Some thoughts on the evolution of Grid and Cloud computing

D. Salomoni
INFN Tier-1 Computing Manager
Davide.Salomoni@cnaf.infn.it

SuperB Computing R&D Workshop - Ferrara, 9-12 March, 2010

Outline

Outline

The compulsory slide on definitions. . .

*By the name **Grid**, we understand a widely distributed computing infrastructure, including hardware resources and the corresponding software tools and services, which allow optimal execution of computational tasks, with appropriate access to the distributed data. The implementation of such an infrastructure is beyond the scope of this document. The Grid is assumed to provide proper authentication and authorisation, transparent access to resources, and overall management of the necessary databases. (HEPCAL, 2002)*

---

*Five attributes that define **Cloud Computing** (Gartner Group, 2009):*

- *Service-Based*
- *Scalable and Elastic*
- *Shared*
- *Metered by Use*
- *Uses Internet Technologies*

The evolution of cloudy promises

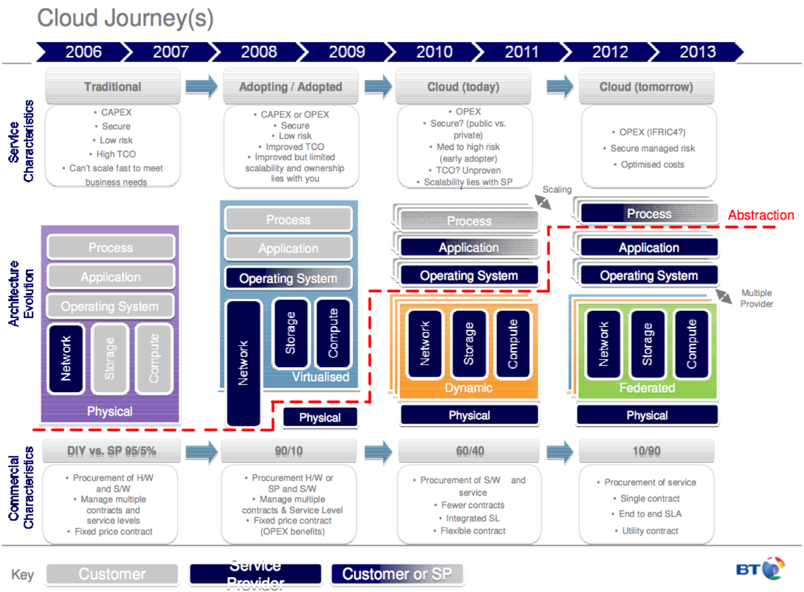

Some Mantras: Service, Cost Savings, Consolidation, Infinity, Utility. All On Demand.

January 27, 2010: The UK government has unveiled a sweeping strategy to create its own internal "cloud computing" system – such as that used by Google, Microsoft and Amazon – as part of a radical plan that it claims could save up to £3.2bn a year from an annual bill of at least £16bn.

> *The key part of the new strategy [...] will be the concentration of government computing power into a series of about **a dozen** highly secure data centres, each costing up to £250m to build, which **will replace more than 500** presently used by central government, police forces and local authorities.*
>
> *By 2015, the strategy suggests, 80% of central government desktops could be supplied through a "shared utility service".*
>
> *The new "cloud" system will not include the security services such as MI5 or MI6, which have their own, separate systems.*

New requirements to existing Grid infrastructures

While Grid interfaces are widely used esp. by large user communities, Cloud computing offers significant advantages for many uses (among them, pay-as-you-go models, simplified access).

Ideally, though, one would like to adopt Cloud services so that:

- Resources are shared between access interfaces (Grid, Cloud, or else).
- Scalability is ensured.
- Existing services and agreements are not required to change.
- Resource center policies and know-how are honored.
- New services can attract both existing and new customers.

These are both key challenges and opportunities for existing Grid infrastructures.

New requirements to existing Grid infrastructures

While Grid interfaces are widely used esp. by large user communities, Cloud computing offers significant advantages for many uses (among them, pay-as-you-go models, simplified access).

