An Introduction to Grid Technology
and the
Utility Data Center

Jem Treadwell
Senior Software Engineer
Hewlett-Packard
Software Global Business Unit
Sven Graupner
HP Labs

HP WORLD 2003
Solutions and Technology Conference & Expo

Agenda

- Introduction to Grid Computing
- The Globus Project
  - Globus Toolkit 2.4
  - OGSA/OGSI & Globus Toolkit 3.0
- Utility Computing and the Adaptive Enterprise (AE)
- HP's Utility Data Center (UDC)
- Utility Data Center and the Enterprise Grid

Agenda

- Introduction to Grid Computing
- The Globus Project
  - Globus Toolkit 2.4
  - OGSA/OGSI & Globus Toolkit 3.0
- Utility Computing and the Adaptive Enterprise (AE)
- HP's Utility Data Center (UDC)
- Utility Data Center and the Enterprise Grid

About Grids

Warren Gretz

- A **grid** infrastructure enables organizations to share distributed resources
- An **electrical grid** distributes **power** to consumers **when** and **where** it is needed
- A **computer grid** makes **computing resources** available **when** and **where** they are needed

Grid Computing: Background & Vision

- Idea emerged in the early nineties:
  - Networks were commonplace
  - The Internet was established
  - A few people began to envisage the use of networks for **utility computing**:
    - Computer resources as ubiquitous, on-demand commodity
- Key technologies & standards:
  - Security (SSL, PKI, Kerberos)
  - File Transfer (FTP)
  - Remote Procedure Calls
  - Resource directories (LDAP)
  - Load-balancing (LSF, PBS)
  - Parallel processing (MPI)
  - Cycle-stealing (Condor, United Devices, SETI@home)
  - Virtual environments (Legion, Avaki)
  - Global addressing - WWW & URIs
- Foster, Tuecke & Kesselman:
  - Published **The Anatomy of the Grid**
  - Started the **Globus Project**
  - Began to formalize the concept

Networks & Grids

- A Grid uses a network infrastructure to:
  - Span boundaries that may be:
    - organizational
    - geographical
    - political
  - Share resources:
    - easily
    - safely

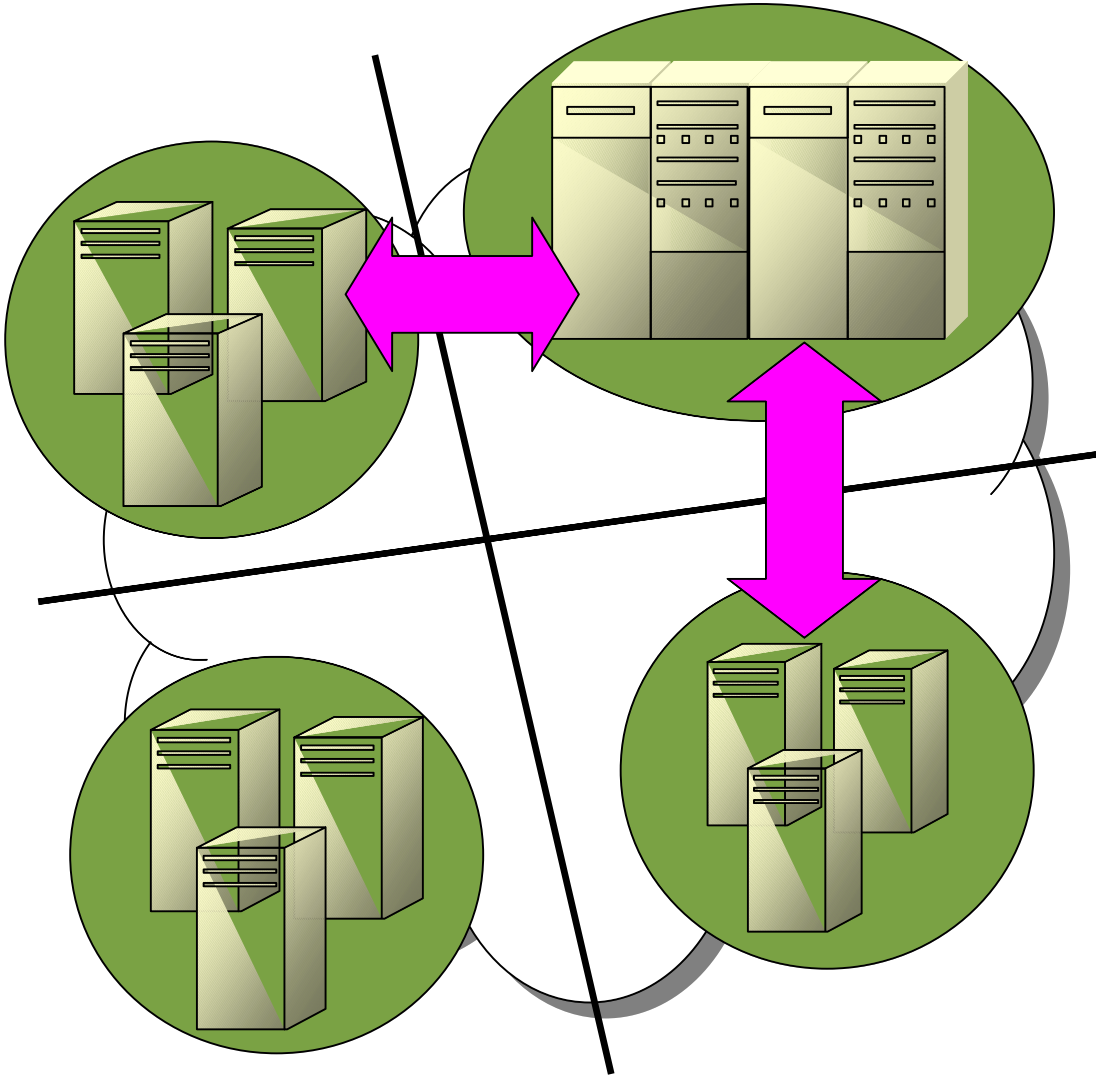

Grids Defined

- A Grid:
  - Coordinates resources that are not subject to centralized control
  - Uses standard, open, general purpose protocols and interfaces
  - Delivers non-trivial qualities of service

    – *Ian Foster, 2002*
- An abstract concept:
  - There is no central "InterGrid"
  - A Grid can be formed simply by using a common infrastructure & protocols to share resources
  - No specific products required
- Primary Grid types:
  - Computational Grid:
    - Primarily for access to computational power
    - e.g. resources for long-running simulations
  - Data Grid:
    - Primarily for access to data
    - e.g. results of scientific observations
  - Grids often have characteristics of both!

Example Grid Uses

- A company has under-utilized resources at various global locations
  - Uses Grid technology to make them available to employees in other locations
- A company has cyclic resource demands for its web service
  - Uses Grid technology to access other internal and external resources at peak periods
- A research lab's simulation takes **months** to run using local resources alone
  - Reduces the time to **days** by joining a research Grid and agreeing to share resources
- A life sciences researcher needs access to remotely stored experimental data
  - Data provider uses Grid technology to provide controlled access to the data
- Engineers in collaborating companies need to share design data and facilities
  - Use Grid technology to provide each other with secure, restricted access to the necessary facilities

Example: TeraGrid

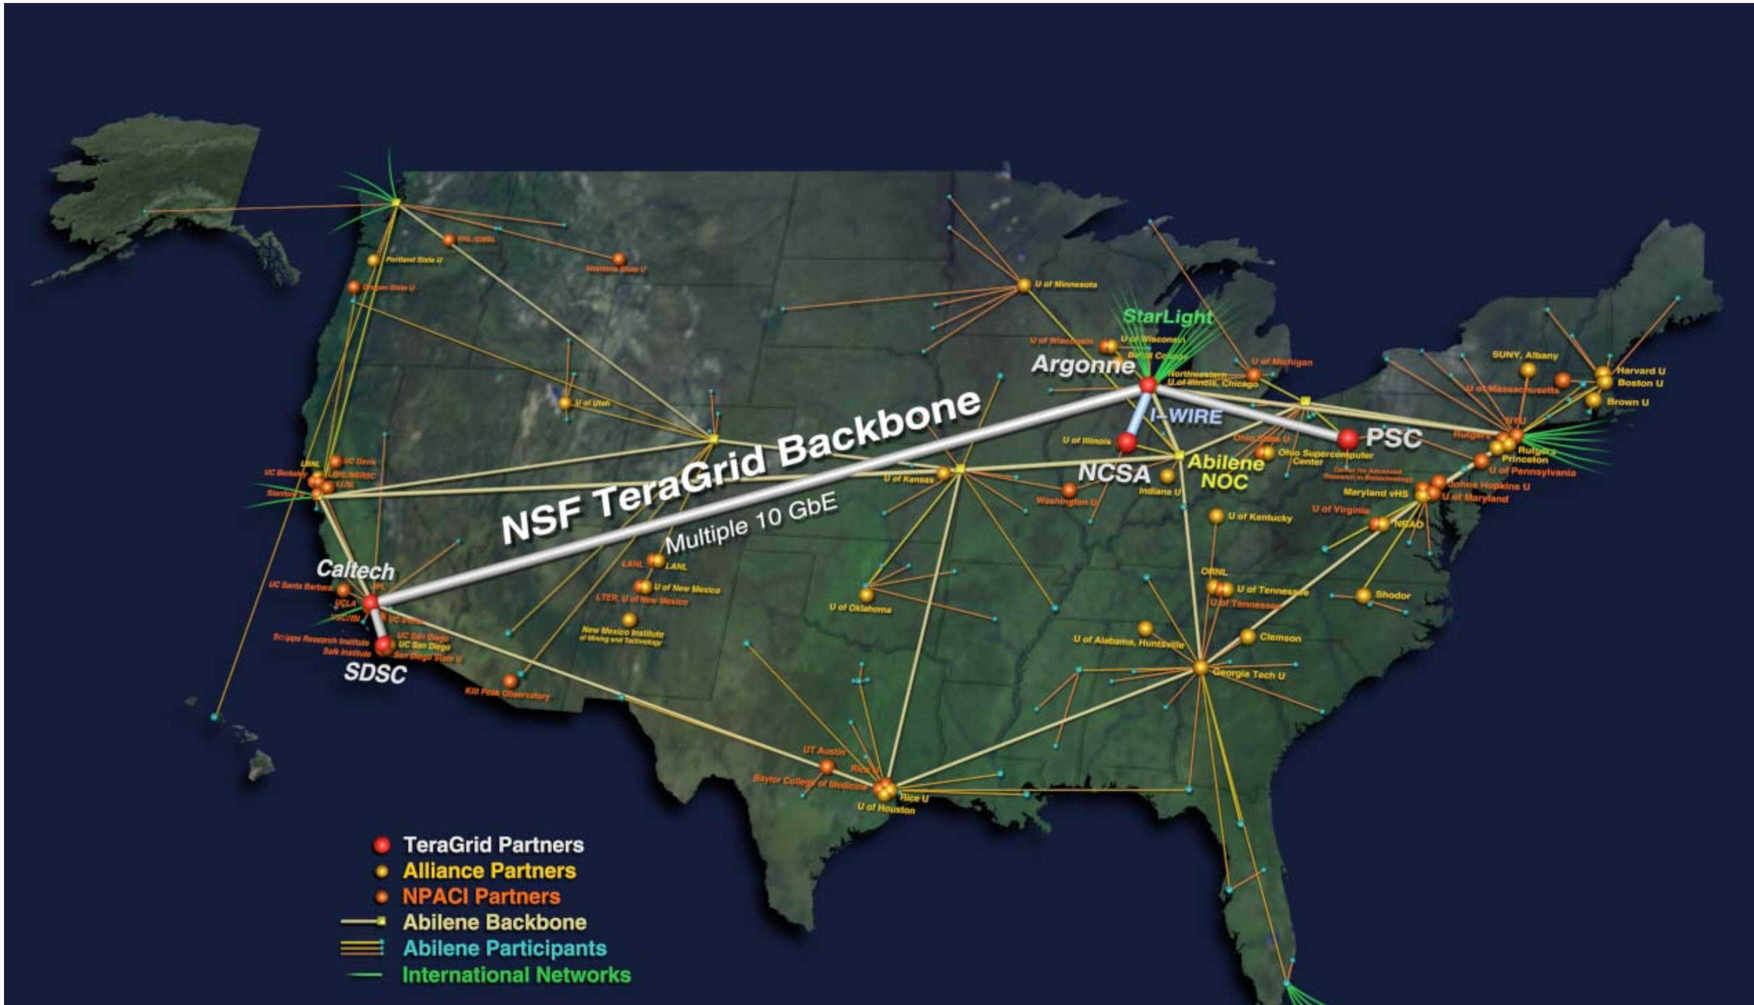

Source: http://www.teragrid.org

Virtual Organizations

- Resources may be shared by many groups or individuals
- Controlled allocation of resources is critical
- Resources may be allocated to "**virtual organizations**"
- A virtual organization (VO) defines:
  - A **set of resources** allocated for the purpose of completing a task
  - The **set of users or organizations** that may access the resources
  - The **set of conditions** under which sharing occurs: time, scope, …
- Properties of a virtual organization:
  - Usually **transient**:
    - Exists only until task is completed
  - May include only a **subset** of a Grid's total resources:
    - A Grid can include multiple VOs
    - VOs can share resources

What is Needed?

- Trust & Cooperation
- Sharable resources:
  - Systems, data, storage, applications, …
- Access to the resources:
  - Common network
  - Common protocols
  - Secure communication
- A way to discover the resources:
  - Information services
- A way to control the sharing:
  - **Who** has access?
    - Authentication
  - **What** can they access?
    - Authorization
  - **When** do they have access?
    - Allocation
- Middleware:
  - Tools
  - Application Programming Interfaces (APIs)
  - Software Development Kits (SDKs)

Some Challenges

- Security:
  - External users must be given access to systems
  - Firewalls
- Support:
  - Availability of remote system administrators
- Application considerations:
  - Architecture (parallelize?)
  - Portability
  - Availability
  - Interoperability
  - Security
  - Data access
  - Licensing model
- System management:
  - Compatibility and interoperability issues
  - User management:
    - Must maintain accounts for a more transient user population
    - Must control access to specific resources (authorization)
  - Maintain/publish accurate information about resources

Agenda

- Introduction to Grid Computing
- The Globus Project
  - Globus Toolkit 2.4
  - OGSA/OGSI & Globus Toolkit 3.0
- Utility Computing and the Adaptive Enterprise (AE)
- HP's Utility Data Center (UDC)
- Utility Data Center and the Enterprise Grid

