
Cloud Management Platform

Service Management for the Cloud



Course Agenda Day 2 – Cloud Service Management

- Cloud Service Management (Gerd Breiter)
- Cloud Customer Projects / Cloud Platform Exploitation (Michael Behrendt)
- Cloud Security (Michael Waidner)

Agenda

- *Introduction*
- *Lifecycle of a Cloud Service*
- *Cloud Computing Reference Architecture*
- *Cloud Computing Management Platform*
- *Selected Management Areas*
 - *Service Automation*
 - *Image Mgmt*
 - *Usage and Accounting*
 - *Cloud and Green IT*
- *Exploitability of the Cloud Management Platform*
- *Multi Tenancy*
- *Design Methodology for the Cloud Management Platform Reference Architecture*
- *Summary*

Some macroscopic observations...

In 2001, there were 60 million transistors for every human on the planet ...

... by 2010 there will be 1 billion transistors per human...

... each costing 1/10 millionth of a cent.

Worldwide mobile telephone subscriptions reached 3.3 billion in 2007

In 2005 there were 1.3 billion RFID tags in circulation...

... by 2010 there will be 33 billion.

One billion camera phones were sold in 2007, up from 450 million in 2006 ...

**An estimated 2 billion people will be on the Web by 2011 ...
... and a trillion connected objects – cars, appliances, cameras, roadways, pipelines – comprising the "Internet of Things."**

Cloud Computing: The next step in the evolution of IT

1. Centralized Computing: 1960 –

- Optimized for sharing, industrial strength, systems management, ...
- Managed by central IT organization
- Back office applications involving transactions, shared data bases, ...
- Mainframes, supercomputers, minicomputers, ...

2. Client/Server: 1985 –

- Optimized for low costs, simplicity, flexibility, ...
- Distributed management across multiple departments and organizations
- Large numbers of PC-based applications
- PC-based clients and servers, Unix, Linux, ...

3. Cloud Computing: 2010 –

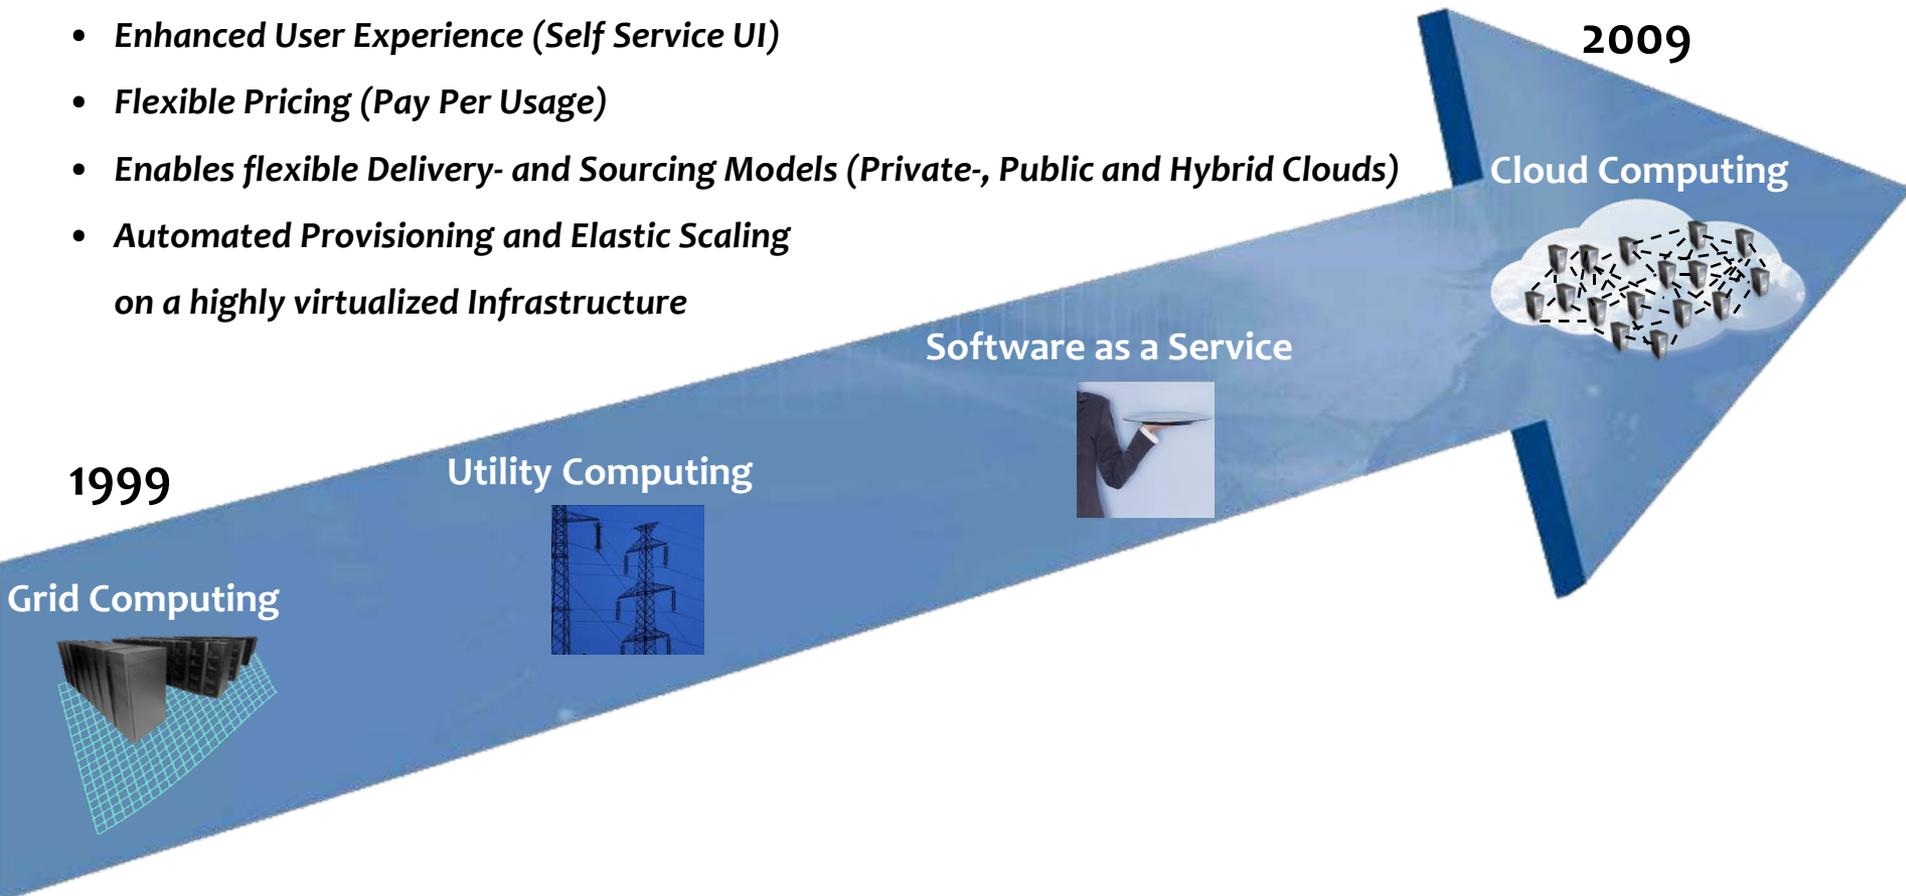
- **New consumption and delivery model**
- Optimized for massive scalability, delivery of services, ...
- Centralized model, hybrid service acquisition models
- Supports huge numbers of mobile devices and sensors
- Internet technology-based architecture

Just like introducing the Client/Server model impacted almost everything we did in IT (operation IT, developing applications, ...), Cloud computing has severe impact on the IT industry

“Clouds will transform the information technology (IT) industry... profoundly change the way people work and companies operate.”

The Economist

- *A new paradigm for consumption and delivery of IT based services*
- *Enhanced User Experience (Self Service UI)*
- *Flexible Pricing (Pay Per Usage)*
- *Enables flexible Delivery- and Sourcing Models (Private-, Public and Hybrid Clouds)*
- *Automated Provisioning and Elastic Scaling on a highly virtualized Infrastructure*



Cloud Computing Layers



Business Process-as-a-Service

Examples



Application-as-a-Service

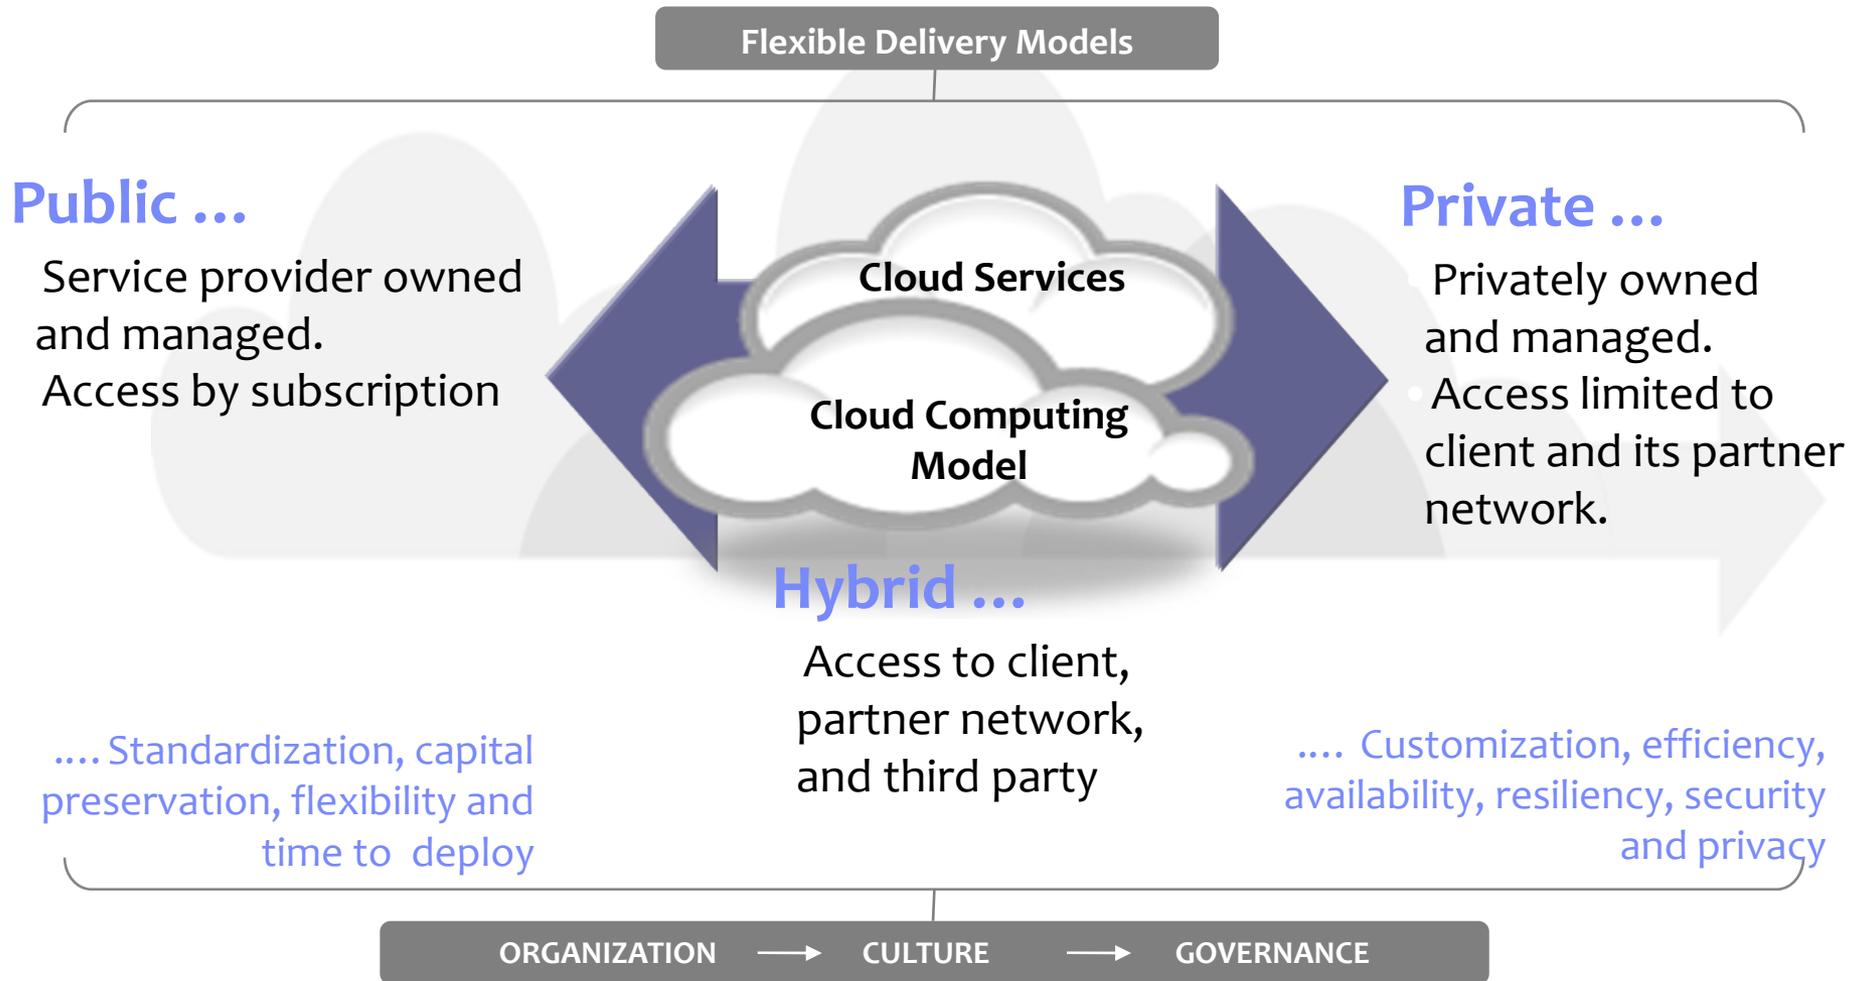


Platform-as-a-Service



Infrastructure-as-a-Service





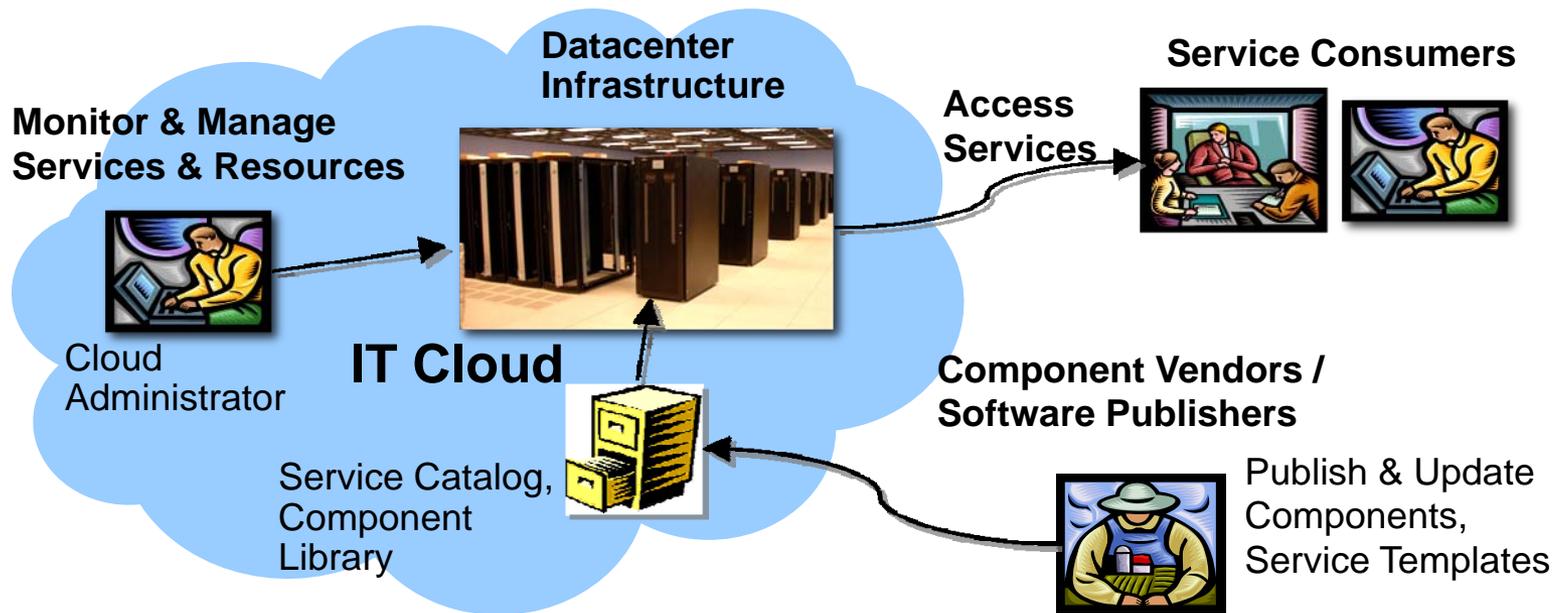
Cloud Computing...

