

Grids and System Virtualization

BOF

Yaron Haviv

yaroh@voltaire.com

CTO, Voltaire

Intellectual Property Policy

- I acknowledge that participation in OGFx is subject to the OGF Intellectual Property Policy.
- Intellectual Property Notices Note Well: All statements related to the activities of the OGF and addressed to the OGF are subject to all provisions of Section 17 of GFD-C.1 (.pdf), which grants to the OGF and its participants certain licenses and rights in such statements. Such statements include verbal statements in OGF meetings, as well as written and electronic communications made at any time or place, which are addressed to: the OGF plenary session,
 - any OGF working group or portion thereof,
 - the GFSG, or any member thereof on behalf of the GFSG,
 - the GFAC, or any member thereof on behalf of the GFAC,
 - any OGF mailing list, including any working group or research group list, or any other list functioning under OGF auspices,
 - the GFD Editor or the GWD process
- Statements made outside of a OGF meeting, mailing list or other function, that are clearly not intended to be input to an OGF activity, group or function, are not subject to these provisions.
- Excerpt from Section 17 of GFD-C.1 Where the GFSG knows of rights, or claimed rights, the OGF secretariat shall attempt to obtain from the claimant of such rights, a written assurance that upon approval by the GFSG of the relevant OGF document(s), any party will be able to obtain the right to implement, use and distribute the technology or works when implementing, using or distributing technology based upon the specific specification(s) under openly specified, reasonable, non-discriminatory terms. The working group or research group proposing the use of the technology with respect to which the proprietary rights are claimed may assist the OGF secretariat in this effort. The results of this procedure shall not affect advancement of document, except that the GFSG may defer approval where a delay may facilitate the obtaining of such assurances. The results will, however, be recorded by the OGF Secretariat, and made available. The GFSG may also direct that a summary of the results be included in any GFD published containing the specification.

OGF Intellectual Property Policies are adapted from the IETF Intellectual Property Policies that support the Internet Standards Process.

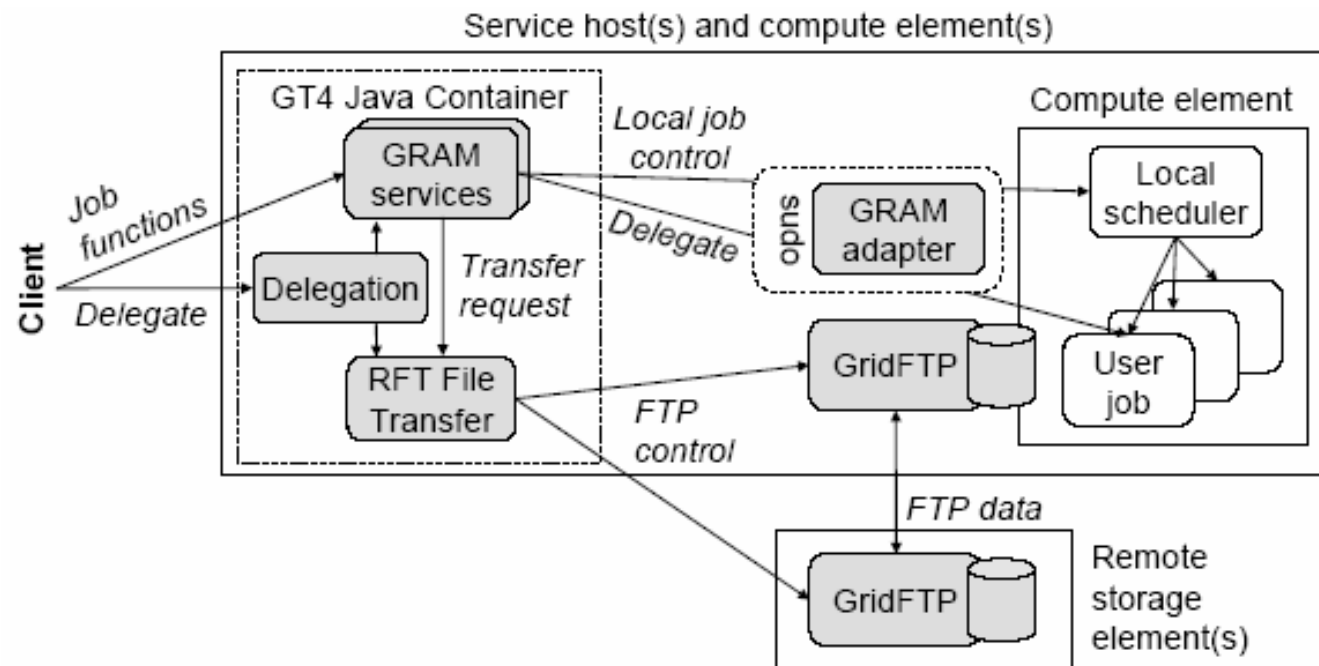
Agenda



- Grids Today
- Server Virtualization and Grids
- Enterprise Grid Challenges
- Data Center Virtualization and Grids
- Call to action and OGF tasks

