

HPSS 4.5 to 5.1 Metadata Conversion

HPSS User Forum

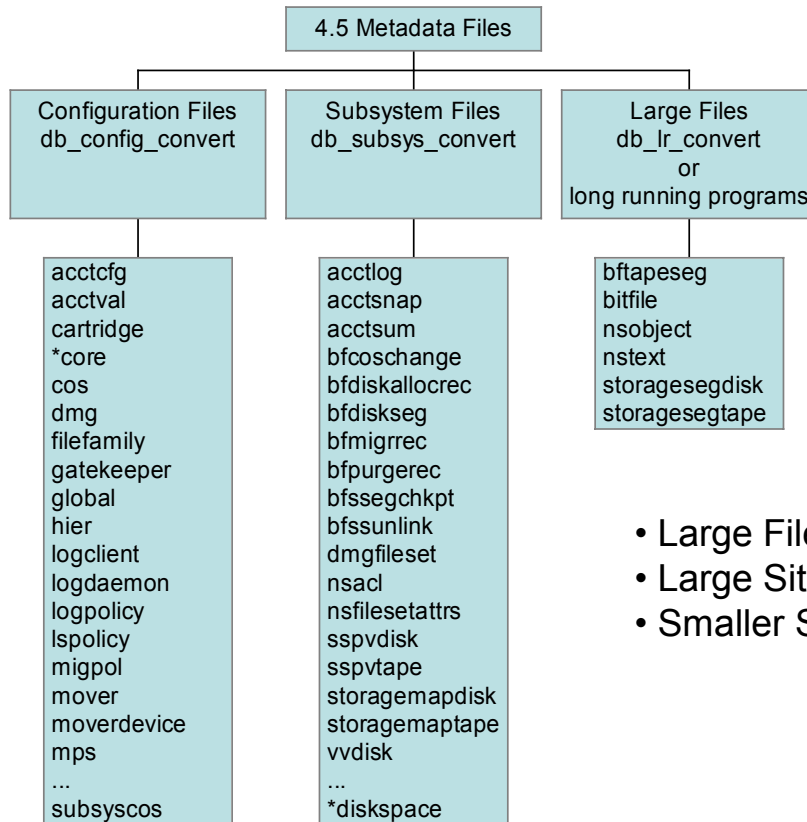
10-12 Jun 2003

rosebud:/hpss/doc/r5.1/plans/conversion/UserForum_conversion.ppt

Agenda

- **Metadata Conversion Design**
- **Features**
- **Constraints and Limitations**
- **Machine Considerations**
- **Performing a Conversion**
- **Conversion Test Results**

Metadata Conversion Design



- Large Files – contain 1 million or more records
- Large Sites – more than 3 million HPSS files
- Smaller Sites – less than 3 million HPSS files

Features

1.) Restart Capability

Long running conversion programs are able to restart without reprocessing records that have been successfully converted.

2.) Flexibility

Programs may be run using 3 scripts, or individually from the command line.

3.) Performance

Long running conversion programs may be run in parallel to decrease overall conversion time.

Constraints and Limitations

1.) Complexity of 4.5 System

The conversion programs are not able to convert 4.5 systems with multiple PVL servers

2.) Restart capability only for long running conversion programs

Configuration and Subsystem conversion programs do not have restart capability.

3.) Remote conversion has not been tested

HPSS 4.5 should be converted to HPSS 5.1 on the same machine. We have not tested machine to machine conversions.

Machine Considerations

1.) Platform supports DB2 v8.1

DB2 must be installed prior to running a conversion.

2.) Additional disk space for DB2

Cannot reuse or reclaim SFS disk space. This may require installation of additional disk for DB2.

DB2 requires more disk space than SFS for the same amount of metadata (about twice as much).

3.) Memory

The minimum amount of memory used in testing was 756 MB (we recommend between 1 to 2 GB).

4.) Consider test machine

This would allow sites to ensure proper DB2 setup and estimate overall conversion time.

Performing a Conversion

1.) Determine the number of databases you need

We recommend one database for configuration metadata (config or cfg) and a separate database for each subsystem (subsys1, subsys2, etc.).

2.) Calculate Disk Space Requirements

We recommend a RAID disk array but at a minimum each database must have:

- 2 separate disks for database logging (mirrored)
- 1 disk for each large table (> 1,000,000 records)
- 1 disk for configuration or subsystem metadata
- 1 disk for database temporary and system catalog storage

3.) Create raw logical volumes for each database

1 log volume, 1 volume for each large table, 1 temp space volume, 1 system catalog volume, 1 config or subsystem volume

Performing a Conversion

4.) Setup and configure DB2 databases, table spaces, and bufferpools

Use the DB2 GUI (control center) to create each database, table space and bufferpool. The create database wizard assists you with most of this process.

