



# *Enabling Scalable Cloud Service Choreographies*

*Marco Aurelio Gerosa*

*Carlos Eduardo Moreira dos Santos (Cadu)*

*Baile Project  
FLOSS Competence Centre  
University of São Paulo (USP)*

*HP Labs - November/2011*

- **Marco Aurelio Gerosa**

- Professor in the Computer Science Dept. of the University of São Paulo. Main areas of interest: Software Engineering and CSCW.

- **Cadu**

- PhD student at University of São Paulo. He is currently studying on how to provide a middleware for ultra-large-scale service choreographies. Prof. Fabio Kon is his advisor.



- **Baile**
  - A joint research project involving the University of São Paulo and HP
- **GSD**
  - Distributed Systems Group
- **CCSL**
  - FLOSS Competence Center
- **DCC**
  - Department of Computer Science
- **USP**
  - University of São Paulo



- Latin America Largest University
- Total area: 76 millions m<sup>2</sup>
- 89,000 students
- 5,800 professors
- 16,000 administrative staff
- 2,300 doctorate degrees per year
- > 25% of Brazilian scientific production



Source: <http://www5.usp.br/en/usp-em-numeros/>

\* USP is the most relevant university in Latin America according to several rankings



- **40 full-time professors**
- **250 undergrads**
- **265 graduate students (190 masters + 75 PhD)**
- **Some research areas**
  - Database systems, distributed systems, software engineering, computer theory, optimization, artificial intelligence, vision and image processing, bioinformatics, musical computing, etc.




- **Distributed systems professors:**





Fabio Kon Alfredo Goldman Marco Gerosa Daniel Batista

- **~ 10 doctoral students**

- **1 postdoc**




- **~ 20 masters students**



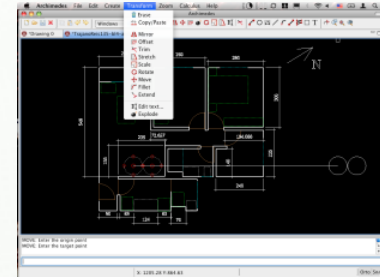
- **~ 10 undergrads**



# Some research projects



- Borboleta (Telehealth with smartphones)
- Agile Methods for Software Development
- Qualipso (Quality in Open Source)
- Groupware Workbench (software components for groupware development)
- Xflow (Mining software repositories)
- InteGrade (Opportunistic Grid Computing)
- BAILE/CHOReOS - Web Service Choreographies



- **Back to the theme:**

- *Enabling Scalable Cloud Service Choreographies*

- **Agenda:**

- Future Internet view
- Service choreographies
- BAILE project
  - A middleware for service choreographies
  - Choreography analysis
  - Prototype demo



- **The growing scale of the Internet demands new theories and technologies.**
- **Initiatives:**
  - NSF Future Internet Design (FIND) initiative:  
<http://www.nets-find.net>
  - European Projects:  
<http://www.future-internet.eu>,  
<http://www.choreos.eu>
  - China Next Generation Internet (CNGI):  
<http://www.cstnet.net.cn/english/cngi/cngi.htm>
  - Japan initiative:  
<http://akari-project.nict.go.jp/eng/overview.htm>
- **No common definition yet**

- **Internet of Content**
  - Content may be combined, mixed or aggregated to generate new content and may be cached or live, static or dynamic, monolithic or modular.
- **Internet of Services**
  - Features are exposed as services that can be integrated into other systems or used to dynamically create new systems.
- **Internet of Things**
  - A global network infrastructure, linking physical and virtual objects through the exploitation of data capture and communication capabilities.

# Some requirements

- **Scalability**
- **Interoperability**
- **Mobility**
- **Awareness and Adaptability**
- **Security, Privacy & Trust**

| <b>Today's internet</b>                 | <b>Toward the Future Internet</b>            |
|---|--|
| 1 billion PCs (2008)                    | 1.78 billion PCs (2013)                      |
| 647 million smartphones (2010)          | 1.82 billion smartphones (2013)              |
| 5 exabytes of data (2005)               | 990 exabytes of data (end of 2012)           |
| 10,000 services (2007)                  | Billions of services                         |
| 10 billion terminals (2010)             | 100 billion terminals (2015)                 |
| Consumer traffic of 12.7 exabytes/month | Consumer traffic of 42 exabytes/month (2014) |

## Today's internet

Islands of interconnected elements

Service/content mashups leading to the provision of new, diverse services by prosumers

Wide-spread usage of smart mobile devices with limited resources

Service-oriented software development is mostly a static process

## Toward the Future Internet

Internet-scale connection of highly heterogeneous elements (vehicles, sensors, mobiles devices, home appliances, etc.)

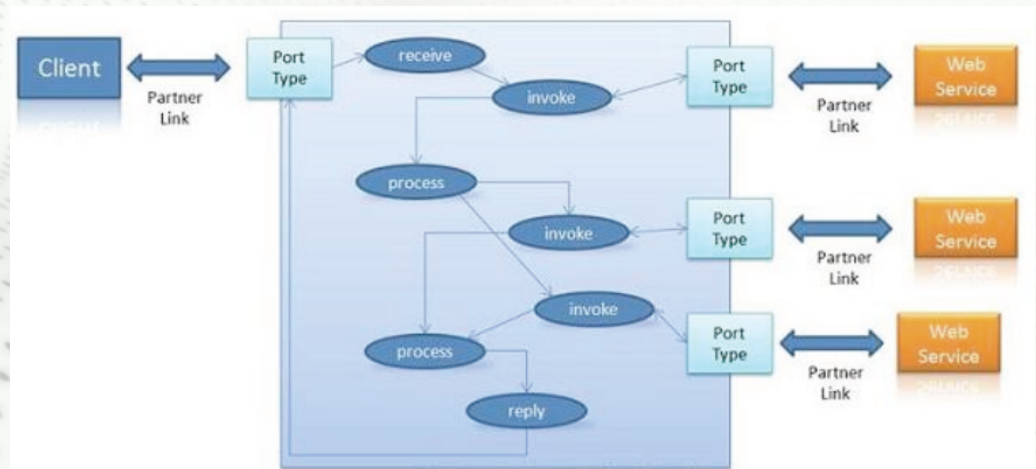
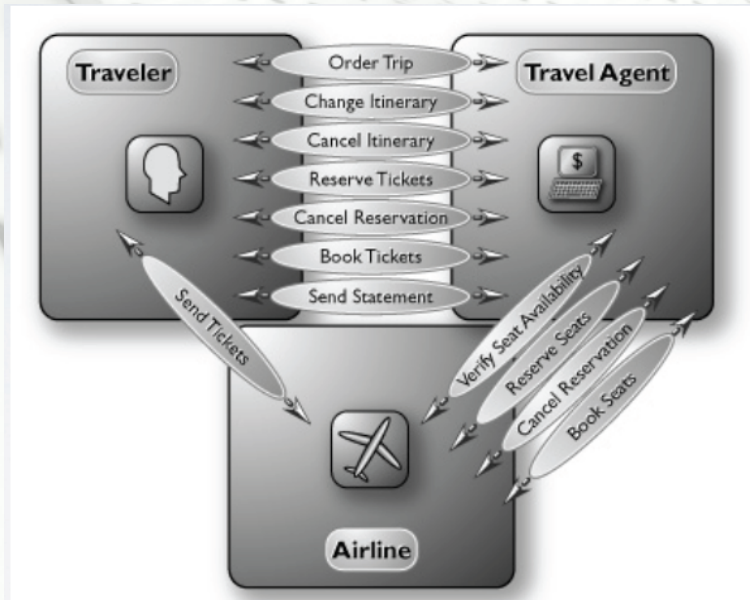
Global-scale services/content mashups creating new services/content with different types and formats

Global scale usage of mobile devices with ever-growing capabilities. The majority of the connected devices will be mobile.

Software development is a completely dynamic user-centric process.

# Approaches for service composition

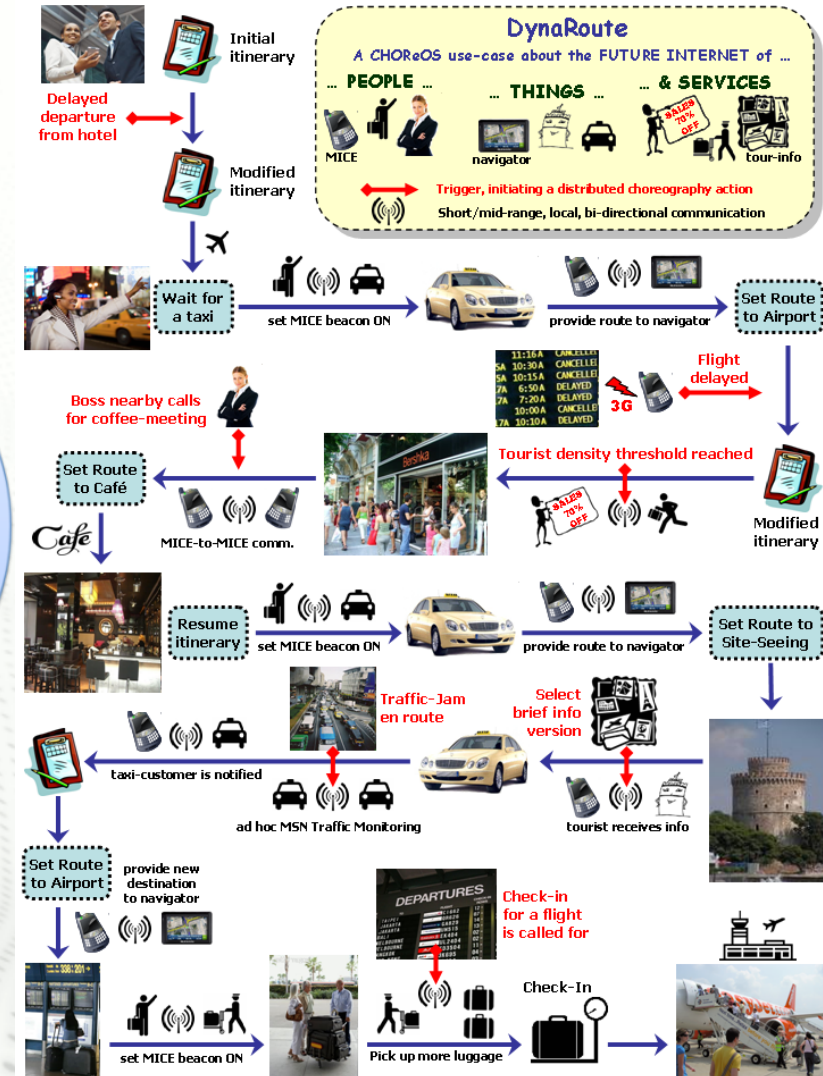
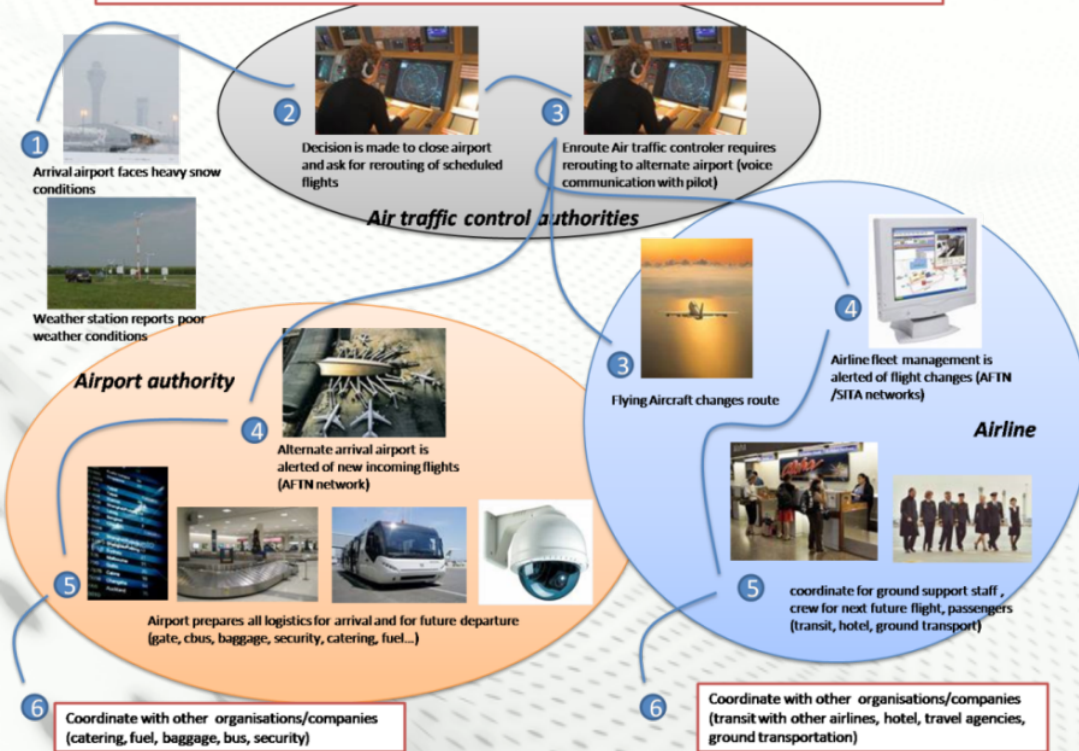
- **Service compositions deal with order and interdependencies among services, the flow of information, transactions etc.**
- **Strategies: service orchestration / choreographies**
  - Service orchestration - centralized control
  - Service choreography - peer-to-peer collaboration without a central node of control



