RESERVOIR – A Service-Oriented Virtualized Platform

Shimon Agassi, Benny Rochwerger, Eliot Salant
{shimon_agassi, rochwer, salant} @il.ibm.com

RESERVOIR Project, IBM Haifa Research Lab

EGEE'08 Conference, Istanbul, *September 22, 2008*

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 215605

What is Cloud Computing?

- Cloud Computing is an approach to a shared **IT infrastructure** in which large pools of computer systems are linked together to provide **IT services.**
- Cloud Computing offers simplified platform for use when needed, lowering costs and energy use.
- Since it accesses "virtual" resources, Cloud Computing is not limited by the power and capabilities of local or remote computers.
- It is the next generation of enterprise data centers, which operate like the Internet, providing extreme scale and fast access to networked users.
- Unlike grid computing, which distributes IT for a specific task, Cloud Computing is used across an entire range of activities.

Source: Willy Chiu, VP IBM Software Group

Cloud Computing is evolving

- **From:**

  Models which provide infrastructure-on-demand and application development,

- **To:**

  A more value-driven model in which the cloud offers scalable, easily accessed services, with an outstanding user experience.

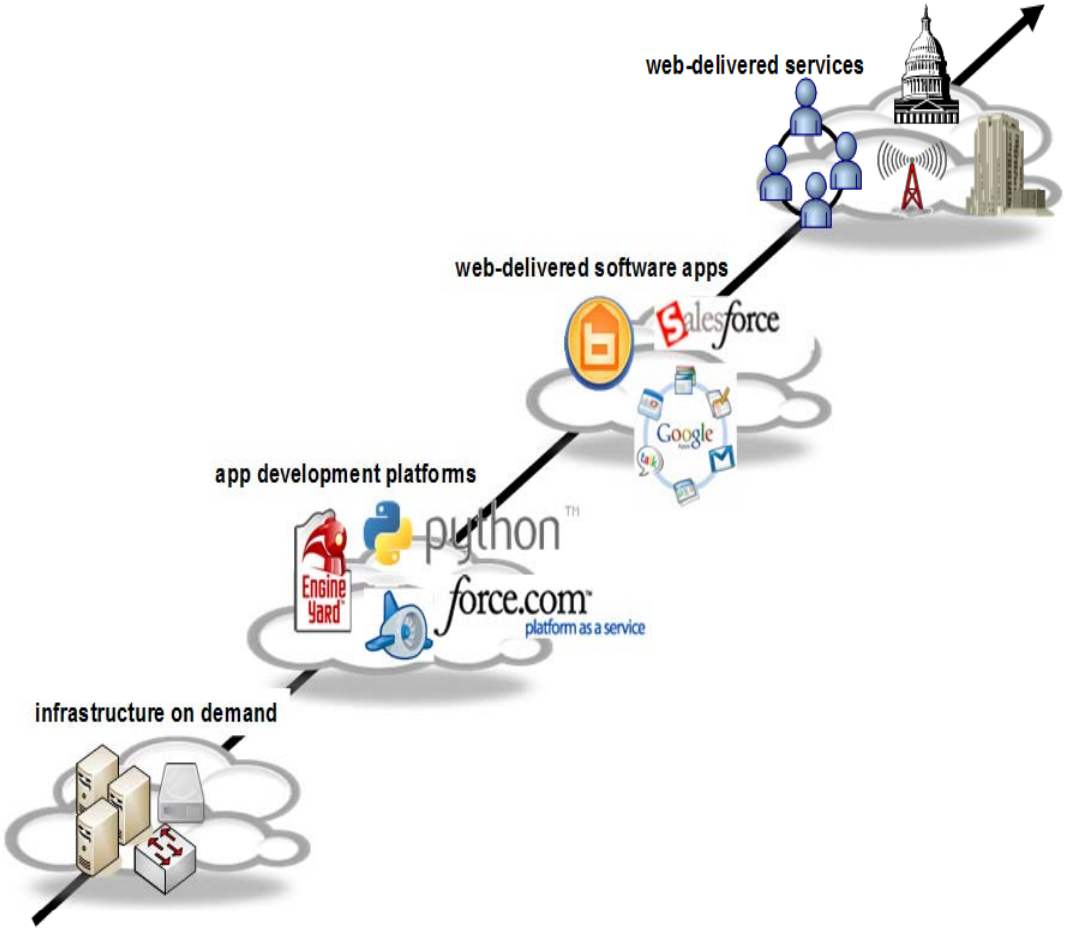

Cloud Players Today

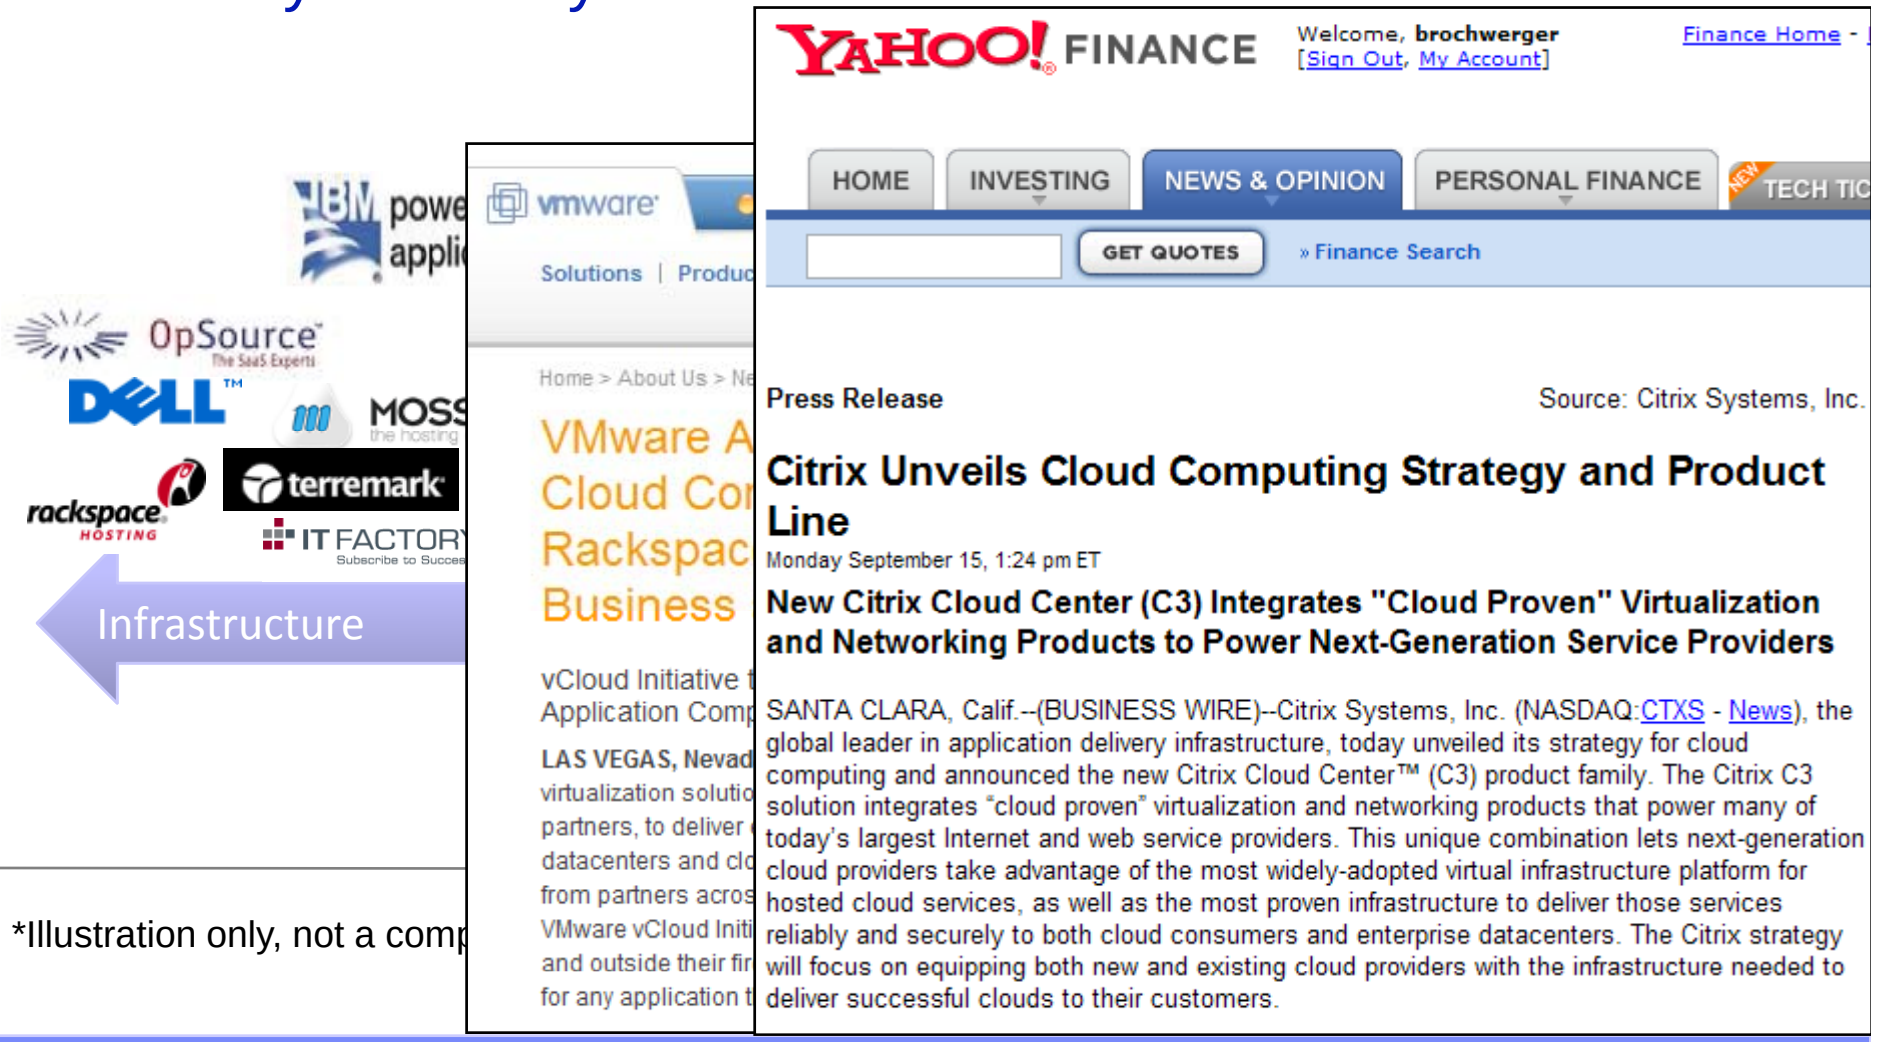

*Illustration only, not a comp

IBM and Cloud Co

Driven by real custome

February 20

2007

IBM
Google
Academic Initiative

IBM
Blue Cloud

Vietnam Innovation Portal

WUXI · CHINA
Wuxi China Clou
Computing Cent

IBM
Joint research initi
with 13 Europe
partners

IBM
Googl
Partner to enha
academic resea
opportunities

On GameS

c|net NEWS.com

Search:

Today on CNET | Reviews | News | Downloads | Tips & Tricks | CNET TV | Compare Prices

Business Tech | Cutting Edge | Green Tech | Wireless | Security | Media | Markets | Personal Te

News Blog

Recent posts on technology, trends, and more

November 15, 2007 6:50 AM PST

IBM floats Blue Cloud computing plan

Posted by Martin LaMonica 2 comments

IBM on Thursday announced an offering called Blue Cloud meant to make it easier to run large-scale applications with massive databases over the Internet.