... is a user experience and a business model

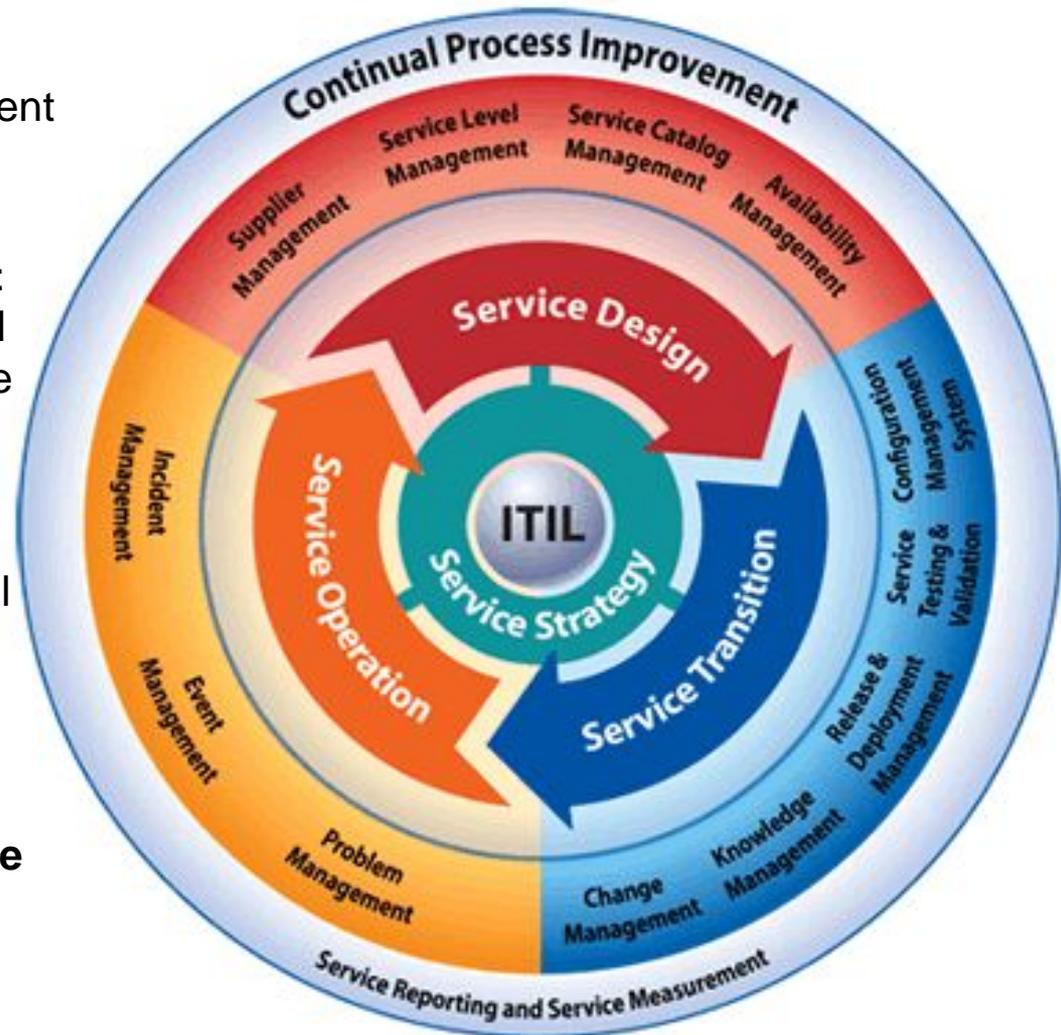
- Cloud computing is an emerging style of computing in which applications, data, and IT resources are provided as services to users over the network.

... is a infrastructure management methodology

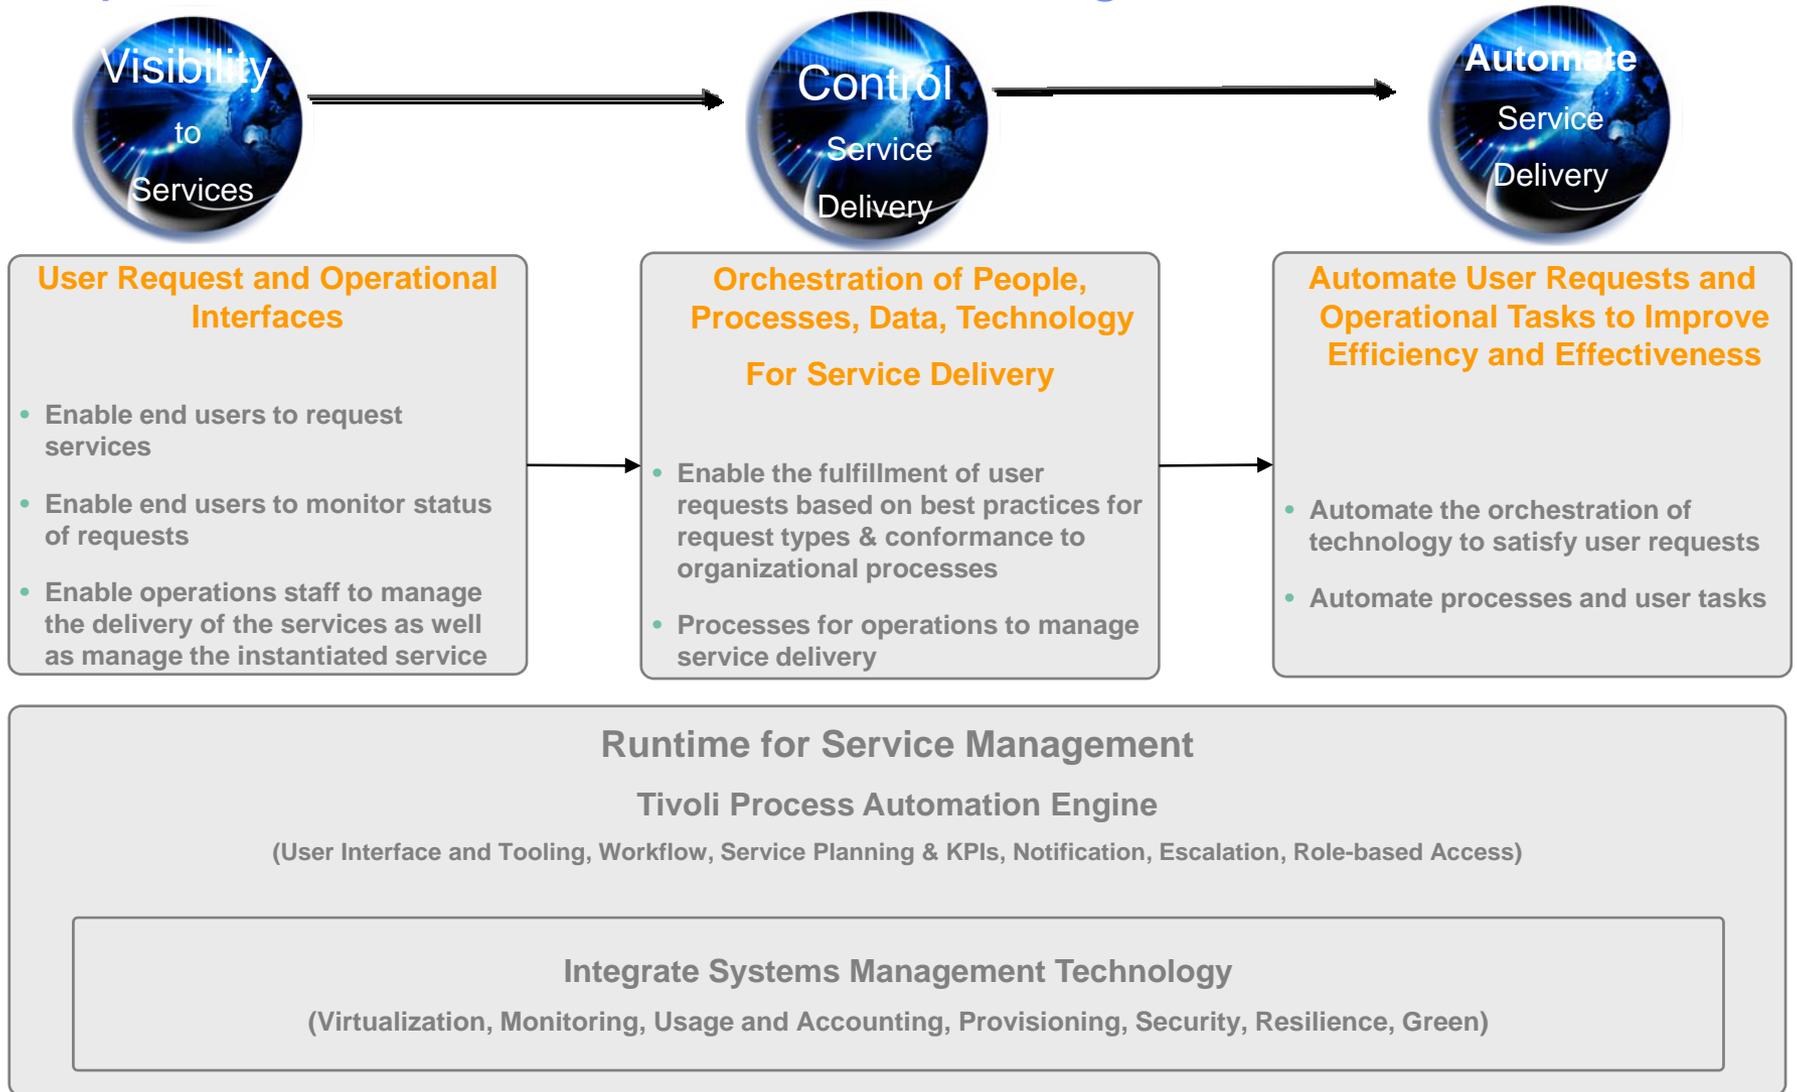
- Cloud computing is a way of managing large numbers of highly virtualized resources such that from a management perspective, they can be automatically aggregated to deliver services.



- The Information Technology Infrastructure Library (ITIL) is a set of concepts and practices for managing Information Technology (IT) services, IT development and IT operations.
- The objective of **Service Management** is the coordination of specific, technical and organizational resources to provide added value to the customer through services.
- Service management according to ITIL is a framework and as such contains all the necessary and specialized organizational capabilities available for the generation of added value to the customer as services.
- Included are **Service Strategy, Service Design, Service Transition, Service Operation** and **Continual Service Improvement**.



Capabilities needed for Service Management



What's so different about cloud-like Service Management? – Changes in orders of magnitude

Capability	From	To
Server/Storage Utilization	10-20%	70-90%
Self service	None	Unlimited
Provisioning	Weeks	Minutes
Change Management	Months	Days/Hours
Release Management	Weeks	Minutes
Metering/Billing	Fixed cost model	Granular
Payback period for new services	Years	Months

