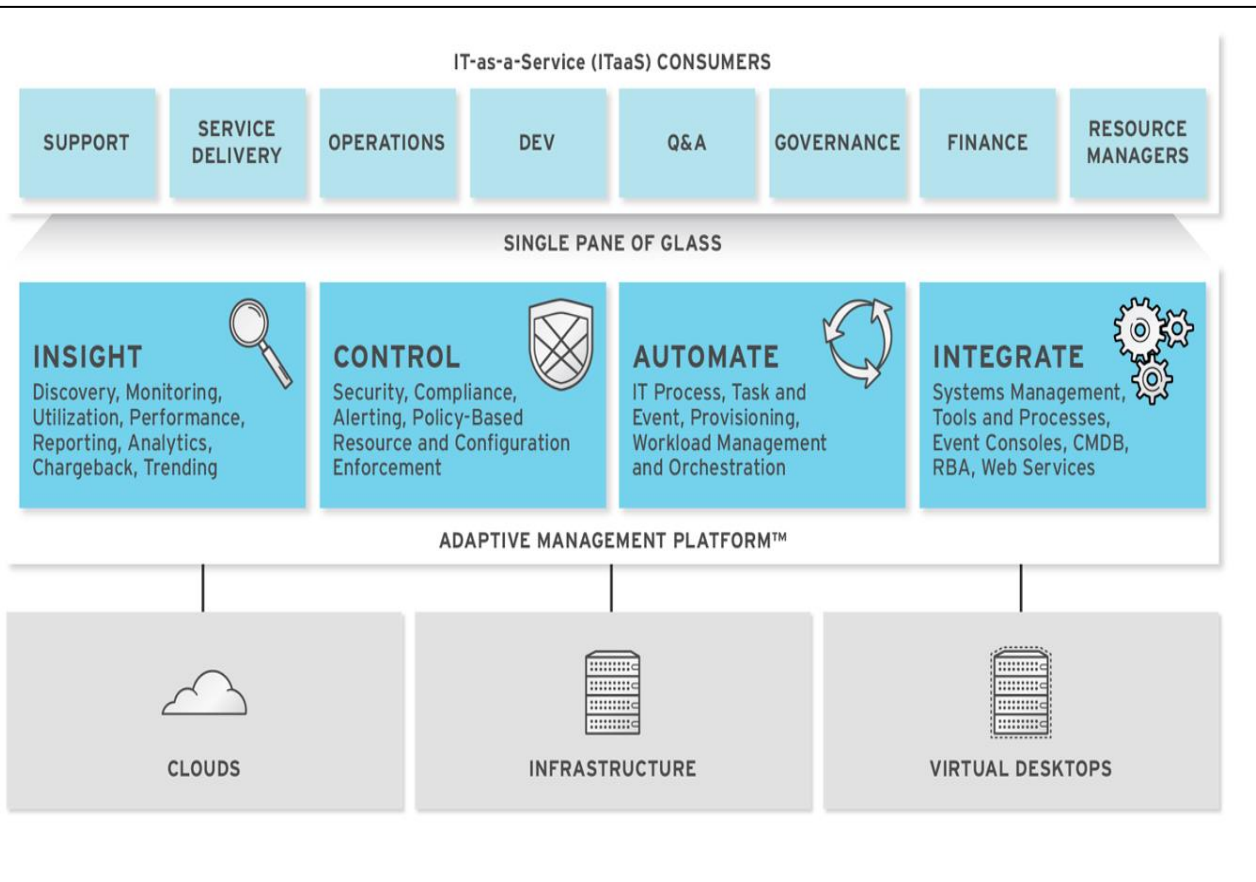
Examples for hybrid cloud management platforms

