Helix Nebula: Securing tomorrow's innovation with today's cloud

Robert Jenkins
CTO CloudSigma

CSA EMEA Congress
September 26th 2012
Amsterdam

CloudSigma: Mission Statement

We deliver ubiquitous cost effective access to flexible computing resources as a utility service.

CloudSigma

CloudSigma: Our Approach

HELIX NEBULA THESCIENCECLOUD

- Complete Data Portability
- Run any OS with full root control (BSD, Linux, Windows etc.)
- No fixed server instance sizes
- Open networking

Helix Nebula Strategic Plan

for a scientific Cloud Computing Infrastructure in Europe

- Establish a sustainable multi-tenant cloud computing infrastructure in Europe
- Initially based on the needs for the European Research Area & space agencies
- Integrate commercial services from multiple IT industry providers

A Collaboration Initiative

HELIX NEBULA
THESCIENCECLOUD

European Commission
& relevant projects

User organisations
Demand-side

European Cloud Computing Strategy

Commercial Service Providers
Supply-side

Bringing together all the stakeholders to establish a public-private partnership

Timeline

- Select flagships use cases
- Identify service providers
- Define governance model

- Deploy flagships
- Analysis of functionality, performance & financial model
- Success Stories

- More applications
- More services
- More users
- More service providers

CloudSigma

Pilot Phase

HELIX NEBULA THESCIENCECLOUD

Explore / push a series of perceived barriers to Cloud adoption:

- **Security**: Unknown or low compliance and security standards
- **Reliability**: Availability of service for business critical tasks
- **Data privacy**: Moving sensitive data to the Cloud
- **Scalability / Elasticity**: Will the Cloud scale-up to our needs
- **Network performance**: Data transfer bottleneck; QoS
- **Integration**: Hybrid systems with in-house / legacy systems
- **Vendor lock-in**: Vendor dependency once data & applications are transferred to the Cloud
- **Legal concerns**: liability, jurisdiction, intellectual property
- **Transparency**: Clarity of conditions, terms and pricing

