



*“Open Source, OGSA Implementation”*

Genesis II:  
*Accessible, Standards Based  
Grid Computing*

Mark Morgan, Andrew Grimshaw

Global Bio Grid Team

University of Virginia



# Responsible Parties

- Andrew Grimshaw
- Mark Morgan
- John Karpovich
- Duane Merrill
- Howie Huang
- Krasimira Kapitanova
- Karolina Sarnowska
- Chris Sosa



# What's the Problem

Our target grid users are unable to, or unwilling to learn new programming languages, coding paradigms, or complicated tooling.

Users want the benefit of the grid, but they want it transparently!



# Abstract

- Sufficient richness of specification exists in OGF and OGSA to produce interesting and useful grids
- Genesis II is an open source, implementation of these specifications, modeled after the Legion grid system, providing both compute and data grid technology
- **By focusing on familiar, traditional abstractions such as files and directories we better serve the target grid user**



# Outline

- Background
- Genesis II
- Summary



# Background

- Specifications cannot exist in vacuum
  - Implementations Vet Specifications
  - Implementation experience shows how spec's interact
- Grids have been around for a while, but adoption remains low – why?
  - Usability
  - Gap between grid designers and grid users

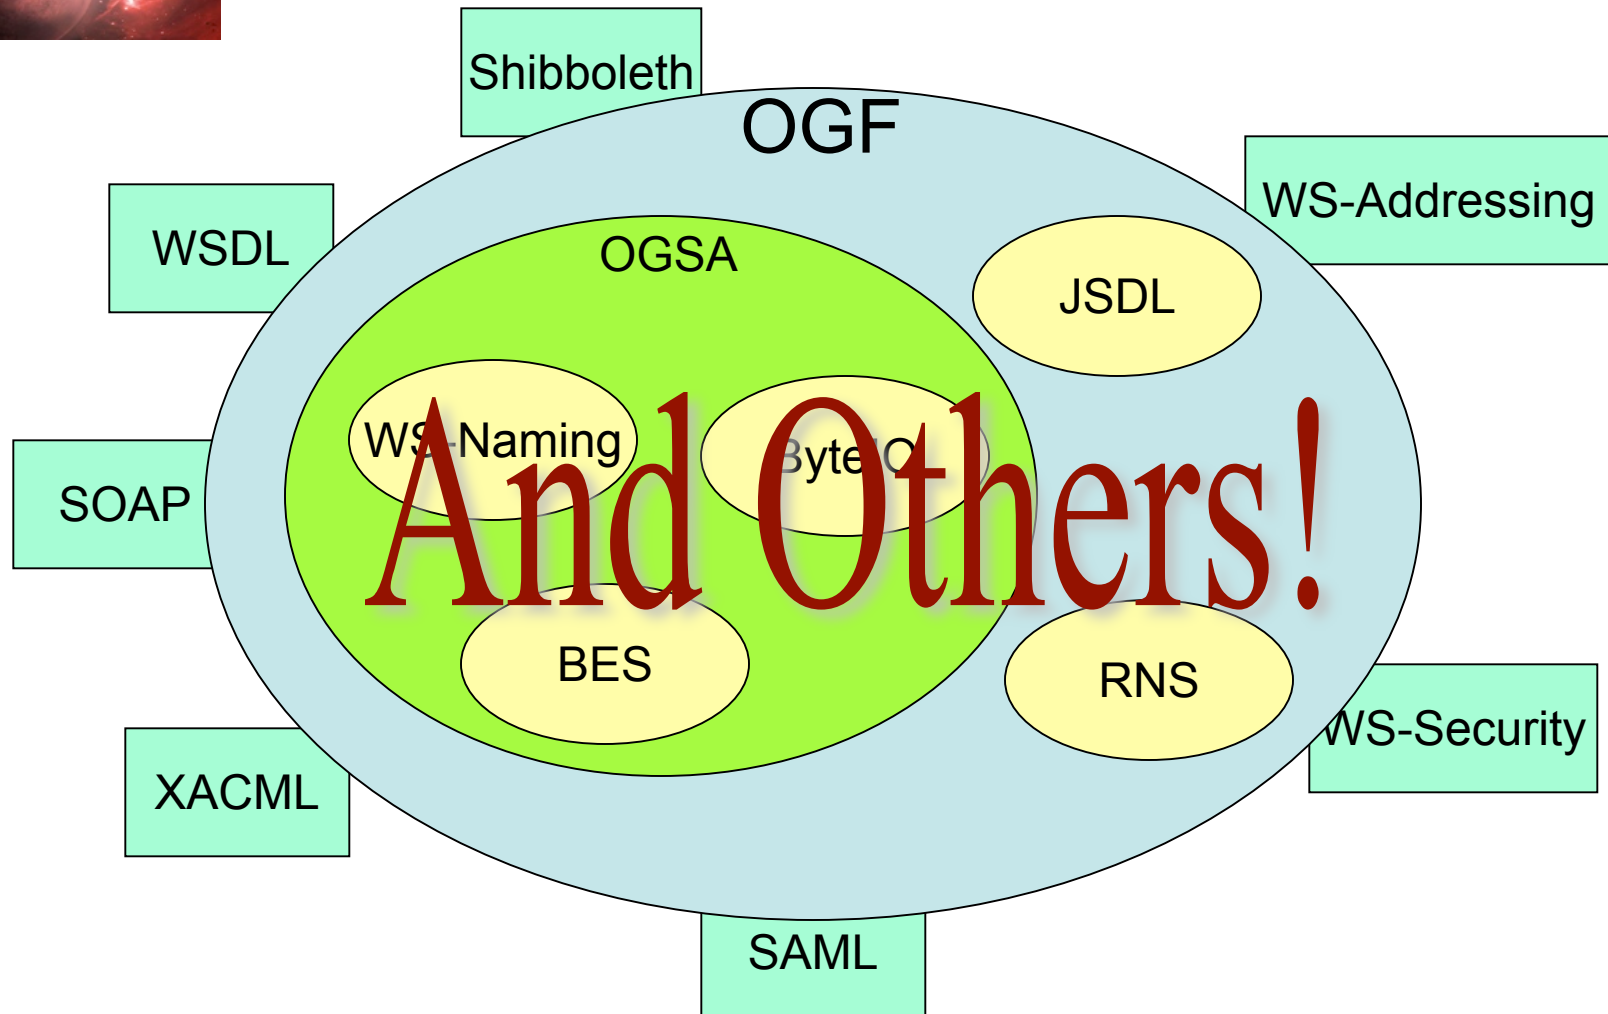


# Genesis II Goals

- Provide an open source, reference implementation of the OGSA and OGSA-related specifications
- Use standards and proto-standards available from the OGF and OGSA to
  - Provide a secure, cohesive system in a production system available to users today!
  - Provide feedback into the OGF process on various standards based on implementation experience
  - Design the system from the ground up with the overriding mantra that **users come first!**