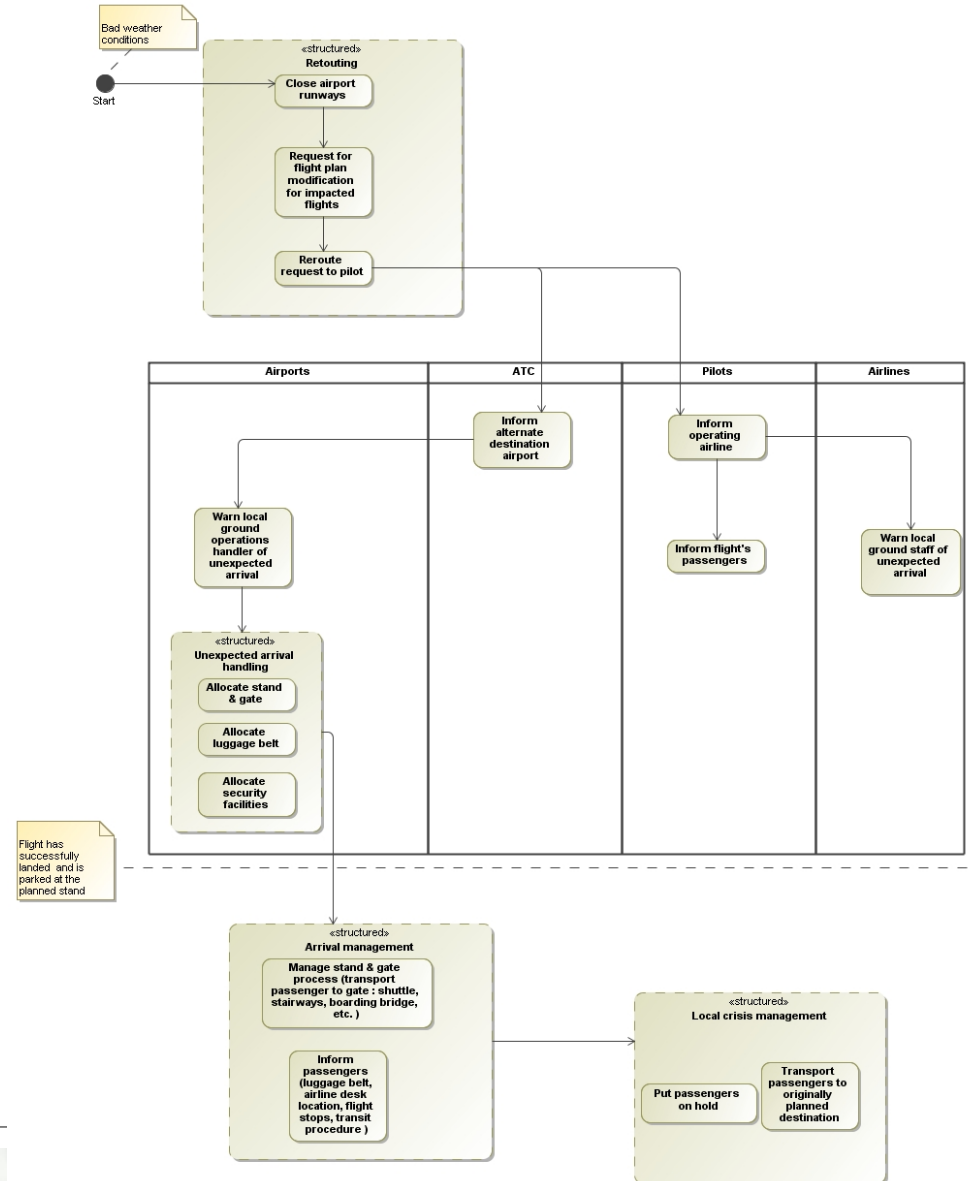
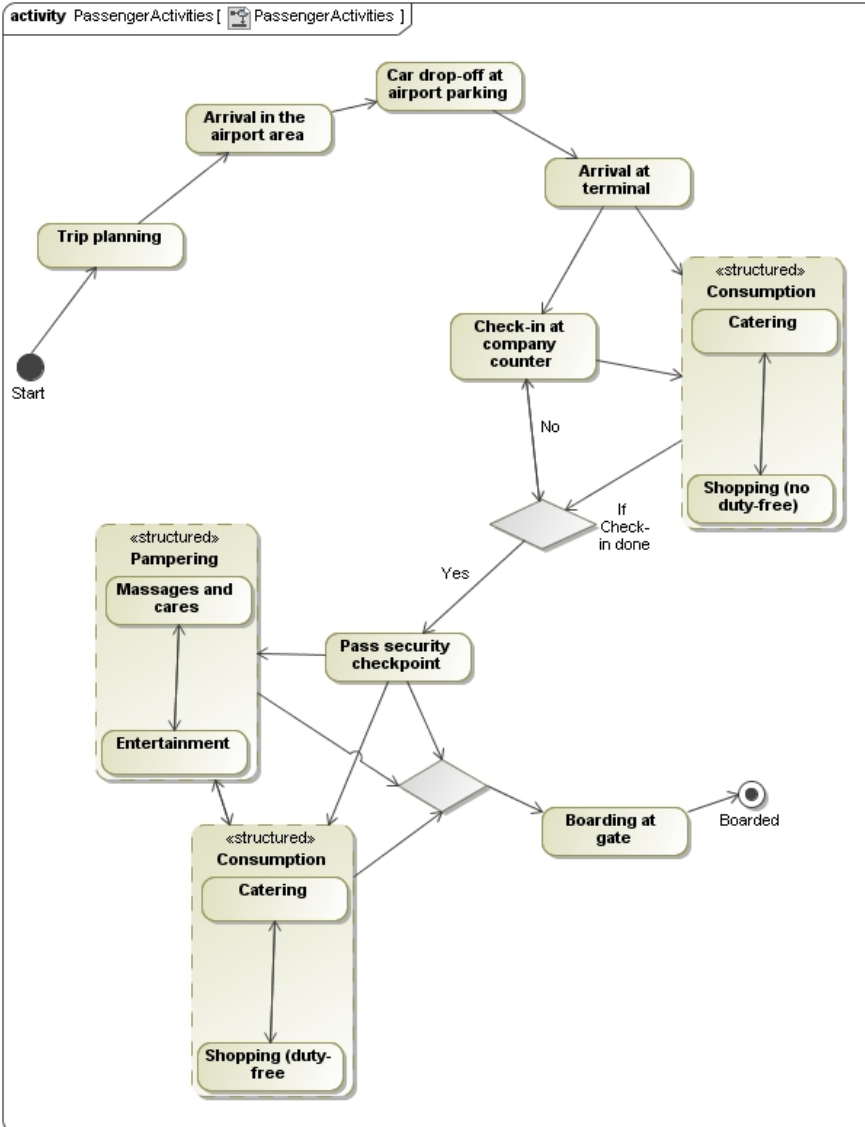
- **Web services adoption has grown during the last years; its usage has become a *de facto* standard for communication among high-level Internet systems.**
- **Orchestration is widely used for composing web services.**
  - But its centralized approach to composition has scalability and reliability problems.
- **Choreographies are Web services compositions organized in a decentralized, distributed manner, with no single point of failures.**
  - Global perspective
  - No central coordination node
  - Cross-enterprise business
  - Roles

# Scenarios for service choreographies

Scenario: Aircraft rerouting due to bad weather conditions – Coordinate consequences for all actors

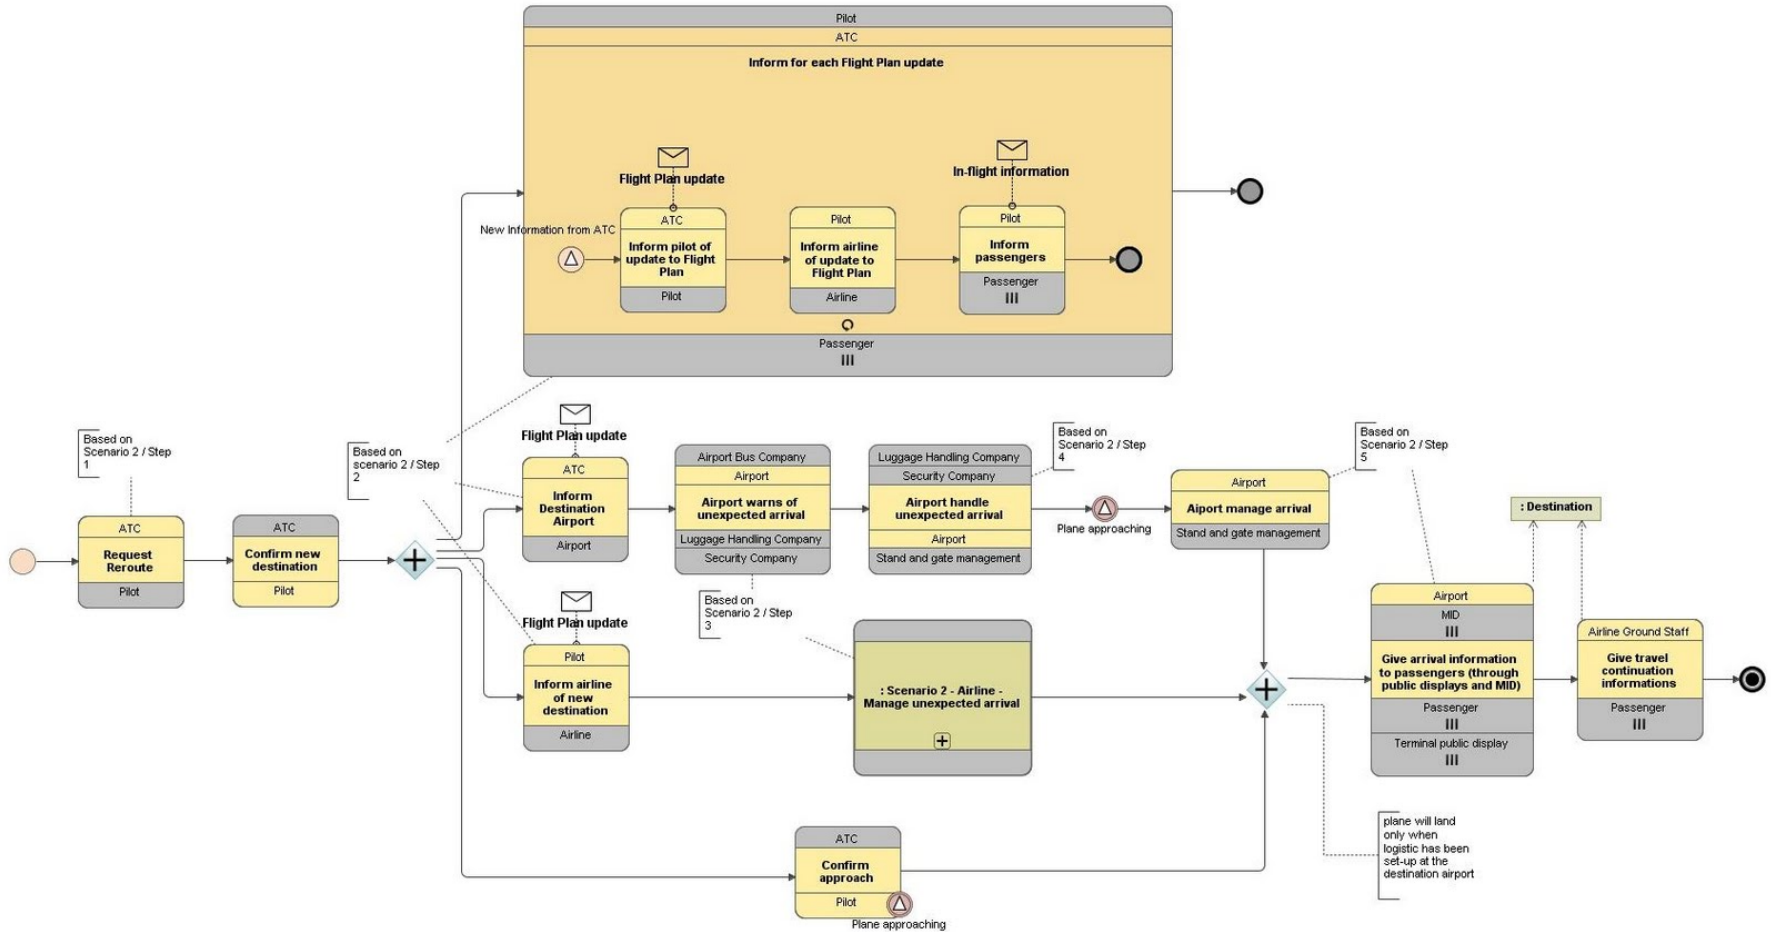






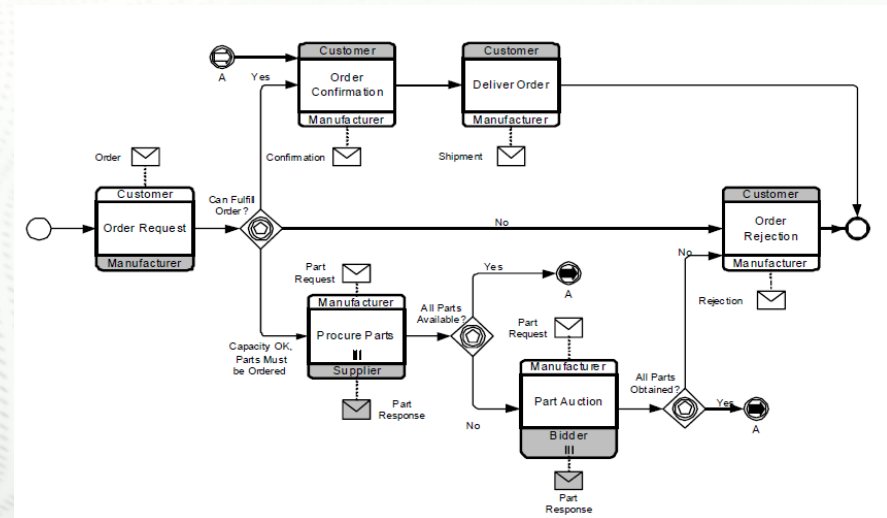


# BPMN2 model



# Scenarios for service choreographies

- Complex workflow involving multiple enterprises
- Interaction among several distributed orchestrations
- Fault tolerance
- Data intensive computing [Baker et al., 2009]



A. Barker, P. Besana, D. Robertson, and J. B. Weissman, "The Benefits of Service Choreography for Data-Intensive Computing," in *Proceedings of the 7th international workshop on Challenges of large applications in distributed environments*, 2009, p. 1--10.

- Choreographies have a good potential to deliver higher fault tolerance, adaptability, configurability, global optimization, and freedom to grow.
- This was identified some time ago (early 2000s) but all standardization initiatives failed up to now (WSCI, WS-CDL). (2011: OMG BPMN2)
- Many companies currently use ad hoc approaches to choreographies developed in-house.
- However, as current ad hoc choreographies get larger and more intricate, they can easily become unmanageable.
- The CHOReOS and BAILE projects will tackle these issues.



