Controls

EPICS Archiver Appliance
Fall 2013

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

LELAND STANFORD JUNIOR UNIVERSITY

Our top 5 Objectives

Controls

- Scale to 1-2 millions PV's
- Fast data retrieval
- Users add PV's to archiver
- Zero oversight
- Flexible configurations on a per PV basis

Murali Shankar
Luofeng Li

Components

Controls

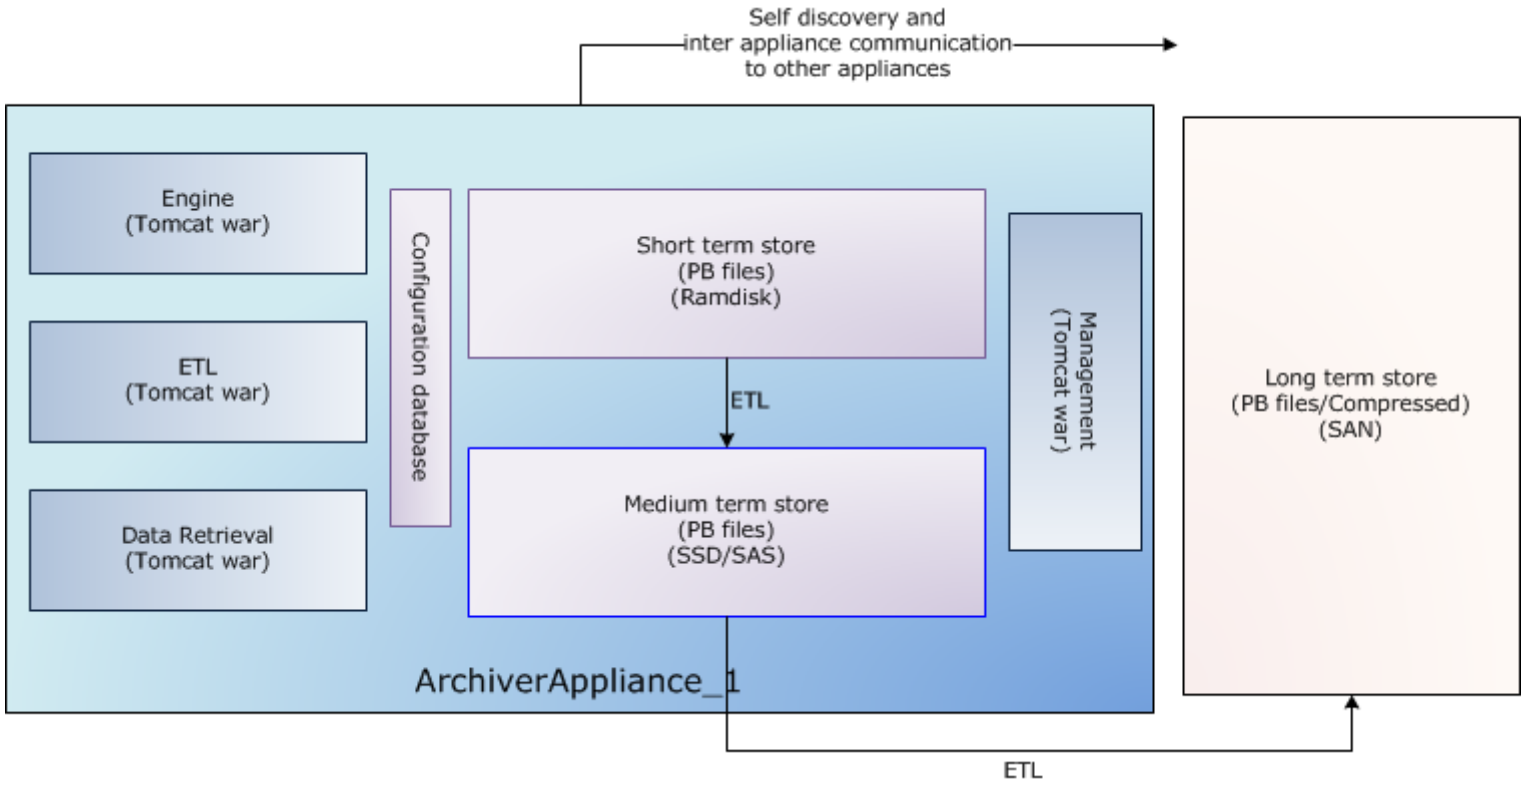

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

Scale by clustering appliances

Controls

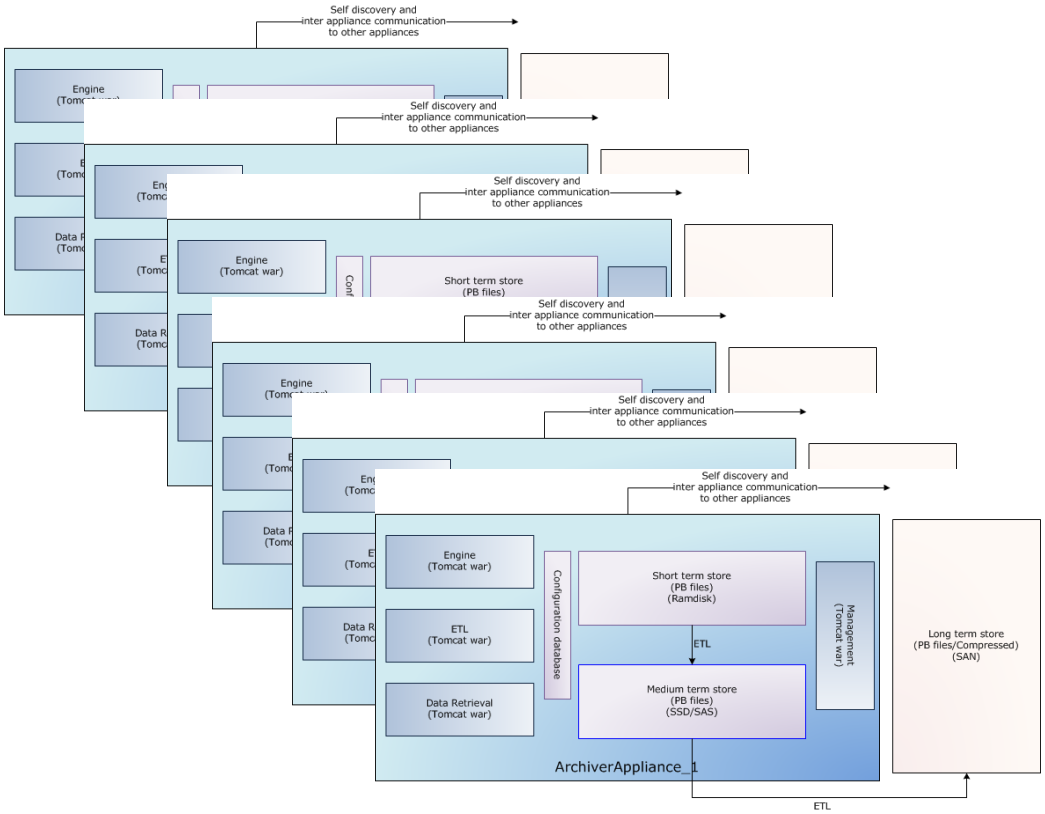

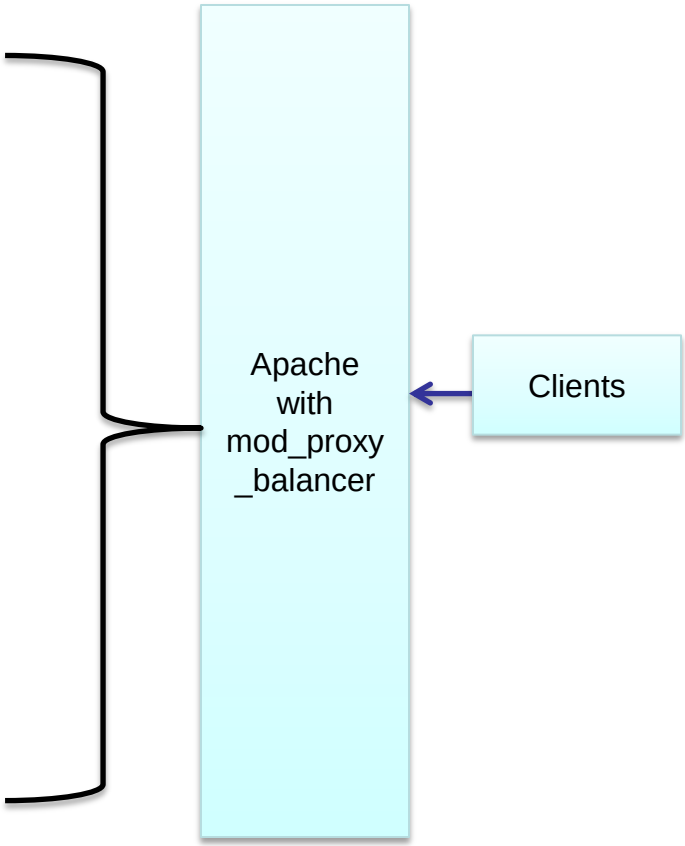

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

LELAND STANFORD JUNIOR UNIVERSITY

Status

Controls

- In production
  - In TestFac since May
  - In FACET since June
  - Testing in LCLS since end of July
    - Running in parallel
    - Should wrap up testing soon.
- Targeted at machine physicists and operators

