



STARFISH

Virtual Global File System

Metadata and Rules Framework for File
Systems and Object Stores
Peter Baer Galvin – pgalvin@starfishstorage.com

What Makes Starfish Unique?



- ★ Speed and Scale
 - ★ Versioning
 - ★ Directory Aggregates
 - ★ Policies, Workflow, File Movement
 - ★ Storage synchronization and migration
-

Speed and Scale



- ★ Fast initial crawl, with faster subsequent crawls
 - ★ Scale to 10s of BILLIONS of files (compare to hundreds of millions with similar tools)
 - ★ Easily track 500M-3Billion file changes/day
 - ★ Parallel simultaneous scans of file systems
 - ★ Crazy-fast with Starfish Agent Systems (Change Monitoring)
-

File Versioning – A Major Differentiator



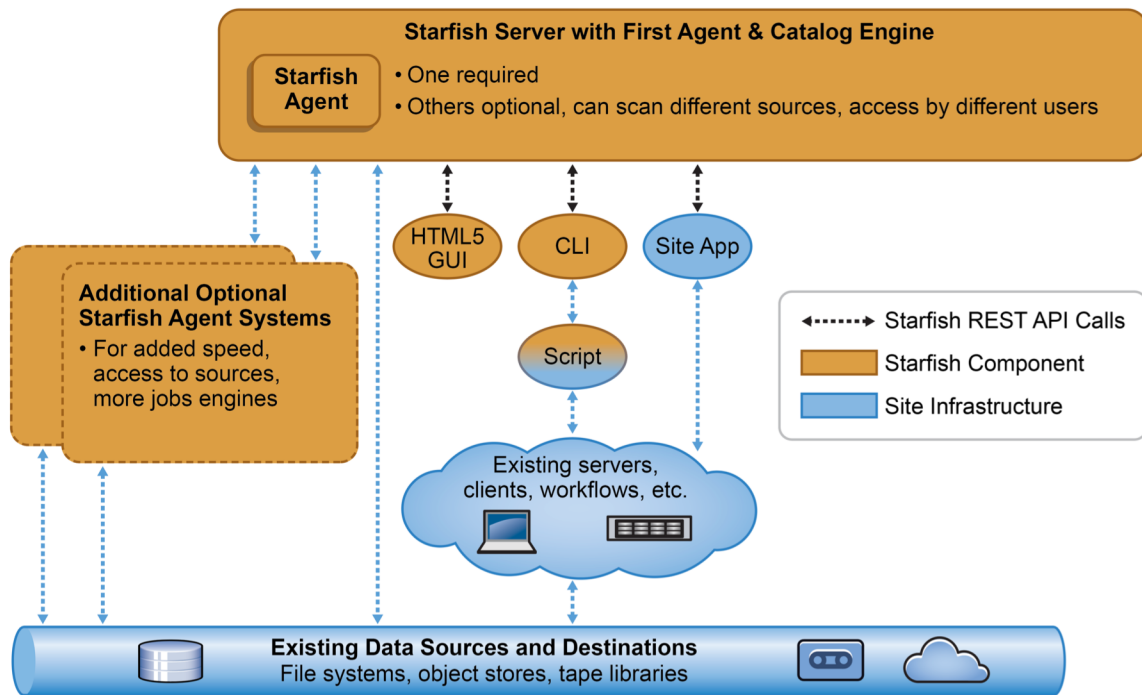
- ★ Preserve file history, changes, locations
 - ★ Save ‘Sticky’ Metadata with the files and directories regardless of file changes or moves
 - ★ Permanent Addressing in the form of path name + time/date
 - ★ Find new file location: Query the catalog with “last known address”.
 - ★ Unified Database: 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.
 - ★ Provenance: Point-in-time representations of file collections with all associated metadata
-

Directory Aggregates



- ★ Instant recursive totaling - akin to instant UNIX DU
 - ★ Information presented for each directory branch
 - File count
 - Total capacity
 - Storage costs
 - ★ Ability to store and aggregate other variables
 - ★ History retained for historical & trend reporting
-

So What is Starfish?