Ideally, though, one would like to adopt Cloud services so that:

- Resources are shared between access interfaces (Grid, Cloud, or else).
- Scalability is ensured.
- Existing services and agreements are not required to change.
- Resource center policies and know-how are honored.
- New services can attract both existing and new customers.

These are both key challenges and opportunities for existing Grid infrastructures.

New requirements to existing Grid infrastructures

While Grid interfaces are widely used esp. by large user communities, Cloud computing offers significant advantages for many uses (among them, pay-as-you-go models, simplified access).

Ideally, though, one would like to adopt Cloud services so that:

- Resources are shared between access interfaces (Grid, Cloud, or else).
- Scalability is ensured.
- Existing services and agreements are not required to change.
- Resource center policies and know-how are honored.
- New services can attract both existing and new customers.

These are both key challenges and opportunities for existing Grid infrastructures.

New requirements to existing Grid infrastructures

While Grid interfaces are widely used esp. by large user communities, Cloud computing offers significant advantages for many uses (among them, pay-as-you-go models, simplified access).

Ideally, though, one would like to adopt Cloud services so that:

- Resources are shared between access interfaces (Grid, Cloud, or else).
- Scalability is ensured.
- Existing services and agreements are not required to change.
- Resource center policies and know-how are honored.
- New services can attract both existing and new customers.

These are both key challenges and opportunities for existing Grid infrastructures.

New requirements to existing Grid infrastructures

While Grid interfaces are widely used esp. by large user communities, Cloud computing offers significant advantages for many uses (among them, pay-as-you-go models, simplified access).

Ideally, though, one would like to adopt Cloud services so that:

- Resources are shared between access interfaces (Grid, Cloud, or else).
- Scalability is ensured.
- Existing services and agreements are not required to change.
- Resource center policies and know-how are honored.
- New services can attract both existing and new customers.

These are both key challenges and opportunities for existing Grid infrastructures.

New requirements to existing Grid infrastructures

While Grid interfaces are widely used esp. by large user communities, Cloud computing offers significant advantages for many uses (among them, pay-as-you-go models, simplified access).

Ideally, though, one would like to adopt Cloud services so that:

- Resources are shared between access interfaces (Grid, Cloud, or else).
- Scalability is ensured.
- Existing services and agreements are not required to change.
- Resource center policies and know-how are honored.
- New services can attract both existing and new customers.

These are both key challenges and opportunities for existing Grid infrastructures.

New requirements to existing Grid infrastructures

While Grid interfaces are widely used esp. by large user communities, Cloud computing offers significant advantages for many uses (among them, pay-as-you-go models, simplified access).

Ideally, though, one would like to adopt Cloud services so that:

- Resources are shared between access interfaces (Grid, Cloud, or else).
- Scalability is ensured.
- Existing services and agreements are not required to change.
- Resource center policies and know-how are honored.
- New services can attract both existing and new customers.

These are both key challenges and opportunities for existing Grid infrastructures.

Examples of currently requested services

Some of the typical *new* service requests:

- Customer-definable software environments. This is actually a feature that finds several uses in "traditional" Grids as well.
- Setting up dynamic pools of virtual servers (e.g., user interfaces, or worker nodes for parallel interactive analysis). More generally, flexibly allocating hardware resources through complex advance-reservation requests.
- Instantiating pre-packaged, ready-to-go services.
- Truly distributed, on-demand, Cloud storage.
- Not everybody "speaks Grid": providing access to distributed, traditional Grid infrastructures as if they were not Grids, also to non-traditional users, like Public Administrations, or to the private sector.

The key problem is one of integration between several access interfaces (Grid, Cloud, or else).

Examples of currently requested services

Some of the typical *new* service requests:

- Customer-definable software environments. This is actually a feature that finds several uses in "traditional" Grids as well.
- Setting up dynamic pools of virtual servers (e.g., user interfaces, or worker nodes for parallel interactive analysis). More generally, flexibly allocating hardware resources through complex advance-reservation requests.
- Instantiating pre-packaged, ready-to-go services.
- Truly distributed, on-demand, Cloud storage.
- Not everybody "speaks Grid": providing access to distributed, traditional Grid infrastructures as if they were not Grids, also to non-traditional users, like Public Administrations, or to the private sector.

