


interactiveGrid

A Resource Management Framework For interactive Grid

Raj Kumar, Vanish Talwar, Sujoy Basu
HP Labs
July 2, 2003
(Contact: raj.kumar@hp.com; 650-857-8104)


July 2, 2003 HP's interactiveGrid Page 1 

interactiveGrid

Grid application domains

```
graph TD; A[Grid application domains] --> B[Batch]; A --> C[Interactive]; C --> D[Graphics]; C --> E[Streaming]; C --> F[CAD/CAM]; C --> G[Video Games]; C --> H[SW development]; C --> I[Application Hosting]; C --> J[Digital Content Creation]; C --> K[Financial markets];
```

vision: Innovate technologies for next generation grid

July 2, 2003 HP's interactiveGrid Page 2 

interactiveGrid

Grid requirements

First Generation- Batch jobs

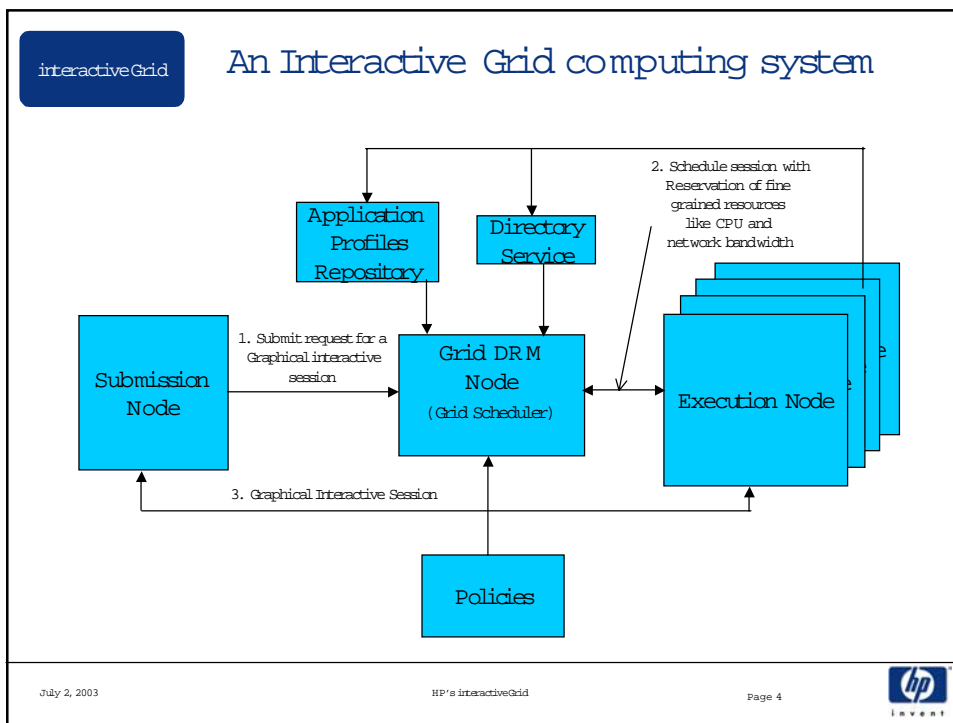
- ✓ Virtualization
- ✓ Authentication
- ✓ Authorization
- ✓ Registry service
- ✓ Discovery service
- ✓ Job submission
- ✓ Job query
- ✓ Notification
- ✓ Accounting
- ✓ Billing
- ✓ Resource reservation
- ✓ Data management
- ✓ Scheduling
- ✓ QoS/SLAs
- ✓ Dynamic accounts

+

Next Generation- Interactive jobs

- ✓ Export remote desktop
- ✓ Security
 - access control
 - account mgmt
- ✓ Session(s) management
 - QoS issues, SLAs
 - Admission control
- ✓ Monitoring services
- ✓ Real-time schedulers
- ✓ Staging data

July 2, 2003
HP's interactiveGrid
Page 3



interactiveGrid

Key architectural concepts

- Virtualized interactive nodes
- Controlled access (desktop, shell, user, superuser)
- Hierarchical Sessions
- Hierarchical Admission control
- Hierarchical agents (sensing/monitoring/aggregating/enforcing/registering,..)
- Application Predictor system
- Real-time scheduling with QoS guarantees
- Dynamic template accounts
- Classes of dynamic accounts
- Persistent user environments
- User data management (affinity scheduling)