The Globus Project

- Based at Argonne National Lab & University of Chicago
- Principals:
  - Ian Foster
  - Steve Tuecke
  - Carl Kesselman (USC/ISI)
- Objectives:
  - Develop and promote standard Grid protocols, APIs & SDKs
  - Collaborate with Grid projects in science and industry
- The Globus Toolkit:
  - Open source Grid middleware toolkit for building:
    - Scalable Grid infrastructures
    - Grid-enabled tools and applications
  - Uses public protocols & APIs
  - Current versions:
    - GT2.4
    - GT3.0

http://www.globus.org

Globus Components

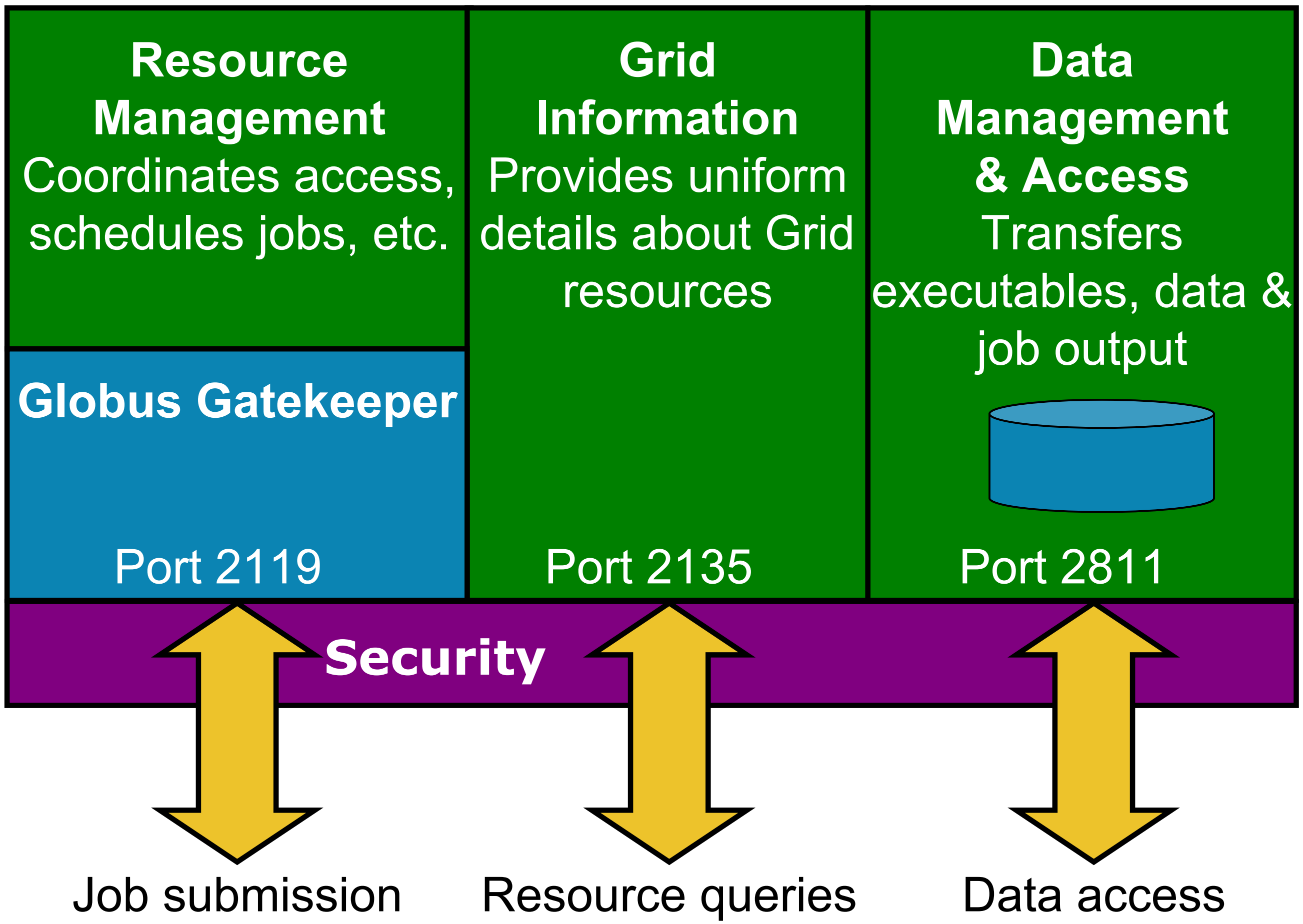

Security

- Security is critical:
  - External users have access
  - Greater risk of unauthorized access
- Must ensure that:
  - Only authorized users have access
  - A users can only access resources allocated to his/her VO
  - Inter-node communication is secure
- Certificate Authority (CA):
  - Verifies user's identity
  - Issues/revokes certificates
  - Grid partners must agree on CAs
  - Must establish CA procedures:
    - Certificate revocation
    - Audit process
- **G**rid **S**ecurity **I**nfrastructure:
  - PKI for credential checking
  - Proxy credentials and delegation for secure single sign-on
  - SSL/TLS for authentication and secure communication
  - Uses OpenSSL APIs & commands
  - Kerberos also supported
- CA facilities:
  - OpenSSL includes CA tools
  - Globus provides a Simple CA package based on OpenSSL

Resource Management

- Resource management system provides **secure** and **controlled access** to Grid resources
- Responsibilities include:
  - Authentication & authorization
  - Reservation & allocation
  - Job submission and monitoring
- **G**rid **R**esource **A**llocation & **M**anagement:
  - Gatekeeper provides access to resources
  - Runs & monitors jobs:
    - Internal “fork” job manager
    - Can interface with LSF, PBS, …
  - **R**esource **S**pecification **L**anguage specifies:
    - Hardware resources
    - Execution environment
  - Co-allocation facility (DUROC) can control parallel jobs
  - Currently no reservation or accounting facilities

Data Management

- Data management is critical to Grid applications:
  - Preparing a job may involve copying the application and data to each node
  - Efficient data access can reduce time of execution and analysis of results
  - Unnecessary transfer of large datasets impacts network
- **G**lobal **A**ccess to **S**econdary **S**torage:
  - GASS API provides transparent access to data in any location
  - Uses URL-encoded file names:
    - `file:///home/jt/my_data`
    - `gsiftp://remote_system.hp.com/~jt/my_data`
  - Suitable for accessing/transferring programs, small datasets, logs, …
- GridFTP:
  - Extends traditional FTP
  - Uses proxy credentials
  - Secure transfer
  - Partial transfers
  - Parallel transfer streams
  - Restarts failed transfers
- Data replication services

Information Services

- Grid users need to know:
  - What nodes are available
  - Platform type
  - O/S name & version
  - Storage capacity
  - CPU count, speed, utilization
  - Installed applications
  - etc.

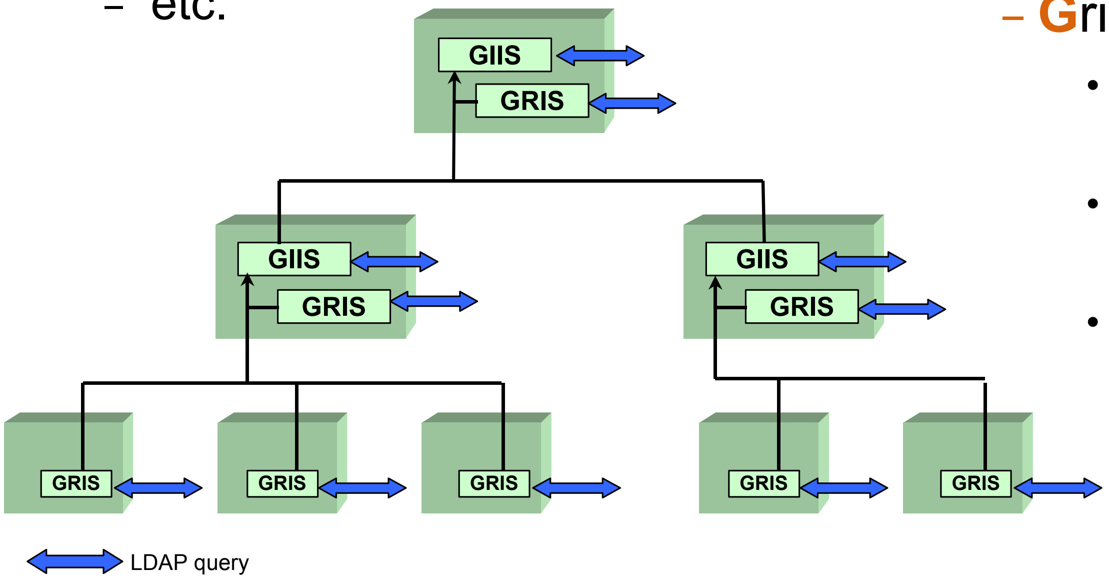

- Monitoring & Discovery Service:
  - Query using LDAP notation
  - Grid Resource Information Service:
    - Local directory server
    - Can be queried directly
    - Reports to a GIIS
  - Grid Index Information Service:
    - Aggregates information from multiple GRIS servers
    - Can report aggregated information to another GIIS
    - GIIS server nodes can be arranged in a hierarchy

Agenda

- Introduction to Grid Computing
- The Globus Project
  - Globus Toolkit 2.4
  - OGSA/OGSI & Globus Toolkit 3.0
- Utility Computing and the Adaptive Enterprise (AE)
- HP's Utility Data Center (UDC)
- Utility Data Center and the Enterprise Grid

Next Generation Grids

- GT2 evolved with focus on "stovepipe" protocols:
  - SSL/TLS/PKI for security
  - LDAP for discovery & monitoring
  - GridFTP for data transfer
  - etc.
- Widely adopted by research community
- Less impact so far in the commercial world

- The Physiology of the Grid, Foster, Tuecke, Kesselman, Nick, 2002
  - **O**pen **G**rid **S**ervices **A**rchitecture
  - Builds on Web services technology
  - Views a Grid as a set of services that can be:
    - extended
    - aggregated
  - Defines the core interfaces and behaviors that Grid services should possess
  - Emphasizes needs of commercial Grid applications

Web Services

- Distributed computing paradigm:
  - Application (not human) oriented
  - Use simple Internet-based protocols
  - Independent of:
    - Platform
    - Programming language
  - Services are typically persistent
  - Scalable:
    - Service instances can be added to meet fluctuating demand
- Key standards:
  - XML – language for representing structured data
  - WSDL – describes Web service interfaces & protocols
  - SOAP – protocol for exchanging XML messages
  - WS-Inspection – conventions for locating service descriptions

OGSA & Grid Services

- Open architecture:
  - Extensible
  - Vendor-neutral
  - Community standards
- Converges Web services & Grid:
  - Add Grid concepts:
    - Lifecycle, state, …
    - Standard interfaces and behaviors
  - Persistent **or** transient
- Enables virtualization & sharing of resources
  - All resources represented by services
  - Provides a common interface for a resource, hides implementation
- **A Grid service is a Web service!**

- **O**pen **G**rid **S**ervices **I**nfrastructure:
  - Specifies an infrastructure for defining & implementing OGSA interfaces & behaviors
  - Defines core services:
    - Base Grid service (required)
    - Service creation & termination
    - Data operations
    - Notification
    - etc.
  - No Grid required!
- Globus 3.x is **an implementation** of OGSA:
  - GT3.0 (Linux) released June 2003
  - Incremental releases planned
  - Functionally compatible with GT2
  - Java-based
  - Larger system footprint

The Grid Service

Interfaces/behaviors + service data = Grid Service

Required:
- Introspection (service data)
- Explicit destruction
- Soft-state lifetime

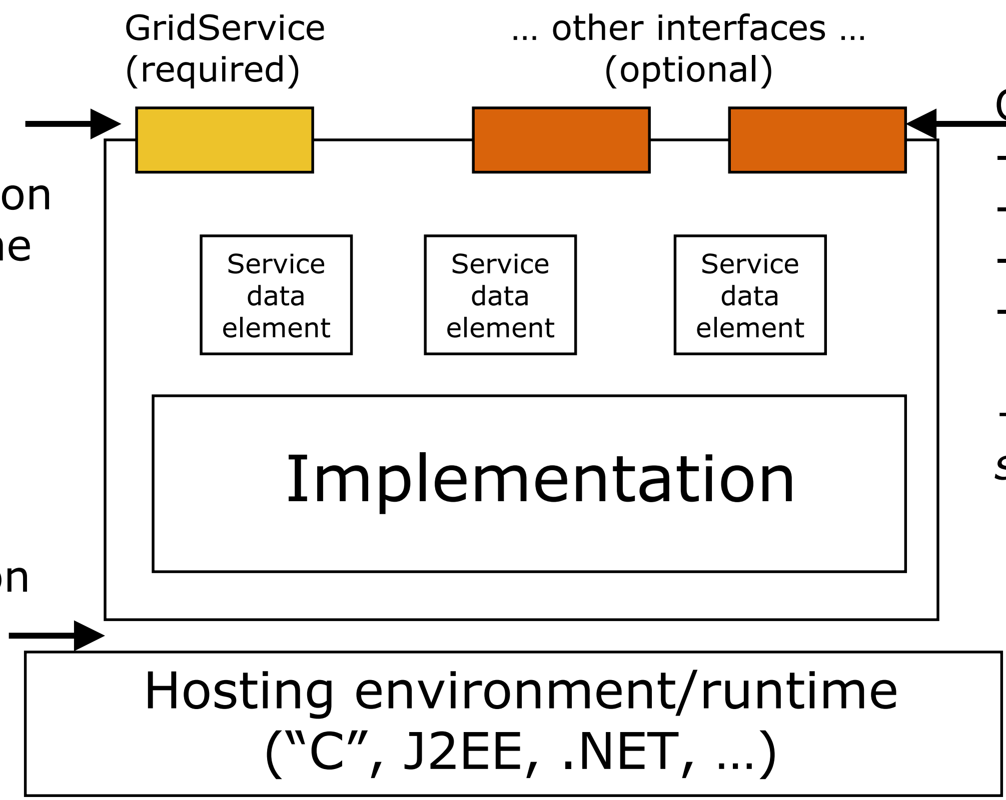

Optional:
- Service creation
- Notification
- Registration
- Collections

+ *application-specific interfaces*

Binding properties:
- Authentication
- Reliable invocation
- Transactions
- QoS

Source: Ian Foster

Globus on HP Platforms

- Globus Toolkit 2.4 is available for:
  - HP-UX on PA-RISC & Itanium
  - Linux on IA32, Itanium & Alpha
  - Tru64 UNIX on Alpha
- Download from:
  - www.globus.org
  - www.hp.com/products/globus
- Globus Toolkit 3.0
  - Currently porting to all HP platforms
  - Watch for announcements!

Agenda

- Introduction to Grid Computing
- The Globus Project
  - Globus Toolkit 2.4
  - OGSA/OGSI & Globus Toolkit 3.0
- Utility Computing and the Adaptive Enterprise (AE)
- HP's Utility Data Center (UDC)
- Utility Data Center and the Enterprise Grid

Driving Factors for Grids in Enterprise IT Environments

- **Goals in enterprise IT environments:**
  - **Agility** – adapt IT infrastructure faster to changing business needs
  - **Collaboration** – enable cooperation and sharing across organizations
  - **Return on Investment** – increase asset utilization and lower operational costs

- **Goals translate into technologies:**
  - **Virtualization** of resources decouples applications from hardware
  - **Federation** of resource pools for sharing and collaboration within and across organizations — **Grid**
  - **Automation** of operation and management

Utility Computing

- Resources become commodities
  - Users access **services**, and are unaware of:
    - Platforms
    - Applications
    - Management
  - Seamless resource sharing across resource pools (data centers)
- Utility system:
  - Manages underlying resources automatically
  - May change resource assignments during operation:
    - Add/remove
    - Scale capacity up/down
    - Migrate applications
- Enables pay-per-use model

Adaptive Enterprise

- Infrastructure responds to changing conditions:
  - Load, failures, business needs
- Partially or fully automated

