

RESERVOIR

Resources and Services
Virtualization without Barriers

Philipp Drum
Sun Microsystems



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.

Agenda

- **Part 1: Project overview**
 - Description
 - Partners
 - Technology
 - Architecture

- **Part 2: Sun contribution**
 - Sun Use Case
 - VJSC

Key Points

- Funded by the European Union
- Budget: 17M Euro
- Seventh Framework Program (Research and Technological Development)
- 13 partners
- 3 Years (Feb 2008 – Jan 2011)
- Visit <http://www.reservoir-fp7.eu/>

Partners

Partner	Role	Comment
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

Buzzwords

- Cloud Computing
- Virtualization
- Large scale applications
- Automatic deployment
- Utility computing
- Service oriented computing

What is it?

Provide revolutionary foundation for a new European infrastructure with transparent and dynamic...

- Management
- Provisioning
- Relocation (virtually “without borders”)
- Utility-like deployment

... of Services

Analogies

- Electrical power delivery: capacity can be shifted to guarantee supply and lower costs
- Roaming cellular communications

Goals

- Minimize over-provisioning of resources
 - Better utilization by dynamic allocation and re-provisioning
 - Break down platform and geography barriers
- Adhere to SLA constraints through intelligent placement and relocation algorithms (cross administrative: may be in different organizations)
- Create standards to allow for interoperability between administrative domains
 - SLA
 - Billing
 - Application meta-data...
- Reducing software cost
- Expediting time to market

Differentiators

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

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.

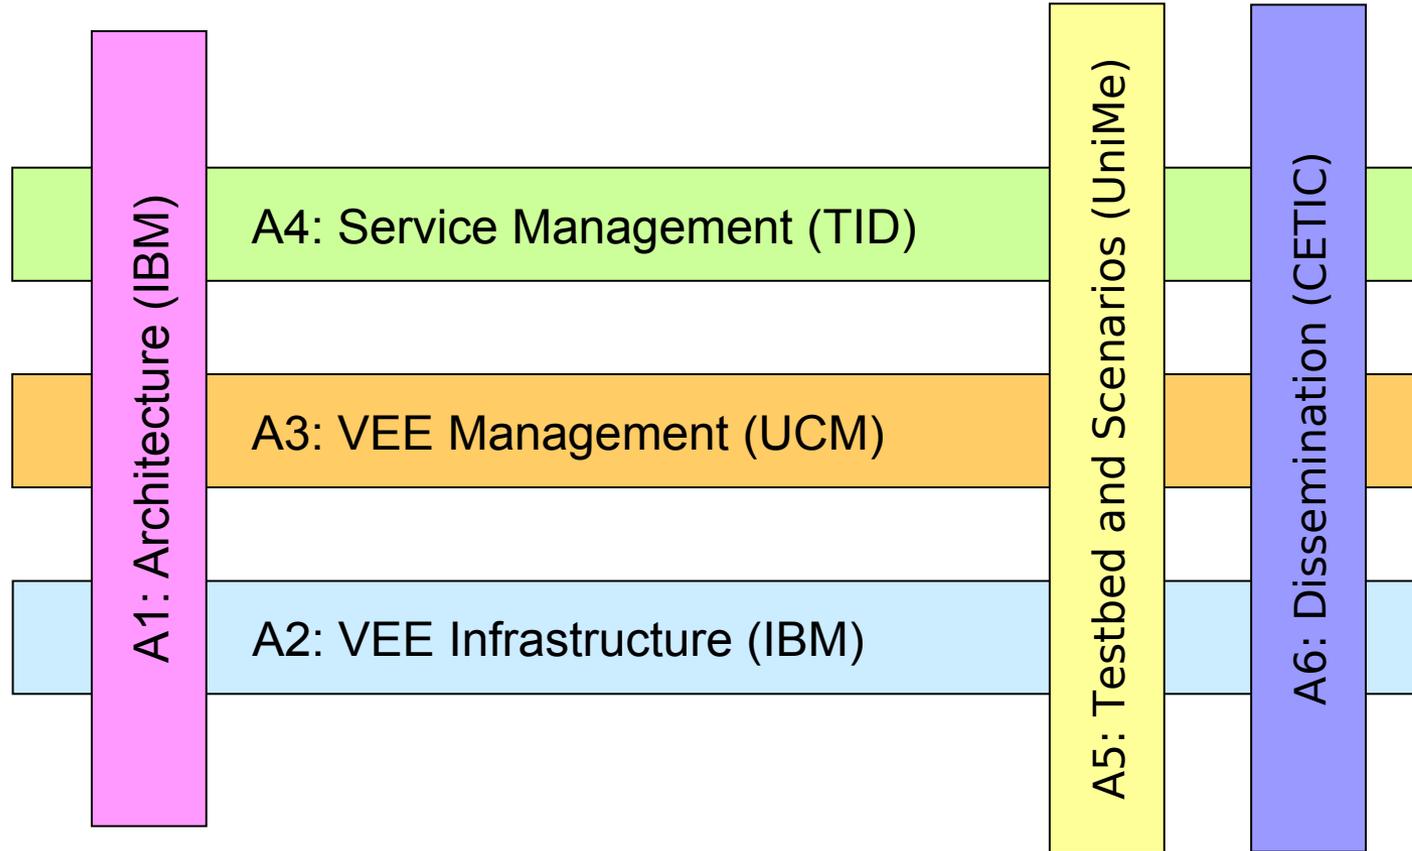
Live migration without borders: Cross geographical, network and management domains

