

A Web Based Data Catalog for Data Access and Analysis

An iteration over the Fermi Gamma-Ray Space
Telescope Data Catalog developed by
The Scientific Computing Applications Group



Brian Van Klaveren
with Tony Johnson

bvan@slac.stanford.edu

tonyj@slac.stanford.edu

CHEP2013 2013-10-04



This Talk

- Overview of the Data Catalog
- Updating the Data Catalog for use by different experiments
- Moving to RESTful APIs for the Data Catalog
- Leveraging HTML5 for a modern user experience

Data Catalog Overview

- Initially developed for Fermi Gamma-Ray Space Telescope (Fermi-LAT)
- Designed to be experiment independent
- Development started 2007, in use since launch 2008
- Written in Java (JSP for web)
- Now in use by other experiments (EXO, LSST)

It is a metadata database for your files.

And that metadata is stored in a virtual hierarchy.

It is not tied to any processing system.

It is not tied to any file system or protocol.

It needs only an RDBMS and servlet container.

More detail

Comprised of several software components:

- APIs for creation, reading, modification, and searching
- A web application for user browsing and interaction
- A daemon which checks file integrity and automating metadata extraction from files
 - Plugins are written for supporting new file types, access protocols, and metadata extraction
- A download manager for reliable download of files and file collections which fetches underlying files

Data Catalog Primitives

The Data Catalog consists of three primitives, all of which can be tagged with user-defined metadata.

- A Dataset is a logical set of data (run, MC, calibration, e.g. file)
 - Can have different versions
 - Can have multiple physical locations
- A Folder is a container. It can contain other folders, groups, or datasets.
- A Group is a Dataset container, wherein the Datasets contained are restricted to a common semantic type. It's basically a macro-Dataset.

Current web application



Fermi LAT Data Catalog

Version: 1.11

[Login](#) | [Jira](#)

Mode: [[Prod](#) | [Dev](#) | [Test](#)]

View: [[Tree](#) . [Data Types](#) . [File Formats](#) . [Messages](#) . [Admin](#) . [Problems](#)]

Show: MC beam test obsolete

[Update](#)

To view private folders you have to [login](#).

Folders

- Data
 - Flight
 - Level1
 - Reprocess
 - P130
 - P200
 - P202
 - CAL
 - DIGI
 - DIGIGAP
 - ELECTRONFT1
 - ELECTRONMERIT
 - EXTENDEDFT1
 - EXTENDEDLS1
 - FILTEREDMERIT
 - FT1
 - FT2
 - FT2NOQUAL
 - FT2SECONDS
 - FT2SECONDSNOQUAL
 - GCR
 - LS1
 - MERIT
 - RECON
 - P300

Folder /Data/Flight/Reprocess/P202 Group MERIT

Created (UTC):	25-Jan-2012 00:53:30
Run Min:	239557414
Run Max:	402560477
Files:	28675
Events:	62,383,754,997
Size:	49.9 TB
Data Type:	MERIT

[List Files](#) . [Download Files](#) . [Skim Files](#) . [Dump file list](#)

Meta-data

Name	Value	Type
sBackup	true	STRING

Aggregate information
(status)

Folders

Groups

Modernizing and Generalizing the Data Catalog

Why?

- So we can reuse it and better maintain it
- Better API and Web App performance

What do we need to support going forward?

- Reduce the Data Catalog to core features
 - Add plugins for experiment-specific features
- Create language-agnostic APIs (we use REST)
- Provide secure access to the Data Catalog over the web
 - Again: Not the underlying files, just info about them
- Modernize the web application (HTML5, AJAX)

REST

(Representational State Transfer)

Or...