# Standards Represented in Genesis II

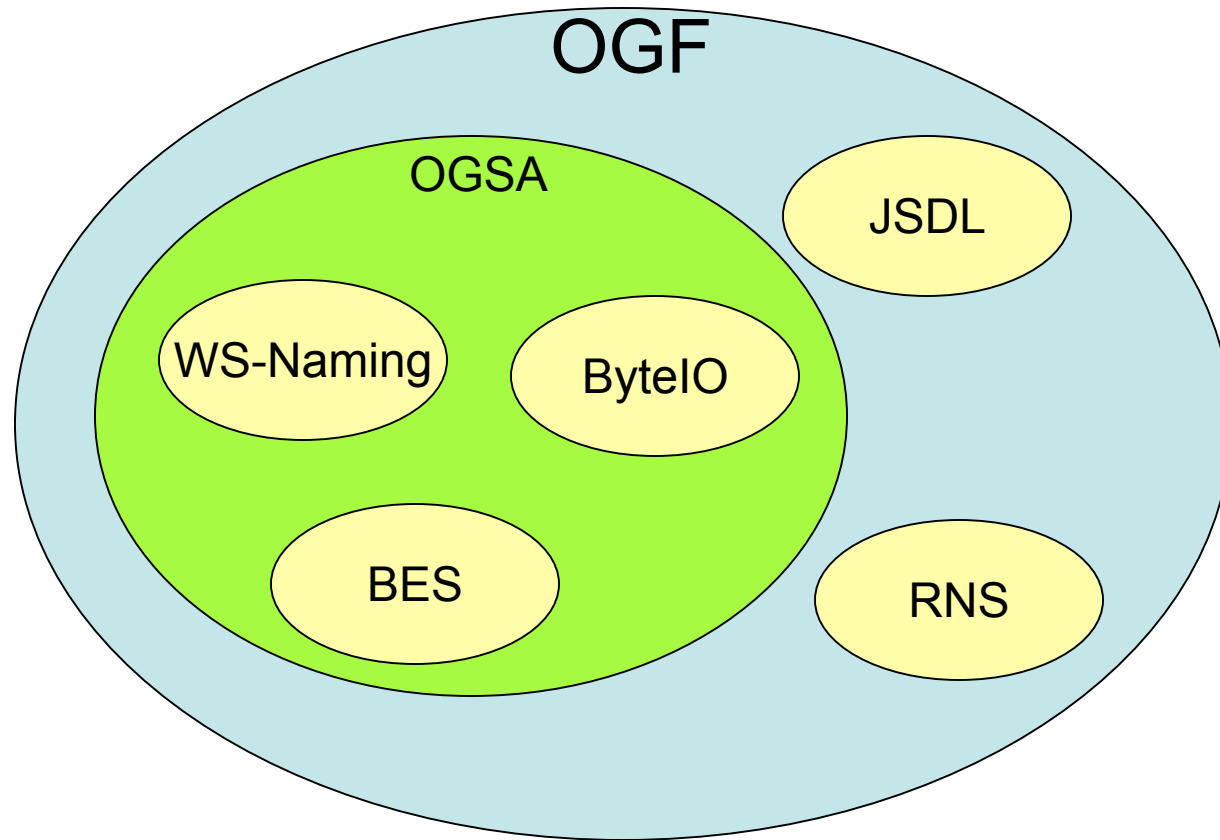


**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"





# Standards Represented in Genesis II

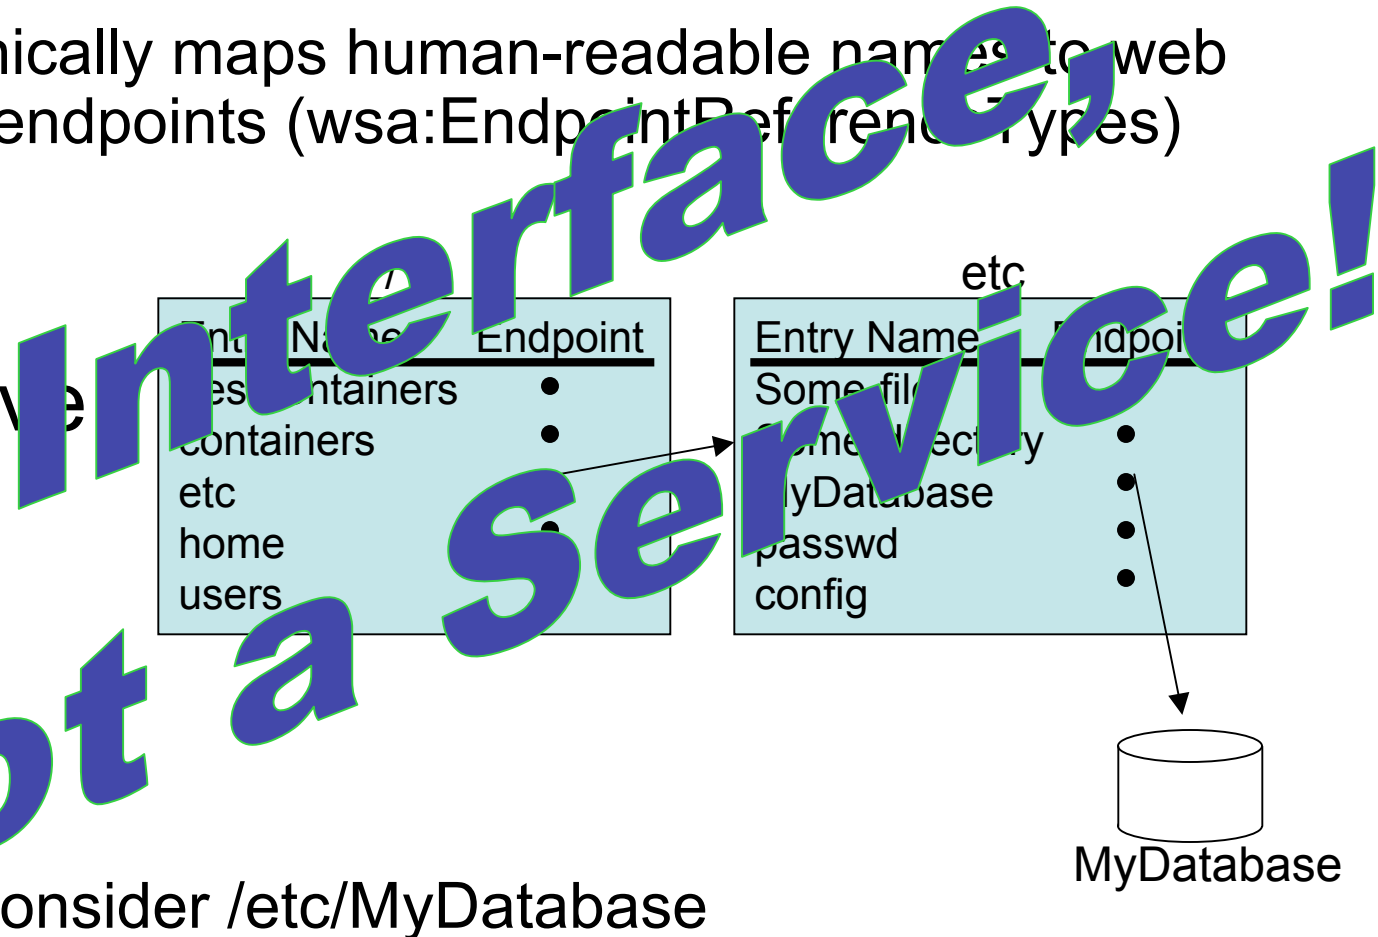




# RNS (Resource Naming Service)

- Hierarchically maps human-readable names to web service endpoints (wsa:EndpointReferenceTypes)

- Add
- Remove
- List





# ByteIO

- Posix-*like* data IO
- Treat a resource as if it were a file
- Familiar operations
- Read, write, seek, truncate, append, etc.
- Ideal for mapping *many* familiar abstractions

**Interface,  
Not a Service!**



# BES (Basic Execution Service)

- Service interface for starting and managing remote compute jobs
- Implementation is not specified
  - Queue
  - Fork/exec
  - Virtualize
  - Etc.\
- Emphasis on Basic



# WS-Naming

- Two components
  - Endpoint Identifiers
  - Endpoint Resolution

```
<EPR>
  <Address>http://tempuri.org/some-service</Address>
  <ReferenceParameters> xs:any </ReferenceParameters>
  <Metadata>
    <EndpointIdentifier> xs:anyURI </EndpointIdentifier>
    <EndpointIdentifierResolver>
      wsa:EndpointReferenceType
    </EndpointIdentifierResolver> *
  </Metadata>
</EPR>
```



# Users First

- A large percentage of a grid's target audience is unable or unwilling to learn new interaction abstractions
- Instead of asking the user to adapt to the grid, we should adapt the grid to the user



# User Abstractions

- One of the most ubiquitous user interaction abstractions is the file system
  - Drag-and-drop
  - Double Click
  - Named pipes
  - /proc filesystem
  - Plan 9
- RNS, ByteIO, and WS-Naming provide the foundation for building these abstractions



# Thought Games

- What would it mean to browse a grid “directory” structure?
- What if you double-clicked on a BytelO resource in it?
- What if you double-clicked on a grid resource representing a database query?
- What about dragging a JSDL document onto a BES container? A scheduler?
- What about “browsing” into a BES container?