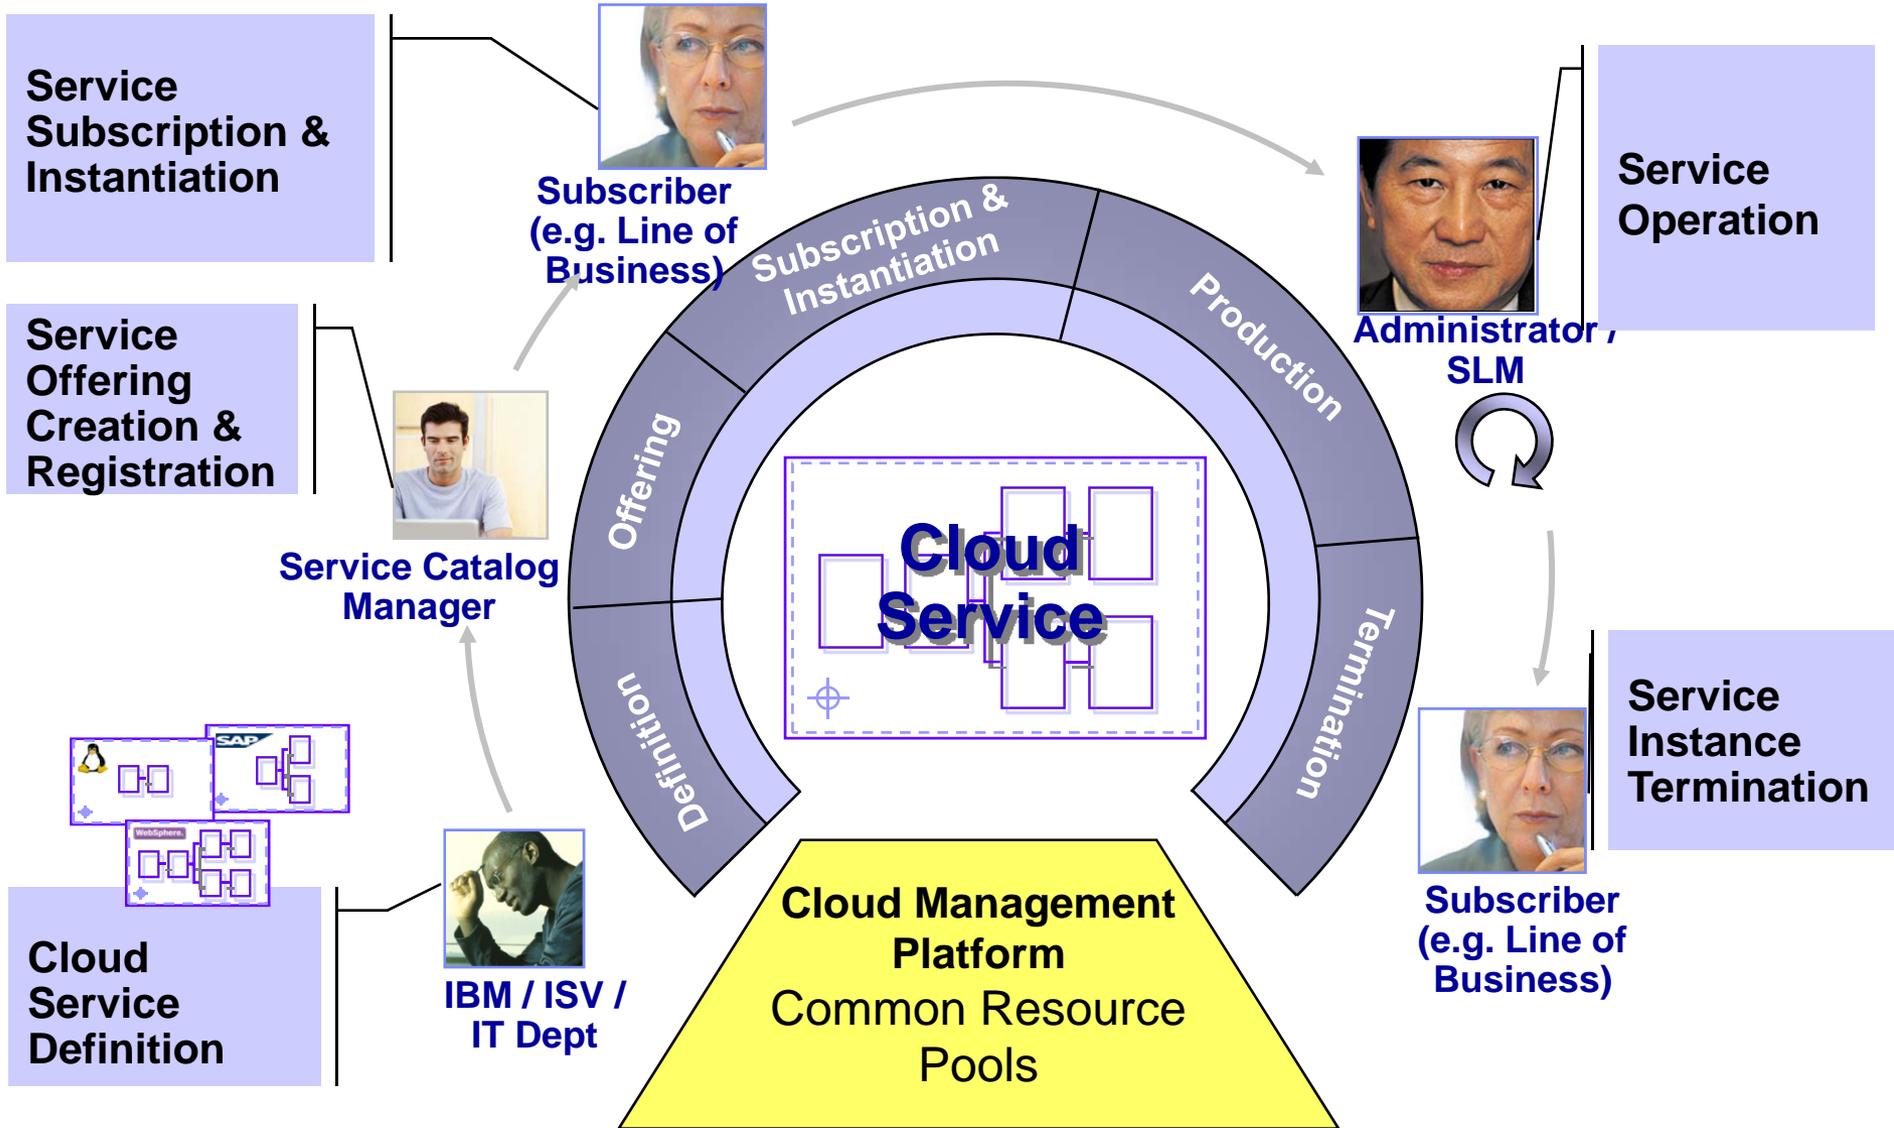
←

Traditional Service Management

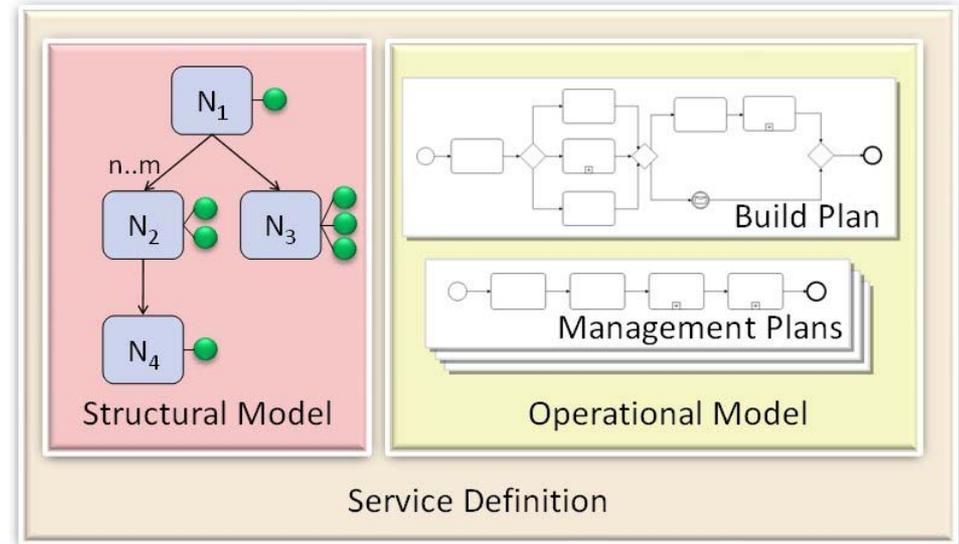
→

Cloud-like Service Management

Lifecycle of a Cloud Service



- **Service Definition provides a model for managing Cloud Services throughout their complete life cycle:**
 - *Initial Deployment* of a service instance
 - *Operational management* of a service instance (e.g. modify capacity, patch management, upgrades, incident and problem management, etc.)
 - *Termination* of a service instance
- **Service Topology Template:** structural model of a service, *components* and their *relationships*, plus *operations* that can be invoked on service components
- **Build- and Management Plans:** Process model of how to *set up*, *manage* and *terminate* a service



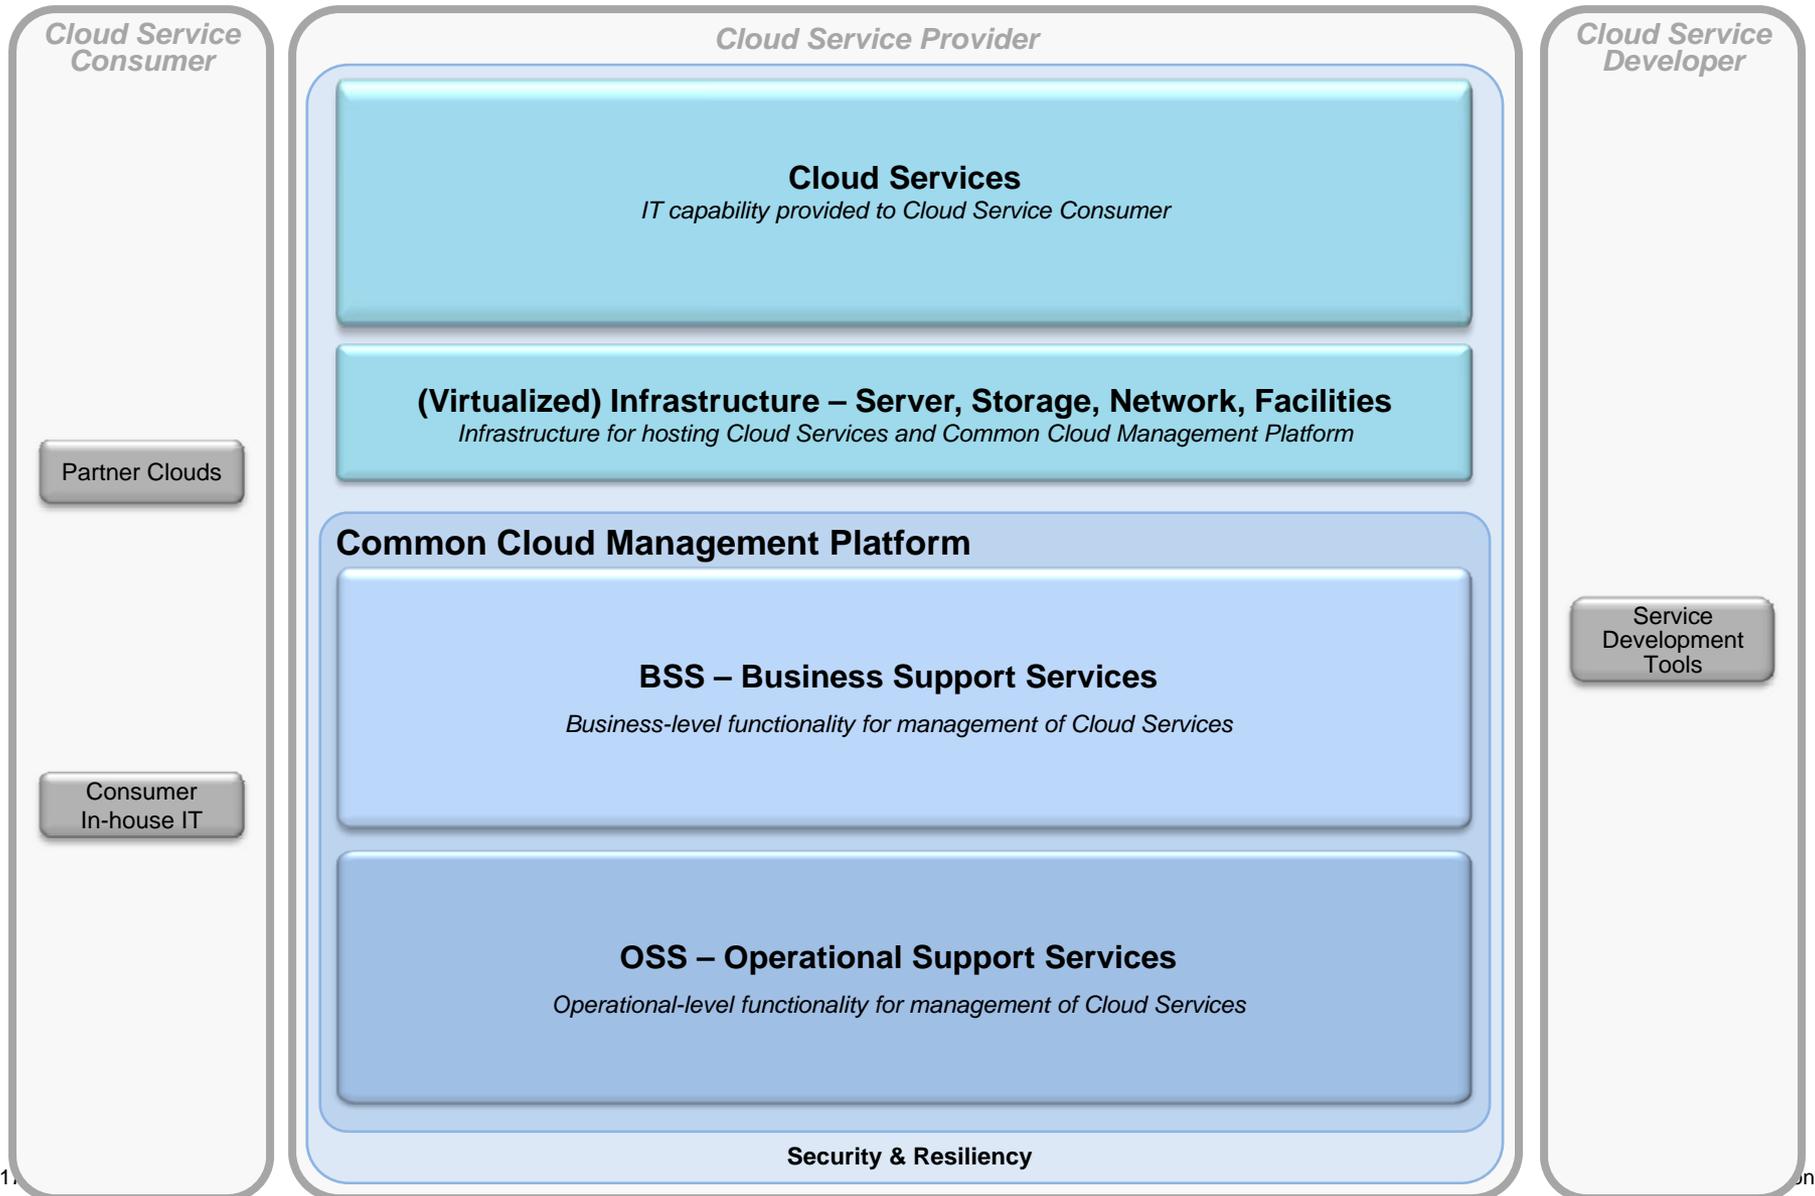
- Initial Systems Instantiation
 - Bring up servers, assign virtual storage, bring-up and configure network components, configure network zones, install, configure and start middleware and apps
- Change resources depending on actual need
 - Provision / de-provision resources, adapt size of virtual resources (CPU, Memory, storage capacity)
 - Enables more efficient use / higher utilization of resources like memory
- Clone a system
 - Create a test system as an exact clone of a production system including DB content
 - Clone could reside in different location
 - Enables temporary creation of non-production systems vs. keeping and maintaining rarely used systems “forever“
- Termination
 - Shut-down all systems and transfer back to free pool
 - Enables easy re-use of resources that are not needed any more

→ All these scenarios are in the context of an existing contract between provider and consumer and its SLAs
→ Metering Records reflect all changes and enable usage based accounting and billing
→ All these scenarios are initiated through one user interface and executed 100% operator-less

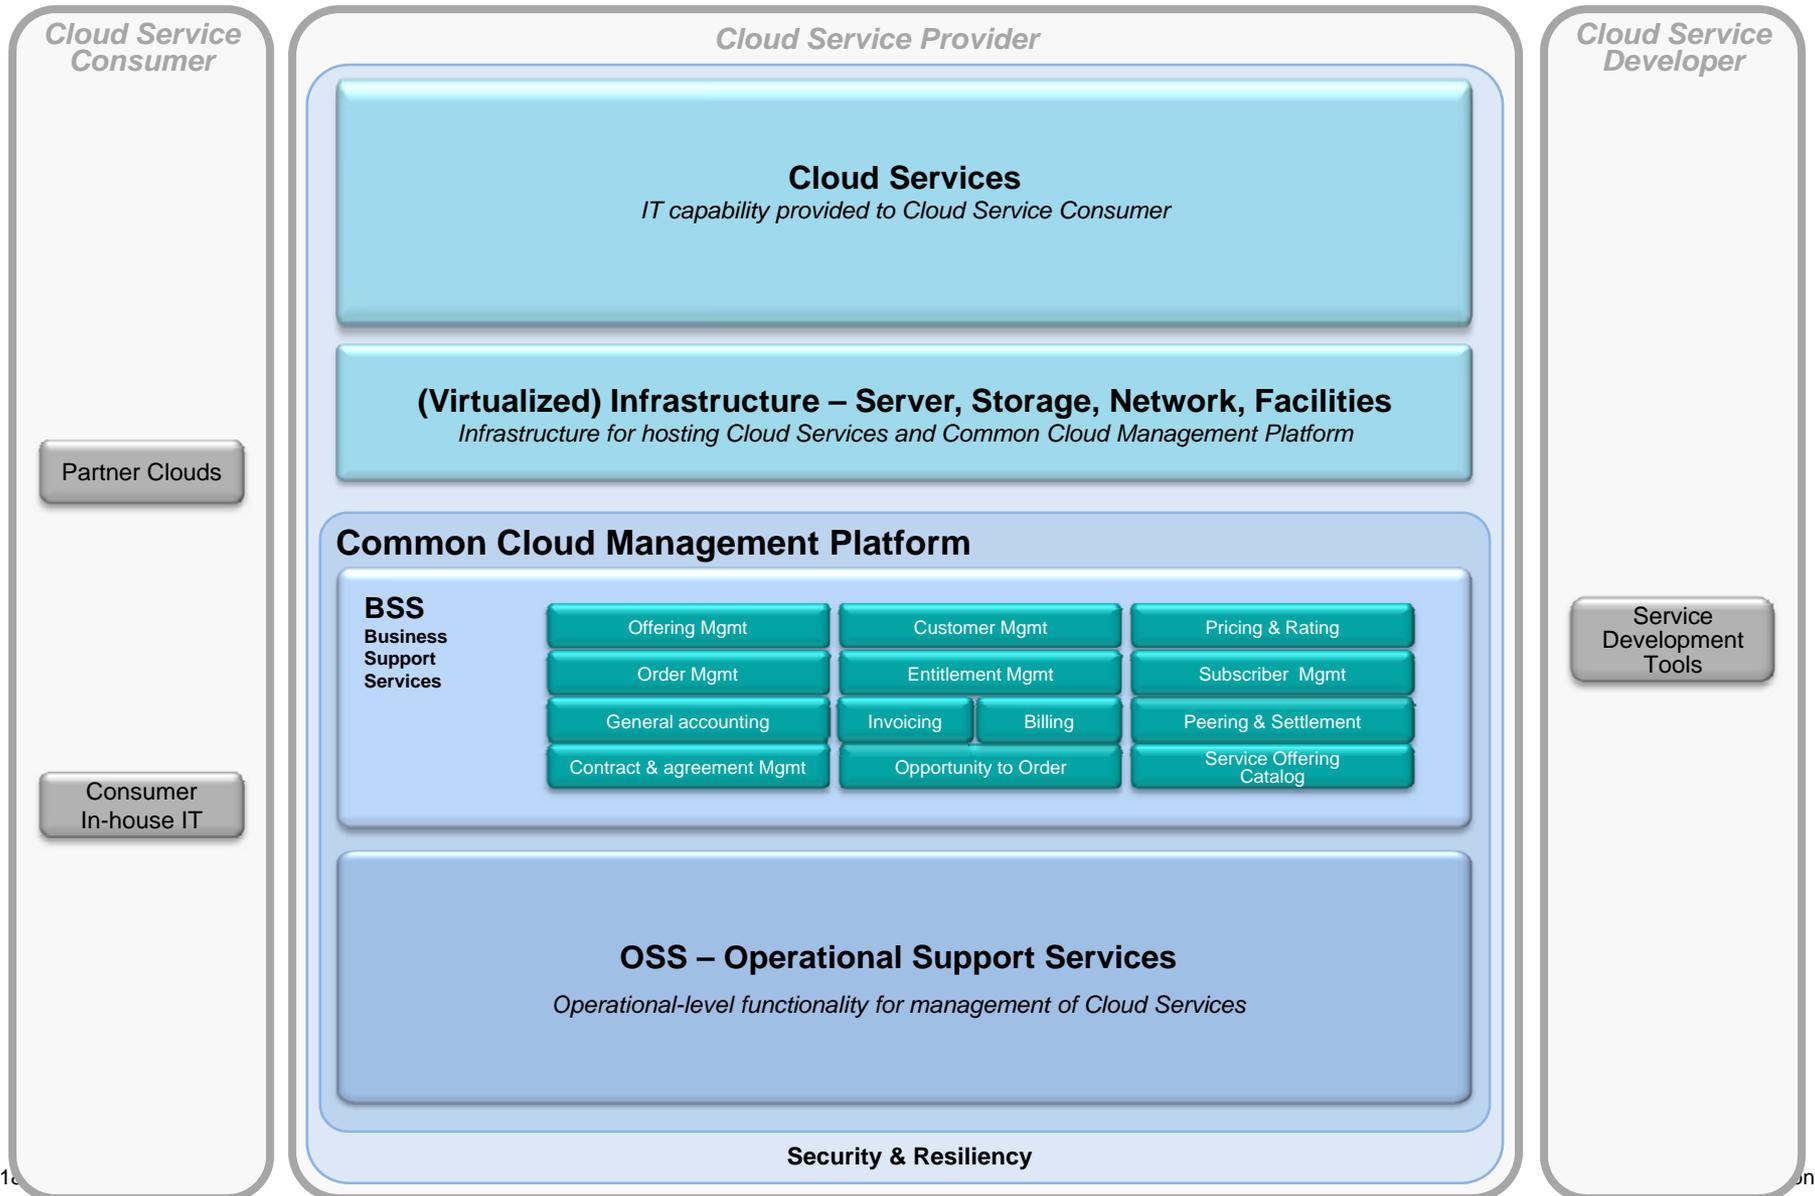
Overview – Cloud Computing Reference Architecture