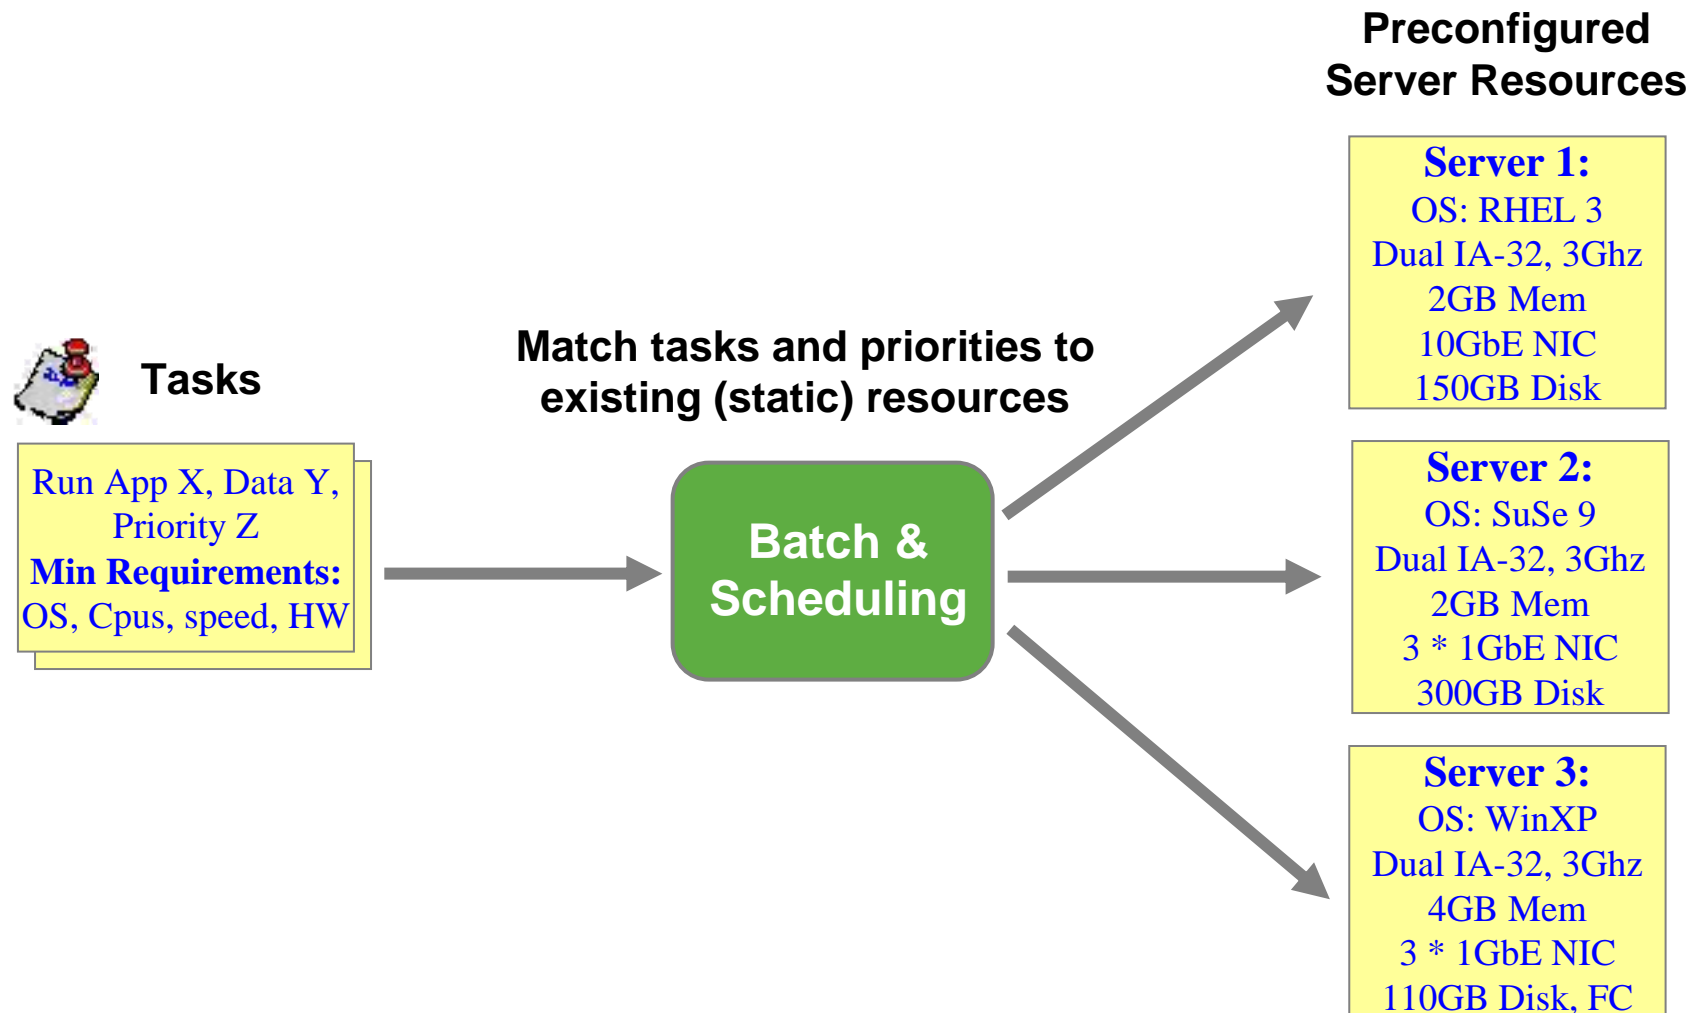
Grids Today, Globus

Globus GRAM Example



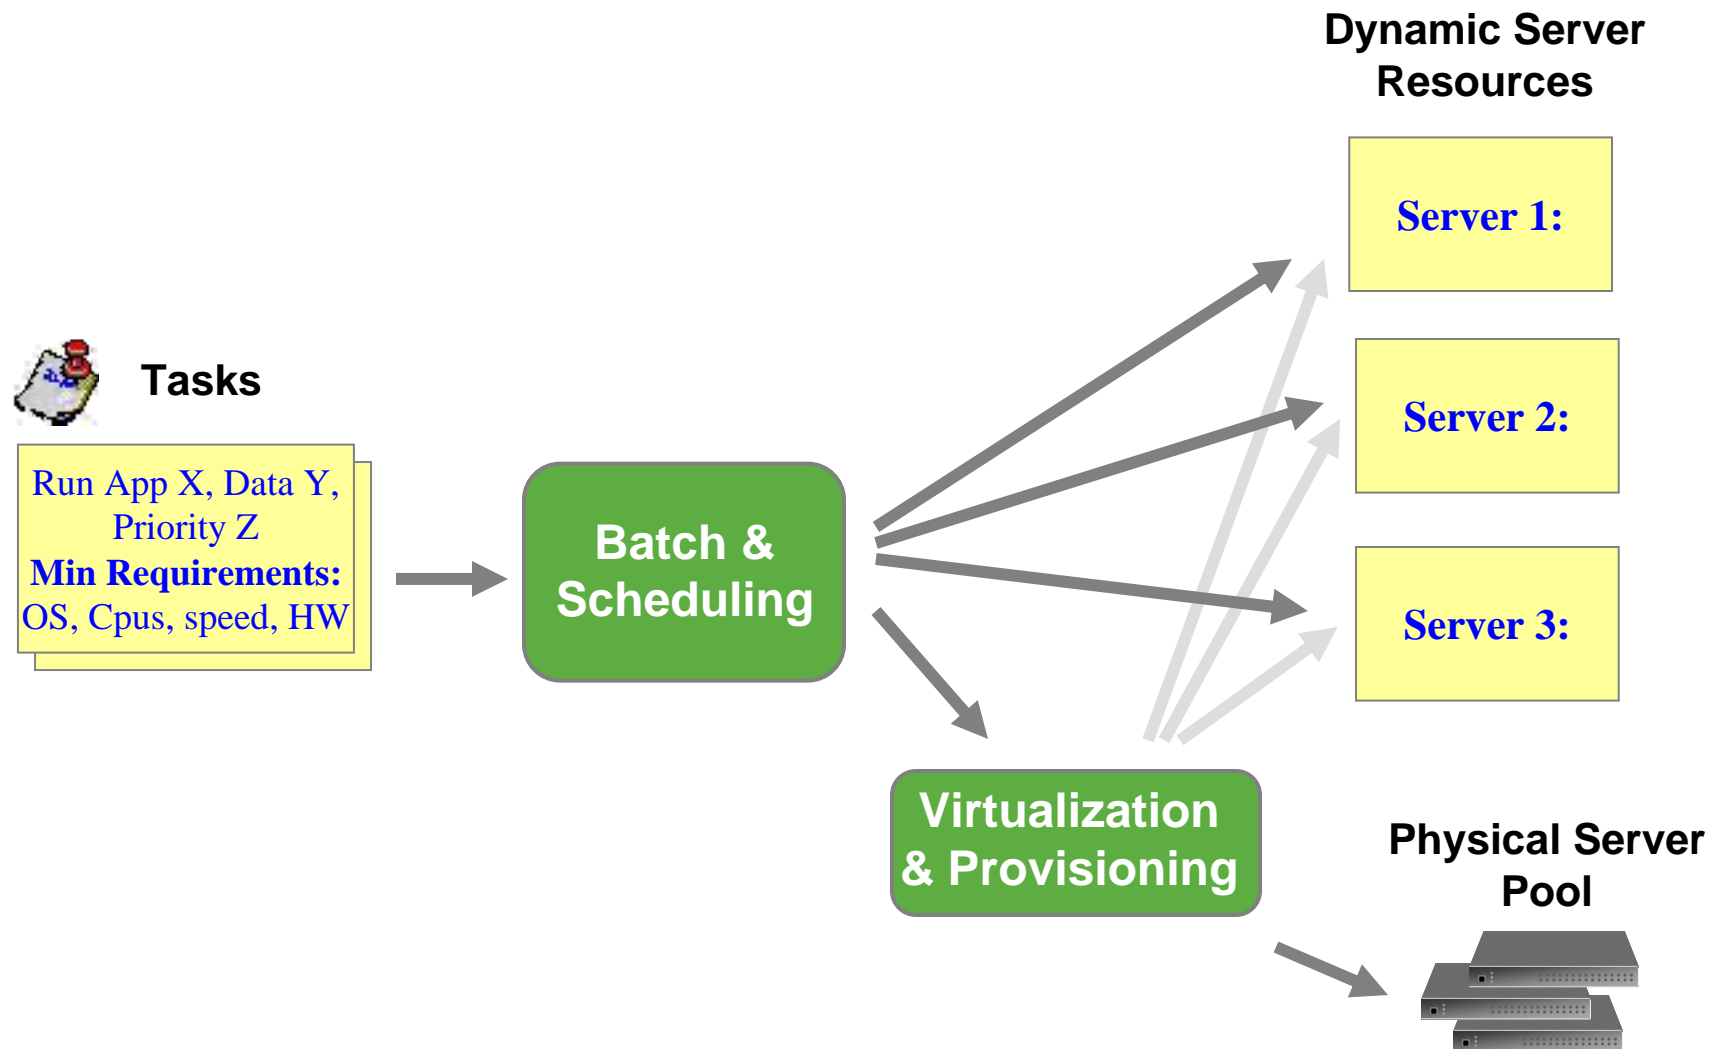
- Client issue a job on a remote machine
- Optional file staging and exporting
- Manage Job lifecycle
- **Assume hosts are live and environment is configured !**

Grids Today, Static Resources



- Servers can be created on-demand using Virtualization
- OS and Applications can be deployed automatically on Servers using server provisioning tools
- Grid job initiation process need to change to allow creation of Virtual servers (vs. using static resources)

Grids with Virtualization, Gen 1



Enterprises Today



Business Line Executive:

“We need to implement Oracle. I need it to be highly available and address our peak month-end load requirements. Make it happen within my existing CapEx budget, but with lower OpEx.”