# Outline

- Background
- **Genesis II**
- Summary



# Genesis II

- Standards based, open source, production level grid system
- Designed with the primary goal of putting users first
- Provides data grid and compute grid technology using secure infrastructure



# Genesis II Specifications

- Written in Java 1.6.x using Jetty 5, Axis 1.4, Apache Derby, and WSS4J
- Pluggable security infrastructure currently supporting both Username/Password profile, and GAML (Genesis II SAML implementation)
- Currently tested on Windows XP and Linux



# Genesis II

- Every service/resource in Genesis II may implement RNS and/or ByteIO (and most do)



# UNIX-like command line interface

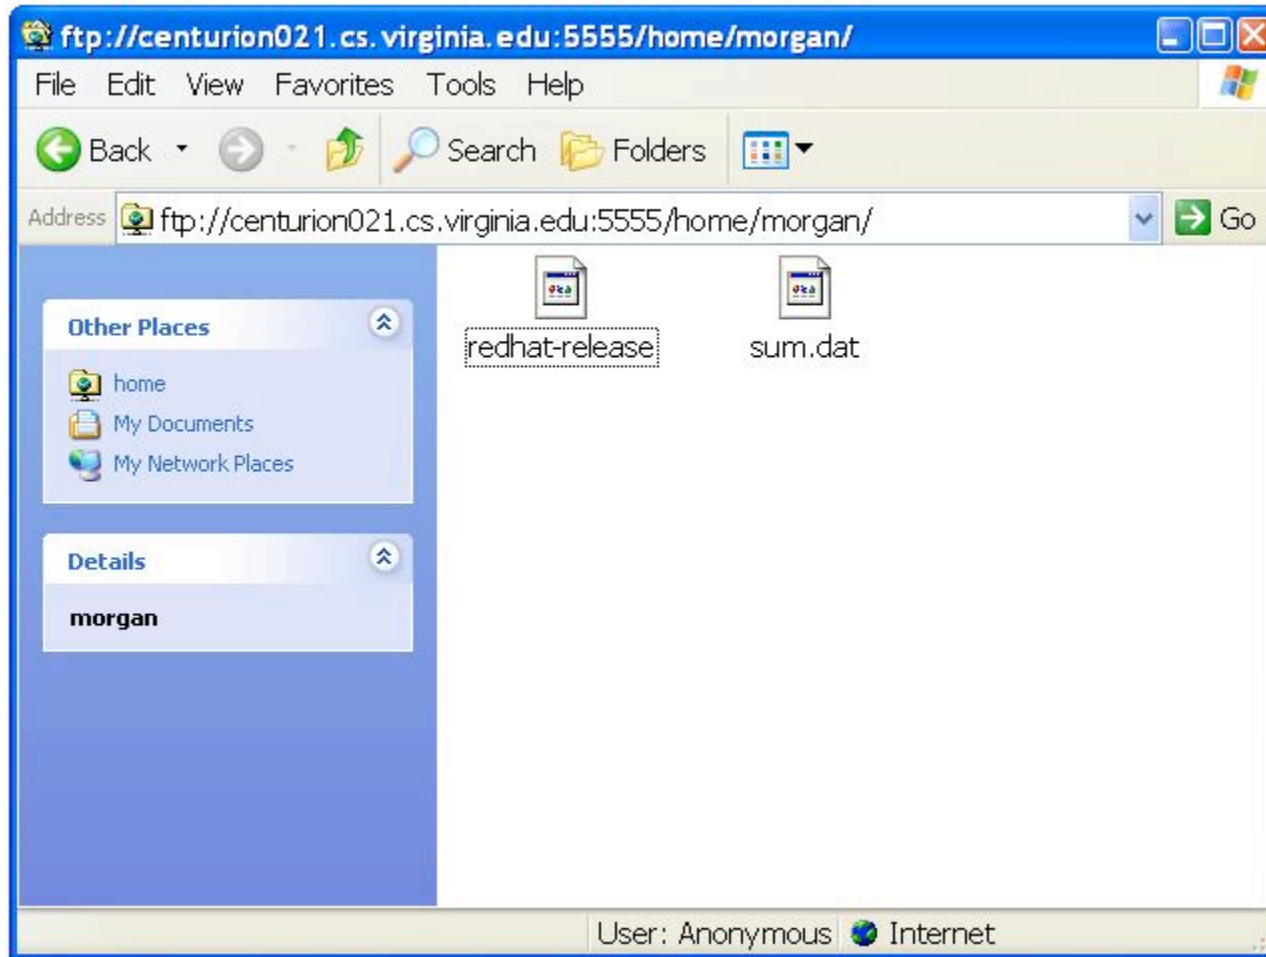
```
X gbg@centurion021:~/OGF/Genesis II
[gbg@centurion021 Genesis II]$ ./grid whoami
GAML: CN=Mark M. Morgan 6, EMAILADDRESS=mmm2a@virginia.edu, OU=UVA Standard PKI
User, O=University of Virginia, C=US
[gbg@centurion021 Genesis II]$ ./grid pwd
/
[gbg@centurion021 Genesis II]$ ./grid cd /home/morgan
[gbg@centurion021 Genesis II]$ ./grid cp --local-src /etc/redhat-release redhat-
release
[gbg@centurion021 Genesis II]$ ./grid ls -la
morgan:
31          redhat-release
3           sum.dat

[gbg@centurion021 Genesis II]$ ./grid cat redhat-release
Fedora Core release 4 (Stentz)
[gbg@centurion021 Genesis II]$ █
```