- **How to support the**
  - Coordination of an ultra large number of services widely distributed and moving on a highly heterogeneous and changing network, composed of heterogeneous devices (from tiny scale sensors/actuators to infrastructure servers).
  - Execution of services with varying load in dynamically evolving applications.
- **Our approach: The use of a middleware to enable scalable service provisioning based on Grid and Cloud computing for service choreographies**

- **Scalable Web Service Choreographies for the Future Internet**

- European Commission funding:

- CHOReOS – <http://www.choreos.eu>



- HP Brasil funding:

- BAILE – <http://ccsl.ime.usp.br/baile>



- **Some research lines:**

1. Development of a cloud-enabled middleware for service choreographies
2. Verification & validation of choreographies
3. Choreographies topology analysis
4. Performance analysis for large scale choreographies

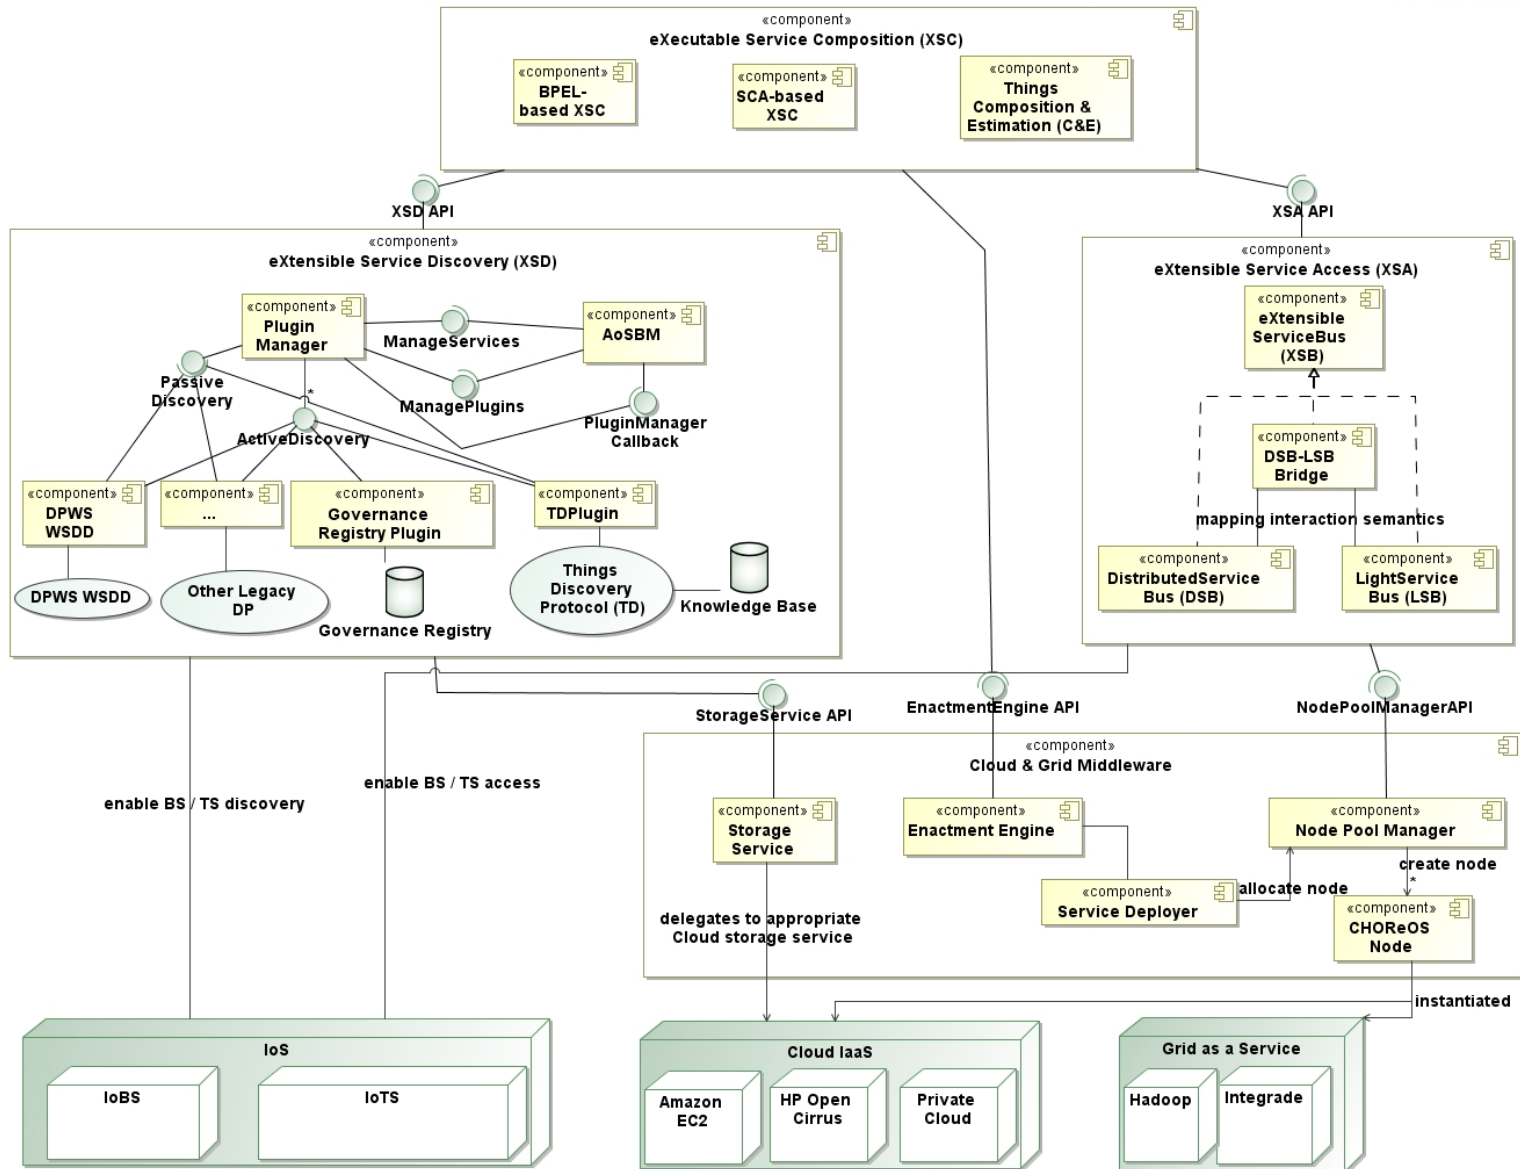


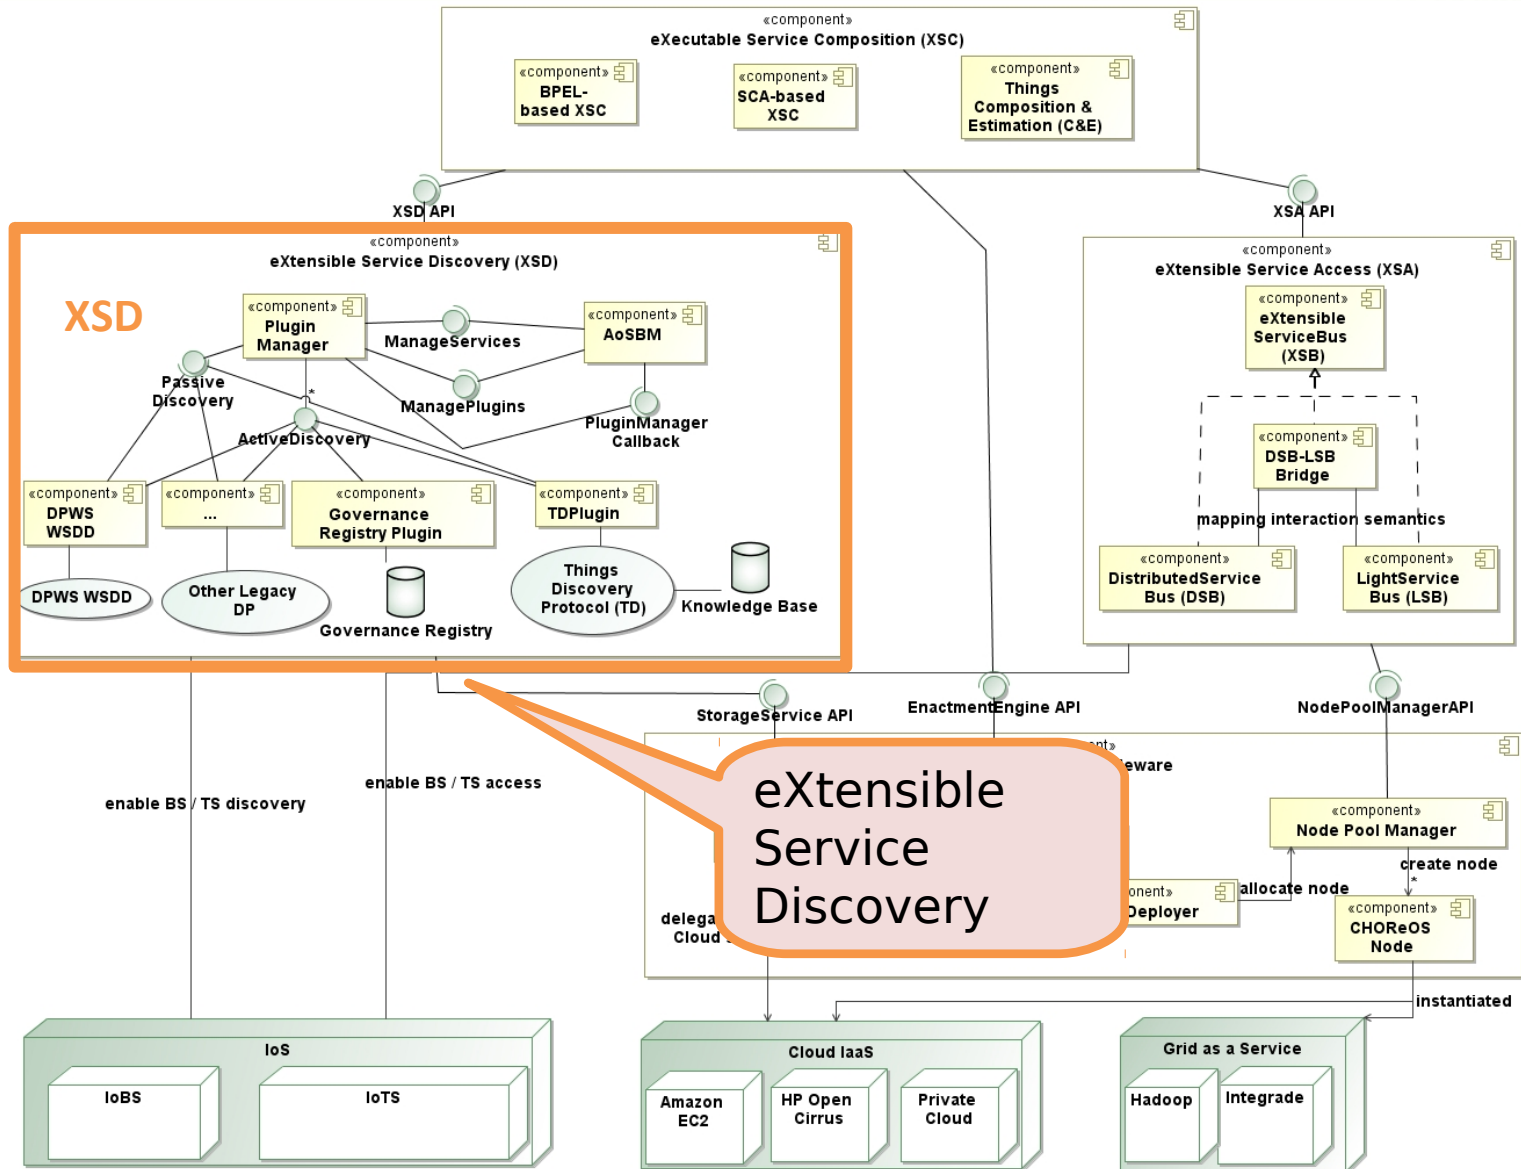
# **(1) Development of a cloud-enabled middleware for service choreographies**