The key problem is one of integration between several access interfaces (Grid, Cloud, or else).

Examples of currently requested services

Some of the typical *new* service requests:

- Customer-definable software environments. This is actually a feature that finds several uses in "traditional" Grids as well.
- Setting up dynamic pools of virtual servers (e.g., user interfaces, or worker nodes for parallel interactive analysis). More generally, flexibly allocating hardware resources through complex advance-reservation requests.
- Instantiating pre-packaged, ready-to-go services.
- Truly distributed, on-demand, Cloud storage.
- Not everybody "speaks Grid": providing access to distributed, traditional Grid infrastructures as if they were not Grids, also to non-traditional users, like Public Administrations, or to the private sector.

The key problem is one of integration between several access interfaces (Grid, Cloud, or else).

Examples of currently requested services

Some of the typical *new* service requests:

- Customer-definable software environments. This is actually a feature that finds several uses in "traditional" Grids as well.
- Setting up dynamic pools of virtual servers (e.g., user interfaces, or worker nodes for parallel interactive analysis). More generally, flexibly allocating hardware resources through complex advance-reservation requests.
- Instantiating pre-packaged, ready-to-go services.
- Truly distributed, on-demand, Cloud storage.
- Not everybody "speaks Grid": providing access to distributed, traditional Grid infrastructures as if they were not Grids, also to non-traditional users, like Public Administrations, or to the private sector.

The key problem is one of integration between several access interfaces (Grid, Cloud, or else).

Examples of currently requested services

Some of the typical *new* service requests:

- Customer-definable software environments. This is actually a feature that finds several uses in "traditional" Grids as well.
- Setting up dynamic pools of virtual servers (e.g., user interfaces, or worker nodes for parallel interactive analysis). More generally, flexibly allocating hardware resources through complex advance-reservation requests.
- Instantiating pre-packaged, ready-to-go services.
- Truly distributed, on-demand, Cloud storage.
- Not everybody "speaks Grid": providing access to distributed, traditional Grid infrastructures as if they were not Grids, also to non-traditional users, like Public Administrations, or to the private sector.

The key problem is one of integration between several access interfaces (Grid, Cloud, or else).

Examples of currently requested services

Some of the typical *new* service requests:

- Customer-definable software environments. This is actually a feature that finds several uses in "traditional" Grids as well.
- Setting up dynamic pools of virtual servers (e.g., user interfaces, or worker nodes for parallel interactive analysis). More generally, flexibly allocating hardware resources through complex advance-reservation requests.
- Instantiating pre-packaged, ready-to-go services.
- Truly distributed, on-demand, Cloud storage.
- Not everybody "speaks Grid": providing access to distributed, traditional Grid infrastructures as if they were not Grids, also to non-traditional users, like Public Administrations, or to the private sector.

The key problem is one of integration between several access interfaces (Grid, Cloud, or else).

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

On Cloud storage, very briefly

The SNIA (Storage Networking Industry Association) has published an interesting paper on Cloud Storage Use Cases. Some examples:

- Unstructured Data Storage (e.g., a cloud drive – think Dropbox)
- Backup to the Cloud
- Archive and Preservation to the Cloud
- Databases in the Cloud (e.g. Google Bigtable)
- Storage for Cloud Computing (e.g. VM images, Application images, Amazon S3)
- Content distribution, i.e. distributing data geographically
- "Intercloud" storage (e.g. federated cloud storage, cloud bursting, hybrid public/private clouds)

The definition of Cloud Storage boundaries and QoS can take many forms. These forms and the related offerings are probably not fully defined even by the industry.

Outline

A practical Grid/Cloud integration example: WNoDeS

The **W**orker **N**odes **o**n **De**mands **S**ervice (WNoDeS) is a software INFN is developing. It is built around a tight integration with a LRMS (a "batch system") and is running in production at the INFN Tier-1 Computing Center. Its main characteristics are:

- Full integration with existing computing resource scheduling, policing, monitoring and accounting workflows.
- On-demand virtual resource provisioning and VLAN support to dynamically isolate Virtual Machines depending on service type / customer requests.
- Support for users to select and access WNoDeS-based resources through Grid, Cloud interfaces, or also through direct job submissions.