**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"



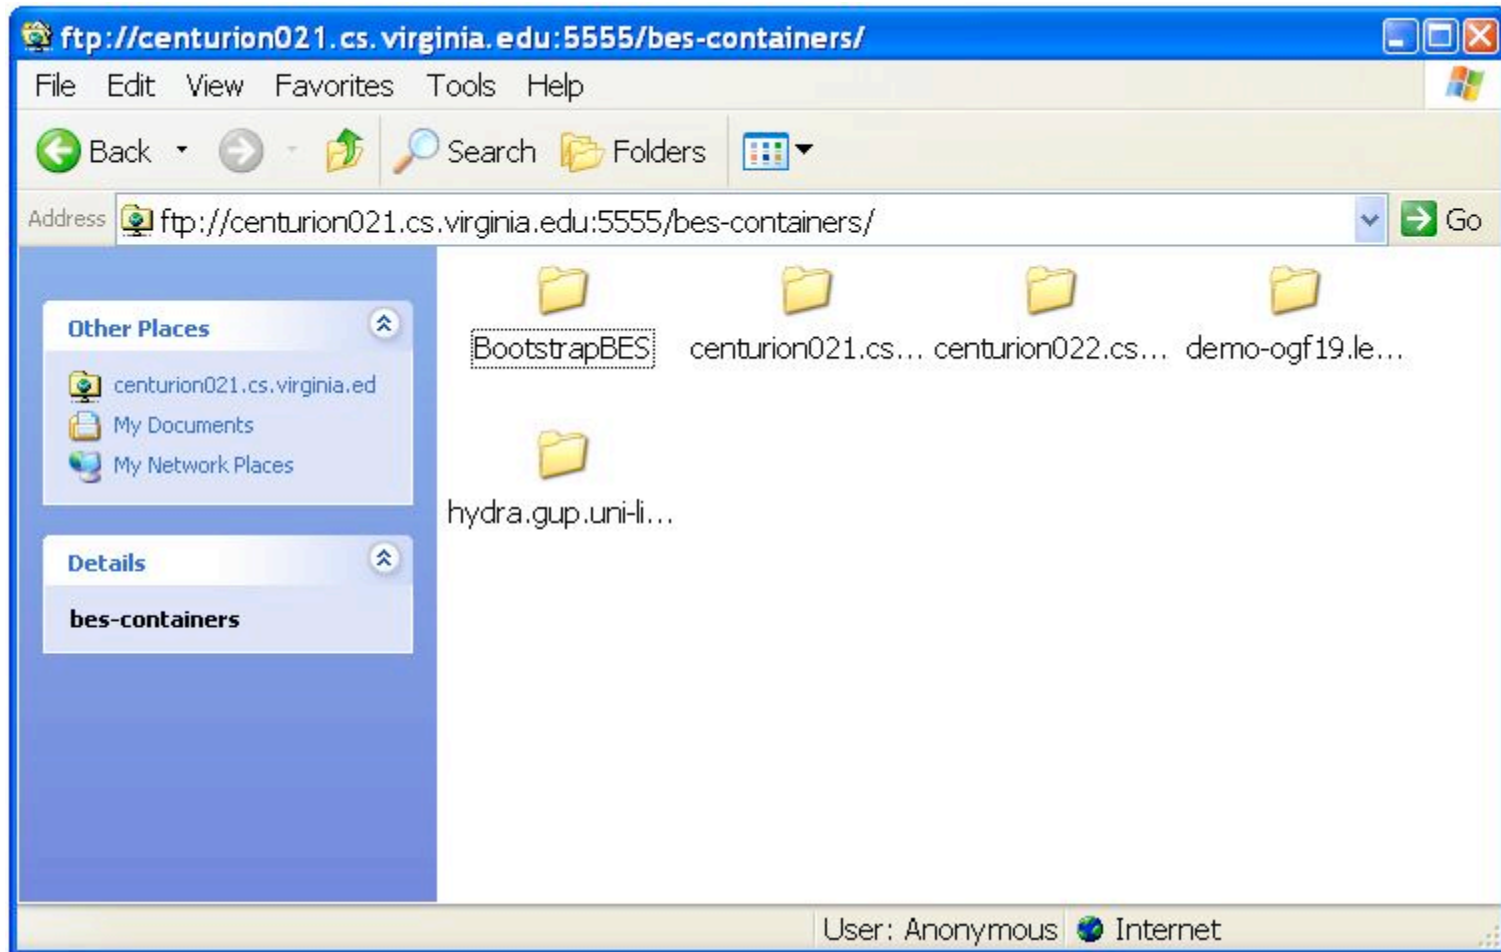
# Filesystem Aware Interfaces



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
“Open Source, OGSA Implementation”



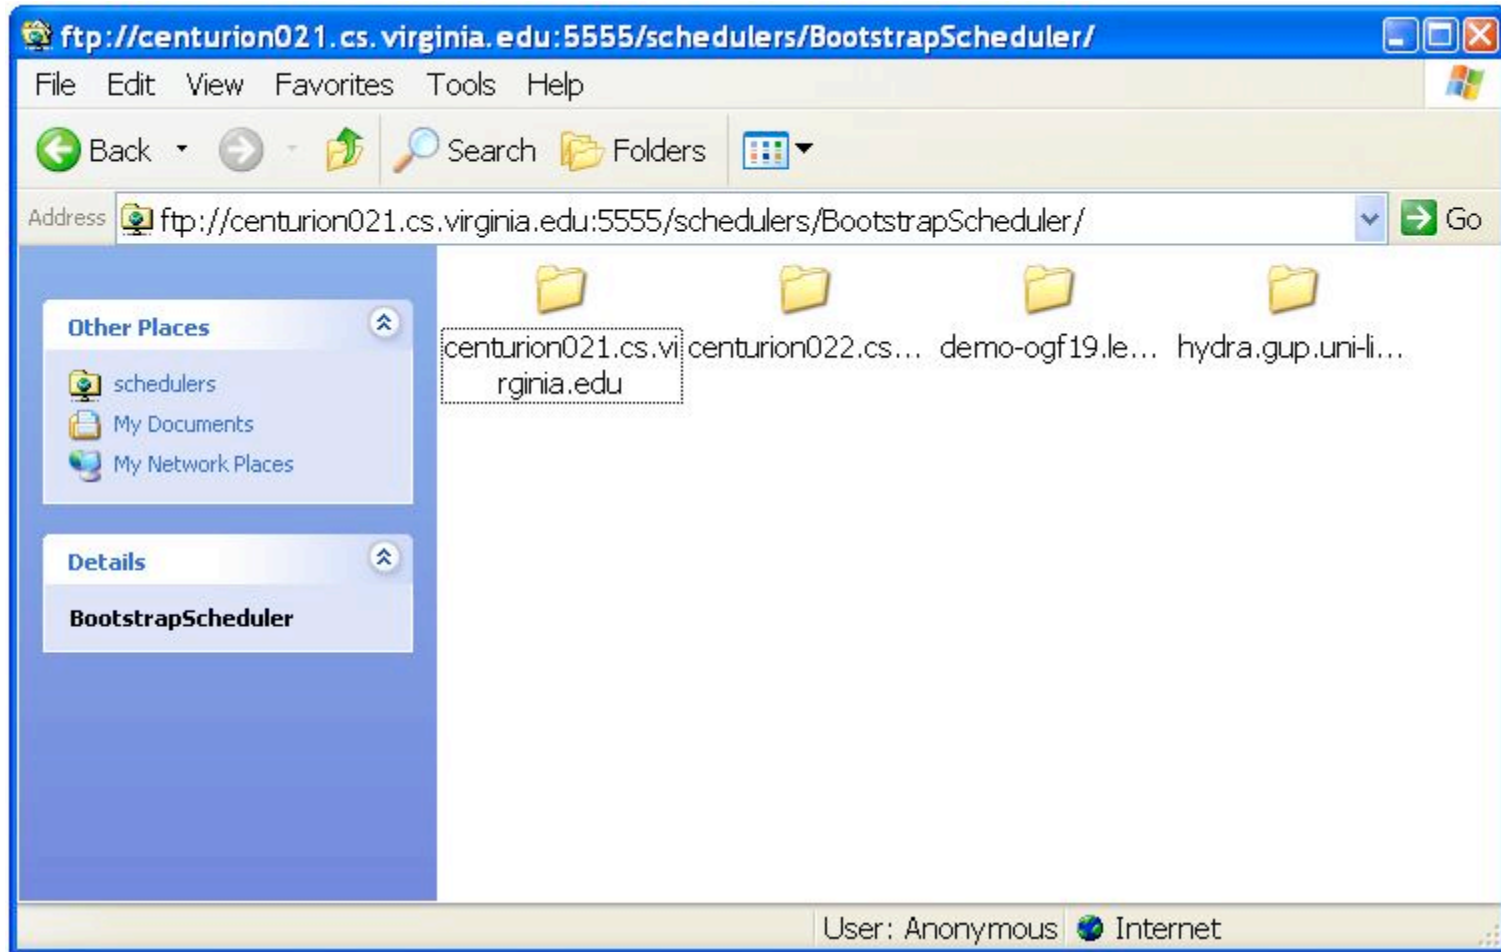
# Using RNS to name non-filesystem components



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
“Open Source, OGSA Implementation”



