



Nirvana
A DIVISION OF GENERAL ATOMICS

US DoD HPCMP Data Grid May 4, 2010

Constantin Scheder, Director and Chief Architect
Nirvana division of General Atomics
858-455-2536, scheder@ga.com



General Atomics & Nirvana

- **1955 – General Atomics (GA) founded San Diego**
Hi-tech systems development:
 - Photonics, lasers, UltraWideBand wireless
 - Nuclear fission and fusion, Uranium
 - Predator and Warrior UAVs
 - Electromagnetic carrier launch systems
- **1985 – GA founds San Diego Supercomputer Center**
- **1995 – Storage Resource Broker developed at SDSC**
- **2000 – Nirvana chartered to commercialize SRB**
- **2002 – 100th Data grid: NASA Remote Data Store**
- **2007 – 150th Data grid: DHS ICE C3**
- **2008 – first commercial implementation: infoUSA**
- **2009 – Largest HPC federation project to date:**
 - 30PB SLM Federation across 6 major DoD Research Facilities



GA is a global enterprise with a successful track record in complex integration projects. SRB is in its 14th year of development and is at the core of the world's most demanding grid implementations.

DoD HPCMP SLM Project

A Challenging Environment

- 6 main centers, over 4,500 users, more than 170 total sites
- 30 PB (dual copy)/ 330M files today with 40% annual growth
- 115 PB in 4 years

An Ambitious Solution

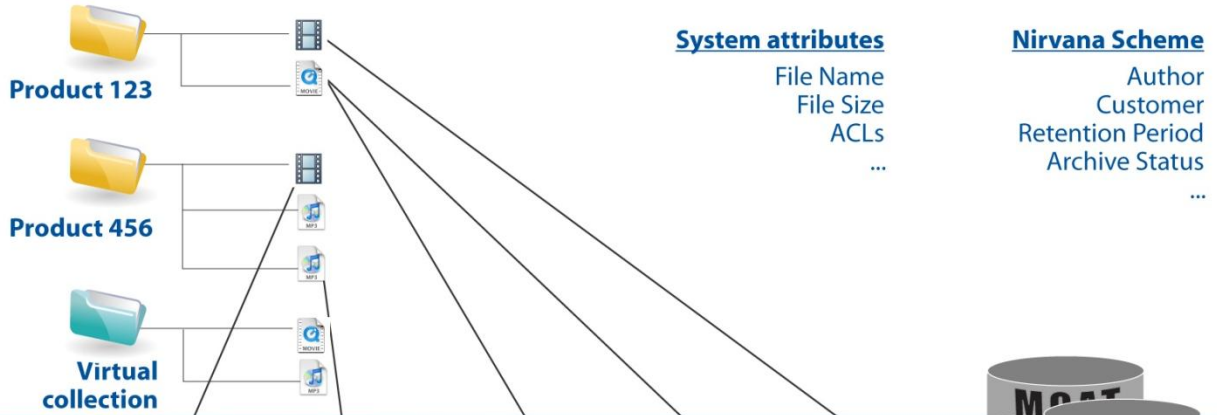
- **Storage Lifecycle Management (SLM) = ILM + HSM**
- **Information Lifecycle Management (ILM)**
 - Metadata for better data management
 - Users decide what to archive and when to create offsite copy
 - Central policies for archival, synchronization, and expiration
 - Virtual data views
 - Global Namespace
- **Hierarchical Storage Management (HSM)**
 - Existing archival infrastructure (software + hardware)
 - Near-real-time file system sync with ILM
- **User Interfaces**
 - Command-line
 - Global file system
 - Web
- **Single-console Administration**