The computing giant said it has dedicated 200 researchers to the project and will have services available as of the first quarter next year.

The first product will be an IBM BladeCenter with Linux-based servers equipped with a suite of "cloud" software.

IBM's Blue Cloud includes a Blade Center rack of Linux servers and data center software for handling data-intensive applications.
(Credit: IBM)

That includes grid-computing software, virtualization tools Xen and PowerVM, and the open-source Hadoop parallel workload-scheduling software originally developed by Yahoo. The suite will also include IBM's Tivoli data center management software for automatically provisioning hardware to adjust for fluctuating computing demand.

The offering came out of internal work IBM did in creating its own clouds--essentially a server installation that is capable of handling very large sets of data and can operate efficiently by drawing on virtualization and other data center management software.

Cloud Computing Benefits

- Innovative business models require a utility-like intelligent infrastructure that embraces complexity to be successful in the competitive, fast-paced, services-based global economy.
  - ECONOMICS: Small up front investment and can be billed by consumption. Reduction of TCO allows clients to pursue operational efficiency and productivity.
  - RISK MANAGEMENT: Small up front commitment allows clients to try many new services faster and choose. This reduces big failure risks and allows clients to be innovative.
  - TIME TO MARKET: Adopt new services quickly for pilot usages and scale quickly to global scale.
  - INFORMATION SOCIETY: Value-added information generated by collection and analysis of massive amounts of unstructured data.
  - UBIQUITOUS SOCIETY: Accessible via a heterogeneous set of devices (PC, phone, telematics..)

Cloud Economics examples

- OnDemand apps like Salesforce.com can be provisioned for as little as $300-500 per subscriber after fully costing hardware, software and service vs. as much as $8,000-10,000/user for OnPremise clientserver apps.
- Merrill Lynch estimates that today's $2 billion market in on-demand applications will expand to a $165 billion market opportunity.

	Company Employees			
	10	100	1,000	10,000
Microsoft Implementation				
Hardware	$1,000	$2,000	$20,000	$200,000
Exchange licenses (50 users/license)	$400	$800	$8,000	$80,000
Data center space	$2,500	$5,000	$50,000	$500,000
Office licenses	$500	$4,500	$40,000	$300,000
IT staff	$50,000	$200,000	$1,000,000	$2,500,000
Total costs	$54,400	$212,300	$1,118,000	$3,580,000
Annual cost per employee	**$5,440**	**$2,123**	**$1,118**	**$358**
Google Implementation				
Premiere edition license	$500	$4,500	$40,000	$350,000
Annual cost per employee	**$50**	**$45**	**$40**	**$35**
Cost differential - Gmail vs. Exchange	**109x**	**47x**	**28x**	**10x**

Source: Merrill Lynch estimates

What are the requirements on cloud infrastructure ?

- Demand puts requirements on scalability, reliability…
  - One billion people worldwide have access to the Web
  - MySpace signs up an average of 300,000 new users **every day** with 65 billion page views per month.
  - 50 million blogs are created at the rate of 2 per second, and growing.
  - Activision/Blizzard – “World of Warcraft” multiplayer online role-playing game has over 10 million monthly-paying subscribers and growing.
  - And what will happen as millions (billions?) of inexpensive sensors (RFID “smart dust”) start connecting to the Web?
- Web 2.0 best practice principles will also drive infrastructure requirements:
  - High availability, systems monitoring and management…
- Growing diversity of services will require
  - Privacy and Security compliance
  - Ease of use and operation
  - Pricing models simplicity, SLA driven

RESERVOIR Project Profile

- **A 3-year**
  project,
  exploratc
- **Initiated**
  from acrc
  standard
- **Designa**
  - NES
    Soft
    focu
    the
- **Public w**
  - http

CNN Money.com
A Service of CNN, Fortune & Money

IBM and European Union Launch Joint Research Initiative for Cloud Computing

February 05, 2008: 12:00 AM EST

Expanding its cloud computing initiative, IBM (NYSE: IBM) today announced that it is leading a joint research initiative of 13 European partners to develop technologies that help automate the fluctuating demand for IT resources in a cloud computing environment.

The 17M Euro EU-funded initiative, called RESERVOIR -- Resources and Services Virtualization without Barriers -- will explore the deployment and management of IT services across different administrative domains, IT platforms and geographies. This cloud computing project aims to develop technologies to support a service-based online economy, where resources and services are transparently provisioned and managed.

SEVENTH FRAMEWORK PROGRAMME

The RESERVOIR Vision - *the technical goals*

- Minimize over-provisioning of resources
  - Dynamic allocation and re-provisioning can get better utilization out of existing resources
  - Break down platform and geography barriers
- Adhere to SLA constraints through intelligent placement and relocation algorithms
  - Address cross administrative domain SLA
    - Domains may be in different organizations
- Create standards to allow for federation of clouds
  - Based on the premise that no single cloud will be able to supply the demand all the time from everywhere,

RESERVOIR - Technical 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

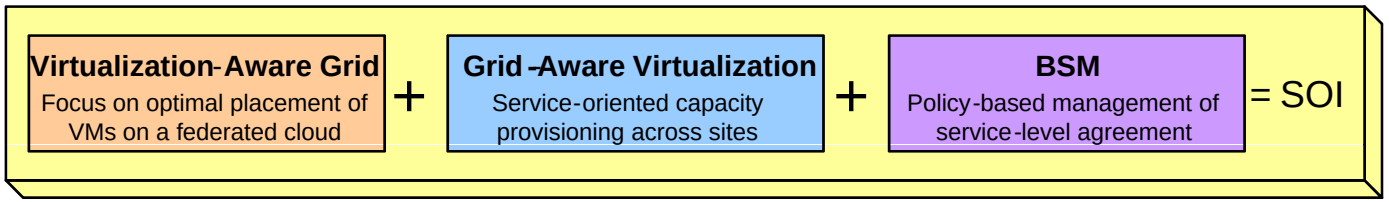

A Service Scenario – Service Definition

Web site service

1. A **Service Provider** uses client tools to generate the **service definition** required from an **Infrastructure Provider**.

- Tier definition (web servers, application servers, databases)
- Required Virtual Execution Environments (VEEs)
- Software
- Images
- Storage
- Network
- Required configuration
- Inter-tier relations
- Required SLOs

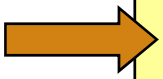

```
<service NEW SERVICE … >
  <tier web-servers … >
    <VEE-requirement … >
    <image … >
    <software … >
    <storage …>
    <network … >
    <configuration …  >
    <tier-QoS … >
  </tier>
  <tier app-servers … >
    …
  </tier>
  <tier DB-servers … >
    …
  </tier>

  <inter-tier-configuration … >
  <service-QoS … >
  …
</service>
```

Service Deployment

2. The Service Provider negotiates and ships the service definition to a **P**rimary **R**ESERVOIR **S**ite (PRS)

3. The PRS automatically deploys the complex service on its own site:

- Configure required storage & network, creates VEEs selecting proper physical resources to meet QoS
- Install required images, software according the service definition
- Apply the required configuration
- Setup the monitoring and billing

RESERVOIR Contribution:

Service definition language enabling automatic deployment of complex services over virtual infrastructure

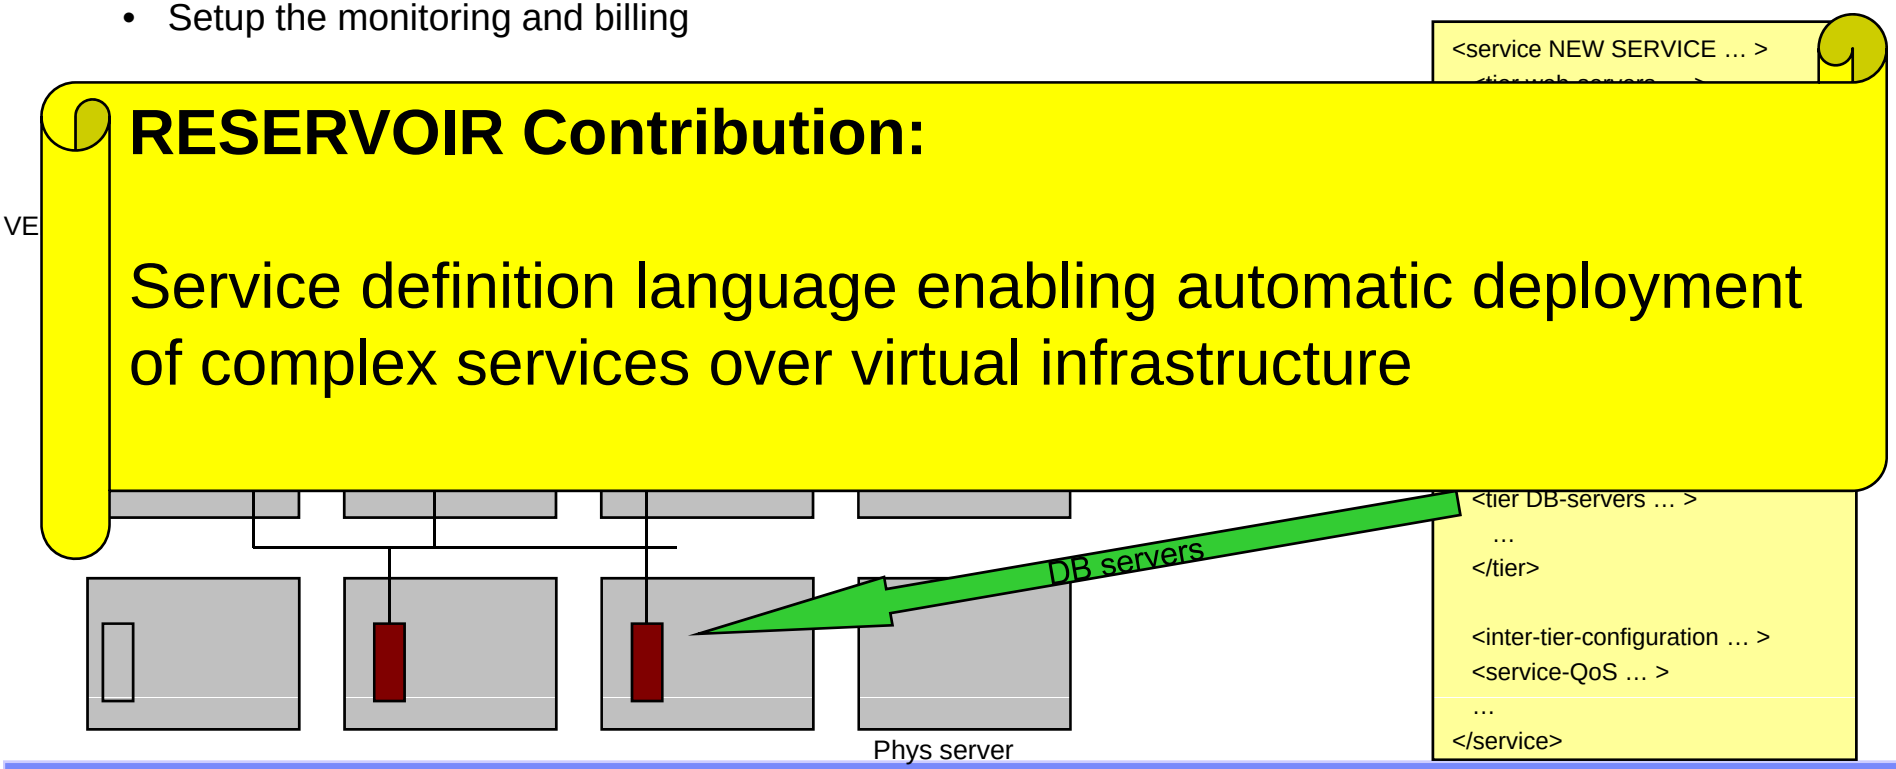

Service Cooperation

4. For HA and assuring the SLA, the PRS negotiates with two other **R**ESERVOIR **S**ites (RSs), and ships the service definition to them

