DDM Troubleshooting Console

(quick evaluation of splunk for such)

Rob Gardner

University of Chicago

ATLAS DDM Workshop

January 26, 2007

Motivation

- A distributed system means:
  - distributed servers
  - distributed errors
  - distributed log files
  - distributed centers
  - distributed people & expertise
- When things go smoothly
  - looks like single, monolithic system, running swimmingly, not a cloud in the sky
- But when things go wrong:
  - one massive (often distributed) headache!

The error messages for SLAC were 'transfer timeout' from SLAC to BNL
I looked at some of those failed files at SLAC. They were all created at Jan 10. All exist in SLAC LRC. **We only have gridftp log of the last several days** so I can't really tell if the transfer ever happened. **If BNL has the DQ2 logs back to Jan 10, that may help.**

I also **saw file lookup error in BNLPANDA log**.
2007-01-24 16:39:03,002 INFO [29157] vuidresolver: query files in dataset for misal1_csc11.005100.JimmyWenu.digit.RDO.v12003107_tid004539_sub96 failed with DQ2 backend exception [Error getting file info [MySQL server has gone away]] For **which dataset**, you saw 'transfer timeout'?

For trouble shooting DQ2 problem, this is a bit **harder to generalize** as it could be caused by many reasons. The **starting place is to look at FTS log entry from DQ2**. If there is a FTS entry, it means that DQ2 at least is trying to move data. Then, the **problem can be the low level** storage/transfer issue (**gsiftp or dcache problem, myproxy, etc**.) If there is no FTS entry (such as no FTS ID), it means that either it can not issue FTS command (FTS problem), or it is not even trying FTS command (DQ2 site service problem). Again, the cloud monitoring with dq2ping dataset should help with dq2 callback (callbacks-error) to start debugging.

But, at the end, **T1 and T2 DQ2 site admin needs to communicate the detailed situation to solve most of the problem** as **no one has access to every service at every machine.**

From the **subscription log file,** the **MySQL error** should be from LRC@SLAC
since SLAC do not have these three files somehow.
From log in Jan 10, 3 files (indeed) missing in the subdataset. Mystery!

```
sh-2.05b$ grep
misal1_csc11.005100.JimmyWenu.digit.RDO.v120031
07_tid004539_sub96
subscriptions.log.14 | grep resolver
2007-01-10 12:23:44,896 INFO [21995] vuidresolver:
dataset
misal1_csc11.005100.JimmyWenu.digit.RDO.v120031
07_tid004539_sub96 empty
removed subscription
```

Hello.

Datasets not being transferred. Can someone **take a look?**

Premise

- Local troubleshooting done with variety of monitoring systems and sensor data - but situations in DDM operations are intrinsically non-local
- Invariably one greps logfiles for specific event stings from more than one host (by more than one person)
  - and then correlating with events from other sources, many times a log file. E.g.:
    - No jobs running at a site, pilots executing & exiting
    - That's because datasets for assigned jobs not available at the site
    - What happened data flow? Where's the bottleneck
    - Check local subscriptions log, FTS-progress log
    - Problem at the Tier1? Problem with central services database?
- Problem is that one doesn't always have access to those files
- And how to correlate that information with other sensors?

