

# SRB in a Production Context

Peter Berrisford, Gordon Brown, Ananta Manandhar, Bonny Strong, Nick White

(CCLRC e-Science Centre)

Roman Olschanowsky (BIRN Project)

## Abstract

The Storage Resource Broker (SRB), a distributed data management product developed by the San Diego Supercomputer Center (SDSC), is now in use throughout the world on a wide range of research projects. This year has seen a rapid increase in the number of UK e-Science projects with SRB as a core grid-enabling element within the I.T. strategy. This includes use for data archiving, particularly as a front-end to mass tape storage, as well a means of providing semantic structure to the data and enabling easy data retrieval. Many of these projects require a production-level service in terms of system availability and reliability, data throughput capabilities, backup facilities and disaster recovery. This paper provides an overview of typical system configurations and describes the special challenges faced when moving SRB to a production footing. Issues discussed include the realities of supporting a production system, including professional management of the software layers upon which SRB depends, and highlighting the need for good quality procedures and planning.

## 1. A Production SRB System

This paper examines the issues to be addressed when hosting production-standard SRB systems for multiple customers. Rather than creating and supporting a complete SRB infrastructure, many projects utilize the services of CCLRC to host key components. These typically include the MCAT (Metadata Catalogue) database and the associated “front-end” MCAT Enabled Server, plus the “SRB-ADS” server where the Atlas Data Store mass tape storage facilities are required.

Setting up and operating a production SRB system to achieve a viable data grid is a non-trivial task. A production system should ideally be robust, reliable, available and secure, have acceptable performance and meet a customer’s functional requirements. Where problems are experienced, service disruption should be kept to a minimum, with good levels of communication.

One way to understand a system and its context is to examine the roles that need to be fulfilled in order to achieve successful implementation and operation. A proposed set of roles is outlined in the following section. The necessary activity and infrastructure requirements are then considered, including:

- System Infrastructure
- Support and Operational Infrastructure
- Initial set-up and Training

This paper focuses on the implementation aspects of a production SRB system, and does not examine different applications of SRB in detail. However, real life examples are used, taken from experiences with projects such as the CMS Data Challenge as well as mature production systems such as BIRN (Biomedical Informatics Research Network). For further examples, see [1] and [2].

## 2. Roles

Several distinct sets of tasks, requiring different skill sets, map to a surprisingly large number of roles. Note that one person may in fact perform more than one role, and one role can be performed by more than one person. The former helps to make the operation of the SRB service more cost-effective, while the latter allows the necessary cover to be maintained when a primary ‘contact’ is unavailable.

### 2.1 System Architect

The System Architect is responsible for design of the SRB system architecture, both the “core” components (discussed later) and the new components required for each new customer. Re-design work may be required at intervals depending upon changes in usage profiles.

### 2.2 Data Grid Engineer

The Data Grid Engineer is the key SRB system implementer. This role is responsible for implementing all SRB-related aspects of the

Architect's design, covering all SRB Administrator tasks, plus:

- Installation and configuration of the MCAT database. Once the DBA has created the database (used across multiple projects) and the schema (one for each project), the Data Grid Engineer (DGE) is responsible for creating, populating and testing the tables, sequences, synonyms and indexes.
- Installation and configuration of the MCAT servers and the SRB-ADS servers (one instance of each server per project).
- Building and maintenance of the monitoring and reporting infrastructure;
- Liaising with the DBA's and System Administrators. This relates to performance tuning, monitoring and reporting in addition to system set-up and problem resolution.
- Problem investigation – liaising with SDSC in the case of bugs.
- System upgrades.
- Communication with the User community in the case of system down time (planned or otherwise), unless handled by the Grid Support Officer.
- Definition and enforcement of naming conventions. This includes all aspects of the grid – user names, domains, resources and DNS entries.
- Documentation of standard procedures and policies. A good example of well-established procedures can be found in [4].

### 2.3 Data Grid Consultant

The Consultant is responsible for enabling the successful addition of a new customer to the service and for helping to ensure that they can make effective use of the service on an ongoing basis. This may be regarded as both an educational and a facilitating role, involving:

- The provision of the necessary training or ensuring that access to the necessary training or documentation is available;
- Consultancy sessions allowing a customer's needs to be clearly identified and the available options to be assessed and a suitable strategy agreed, i.e. how the product should be used.
- As well as assisting in terms of application design, the Consultant helps in terms of Customer infrastructure design – covering issues such as fault tolerance and backup strategies.