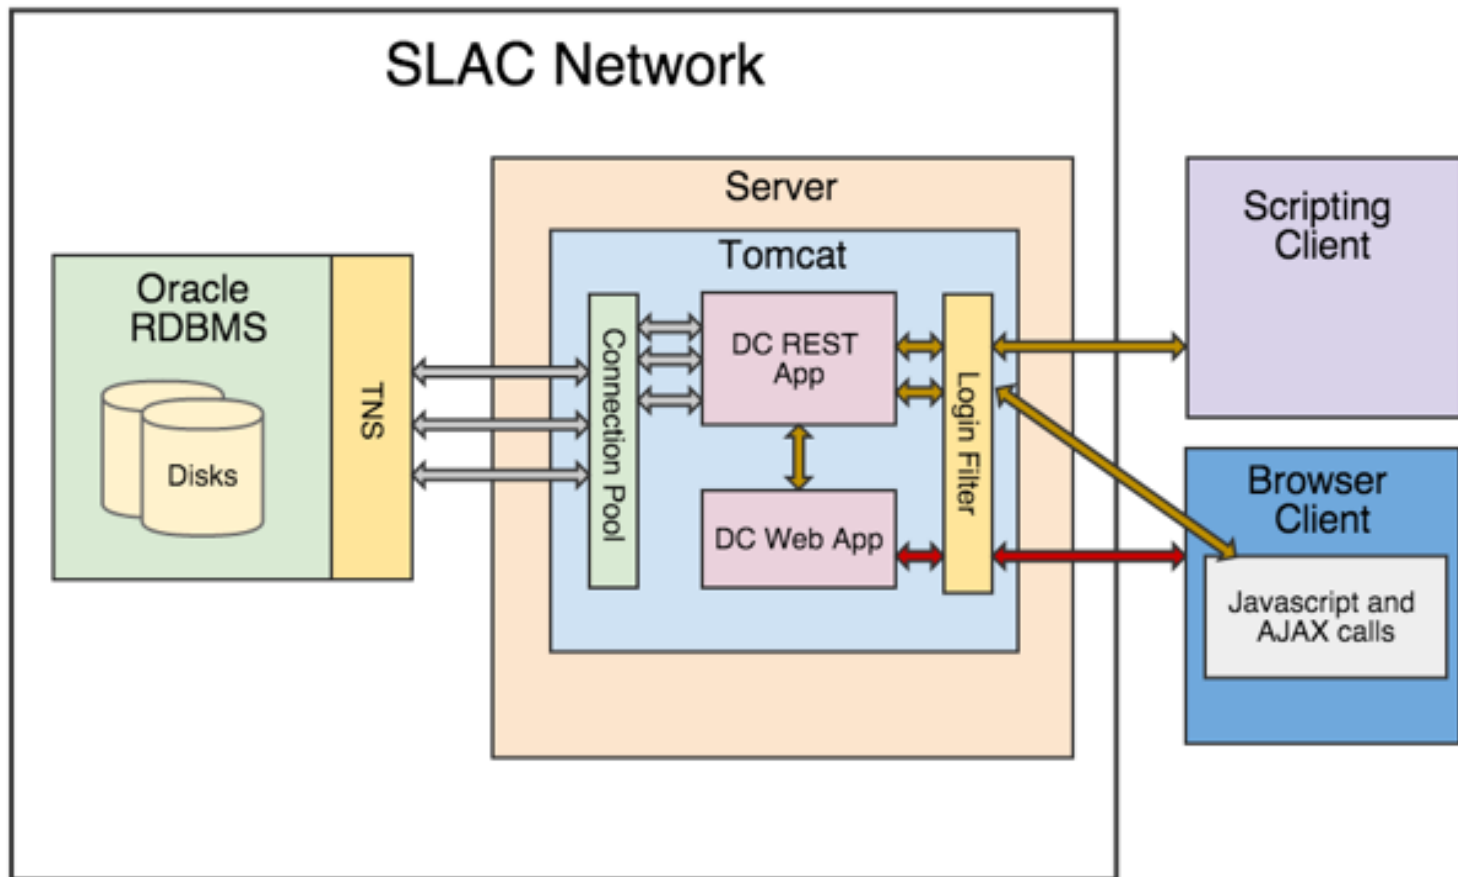
Stateless Create/Read/Update/Delete
Implemented over HTTP
using Resources (not methods)

RESTful API

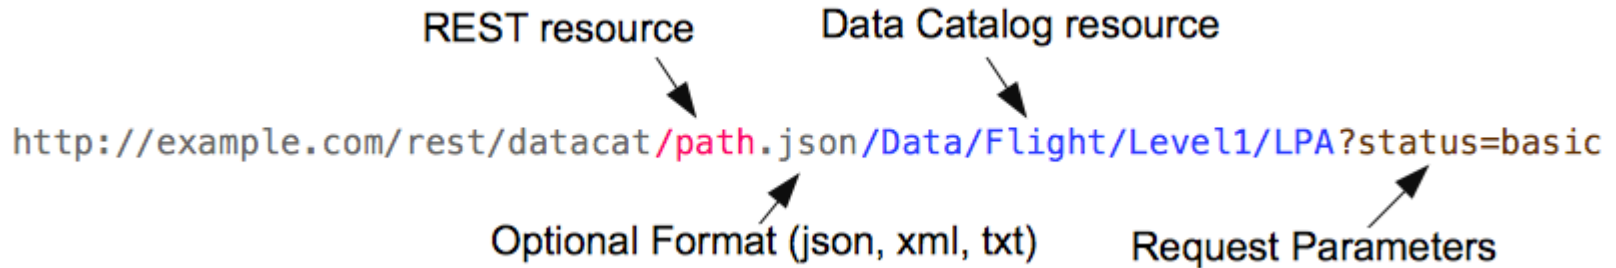
A scalable, supportable way to supporting [C/C++, Java, Javascript, Python, Perl, Ruby, etc...] is to use a RESTful API over HTTP.

