Cloud Computing

Orran Krieger

Senior Staff Engineer, VMware, Inc.

Lead Kvetcher VMware Cloud team

Adjunct Professor CMU

vmware

Outline

- **Background on VMware**
- **Cloud computing today**
- **What should it be**
- **The vCloud initiative**
- **Research ideas**

Evolution of Virtualization

PHYSICAL

vmware

Evolution of Virtualization

PHYSICAL **VIRTUALIZED**

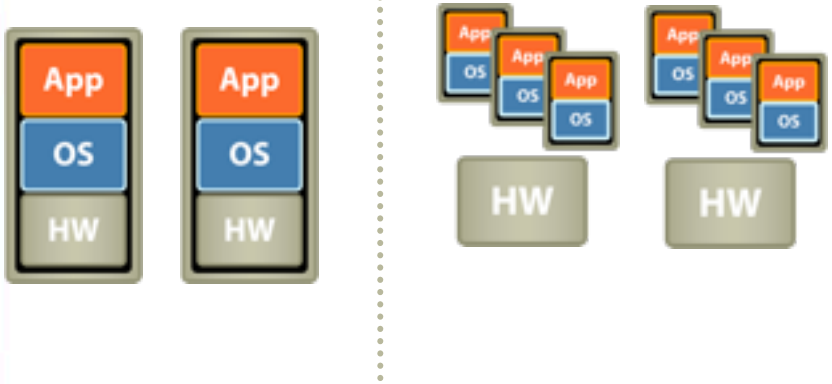

vmware

Evolution of Virtualization

PHYSICAL **VIRTUALIZED** **POOLED**

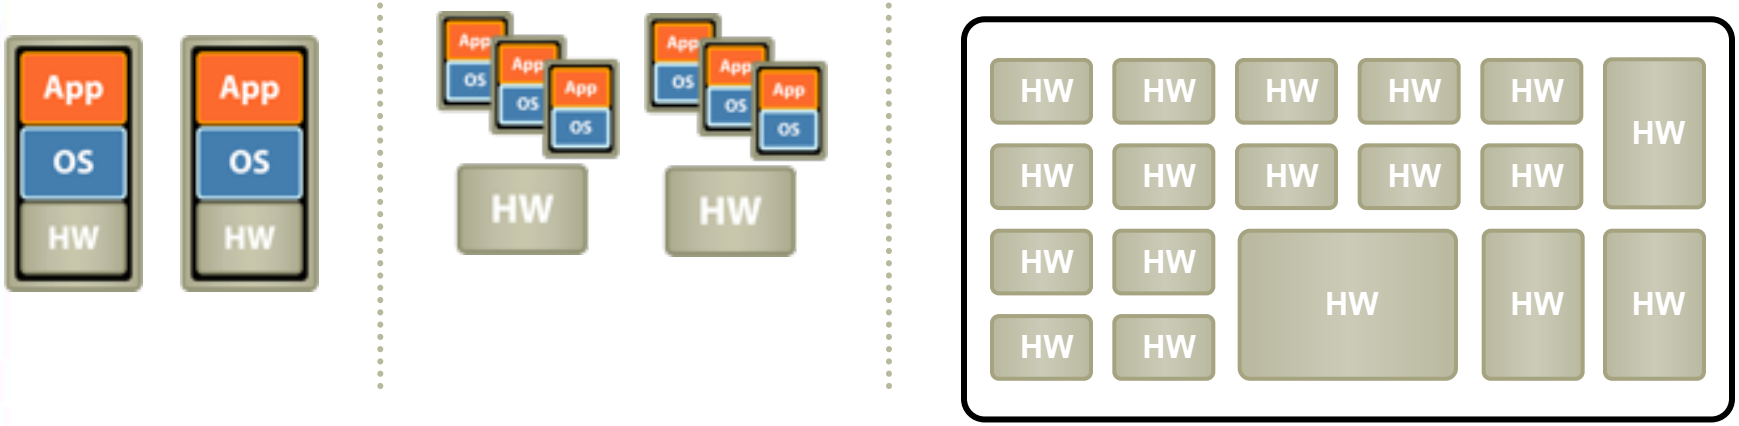

> Logical Resource Pooling (RP)
> Distributed Resource Scheduler (DRS)

vmware

Evolution of Virtualization

POOLED

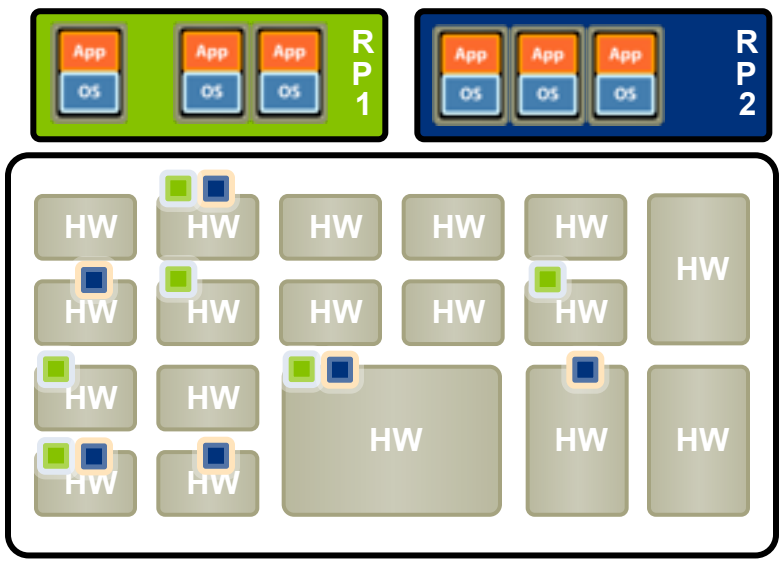

> Logical Resource Pooling (RP)
> Distributed Resource Scheduler (DRS)

vmware

Evolution of Virtualization

POOLED

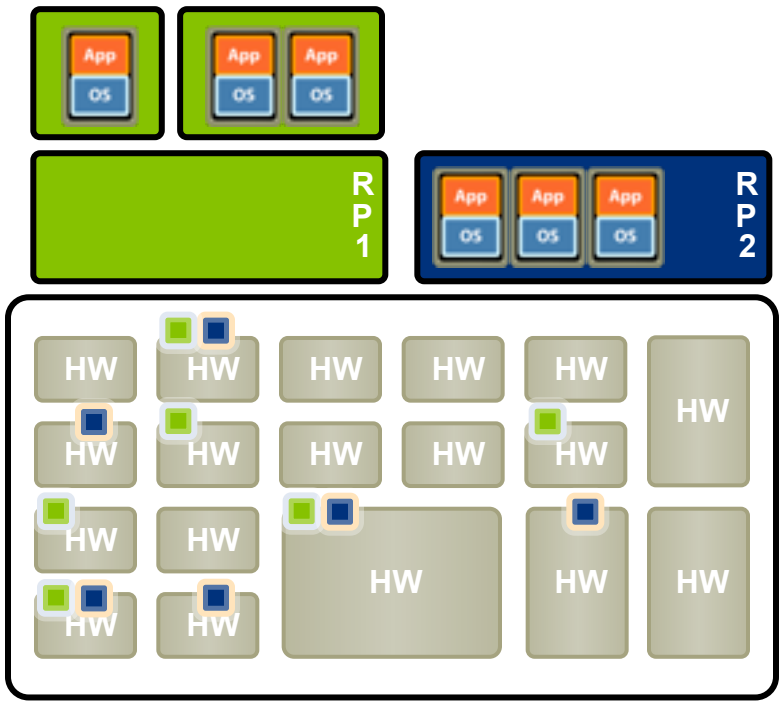

> Logical Resource Pooling (RP)
> Distributed Resource Scheduler (DRS)

vmware

Power Savings with DPM

- 5 hour VMmark run
- 110 VMs, 4 server cluster

vmware

Power Savings with DPM

- 5 hour VMmark run
- 110 VMs, 4 server cluster

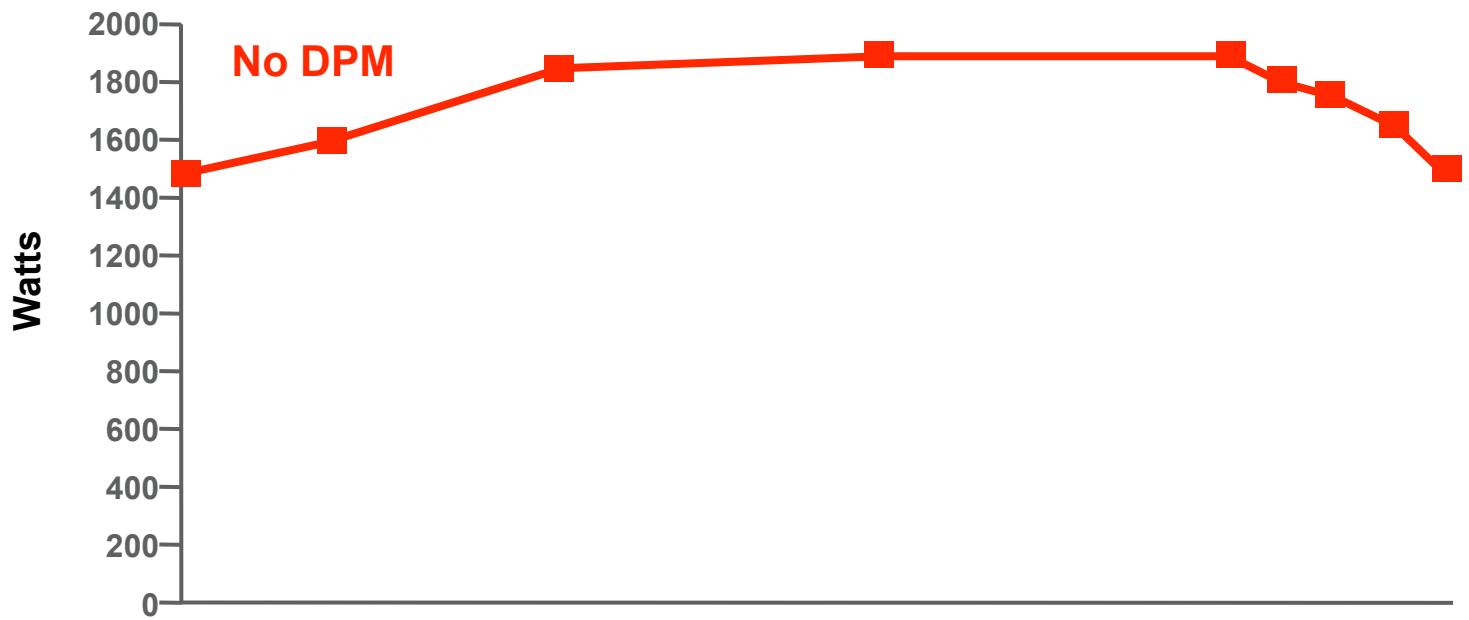

vmware

Power Savings with DPM

- 5 hour VMmark run
- 110 VMs, 4 server cluster

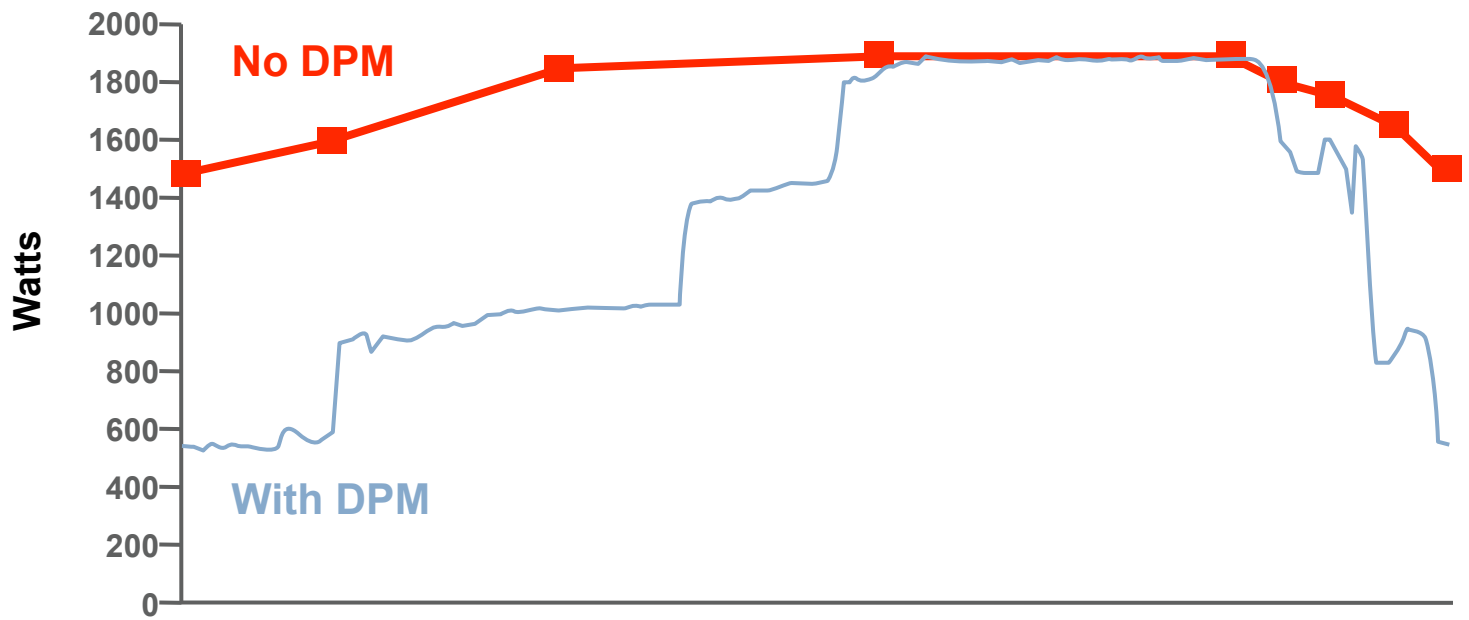

50% Savings during VMmark

vmware

Distributed Virtual Switch

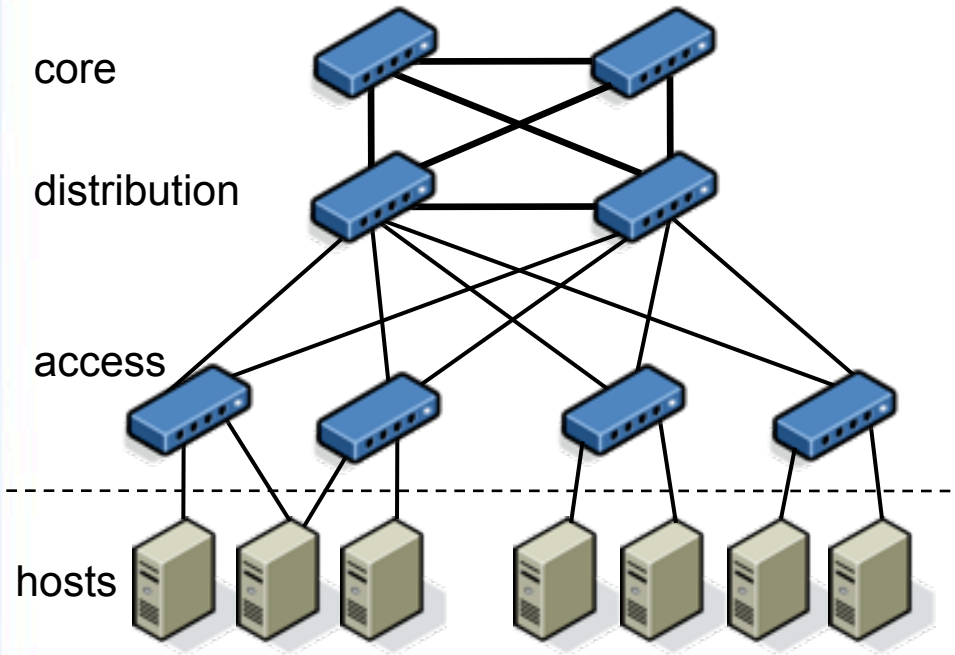

vmware

Distributed Virtual Switch

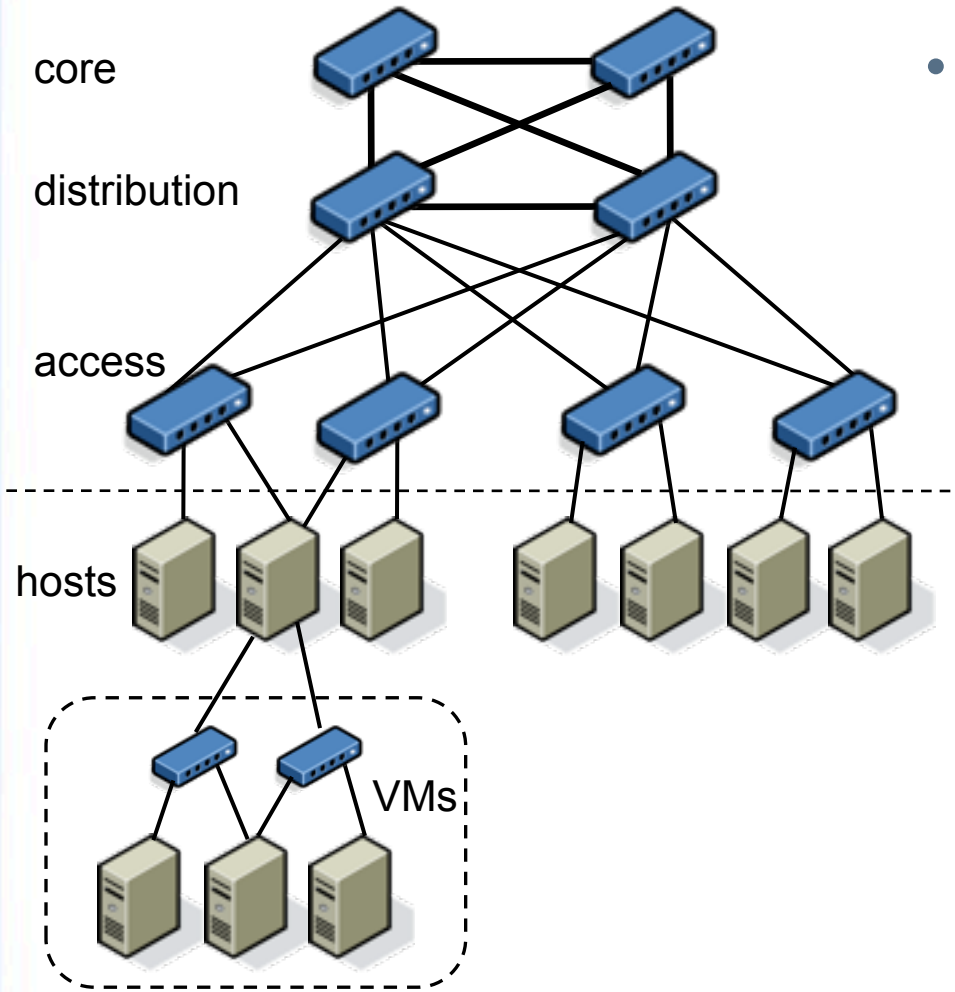

- **Network-centric view: virtualization blurs the host/network boundary**

vmware

Distributed Virtual Switch

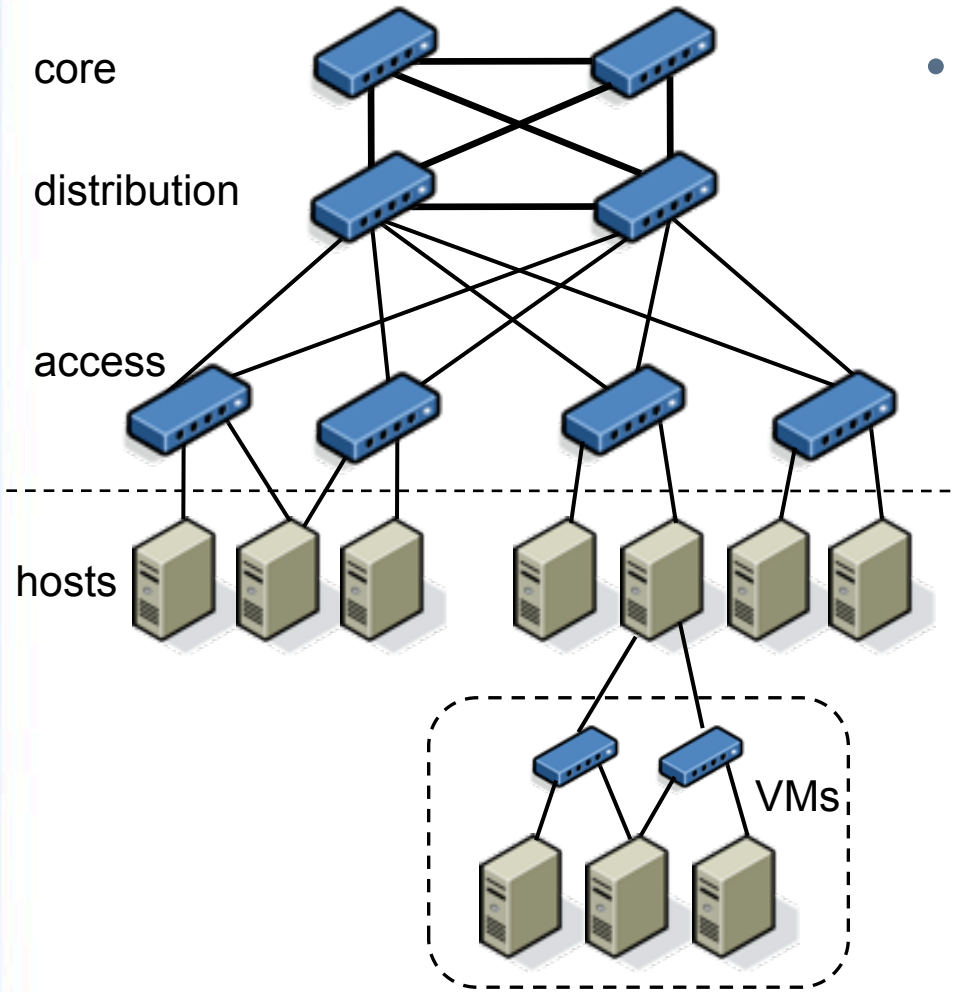

- **Network-centric view: virtualization blurs the host/network boundary**
  - Virtual network switches and topologies within a host
  - VMs can move anywhere
  - The virtual network switch is a new access layer.