The WNoDeS focus is on *Everything as a Service*, where *Everything* may be hardware, software, data, platform, infrastructure.

A practical Grid/Cloud integration example: WNoDeS

The **W**orker **N**odes **o**n **De**mands **S**ervice (WNoDeS) is a software INFN is developing. It is built around a tight integration with a LRMS (a "batch system") and is running in production at the INFN Tier-1 Computing Center. Its main characteristics are:

- Full integration with existing computing resource scheduling, policing, monitoring and accounting workflows.
- On-demand virtual resource provisioning and VLAN support to dynamically isolate Virtual Machines depending on service type / customer requests.
- Support for users to select and access WNoDeS-based resources through Grid, Cloud interfaces, or also through direct job submissions.

The WNoDeS focus is on *Everything as a Service*, where *Everything* may be hardware, software, data, platform, infrastructure.

A practical Grid/Cloud integration example: WNoDeS

The **Worker Nodes on Demands Service** (WNoDeS) is a software INFN is developing. It is built around a tight integration with a LRMS (a "batch system") and is running in production at the INFN Tier-1 Computing Center. Its main characteristics are:

- Full integration with existing computing resource scheduling, policing, monitoring and accounting workflows.
- On-demand virtual resource provisioning and VLAN support to dynamically isolate Virtual Machines depending on service type / customer requests.
- Support for users to select and access WNoDeS-based resources through Grid, Cloud interfaces, or also through direct job submissions.

The WNoDeS focus is on *Everything as a Service*, where *Everything* may be hardware, software, data, platform, infrastructure.

A practical Grid/Cloud integration example: WNoDeS

The **W**orker **N**odes **o**n **De**mands **S**ervice (WNoDeS) is a software INFN is developing. It is built around a tight integration with a LRMS (a "batch system") and is running in production at the INFN Tier-1 Computing Center. Its main characteristics are:

- Full integration with existing computing resource scheduling, policing, monitoring and accounting workflows.
- On-demand virtual resource provisioning and VLAN support to dynamically isolate Virtual Machines depending on service type / customer requests.
- Support for users to select and access WNoDeS-based resources through Grid, Cloud interfaces, or also through direct job submissions.

The WNoDeS focus is on *Everything as a Service*, where *Everything* may be hardware, software, data, platform, infrastructure.

A practical Grid/Cloud integration example: WNoDeS

The **W**orker **N**odes **on** **De**mands **S**ervice (WNoDeS) is a software INFN is developing. It is built around a tight integration with a LRMS (a "batch system") and is running in production at the INFN Tier-1 Computing Center. Its main characteristics are:

- Full integration with existing computing resource scheduling, policing, monitoring and accounting workflows.
- On-demand virtual resource provisioning and VLAN support to dynamically isolate Virtual Machines depending on service type / customer requests.
- Support for users to select and access WNoDeS-based resources through Grid, Cloud interfaces, or also through direct job submissions.

The WNoDeS focus is on *Everything as a Service*, where *Everything* may be hardware, software, data, platform, infrastructure.

Integrating Cloud services

Access to Cloud services happens in WNoDeS through the Open Grid Forum Open Cloud Computing Interface (OCCI). WNoDeS implements a subset of the OCCI API, using X.509 authentication and exposing a REST interface.

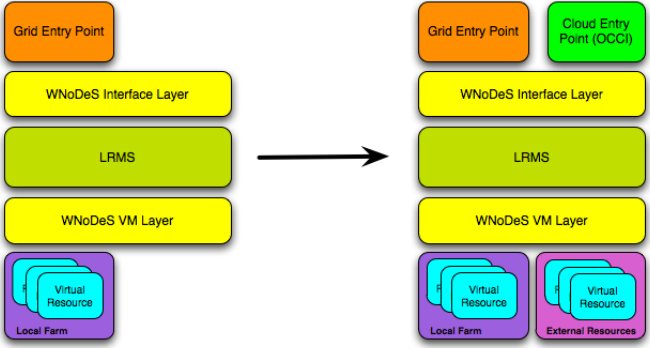

Access to resources is then polymorphic, but what about integration with an existing distributed infrastructure like EGEE?