TestFAC

Controls

Ran it for a month without restarting/upgrading

EPICS Archiver Appliance for Test Facilities

Home Reports Metrics Storage Appliances Integration

25 | Page 1 of 1

Instance Name	Status	PV Count	Connected	Mgmt Uptime
testfac-archapp	Working	27489	27182	1 month4 days 2:14:52

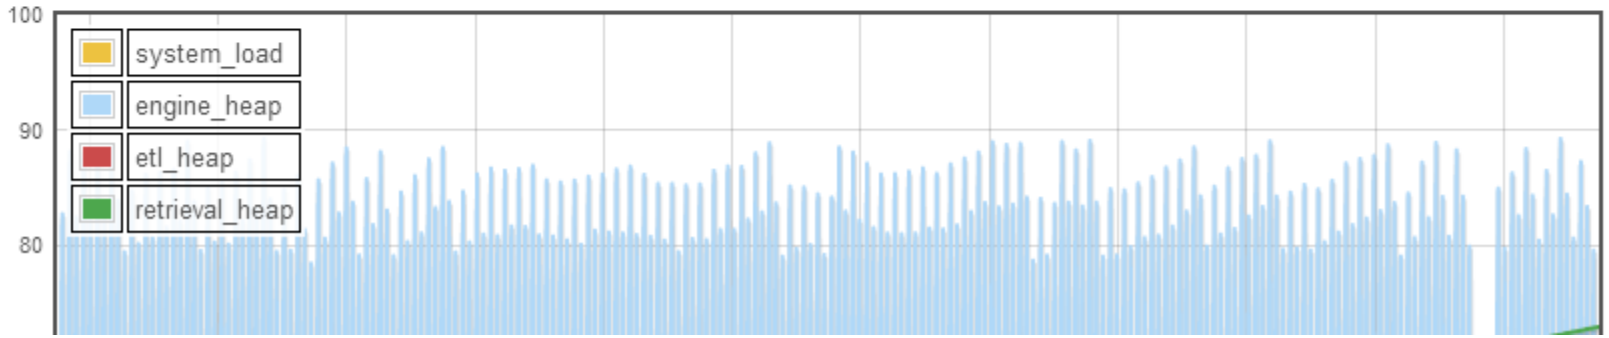

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

TestFAC

Controls

EPICS Archiver Appliance for Test Facilities

Home | Reports | Metrics | Storage | Appliances | Integration

25 | Page 1 of 1

Instance Name	Status	PV Count	Connected	Event Rate	Data Rate (GB/day)	Engine write thread(s)	Max ETL(%)
testfac-archapp	Working	27489	27463	296.85	0.58	0.12	0

Total number of ETL(0»1) runs so far	191
Average time spent in ETL(0»1) (s/run)	10

Total number of ETL(1»2) runs so far	24
Average time spent in ETL(1»2) (s/run)	26

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

FACET

Controls

EPICS Archiver Appliance for FACET

Home Reports Metrics Storage Appliances Integration

25 | Page 1 of 1

Instance Name	Status	PV Count	Connected	Event Rate	Data Rate (GB/day)	Engine write thread(s)	Max ETL(%)
facet-archapp	Working	27266	26492	211.53	1.27	0.11	0

Engine writes every 10 seconds

Total number of ETL(0»1) runs so far	703
Average time spent in ETL(0»1) (s/run)	20

ETL 0->1 runs once/hour

Total number of ETL(1»2) runs so far	87
Average time spent in ETL(1»2) (s/run)	97

ETL 1->2 runs once/day*

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

LCLS

EPICS Archiver Appliance for LCLS

Home Reports Metrics Storage Appliances Integration

25 Page 1 of 1

Instance Name	Status	PV Count	Connected	Event Rate	Data Rate (GB/day)	Engine write thread(s)	Max ETL(%)
lcls-archapp01	Working	38166	37910	2,682.2	9.83	0.53	1
lcls-archapp02	Working	48574	48347	2,974.83	6.76	0.51	2
lcls-archapp03	Working	60147	59611	3,222.1	10.62	0.78	3

Here are the some detailed metrics of the appliance **lcls-archapp01**

Murali Shankar
Luofeng Li

LCLS channel count

Controls

Connected PV count	37910
Total channels	326486

Connected PV count	48347
Total channels	413974

Connected PV count	59611
Total channels	513142

- PVs being archived = 145,949
- We plan to automatically archive (when available)
  - HIHI,HIGH,LOW,LOLO,LOPR,HOPR,DRVH,DRVL
- Total EPICS fields we are archiving = 1,253,602
- This also shows up on the IOC in casr()

```
sioc-sys1-ml00>casr
Channel Access Server V4.13
Connected circuits:
TCP 134.79.151.21:58258(lcls-prod01): User="laci", V4.13, 26030 Channels, Priority=0
TCP 172.27.72.23:59458(facet-srv02): User="fphysics", V4.11, 119 Channels, Priority=0
```

Typical ETL performance

Controls

Total number of ETL(0»1) runs so far	114
Average time spent in ETL(0»1) (s/run)	66
Average percentage of time spent in ETL(0»1)	1.82
Avg time spent by getETLStreams() in ETL(0»1) (s/run)	1
Avg time spent by prepareForNewPartition() in ETL(0»1) (s/run)	0
Avg time spent by appendToETLAppendData() in ETL(0»1) (s/run)	5
Avg time spent by commitETLAppendData() in ETL(0»1) (s/run)	0
Avg time spent by markForDeletion() in ETL(0»1) (s/run)	0
Avg time spent by runPostProcessors() in ETL(0»1) (s/run)	0
Avg time spent by executePostETLTasks() in ETL(0»1) (s/run)	0
Estimated bytes transferred in ETL (0»1)(GB)	46.46

Total number of ETL(1»2) runs so far	15
Average time spent in ETL(1»2) (s/run)	410
Average percentage of time spent in ETL(1»2)	1.48
Avg time spent by getETLStreams() in ETL(1»2) (s/run)	7
Avg time spent by prepareForNewPartition() in ETL(1»2) (s/run)	0
Avg time spent by appendToETLAppendData() in ETL(1»2) (s/run)	351
Avg time spent by commitETLAppendData() in ETL(1»2) (s/run)	0
Avg time spent by markForDeletion() in ETL(1»2) (s/run)	1
Avg time spent by runPostProcessors() in ETL(1»2) (s/run)	0
Avg time spent by executePostETLTasks() in ETL(1»2) (s/run)	0
Estimated bytes transferred in ETL (1»2)(GB)	40.15