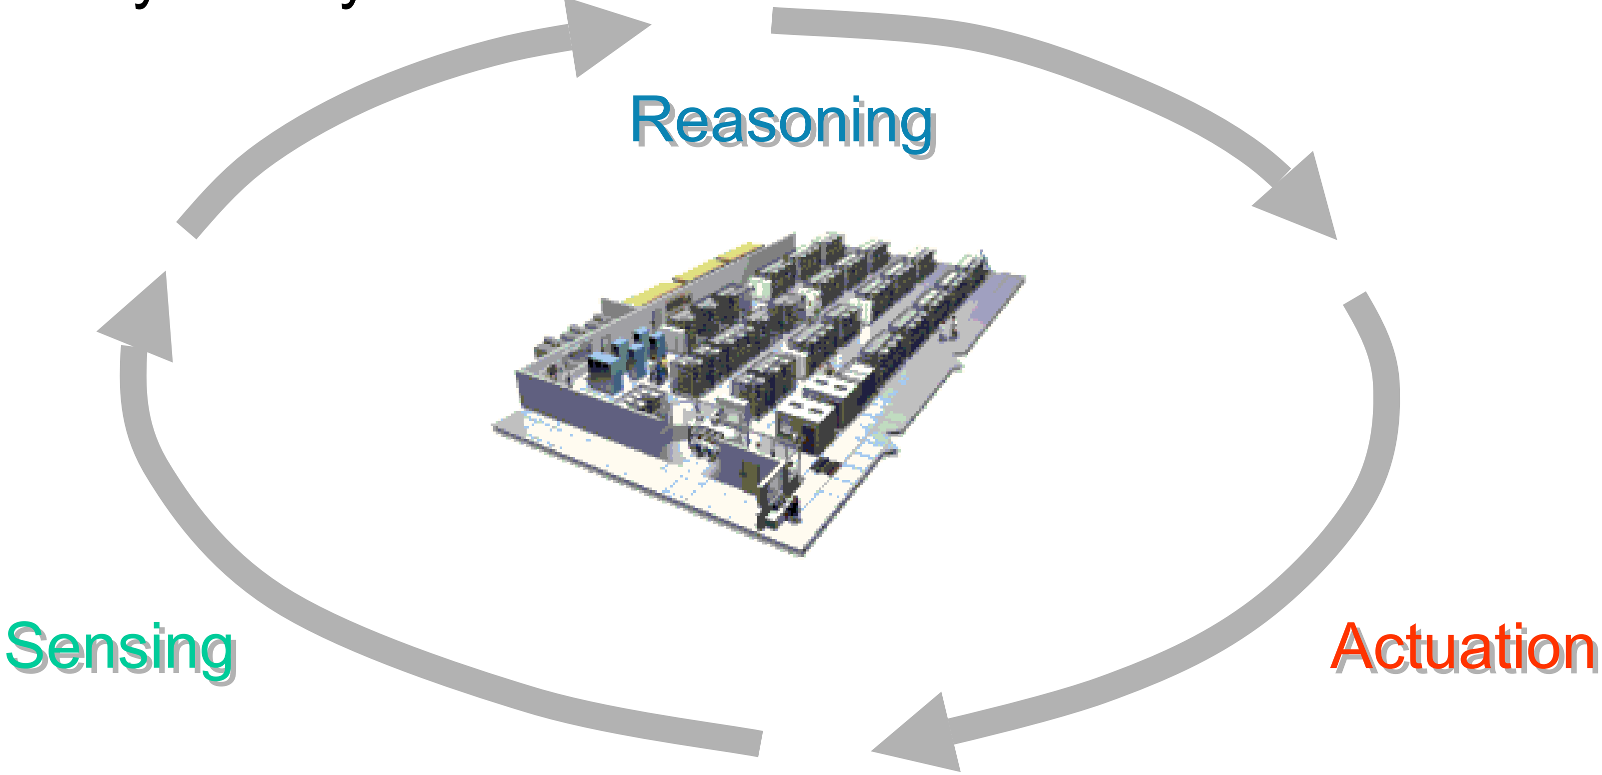

Agenda

- Introduction to Grid Computing
- The Globus Project
  - Globus Toolkit 2.4
  - OGSA/OGSI & Globus Toolkit 3.0
- Utility Computing and the Adaptive Enterprise (AE)
- HP's Utility Data Center (UDC)
- Utility Data Center and the Enterprise Grid

HP's Utility Data Center

HP UDC at the Palo Alto Research Labs. 12/02

UDC Components

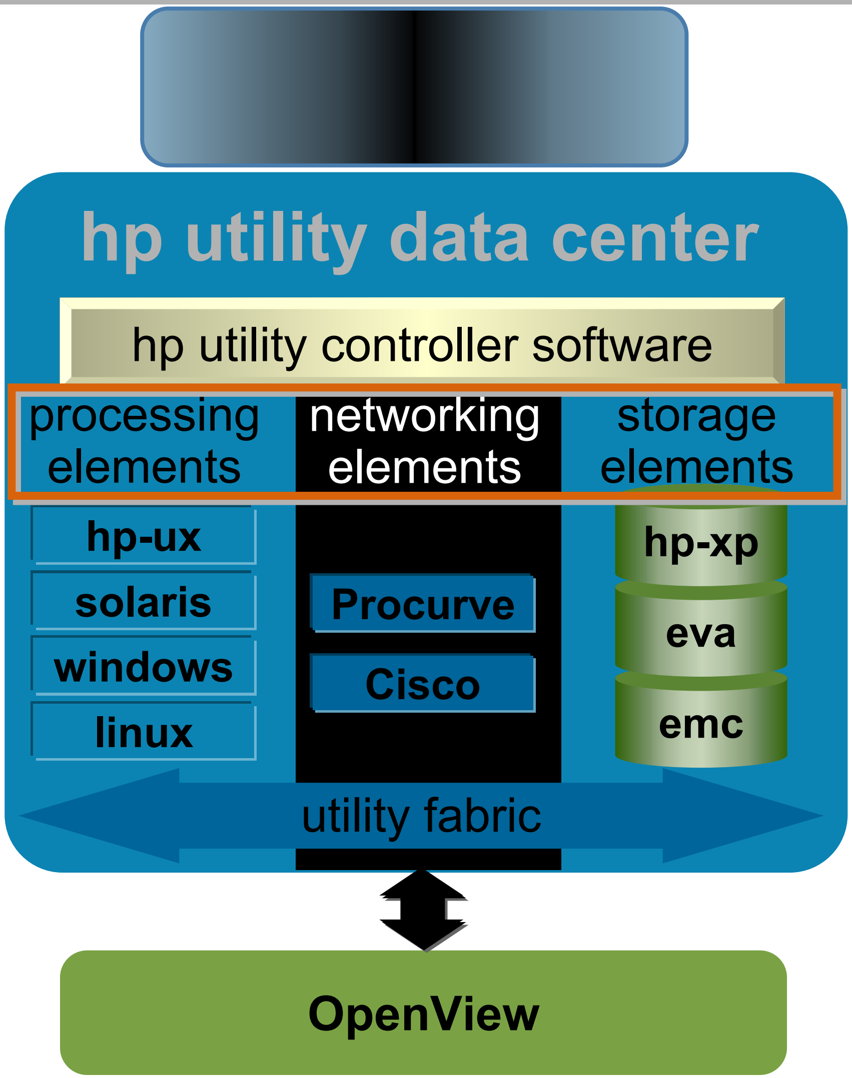

- Utility Controller Software
  - Manages service templates
  - Controls resources
  - Reacts to workload changes and failures
- Virtual Server Pools
  - Heterogeneous server environments
  - HP servers optimized for UDC
  - Protect your current investment
- Virtual Network Pools
  - Standards-based VLANs
  - Flexible and robust network infrastructure
- Virtual Storage Pools
  - HP XP and EVA storage offer flexible 'network-based' virtualization
  - EMC Symmetrix
  - Integration with OpenView for storage management

Creating a Service with the Utility Data Center

1. Architect new service:

- business case
- service growth projection
- SLO requirement
- availability
- security needs
- time to implement

3. Create the service

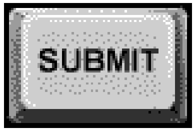

- automatically find and allocate resources
- auto-configure network and storage
- auto-configure firewall & load balancers
- auto-configure and boot the servers

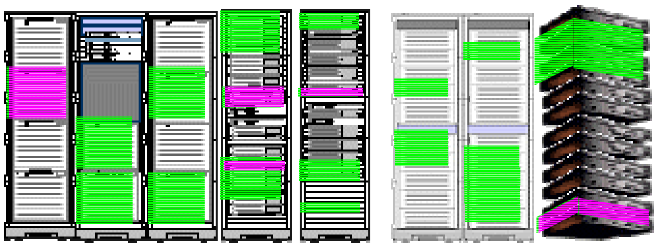

available

new service added

2. Build a service template:

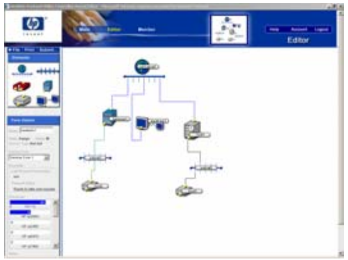

Svc "X"

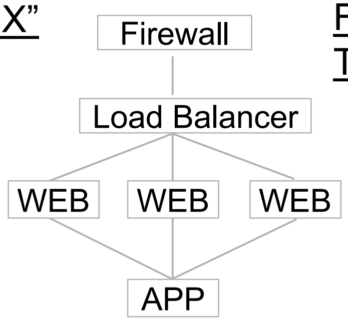

Farm "A" Template

Utility Controller Portal

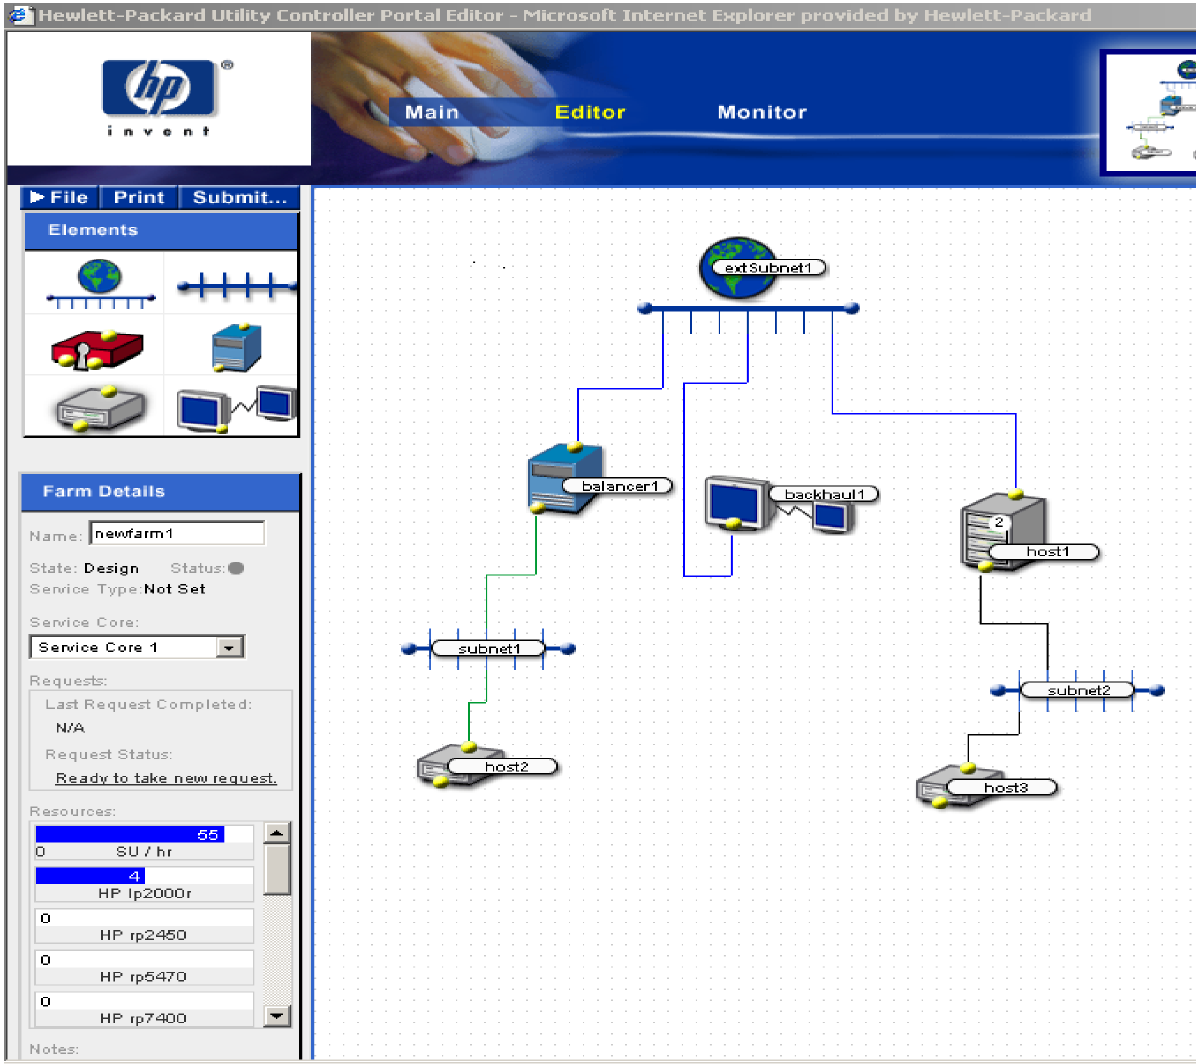

- Web-based interface
- Manages all UDC resources
- Enables configuration of new IT Services
- Enables changes to service definitions
- Automatic fault mgt. and rectification
- Tracks usage

Utility Data Center in Action

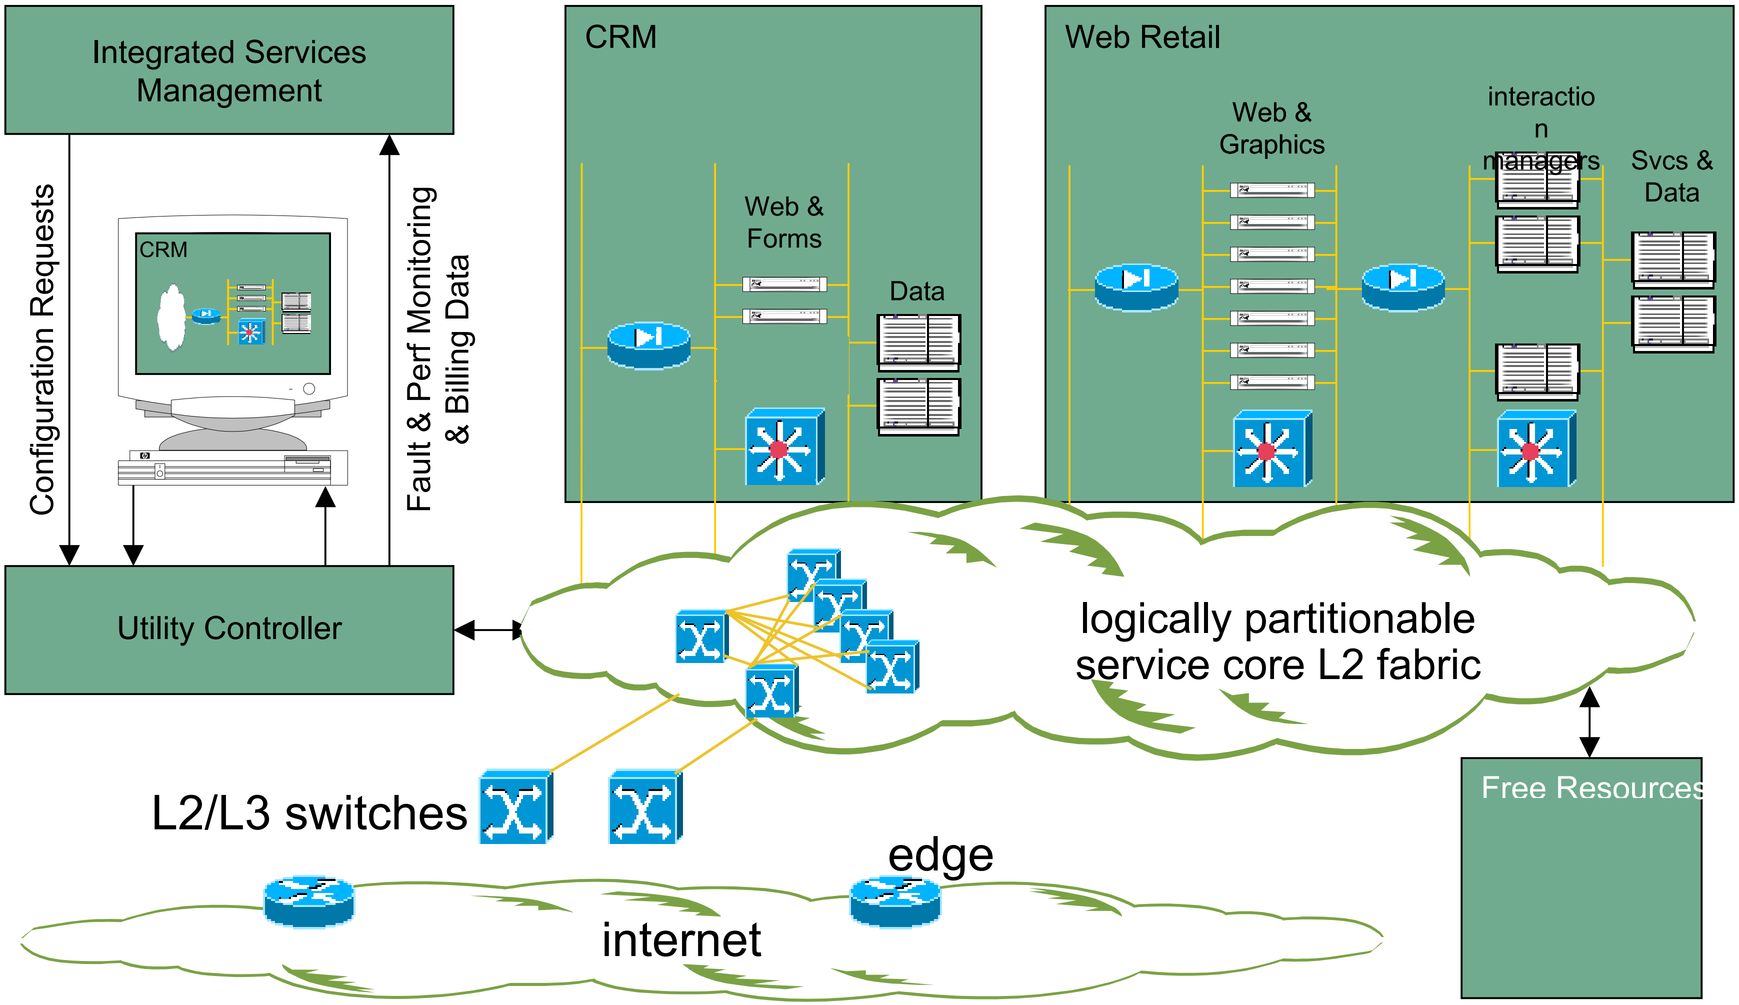

Summary: HP Utility Data Center – enabling Adaptive Enterprise

HP Utility Data Center is a complete solution for virtualizing data center environments:

1. Resources are wired once to support their virtual, flexible allocation and reallocation
2. New applications and systems can be ignited within minutes
3. Utilization of servers, storage and network are significantly improved
4. Resources are virtualized, and optimize themselves to meet your service level objectives
5. Administrative and operational overhead is minimized and opportunities for error reduced

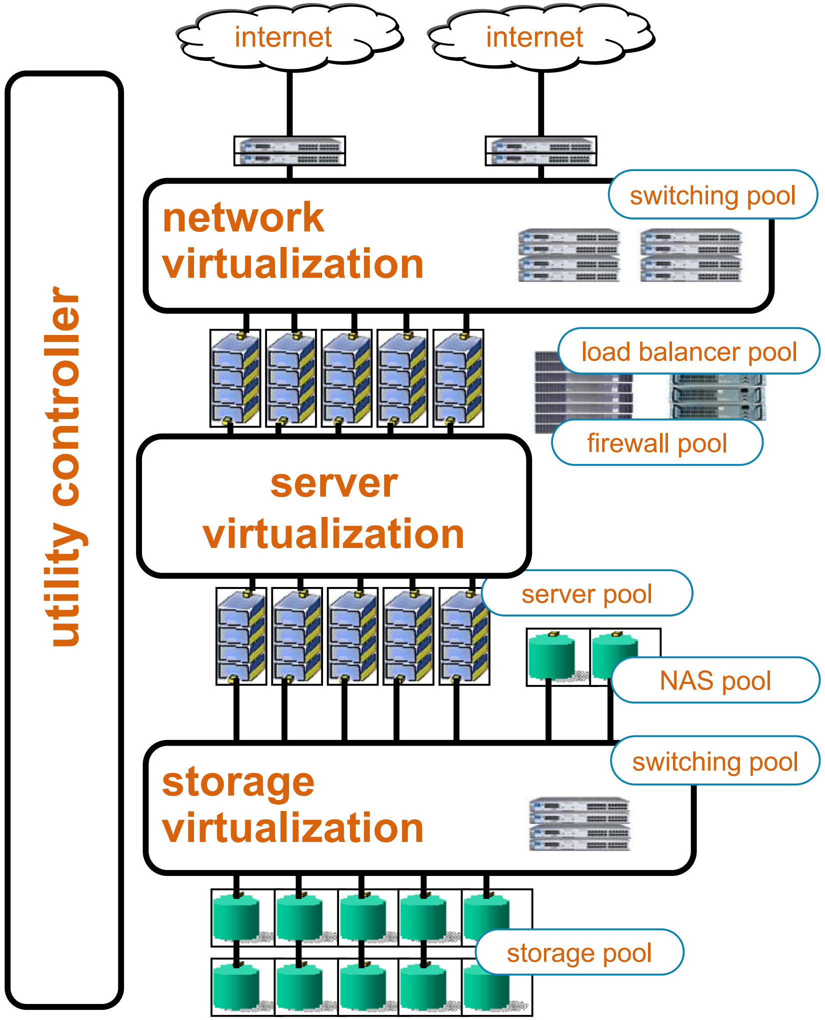

Agenda

- Introduction to Grid Computing
- The Globus Project
  - Globus Toolkit 2.4
  - OGSA/OGSI & Globus Toolkit 3.0
- Utility Computing and the Adaptive Enterprise (AE)
- HP's Utility Data Center (UDC)
- Utility Data Center and the Enterprise Grid

UDC and the Enterprise Grid

- Goal:
  - Give users access to resources in **multiple UDC resource pools**
  - Enable enterprise-wide:
    - Collaboration
    - Resource sharing

Note: Slides in this section illustrate a research application, and do not represent a product.

UDC & Grid: Mutual Benefits

UDC brings to the Grid:

- Sophisticated, flexible resource topologies
- Hardware-enforced isolation and protection
- Virtualized resource environment

The Grid brings to UDC:

- Ability to co-allocate resources from multiple UDCs
- Cross-organizational user credentials
- Uniform service model (OGSI)
- Community and standards focused on sharing of resources

Utility Data Center

Utility Data Center =
programmable pool
of data center resources

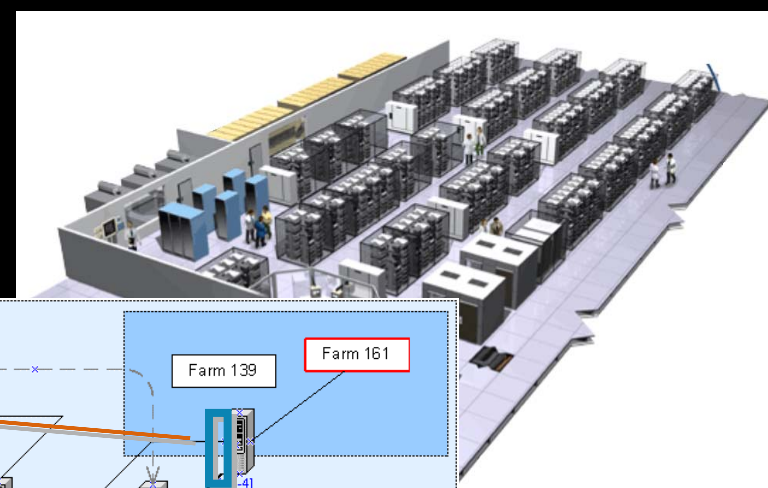

External UDC/XML
Control Interface

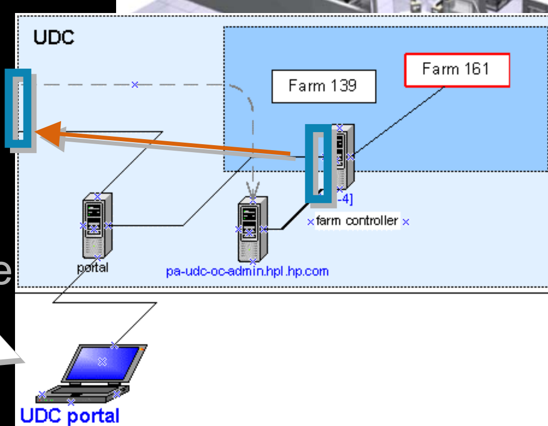

Operator interface

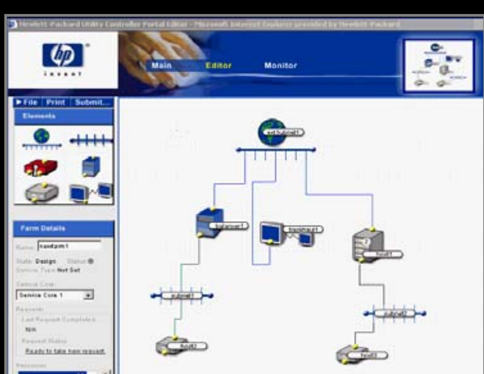

Grid Interface for UDC

Utility Data Center =
programmable pool
of data center resources

OGSI Interface

UDC GRAM

UDC/XML
Control Interface

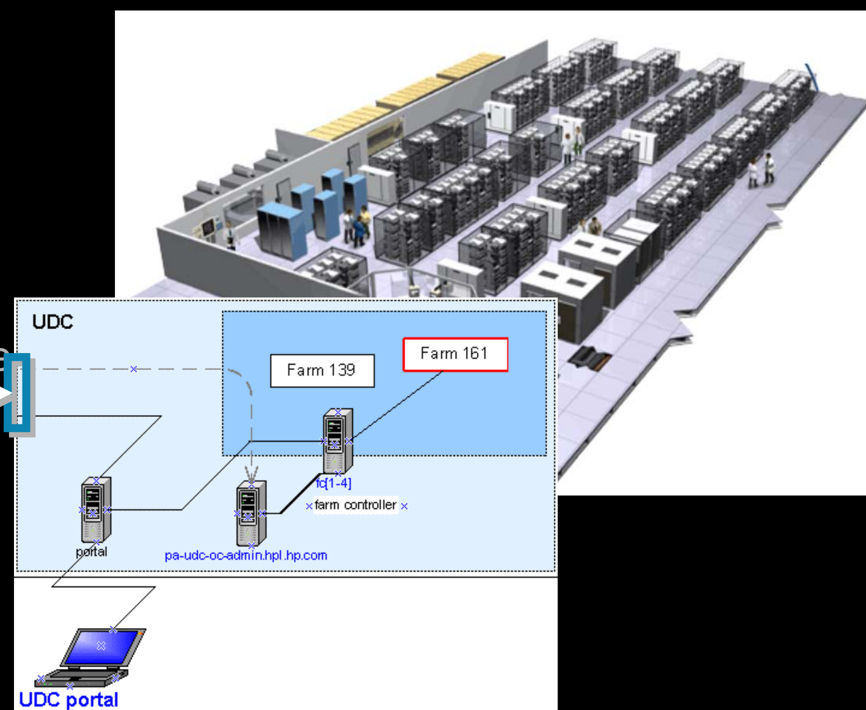

UDC GRAM =
Globus Gatekeeper +
UDC Adapter

GRAM = Grid Resource Allocation Manager

UDC in an Enterprise Grid

User Interface
(e.g. Resource Topology Designer)