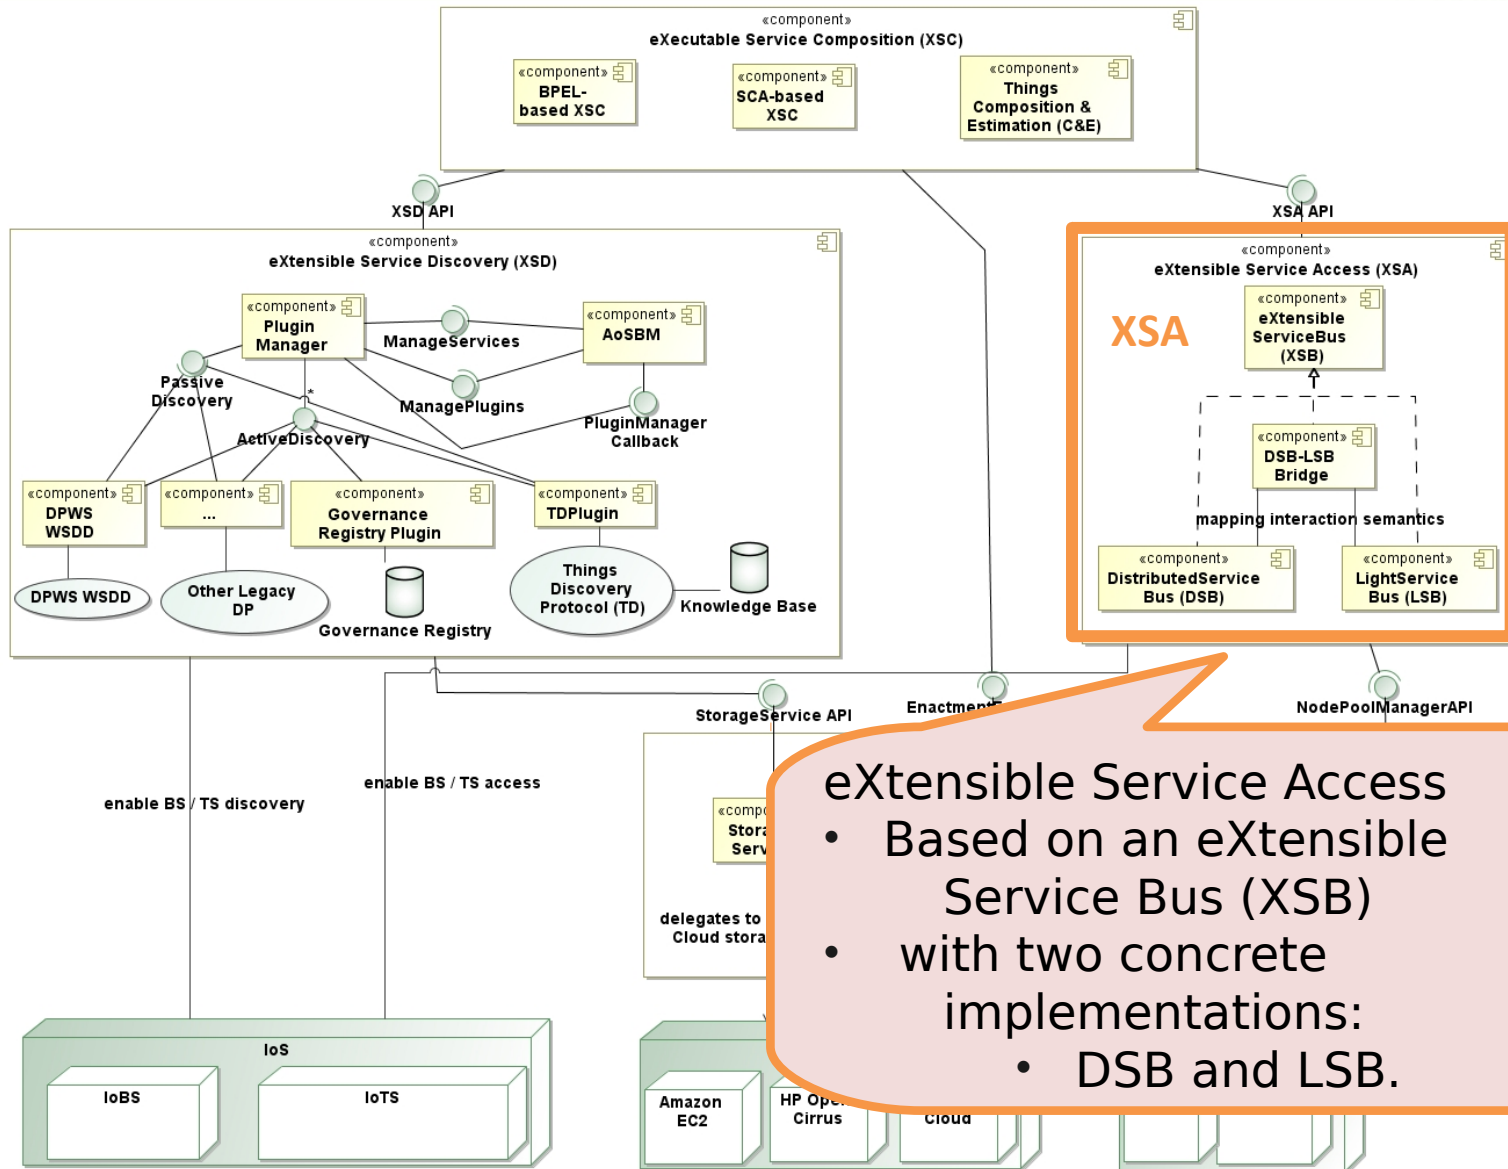


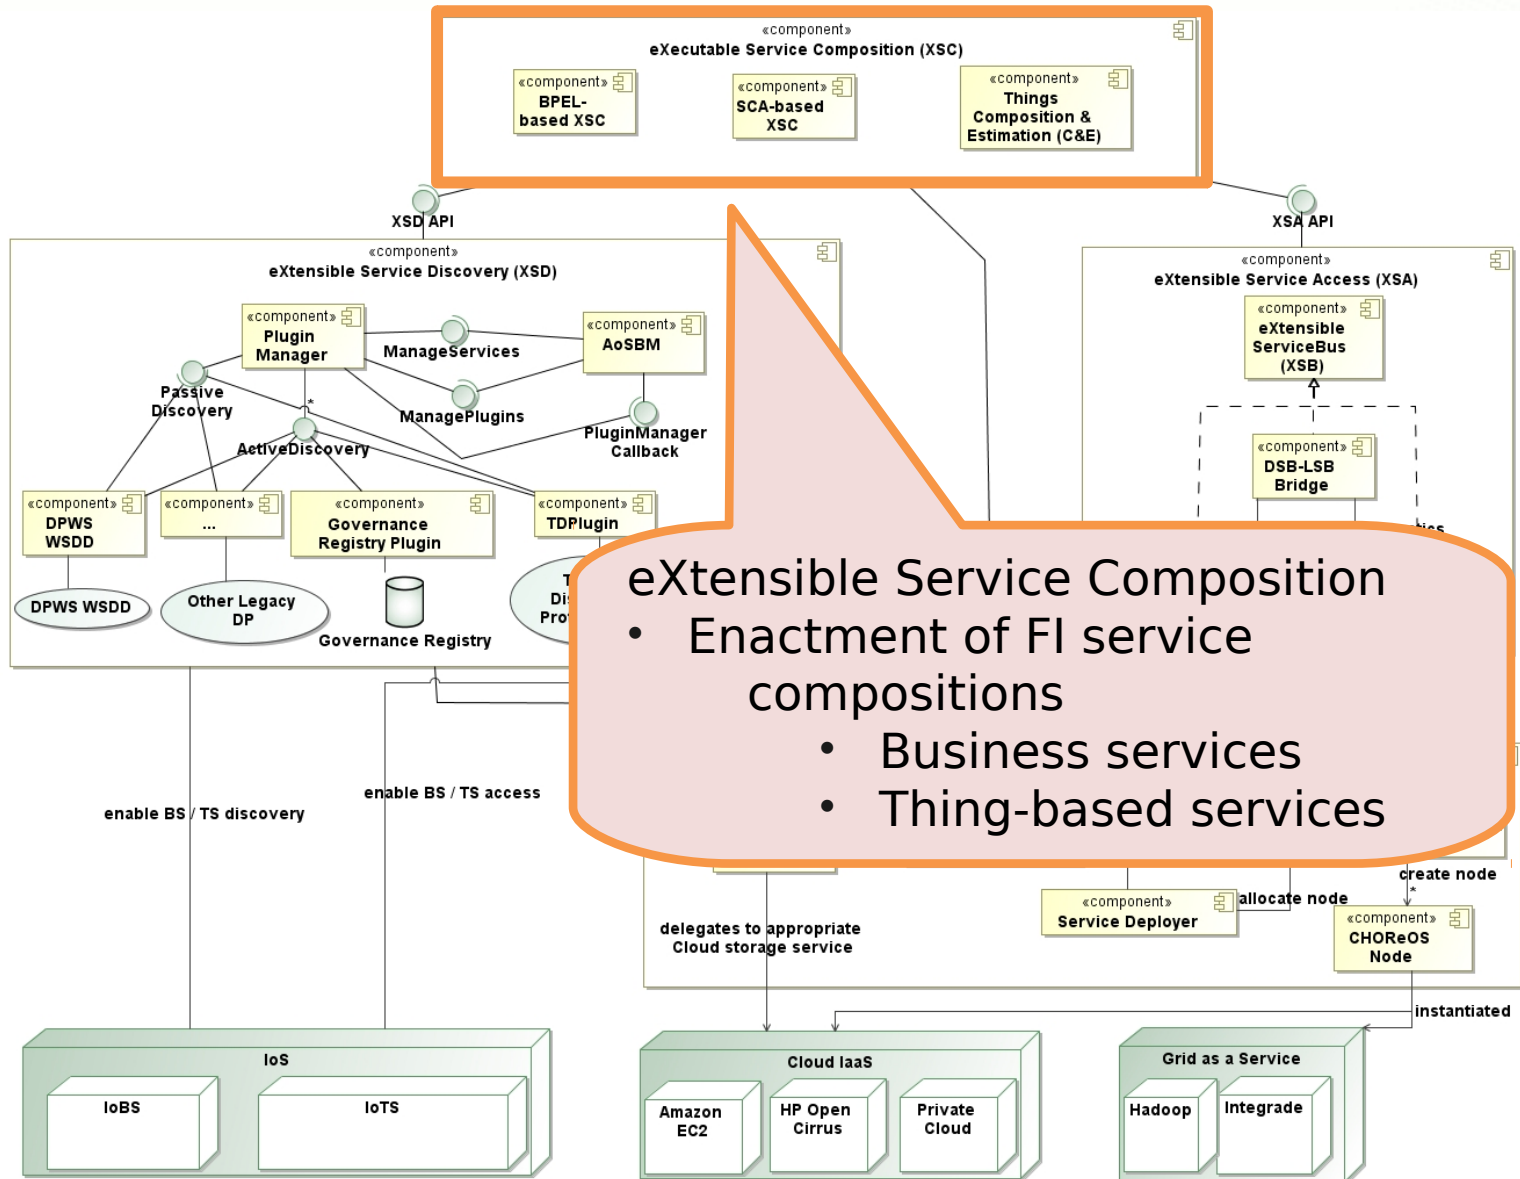
- **1 - Compile a choreography**
  - Given a BPMN2 Choreography description and the description of the web services generate the executable bytecode for this choreography.
- **2 - Enactment of a Choreography**
  - Given a choreography executable bytecode, enacts that choreography by (1) creating the required web services, (2) deploying the web services whose executable code is available, (3) assigning web services to the roles of the choreography, (4) enabling the choreography to receive messages.
- **3 - Instantiation of Choreographies in a Cloud Infrastructure**
  - Instantiate the choreography middleware and the nodes of a choreography in virtual machines of a given cloud infrastructure, e.g., Amazon EC2, Open Cirrus, or OpenNebula. This must be performed in a way that maximizes the QoS as perceived by the final users and minimizes resource consumption.

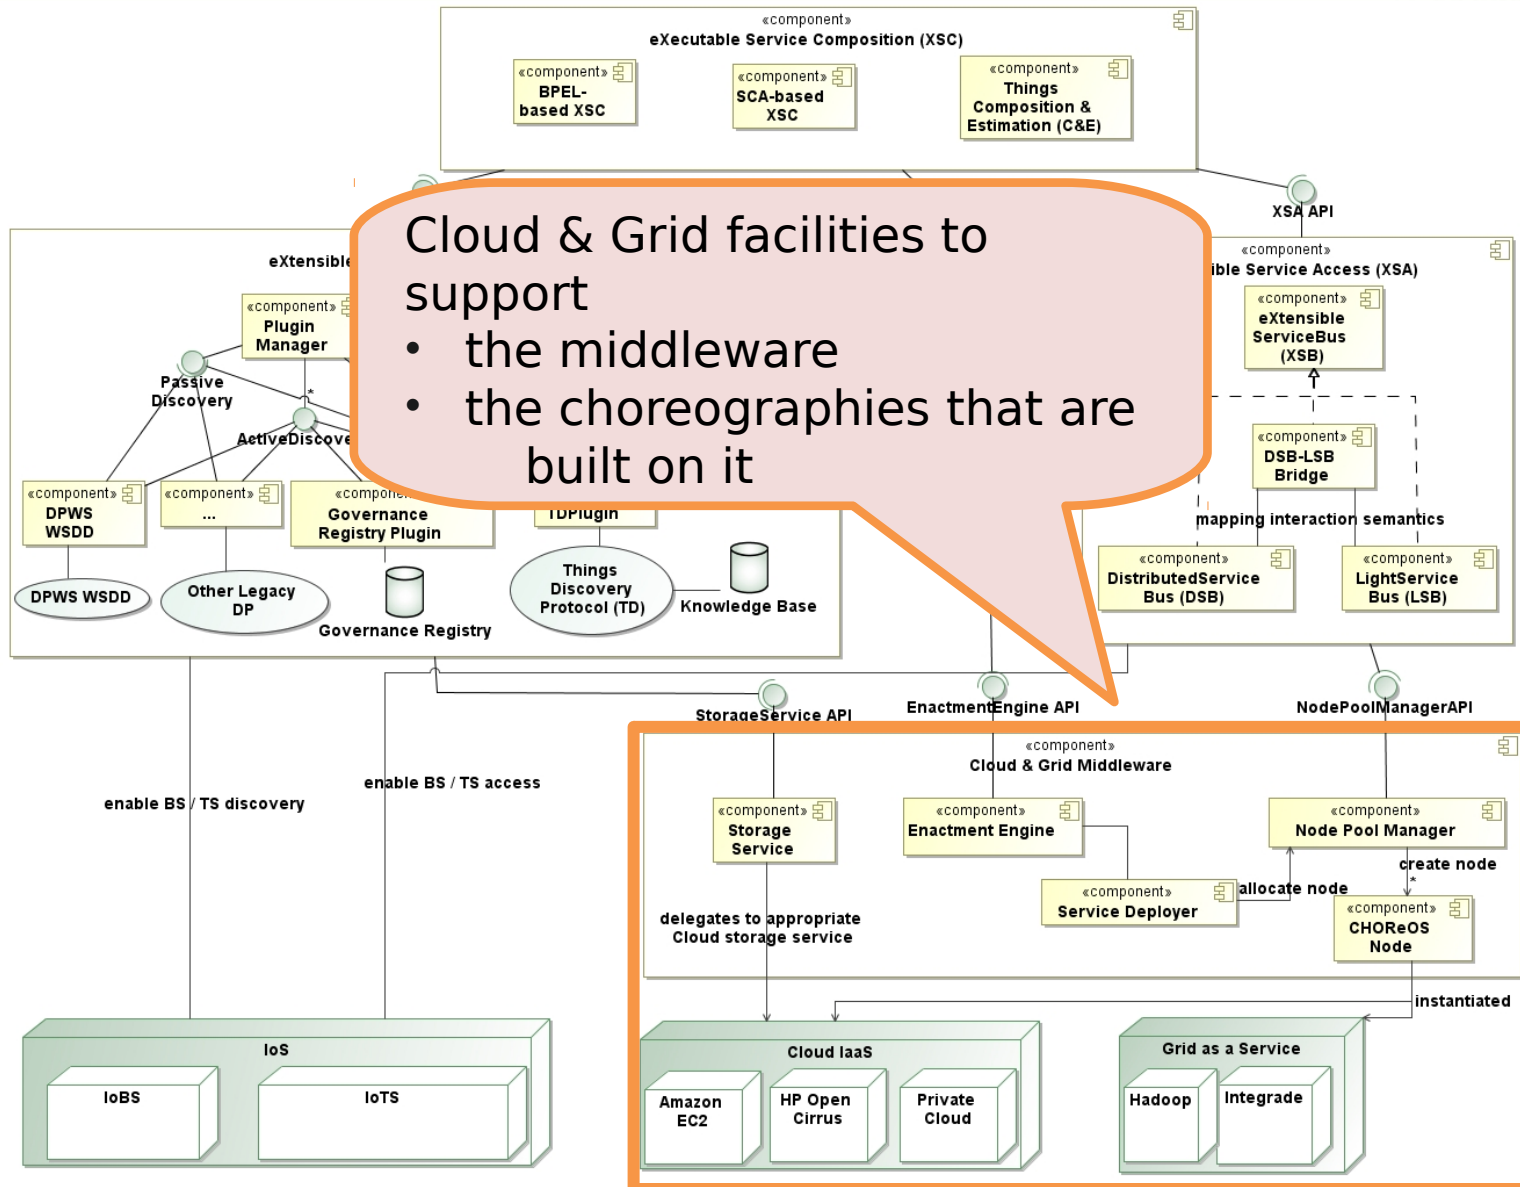
- Our goal in the BAILE project is to study how scalable a choreography middleware can get.
- We aim at developing choreographies with
  - tens to thousands of services
  - running on hundreds to thousands of nodes
  - servicing thousands to millions of users
  - running computationally/data intensive tasks
- This involves efforts in
  - Software Architecture and Engineering
  - Investigation of novel methods for creating, managing, and processing choreographies