vmware

Distributed Virtual Switch

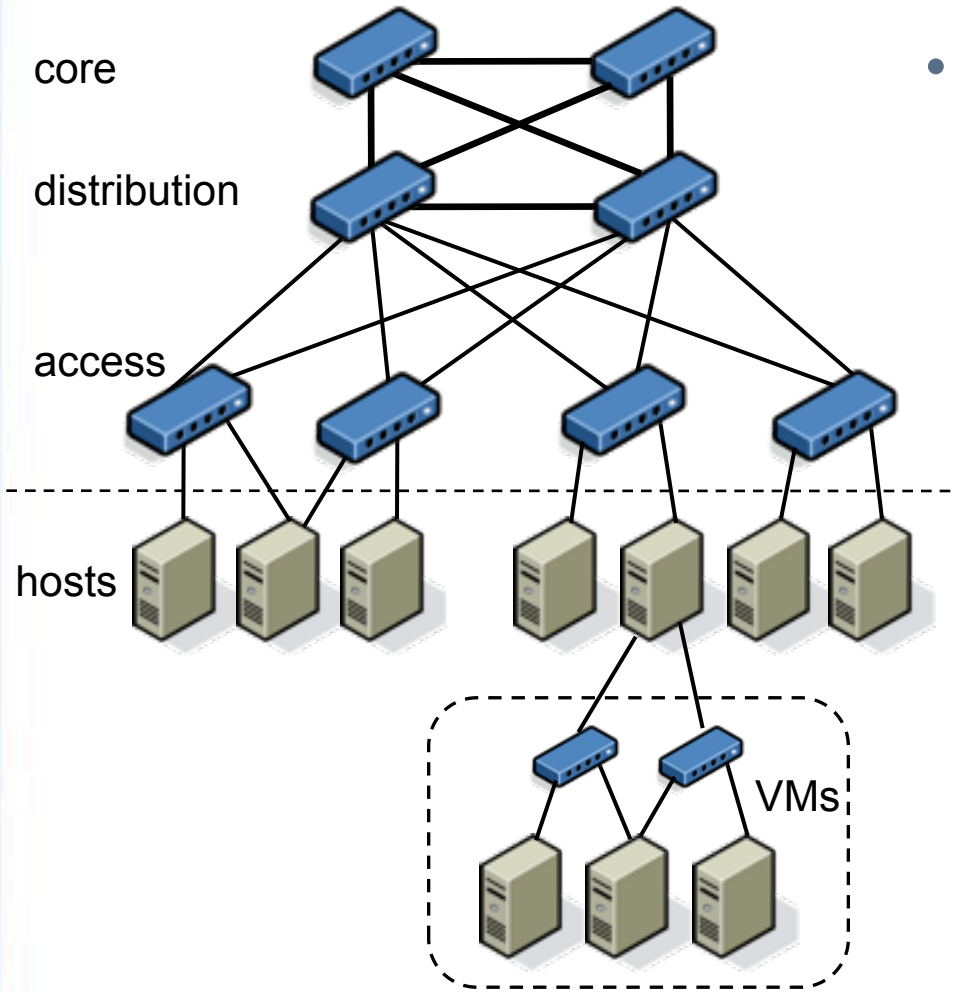

- **Network-centric view: virtualization blurs the host/network boundary**
  - Virtual network switches and topologies within a host
  - VMs can move anywhere
  - The virtual network switch is a new access layer.
  - Network administrators can no longer rely on physical port-based access control

vmware

Distributed Virtual Switch

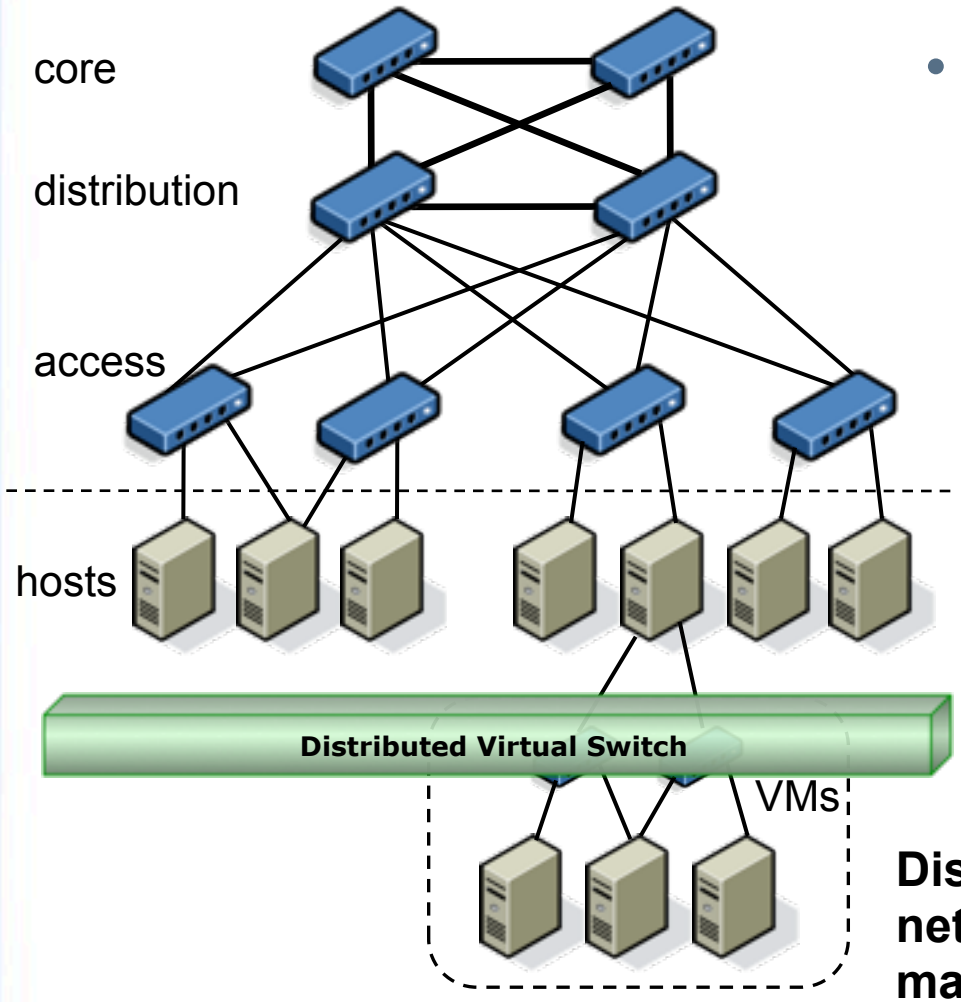

- **Network-centric view: virtualization blurs the host/network boundary**
  - Virtual network switches and topologies within a host
  - VMs can move anywhere
  - The virtual network switch is a new access layer.
  - Network administrators can no longer rely on physical port-based access control

Distributed virtual switch to extend network access control and management to virtual environments

vmware

Availability

Hot-plug resources

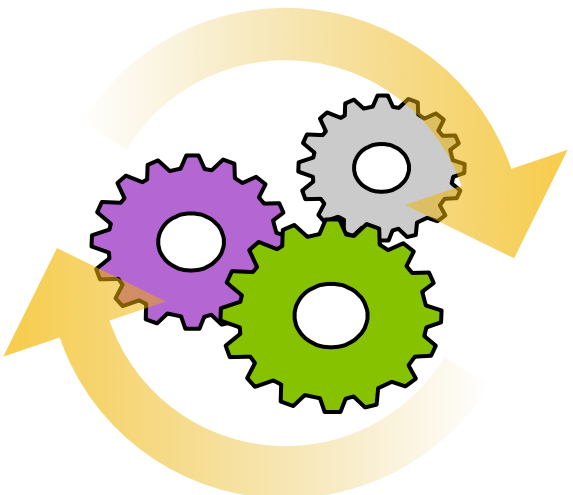

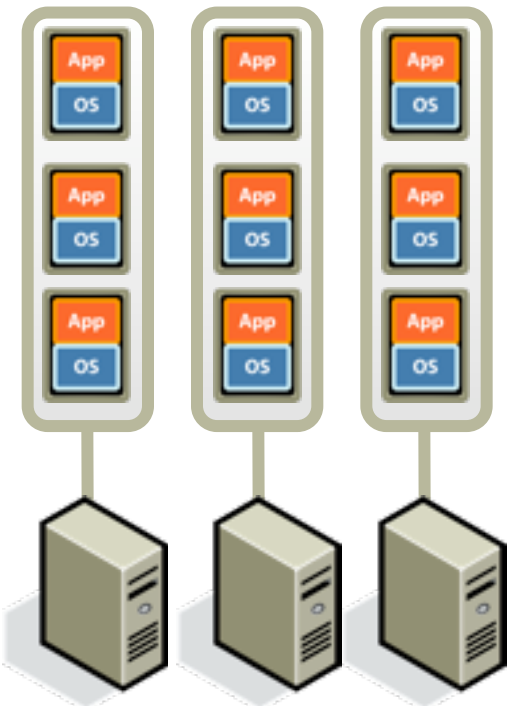

vmware

Availability

Hot-plug resources

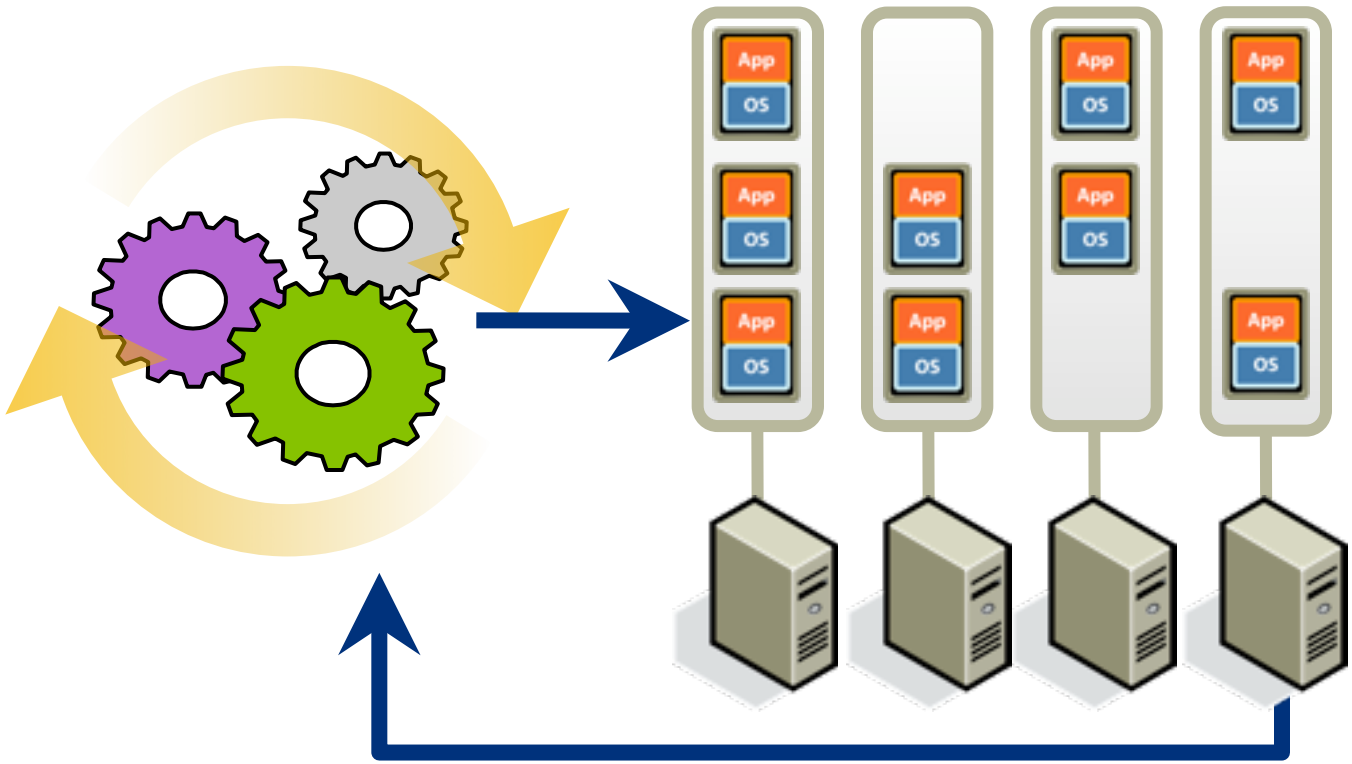

vmware

Availability

Hot-plug resources
> Add/remove capacity on demand
> Improve application availability

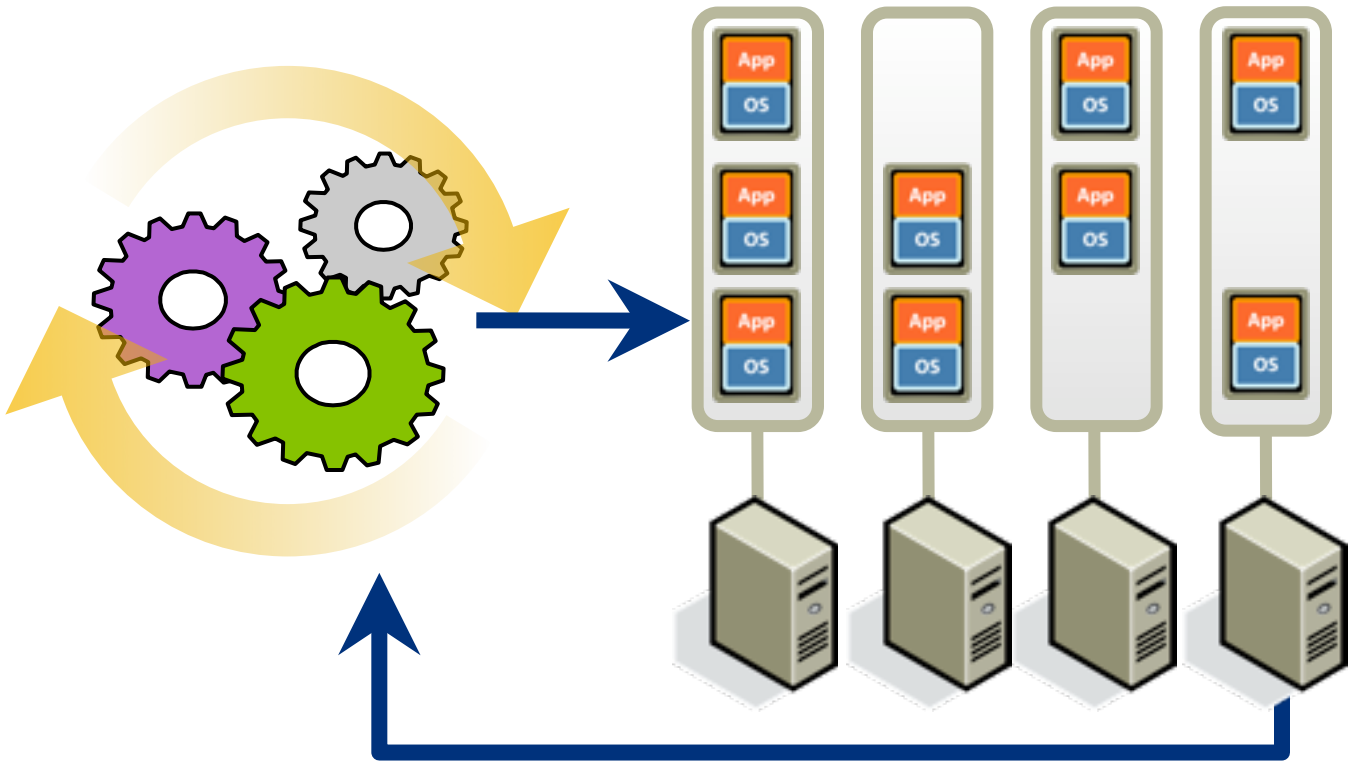

vmware

Availability

Hot-plug resources
> Add/remove capacity on demand
> Improve application availability