- 1. The IBM Cloud Computing Reference Architecture (CC RA) is structured in a modular fashion (similar to the SOA Reference Model)**
 - On its highest level of abstraction, it defines a base set of architectural elements, which are refined to the next level of detail
 - This modular approach allows refinement of the CC RA architectural elements independent from each other by the respective SMEs.
- 2. The IBM Common Cloud Management Platform Reference Architecture (CCMP RA) is the reference architecture for the CCMP being one fundamental architectural elements of the IBM CC RA.**

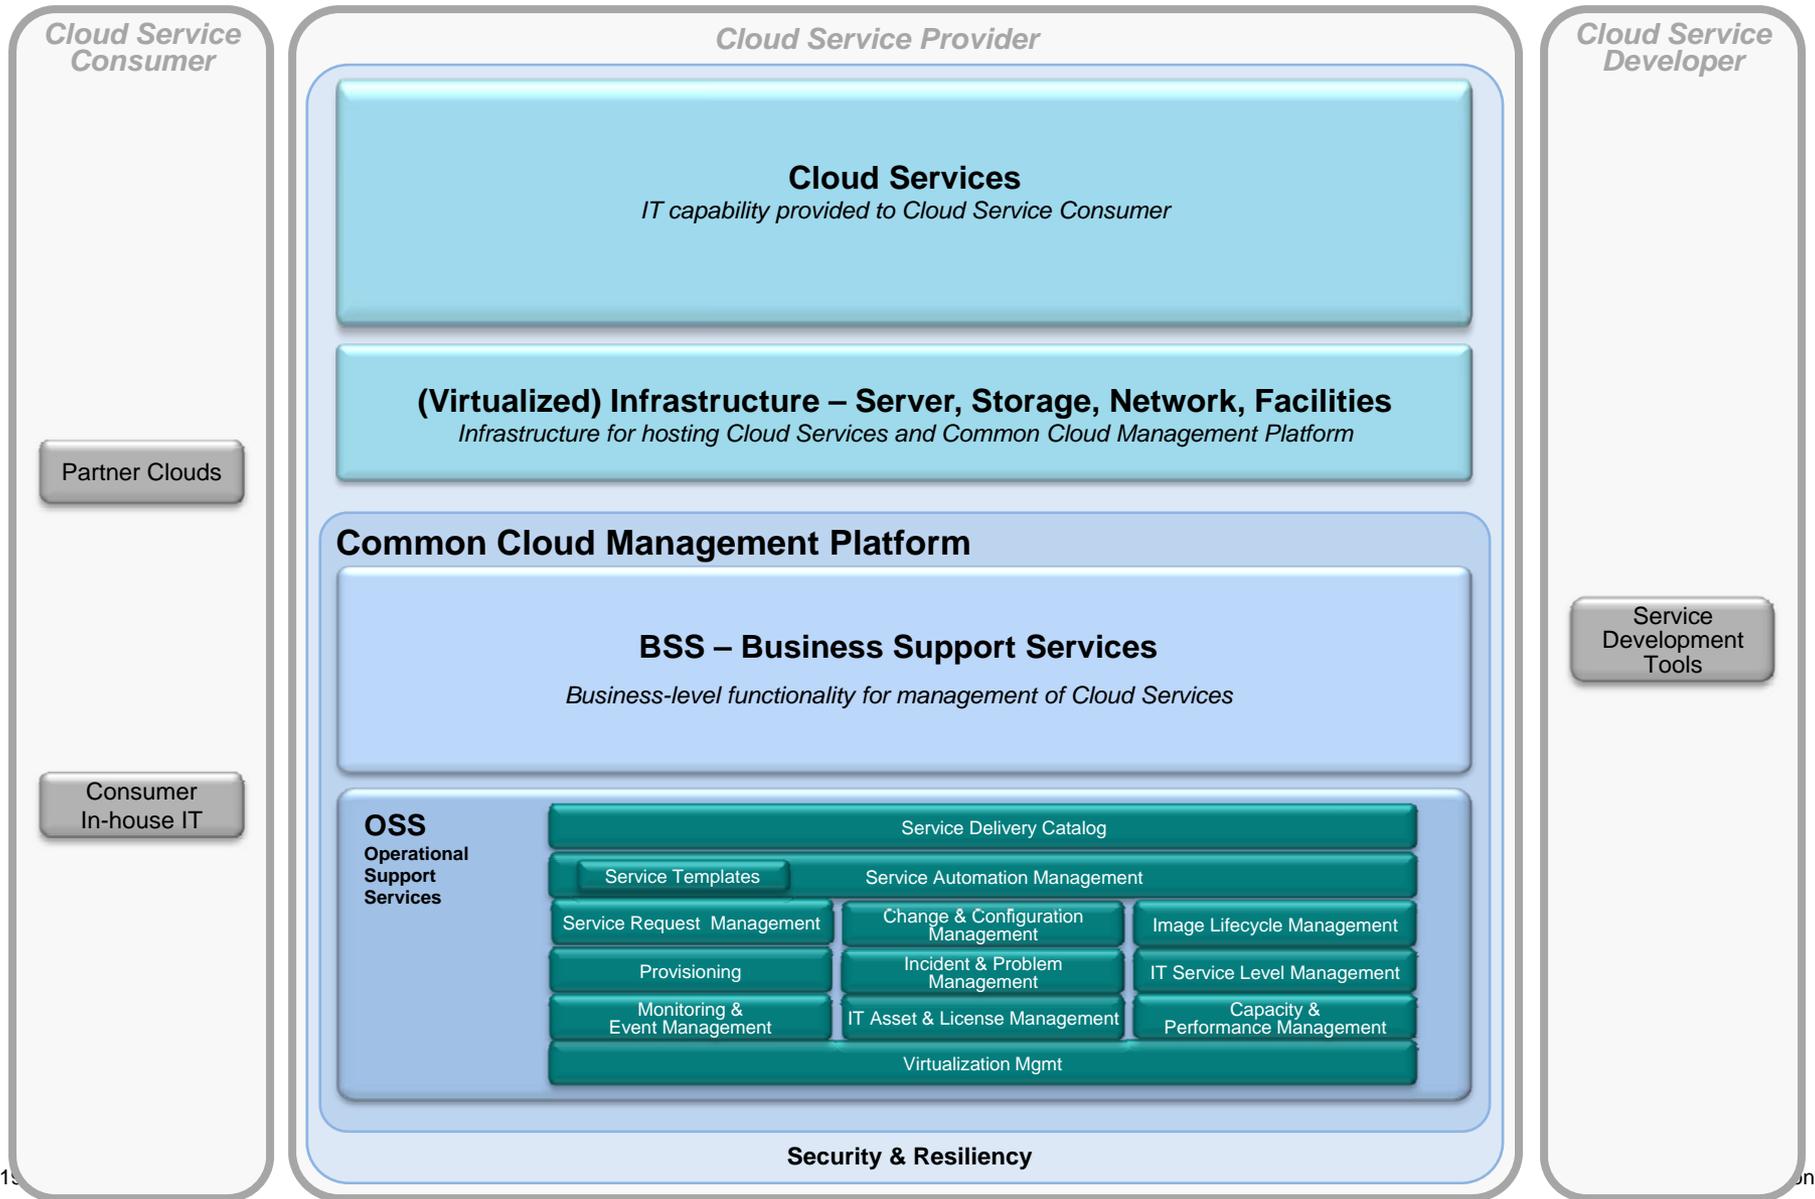
Cloud Computing Reference Architecture (CC RA) – Overview



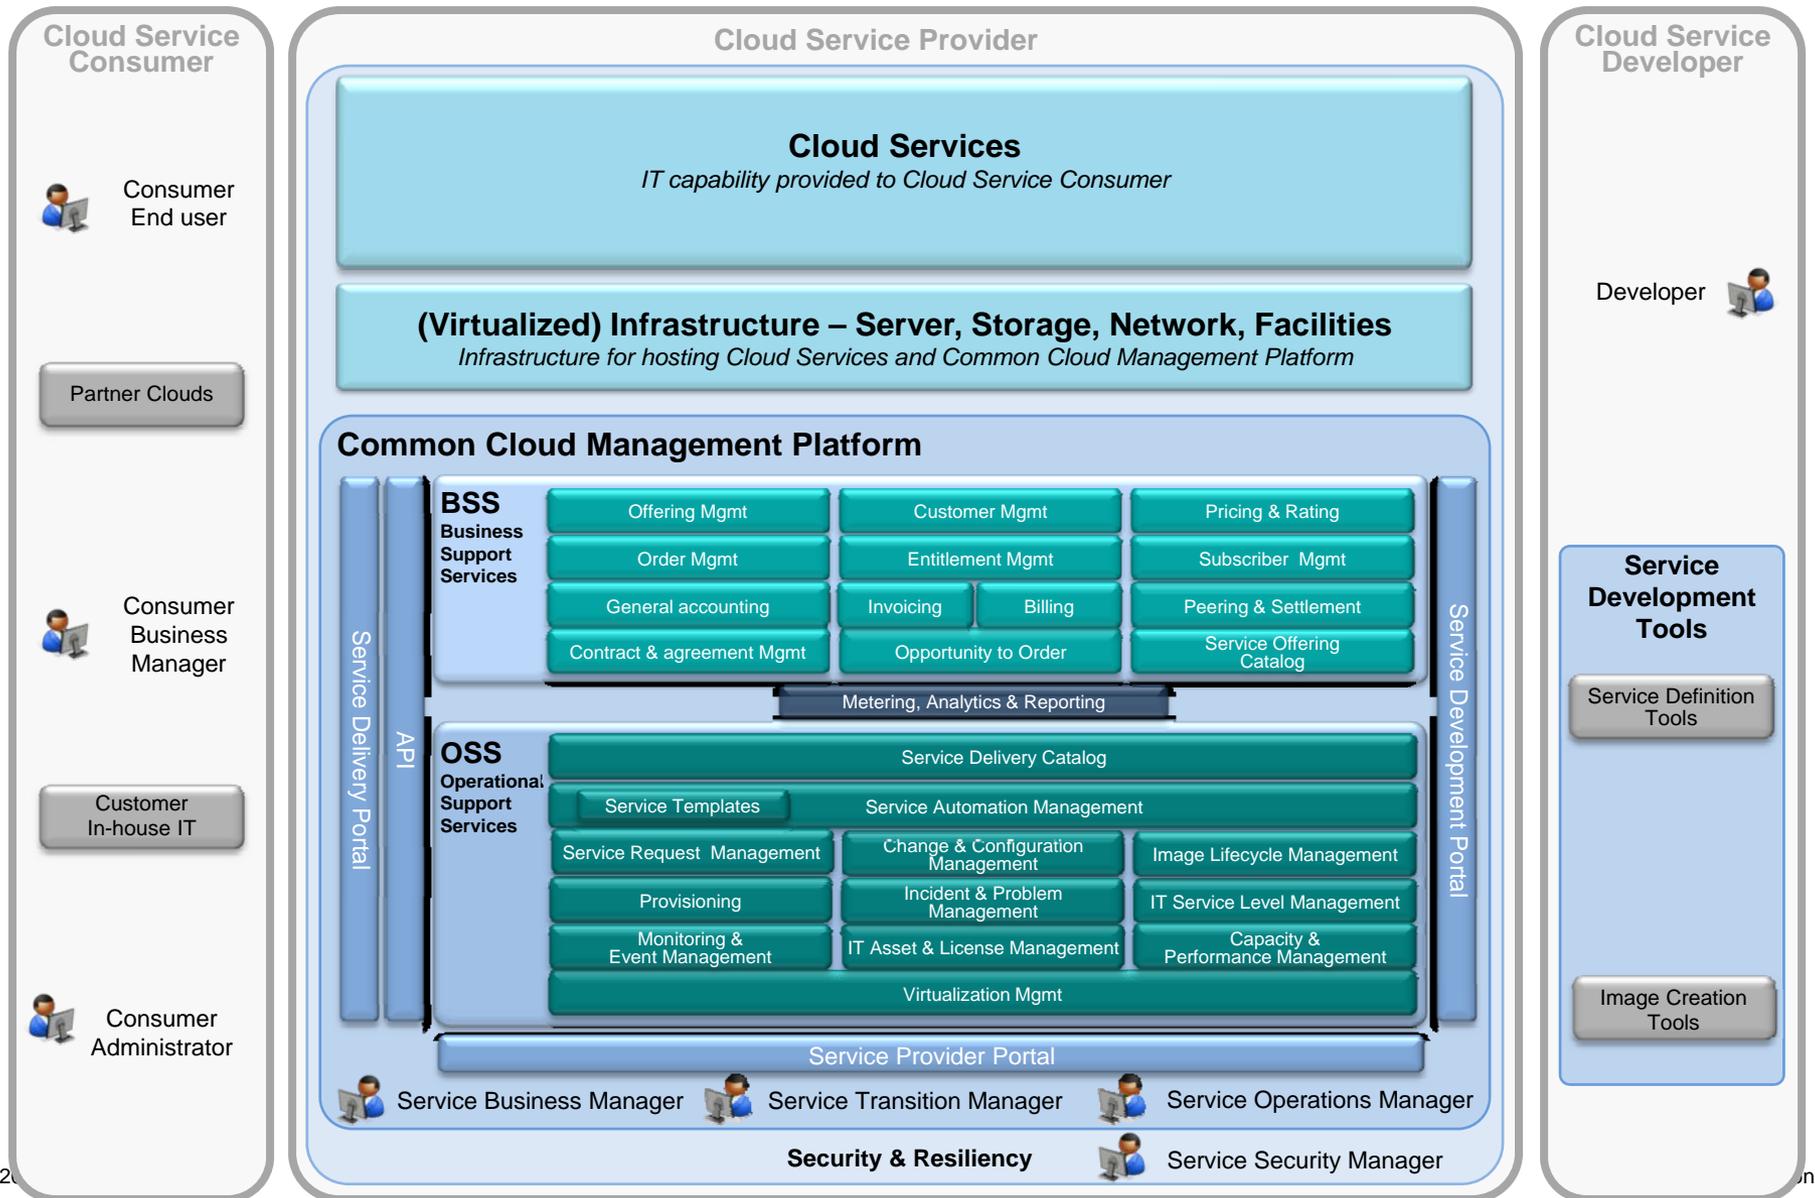
Cloud Computing Management Platform Reference Architecture – BSS Details