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

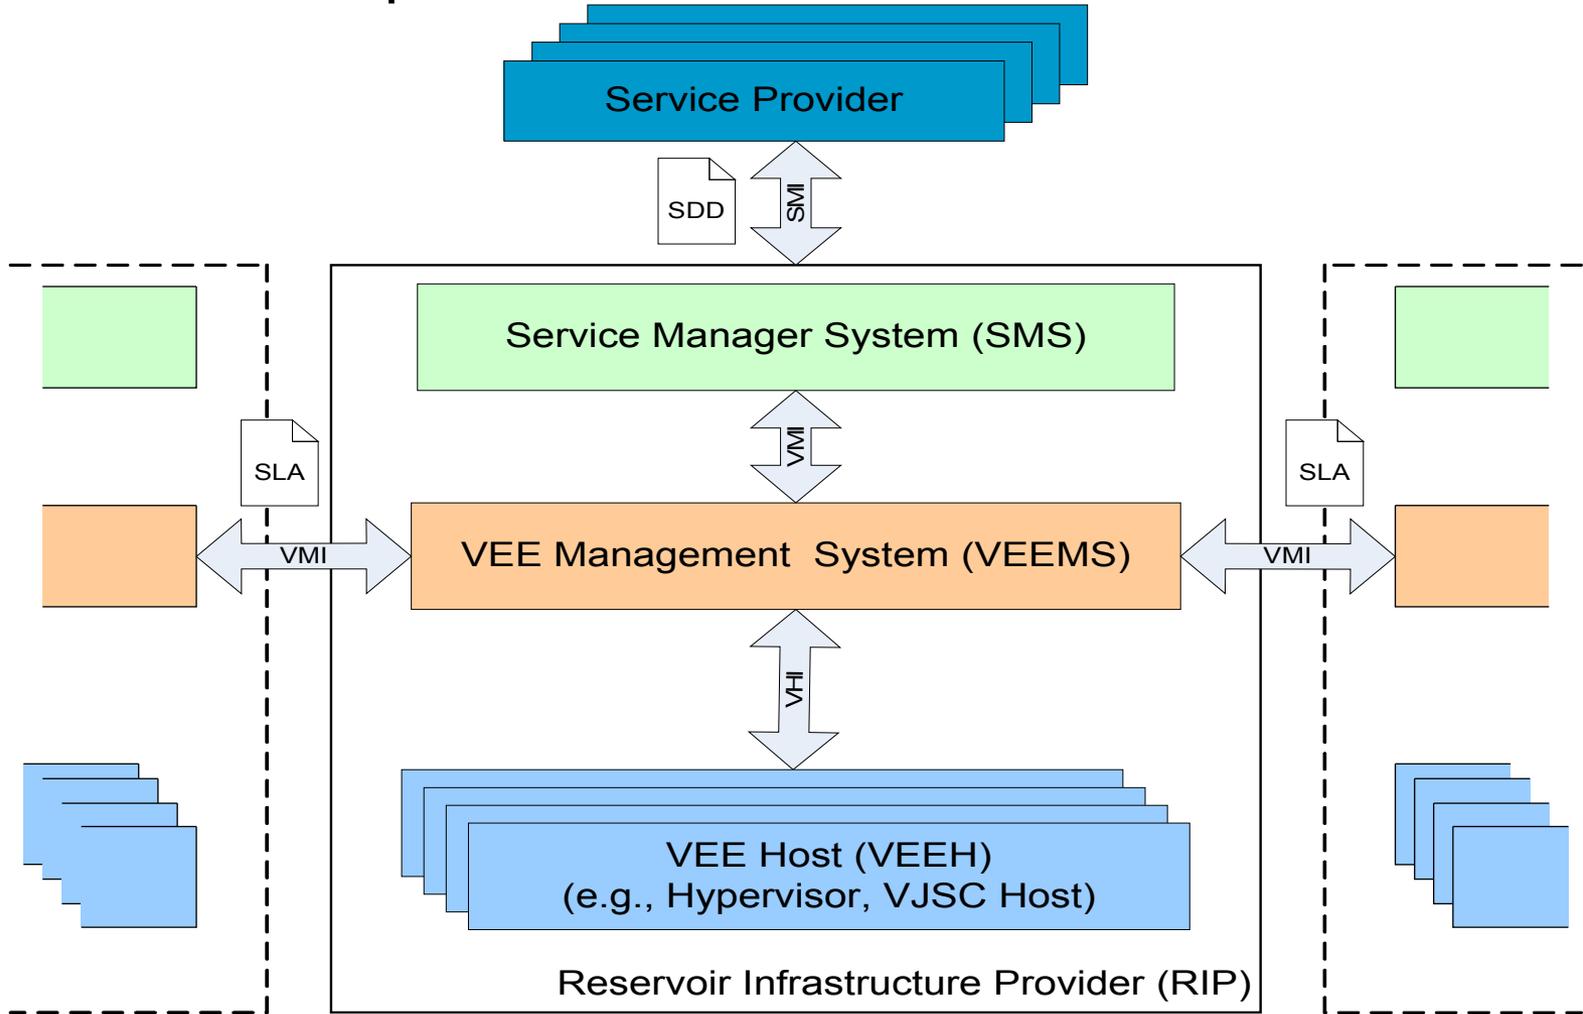
Organization

- **A1 Global technical assembly:** Definition of the overall architecture, Integration plan of all layers, Collaboration with other FP7 Projects
- **A2 Virtual Execution Environment (VEE) infrastructure:** Performance improvements, optimized I/O, migration technology, Virtual Java Service Containers, RESERVOIR integration
- **A3 VEE management:** Provisioning and supervision, placement policies, federation of management domains
- **A4 Service Management:** Service definition, lifecycle management, accounting, payment and billing.
- **A5 Use cases and testbed development**
- **A6 Market Analysis, exploitation plan, publishing, public promotion, standardization and training**

Organization (cont)