Integrating Cloud services

Access to Cloud services happens in WNoDeS through the Open Grid Forum Open Cloud Computing Interface (OCCI). WNoDeS implements a subset of the OCCI API, using X.509 authentication and exposing a REST interface.

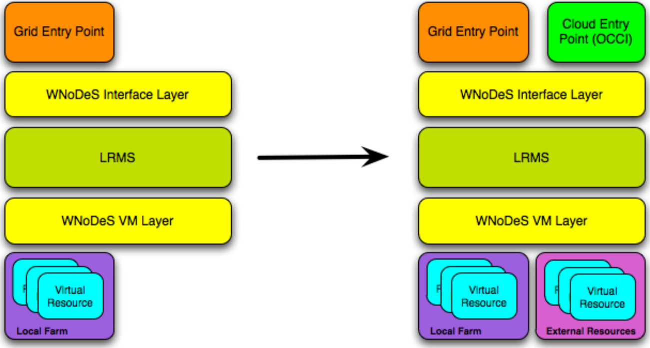

Access to resources is then polymorphic, but what about integration with an existing distributed infrastructure like EGEE?

Integrating authentication mechanisms

X.509-based access, widely adopted in Grids, is by no means the only authentication mechanism in use. Providing access to services for an expanded customer base means also the need to cater for several authentication methods.

The WNoDeS project is developing an authentication gateway to map several authentication mechanisms (Kerberos, Shibboleth) to a dynamically-assigned, short-lived X.509 personal certificate.

Authentication gateway advantages

- Authentication for the WNoDeS software framework converges around a single method (X.509).
- X.509 is the authentication mechanism used to access the Grid. Being able to generate a dynamic X.509 certificate opens up the possibility for e.g. users of Cloud services to access Grid (e.g. EGEE) resources.

Integrating authentication mechanisms

X.509-based access, widely adopted in Grids, is by no means the only authentication mechanism in use. Providing access to services for an expanded customer base means also the need to cater for several authentication methods.

The WNoDeS project is developing an authentication gateway to map several authentication mechanisms (Kerberos, Shibboleth) to a dynamically-assigned, short-lived X.509 personal certificate.

Authentication gateway advantages

- Authentication for the WNoDeS software framework converges around a single method (X.509).
- X.509 is the authentication mechanism used to access the Grid. Being able to generate a dynamic X.509 certificate opens up the possibility for e.g. users of Cloud services to access Grid (e.g. EGEE) resources.

Integrating authentication mechanisms

X.509-based access, widely adopted in Grids, is by no means the only authentication mechanism in use. Providing access to services for an expanded customer base means also the need to cater for several authentication methods.

The WNoDeS project is developing an authentication gateway to map several authentication mechanisms (Kerberos, Shibboleth) to a dynamically-assigned, short-lived X.509 personal certificate.

Authentication gateway advantages

- Authentication for the WNoDeS software framework converges around a single method (X.509).
- X.509 is the authentication mechanism used to access the Grid. Being able to generate a dynamic X.509 certificate opens up the possibility for e.g. users of Cloud services to access Grid (e.g. EGEE) resources.