Setup

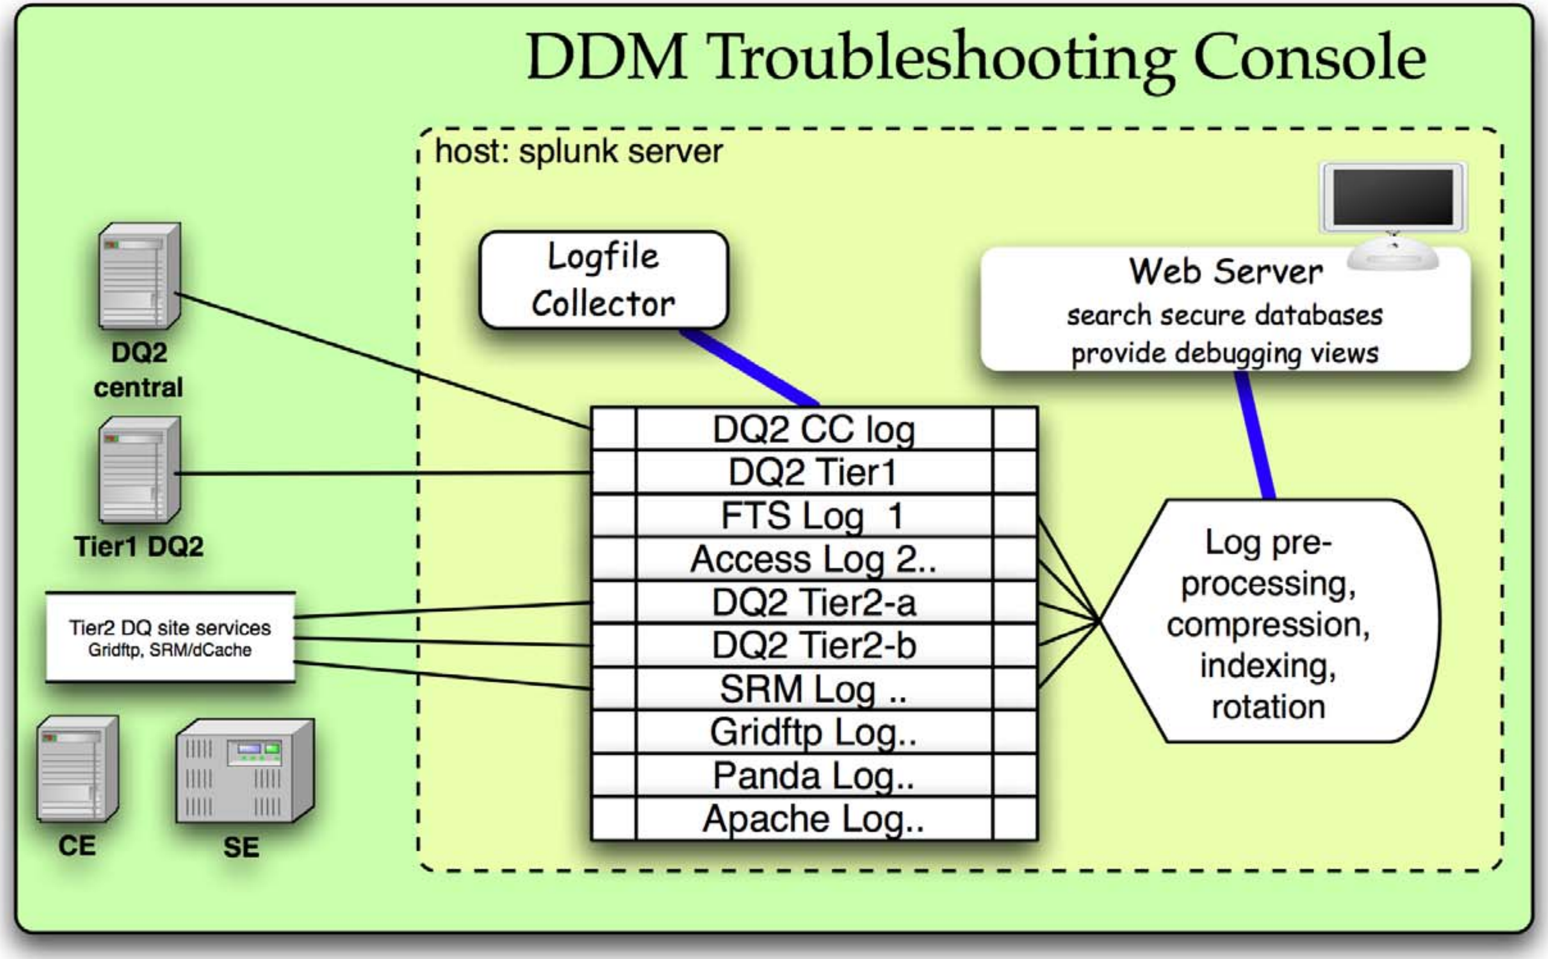
DDM Troubleshooting Console
host: splunk server
Logfile Collector
Web Server
search secure databases
provide debugging views
DQ2 central
Tier1 DQ2
Tier2 DQ site services
Gridftp, SRM/dCache
CE
SE
DQ2 CC log
DQ2 Tier1
FTS Log 1
Access Log 2..
DQ2 Tier2-a
DQ2 Tier2-b
SRM Log ..
Gridftp Log..
Panda Log..
Apache Log..
Log pre-processing, compression, indexing, rotation

Best of breed solutions?

- Research area in distributed computing
- ARDA console and new LCG monitoring group addressing similar issues
- Splunk-based system - focus on troubleshooting info
  - attractive features
    - indexed, searchable tags (Google-like) database
    - recognizes many industry message formats (Cisco, Unix syslog, etc)
    - company young, but seems viable
    - little development effort, mostly an integration
  - drawbacks
    - commercial product - must weigh license cost against human effort (development, and operational troubleshooting)

Example source types

- Dq2 log files
- SRM and Gridftp, Gram, MyProxy
- MySQL logs
- Job manager - PBS, Condor
- Panda service log files
- Pilot submitter logs
- Apache
- Network
- syslog

Example event types

- DQ2 related
  - Tier 2 Site service events
  - Tier 1 DQ2 server events
  - Tier 0 Central server events
- Backend site services
  - Authentication failures
  - Storage services, dCache monitors
  - File system crashes, kernel panics
  - Correlate with job manager messages
- Panda server
  - Correlate with production dataflow

