ELK, RabbitMQ, Collectd and Freeboard in the wild at NERSC.

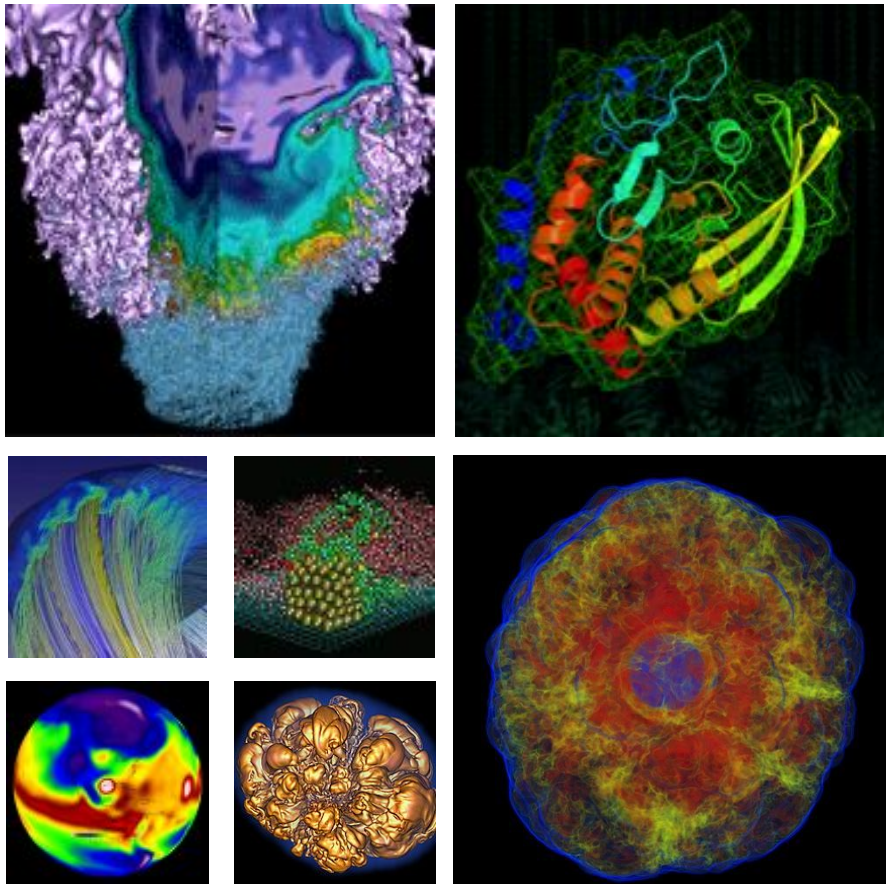

Thomas Davis (tadavis@lbl.gov)
Cary Whitney (clwhitney@lbl.gov)

10/13/15

Nobel Prize in Physics 2015

Arthur B. McDonald, Queen's University (SNO)
Takaaki Kajita, Tokyo University (Super Kamiokande)

Scientific Achievement

The discovery that neutrinos have mass and oscillate between different types

Significance and Impact

The discrepancy between predicted and observed solar neutrinos was a mystery for decades. This discovery overturned the Standard Model interpretation of neutrinos as massless particles and resolved the "solar neutrino problem"

Research Details

The Sundbury Neutrino Observatory (SNO) detected all three types (flavors) of neutrinos and showed that when all three were considered, the total flux was in line with predictions. This, together with results from the Super Kamiokande experiment, was proof that neutrinos were oscillating between flavors and therefore had mass

A SNO construction photo shows the spherical vessel that would later be filled with water.

NERSC helped the SNO team use PDSF for critical analysis contributing to their seminal PRL paper. HPSS serves as a repository for the entire 26 TB data set.

Q. R. Ahmad et al. (SNO Collaboration). Phys. Rev. Lett. 87, 071301 (2001)

Old System @ OSF

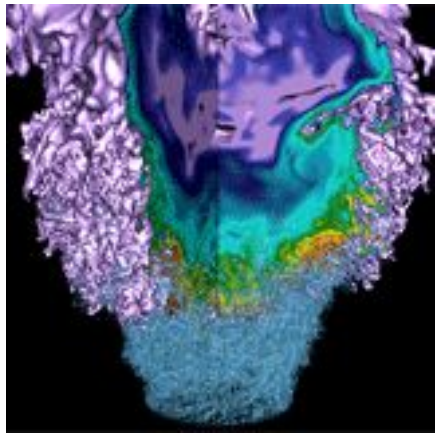

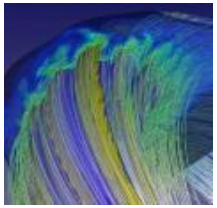

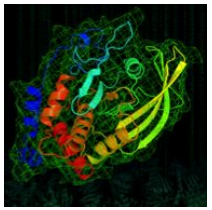

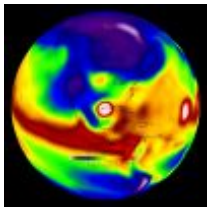

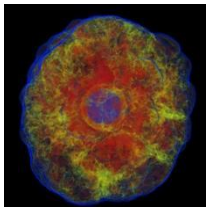

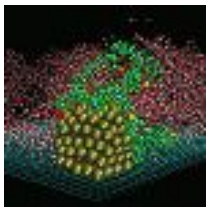

Old System

MySQL based

60+ billion items makes for an unmanageable db.

Dashboard is fixed.

End user cannot make any changes at all.

Changes in the center

Needs to programmed into the dashboard.

Synapsense

Only does temperatures/flow/basic power

Update rate is 6-7 minutes due to battery life.

Vendor does the updates

Adds delays and costs.

PHP & Perl based

Yea, let's live it up 90's style.

Controls & Protocols

Bizarre world of protocols

Seabus, Modbus/RTU, Modbus/TCP, SNMP, BACNET, Metasys, Incom, Johson Controls Metasys2.

Gateway for all but Metasys2 was installed.

Security was not a priority in these systems.

Three different Building Management Systems

Two have been joined together

- Yard Plant and Roof Cooling Towers
- BACNET based
  - Insecure protocol.

One is Johnson Controls Metasys2 based

- Win98 based
- No way to collect any information from it.

Collection rates are slow.

Anything faster than 5 minutes per point is ***hard***

Logs & metrics

No scalable interface to syslogs

Splunk is available, but it gets overloaded..

Batch queue stats

Available, but in a different area…

Node metrics

Ganglia on some systems, but not integrated.

Some nodes there is no idea what is going on at all.

Lustre, GPFS stats

Lustre performance monitoring is available..

But it's cranky at times..

No real collection of GPFS stats at this time.

Can cause performance problems if not done right.

New @ CRT

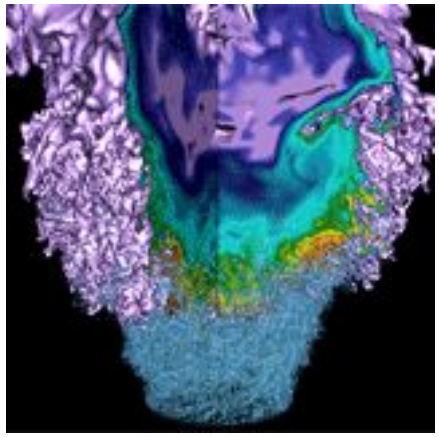

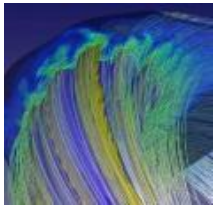

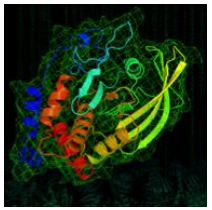

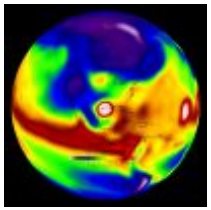

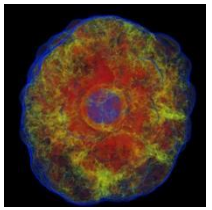

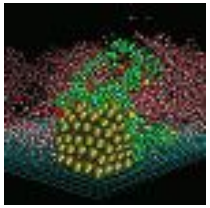

NERSC CRT

CRT is brand new Center

- Computational Research Center
- Both machine room and office space.

