Cloud Computing Architecture and Strategy

Gerd Breiter
IBM Distinguished Engineer
gbreiter@de.ibm.com

06/08/2010

Agenda

- *Introduction*
- *Cloud Computing Reference Architecture*
  - *Cloud Computing Management Platform*
  - *Selected Management Areas*
- *Hybrid Clouds*
- *Customer Projects*
- *Standardization Efforts*
- *Summary*
- *References*

The world is getting smarter…

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

Cloud Computing…

The Industrialization of IT…

….leverages virtualization, automation, standardization and self service to free up operational budget for new investment

… allowing you to optimize new investments for direct business benefits

Cloud Computing Layers

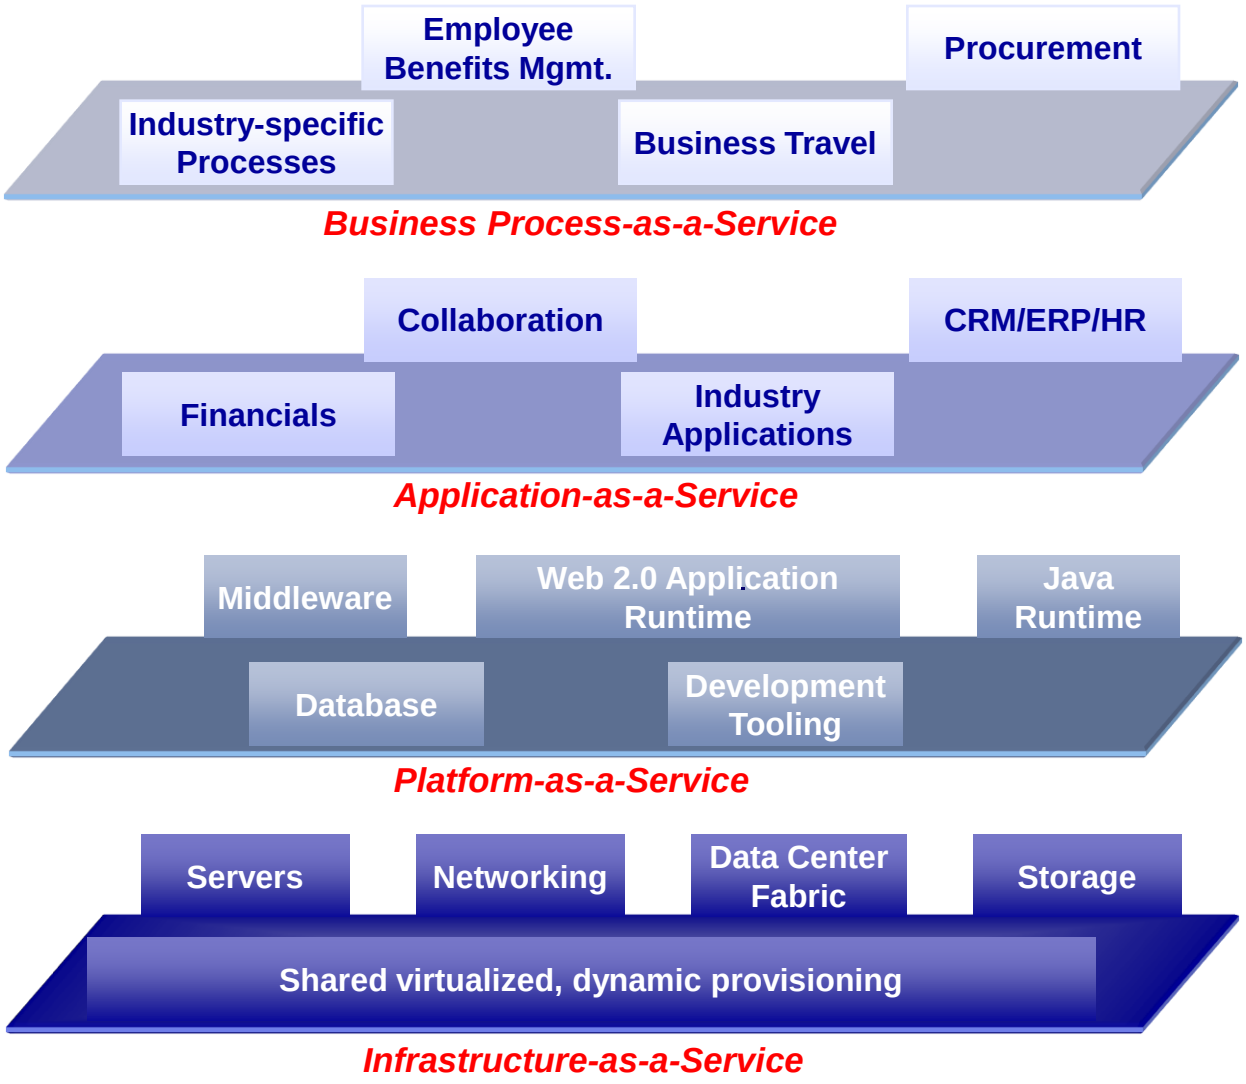

Examples

Cloud Computing Delivery Models

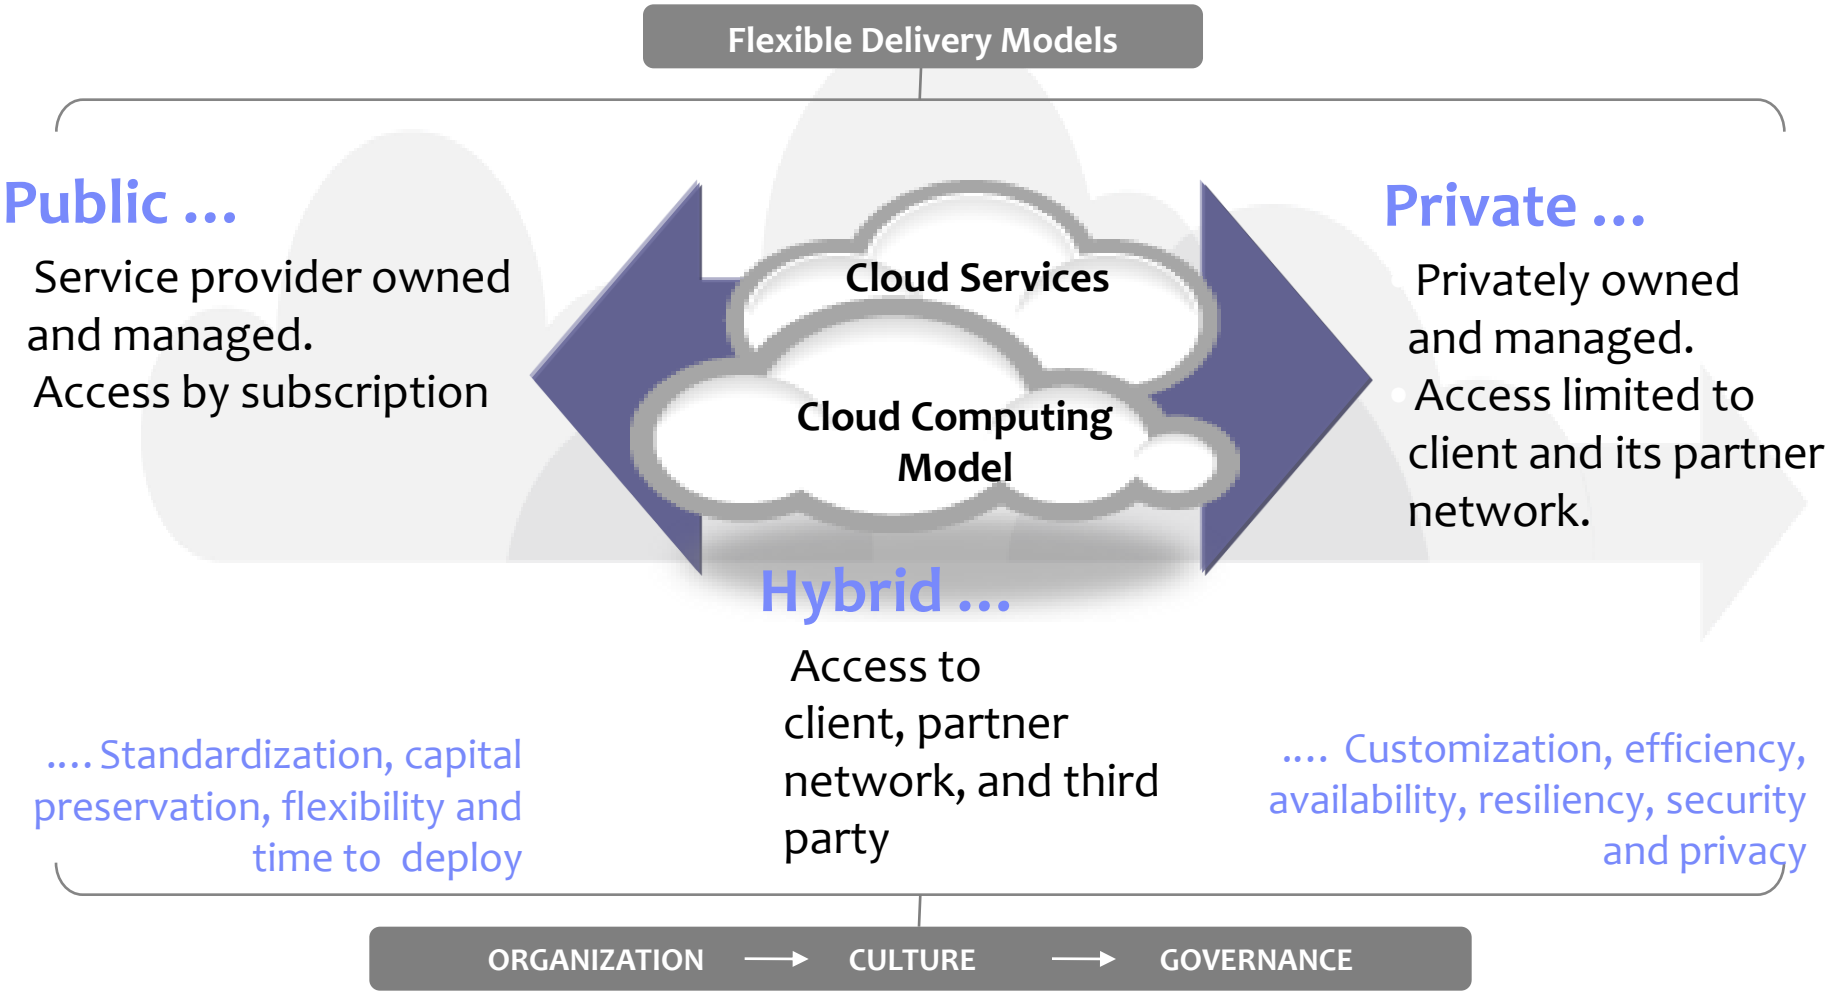

What's so different about cloud-like Service Management? – Changes in orders of magnitude