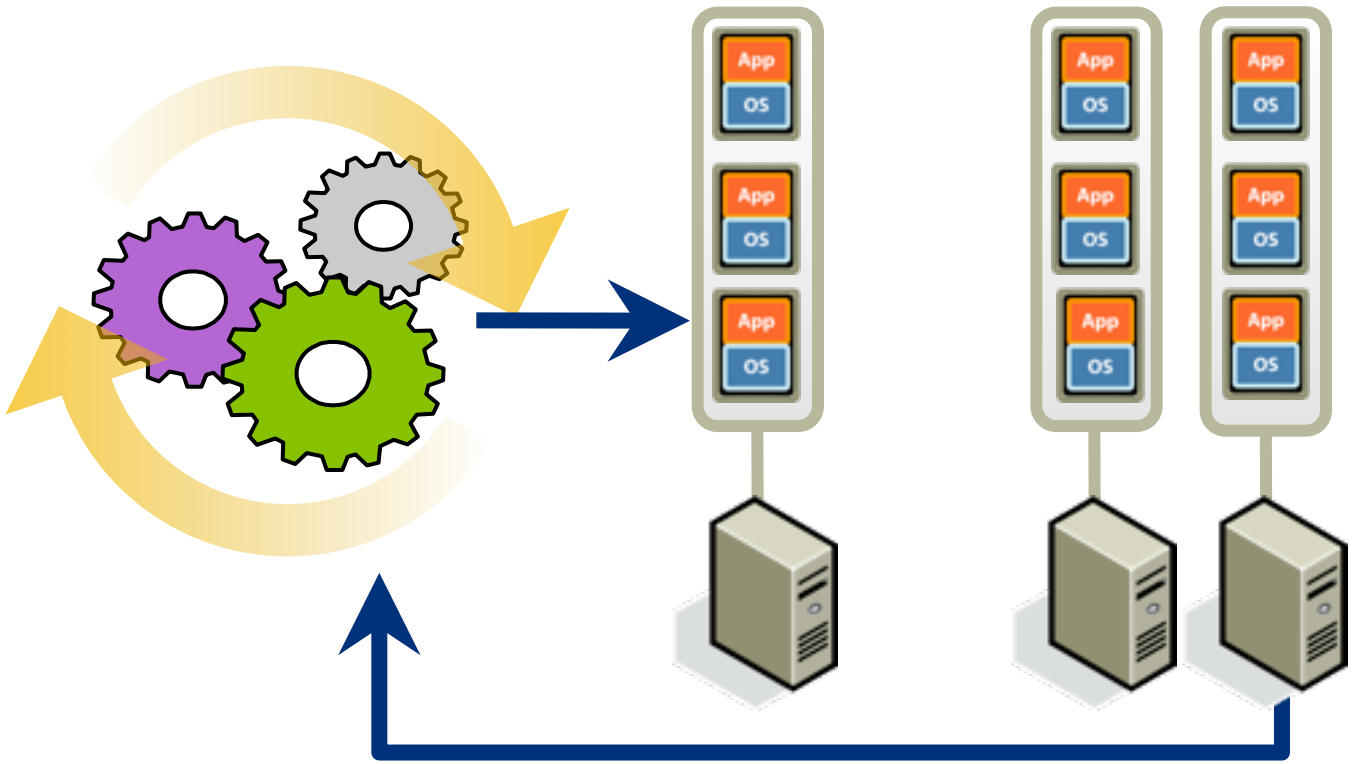

vmware

Site Recovery Manager

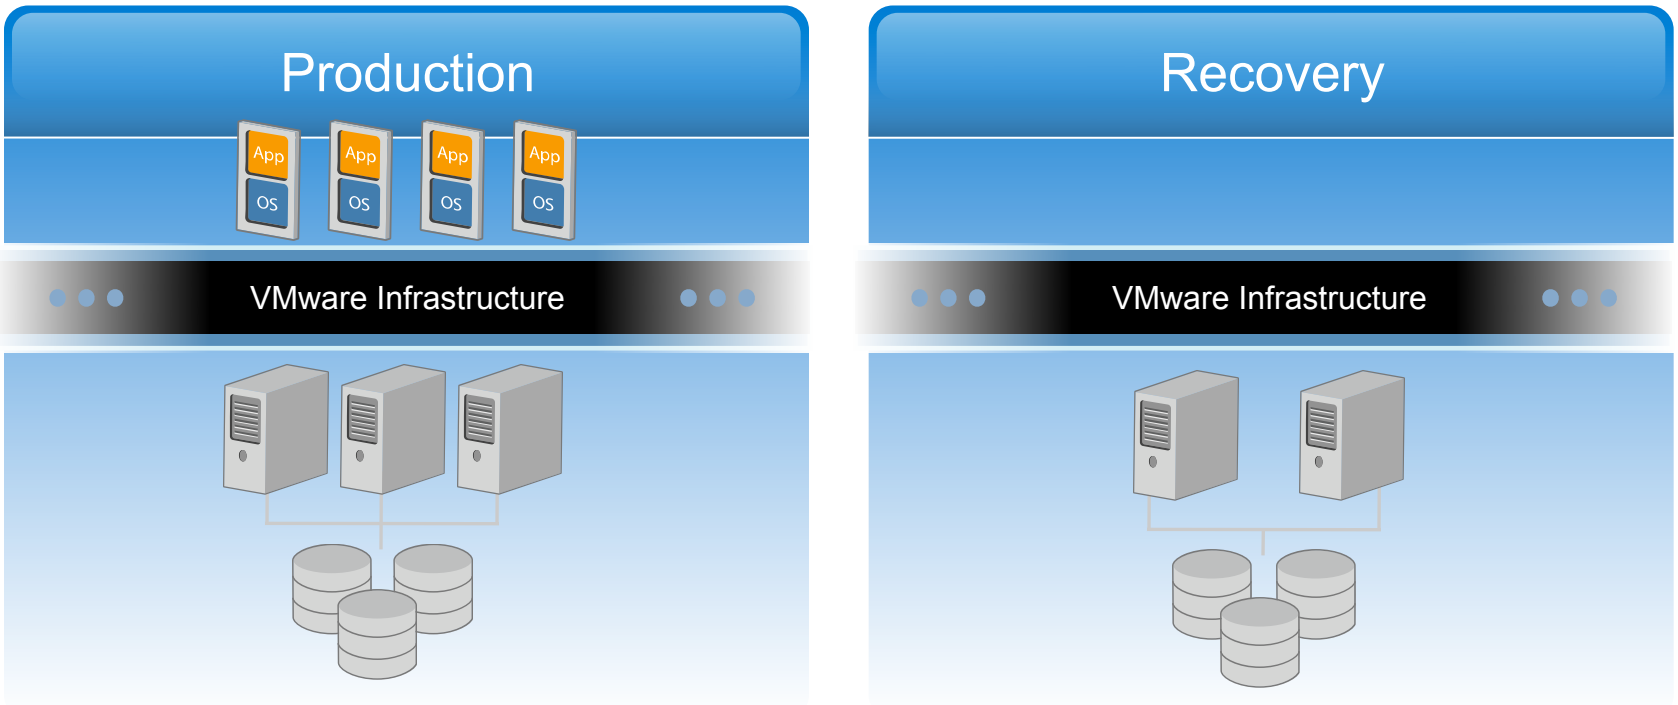

Site Recovery Manager transforms disaster recovery

vmware

Site Recovery Manager

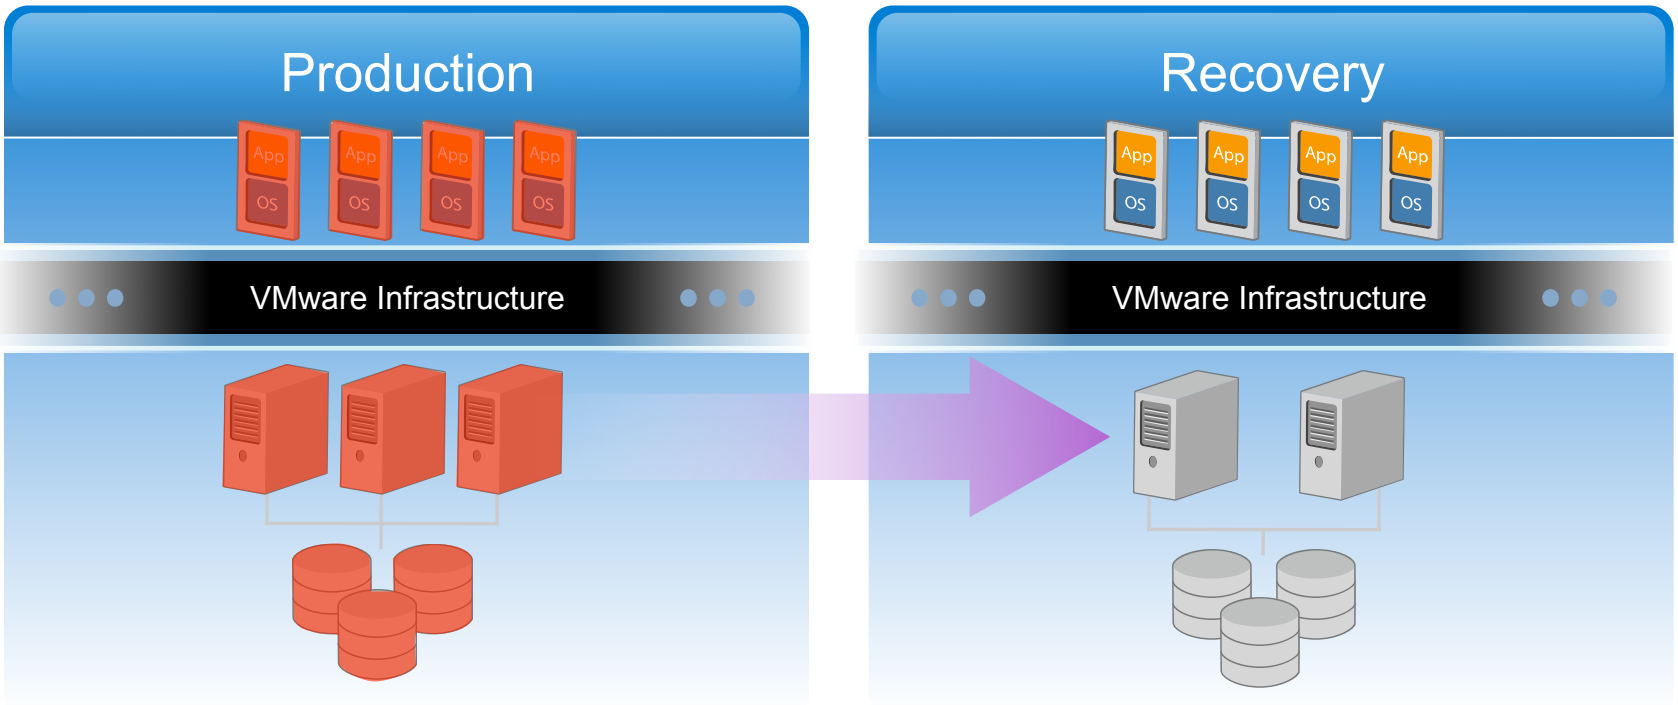

Site Recovery Manager transforms disaster recovery

vmware

Security: VMware VMsafe

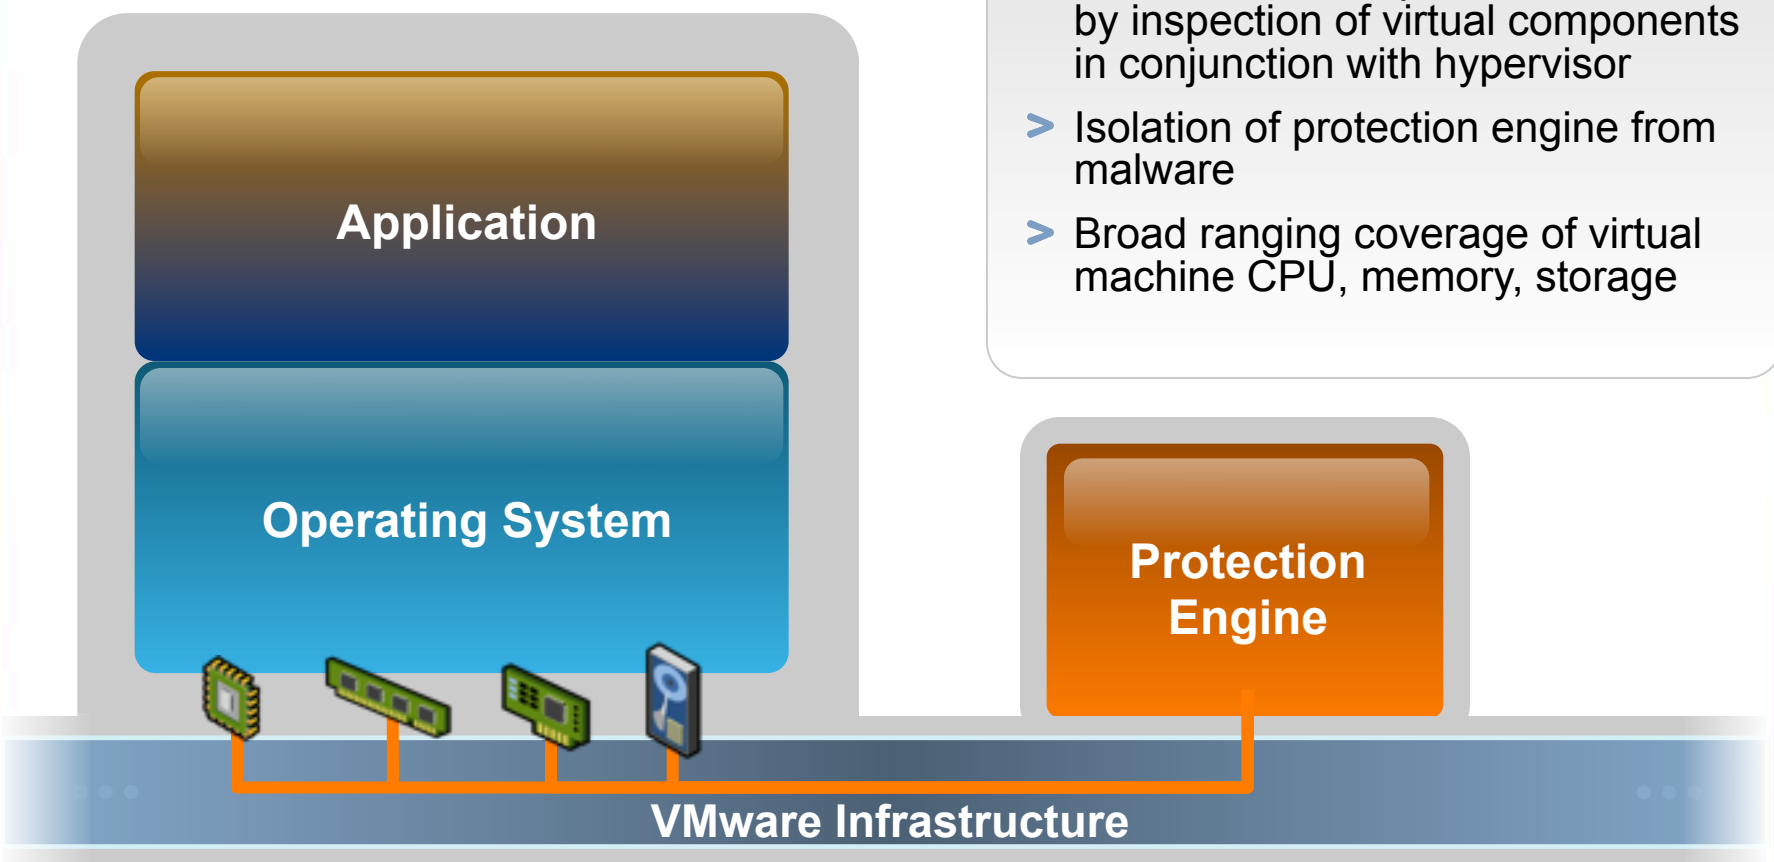

- API that enables protection of VMs by inspection of virtual components in conjunction with hypervisor
- Isolation of protection engine from malware
- Broad ranging coverage of virtual machine CPU, memory, storage

vmware

vApp: A new type of Software

- **Properties**
  - Policy-based operations
  - Multi-tier
  - Distributed as an OVF package
- **Built by:**
  - ISVs / Virtual Appliance Vendors
  - By internal IT shops
  - By IT administrators
  - SI/VARs

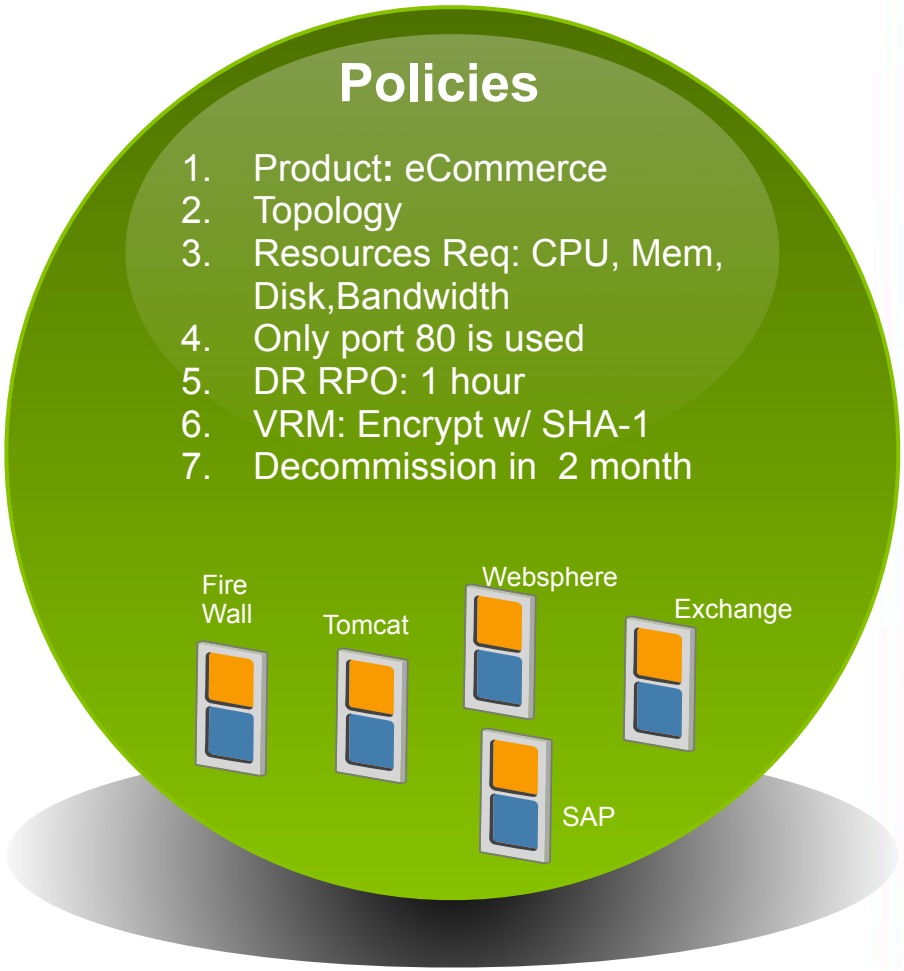

vmware

Desktops

Virtual Desktop Infrastructure (VDI)

Centralized
Virtual Desktops

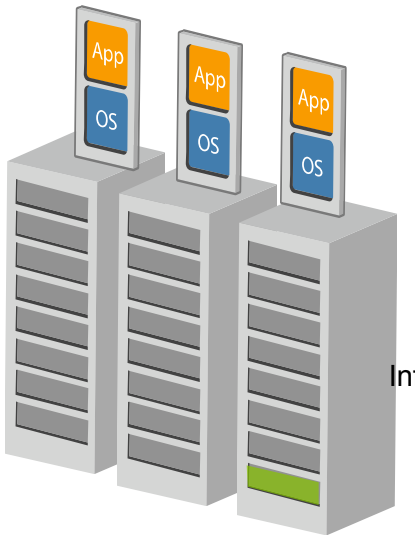

VMware
Infrastructure 3

vmware

Desktops

Virtual Desktop Infrastructure (VDI)

- Full desktops run as VMs in the data center
- Connect securely to desktop from anywhere
- Centralized management
- More efficient resource usage
- Higher availability
- Online and offline mode