Murali Shankar
Luofeng Li

System performance

Controls

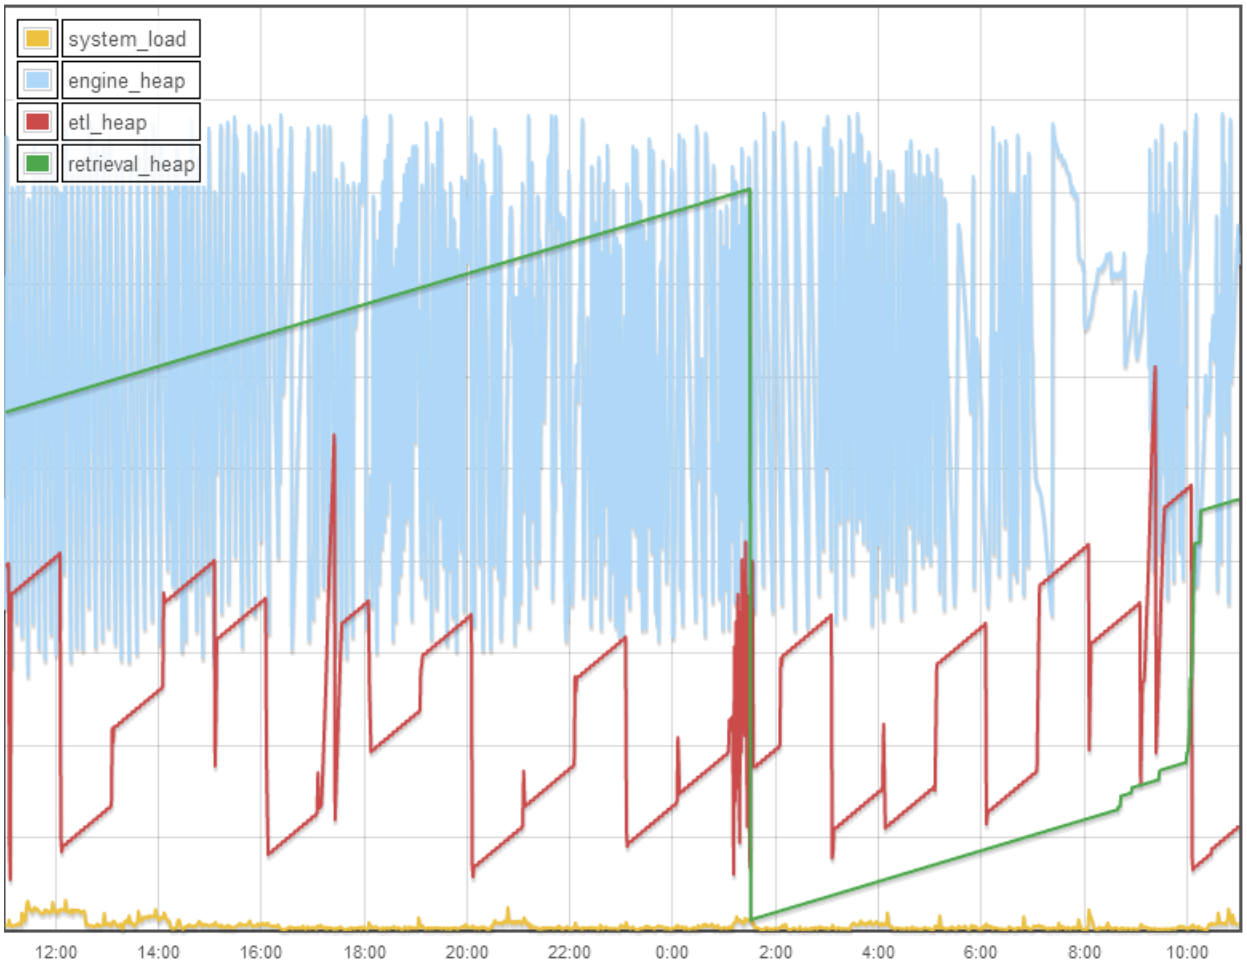

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

LELAND STANFORD JUNIOR UNIVERSITY
DIE LUFT DER FREIHEIT WEHT
1891

LCLS deployment

Controls

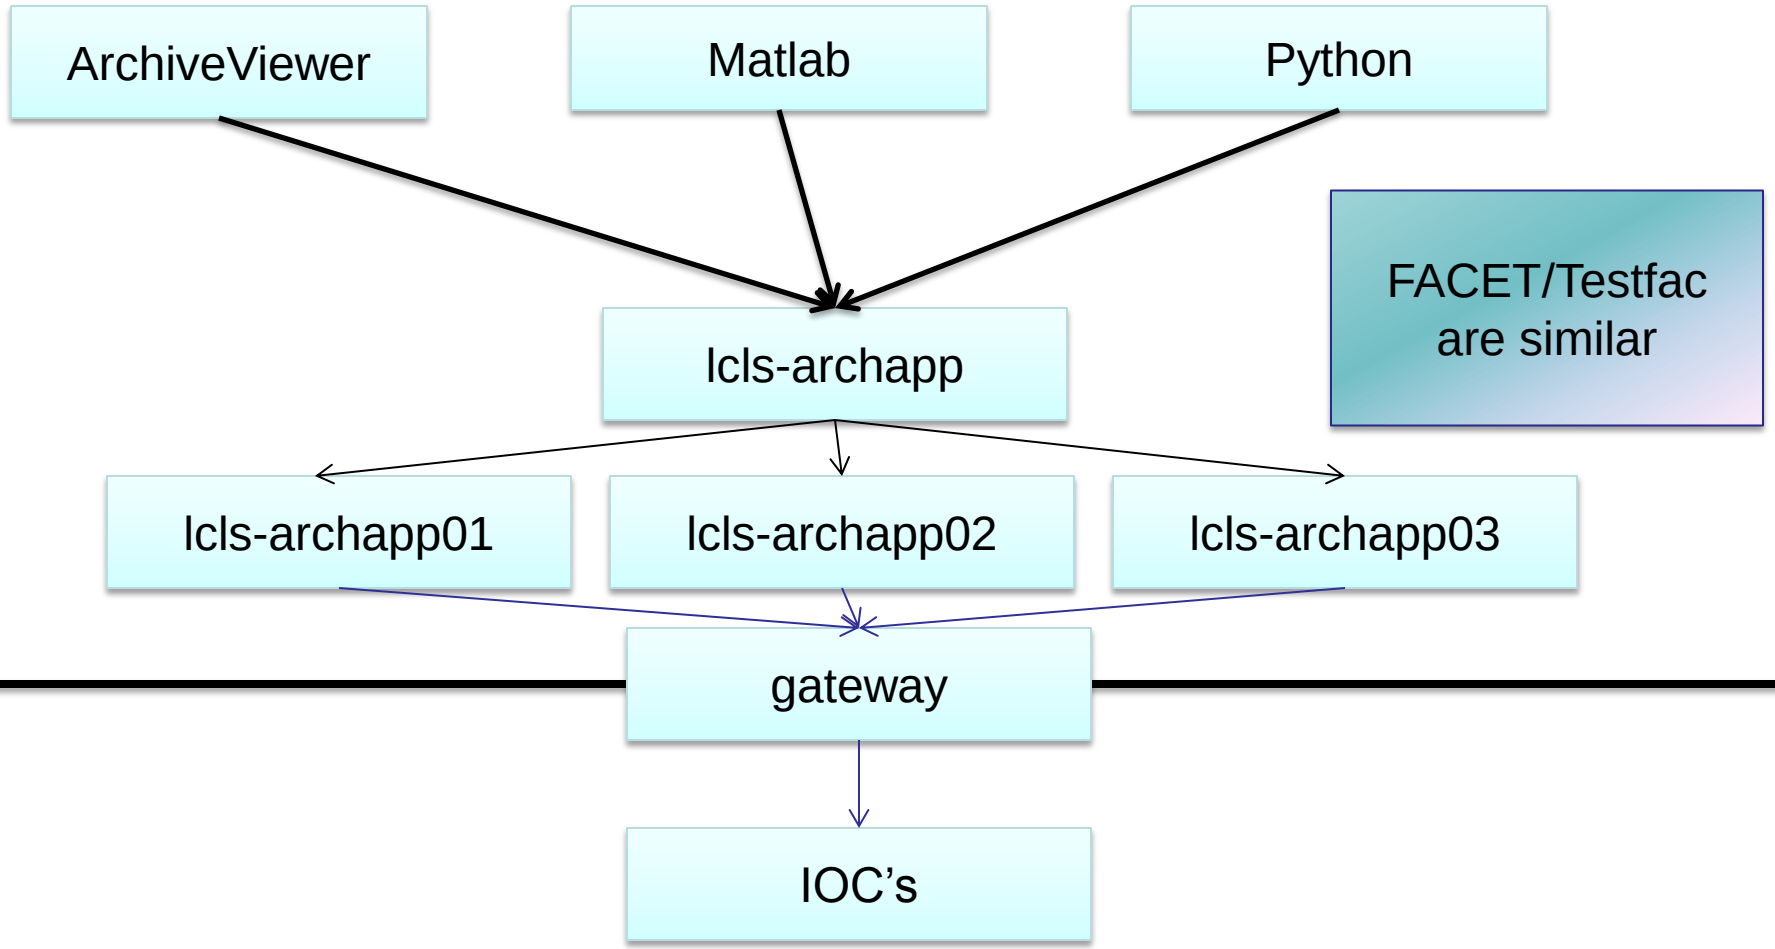

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

Clients

Controls

- ArchiveViewer
  - Minimal changes – same client talks to ChannelArchiver and appliance.
  - Some changes to support export
- CSS Databrowser
- Matlab
- Python