### 2.4 Data Grid Application Programmer

The programmer may be regarded as an advanced user; however it would probably be more accurate to view this role as one that extends the core system functionality to enhance the manner in which the 'real' users may perform their work. This includes the writing of shell, Perl or Python scripts or coding against the API's provided, for example coding in C or using the Jargon (Java) interface.

### 2.5 SRB Administrator

The SRB Administration role represents a subset of the Grid Engineer role.

The majority of SRB service customers will wish to have their own SRB Administrator to set up SRB Storage Servers and perform administrative tasks such as user and resource creation and maintenance. Appropriate training is essential if Data Grid Engineer intervention to resolve problems is to be avoided. Once proper training has been received, the DGE can assign the Sys Admin role to the relevant SRB user.

The SRB Administrator role covers:

- Installation and configuration of SRB Storage Servers;
- Assisting with the installation and configuration of SRB Client tools;
- SRB User and Resource maintenance;
- Maintenance and care of the SRB vaults and cache spaces;
- Scheduling of moves or copies to archive. Such tasks may be automated to some degree;
- Determining who is allowed to use which resource;
- Determining how to provide long and short-term cache space;
- Ensuring that suitable backup and fault tolerance strategies are in place;
- Retrieving data back out of archive to cache, in order to optimise performance.
- Recovery activities such as resolving "dangling files";
- Disaster Recovery activities such as performing restores. DGE involvement may be required.

Potentially, an SRB Administrator may also perform a System Administration role, with responsibility for tasks such as system backup (for local SRB Vaults), or may be biased more towards the user or Application Programmer roles.

## 2.6 SRB User

The SRB User is the reason the SRB service is provided. Each Customer user base typically makes use of a number of SRB Client tools (i.e. S-Commands, inQ and / or mySRB), plus interfaces developed by Data Grid Application Programmers.

## 2.7 Grid Support Officer

Where an SRB Service customer has I.T. expertise available, and assuming that the necessary training and documentation are available, front-line support may well be handled by the customer directly. This would handle day to day issues such as users mistyping S-Commands.

The next line of support will typically be via a help desk or a special support email address. The Grid Support Officer (GSO) is responsible for:

- Analysing requests
- Responding directly where the issue can be resolved
- Redirecting other requests to a Data Grid Engineer or Consultant or directly to a System Administrator or DBA where this is obviously required. Communication issues are discussed later.
- Ensuring that cases are dealt with.

Ideally the Grid Support Officer will be able to perform tasks such as reviewing the srbLog file. More training and experience equates to a higher proportion of calls that may be handled directly.

This operational support role can be proactive as well as reactive. Assuming a suitable system of monitoring is in place, the GSO can review the health of the system at set intervals in order to determine whether loss or degradation of service has occurred. This enables corrective action to be initiated as early as possible. Potentially, notification of service degradation or loss can be automated.

## 2.8 DBA

Given that SRB relies on a DBMS for the MCAT database, and given that Oracle is usually used for production systems, a properly qualified Database Administrator is essential. The DBA is responsible for:

- MCAT database creation and maintenance.
- Schema creation, one per project.
- Database instance configuration.
- Database back-up.
- Replication if a failover system is in use.

- Installation and maintenance of Oracle Management Server, plus associated Oracle Enterprise Manager for monitoring and automatic failure notification.
- Installation of Oracle Client software on separate MCAT server machines
- Installation and maintenance of Oracle Real Application Clusters (RAC) and Oracle Cluster File System if in use.
- Reconfiguration and tuning to provide optimal performance.
- Provision of SQL scripts to generate statistics relating to SRB usage – this is carried out in conjunction with the Data Grid Engineer (DGE). In some production environments, the DGE is entirely responsible for all aspects of generating statistics and will actually write all SQL.
- Provision of support to DGE or Consultant and / or Grid Support Officer in the case of SRB system problems that are / may be database related.

## 2.9 System Administrator

While not normally directly involved with the SRB software, without good quality System Administration support the SRB service is likely to suffer. The System Administrator is responsible for the hardware, the operating system and the network. It is necessary for this role to work closely with the DBA and the Grid Engineer.