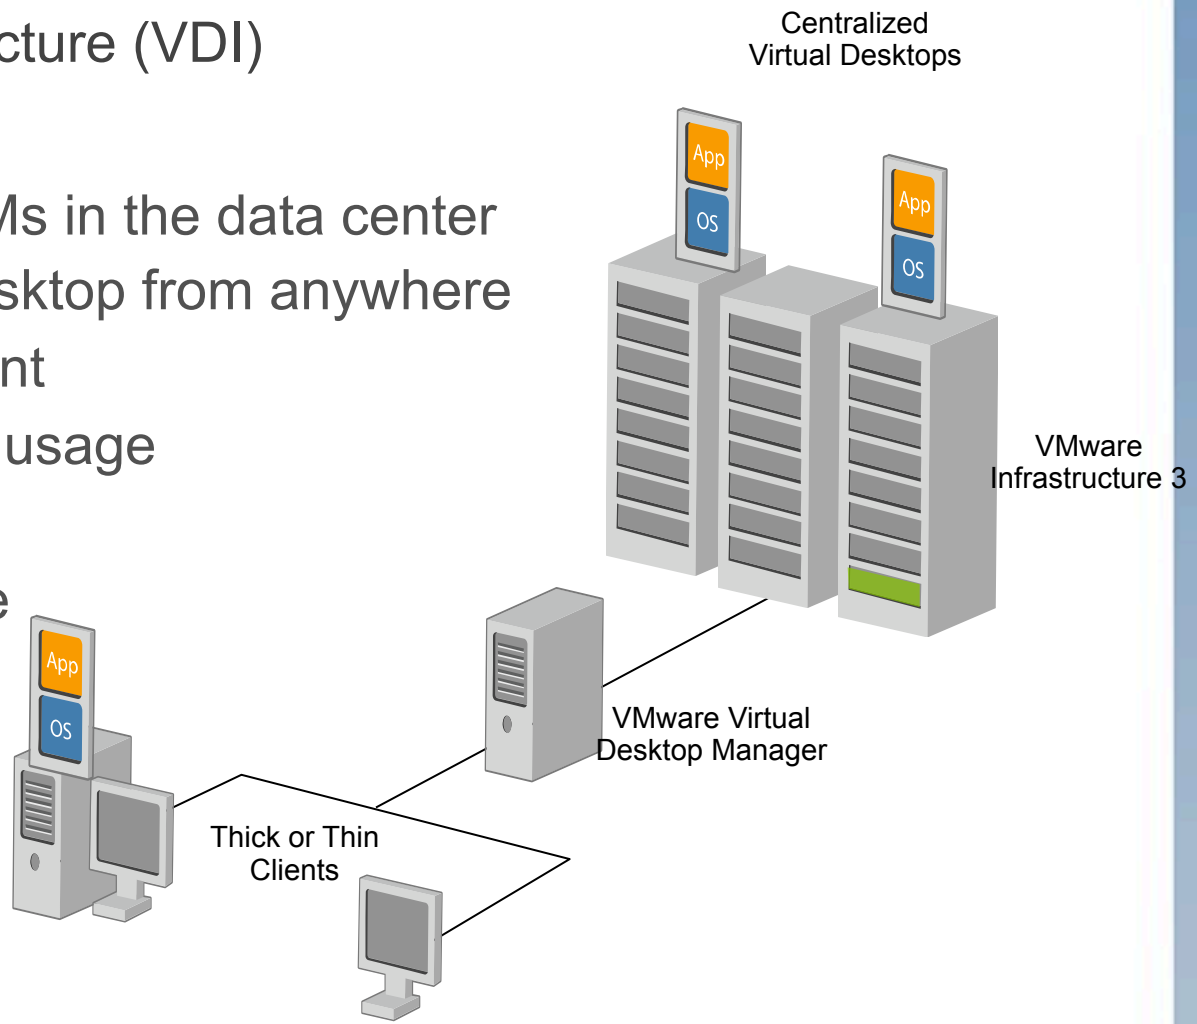

vmware

Evolution of Virtualization

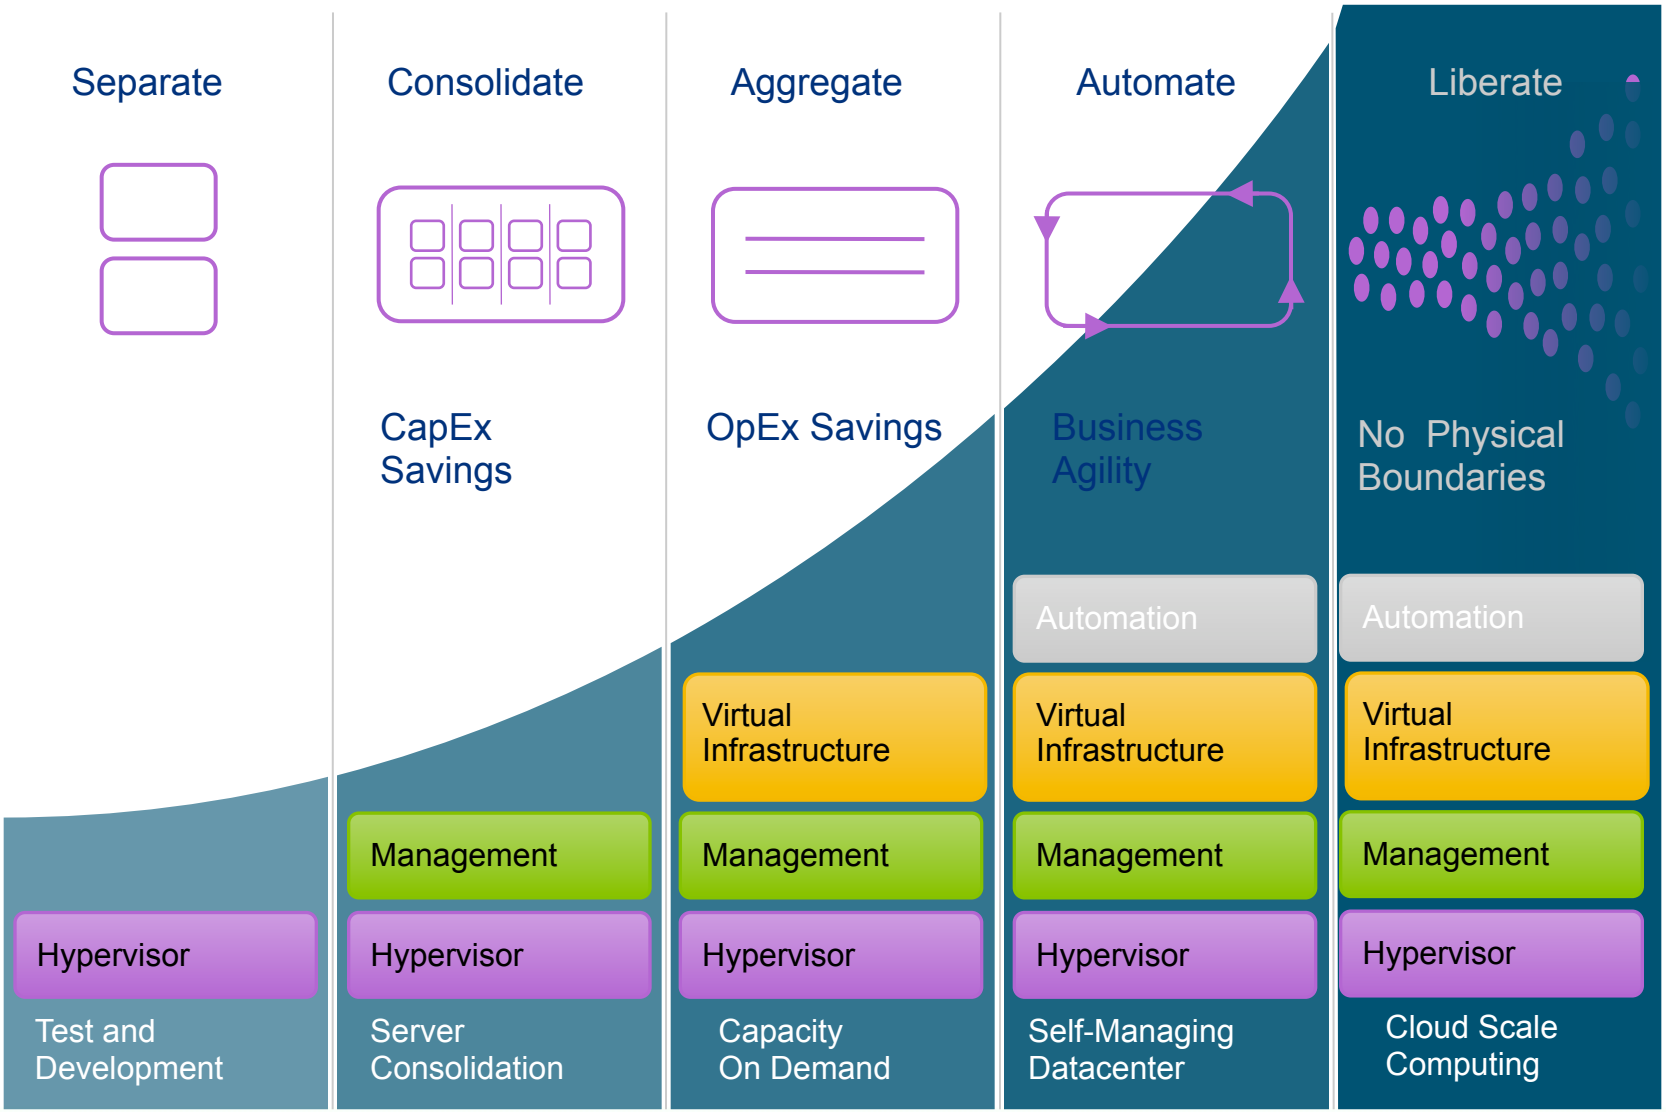

vmware

The cloud started with SaaS/Web

Traditional software model

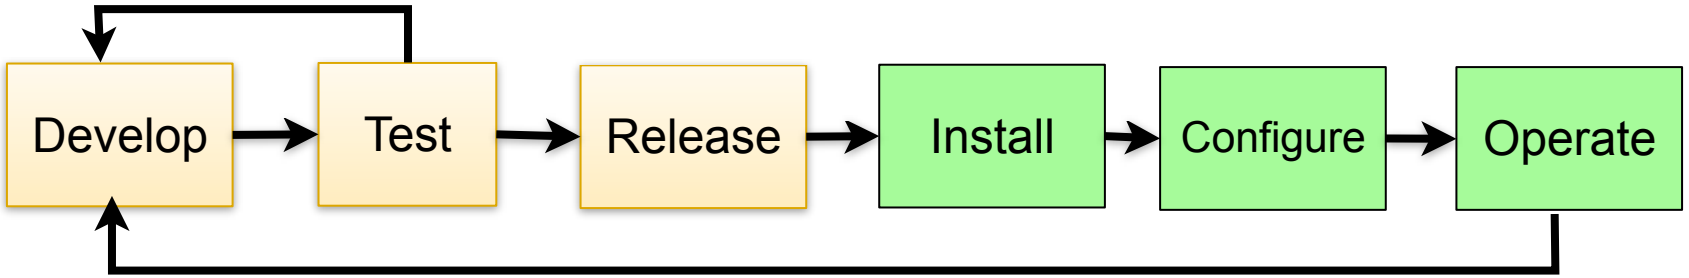

SaaS model

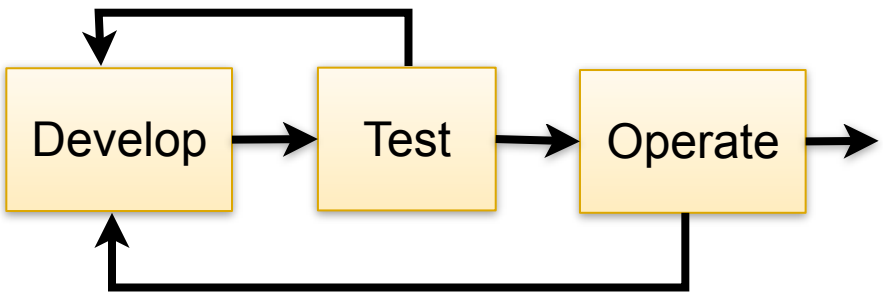

vmware

The reality: success is very hard

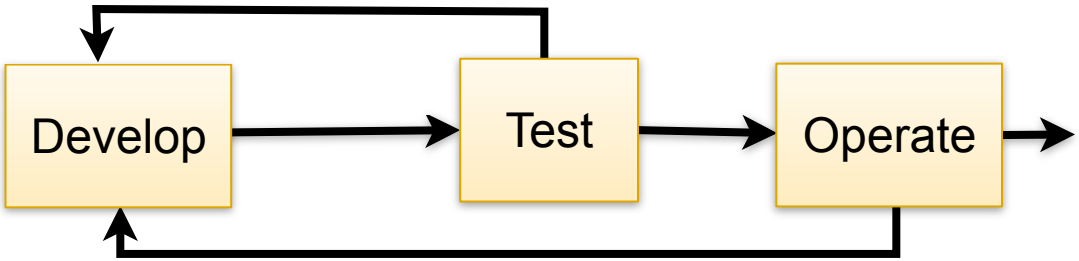

vmware

The reality: success is very hard

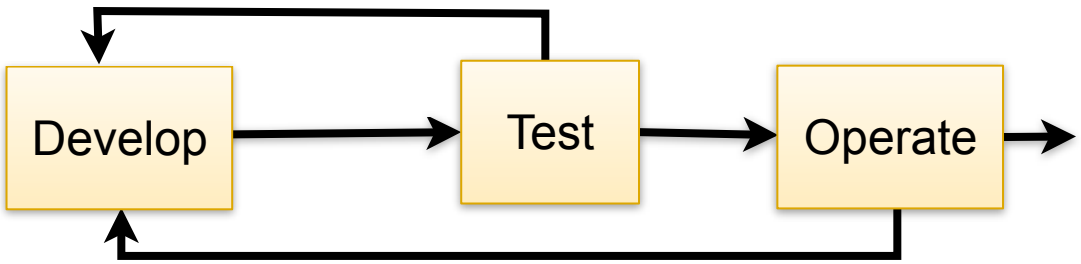

- **Most SW companies don't have the in-house skill to operate at scale:**
  - e.g. loadbalancer, network config, security, disaster recovery, …& 70% of investment spent not enhancing application

vmware

The reality: success is very hard

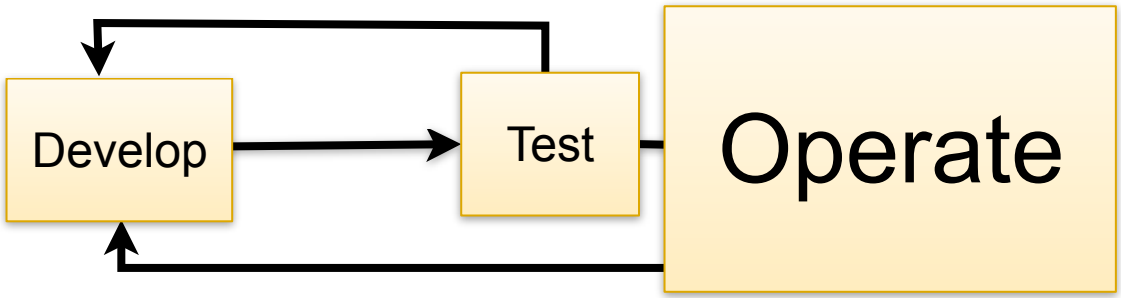

- **Most SW companies don't have the in-house skill to operate at scale:**
  - e.g. loadbalancer, network config, security, disaster recovery, …& 70% of investment spent not enhancing application

The reality: success is very hard

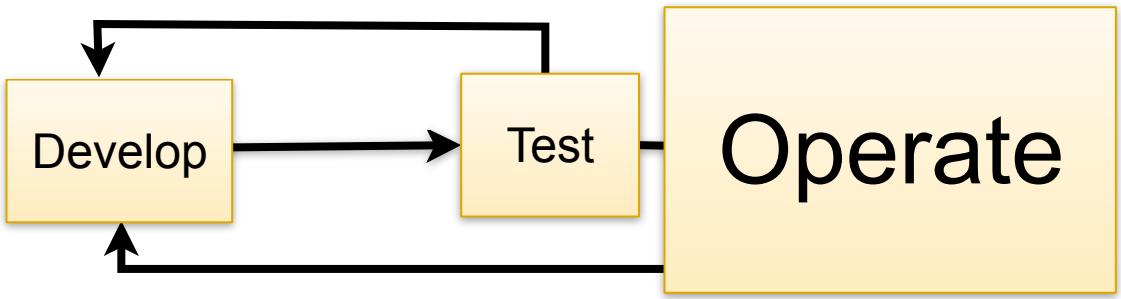

- **Most SW companies don't have the in-house skill to operate at scale:**
  - e.g. loadbalancer, network config, security, disaster recovery, …& 70% of investment spent not enhancing application
- **Enormous investment required in application level to scale.**

The reality: success is very hard

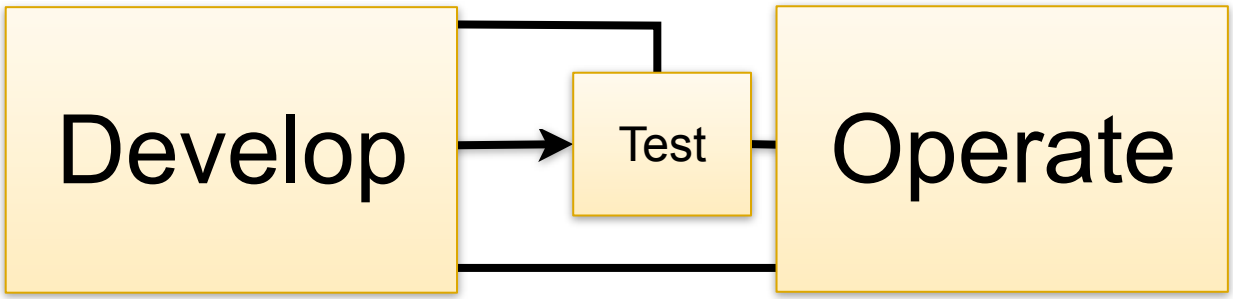

- **Most SW companies don't have the in-house skill to operate at scale:**
  - e.g. loadbalancer, network config, security, disaster recovery, …& 70% of investment spent not enhancing application
- **Enormous investment required in application level to scale.**

The reality: success is very hard

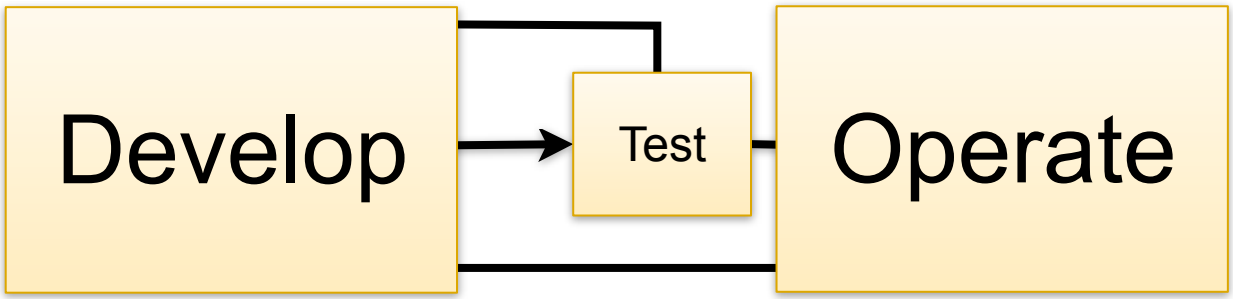

- **Most SW companies don't have the in-house skill to operate at scale:**
  - e.g. loadbalancer, network config, security, disaster recovery, …& 70% of investment spent not enhancing application
- **Enormous investment required in application level to scale.**
- **So, successful SaaS vendors started building re-usable platforms…**

Cloud offerings

- **Software as a Service (SaaS):**
  - Application hosted in the cloud.
  - e.g., gmail, google apps, salesforce.com, wikipedia…
- **Platform as a Service (PaaS):**
  - Application development environment and runtime hosted in cloud.
  - e.g., engine yard, Google App Engine, Force.com
- **Infrastructure as a Service (IaaS):**
  - Developers/administrators obtain general computing, storage, messaging services…
  - e.g., terramark, Amazon AWS, Mosso