One stop troubleshooting..(ideal)

- Repository of data sources and event types in one or set of indexed, searchable database
- Secure access for DDM operations team
- Create and store standard queries

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

splunk server - data sources, saved queries

- dq2 site service logs at Tier1 and Tier2
- panda server log
- tier2 condor job manager
- tier2 dq2 backend - gridftp

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Error search, since: NNN minutes ago

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

distribution of subscription events

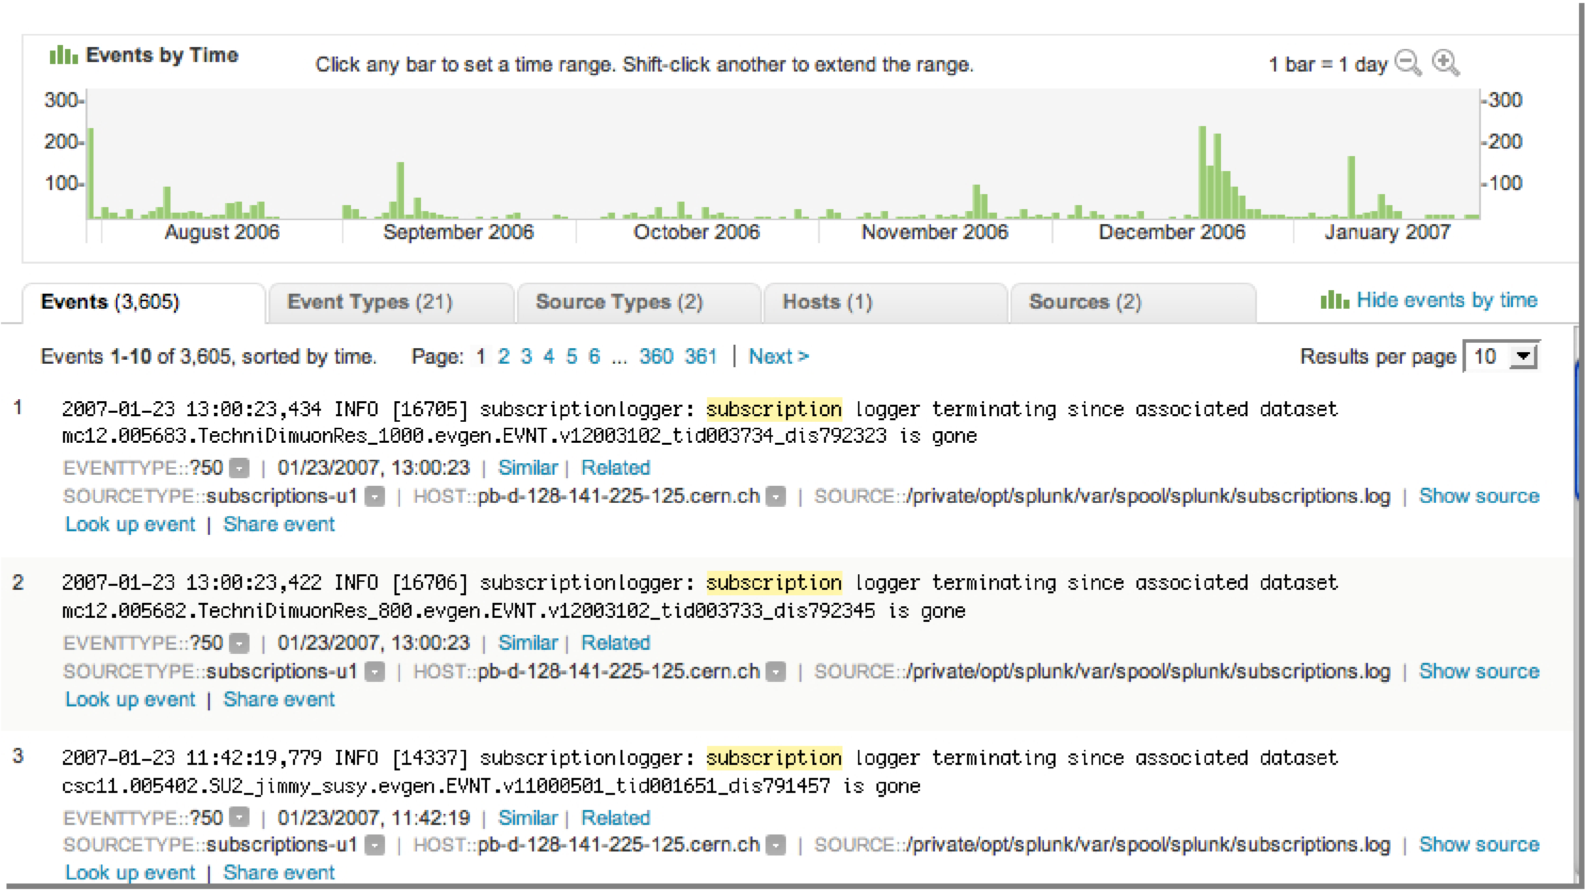

Events 1-10 of 3,605, sorted by time.

```
1  2007-01-23 13:00:23,434 INFO [16705] subscriptionlogger: subscription logger terminating since associated dataset
   mc12.005683.TechniDimuonRes_1000.evgen.EVNT.v12003102_tid003734_dis792323 is gone
```