Minimal changes

Controls

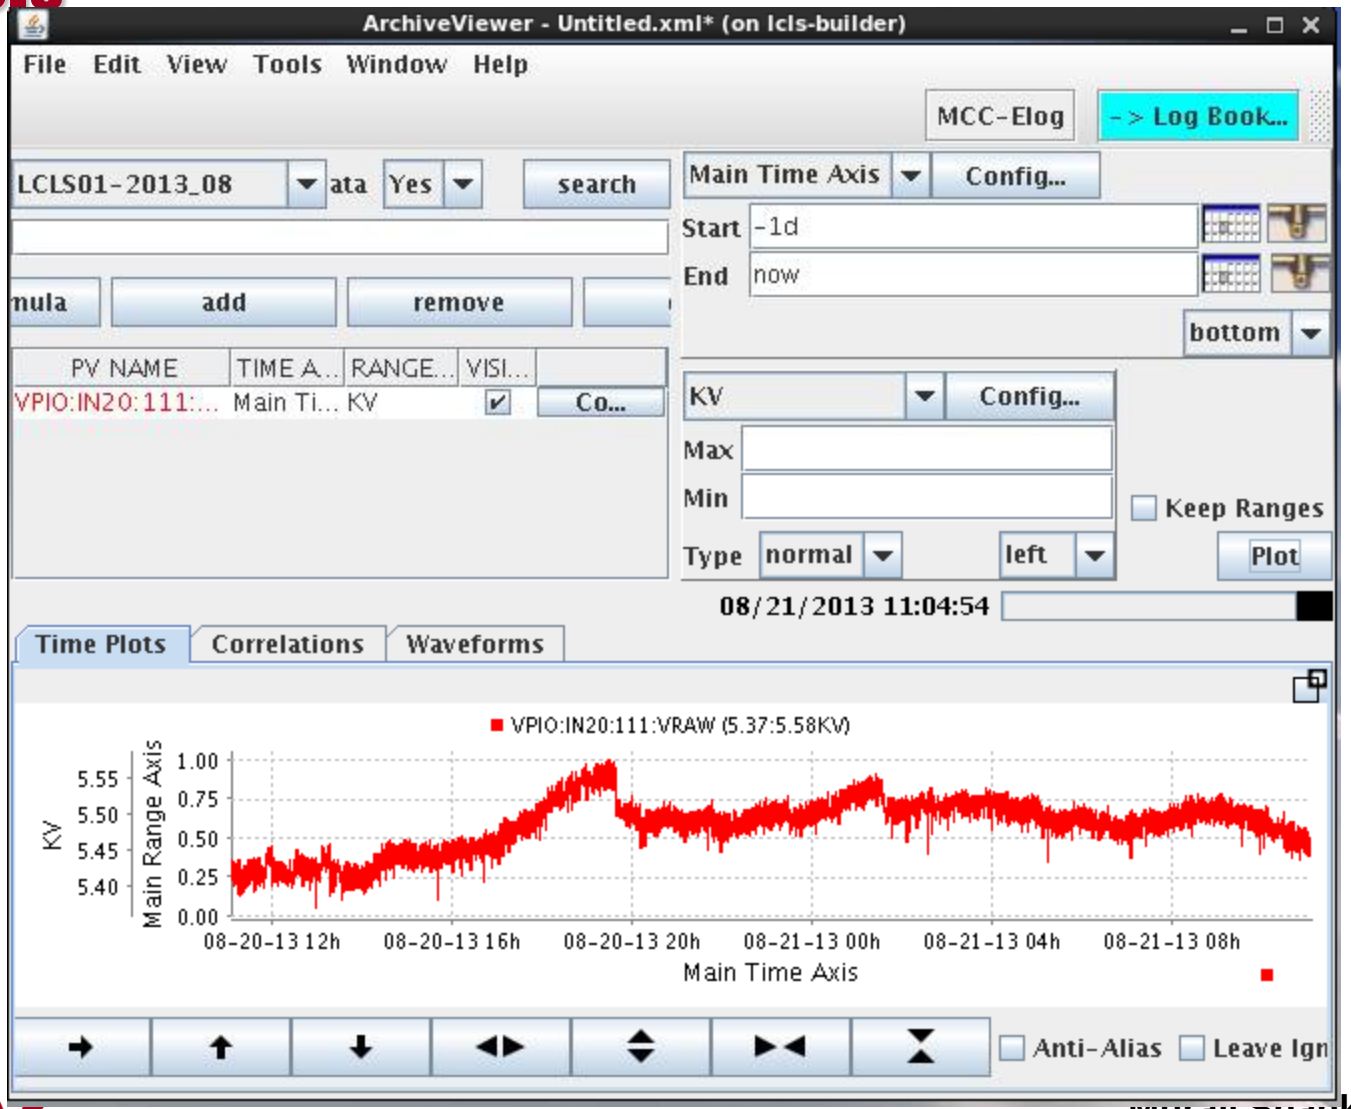

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

Minimal changes

Controls

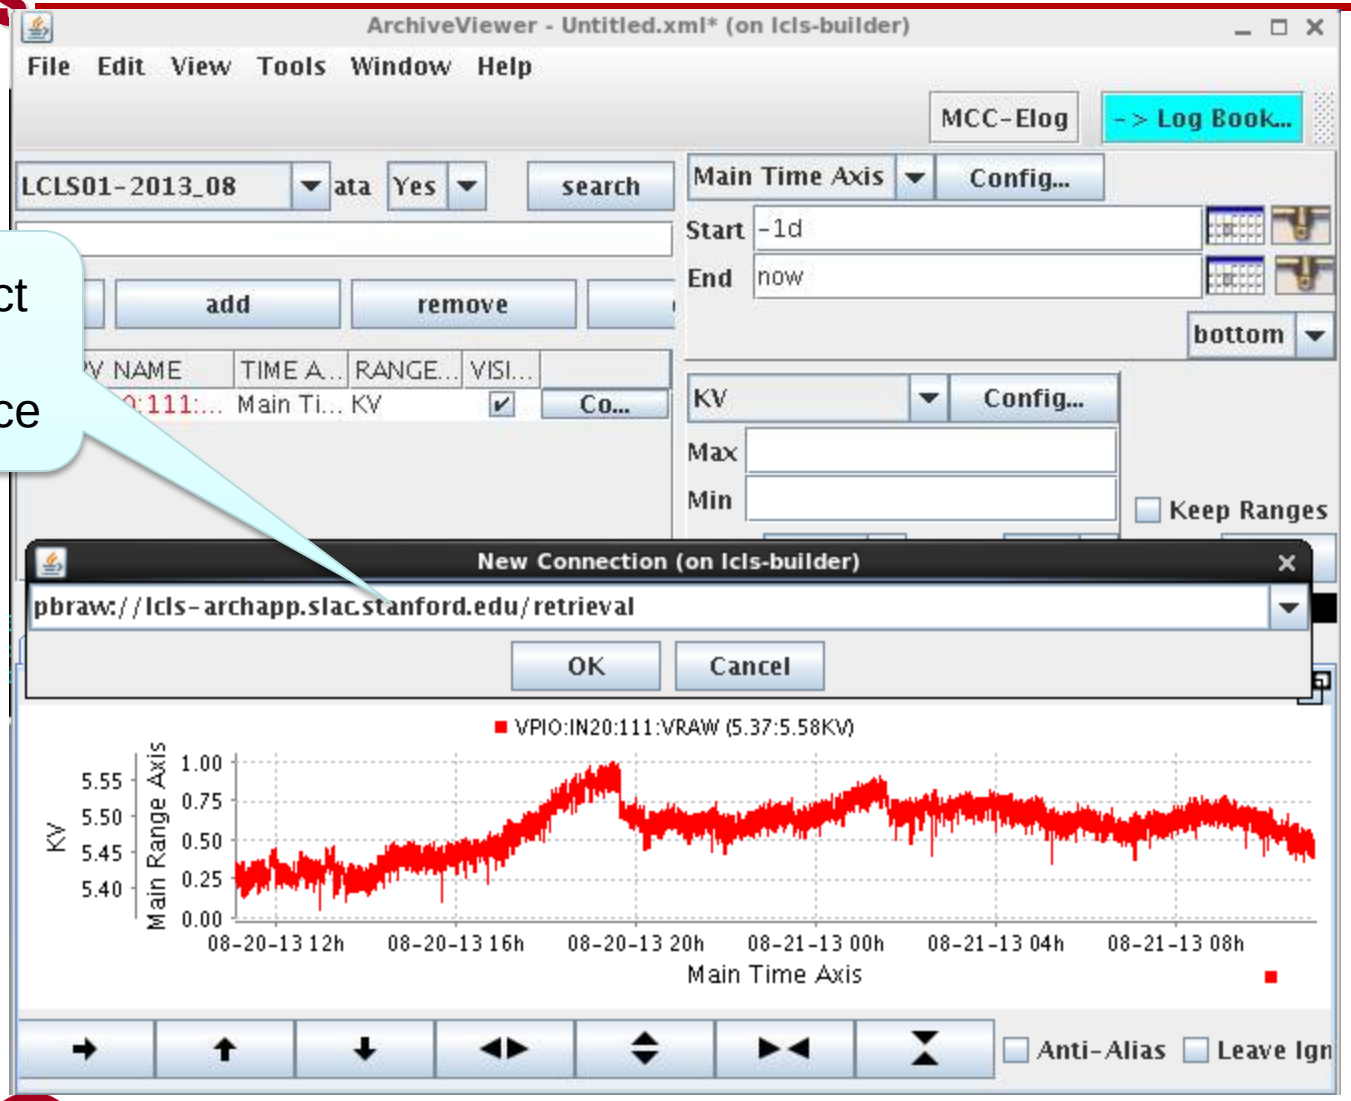

Connect to appliance

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

Minimal changes

Controls

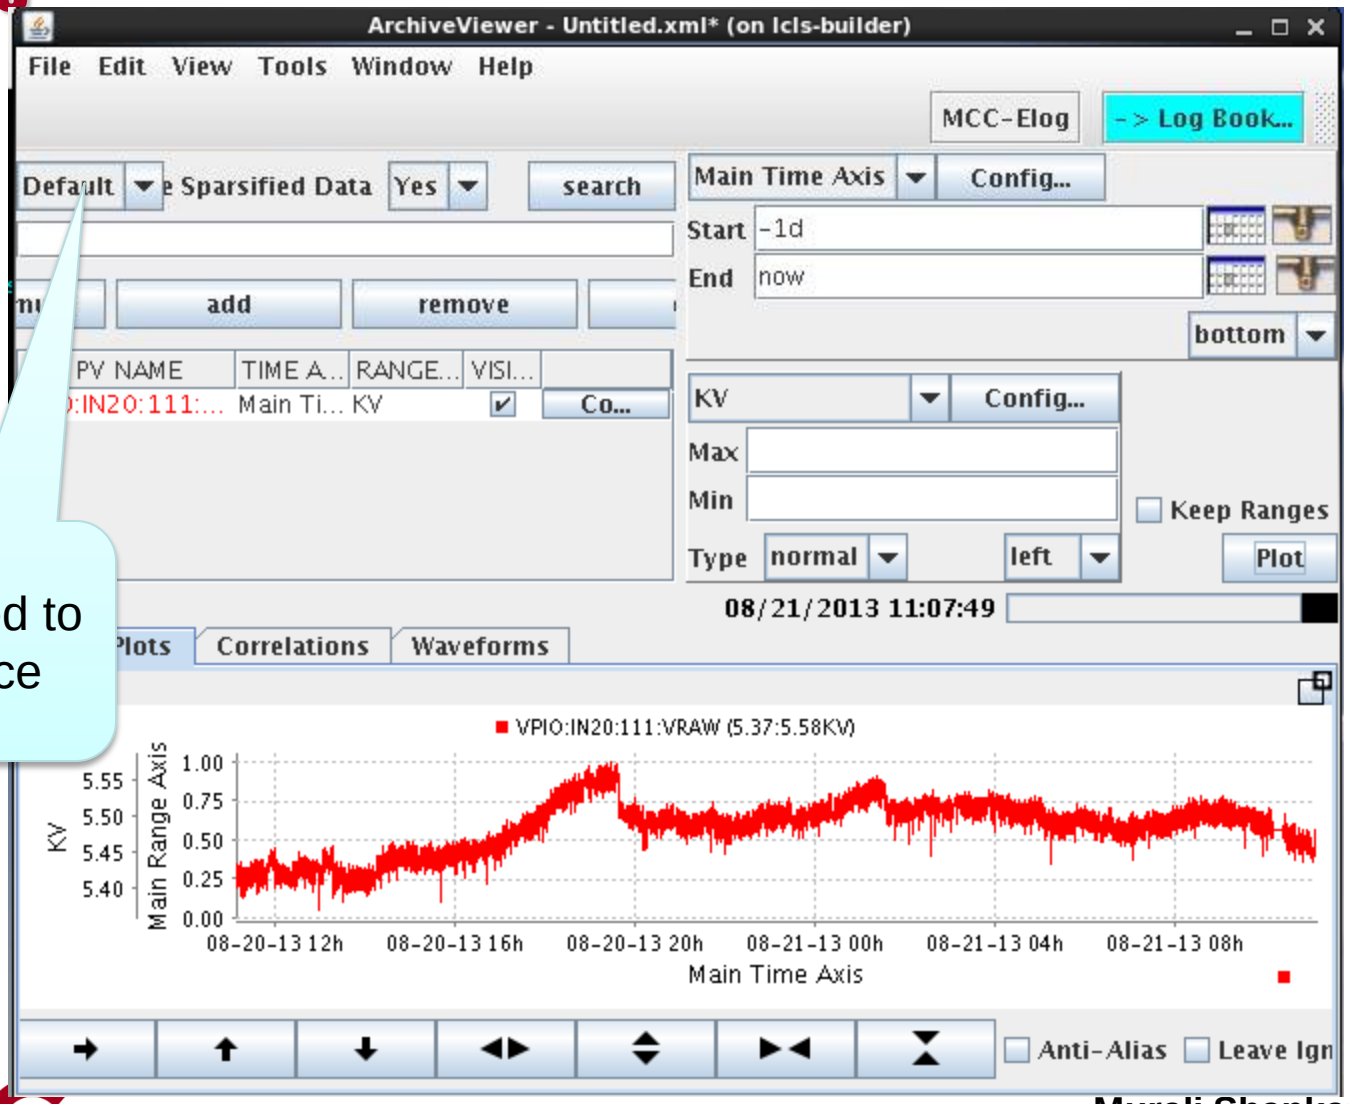

Connected to appliance

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

Response times

Controls

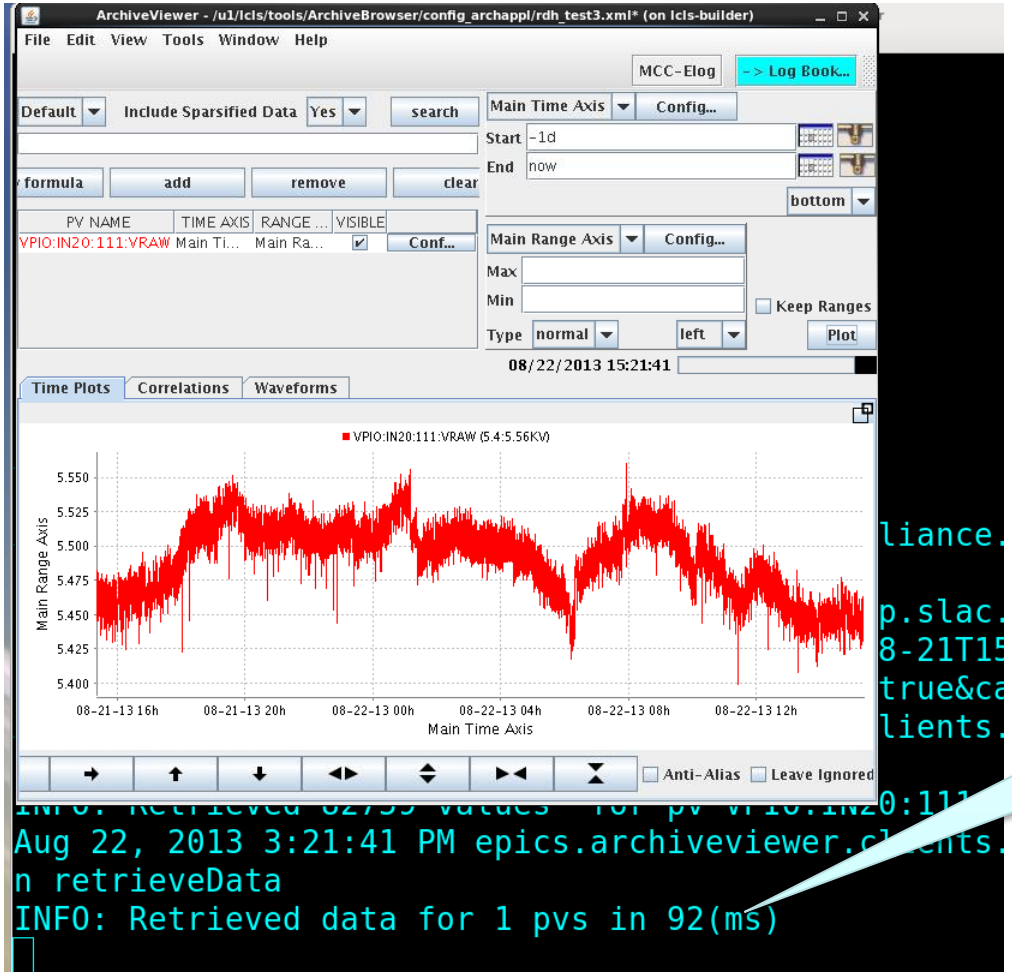

1 days worth of 1Hz DBR_DBL< 500ms

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

Transparent Chnl Arch proxy

Controls

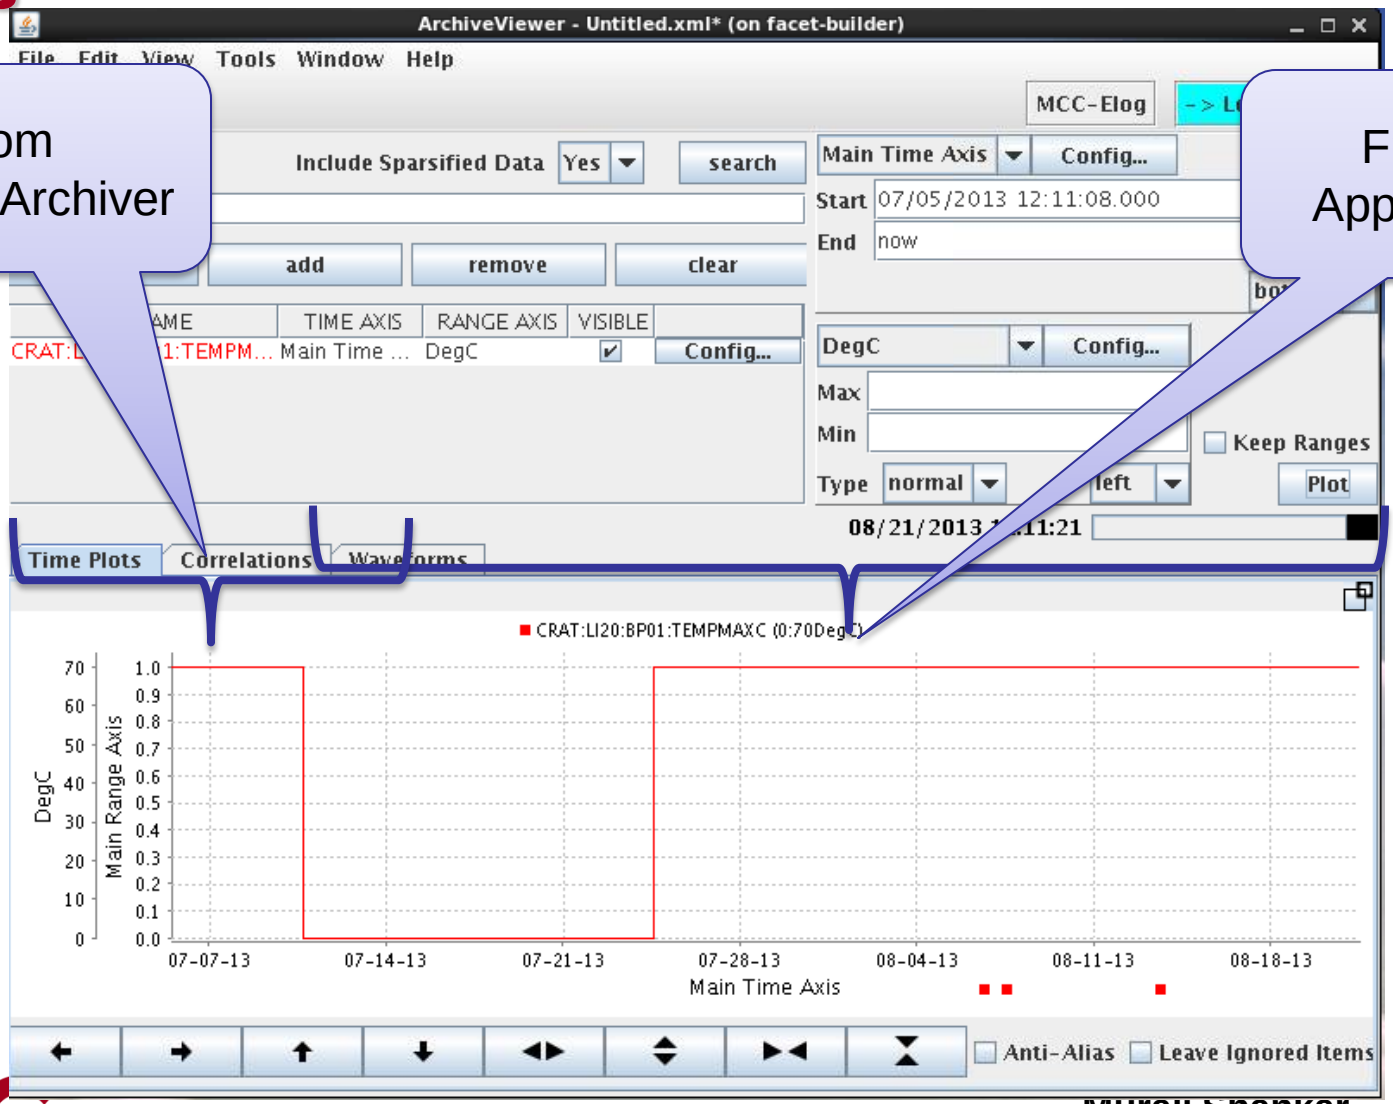

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

CSS Databrowser

Controls

- Extension point implemented but needs to be updated to cater to interface changes
- Thanks, Kunal

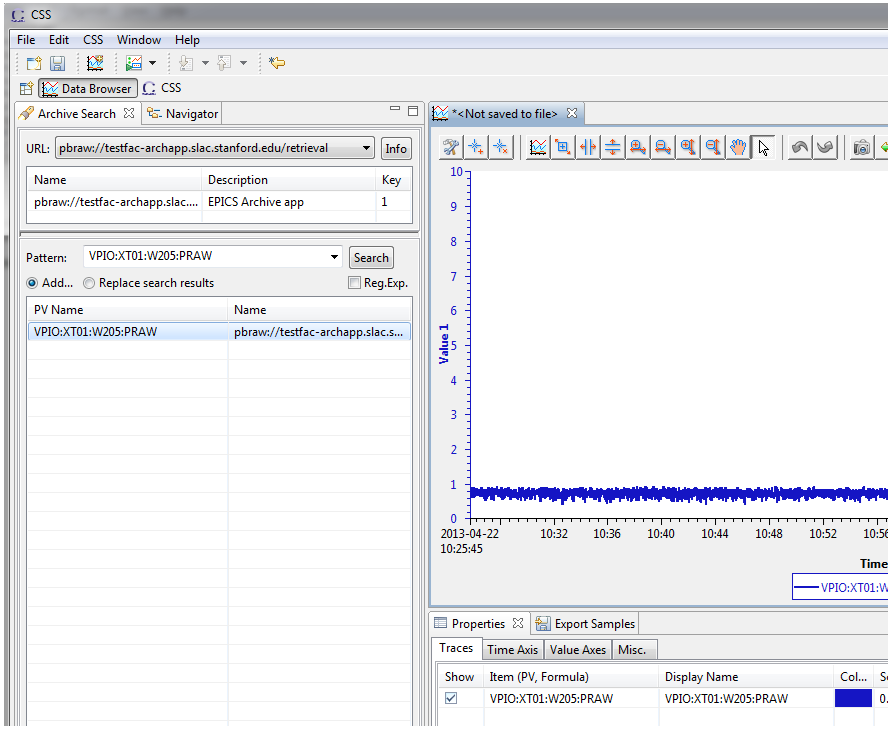

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

Post processing

Controls

- Using Apache Commons Math

Operator	Desc
`firstSample`	Returns the first sample in a bin. This is the default sparsification operator.
