



A Virtual Object Store

and



Yet Another HSM

An update from the GA Storage group

How it all started?

- The Storage Software Group at General Atomics has a long tradition in data handling
 - Going back to 1995 when we ran the San Diego Supercomputing Center
- Main product is Nirvana, a complete (virtual) object store solution
 - Storage abstraction, federation and tiering
 - Metadata repository, extraction and handling
 - A policy/ILM engine
 - Audit trail

<http://www.ga.com/nirvana>

A Virtual Object Store

- Nirvana operates on top of existing storage
- Nirvana does not provide its own filesystem solution
 - There are many available, why re-invent the wheel?
- Nirvana can operate on top of anything that supports
 - File paths
 - Basic file permissions
 - Data ingest and retrieval



At least in theory, we provide implementation for a fixed set

What was the problem?

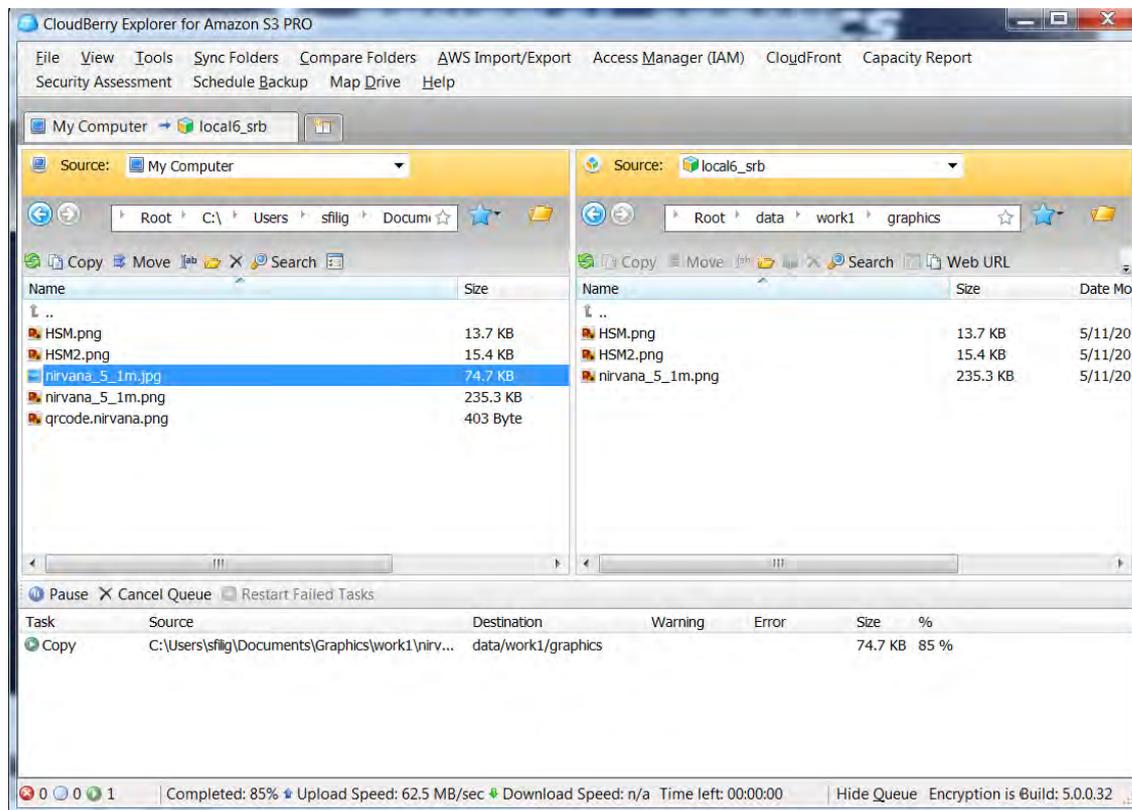
- Nirvana was very capable, but lacked standard APIs
 - Hard to integrate in existing workflows
 - Even though it had a rich set of proprietary tools, portals and APIs
- So we went ahead and added
 - A standard object storage API, implementing the S3 REST API
 - A standard POSIX interface, and we called it EasyHSM

Nirvana with S3

- The S3 REST protocol has become the de-facto object storage standard
 - Based on published AWS S3 documentation
 - Many client tools and libraries available
- Since Nirvana is essentially an object store, it was a natural fit
- Currently provided as an optional component
 - Can be installed on dedicated nodes, or alongside the other Nirvana daemons

Nirvana with S3 (cont)

