

The Access Grid Toolkit

ESnet Collaboration Workshop

10/27-29/2004

The Futures Laboratory



Access Grid Project Goals

- Enable Group-to-Group Interaction and Collaboration
 - Connecting People and Teams via the Grid
- Improve the User Experience: Go Beyond Teleconferencing
 - Provide a Sense of Presence
 - Support Natural Interaction Modalities
- Use Quality but Affordable Digital IP Based Audio/video
 - Leverage IP Open Source Tools
- Enable Complex Multisite Visual and Collaborative Experiences
 - Integrate With High-end Visualization Environments
 - ActiveMural, Powerwall, CAVE Family, Workbenches
- Build on Integrated Grid Services Architecture
 - Develop New Tools Specifically Support Group Collaboration



Our Approach

- Attack Research Questions in the context of real world experience
 - Build up a critical mass of groups using the AG Platform
 - Involve multiple groups in trying new ideas and evaluation
- Build Working Infrastructure as well as Prototype Software
 - Argonne has five working AG nodes under development
 - New Software is used weekly/Daily as part of standard nanocruises
- Involve multiple groups in deployment, use and research
 - Active collaborations with over a dozen groups working on AG technology



Release software early and often (use open source model)

Why this approach?

- We build based on standards, open architectures, clear designs.
- H.323 too closed, broken service model
 - MCU Cost driven by profit model, not technology
- SIP is/was too volatile, broken service model
- Web Services are something that looked interesting
 - They still do, as an open standards effort
 - Grid Services (Computing, Data Storage, Instruments) can just “plug in”...soon

Collaboration Technologies:

- H.323 → SIP
- ConferenceXP
- Live Communications Server
- RTCommunication/MSN Messenger

Others:

- Application Layer Networking Blackhole
- Streaming Media Morass
- Firewalls **will** kill collaboration
 - Publish/Subscribe model


Group-to-Group Interaction is Different

- Large-scale scientific and technical collaborations often involve multiple teams working together
- Group-to-group interactions are more complex than than individual-to-individual interactions
- The access grid project is aimed at exploring and supporting this more complex set of requirements and functions
- The access grid will integrate and leverage desktop tools as needed

Some Access Grid Active Research Issues

- Scalable wide area communication
 - Evolution of multicast related techniques, and time shifting issues
- Scoping of resources and persistence
 - Value of spatial metaphors, security models
 - Virtual Venues, synchronous and asynchronous models
- Improving sense of presence and point of view
 - Wide Field Video, Tiled Video, High-resolution video codecs
- Network monitoring and bandwidth management
 - Beacons and network flow engine
- Role of Back-channel communications
 - Text channels and private audio
- Recording and playback of multistream media

Access Grid Architecture



- The **Community** provides a Public service for discovery
- Application S
 - Network Serv
 - Virtual Venue

What is the Access Grid?

- **Community Services**

- Management of the Community Resources

- **Virtual Venues**

- Places where users collaborate

- **Network Services**

- Advanced Middleware

- **Virtual Venues Client**

- User Software

- **Nodes**

- Shared Nodes
 - Administratively scoped set of resources

- **Resources**

- Provide capabilities

Users collaborate by sharing:

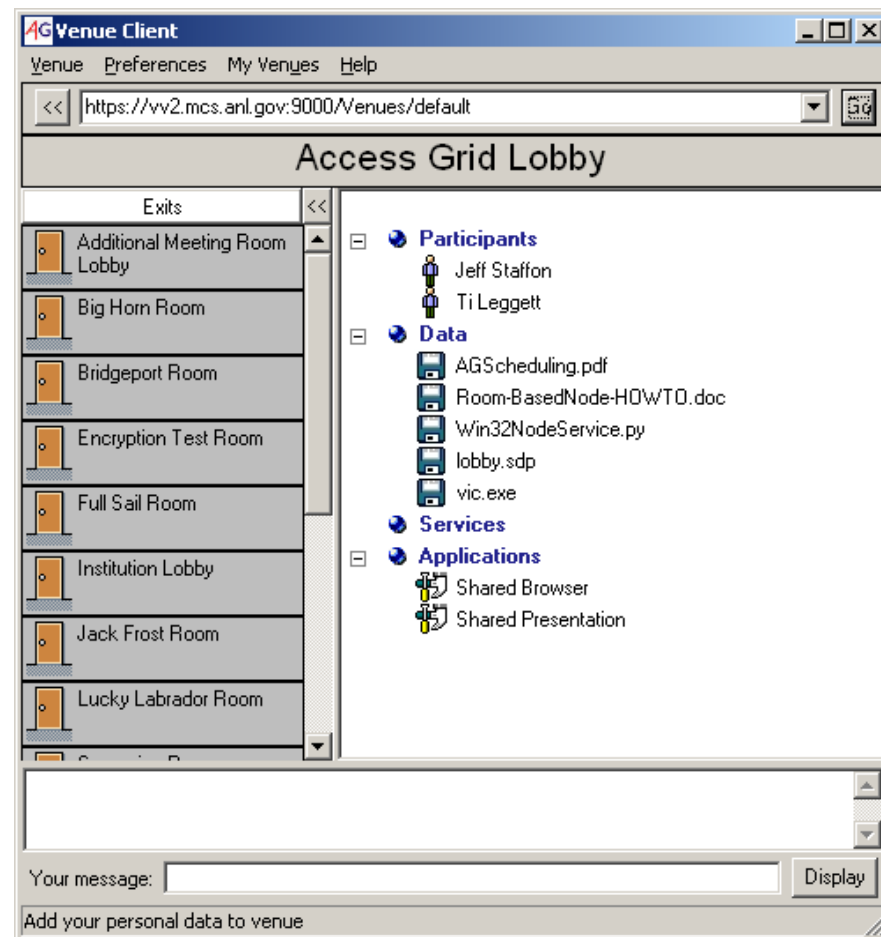
- Data
- Applications
- Resources

- **Personal Nodes**

- User scoped set of Resources

Virtual Venues Client

- Enable face-to-face meeting activities
- What can be done:
 - Sharing Data
 - Shared Applications
- Applications:
 - Distributed PowerPoint
 - Shared Web browser
 - Whiteboard
 - Voting Tool
 - Question & Answer Tool
 - Shared Desktop Tool
- Integrate legacy single-user apps



Community Services

- User Management Web Service Specification
 - CA Web Service
 - CA Cert Bundles
 - Request Interface
 - Admin Interface
 - Online CA Support
 - Jabber Account Integration
 - Publish/Subscribe Registries for:
 - Venue Servers
 - Node Service
 - Shared Application Clients
 - Application Service Clients
- No Authorization Services

Virtual Venues

- What is a Virtual Venue?
 - A Virtual Venue is a virtual space for people to collaborate
- What do Virtual Venues provide?
 - Authorization Information
 - Coherence among Users
 - Venue Environment, Users, Data
 - Client Capabilities Negotiation
 - List of Available Network Services
 - Keep track of resulting Stream Configurations
 - Applications and Services
 - Connections to other Venues
- Virtual Venues have two interfaces
 - Administrative – Venue Management Software
 - Client – Virtual Venue Client Software

Network Services

- Network Services
 - Provide a middleware layer for enabling the richest collaborations
 - Are invisible to Venues Clients, used by Virtual Venues
 - Primarily Transform streaming data
 - Can be anywhere on the network
 - Can be composed to build complex solutions:
 - Venue Audio Stream → Audio Transcoder → Audio to Text → Two-Way Pager
 - Two-Way Pager → Text to Audio → Audio Transcoder → Venue Audio Stream
- Network Services provide opportunities for third party developers
- ANL is working on Network Services for
 - Audio Transcoding
 - Audio Mixing
 - Stream Selection