- ManageIQ
- Redhat Cloudforms
- Morpheus
- Cloudbolt

ManageIQ

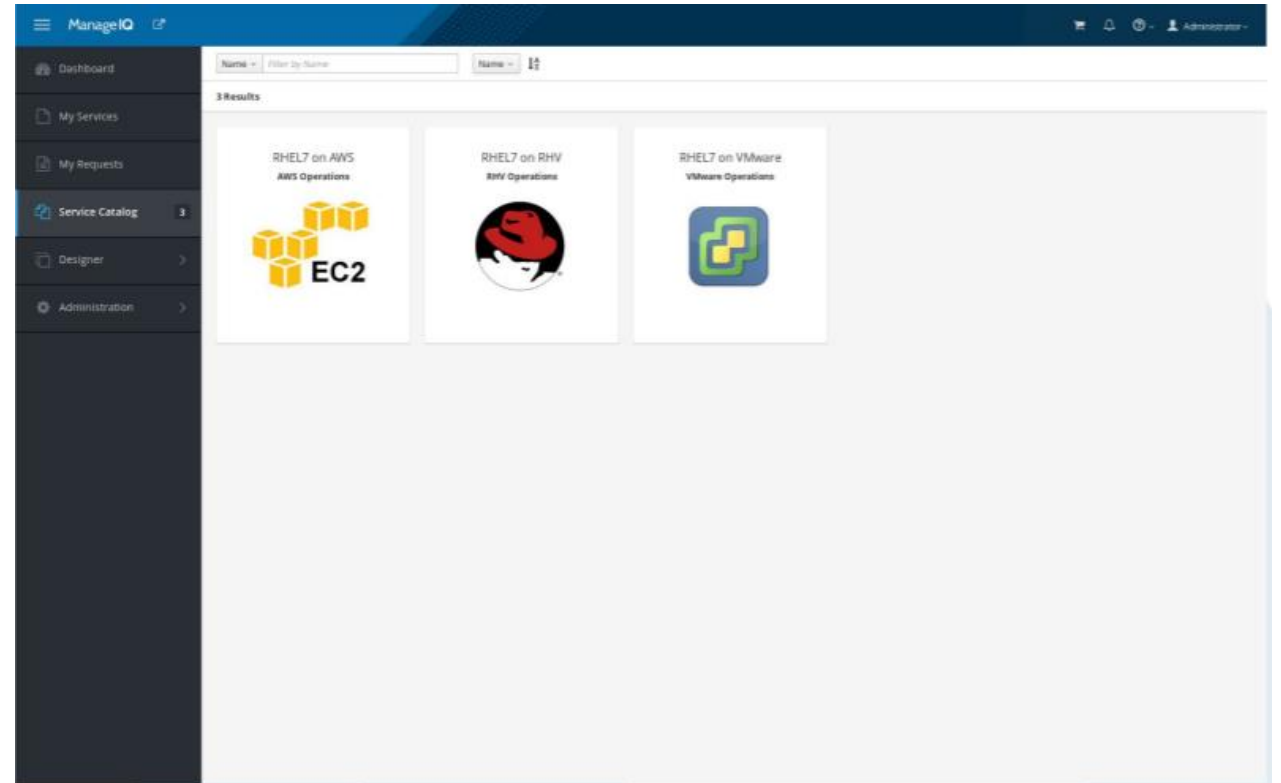
- It's hybrid cloud management tool that helps to discover, optimize and control the hybrid cloud environments.
- Manage containers, virtual machines, networks and storage in a single platform.
- It was started as a open source project in 2006 and later acquired by Red hat.
- Supported by Xlab, Google, Lenovo and red hat, etc,.
- It supports all the major public cloud service providers such as aws, azure, etc,.

ManageIQ Architecture



ManageIQ- Self service delivery

- Create service delivery catalogs for users to choose the services that they need to deploy.
- Shopping cart functionality allows multiple services to be requested at one time.
- Service requests can be routed for approval.