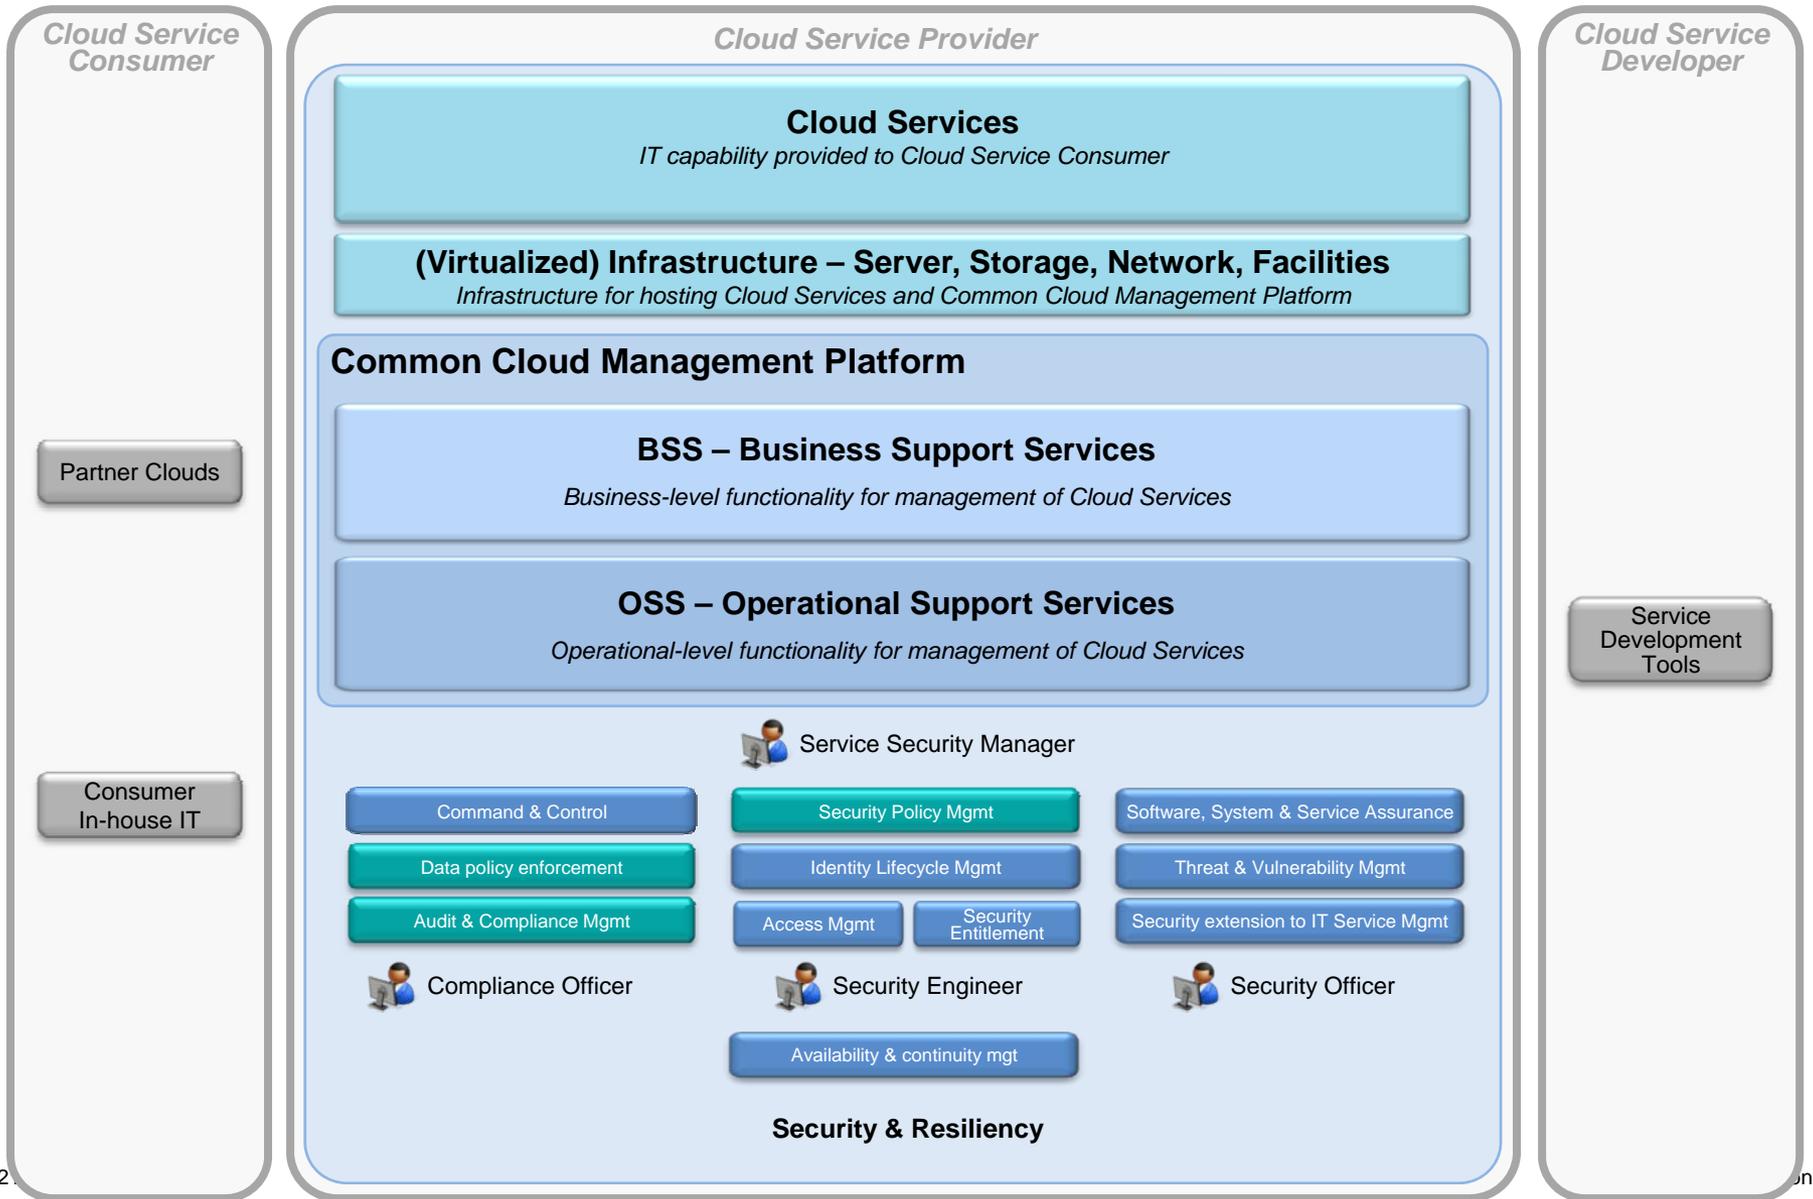
CCMP RA – OSS Details



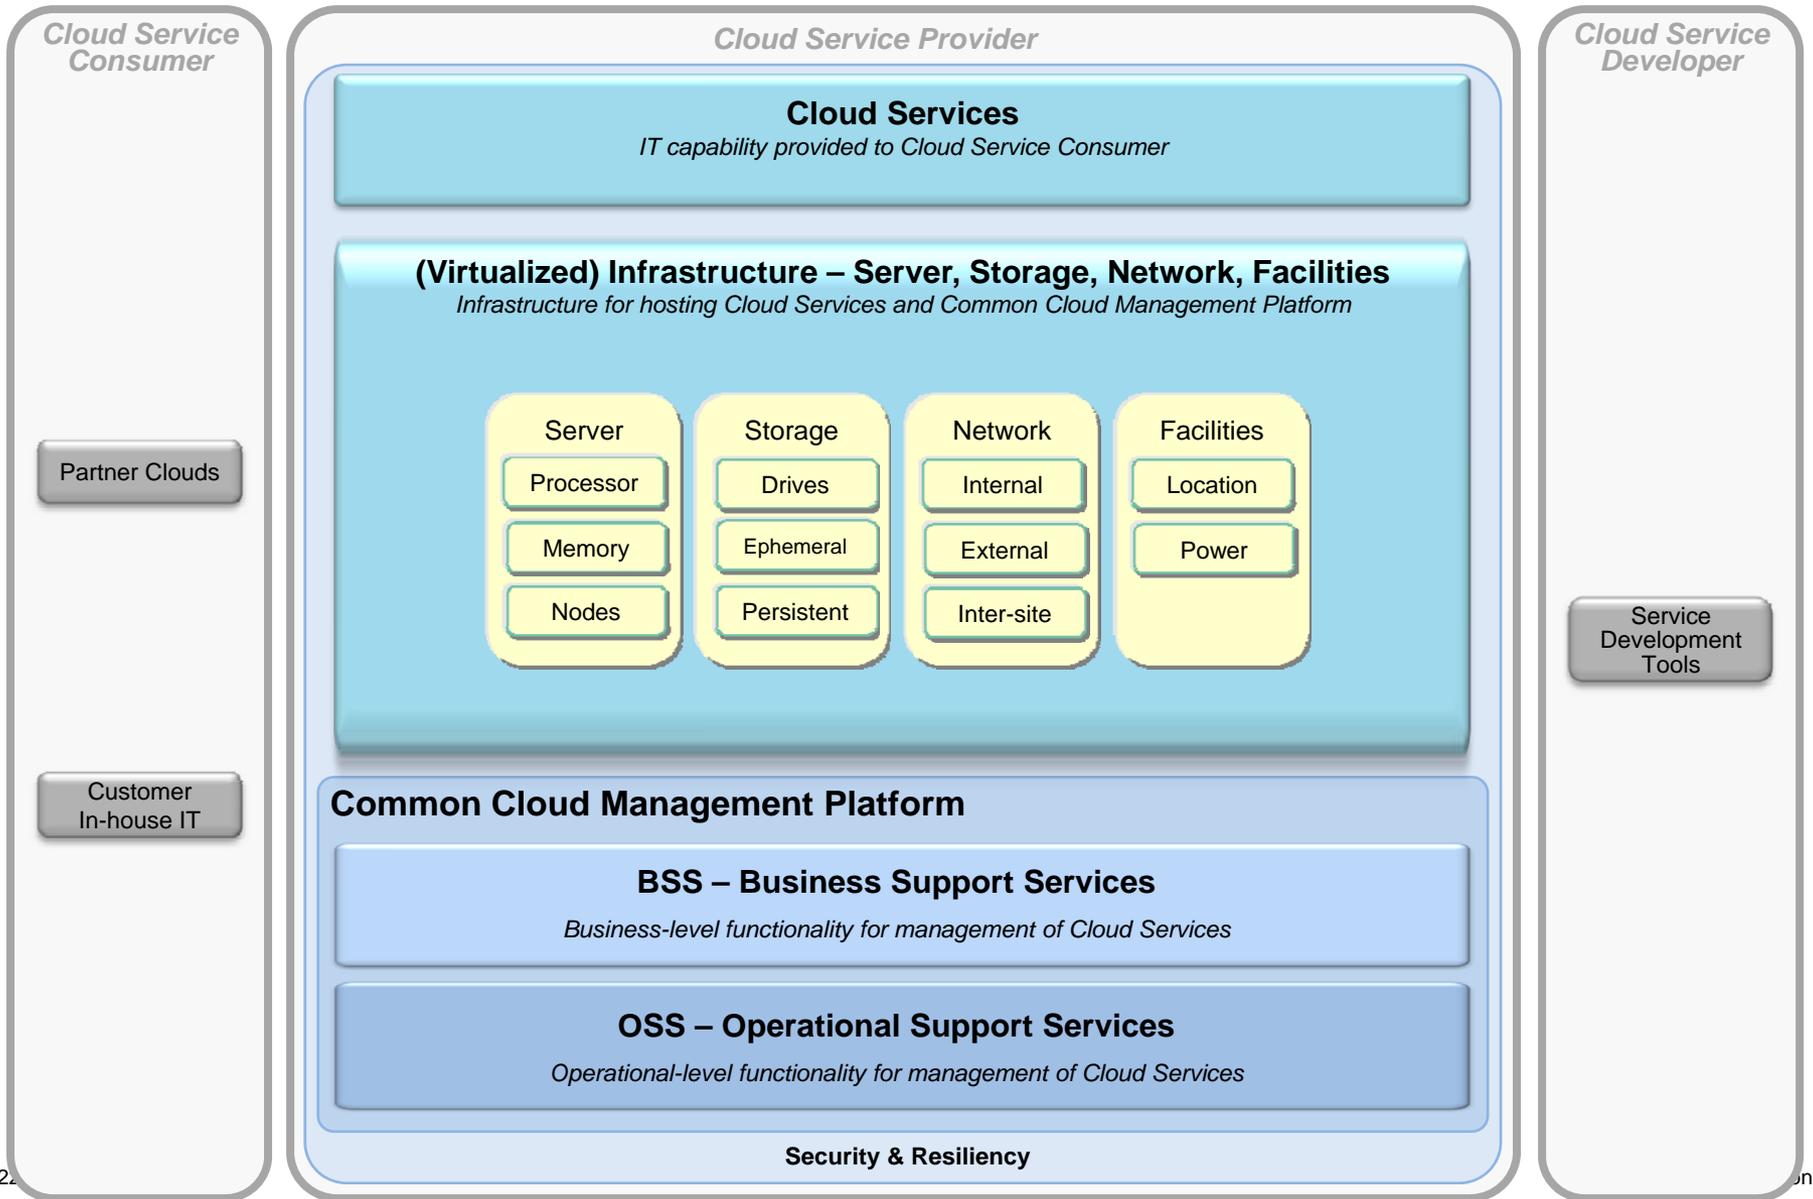
CC RA – CCMP Details



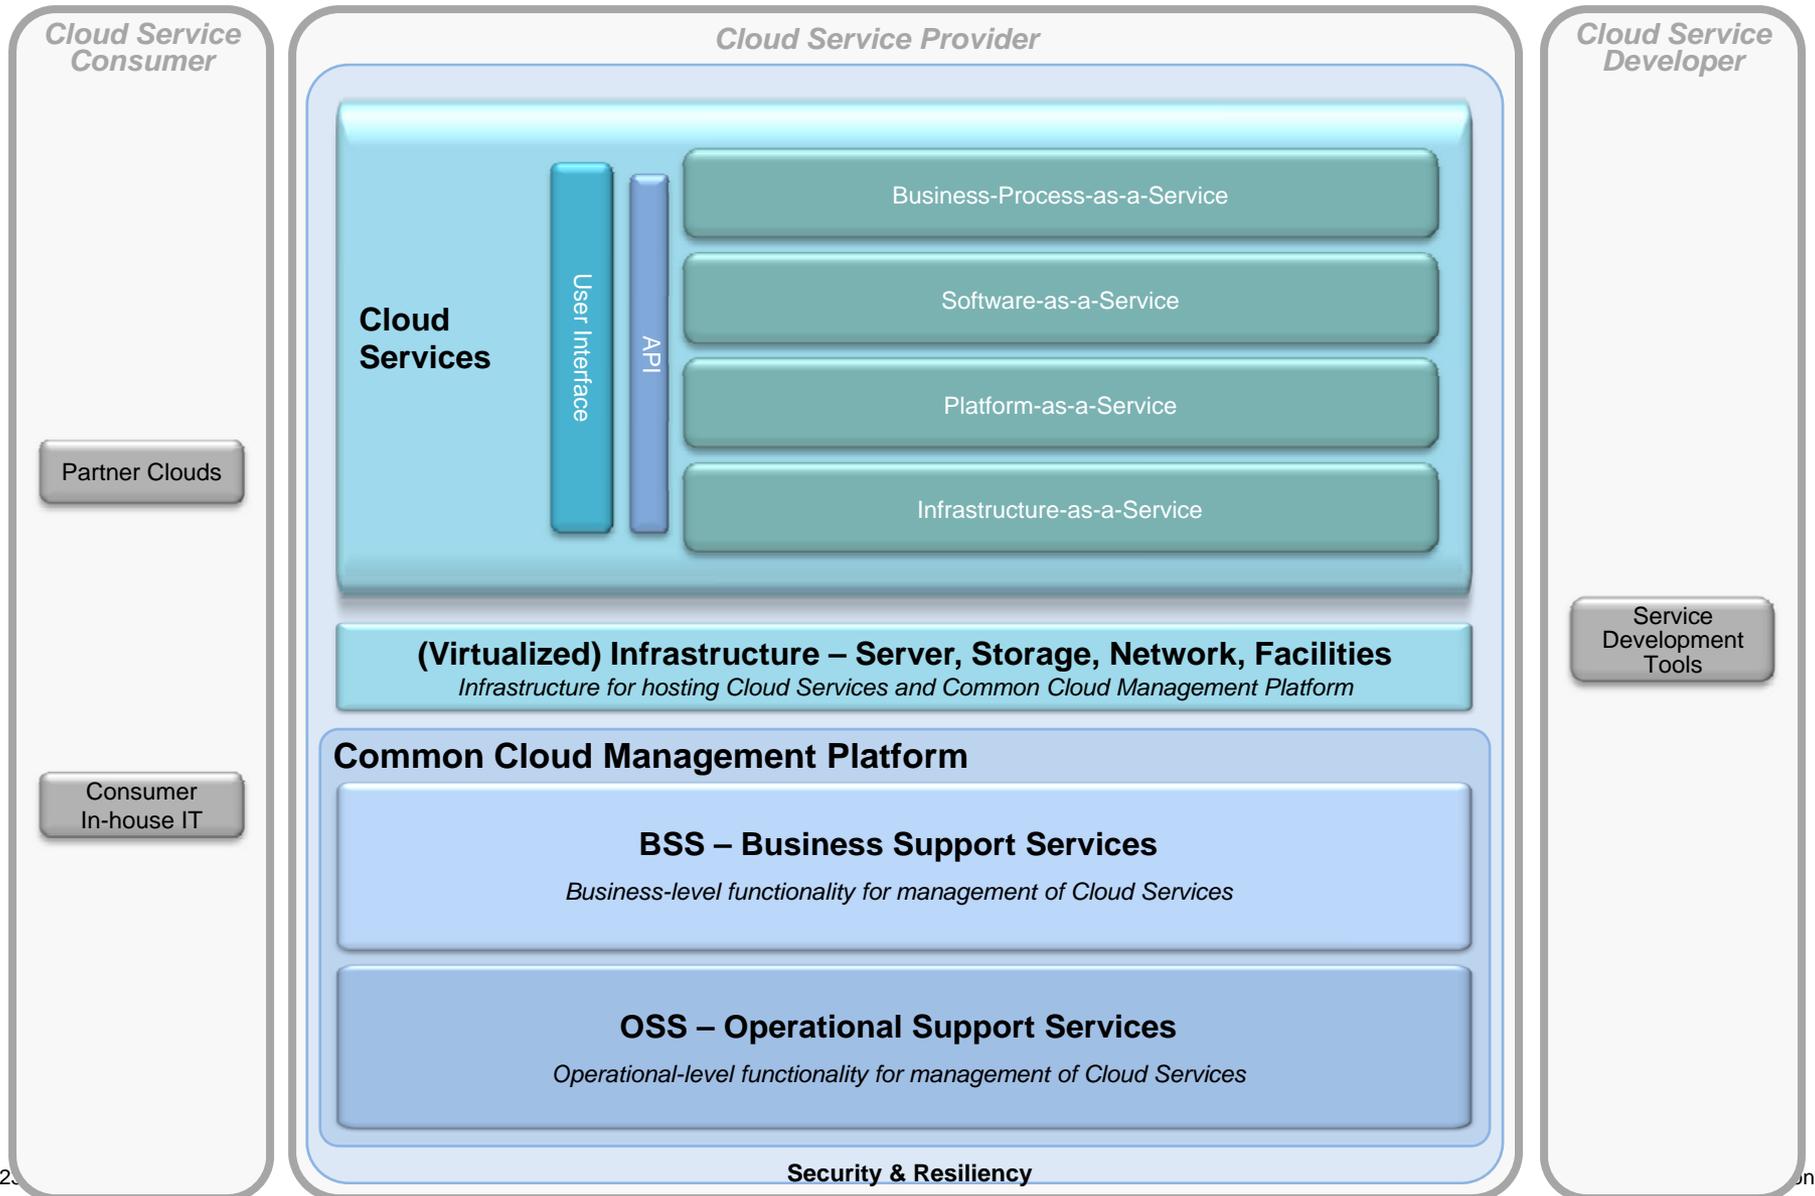
CCMP RA – Security & Resiliency Details



CC RA – Virtualized Infrastructure Details



CC RA – Cloud Services Details



Virtualized Resource Management

- Deploy cloud services on virtualized resources
- Manage virtual resources

Service Automation Management

- Interpret and Execute Build- and Management Plans
- Orchestrate Management Componentry

Heat and Power Management

- Control Energy Consumption

Image Management

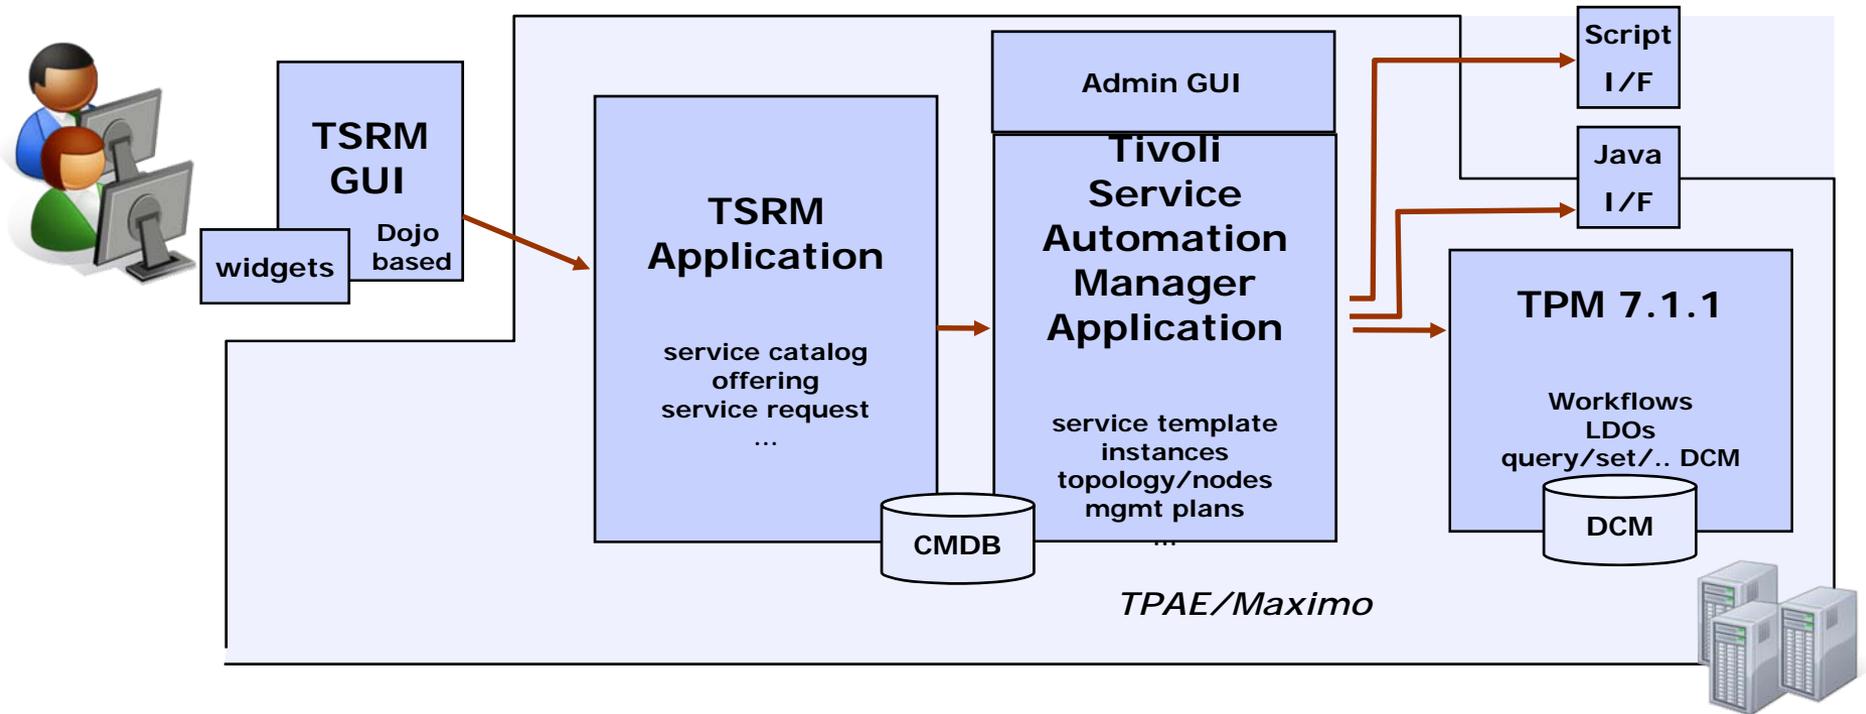
- Design, build and manage images for cloud services

Security

- Design for Multi-Tenancy
- Protect assets through Isolation, integrity, image- risk and compliance management

Usage Metering and Accounting