IT Executive:

“We need to address workloads ranging from zero to 32 processors. We need to connect to the HR, Manufacturing and Finance networks and have between 1 and 15 TB of on-line storage available”.



Network Admin Team

NIC & IO, DHCP, DNS, Switches, Routers, Load-Balancers, Security



Storage Admin Team

HBA, SAN, RAID, Zoning, Backup



Server Admin Team

Server, OS, Cluster, Application, Middleware

Enterprise, More Challenges



- User language is still much higher level than used by configuration tools today, require manual and expert interaction
- Applications require different balance of compute, IO, network resources, most tools focus on compute and have no visibility to networks, IO, and storage
- Applications may consist of multiple tiers
- Resources need to migrate, resize, fail-over without breaking the application or require user intervention
- Applications connect to enterprise resources

Storage Example



- **User:** I need 10GB of fast storage, with RAID5 & Multi-path for my DB application
- **IT:** create a 10GB volume in SAN, configure LUNs, Mapping, Zoaning, HBA, etc.
- How does current models address mapping user requirements to resources and configuration ?
- How is integration between server and storage virtualization technologies achieved ?
- When a server migrates who reestablishes the SAN connectivity ?

Connectivity Example



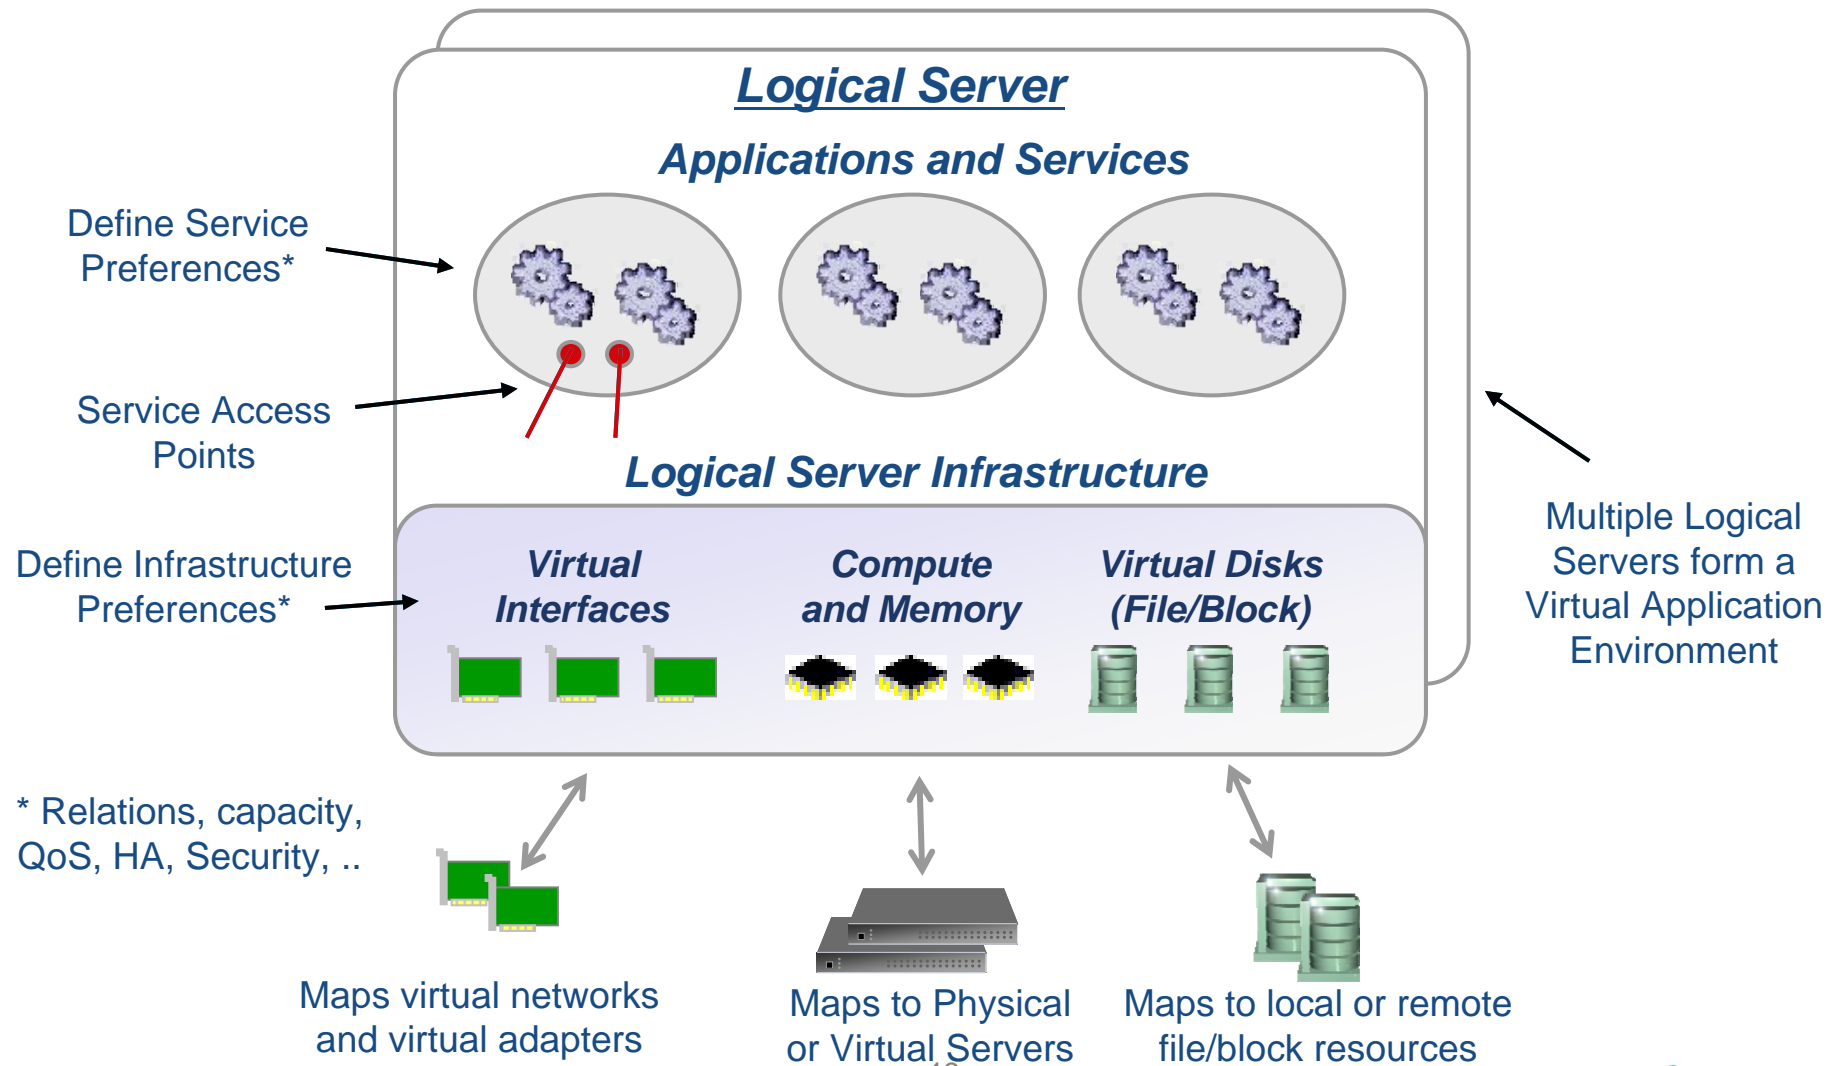
- **User:** my application consists of web, db, and app servers, the DB connects to HR, has a private low latency cluster network, and the web server connects to the internet, the admin access all via mng vlan
- How can a tiered application be described ?
- How is the overall switching described, virtualized, and provisioned to satisfy requirements and QoS ?
- Who verifies end to end connectivity when resources migrate, fail-over, extend ?
- How is security enforced ?