Access Grid Nodes

Access Grid 2.0 reference platforms:

1. **Advanced Node** – Tiled Display, Multiple Video Streams, Localized Audio
2. **Room Node** – Shared Display, Multiple Video Streams, Single Audio Stream (AG 1.x Node)
3. **Desktop Node** – Desktop Monitor, Multiple Video Streams, Single Audio Stream
4. **Laptop Node** – Laptop Display, Single Video Stream, Single Audio Stream
5. **Minimal Node** – Compact Display, Single Video Stream, Single Audio Stream

What Hardware?

- Cameras, Microphones, Speakers, Display, Input Devices
- Get Audio Correct!

What Platforms?

- Mac OS X, Windows XP, Linux

• Access Grid Nodes

- Comprise a set of collaboration resources
- Expose those resources through Node Services

• Basic Node Services include:

- Audio & Video Services
- Network Performance Monitoring Service
- Network Reliability/Fallback Service

• Extended Node Services could be:

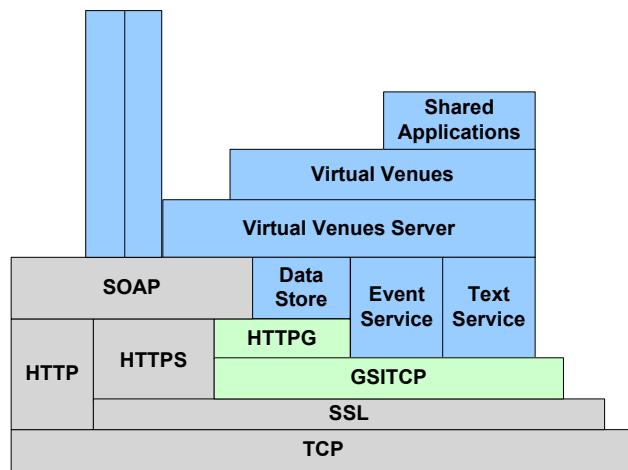
- Display Service with enhanced layout control
- Video Service supporting new CODECs
- Automatic performance adaptation
- Application Hosting Service

Software Requirements?