EVENTTYPE::?50 | 01/23/2007, 13:00:23 | Similar | Related
SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-225-125.cern.ch | SOURCE::/private/opt/splunk/var/spool/splunk/subscriptions.log | Show source
Look up event | Share event

```
2  2007-01-23 13:00:23,422 INFO [16706] subscriptionlogger: subscription logger terminating since associated dataset
   mc12.005682.TechniDimuonRes_800.evgen.EVNT.v12003102_tid003733_dis792345 is gone
```

EVENTTYPE::?50 | 01/23/2007, 13:00:23 | Similar | Related
SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-225-125.cern.ch | SOURCE::/private/opt/splunk/var/spool/splunk/subscriptions.log | Show source
Look up event | Share event

```
3  2007-01-23 11:42:19,779 INFO [14337] subscriptionlogger: subscription logger terminating since associated dataset
   csc11.005402.SU2_jimmy_susy.evgen.EVNT.v11000501_tid001651_dis791457 is gone
```

EVENTTYPE::?50 | 01/23/2007, 11:42:19 | Similar | Related
SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-225-125.cern.ch | SOURCE::/private/opt/splunk/var/spool/splunk/subscriptions.log | Show source
Look up event | Share event

Filter by dataset

eg. csc11 pythia minbias subscription events

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

..and replicaregister complete

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

.and finished transfers on a given day

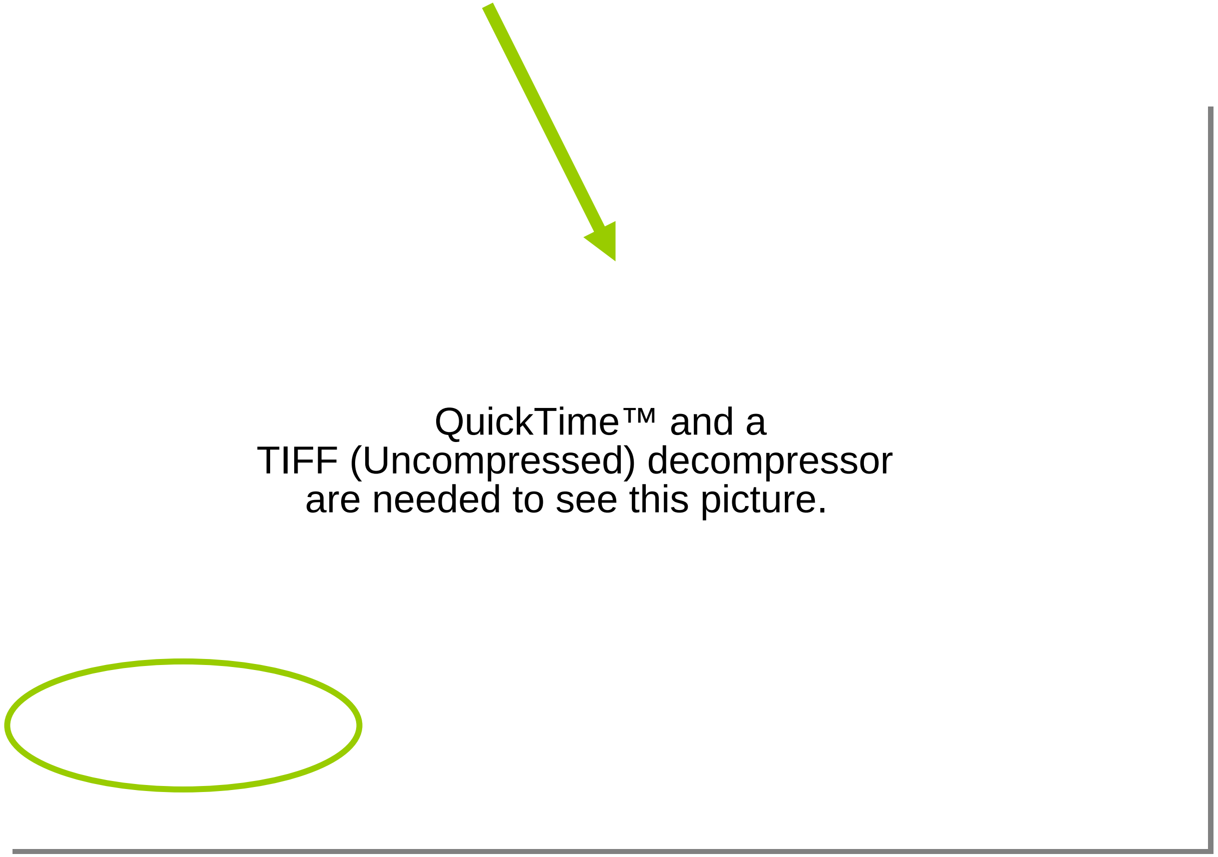

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

filter by subscription event type

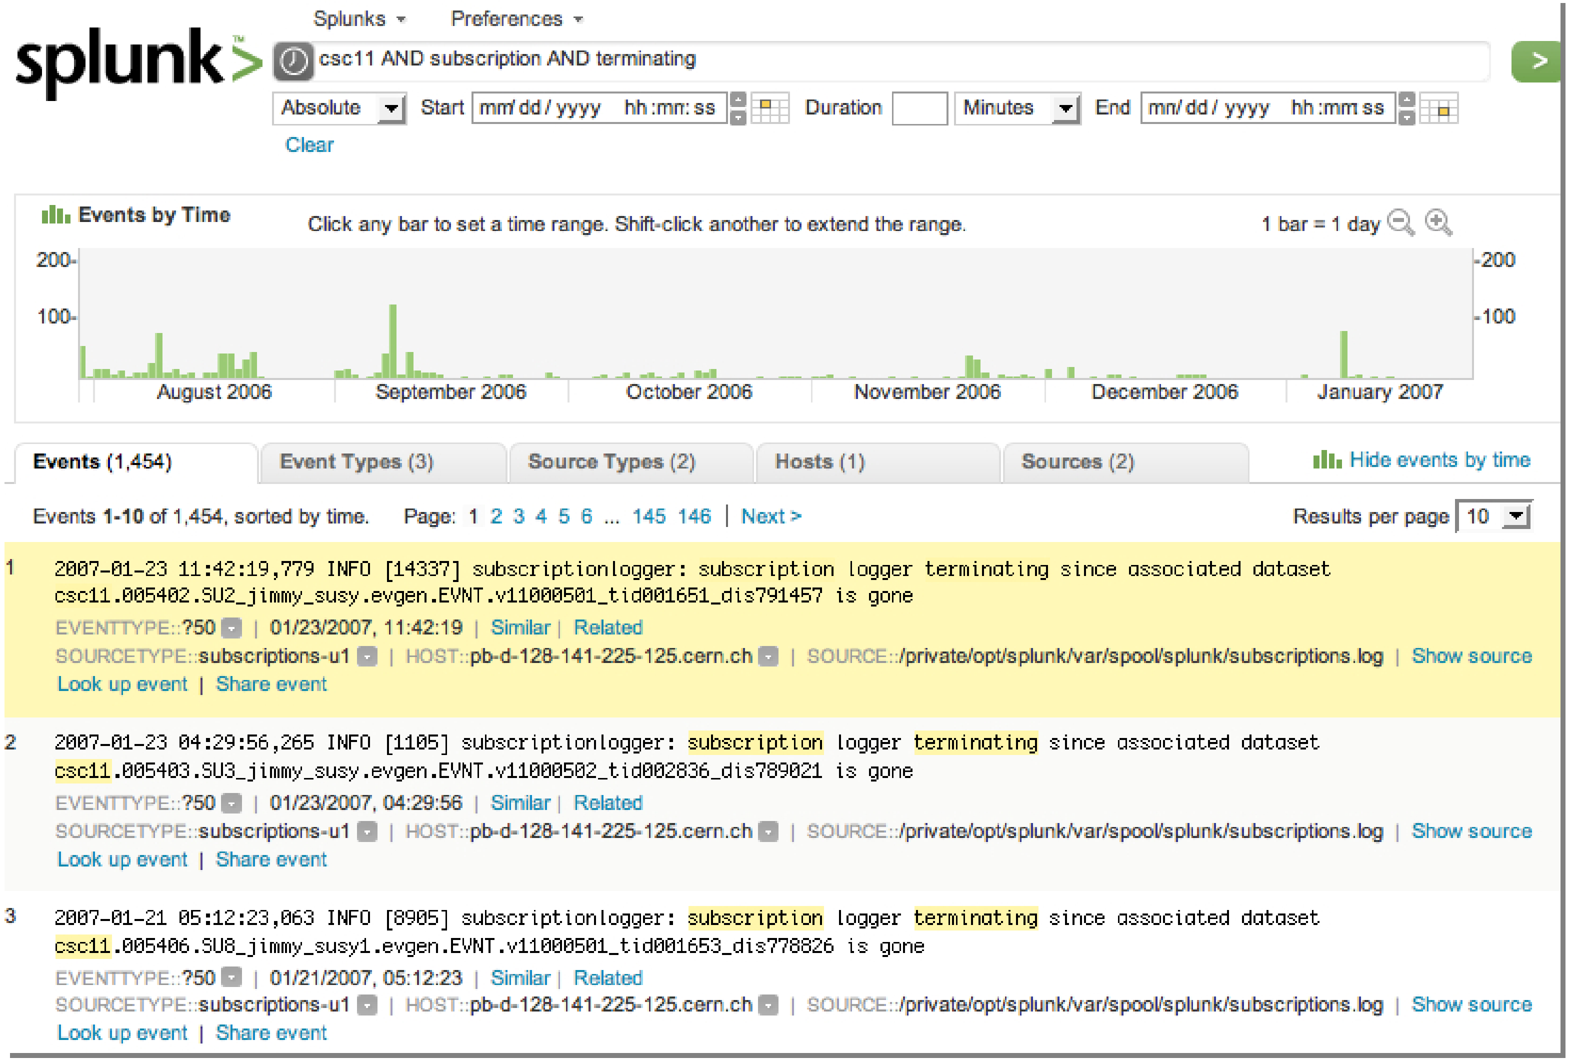

csc11 AND subscription AND terminating

Events 1-10 of 1,454, sorted by time. Page: 1 2 3 4 5 6 ... 145 146 | Next >

```
1  2007-01-23 11:42:19,779 INFO [14337] subscriptionlogger: subscription logger terminating since associated dataset
   csc11.005402.SU2_jimmy_susy.evgen.EVNT.v11000501_tid001651_dis791457 is gone
```

EVENTTYPE::?50 | 01/23/2007, 11:42:19 | Similar | Related
SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-225-125.cern.ch | SOURCE::/private/opt/splunk/var/spool/splunk/subscriptions.log | Show source
Look up event | Share event

```
2  2007-01-23 04:29:56,265 INFO [1105] subscriptionlogger: subscription logger terminating since associated dataset
   csc11.005403.SU3_jimmy_susy.evgen.EVNT.v11000502_tid002836_dis789021 is gone
```

EVENTTYPE::?50 | 01/23/2007, 04:29:56 | Similar | Related
SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-225-125.cern.ch | SOURCE::/private/opt/splunk/var/spool/splunk/subscriptions.log | Show source
Look up event | Share event

```
3  2007-01-21 05:12:23,063 INFO [8905] subscriptionlogger: subscription logger terminating since associated dataset
   csc11.005406.SU8_jimmy_susy1.evgen.EVNT.v11000501_tid001653_dis778826 is gone
```

EVENTTYPE::?50 | 01/21/2007, 05:12:23 | Similar | Related
SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-225-125.cern.ch | SOURCE::/private/opt/splunk/var/spool/splunk/subscriptions.log | Show source
Look up event | Share event

(csc11 AND subscription AND terminating)

