
The Clouds Are Coming! Are We Ready?

Gerrit Huizenga

gerrit@us.ibm.com

Cloud Market Place: The Clouds are Coming!

- The Cloud Market is growing extremely rapidly
 - Some assessments suggest a 2011 total market share in excess of \$100B
 - And growth through 2015 to over \$300B
- Market view is very broad, including
 - Public Cloud – infrastructure expense
 - Public Cloud – revenue generated
 - Private Cloud – infrastructure build out
 - Private Cloud – software and licensing
 - Includes labor costs and management
- Also includes a number of established and emerging sub-sectors such as:
 - Cloud based email (aka gmail)
 - Sales Force Automation (SaaS) such as Salesforce.com
 - Public Cloud Storage
 - Public Cloud Computing
 - Private Cloud Computing
 - Platform as a Service
 - Cloud Collaboration Services
 - Database Platform as a service
 - “Big Data” Cloud solutions
 - DevOps
 - And many, many more


A user experience and a business model

- Cloud computing is an emerging style of IT consumption and delivery in which applications, data, and IT resources are rapidly provisioned and provided as standardized offerings to users over the web in a flexible pricing model.

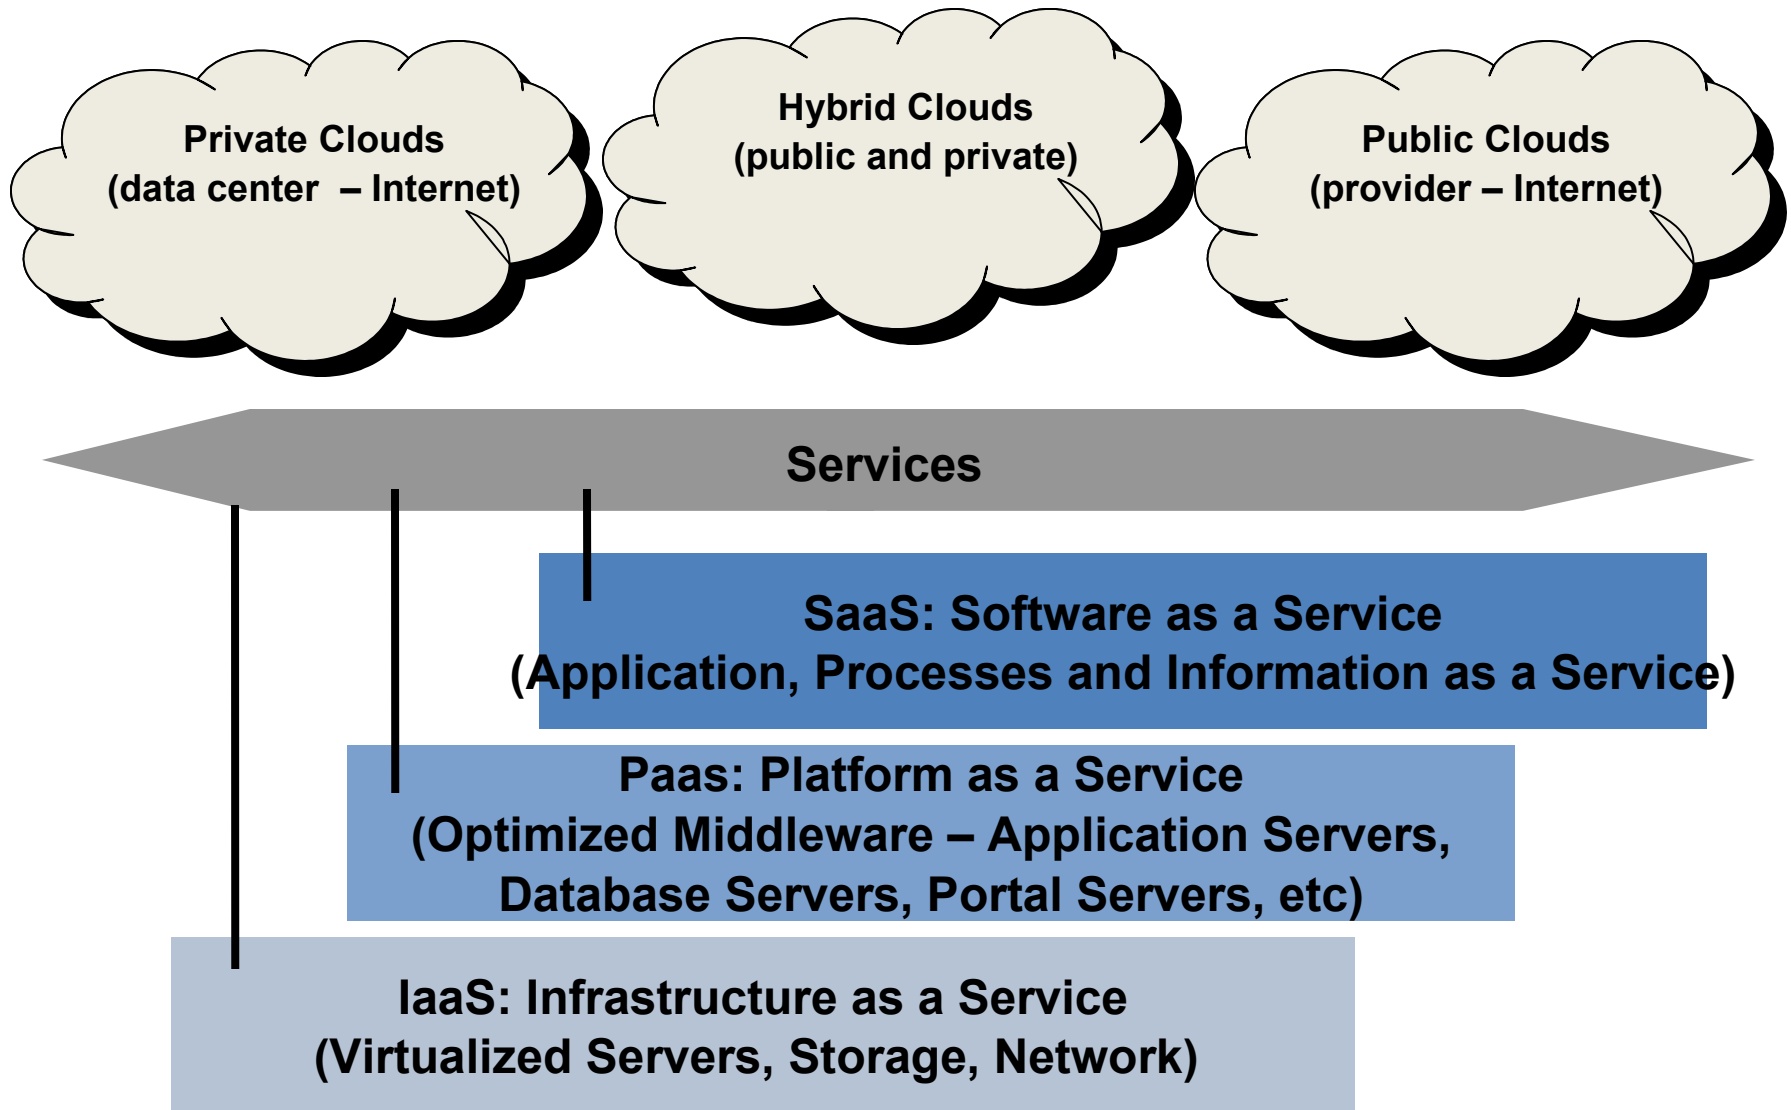
An infrastructure management and services delivery methodology

- Cloud computing is a way of managing large numbers of highly virtualized resources such that, from a management perspective, they resemble a single large resource.

