RESERVOIR:

Integrating Virtualization and Grid Technologies for Federated Cloud Computing

David Breitgand, Benny Rochwerger, IBM Haifa Research Lab, Israel

The research leading to these results is partially supported by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 215605.

Cloud Computing:

A style of computing where massively scalable IT-enabled capabilities are delivered as a service to external customers using Internet technologies.

Premise:

No single cloud can create a seemingly infinite infrastructure capable of serving massive amounts of users at all times, from all locations

RESERVOIR:

Investigate technologies for advanced Cloud Computing

Focus on technologies that enable to build a federation of cooperating computing clouds

Project Profile

- 3 Years EU FP7 project started in February
- Budget: € 17 million
- 13 partners from across industry, academia and standards bodies
- Selected as a NESSI strategic project
  - The Networked European Software and Services Initiative (NESSI) is an industrial consortium focusing on advancing research in the area Services Architectures and Software Infrastructures
- Public web site
  - http://www.reservoir-fp7.eu/

RESERVOIR in a Nutshell

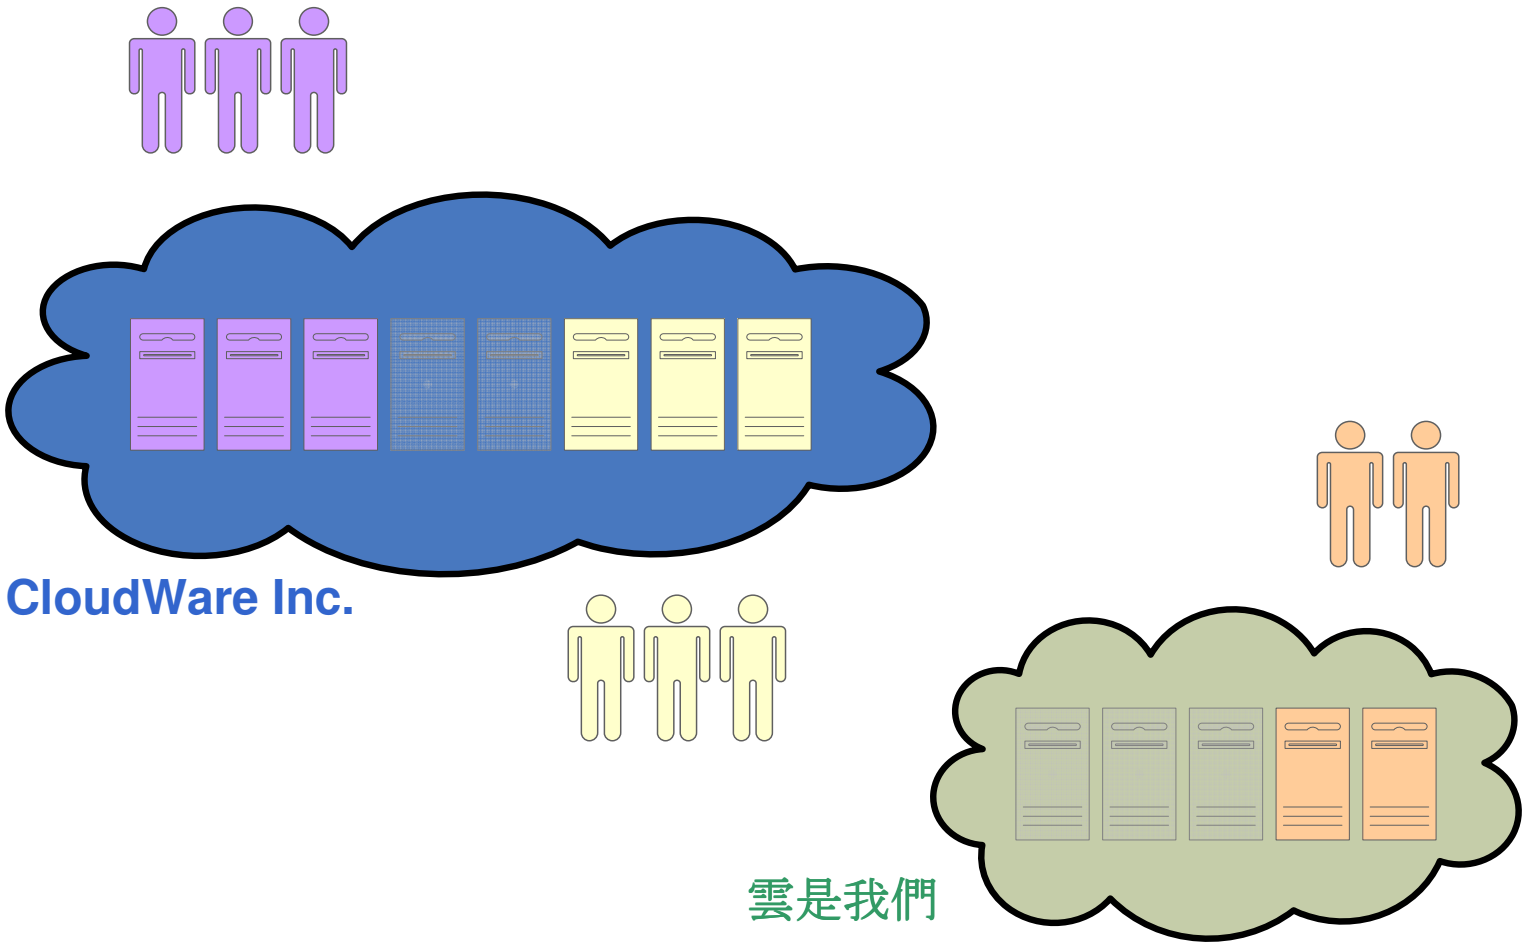

RESERVOIR in a Nutshell

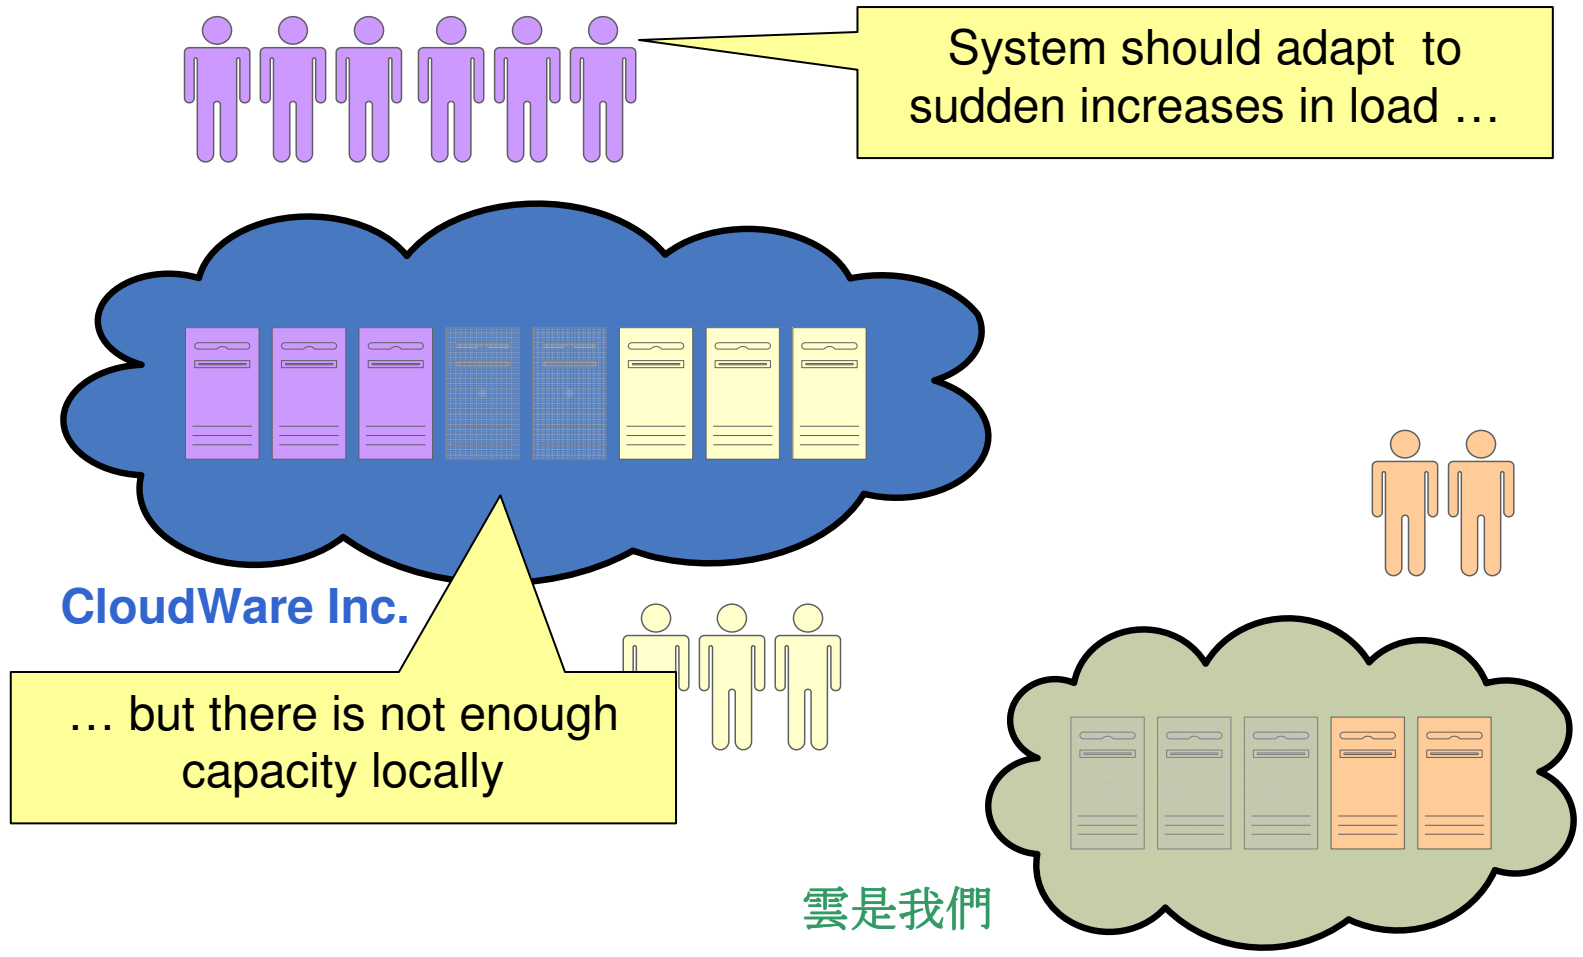

RESERVOIR in a Nutshell

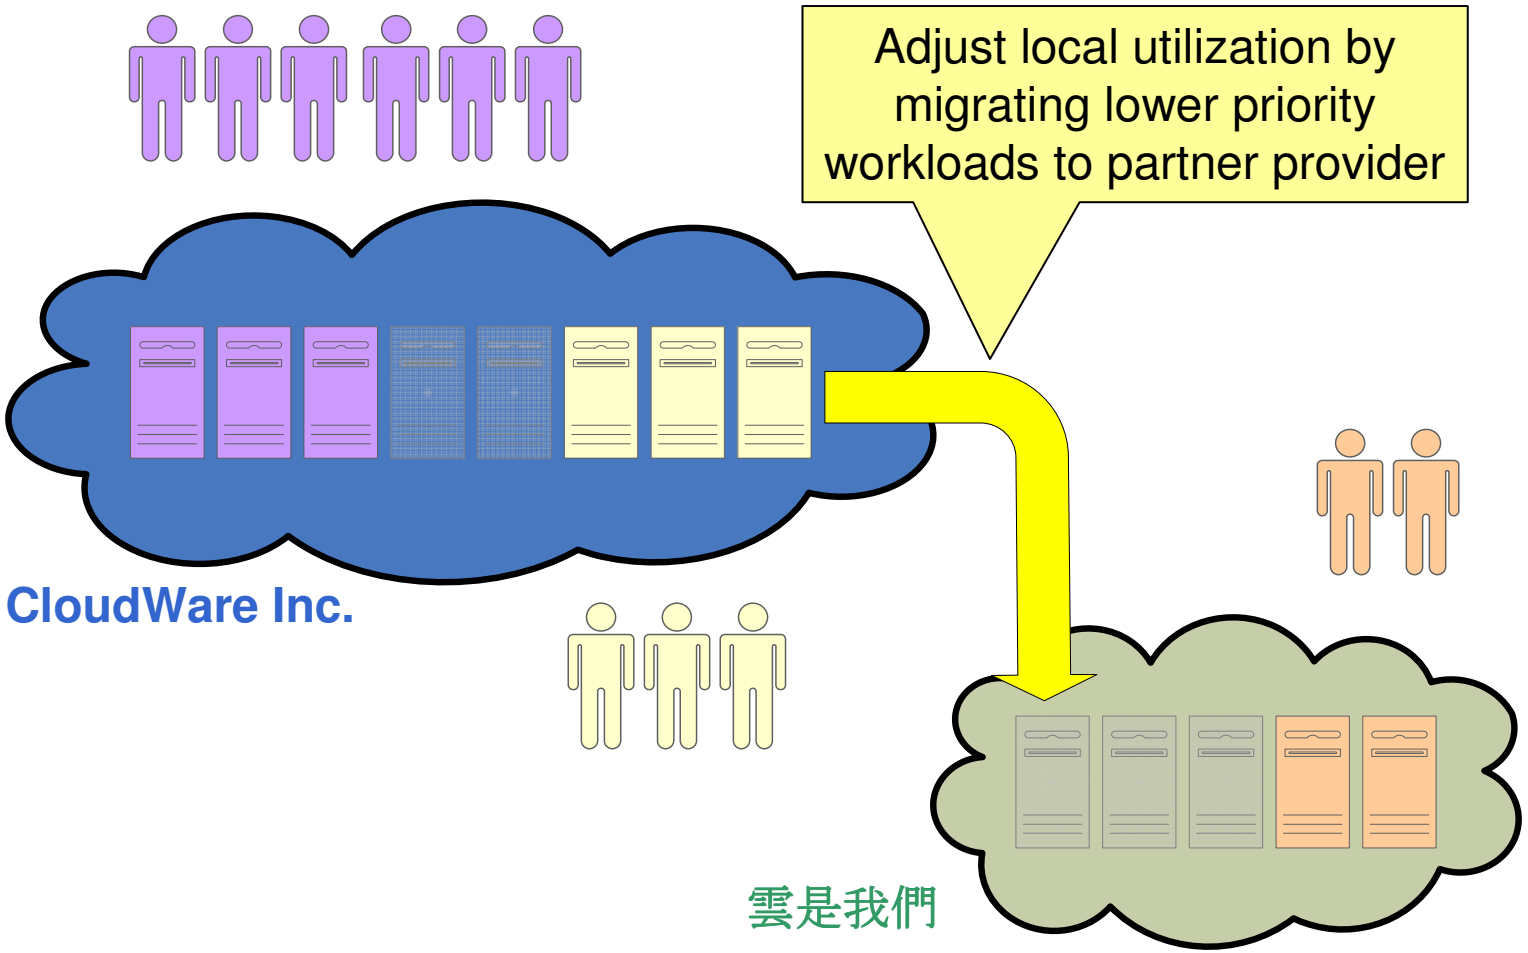

RESERVOIR in a Nutshell

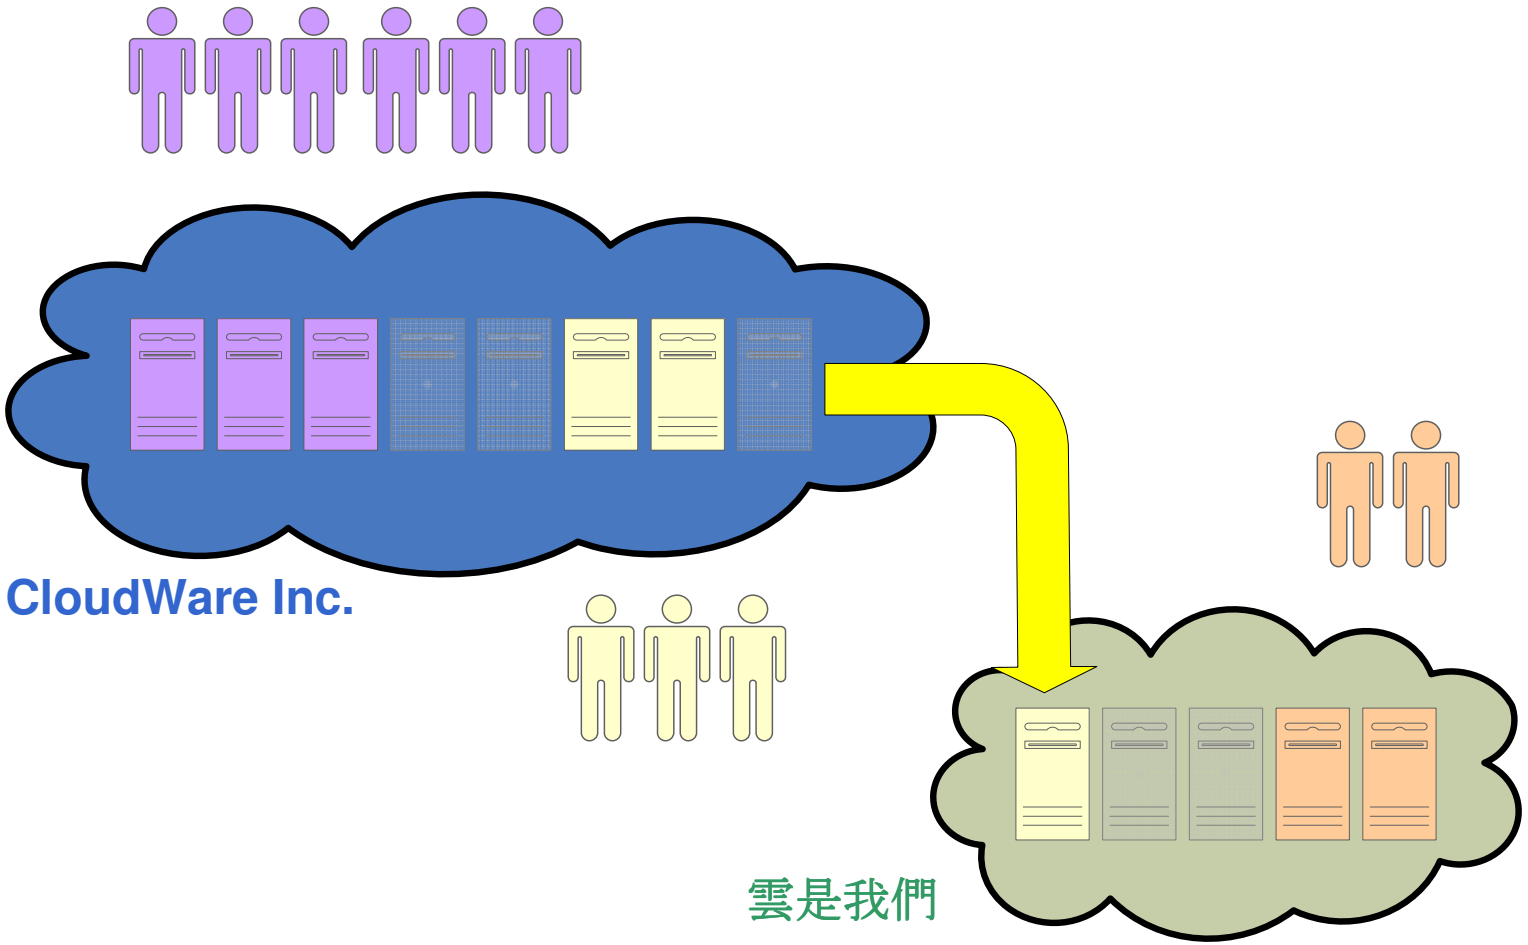

RESERVOIR in a Nutshell

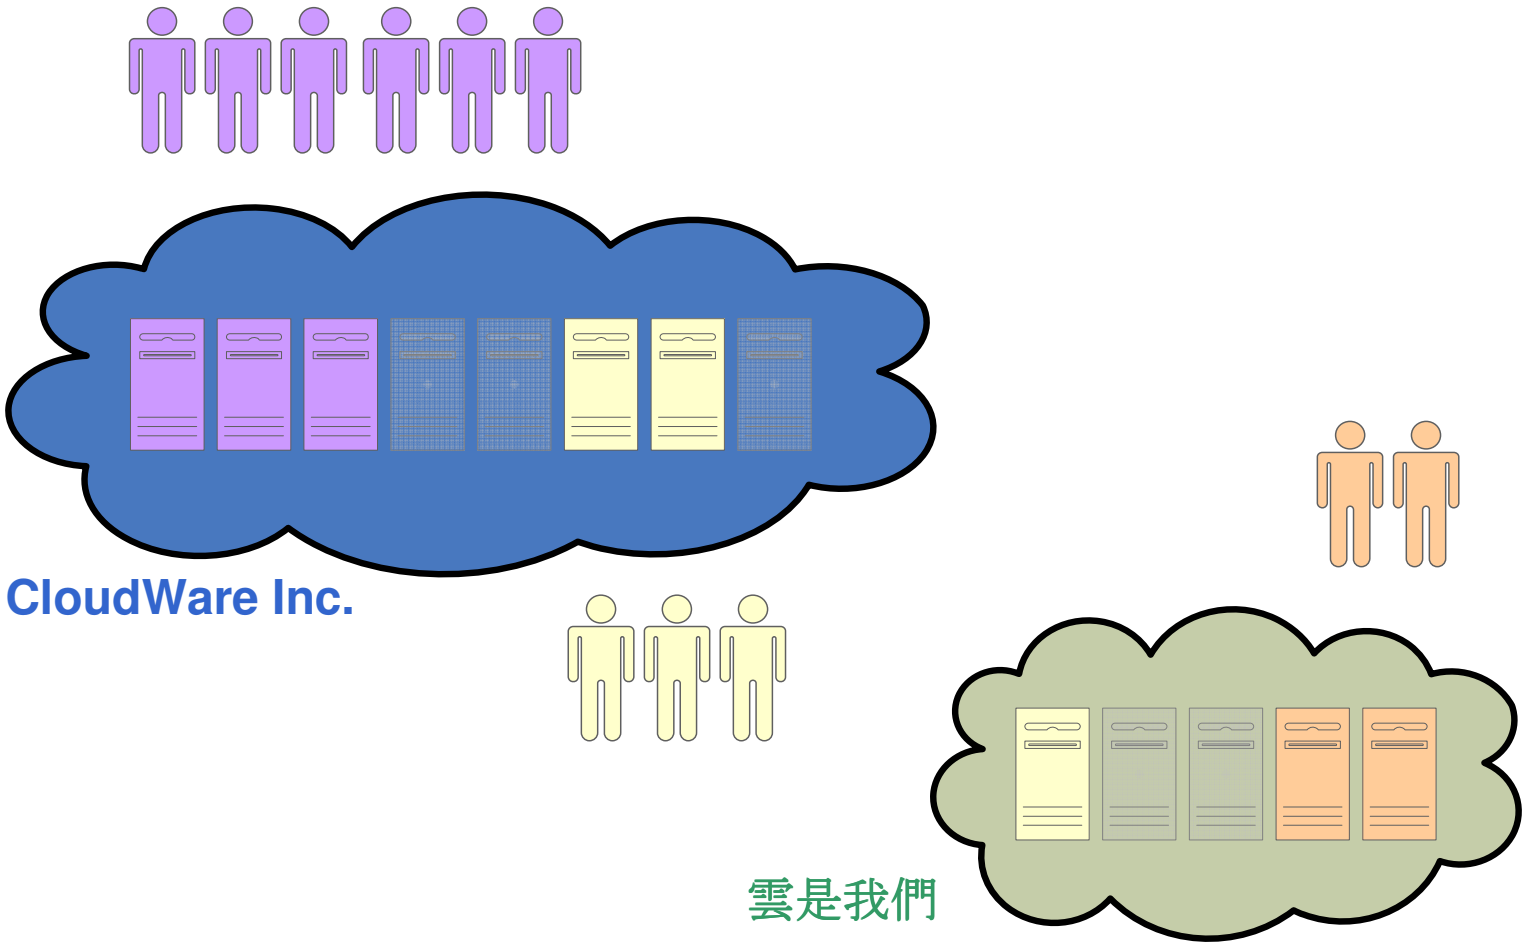

RESERVOIR in a Nutshell

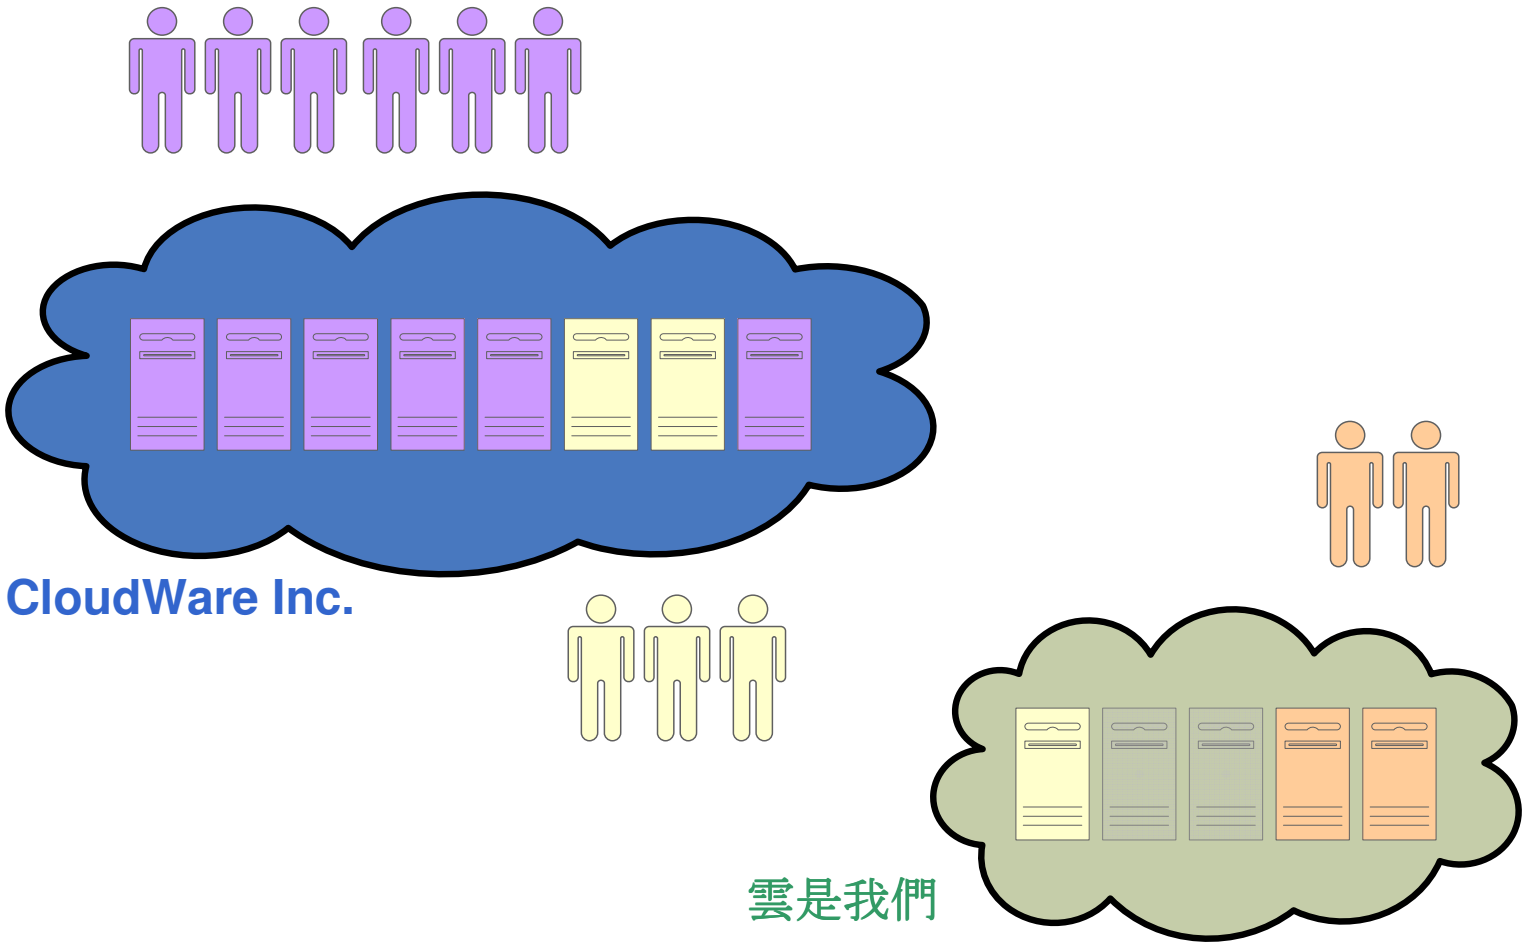

RESERVOIR in a Nutshell

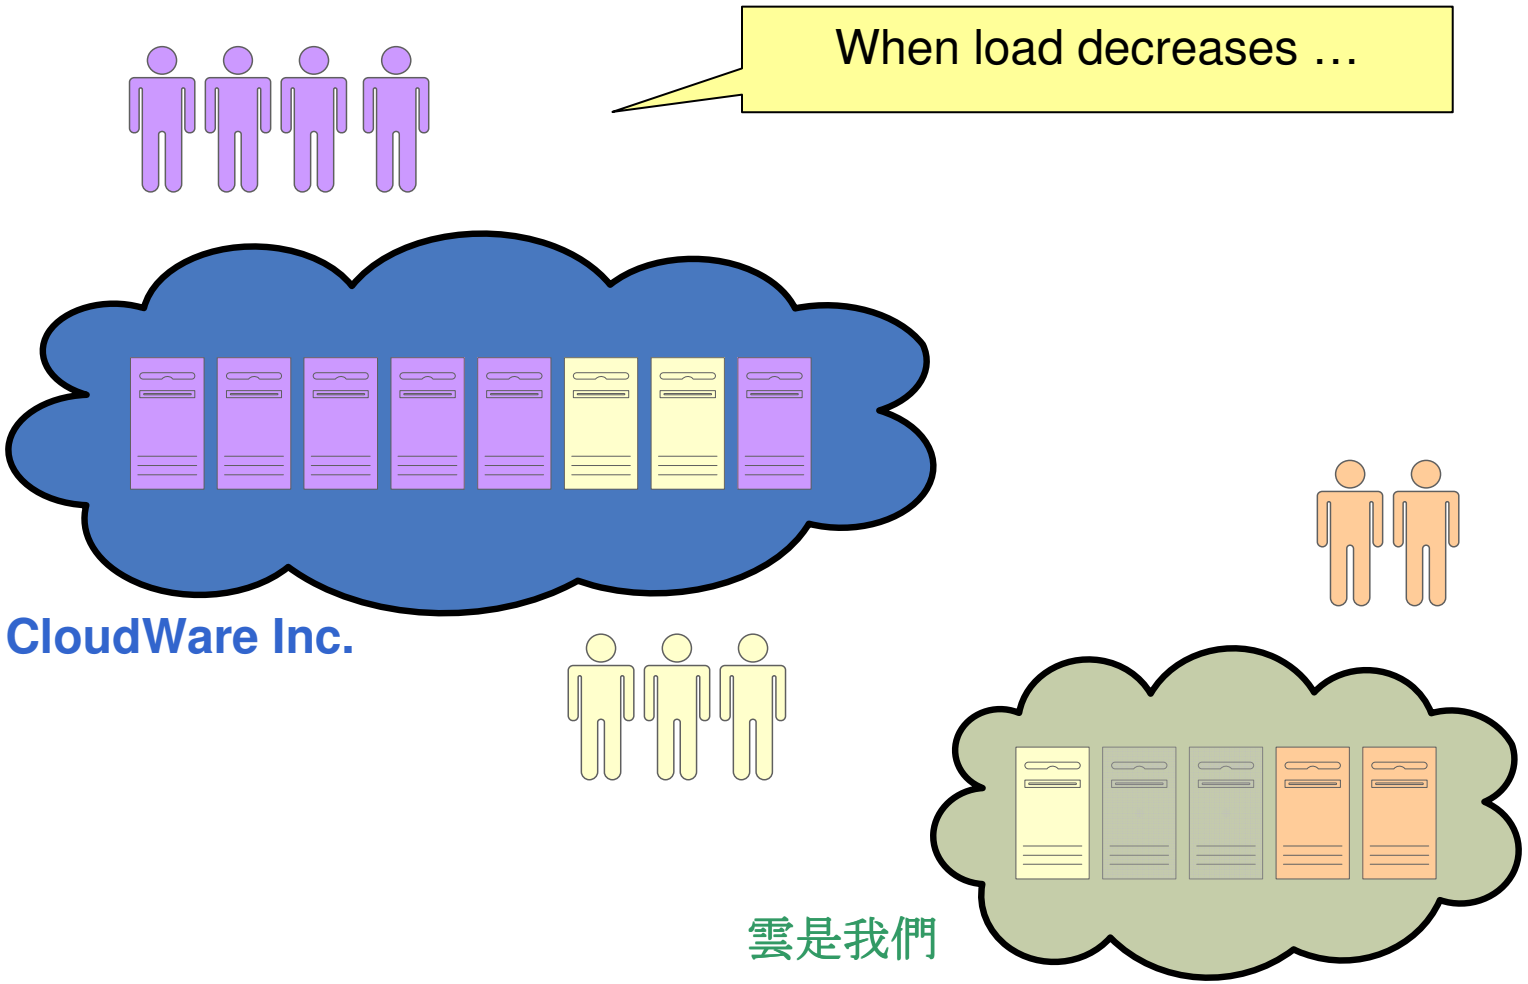

RESERVOIR in a Nutshell

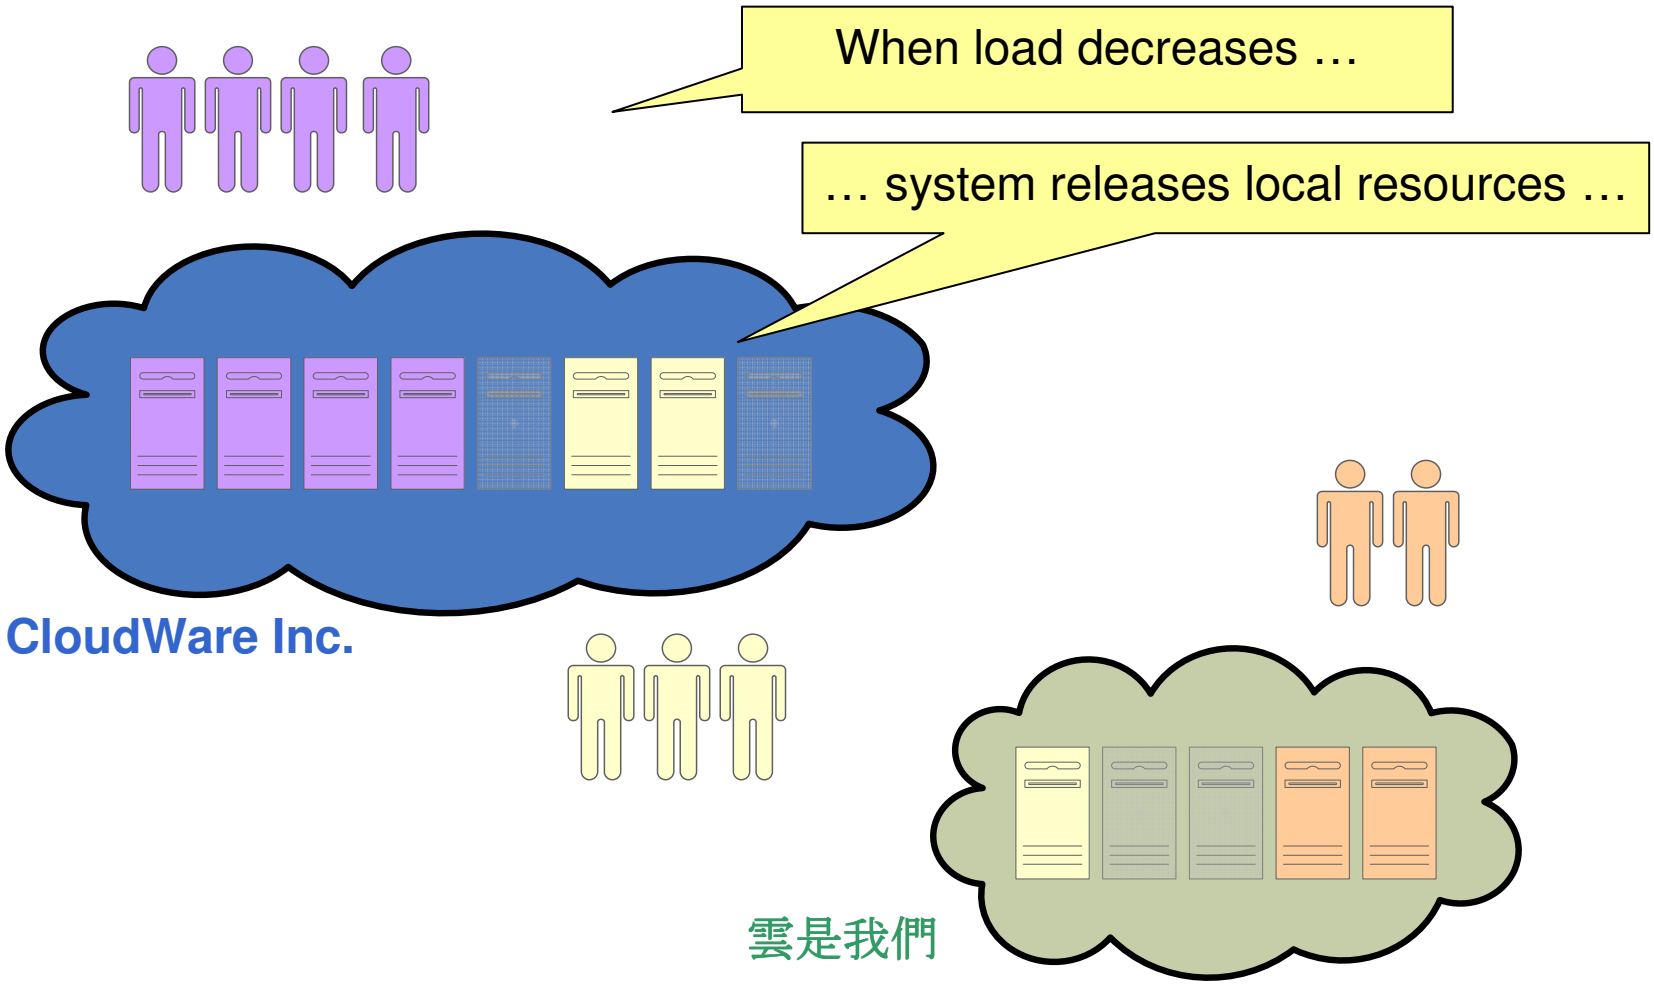

RESERVOIR in a Nutshell

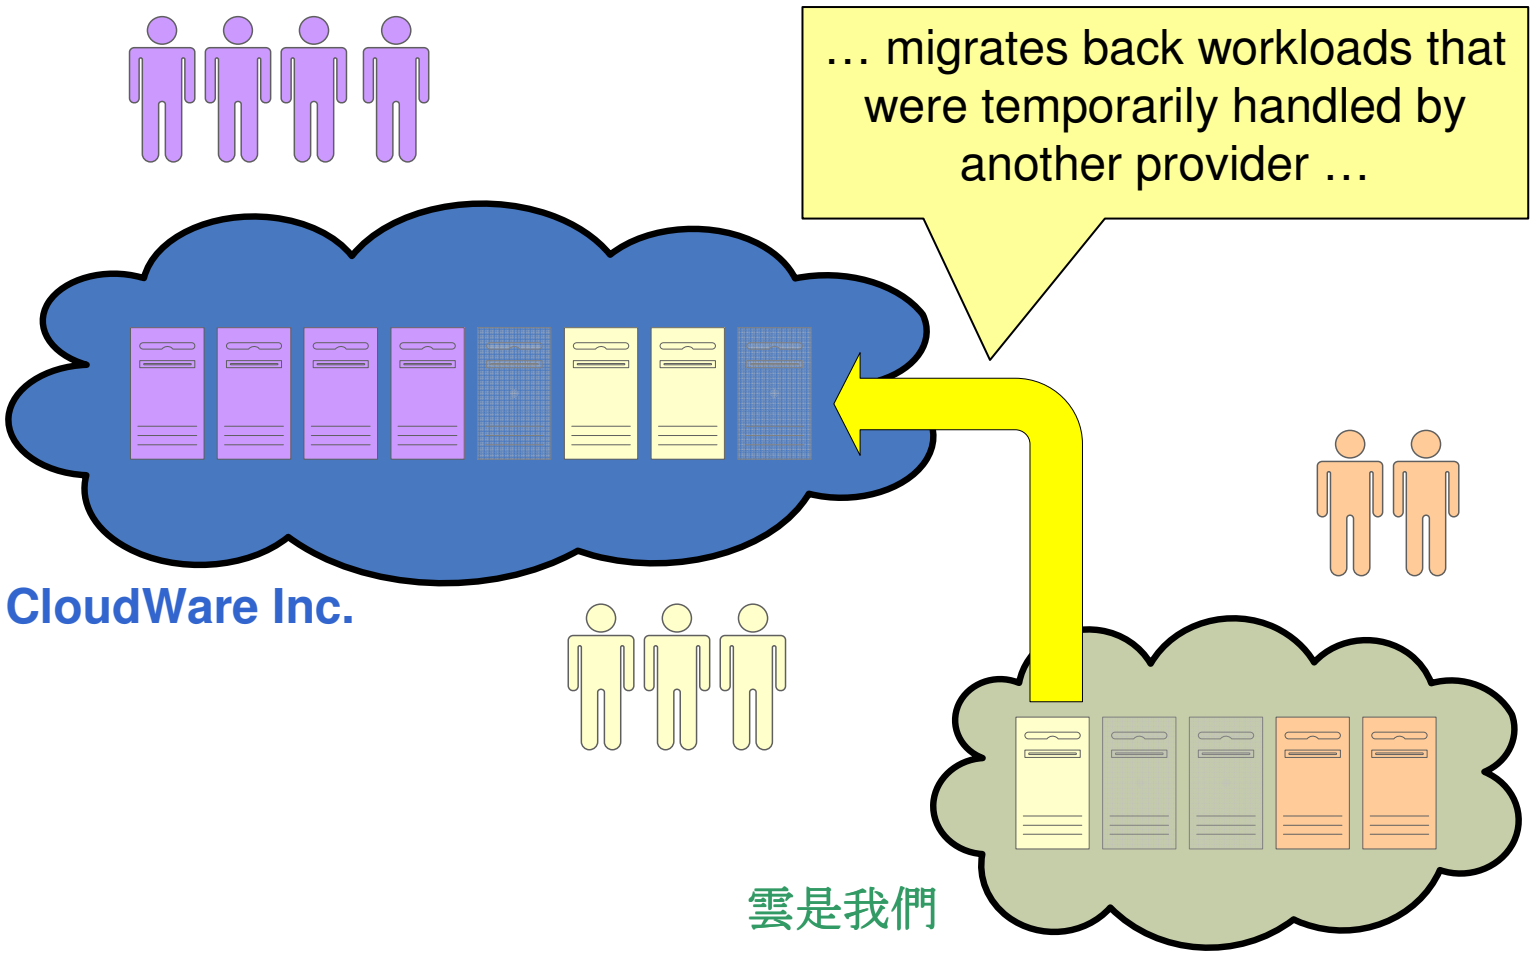

RESERVOIR in a Nutshell

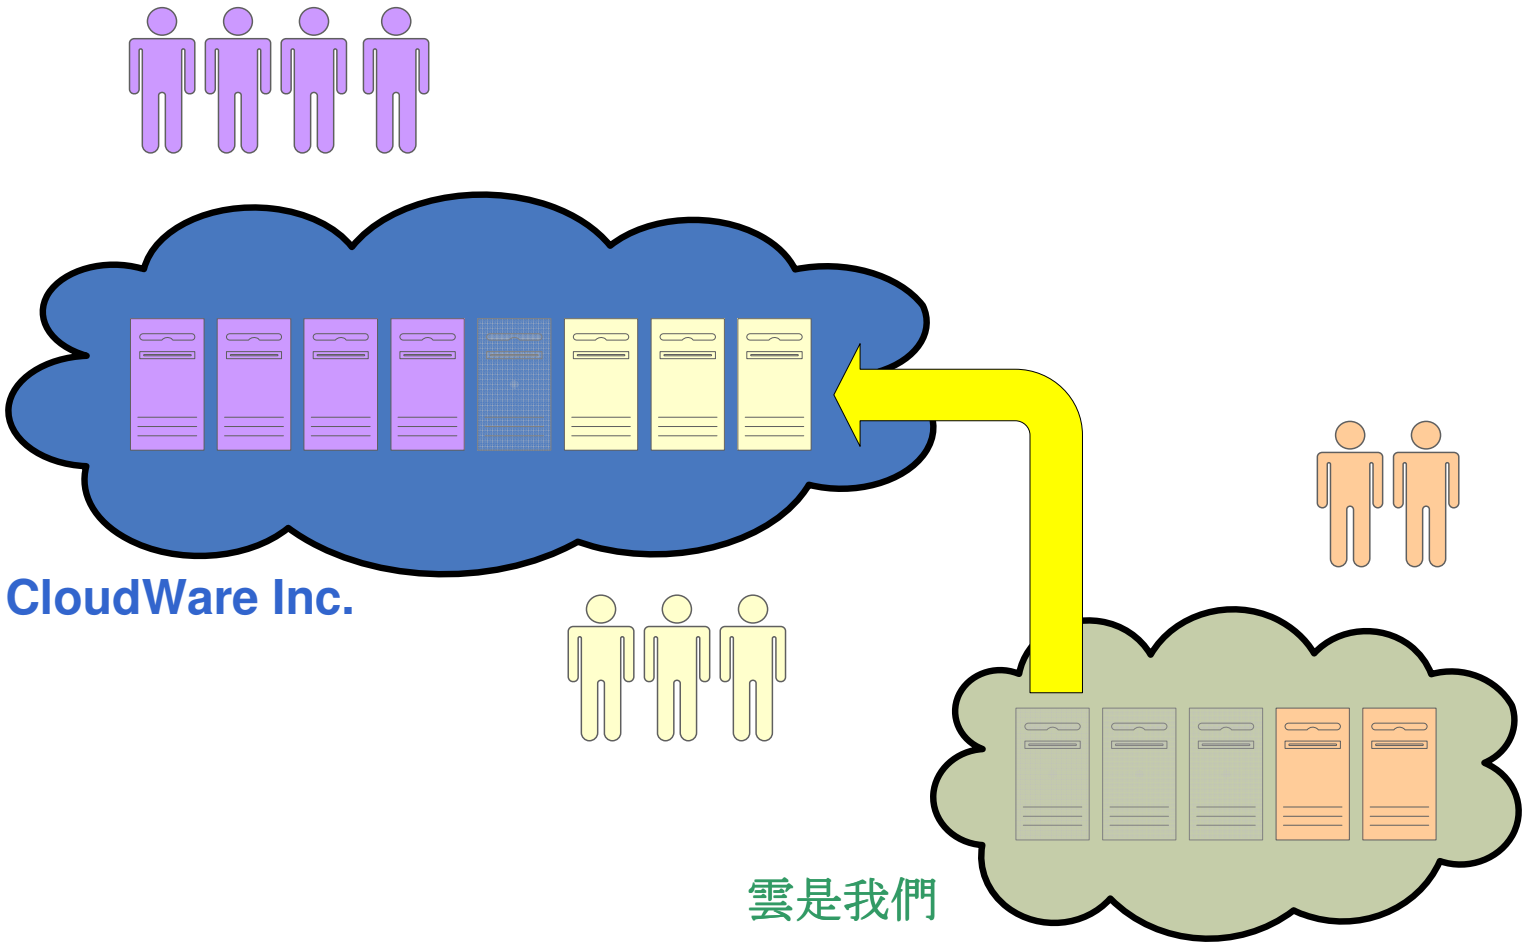

RESERVOIR in a Nutshell

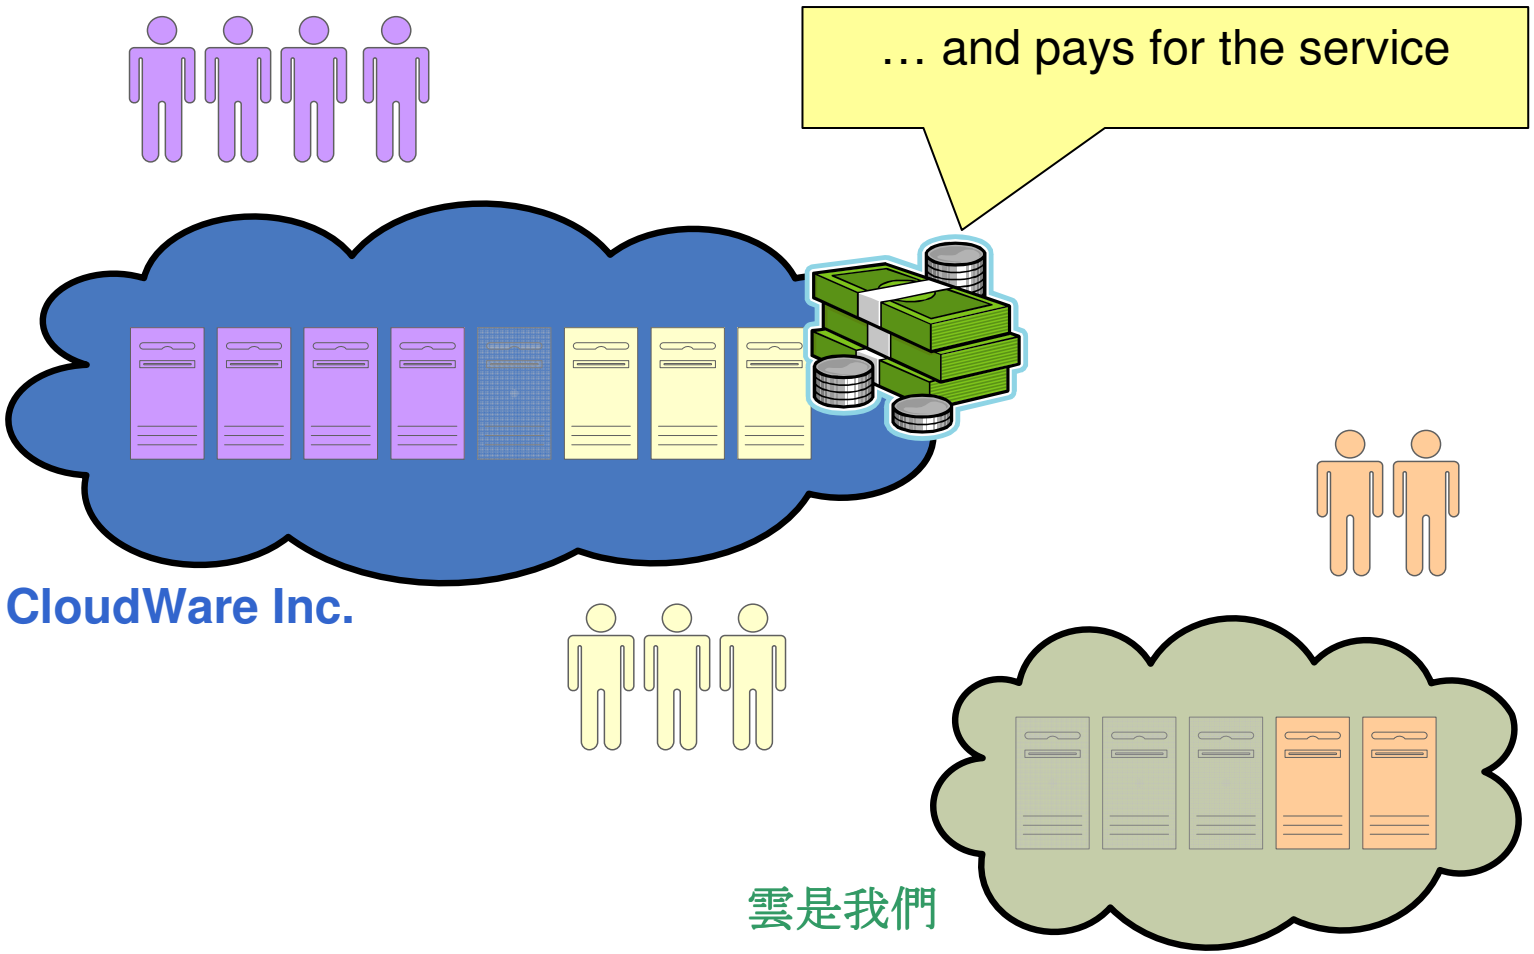

CNN

A Service of CNN, For

IBM and Researc

February 05, 2008:

Expanding its clou
announced that it i
European partners
fluctuating demand
environment.

The 17M Euro EU-f
Resources and Se
explore the deploy
different administrative domains, IT plat
cloud computing project aims to develo
service-based online economy, where
transparently provisioned and managed

