

# HPSS 5.1 Overview

HPSS User Forum

June 5, 2001

David Fisher

## The focus of HPSS 5.1 is infrastructure replacement

- 5.1 is based on 4.3 with SFS replaced by a commercial relational database. All HPSS metadata will be managed by the RDB.
- In 4.3, SFS is the metadata transaction coordinator
- In 5.1, all metadata transactions will be coordinated by the RDB

## How will this be done?

- Merge core servers
- Switch to common infrastructure
- Create DB Schema
- Replace MMLIB
- Replace transaction logic

## Step 1 - Merge core servers to eliminate distributed transactions

- Transactions that modify metadata are distributed among NS, BFS, TSS and DSS.
- But distributed transactions require an external transaction monitor, Encina.
- To eliminate distributed transactions, NS, BFS, TSS, DSS and parts of the Client API will be merged into a single process, the HPSS Kernel.

## Step 2 - Use common infrastructure in the Kernel

- Some of these common subsystems are:
  - ◆ Logging
  - ◆ Real-Time Monitoring
  - ◆ Signal handling
  - ◆ SSM driver

## Step 3 - Create DB schema based on SFS schema

- Mostly a process of mapping HPSS metadata structures to DB tables. Both SFS and DB tables are flat, so conversion is easy for most tables.
- Name Server MD will be a new design

## Step 4 - Replace SFS MMLIB with DB MMLIB

- Old MMLIB is very SFS specific.
- New MMLIB will perform DB functions similar to old MMLIB, but will be generic, not specific to a vendor's DB.
- Will include transaction management tools.

## Step 5 - Replace transaction logic

- Tran-C and TX transaction interface will go.
- Replaced with simple HPSS transaction management functions:
  - ◆ Start, commit, rollback, callback

## Technical considerations

- Architecture changes
- Server config changes
- Portable implementation
- Performance
- Conversion
- Administration

## Other than new Kernel, system architecture unchanged

- MPS, PVL, Movers, SSM, LS, GK, Loggers and utilities remain separate processes, but will read/write DB, not SFS.
- Storage device and media configuration unchanged
  - ◆ COS, Hierarchies, Storage Classes
- Subsystem concept remains
  - ◆ Subsystem is Kernel and MPS
  - ◆ Each subsystem will have its own set of DB tables

## Server configuration changes

- No configuration files for NS, BFS, SSs, all moved to Kernel.
- No SFS files in server configs.
- DB tables will have fixed names.
- Continue to use DCE for RPCs and security.

## Server configuration changes

- Kernel will display a single managed object containing:
  - ◆ Summary of NS, BFS, SS config stuff
  - ◆ Major statistics
- Kernel specific config will have no statistics, just startup parameters.

## Portable implementation

- MMLIB will provide a generic interface for DB and transaction operations.
- The ability to host HPSS on more than one vendor's RDB is a design goal.
- But, our plans at this time are to release 5.1 on a single vendor's RDB.

## Performance

- Benchmarked DBs run 8X faster than SFS
- Kernel faster – fewer RPCs, more shared data.

## Conversion

- Convert 10M File HPSS in 24 hours or less
- Some experimentation to be done
- Checkpoint/restart
- Metadata checker

## Administration

- HPSS admins will learn DBA techniques
- Wealth of knowledge and help available
  - ◆ Books, magazines, newsgroups, mailing lists, colleagues
- Most HPSS admin knowledge and skills transfer to 5.1
- Ad-hoc queries become easy

## Schedule

- DB Benchmark complete: April 01
- Design documents: June 01
- Kernel infrastructure: Aug 01
- Rehost over DB MMLIB: Feb 02
- Integrate SSM: Feb 02
- Convert other servers and utilities: Dec 01
- Build conversion tools: Mar 02
- System test, documentation & packaging: Dec 02