`firstFill`	Similar to the firstSample operator with the exception that we alter the timestamp to the middle of the bin and copy over the previous bin's value if a bin does not have any samples.
`lastFill`	Similar to the firstFill operator with the exception that we use the last sample in the bin.
`mean`	Returns the average value of a bin. This is computed using SummaryStatistics and is `SummaryStatistics.getMean()`
`median`	Returns the median value of a bin. This is computed using DescriptiveStatistics and is `DescriptiveStatistics.getPercentile(50)`
`std`	Returns the standard deviation of a bin. This is computed using SummaryStatistics and is `SummaryStatistics.getStandardDeviation()`
`jitter`	Returns the jitter (the standard deviation divided by the mean) of a bin. This is computed using SummaryStatistics and is `SummaryStatistics.getStandardDeviation()/SummaryStatistics.getMean()`
`ignoreflyers`	Ignores data that is more than the specified amount of std deviation from the mean in the bin. This is computed using SummaryStatistics. It takes two arguments, the binning interval and the number of standard deviations (by default, 3.0). It filters the data and returns only those values which satisfy `Math.abs(val - SummaryStatistics.getMean()) <= numDeviations*SummaryStatistics.getStandardDeviation()`
`variance`	Returns the variance of a bin. This is computed using SummaryStatistics and is `SummaryStatistics.getVariance()`
`popvariance`	Returns the population variance of a bin. This is computed using SummaryStatistics and is `SummaryStatistics.getPopulationVariance()`
`kurtosis`	Returns the kurtosis of a bin - Kurtosis is a measure of the peakedness. This is computed using DescriptiveStatistics and is `DescriptiveStatistics.getKurtosis()`
`skewness`	Returns the skewness of a bin - Skewness is a measure of the asymmetry. This is computed using DescriptiveStatistics and is `DescriptiveStatistics.getSkewness()`

Who defines functions?

Controls

- Defined by ops and physicists
  - For example, to calculate beam energy/second delivered to a given area
    - lastFill_1(Stopper) * lastFill_1(ChargeInBunch) * lastFill_1(BunchRepRate)