5. Each RS deploys the service (according to the contracted resources) in its site similar to what the PRS did

6. The PRS and RS1, RS2 maintain a service cooperation

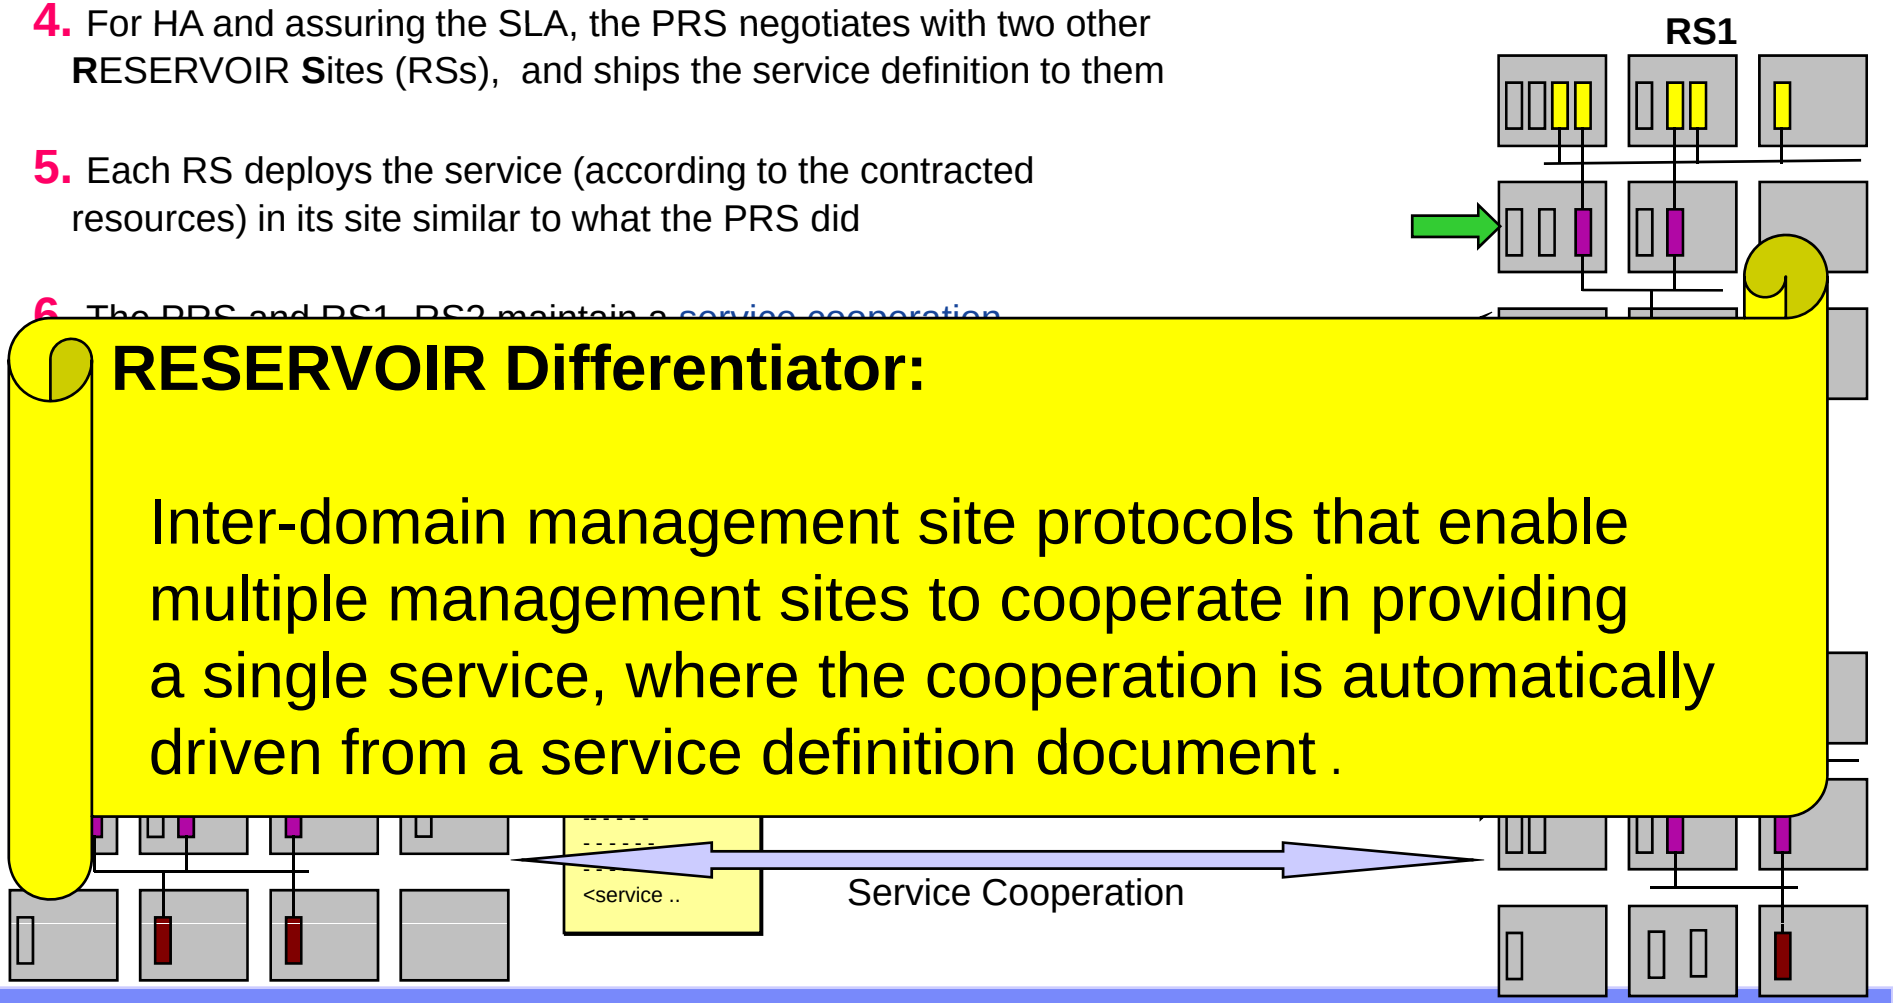

RESERVOIR Differentiator:

Inter-domain management site protocols that enable multiple management sites to cooperate in providing a single service, where the cooperation is automatically driven from a service definition document .

Service Optimization with Live VM Migration

7. Hit rate goes down in the PRS site, and goes up in RS1.

8. PRS and RS1 coordinate VEE migration, employing the RS-RS protocol

9. PRS evacuates VEE, migrating them to RS1, and powers off its free servers
   - Live migration to maintain application servers' states and client connections