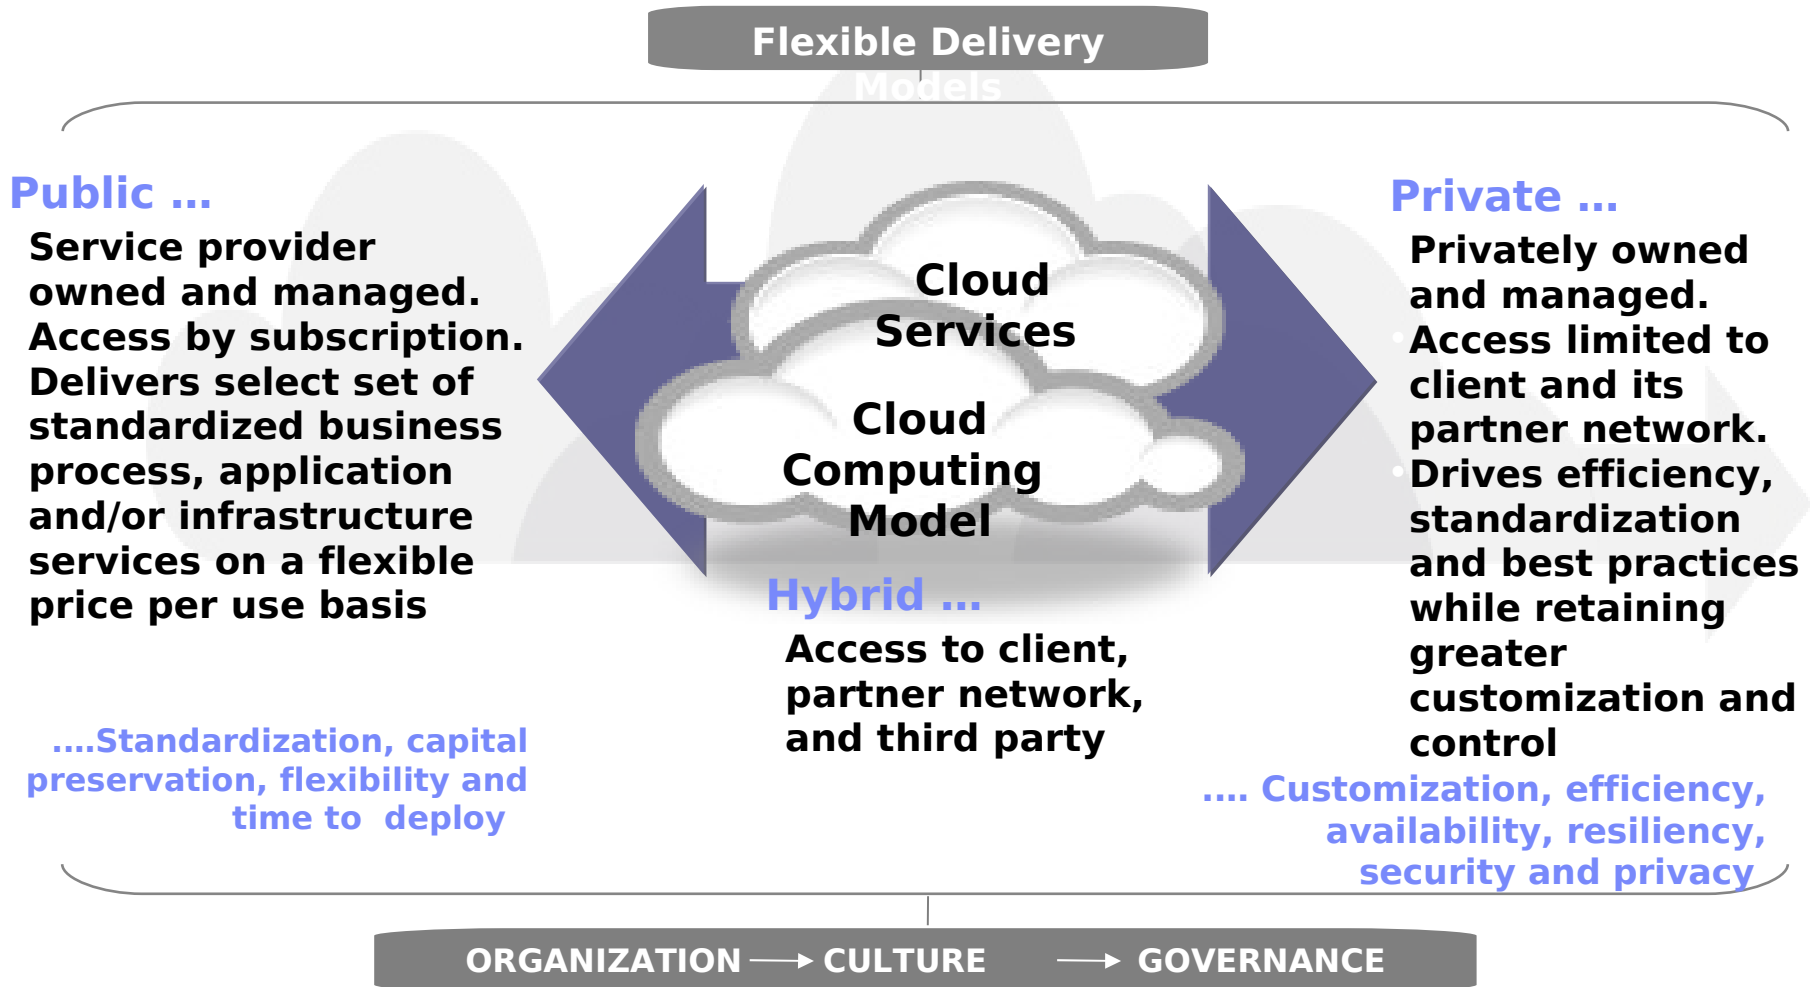
This can then be used to deliver services with elastic scaling.