Interactive Grid: A User's interaction

Global Admission
1. Submit request for graphical interactive session
2. Allocate dynamic accounts: dyn004
3. Session ends
4. Schedule session
5. Global session
6. Per application session
7. Session ends
8. Save User Env

Per application admission
Enforce QoS guarantees
Based on SLAs

July 2, 2003

HP's interactiveGrid

Page 5

interactiveGrid

Resource Management Framework on execution Node(s)

1. Admission Control System
2. Resource Management Agents
3. Application Predictor System

Applications

Admission Control System

Application Predictor System

Resource Management Agents

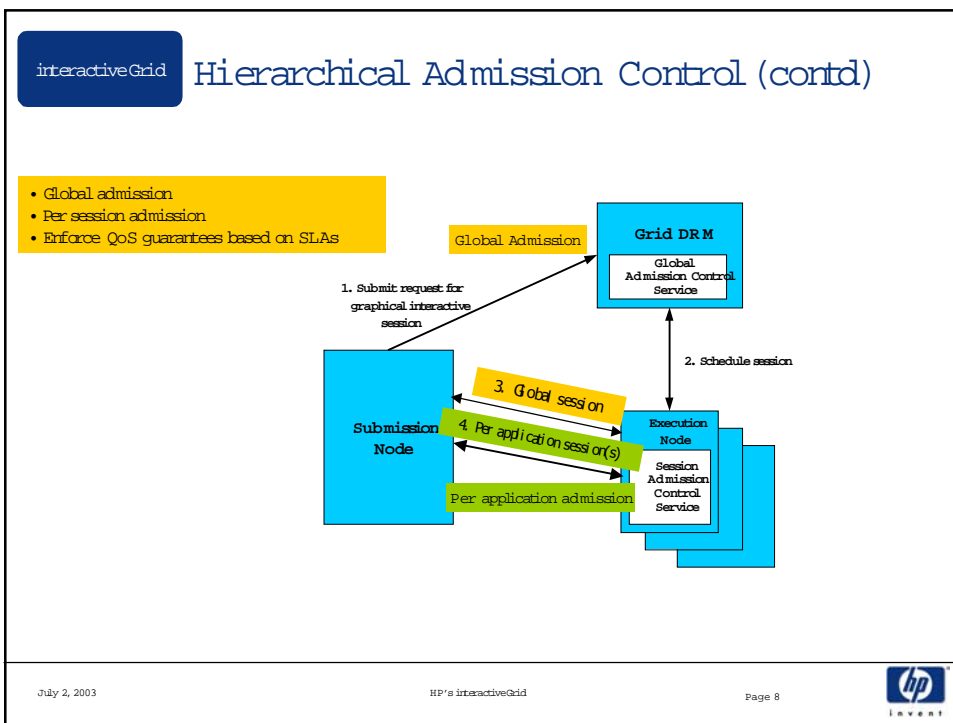
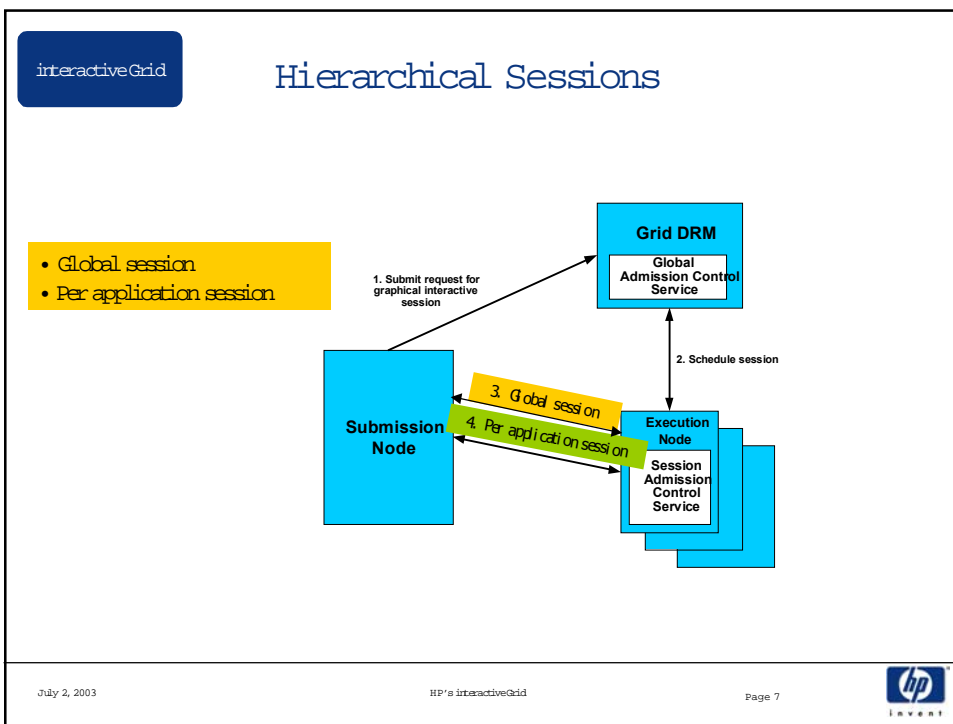
Resource Management Framework