5.) Run `hpss_managetables` to create tables

Menu driven utility that performs just like `managesfs`. Allows you to create tables and place them in any of the table spaces you've configured. Also, allows you to place table indexes in a separate table space from table data if you so desire.

6.) Run the collect configuration information script `db_convert_collect_info`

This script will create 3 text files that will enable the conversion programs to understand how DB2 is configured and provide the necessary information to merge the bitfile, name and storage servers into a core server.

Performing a Conversion

- 7.) Start conversion by running the configuration conversion script *db_config_convert*

This script converts all non-subsystem metadata.

- 8.) Run the subsystem conversion script *db_subsys_convert*

This script converts all subsystem metadata except large tables (bftapesegment, bitfile, nsubject, nstext, storagesegdisk, and storagesegtape). This program must be run once for each subsystem.

- 9.) Run long running conversion script *db_lr_convert* or run each long running program in parallel. The program(s) must be run once for each subsystem.

The script will run conversion programs in serial and will take longer than running the programs in parallel. We recommend running the programs in parallel. The individual programs have restart capability, the *db_lr_convert* script does not have restart capability.

Performing a Conversion

10.) Cleanup and reconfigure DB2

We recommend you run the DB2 RUNSTATS and REORGCHK utility on each table. Might consider changing the log file size and a few other database configuration parameters (details in Chapt 14 of the HPSS 5.1 Admin Guide).

11.) Run dce conversion utility and conversion verification utilities

We have included a dce conversion utility which creates the necessary dce entries for your new core server. We have also provided four utilities to check newly converted 5.1 metadata against 4.5 metadata. This will ensure that metadata was not unintentionally altered during the conversion process.

Conversion Test Results

- **Used relatively inactive 44p170 (AIX) with 2 CPUs, 1 GB memory, and 6 disk array (Clarion) running conversion on imported SFS production data (3,000,000 files).**

Configuration Conversion = 2 min, less than 2 MB disk space

Subsystem Conversion = 10 min, less than 2 MB disk space

Long Running Conversions in parallel

bftapeseg (6,285,718 recs) = 55 min, 1.2 GB

bitfile (2,915,551 recs) = 30 min, 1.0 GB

nsubject (3,707,615 recs) = 108 min, 2.0 GB

nstext (3,059,569 recs) = 20 min, 3.3 GB

storagesegtape (5,779,962 recs) = 110 min, 3.0 GB

storagesegdisk (376,923 recs) = 9 min, 150 MB

SFS Disk Space = 5.678 GB

DB2 Disk Space = 10.697 GB

Conversion Test Results

- **Used relatively inactive H50 (AIX) with 4 CPUs, 1 GB memory, and 10 disk array (Clarion) running conversion on SFS production data (7,000,000 files).**

Configuration Conversion = 2 min, less than 2 MB disk space

Subsystem Conversion = 25 min, less than 2 MB disk space

Long Running Conversions in parallel

bftapeseg (10,597,423 recs) = 6 hrs 24 min, 2.4 GB

bitfile (7,089,681 recs) = 6 hrs 34 min, 2.7 GB

nsubject (7,663,513 recs) = 6 hrs 52 min, 8.4 GB

nstext (6,022,705 recs) = 1 hr 54 min, 8.1 GB

storagesegtape (10,597,168 recs) = 11 hrs 57 min, 5.9 GB

storagesegdisk (373,782 recs) = 31 min, 163 MB

SFS Disk Space = 14.277 GB

DB2 Disk Space = 27.663 GB

Conversion Test Results

- **Used relatively inactive IBM Winterhawk node (AIX) with 2 CPUs, 756 MB memory, and 4+1 Raid 5 disks running conversion on SFS production data (8,500,000 files).**

Configuration Conversion = 1 min, less than 2 MB disk space

Subsystem Conversion = 28 min, less than 2 MB disk space

Long Running Conversions in parallel

bftapeseg (12,582,048 recs) = 2 hrs 47 min, 2.5 GB

bitfile (8,522,065 recs) = 6 hrs 1 min, 3.1 GB

nsubject (8,792,049 recs) = 2 hrs 56 min, 9.2 GB

nstext (1,370,118 recs) = 3 min, 1.8 MB

storagesegtape (12,582,056 recs) = 7 hrs 1 min, 5.3 GB

storagesegdisk (996,697 recs) = 22 min, 424 MB

DB2 Disk Space = 22.324 GB

Questions?