Example Amazon IaaS

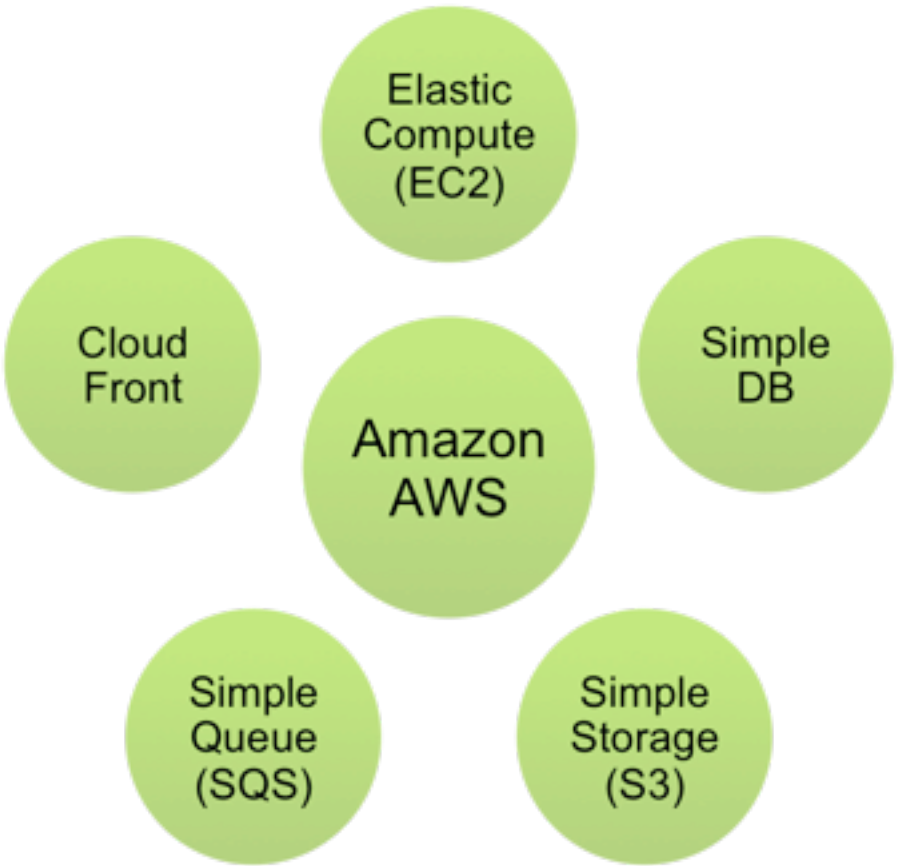

vmware

Example

- **Animoto April 2008: Peak EC2 instances:**
  - Mon 50, Tues 400, Wed 900, Friday 3400

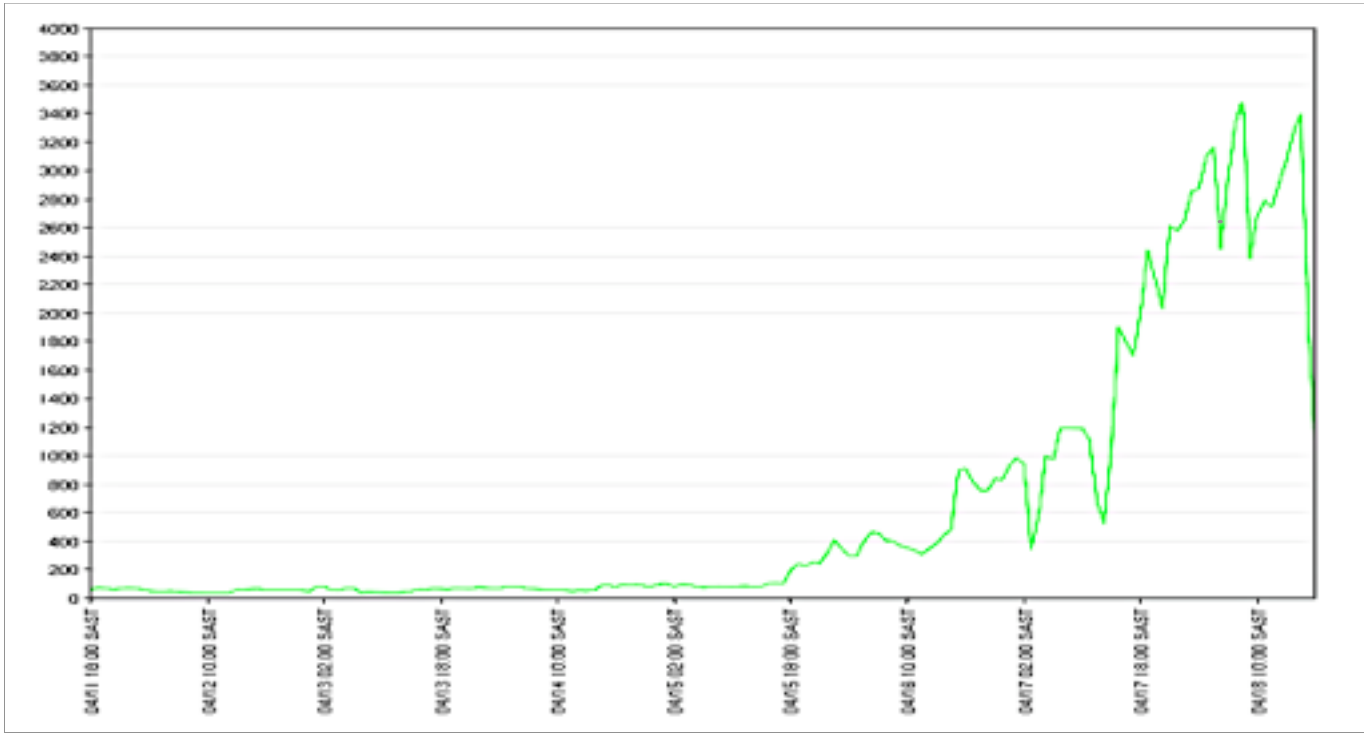

vmware

The problems with this evolution

- **Vertical offered by a single vendor**
  - Just a few hundred/thousand developers enhancing offering
  - No on-premise offering for enterprise/university/SaaS vendor
  - Who wants to trust a single company?
- **Limited largely to web applications:**
  - Limited support legacy, HPC, hosted client, grid,
- **Disintermediates OEMs, infrastructure vendors, traditional management stacks… (VARs for SaaS)**
- **Ignores the key value in the virtualization needed for general purpose workloads: over provisioning, SRM, DRS, DPM, OVF, SVI …**

Disclaimer…

Disclaimer…

- **At this point, there is the mandatory… okay there is a lot of hype here, but…**

Disclaimer…

- **At this point, there is the mandatory… okay there is a lot of hype but…**

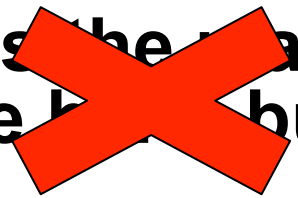

Disclaimer…

- **At this point, there is the mandatory… okay there is a lot of hype but…**

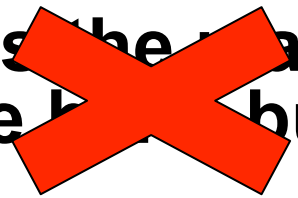

Disclaimer…

- **At this point, there is the mandatory… okay there is a lot of hype but…**

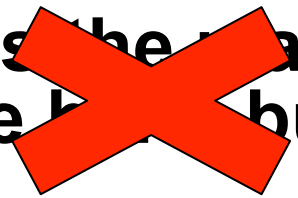

- **The transformation is more profound that is yet understood, cloud done right will:**
  - be used for all applications,
  - change how we deploy and develop applications,
  - enable new OSes, new programming models, new servers, new storage solutions
  - enable new markets for computer services
  - …, in other words, its gonna change everything

What do we really want

What do we really want

Original vision of Utility/grid computing:

What do we really want

Original vision of Utility/grid computing:

"If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility... The computer utility could become the basis of a new and important industry."

What do we really want

Original vision of Utility/grid computing:

"If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility... The computer utility could become the basis of a new and important industry."

John McCarthy, MIT Centennial in 1961

What do we really want

Original vision of Utility/grid computing:

"If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility... The computer utility could become the basis of a new and important industry."

John McCarthy, MIT Centennial in 1961

What do we really want

Original vision of Utility/grid computing:

"If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility... The computer utility could become the basis of a new and important industry."

John McCarthy, MIT Centennial in 1961

Virtualization converts computation into a fungible commodity

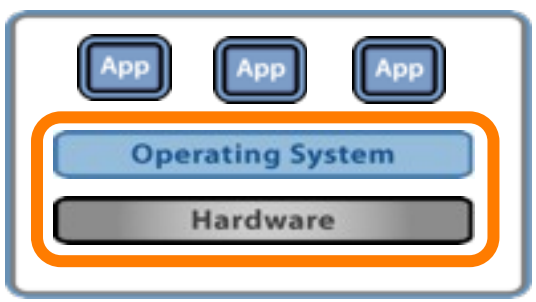

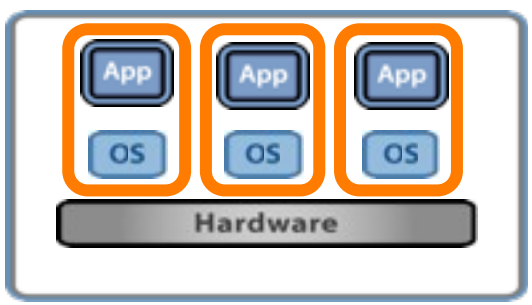

vmware

Why would this be transformative

- **Nicholas Carr:**
  - “As with the factory-owned generators that dominated electricity production a century ago, today's private IT plants will be supplanted by large-scale, centralized utilities.”
- **Gets rid of key impediments to innovation:**
  - Virtual appliance model for distributing installing applications.
  - Avoids need for broad HCL, OS support, …
  - Availability of massive capacity on demand.
- **Enables long-tail in SW**

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

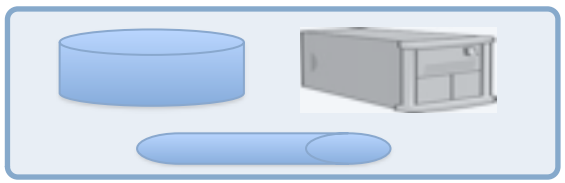

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

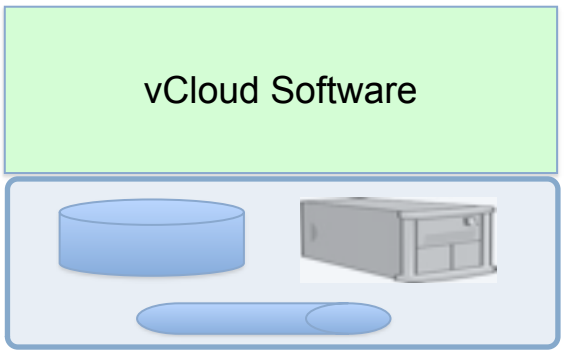

vmware

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

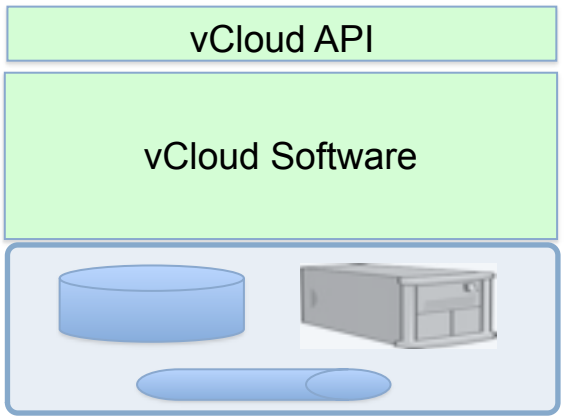

vmware

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

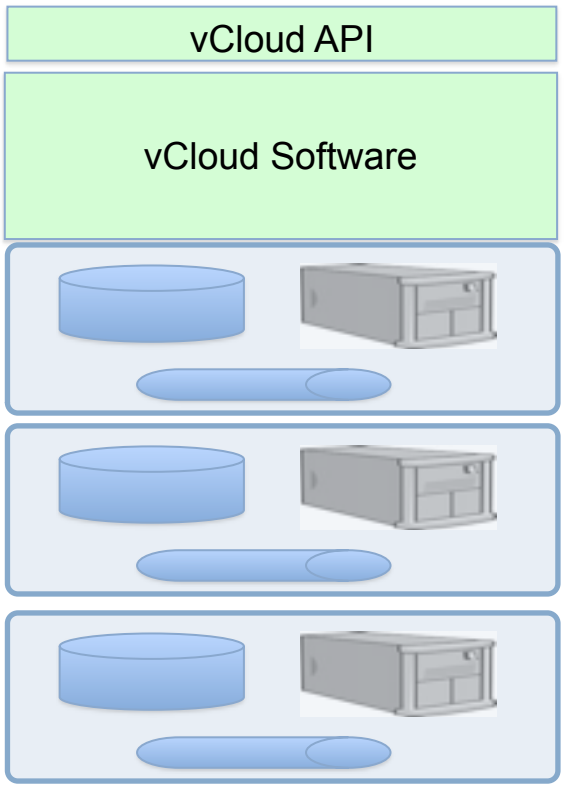

vmware

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

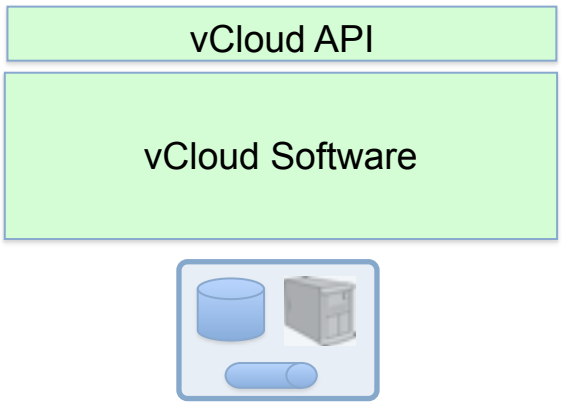

vmware

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

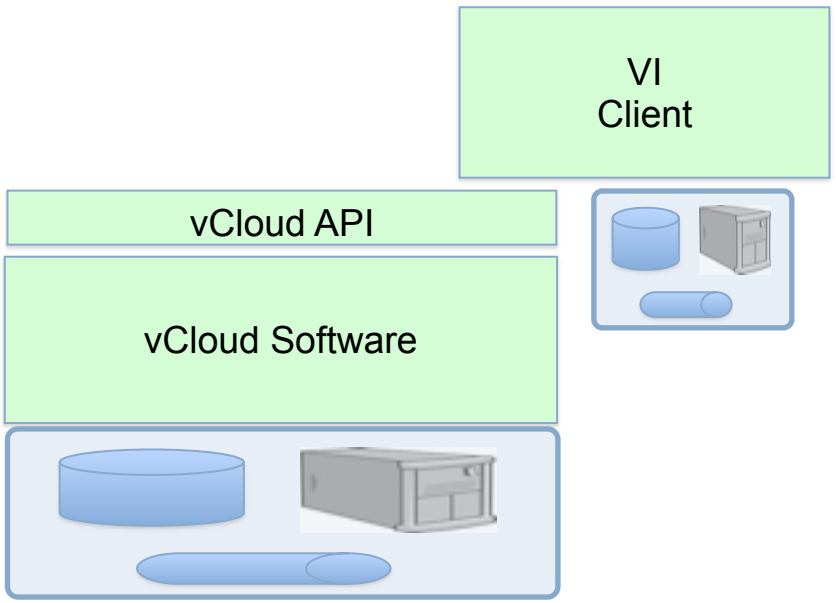

vmware

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

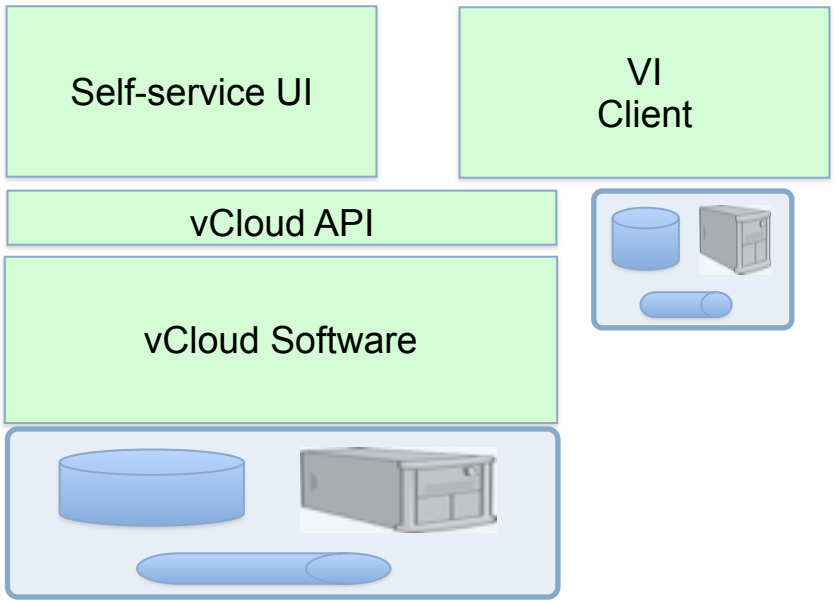

vmware

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

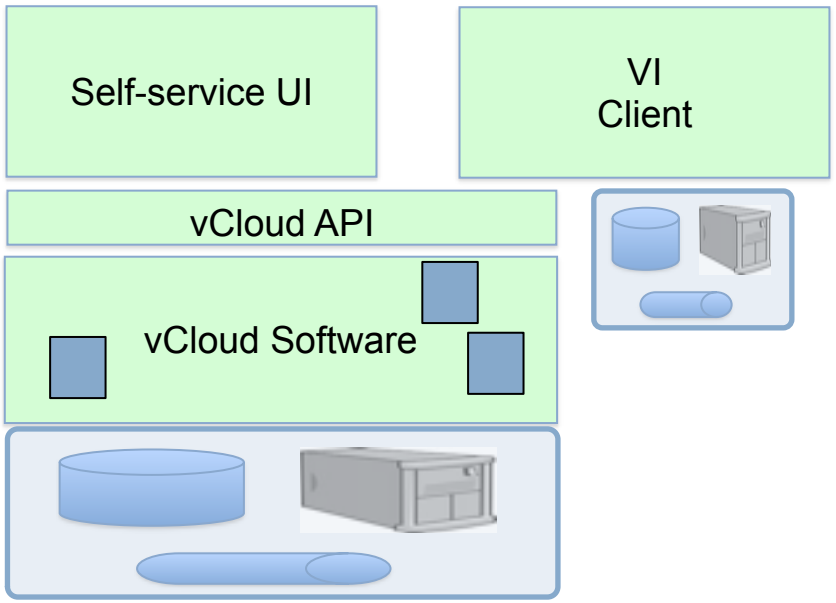

vmware

VMware's vCloud initiative

- **Provide SW so that anyone can easily create cloud: e.g., SP, University, Enterprise**
- **Supports all applications**
- **Controlled programmatically:**
  - From VMs, from capacity owner, from portal
  - OVF based API
- **Can scale up to massive data centers.**
- **Can scale down to small department.**
- **Can be accessed from VI Client**
- **New end user interface provide simple self service experience.**
- **Enable broad partner and research collaboration:**
  - Researchers can replace any part of the service.
  - Researchers can replace the entire implementation and clone the API