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

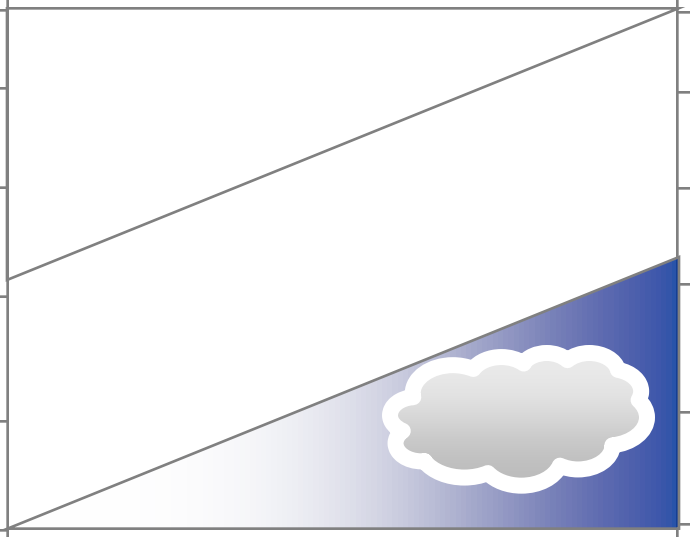

Traditional Service Management

Cloud-like Service Management

Lifecycle of a Cloud Service

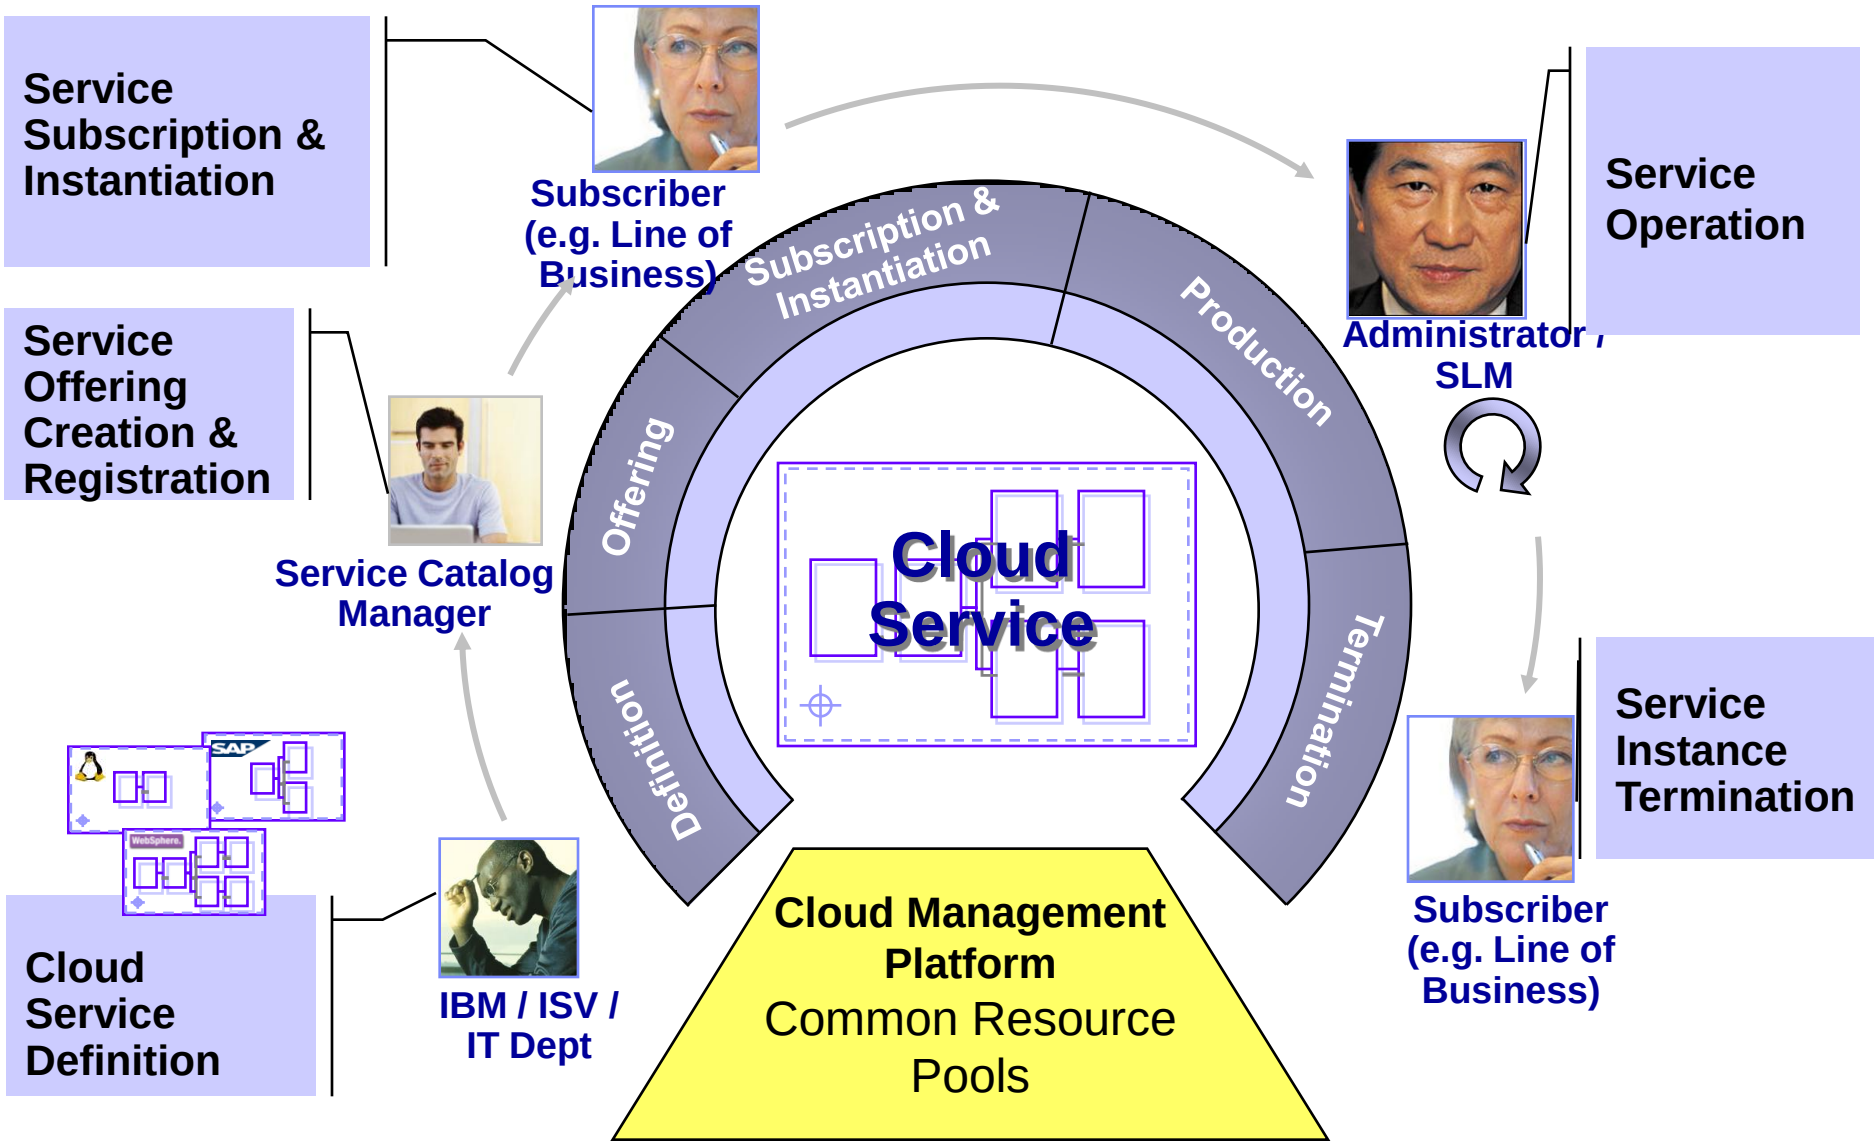

Agenda

- *Introduction*
- *Cloud Computing Reference Architecture*
  - *Cloud Computing Management Platform*
  - *Selected Management Areas*
- *Hybrid Clouds*
- *Customer Projects*
- *Standardization Efforts*
- *Summary*
- *References*