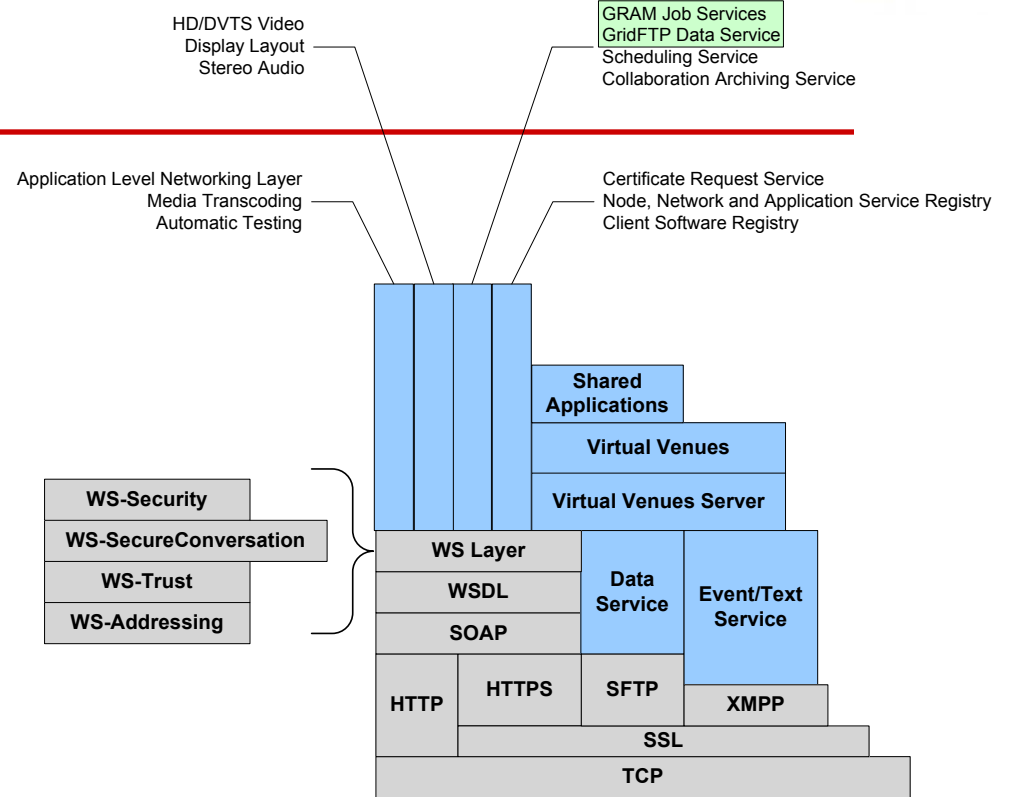
- Python 2.3, wxPython




Access Grid Technology Tower




Access Grid Toolkit 2.X
Services Infrastructure



Access Grid Toolkit 3.X
Services Infrastructure



Access Grid Toolkit
Media Software



Access Grid Toolkit Security Software

Access Grid Security Overview

Requirements

- Identify users, authenticate them with a trusted authority
- Authorize their access to the resources they request
- Provide them privacy and secure access to their applications and data

Stream Security

- Current vic / rat support AES/Rijndael encryption
- Media Key distribution via venues services mechanisms
- Existing concerns:
 - Are keys recoverable (in face of many gigabytes of encrypted data)
 - Would rekeying at random intervals help?
- One possible solution: IETF Secure RTP draft

Security in the AG

Replacing the AGTk 2.X use of X.509 Certificate with the more general notion of a “Credential”

- Credentials could be:
 - X.509 Certificates (Initially)
 - Username/Password/Token or Cookie
 - Kerberos Tickets
- Every user and service must have a credential
- Communications use SSL, if possible
- SSL provides confidentiality

Our current practice,

- Each user has an identity certificate, issues by a CA:
 - Access Grid Developer’s CA (integrated with software)
 - Other CAs
 - DOE Science Grid
 - NCSA
 - EUROGRID
 - Verisign, Thawte, ...

Role-based Access Control

- *Permissions*
 - associated with *roles*
 - defined in the context of some resource
- *Users*
 - assigned to roles based on some policy
 - identified by a credential

Access Grid vs. Commercial Desktop Tools

- AG targets beyond the desktop
 - large format multi-screen for AG Global Channels
 - room scale hands free full-duplex audio
- AG uses dedicated hardware
 - multiple machines, separation of function XP, Linux
- AG software is Open Source
 - extends and builds on community tools
- AG environment is integrated with Grid services
 - extensible framework designed to “plug-into” the Grid
- AG development is a Community Effort
 - you are welcome to join in the fun!!



Access Grid Past

First Generation AGTk 1.X:

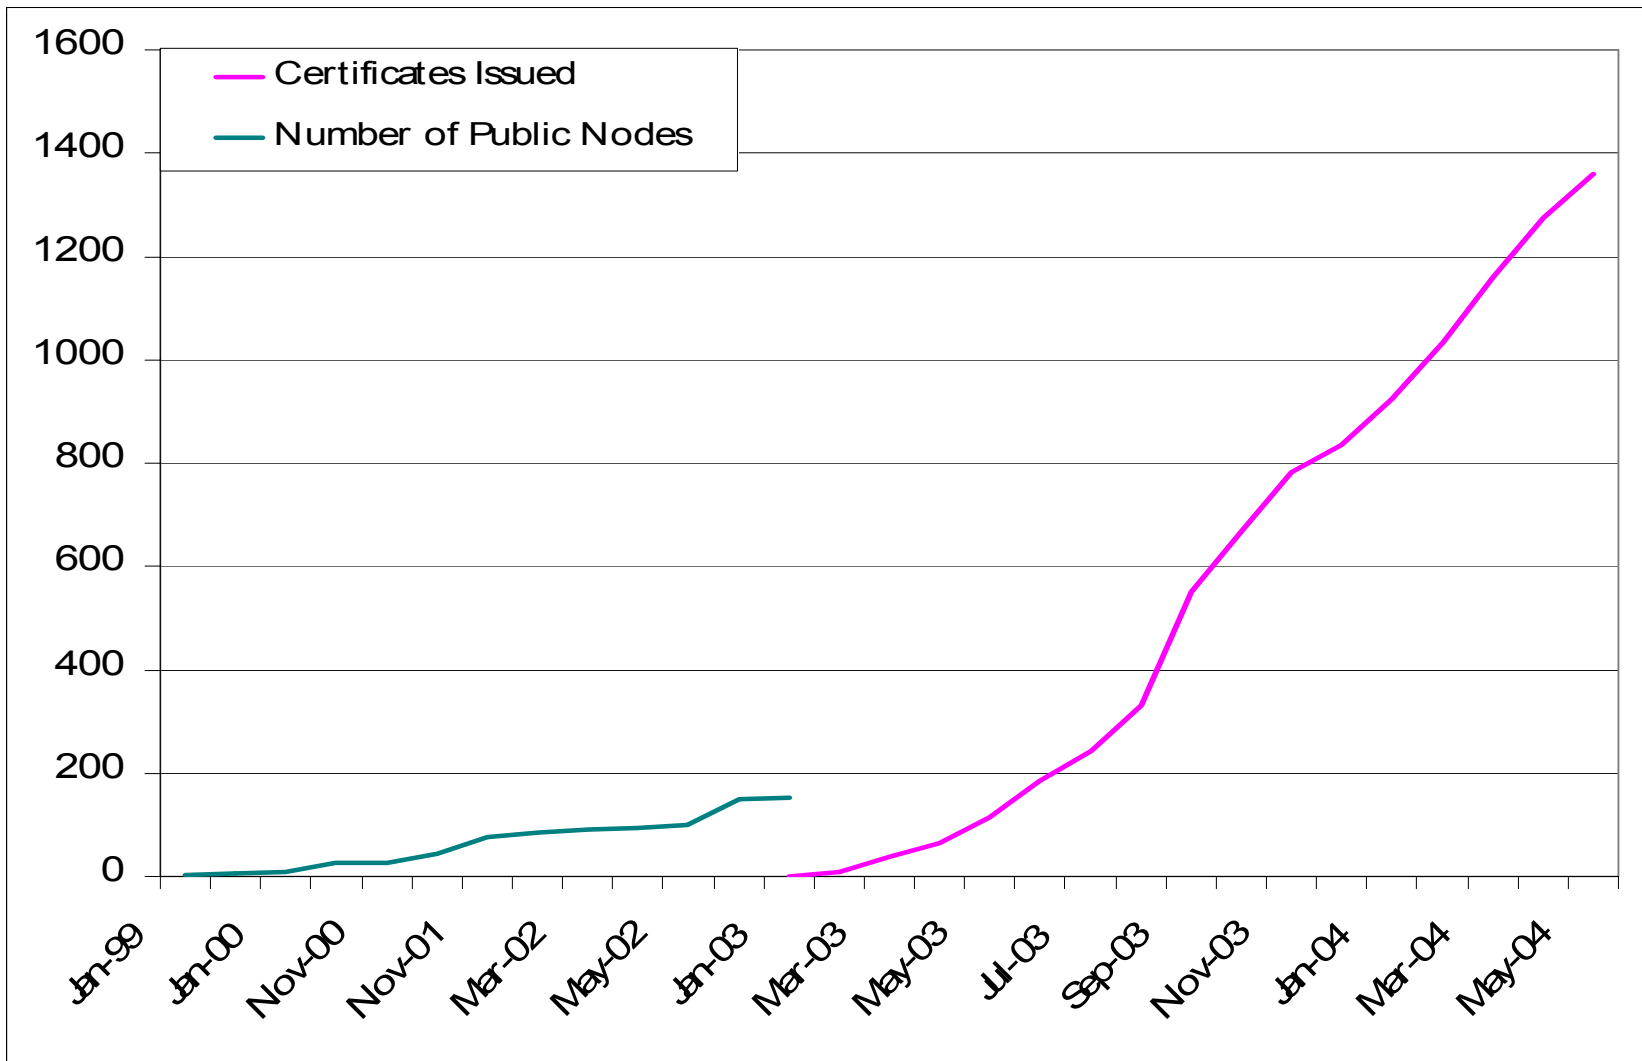
- Simple integration of mbone tools, via apache plug-ins
- 150-250 node installations (each @ \$50k)
 - required multicast
- No notion of users, only nodes
- Limited security and unicast fallback