vmware

Solutions | Produ

Communities Virtual Appliances Store Support

Home > About Us > N

VMware A
Cloud Co
Rackspac
Business

vCloud Initiative
Application Con

LAS VEGAS, Neva
virtualization solut
partners, to delive
datacenters and c
from partners acro
VMware vCloud In
and outside their f
for any application

YAHOO! FINANCE

Welcome, **brochwerger**
[Sign Out, My Account]

Finance Home -

HOME | INVESTING | NEWS & OPINION | PERSONAL FINANCE | NEW TECH TIC

GET QUOTES » Finance Search

Press Release Source: Citrix Systems, Inc.

Citrix Unveils Cloud Computing Strategy and Product Line

Monday September 15, 1:24 pm ET

New Citrix Cloud Center (C3) Integrates "Cloud Proven" Virtualization and Networking Products to Power Next-Generation Service Providers

SANTA CLARA, Calif.--(BUSINESS WIRE)--Citrix Systems, Inc. (NASDAQ:CTXS - News), the global leader in application delivery infrastructure, today unveiled its strategy for cloud computing and announced the new Citrix Cloud Center™ (C3) product family. The Citrix C3 solution integrates "cloud proven" virtualization and networking products that power many of today's largest Internet and web service providers. This unique combination lets next-generation cloud providers take advantage of the most widely-adopted virtual infrastructure platform for hosted cloud services, as well as the most proven infrastructure to deliver those services reliably and securely to both cloud consumers and enterprise datacenters. The Citrix strategy will focus on equipping both new and existing cloud providers with the infrastructure needed to deliver successful clouds to their customers.

Goals of the Reservoir Project

- **Develop technologies for advanced Cloud Computing**
  - Provide a software architecture where resources and services can be transparently and dynamically managed, provisioned and relocated like utilities – virtually "without borders"
    - Capabilities of service mobility and migration
- **Premise: No single cloud can create a seemingly infinite infrastructure capable of serving massive amounts of users at all times, from all locations**
  - Federation of clouds
  - Leverage the diversity factor to achieve economies of scale
  - Leverage locality
- **Envisioned Impact**