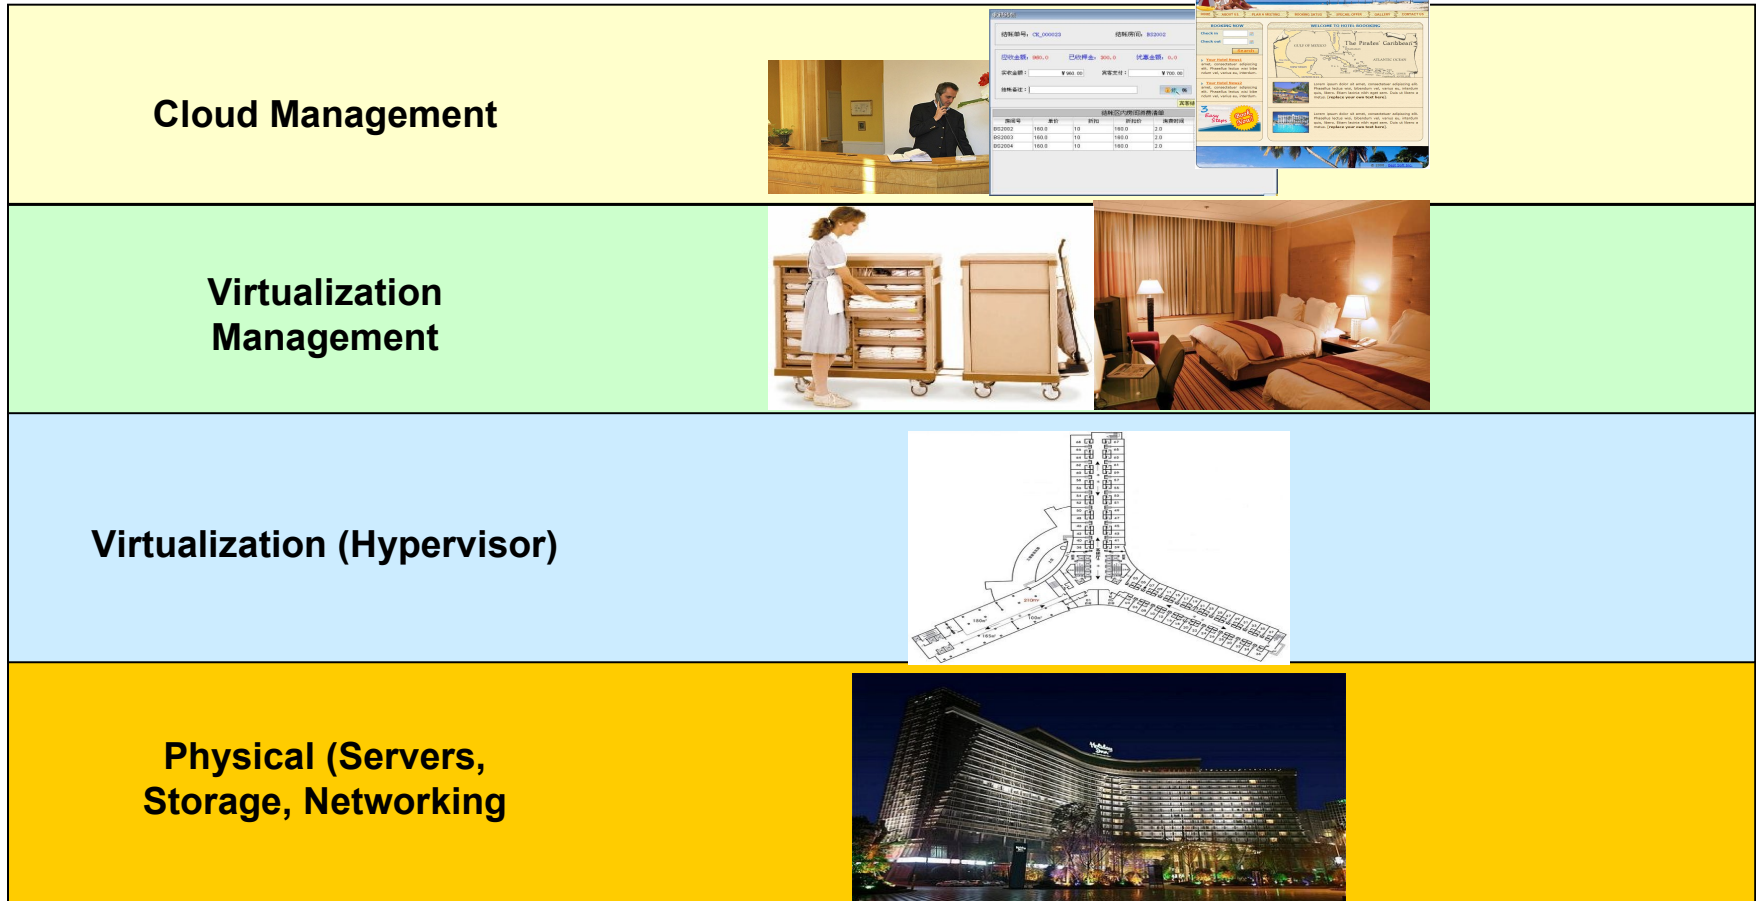
Cloud Computing Deployment and Services



For the enterprise, Cloud computing enables new services sourcing and value models



Core Commonalities in all Cloud Models



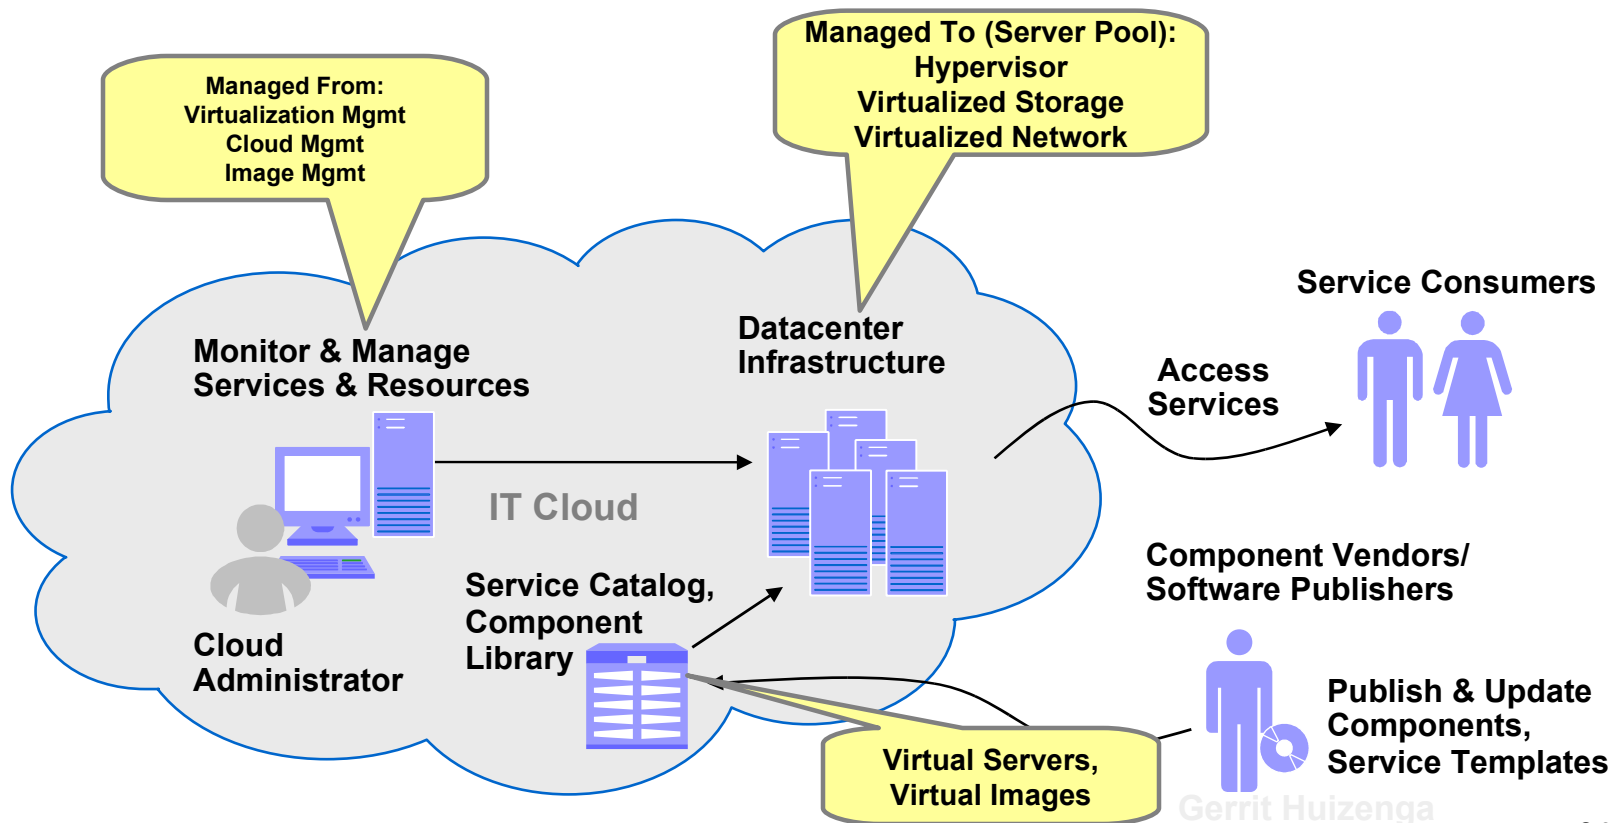
Cloud solutions add End User front end for selecting workloads (think airline self-service check-in terminals)
 Service catalog of choices that user can request (e.g. airline or hotel selections for car type, room type, etc.)