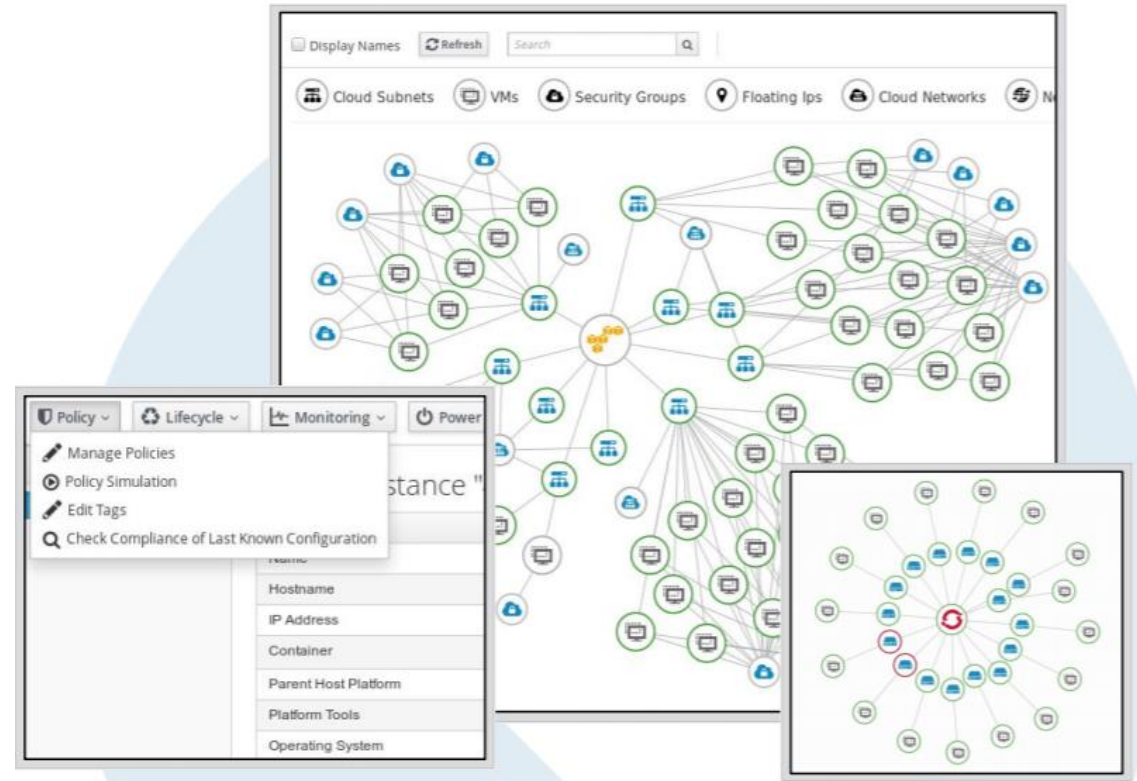
ManageIQ- Automated provisioning

- Automatically deploys and configures requested services on any infrastructure platform.
- Automation steps can be codified in Ansible playbooks or natively in ManageIQ.
- Integration to external IT systems allows ManageIQ to automate all process steps.



ManageIQ- Root cause analysis

- View instance performance and resource usage over time to pinpoint problem initiation.
- Quickly compare system state against known good state or other systems.
- Navigate across relationships and drill down infrastructure layers to identify underlying causes.



ManageIQ- Policy enforcement

- Continuous discovery and deep SmartState inspection of virtual instances.
- Policy violations can raise alerts or be remediated automatically.
- Policy can be applied uniformly or based on virtual instance criteria.