- ★ Hardware agnostic software
- ★ Out of band scanning - keep existing workflows
- ★ Runs on common Linux distros
- ★ Simple licensing model allows for multiple systems, Scale out to multiple Catalogs
- ★ **Crawls any POSIX-based file system**
- ★ Direct integrations being built for:
 - Lustre
 - Isilon
 - GPFS
 - Netapp
 - ZFS
- ★ Command-line, GUI and communications via REST APIs
- ★ Starfish Server and optional agent systems can be virtual or physical

What's in our Catalog?



- ★ POSIX Stat() metadata and S3-type object store
 - ★ Extended attributes and non-POSIX system metadata
 - ★ Permissions
 - ★ Additional, user added metadata for directories and files
 - Simple tags
 - Key-value pairs
 - ★ Directory-level pre-staged totals
 - ★ Version history of files and directory tree
-

What makes our database magical?

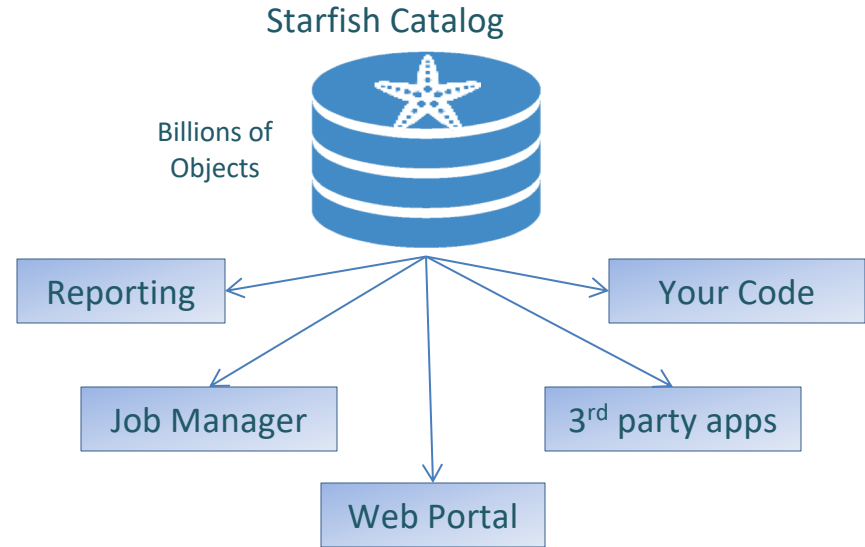


- ★ More metadata = More meaningful reports
 - Chargeback/showback, churn, growth...
 - ★ More metadata = More specific policies and rules
 - Attributes drive policies
 - ★ Simple Triggers
 - Metadata changes trigger processes
 - ★ Instant Search
 - ★ Unique IDs for all files
 - ★ **Did we mention VERSION HISTORY?**
-

What can you do with the catalog?