OGSI Service Interface

Grid Interface:
- discover
- request
- manage resources

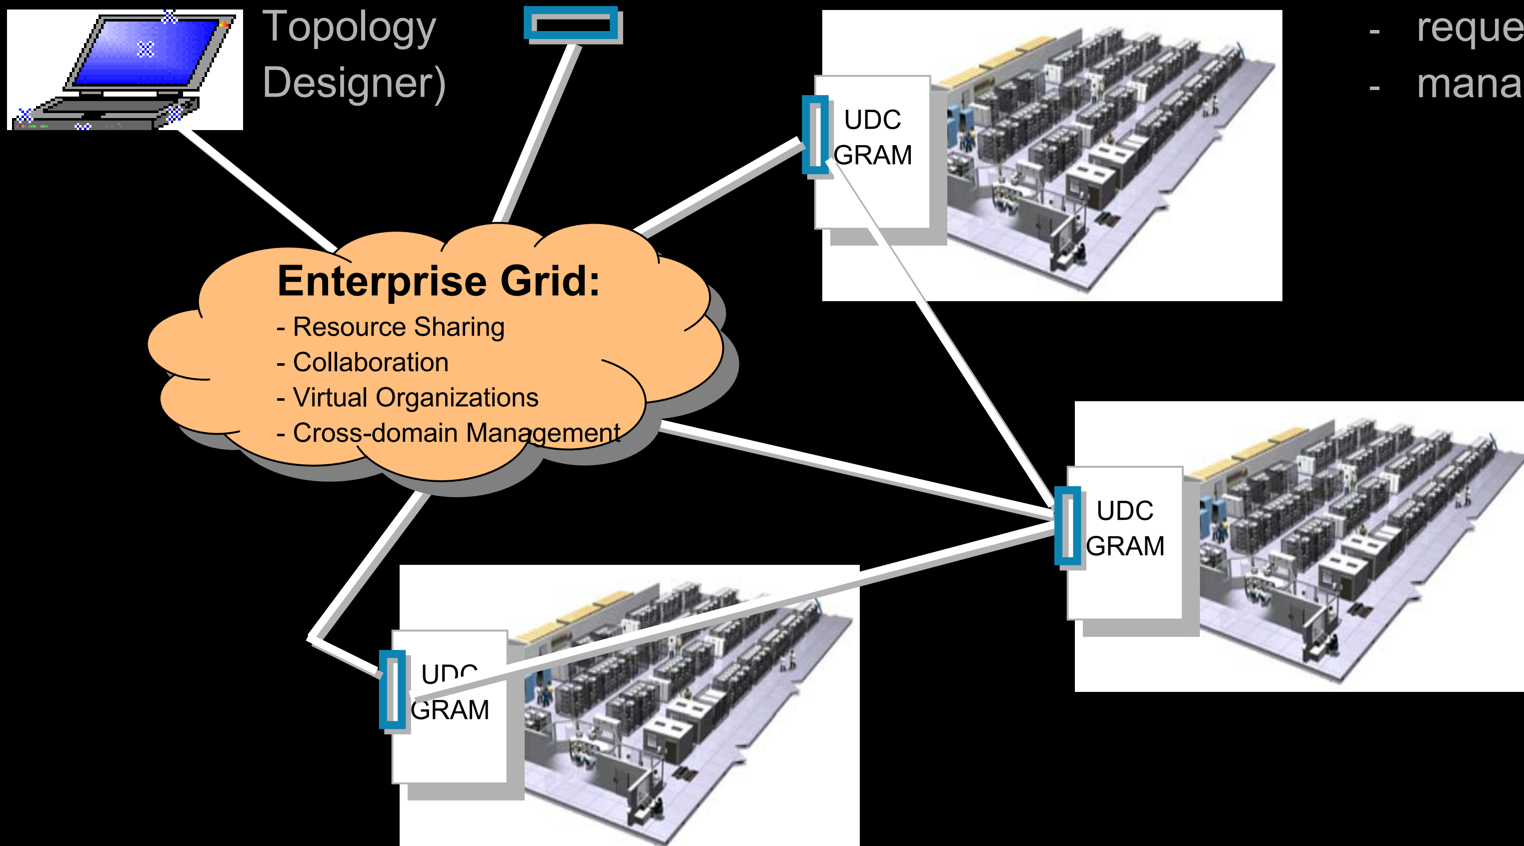

GRAM = Grid Resource Allocation Manager

Resource Request Construction

1. Designing a Resource Topology

2. Generating RSL for this topology

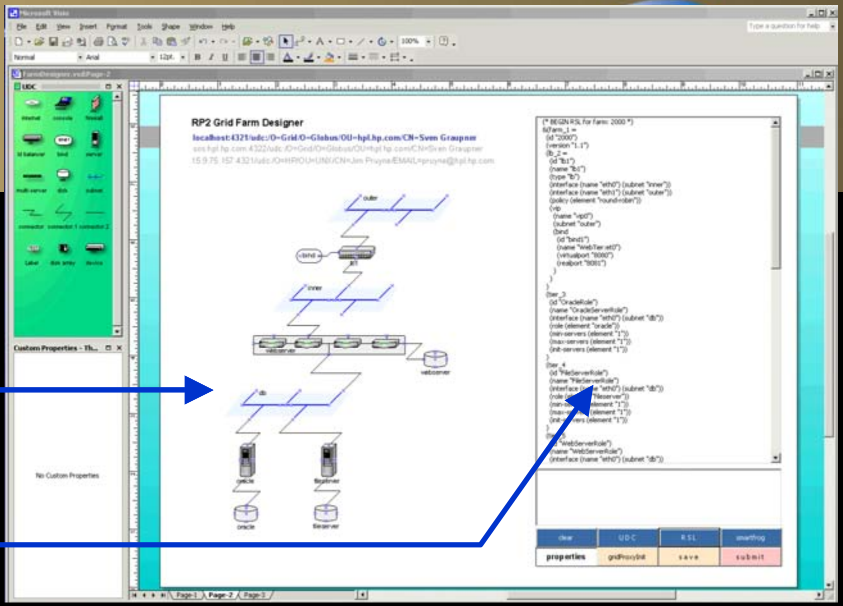

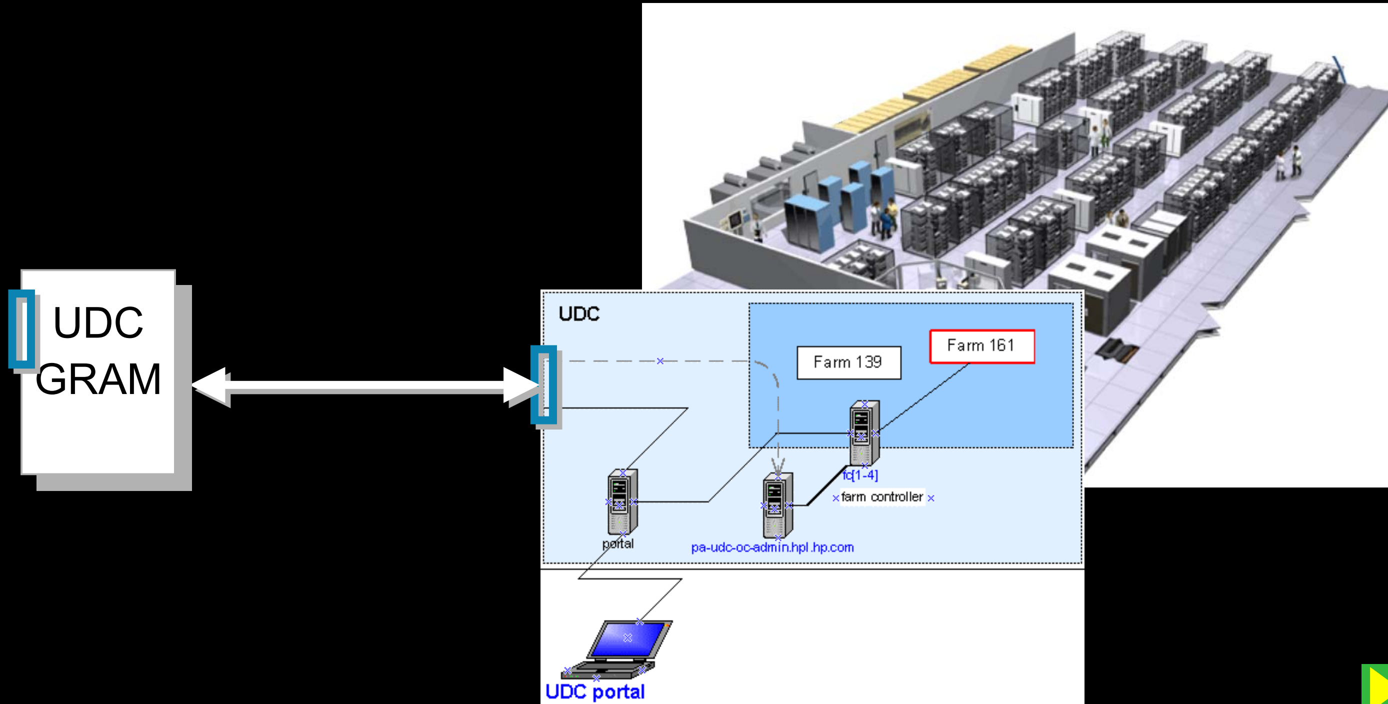

Resource Request Submission

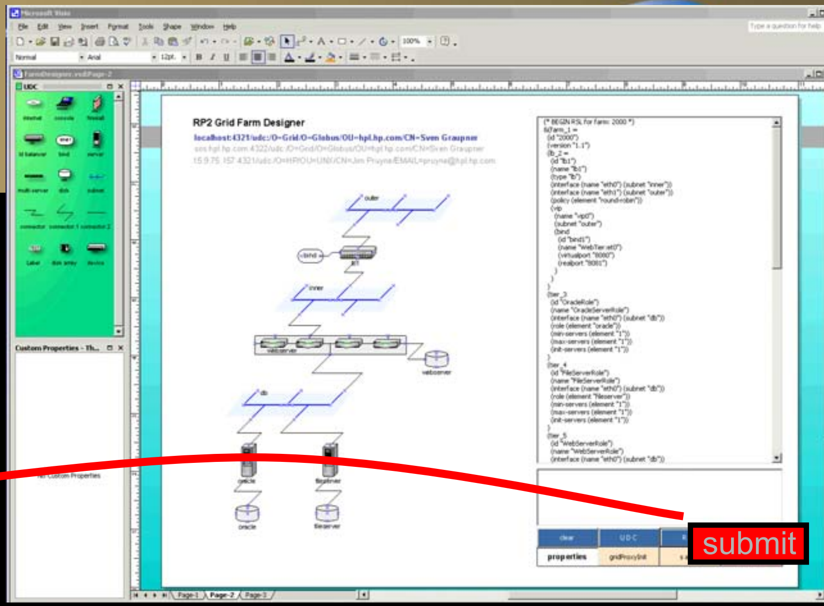

3. Submitting the Resource Topology to the UDC

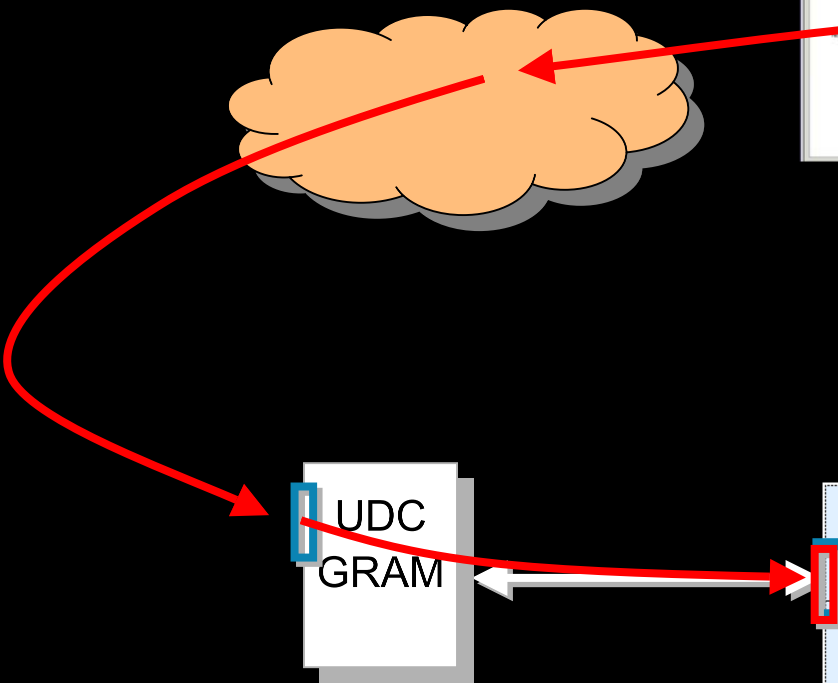

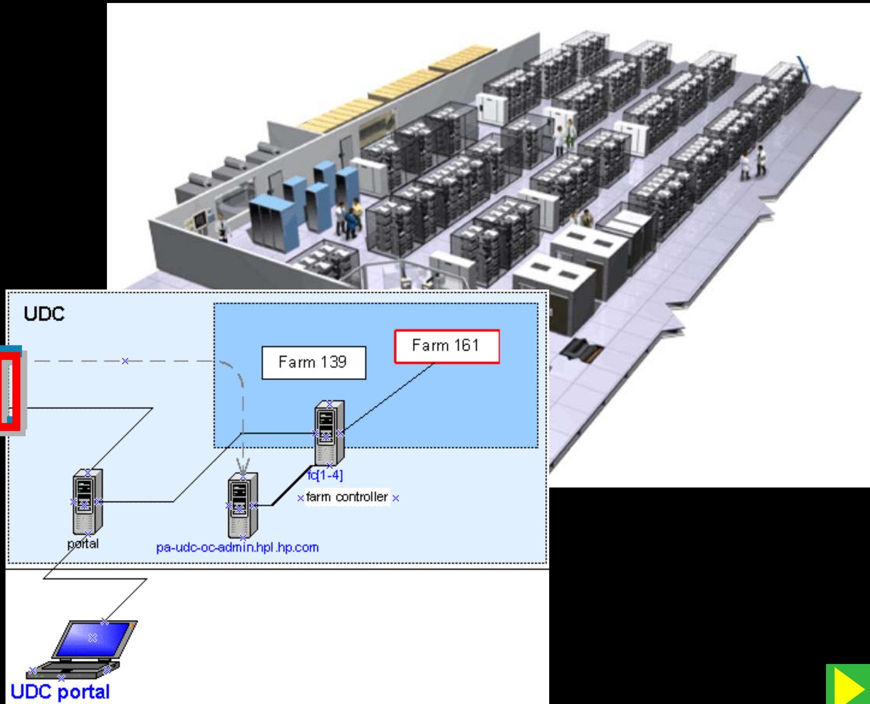

Resource Request Rendering as "Farm"