Hybrid cloud- Case study with Aneka

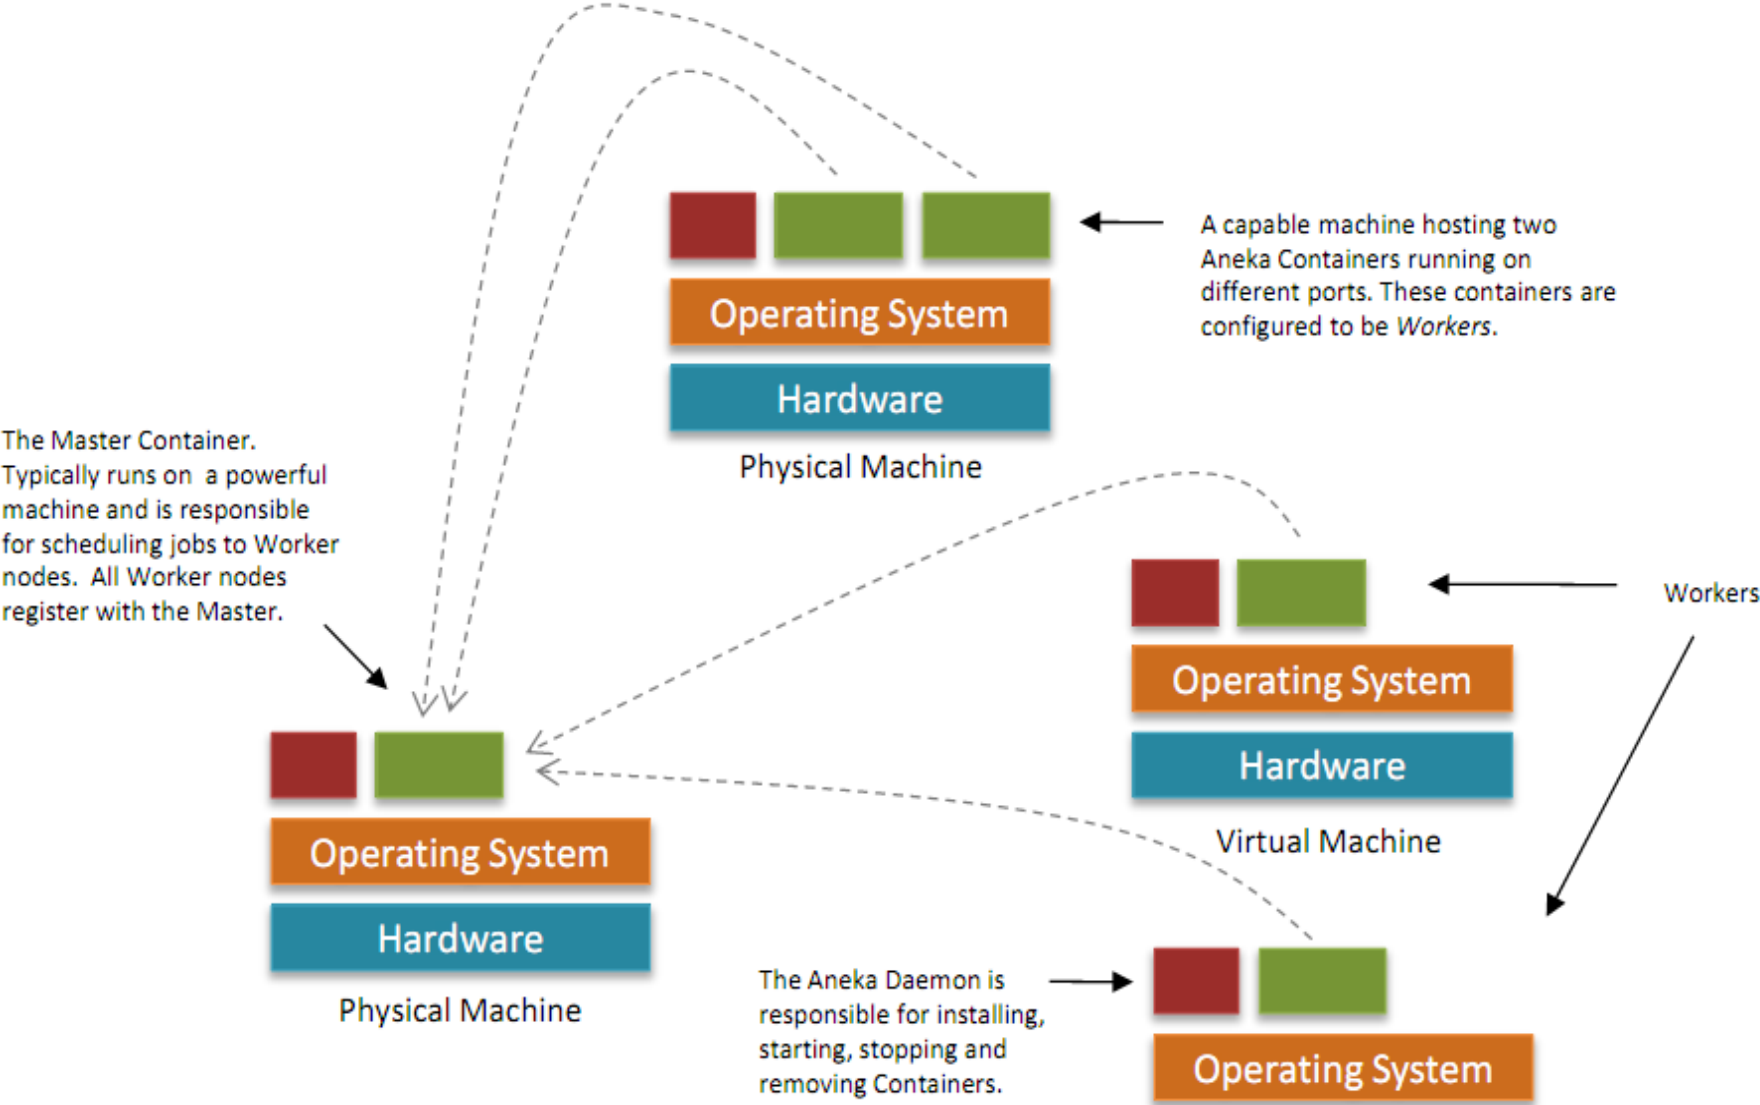
Hybrid cloud research challenges

- Resource management and scheduling in hybrid clouds
 - Metrics: QoS, SLA, Failure, profit
- Interoperability challenges
- Data replication and sharing
- Security and privacy
- Disaster recovery and high availability

Aneka – Tool for Cloud Computing research


- Aneka is a .net based cloud application development platform developed by Manjrasoft Pty Ltd Melbourne Australia (Spin-off company of University of Melbourne)
- It supports deployment of three types of application
 - High Throughput Computing (Task programming API)
 - High Performance Computing (Thread Programming API)
 - Data intensive computing (Map Reduce Programming API)
- Aneka can be deployed in heterogeneous infrastructure
 - Private Clouds, Hybrid Cloud, Local desktop clusters, Virtualization clusters
- Develop custom resource scheduling algorithms
- It has ability to connect to multicloud environments
 - AWS, Azure, openstack