The WNoDeS authentication gateway

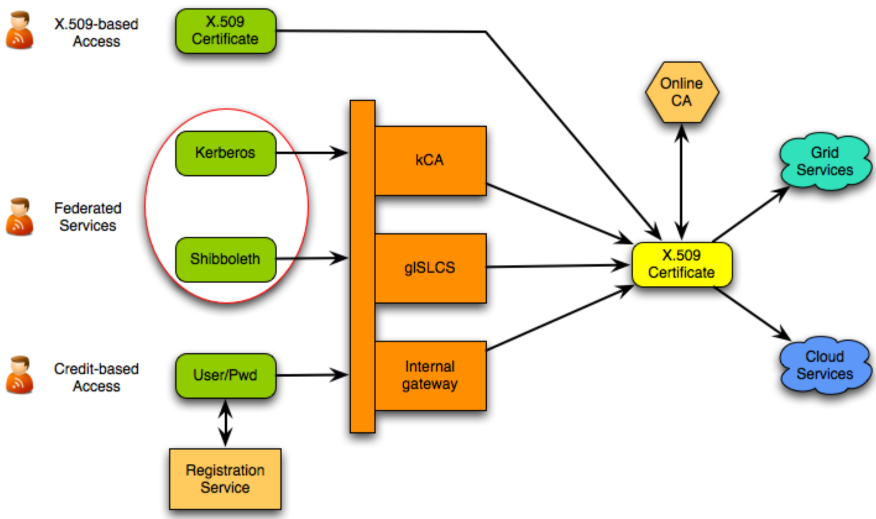

WNoDeS: overall architectural framework

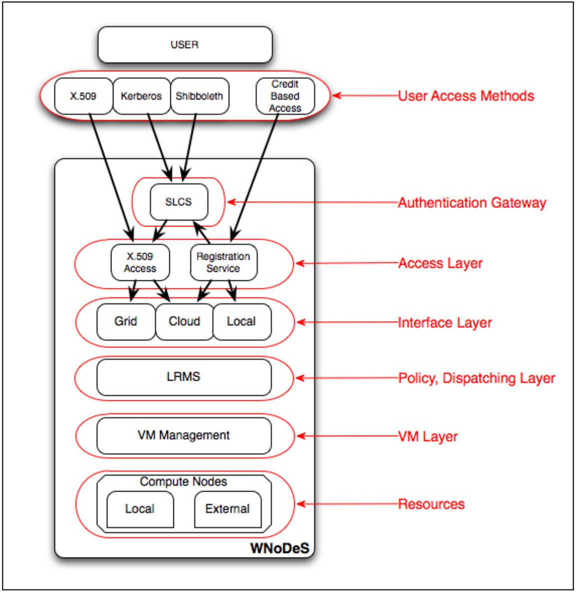

Key points

1. On-demand virtual service provisioning
2. Flexible, integrated scheduling policies
3. Multiple access interfaces
4. Multiple authentication methods
5. Integrated access to existing (e.g., Grid) infrastructures
6. Access to external resources

Outline

WNoDeS: status

The WNoDes framework is running at the INFN Tier-1

- It is in production with currently ∼1400 on-demand Virtual Machines, O(10) supported Virtual Images, serving 20 different user communities; on average, more than 20,000 jobs are executed each day through WNoDeS.
- The plan is to have ∼4000 Virtual Machines by April 2010 and progressively integrate all Tier-1 resources.
- Distributed selection of VM images works either statically (on a customer-community basis), or dynamically (per-user), through standard Grid job submission commands.
- A pilot Cloud service is in place to allocate on-demand resources through OCCI.

A few questions (1)

- Do we still need or want an e-Infrastructure for scientific computing? Do *you* care?

The Economist

Log in | Register | My account

Home | World | Business & finance | Science & technology

Tech.view

Cloudy with a chance of rain

Few companies are ready to accept cloud computing

Mar 5th 2010 | From *The Economist* online

A recent poll by CommVault identified the following as the main obstacles or worries for cloud adoption:

- Security and privacy
- Reliability
- Cost
- Scalability (are there *really* infinite resources?)

Can we formalize what our *difference* is?

A few questions (2a)

- Do we fully understand the economic implications of each choice? See the article by Folker Meyer of ANL, "Genome Sequencing vs. Moore's Law: Cyber Challenges for the Next Decade" (2006).