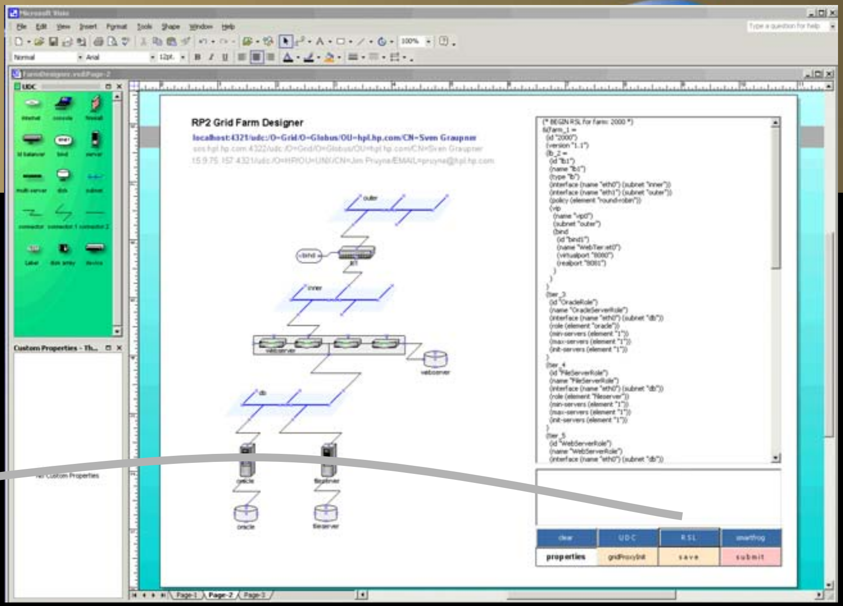

4. Instantiate the farm in the UDC

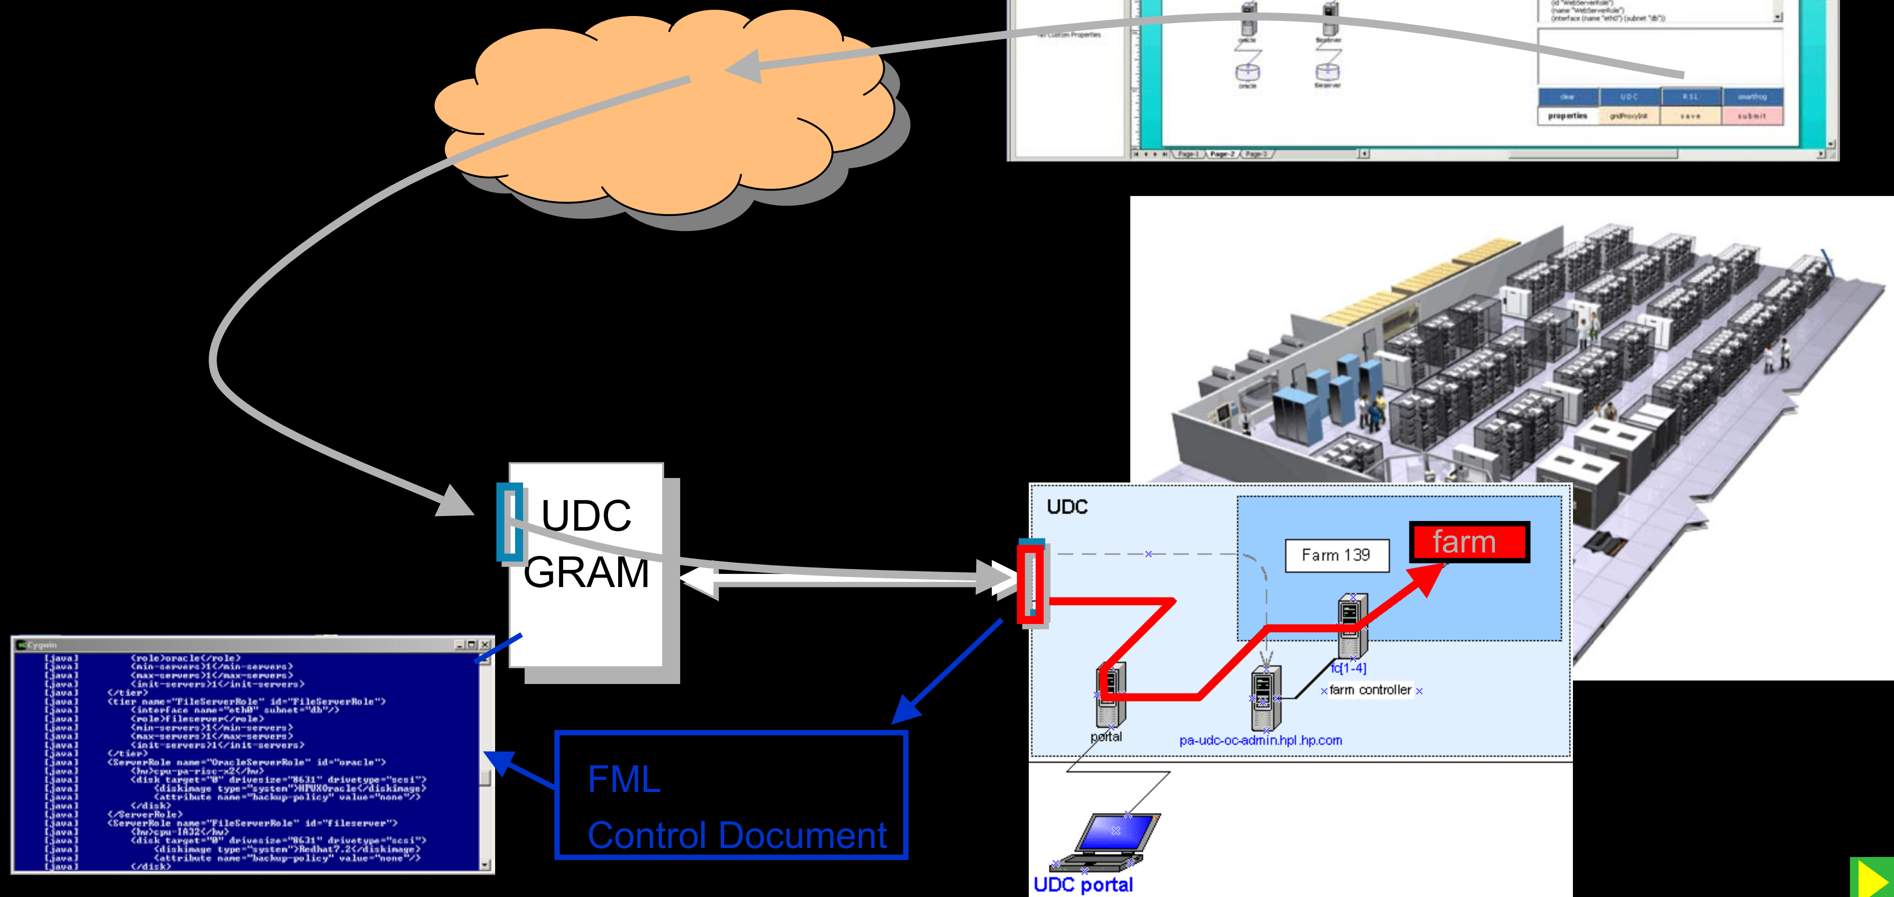

Grid Resource Topology Designer

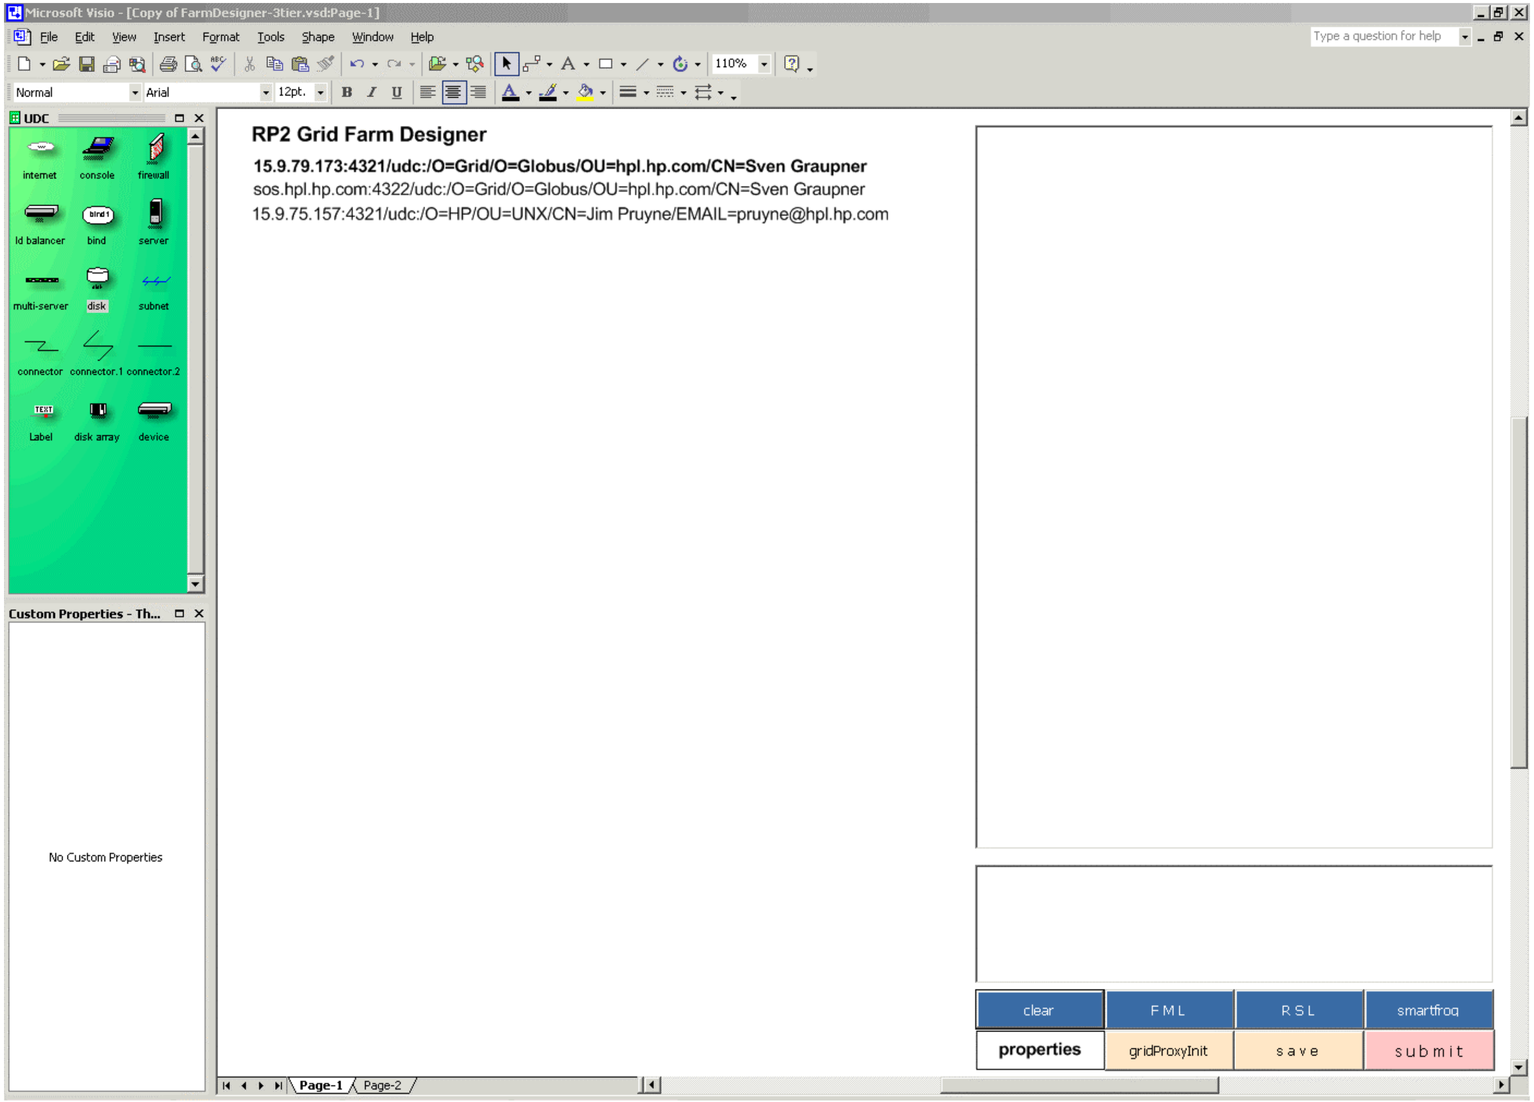

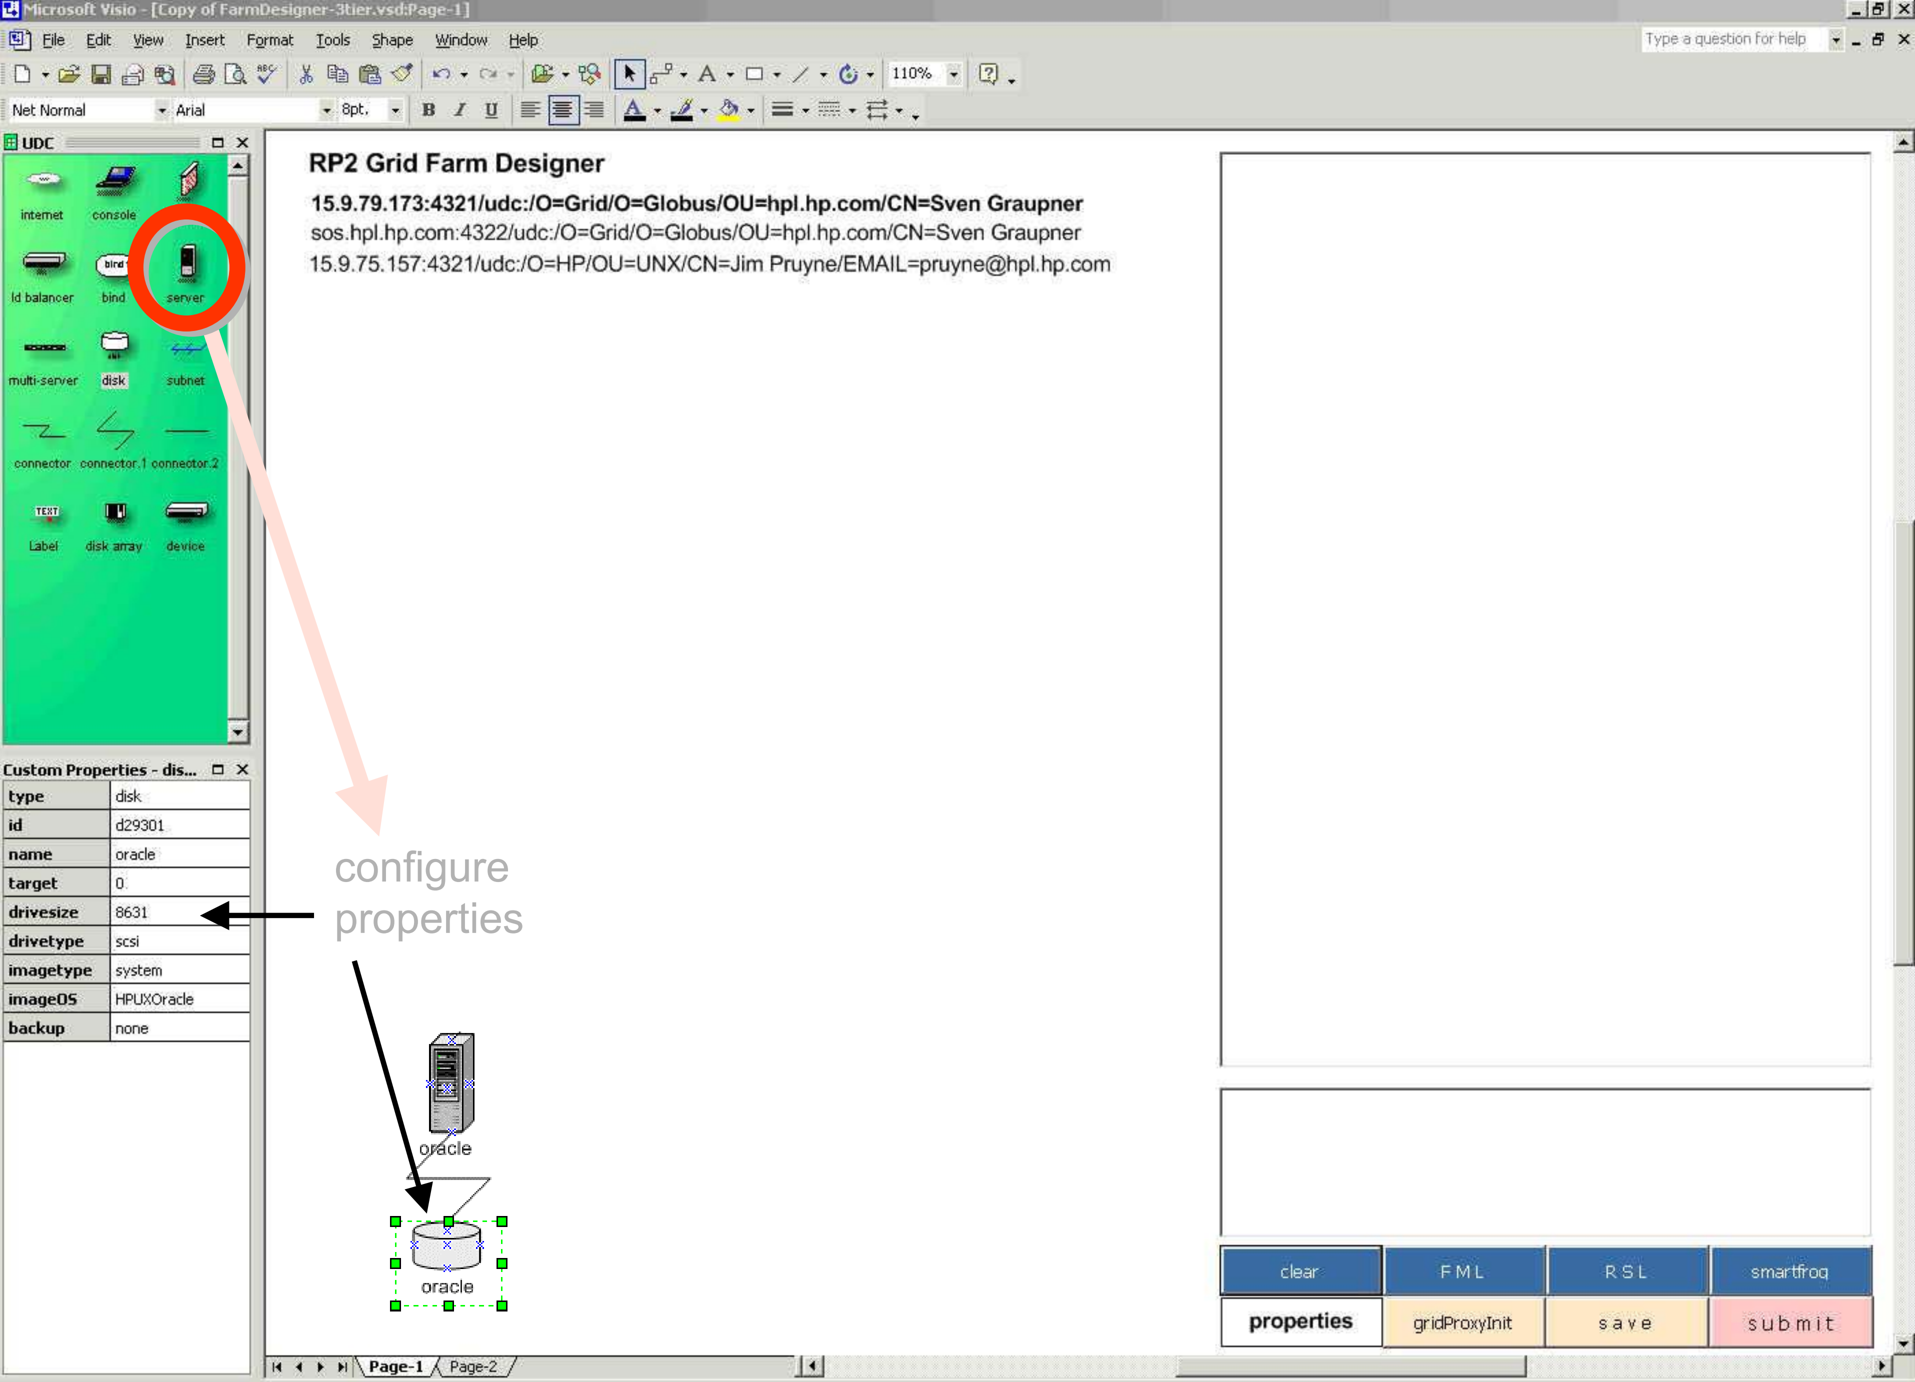
Microsoft Visio - [Copy of FarmDesigner-3tier.vsd:Page-1]
File Edit View Insert Format Tools Shape Window Help
UDC
internet
console
ld balancer
bind
server
multi-server
disk
subnet
connector
connector.1
connector.2
Label
disk array
device
Custom Properties - dis...
type disk
id d29301
name oracle
target 0
drivesize 8631
drivetype scsi
imagetype system
imageOS HPUXOracle
backup none
RP2 Grid Farm Designer
15.9.79.173:4321/udc:/O=Grid/O=Globus/OU=hpl.hp.com/CN=Sven Graupner
sos.hpl.hp.com:4322/udc:/O=Grid/O=Globus/OU=hpl.hp.com/CN=Sven Graupner
15.9.75.157:4321/udc:/O=HP/OU=UNX/CN=Jim Pruyne/EMAIL=pruyne@hpl.hp.com
configure properties
oracle
oracle
clear
FML
RSL
smartfroq
properties
gridProxyInit
save
submit
Page-1
Page-2