  - Optimize service delivery, relieving service consumers from awareness of IT attributes while providing QoS and security guarantees
  - Increase the competitiveness of the European economy by introducing a basis for a new network-centric, service-oriented infrastructure which will support highly-effective execution of tasks, services and business models across disparate locations and platforms

Approach

- Focus on technologies that enable to build cooperating computing clouds
  - Connect computing clouds to create an even bigger cloud
- The Service Oriented Infrastructure (SOI) equation:
  - Start with grid computing concepts
    - Resource sharing across organizations and geographies
  - Add virtualization technologies
    - Use of virtual machines as the basic unit of work
  - Drive the system by new techniques for business service management

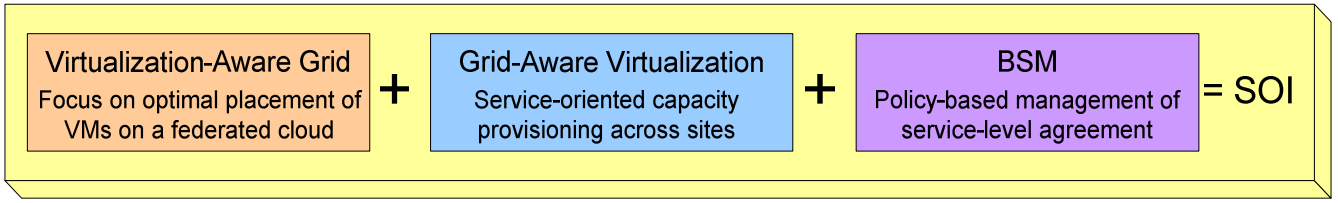

Design Driven by Real-World Scenarios

- Scenario 1: SAP business application (SAP)
  - Business application oriented use cases and the opportunities to execute them on a flexible infrastructure
- Scenario 2: Telco application (TID)
  - Hosting web sites that deals with massive access (e.g., the Olympics games)
  - High degree of personalization and support for mashups
- Scenario 3: Utility computing (Sun)
  - Deploy arbitrary operating system and application stacks on remote resources. Provide secure and seamless access to them. Adjust resource allocation on-demand without the end user noticing disruption of service
- Scenario 4: eGov application (Thales)
  - Automatic adjustment of resources and domains cooperation

The RESERVOIR Architecture

- Monitors SLA compliance
- Enforces SLA compliance by adjusting **application capacity:**
  - Service Components Sizing (VEEs)
  - Service Tiers sizing

Service Provider supplies:
- Pre-packed Service Components (VEEs)
- Service Definition

SLAs are established b/w Service Provider and Service Manager

- Optimizes **VEE placement** subject to constraints
- Deals with site federation

- Translates **generic commands** into specific virtualization platform commands
- Creates and maintains isolated **virtual networks**
- Enables efficient and secure access to **remote storage**
- Performance **optimizations**