- Nirvana accessed through Cloudberry Explorer

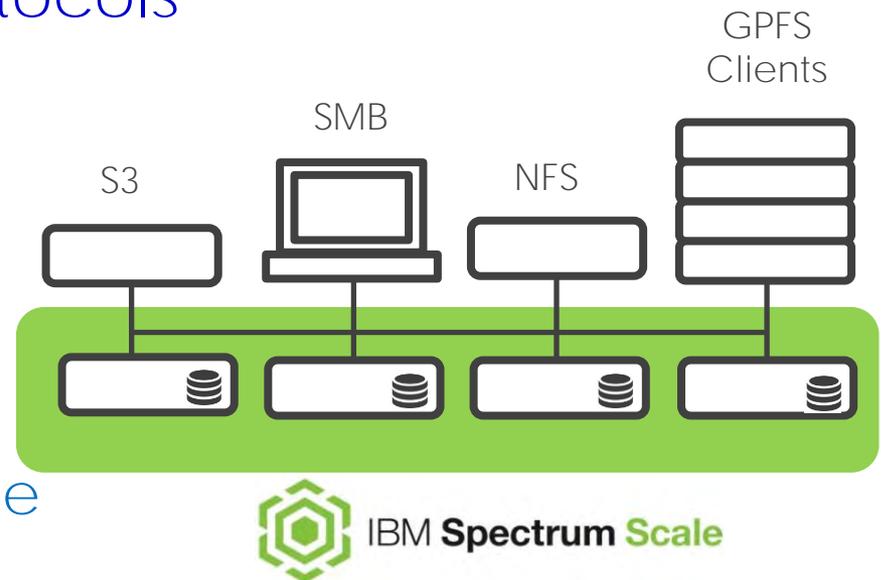


The POSIX interface

- Nirvana natively does not provide a fully POSIX compliant native API
 - A natural restriction due to support for many different backend storage solutions
- Developing a fully featured caching filesystem not feasible
- We looked for an existing filesystem that is HSM enabled out of the box
 - IBM Spectrum Scale (aka GPFS) seemed the best

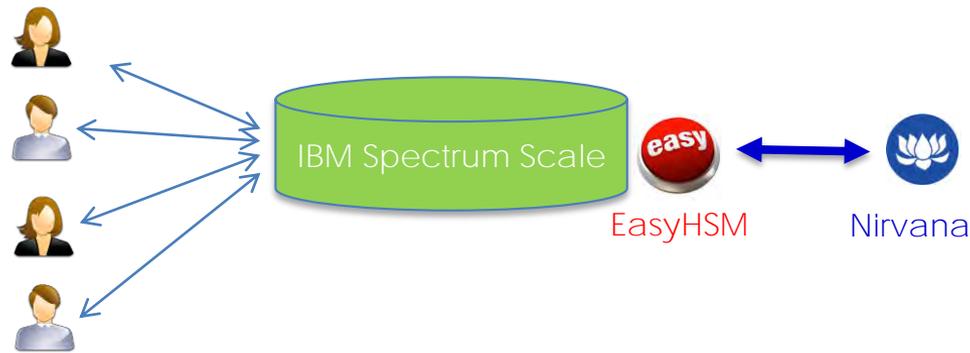
The IBM Spectrum Scale Advantage

- A mature, very reliable storage system
 - With solid backing from IBM
- Highly scalable
 - Originally designed for HPC applications
- Multiple access point protocols
 - High performance Native Client Drivers
 - Traditional NFS, SMB/CIFS
 - Even an S3 interface
- Natively HSM-enabled
 - Standard DMAPI interface
 - Rich ILM capabilities



EasyHSM as a Nirvana POSIX frontend

- EasyHSM designed as a Hierarchical Storage Management (HSM) solution for IBM Spectrum Scale (GPFS)



- GPFS provides the namespace management
 - And hot data storage management
- EasyHSM provides the tiering logic
 - Between GPFS storage and external storage

Why yet another HSM?

- All available HSM solutions were cost prohibitive
 - They employ capacity based licensing!



EasyHSM is not free either,
but we charge a flat fee

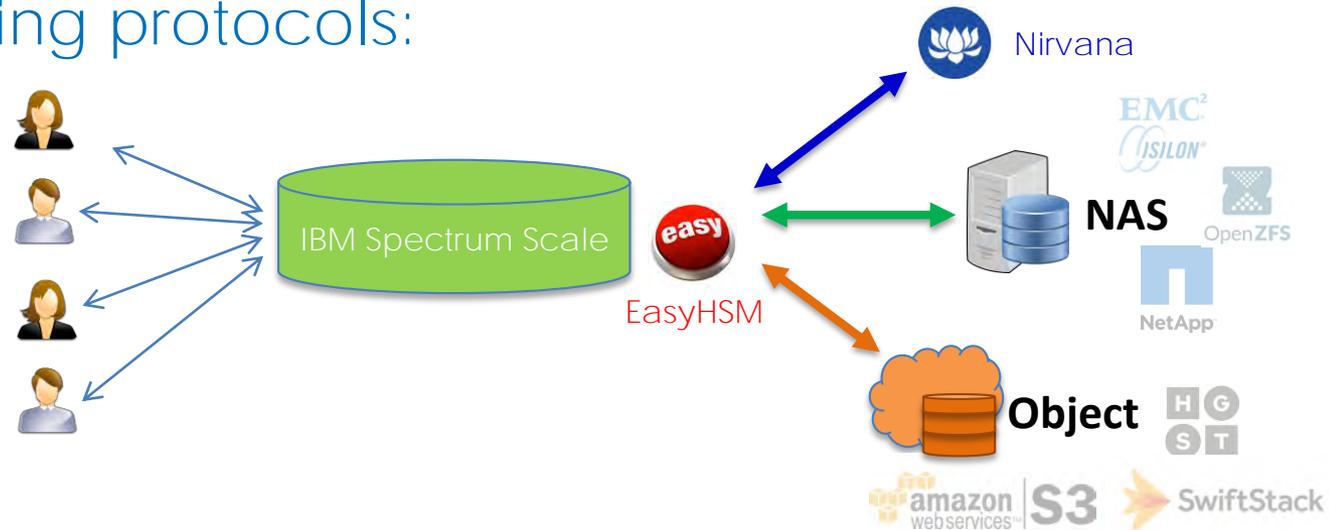
- Plus, none would talk the native Nirvana protocol