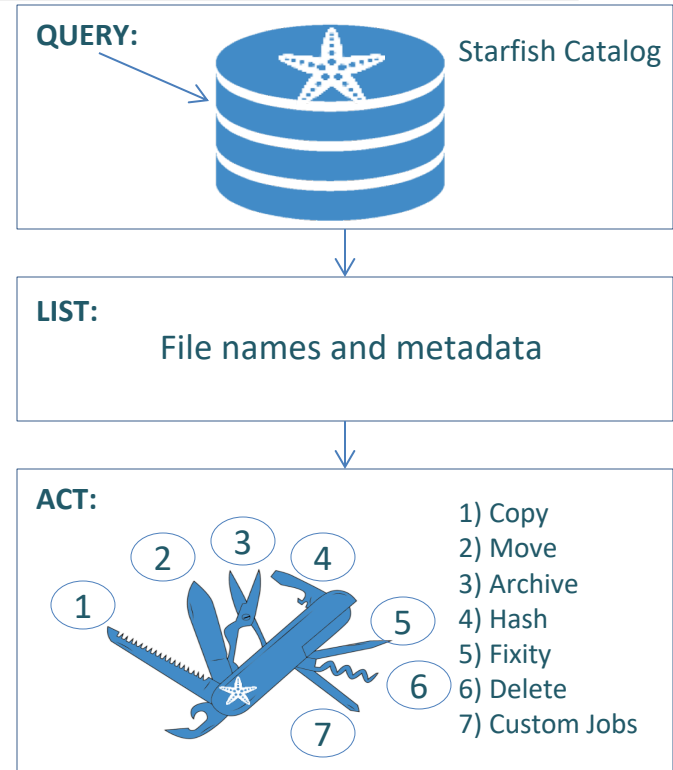
- ★ Web portal access
- ★ Browse, search, view the database
- ★ Granular reporting and analytics
- ★ Rules-based storage administration
- ★ Copy, move, delete
- ★ Get / put to object store
- ★ Extract metadata, calculate hashes
- ★ Customize jobs, actions against files
- ★ Run jobs against the database using your code, ours, or 3rd party software



Starfish in Action – Use Cases



- ★ Updating video codecs
- ★ Move aging files to cloud
- ★ Replication
- ★ Auto-delete scratch older than X days
- ★ Ransomware discovery
- ★ Find lost directories
- ★ Find files that shouldn't exist (mp3, .vid)
- ★ Remove files associated with expired licenses
- ★ Export control
- ★ Discover extremely deep directories
- ★ File workflow management via tagging
- ★ Chargeback/Showback/Shameback
- ★ Event driven alerts



Browser View, with Tags and Details



STARFISH
Dashboard **Browser** Jobs Scans
4.0.4247+3747c3f+gui_284

VOLUMES

ENRON
1.75 GiB/498 GiB

SAMPLE
30.29 GiB/469.46 GiB

DEMOS
53.26 MiB/499.7 GiB

DEMO3
53.26 MiB/499.7 GiB

DEMO6
53.26 MiB/499.7 GiB

DEMO4

SEARCH

Advanced search

CURRENT VIEW FILTERS

ACCESS TIME

>10 years 0

MODIFICATION TIME

>10 years 0

DIRECTORY SIZE max: 12.92 GiB

0 B 12.92 GiB

FILE SIZE max: 5.38 GiB

0 B 5.38 GiB

Reset filters

DATA AS OF

< >

📁 / sample / sample-file-system (i)

File ↑	Size	Owner	Count	Cost	Modified
📁 4 files	10.17 GiB		4	0.00	
📁 Bthread <input type="text" value="archive-requested"/>	28.32 MiB	root	115	0.28	2018-01-26 15:07
📁 Cthead <input type="text" value="demo-for-dave"/>	28.32 MiB	root	115	0.28	2017-09-27 17:21
📁 Dthead <input type="text" value="is-this-needed"/>	56.64 MiB	root	230	0.55	2018-02-13 12:36
📁 Fthead	28.32 MiB	root	115	0.28	2018-01-26 16:08
📁 growing <input type="text" value="too-big"/>	375.45 MiB	root	100,131	3.67	2018-02-13 12:35
📁 igdiscover-testdata-b135a13a3595	2.66 MiB	root	9	0.03	2017-06-29 04:35
📁 largenumfiles <input type="text" value="jim"/>	5.39 MiB	root	199,003	0.05	2018-02-19 16:22
📁 materials <input type="text" value="keep-until-jan-2023"/>	53.26 MiB	root	304	0.52	2009-06-03 12:00
📁 materials2 <input type="text" value="instrument-D1"/> <input type="text" value="owned-by-fred"/> <input type="text" value="process-c8"/> <input type="text" value="process-d8"/>	53.26 MiB	root	304	0.52	2017-09-28 16:17
📁 medical <input type="text" value="lab-biology-smith"/> <input type="text" value="PCI"/> <input type="text" value="PII"/>	12.92 GiB	root	380,057	129.22	2017-09-28 14:26