High level view of Aneka



Aneka API's and its working

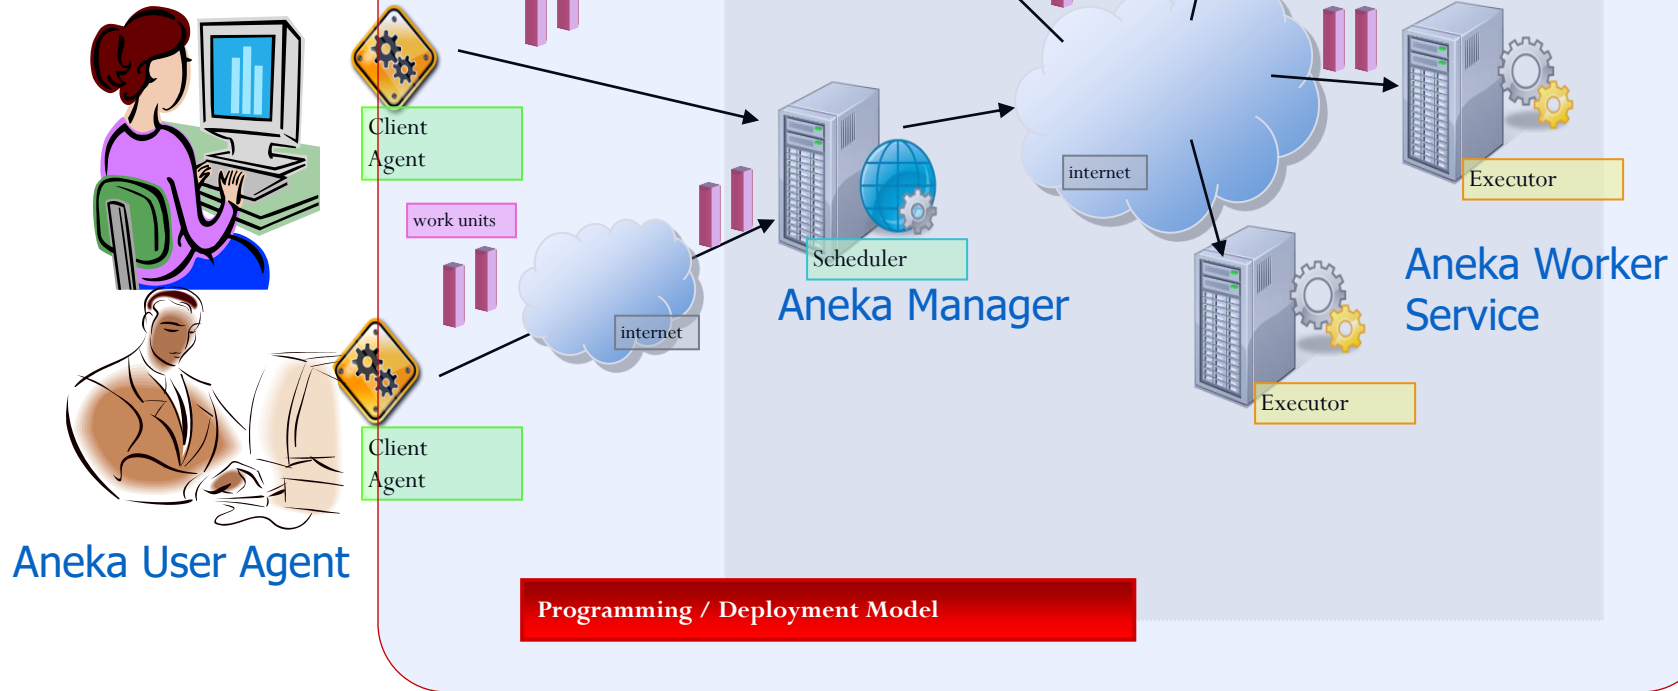
- Aneka Thread API
- Aneka Task API
- Aneka Map rReduce API