So it's time for a clean slate.

- Much of what is at OSF will not be transferring to CRT.
- Older PDU's, power strips, racks will end up in surplus.
  - **Standardize as much as possible!**
    - No rack gets less than 50 amps of power.
    - Cable for 60 amps of power
- Narrow down the protocols
- Establish security protocols up front.
- New procedures, new processes.
  - Examine ***everything.***

NERSC CRT continued

CRT has a single BMS

- ALC based system.
  - BACNET based interface.

CRT was designed for a very low PUE (<1.1)

- To achieve this, CRT has NO chillers.
- Tower water into heat exchangers for the Crays.
- Outside air for the rest of the room.

True hot/cold aisle design in the common area.

- Doors on the hot aisle.
- Blanking panels to fill all holes.
  - Temperature monitoring of both the front, rear and top of racks to detect problems.

CRT Continued

One PDU design across the board

- Emerson/Liebert 300 kva PDU with Modbus/TCP network card.
- Monitors down to each pole on a breaker.
- Only two types of plugs
  - 50 amp and 60 amp provisioning
  - All cables are for 60 amps.
    - Minimize needs to pull cables.

Limited Number of Protocols

- Onewire, BACNET, Modbus/TCP, Modbus/RTU and SNMP.
- Security is part of the design when possible.
  - Private IP or multiple firewalls.
  - Separate networks.
- UPS power to certain parts to keep things up as long as possible.

CRT Continued.

Large number of collection points.

- Several thousand meters just for power
  - Medium & Low Voltage Switchgear
  - 1200 amp Panel Boards
  - Liebert/Emerson PDU's
  - Raritan PDU strips.
- Lots of temperatures sensors.
  - 10 per rack
    - 4 per set of doors
    - 2 on top
  - Onewire network based.
  - 22 sensors on the ceiling in the common area.

Need to design for a minimum collection rate of 10k points per second, just for the environmental.

CenterWide monitoring

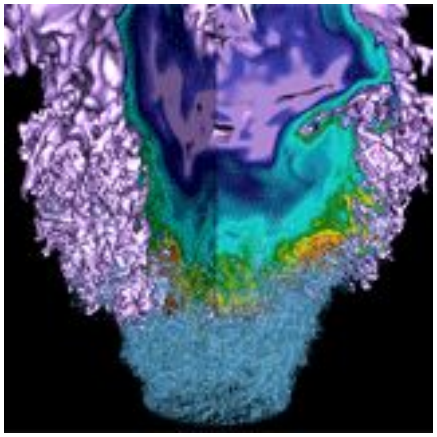

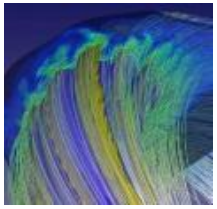

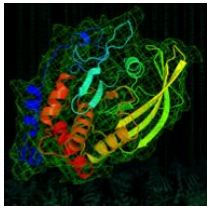

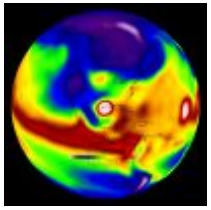

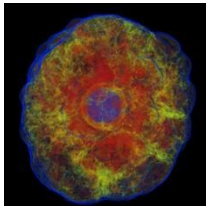

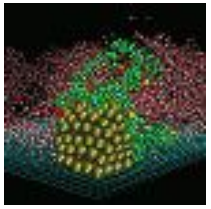

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Center Wide Monitoring

Need to answer the 5 W's.

Who, What, Why, Where and When.

Need to combine logs and metrics; ie:

- System logs
- Node performance
  - Login, batch, other nodes
- Cray environmentals
  - Power consumption.
    - Efficiency is big concern.
  - Internal and External Temperature
- External Environmentals can affect the systems.
  - We need to identify these, and minimize those effects.

Project Goals

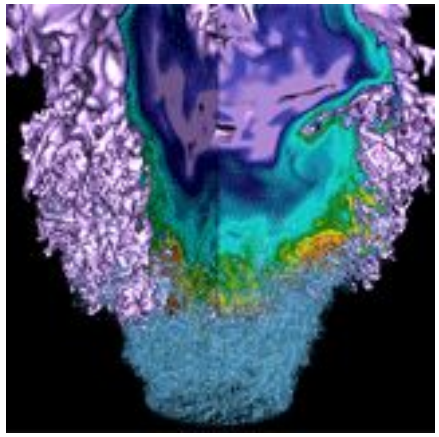

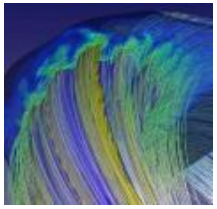

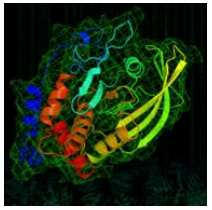

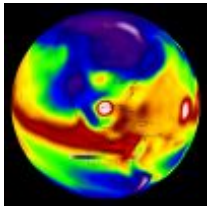

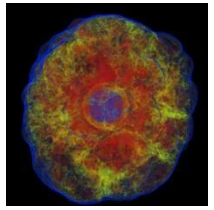

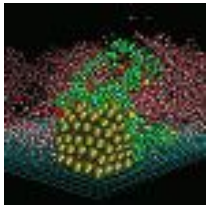

Project Goals

Scalable

Reliable

Maintainable

Composable

Different users, different data problem.

What User A wants to see, is not what User B wants to see.

Project Goals (continued)

Instrument *EVERYTHING.*

- OS & programs
  - Glue code performance and logs are sent into the system.

Get *ALL* logs

- Batch
- Syslog
- SEDC, etc, etc..

Loose *NOTHING.*

- Everything that is captured must make into a storage system.
- Once in the primary storage system the data must be able to be manipulated into other storage/systems.
- Provide a ***FAST*** interface when possible.
- Real-time interfaces.

Project Goals (continued)

Look at what Netflix, Google, Facebook, Twitter are using and try to leverage their technology.

They already scale to volumes we will see.

Common programming language where possible.

Try for Python v2.x, v3.x in the future for the glue code.

Go, Java, Ruby/JVM, erlang, C are in use in the system.

Common data format.

Allows easy addition or replacement of collectors and clients.

Simplifies interfaces

This does however place a larger burden on transport.

Web based client

must support tunneling or proxying

No flash or java applets.

HTML5 based

SSH w/SOCKS5 & PAC file must work.

Project Goals (continued)

Four types of consumers of the data.

System Administrators who want to see the immediate data and correlations.

Scientific Researchers who want to see their job performance and resource usage.

Researchers who want to see general and/or specific data about the systems and/or jobs for the planning of future systems. This requires save all the data for the life of the system or beyond.

Management who would like overall views of resource utilization.