Nirvana Global Namespace & Metadata



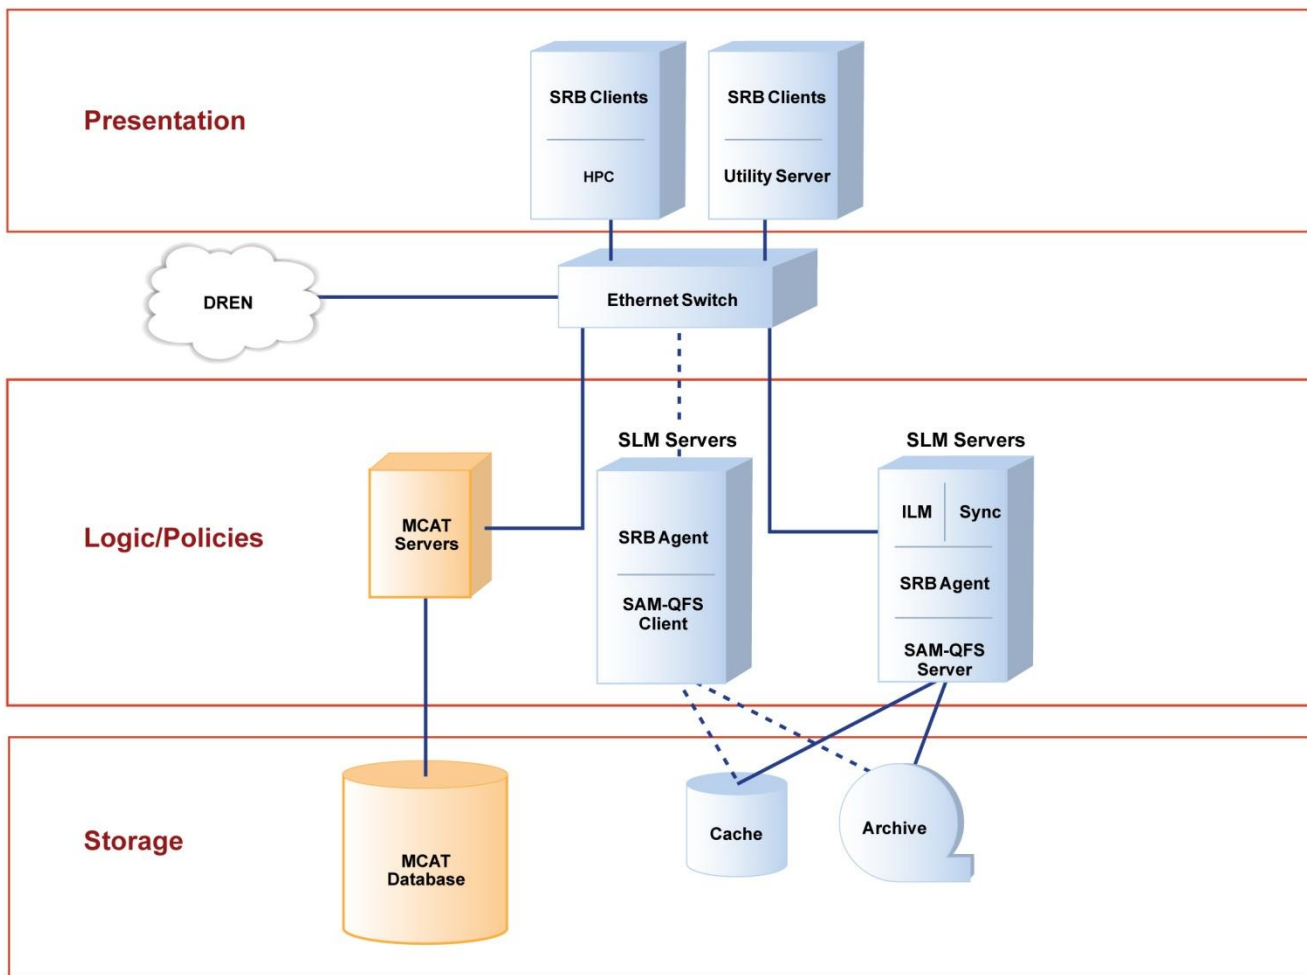
Global Namespace



Metadata Catalog



DoD HPC Component Interaction



Planned Approach: DoD HPCMO SLM Project

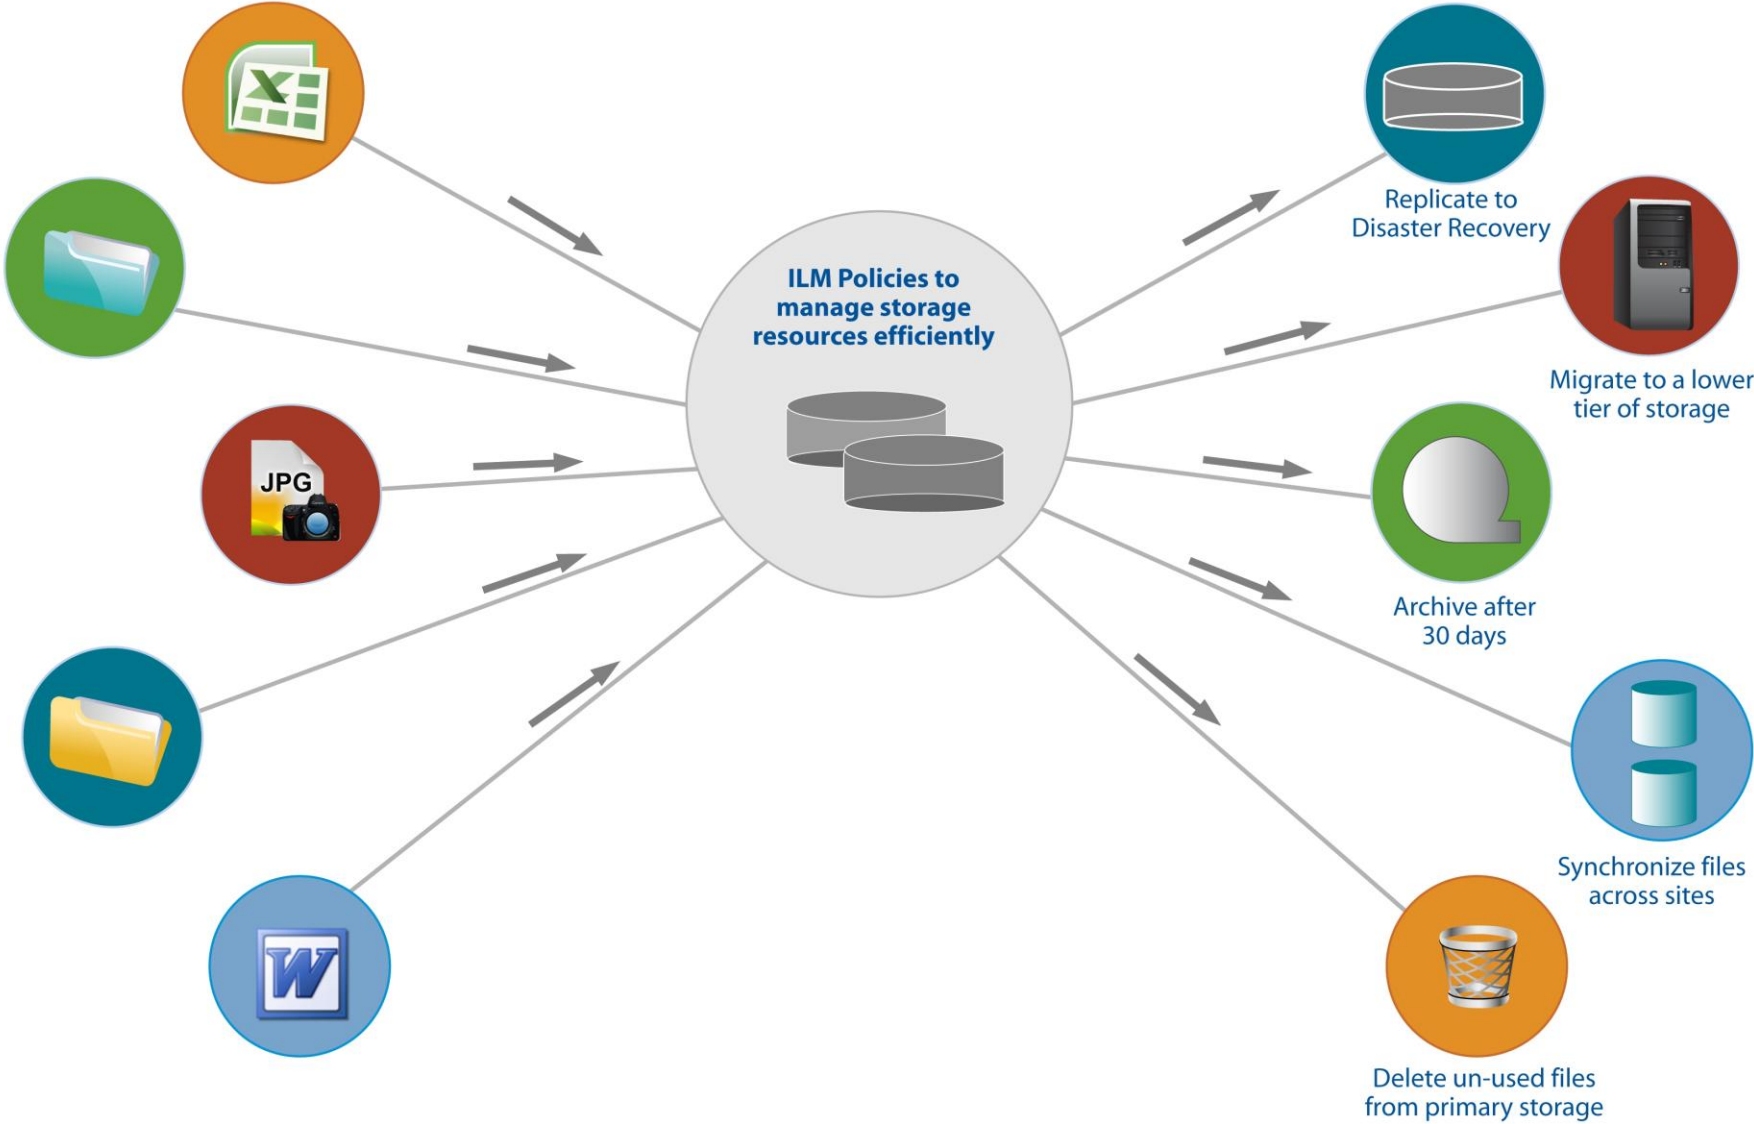
Project Plan: August '09-August '10

- Nirvana SRB + Sun Hardware & Software
- Design meetings and Administrator training with key stakeholders and Center Reps
- Stand up and document test-beds: ARL & ARSC
- Ensure security compliance DoD C&A requirements
- Test & Acceptance
- Implementation

Nirvana SRB manages ILM metadata attributes and orchestrates SAM-QFS across file systems & sites.