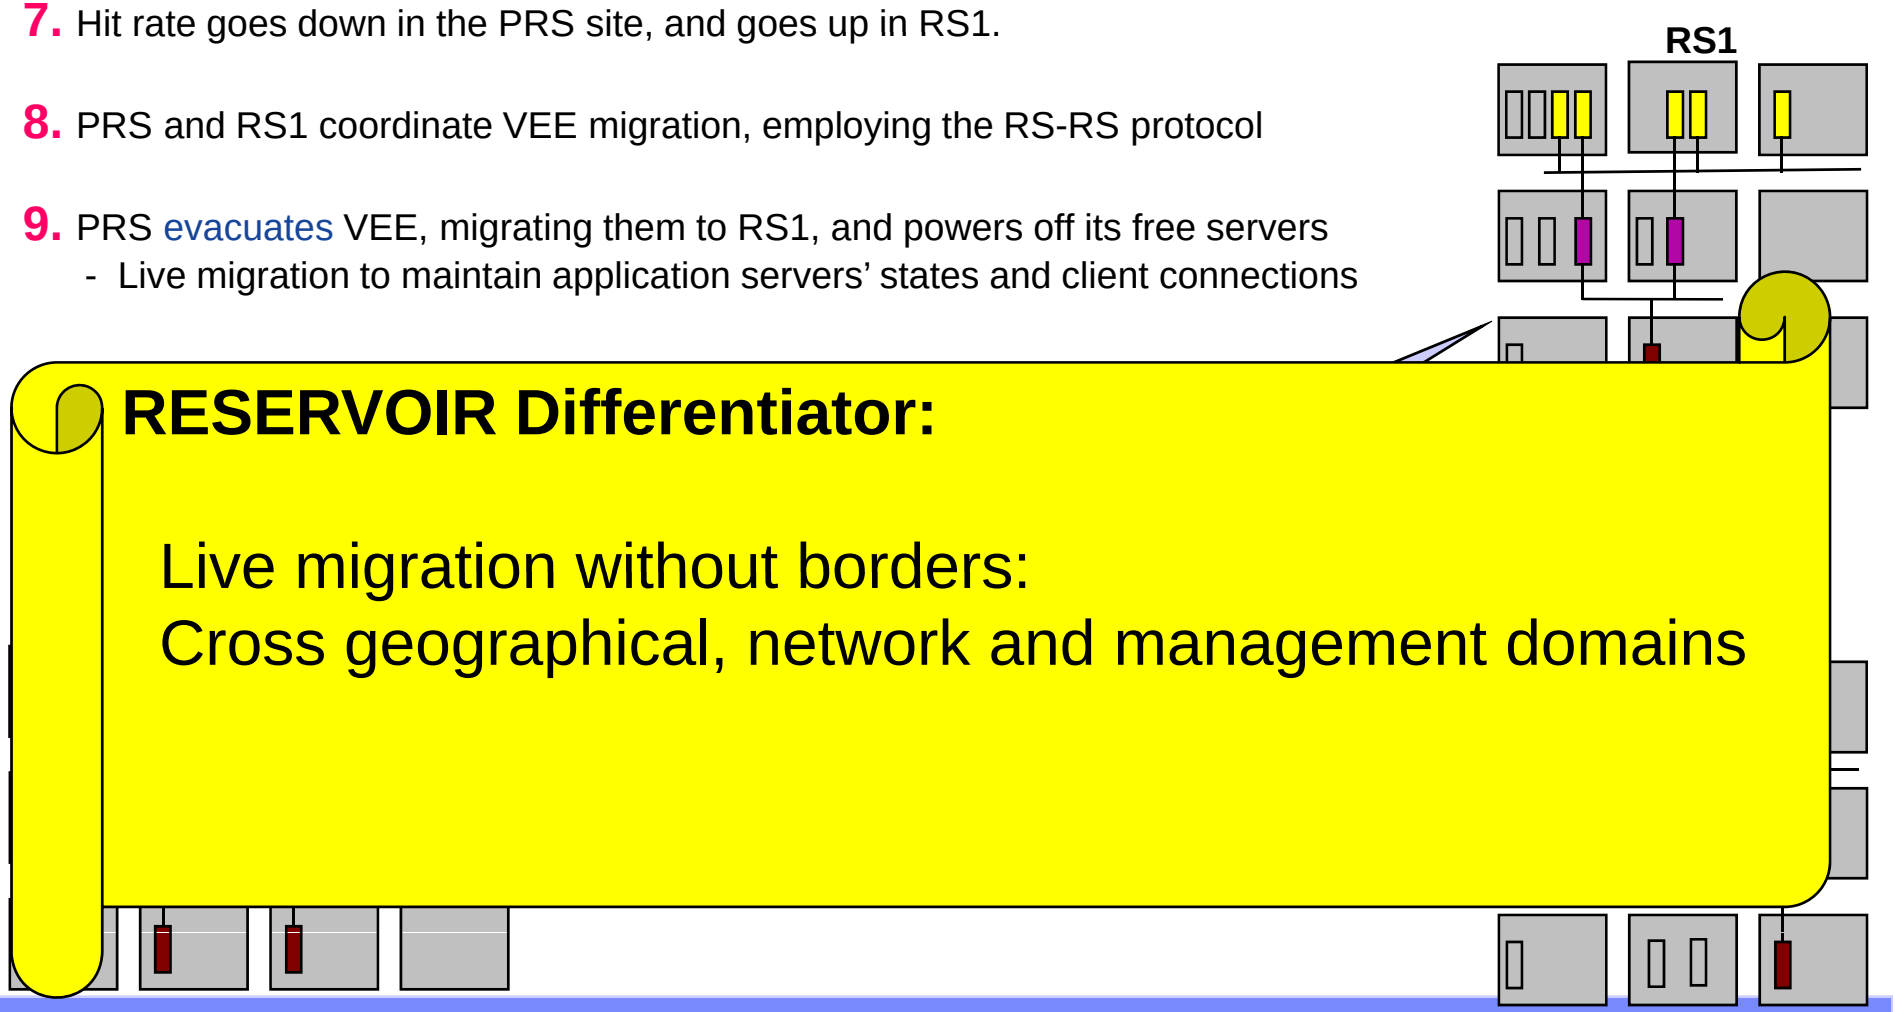

RESERVOIR Differentiator:

Live migration without borders:
Cross geographical, network and management domains

On Demand Service Expansion

10. Load increases and PRS realizes that the available resources at the 3 sites are not enough

11. PRS negotiates with additional RS3, and ships it the service definition

RESERVOIR Differentiator:

The ability to dynamically hire additional 'service power' from a new management site, fully automated, using the service definition language and the inter-domain site protocols

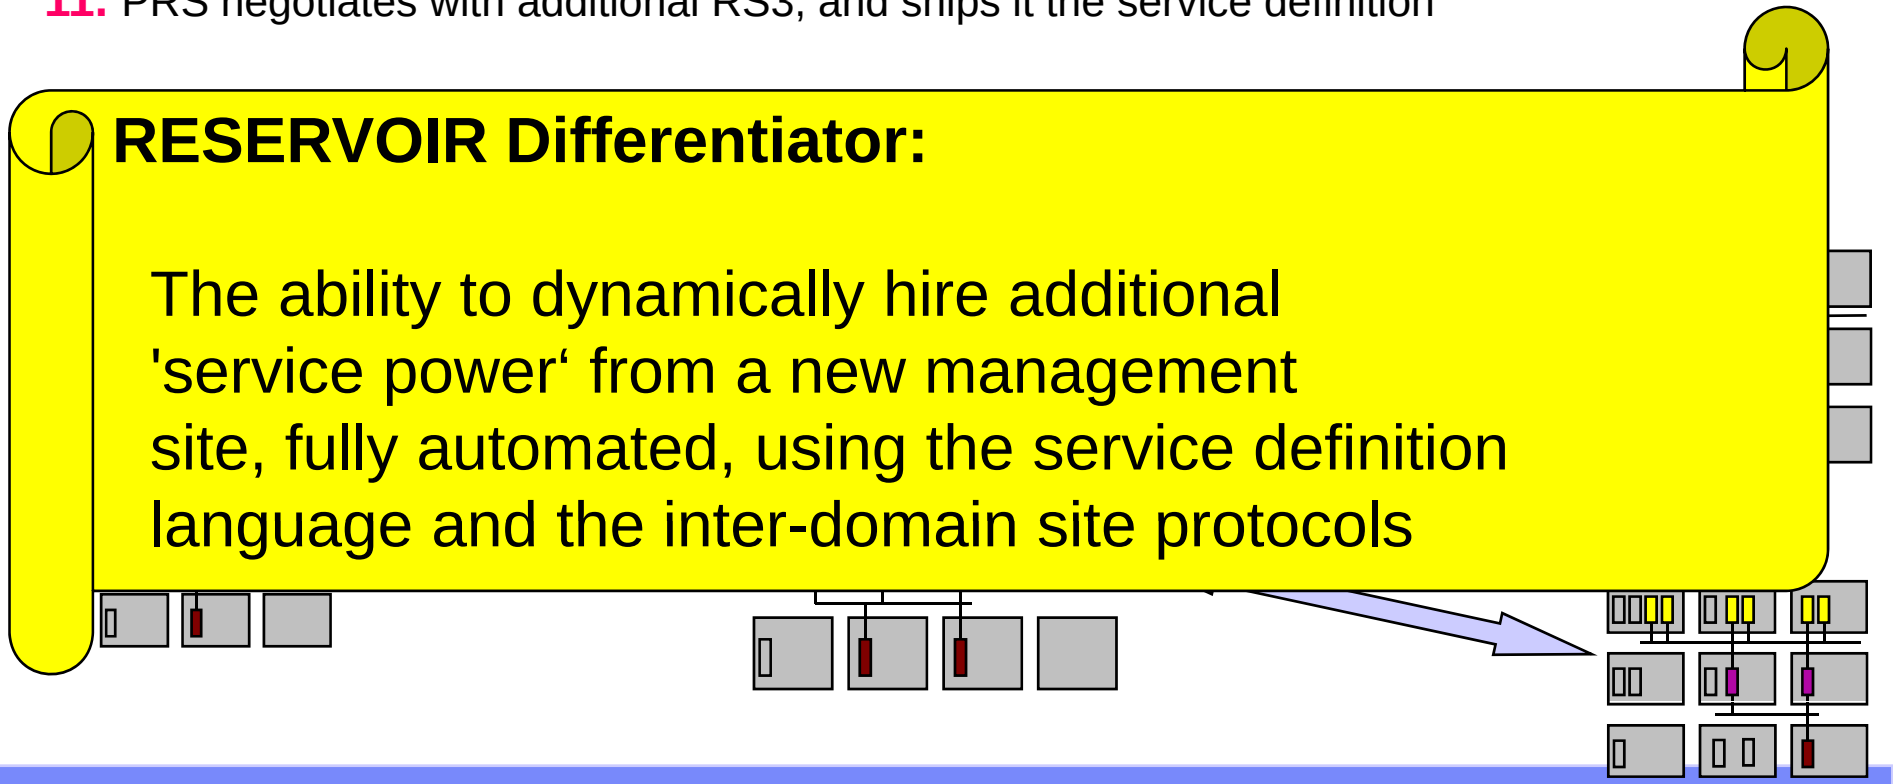

Summary

- Cloud Computing is evolving from Internet-wide IT Resource Sharing into Internet-wide diverse Service Delivery.
- Internet-wide diverse Service Delivery requires Scalable, Highly Available, Simple to use and operate, Secured, Affordable and Manageable infrastructure.
- RESERVOIR is a very aggressive 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
- RESERVOIR Status
  - Architectural spec. published (2Q08)
  - Final stages of first year prototype design
  - Initial testbed in place
    - Sites in UNIME (Italy), UCM (Spain) and IBM (Israel)
  - On track to first official demo (1Q09)

There are no stupid questions

or stupid answers

Backup Slides

Trends Driving Cloud Computing

Infrastructure Technologies: Virtualization, Automation, SLAs

Application Technologies: Grid, MapReduce, Hadoop, SOA, Web 2.0

Data Intensive Applications: From massively parallel (e.g. Google) to large data files (e.g. You Tube)

Computing & Network Appliances: Special servers designed to handle specific tasks are blurring the lines between Network and Data Center

Open IT: Open Technologies, APIs, protocols, data formats, software platforms / data (e.g. Creative Commons, Open Data License)

Business Agility: Enter new markets, Deploy new application services. Stay ahead of compeition.

Networking: Growth in Internet bandwidth enabling ubiquitous connectivity. Increased reliability and functionality embedded in the network.

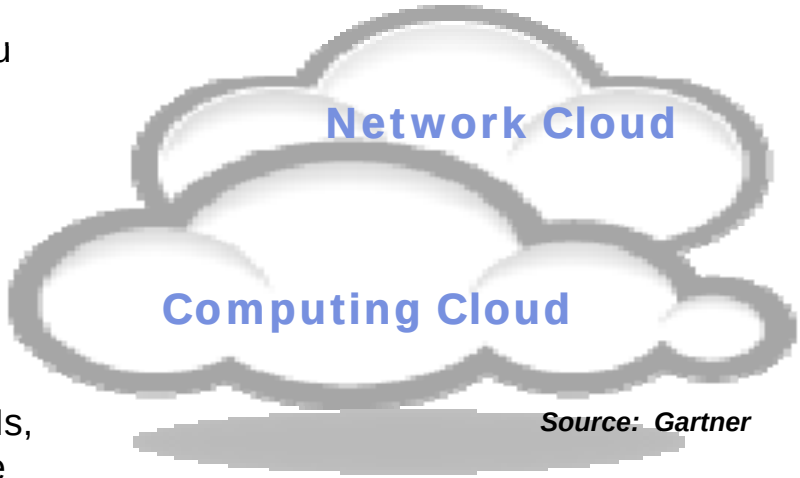

Source: Gartner

Utility Computing: Get as much computing power as you need when you need it, pay for only what you use.

Industrialization of IT: Standardization, and commoditization (e.g email). Falling costs of storage.