Access Grid Present

Second Generation, AGTk 2.X:

- Much more complex software structure, leverages Globus Toolkit
- Web Services based, but not Grid Services
 - Interoperability issues
- Moved to individual user credentials, not node credentials, numbers have grown to ~ 2500 credentials
- Richer Services Layer
- Community Development, Packaging
- Broader platform support
- Security enabled from the ground up
- Multicast still an issue
- Demonstrated HD and DVTS Video Services, this year

Access Grid Growth



AGTk 2.X Lessons (and Impact)

- AGTk 2.X has not been as stable as AGTk 1.X
- Four things changed from 1.X → 2.X:
 - Underlying Technology
 - The Development Team
 - The Timeline
 - System Complexity
- These four factors caused a rough first 9-12 months of AGTk 2.X, but it has now settled down *mostly*

AGTk 2.X Commercialization Effects

- Multiple companies have tried to commercialize the AG software
- Most started during AGTk 1.X
- AG 2.X didn't not "help" the commercialization
- The commercial solutions are AGTk 1.X based
- There has been a failure of interop between the commercial and research
- There have been efforts to re-converge the solution, and a commitment from the commercial providers to have interop by September, 2005

Access Grid Future

Third Generation, AGTk 3.X:

- Completion of the Toolkit, harden existing services
- Web Services Based, BP1.0 Compliance
- Security enabled, backing off from GSI to stock SSL
 - Include an integrated “KeyChain” for managing multiple credentials
- Application Level Networking (in the Network Services Middleware) to address Multicast issues
- IPv6 Support integrated


Beyond AGTk 3.X:

- Identify interesting areas where collaboration can accelerate science, then use the AGTk as a vehicle to extract that potential speed-up
- Investigate interesting “high-end” and “low-end” node environments
- Refine Core interfaces