Linux as the Platform of the Cloud (Are We Ready?)

- Rumor is that some have called Linux the Platform of the Cloud
 - Is that realistic?
 - Sadly – NO! (at least not yet!)
- Today the Market Leader is truly VMware
 - VMware ESXi is the primary virtualization platform
 - vCenter is at the forefront of virtualization management
 - vCloud is becoming the most common Cloud Management platform
- In the Public Cloud
 - Amazon has been a Xen hypervisor based provider
 - And uses a home-grown virtualization management stack
 - With their own “Cloud-management” stack build in/on-top of their stack
- Microsoft's HyperV and System Center Management
 - Newer player but moving in rapidly
 - MS using its broad client base and aggressive pricing to capture Cloud mindshare
- Oh, and probably 25 more or so...
 - Cloud.com CloudStack, Surgient Platform, Enomoly ECP Ca/3Tera Applogic, Abiquo, Nimbula Director, Citrix XenCloud, Joyent Smart OS, Platform ISF, etc. etc.
- And don't forget Open Source!
 - Eucalyptus, OpenStack, Ubuntu Enterprise Cloud, Nebula, Open Nimbus, Piston (OpenStack)
- So many options, and the Open Source stacks are still in the very small minority!

Within the Cloud, the hardware and cloud related software is divided into three components:

- **Managed From:** Hardware used to support the Cloud Management Stack
- **Managed To:** (usually) Virtualized Hardware
- **Virtual Servers, aka Virtual Images**