Aneka job submission

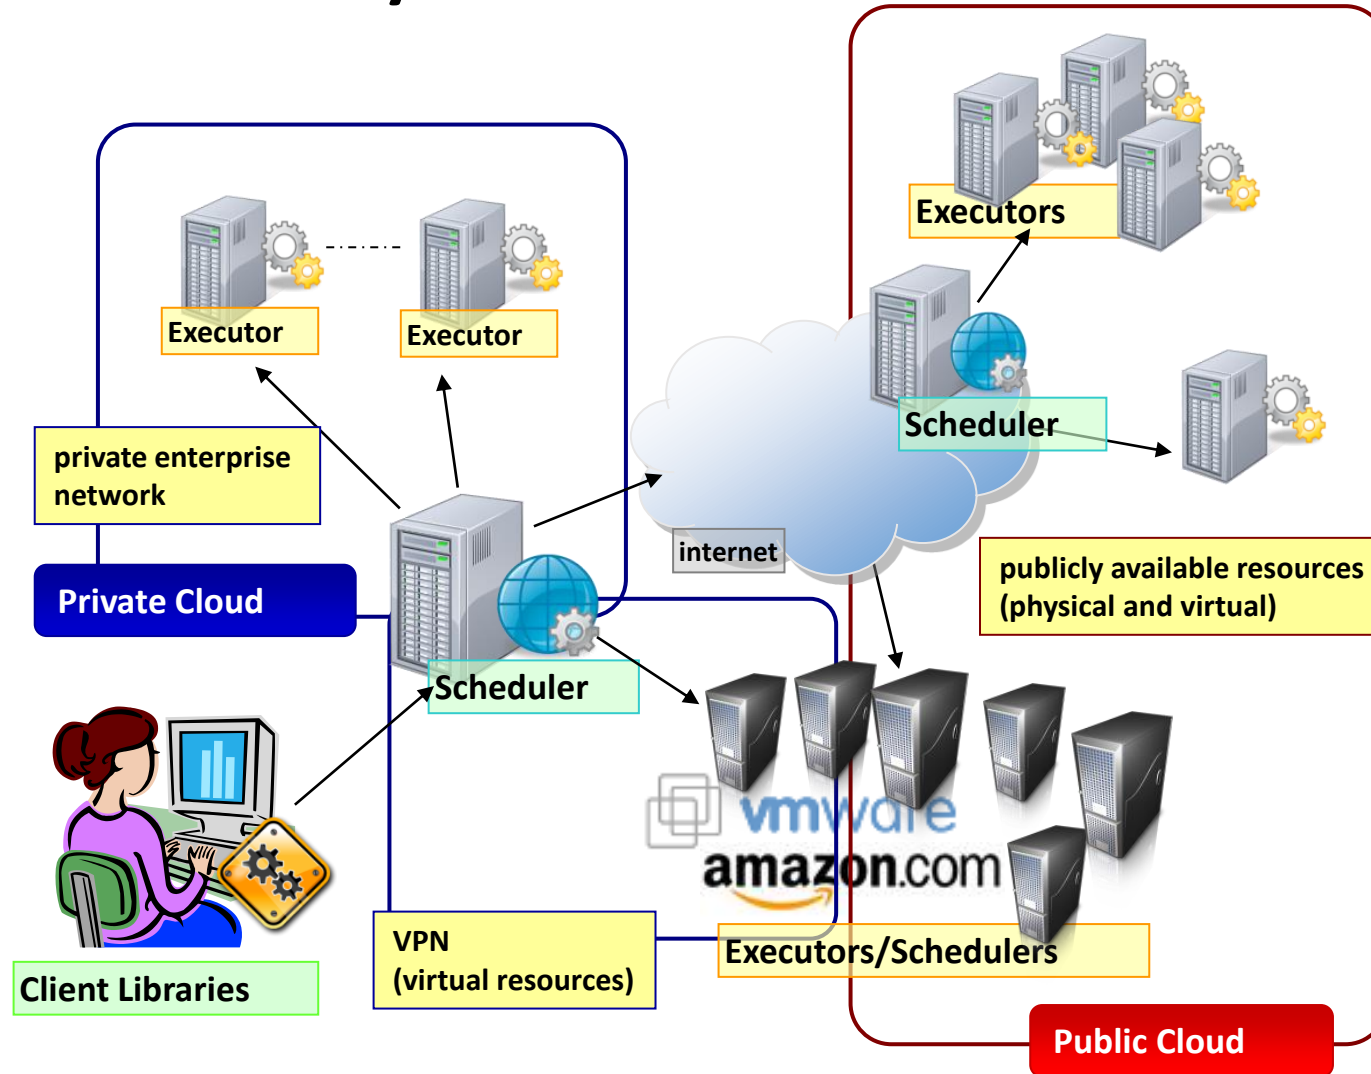
```
public DumbTask: ITask
{
    ...
    public void Execute()
    {
        ...
    }
}

for(int i=0; i<n; i++)
{
    ...
    DumbTask task = new DumbTask ();
    app.SubmitExecution(task);
}
```