Stats Details

History

Cost 3.67\$

Subdirectories 148

Files 99,983

TOP 5: EXT | USERS | GROUPS

Reporting



User Size Change Rate



Minimum delta size (GB) Minimum percent change Number of days to look back

	USER_NAME	VOLUME	START DATE	END DATE	PERCENT DELTA	PREVIOUS SIZE GB	CURRENT SIZE GB	DELTA SIZE GB
Zoom	[redacted] 2	[redacted] 6	10/27/17	11/06/17	194%	16,355.40	48,010.30	31,654.80
Zoom	[redacted] 7	[redacted] 5	10/27/17	11/06/17	248%	12,548.40	43,710.40	31,162.00
Zoom	[redacted] 4	[redacted] 6	10/27/17	11/06/17	13%	228,925.10	257,855.20	28,930.10
Zoom	[redacted] 3	[redacted] 6	10/27/17	11/06/17	5%	543,425.00	572,349.00	28,924.10
Zoom	[redacted] 6	[redacted] 7	10/27/17	11/06/17	72%	24,585.90	42,220.50	17,634.60
Zoom	[redacted] 6	[redacted] 6	10/27/17	11/06/17	15%	108,806.00	125,234.80	16,428.80
Zoom	[redacted] 29	[redacted] 6	10/27/17	11/06/17	50%	31,191.10	46,702.10	15,511.00
Zoom	[redacted] 9	[redacted] 5	10/27/17	11/06/17	57%	24,822.20	38,901.50	14,079.30
Zoom	[redacted] 7	[redacted] 7	10/27/17	11/06/17	5%	242,068.00	255,140.90	13,072.80
Zoom	[redacted] 34	[redacted] 6	10/27/17	11/06/17	46%	25,881.20	37,687.50	11,806.40
Zoom	[redacted] 7	[redacted] 5	10/27/17	11/06/17	320%	3,633.90	15,245.90	11,612.00
Zoom	[redacted] 1	[redacted] 7	10/27/17	11/06/17	12%	82,481.90	92,033.40	9,551.50
Zoom	[redacted] 8	[redacted] 5	10/27/17	11/06/17	2914%	322.00	9,704.20	9,382.20
Zoom	[redacted] 16	[redacted] 6	10/27/17	11/06/17	3%	240,897.70	248,804.70	7,907.10
Zoom	[redacted] 32	[redacted] 6	10/27/17	11/06/17	37%	17,164.10	23,591.60	6,427.50

Reporting



File Age Report

Last access by user

Setting size filter will filter all views (size, count, cost). Note that totals will reflect the filtered data. To show volume-wide totals, set size filter to 0.

Min size per user (GB)

Col Heatmap

Sum

volume

user_name

		size (GB)	count	uid	group_name	cost	mtime_age	gid				
		atime_age										
volume	user_name	atime_age	Previous Months: 0-1	Previous Months: 1-3	Previous Months: 3-6	Previous Months: 6-12	Previous Years: 1-2	Previous Years: 2-3	Previous Years: > 3	Totals		
curie2	9999		10.00	24.00	63.00	17,745.00	3,795.00	17,568.00	297,372.00	336,677.00		
dgadsv	root		934.00	3,231.00	3,203.00	24,790.00	11,385.00	39,347.00	105,822.00	188,718.00		
etc	root		83.00				534.00			617.00		
mailun	root					92.00				92.00		
e_data1	1201			27,177.00	501,830.00					509,007.00		
	1204				150,313.00					150,313.00		
	1205			3,083.00	4,716.00					7,799.00		
	1206				27,191.00					27,191.00		
	4294967294				79,923.00					79,923.00		
	root			30,248.00		17.00				30,265.00		
Totals			1,027.00	63,761.00	827,263.00	42,633.00	15,714.00	57,015.00	403,194.00	1,410,606.00		

Updated: a few seconds ago

