



<http://www.grid-support.ac.uk>



<http://www.ngs.ac.uk>

Introduction to The Storage Resource Broker



JISC





Policy for re-use



- This presentation can be re-used for academic purposes.
- However if you do so then please let training-support@nesc.ac.uk know. We need to gather statistics of re-use: no. of events, number of people trained. Thank you!!



Acknowledgements



- This tutorial selects slides from several sources, specifically from talks given by Wayne Schroeder (SDSC) and Peter Berrisford (RAL)



Goal



- Introduce use of the SRB for distributed file management on the NGS
 - This is the focus of the practical that follows
- To explore further:
 - <http://www.sdsc.edu/srb/> and http://datacentral.sdsc.edu/user_guide.html
 - For a full SRB tutorial, at NIEES last January, see: <http://www.niees.ac.uk/events/srb2006>



What is SRB?



- Storage Resource Broker (SRB) is a software product developed by the San Diego Supercomputing Centre (SDSC).
- Allows users to access files and database objects across a distributed environment.
- Actual physical location and way the data is stored is abstracted from the user
- Allows the user to add user defined metadata describing the scientific content of the information



What is SRB?



The SRB is an integrated solution which includes:

- a logical namespace,
- interfaces to a wide variety of storage systems,
- high performance data movement (including parallel I/O),
- fault-tolerance and fail-over,
- WAN-aware performance enhancements (bulk operations),
- storage-system-aware performance enhancements ('containers' to aggregate files),
- metadata ingestion and queries (a MetaData Catalog (MCAT)),
- user accounts, groups, access control, audit trails, GUI administration tool
- data management features, replication
- user tools (including a Windows GUI tool (inQ), a set of SRB Unix commands, and Web (mySRB)), and APIs (including C, C++, Java, and Python).

SRB Scales Well (many millions of files, terabytes)

Supports Multiple Administrative Domains / MCATs (srbZones)

And includes SDSC Matrix: SRB-based data grid workflow management system to create, access and manage workflow process pipelines.



SRB Scalability



Storage Resource Broker (SRB)

Data brokered by SDSC instances of SRB**

Project Instance	As of 7/24/2003			As of 9/12/2003			As of 10/01/2003			As of 11/14/2003		
	Data_size (in GB)	Count (files)	Users	Data_size (in GB)	Count (files)	Users	Data_size (in GB)	Count (files)	Users	Data_size (in GB)	Count (files)	Users
NPACI	6,050.00	2,317,368	367	8,350.00	2,903,386	372	8,570.00	2,946,526	375	8,822.00	2,995,432	377
Digsky	46,100.00	5,719,025	68	46,100.00	5,719,025	68	46,100.00	5,719,025	68	42,786.00	6,076,982	69
DigEmbryo	720.00	45,365	23	720.00	45,365	23	720.00	45,365	23	720.00	45,365	23
HyperLter	215.00	5,097	27	215.00	5,097	28	215.00	5,097	28	215.00	5,097	28
Hayden	7,078.00	59,399	142	7,130.00	59,781	158	7,830.00	59,983	158	7,835.00	60,001	168
Portal	968.00	27,250	316	1,133.00	31,717	339	1,141.00	32,396	342	1,244.00	34,094	352
SLAC	1,790.00	254,974	43	1,790.00	254,974	43	1,790.00	254,974	43	2,108.00	294,149	43
NARA/Collection	52.80	79,195	51	53.10	79,112	55	53.90	79,118	55	67.00	82,031	56
NSDL/SIO Exp	232.00	15,809	23	315.00	27,170	23	418.00	39,253	23	603.00	87,191	26
TRA	90.60	2,385	25	92.00	2,378	26	92.00	2,387	26	92.00	2,387	26
LDAS/SALK	498.00	9,858	60	737.00	12,898	66	767.00	12,906	66	824.00	13,016	66
BIRN	121.00	237,283	138	273.00	675,531	145	382.00	2,048,132	156	389.00	1,084,749	167
AfCS	95.30	18,762	20	99.00	19,714	20	102.00	20,260	20	107.00	21,295	21
UCSDLib	1,084.00	138,415	29	1,085.00	138,421	29	1,085.00	138,421	29	1,085.00	138,421	29
NSDL/CI	278.00	993,886	113	379.00	2,596,090	114	379.00	2,596,090	114	465.00	2,948,903	114
SCEC	12.60	18,660	38	7,561.00	1,249,144	39	9,680.00	1,561,396	40	12,274.00	1,721,241	43
TeraGrid	623.00	36,508	1,978	1,664.00	47,644	1,942	1,745.00	49,106	2,073	10,603.00	433,938	2,229
TOTAL	66,008.30	9,979,239	3,461	77,696.10	13,867,447	3,490	81,069.90	15,610,435	3,639	90,239.00	16,044,292	3,837
	66 TB	9.97 million	3 thousand	77 TB	13.9 million	3 thousand	81 TB	15.6 million	thousand	90 TB	16 million	3 thousand

** Does not cover data brokered by SRB spaces administered outside SDSC.

Does not cover databases; covers only files stored in file systems and archival storage systems

Does not

cover shadow-





Storage Resource Broker




User sees a virtual filesystem:

- Command line (S-Commands)
- MS Windows (InQ) →
- Web based (MySRB).
- Java (JARGON)
- Web Services (MATRIX)



Filesystems in different administrative domains

How SRB Works



- 4 major components:
 - The Metadata Catalogue (MCAT)
 - The MCAT-Enabled SRB Server
 - The SRB Storage Server
 - The SRB Client



SRB on the NGS



- SRB provides NGS users with
 - a virtual filesystem
 - Accessible from all core nodes and from the “UI” / desktop
 - (will provide) redundancy – mirrored catalogue server
 - Replica files
 - Support for application metadata associated with files
 - fuller metadata support from the “R-commands”



Practical Overview



- Use of the Scommands
 - Commands for unix based access to srb
 - Strong analogy to unix file commands
- Accessing files from multiple (two) sites