The roles described above are primarily technical. From a non-technical perspective, the management and reporting aspects need to be considered, with external 'stakeholders' such as funding bodies forming an important part of the system context. The importance of ensuring adequate funding and resource levels is addressed later in the paper.

## 3. System Infrastructure

In order to deliver an effective SRB service to all customers, it is necessary to design a system architecture that takes into account current and future needs. As indicated in the previous section, this is the job of the System Architect. Two key views architecturally are:

- The software stack for the MCAT database and MCAT Enabled Server machines;
- An overview of the entire SRB system infrastructure.

### 3.1 The Software Stack

Figures 1 and 2 illustrate the software stack chosen for the CCLRC SRB service.

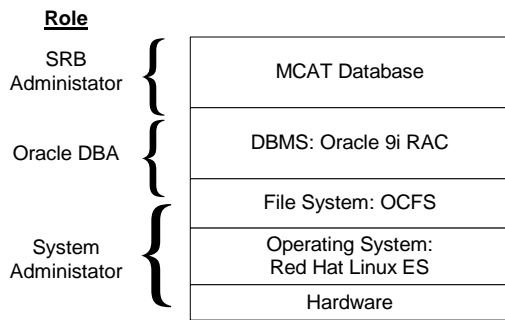


Figure 1: MCAT Database Server

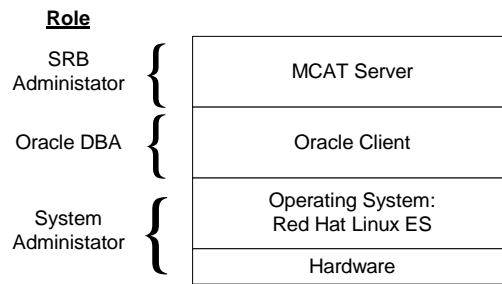


Figure 2: MCAT Enabled Server (MES)