7

Governance Model

Proof of Concept stage

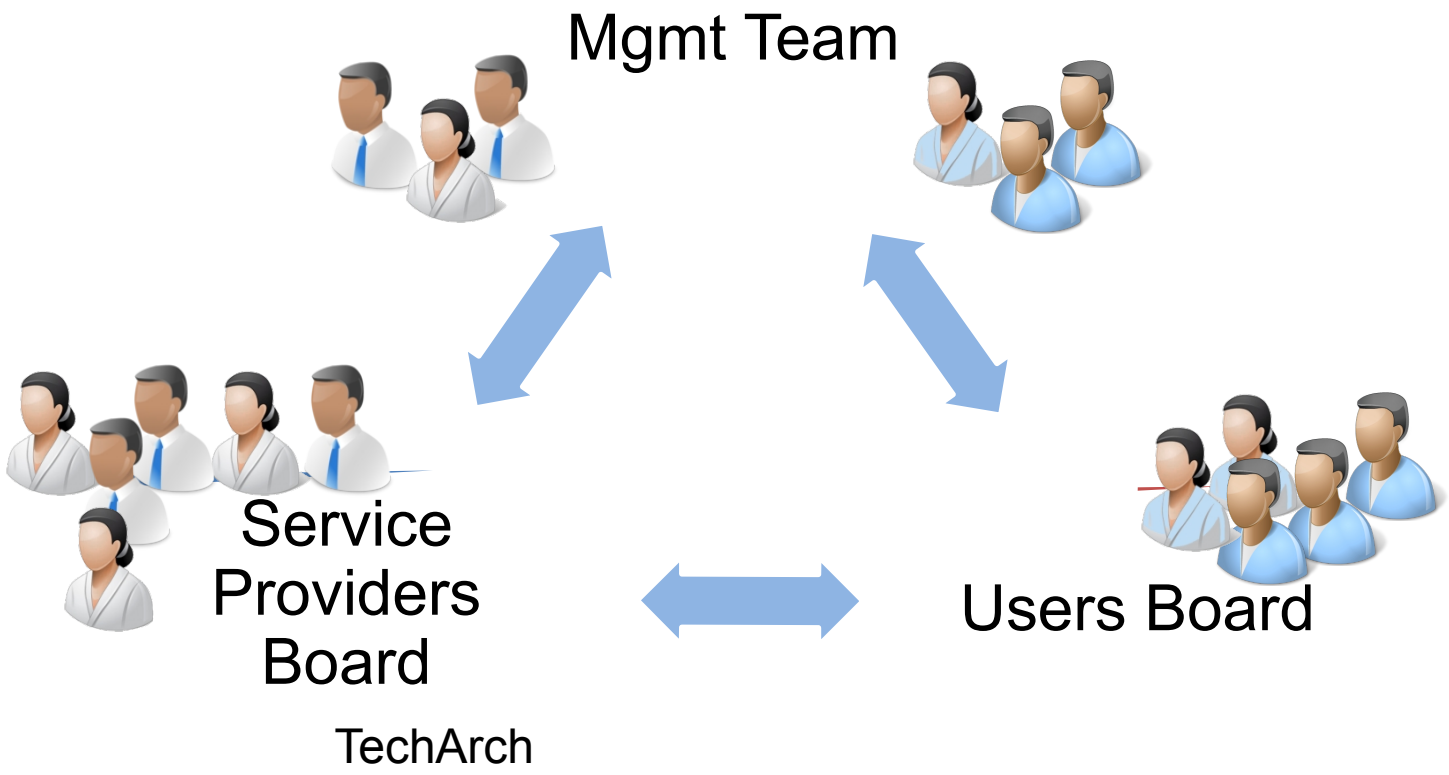

Activities covered by NDA

Initial flagships use cases

ATLAS High Energy Physics Cloud Use	Genomic Assembly in the Cloud	SuperSites Exploitation Platform
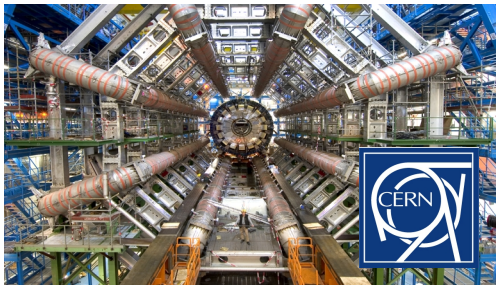	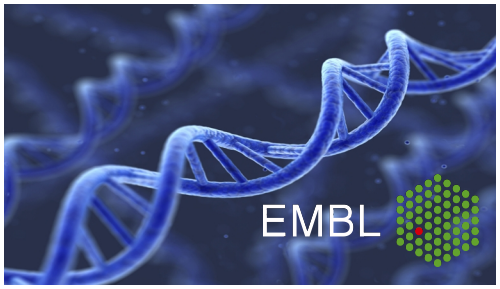	
To support the computing capacity needs for the ATLAS experiment	A new service to simplify large scale genome analysis; for a deeper insight into evolution and biodiversity	To create an Earth Observation platform, focusing on earthquake and volcano research

- **Scientific challenges with societal impact**
- **Sponsored by user organisations**
- ***Stretch* what is possible with the cloud today**