Design

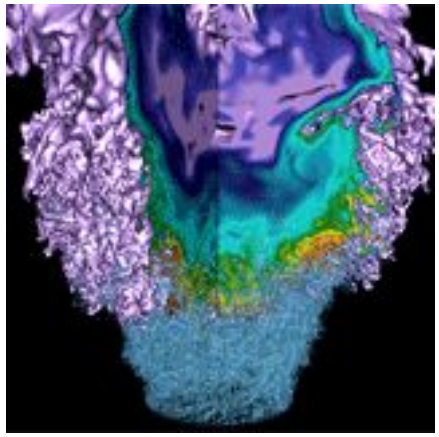

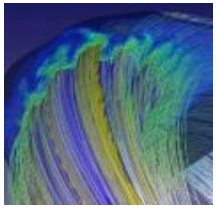

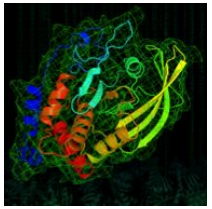

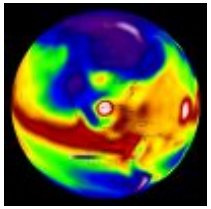

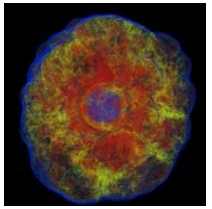

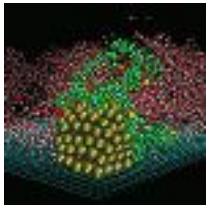

New System at CRT.

ELK stack

Elasticsearch/Logstash/Kibana

RabbitMQ

AMQP/MQTT/STOMP-WS

Collectd

Plugins! Plugins!

Glue Scripts pushing data into RabbitMQ

BACNET, MODBUS, SNMP, OneWire

Freeboard

Pluggable API for widgets and data sources.

Mozilla Heka

Stream processing.

Yes, you can cross the streams!

Hardware

32 nodes

- Supermicro Fat Twin/High Density 4U units
  - 8 nodes per 4U of space.
- 10G Ethernet for clustering
  - Data movement
  - Virtual Machine movement
- 64GB per node
  - Elasticsearch/Logstash are Java based, and heap sizes beyond 32G are problematic.
- 540G SSD
- 2x2TB hard drive

Sensor Network

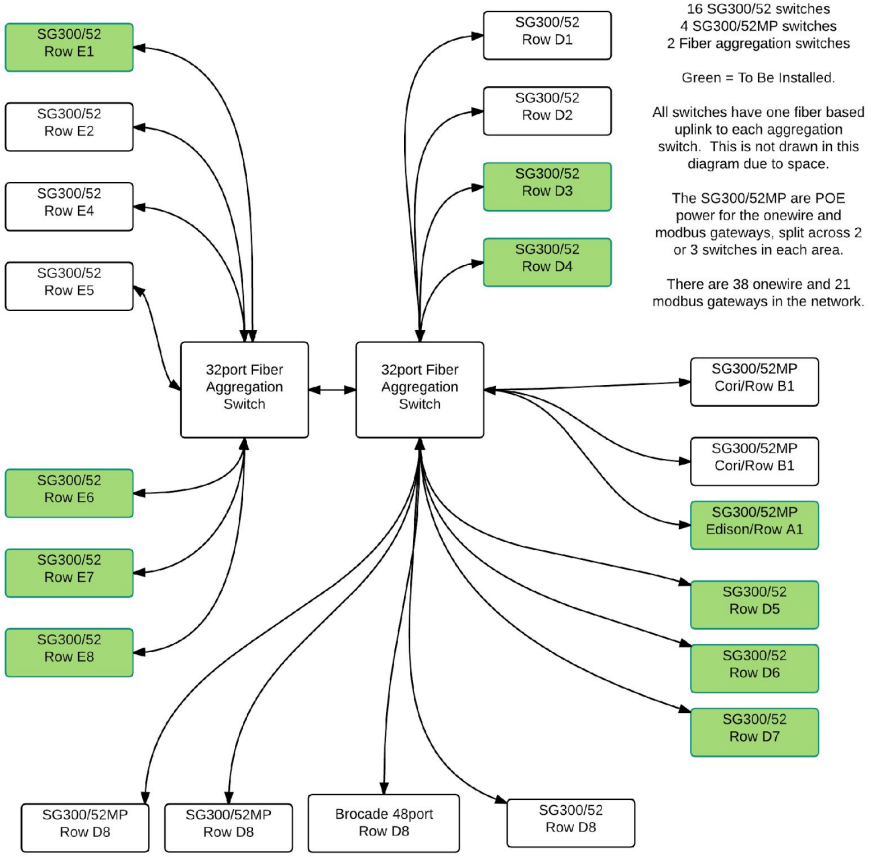

How it all works.

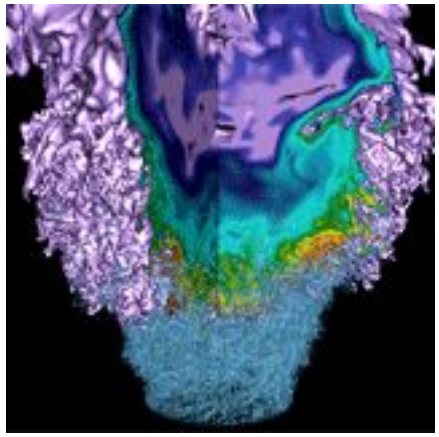

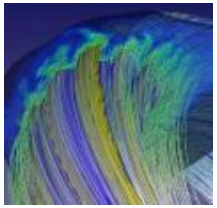

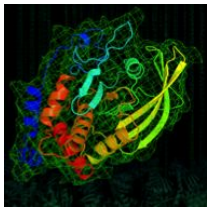

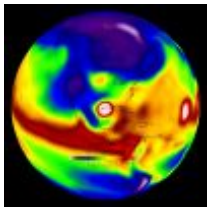

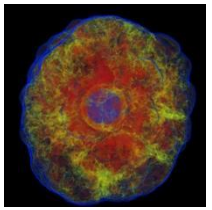

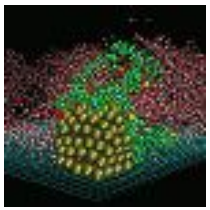

Flow Diagram.