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

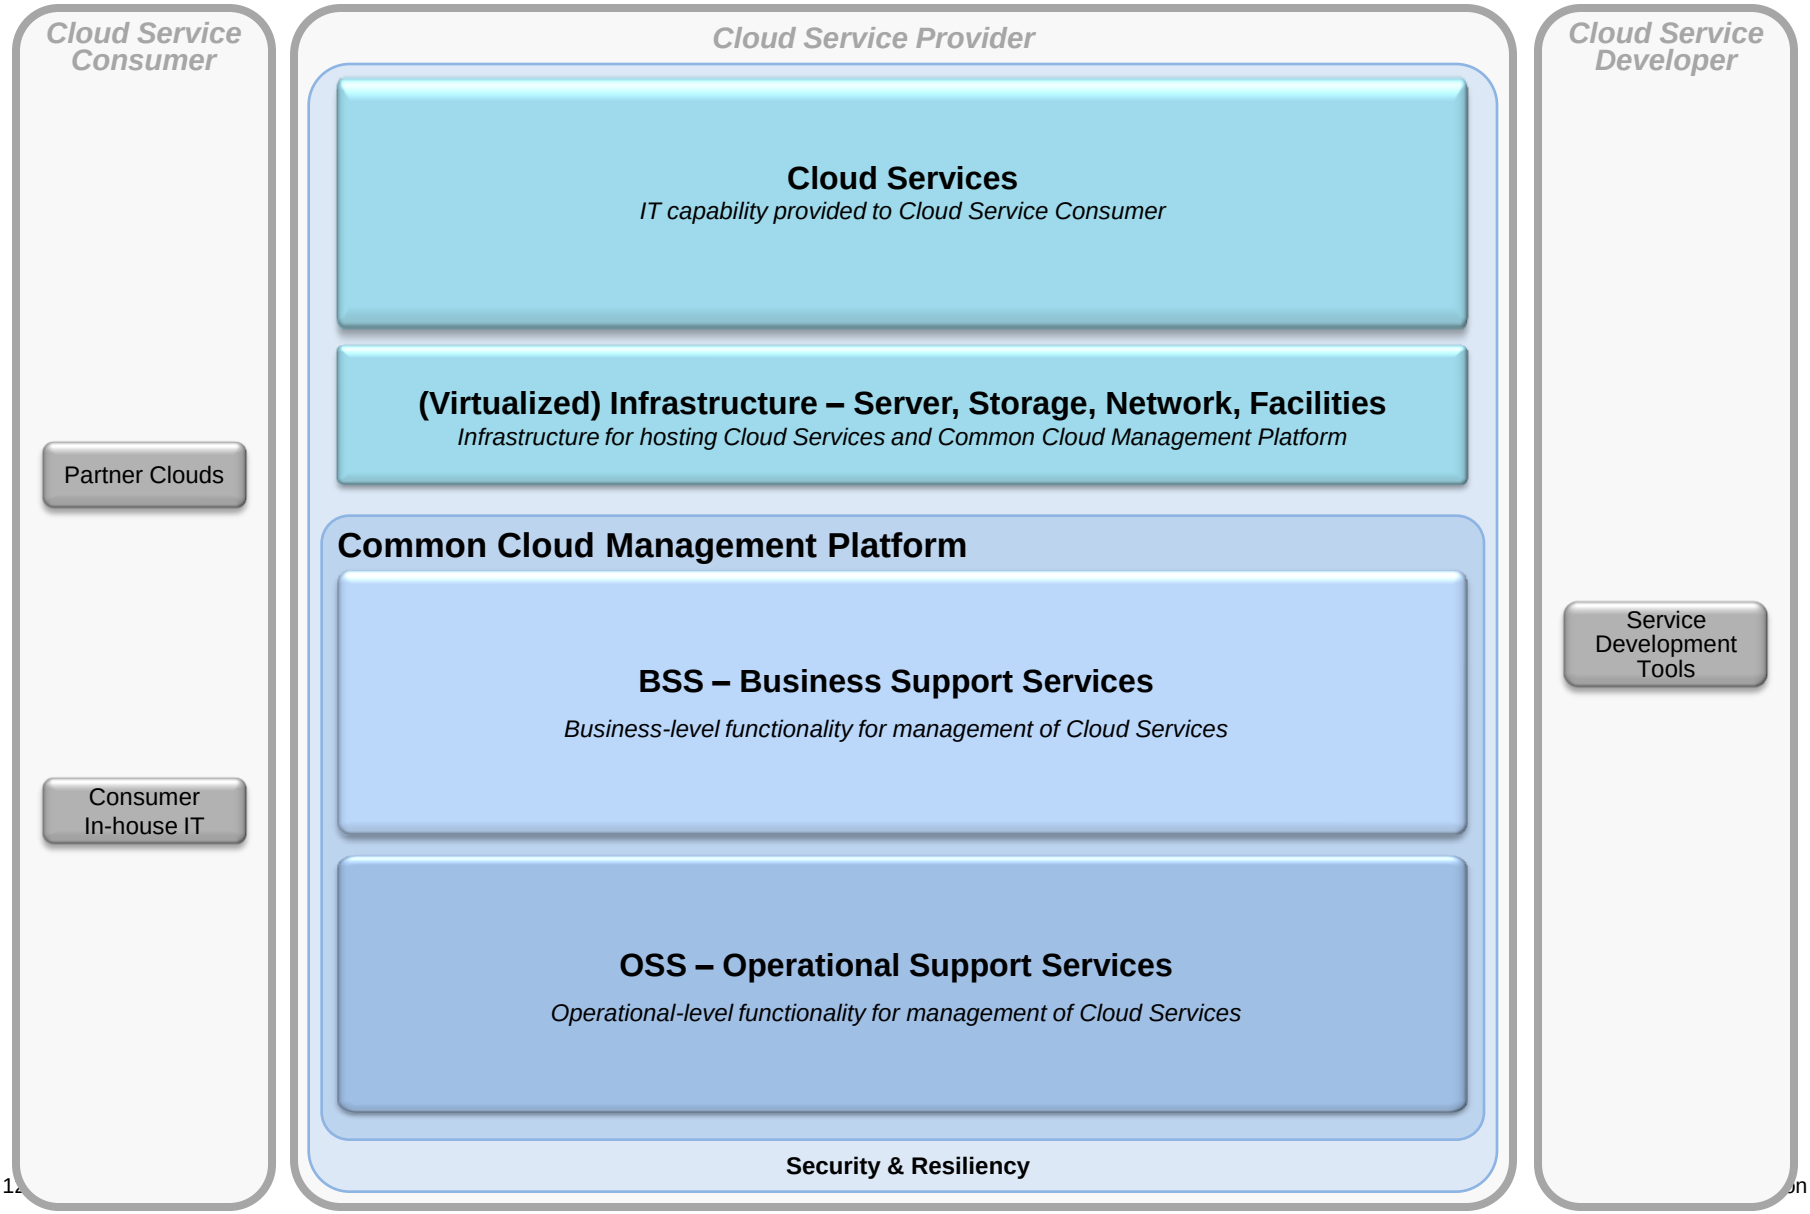

Common Cloud Management Platform RA - Details

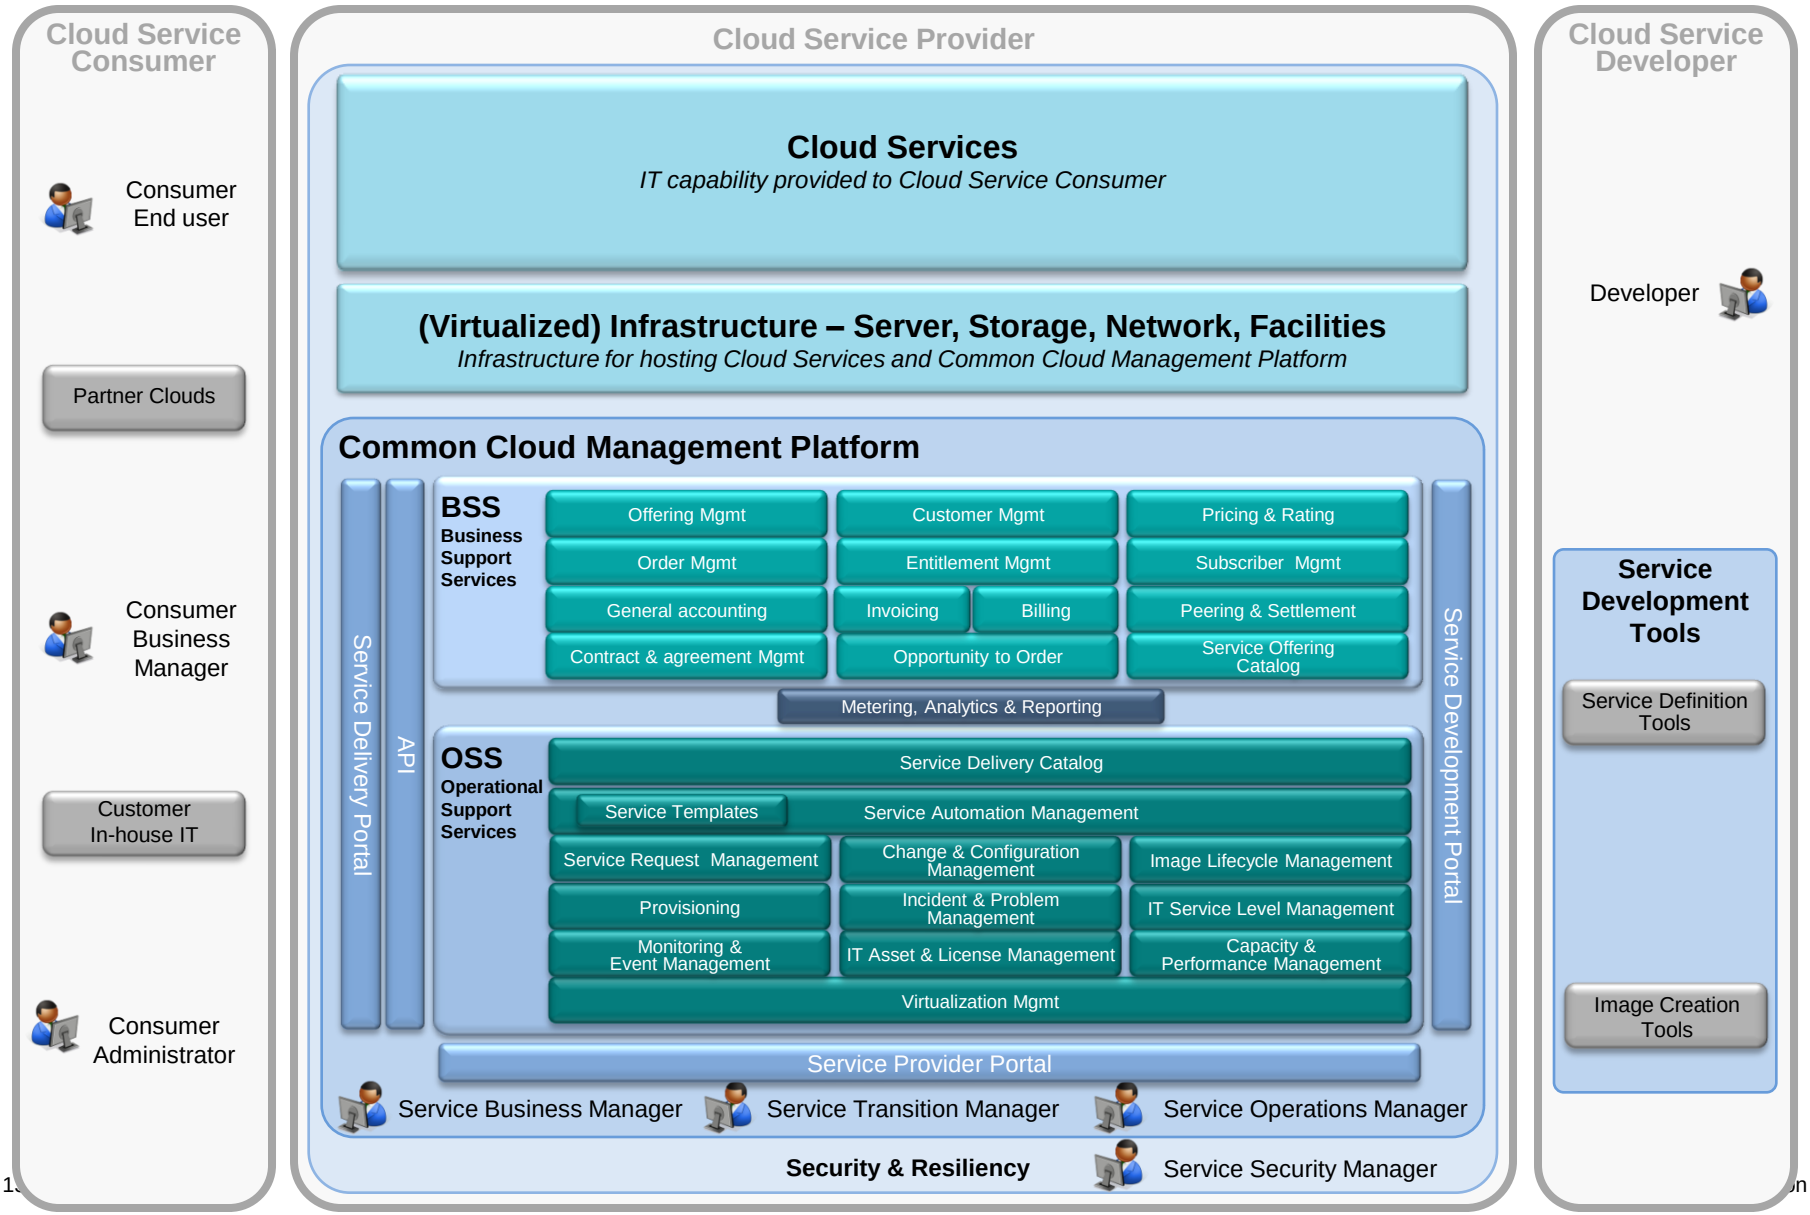

Cloud Management Platform – Selected Management Problems

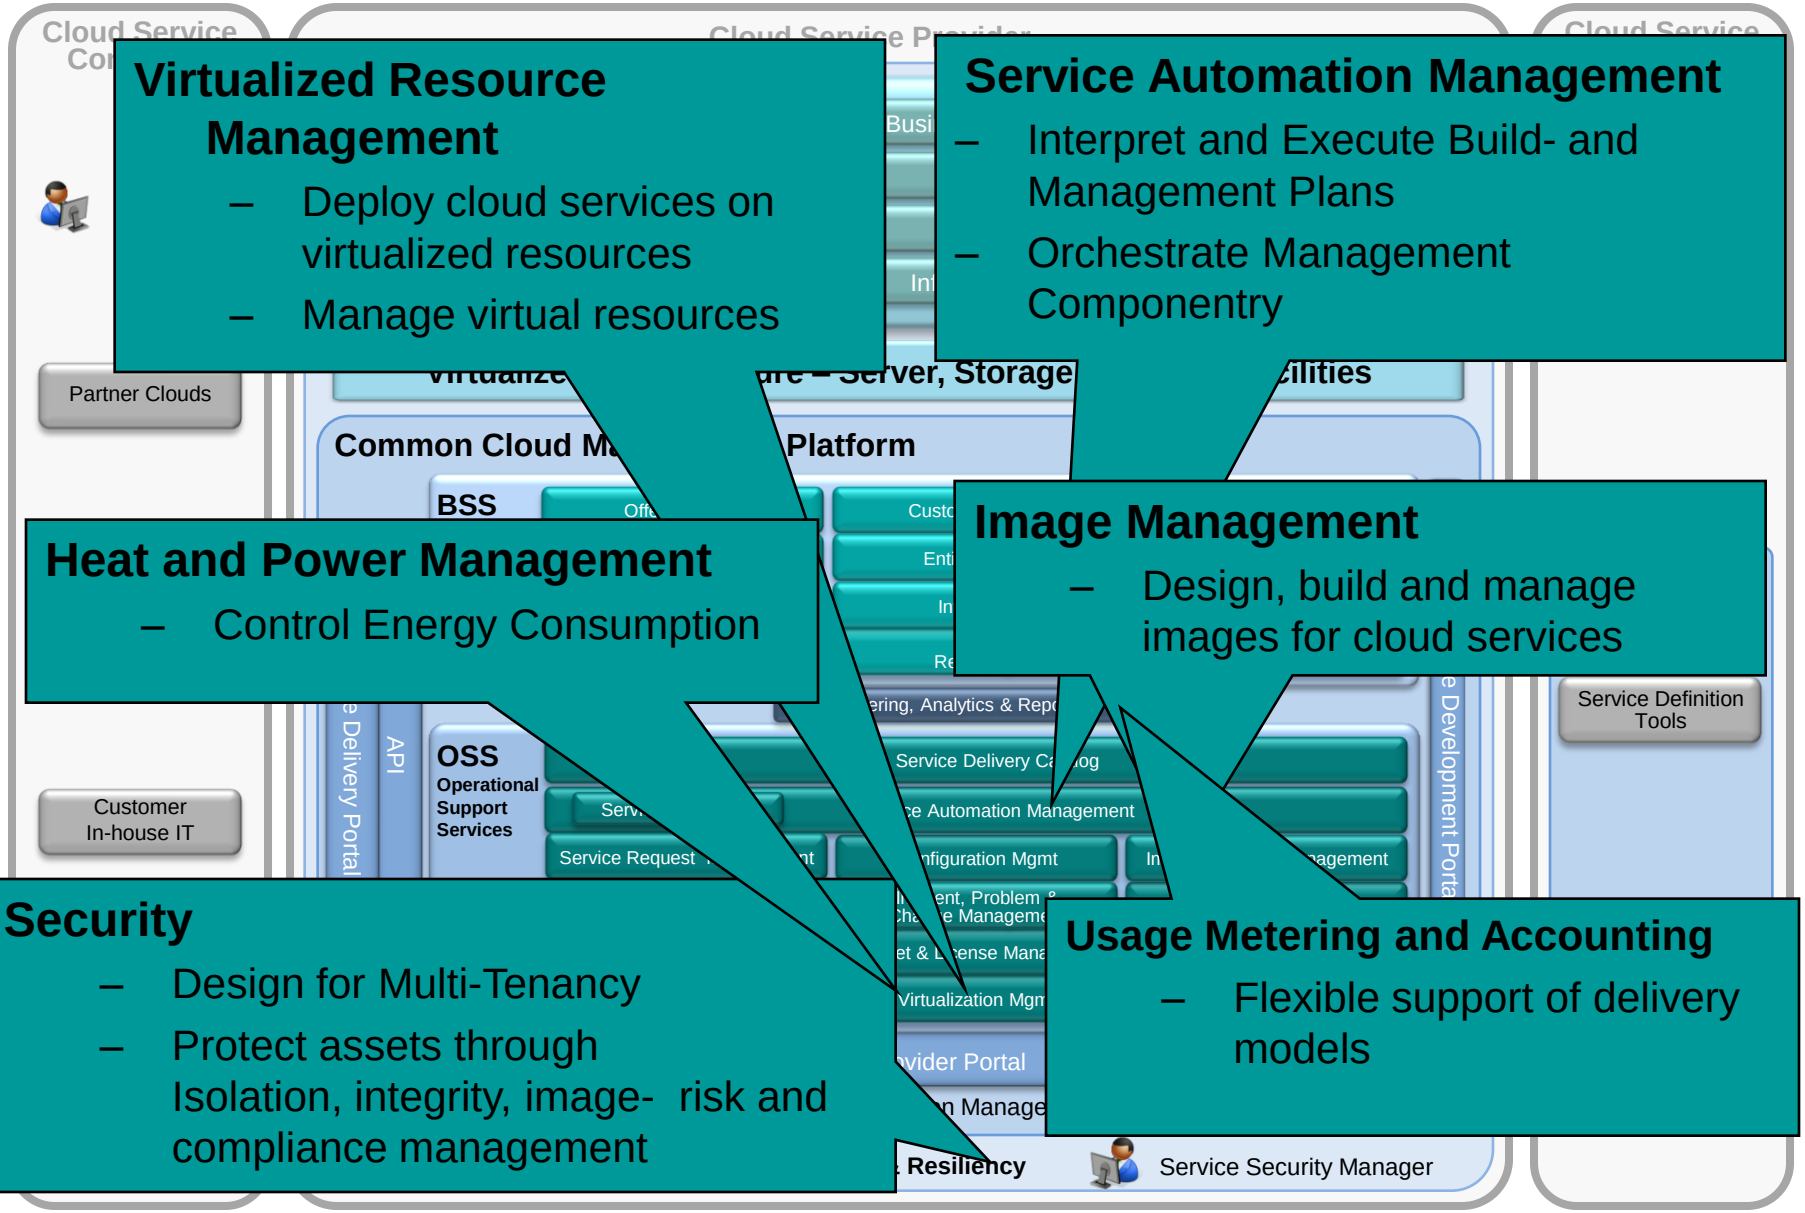

Typical Cloud Management Platform Middleware Stack

Workloads

- Service measurement
- Service reporting
- Usage accounting
- Auditing and controls

Tivoli Service Automation Layer

- Automate process of instantiating and managing a distributed IT environment.