CloudSigma

CERN Grid Usage

HELIX NEBULA THESCIENCECLOUD

CPU used at Tier 1s + Tier 2s (HS06.hrs/month) – last 12 months

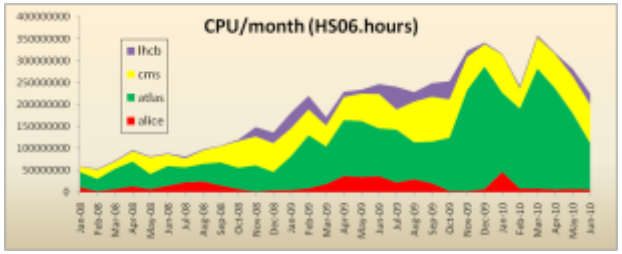

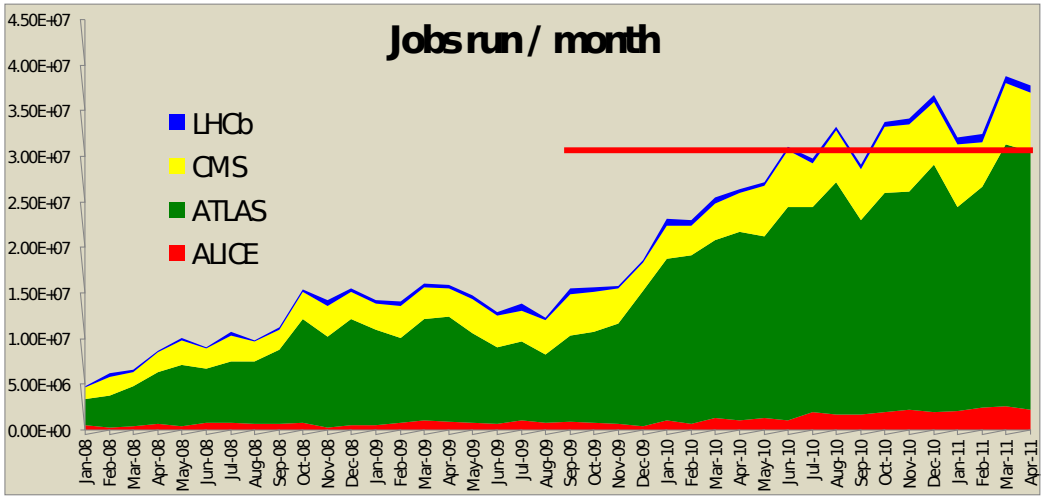

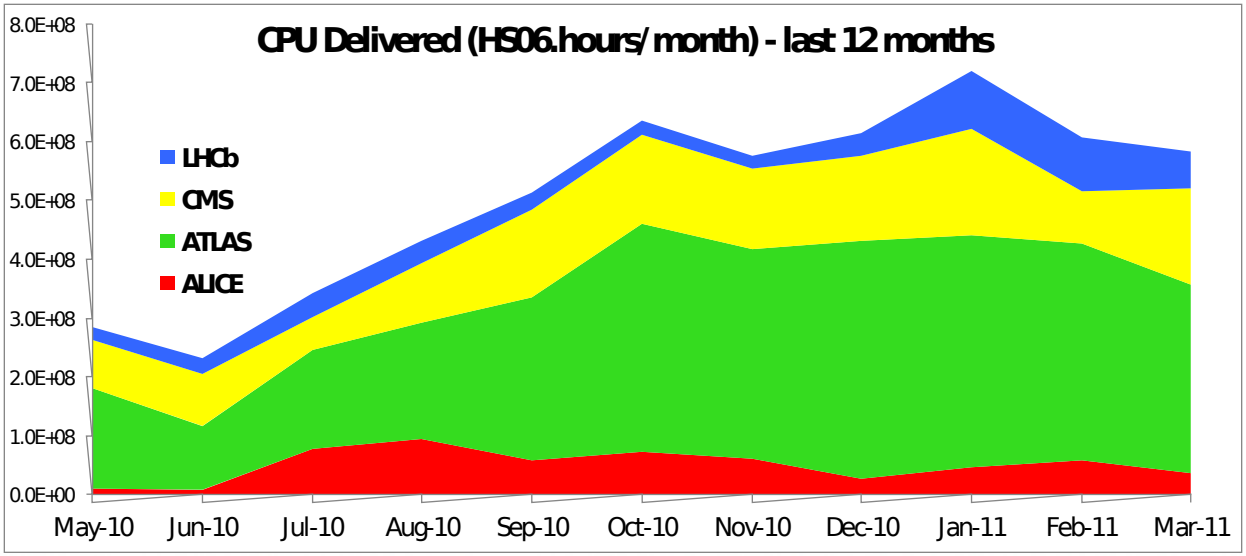

At the end of 2010 we saw all Tier 1 and Tier 2 job slots being filled

CPU usage now >> double that of mid-2010 (inset shows build up over previous years)

In 2010 WLCG delivered ~ 80-100 CPU-millennia!

CloudSigma

ESA: SSEP & GeoHazard Supersites

HELIX NEBULA THESCIENCECLOUD

- The SuperSites Exchange Platform (SSEP) flagship is based on the **GeoHazard Supersites** initiative.

- The GeoHazard Supersites initiative aims to increase the **scientific knowledge of the Earth** through a **global partnership** of projects, organizations and scientists.

- The GeoHazard Supersites goal is to advance the **scientific understanding** of the physical processes controlling **earthquakes**, **volcanic eruptions**, and unrest episodes as well as those driving tectonics and Earth surface dynamics by:
  - improving geohazard monitoring through the **combination of in-situ and space-based data**,
  - **facilitating the access** to data relevant for fundamental research and geohazard assessment.