Mobility: Explosion of form factors, cell phones/connected devices, Proliferation of sensors

New Business Models: Advertising, Services, Subscription

Web Applications and Platforms: Mashable applications and services built on Web Oriented Architecture (e.g. REST, RSS/ATOM)

Data Center Pressures: Growing costs of power and space, server sprawl

The Reservoir Architecture

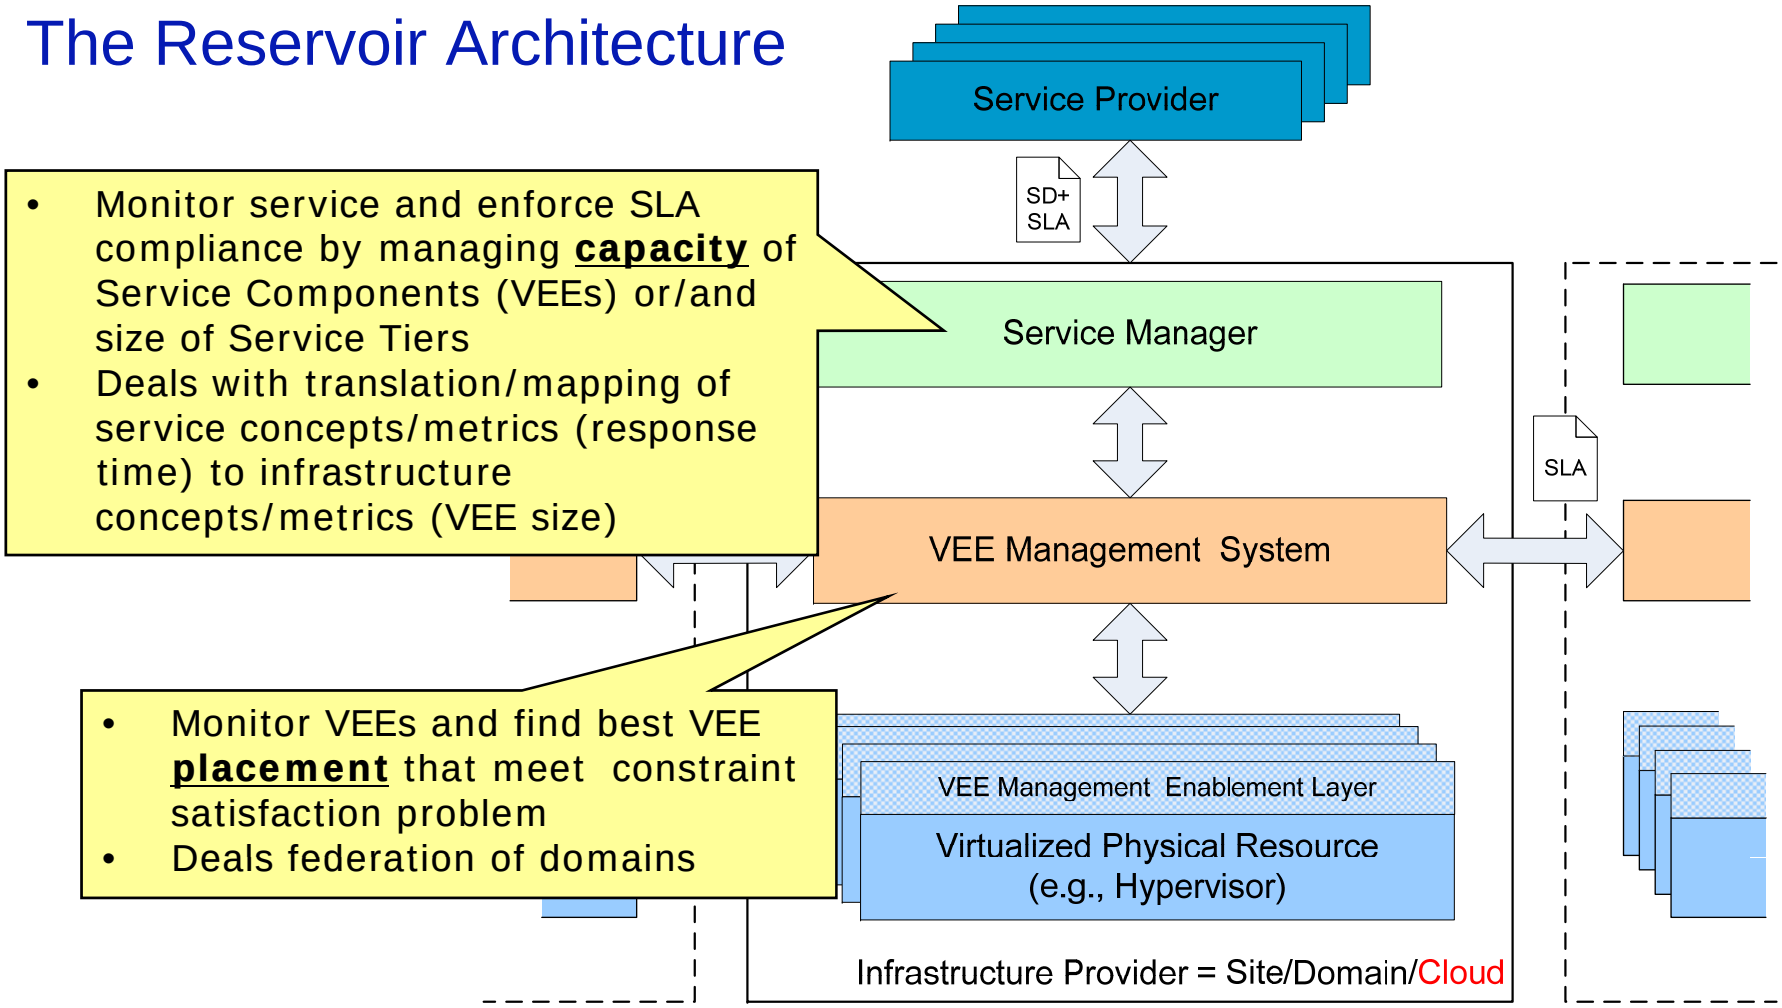

- Monitor service and enforce SLA compliance by managing **capacity** of Service Components (VEEs) or/and size of Service Tiers
- Deals with translation/mapping of service concepts/metrics (response time) to infrastructure concepts/metrics (VEE size)

- Monitor VEEs and find best VEE **placement** that meet constraint satisfaction problem
- Deals federation of domains