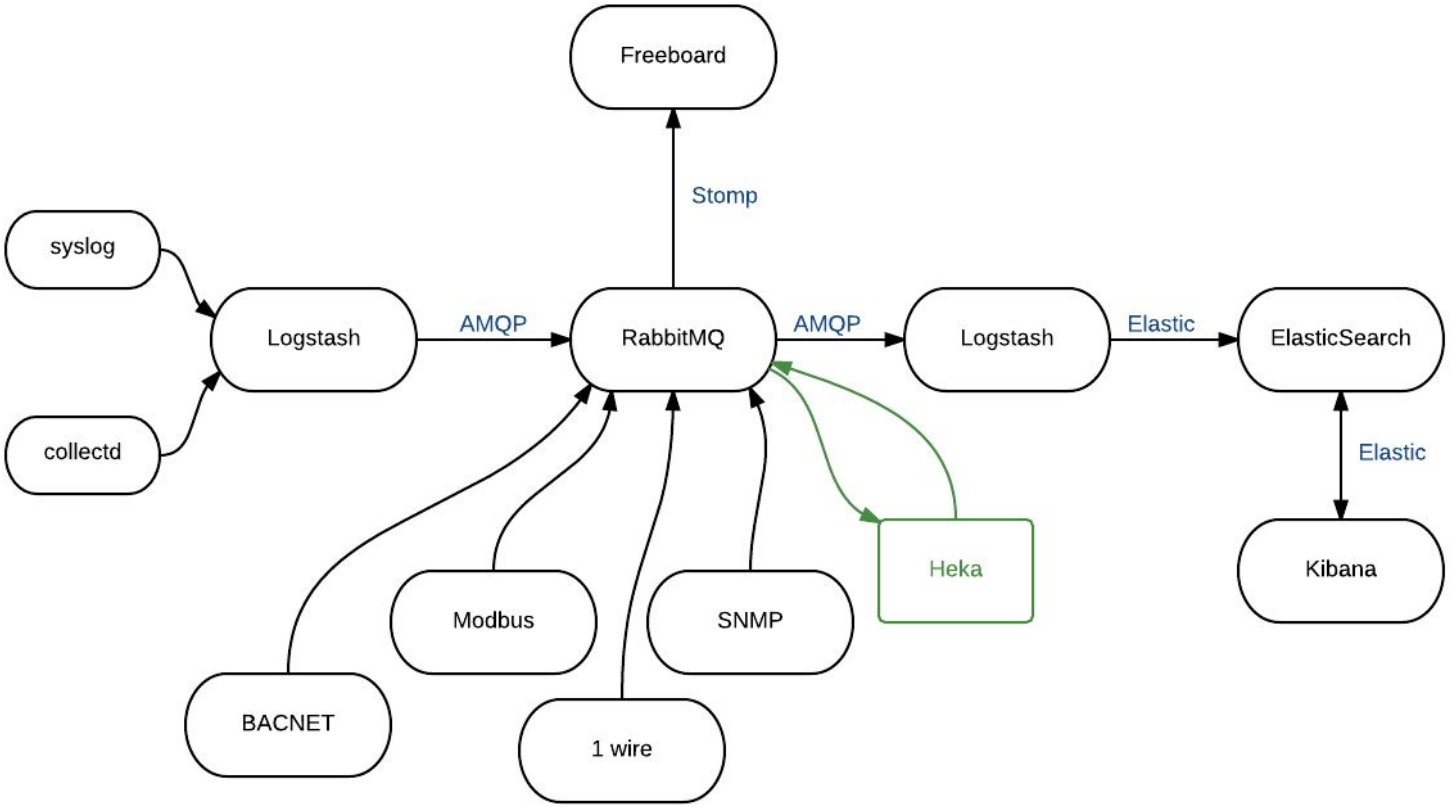

Kibana

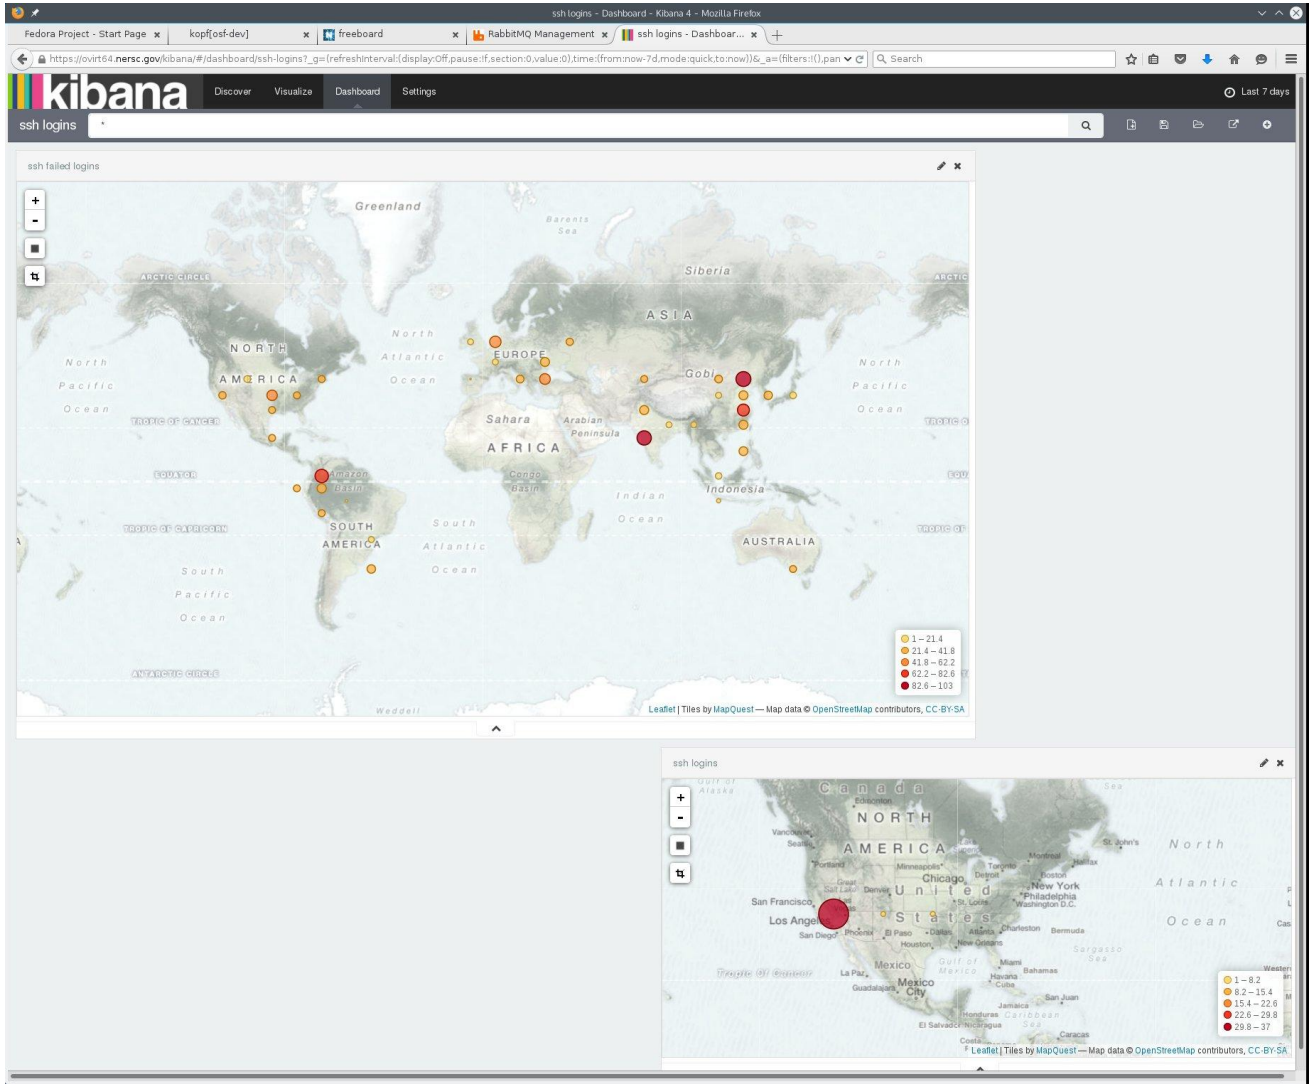
ssh logins - Dashboard - Kibana 4 - Mozilla Firefox
kibana
Discover
Visualize
Dashboard
Settings
Last 7 days
ssh logins
ssh failed logins
Greenland
NORTH AMERICA
EUROPE
ASIA
Siberia
Gobi
Sahara
Arabian Peninsula
AFRICA
Indonesia
AUSTRALIA
SOUTH AMERICA
Amazon Basin
Congo Basin
1 – 21.4
21.4 – 41.8
41.8 – 62.2
62.2 – 82.6
82.6 – 103
Leaflet | Tiles by MapQuest — Map data © OpenStreetMap contributors, CC-BY-SA
ssh logins
Canada
NORTH AMERICA
United States
Mexico
Mexico City
San Francisco
Los Angeles
Chicago
New York
1 – 8.2
8.2 – 15.4
15.4 – 22.6
22.6 – 29.8
29.8 – 37
Leaflet | Tiles by MapQuest — Map data © OpenStreetMap contributors, CC-BY-SA