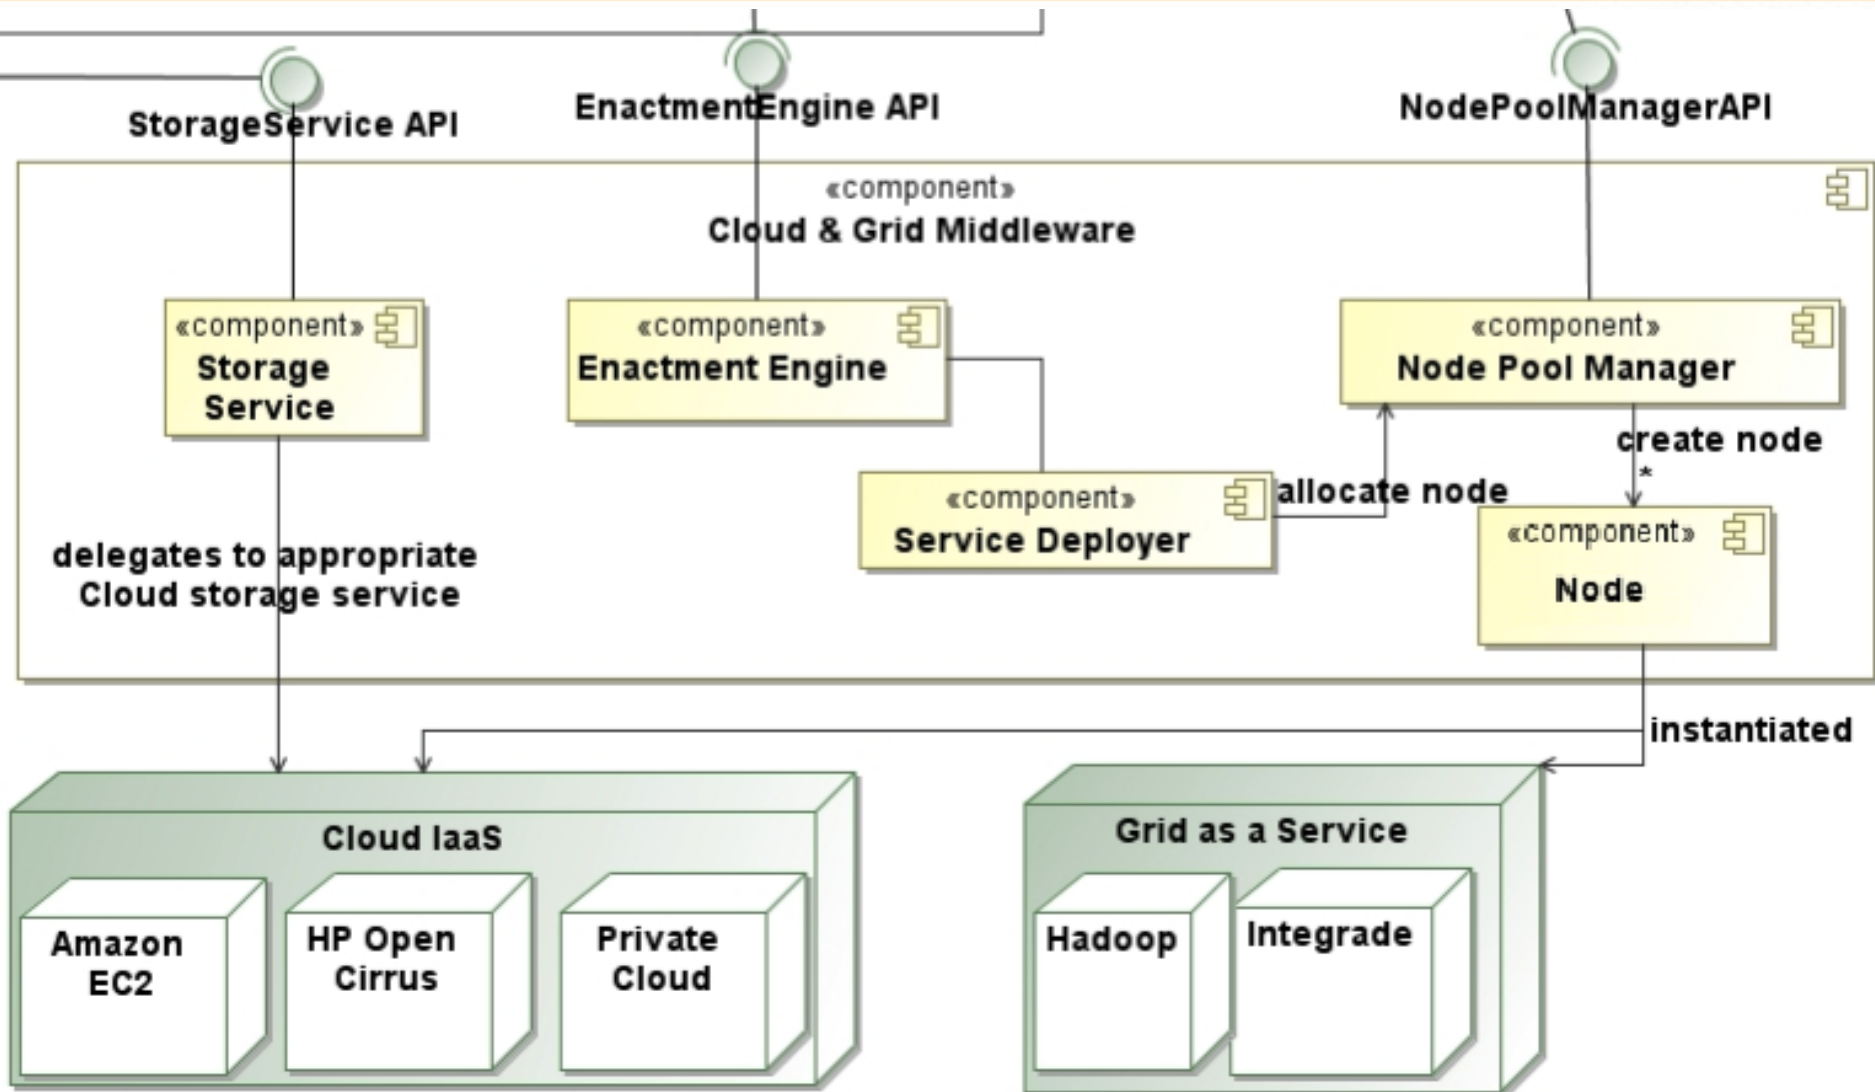


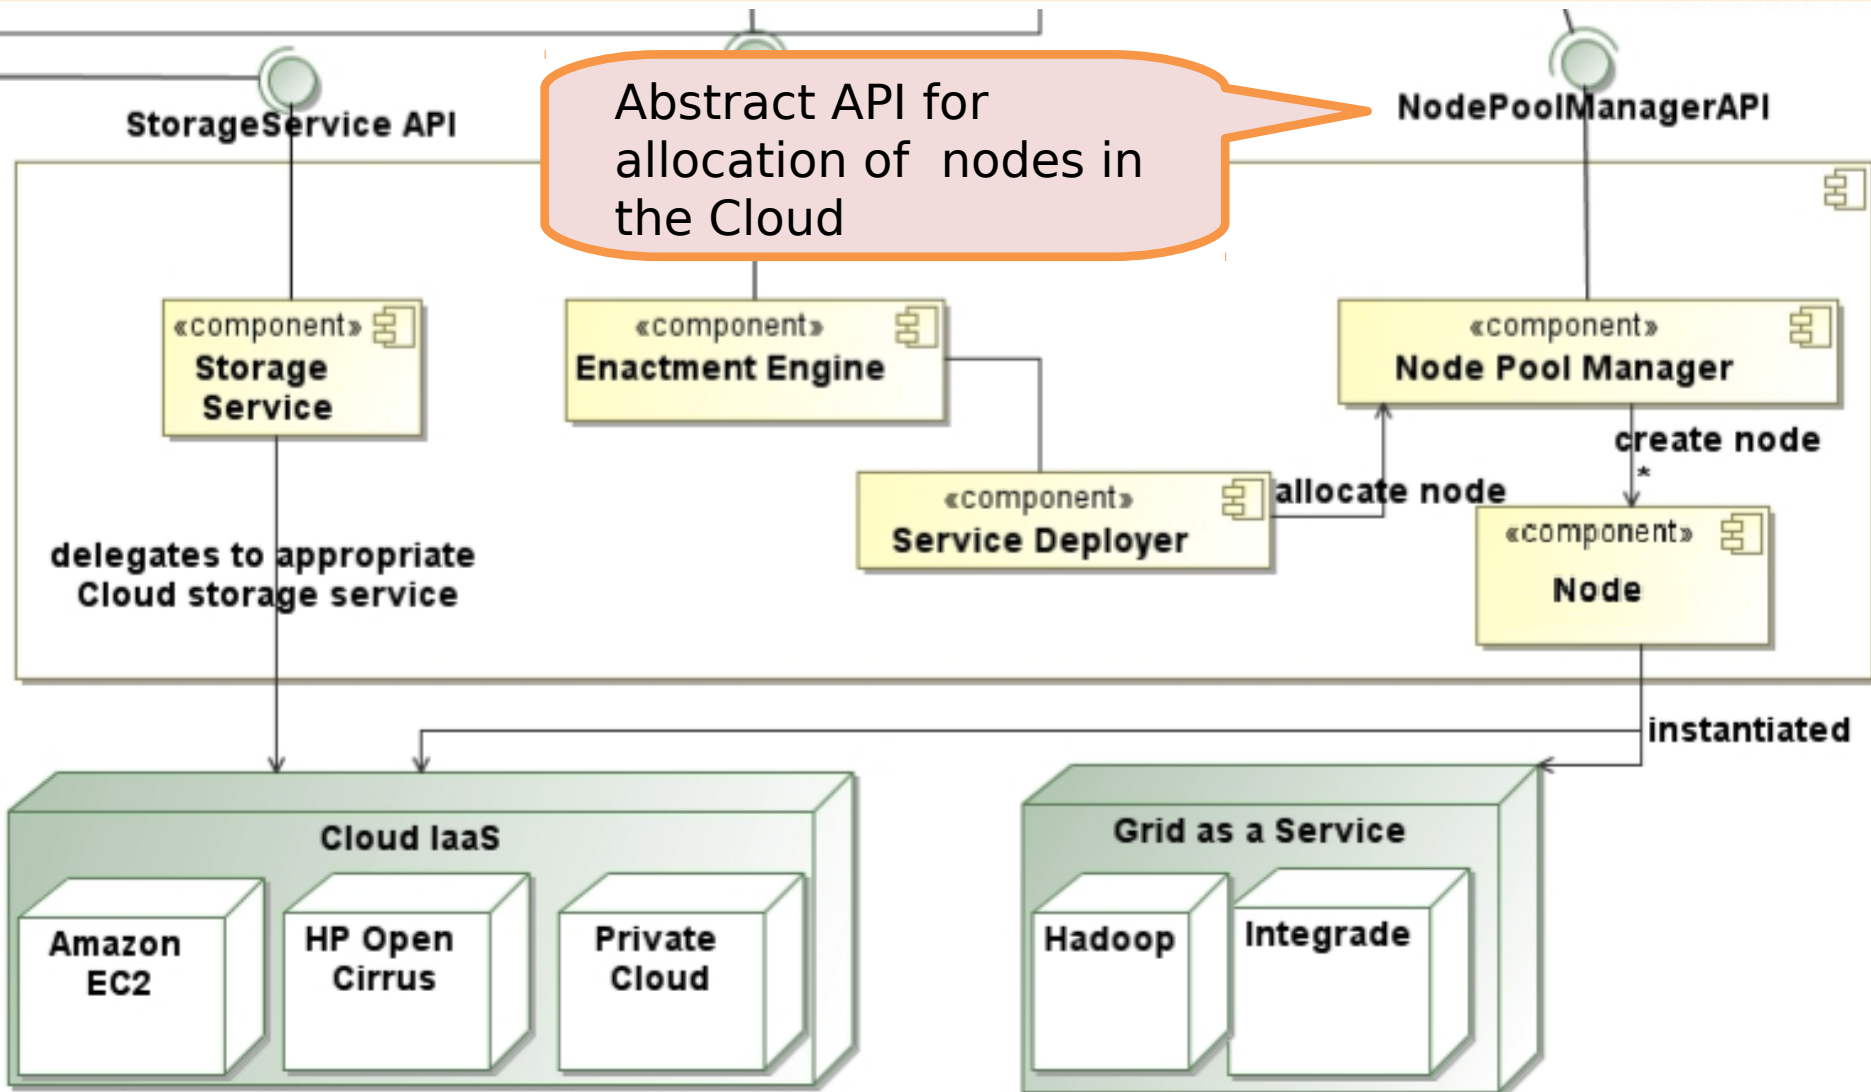
- **Cloud**: to execute large quantities of services, replicate services, and balance the load from millions of requests from thousands of users
- **Grid**: for CPU-intensive or data-intensive applications

## Two-way contribution:

1. Cloud will provide a means to support scalable choreographies
2. Higher-level abstraction for the execution of complex, distributed, service compositions







StorageService API

EnactmentEngine API

NodePoolManagerAPI

«component»  
Cloud & Grid Middleware

«component»  
Node Pool Manager

«component»  
Node

allocate node

create node  
\*

instantiated

- Multiple infrastructure support:
- Public Clouds: Amazon EC2, GoGrid, ...
  - Private Clouds: OpenNebula, OpenStack, ...
  - Open Cirrus testbed