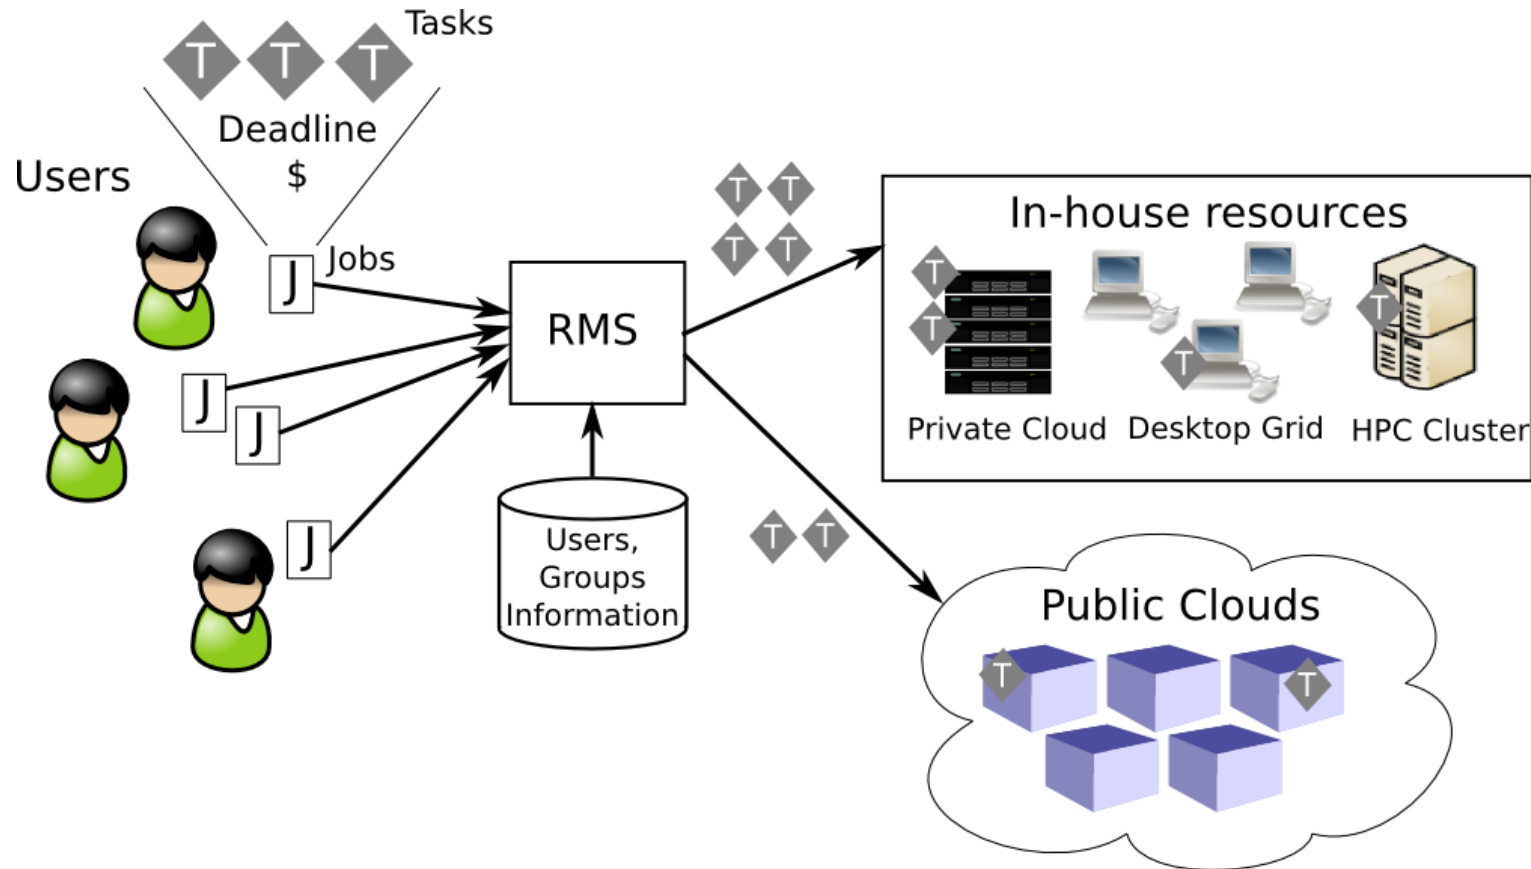
Aneka support for hybrid cloud

- XenServer Pool
 - Provisioning over private Cloud managed by Xen Server
- VMWare Pool
 - Provisioning over private Cloud managed by VMWare
- Amazon EC2 Pool
 - Provisioning over public Cloud provider: Amazon EC2