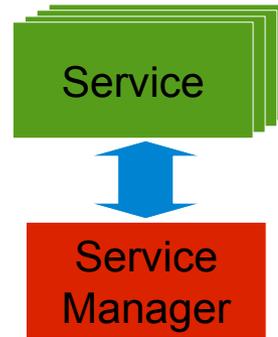
Architectural Perspective



Architecture



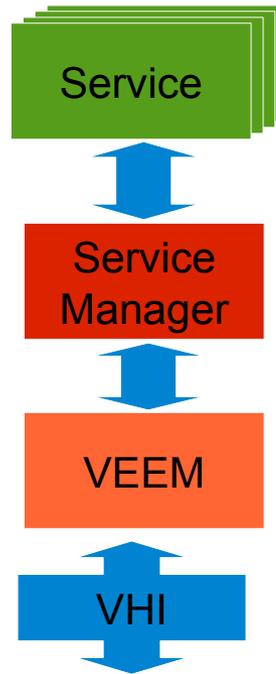
Architecture



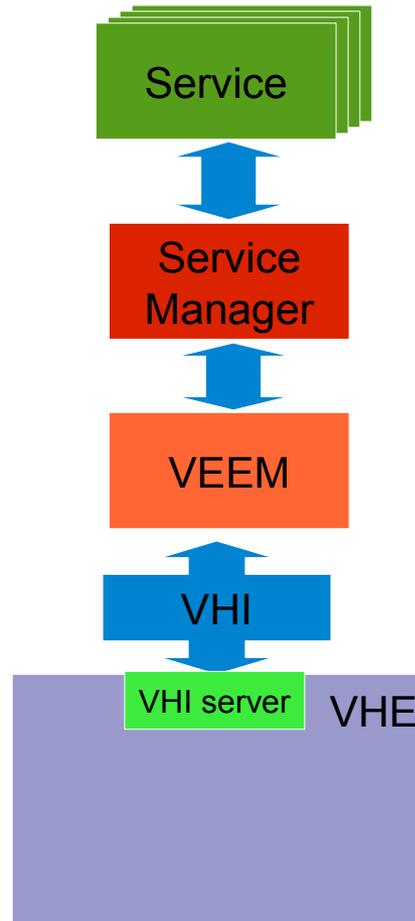
Architecture



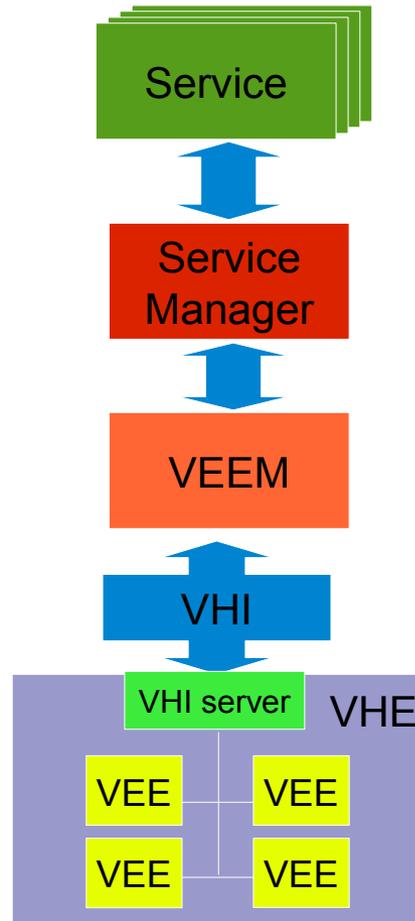
Architecture



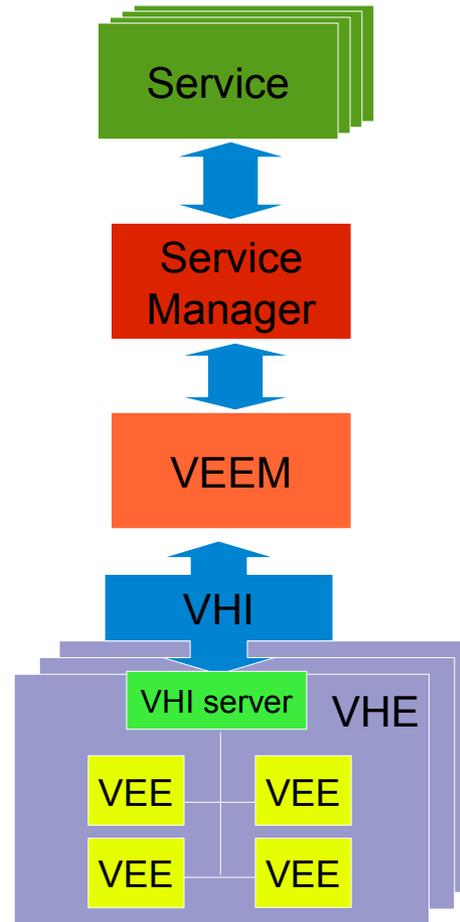
Architecture



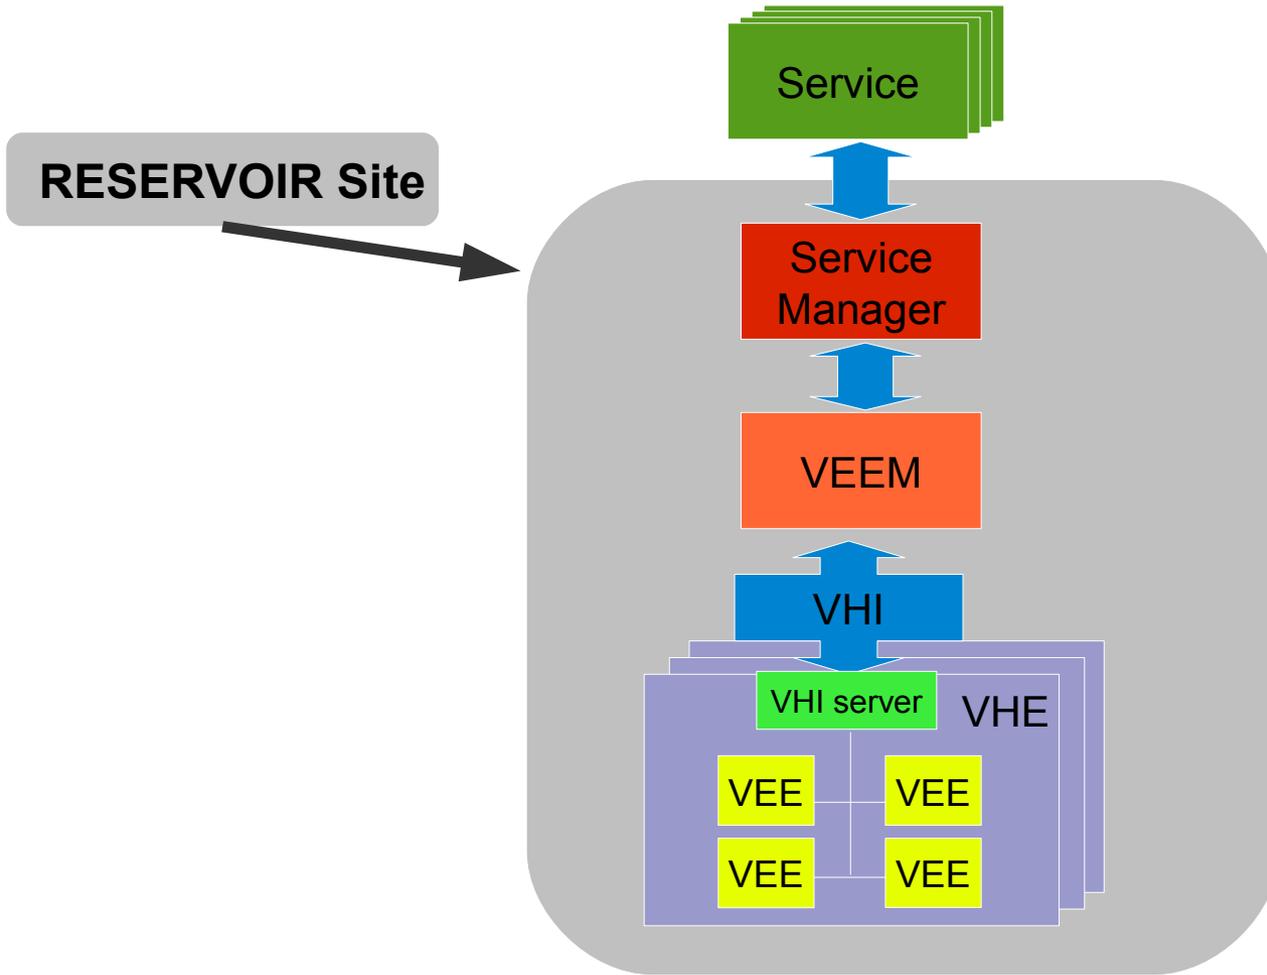
Architecture



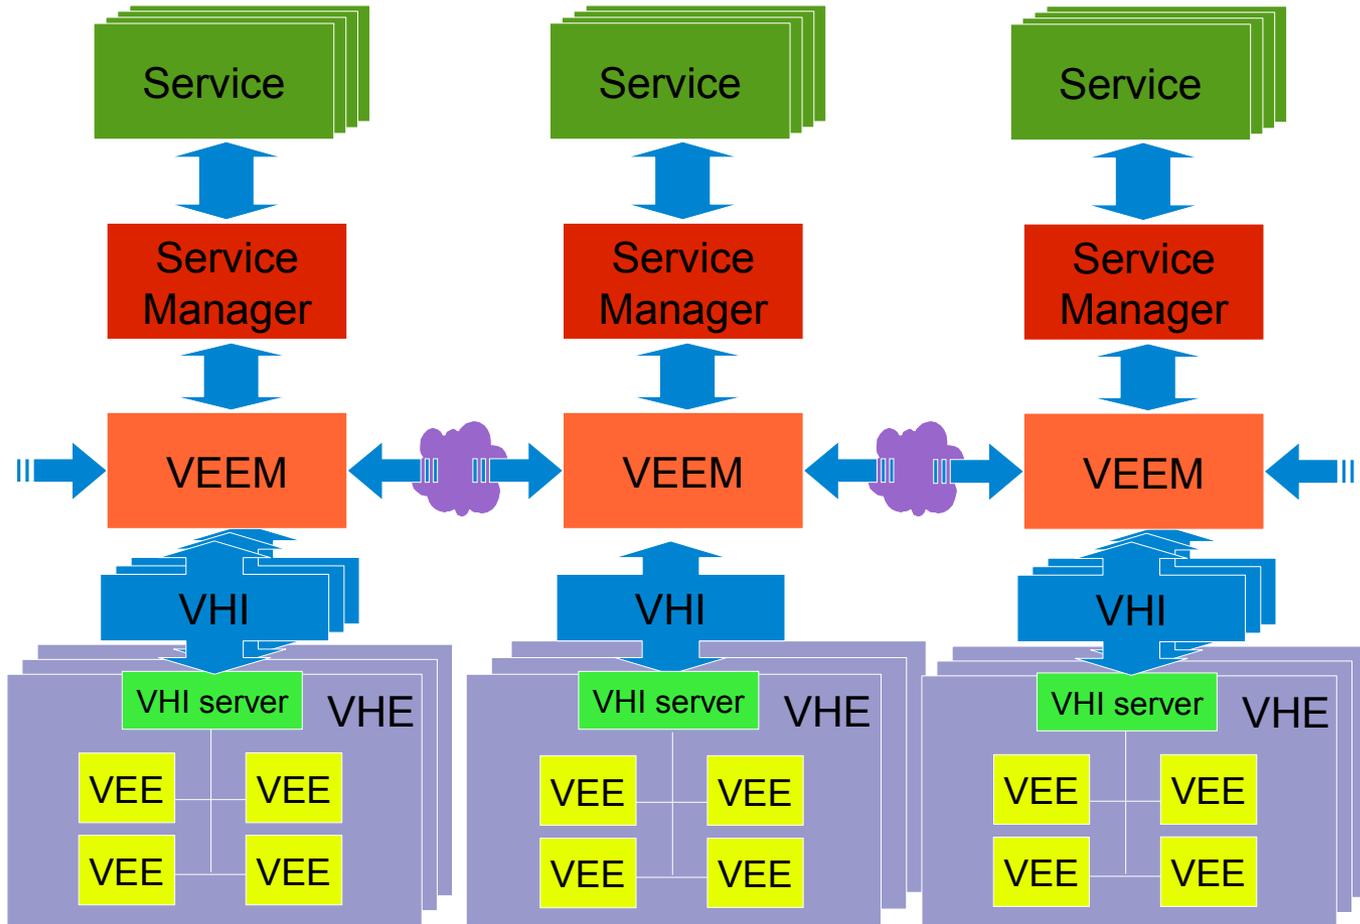
Architecture



Architecture



Architecture



The RESERVOIR envisioned highly dynamic and cooperative infrastructure relies on:

- A service definition language that captures in a high level language the functional and infrastructure requirements of the service (including servers, images, network, storage, inter-tier relations and QoS requirements)
- An abstraction layer that separates implementation details from the high level automation system that is responsible for the provisioning, monitoring and reallocation of resources
- Inter-domain protocols that enable multiple management sites to cooperate in providing a single service, where the cooperation is automatically driven from a service definition document (fully automated cross-domains SLA management)