# Using RNS to name non-filesystem components

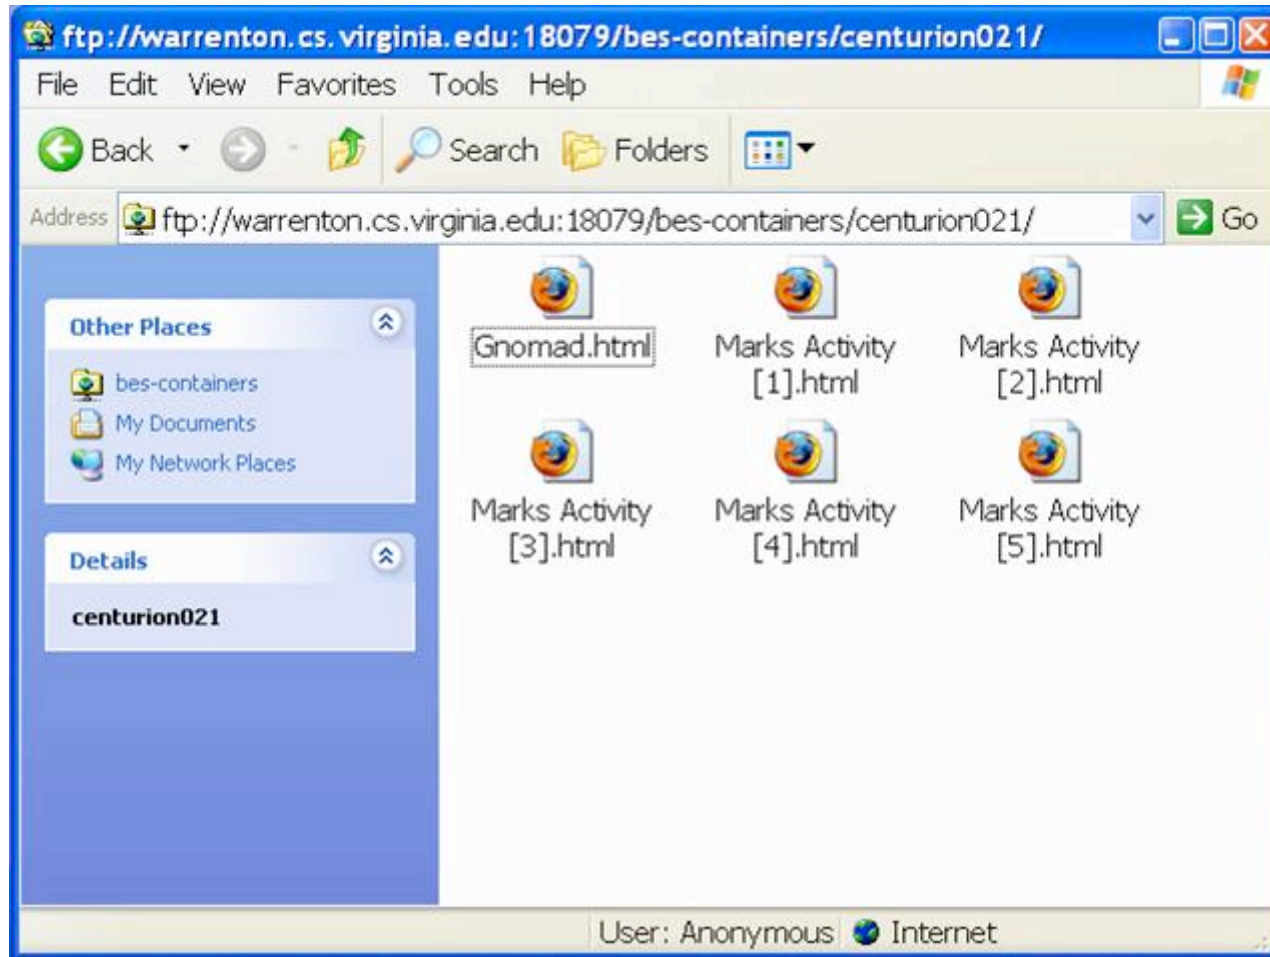


**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"





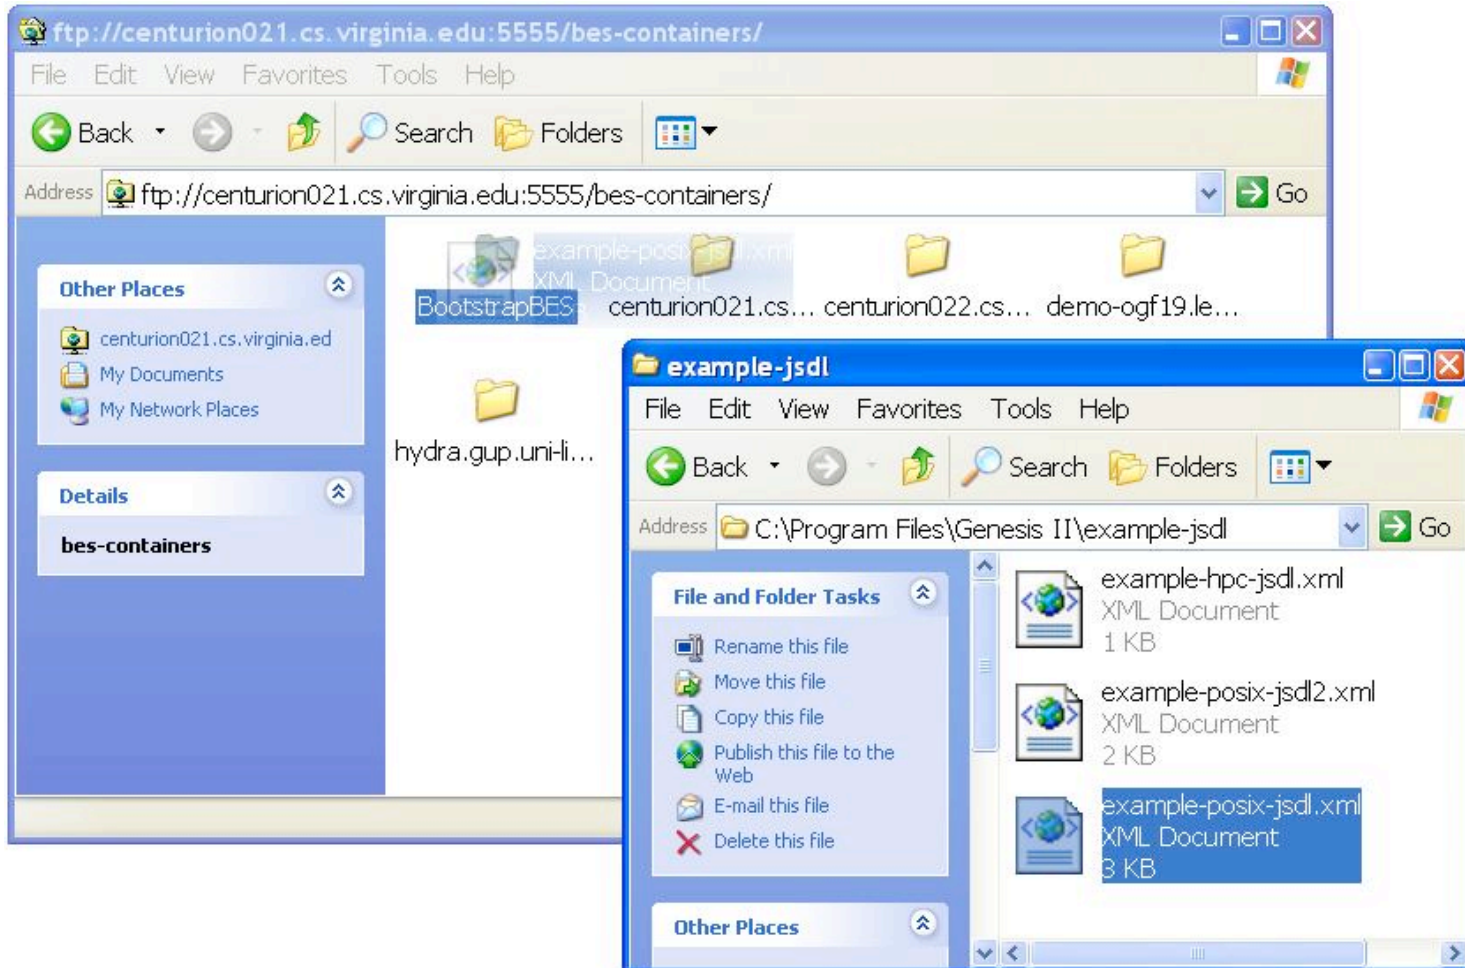
# Genesis II's BES implements RNS too!



