Starfish

A Side-band Database for HPC and Archival Storage Systems

Part of a larger solution for managing the life cycle of scientific research from creation through publication and reuse.

Our Core Technology: Sync File System to Database

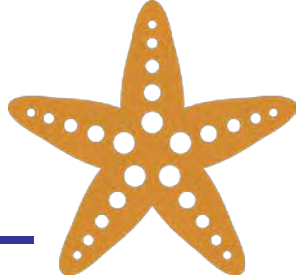

- **Imagine if it were easy to keep your file systems synchronized with a database?**
  - All of the stat() metadata that comes off POSIX
  - Additional file system metadata such as in GPFS
- **Imagine if you could add tags or key-value pairs to the records that represent files and directories.**
- **Imagine if the database kept version histories of the directory tree and individual files.**
- **Imagine if the database pre-staged some common aggregate values up and down the directory tree?**
  - Total files, total capacity, etc.

What Would You Use The Database For?

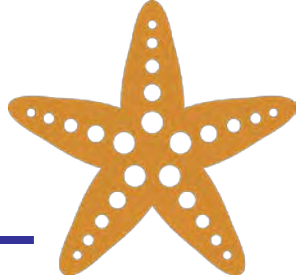

- **Reporting**
  - Better reports enabled by extensible metadata
- **Running Scripts / Feeding batch processes**
- **Data migration workflows**
  - Migration
  - HSM
  - Backup/Restore
  - Check-in / Check-out
  - Move to and from object store
- **Single namespace user portal**
- **Resolving broken links and finding lost files**
- **Calculating and storing hashes**
  - Fixity checking, duplicate file detection, content addressing

What Makes the File System Catalog Awesome?

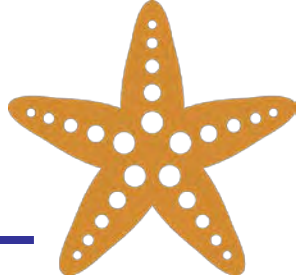

- **Massively scalable**
  - Handles billions of files
  - Multi-threaded and multi-host for greater parallelization
  - Highly tunable and configurable
- **Agents for specific file systems**
  - Agents capture file system events reducing the need to crawl and compare
  - Agents capture device-specific metadata
- **Metadata persists as files and directories move around**
  - Add tags and key-value pairs to files and directories
  - Directory-level metadata can be inherited down the tree
  - Metadata is retained even when file system objects are moved and renamed.
- **Version histories**
  - We track version changes of individual files
  - We keep a version history of the directory tree

Versioning – Critical Feature

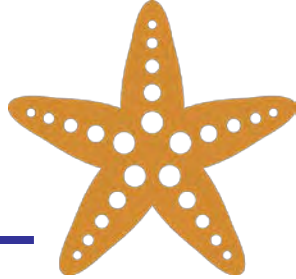

- **Backup/ Restore**
  - Replaces enterprise backup software
- **Permanent Addressing**
  - A digital object has a permanent address in the form of path name + time/date
- **Find missing files**
  - Query the catalog with the “last known address". Find out where the file is now.
- **Virtual HSM**
  - Individual files can be removed from the POSIX name space and moved to lower cost storage while retaining the file record in the virtual namespace.
- **Checkpoint**
  - Retain a collection of files at a point in time
- **Provenance**
  - Point-in-time representations of file collections provide a foundation for data provenance

The Grand Vision

Publication/Preservation
(Librarians, Archivists, Curators)

Open Links / DOIs
Metadata Extraction
Curation Workflows
Version Controls
Access Controls
Fixity Checks

Content Creation
(Scientists, Engineers, Artists)

Metadata Tagging
Workflow Automation
Data Management Plans
Open Access
Data Reusability
Collaboration

IT Operations
(Storage & Backup Administrators, IT Governance)

Data Movement
Tiered Storage
Backup Restore
Data Migration

Governance
Permissions Management
Auditing
Chargeback / Show-back

Reporting
Capacity Planning
Aging / Utilization
File System Analysis

Bragging Rights

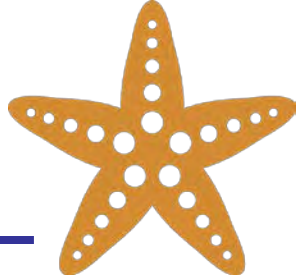

- **Easy to install - 10 mins for core system**
  - Major components discover themselves.
  - Upgrades invoked by a single command from CLI
- **Largest single installation: 8+ billion files**
- **Scanning at a rate of 2.8 billion files per day**
  - 30,000+ file system events per second
- **51 sites using the software as of May 2017.**
  - Most are top tier data centers and/or household names
- **Multi-phase duplicate checking at 1.7PB/day**