

# DMF and PostgreSQL

Joseph C. Nemeth

Technical Lead

# RDM (RAIMA Data Manager)

- Record-based database maintained by SGI
  - Non-scalable design
  - Hard limit of 4 billion indexed entries
  - Snapshots involve physical copy of table data
  - No support for foreign keys or table joins
  - Limited ACID compliance (Atomicity, Consistency, Isolation, Durability)

# PostgreSQL

- Industry-standard database
  - Scalable design (size and performance)
    - Supports database clustering
    - 32 TB limit on table size, supports sharding
  - Logical snapshots of all tables
  - Foreign keys and table joins supported
  - Fully ACID compliant (Atomicity, Consistency, Isolation, Durability)
  - Industry standard query language

# Additional Advantages

- Row locking maintains DB availability during normal DMF operations
- Supports parallel DB dumps during normal DMF operations
- Logical snapshots avoid redundant data transfer
- Eliminates need for dmdbcheck of dumped rows
- Eliminates manual dependency checks of foreign keys

# Implementation

- dmscanfs enhanced to store paths, handles, and bfids in the core database
- dmhdelete enhanced to use direct logical snapshots of the database
- dmaudit enhanced to avoid redundant data copying and scale to larger data sets
- Misc user tools like dmvoladm, dmdump all support the new database using the old query formats

# Optimization

- DB is normalized to prevent duplication of column values across tables
  - object attributes like bfid, inode, and size are stored in a single table
  - chunk table refers to object attributes by foreign keys
- All library server tables stored in a single database cluster with logical namespaces
- Large columns are automatically stored in compressed format
- Indexes are maintained on frequently-used columns

# Deployment

- DMF version 6.4 (spring 2015)
  - Supports RAIMA
- DMF version 6.9 (spring 2015)
  - Same as 6.4
  - Supports PostgreSQL
  - Contains tools to convert RAIMA to PostgreSQL
- DMF version 7.0 (fall 2015)
  - Supports PostgreSQL
  - Contains tools to convert RAIMA to PostgreSQL