Virtualized Infrastructure Layer

- Virtualized resources
- Virtualized aggregation
- Physical infrastructure

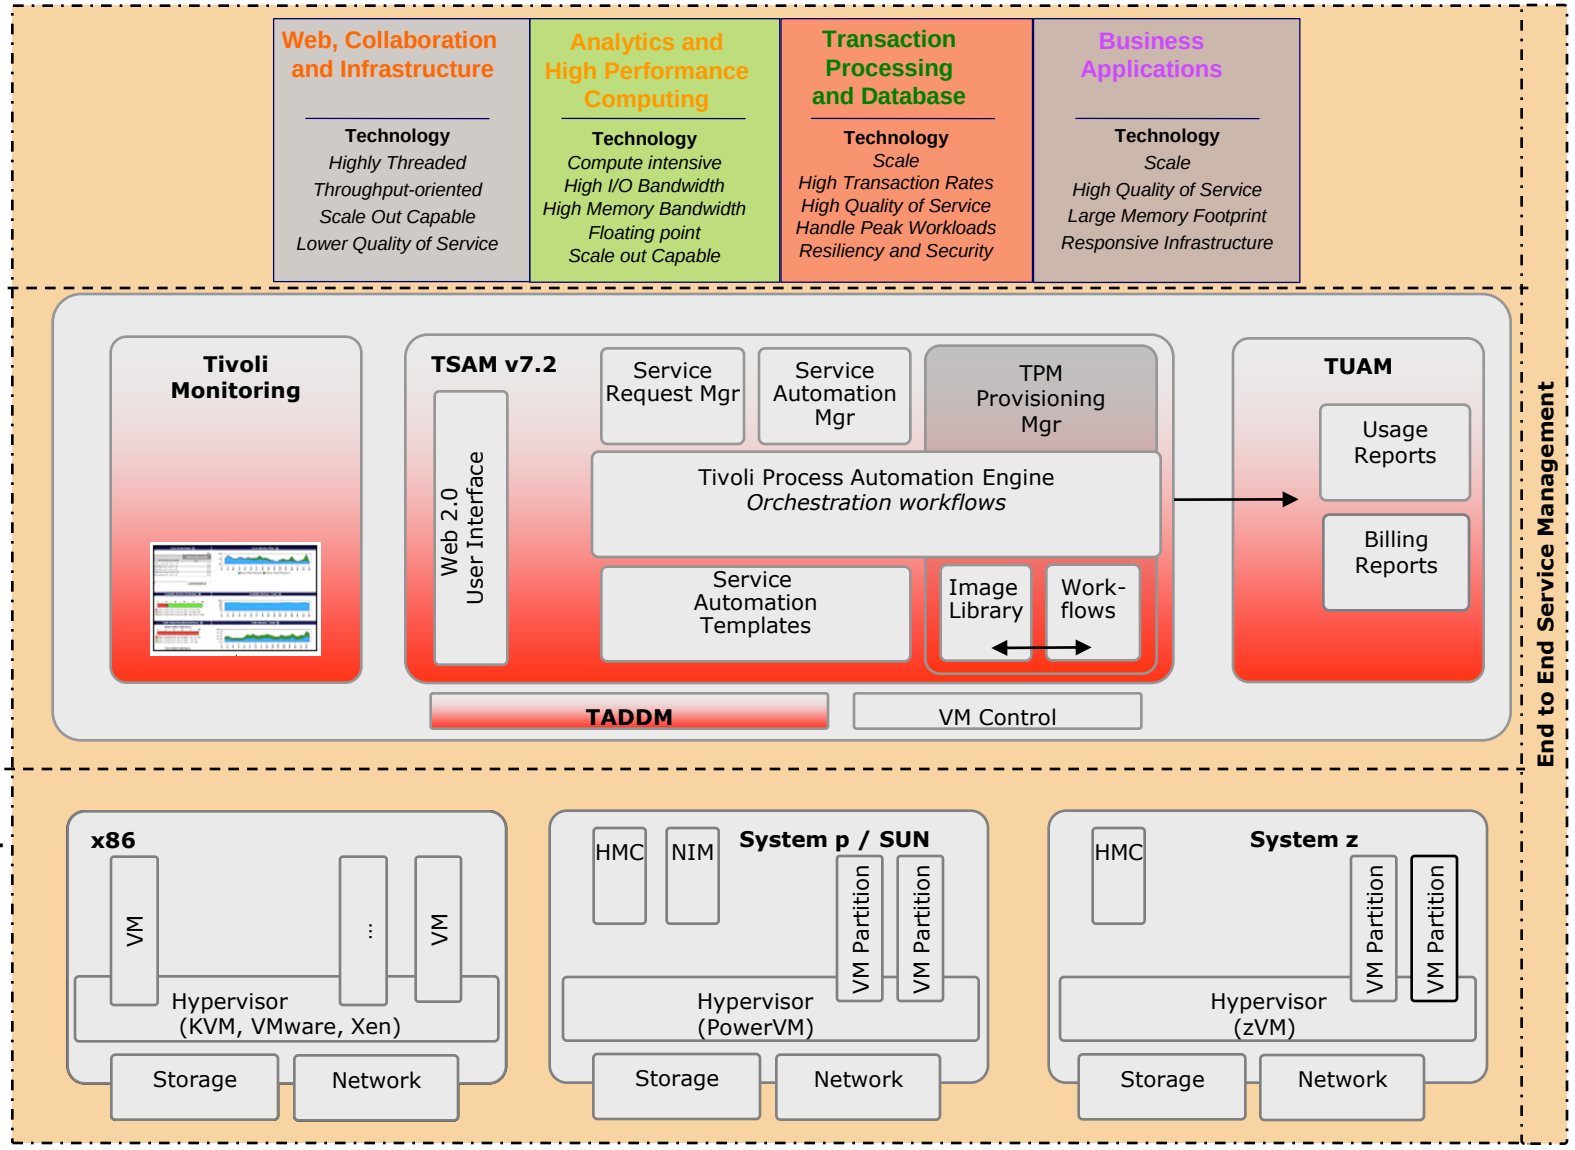

Traditional Data Center Management vs. "Cloud-like" Management

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!

Agenda

- *Introduction*
- *Cloud Computing Reference Architecture*
  - *Cloud Computing Management Platform*
  - *Selected Management Areas*
- *Hybrid Clouds*
- *Customer Projects*
- *Standardization Efforts*
- *Summary*

Hybrid Cloud Management, Security and Integration

- From the Enterprise Client's perspective:
- ***Management of workloads running off-premise on clouds***
  - *Management of software applications and services (monitoring, events, availability, performance)*
  - *Service Request Management (governance of service provisioning)*
  - *Dashboard for service visibility*
- ***Security for Hybrids***
  - *Control security and resilience of services (identity management, compliance, isolation)*
- ***Integration of applications & data***
  - *On-premise to off-premise business application connectivity & governance*
  - *Information exchange and data integration across the enterprise and clouds*
- ***Application and Workload migration workbench***
  - *Tools to support the migration of workloads to the cloud*

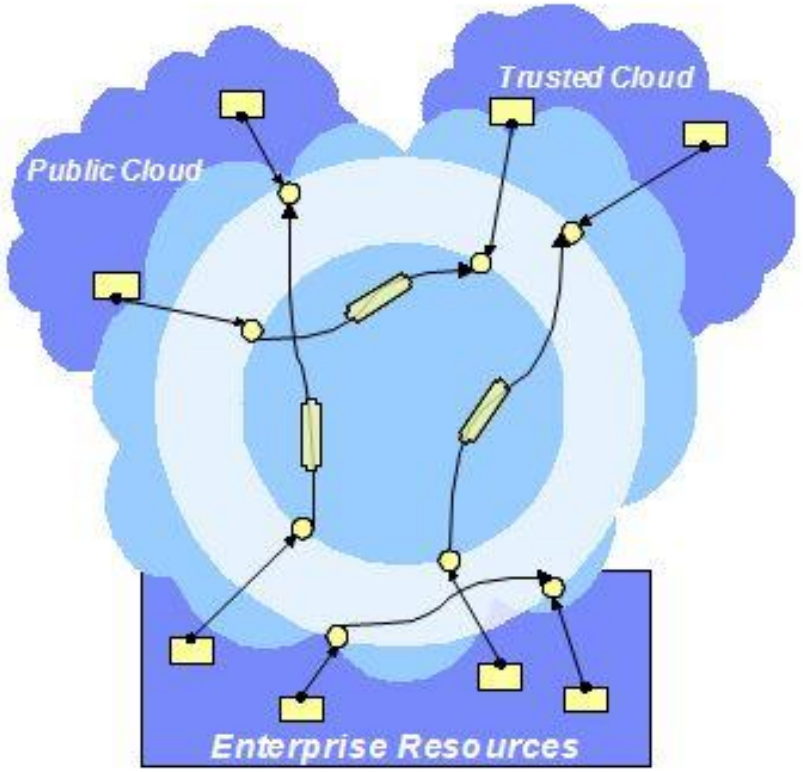

Initial focus for 'Hybrid Cloud':
'Provide clients the ability to manage and integrate workloads and resources on a cloud with their existing processes, management and business systems.'