- Flexible support of delivery models



GUI

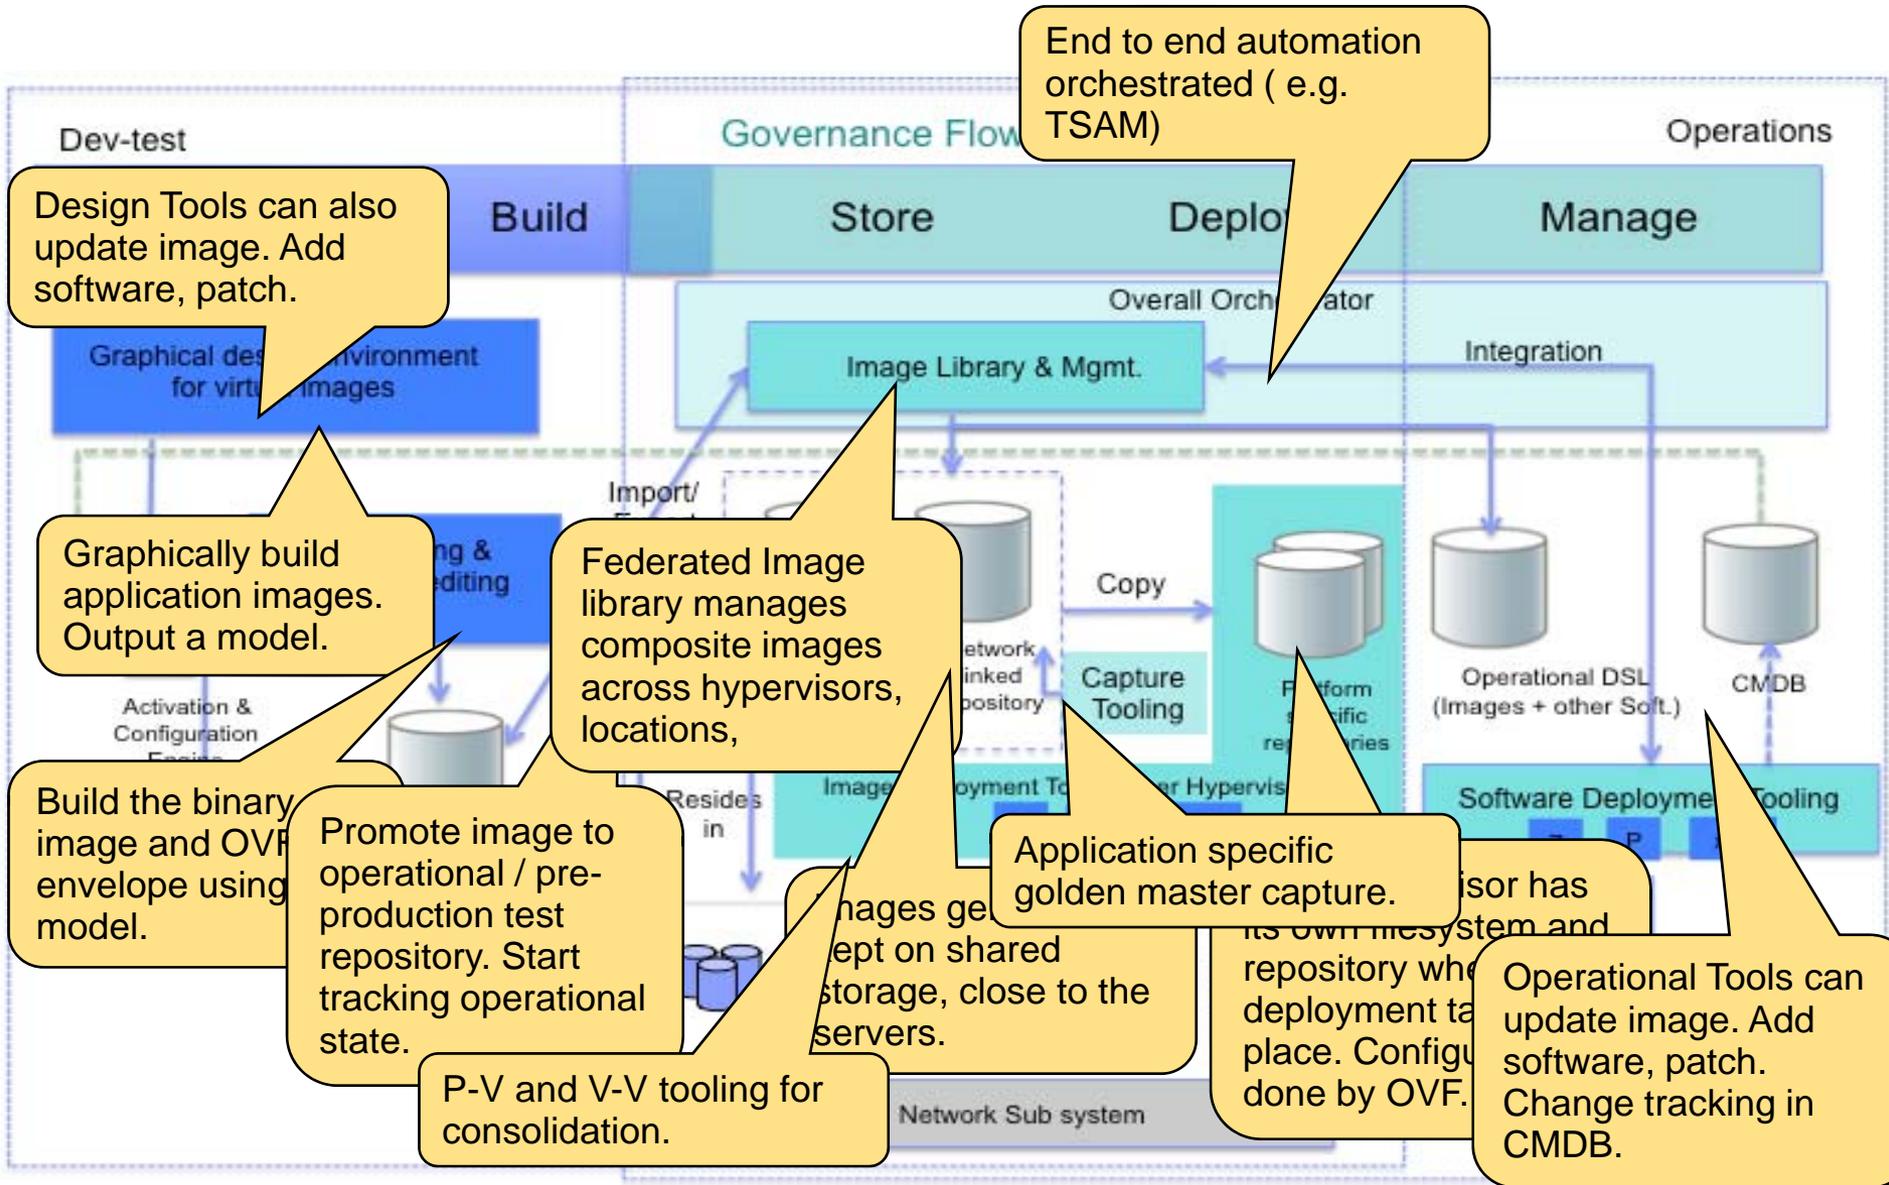
- Interaction with end user
- Collect parameters for management plans

SRM

- Prepare service request from given input parameters
- Perform reservation of resources
- Approval and notifications on business level

Tivoli Service Automation Mgr TPM

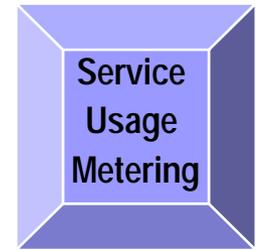
- Topology definition
- Orchestration by management plans
- Management plan definition
- Management plan execution - push down on eg. TPM (or Script)
- Approval and notifications on technical level (admin)
- Situation governance incl. error handling by admin
- Work assignments on admin level (“inbox”)
- Management plan fulfillment by executing TPM workflows/LDOs ... or native scripts ... or Java based actions ... or manual tasks
- Change resource state



What can be metered?