RESERVOIR's Service Oriented Infrastructure (SOI): Grid Computing

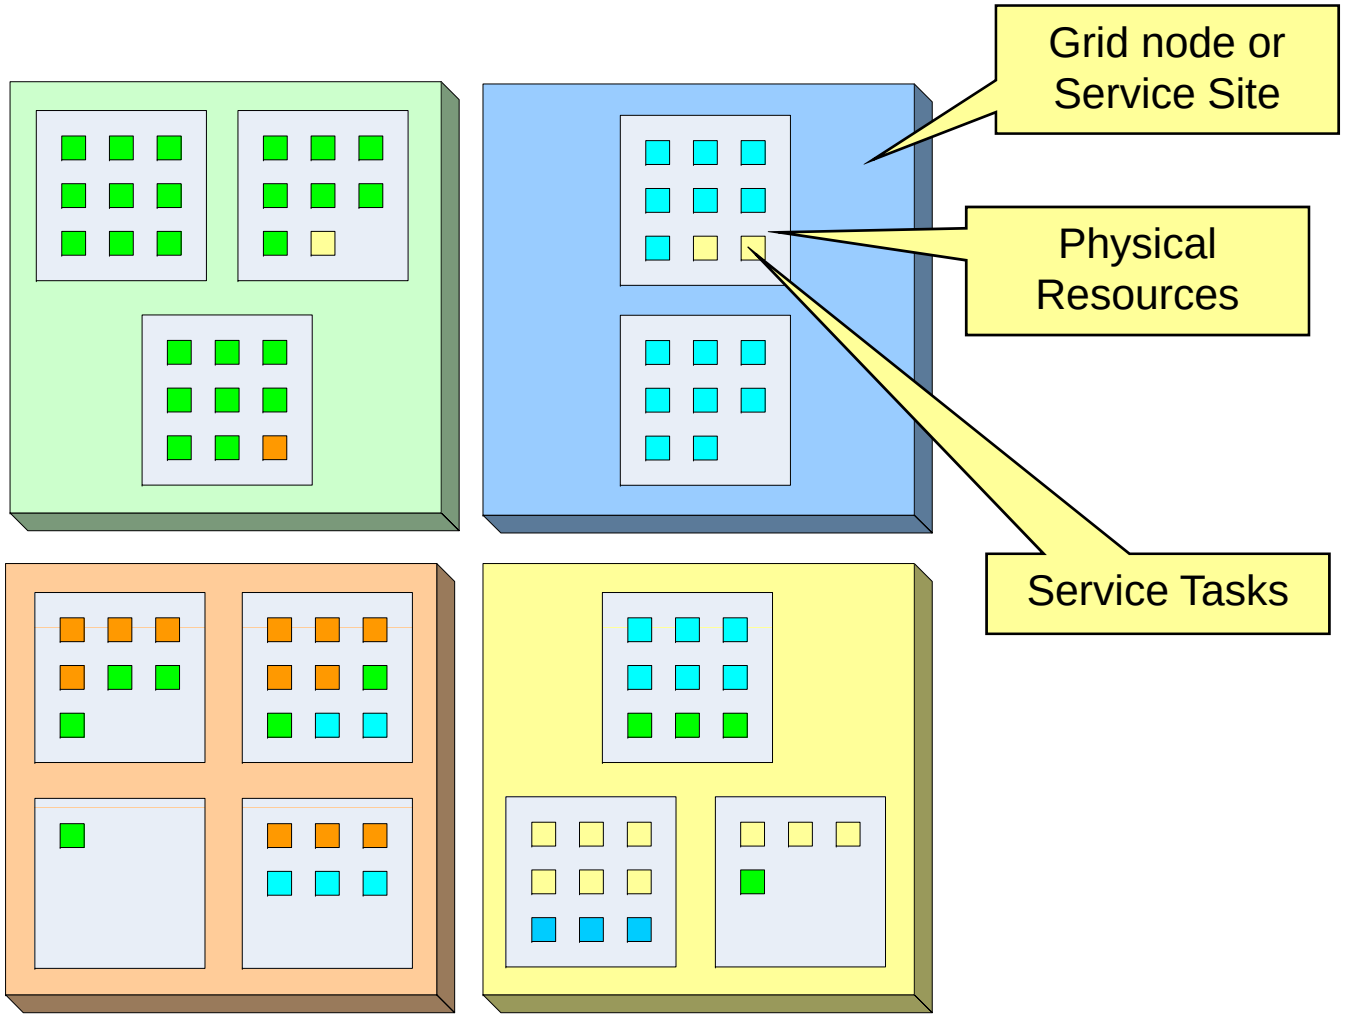

RESERVOIR's SOI: Grid Computing + Virtualization

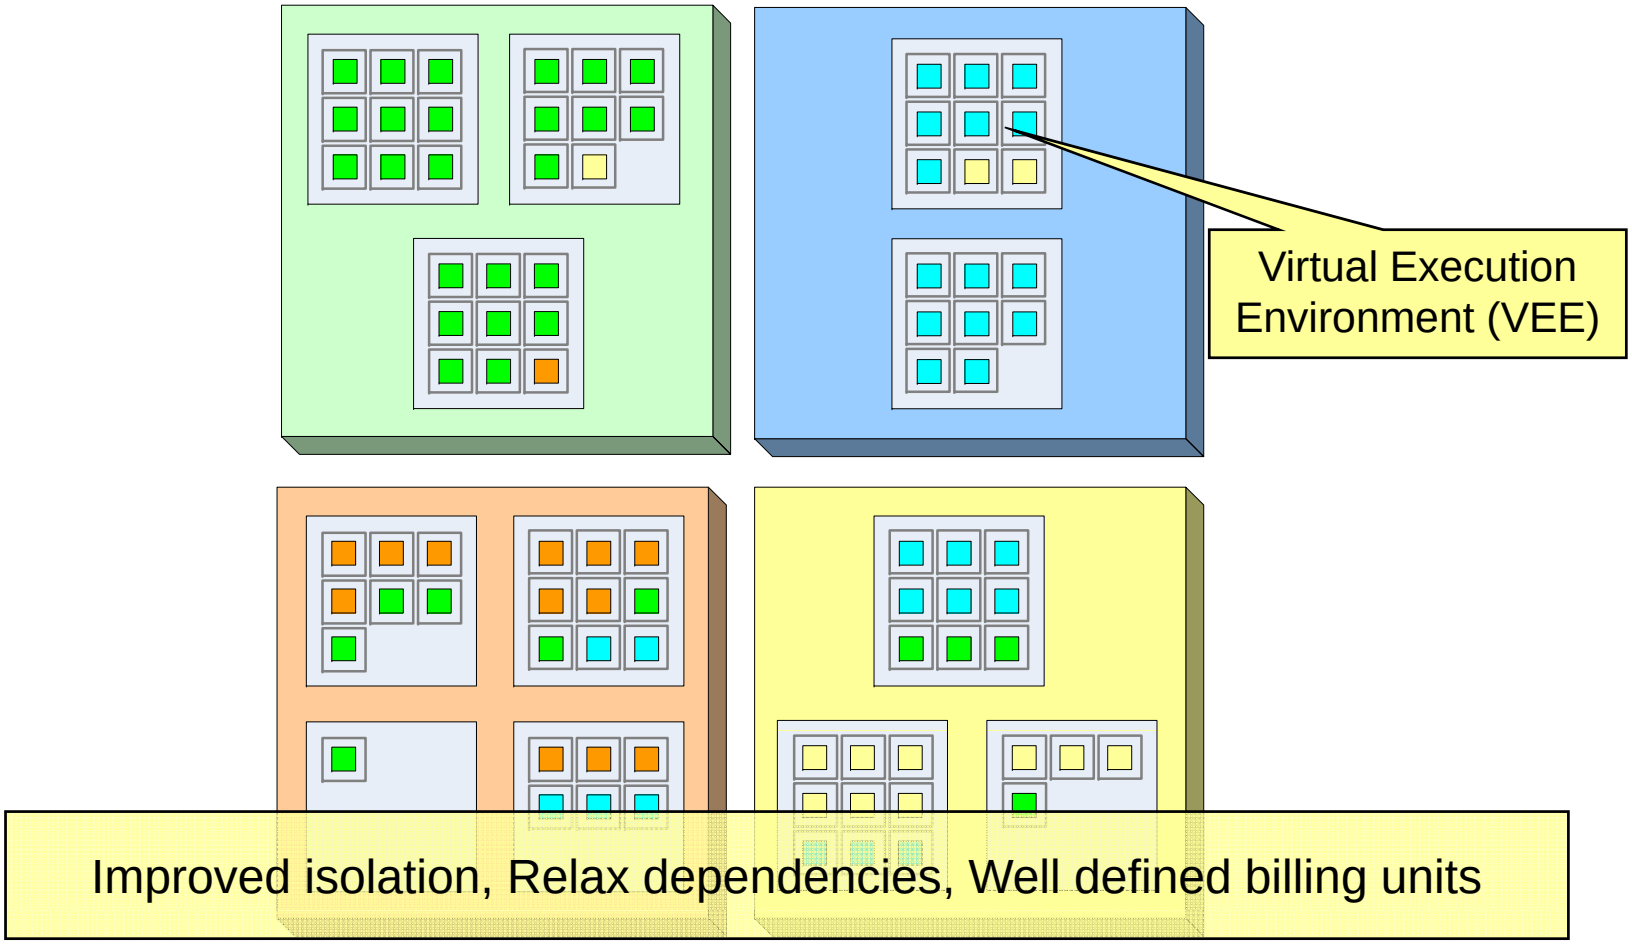

RESERVOIR's SOI: Grid Computing + Virtualization + BSM

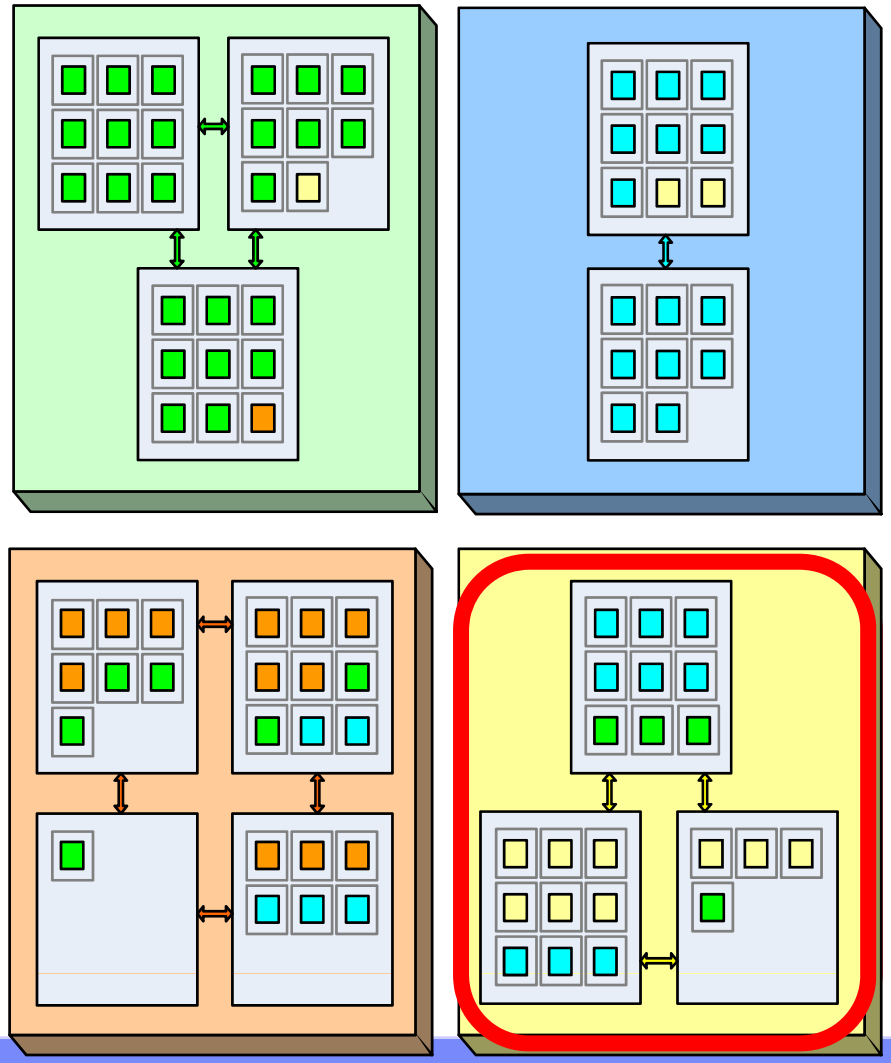

Policy 1:
If possible keep VEEs from the same organization in the same physical box

RESERVOIR's SOI: Grid Computing + Virtualization + BSM

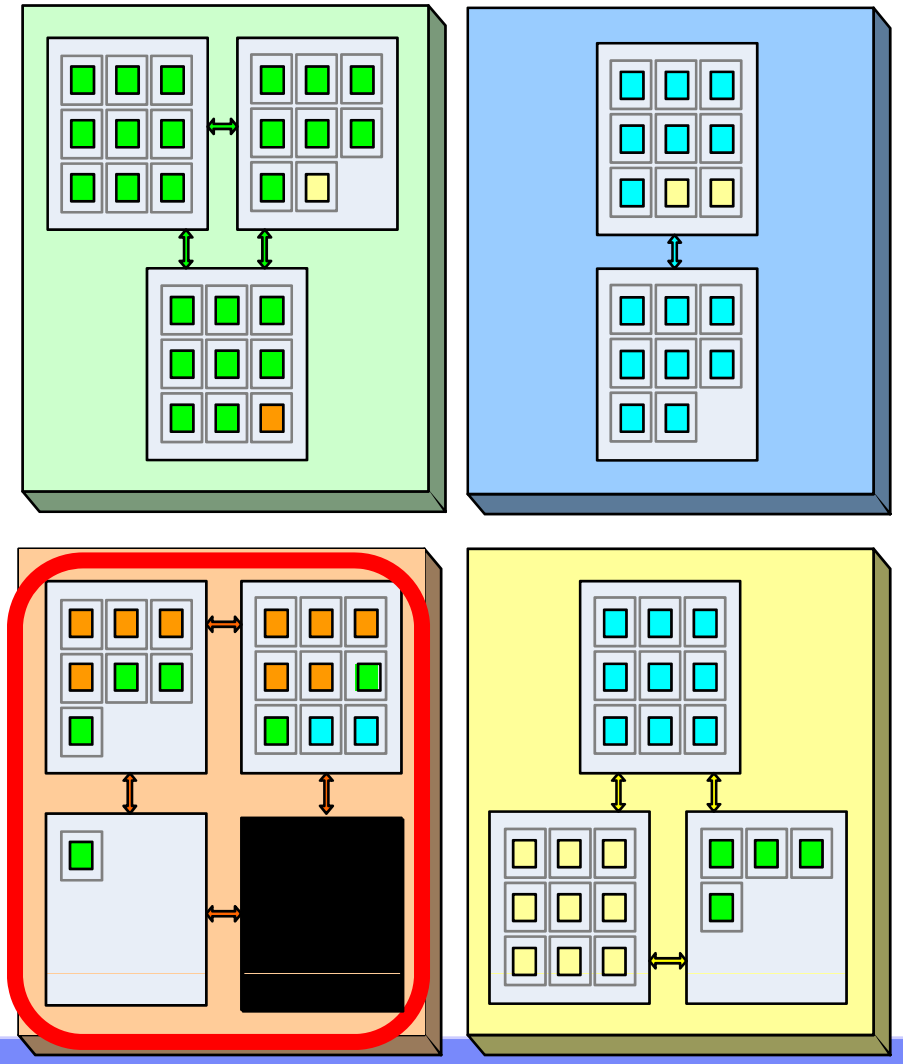

Policy 1:
If possible keep VEEs from the same organization in the same physical box

Policy 2:
Turn off underutilized physical boxes

RESERVOIR's SOI: Grid Computing + Virtualization + BSM

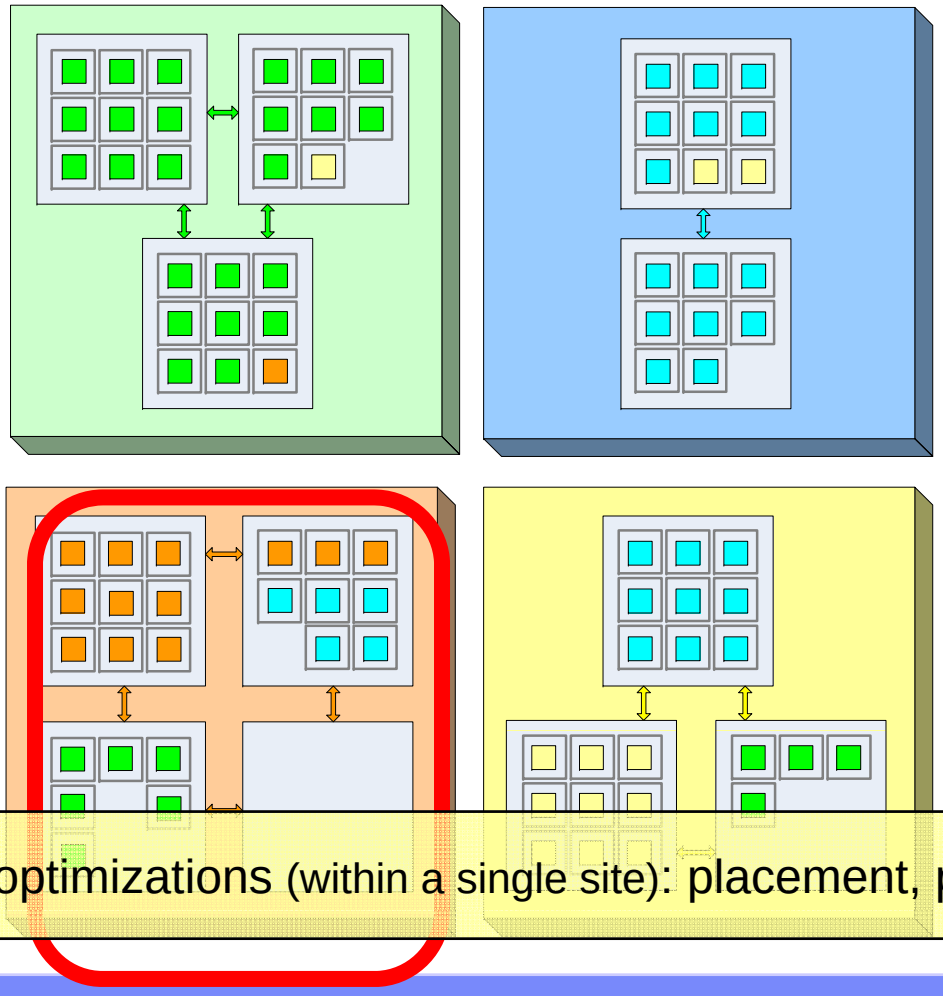

Policy 1:
If possible keep VEEs from the same organization in the same physical box

Policy 2:
Turn off underutilized physical boxes

Local optimizations (within a single site): placement, power, etc.