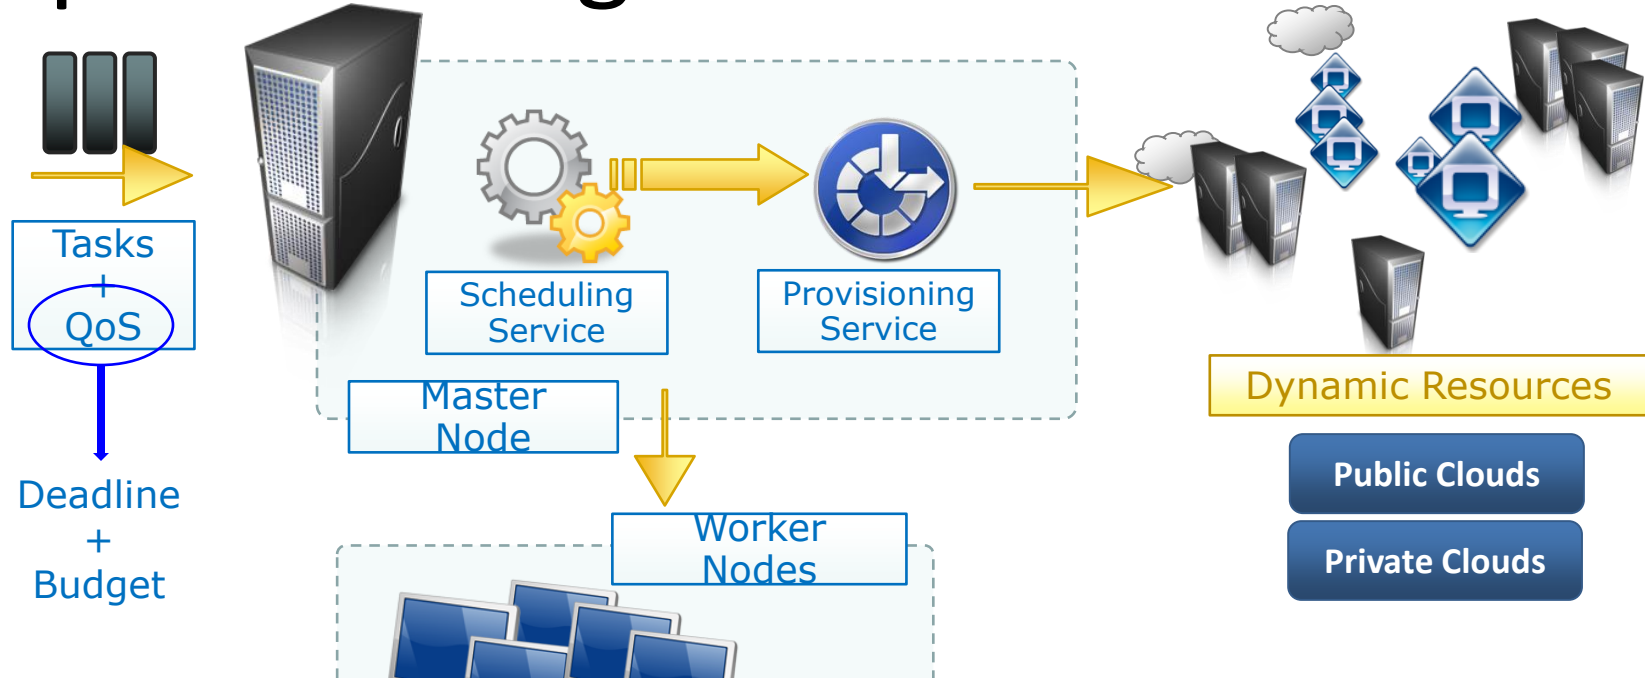


Hybrid Cloud Case study: Aneka's dynamic provisioning with Amazon EC2

Use Case : Scheduling problem based on cost and time



Hybrid Cloud Case study: Aneka's dynamic provisioning with Amazon EC2



AWS Spot Instances, Small

	Deadline	Execution time	Cost
Static	80 min	70 min	U\$ 0,06
	70 min	60 min	U\$ 0,12
	60 min	55 min	U\$ 0,15
	50 min	50 min	U\$ 0,15
	40 min	35 min	U\$ 0,30

Advantages of hybrid clouds

- Flexibility : The ability to distribute workloads across public and private environments based on security, efficiency, and cost.
- Scalability: Dynamically the resources based on the demand spike and release when not required.
- Reliability: Less possible downtime even when demand spikes.
- Security: Specific workload distribution of sensitive data on private clouds and non sensitive data on public clouds.
- Affordability: Need not to purchase and manage new resources to handle short spike in demand.

Advantages of hybrid clouds

- Continuity - Business continuity is not interrupted during a failure or disaster, as data is still accessible with little or no downtime.
- Opportunity: New prototypes or new applications can be tested and deployed rapidly.
- Accessibility: Any time, any where

Disadvantages of hybrid cloud

- Possibility of vendor lock in
- Service integration issues
- Interoperability issues
- Transparency
- Security and privacy issues
- Data locality issues

Benefits of hybrid clouds

- Flexibility and agility
- Elasticity
- Self-service
- Faster delivery of new products and services
- Cost control
- Avoidance of lock-in
- Access to the latest technology

Multi Clouds

- In multi cloud environment, enterprises will utilize the services from multiple cloud providers and private cloud could be an optional
- A hybrid cloud becomes, multi-cloud when there are more than one public cloud service combined with private cloud resources.

Elements	Hybrid Cloud	Multi-Cloud
Private + public clouds	Yes, always	Optional
Multiple public clouds	Optional	Yes, always

- Gartner survey says, 81% enterprises reported working with multi clouds
- Higher Pros as compared to hybrid clouds in service delivery, access and deployment

Practical session

- In the practical session your going to learn the following tools for hybrid cloud solution
 - ManageIQ
 - Openshift

References

- <https://azure.microsoft.com/en-in/overview/azure-stack/>
- <https://aws.amazon.com/hybrid/>
- <https://www.vmware.com/topics/glossary/content/hybrid-cloud>
- <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-146.pdf>
- www.buyya.com
- <https://www.harbourit.com.au/blog/everything-you-need-to-know-about-hybrid-cloud/>
- Buyya, R., Srirama, S. N., Casale, G., Calheiros, R., Simmhan, Y., Varghese, B., ... & Toosi, A. N. (2018). A manifesto for future generation cloud computing: Research directions for the next decade. *ACM computing surveys (CSUR)*, 51(5), 1-38.