- Creation/termination of a Cloud Service instance, e.g. duration
- Assignment of resources to a Cloud Service instance over a given period of time
- Capacity changes to resources assigned to an Cloud Service instance

✓ **Metered by Management Infrastructure**



- Resource utilization
 - E.g CPU usage, Memory Usage
- Resource consumption
 - E.g on a per user, per transaction basis

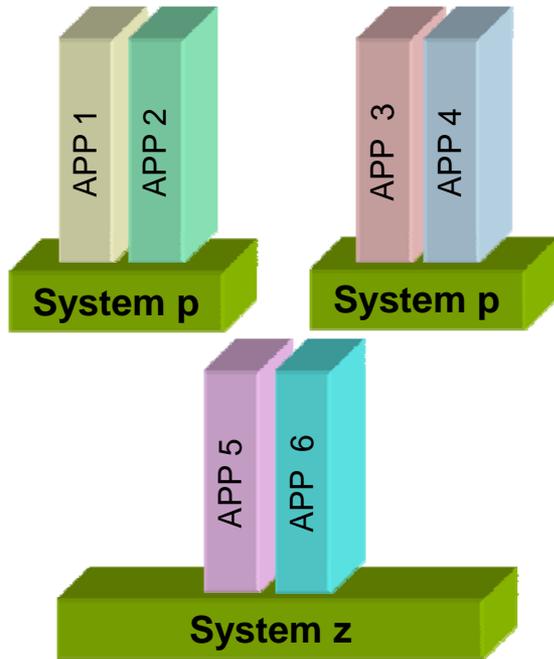
✓ **Requires instrumentation/collectors on each resource**



Accounting: What should be metered?



- Definition of an Accounting Model:
 - Business Metrics
 - Which „entities“ should a consumer be billed for?
GB per day
 - What defines price and/or cost?
Number of disks, servers, SW license
 - Relationship of Business Metrics to IT Metrics
 - What can be collected with „reasonable“ effort ...
Number of GB assigned, max number of GB used every hour, every 24 hours
 - ... and be mapped to Business Metric
 - Define of cost and/or price (rates) for above „entities“
GB allocated per day: 5 \$
GB used per day: 1 \$
 - Define Accounting Structure
 - Assign Usage to Application, to Organizations, to Services/Subscriptions



Holistic Management of Power and Thermal Issues.

Combination of Hibernation, Power Off etc. with Monitoring, Workload Balancing and Provisioning

Automate Energy Control

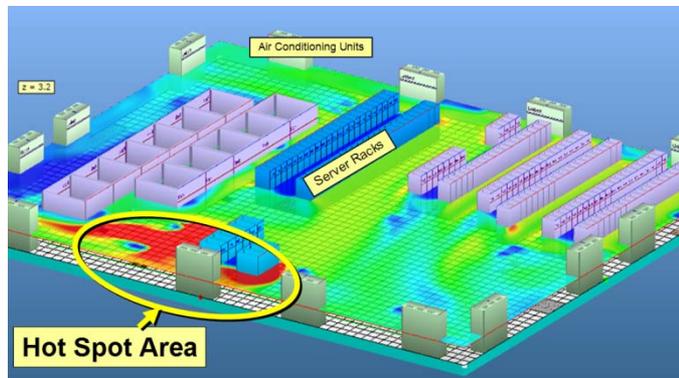
- Policy-based automation

Control Energy Consumption

- Consolidate and migrate workloads

Hot Spot Avoidance

- Physical Movement of Workloads



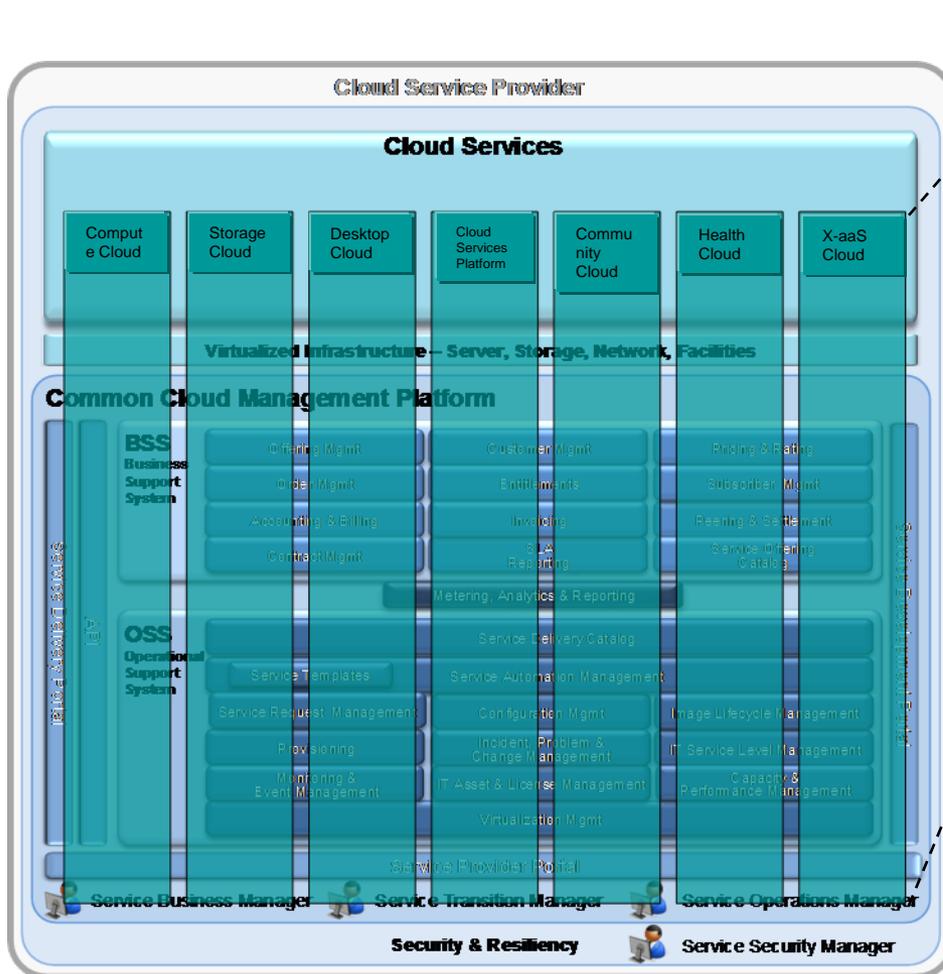
Traditional Data Center Management vs. “Cloud-like” Management

The overall objective of Cloud-managed data centers is to **automate any type of task or situation** (by reducing manual intervention) for **increasing flexibility** and **reducing operational expenses**

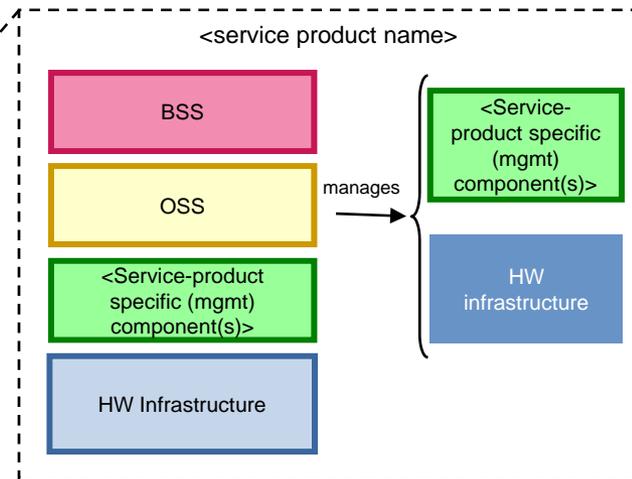
Core Metrics	Traditionally managed Data Center	“Cloud-managed” data center
Admin/Server ratio → Costs	1:50 – 1:100	1:100’s – 1:1000’s
Time to provide new service instances & changing them → Flexibility	Days / weeks	Hours / minutes / seconds

Core Disciplines
IT Management approach
Administration Tasks
Problem handling
Service Consumer <-> Service Provider interaction

For Cloud-like efficiencies and flexibility, it is not sufficient to have the right technology, but to also use it in the right way!



Service products used in diagram only for illustration purpose, not normative



- BSS, OSS & UI/Portals provide a set of management functionalities (or services) commonly needed by cloud service implementers
 - BSS, OSS and UI to be leveraged by cloud service implementers
- In general, maximum CCMP exploitation by cloud service offerings is required for
 - Reduced Time-to-market for new cloud service offerings
 - Reduced overall management cost
 - Simplify deployment of cloud services on top of other cloud services
 - Building on common existing (compliance) certifications, existing integration with billing systems, common UI technology / look & feel, etc.
- BSS, OSS & UI can be configured / exploited through different plug-in mechanisms

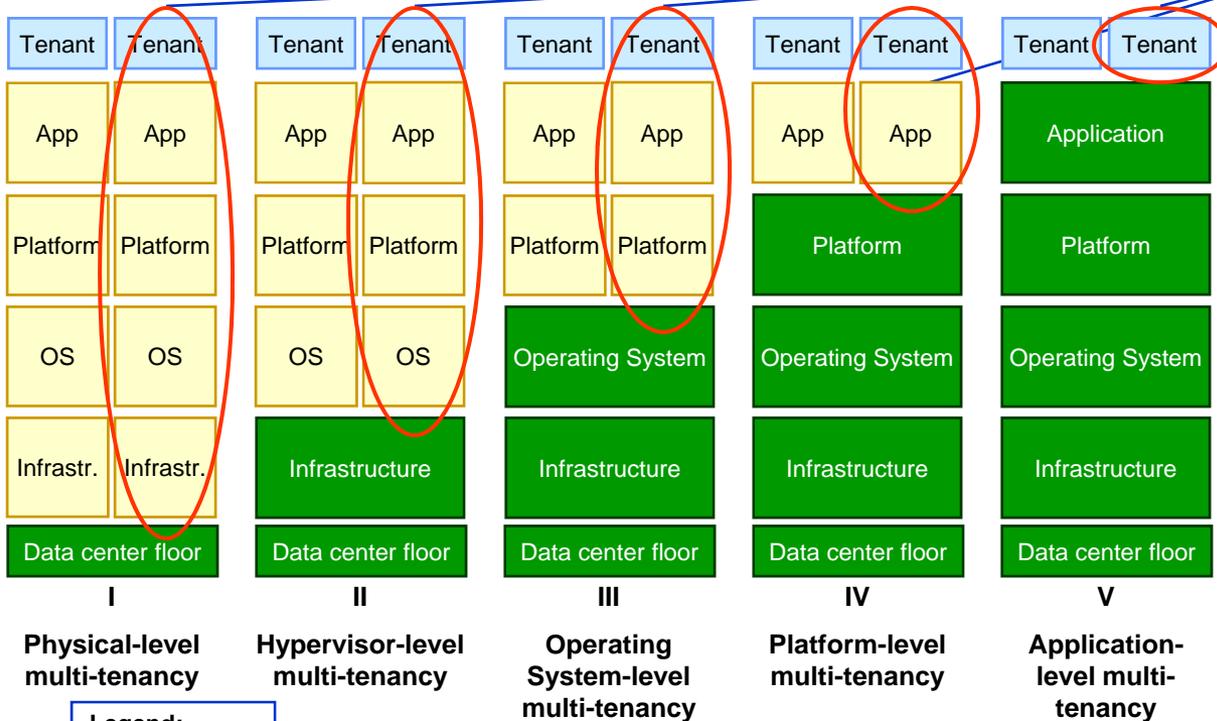
Multi-tenancy in Cloud environments

In a cloud environment, there are two primary technology stacks where multi-tenancy is relevant:

- The management environment
- The managed environment (Infrastructure, Platform or application that is provided as a service)

The amount of work required for setting up a new tenant depends on where the “multi-tenancy point” sits within the technology stack.

The higher the multi-tenancy point, the less effort is required for setting up a new tenant (because more underlying technology is shared). Conversely, the higher the multi-tenancy point the more resources can be shared amongst tenants.



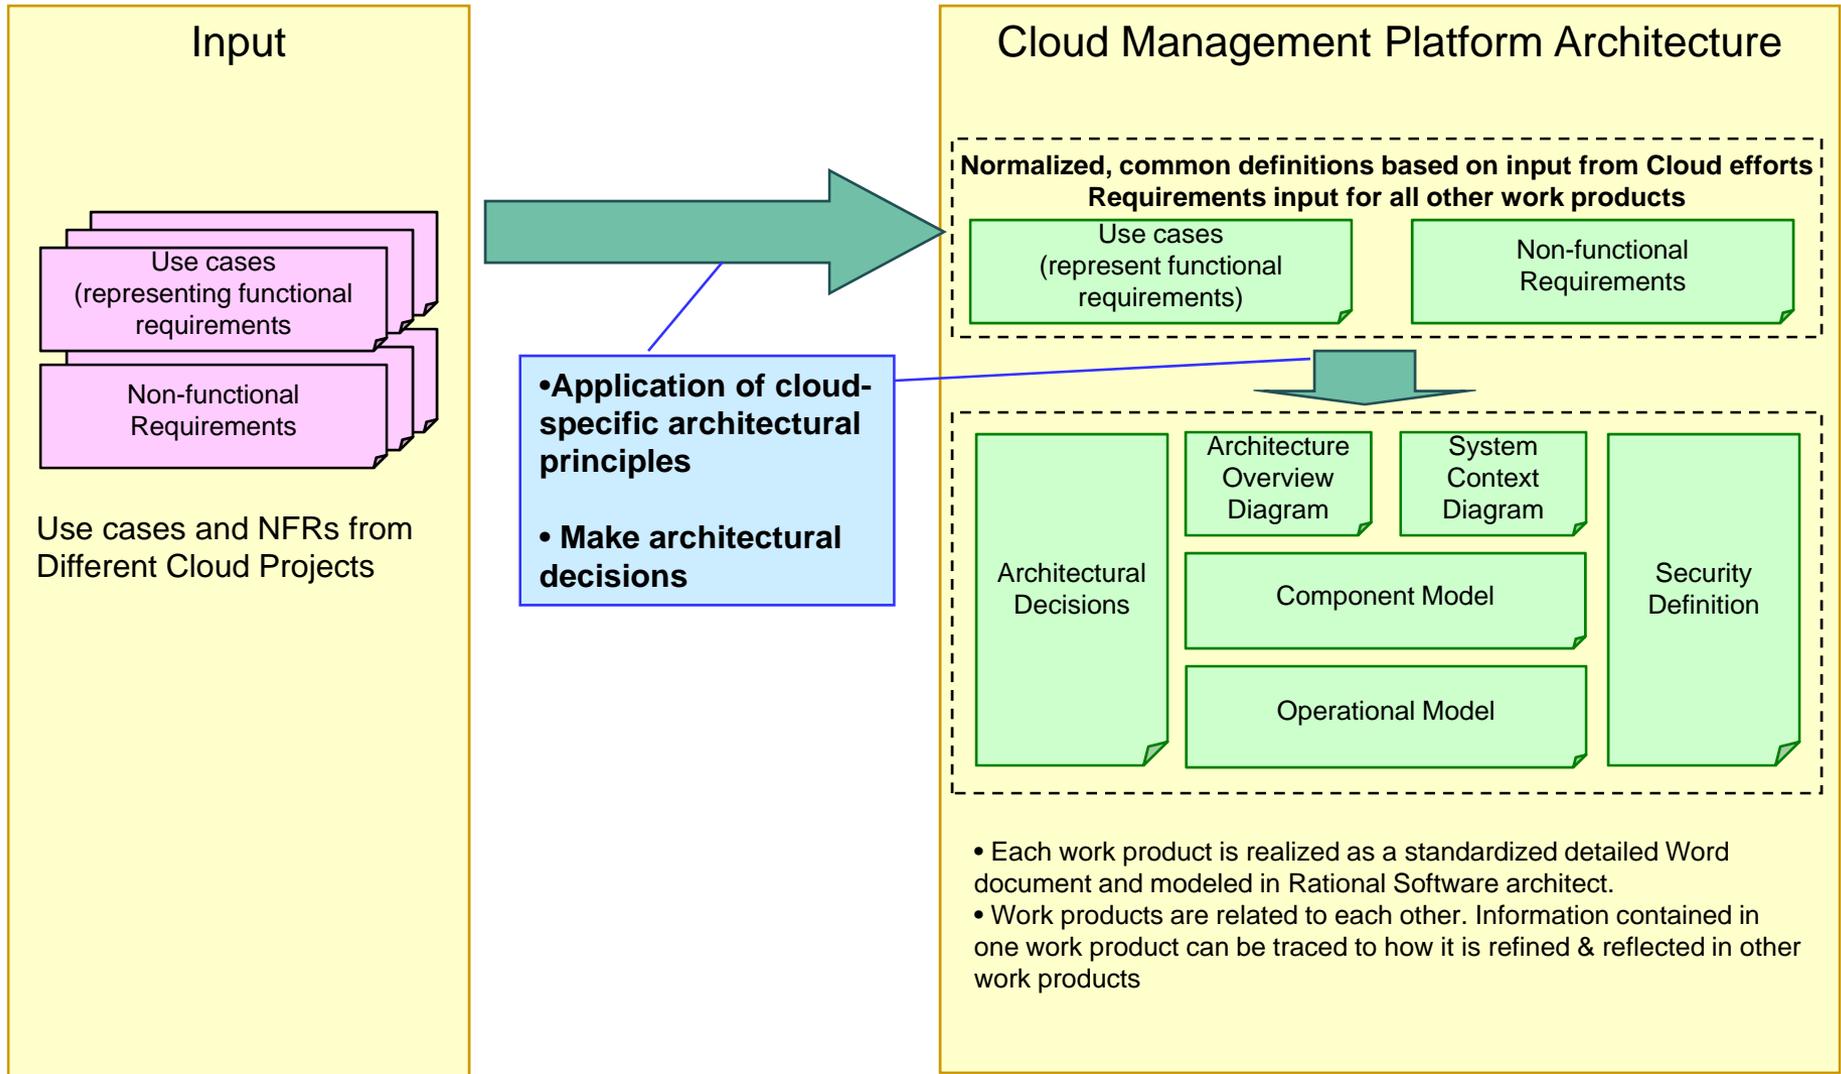
Multi-tenancy: Shared vs. dedicated

In multi-tenant environments shared layers must behave as if they were setup in a dedicated fashion.