- The capability of creating fully isolated virtual organizations spread across geographies and management domains
- The flexibility of placing and relocating service instances on resources anywhere even across geographies and management domains

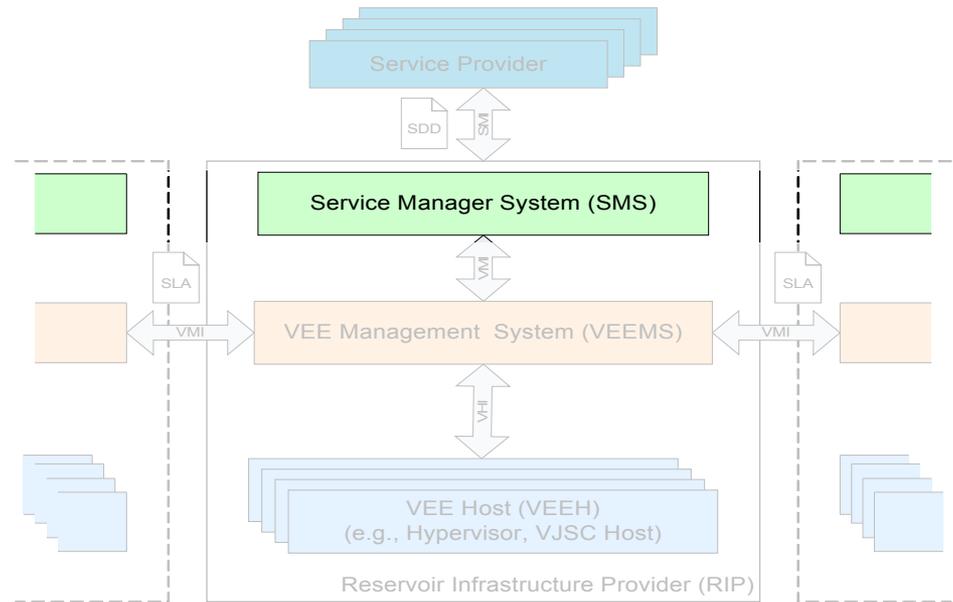
The Service Manager

Instantiation of the service application

- requesting the creation
- configuration of VEEs for each service component

Compliance to the Elasticity rules

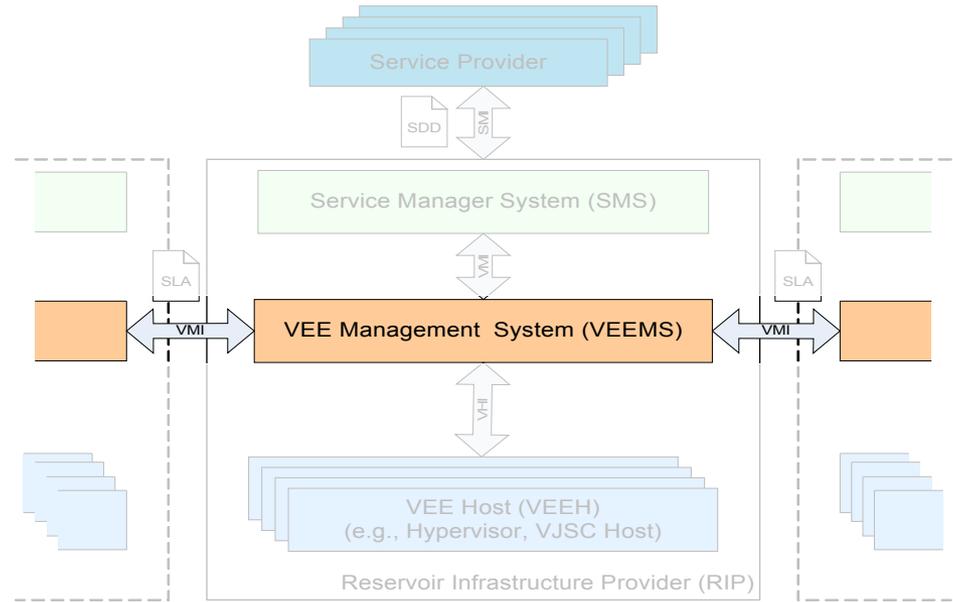
- monitoring the service application
- triggers up and down scaling



The VEE Manager

Placement of VEEs into VEE hosts requests from the Service Manager

- create VEEs
- resize VEEs
- placement (including remote sites)
 - affinity or anti-affinity
 - security
 - costs
- responsible for the federation of remote sites

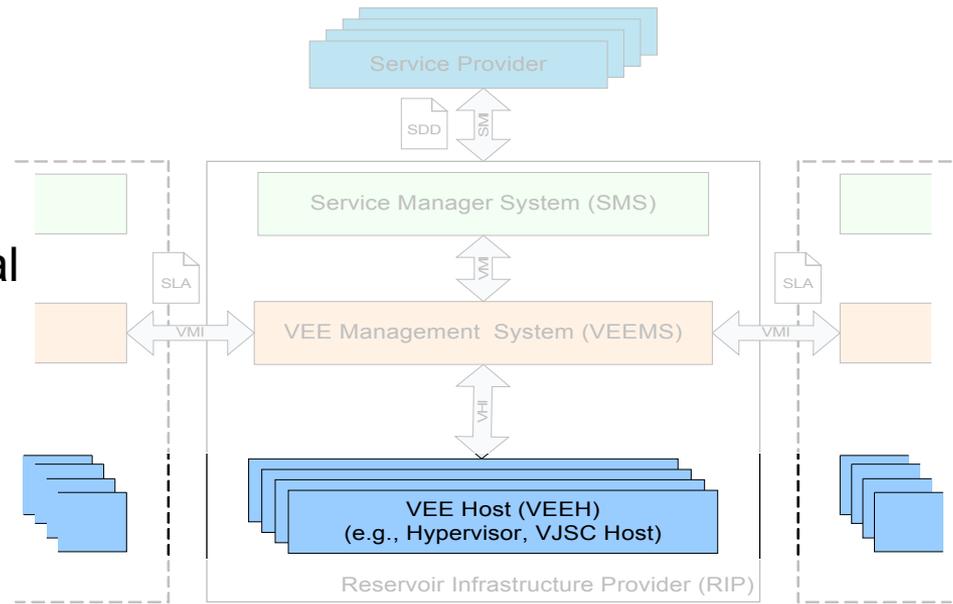


The VEE Host

Abstraction layer between the logical processing of the system and the actual virtualized resource.