Cloud IaaS

Amazon  
EC2

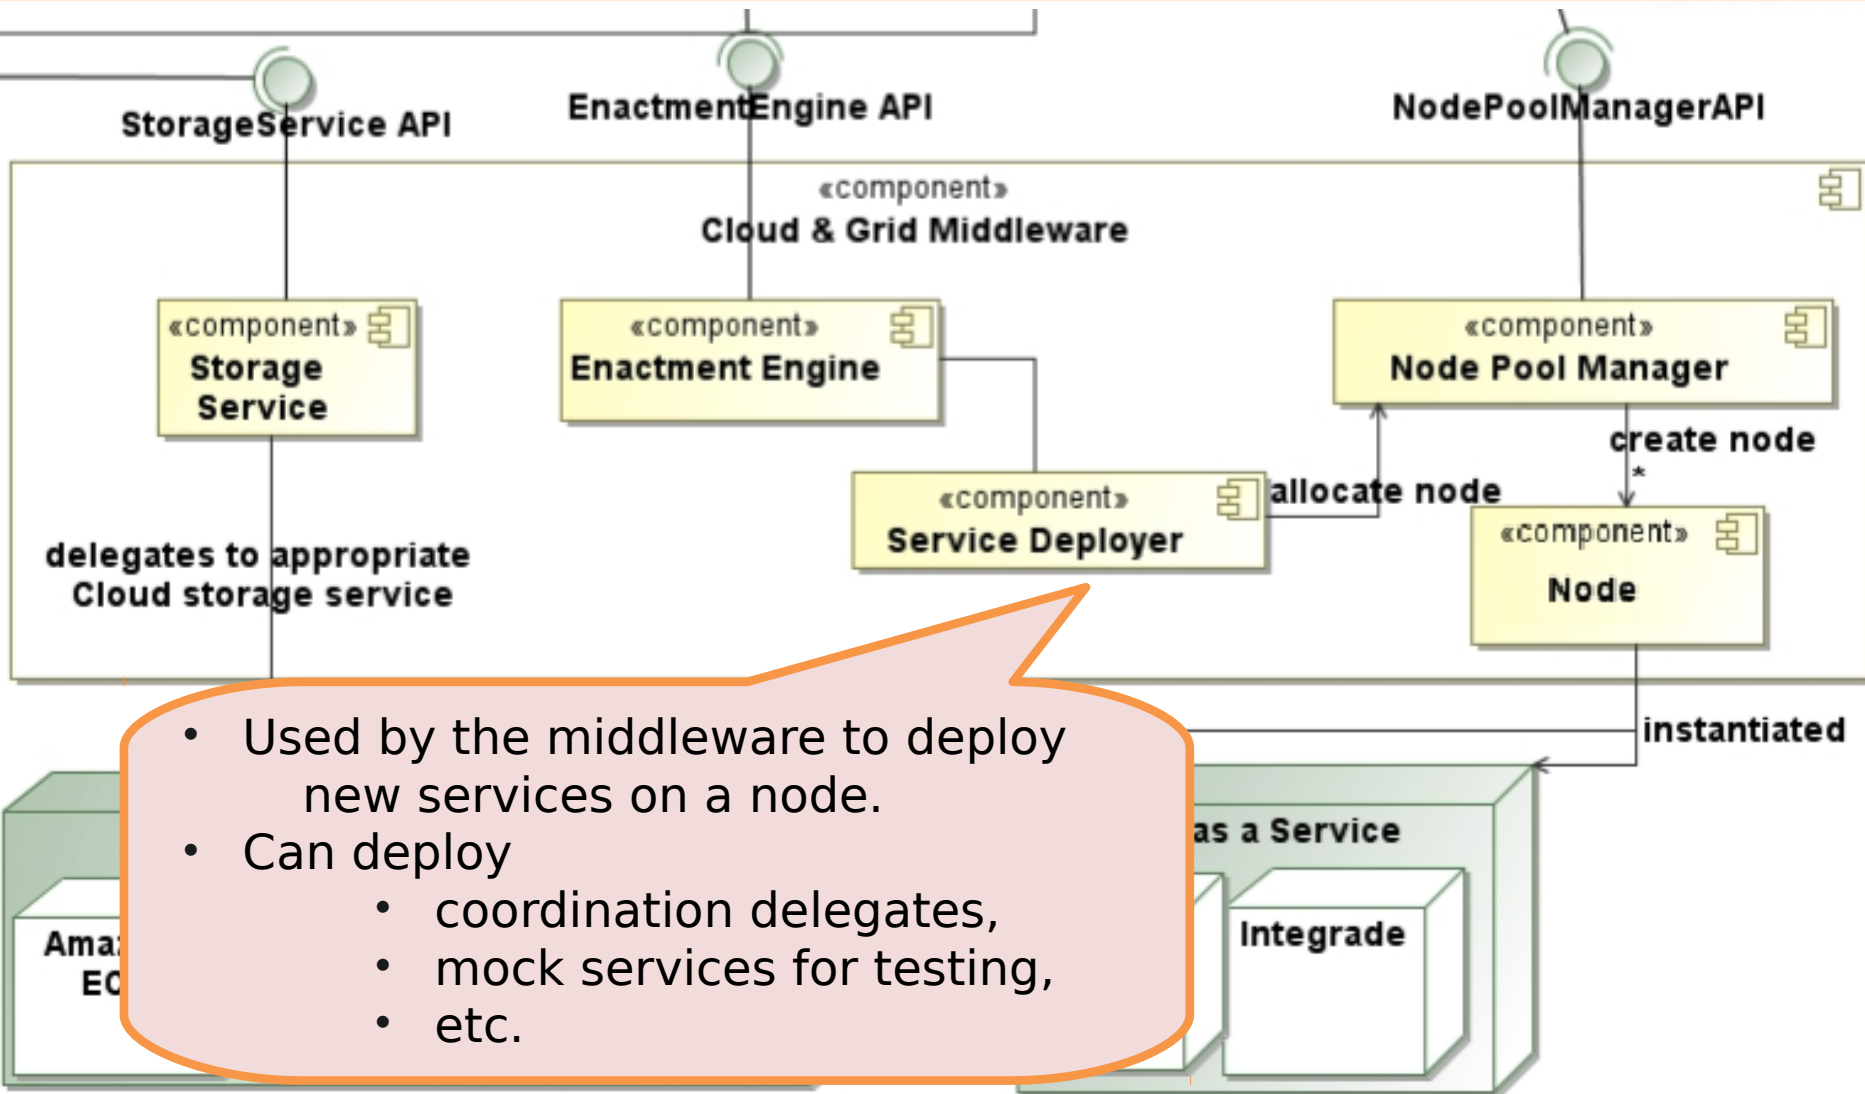
HP Open  
Cirrus

Private  
Cloud

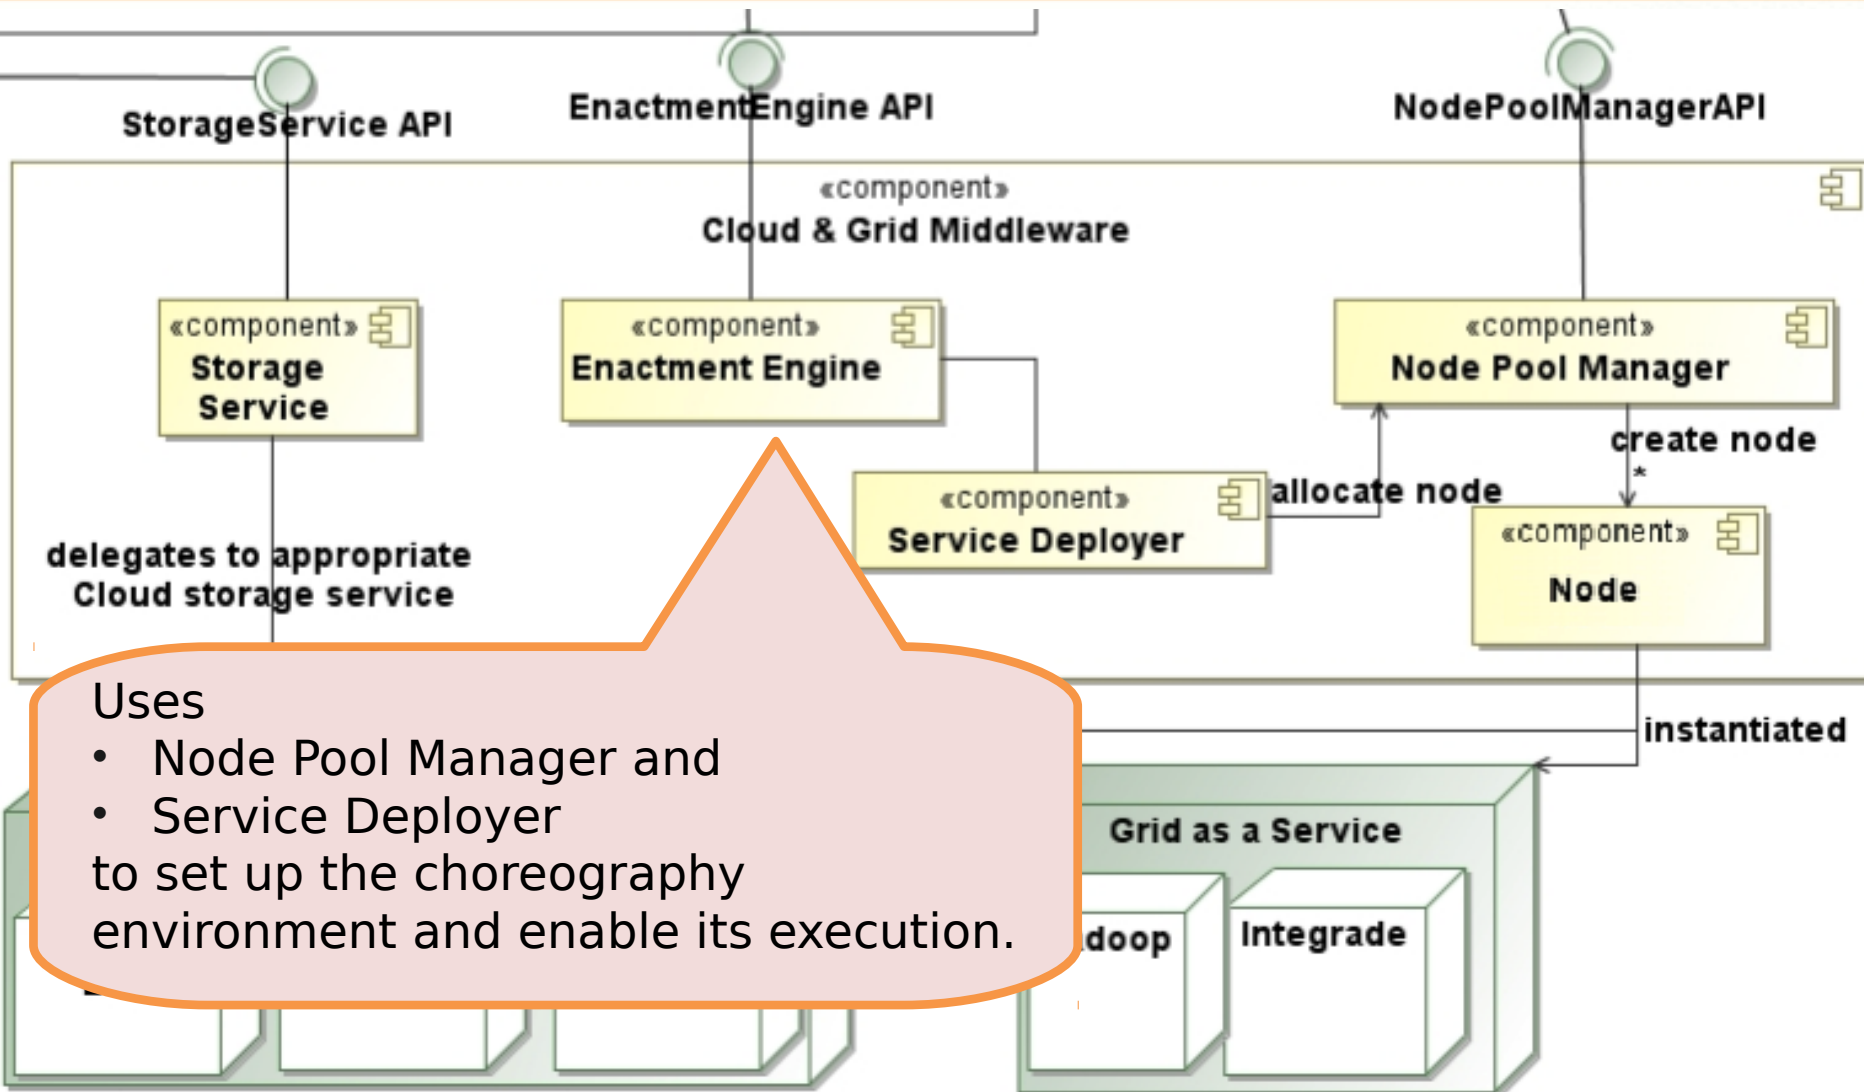
Grid as a Service

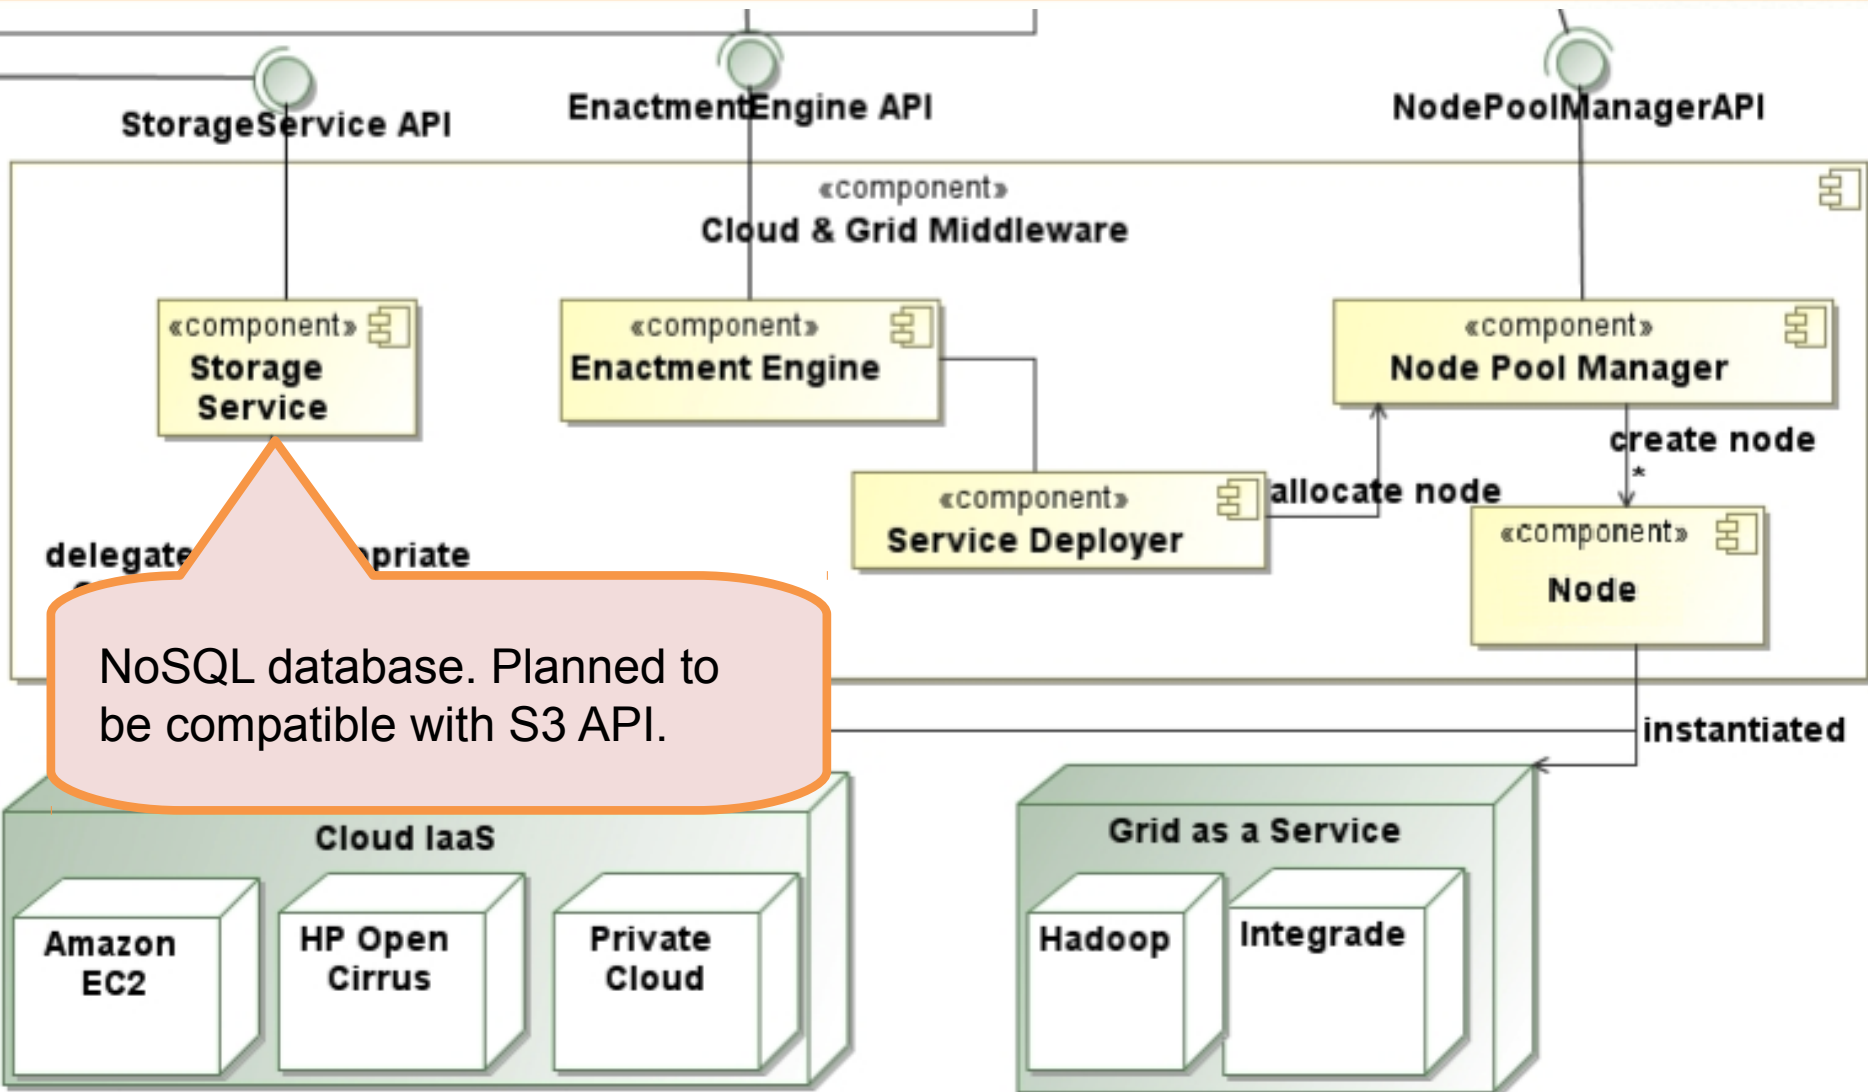
Hadoop

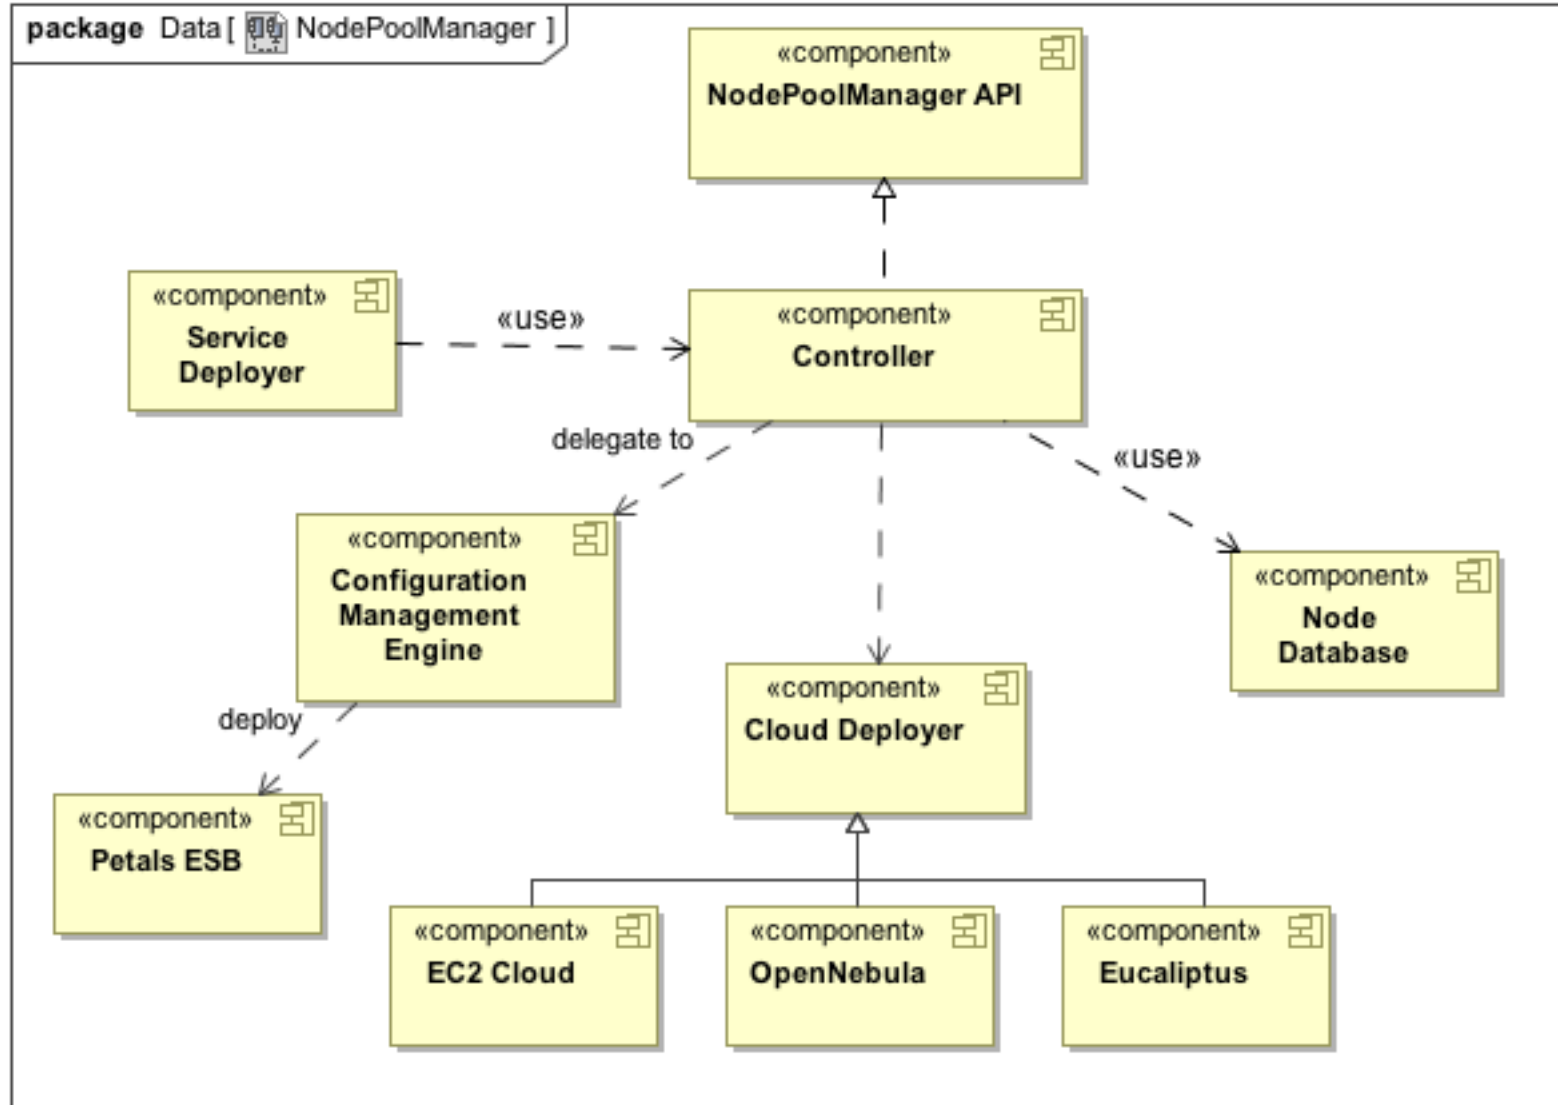
Integrate



- Used by the middleware to deploy new services on a node.
- Can deploy
  - coordination delegates,
  - mock services for testing,
  - etc.







## Implementation Status

- **Node Pool Manager implementation for Amazon EC2**
- **Automatic deployment of DSB on Amazon and Open Cirrus**





## Node Pool Manager Demo



# **(2) Verification & validation of choreographies**

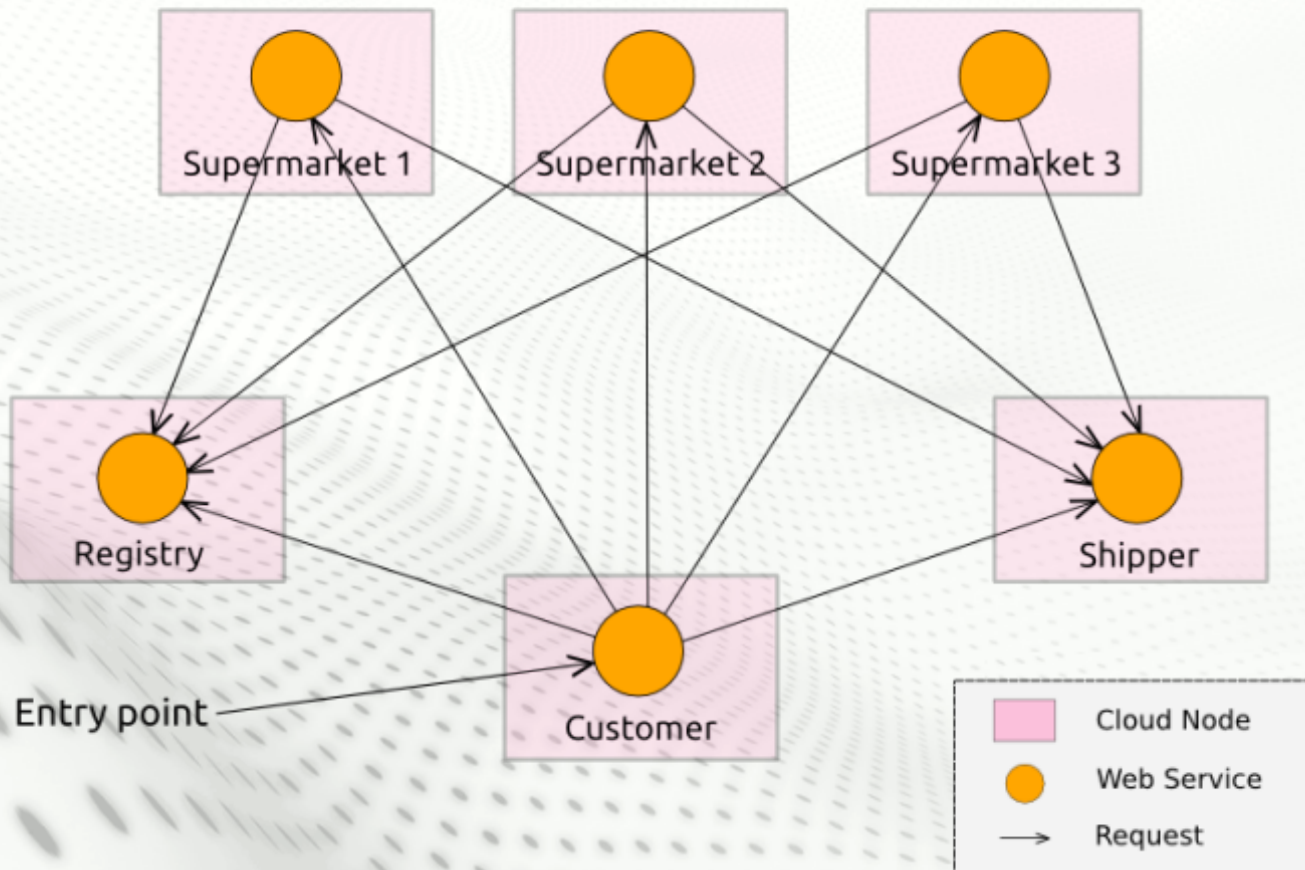
- **A framework and tool for automated testing of choreographies.**
- **With this, we intend to:**
  - abstract the choreography and its elements to simplify coding of automated tests;
  - implement a conversation monitoring service to support integration tests;
  - mock web services functionalities;
  - Invoke SOAP and REST services dynamically
  - Support scalability testing of web services and choreographies

**The example choreography implements a distributed shopping service:**

- 1. A customer provides to the choreography a shopping list;**
- 2. The price of each item list is queried in different supermarket services to find which one has the lowest price;**
- 3. The choreography returns to the customer the total cheapest price of its list and provides features for purchasing and delivering the items.**