Addressing The Gaps




- Need to Virtualize & provision IO, communication and storage not just servers
- Need to create a higher level definition language and object model addressing:
 - High-level user requirements and QoS
 - Compute, storage, IO, communication
 - Multi-tiered applications, clusters
 - Relations to enterprise resources/networks

The Extended Server Model



The Virtual Network Model



- A Network is a collection of Interfaces (end-points)
 - Each Logical Server may have multiple interfaces
- Interfaces deliver multiple network services
- Links connect networks, interfaces, and services
- Routers manage traffic and connect networks

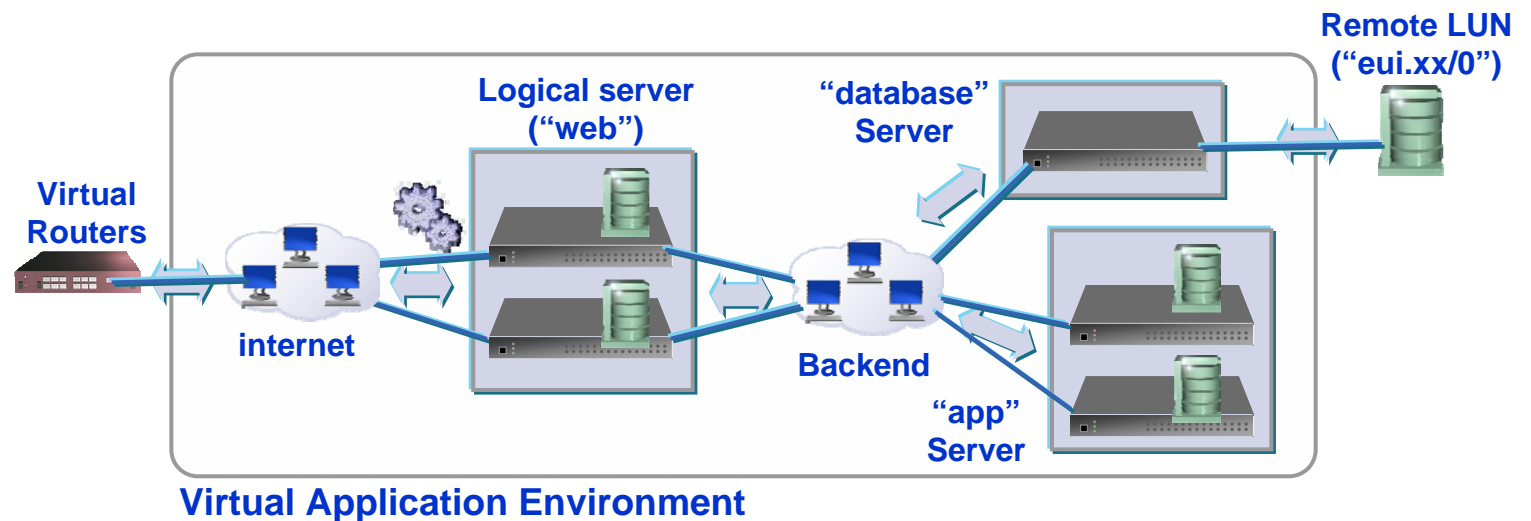
Application Environment Model



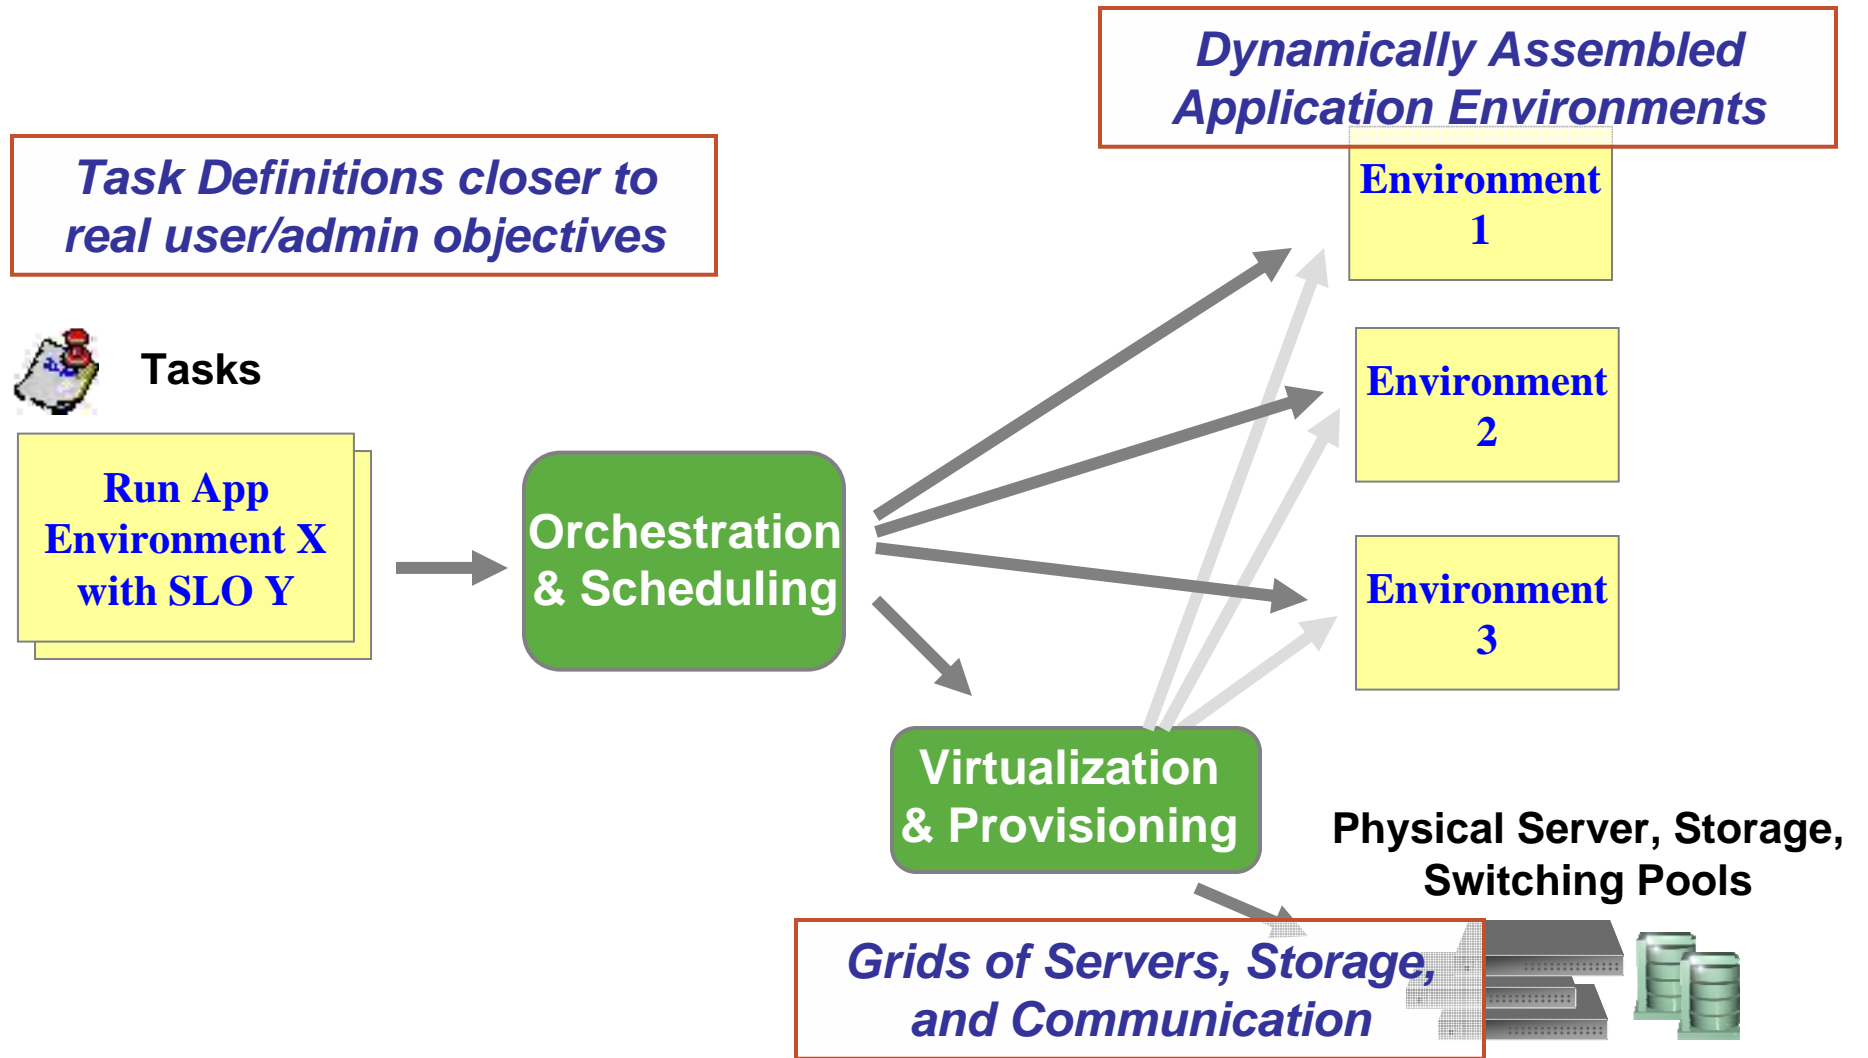
1. Define application environment and internal/external relations in a resource independent way
2. Assign/resize resources independently as needed