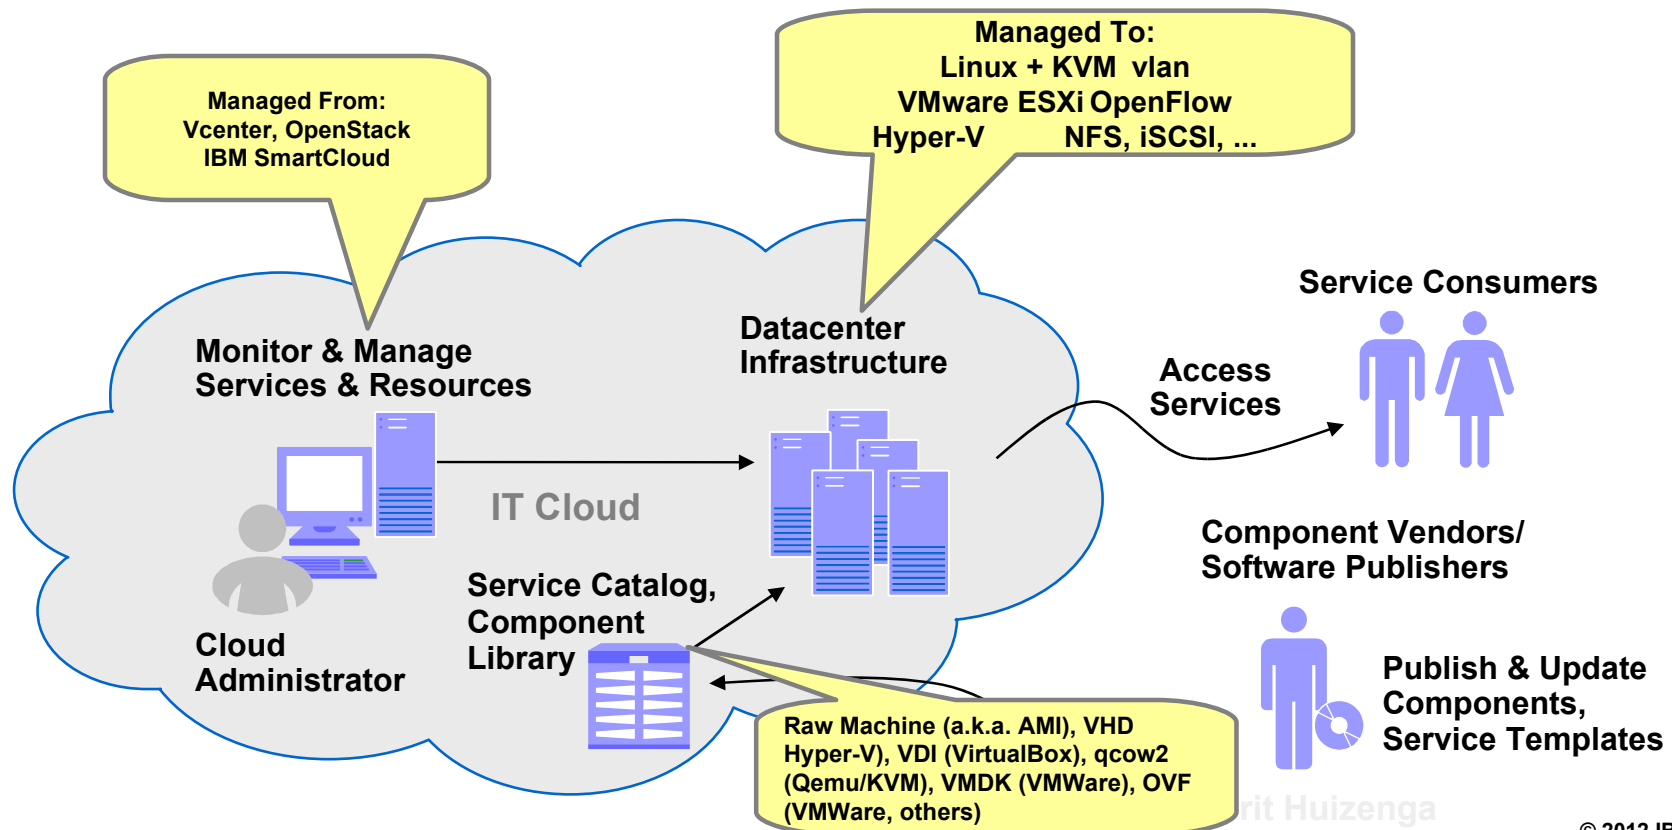


Managed To vs. Managed From

- Cloud Includes a Software Management Stack and a set of managed (virtualized) hardware
- It is important to distinguish between the management stack running on physical (or virtual) hardware and...
 - The virtualized managed hardware
- The Management Stack
 - Often runs on non-virtualized hardware
 - Includes a management stack such as vCenter, vCloud, OpenStack, etc.
 - Is completely independent of the managed hardware (at least in general)
- The Managed Servers
 - Are most often virtualized
 - Are managed by the management stack as a “server pool” (many names for this)
 - Can be re-imaged on demand
 - Are independent of the management stack
 - Could be multiple pools of diverse hardware, hypervisors, etc., managed by a common management stack
- And.. Consider the contents of the virtual servers running on the managed servers...

Within the Cloud, the hardware and cloud related software is divided into two components:

- **Managed From:** Linux or Windows, vCenter, OpenStack, vCloud, etc.
- **Managed To:** (usually) Virtualized Hardware
- **Virtual Servers, aka Virtual Images**



Public Cloud vs Private Cloud

- Public Cloud
 - **Managed From** and **Managed To** are usually hidden by the Cloud Service Provider
 - **Virtual Servers** can still contain Linux/Open Source
 - Service Provider may choose proprietary, home grown, or open management stack
 - Service Provider chooses they hypervisor on the **Managed To** nodes
 - Hypervisor is “mostly” hidden from the end user
 - Hypervisor often hides any hardware differentiation
- Private Cloud
 - Cloud owner selects **Managed From, Managed To** components
 - Cloud owner typically owns process for creating **Virtual Servers**
 - Cloud owner may choose to augment cloud management software to address any special concerns, e.g. security models, backups, performance, enterprise integration, etc.
 - Cloud owner chooses the hypervisor for their managed nodes – and more often is choosing/evaluating multiple hypervisors
 - Cloud owner more often selects hardware based on business needs

Where does Linux and Open Source fit in today?

- Managed From
 - Linux is almost non-existent here
 - Primary “Competition” is:
 - vCenter, vCloud on Windows
 - Pretty much all other competitors are a distant #2 at best
 - OpenStack is a promising alternative on the horizon, plus Eucalyptus, cloud.com CloudStack, CrowdDirector, IBM SmartCloud solutions, Surgient Platform, Enomoly ECP, Abiquo, Nimbula Director, Citrix XenCloud, Joyent Smart OS, RHEV, oVirt, CA/3Tera Applogic, Abiquo, Platform ISF, RightScale, etc
- Managed To
 - Linux / KVM is an occasional hypervisor
 - Used in the IBM Public Cloud (SmartCloud Enterprise) as an example
 - This is the layer sometimes referred to as the “Platform for the Cloud”
 - Most often VMware ESXi, Hyper-V, Xen, etc.
- Virtual Servers
 - A strength of Linux
 - Hard to measure how many virtual servers use Linux as the base operating system
 - But very easy/common in Amazon EC2, IBM SmartCloud Enterprise, OpenStack configurations, etc.

Linux and KVM on the Managed To Servers

- Linux + KVM is clearly qualified and ready to be the hypervisor in the Cloud today
 - Performance is competitive
 - Features are generally competitive
 - Linux does reasonably well in the key metric of \$/VM-hour
 - Often measured in terms of virtual servers per core
 - Measured consistently at the top on benchmarks such as SpecVirt
 - Features inherited from Linux bring a wealth of capability
 - But many workloads and administrators expect same flexibility as with real hardware
 - Strong virtual server isolation
 - Assurance of tamper-free/tamper-evident security
 - Data isolation through encryption
 - Same flexibility in setting up networks as with physical cable access to switches
 - Easy access to storage as part of deployment of a workload
 - Long list of more specific challenges
 - Overcommit of key resources: processor, memory, network, storage
 - Attaching more vlans to workloads and isolated networks spanning LAN segments
 - Long term job scheduling for virtual servers
 - Management hooks (e.g. libvirt, VDSM, oVirt, CIM agents)
 - SELinux isolation, TPM/vTPM, attestation
 - Addressing virtualization overhead and impacts to performance in server, storage, networking
 - Specifying and enforcing SLO characteristics (availability, co/anti-co-location, restart on fail, etc.)