Freeboard

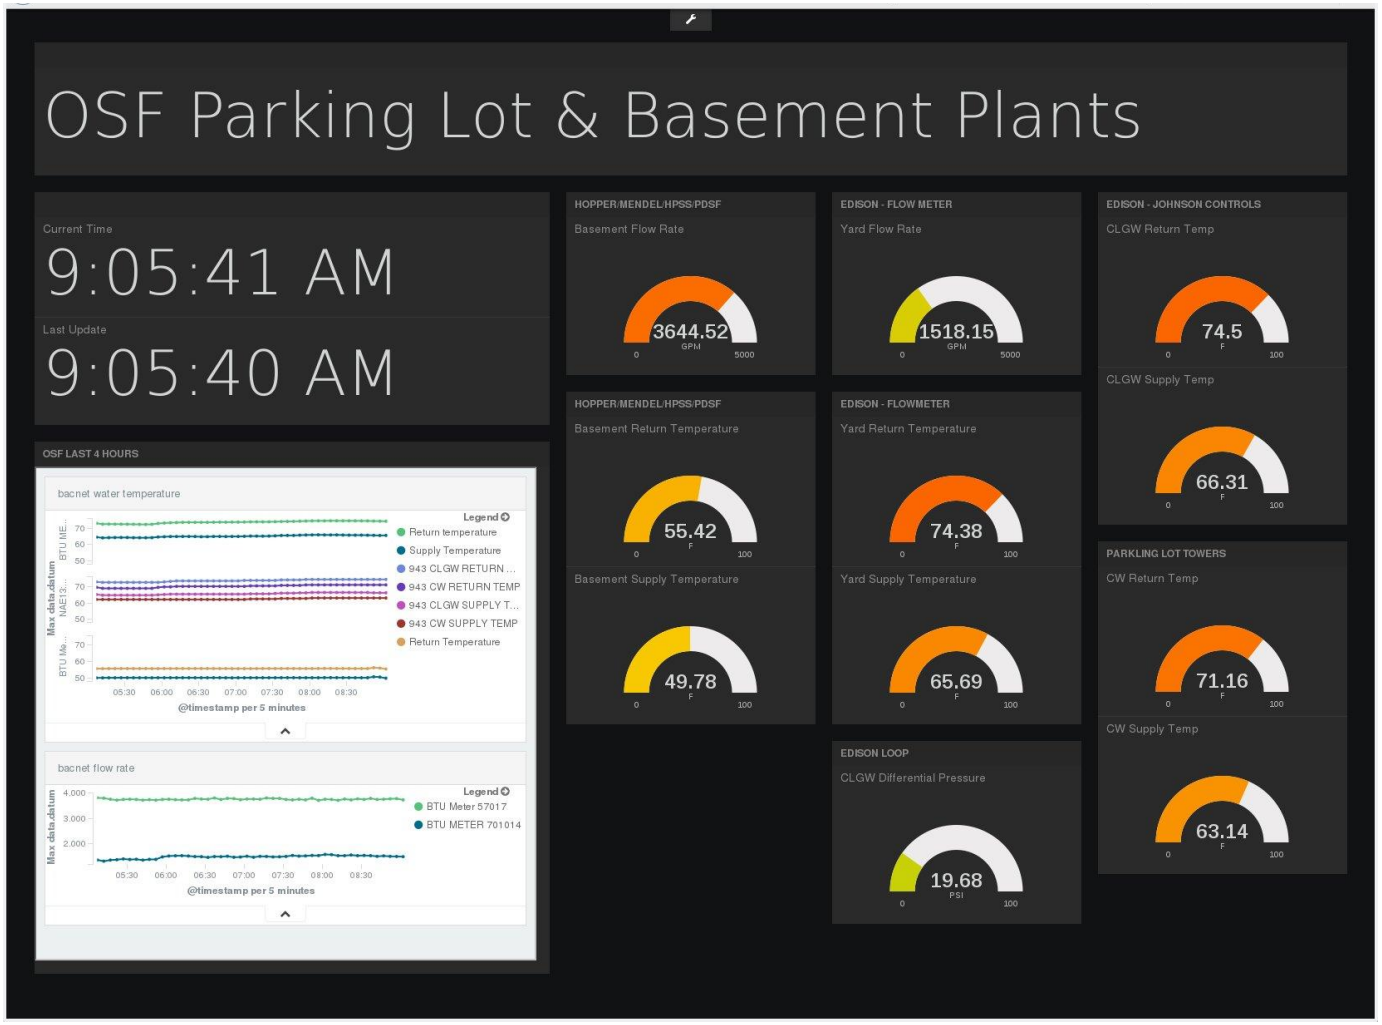
OSF Parking Lot & Basement Plants
Current Time
9:05:41 AM
Last Update
9:05:40 AM
OSF LAST 4 HOURS
bacnet water temperature
Legend
Return temperature
Supply Temperature
943 CLGW RETURN ...
943 CW RETURN TEMP
943 CLGW SUPPLY T...
943 CW SUPPLY TEMP
Return Temperature
@timestamp per 5 minutes
bacnet flow rate
Legend
BTU Meter 57017
BTU METER 701014
@timestamp per 5 minutes
HOPPER/MENDEL/HPSS/PDSF
Basement Flow Rate
3644.52
GPM
HOPPER/MENDEL/HPSS/PDSF
Basement Return Temperature
55.42
F
Basement Supply Temperature
49.78
F
EDISON - FLOW METER
Yard Flow Rate
1518.15
GPM
EDISON - FLOWMETER
Yard Return Temperature
74.38
F
Yard Supply Temperature
65.69
F
EDISON LOOP
CLGW Differential Pressure
19.68
PSI
EDISON - JOHNSON CONTROLS
CLGW Return Temp
74.5
F
CLGW Supply Temp
66.31
F
PARKLING LOT TOWERS
CW Return Temp
71.16
F
CW Supply Temp
63.14
F