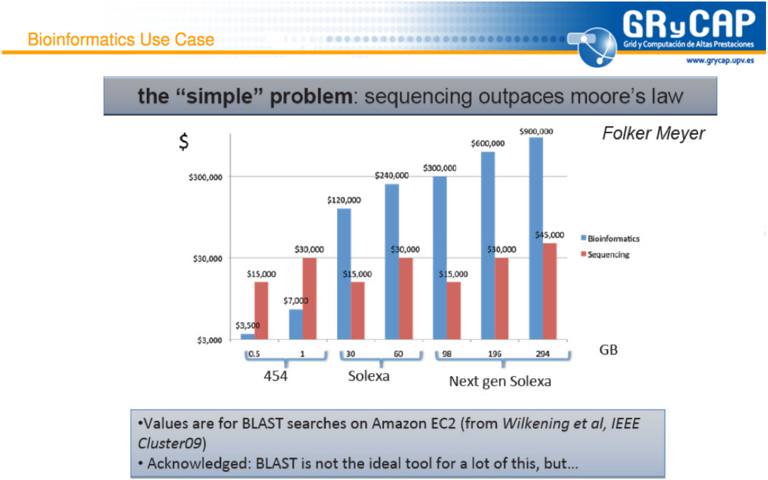

A few questions (2b)

For example, elasticity is a great thing, but. . .

Cloud Elasticity Could Make You Go Broke

March 11th, 2009 · 13 Comments

Ever had a mobile phone and get a bill that was way, way more than you expected? You know what I mean. The day that bill for 700 dollars comes in and your eyes bug out of your head because you could swear (and in fact you do swear – at the customer service rep) that you could not possibly have exceeded your plan minutes? Or maybe you "pay as you

(Daryl Plummer, Gartner) – i.e., **On-Demand Overspending:**

A few questions (3)

- What do we make of the existing, considerable know-how in middleware, software development, and in general of the investments made in developing and running the current production grid infrastructures? How could a new computing model integrate into (and perhaps shape) projects like EGI, or EMI, or the National Grid/Cloud Infrastructures?
- Do we still need or want a multi-tiered computing model, with – say – N>1 or N>2 in Tier-N? Do we still need per-experiment, per-research, per-subgroup resources? Think also of overhead, critical mass, and expected availability/reliability of smaller sites.
- The computing/business world is undoubtedly showing, like always, a changing (and cloudy) scenario, e.g. for what regards service provider offerings. Can we start exploring or asking for or testing "novel" requirements *now*?

A few questions (3)

- What do we make of the existing, considerable know-how in middleware, software development, and in general of the investments made in developing and running the current production grid infrastructures? How could a new computing model integrate into (and perhaps shape) projects like EGI, or EMI, or the National Grid/Cloud Infrastructures?
- Do we still need or want a multi-tiered computing model, with – say – N>1 or N>2 in Tier-N? Do we still need per-experiment, per-research, per-subgroup resources? Think also of overhead, critical mass, and expected availability/reliability of smaller sites.
- The computing/business world is undoubtedly showing, like always, a changing (and cloudy) scenario, e.g. for what regards service provider offerings. Can we start exploring or asking for or testing "novel" requirements *now*?

A few questions (3)

- What do we make of the existing, considerable know-how in middleware, software development, and in general of the investments made in developing and running the current production grid infrastructures? How could a new computing model integrate into (and perhaps shape) projects like EGI, or EMI, or the National Grid/Cloud Infrastructures?
- Do we still need or want a multi-tiered computing model, with – say – N>1 or N>2 in Tier-N? Do we still need per-experiment, per-research, per-subgroup resources? Think also of overhead, critical mass, and expected availability/reliability of smaller sites.
- The computing/business world is undoubtedly showing, like always, a changing (and cloudy) scenario, e.g. for what regards service provider offerings. Can we start exploring or asking for or testing "novel" requirements *now*?

That's it – It's Shoppable (although perhaps not yet shippable)

[Cloud computing is] nothing more than a faddish term for the established concept of computers linked by networks. A cloud is water vapor. (Larry Ellison, co-founder and CEO, Oracle Corporation, September 2009)

The truth is rarely pure and never simple. (Oscar Wilde, The Importance of Being Earnest, 1895)

Thanks!

e-mail: `Davide.Salomoni@cnaf.infn.it`