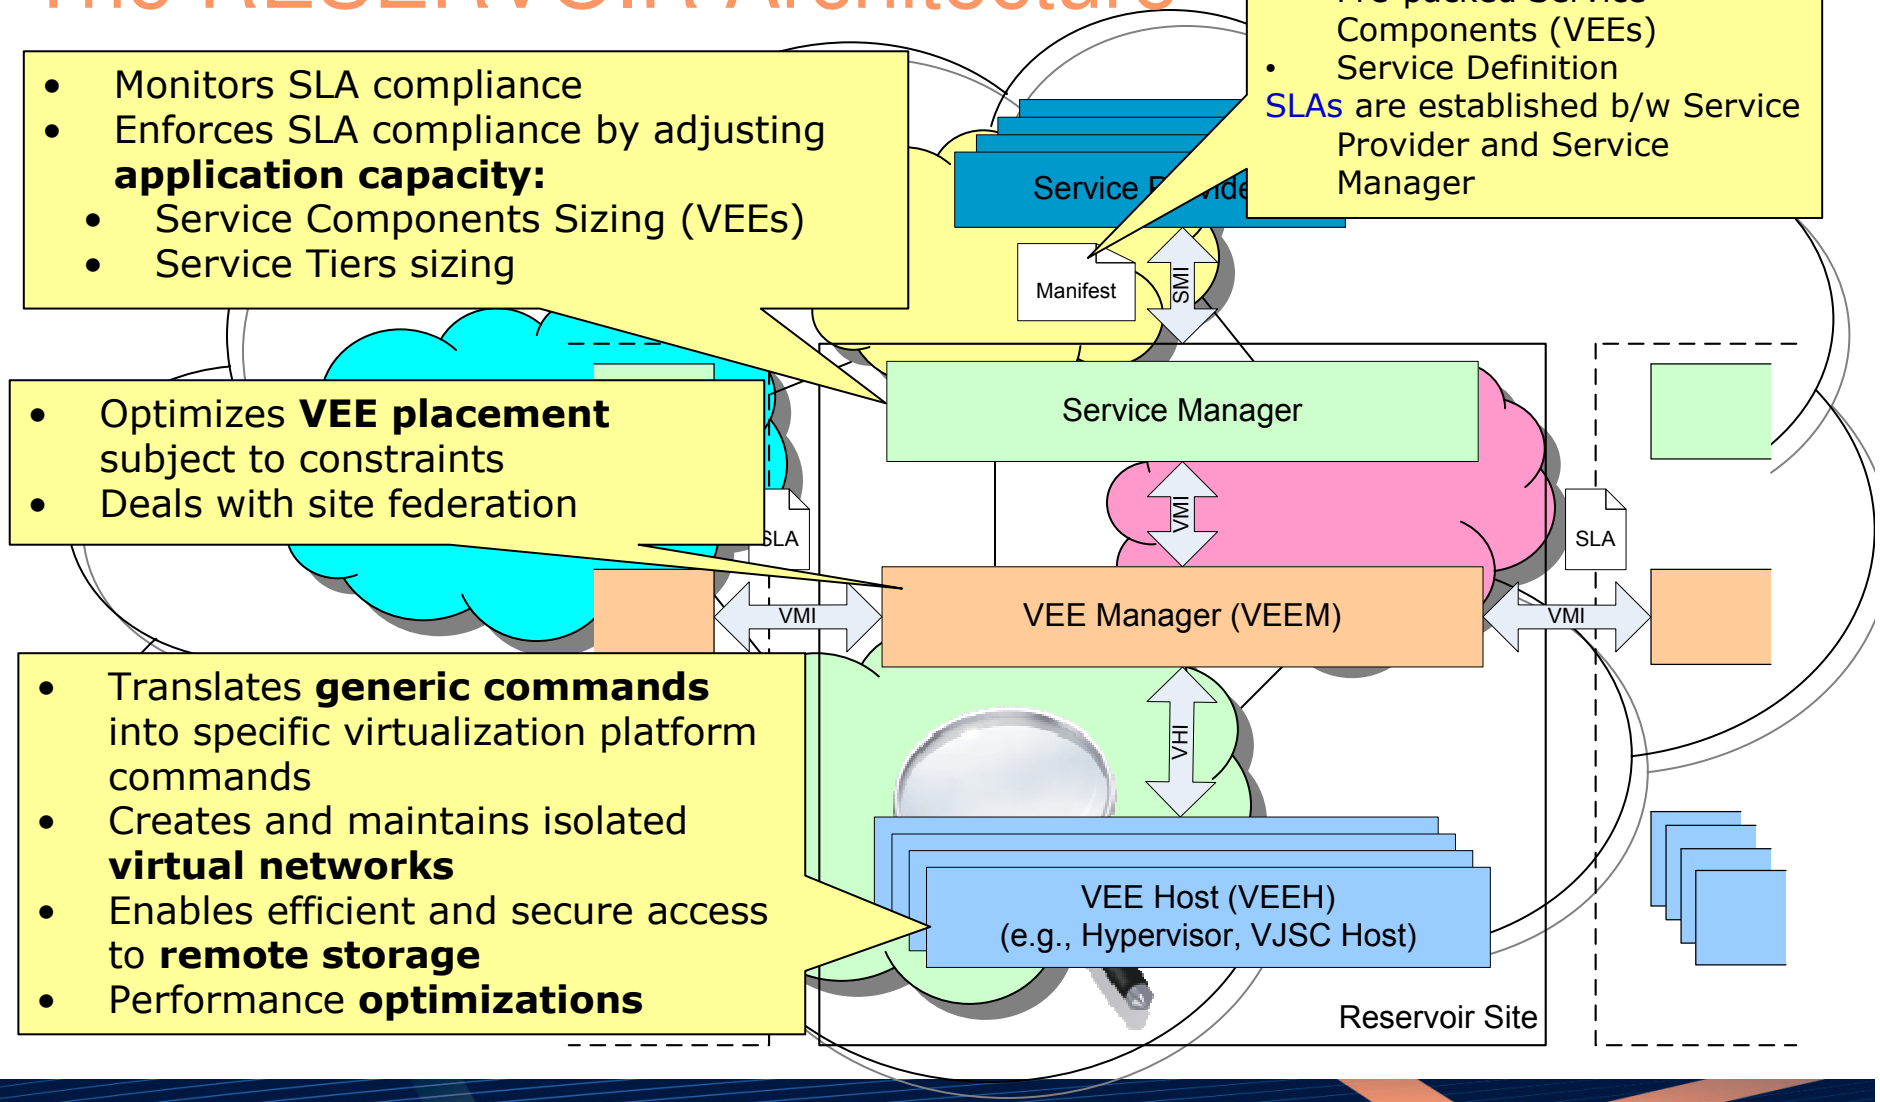

IBM Focus Areas

- Manageability
  - Extend libvirt, e.g., with richer networking model and asynchronous event notification
- Internet-scale network virtualization
  - Distributed Layer 2 switch over multiple physical network topologies including cross-site
  - Compatible networking model across hypervisors (KVM and PowerVM)
- Internet-scale image management and storage virtualization
  - Hierarchical repositories
    - Master image repositories spread across cloud
    - Instance images are created at a repository accessible by hosts
- **Admission control and placement optimization**
  - **Maximize resource overbooking while protecting SLAs by controlling admission of new services through statistical multiplexing**
- Performance
  - PowerVM modifications driven by end-to-end performance of real world workload (SAP)

Admission Control and Placement Optimization

- Over-booking of resources in RESERVOIR while protecting SLAs and minimizing other costs
- Statistical admission control (for SLA protection with thin over-provisioning)
- Smart placement (to optimize operational costs)

Elastic Service (concept review)

- Traditional commercial services are sized *statically*, i.e., with respect to the *maximal* projected demand:
  - Results in either over-provisioning or under-provisioning
- Elastic service is sized *adaptively* to the actual observed workload
- In RESERVOIR the concept of elastic service is realized by means of
  - Varying the number of VEE instances
  - Varying the sizing of the VEE types
- Elasticity of the service is subject to constraints:
  - Minimal and maximal number of VEE instances of specific types
  - Minimal and maximal allowed sizing per VEE type

Main Messages

- Each RESERVOIR site wishes to maximize its own benefit via resource over-booking while protecting SLAs
- With *statistical multiplexing,* the total capacity of the shared (i.e., multiplexed) resource is allowed to be smaller than the maximum total demand while the probability of congestion is controlled
- Equivalent capacity is capacity required to keep the probability of congestion below a pre-defined threshold value
- We separate between two complimentary components:
  - Admission control (to guarantee SLA adherence): performed jointly by SM and VEEM
  - Continuous placement optimization (to minimize operational costs): performed by VEEM

Service Admission Decision (summary)

- When admitting a new service into the system, Service Manager appreciates an impact on the SLAs of the already deployed services:
  - Admit the service if after deploying it to the cloud (using local and remote resources) the probability of congestion does not exceed the *acceptable risk level* defined for this site
  - Reject - otherwise