CloudSigma

EMBL: Genomic Sequencing

HELIX NEBULA THESCIENCECLOUD

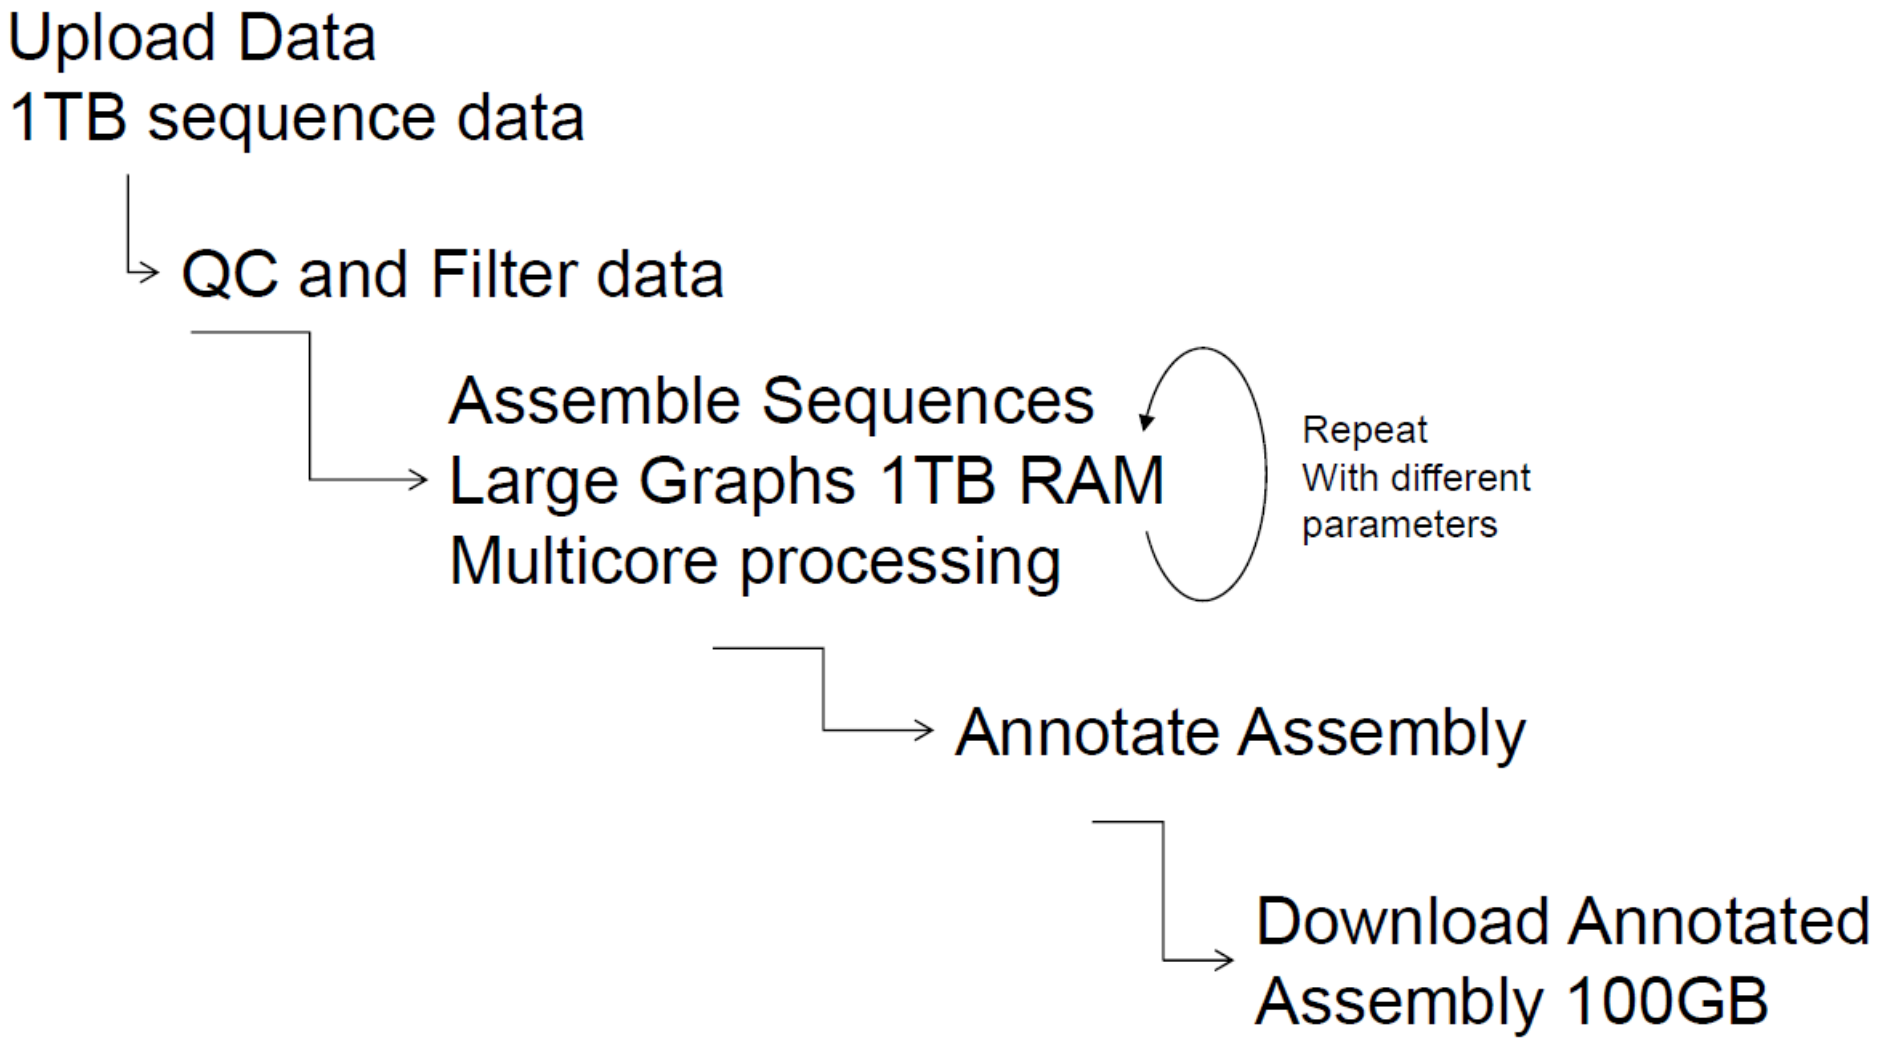

CloudSigma

Flagship Requirements

HELIX NEBULA THESCIENCECLOUD

	ATLAS H.E.P. Cloud Use (CERN)	Genomic Assembly in the Cloud (EMBL)	SuperSites Exploitation Platform (ESA/CNES/DLR)
Scientific goal/society impact/photogenic	•	•	•
Scale of resources used	•	•	
Federation/Aggregation of datasets		•	•
Long-term archiving of data			•
On-demand processing	•	•	•
Impact on community & benefits	•	•	•
Potential increase of users	•	•	•
Interoperability	•	•	•
Data security	•	•	•
Maturity	•	•	•
Access to license-controlled sw			•

CloudSigma

Flagship use cases Participating Suppliers in Proof of Concept stage

HELIX NEBULA THESCIENCECLOUD

CloudSigma

Atos

interoute
from the ground to the cloud

logica
be brilliant together

terradue 2.0

sixsq

T-Systems

the SERVER LABS
the IT architects

CloudSigma

Lessons from POCs

HELIX NEBULA THESCIENCECLOUD

- No common deployment models across participating cloud providers
- Networking set-ups vary considerable
- Demand-side institutions have their own legacy requirements
- Reality is a many to many model for supply to demand side interactions

CloudSigma

Security User Stories

HELIX NEBULA THESCIENCECLOUD

As a customer,

- I can set-up firewall rules for each instantiated VM such that I can control access to my cloud resources
- I can set-up firewall rules between different cloud providers in order to establish transparent networks
- I can manage my instance keys such that each have pre-assigned firewall rules, which I can specify when provisioning instances or modify while the instances are running
- I have a mechanism to rotate login credentials and instance keys such that different instances are started with different credentials and keys thus improving protection and security of my Vms
- I can encrypt data in movement as well as persistent data
- I can use the same login credentials to access all the services and cloud providers using SSO