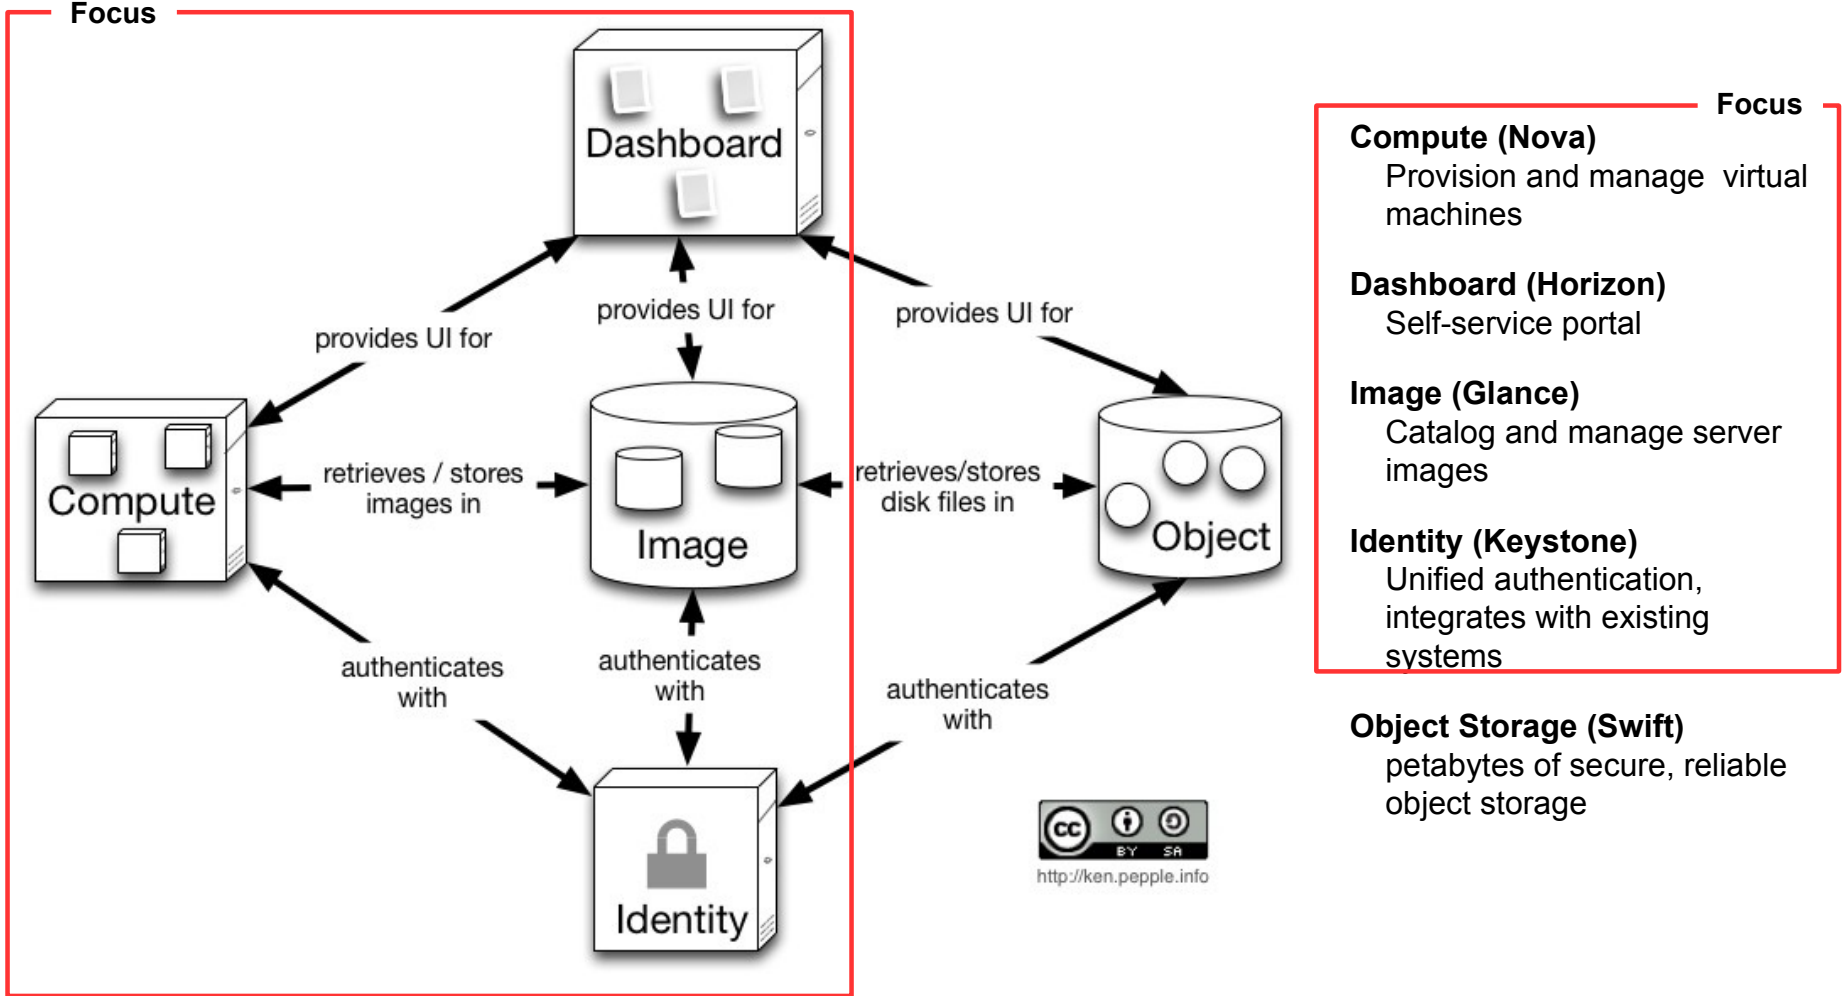
Linux/KVM and The Management Stack

- Linux is probably not a key differentiator
- But Open Source could be a viable player
 - Leading contender is probably OpenStack
 - RHEV and oVirt are also interesting components
- Two ways to look at Linux/KVM and Open Source and the Management Stack
 - What is installed on the Management Nodes
 - How do the Management Nodes interact with the Managed Nodes?
- Management Nodes may be some combo of:
 - Windows, vCenter/vCloud,
 - Linux, RHEV/oVirt/OpenStack/etc.
- Interactions need to cover everything from:
 - End User and Administrator front end GUI/REST API/CLI/etc.
 - Workload placement (which node is least busy, enables appropriate SLA, energy awareness, specific server pool, ...)
 - Virtual Server configuration
 - Post-install software configuration
 - Migration management
 - Workload snapshots
 - Metering and Monitoring
- Standardization is fairly weak today

Cloud Management Scope

- Generic taxonomy typically includes:
 - Server Management
 - Storage Management
 - Network Management
 - Image Management
- Some models break out a concept of OSS vs. BSS
 - Operational Support Services cover activities specific to day to day operation of the cloud
 - Business Support Services cover activities related to customer management, billing and metering, order management, etc.
 - OSS is often provided locally
 - BSS is often globally centralized for a distributed Cloud