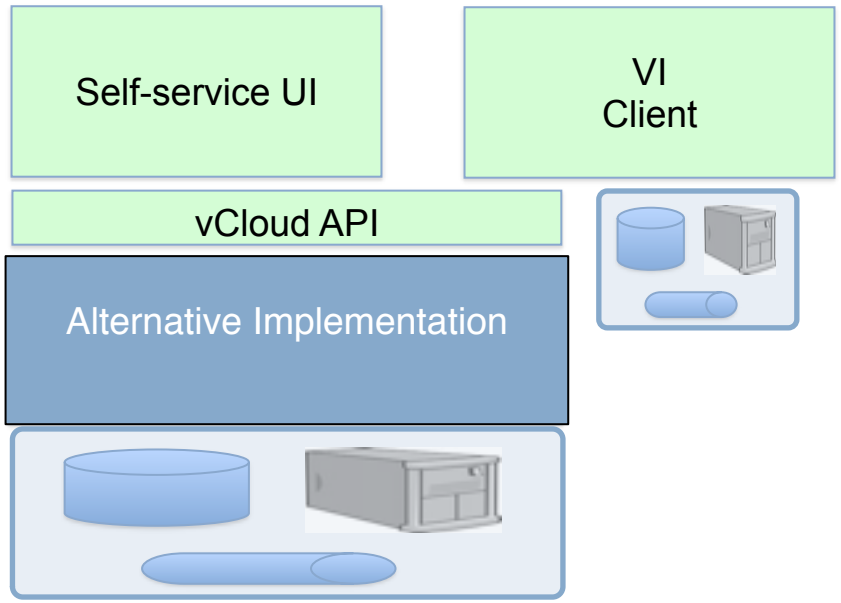

vmware

Key abstractions vCloud

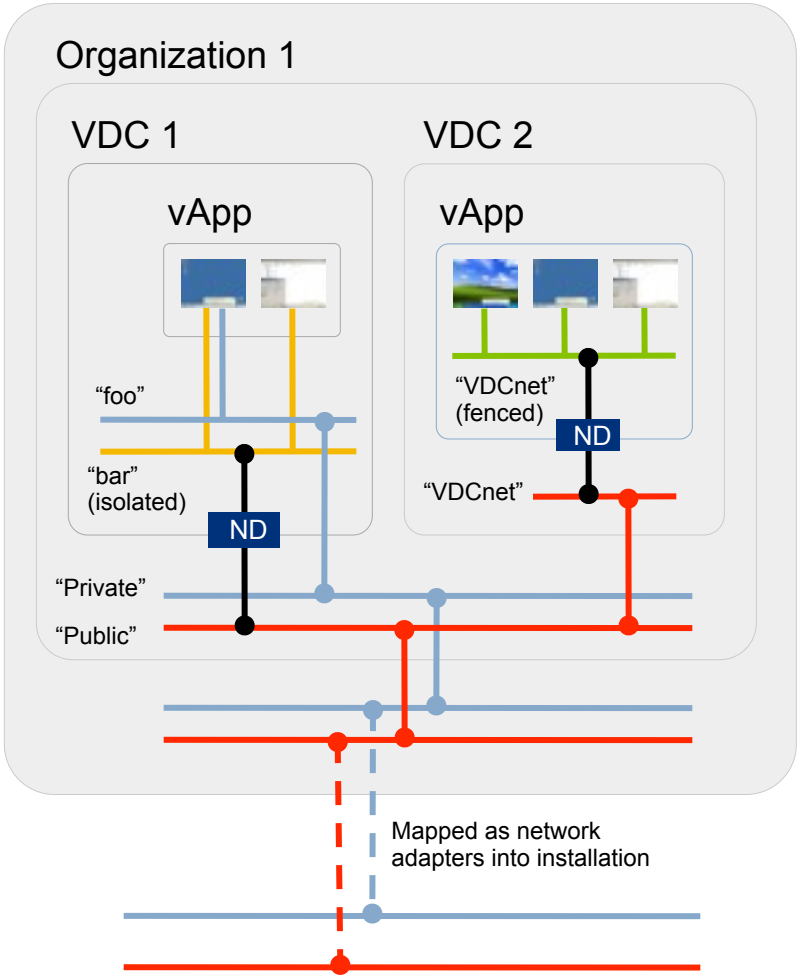

- **Organization with users/ roles**
- **Cloud to admin on his own terms:**
  - data centers
  - L2 networks
  - Persistent VMs
  - SLA on vApp
  - overprovisioning
  - ...
- **OVF based REST API:**
  - multi-tiered applications
  - serialized using OVF
  - configuration via OVF

vmware

Key abstractions vCloud UI & API

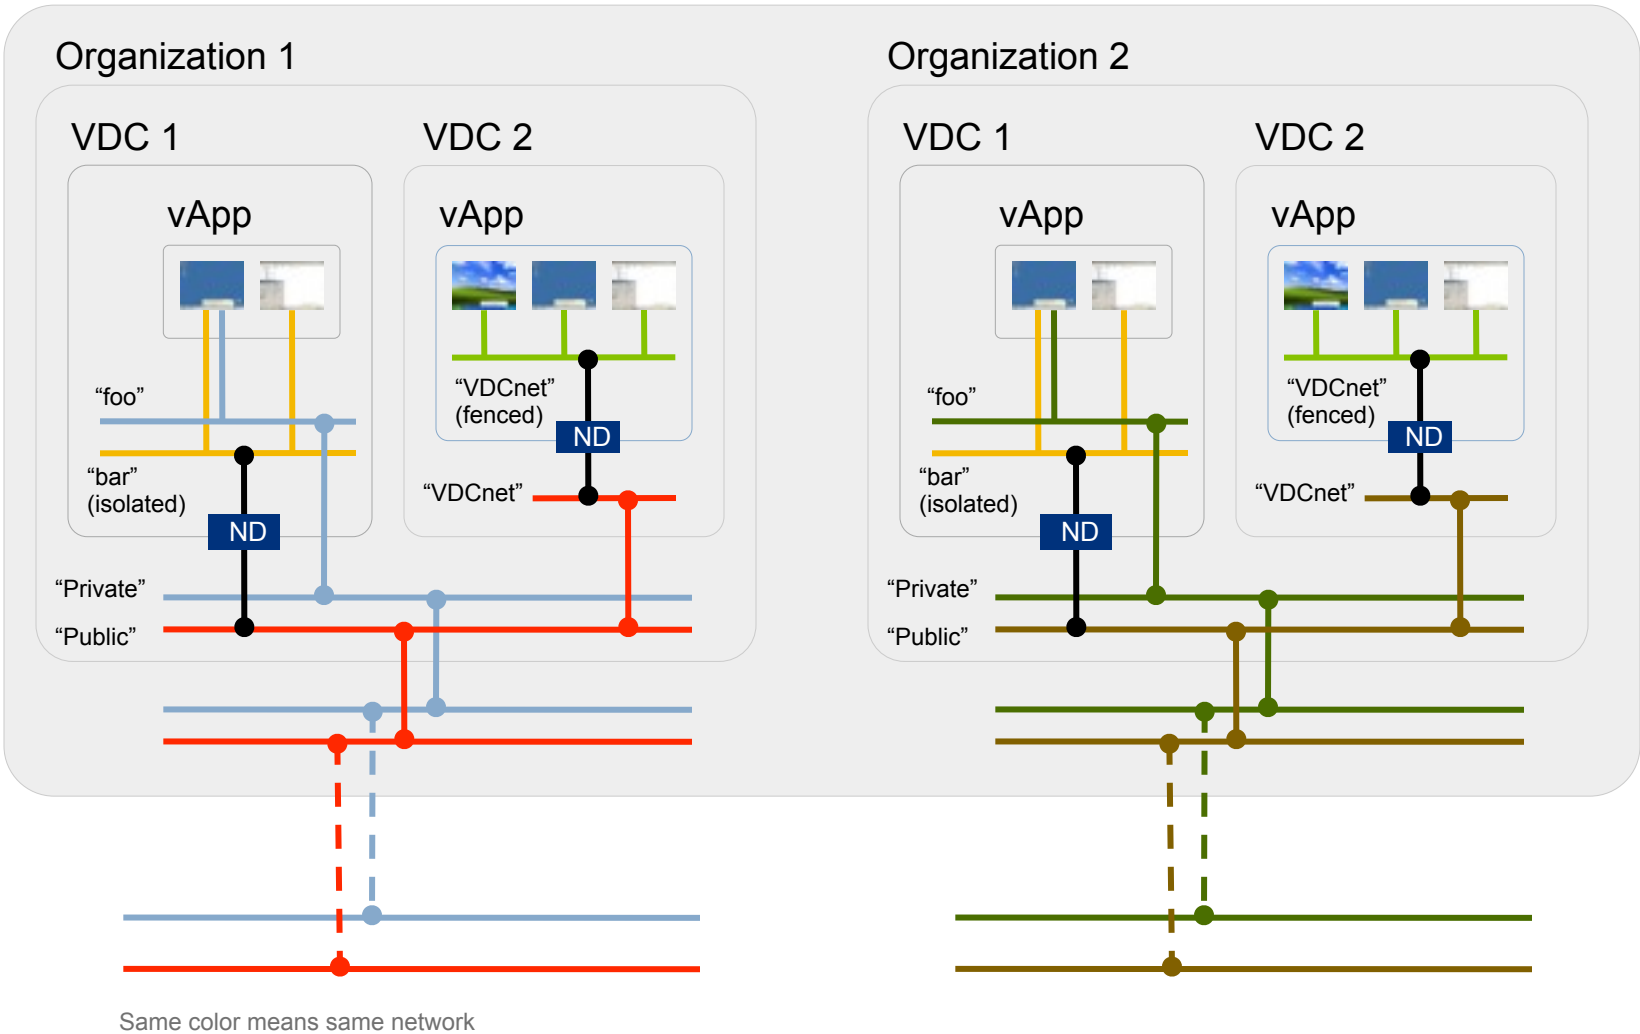

Same color means same network

vmware

It should be easy, VMware already has

- **Rich service provider, OEM, ISV partner ecosystem**
- **API for controlling virtualization that has become a defacto standard**
- **Ability to deal with large numbers of hosts as a pool of resources**
- **Support for backup, DR, resource management, power management, …**
- **Rich community of users**
- **Rich user interface loved by administrators**
- **…**

Not so fast…

- **Problems with our existing technology:**
  - Core abstractions exposed physical as well as virtual
  - Scaling up to 10000 hosts & 100K VMs fundamentally different
  - Security more of a issue in multi-tenancy environment

- **Challenges with our approach:**
  - Need to scale down
  - Need to enable partner ecosystem:
    - No one circumscribed approach.
    - Need to release SW to others to install, configure…
    - Need to enable others to innovate.

Difference between small & large scale

Difference between small & large scale

Difference between small & large scale

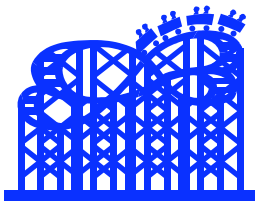

Difference between small & large scale

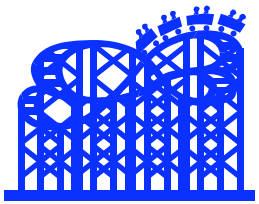

Difference between small & large scale

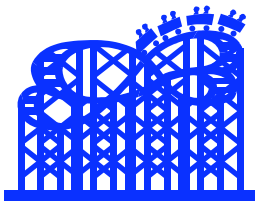

Difference between small & large scale

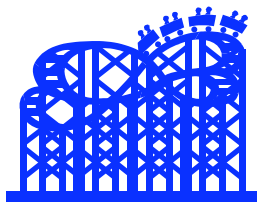

Difference between small & large scale

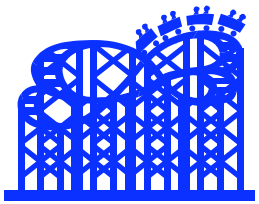

Difference between small & large scale

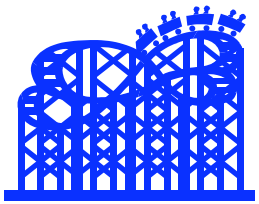

Scale up versus down

- **Large scale service: Service Oriented Architecture**
  - Each service totally independent with own DB, own set of machines… all communication through interface:
  - Advantages
    - Fails independently & can identify failures
    - Own reliability based on requirements of data
    - Can evolve services independently
    - Can scale them independently
- **Small scale: shrink wrapped software**
  - Single general purpose piece of software
  - Advantages:
    - Minimize overhead
    - Simplified installation

Modular design

Large Scale

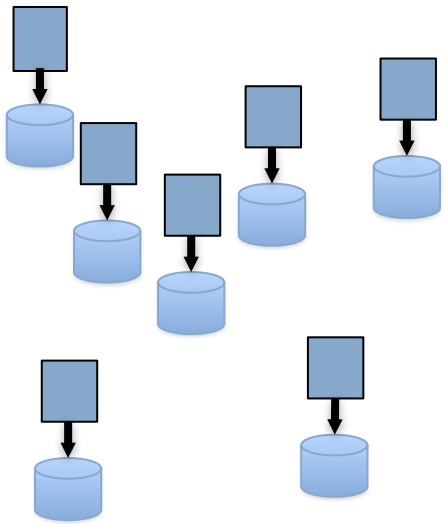

Modularity also enables extensibility

Modular design

Large Scale

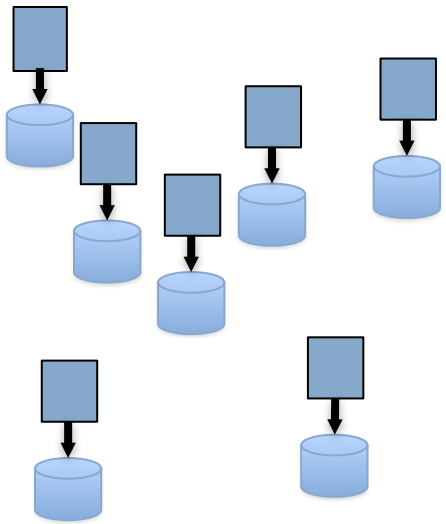

Small Scale

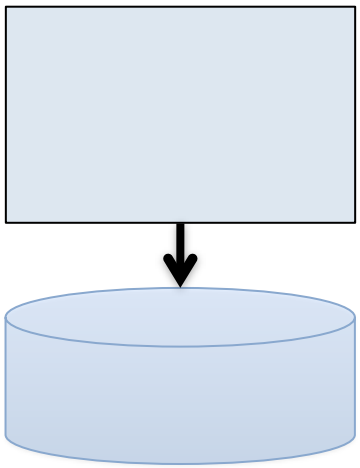

Modularity also enables extensibility

Modular design

Small Scale

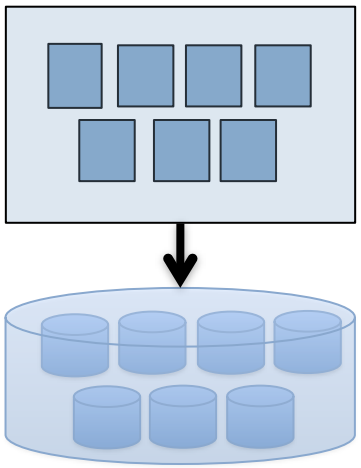

Modularity also enables extensibility

Key Technologies

REST API : HTTP based resource oriented interface;
All the characteristics of the WWW

- **Extensible without breaking client.**
- **Client only has to know about what it cares about.**
- **Can route, proxy, cache**

Spring: Standard component framework

- **Injects dependencies and wires together Spring beans**
- **Forces programmer into maintainable design pattern; isolates dependencies**

OSGI: Standard dynamic module framework

- **Global registry of interfaces to instances**
- **Dynamically load, unload, start, stop bundle**

Other technology

- **JMS publish/subscribe messaging bus isolates end points**
- **Hibernate simplifies DB code & DB independence**