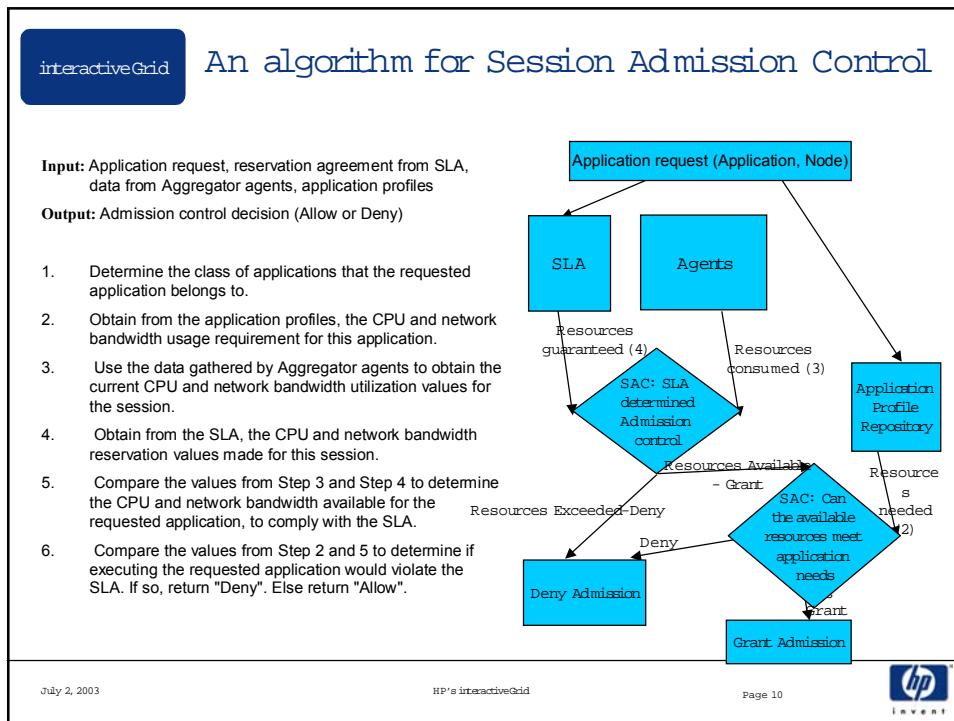
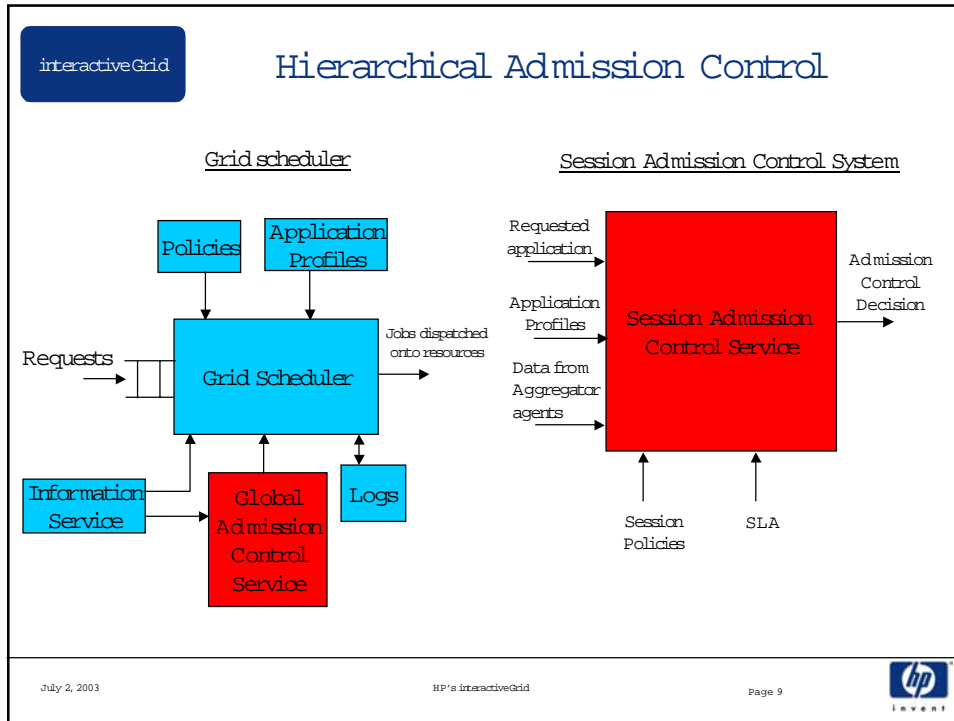
Operating System