The virtualized resource can host a certain type of VEE:

- Xen domains
- VirtualBox or VMWare
- VJSC

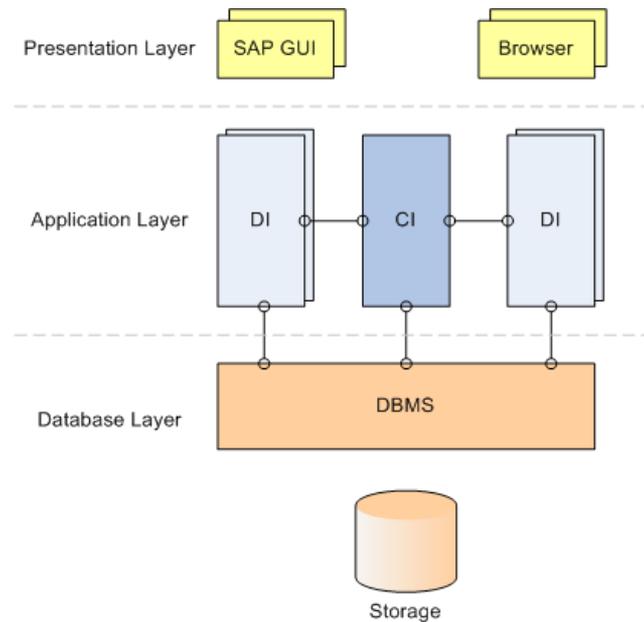


Architecture notes

- Clear separation of service provider and infrastructure provider
 - Service Provider knows about business aspects
 - Infrastructure provider knows about resources
 - The link is the service manifest and SLAs
- Virtualization as a key technology

Use Case Scenario 1: SAP business application (SAP)

- Business orientated use case
- Classical multi tier application (DB layer, application layer, presentation layer)
- State-oriented workload spans multiple tiers
- Applications:
 - SAP Netweaver
 - SAP Web Application Server
 - SAP MaxDB



Use Case Scenario 2: Telco application (TID)

- Web sites with massive access
 - Olympic events
 - Ads while the superbowl
- End user specific customization of web pages
 - RESERVOIR could move resources to location near the end user
 - Timezones could be taken into account
- VJSC as platform

Use Case Scenario 3: Utility Computing (SUN)

- Running the Sun Grid Engine as a service inside the Reservoir framework
- On demand resource provisioning
 - Fast and flexible response to changing requirements is key in many business areas
 - example DCC or genome computations
- No over-sizing of the grid
- Large, cross site grids
- Lower operational costs
- Access to Utility Computing Resources: Deploy arbitrary operating system and application stacks on remote resources. Provide secure and seamless access to them. Grow, shrink, destroy and recreate the resources without the end user noticing disruption of service.

Use Case Scenario 4: eGovernment application (Thales)

- Application from Thales for the French government
- Service is “cadastre” (land registry), registry service in every municipality
- Real-estate tax and special transitions
- stringent requirements for availability, scalability, reliability and performance

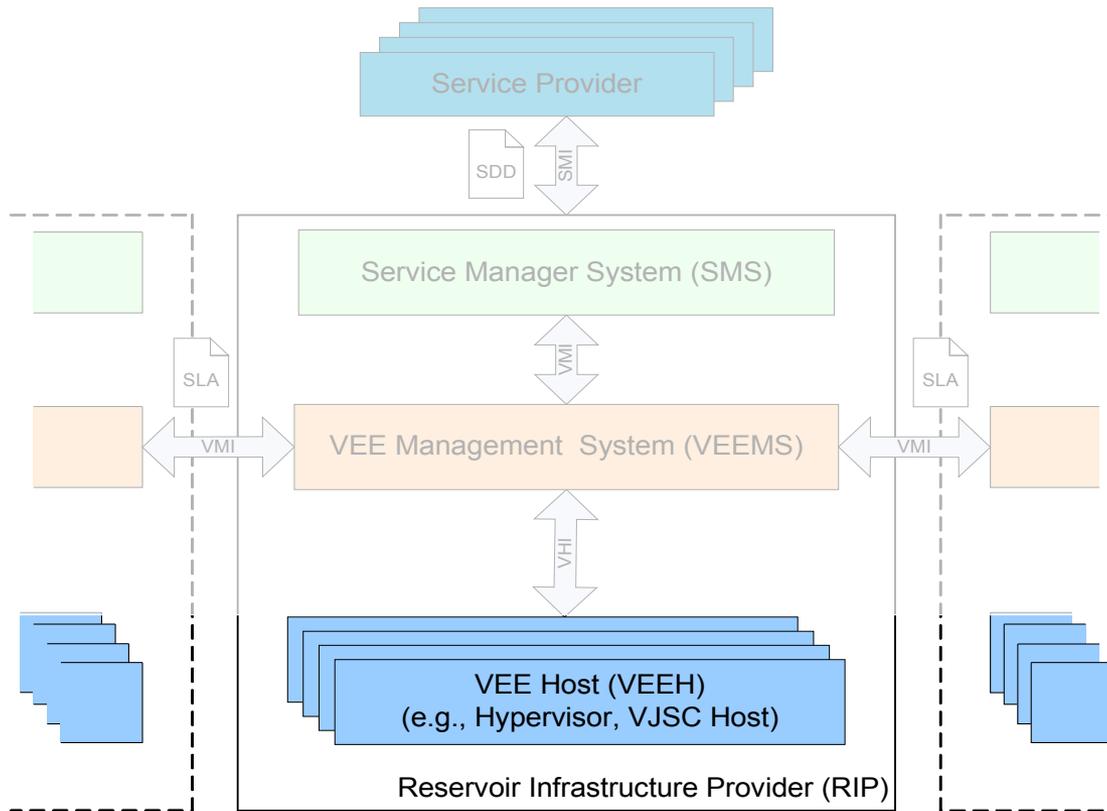
The Others

- Google App Engine
- Amazon EC2
- Flexiscale
- Limelight
- Salesforce.com
- Facebooks apps
- Oracle On Demand
- Citrix Cloud center
- Cisco Vframe
- ConVirt
- Enomaly
- Flexiscale
- Eucalyptus
- Elastra
- Rightscale
- LongJump
- Joyent Accelerator