mean_600 in ArchiveViewer

Controls

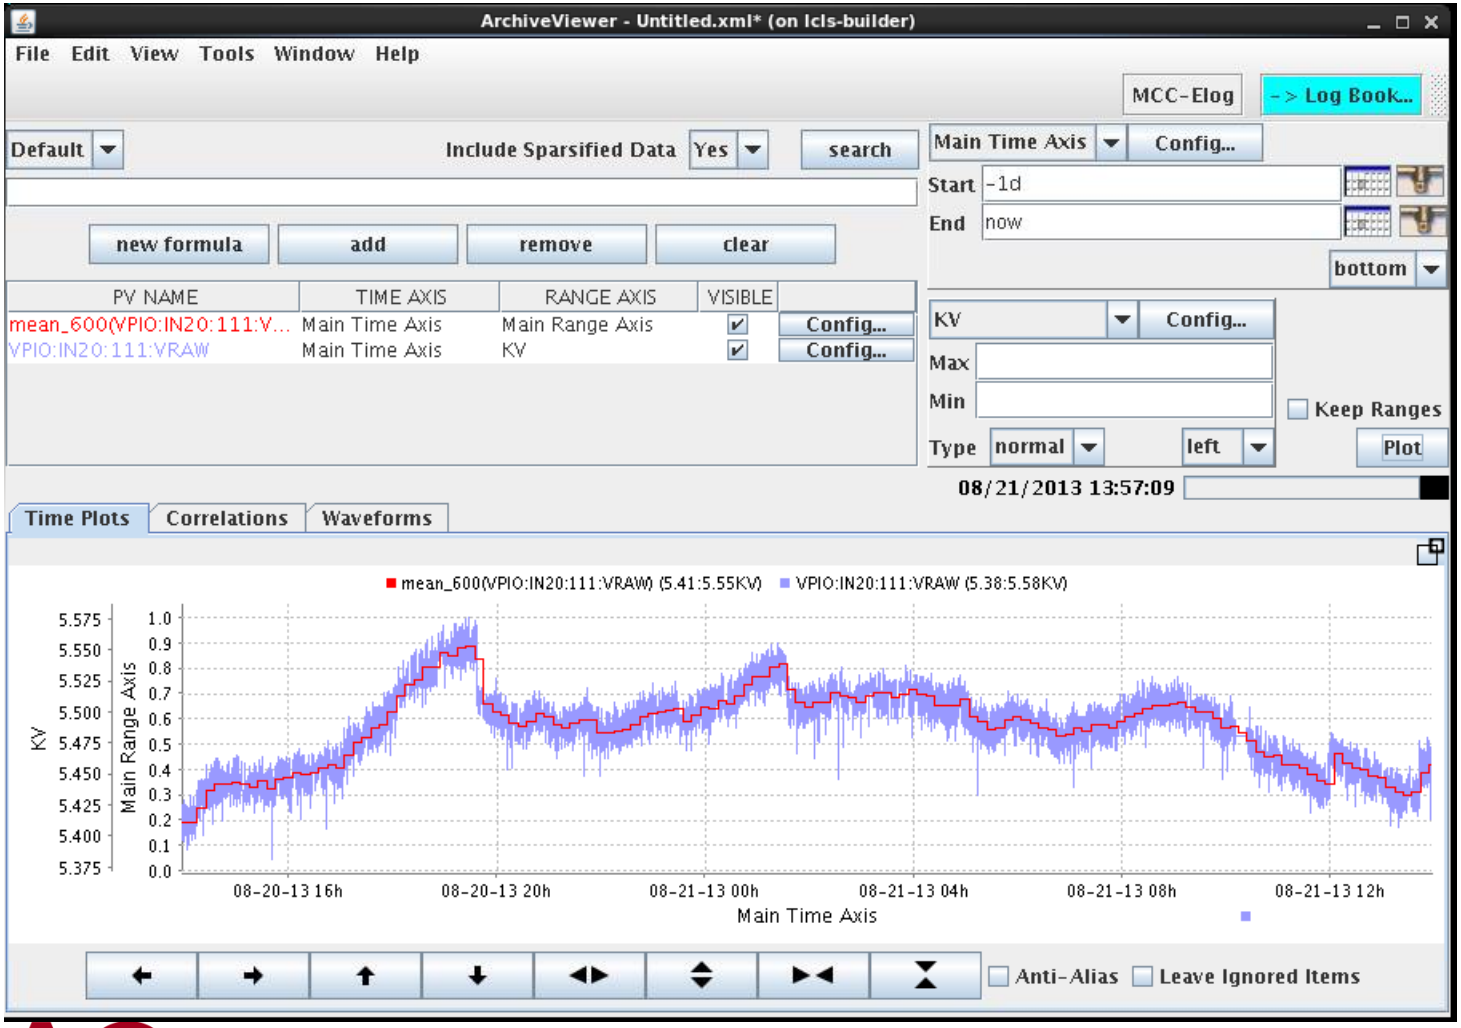

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

Controls

mean_600 Timestamps

Timestamp	Value
Aug/20/2013 15:15:00 PDT	5.439405171
Aug/20/2013 15:25:00 PDT	5.444522275
Aug/20/2013 15:35:00 PDT	5.4382804
Aug/20/2013 15:45:00 PDT	5.445006784
Aug/20/2013 15:55:00 PDT	5.446741128
Aug/20/2013 16:05:00 PDT	5.450580706
Aug/20/2013 16:15:00 PDT	5.449584914
Aug/20/2013 16:25:00 PDT	5.45015255

mean_600 response

Controls

30 days worth of mean_600(1Hz DBR_DBL) < 1 s

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

mean_600 multiple PVs

Controls

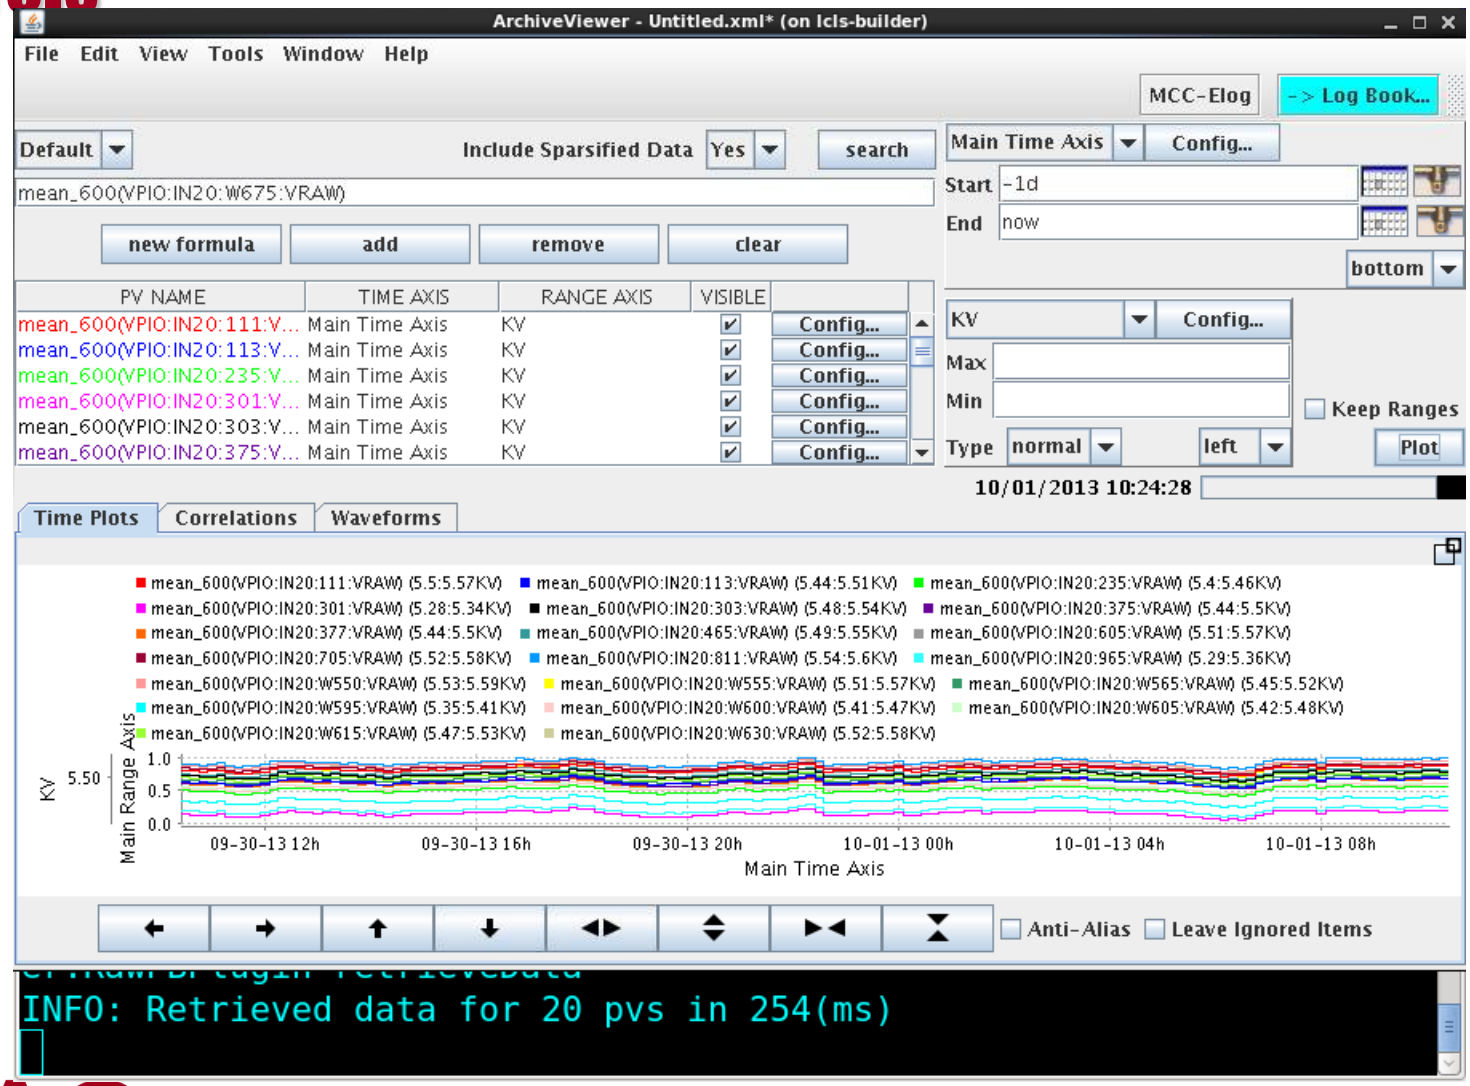

Murali Shankar
Luofeng Li

SLAC NATIONAL ACCELERATOR LABORATORY

Post Processing 2

Controls

- However, Matlab is **still** weapon of choice.
  - Directly get data into Matlab, Python.
  - Appliance serves **.mat** files directly
  - Easy to support other tools (R?)