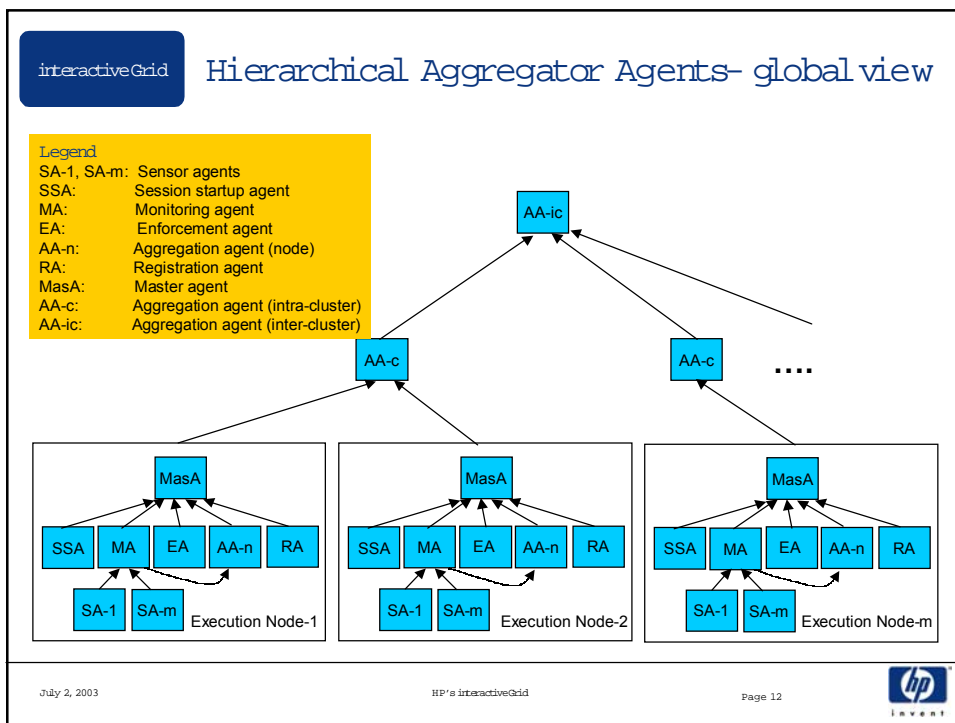
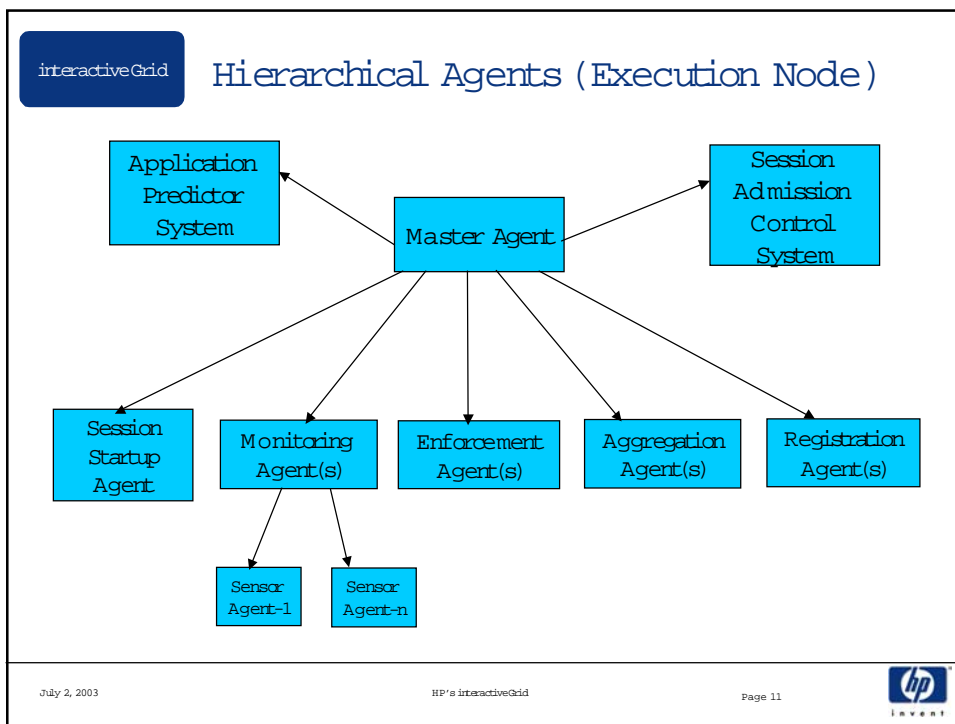
July 2, 2003