Architecture/Implementation

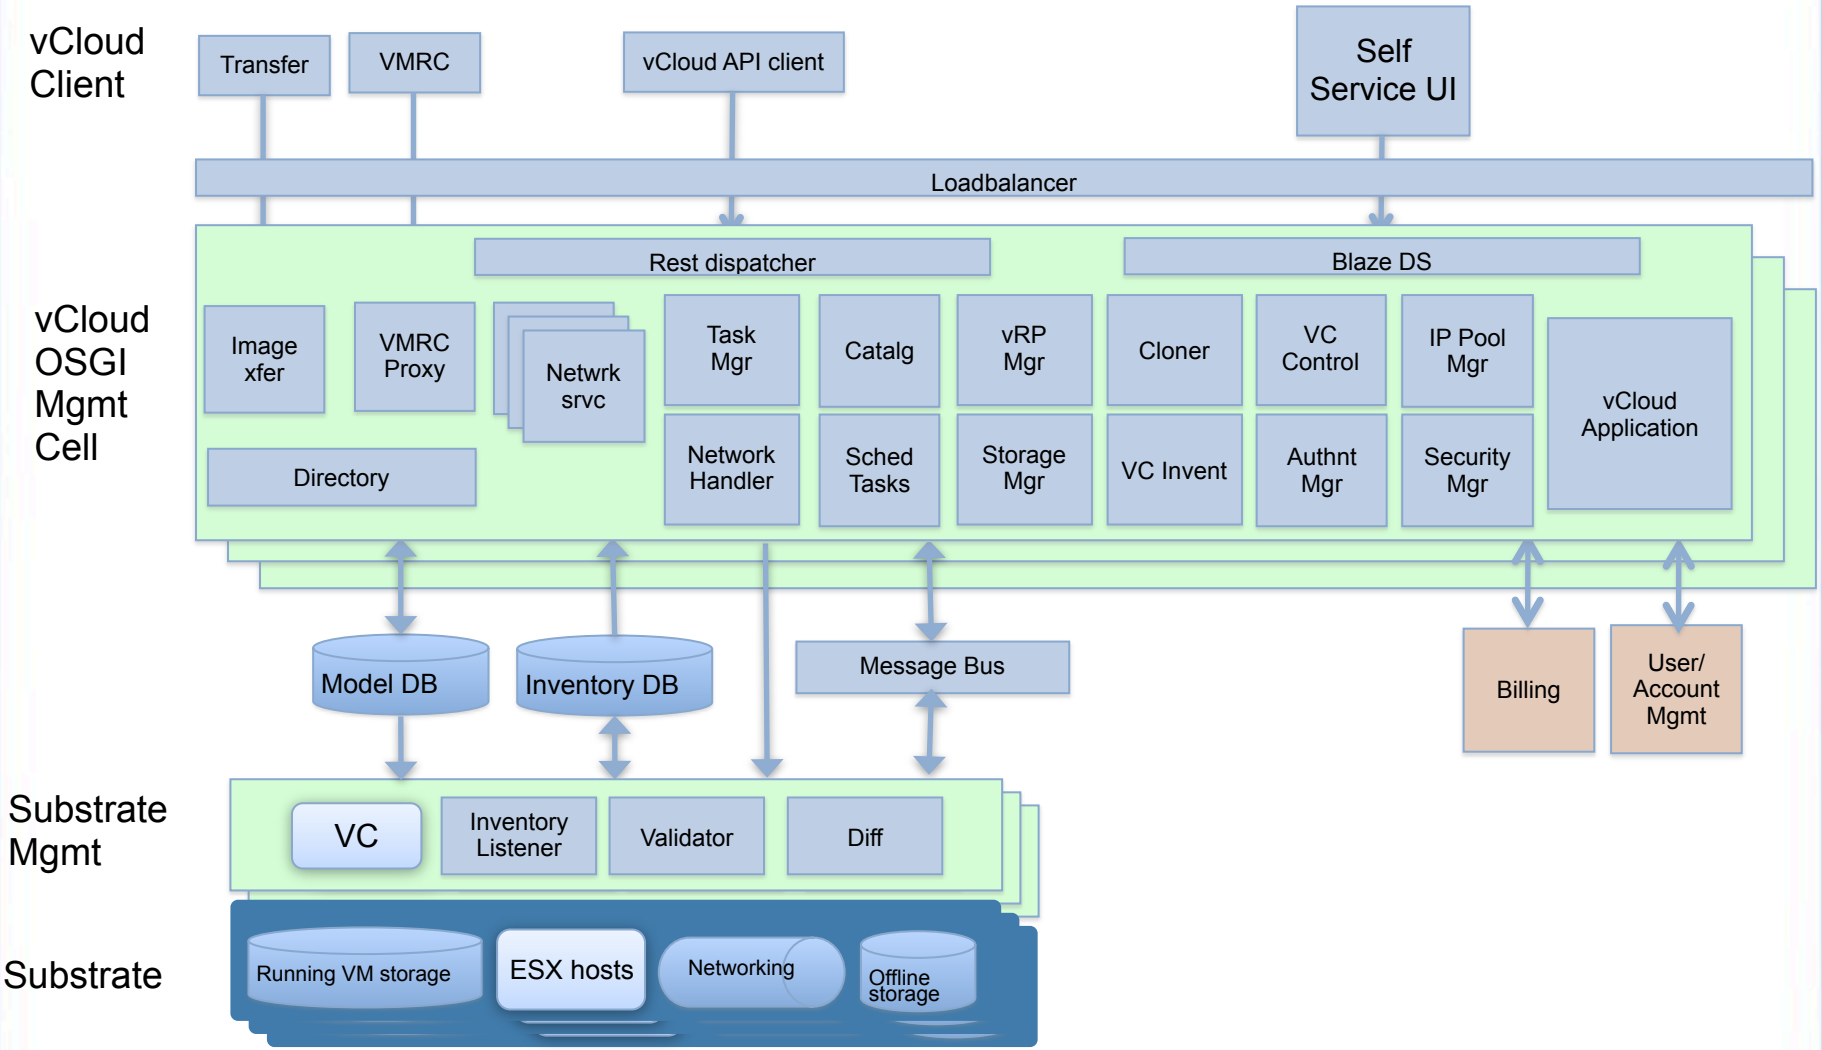

vmware

Architecture/Implementation

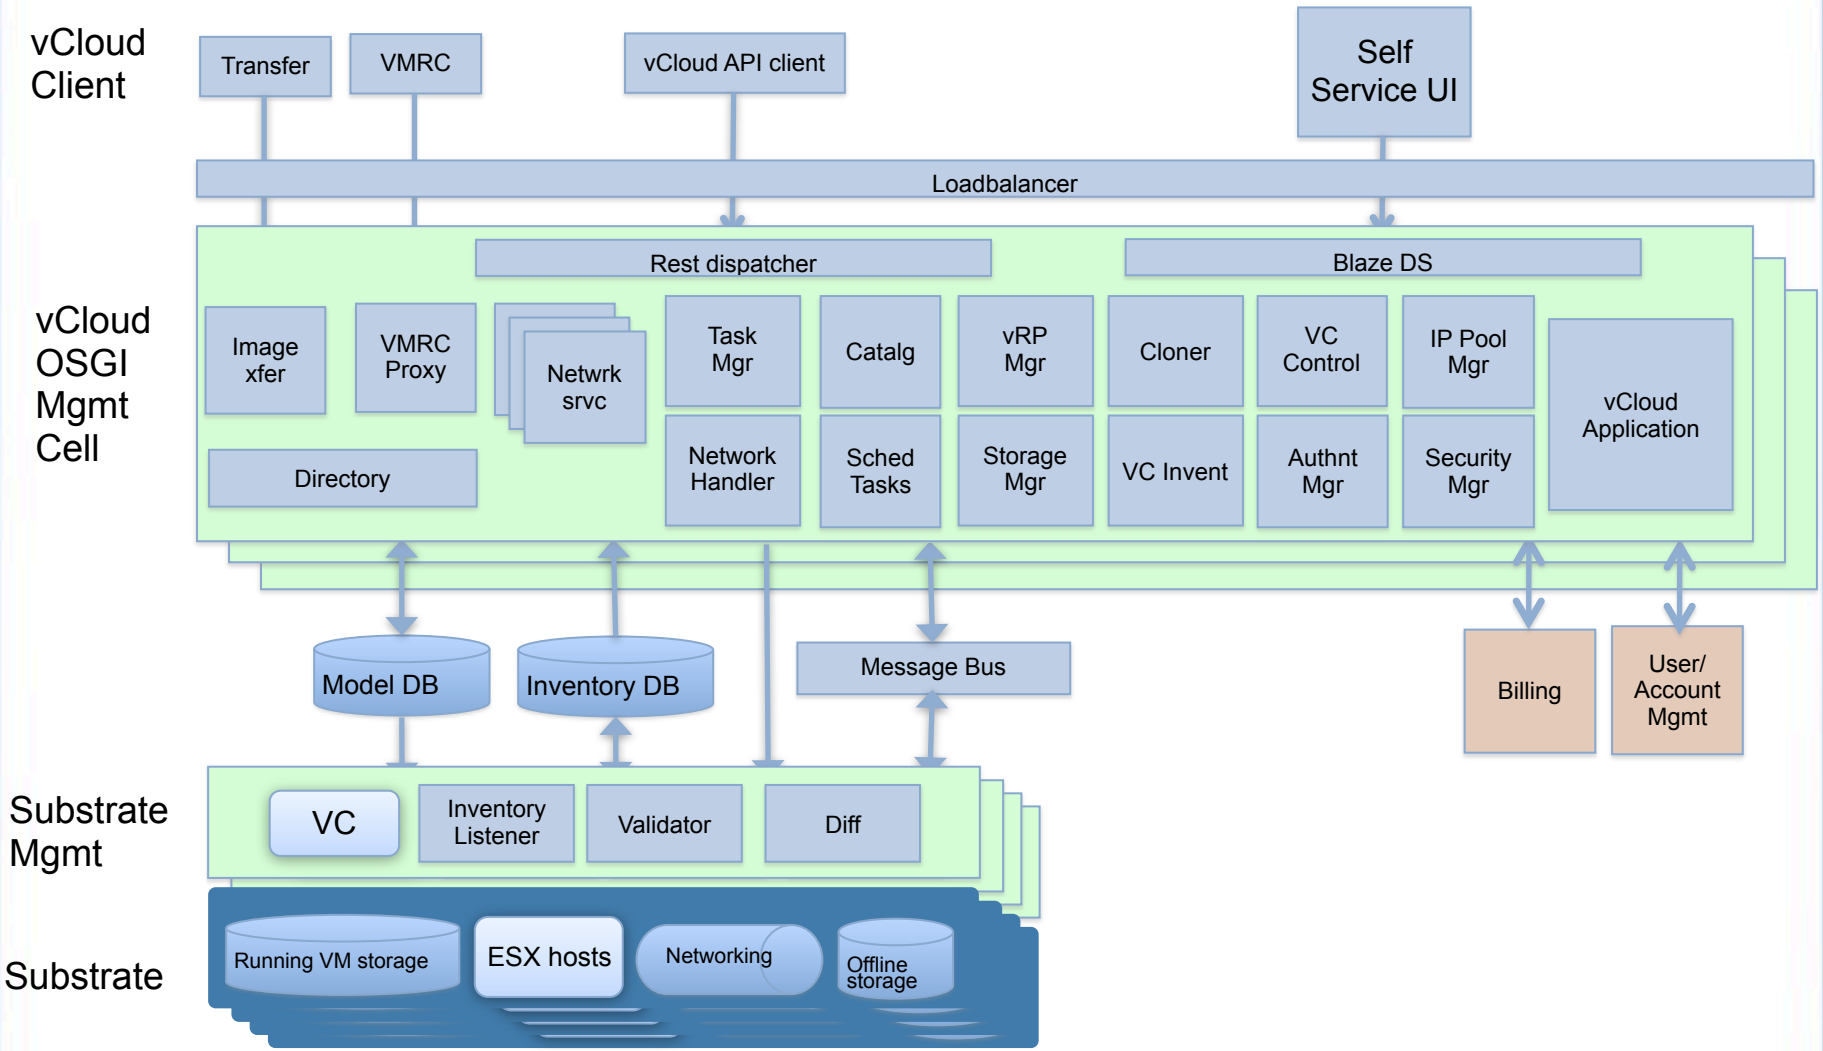

vmware

Architecture/Implementation

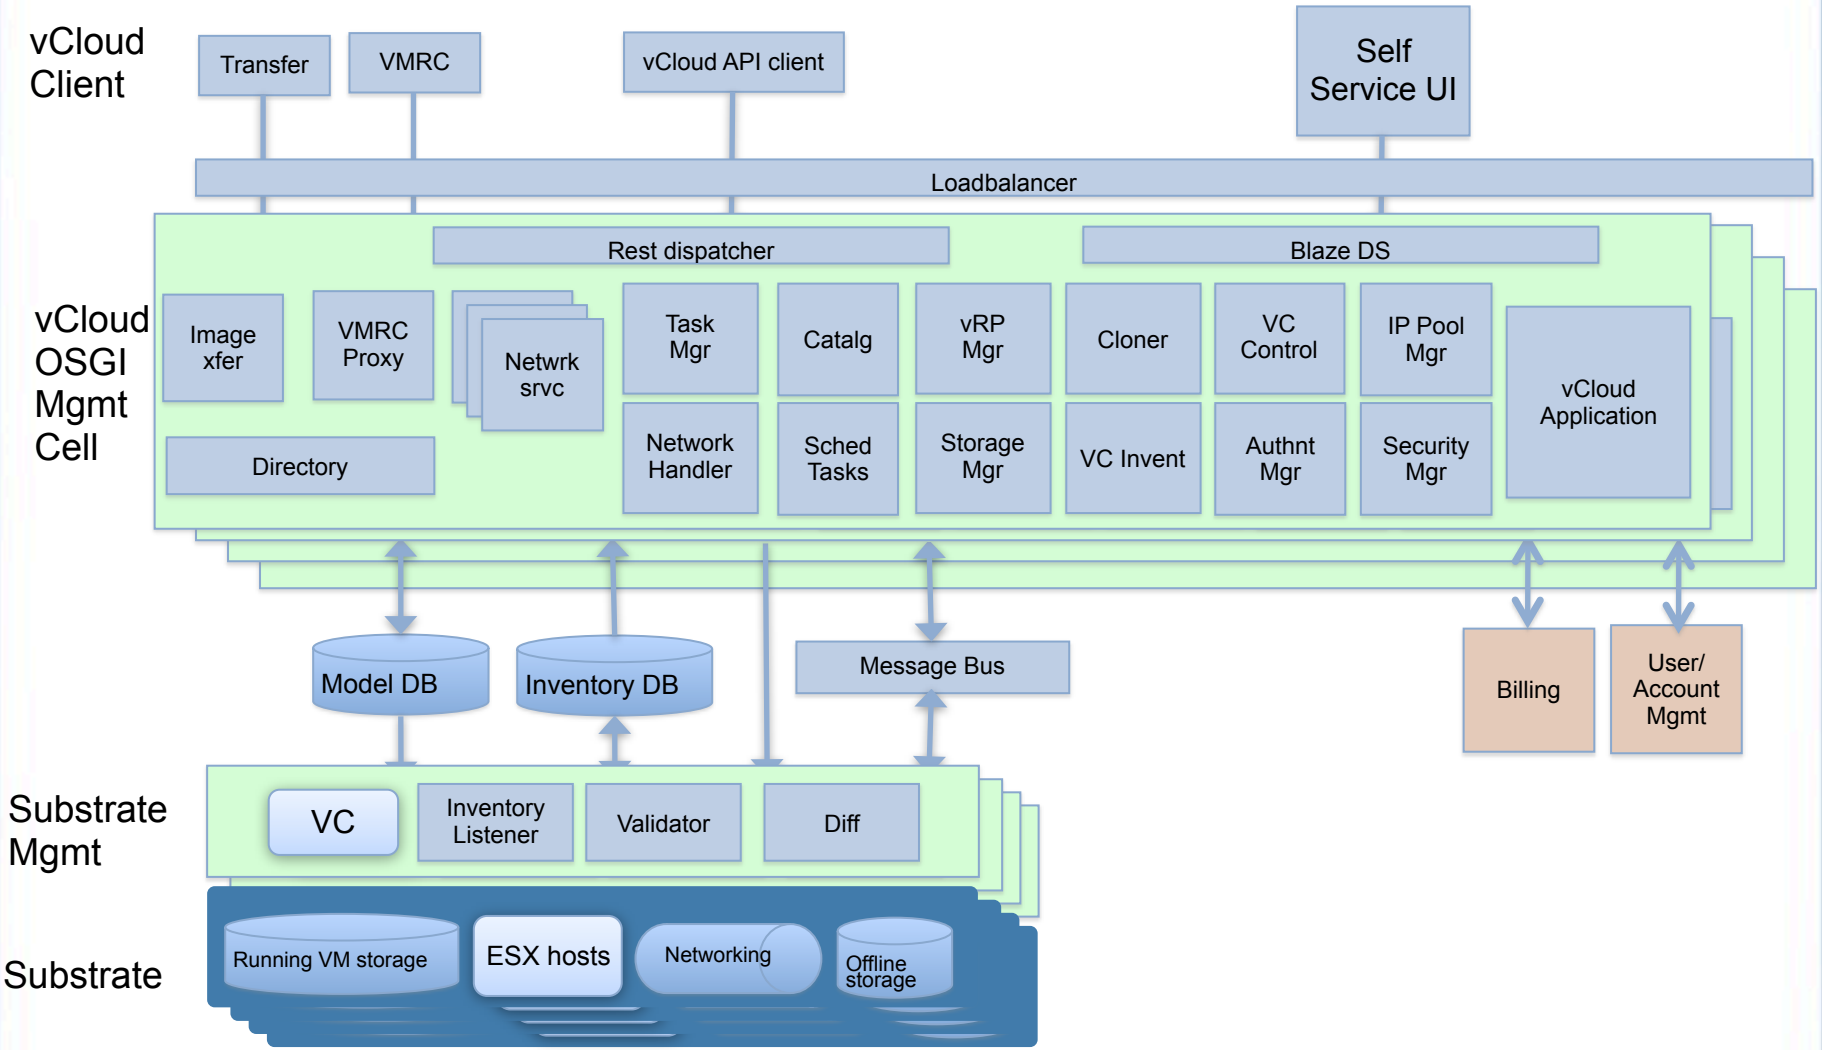

vmware

Architecture/Implementation

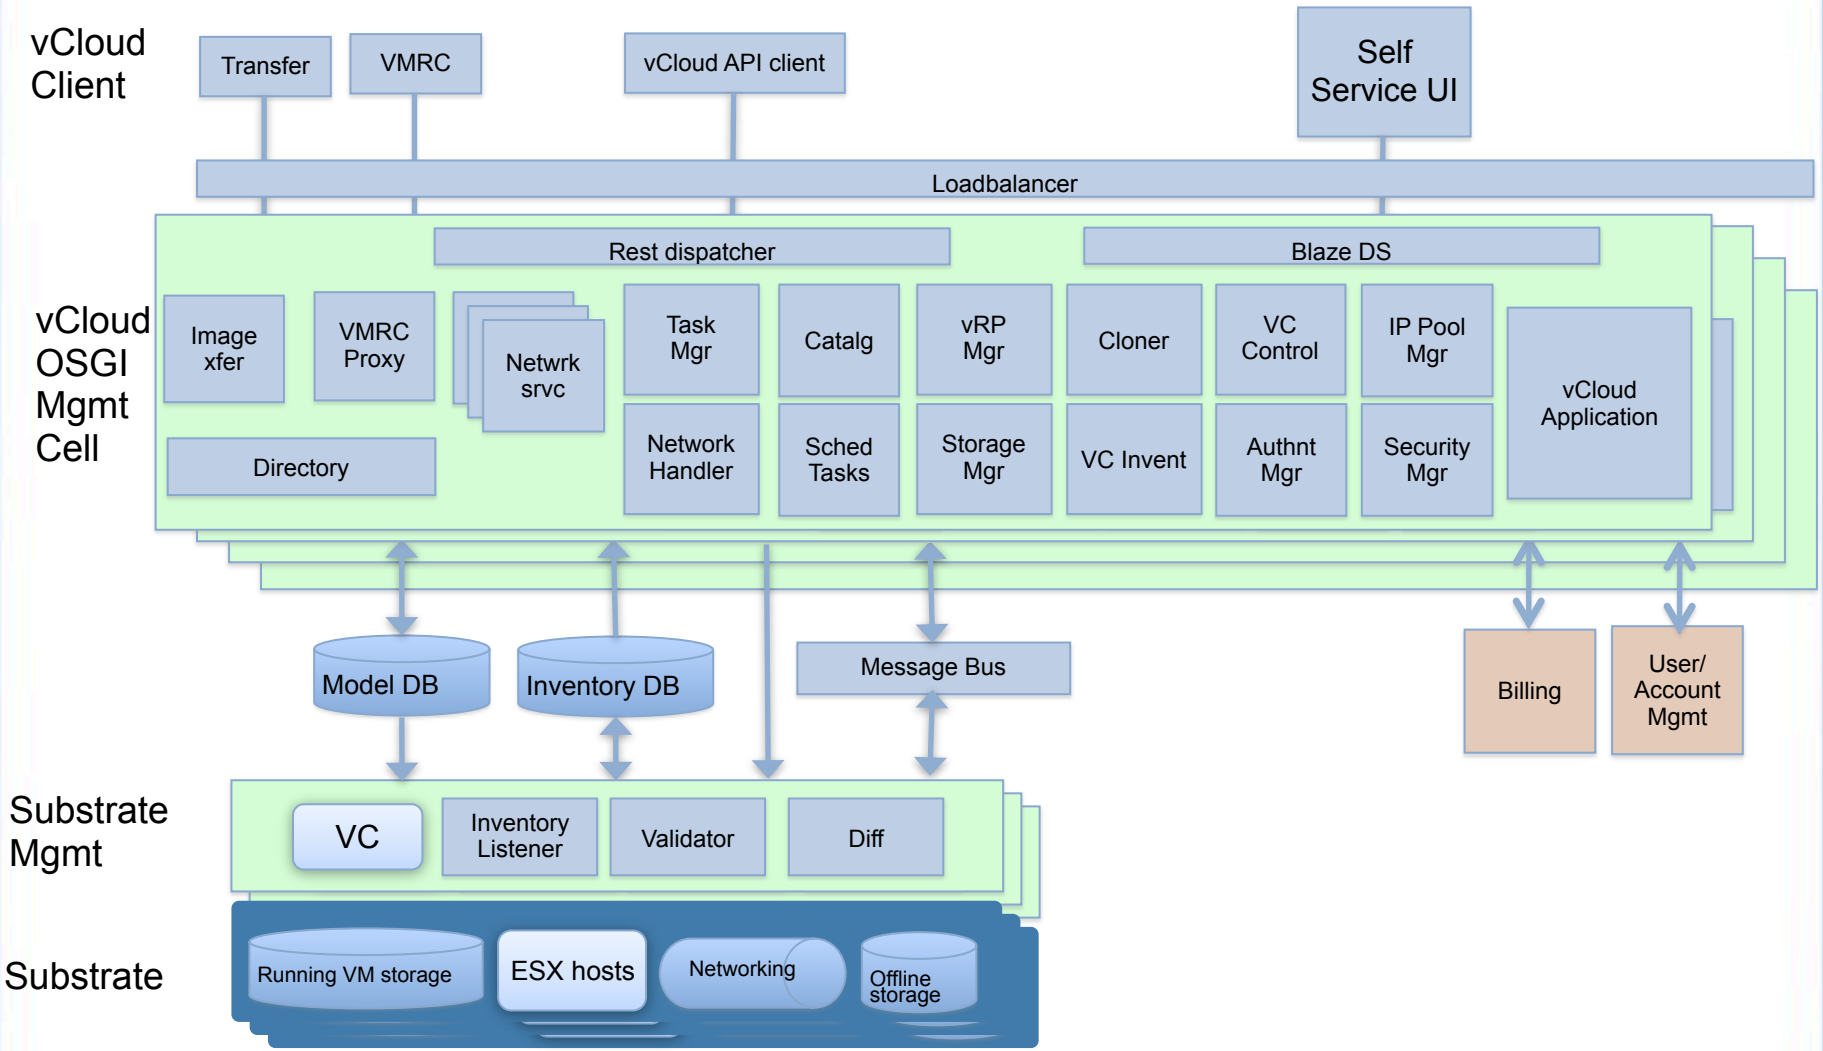

vmware

Architecture/Implementation

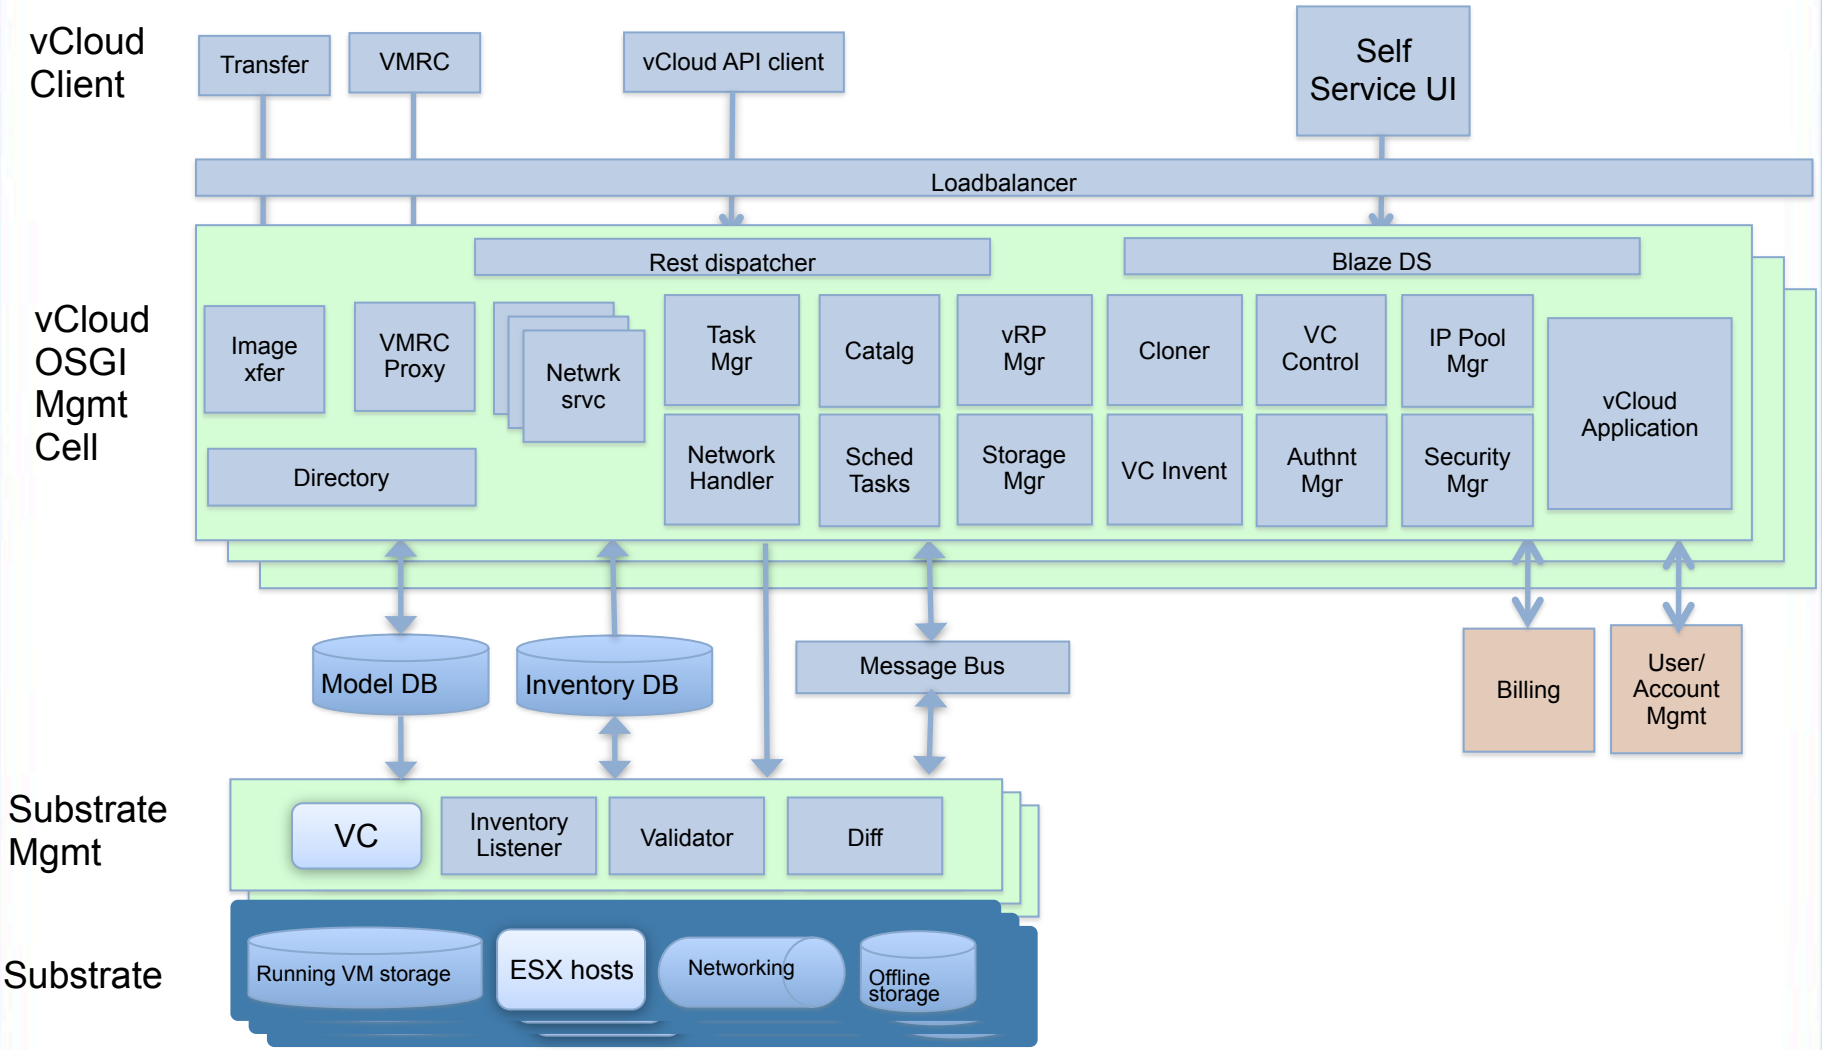

vmware

Architecture/Implementation

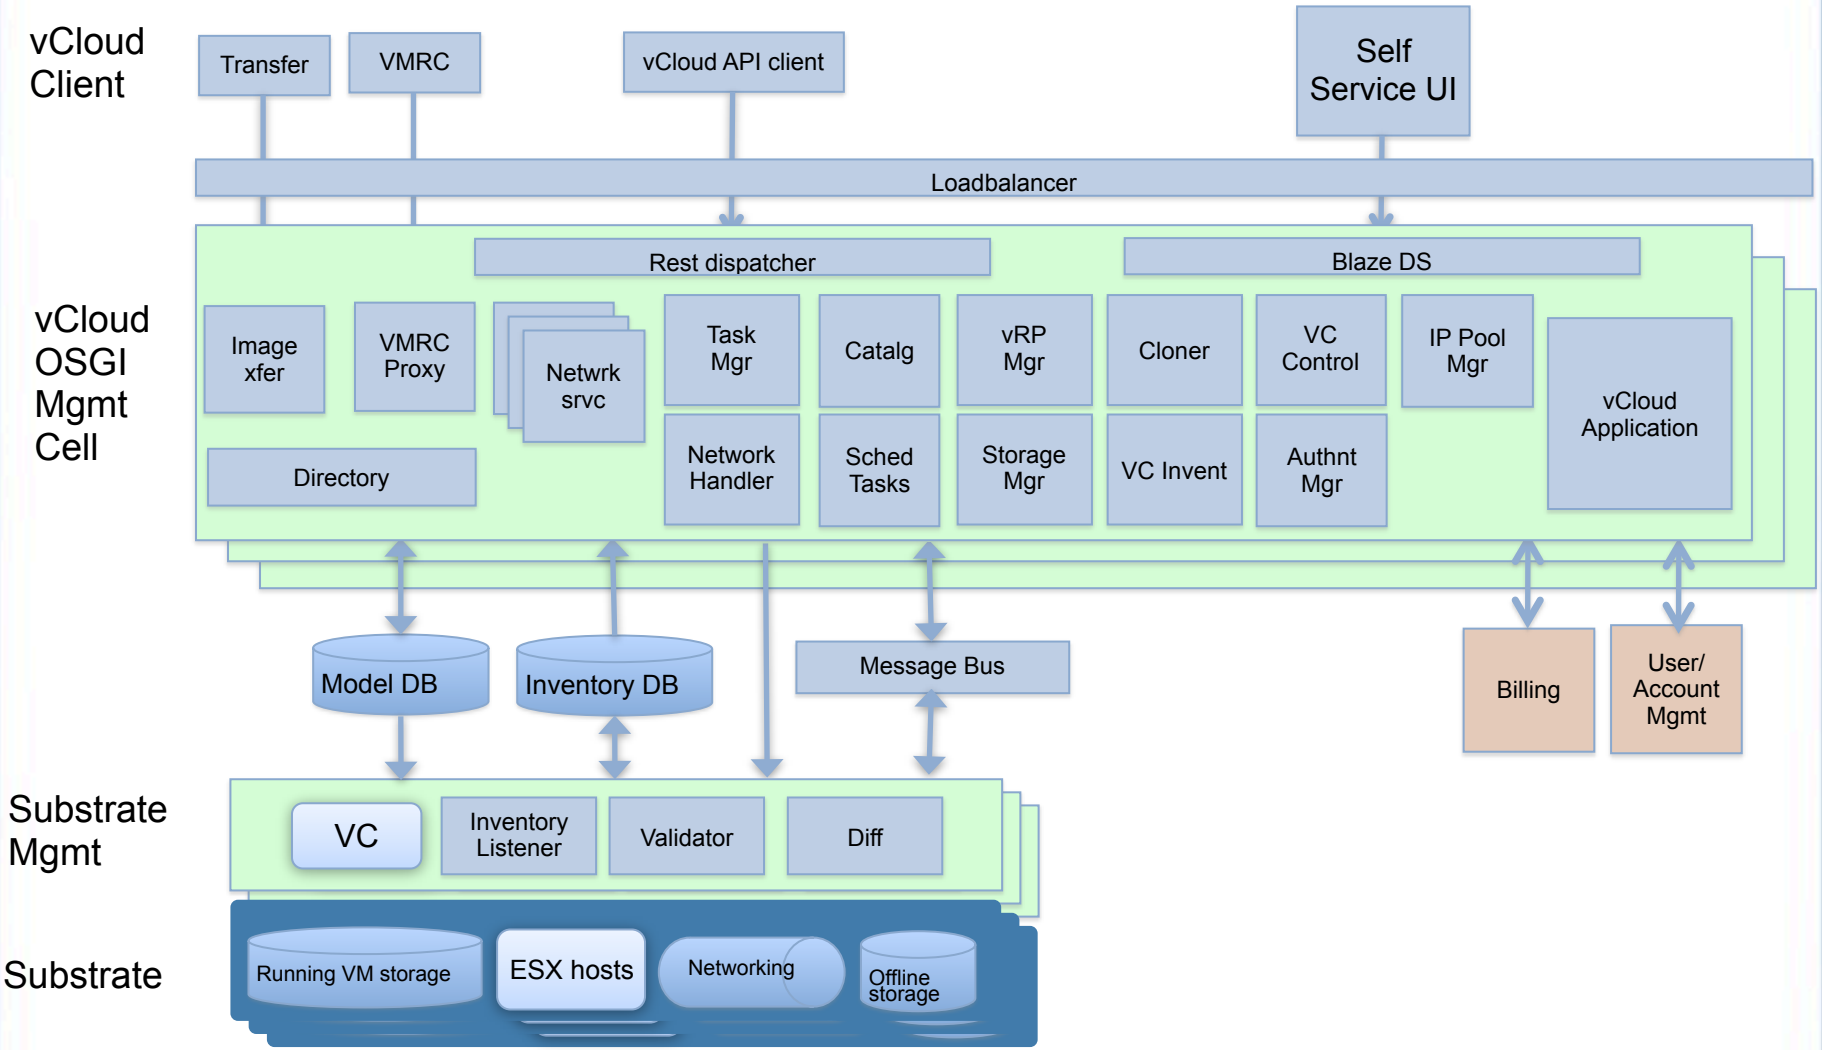

vmware

Architecture/Implementation

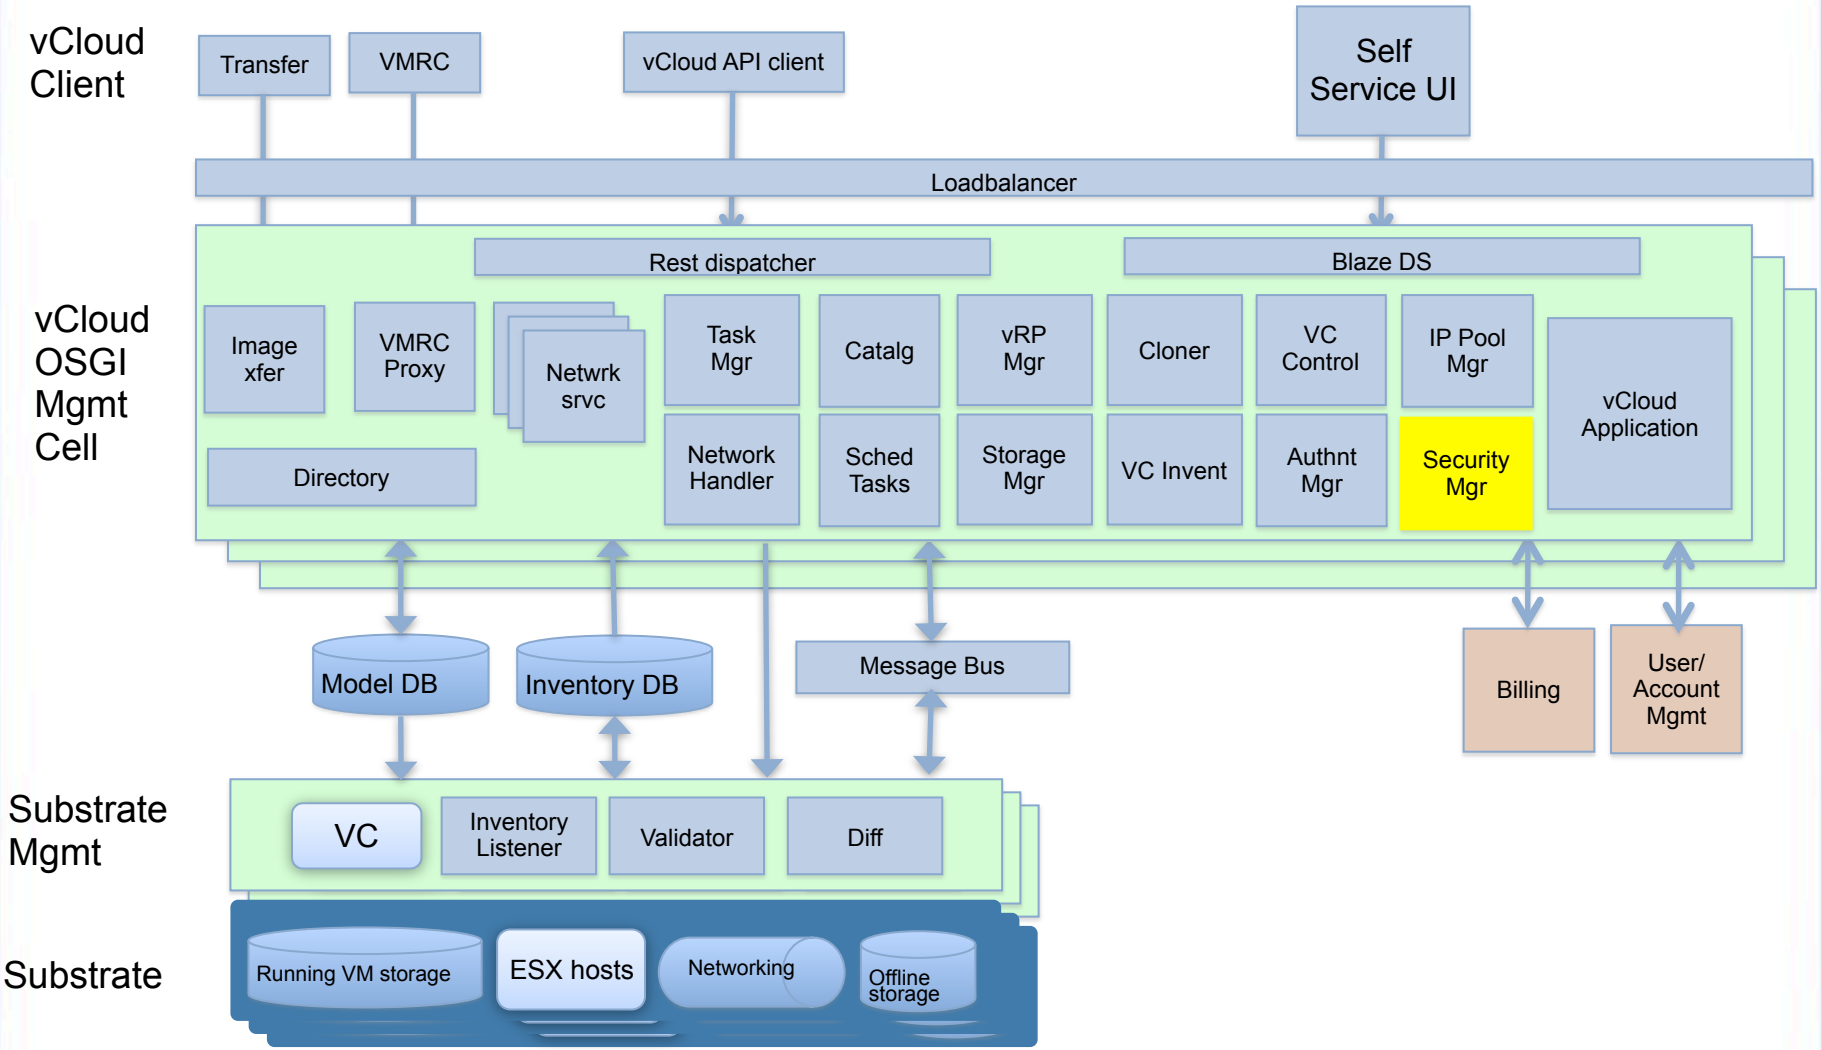

vmware

Research directions

- **New operating systems:**
  - e.g., for Java, security services, clustering services, HPC/ multi-core, ...
  - library OS, multi-core, accelerator to OS, sharing
- **Fungible computing:**
  - performance, cost to user, cost to SP
- **Grid/HPC/batch scheduling, e.g., snowflock**
- **Federation:**
  - long tail for data de-duplication, encryption/security for data, trading floor/futures market, vmotion to user
- **What changes when your desktop is in the cloud?**
  - Mom's sysadmin, Fusion on steroids, disaster recovery...

Concluding reports

- **Cloud computing is going to be transformative to our industry**
- **VMware building a platform so that anyone can play**
- **There are a wealth of research opportunities**
- **We will be providing the SW to universities this year**
- **This is just the start...**