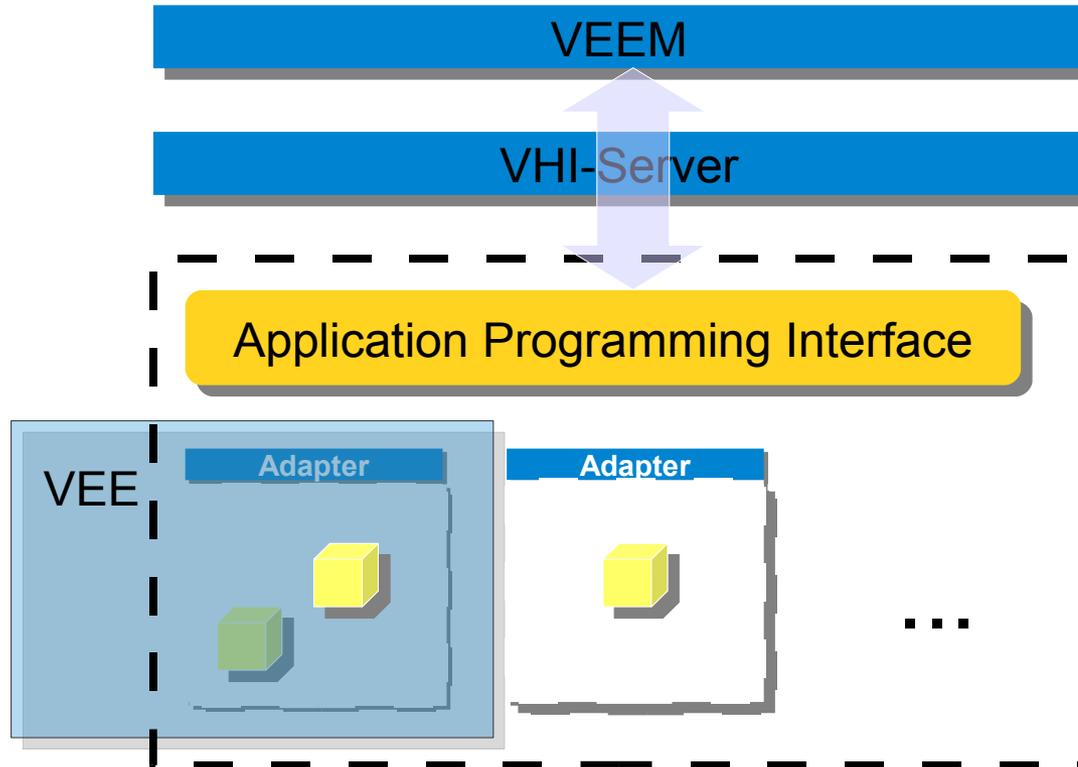
Agenda

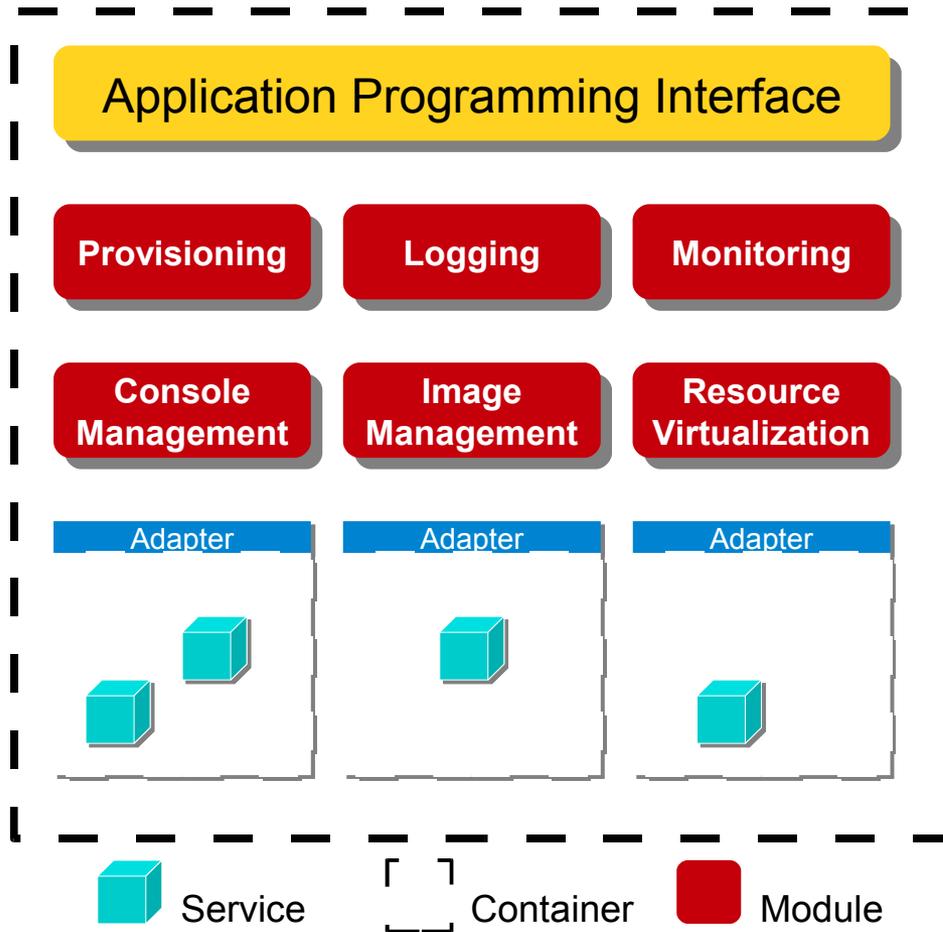
- Part 1: Project overview
 - Description
 - Partners
 - Technology
 - Architecture
- Part 2: Sun contribution
 - Sun Use Case
 - VJSC