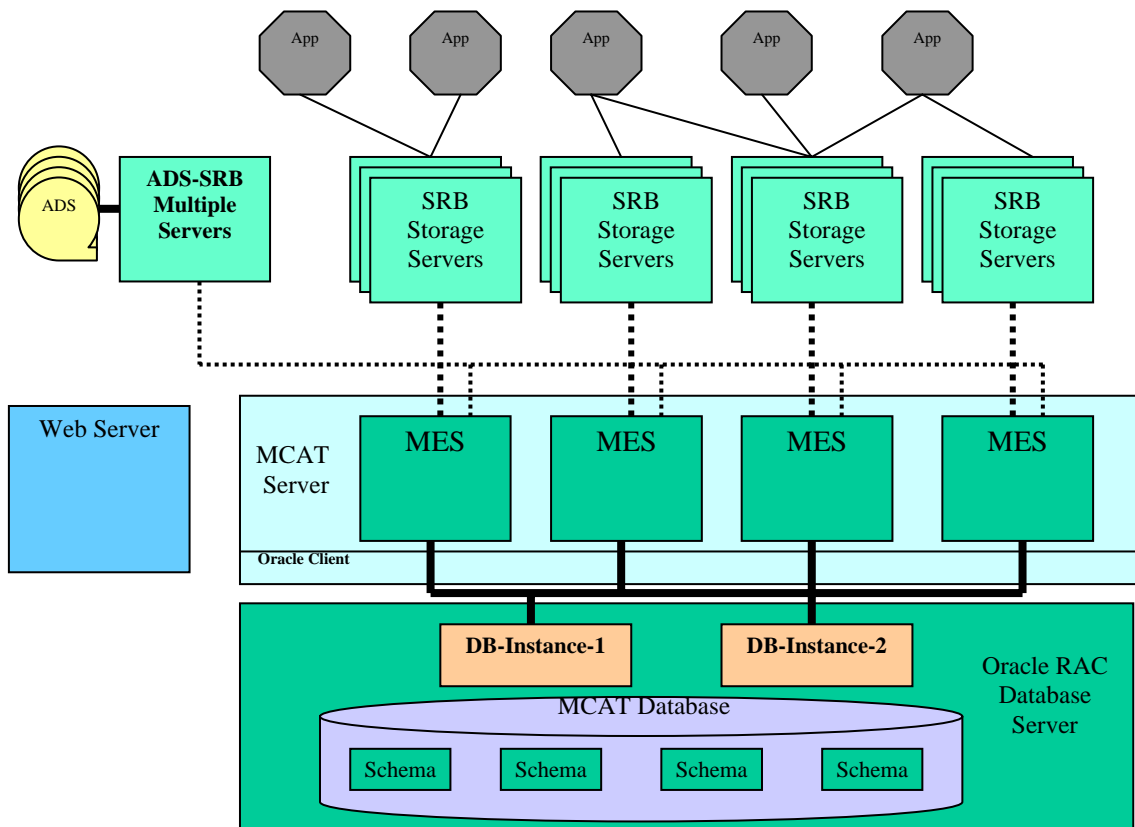


Figure 3: CCLRC SRB Service System Infrastructure

Figure 3 shows the overall system. This configuration is fairly unique and has both advantages and disadvantages. Both Solaris and Linux are fully supported as an SRB platform, with other major projects such as BIRN also using Red Hat Linux. Oracle 9i can also be regarded as the standard database platform for production SRB services.

However, few people are using Oracle Real Application Clusters (RAC) currently. One key advantage of clustering the database instances is the enhanced resilience of the system in addition to potential performance benefits. Load should be fairly well balanced between the

nodes in normal circumstances, with the system still able to function if one node dies.

One potential drawback is the relative immaturity of the software. The Red Hat / Oracle RAC combination, particularly with the use of the Oracle Cluster File System (OCFS), has proved to be relatively unreliable at times, as discussed later in relation to the CMS Data Challenge project. The BIRN project is planning to move from a single Oracle database instance to RAC in the coming months, however raw disk is to be used in preference to OCFS. The advantage of OCFS is that files can be viewed externally to the database server.

The decision has also been taken at CCLRC to run the MCAT Enabled Server (MES) instances on a separate machine from the database servers, such that each machine has a dedicated role and can be tuned appropriately. Many production systems run the MES instances on the same machine as the MCAT database.

The CCLRC approach requires a fast link between the database and MCAT Server machines. A brief period running the MES machine and MCAT database at geographically separate locations clearly demonstrated a significant performance penalty. This has clear implications for any failover system – the MES and MCAT database machines both need to failover together.

### 3.2 The ADS Interface

Most projects have a requirement for access to the Atlas Tape Store. In the same way that multiple instances of the MES are run on a single machine, each project has its own instance of the SRB-ADS server on an ADS front-end machine. All SRB Storage servers and the MES and SRB-ADS server for a particular project share the same port number.

A large number of small files have proved to be a common feature across the majority of projects. While this has performance implications for any Data Grid system, it causes particular problems for mass tape storage systems. The use of SRB containers is a key strategy in avoiding many of the problems experienced during the CMS Data Challenge. The SRB-ADS server instances will have an SRB disk vault which will act as a disk cache front-end to the tape store. Once a container has been filled, the system is configured to rename the completed container and write it to the tape store as a background task. A fresh container with the original name is immediately available.

### 3.3 The Customer Infrastructure

Most customers will set up and maintain their own SRB Storage Servers with local disk vaults, plus all client tools. Each server will share the port number associated with that customer's MCAT Enabled Server instance. The number of SRB servers will depend upon customer needs, partly based upon the number of geographically separate sites. Fault tolerance and disaster recovery for those machines / SRB vaults will normally be the responsibility of the customer.

Typically, data that can be regarded as still in active use will be retained on the customer's disk vaults, with transfer to the ADS once archival status is reached. The degree of

replication will again be based upon customer needs. For instance where a minimum of two copies are always required, a logical resource may be configured with two physical resources, each automatically receiving a replica of the data (see figure 4).