# Sample Choreography

The choreography is deployed on 6 nodes of Open Cirrus.



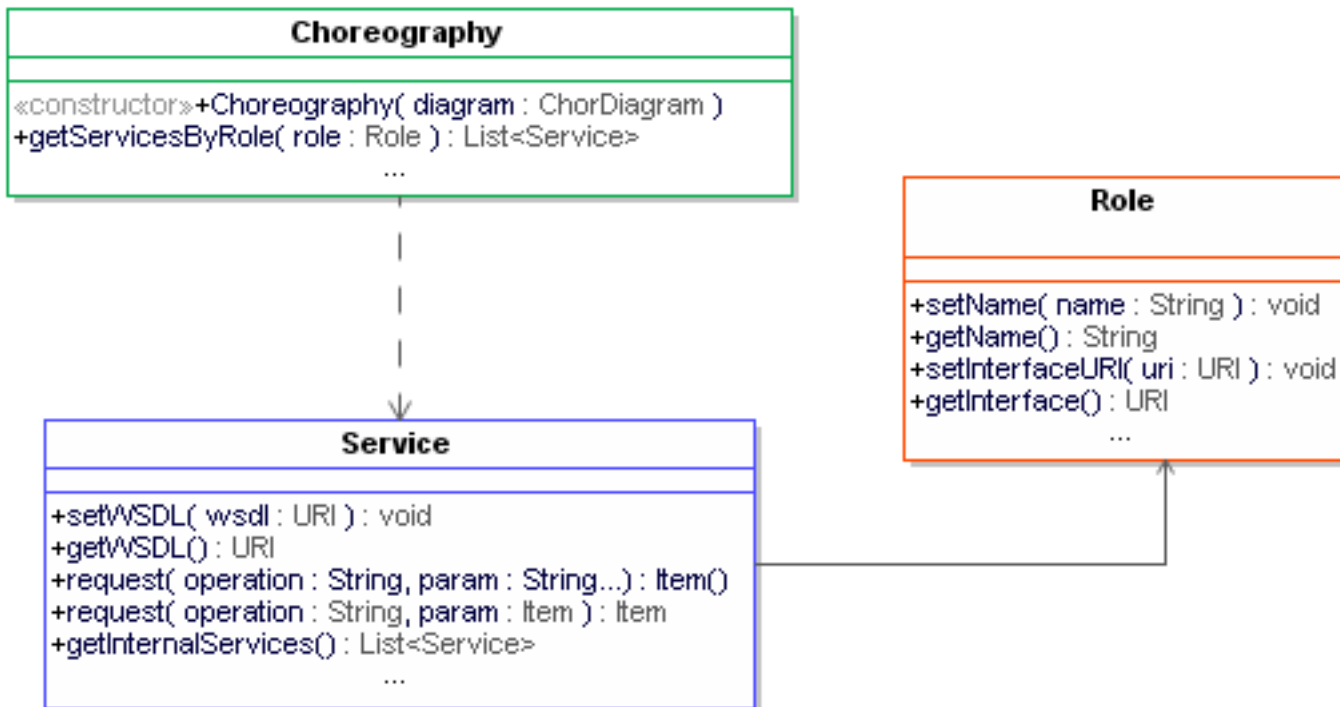
## Assessment

- 1. “Manually” enacted a choreography (with scripts).**
- 2. With the help of *Rehearsal* TDD framework, we perform 5 tests to verify if the choreography is running correctly (1 compliance and 4 acceptance tests).**
- 3. In the future, we will also be able to perform scalability tests to verify if the choreography continues to work correctly, meeting QoS requirements, as its scale grows.**



## Abstraction of Choreography

Interact with the choreography elements (roles, services, messages) through Java objects.



## Test set up

```
// Choreography creation
```

```
Choreography futureMarket = new Choreography();
```

```
// role creation: new Role("<role name>", "<role wsdl">");
```

```
Role supermarket = new Role("supermarket", "supermarket?wsdl");
```

```
// service creation
```

```
Service store1 = new Service();
```

```
store1.setWSDL("store1?wsdl");
```

```
store1.addRole(supermarket);
```

```
// adding the services and roles into the choreography
```

```
futureMarket.addRole(supermarket);
```

```
futureMarket.addService(store1);
```



## Service Role Compliance Testing

Apply a generic test suite to verify if the participants are playing a role properly.

**SMRoleTest.class** => JUnit test cases that must be applied in all participants that want to play the supermarket Role.