HP's interactiveGrid

Page 6





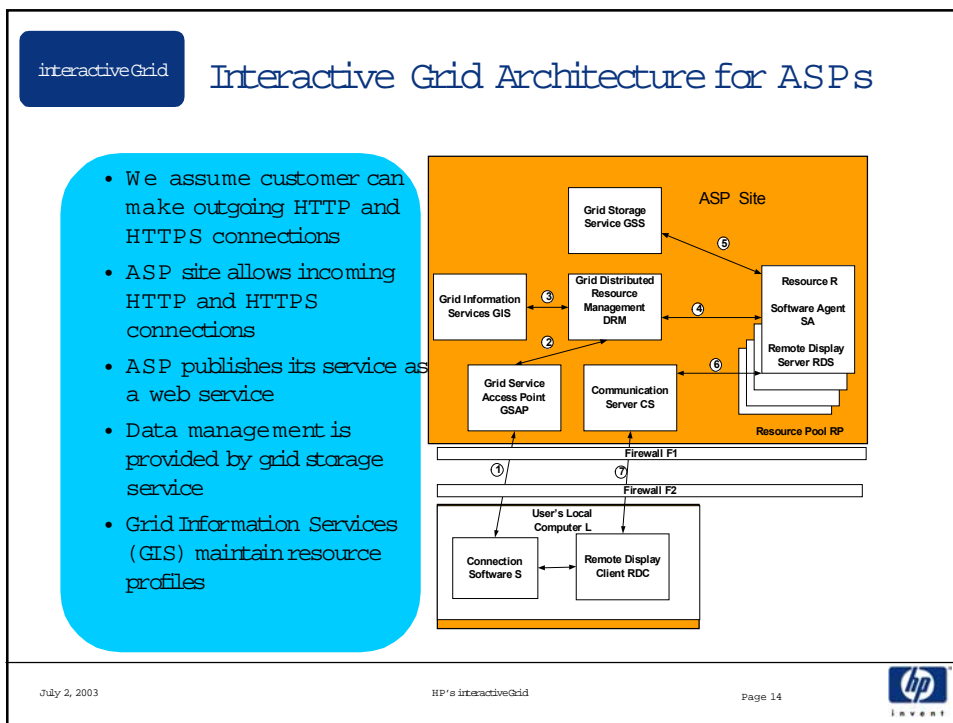


interactiveGrid

An example of application profiles

Application	Acceptable Frame Rate	CPU Requirement (Remote Display Server)		CPU Requirement (Application)		Network Bandwidth Requirement (Remote Display Server)	
		Low	Allowed	Low	Allowed	Low	Allowed
Engineering	~10 frames/sec	8 %	10 %	8 %	10 %	15 %	20 %
Video	~30 frames/sec	20 %	25 %	20 %	25 %	35 %	40 %
Games	~5 frames/sec	3 %	5 %	3 %	5 %	10 %	15 %

July 2, 2003 HP's interactiveGrid Page 13



interactiveGrid

De mo: Login screen

July 2, 2003
HP's interactiveGrid
Page 15