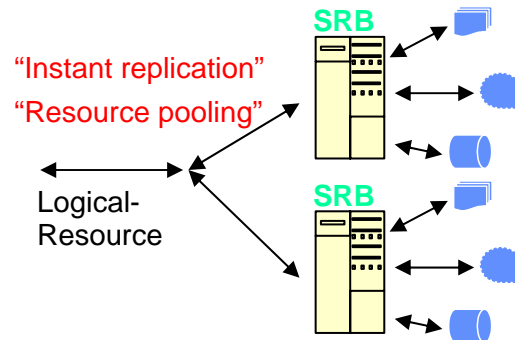


Figure 4: Use of Logical Resources

### 3.4 A Failover System

Although "federated SRB" (version 3.x) is to be used for the CCLRC service, there will initially be one Oracle database for the MCAT (with each customer having their own schema) and one "application server" running multiple instances of the MES. These both potentially represent single points of failure.

Oracle DataGuard will therefore be used to replicate the MCAT database to a geographically distant location (at Daresbury) to enable a failover system. However, as already indicated, the MCAT Enabled Servers also need to failover to the same location. The BIRN system provides one example of how to achieve this. When the Oracle system identifies that there is a problem, an 'indicator' on a website is set reflecting that a new MES machine is to be used. Each SRB server polls this indicator at five minute intervals – if it has been set then the SRB server is reconfigured to use the new MCAT Enabled Server.

### 3.5 Non-Production Environments

The lack of a pre-production environment during the CMS Data Challenge and following periods of SRB testing made life particularly difficult. The plan for the CCLRC service is to have a test / development system and a pre-production system in addition to the production environment. It is not important to have a matching system configuration or similar machine specifications for a test / development system, however the pre-production system, while it does not have to be full production spec, should have a similar configuration. Without this, final system or integration testing,

or testing to determine the precise cause of unusual behaviour in the production system, need to be carried out in production.

Before new software versions are moved to production, regression testing should ideally be performed. This does however require the ongoing maintenance of test scripts, a task that, along with documentation, tends to be bypassed. For the BIRN project, testing of SRB version 3 has been ongoing for a number of months within a test environment prior to a move to production.

### 3.6 Software Version Management

All machines connected to the BIRN system are required to run software that verifies software versions, from operating system to SRB version. CVS (Concurrent Versions System) is in use for version management. From an overall management perspective, it is easier to have all machines running the same version of SRB; however, it is possible to run older versions on some SRB servers as long as the interface presented by the utilized APIs has not changed.

## 4. Support Infrastructure

This category covers a multitude of areas that need to be addressed to achieve a production-level service. It should be noted that the creation and enhancement of a production infrastructure is an ongoing activity, requiring a number of years to attain maturity. A good example of this is the BIRN project, where through a process of evolution an extremely effective support, monitoring and reporting infrastructure has been created (see [3]).

### 4.1 Support Areas

It is important for any production service to determine the scope of support, to clearly identify what can and cannot be done based upon available levels of resources / funding.

Having said this, there are certain areas that clearly need to be supported:

- The hardware, operating systems and network across the system infrastructure – division of responsibility between the Customer and the hosting service needs to be clearly defined. Firewall issues potentially fall within this category as well.
- Oracle.
- SRB system: including performance issues and service availability (local or ‘global’).
- SRB: User queries – straightforward queries should be handled by front-line

support or pre-empted through appropriate training

- SRB: changes to service (resource updates, etc.) – this may well be handled by the customer’s SRB Administrator
- Extended application functionality – such as Data Portals.

### 4.2 Requesting Support

This issue has already been addressed from the support roles perspective in section 2. A mechanism for registering support calls is required, typically either via email or a web-based helpdesk facility. Where email is used, a single support email address is recommended as against, for instance, separate email addresses for SRB database or system admin related queries. This is because it is not always clear as to where the problem lies and centralised control can be maintained.

Standard helpdesk software should be considered across a range of services – this helps to reduce support costs.

Hours of support and response times (for looking at the problem) need to be agreed. Effective cover in terms of availability of relevant staff needs to be planned and controlled. Effective communication and coordination of effort between teams is required.

### 4.3 Training and Documentation

Training and documentation should ideally be available for all SRB-related roles and use cases (e.g. digital library, data grid, persistent archive, and data flow pipelines). Good quality, appropriate training and documentation can lower support costs and may allow additional tasks to be handed over to the customer, e.g. front-line support.

### 4.4 Monitoring and Reporting

The implementation of monitoring relates to system infrastructure and requires the involvement of the System Administrator, DBA and Data Grid Engineer. The benefits of the monitoring relate to support and management of the SRB system.