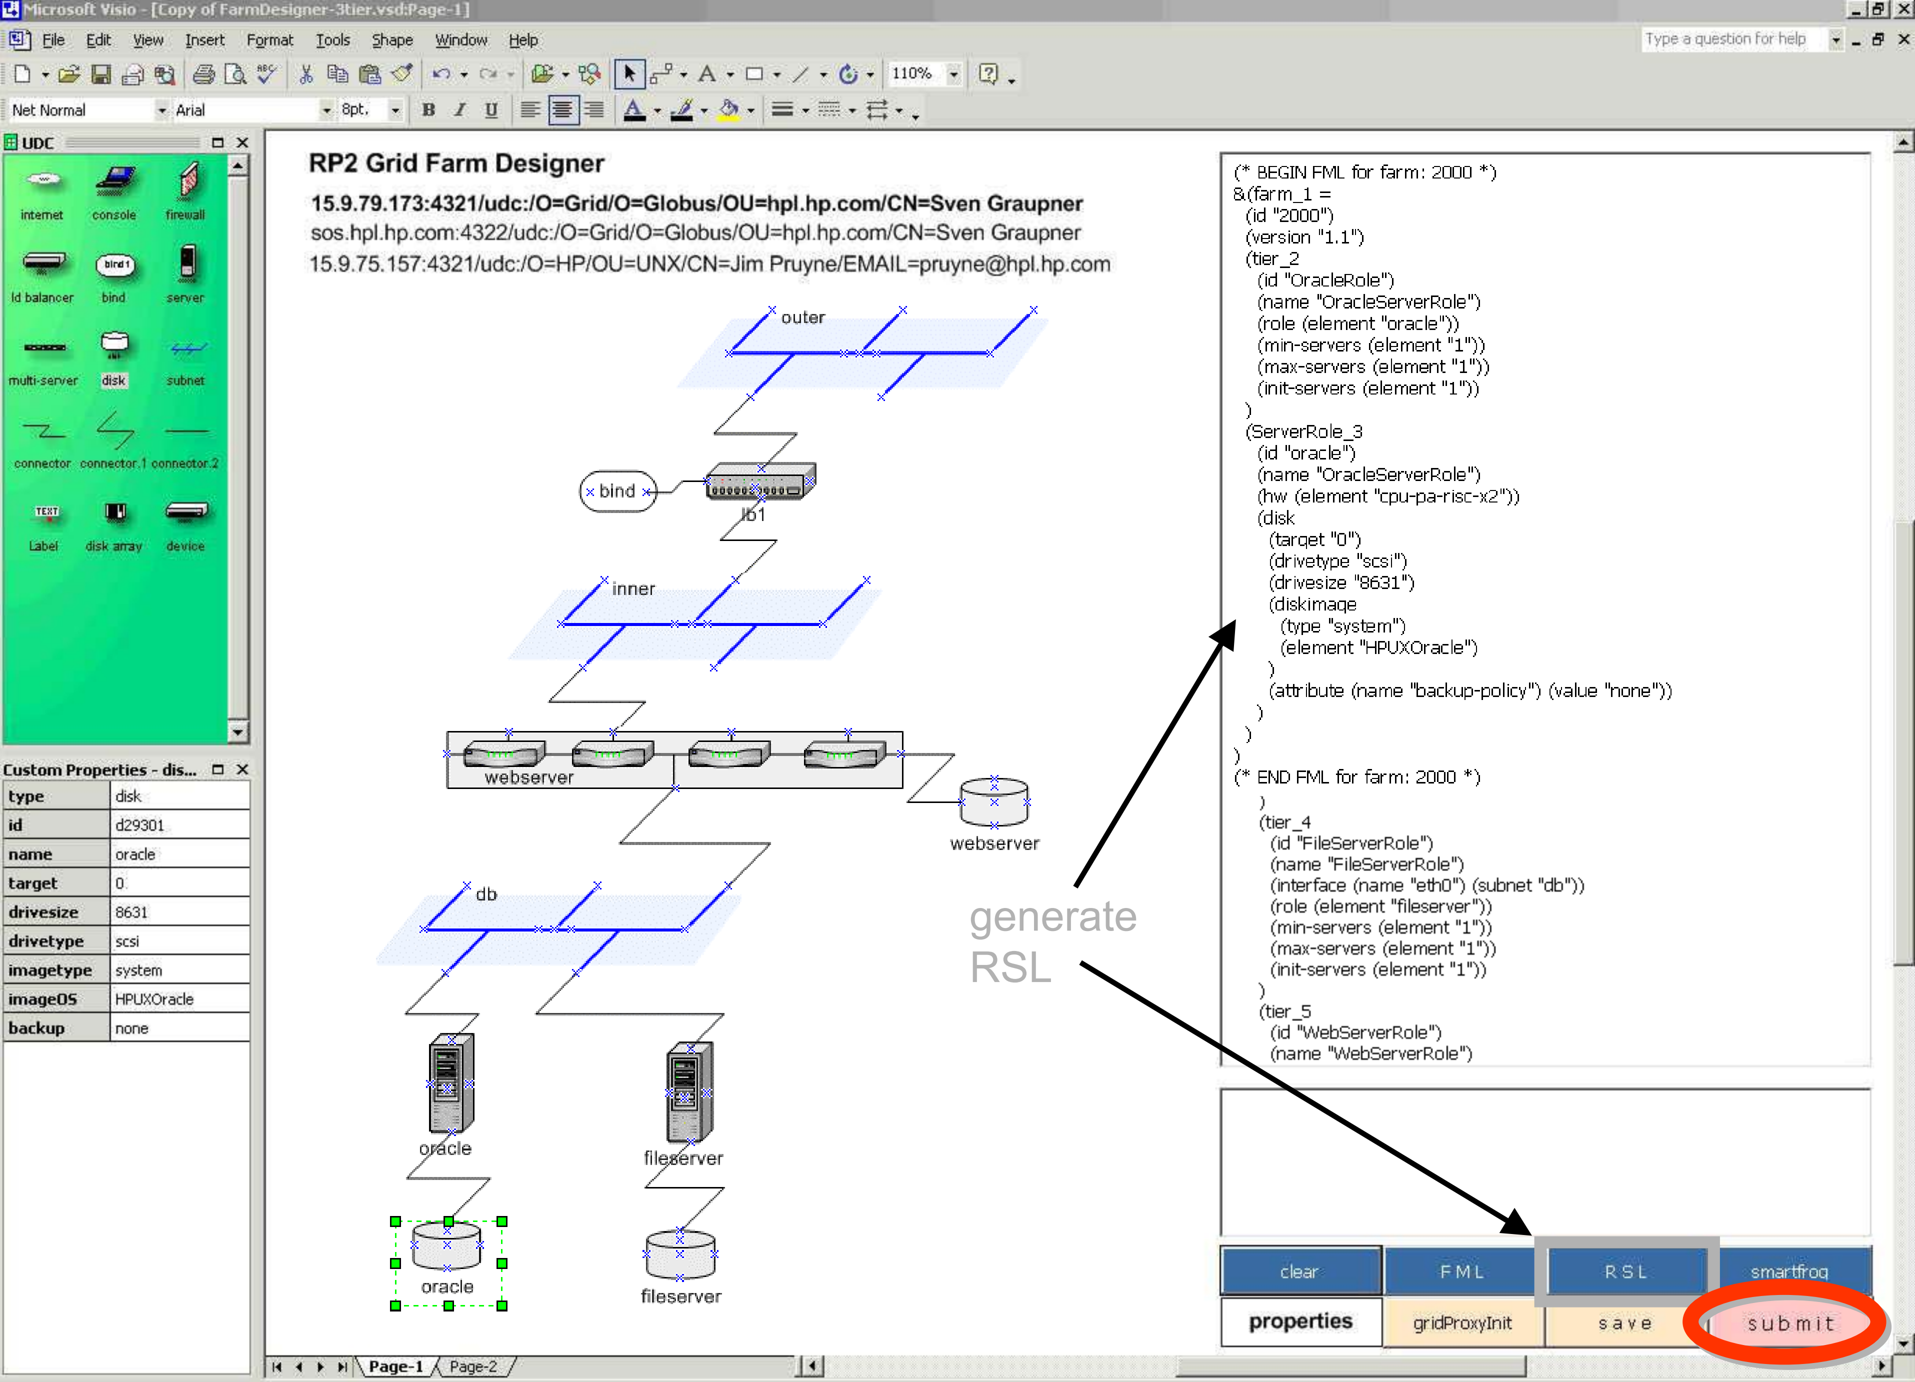
Microsoft Visio - [Copy of FarmDesigner-3tier.vsd:Page-1]
File Edit View Insert Format Tools Shape Window Help
UDC
internet
console
firewall
ld balancer
bind
server
multi-server
disk
subnet
connector
connector.1
connector.2
Label
disk array
device
Custom Properties - dis...
type disk
id d29301
name oracle
target 0
drivesize 8631
drivetype scsi
imagetype system
imageOS HPUXOracle
backup none
RP2 Grid Farm Designer
15.9.79.173:4321/udc:/O=Grid/O=Globus/OU=hpl.hp.com/CN=Sven Graupner
sos.hpl.hp.com:4322/udc:/O=Grid/O=Globus/OU=hpl.hp.com/CN=Sven Graupner
15.9.75.157:4321/udc:/O=HP/OU=UNX/CN=Jim Pruyne/EMAIL=pruyne@hpl.hp.com
outer
bind
lb1
inner
webserver
webserver
db
oracle
fileserver
oracle
fileserver
generate RSL
(* BEGIN FML for farm: 2000 *)
&(farm_1 =
(id "2000")
(version "1.1")
(tier_2
(id "OracleRole")
(name "OracleServerRole")
(role (element "oracle"))
(min-servers (element "1"))
(max-servers (element "1"))
(init-servers (element "1"))
)
(ServerRole_3
(id "oracle")
(name "OracleServerRole")
(hw (element "cpu-pa-risc-x2"))
(disk
(target "0")
(drivetype "scsi")
(drivesize "8631")
(diskimage
(type "system")
(element "HPUXOracle")
)
(attribute (name "backup-policy") (value "none"))
)
)
)
(* END FML for farm: 2000 *)
)
(tier_4
(id "FileServerRole")
(name "FileServerRole")
(interface (name "eth0") (subnet "db"))
(role (element "fileserver"))
(min-servers (element "1"))
(max-servers (element "1"))
(init-servers (element "1"))
)
(tier_5
(id "WebServerRole")
(name "WebServerRole")
clear
FML
RSL
smartfrog
properties
gridProxyInit
save
submit
Page-1
Page-2