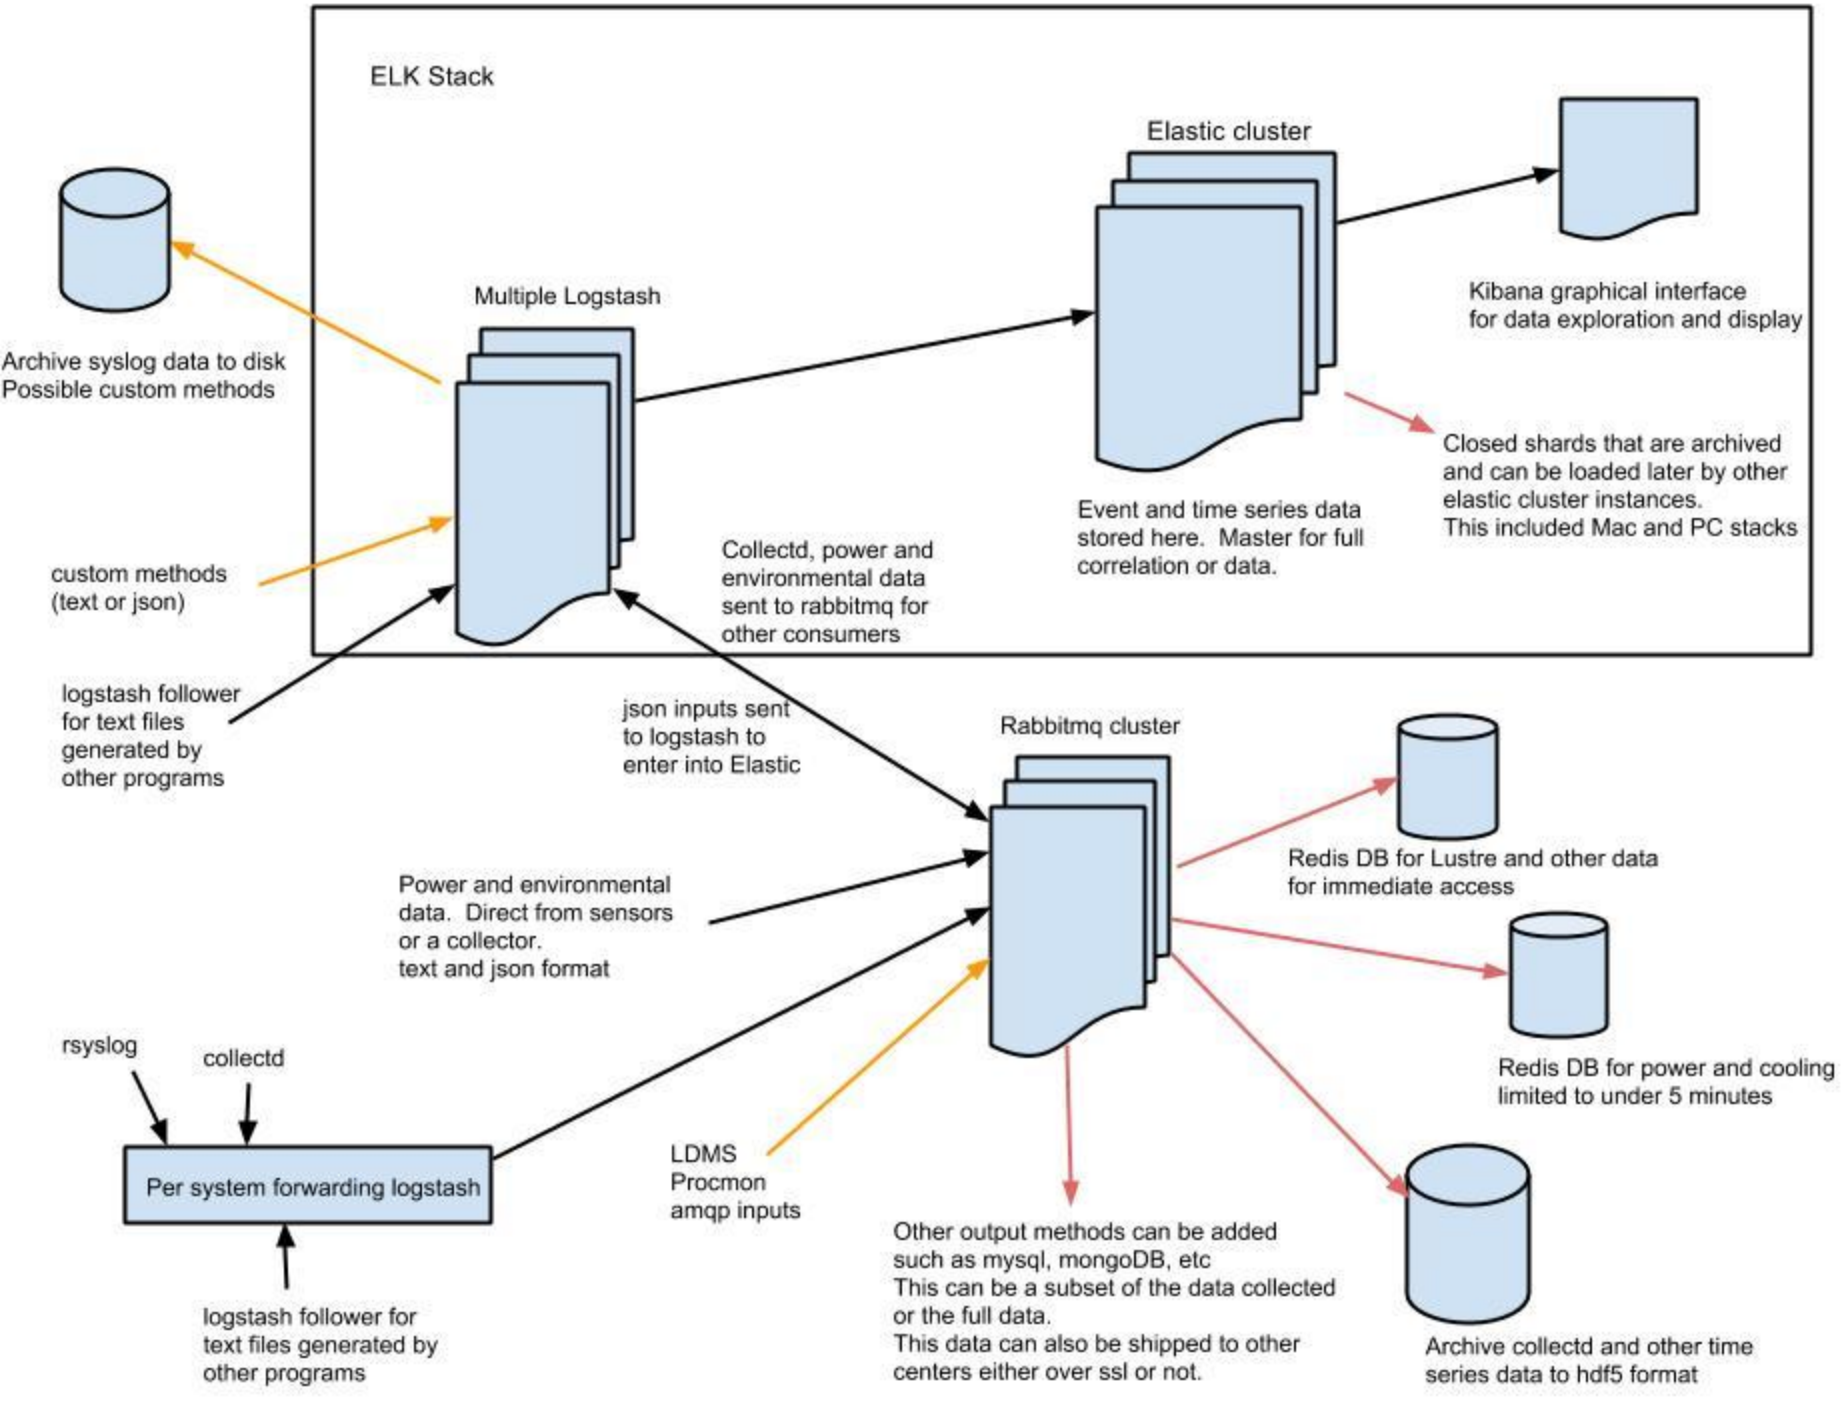
ELK Stack
Elastic cluster
Kibana graphical interface
for data exploration and display
Multiple Logstash
Archive syslog data to disk
Possible custom methods
Closed shards that are archived
and can be loaded later by other
elastic cluster instances.
This included Mac and PC stacks
Event and time series data
stored here. Master for full
correlation or data.
custom methods
(text or json)
Collectd, power and
environmental data
sent to rabbitmq for
other consumers
logstash follower
for text files
generated by
other programs
json inputs sent
to logstash to
enter into Elastic
Rabbitmq cluster
Redis DB for Lustre and other data
for immediate access
Power and environmental
data. Direct from sensors
or a collector.
text and json format
Redis DB for power and cooling
limited to under 5 minutes
rsyslog
collectd
Per system forwarding logstash
LDMS
Procmon
amqp inputs
Other output methods can be added
such as mysql, mongoDB, etc
This can be a subset of the data collected
or the full data.
This data can also be shipped to other
centers either over ssl or not.
Archive collectd and other time
series data to hdf5 format
logstash follower for
text files generated by
other programs

Further Work.