IBM + Cast Iron combines enterprise level scalability and support with rapid on & off premise application integration

Today

- Separate technologies to manage application integration requirements
- Fragmented infrastructure/device sprawl
- Duplicate integration processes

IBM + Cast Iron

- Single, integrated platform for on- to on-, on- to off-, and off- to off-premise application integration
- Uniform infrastructure
- Shared application integration processes

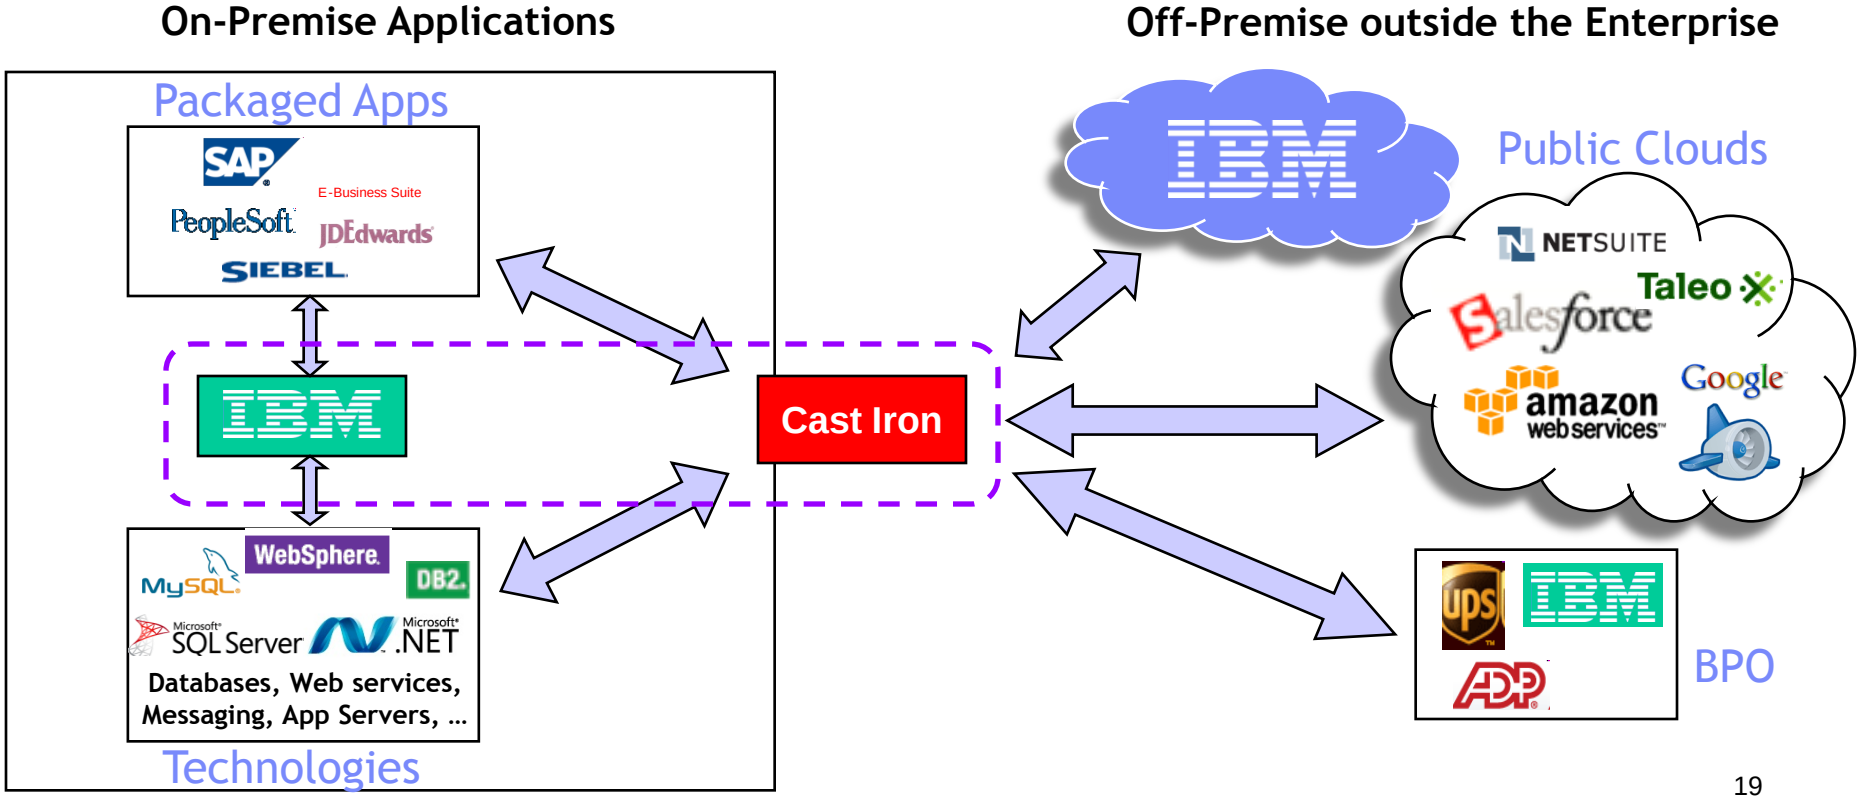

Agenda

- *Introduction*
- *Cloud Computing Reference Architecture*
  - *Cloud Computing Management Platform*
  - *Selected Management Areas*
- *Hybrid Clouds*
- *Customer Projects*
- *Standardization Efforts*
- *Summary*

Emerging Customer Patterns

Self Service Provisioning

Compelling entry-point into Cloud Computing, particularly for Development /Test environments,

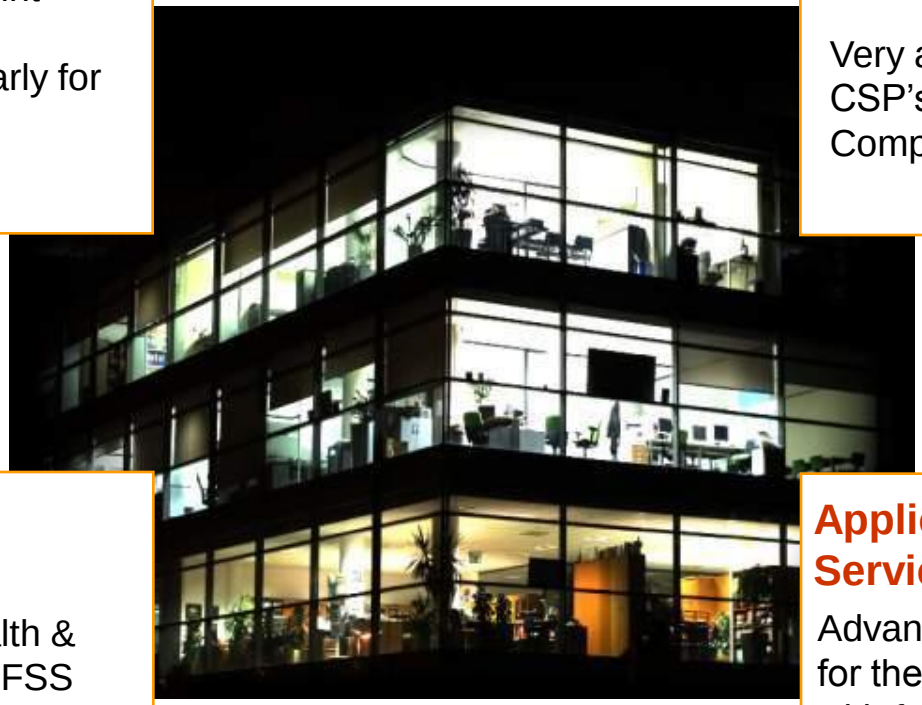

Cloud Service Delivery Platform

Very active with CSP's, Telco's. High Competition

Analytics

Heavy interest in Health & Pharma, emerging in FSS

Application / Platform Service

Advanced enterprises looking for the "big bang" of Cloud, with focus on increasing & optimizing existing infrastructure utilization

Self Service Delivery Project – Financial Customer

Low-cost, low-touch self-enablement server provisioning system that leverages automation around virtualized server and storage infrastructure

Primary Focus Areas	*Supporting Capability*	*Key Metrics*
Improve Efficiency	**Self-service portal with automated provisioning** *Move from traditional high touch provisioning model to a self-service, full-lifecycle, reservation model with automated provisioning*	time to market consistency flexibility server/admin ratio
Improve Efficiency	**Management of the full lifecycle of a server** *Systems can be reserved, provisioned and de-provisioned based on schedule and capacity*	systems utilization systems capacity
Improve Quality	**Image management** *Temporarily restore servers for further testing*	time to market flexibility systems capacity
•AIX LPARs on IBM p5/6 *•Linux and Windows images on x86*	**Policy management and governance** *Consistency of server provisioning and configuration. Flexibility and control over request/approval workflows, resource assignment, utilization and capacity, and cost allocation*	consistency visibility and control systems capacity

IBM Tivoli Development Cloud

Business Background

- IBM Tivoli Development Services (TDS) organization provides IT services for Tivoli and other organizations in IBM Software Group and Research
- IT Footprint had expanded to 24 labs through growth and acquisitions, creating inefficiencies and increased expense
- Plan to exploit Tivoli capabilities to:
  - More effectively manage resources and IT services in the cloud
  - Innovate new business services through process transformation

Cloud Business Benefit

- Transformed business and IT processes
- Improved competitiveness through faster time to value and enhanced productivity
- Avoided $4.8M in capital expense and $3.1M in operational expense in 2009 through consolidation, virtualization and automation
- Consolidated 5 of 24 labs, reduced physical space by 8% and built capacity for 1200 virtual machines.

Solution Overview

- IBM Tivoli Development Cloud implemented with Tivoli Service Automation Manager, Tivoli Provisioning Manager, IBM Tivoli Monitoring, Storage Productivity Center, OMNIbus, Tivoli Business Service Manager, Tivoli Data Warehouse, Tivoli Performance Analyzer
- Infrastructure includes KVM,VMWare and Hyper-V based virtualized images on IBM System X hardware

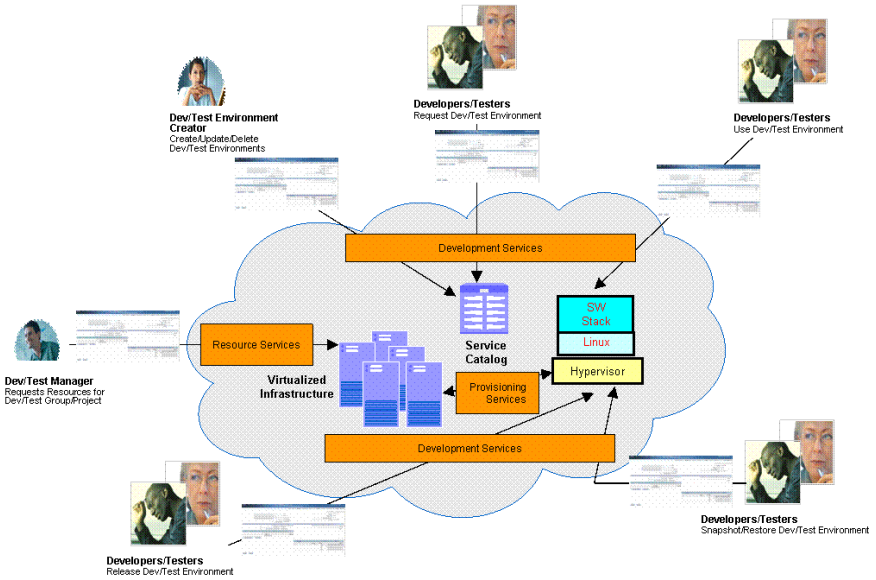

SKT's vision of Cloud Computing

Cloud Computing platform needs to be deployed that enables mobile content providers and business partners with a mobile service idea to develop, test and commercialize new services quickly and easily.