- **Minimized operational risk to the Program**
 - Built on existing architecture
 - Existing files remain on same disk and tape media
 - Familiar SAM-QFS structure remains intact
- **COTS-based solution with ongoing development can be leveraged for other government and commercial projects**
- **Integrates with existing security infrastructure**
- **Ease of transition to ILM system**
 - Business as usual plus enhancements to data management within one year

Nirvana SRB Information Lifecycle Management



Questions

Constantin Scheder

858-455-2536, scheder@ga.com

ga.com/nirvana

Backup Slides

End User Data Views with Nirvana SRB

Selection of versatile SRB

interfaces includes:

- Windows virtual disk
- Samba/CIFS shares for Windows
- NFS mounts for UNIX, Linux
- WebDAV Gateway for Web Folders
- Preload Library (virtual mounts)
- Web Services: SOAP, REST
- Java Client & Java Admin
- Web (IE, Firefox, Google Chrome, Safari, iPhone, Droid, Blackberry)
- Scommands, Acommands (CLI)
- SDK (Microsoft .NET, C & Java)
- Automation Tools: ILM, Sync, Network Stream & Metadata Daemons



For Administrators – Simple Data Management

The screenshot displays the SRB Java Admin web interface. The window title is "SRB Java Admin (super@root)". The menu bar includes "File", "Edit", "View", "Actions", and "Help".