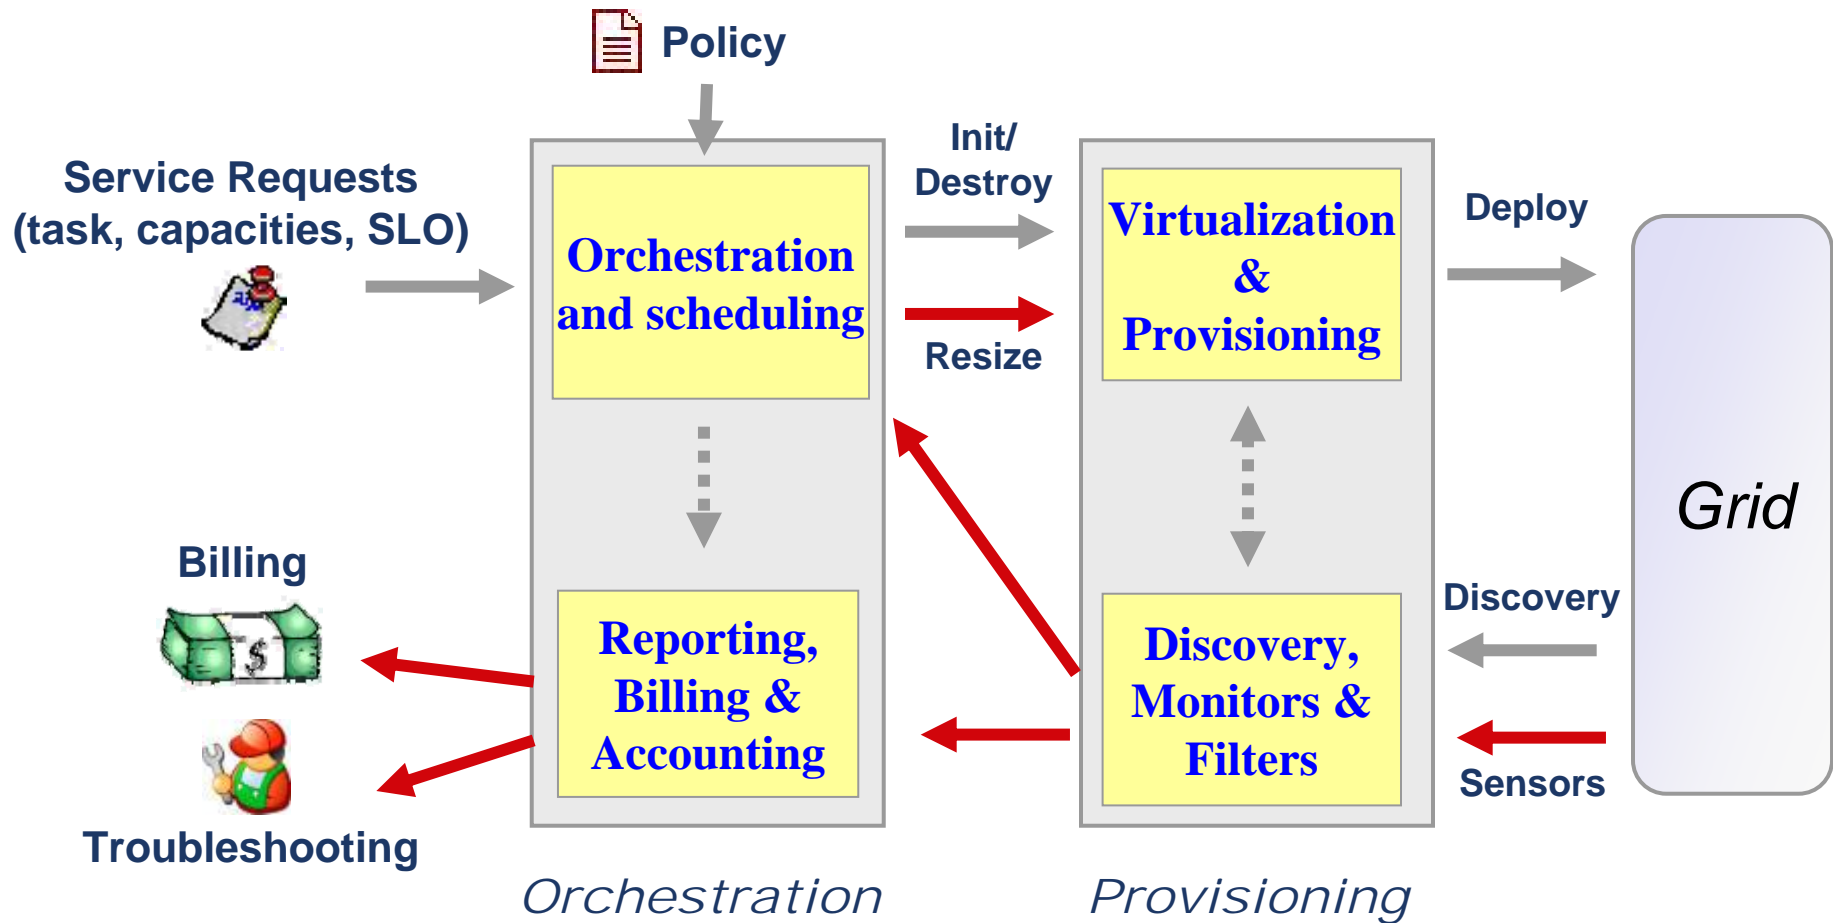
Unassigned resources pool



Future Enterprise Grids



Complete Grid Life Cycle



Call to Action



- Interoperability is key to Grid success
- Need to address interoperability between different Grid elements
 - Orchestrators & schedulers
 - Provisioning & virtualization tools
 - Infrastructure/resources

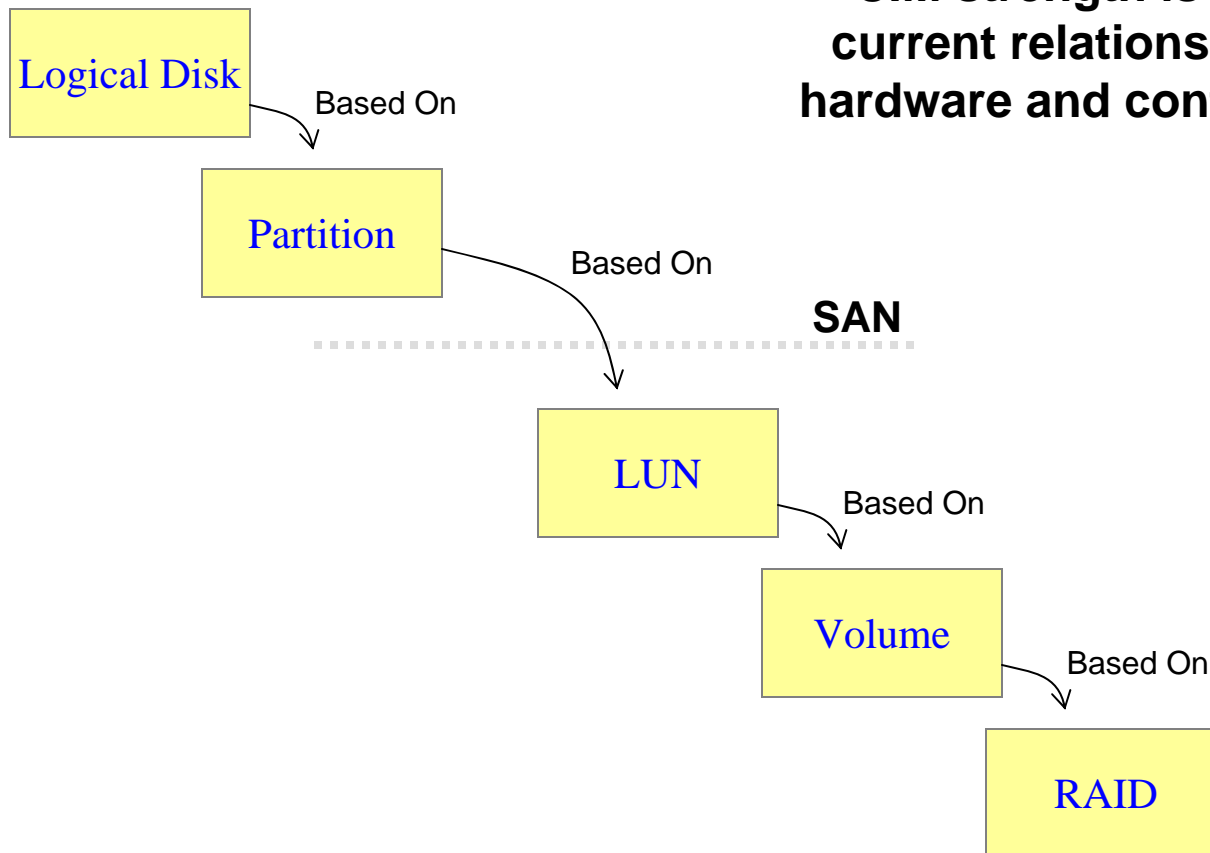
Proposed OGF Work Item



- Standardize efficient definition language and model that address enterprise requirements
 - Focus on user objectives (vs. HW models)
 - Has holistic view (servers, communication, storage, IO, multi-tiered)
 - Be extensible, relational, resizable, technology and resource independent
 - Accessible via a web service

Backup

CIM Model Example



CIM strength is in describing the current relations between different hardware and configuration elements