Business Needs

"Strengthen the Competitiveness of the SKT Internet Service & Create new business opportunities for Platform service"

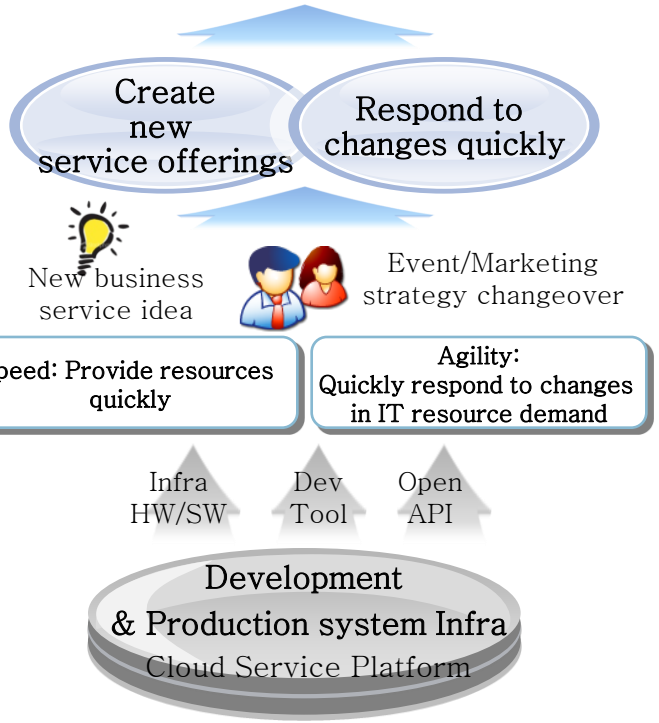

Project Objective

-Provide Better and flexible service to users (CP/BP), enabling self-service request and delivering services more rapidly
To leverage CP/BP who has a new business service ideas
-Reduce cost for operations & management and for new investment

- Improve time to market – react to deliver a new IT service quickly, decrease time to deploy systems for new service offerings
- Lower development cost – increase resource utilization and reduce labor costs
- Find new revenue/profit streams thru embrace a new business service ideas of CP/BP quickly.

United States Air Force
Mission Oriented Cloud Architecture

Business Background

- The United States Air Force (USAF) provides aerial, space and cyber warfare for the United States Armed Forces. The USAF consists of 10 major commands, 100 military bases, and 700,000 personal worldwide

Business Benefit

- IBM will provide research, design and demonstration a secure cloud computing infrastructure for the USAF.
- IBM is helping the USAF understand how to manage, monitor, and secure the information flowing through the USAF, Department of Defense and other intelligence agency networks.
- IBM will demonstrate an unprecedented level of security, network resiliency to the USAF networks.
- The resulting architecture will provide the USAF with an advanced level of “Situational Awareness” by implementing sensors, monitors, detection devices, security policy management, compliance management, and advanced analytic stream processing.
- The new cloud architecture will reduce the time it takes to respond to cyber threats by leveraging automated mission prioritized workload and capacity management systems.

Solution Overview

- Demonstration of a security focused cloud computing architecture that can manage, monitor and secure the information flowing through the Air Force network.
- Advanced analytic processing from InfoSphere Streams coupled via sensors, monitors, and other detection devices
- Automated mission prioritized capacity management
- Real-time situational awareness of the cloud environment
- Policy based security compliance reporting and enforcement
- IBM hardware – System x , BladeCenter, DataPower, ISS Proventia
- IBM software – Tivoli, Rational, WebSphere and InfoSphere

IBM Technology Adopter's Portal (IBM TAP)

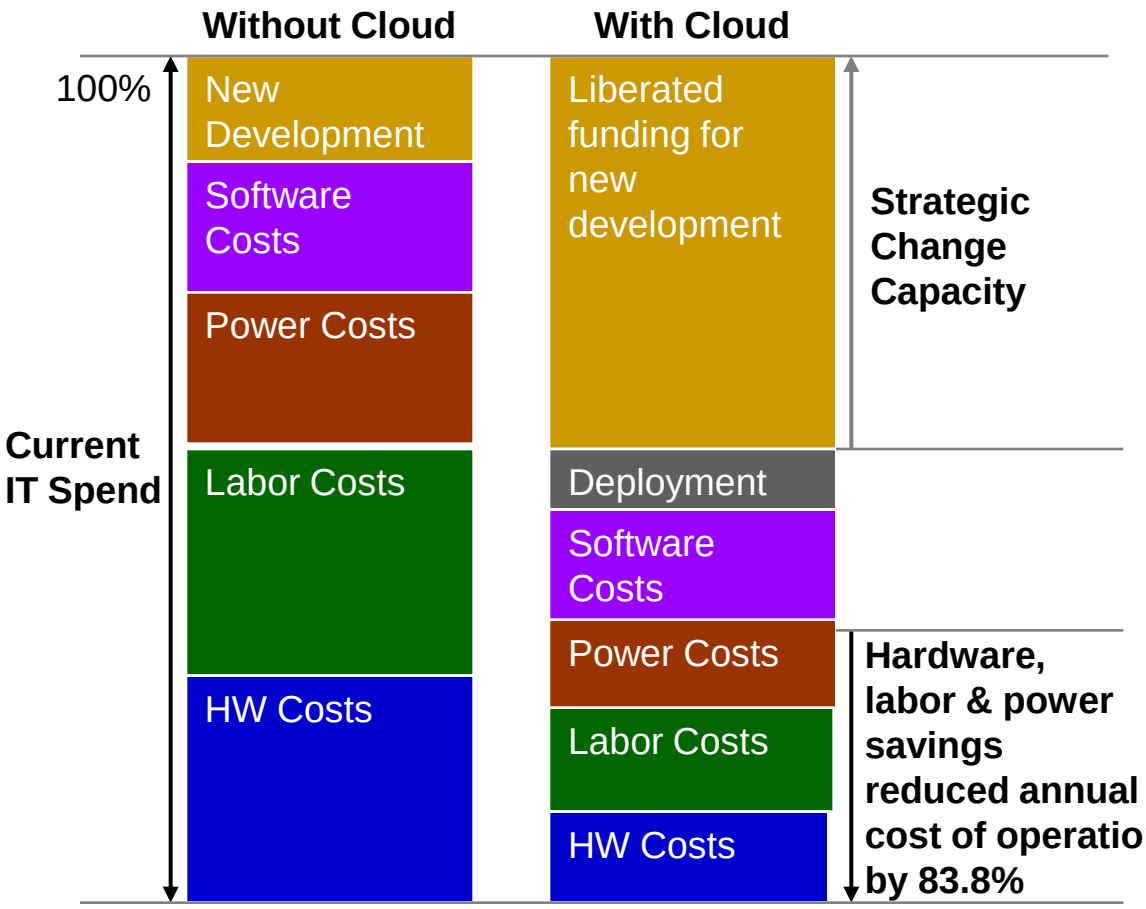

- Innovation Cloud for 100,000 Subscribers
- Reduced Capital Expenditure
  - Reduce from 488 servers to 55
- Reduced Operations Expenditure
  - Reduce from 15 admins to 2
- Additional Benefits:
  - Enhanced customer service
  - Less idle time
  - More efficient use of energy
  - Acceleration of innovation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

Agenda

- *Introduction*
- *Cloud Computing Reference Architecture*
  - *Cloud Computing Management Platform*
  - *Selected Management Areas*
- *Hybrid Clouds*
- *Customer Projects*
- *Standardization Efforts*
- *Summary*

Service Definitions and Service Instances

Service Definition

including a topology model for the services, and build- and management plans

Service Instances

created from template definitions in a Service Definition, representing one deployed service

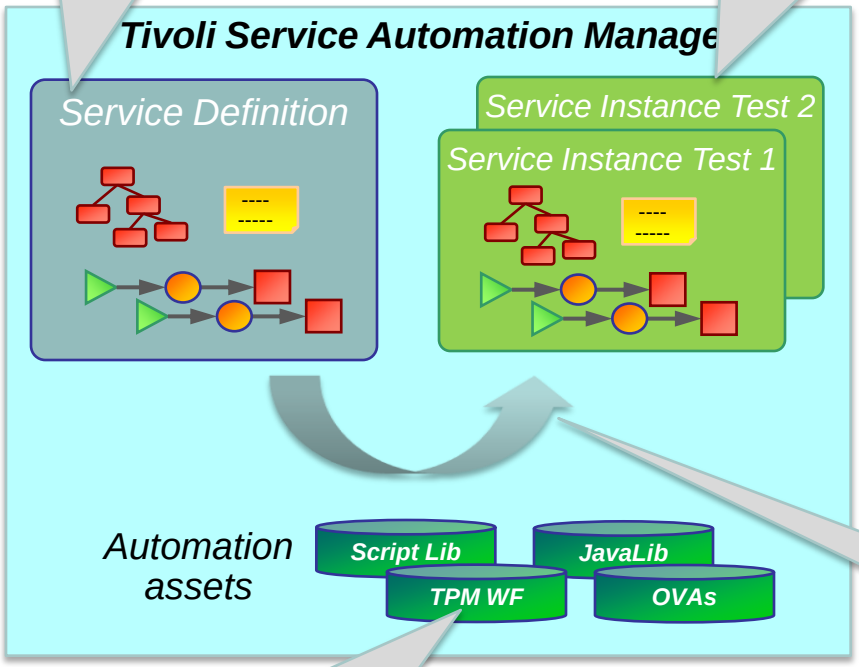

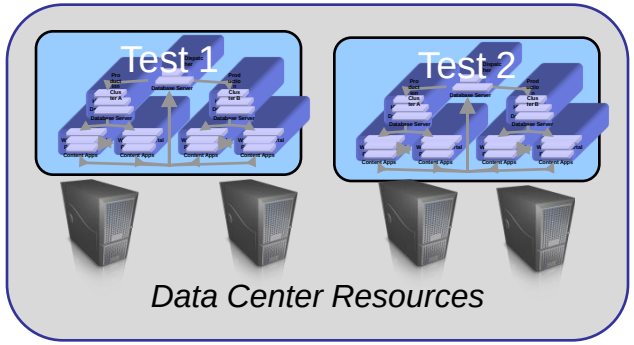

Automation Assets

leveraged by build- and management plans, e.g. OVF images, TPM workflows, scripts, ...