Going beyond Nirvana

- During EasyHSM development, it become clear that there was very little Nirvana-specific code
 - So we decided to make EasyHSM generic

- EasyHSM now supports the following external/tiering protocols:

- POSIX
- S3 REST
- Nirvana



EasyHSM vs Competition

- EasyHSM is the only solution that allows tiering using multiple external protocols (both POSIX and Object)
- EasyHSM does a straight mapping between GPFS path and external storage path/URI
 - Minimal massaging, e.g. adding protocol and IP
 - No obfuscation, easy to interact with external tools
- Scalable, low overhead implementation
 - Linear scaling over multiple hardware nodes
 - Not in the path for hot data
- Pure software solution
 - Does not need dedicated hardware
- Fixed price licensing (not capacity based)

For more information

- EasyHSM home page:
<http://www.ga.com/easyhsm>
- Nirvana home page:
<http://www.ga.com/nirvana>

If you just want to know more, feel free to

- Contact me at
Igor.Sfiligoi@ga.com
- Contact the program manager at
Robert.Murphy@ga.com

Other Recent Nirvana Changes

- Dramatic scrubbing/reconciliation speedup
 - System attributes: observed 1M files/minute (16k/s)
 - JPG scrubbing: observed 60k files/minute (1k/s)
- New administration interface

The screenshot displays the Nirvana Administration web interface. The top navigation bar includes the Nirvana logo, the title 'Administration', and the user 'Logged in as Admin'. The breadcrumb trail shows 'Web admin > Resources > Physical > cabx3 >'. A left sidebar contains navigation icons for Users, Domain, Location, Resources (highlighted), Daemons, and Schema. The main content area is titled 'Information' and shows the following fields:

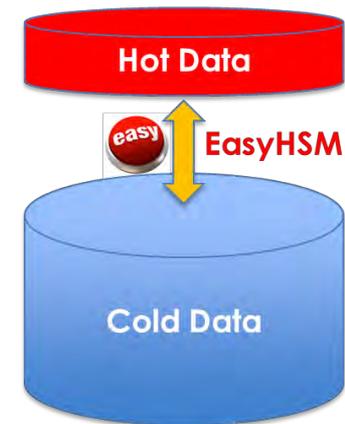
Resource Name:	<input type="text" value="cabx3"/>
Resource Path:	<input type="text" value="/drive/c/a/b/x3/?COLLECTI"/>
Register:	<input checked="" type="checkbox"/>
Unregister:	<input checked="" type="checkbox"/>
Location:	<input type="text" value="martinvm@nir"/>
Class:	<input type="text" value="Archival"/>
Type:	<input type="text" value="disk file system"/>

Below the 'Information' section is the 'Configuration: Disk File System' section with the following fields:

DEBUG_LEVEL:	<input type="text" value="0"/>
LOG_FILE:	<input type="text" value="disk.log"/>
LOG_FILE_SIZE:	<input type="text" value="100000000"/>
LOG_FILE_GROUP:	<input type="text" value="Log_file_group"/>
LOG_FILE_MODE:	<input type="text" value="Log_file_mode"/>
RESOURCE_CAPABILITY:	<input type="text" value="Read & Write"/>

Why use EasyHSM?

- Not all data needs to be in the fast, GPFS layer
 - But all files must be visible to the users
- EasyHSM provides an effective tiering solution
 - Uses HSM file stubs to link to external storage
 - Policy driven explicit movement of data between GPFS and external storage
 - Transparent high performance retrieval of data into the GPFS layer on user access



EasyHSM components

- Tools for migrating data from GPFS to external storage
- Tools for freeing disk space on GPFS (stubbing)
 - Invisible to final users
- A daemon for transparent recall on access
 - And handling file removal
- Tools for explicit recall from external storage to GPFS
- Support scripts for easy integration with GPFS Policy Manager (ILM)