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
*"Open Source, OGSA Implementation"*



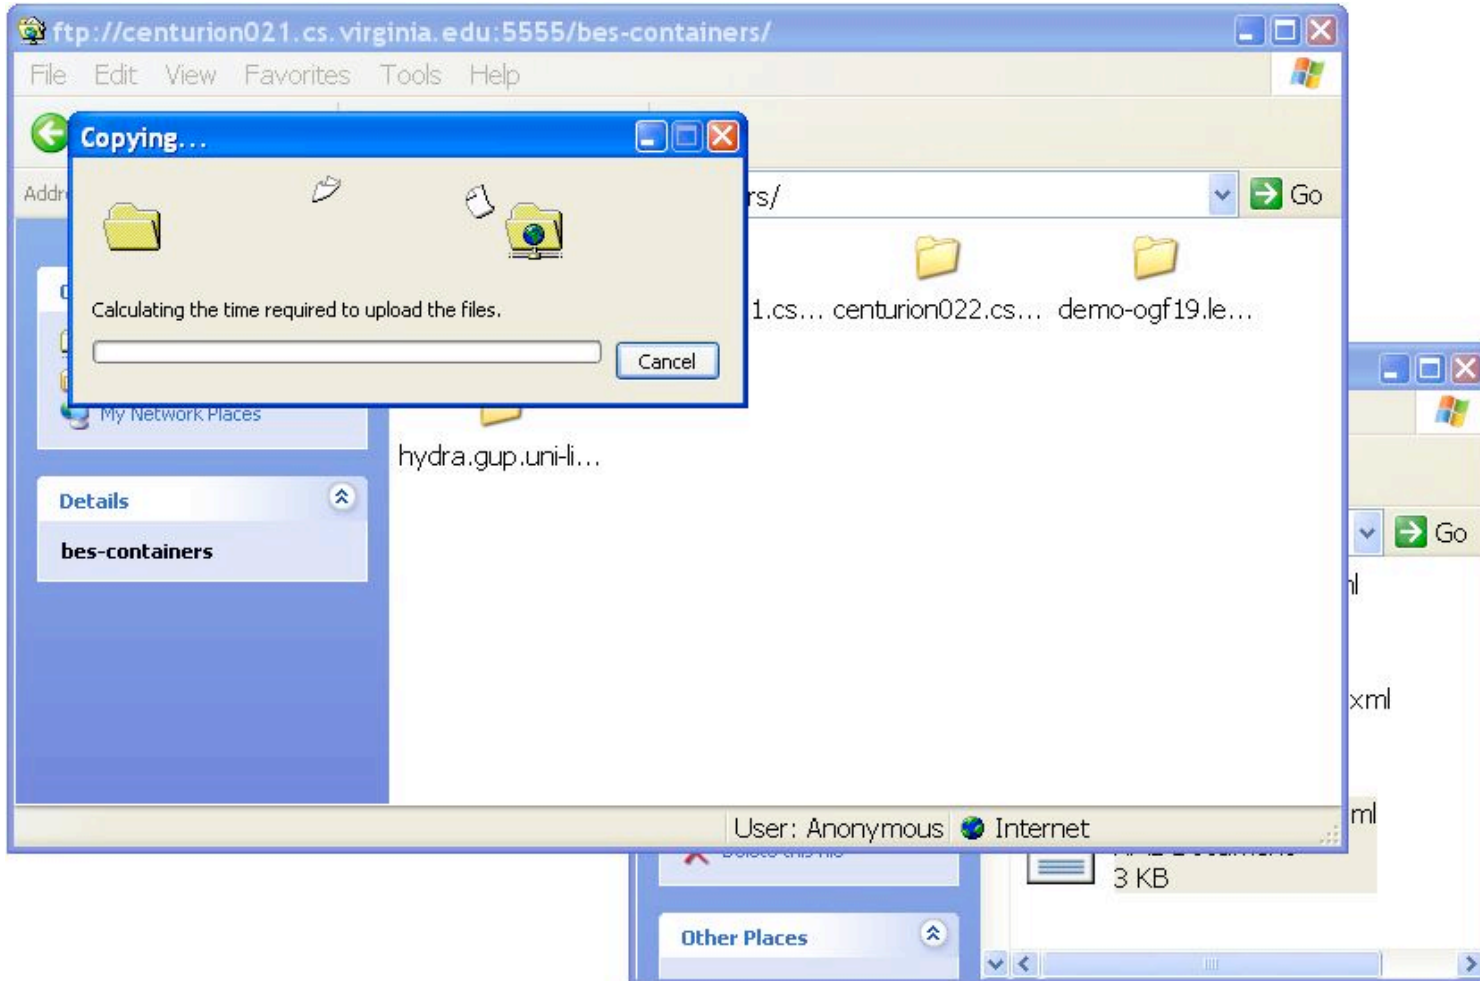
# Genesis II's BES even implements ByteIO!



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"



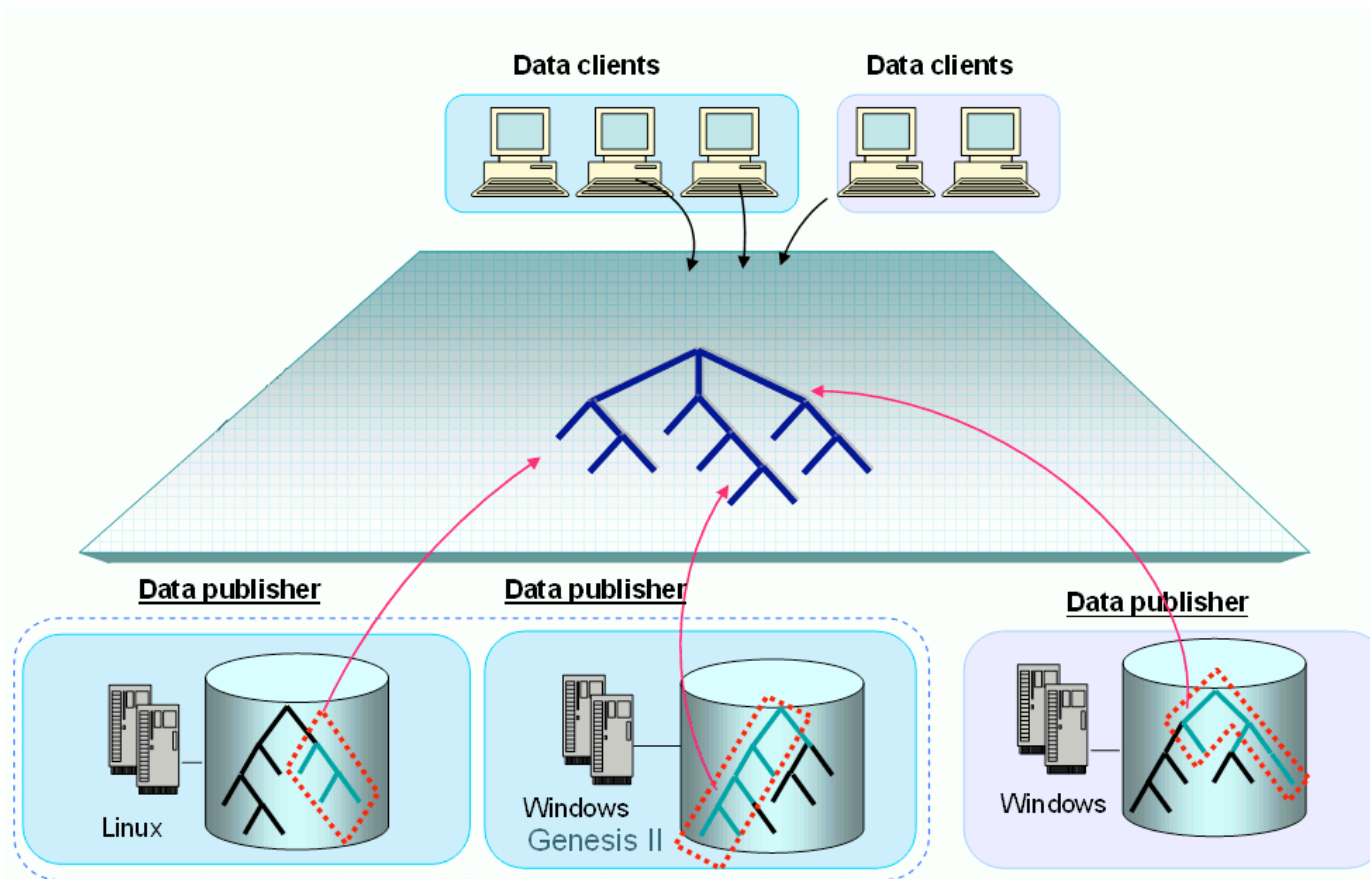
# Genesis II's BES even implements ByteIO!