Instantiation

of services from a Service Definition, parameterized through user input filling in point of variability of the Service Definition

Service Definition Overview

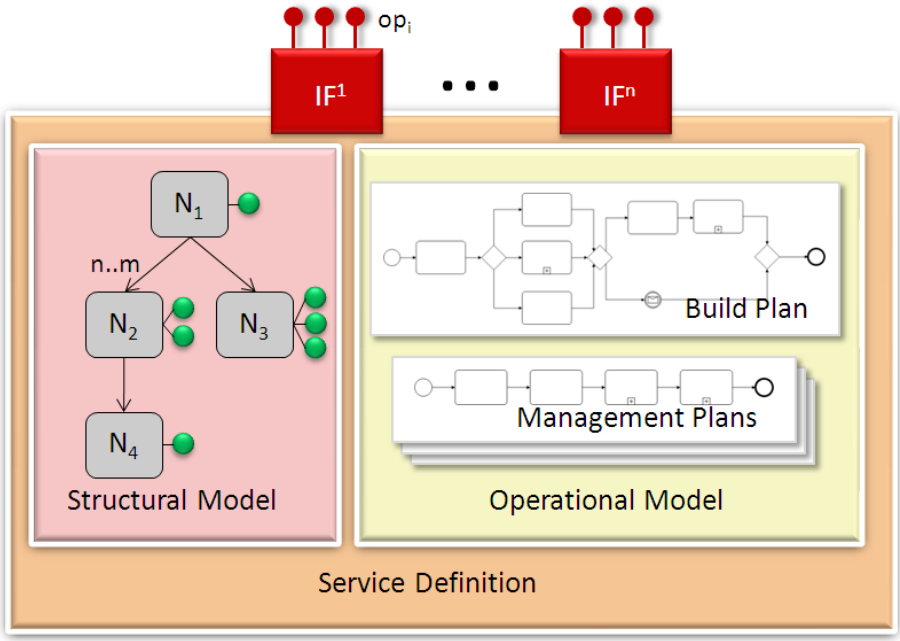

- **Service Definition provides a model for managing Cloud Services throughout their complete lifecycle:**
  - Initial Deployment of a service instance
  - Operational management of a service instance (e.g. modify capacity, patch management, upgrades, incident and problem management, etc.)
  - Termination of a service instance
- **Service Topology Template:**
  - Structural model of a service, i.e. its components and their relationships
  - Includes operations that can be invoked on service components as the basis for instrumentation
- **Build- and Management Plans:**
  - Process model of how to set up, manage and terminate a service
  - Plans are represented in BPMN
- **Interfaces describing the CRUD operations that can be executed on Cloud Service Definitions and Cloud Service Instances**
  - For orchestration of the service, and for creation and management of composed services (Hybrid Clouds)

Agenda

- *Introduction*
- *Cloud Computing Reference Architecture*
  - *Cloud Computing Management Platform*
  - *Selected Management Areas*
- *Hybrid Clouds*
- *Platform Exploitability and Customer Projects*
- *Standardization Efforts*
- *Summary*

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 sucessfully 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!

Agenda

- *Introduction*
- *Cloud Computing Reference Architecture*
  - *Cloud Computing Management Platform*
  - *Selected Management Areas*
- *Hybrid Clouds*
- *Platform Exploitability and Customer Projects*
- *Standardization Efforts*
- *Summary*
- *References*

References

- Articles downloaded from the Internet
  - Michael Armbrust et al., Above the Clouds: A Berkeley View of Cloud Computing, Feb. 2009
    - http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf
  - Cloud Computing: Platform as a Service. InformationWeek Analytics, October 2, 2009
  - Luiz André Barroso and Urs Hölzle, The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines, Synthesis Lectures on Computer Architecture, 2009, http://www.morganclaypool.com/doi/pdf/10.2200/S00193ED1V01Y200905CAC006?cookieSet=1
  - Scott Crowder, Introduction to Workload Optimized Approach & Workload Market Segmentation, IBM White Paper, December 2009
  - David Chappell, A short introduction to Cloud, http://www.davidchappell.com/CloudPlatforms--Chappell.pdf
  - David Chappell, Cloud Platforms Today: A Perspective, April 2009 http://www.davidchappell.com/CloudPlatformsToday--APerspective--Chappell.pdf
  - Jeffrey Dean and Sanjay Ghemawat, MapReduce: Simplified Data Processing on Large Clusters,
    - *labs.google.com/papers/**mapreduce**-osdi04.pdf*
  - DeCandia et al. Dynamo: Amazon's highly available key-value store, SOSP 2007, http://portal.acm.org/citation.cfm?id=1294281&dl=ACM&coll=ACM&CFID=47859964&CFTOKEN=98797782
  - European Network and Information Security Agency (ENISA), Cloud Computing, Benefits, risks and recommendations for information security, Nov 2009 (http://www.enisa.europa.eu)
  - Gregor Hohpe, Programming the Cloud, November 2009, http://www.enterpriseintegrationpatterns.com/docs/HohpeProgrammingCloudKeynote.pdf
  - Anna Liu, Architecting Cloud Applications – the essential checklist, AAF Keynote 2009,
  - National Institute of Standards and Technology, Definition of Cloud Computing, http://csrc.nist.gov/groups/SNS/cloud-computing/
  - Ning Duan et al., Tenant Behavior Analysis in Software as a Service Environment, ICSOC 2009
  - Daniel Nurmi et al., The Eucalyptus Open-source Cloud-computing System, http://www.cca08.org/papers/Paper32-Daniel-Nurmi.pdf
  - Open Cloud Manifesto, http://www.opencloudmanifesto.org/
  - OpenNebula.org – Various papers
  - B. Rochwerger et al., The Reservoir Model and Architecture for Open Federated Cloud Computing, IBM Journal of Research and Development, April 2009 http://www8.cs.umu.se/~elmroth/papers/ibmjrd2009.pdf
  - Werner Vogels, Eventually Consistent, ACM Queue, October 2008
  - Ying Huang et al., A Framework for Building a Low Cost, Scalable and Secured Platform for Web-Delivered Business Services, IBM Systems Journal, November 2009
- Company Web Sites: Amazon, Microsoft, Google, IBM, Salesforce.com
- Gregor Hohpe, Bobby Woolf, Enterprise Integration Patterns, Addison-Wesley 2004
- Kristof Kloeckner, Middleware for Distributed Systems, Lecture Notes 2004
- Kristof Kloeckner, The IBM Cloud Agenda, White Paper 2009
- George Reese: Cloud Application Architectures, O'Reilly 2009
- John W. Rittinghouse, James F. Ransome, Cloud Computing. Implementation, Management and Security, CRC Press 2009
- Andrew Tanenbaum, Maarten van Steen: Distributed Systems. Principles and Paradigms, Prentice-Hall 2009
- Rich Schiesser: IT Systems Management, Prentice-Hall 2002
- Jim Rymarczyk, Virtualization, Pre-Print 2009
- Tivoli Service Automation Manager Solution Guide

References & Reading List: General Information Security

- [Allan 10] Allan, D., Hahn, T., Szakal, A., Whitmore, J. and Buecker, A. Security in Development: The IBM Secure Engineering Framework. IBM RedGuide, New York, 2010.
- [Anderson 08] Anderson, R. Security Engineering: A Guide to Building Dependable Distributed Systems (2nd ed). Wiley, Indianapolis, 2008.
- [Bishop 02] Bishop, M. Computer Security, Art and Science. Addison-Wesley, Boston, 2002.
- [BSI 100] BSI-Standard 100-1-4. IT Grundschutz. Bundesamt für Sicherheit in der Informationstechnik, Bonn, 2009.
- [Buecker 09] Buecker, A. et. al. Introducing the IBM Security Framework and IBM Security Blueprint to Realize Business-Driven Security. IBM RedGuide, New York, 2009.
- [Cheswick 03] Cheswick, W., Bellovin, S. and Rubin, A. Firewalls and Internet Security: Repelling the Wily Hacker (2nd ed). Addison-Wesley, 2003.
- [Crawford 10] Crawford, S. High Performers and Foundational Controls: Building a Strategy for Security and Risk Management. EMA White Paper, 2010.
- [Gasser 88] M Gasser: Building a Secure Computer System. Van Nostrand Reinhold, New York, 1988.
- [Goldberg 74] Goldberg, R. P. Survey of Virtual Machine Research. IEEE Computer June. 1974, pp 34-45.
- [IBM 10] IBM X-Force 2009 Trend and Risk Report. IBM, Armonk, February 2010.
- [ISO 27001] ISO/IEC 27001. Information Security Management System. ISO 2005.
- [McGraw 06] McGraw, G. Software Security: Building Security In. Addison-Wesley, 2006.
- [Menezes 96] Menezes, A., van Oorschot, P. and Vanstone, S. Handbook of Applied Cryptography. CRC Press, 1996.
- [OWASP 10] OWASP Top 10 – 2010. The Ten Most Critical Web Application Security Risks. OWASP Foundation, 2010.
- [Stallings 10] Stallings, W. Network Security Essentials: Applications and Standard (4th ed). Prentice Hall, Upper Saddle River, 2010.

References & Reading List: General Cloud Security

- [CSA 10] Top Threats to Cloud Computing. Cloud Security Alliance (CSA), 2010.
- [ENISA 09] Cloud Computing: Benefits, Risks and Recommendations for Information Security. European Network and Information Security Agency (ENISA), 2009.
- [ENISA 09a] Cloud Computing Information Assurance Framework. European Network and Information Security Agency (ENISA), 2009.
- [Heiser 09] Heiser J. and Nicolett M. Assessing the Security Risks of Cloud Computing. Gartner Research, 2008
- [IBM 09] IBM Point of View: Security and Cloud Computing. IBM, New York, 2009.
- [Jericho 09] Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration. The Jericho Forum, 2009.
- **[Mather 09] T. Mather, S. Kumaraswamy, S. Latif. Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance. O'Reilly, Gravenstein, 2009.**
- [Mell 09] Mell, P. and Grance, T. Effectively and Securely Using the Cloud Computing Paradigm. In ACM Cloud Computing Security Workshop, Chicago, 2009.

Thank you!

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

Or contact me at:
gbreiter@de.ibm.com