CloudSigma

Security Aims

HELIX NEBULA THESCIENCECLOUD

- Each cloud provides its own physical security and policies. The consortium has security policy with minimum requirement plus transparency on policies by supplier
- Perimeter security configurable across clouds using a service enabling framework (more later)
- Enable transfer of data & drive/server images between clouds securely and transparently
- Ability to security harden server images across providers using a standardised interface & methodology
- Centralised security monitoring and incident response services

CloudSigma

Federation Architecture

HELIX NEBULA THESCIENCECLOUD

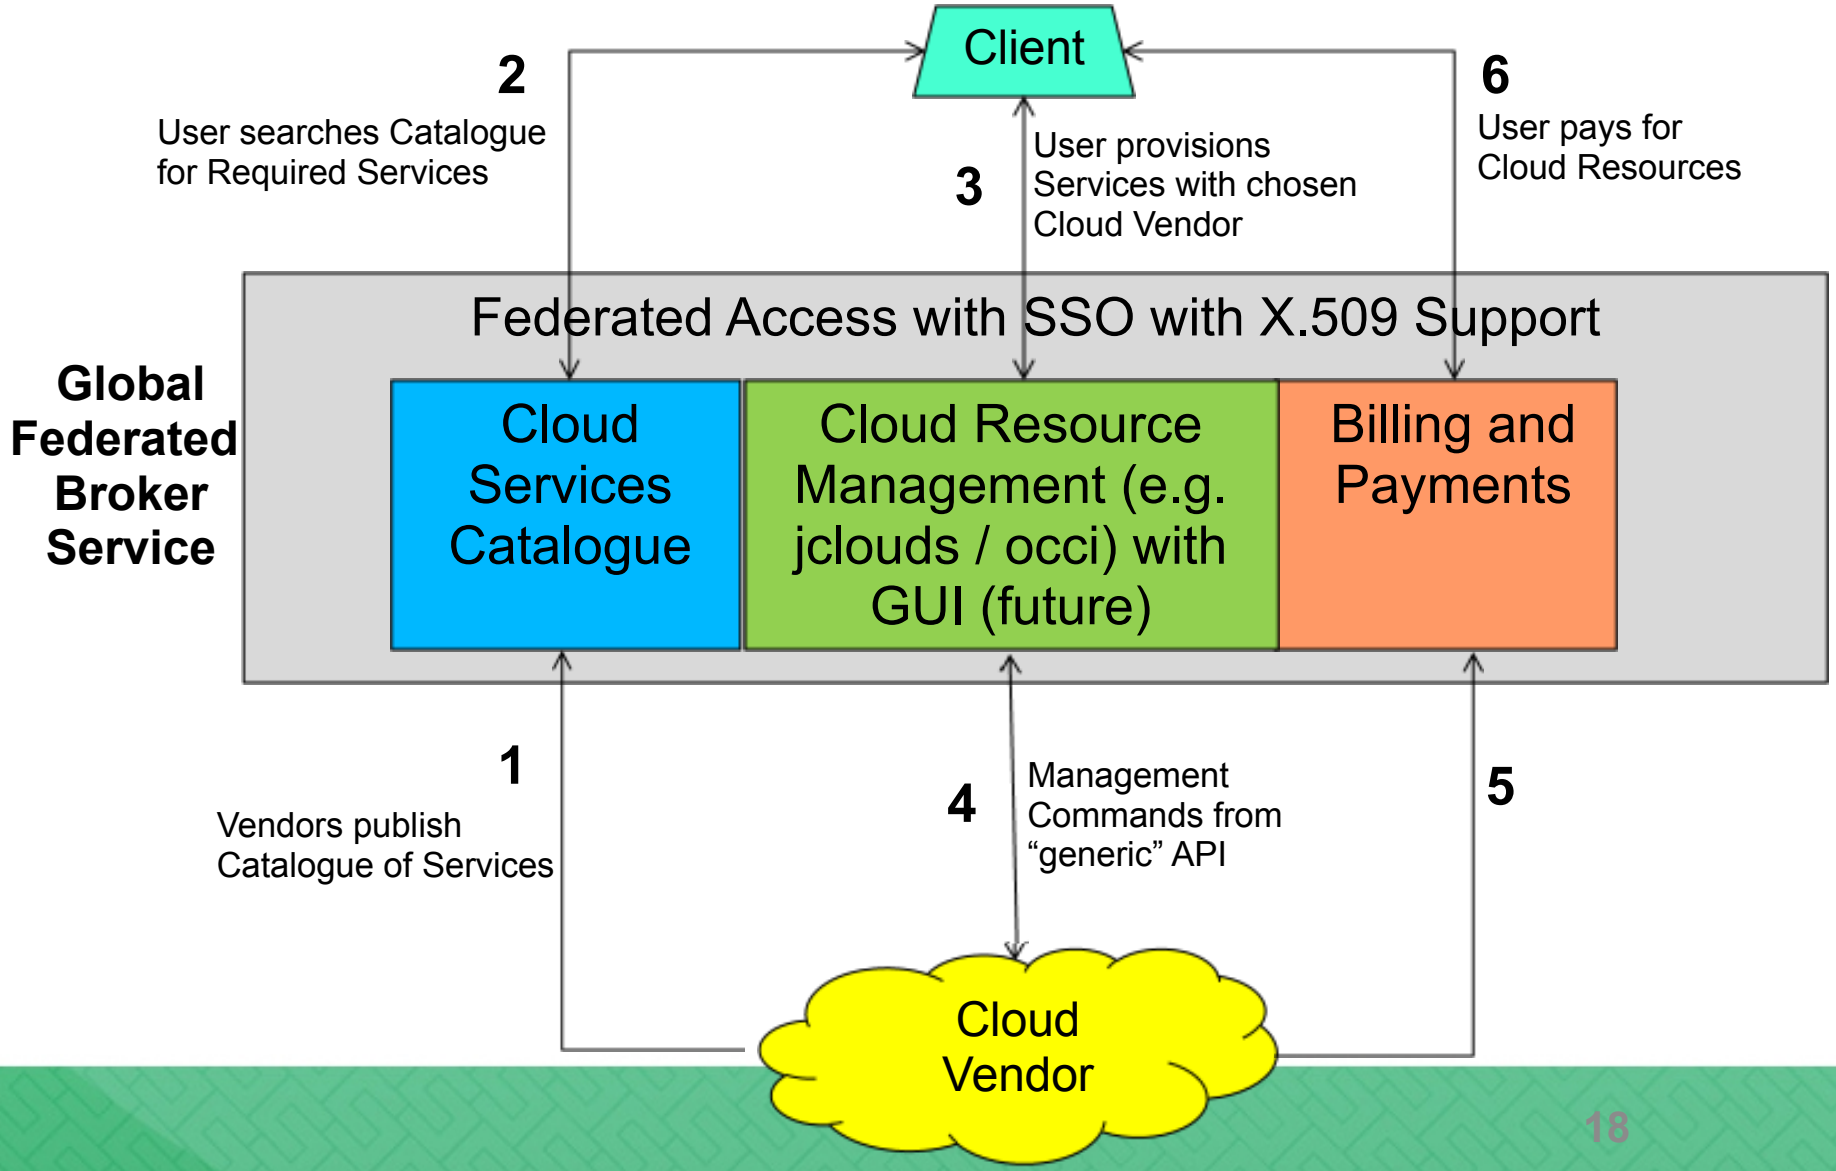

CloudSigma

Federation Architecture

HELIX NEBULA THESCIENCECLOUD

- A platform hosting the Global Federated Broker would offer 3 core services
  - **Cloud Services Catalog**, populated using some common standards-based language, this would be the repository for a vendor to list services, costs, SLA's, etc.
  - **Cloud Resource Management**, this would be a common API (RESTful) and/or a GUI to allow the user to provision and manage cloud compute resources over a group of cloud providers.
  - **Billing and Payments,** here a common set of processes would come into play to allow vendors to bill users for services consumed, there would be some B2B aspects to the delivery of the invoice (e.g. EDI)

- The Broker would need to be:
  - Distributed
  - Highly Available
  - Vendor agnostic
  - Standards based and compliant

Technical guidance and execution driven by dedicated technical architecture committee.