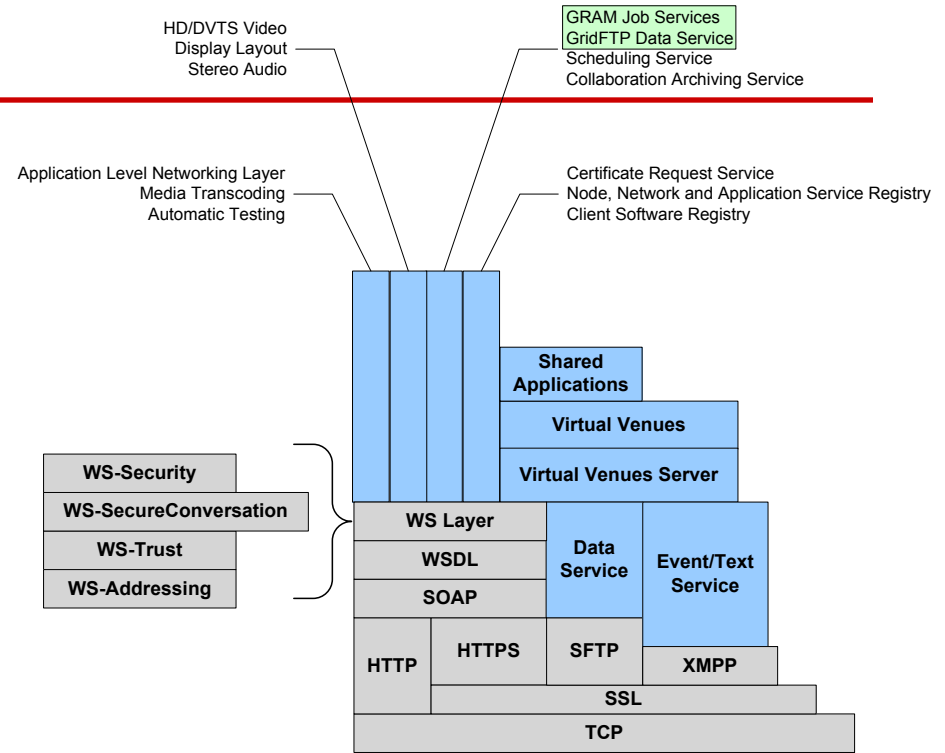
AGTk 3.X Plan (Lessons Learned)

- AGTk 3.X needs to be more stable than AGTk 2.X
- The things changing from 2.X → 3.X:
 - Underlying Technology
 - The Development Team
 - The Timeline
 - System Complexity
- We are becoming more conservative on what software we rely on and ship with the Toolkit.


AGTk Technology Changes 2.X → 3.X




Access Grid Toolkit 2.X Services Infrastructure



Access Grid Toolkit 3.X Services Infrastructure



Access Grid Toolkit Media Software



Access Grid Toolkit Security Software



Access Grid: Moving from Present → Future

- Encourage Infrastructure Support for Collaboration
 - Get others to support:
 - Multicast as a production capability
 - Venue Servers, Bridge Services, Account Mgmt Services
 - Help lines, troubleshooting, “on-call” for transient multicast issues
- Examples
 - Internet2 Commons
 - ES Net Collaboration Services

Considerations for ES Net?

- Collaboration is complicated, it's not just videoconferencing any more.
- There are *no technological* solutions to **social** problems.
- Products are not static, H.323 is moving toward SIP, that rollover may not happen for 3-5 years, but it will happen.
- The product side is much more stable than the research side.
- What are the products and services that will be part of the ES net “solution portfolio”?
 - It needs to be a diverse set, c.f. hammer/nail problem
- What integrated services layer can/will ES Net provide for all the services it offers?
 - Scheduling?
 - Accounting?

Questions for ES Net?

- What's the vision for the future?
- Is ES net?
 - a) A Production Network,
 - b) A Research Network,
 - c) All of the Above
- What role (if any) does ES Net want to play in the adoption and deployment of research solutions into production solutions?
- How will ES Net bridge the administrative domains of the constituent facilities? Can it? Is it interested?
- What resources are required to effectively deploy and support products, services and infrastructure for collaboration?

Recommendations for ES Net

- Maintain a Technology Agnostic Disposition
- Offer an integrated product and services solution
- Greedily, but with a high standard, look for DOE developed technologies that can enhance products and service portfolio provided by ES net.
- Build strategic alliances with places like Internet2 to provide “cross-realm” relationships to enable easier collaboration with strategic partners (DOE and EDU), based on common products and services

Concrete Access Grid Related Suggestions

- Setup Certificate Authorities to be easier to use (as the AG has)
- Provide Access Grid Venues Services, with scheduling services
- Provide Access Grid Network Services (Bridging, etc)
- Establish a tiered service model incorporating lab-based groups, with ES net at the top, then labs, then research groups