- This requires “multi-tenancy support” on the respective technology layer (infrastructure/OS, platform or application)
- If the implementation of a technology layer doesn’t inherently support multi-tenancy (i.e. being able to behave within on installation as if there were multiple dedicated ones), then it must be setup in a dedicated fashion.

Legend:

- Dedicated
- Shared

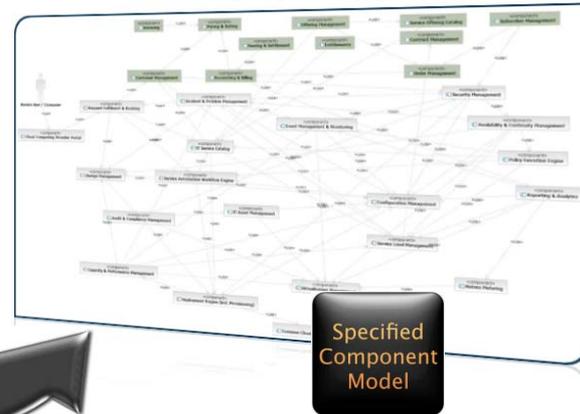


CCMP RA Component Model

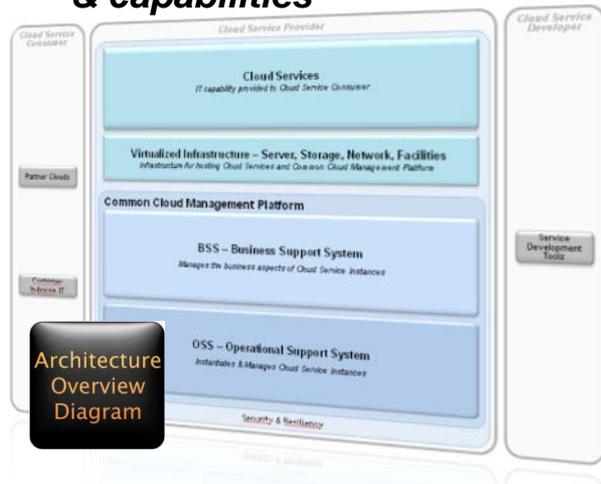
CCMP BSS components & functions



CCMP component

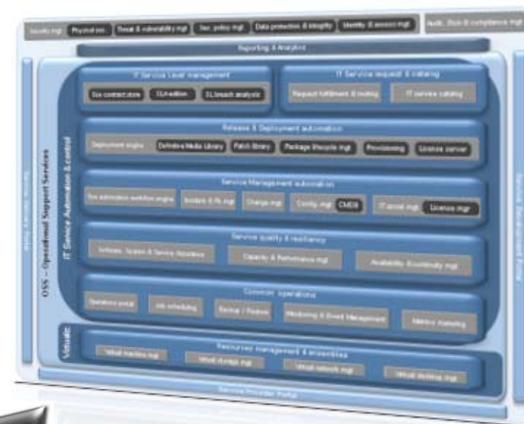


CCMP services & capabilities

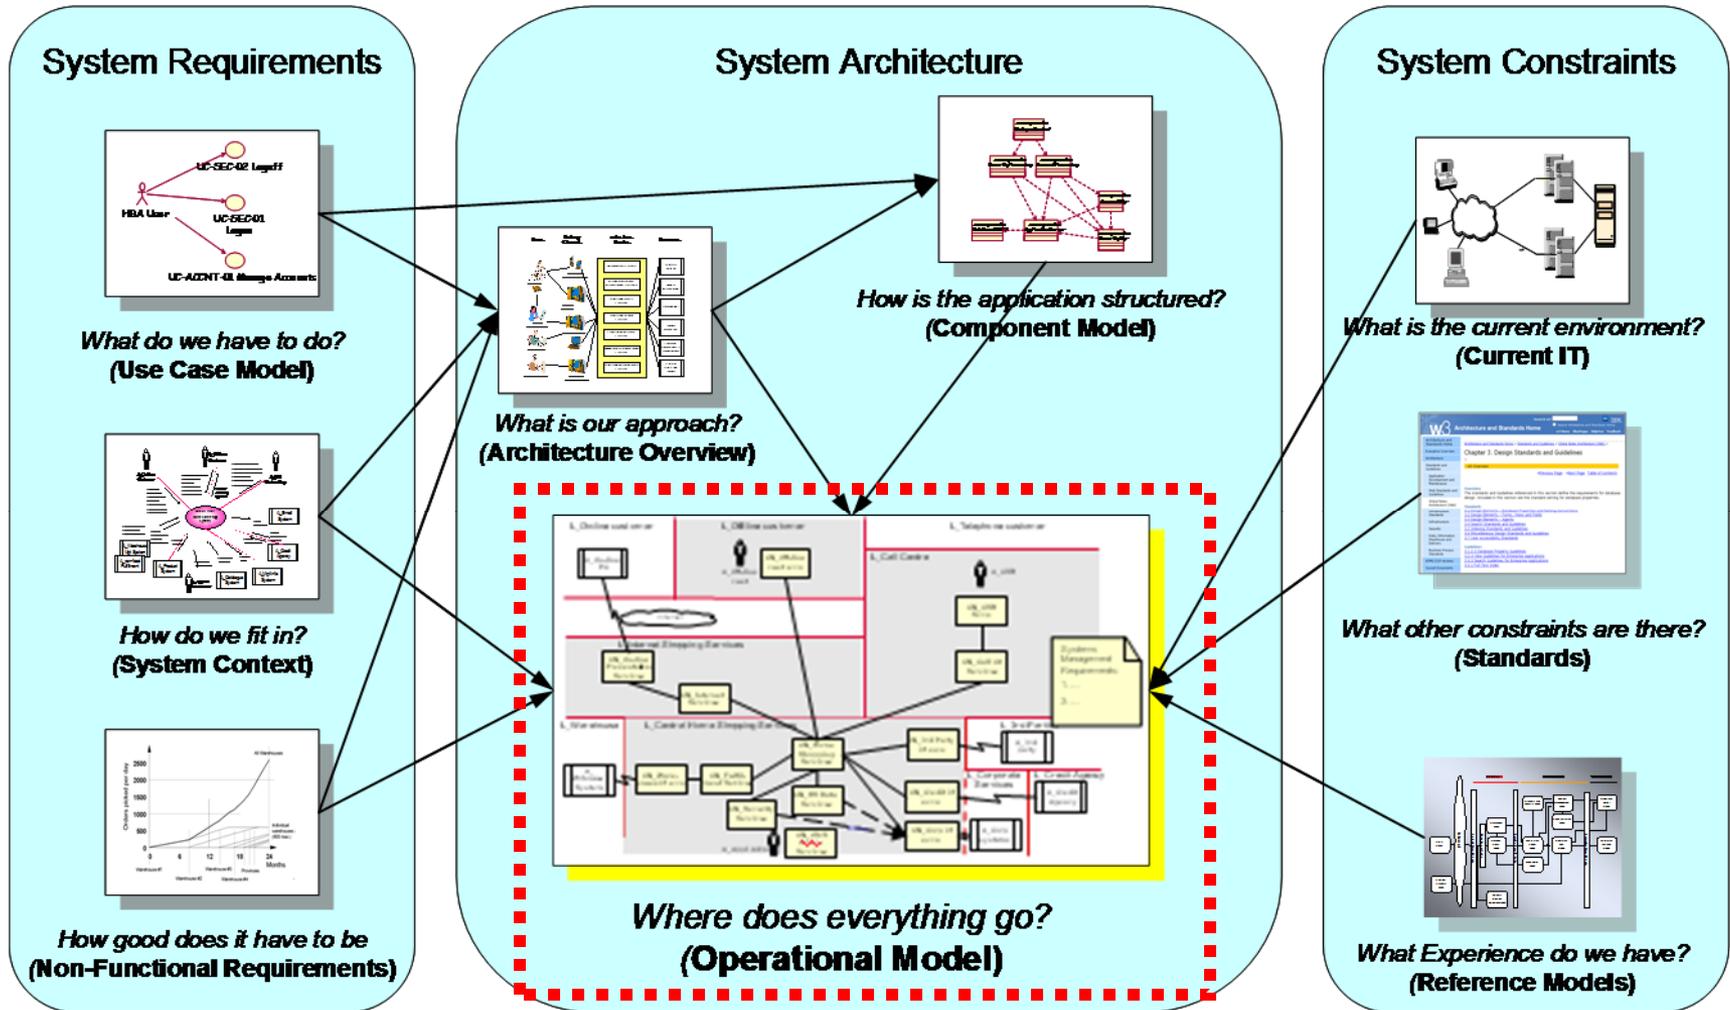


Conceptual Component Model

CCMP OSS components & functions



Operational Modeling helps ensure the IT infrastructure's non-functional requirements are delivered, within all constraints.



China - Cloud computing center

Business Goals

- Build software parks with Cloud Computing Centers to support emerging software companies in China.
- Offers Chinese software companies the ability to tap into a virtual computing environment to leapfrog their development activities.
- Accelerates transformation to a service-led economy.

Lessons Learned

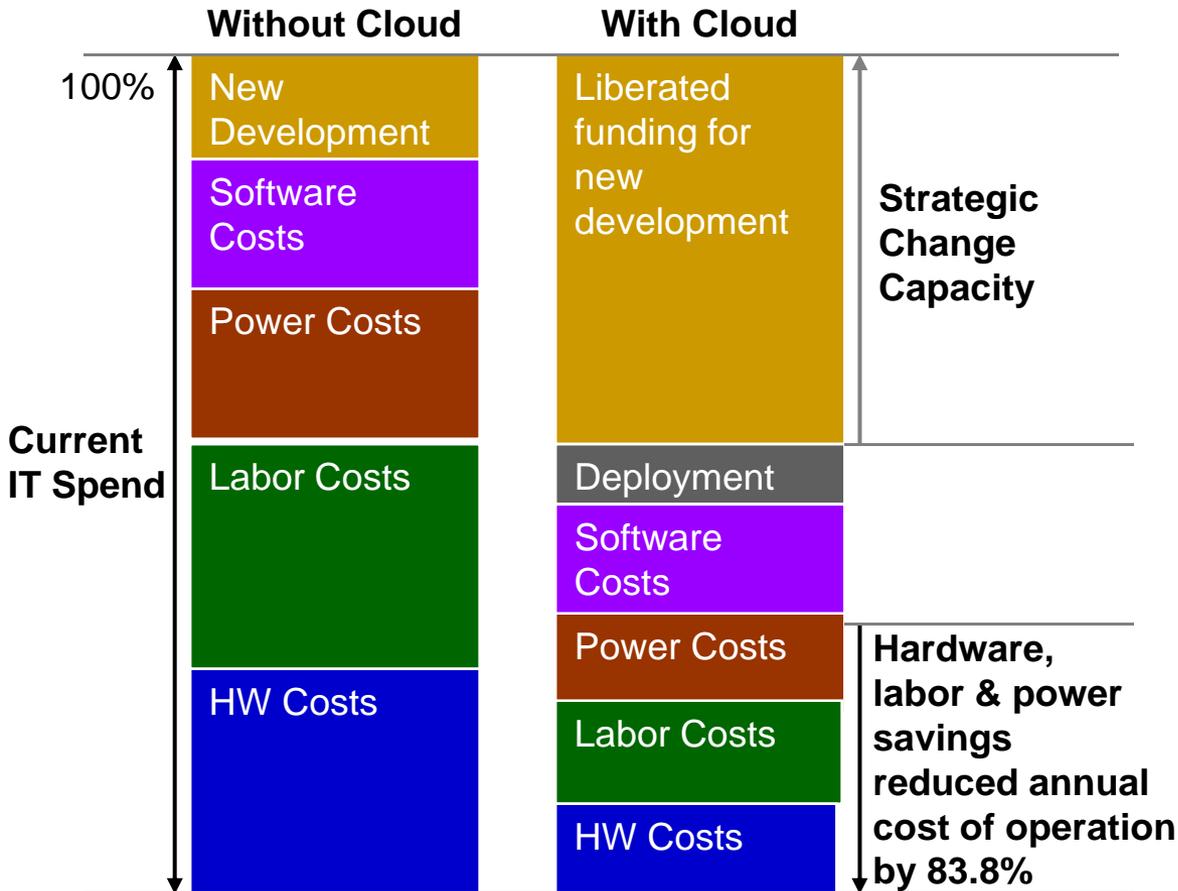
- Working closely with customer to identify short and long term business goals.
- Combining IBM hardware, software and GTS services provides a powerful range of capabilities, which drive significant customer business benefits.
- Incremental and phased approach to establish early wins with balance risk.

Solution Overview

- A cost efficient public cloud enabled by IBM technology and services.
- A shared facility, providing each company in the Software Park with its own virtual data center.
- Accelerates development and test cycles through quick resource on-boarding.
- Managed with IBM Tivoli systems management products (ITM, TPM, ..).
- Hardware – IBM System x, System p and BladeCenter.
- Software images offered to end users include IBM Rational products, WebSphere Application Server and DB2.



IBM Technology Adopter's Portal (IBM TAP)



- Innovation Cloud for 100,000 Subscribers
- Reduced Capital Expenditure
 - Reduce from 488 servers to 55
- Reduced Operations Expenditure
 - Reduce from 15 admins to 2
- Additional Benefits:
 - Enhanced customer service
 - Less idle time
 - More efficient use of energy
 - Acceleration of innovation projects

Summary

- Cloud Computing is a disruptive change to the way IT services are delivered... it is about shifting to the third compute model in the evolution of IT
- Service Lifecycle Management based on a Dynamic Infrastructure is the foundation for managing Clouds
- A solid Cloud Computing Architecture is required to successfully and economically manage Clouds
 - Open standards based architecture for the buildout of private, public and hybrid Clouds
 - Management of IaaS-, PaaS- and SaaS Clouds
 - Build for seamless integration into existing customers environment
- The Journey to Cloud requires an integrated and orchestrated approach
- Customers are adopting Cloud Computing today
 - Adoption often starts in the Development- and Test Environments
- The Benefits of Cloud Computing are real!



Thank you!

**For more information, please visit:
ibm.com/cloud**