SOI: Grid Computing + Virtualization + BSM – Boundaries

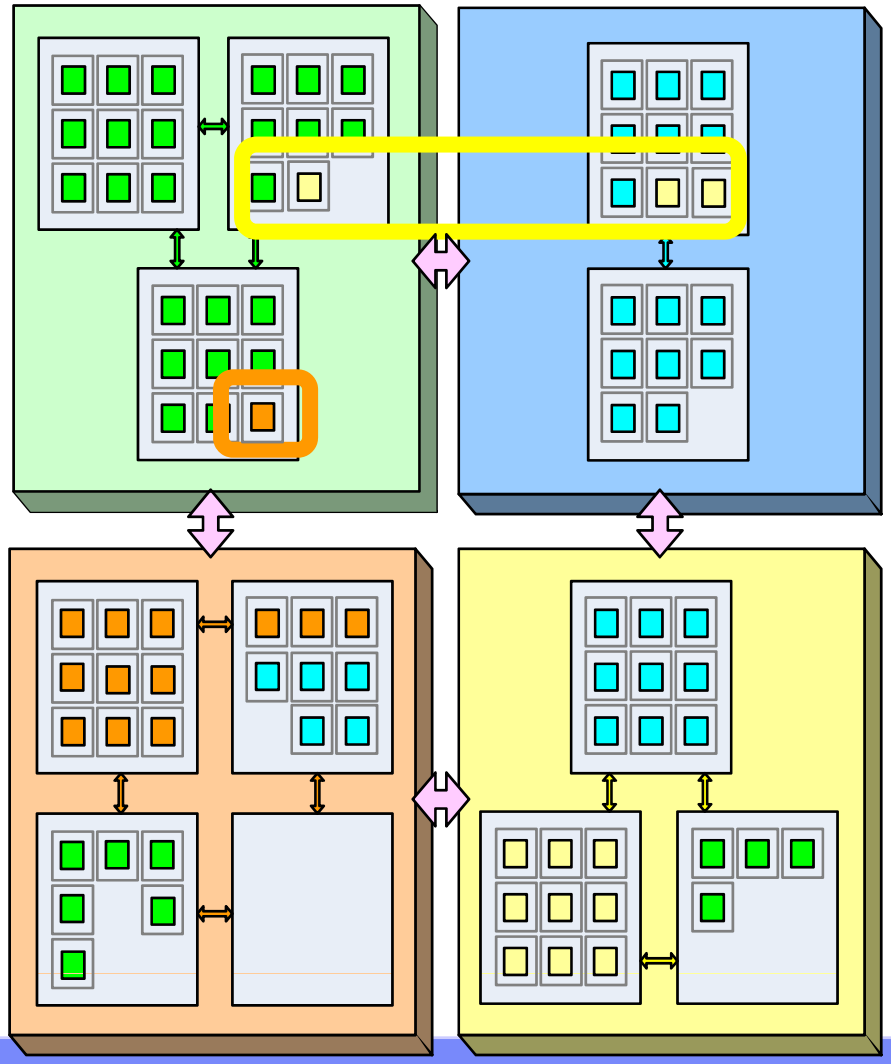

Policy 3:
If possible keep VEEs in “owning” organization

SOI: Grid Computing + Virtualization + BSM – Boundaries

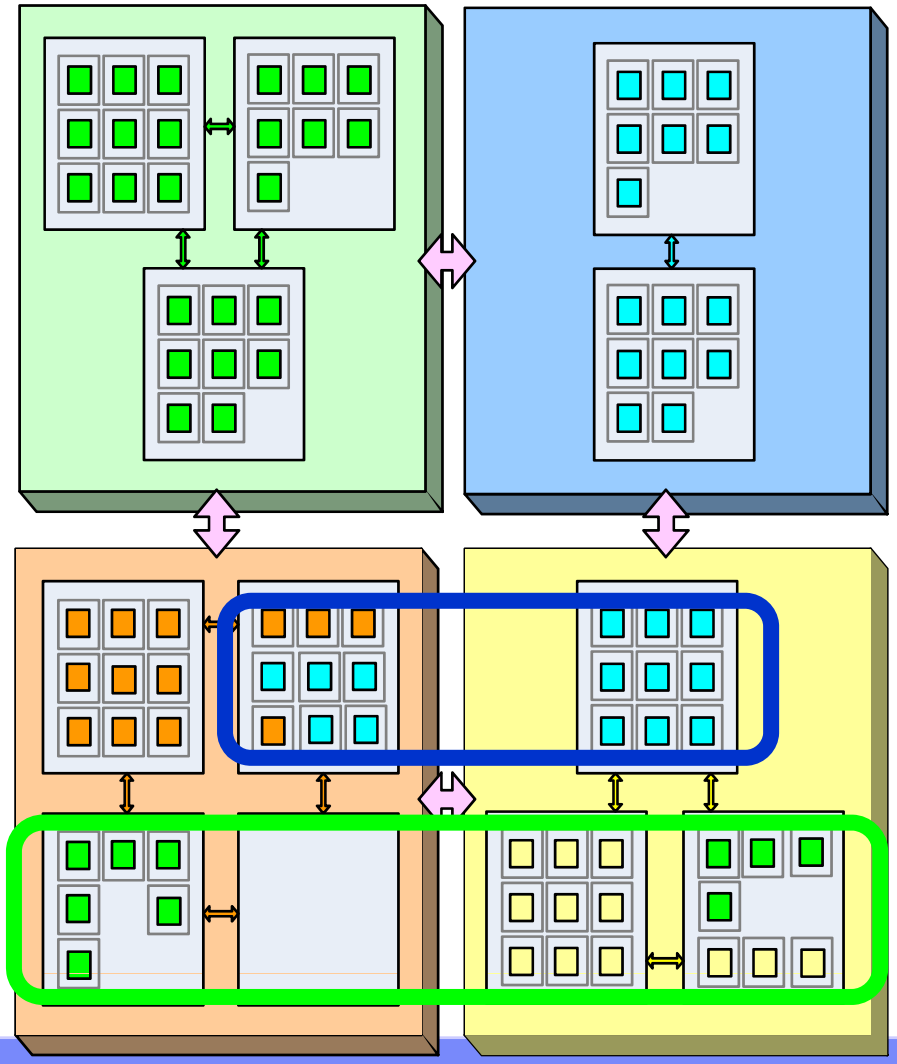

Policy 3:
If possible keep VEEs in “owning” organization

Policy 4:
If possible keep VEEs in least number of external organizations

SOI: Grid Computing + Virtualization + BSM – Boundaries

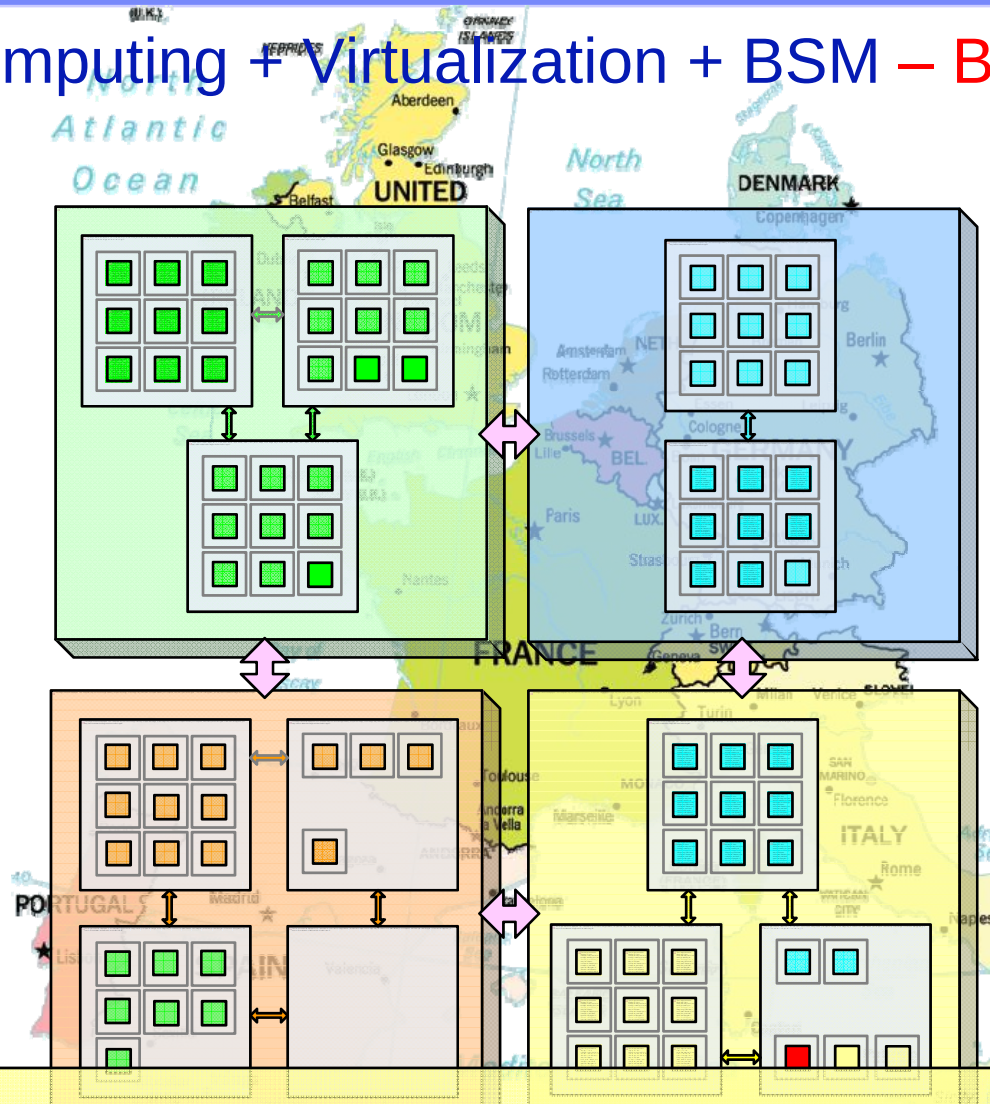

Policy 5:
"Follow" the service customer

Migration across sites ➔ **Global optimizations**: placement, cost, bandwidth, etc.