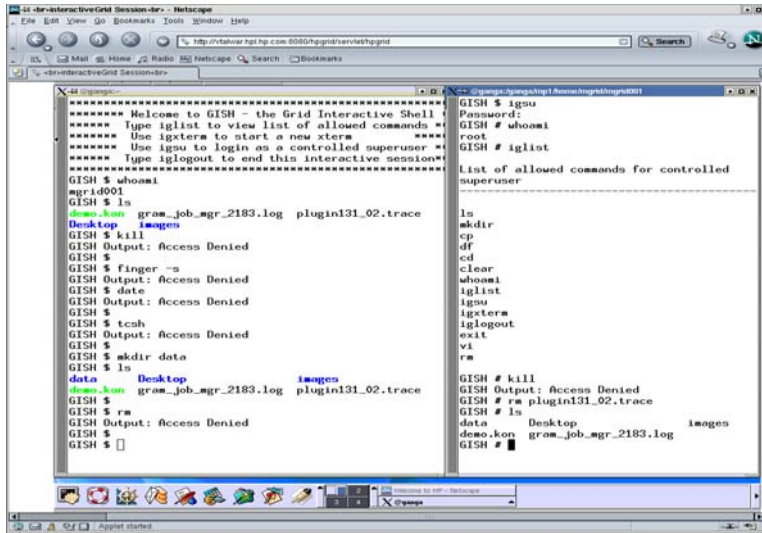
interactiveGrid

De mo: Job Submission

July 2, 2003
HP's interactiveGrid
Page 16

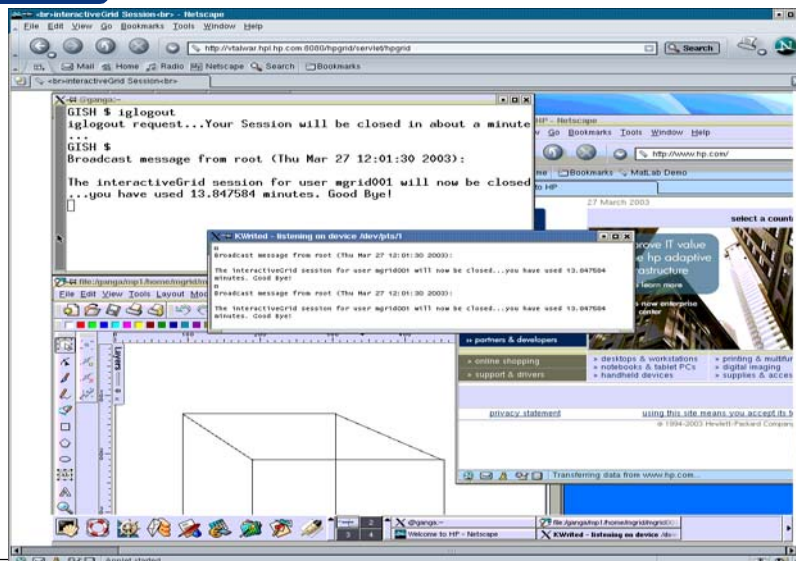
interactiveGrid

Demo: Session Progress



interactiveGrid

Demo: Session End



interactiveGrid

Design Goals


- Support Heterogeneous platforms
- Extend general purpose tools
- Minimal changes to existing software
- Extensible and modular
- Address needs of graphics and multimedia applications
- Support self-managing capabilities
- Work for all application types
- Interoperable with other complementary grid solutions
- Policy driven