..and find its context

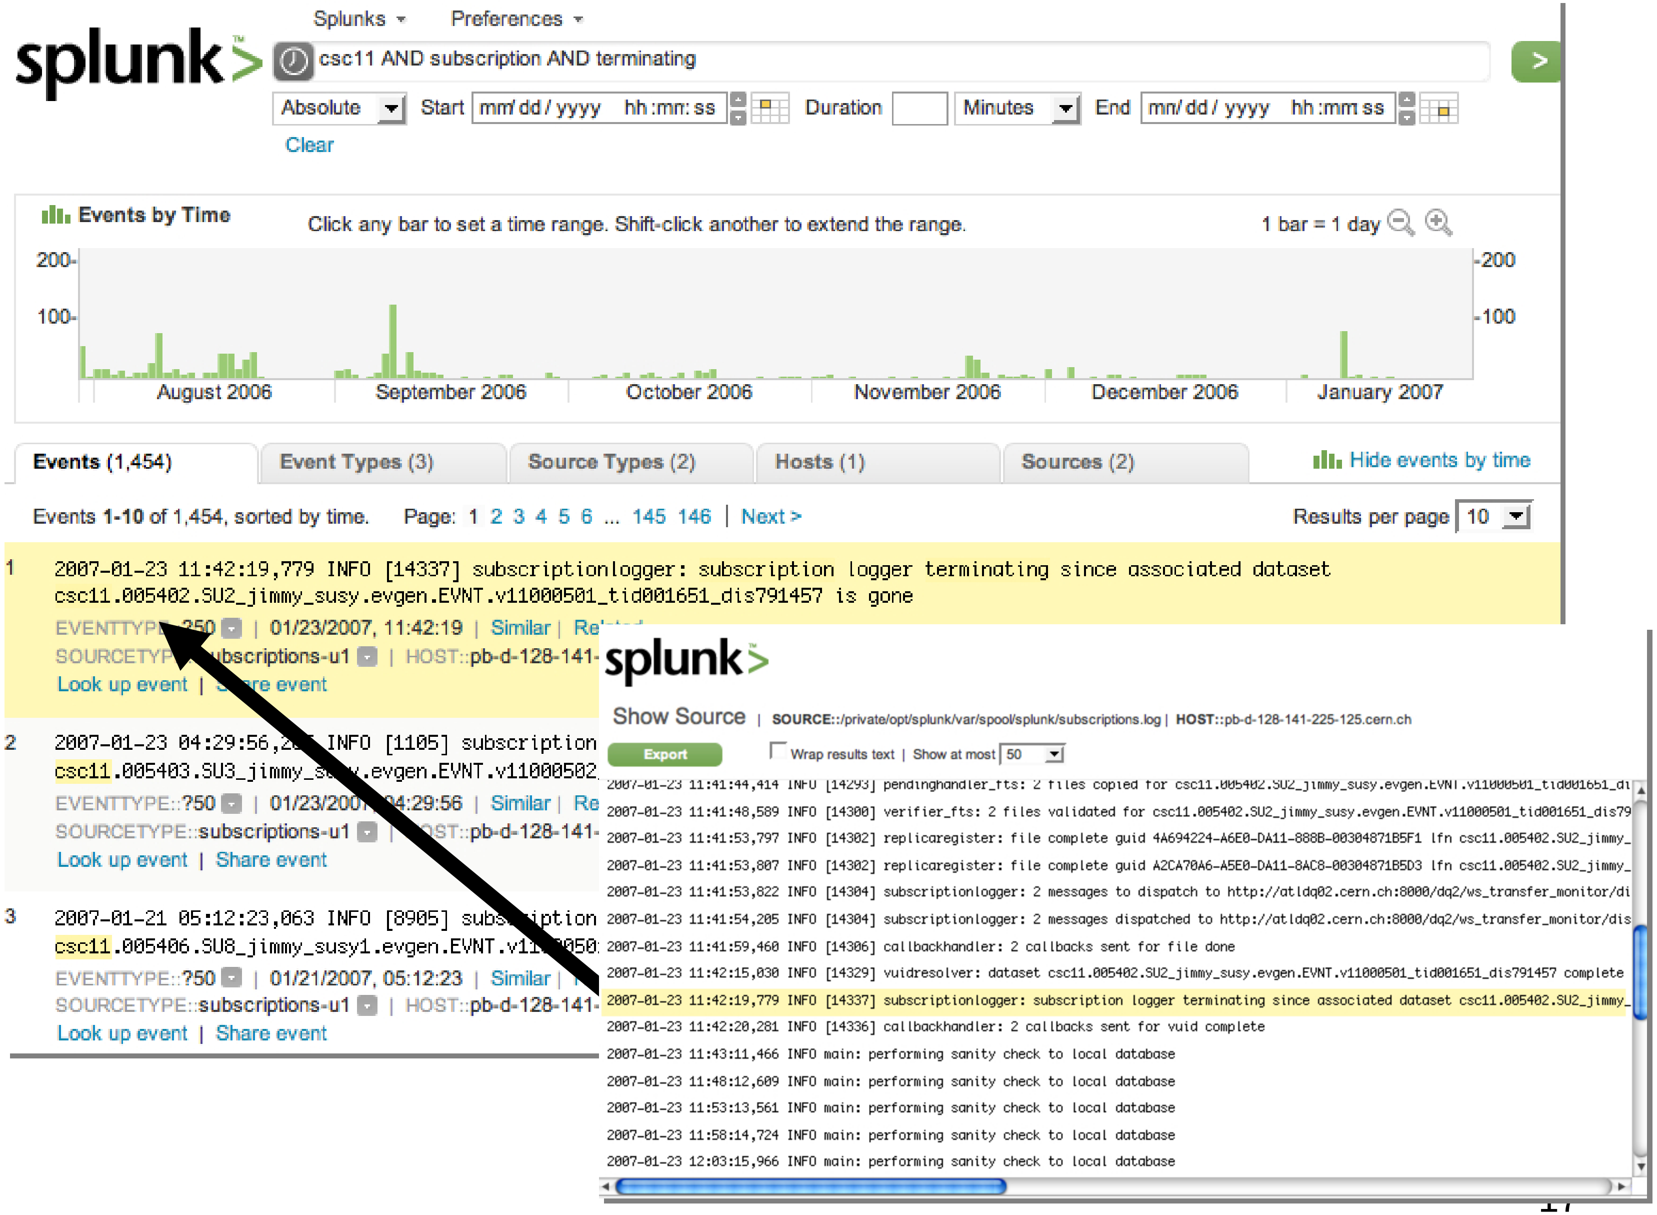

csc11 AND subscription AND terminating

Events 1-10 of 1,454, sorted by time.

```
1  2007-01-23 11:42:19,779 INFO [14337] subscriptionlogger: subscription logger terminating since associated dataset
   csc11.005402.SU2_jimmy_susy.evgen.EVNT.v11000501_tid001651_dis791457 is gone
   EVENTTYPE::?50 | 01/23/2007, 11:42:19 | Similar | Related
   SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-

2  2007-01-23 04:29:56,2 INFO [1105] subscription
   csc11.005403.SU3_jimmy_s .evgen.EVNT.v11000502
   EVENTTYPE::?50 | 01/23/2007, 04:29:56 | Similar | Re
   SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-

3  2007-01-21 05:12:23,063 INFO [8905] subscription
   csc11.005406.SU8_jimmy_susy1.evgen.EVNT.v11 050
   EVENTTYPE::?50 | 01/21/2007, 05:12:23 | Similar |
   SOURCETYPE::subscriptions-u1 | HOST::pb-d-128-141-
```

SOURCE::/private/opt/splunk/var/spool/splunk/subscriptions.log | HOST::pb-d-128-141-225-125.cern.ch

```
2007-01-23 11:41:44,414 INFO [14293] pendinghandler_fts: 2 files copied for csc11.005402.SU2_jimmy_susy.evgen.EVNT.v11000501_tid001651_di
2007-01-23 11:41:48,589 INFO [14300] verifier_fts: 2 files validated for csc11.005402.SU2_jimmy_susy.evgen.EVNT.v11000501_tid001651_dis79
2007-01-23 11:41:53,797 INFO [14302] replicaregister: file complete guid 4A694224-A6E0-DA11-888B-00304871B5F1 lfn csc11.005402.SU2_jimmy_
2007-01-23 11:41:53,807 INFO [14302] replicaregister: file complete guid A2CA70A6-A5E0-DA11-8AC8-00304871B5D3 lfn csc11.005402.SU2_jimmy_
2007-01-23 11:41:53,822 INFO [14304] subscriptionlogger: 2 messages to dispatch to http://atldq02.cern.ch:8000/dq2/ws_transfer_monitor/di
2007-01-23 11:41:54,205 INFO [14304] subscriptionlogger: 2 messages dispatched to http://atldq02.cern.ch:8000/dq2/ws_transfer_monitor/dis
2007-01-23 11:41:59,460 INFO [14306] callbackhandler: 2 callbacks sent for file done
2007-01-23 11:42:15,030 INFO [14329] vuidresolver: dataset csc11.005402.SU2_jimmy_susy.evgen.EVNT.v11000501_tid001651_dis791457 complete
2007-01-23 11:42:19,779 INFO [14337] subscriptionlogger: subscription logger terminating since associated dataset csc11.005402.SU2_jimmy_