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"



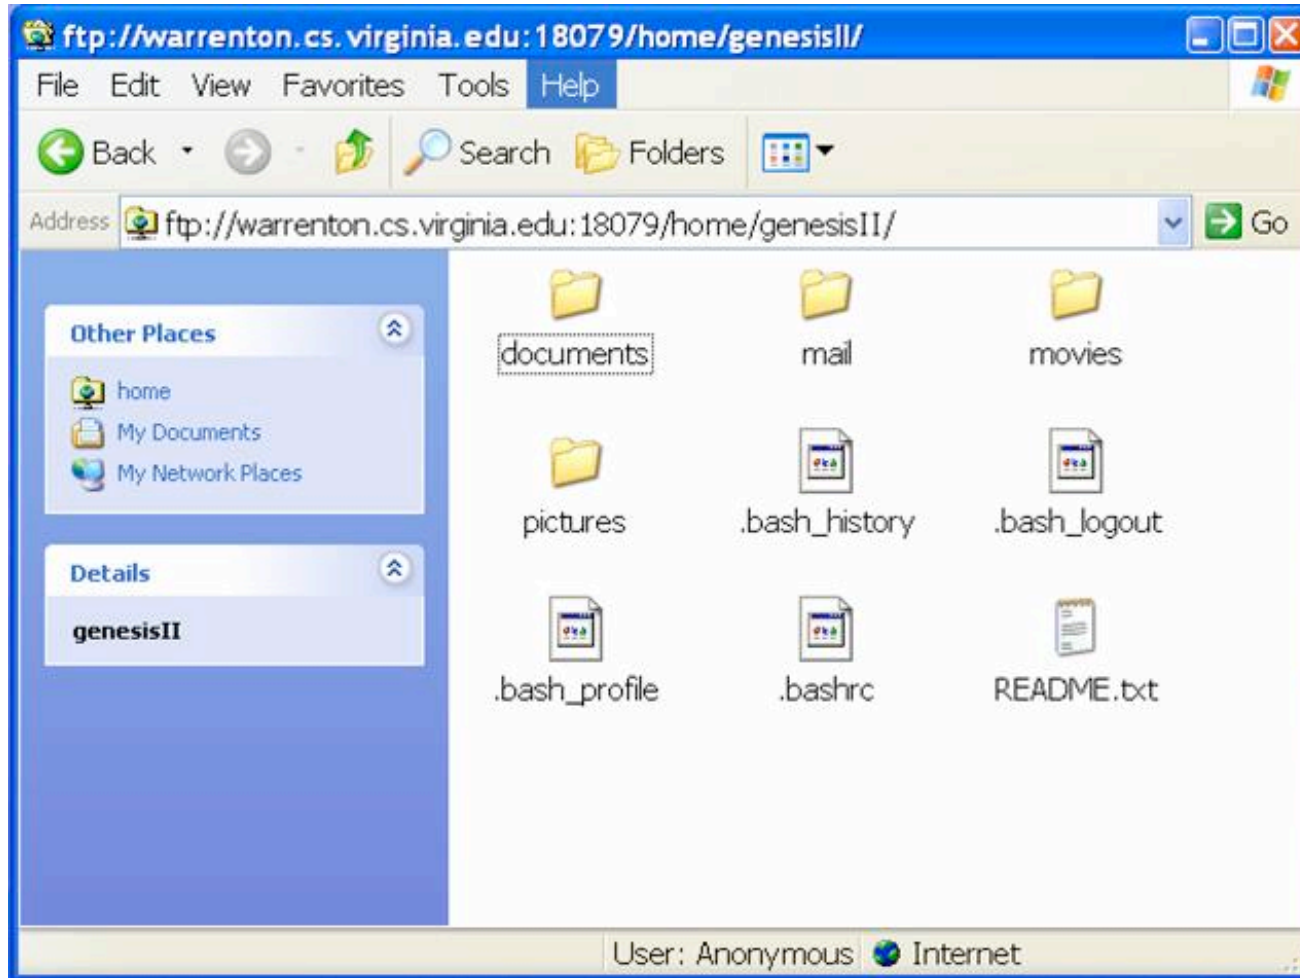
# Export Directory



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"



# Export Directory



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"

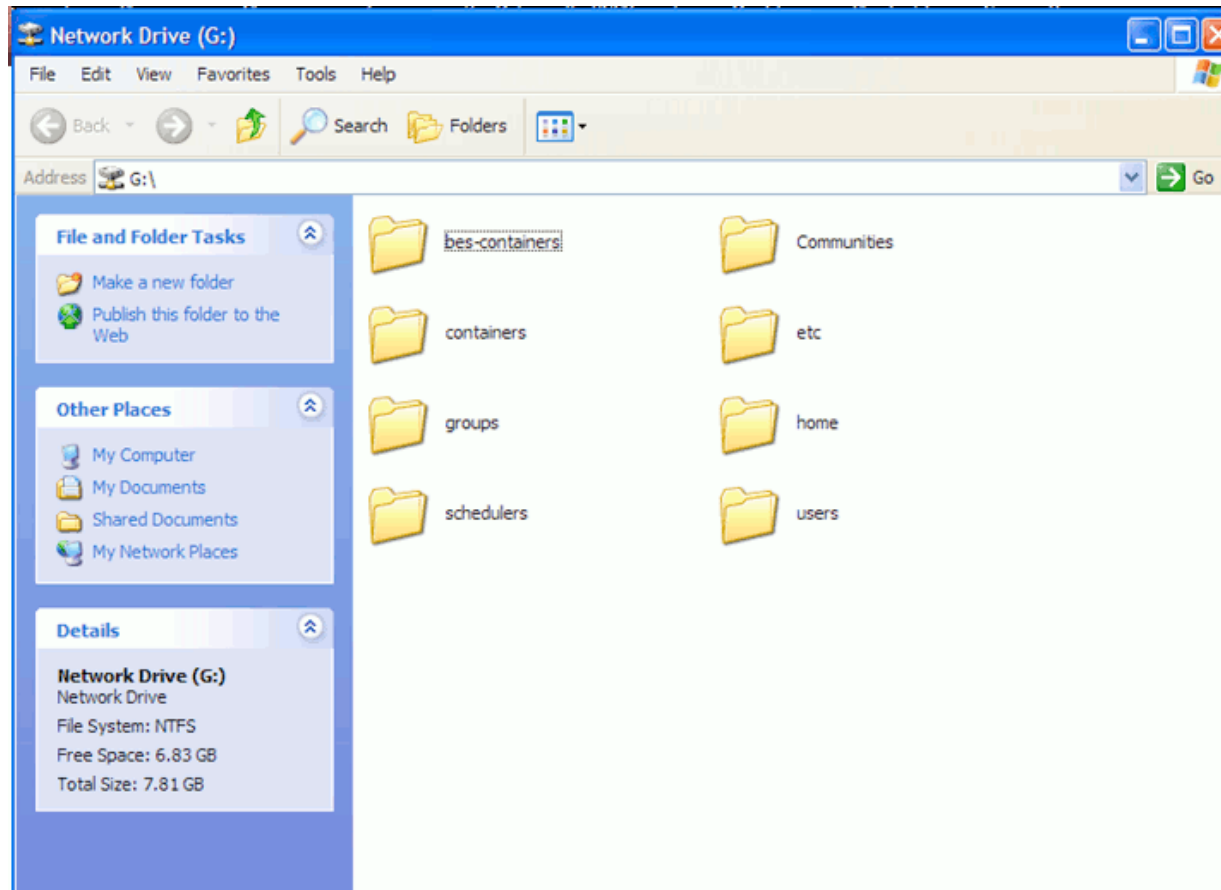


# WS-Naming and Genesis II

- Extensive use of Endpoint Identifiers
- Simple Resolvers
- Future Work
  - Security Modifications
  - Fault Tolerance
  - Performance
  - QoX
  - Etc.