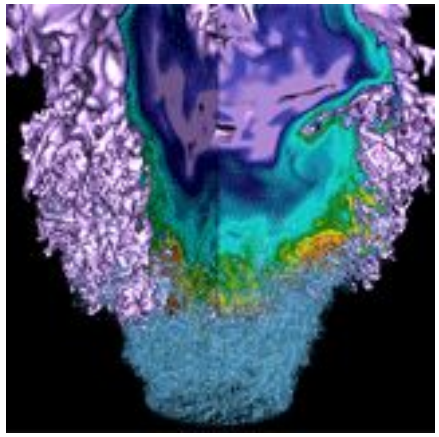

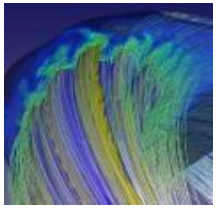

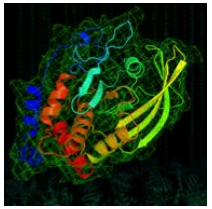

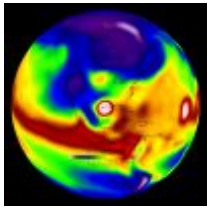

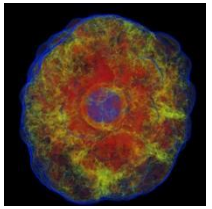

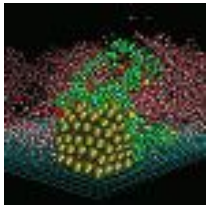

The Great ToDo List.

Dashboards

Always more work here. How do we progress past the pie chart? More complex data. Can we use the human ability of pattern matching?

Stream Processing

Add value to the data coming in. Take action with what we see. Nagios replacement and security actions.

AI, can the machine learn to handle situations on its own?

Data Cleanup/Anonymizer

Everybody wants it but how do we keep private information private?

Would not like IP topology to get out but still need IP knowledge to do correlations.

Workshop? Working Group? Berlin?

More Information.

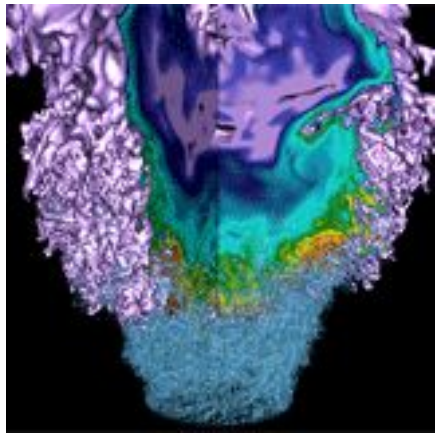

NERSC

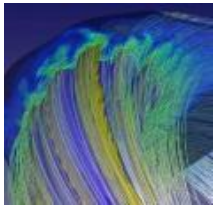
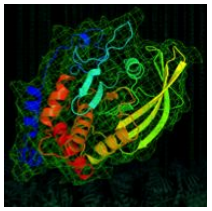
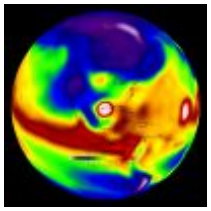
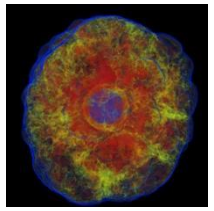
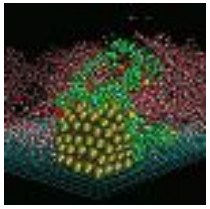

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Reading list

Metrics 2.0

http://metrics20.org/

Metrics need to self-describe when possible.

Software Defined Buildings

https://github.com/SoftwareDefinedBuildings/pybacnet

Needs some updating, but it does work.

Logstash Book

http://logstashbook.com/

Mozilla Heka

https://github.com/mozilla-services/heka

Mbusd (Modbus/tcp to Modbus/RTU daemon)

https://github.com/3cky/mbusd

National Energy Research Scientific Computing Center

Elasticsearch

Timeseries Logging engine with Search.

Cluster driven

- Shards
- Replicas
- Indexes

API is HTTP/JSON based.

No security in the free version.

Java based with Apache Lucene as db engine.

Logstash

Swiss Army knife of logging.

Can transform data from one format to another.

Normalize date/times.

Can accept binary data and convert to JSON.

Supports read/writes to AMQP.

Kibana

Web based visualization tool.

Graphs, Mapping, Tables

User driven.

User can define their own graphs and dashboards.

Exploring data can be interesting.

RabbitMQ

Buffers/Routes between Logstash, collectors into Elasticsearch

Swiss Knife of Protocols

- AMQP
- STOMP-WS
- MQTT

Centerwide secure transport.

- SSL
- LDAP Account(s)
  - But do not use for large volume of transactions.
    - Does not cache or pool connections..

Collectd

Node metrics

Also used to collect information on

- RabbitMQ plugin.
- Elasticsearch plugin.
- Slurm plugin.
- Working on a LMT plugin.

Commonly used

- One collector per node when possible.
  - Complaints of problems in the past however.
  - Need to insure it does not cause problems.

Large support community.

- But this does not translate into mission critical support..

Freeboard

IoT dashboard.

Real-time display.

No accessing of any history.

Taps the streams, not the db.

Uses STOMP-WS interface of RabbitMQ.

Also has plugin interface.

Data Sources.

MQTT, STOMP via WebSocket, JSON, etc..

Widgets.

Several different graphing packages, ie D3.

Mozilla Heka

Stream processing

- Calculate a real-time PUE from the streams.
- Realtime Temperature sensor correction.
- Anomaly detection is another area to explore

Script language is Lua

- Used in other projects.
- Libraries and snippets of code available

Glue Scripts

BACNET

- Flow meters
- Building Management System (BMS)

MODBUS

- TCP and RTU
- Smart Gateways
- GE EPM4500, ION Meters, GE Smartbreakers
- Particle Sensors

SNMP

- PDU strips, others.

Onewire

- Temperature and pressure sensors

Be aware of Value Engineering - GE4500/Specta Meters.

Somehow, metering designed for single phase applications was installed.

Most 1200/800 amp panelboard BCM's are for Apartment/Condo installations, NOT HPC

GE – single phase and 3 phase power metering.

Eaton – just 3 phase power, no currents, no modbus – this is Hopper's cabinet breakers

Siemens – best, everything, but need proprietary software to configure – this is Edison's breakers at OSF.

Ovirt

Ovirt is Redhat's RHEV

- Vmware clone.
- Opensource however.

Provides virtual machine services.

Level of HA for maintenance

- Allows VM's to be migrated from one node to another node.

Options!

Glusterfs, iSCSI, NFS are supported.

Mesos/Marathon/Chronos

Supports Docker containers.

Marathon for node scheduling

- Used for data collection and other long running jobs.
- Modbus/TCP collectors can be scheduled
  - If they die for some reason, Marathon will restart them.

Chronos for cron jobs

- Jobs that must run but need HA to do so
  - IE, rollups.
  - Elasticsearch data movement triggers.

Data Collection

Onewire

Temperature, Humidity, and Pressure sensors.

from iButtonlink.com

SNMP

Raritan PDU's and other misc devices.

MODBUS

ION Power Meters in Substations

Both medium and low voltage stations.

GE Breakers in low voltage substations are attached to the ION's.

GE EPM4500/Quadlogic Metering

MODBUS/RTU via 2 wire RS485

2 boxes per panel

Each box has split RS485 cabling.

Lighthouse particle sensors

MODBUS/TCP based.

Establish a baseline in the center for HPSS tape archives.