OpenStack is comprised of five core projects that form a complete IaaS solution

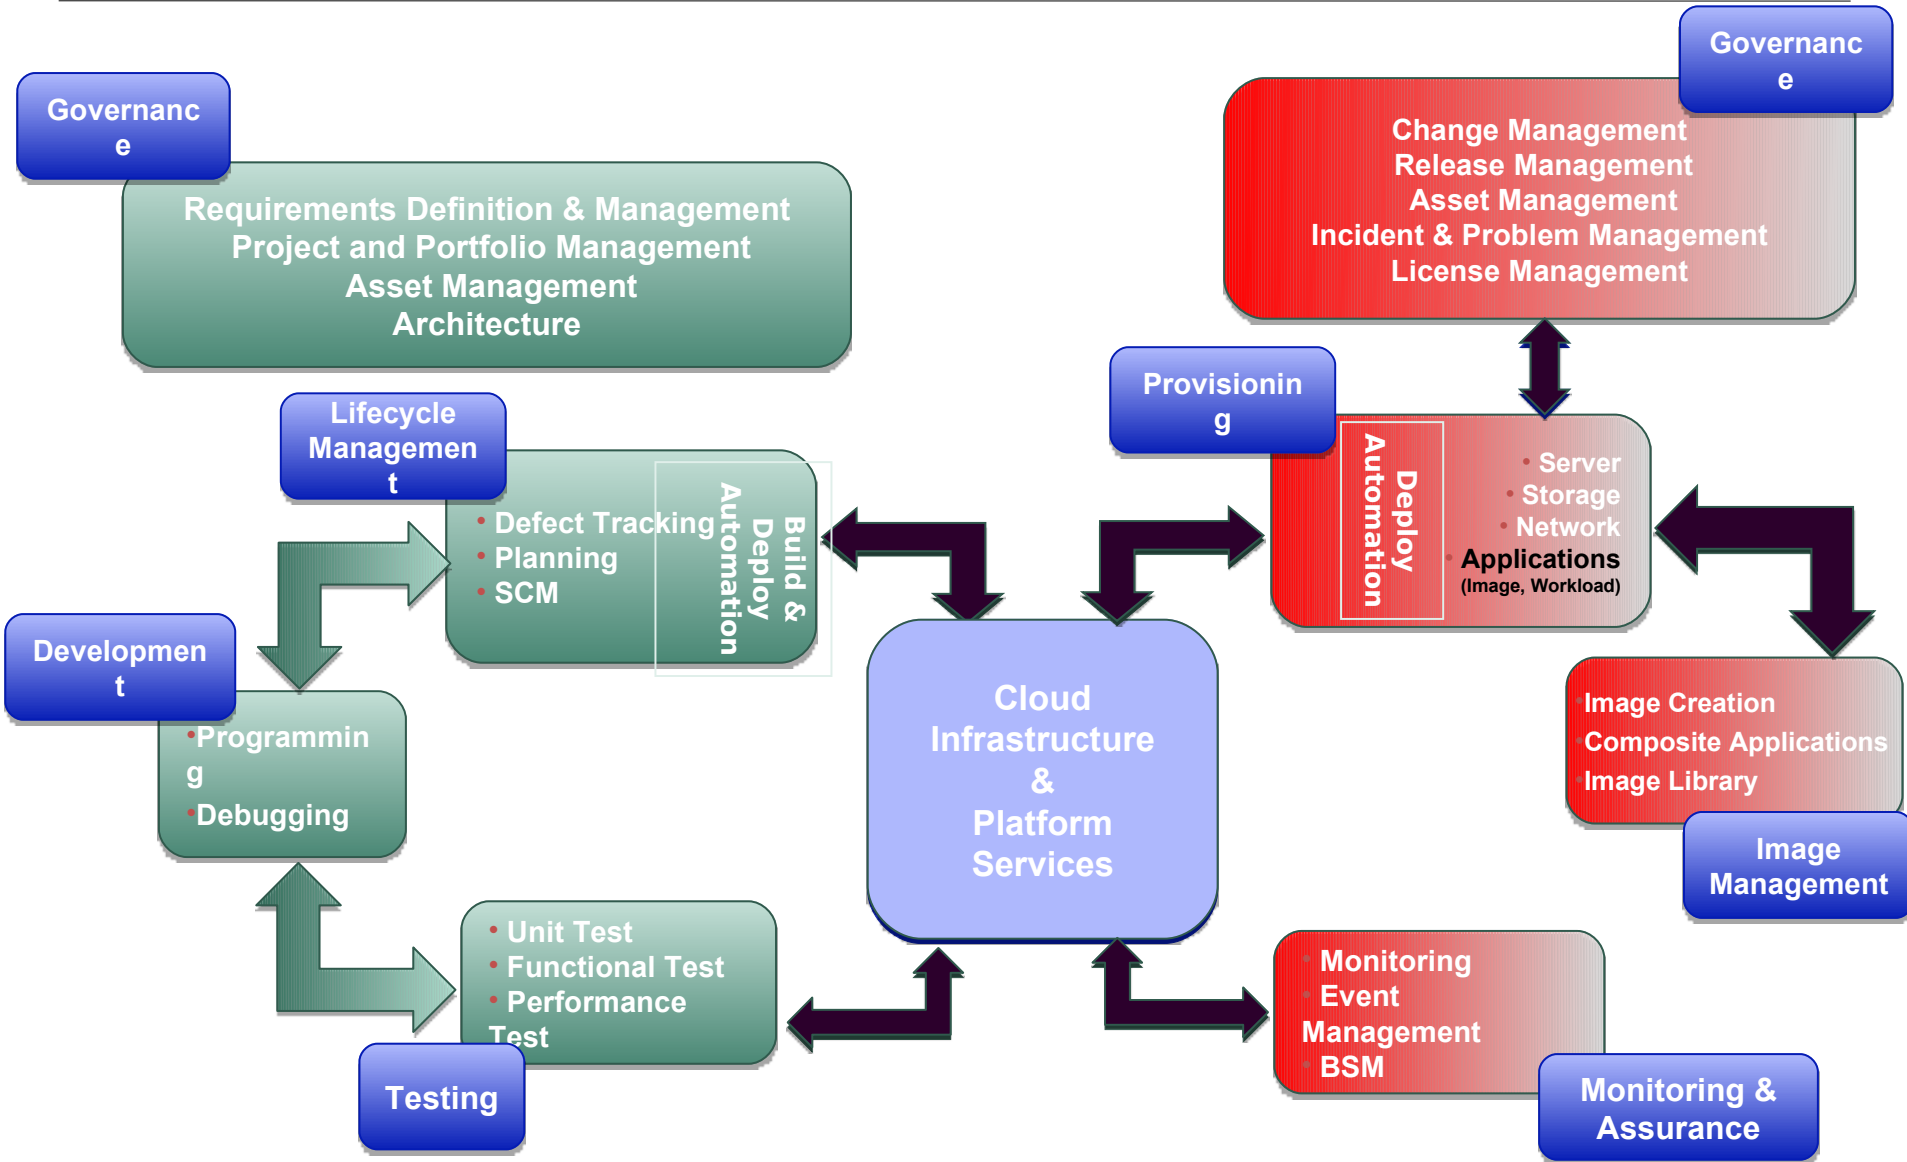


Sources:
<http://ken.pepple.info/openstack/2012/02/21/revisit-openstack-architecture-diablo/>