2007-01-23 11:42:20,281 INFO [14336] callbackhandler: 2 callbacks sent for vuid complete
2007-01-23 11:43:11,466 INFO main: performing sanity check to local database
2007-01-23 11:48:12,609 INFO main: performing sanity check to local database
2007-01-23 11:53:13,561 INFO main: performing sanity check to local database
2007-01-23 11:58:14,724 INFO main: performing sanity check to local database
2007-01-23 12:03:15,966 INFO main: performing sanity check to local database
```

Other features

- Filtering through boolean queries on indexed tags
- Zoom on time windows to localize events
- Link splunk servers together

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Development issues

- Log file identification decisions
  - which, where, how often: management issues
  - crude estimates of storage requirements and methods
  - pre-processing for optimal search formats, storage economy
- Configuration for event types
  - on splunk side, create event-type configurations
- Access and collection
  - dq2 servers run http, others may be available via g-u-c
  - collector on troubleshooting host; cache area for uploads
  - indexes quickly: grab, push to cache, splunk automatically consumes, handles .gz okay
- Analysis
  - timestamp checking, event type definition and correlation, export formats and reporting, documentation-sharing
    - **(e.g.. splunk --> wiki)**

Next Steps

- In the next month - do a feasibility study on OSG
- See if this thing is useful or not; cross-checks needed to establish integrity. operational experience - does it fit?, stable?
- Write event-type customizations for:
  - dq2 site services 0.2 within a cloud, tier1 and tier2
  - dq2 central services 0.2
  - check out the new dq2 logs for 0.3
  - Panda
  - Condor: job manager and pilot submitter
  - Gram, Gridftp
  - MySQL
- Devise simple aggregation, pre-processing system
- Create “splunks” - specific queries into the collection
- Try out in US ATLAS cloud - results by March

DDM documentation

- http://indico.cern.ch/materialDisplay.py?contribId=95&amp;sessionId=5&amp;materialId=0&amp;confId=3738