UDC in an Enterprise Grid

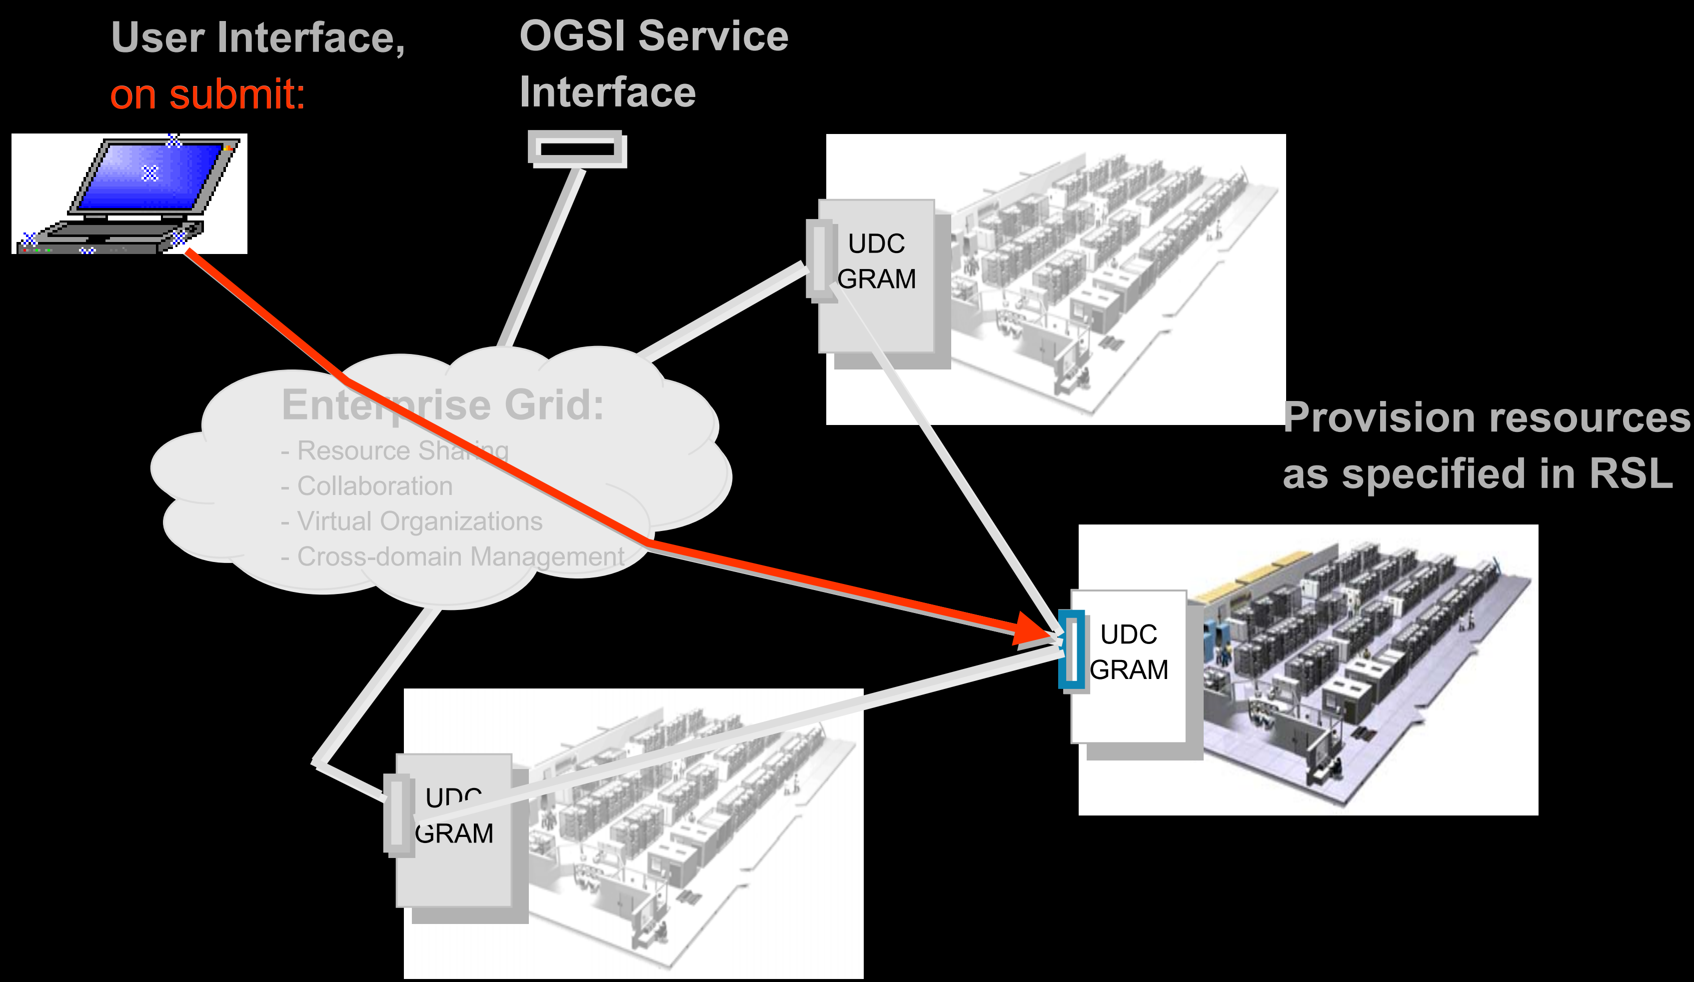

GRAM = Grid Resource Allocation Manager

UDC in an Enterprise Grid

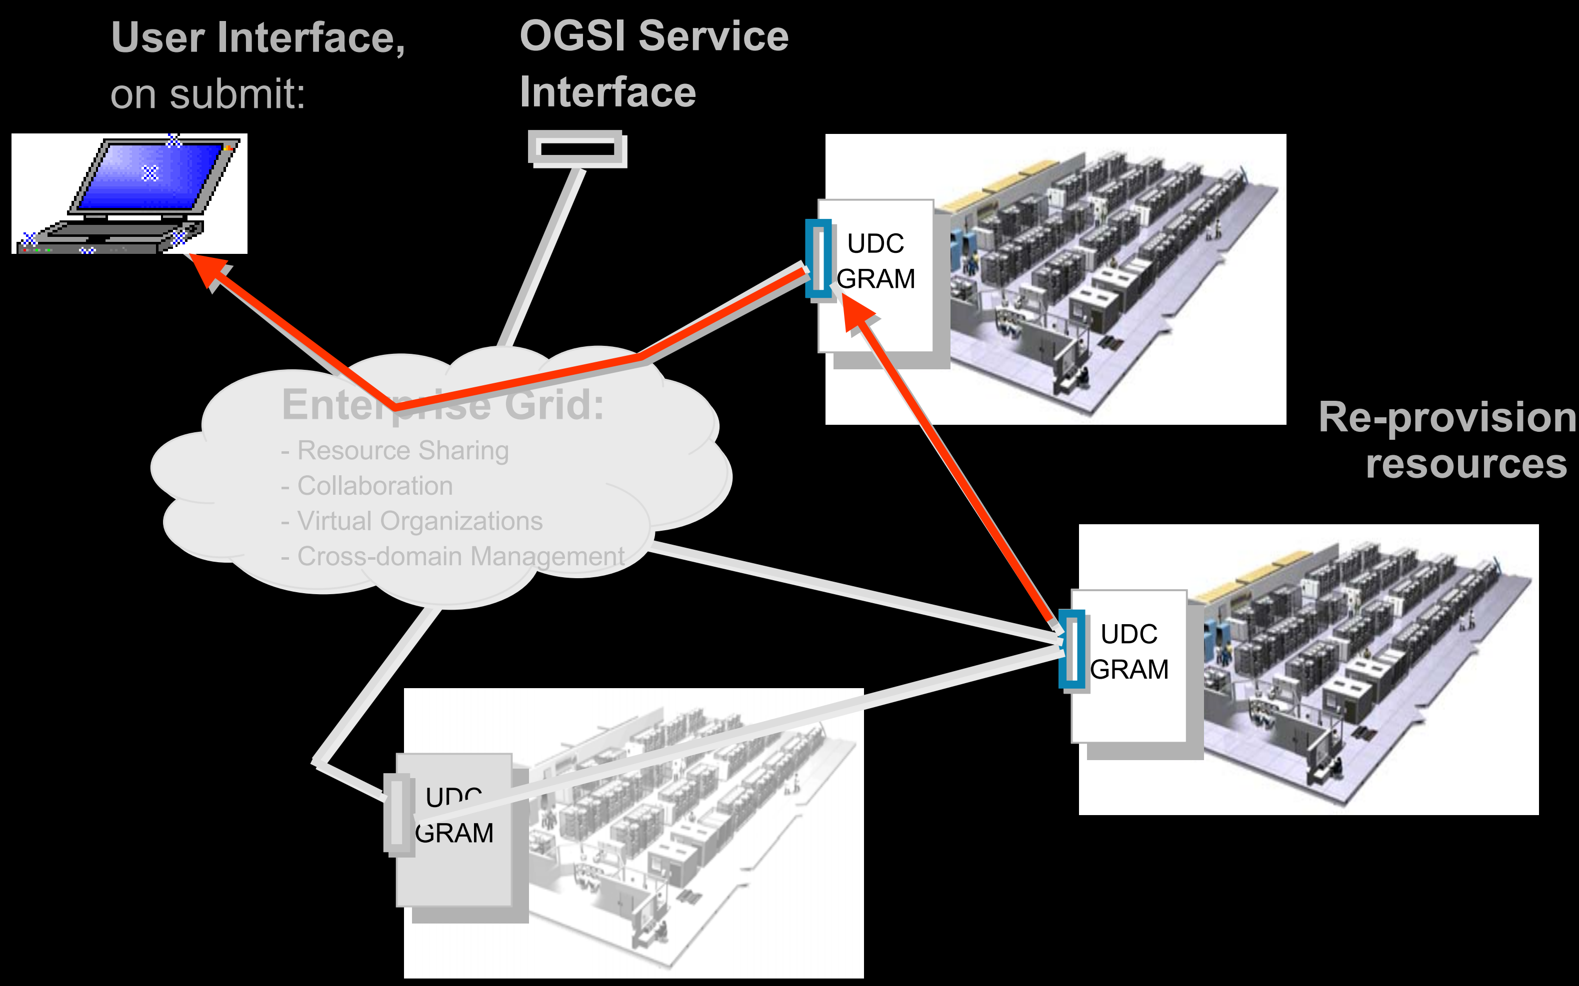

GRAM = Grid Resource Allocation Manager

UDC & Enterprise Grids: Summary

- Goals of **agility**, **collaboration** & **RoIT** can be realized through **virtualization**, **federation** & **automation**:
  - Enabled by Grid technology
  - Grid core: OGSA/OGSI
- Grid technology is moving into Enterprise IT landscapes
- Adaptive Enterprise requires automated management
- UDC adds value with its hardware-supported capability for:
  - Resource virtualization
  - Automated resource management
  - Component isolation

For information on the UDC/Grid demonstration contact Sven.Graupner@hp.com

The Global Grid Forum (GGF)

- Standards-setting body:
  - Working groups
  - Research groups
- Mission:
  - Promote and support development, deployment, implementation of Grid technologies and applications
- Membership:
  - 40+ countries
  - 400+ organizations
  - 2500+ mail-list subscribers
- Three meetings per year:
  - GGF7: March 2003, Tokyo
  - GGF8: June 2003, Seattle
  - GGF9: October 2003, Chicago (working sessions only)
  - GGF10: March 2004, Frankfurt
  - GGF11: June 2004, Hawaii
- 500+ people per meeting
- HP is a Platinum sponsor and active participant in GGF

http://www.ggf.org

References

- HP Grid external web sites:
  http://www.hp.com/products/globus
  http://www.hp.com/techservers/grid
- HP Utility Data Center (UDC):
  http://www.hp.com/go/hpudc
- The Globus Project:
  http://www.globus.org
- The Anatomy of the Grid:
  http://www.globus.org/research/papers/anatomy.pdf
- The Physiology of the Grid (Open Grid Services Architecture):
  http://www.globus.org/research/papers/ogsa.pdf
- The Global Grid Forum:
  http://www.ggf.org

HP WORLD 2003
Solutions and Technology Conference & Expo

Interex, Encompass and HP bring you a powerful new HP World.

interex
shared knowledge • shared power

encompass
AN HP USER GROUP

hp
invent

Backup Slides

Example: European Data Grid

- Funded by the European Union
- Intensive computation and analysis of shared scientific databases
- Applications: high energy physics, biology, earth observations
- Nodes in fifteen countries
- http://www.edg.org

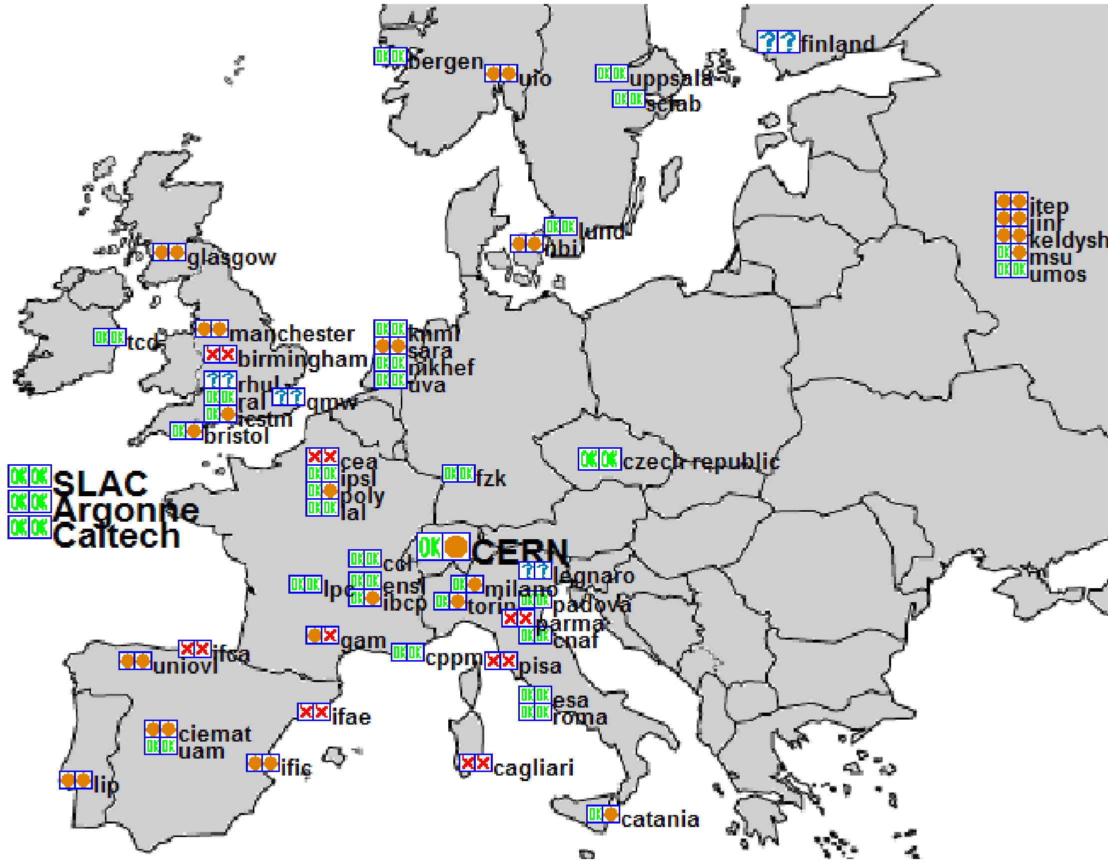

Virtual Organization Example

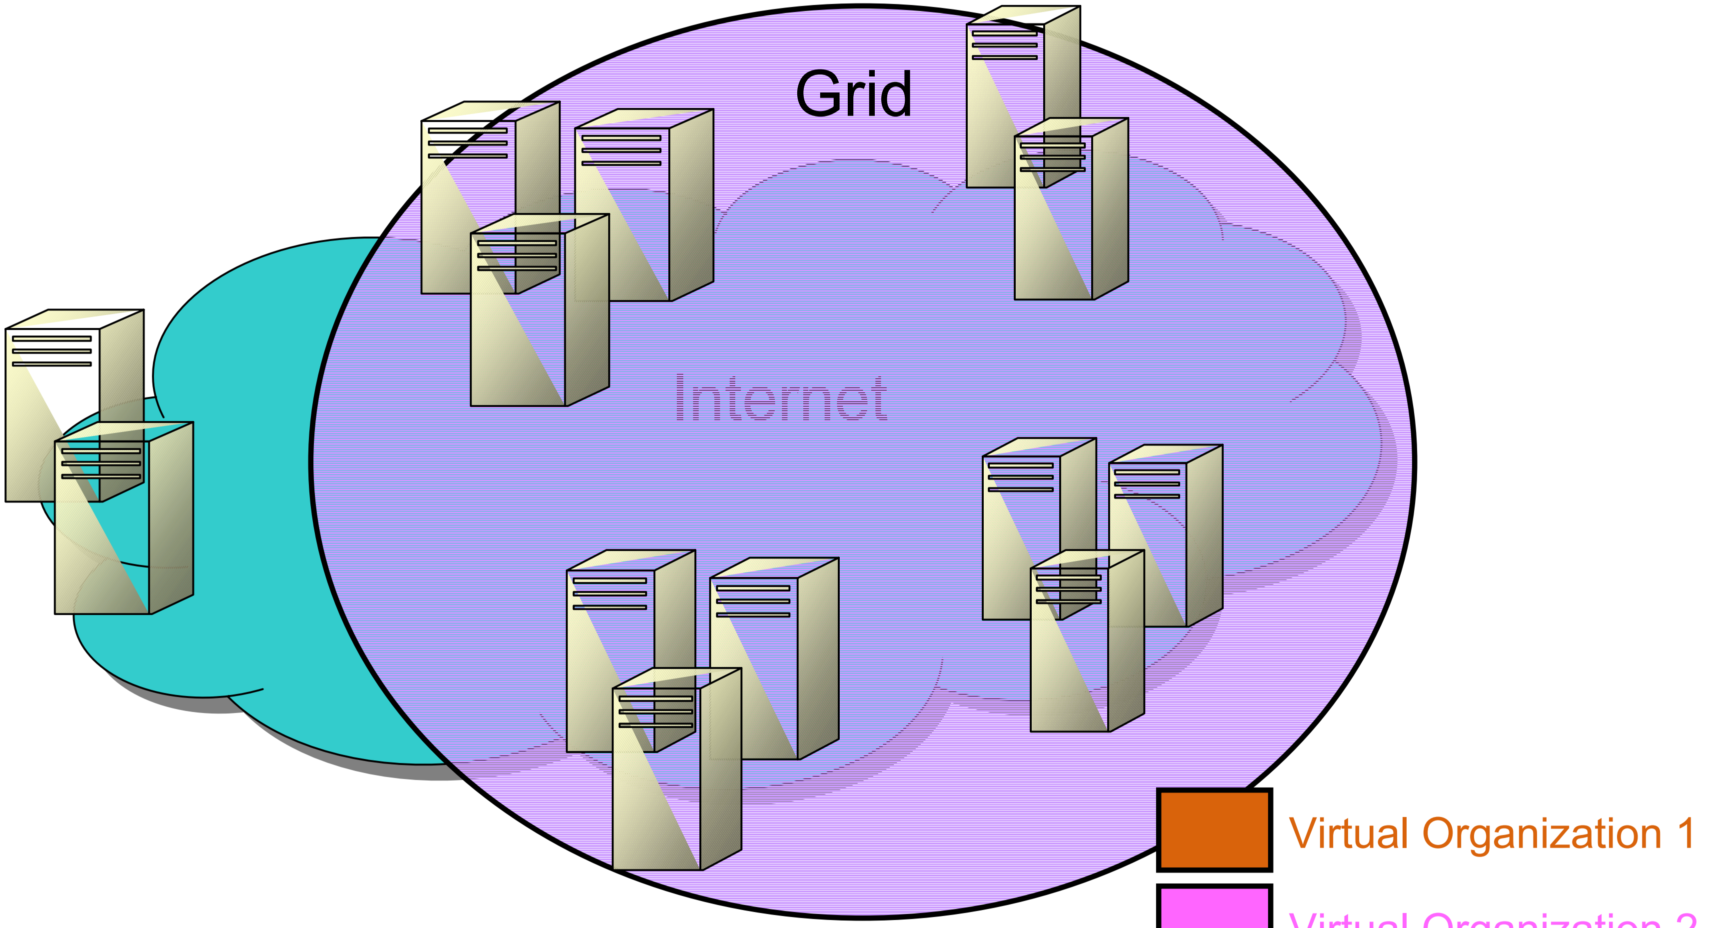

Grid Applications

- Portability:
  - App must run on multiple platforms
  - Use common facilities:
    - databases, event systems, …
- Architecture:
  - Some apps may need to be parallelized
- Licensing:
  - May require different model
- Security:
  - Authentication & authorization
  - May need secure communication
  - Privileges, tempfiles, setuid, …
- Interoperability:
  - Grid applications may need to deal with different data standards:
    - Endianness, floating point format, data size, …
- Accessibility:
  - Install app on every node
  - Stage it at run-time
- Data Access:
  - Local app can NFS-mount
  - Grid app may have to:
    - Transfer input to local storage
    - Transfer output back to the user