From a system administration perspective, monitoring addresses the following types of issues:

- Is the system alive?
- Has there been a system crash or is there a network problem?

- Notification of unexpected increases in load – CPU, memory usage, size of log files, number of open sockets.
- Is the system running out of disk space?

The DBA can implement monitoring using tools such as Oracle Management Server / Oracle Enterprise Manager, plus SQL scripts. This provides notification in case of problems plus provides key information for reporting purposes.

The Data Grid Engineer is responsible for the SRB-level monitoring and reporting. Areas to be considered:

- Is the system running OK? This includes performance of requests, inconsistency of performance, hanging requests, build up of resources that should have died.
- Where a certain level of use (e.g. in terms of connections between the MCAT server and MCAT database, number of concurrent commands, data transfer rates) could indicate that the system is unusually heavily loaded, and therefore runs a higher risk of some type of problem occurring, automatic notification or at least some form of logging should be considered.
- Use of the SRB Spcommand proxy S-Command enables useful system status information to be obtained (the BIRN website [3] provides an excellent example of this).

Strong consideration should be given to the setting up of a website that clearly shows current system status, and which provides detailed reporting on all aspects of the SRB service.

#### 4.5 Operational Issues

This section concludes with a brief synopsis of issues to be considered from a service operation perspective.

Key activities and concerns in this area include:

- System upgrades and maintenance (including the application of patches). The use of RPM's is recommended.
- Pre-production testing of new software versions.
- Level of resource availability.
- Cover for training, holidays, sickness, or key personnel leaving.
- Formal procedures for backups, change control, event notification and disaster recovery.

In the case of disaster recovery, a comprehensive disaster recovery plan should be prepared. Customers need to be aware of

boundaries of responsibilities for backups and recovery. Failover systems across different sites, as discussed earlier, appropriate use of replication and the use of RAID should be considered.

### 5. CMS Data Challenge: The Lessons

This paper has discussed many of the key issues that need to be addressed in order to provide a production-standard SRB service, and has made many recommendations. This does not mean that recent services have always run smoothly.

While there have been significant success stories such as the E-Minerals project, the CMS Data Challenge cannot be regarded as an outstanding success in the context of the provision of a production service; however it has provided invaluable experience and will help to ensure that forthcoming services can meet the needs of the customer.

Key lessons learned:

- Maturity of software – there is a significant risk associated with adopting immature software, particularly if it makes up the majority of the software stack. If the actual combination of software is unique, you have the honour of effectively testing the software.
- Architectural decisions – as with the choice of software, if the architectural configuration is unique, or is rarely used, then the risk of problems is increased.
- Test / pre-production environments – the lack of a realistic environment on which to perform testing means testing occurs in a production environment.
- Sufficient control over the environment is essential, as is the ability to access key machines. The issue of machine ownership needs to be carefully considered.
- Many small files can cause significant performance problems. The use of containers can significantly improve matters.
- Effective SRB application design is essential to achieve good performance and prevent system-breaking loads. Coding against the API's should be considered, as should use of bulk operations. The forthcoming SRB 'shell', avoiding the need for separate connections for each S-Command should help matters.
- Commitment of necessary resources – people with the necessary skill sets need to be officially available, with a realistic

proportion of their time allocated to the relevant roles.

- Importance of all layers of the software stack – if any one part of the software stack causes problems, the knock-on impact can be serious. Any grid software is likely to rely heavily on the underlying infrastructure, and SRB is no exception.
- Appropriate support contracts are necessary to try and achieve swift resolution of problems, although significant work may be involved in order to allow that resolution to be achieved.

With the rapid growth of SRB usage for UK e-Science projects, the UK now has the opportunity to help define best practices for the provision of production grid services.

## 6. References

[1] “Real Experiences with Data Grids – Case Studies in using the SRB”,

(<http://www.npaci.edu/DICE/Pubs/hpcasia2002.pdf>)

[2] “Storage Resource Broker – Managing Distributed Data in a Grid”,

*Computer Society of India Journal*,

(<http://www.npaci.edu/DICE/Pubs/CSI-paper-sent.doc>)

[3] “BIRN – For System Managers”,

(<http://www.nbirn.net/Resources/SystemMgrs/index.htm>)

[4] “BIRN Governance – Participation, Policies & Procedures”