July 2, 2003 HP's interactiveGrid Page 19 

interactiveGrid

Related Work

- Most of the current work is for batch jobs
- CrossGrid
- Other work
 - Punch
 - Quorum
 - Network Weather Service
- Datagrid Dynamic Accounts

July 2, 2003 HP's interactiveGrid Page 20 

Conclusions

- Interactive Grids
 - Next generation grids supporting interactive application domains
 - Virtualized interactive nodes
 - Controlled access (desktop, shell, user, superuser)
 - Hierarchical Sessions
 - Hierarchical Admission control
 - Hierarchical agents (sensing/monitoring/aggregating/enforcing/registering..)
 - Application Predictor system
- Testbed implementation extending Globus Toolkit 2.0



interactiveGrid

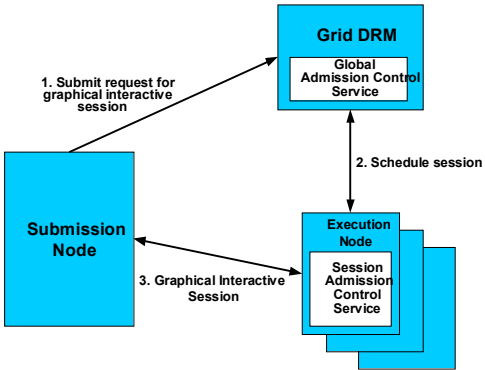
Backup slides

July 2, 2003 HP's interactiveGrid Page 23 


interactiveGrid

Controlled Access

- Controlled Desktop
- Controlled Shell
- Controlled User Account
- Controlled Super User account



```
graph TD; Submission[Submission Node] -- "1. Submit request for graphical interactive session" --> Grid[Grid DRM Global Admission Control Service]; Grid -- "2. Schedule session" --> Execution[Execution Node Session Admission Control Service]; Execution -- "3. Graphical Interactive Session" --> Submission;
```

July 2, 2003 HP's interactiveGrid Page 24 

interactiveGrid

Dynamic Template Accounts

- Pool of Dynamic or Template accounts
- Classes of Dynamic accounts
- Allocate dynamic accounts at session start
- Retrieve dynamic accounts at session end
- Map user requests to appropriate class of dynamic accounts
- Allows scalability
- Delegate client credentials to allocated node

The diagram illustrates the workflow for dynamic template accounts. It features three main components: a Submission Node, a Grid DRM (Global Admission Control System), and an Execution Node (Session Admission Control System). The process is numbered 1 through 7:

1. Submit request for graphical interactive session (Submission Node to Grid DRM)
2. Allocate dynamic accounts: dyn004 (Grid DRM self-loop)
3. Schedule session with dyn004 (Grid DRM to Execution Node)
4. Global session (Submission Node to Execution Node)
5. Per application session (Submission Node to Execution Node)
6. Session ends (Execution Node to Submission Node)
7. Return dyn004 to the pool (Execution Node to Grid DRM)

July 2, 2003

HP's interactiveGrid

Page 25

interactiveGrid

Persistent User Environment

- Load default user environment for new users
- Load saved persistent environment for old users
- Store user environment after the global session ends
- Provides a more familiar environment in a Grid context, and enhances productivity
- A user can store multiple environments

The diagram illustrates the workflow for a persistent user environment. It features four main components: a Submission Node, a Grid DRM (Global Admission Control System), an Execution Node (Session Admission Control System), and a User Environment Repository. The process is numbered 1 through 8:

1. Submit request for graphical interactive session (Submission Node to Grid DRM)
2. Load User Env (Grid DRM to Execution Node)
3. Store User Env (Execution Node to User Environment Repository)
4. Global session (Submission Node to Execution Node)
5. Per application session (Submission Node to Execution Node)
6. Session ends (Execution Node to Submission Node)
7. Store User Env (Execution Node to User Environment Repository)
8. Session ends (Execution Node to Grid DRM)

July 2, 2003

HP's interactiveGrid

Page 26