Virtual Java Service Container

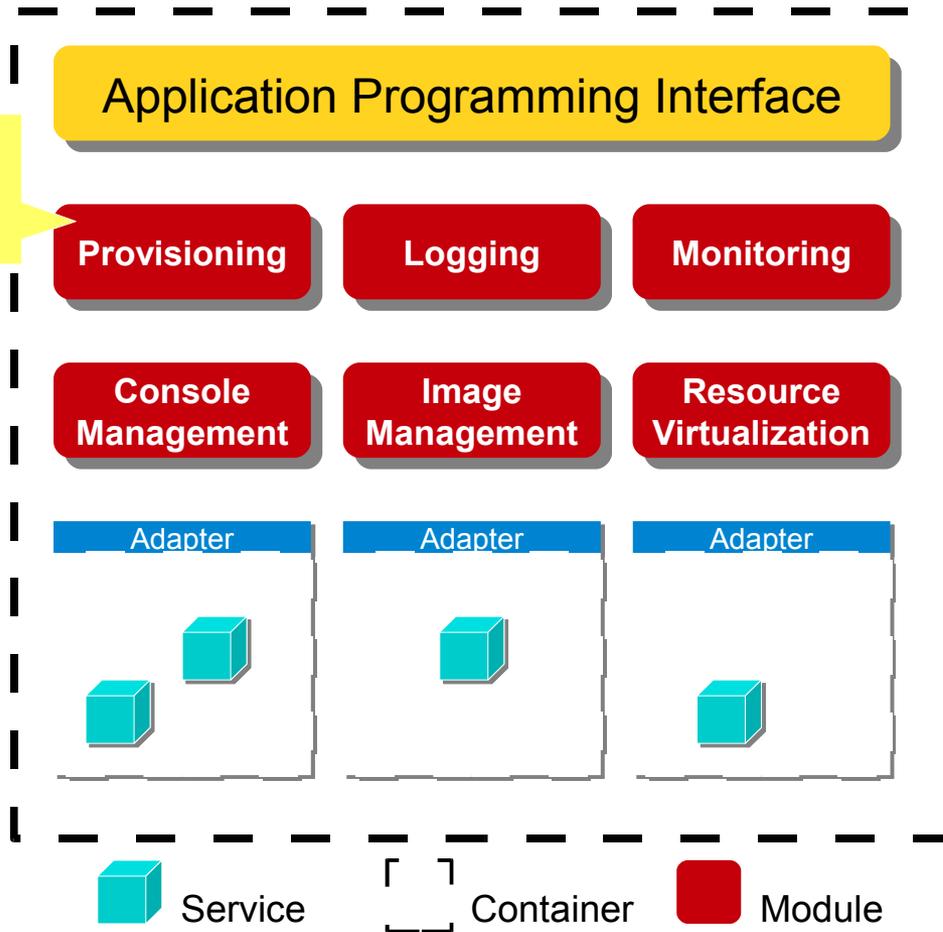


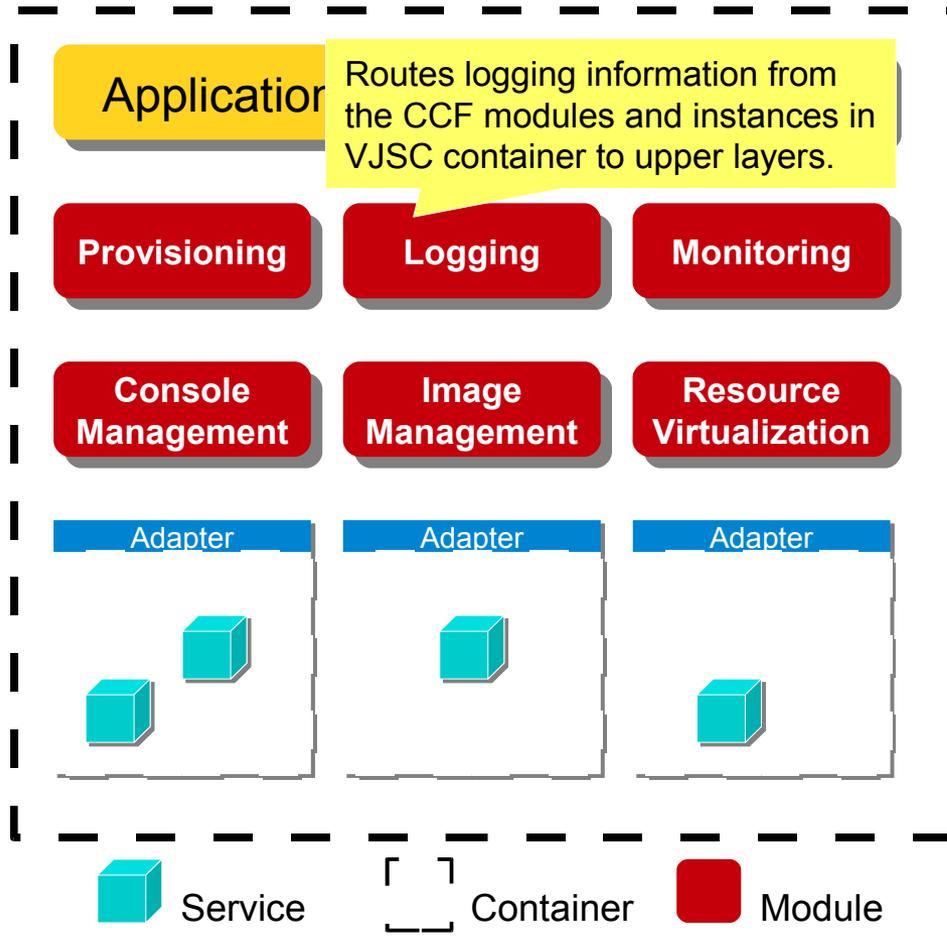
Virtual Java Service Container

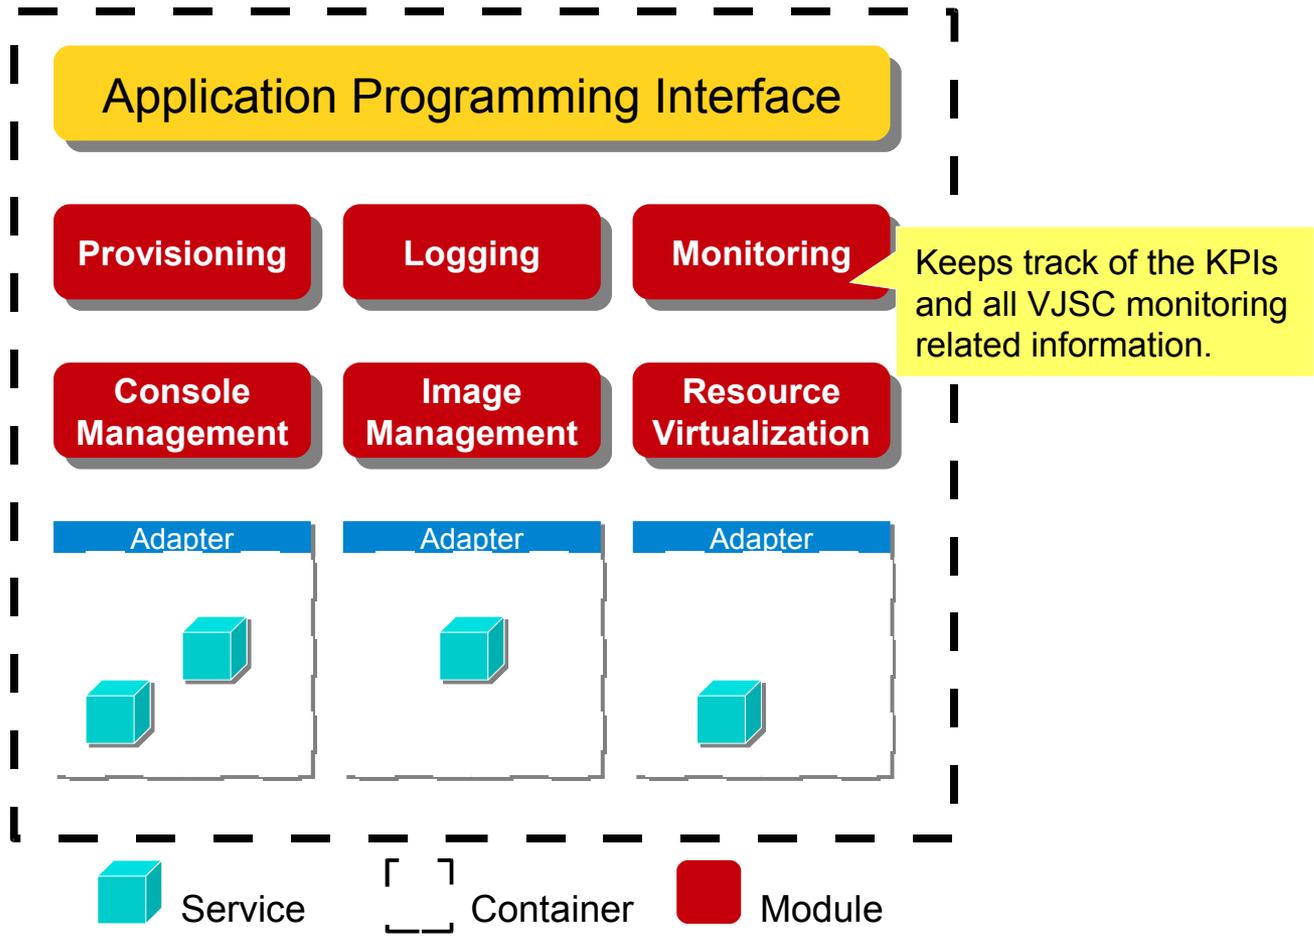


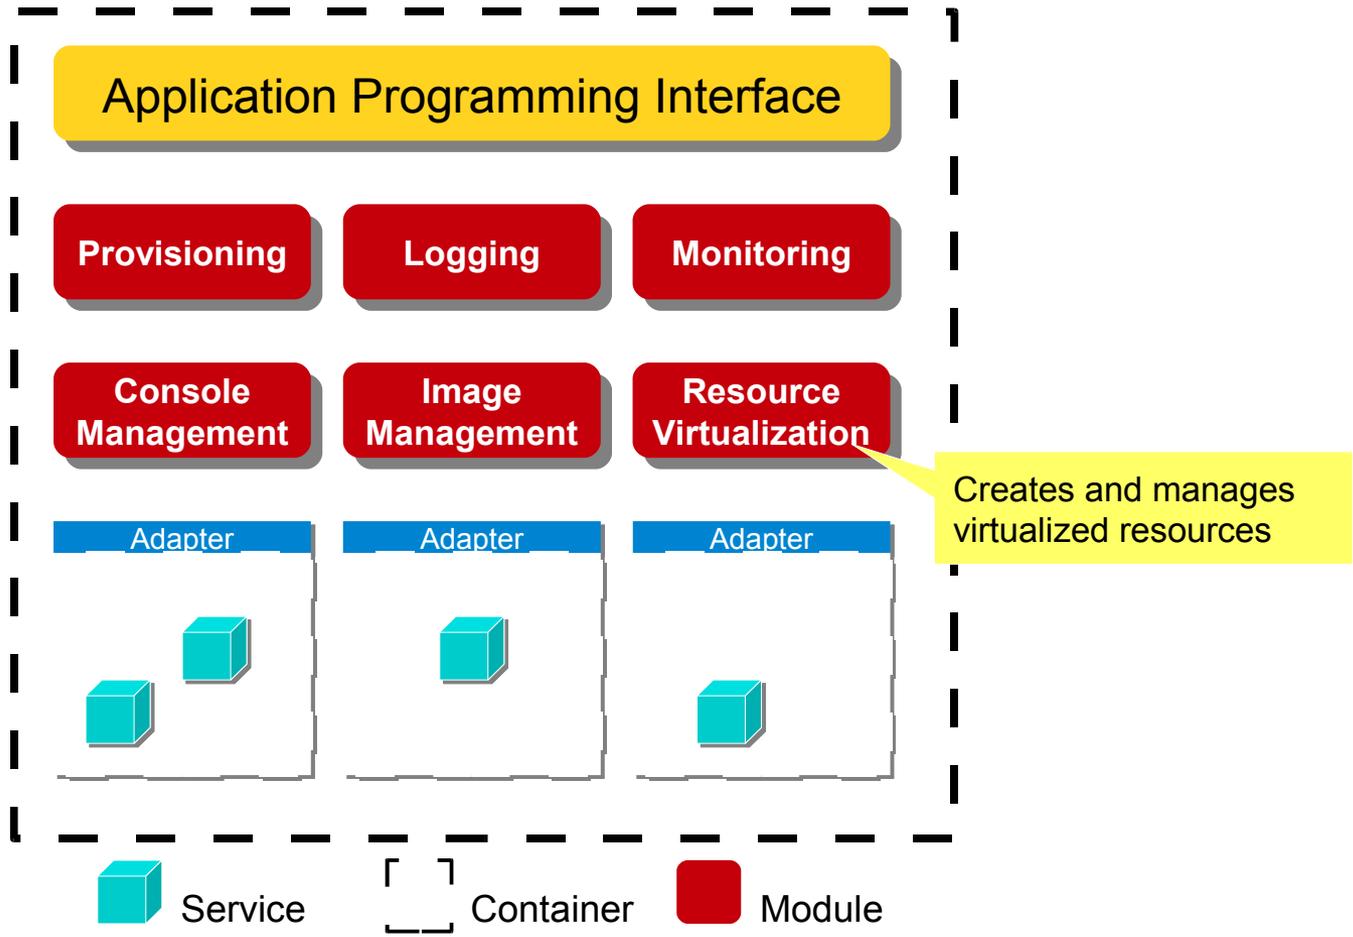


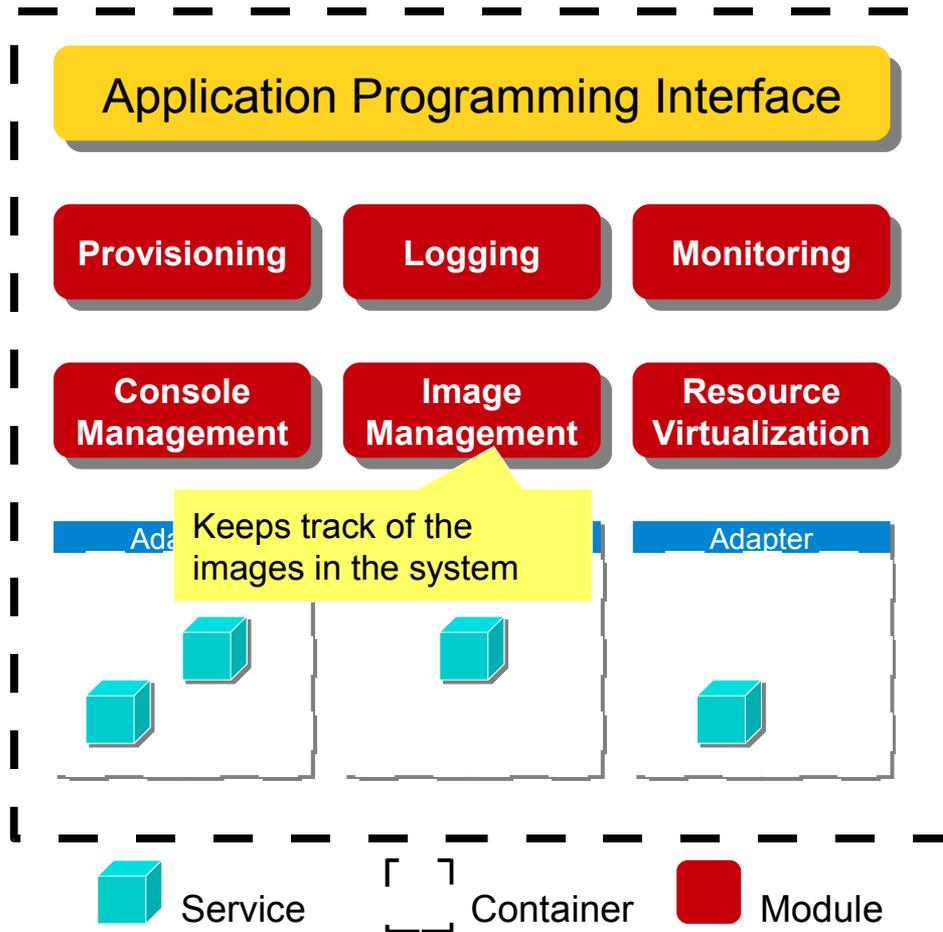
Manages the life cycle state of VJSCs and the deployed Services.











Application Programming Interface

Provisioning

Logging

Monitoring

Console Management

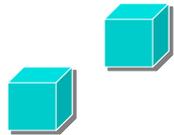
Image Management

Resource Virtualization

Adapter

Adapter

Adapter



 Service

 Container

 Module

Provides access to management consoles provided by the CCF and instances of VJSC containers.

Thank You!

Questions?

Visit <http://www.reservoir-fp7.eu/>