The left sidebar shows a tree view of the system structure:

- Federation
 - Users / Groups / Domains
 - root
 - public
 - nirvanaware.com
 - super
 - Groups
 - Locations / Resources
 - Locations / Physical Resources
 - root
 - public
 - nirvanaware.com
 - engineering
 - testing
 - marketing
 - docctr
 - p193182.5625
 - Logical Resources
 - Cluster Resources
 - Daemons / Executables
 - Daemons
 - Silm
 - Ssync
 - Smeta
 - Sstream
 - Schemes
 - Tokens

The main content area is titled "Modify Location 'p193182.5625'". It features several tabs: "Information *", "Authentication *", "Configuration", "Access", "Alias", and "Filter". The "Information *" tab is active, showing the following fields:

- Location Name *: p193182.5625
- Host Name *: p193182.sd.gat.com:5625
- Contact Information: [Empty field]
- Phone Number: [Empty field]
- Email: [Empty field]
- Parent Location *: Locations / Physical Resources
 - root
 - public
 - nirvanaware.com
 - engineering
 - testing
 - marketing
 - docctr
 - p193182.5625

The status bar at the bottom indicates: "Status: Your home collection is "/home/super.root"."

SRB 2010 Supported Systems

- **Operating Systems (32 & 64bit)**
 - AIX, HP-UX, Linux, Mac OS-X, SGI Irix, Sun Solaris, Windows
- **File Systems**
 - All file systems mountable on all supported operating systems including cluster file systems.
- **Archives**
 - [SAM-QFS](#), EMC Centera, Honeycomb/STK 5800, HPSS, DiskXtender 2000, DiskXtender UX (formerly Unitree), AMASS, SCSI & FC Tape Libraries
- **Relational & Large Object Databases**
 - DB2, MS SQL, Oracle, Sybase, Postgres, [MySQL](#), [Amazon RDS \(cloud\)](#)
- **Gateways**
 - [Virtual Disk](#), CIFS, NFS, GridFTP, WebDAV, Preload Library
- **Additional Drivers**
 - HTTP, [HTTPS](#), [Amazon S3 \(cloud\)](#), FTP, [FTPS](#), [SFTP](#), [SCP](#)
 - Documentum (Content management)
 - MetaCarta (Full-text geospatial indexing)



Panel Questions & Answers

- How does size of collections (# files/ amount of data) affect the infrastructure design?
 - # files requires separate distributed & clustered metadata catalogs
 - Size requires distributed tiered storage infrastructure
 - Only cost-effective solution is tape
 - Users demand and benefit from metadata and cross-center Global Namespace
- Are data sets inherently distributed across multiple repositories?
 - Distributed
- Does large scale imply a homogeneous collection (uniform data format and uniform descriptive semantics)
 - No, 4500 users, hundreds of different projects; however uniform metadata schemes
- Does large scale imply use of collections to manage metadata?
 - yes, nature of Nirvana SRB → collection-based
- What properties are enforced by the data management system?
 - Integrity, authenticity, organization, access control (DAC, MAC), quotas, retention, resiliency, reliability, usability, supportability
- Can the properties of the data management system be validated periodically?
 - Yes, self-consistent database; verification/ monitoring tools;
- How do you repair problems that are found?
 - Probably SRB, very small chance that this could be sensitive
- How do you minimize the amount of labor needed to maintain the data?
 - Automation, virtualization, training, support, monitoring and debugging tools
- What services are provided for manipulation and analysis of the data?
 - Probably ties into our utility server concept. We will mean something different than what is used for large earth sciences or biomedical repositories