Runtime fields

Controls

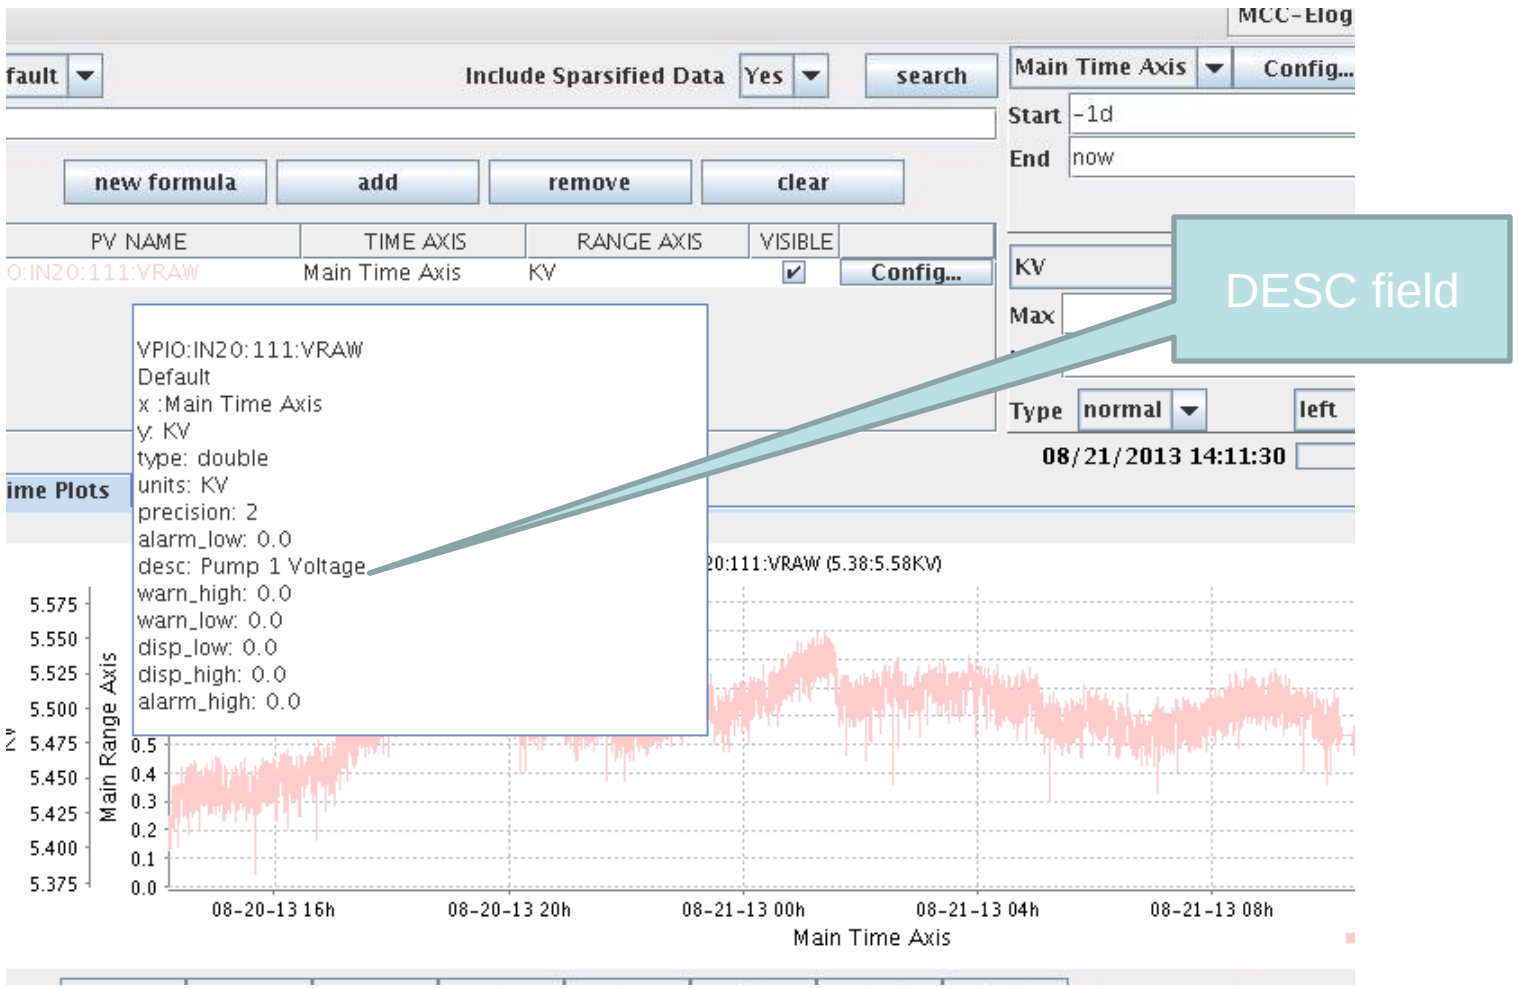

Murali Shankar
Luofeng Li

Users add PVs

Controls

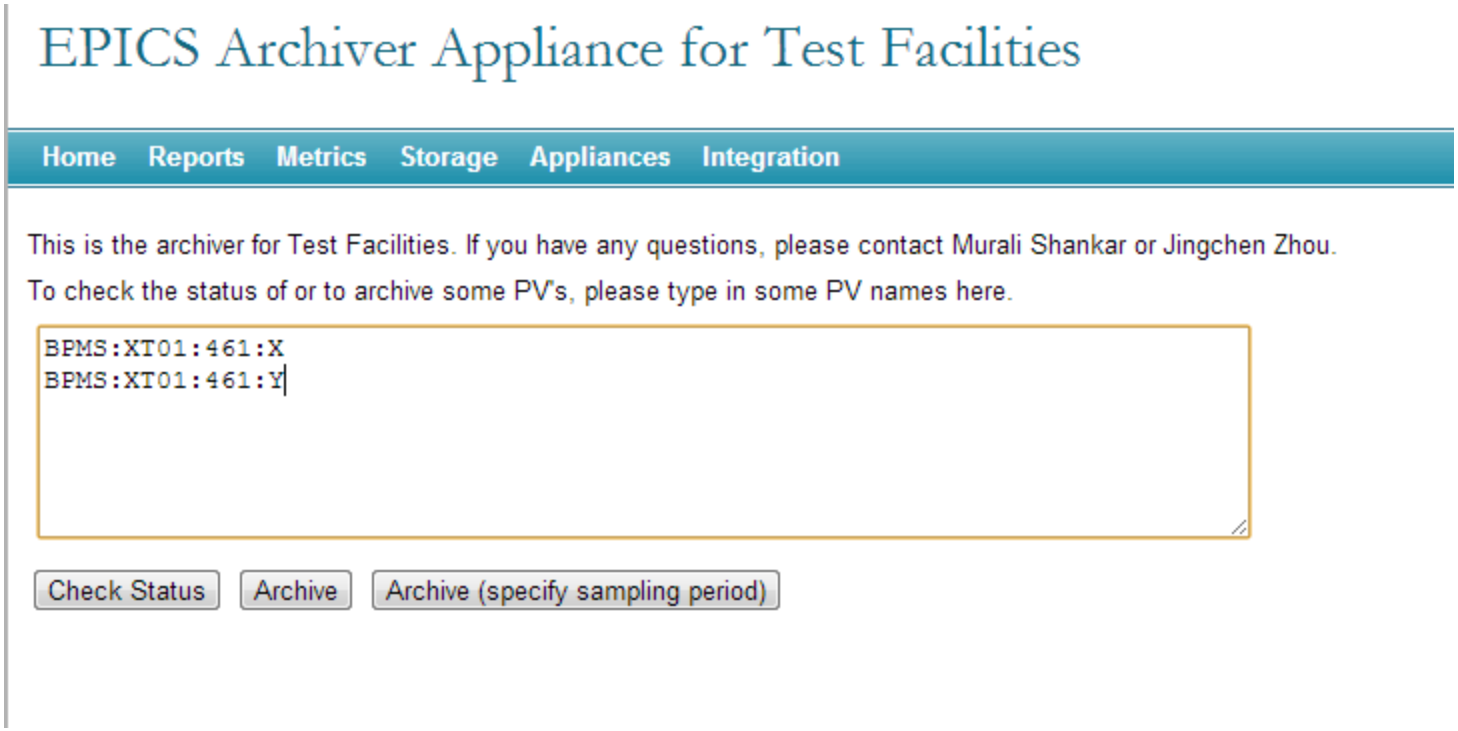

SLAC NATIONAL ACCELERATOR LABORATORY

Murali Shankar
Luofeng Li

Aliases

Controls

- Support for EPICS aliases
  - Use the .NAME field to determine “real” name
- Support for archiver only aliases
  - “Retire” old PV names
  - No UI yet, BPL only
    - If folks are interested, will add to UI
- Thanks, Emmanuel!

Type changes

Controls

- Occasionally, PVs change type
  - Spurious type changes
- Options
  - Rename PV to keep old data
  - Convert if possible
- Suggestion from Ralph (thanks) to indicate to user when this is the case.
  - Takes care of EGU changes.

Reconnect times

Controls

- Reconnect times on IOC reboot
- Reconnect times if engine crashes
- Reconnect times if gateway crashes
- Improving this
  - Reduced reconnects to ~ 4 minutes in dev.
  - Have a backup where we pause/resume the PV
- CAJ/JCA

Current focus

Controls

- Startup times vs CA search storms
- Stability
  - Monitoring scripts
- Policies for waveforms
- Imports of large numbers of PVs
- CAJ

Plans

Controls

- Decimation
- More operators – RMS
- Import data
- Improve retrieval performance
- ChannelFinder integration
- NIO2 subinterface

Quickstart

Controls

- Easy to try it out
  - Setup EPICS environment variables.
  - Download tar.gz and untar
  - Download tomcat
  - Run shell script

```
sh-4.1$
sh-4.1$ ./quickstart.sh apache-tomcat-7.0.27.tar.gz
```

Installation

Controls

- Create appliances.xml
- Create policies.py
- Install script that installs on one appliance
  - Expect most people will develop their own install/upgrade scripts.
  - We have our own install/upgrade scripts.

```
sh-4.1$
sh-4.1$ ./install_scripts/single_machine_install.sh
```


Murali Shankar
Luofeng Li

Documentation

Controls

- Hosted on sourceforge
- Google EPICS Archiver Appliance
- Try it out with the quickstart

Questions

Controls