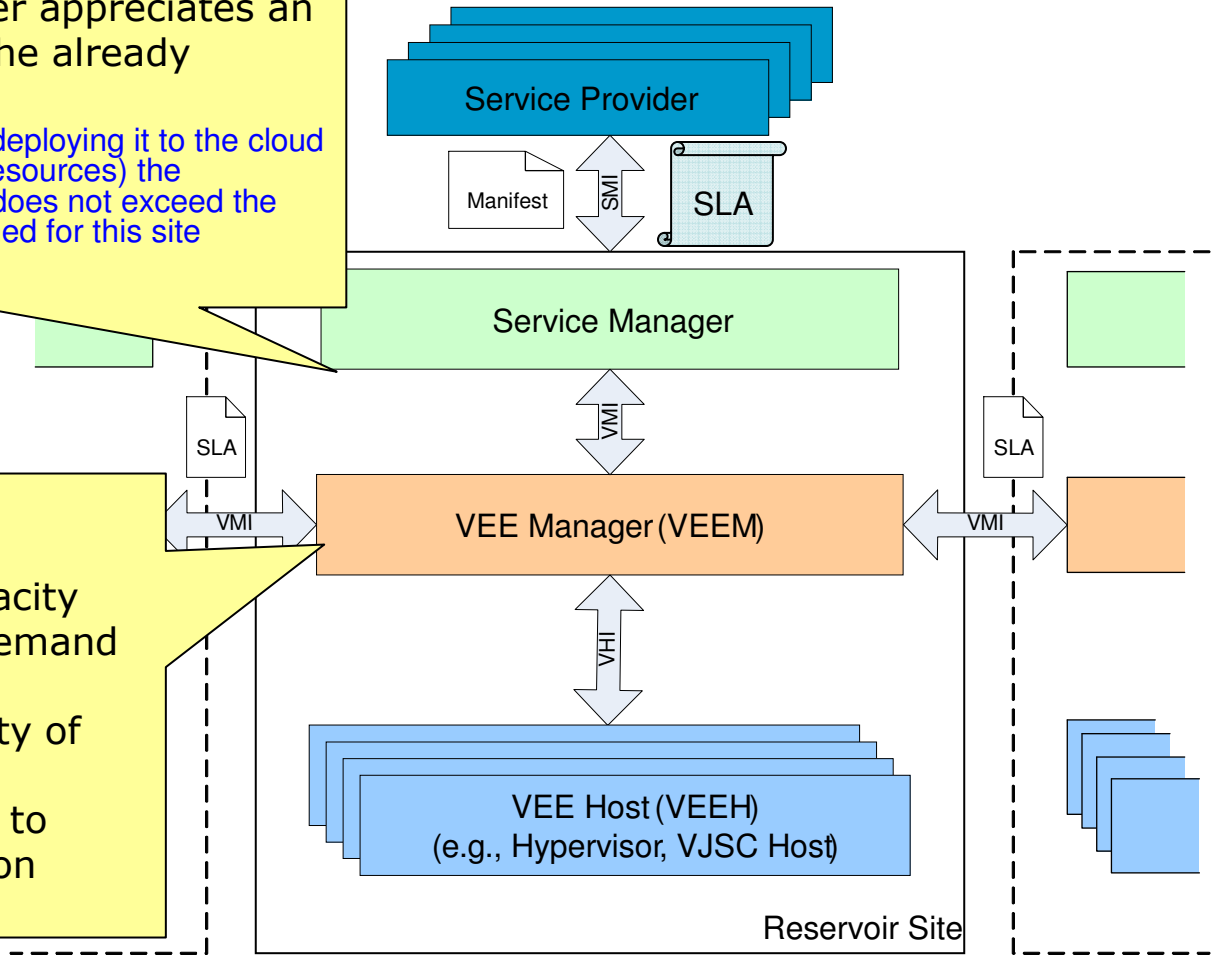

- VEEM knows:
  - Historical demand
  - Current residual capacity
  - Maximal projected demand of the new service
- May find out: availability of remote resources
- Helps Service Manager to make admission decision

General BSM-aligned Optimization Framework

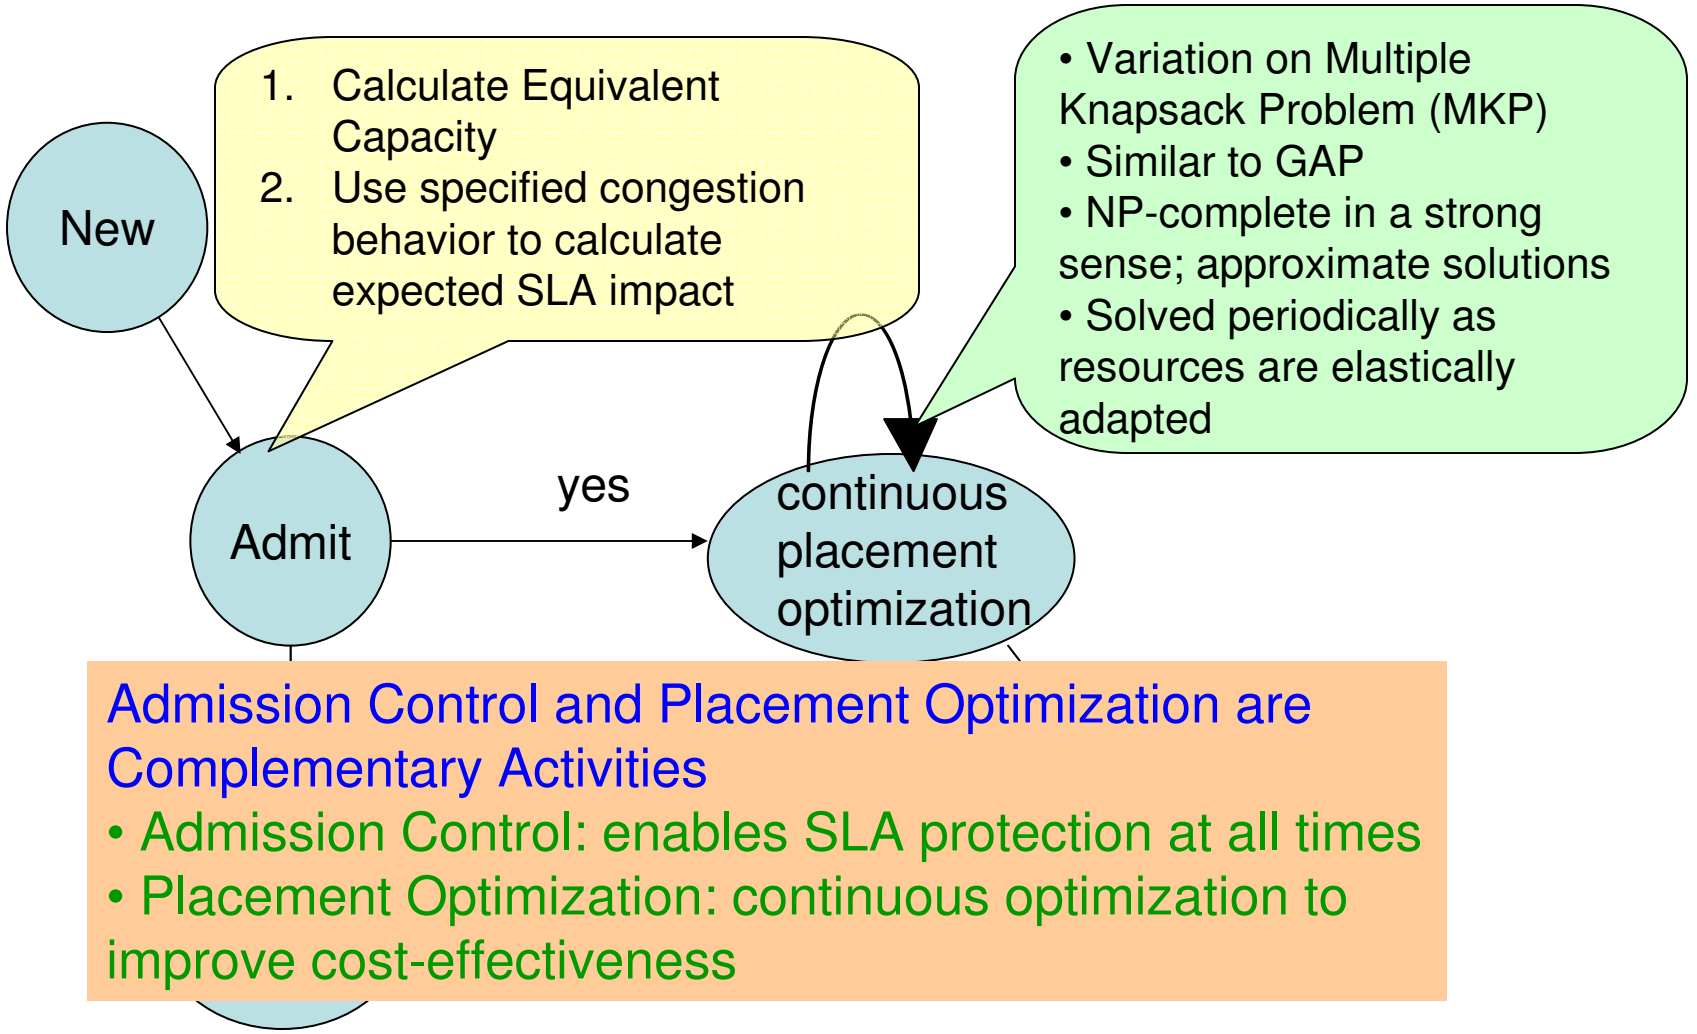

Admission Control and Placement Optimization are Complementary Activities

- Admission Control: enables SLA protection at all times
- Placement Optimization: continuous optimization to improve cost-effectiveness

Summary

- Very ambitious project to create the next generation infrastructure for services
  - Bridge the gap between the services and infrastructure worlds
  - Focus on technologies that enable to build cooperating computing clouds
  - Explore, merge and extend technologies
    - Grid computing concepts (large scale federation)
    - Virtualization
    - Business Service Management
  - Architecture principles:
    - Autonomy of sites
    - Autonomy of management layers within sites
    - Maximizing local value through collaboration while retaining management control
- Status
  - Architectural spec. published (2Q08)
  - Final stages of design of first year prototype
- Additional information:
  - B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Caceres, M. Ben-Yehuda, W. Emmerich, F. Galan, *"The RESERVOIR Model and Architecture for Open Federated Cloud Computing*", IBM Systems Journal, 2009, *to appear*
  - The RESERVOIR Seed Team, "RESERVOIR - An ICT Infrastructure for Reliable and Effective Delivery of Services as Utilities", IBM Technical Report, H-0262, http://domino.watson.ibm.com/library/CyberDig.nsf/papers/A44F6256BB697FCE852574E10052DDEE