- We will let the user determine the output they want
 - JSON, XML, text
 - Made easy thanks to the help of Jersey/JAX-RS
- HTTP/XML/JSON is well supported in nearly all languages
- We get a single point of API access = One API to develop
 - Easier to manage resource use
 - If needed, scaling horizontally is easy
- The Web Application is driven off the same REST API

RESTful API – What does this look like?



RESTful API – Resources



JSON
Output:

```
{
  "Stype": "folder",
  "name": "LPA",
  "path": "/Data/Flight/Level1",
  "metadata": {
    "entry": [
      { "key": "nKeyData", "value": 10 },
      { "key": "eventComponents", "value": "reco=RECON,digi=DIGI,merit=MERIT" },
      { "key": "FitsSkim", "value": "Level 1 LPA data" }
    ]
  },
  "stat": {
    "Stype": "stat",
    "datasetCount": 0,
    "groupCount": 75,
    "folderCount": 0,
  },
  "description": "Output from Level 1 processing of on-orbit data"
}
```

RESTful API – /datasets (Dataset records)

Listing one dataset:

- GET /datasets/Repro/P2/Overlay/r0239557414.root

Searching for datasets:

- GET /datasets/Repro/P2?filter=nRun=239557414+AND+quality+=+'GOOD'

Register a dataset

- POST /datasets/Data/MC/Co60?name=r0123.root&format=ROOT&type=MC \&site=SLAC&fspath=/afs/slac/g/mc/r0123.root

Modify Dataset - add/modify new versions, locations, or dataset metadata

- PUT /datasets/Data/MC/Co60/r0123.root
[{"key": "nRun", "number": 123}, {"key": "source", "string": "cobalt"}]

DELETE - Again, deleting the DC record, not the underlying file(s)

- DELETE /datasets/Data/MC/Co60/r0123.root

/children


- Returns the children (with option to omit datasets) of any container object, with optional paging
- Useful for building lightweight browsing interfaces

/groups, /folders

- Used for creating, modifying, or deleting groups and folders in the same way datasets is used for datasets.

Searching

Search expressions are composed of terms and conditions. The parser creates a syntax tree and translates that to SQL or defers to a plugin.



The parser is aware of three types of terms:

- A primitive's field name or an alias of that name (these are built-ins)
 - name = "run_012345.fits"
- A primitive's metadata key
 - sIntent IN ('nomSciOps', 'nomSciOps_diagEna')
- Or a plugin. In this case, "exp" is the plugin, and "quality" is the term
 - exp.quality IN ('GOOD', 'GOLDEN')

Searching - Conditions and examples

Use words or symbols. Conditions map to SQL:

- <, >, <=, >=, |=|=, !=|~=|<>
- LT, GT, LTEQ|LE, GTEQ|GE, EQ|EQUAL, NEQ|NE|NOT EQUAL
- [NOT] IN
 - nRun IN (1234, 5678, 2468)
- [NOT] NULL
 - runQuality IS NOT NULL
- [NOT] EXISTS
 - runQuality EXISTS
- AND | “&&”
 - (runQuality NEQ 'GOOD' AND nRun EQ 1234)
- OR | “||”
 - (createDate LT d'2013-10-14T17:57Z' OR nRun GT 9876)

Searching - Plugins

A searching plugin will know how to relate to a Data Catalog object to extra, experiment-specific data (which is usually another table).

For example, when given the search expression:

```
exp.quality IN ('GOOD', 'GOLDEN') AND exp.runType EQ 'data'
```

The following will happen:

- The search parser understands to get the “exp” plugin
- The parser hands the plugin the current SQL statement, and the expression `quality IN ('GOOD', 'GOLDEN')`
 - Plugin performs any necessary setup to the SQL statement
 - Plugin returns a SQL expression to parser, if necessary
- Repeat the process for `runType EQ 'data'`

Security and Authentication

Security is handled two different ways

1. Session/Secure Cookies and Single Sign-on

- Ideal when requests originate from web app/browsers
- Possible to use via a script, but requires login every 8 hours

2. Hash Message Authentication Code (HMAC)

- Relies on secret private key to generate a cryptographic hash of the HTTP request, ideal for batch jobs
- Tools supported natively in many languages (Java, Python, etc...)
- Heavily used in industry

Modernizing our web application

- Move towards ajax/HTML5 support and frameworks
 - jQuery for general Javascript
 - Bootstrap for CSS
- Move to lazy loading for things that take a long time
 - Getting the # of files in 30 groups, each with 20k+ datasets, takes a while. Getting the byte size, event count, etc... longer.
 - So, let's load what we can in ~5 seconds, let the web app finish the rest asynchronously

HTML5 – New Layout with Bootstrap



Fermi LAT Datacatalog

Root / Data / Flight / Overlay / Overlay-maker4

Paginate?