# Other Filesystem Aware Interfaces



**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
"Open Source, OGSA Implementation"




# OGRSH (Open Grid Shell)

- Classic software interposition library or shim

Legacy  
Application

Standard  
Shared  
Libraries








# OGRSH (Open Grid Shell)

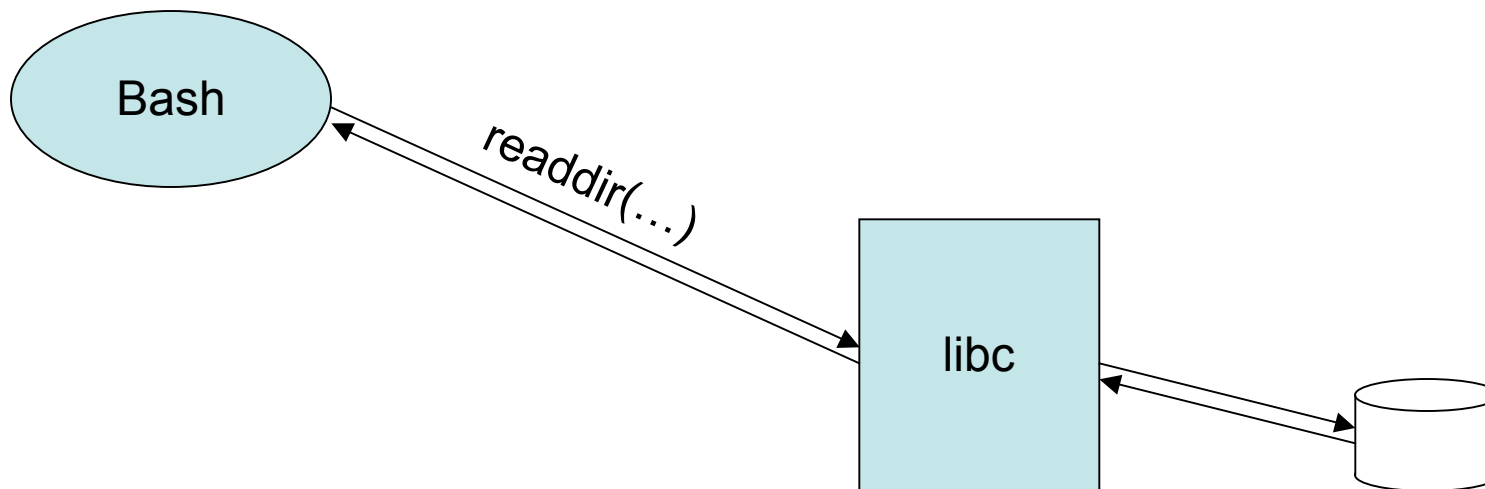
- Classic software interposition library or shim





# OGRSH (Open Grid Shell)


- Classic software interposition library or shim





# OGRSH (Open Grid Shell)

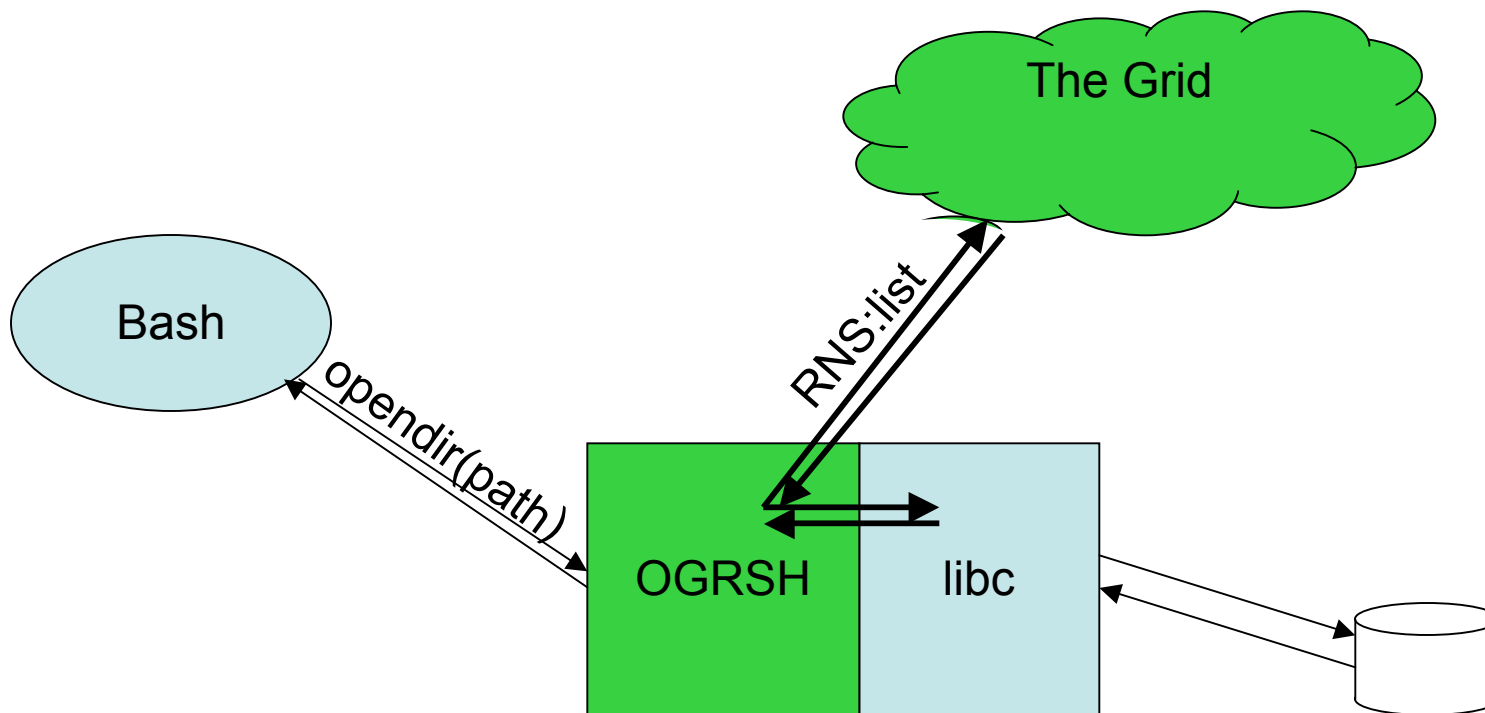
- Classic software interposition library or shim





# OGRSH (Open Grid Shell)

- Classic software interposition library or shim





# OGRSH

**Insert Picture Here**



# Outline

- Background
- Genesis II
- **Summary**



# Summary

- Potential Grid users want the benefits of the grid without the pain.
- Grid uptake therefore is closely tied to usability
- As some systems have demonstrated in the past, users are better at learning new semantics than new syntax.
- Genesis II leverages this by providing the familiar syntax or abstractions of file systems to perform “everyday” grid activities.



# Genesis II

## *Take-away messages*

- Open Source implementation of the OGF and OGSA standards available
- **Sufficient body of standards exist with which to build interesting, useful grid systems**
- Very active project!
- Information and Download Page
  - <http://vcgr.cs.virginia.edu/genesisII>
- Forum
  - <http://www.cs.virginia.edu/forums/viewforum.php?f=26>





# Questions?

<http://vcgr.cs.virginia.edu/genesisII>

**Genesis II** (<http://www.cs.virginia.edu/~vcgr>)  
*“Open Source, OGSA Implementation”*



# Security Design Goals

- **Pluggable security modules (at run time)**
- Secure communication: Authentication, Confidentiality, Integrity, Authorization
- Standards-based:
  - Follow OGSA AuthN & AuthZ models as they develop
  - Leverage standard protocols, specifically Username/Password, X.509, and SAML WS-Security token profiles



# Authorization

- Subject identity (LDAP, NIS, X.509, etc.) is federated via Attribute Authorities (AA)
- Credentials are signed, “holder-of-key” SAML attribute assertions (e.g., identity, roles, capabilities, etc.)
- Credentials may be delegated, but have limited lifetimes
- Attribute Authorities implement the SAML Authentication Request profile (and optionally the Assertion Query and Request profile for a pull-style model enabling Shibboleth-like anonymity)



# Authentication, Integrity, and Confidentiality

- Asymmetric public-key cryptography
- Identity via trust hierarchies: X.509 digital certificates
- Resource identity is correlated to digital certificates by embedding WS-Naming EPIs as a X.509 non-critical extension
- Digital certificate distribution via EPRs: certificate-chains embedded as non-critical EPR metadata
- Server-side authentication (as opposed to mutual authentication which restricts privacy, anonymity, delegation)
- For message integrity, the caller can generate a keypair



# Authentication, Integrity, and Confidentiality