CloudSigma

Federation Architecture

HELIX NEBULA THESCIENCECLOUD

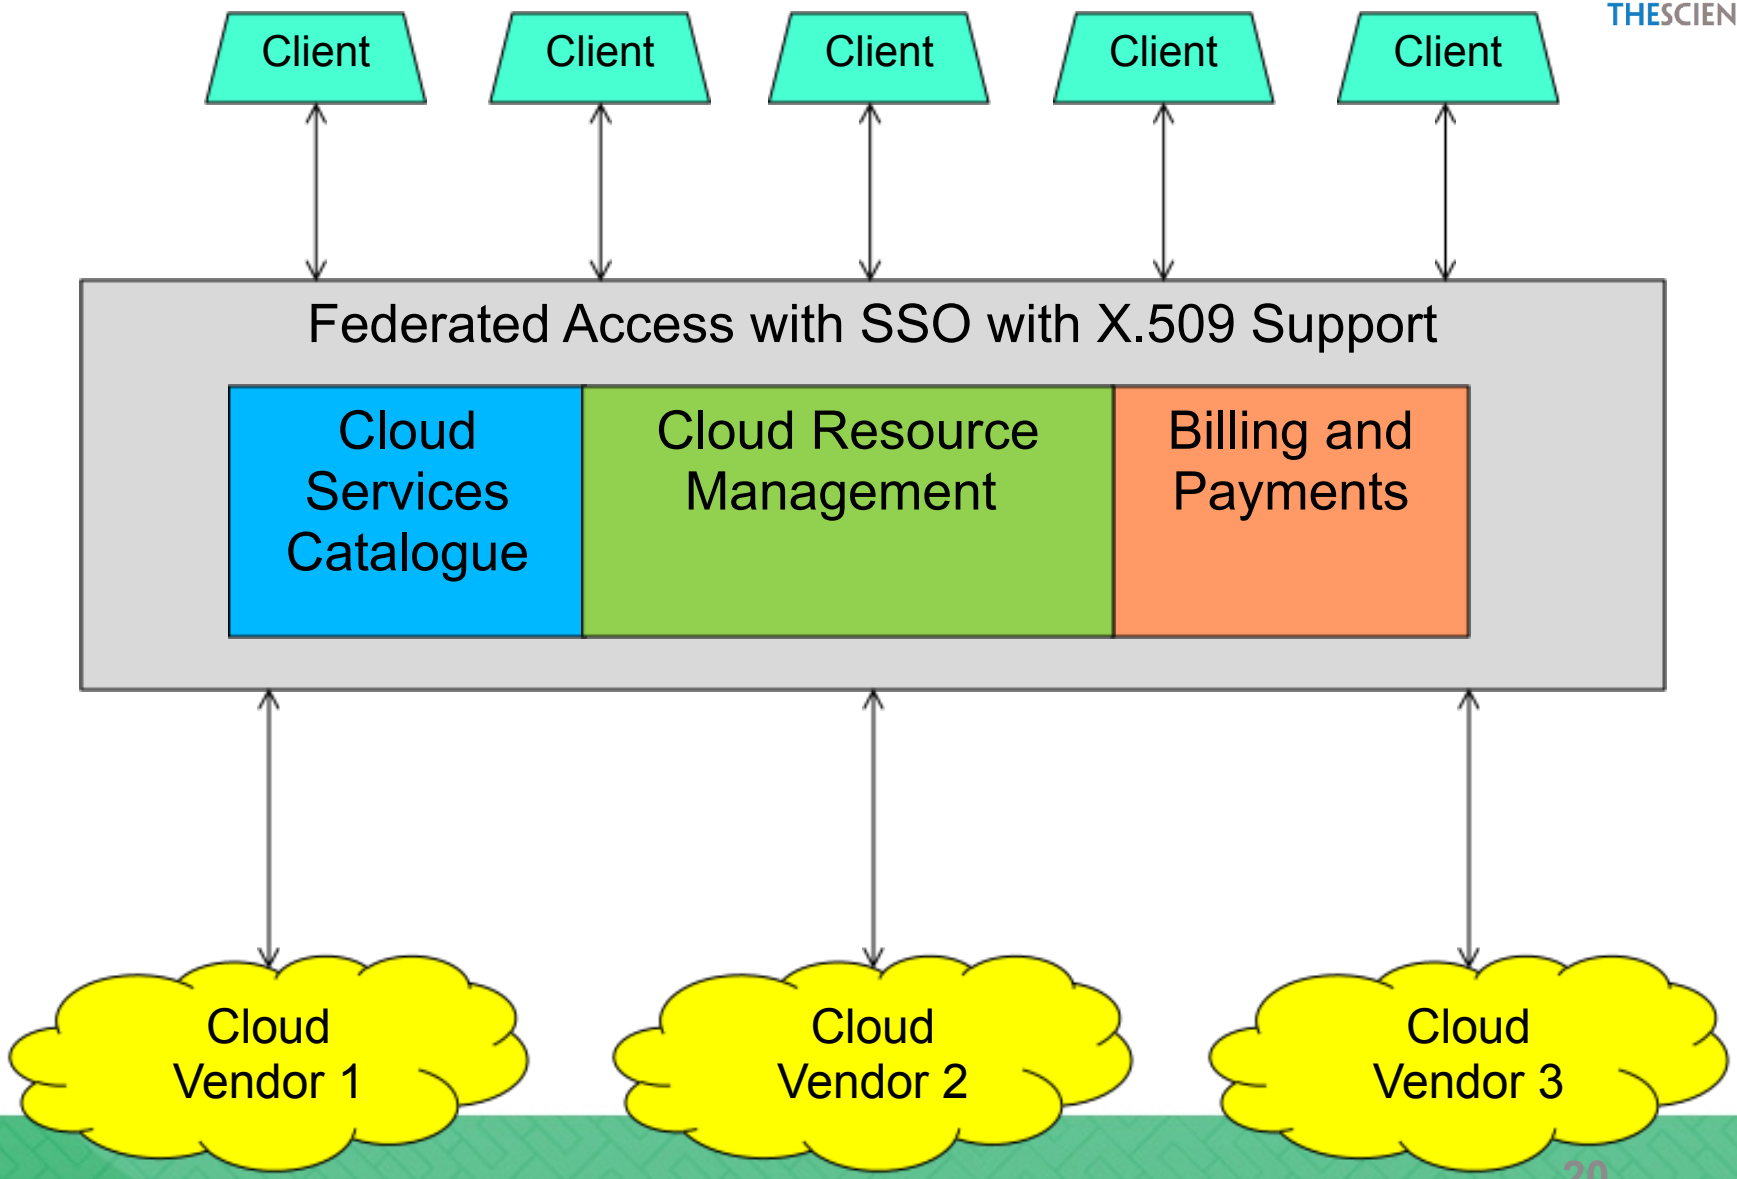

CloudSigma

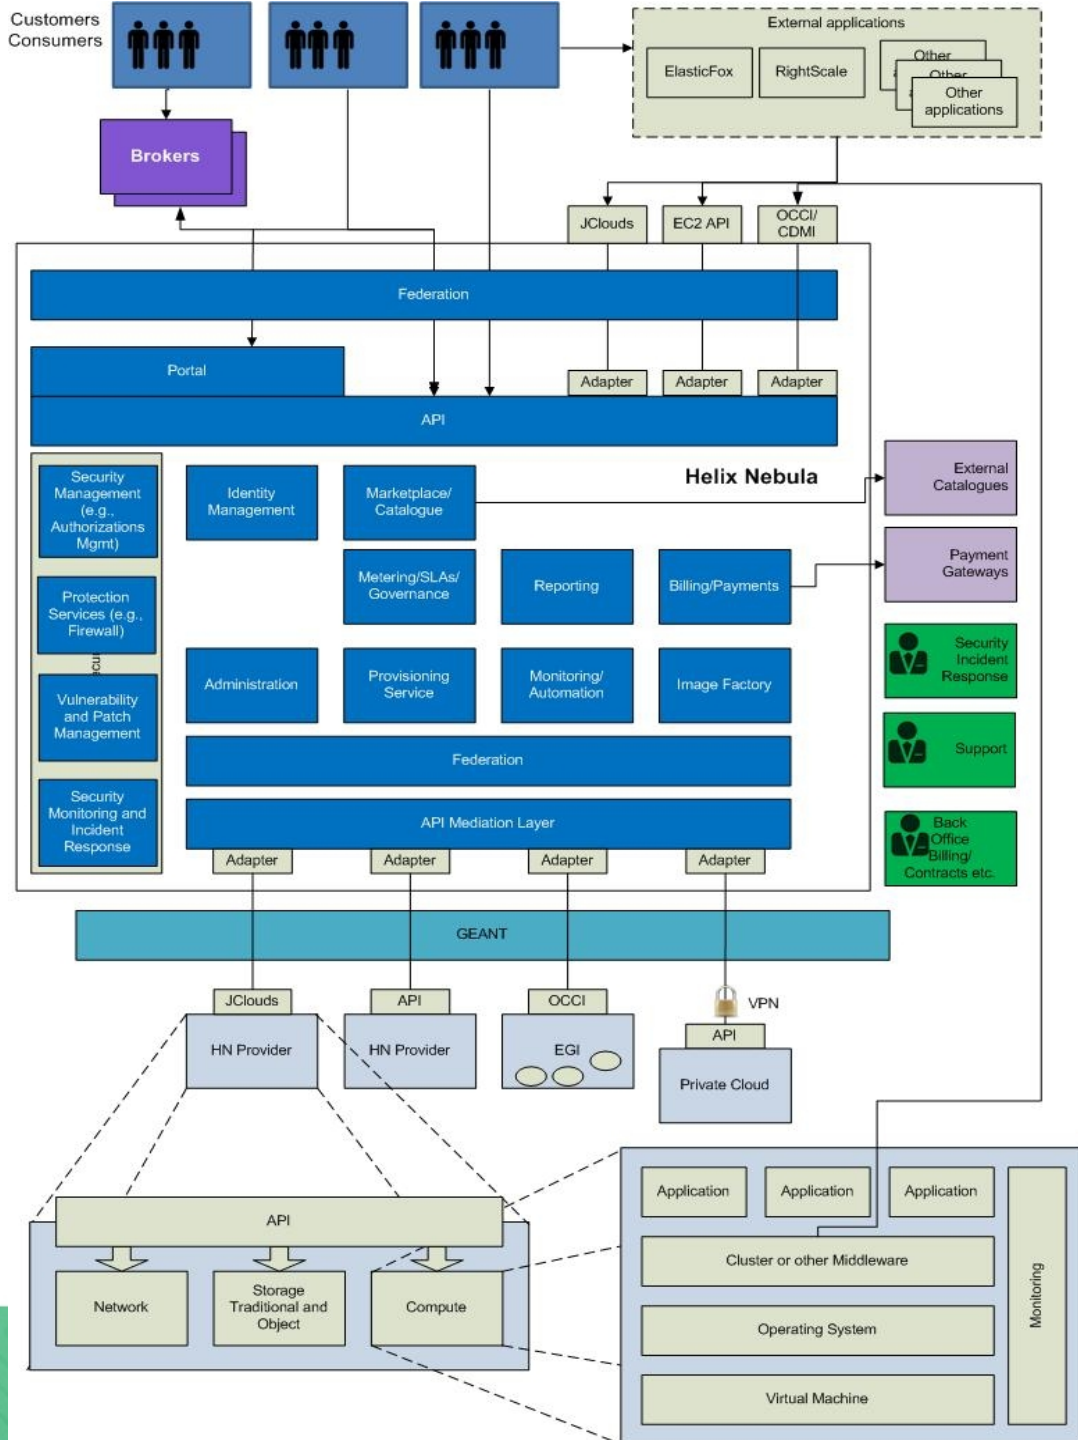
Customers
Consumers
External applications
ElasticFox
RightScale
Other applications
Brokers
JClouds
EC2 API
OCCI/ CDMI
Federation
Portal
Adapter
Adapter
Adapter
API
Helix Nebula
Security Management (e.g., Authorizations Mgmt)
Identity Management
Marketplace/ Catalogue
External Catalogues
Metering/SLAs/ Governance
Reporting
Billing/Payments
Payment Gateways
Protection Services (e.g., Firewall)
Security Incident Response
Administration
Provisioning Service
Monitoring/ Automation
Image Factory
Vulnerability and Patch Management
Support
Federation
Security Monitoring and Incident Response
API Mediation Layer
Back Office Billing/ Contracts etc.
Adapter
Adapter
Adapter
Adapter
GEANT
JClouds
API
OCCI
VPN
HN Provider
HN Provider
EGI
API
Private Cloud
API
Application
Application
Application
Network
Storage Traditional and Object
Compute
Cluster or other Middleware
Monitoring
Operating System
Virtual Machine

Release plan

Release 1:

- Full SSO
- Common API and image library
- Allow full firewall set-up & configuration
- Basic credential and key rotation functionality
- Security monitoring and incident response performed manually
- GEANT connectivity

CloudSigma

Release plan

HELIX NEBULA THESCIENCECLOUD

Release 2:

- Key ring management
- Automated Credential/Key Rotation
- Automated Security Monitoring & Incident Response
- Vulnerability & Patch Management


CloudSigma
HELIX NEBULA
THESCIENCECLOUD