**Compliance test** => through the **assertRole** feature, the tests are applied:

```
@Test
public void servicesMustBeCompliantWithTheSupermarketRole() {
    List<Service> supermarkets = futureMarket
        .getServicesForRole("supermarket");

    for (Service service : supermarkets)
        assertRole(supermarket, service, SMRoleTest.class);
}
```

## Acceptance Testing

Test cases for assessing the choreography workflow from the end-user perspectives:

### Register supermarket

receives a supermarket endpoint and register it into the registry service



### Get price of product list

receives a customer product list and returns the cheapest total price and an order id for purchasing



### Purchase

receives a customer order id and returns the shipper name



### Get delivery data

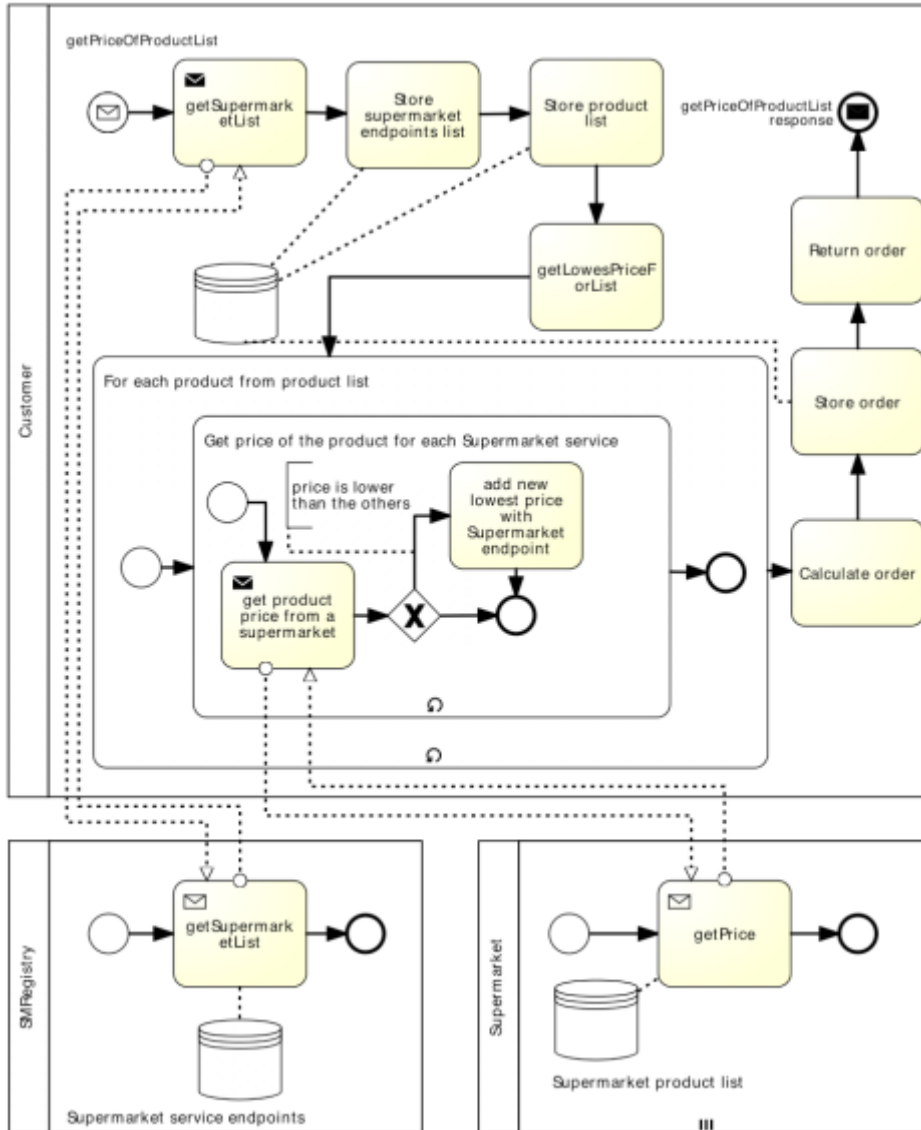
receives a customer order id and a shipper name and returns the delivery data information (day and time)



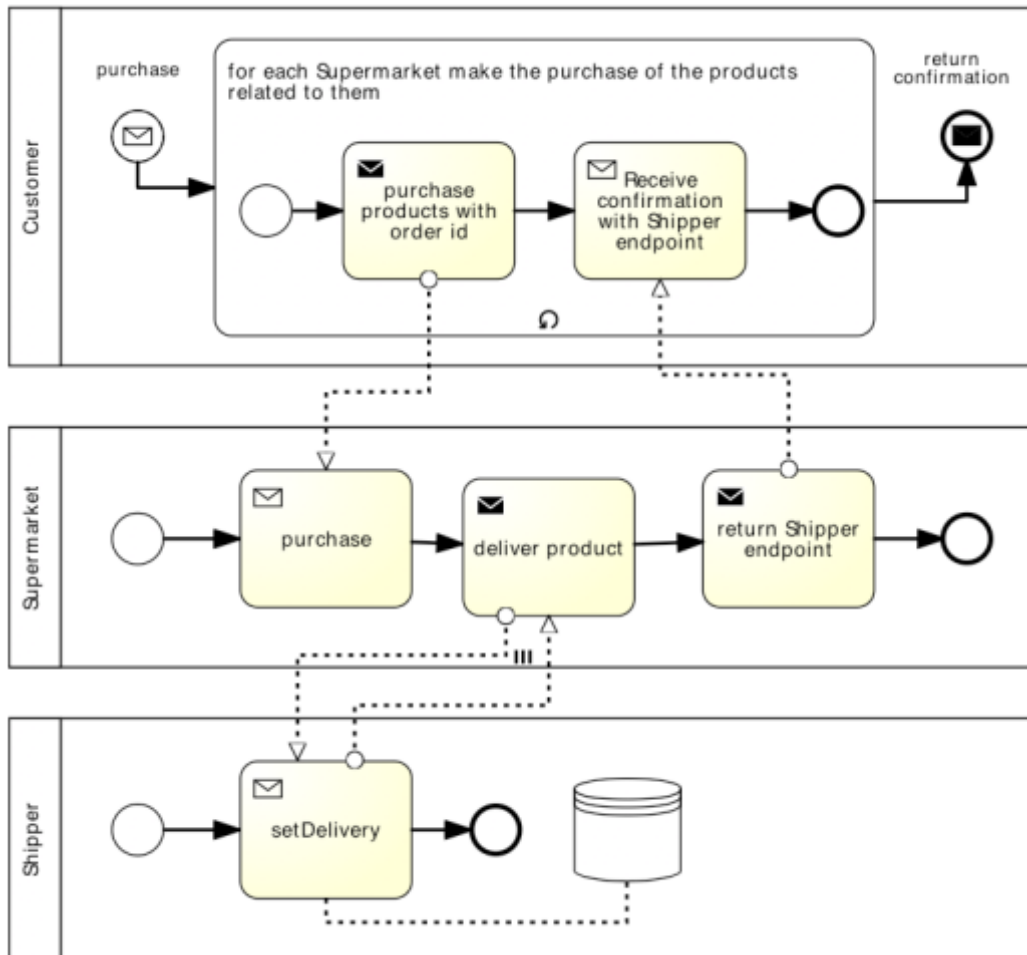
## Choreography Testing Demo



# Sample Choreography



**BPMN2 Model describing the *getPriceOfProductList* operation of the Future Market Choreography**



**BPMN2 Model describing the *purchase* operation of the Future Market Choreography**


- **Using the Node Pool Manager, cloud nodes can be manipulated and the DSB can be deployed easily**

## Next Steps

- **Use a different cloud provider**
- **Development of the Service Deployer and Enactment Engine**
- **Scalability testing**
- **Choreography QoS/SLA management**



# (3) Choreographies topology analysis



- **Services choreography should be designed taking into account the stability and interdependencies of its parts**
- **Some metrics (mainly from SNA - Social Network Analysis)**
  - Degree centrality (number of links incident upon a node)

$$C_{D-}(v) = \frac{\text{deg}^-(v)}{n-1} \quad C_{D+}(v) = \frac{\text{deg}^+(v)}{n-1}$$

$$C_D(G) = \frac{\sum_{i=1}^{|V|} [C_D(v^*) - C_D(v_i)]}{(n-1)(n-2)}$$

- Eigenvector centrality (neighbors and their neighbors)
- Page Rank

---

**Algorithm 1:** Accelerated power method for eigenvector centrality calculation

---

**Input** : An adjacency matrix  $A_{i,j}$ , where  $A_{i,j} = 1$  if the  $i^{\text{th}}$  node is adjacent to the  $j^{\text{th}}$  node, and  $A_{i,j} = 0$  otherwise

**Output:** Eigenvector centrality value for all graph nodes

- 1 Set  $C_E(v_i) = 1$  for all  $i$ ;
  - 2 Compute  $C_E^*(v_i) = \sum_j A_{i,j} * C_E(v_j)$ ;
  - 3 Set  $\lambda$  equal to the square root of the sum of squares of each  $C_E^*(v_i)$ ;
  - 4 Set  $C_E(v_i) = C_E^*(v_i)/\lambda$  for all  $i$ ;
  - 5 Repeat lines 2 to 4 until  $\lambda$  stops changing;
- 

$$PR(u) = \sum_{v \in B_u} \frac{PR(v)}{L(v)}$$



- **Betweenness centrality**

- The chance of a node to be in a shortest paths between other 2 vertices  $C_B(v) = \left( \sum_{s \neq v \neq t \in V} \frac{\sigma_{st}(v)}{\sigma_{st}} \right)$

- **Closeness Centrality**

- The mean shortest path between  $v$  and all other reachable vertices

$$C_C(v) = \frac{|J_v|/(n-1)}{\sum_{t \in J_v} d_G(v, t)/|J_v|}$$

## • Overall Stability

- Average propagation level of changes performed in the nodes of a network

$$AvgImpact(G) = \left[ \sum_{v \in V} \gamma(v) \right] / n^2$$


$$Stability(G) = 1 - AvgImpact(G)$$

Given a set of services that realize the choreography, choose the ones that maximize choreography overall stability



## • Butterfly services

- Likely to cause large ripple-effects
- Core (prominent) services of the choreography
  - Identify, monitor, backup
  - Unit-test and integration tests are advisable

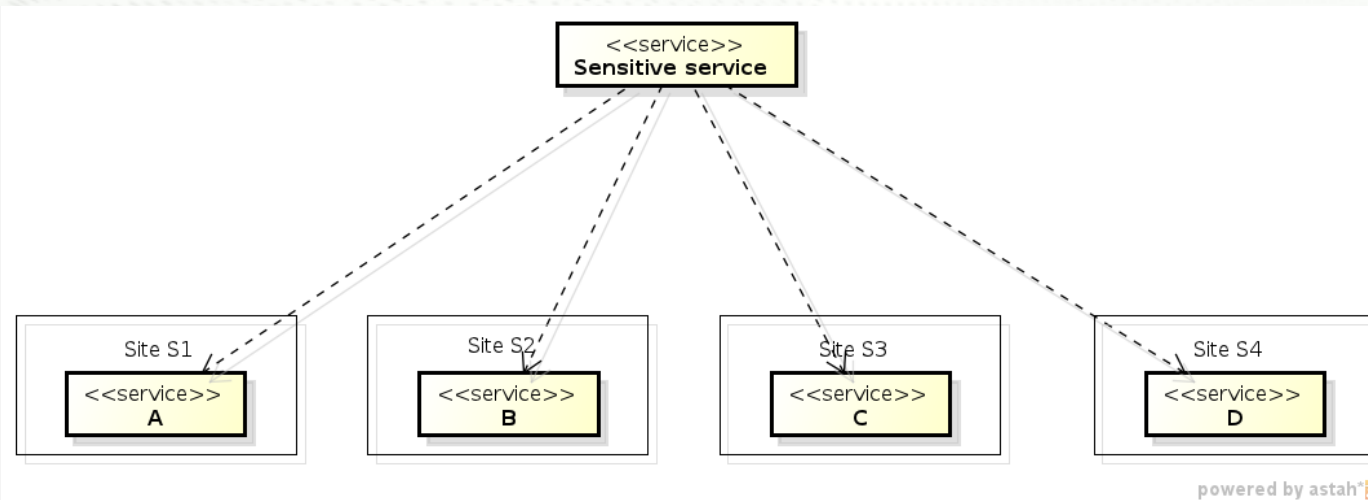


powered by astah



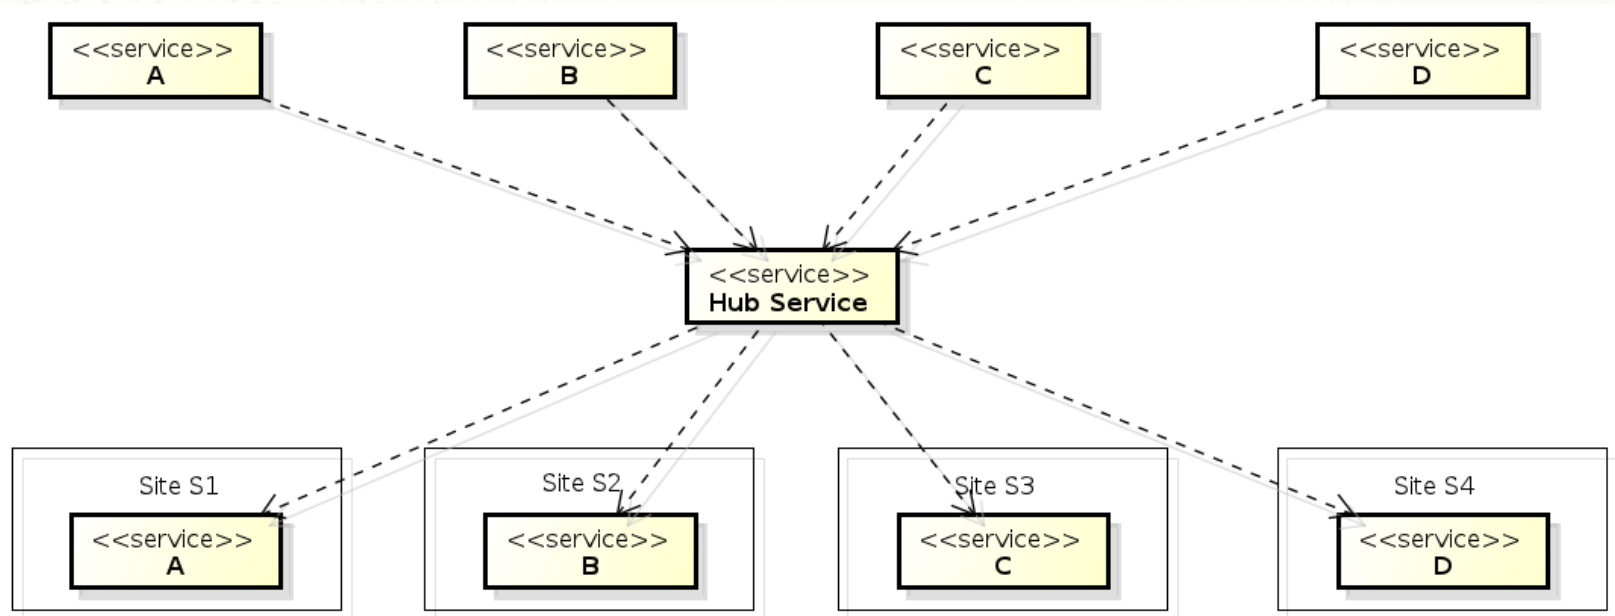
- **Sensitive services**

- Frequently suffer from ripple-effect
- Limit the occurrences of sensitive services during (re)synthesis to improve choreography maintainability and evolvability



- **Hub services (breakable + butterfly)**

- Change propagator – Amplifies changes in a choreography



powered by astah

When choosing between equivalent services, do a dependency analysis and point/exclude possible sensitive, butterfly and hub services

## Baile Project IME/USP

- **Visit us:**

- <http://ccsl.ime.usp.br/baile>

- **Write us:**

- Marco Aurélio Gerosa ([gerosa@ime.usp.br](mailto:gerosa@ime.usp.br))
- Carlos Eduardo Moreira dos Santos ([cadu@ime.usp.br](mailto:cadu@ime.usp.br))

## Some references

- Josuttis, N.: SOA in Practice. O'Reilly Media, Inc. (2007)
- Kokash, N., D'Andrea, V.: Service oriented computing and coordination models. In: Proceedings of Challenges in Collaborative Engineering Workshop, Sopron, Hungary, pp. 95-103. Citeseer (2005)
- Valérie Issarny et al., "Service-oriented middleware for the Future Internet: state of the art and research directions," Journal of Internet Services and Applications 2 (May 25, 2011): 23-45.