The Larger Cloud Ecosystem



Development & Operational Lifecycle Integration towards DevOps



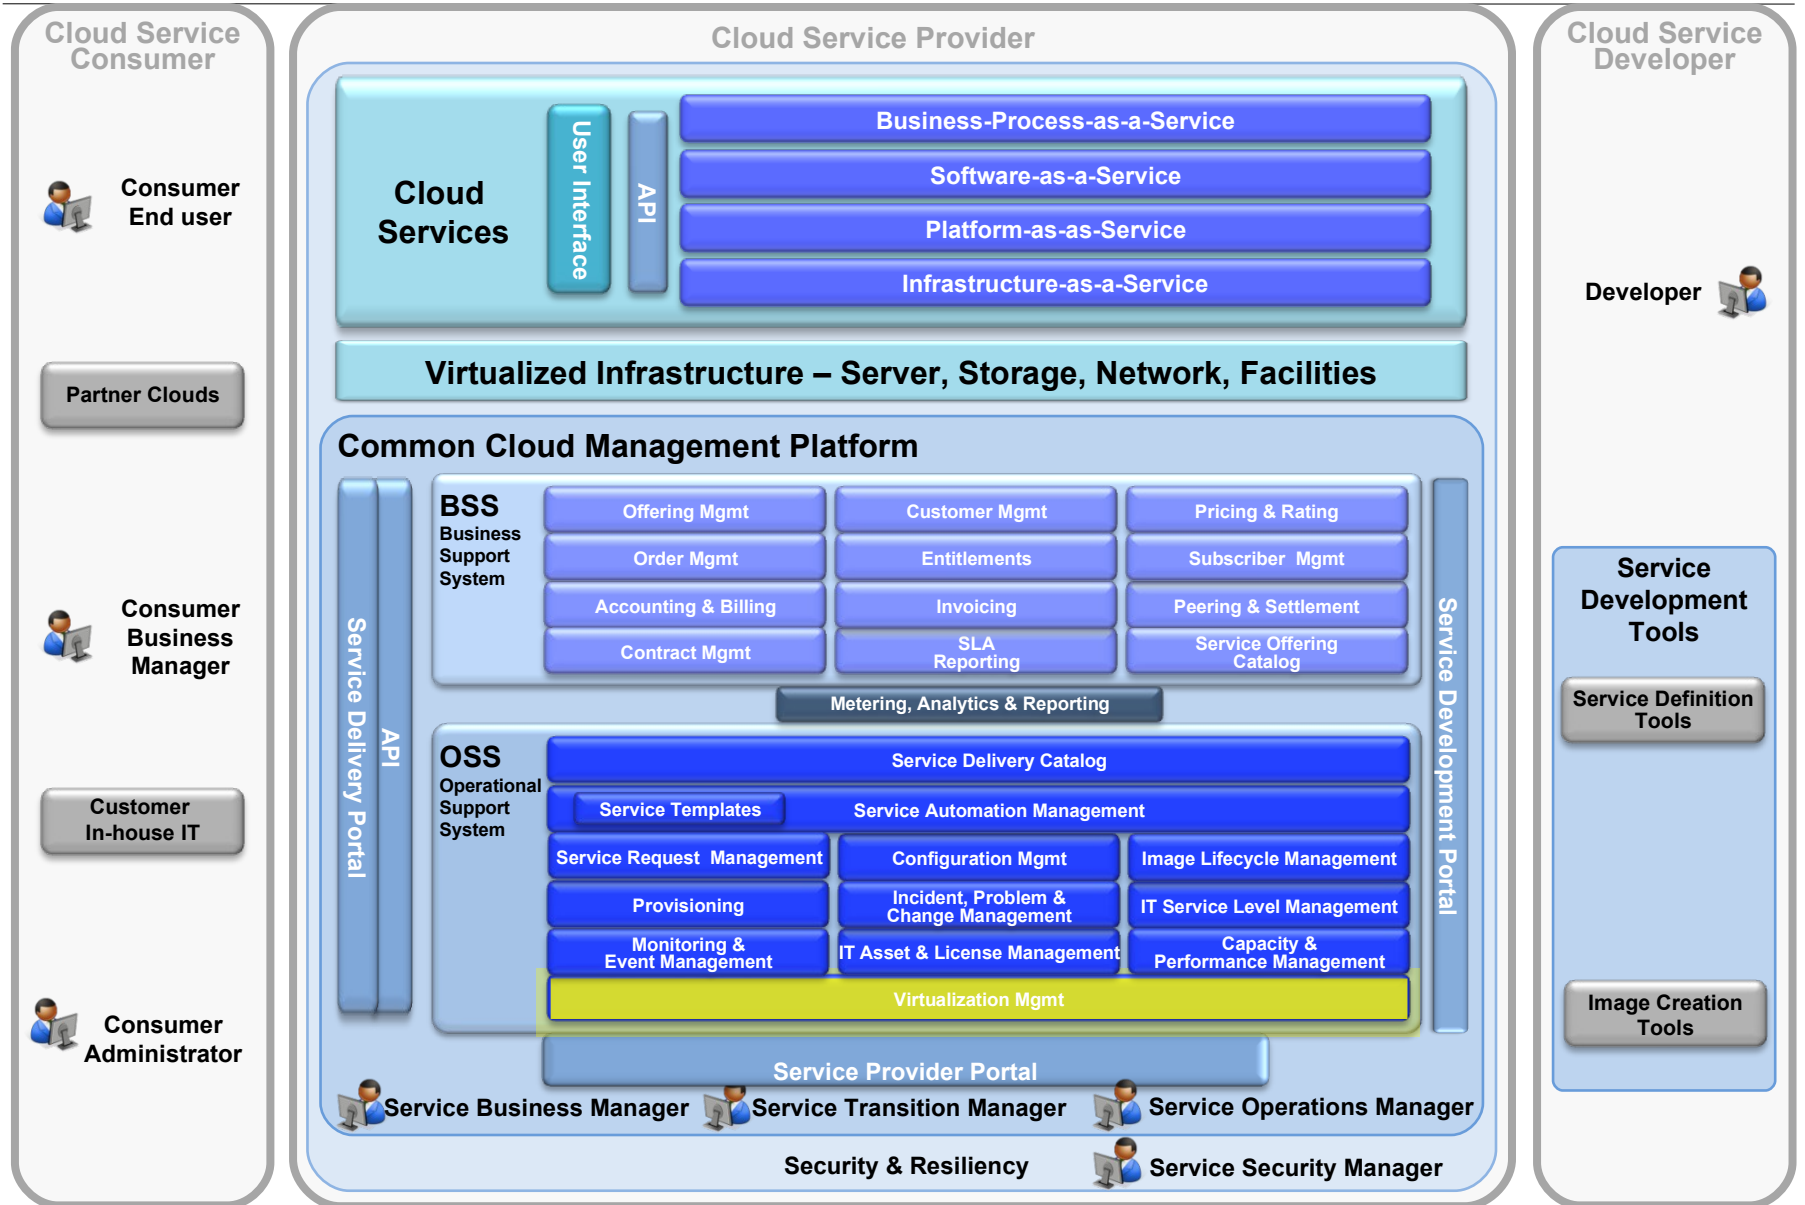
Gerrit Huizenga

Building a cloud Ecosystem

- Hardware, Virtualization, Virtualization Management, Cloud Management
- Management Standards
- Billing & Metering
- Performance Monitoring
- Security Management
- Multi-Tenancy Support
- Managed Service Providers & Business Partners
- Support for PaaS, SaaS, Business-Process as a Service environments
 - All should be able to build on top of a good IaaS solution
 - e.g. Workload Deployer, Complex workload management, end-to-end SLA and performance monitoring, etc.
- (Full OSS list)
- (Full BSS list)

Example Reference Architecture

CCMP Reference Architecture: Architecture Overview Details



The Open Cloud Ecosystem

The Open Virtualization Alliance

Comprised of leading virtualization, datacenter, and cloud solution providers, the Open Virtualization Alliance's goal is to help:

- Increase overall awareness and understanding of Kernel-based Virtual Machine (KVM).
- Foster the adoption of KVM as an open virtualization alternative to proprietary solutions.
- Accelerate the emergence of an ecosystem of third-party solutions around KVM.
- Encourage interoperability, promote best practices, and highlight examples of customer successes.
- <http://www.openvirtualizationalliance.org/>

The Open Cloud Manifesto

- <http://www.opencloudmanifesto.org/>
- Discussion Groups, Use Cases, etc.

DMTF

- Open Standards evolving around cloud – OVA, TOSCA, Cloud

IEEE

- Conferences, Cloud Architecture, Definitions

Linux Foundation: CloudOpen

Summary

- Clouds are a huge segment of the market and that segment is growing rapidly
- Linux and Open Source have 3 major areas in which they may factor
 - Management Stack, Virtualization Layer, Virtualized Servers
- Lots of players in the management stack today
 - Getting close to time for consolidation, elimination of weaker players, etc.
 - Open Source could be a big beneficiary of any consolidation
- Ecosystem is very large
 - Lots of places where standards and open source projects are likely to play