Name			
Livetime	3	0	0
McIlwainL	3	0	0
McIlwainL_Empty	0	2	0
MERIT			2430
OVERLAY			2430

Container Browser

Datasets Browser

Container Status

Version: . Help

Info for group /Data/Flight/Overlay/Overlay-maker4 / MERIT

Datasets	Total Size	Events
2,430	5.9 GB	21,182,240

List of datasets for /Data/Flight/Overlay/Overlay-maker4 / MERIT

Name	Type	Size	Created
277377731-000000	MERIT	2.1 MB	2010-05-06 14:41 PDT
277395696-000003	MERIT	2.3 MB	2010-05-06 14:47 PDT
277389718-000002	MERIT	2.3 MB	2010-05-06 14:45 PDT
277613579-000041	MERIT	1.4 MB	2010-05-11 09:47 PDT
277445212-000012	MERIT	2.7 MB	2010-05-11 10:17 PDT
277699512-000056	MERIT	1.5 MB	2010-05-11 10:47 PDT
278037513-000115	MERIT	1.4 MB	2010-05-11 10:52 PDT
277905666-000092	MERIT	2.3 MB	2010-05-11 11:08 PDT
278213578-000146	MERIT	2.4 MB	2010-05-11 11:11 PDT
278255407-000153	MERIT	2.6 MB	2010-05-11 11:17 PDT
278465524-000190	MERIT	2.4 MB	2010-05-11 11:21 PDT
279028606-000290	MERIT	2.9 MB	2010-05-11 12:03 PDT
278403680-000179	MERIT	2.2 MB	2010-05-11 12:05 PDT
277955445-000101	MERIT	2.5 MB	2010-05-11 12:12 PDT

HTML5 – Lazy and On-Demand loading



Fermi LAT Datacatalog

Root / Data / Flight / Level1 / LPA

Paginate?

Name		
ACDPEDSALARM		27933
ACDPEDSANALYZER		29427
ACDPLOTS		29427
CAL		29427
CALGAINSSALARM		27933
CALGAINSANALYZER		29427
CALHIST		29427
CALHISTALARM		28593
CALPEDSALARM		27933
CALPEDSANALYZER		29427
CALTREND		29427
COMPAREDFM		17131
DIGI		29427
DIGIGAP		23196
DIGIHIST		29427
DIGIHISTALARM		29427
DIGITREND		29427
DIGITRENDALARM		
ELECTRONFT1		
ELECTRONFT1BADGTI		
ELECTRONMERIT		
EXTENDEDFT1		12296
EXTENDEDLS1		12296

List of datasets for /

Name	Type	Size	Created
------	------	------	---------

← Lazy/Async Loading of Status

HTML5 – Responsive Design (thanks to Bootstrap)

Bootstrap gives us the additional feature of a fluid layout.

The web app can transform its presentation depending on device or screen size.



Fermi LAT Datacatalog

Version: . Help

Root / Data / Flight / Overlay / Overlay-maker4

Paginate?

	Name			
	Livetime	3	0	0
	McIlwainL	3	0	0
	McIlwainL_Empty	0	2	0
	MERIT			2430
	OVERLAY			2430

Layout is stacked on narrower screen/viewport

Collapsible List

Info for group /Data/Flight/Overlay/Overlay-maker4 / MERIT

Datasets	Total Size	Events
2,430	5.9 GB	21,182,240

List of datasets for /Data/Flight/Overlay/Overlay-maker4 / MERIT

Conclusion

- The Data Catalog we've made for Fermi-LAT is very successful and used heavily across several experiments
- Generalizing the core and adding a plugin architecture allows us to support the needs of different experiments going forward
- RESTful interfaces are a good way to support many languages without repeating yourself, and a good way to support off-site use of our Data Catalog
- RESTful interfaces also make it easy to create responsive and adaptable human interfaces for data browsing and retrieval
- HTML5 technologies coupled with RESTful interfaces can enable us to create responsive web pages that scale

A bit more...

We are looking for other experiments who may be interested in using this!

- Current version has been in production for 5 years
- The complete rewrite is still 3-6 months out.
- The REST api has is used by EXO in production.
- Optimized for use with Oracle, but can work in MySQL, PostgreSQL
- More work is planned for the Download Manager
 - Optimized for xrootd at the moment

Informal Performance/benchmark notes:

- /path resource handled 600 concurrent connections at 1500 req/sec
- 4x2.6GHZ 16GB VM
- A SQL composing framework allows us to generate dynamic, optimized SQL at query time.