RESERVOIR - Virtualize the Network

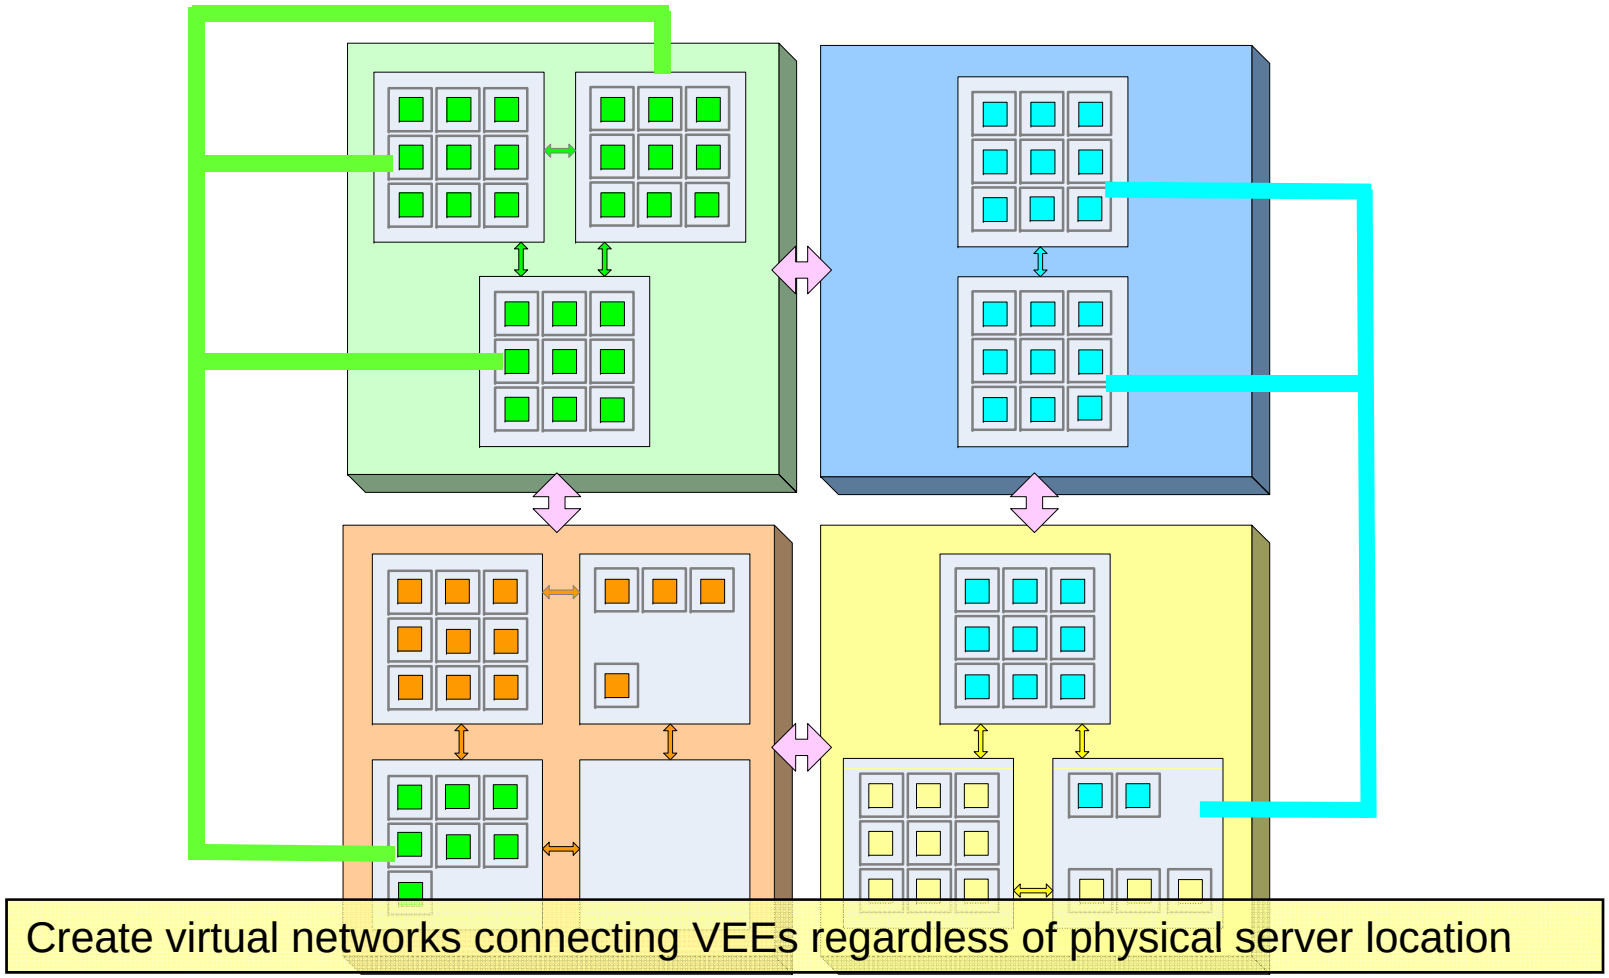

RESERVOIR - Virtualize the Network and the Storage

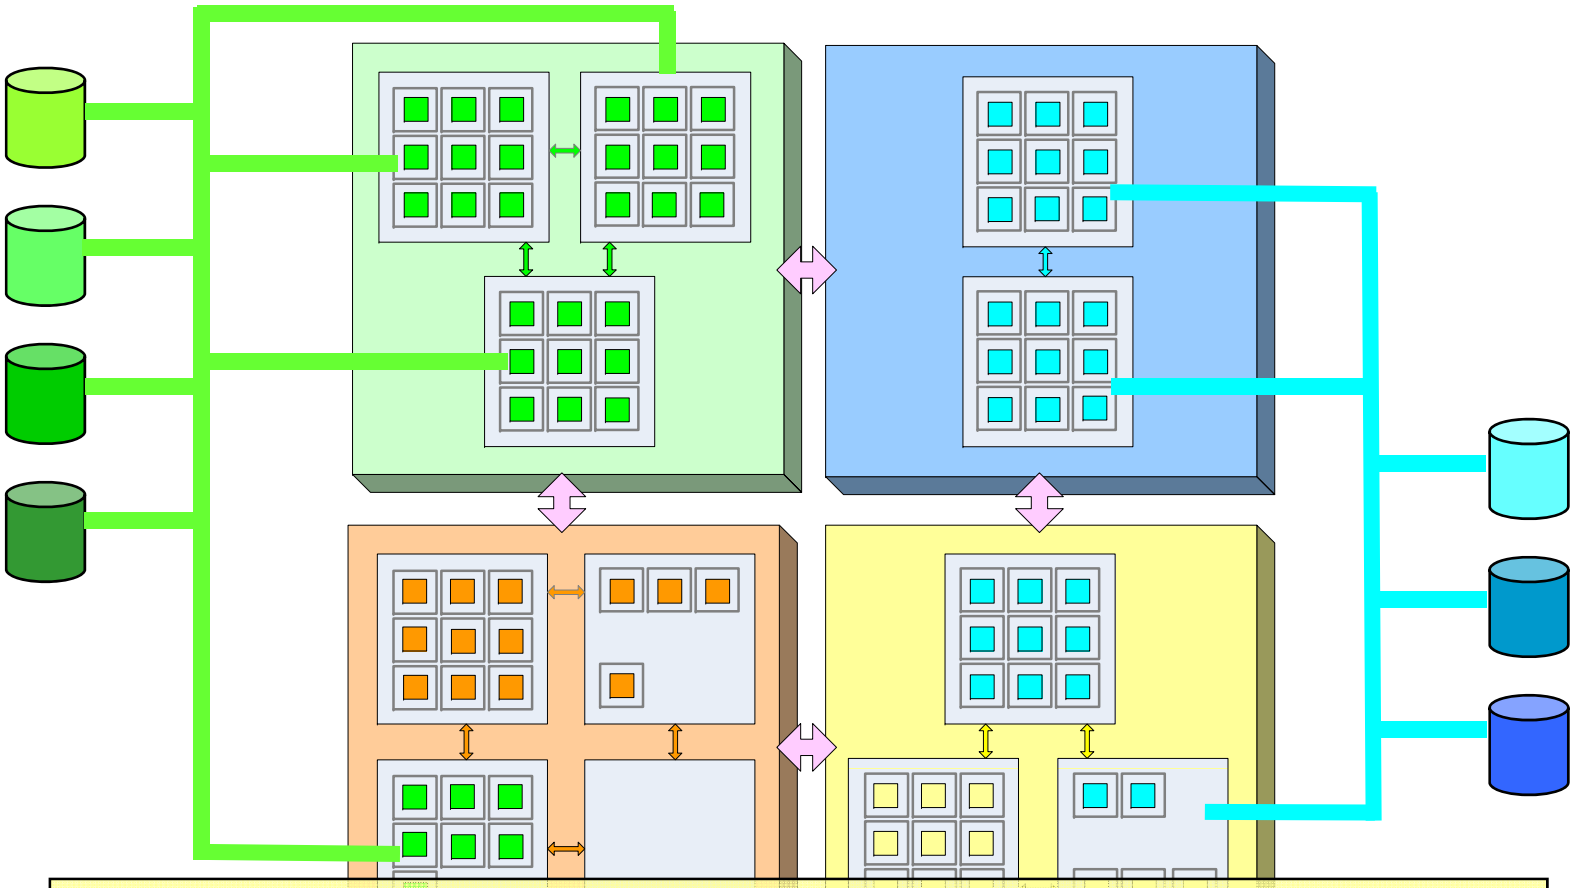

Enable secure and efficient access to relevant data, and ensure data consistency, regardless of storage location

The RESERVOIR Consortium

Partner	Role	Details
IBM HRL	Technology	Project Lead, Virtualization/SOA Infrastructure
Telefonica I+D	Technology	Service Technology, Billing Infrastructure
UCM	Technology	Grid, Dynamic Allocation Technology
Thales	Technology	Security, Virtualization Infrastructure, Hosting
SAP	Use-Cases	Use-Cases, Contribution to Requirement an Standards
Sun Microsystems	Use-Cases + Tech	Contribution to Standards, Java Services, Monitoring
DATAMAT	Technology	Service Management Technologies
University Lugano	Technology	Partner, Monitoring and SLA Management
University UMEA	Technology	Monitoring, Measuring and Billing Technology
University Messina	Technology	Grid Experience, Testbed Development,
UC London	Technology	Virtualization Technology
CETIC	Technology	Security
OGF	Standardization	Grid and Virtualization Standards

The RESERVOIR use cases

- We have four use cases from industrial partners
  - SAP – Business application
    - Multi-tenant service delivery for SMBs in a data center
      - Managing thousands of different service components that comprise a variety of service applications executed by thousands of VEEs.
    - Deployment of a business application with one click
      - Deployment based on Service Manifest
    - Relocation of a multi-tiered business application
  - Sun –
    - Utility computing (example: digital content creation – such as computer generated animated films or special effects), or Web 2.0 application
      - Test performance under the following conditions:
        - Frequent change of resource use
        - Unpredictable loads
        - Pay as you go use.

Use cases cont.

- Telefonica – Telco application
  - Test a mash-up application on top of the RESERVOIR infrastructure
    - Lifecycle management in a highly dynamic environment
    - Accounting, billing and business management from different gadget providers.
- Thales – e-Government
  - Three-tiered application
    - Tests handling dynamic loads – cyclic demands on a user application to reflect hourly/seasonal peaks
    - Maintenance scenarios to physical resources
    - Application manageability – QoS tradeoffs, large number of simultaneous connections etc.