Smart Modbus Gateways

Data collection is done on the gateway

Better fault tolerance

Better performance

Better price.

More stats of what is going on.

- Serial line information can be collected.
- Firewalling of ipv4 and ipv6
- IPv6 ready.
- Multiple serial port(s) are supported.
  - Allows splitting of loops, reducing data acquisition times.

Subpanel networking

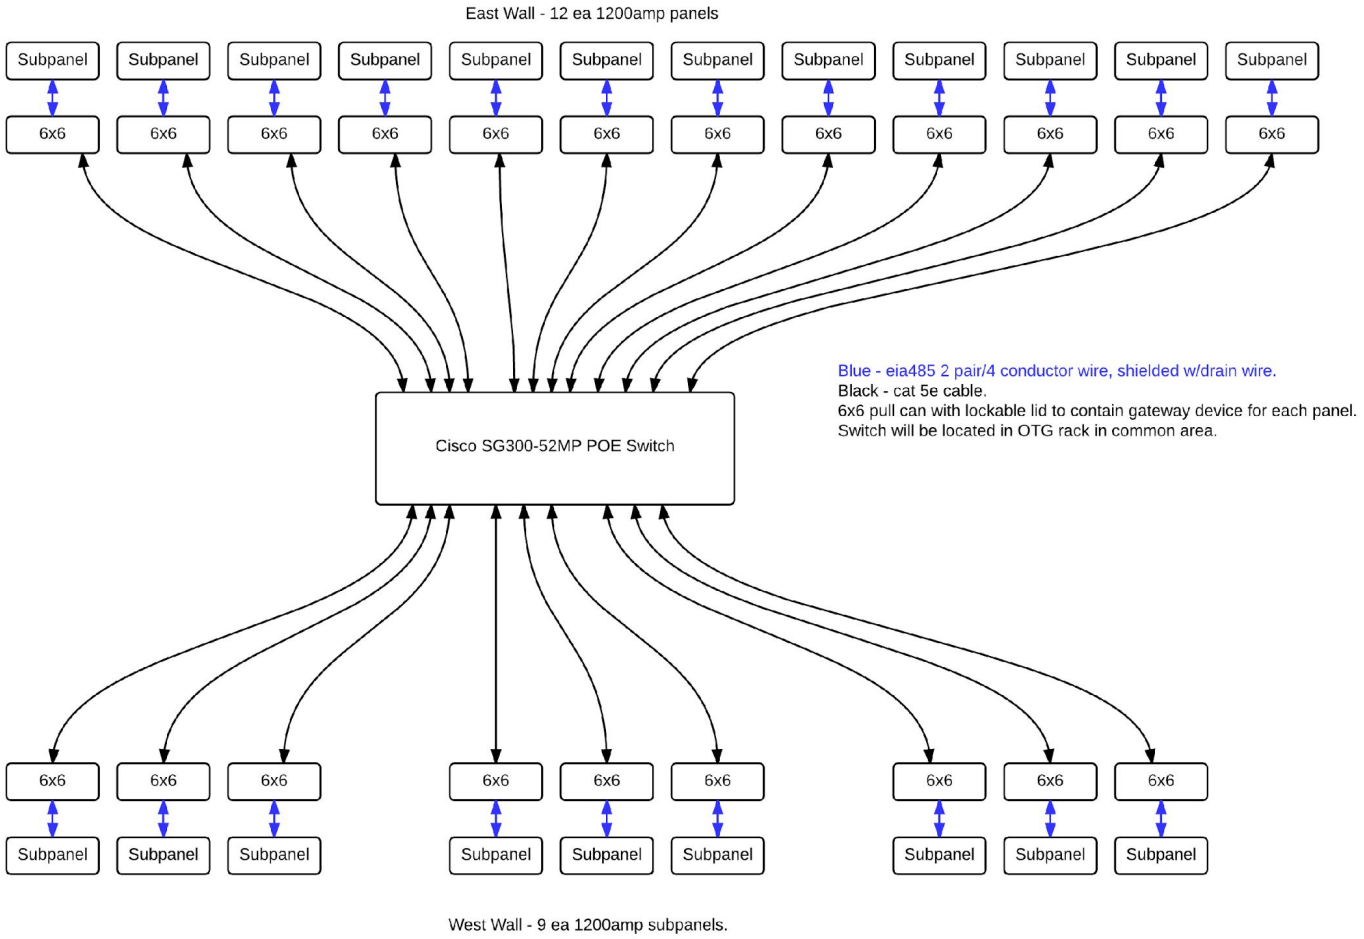

Smart Modbus Gateways

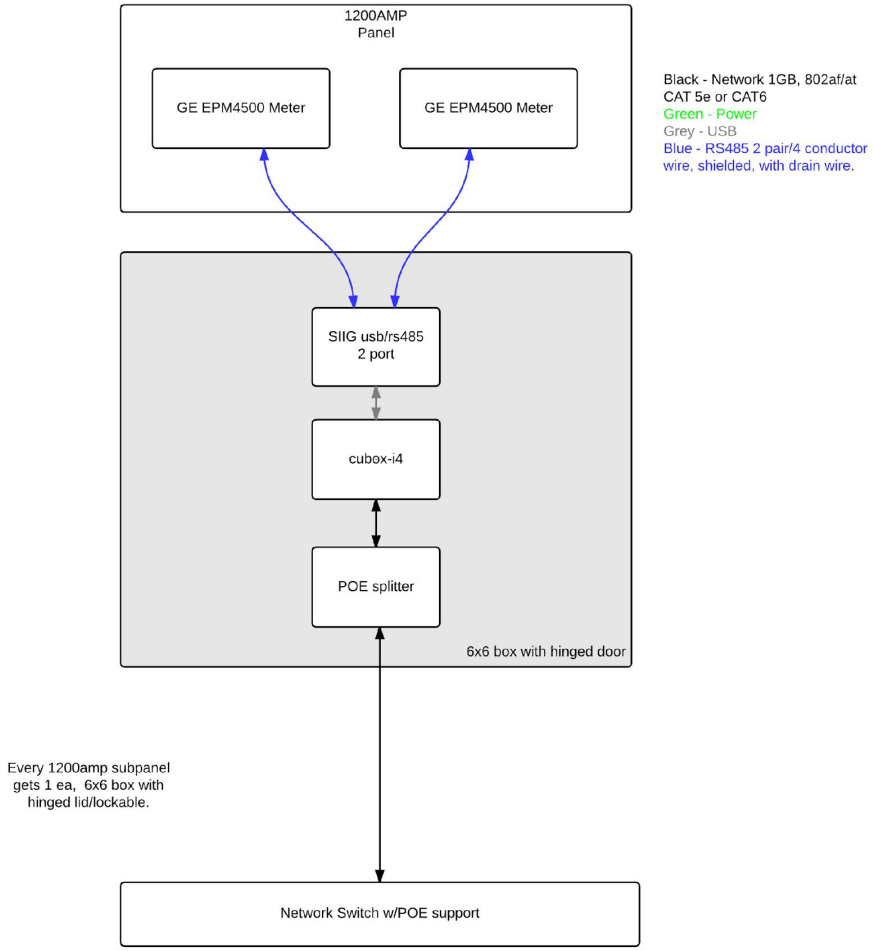

Data Flow in system

Logstash does conversion of logs, collectd information into JSON and inserts into RabbitMQ/AMQP.

Glue scripts inserts collected data in JSON format to RabbitMQ/AMQP.

Performance data is also collected and inserted to RabbitMQ.

Logstash is then used to pull information from RabbitMQ/AMQP and inserted into Elasticsearch

Also creates the appropriate index(es) in Elasticsearch.

Centerwide monitoring.

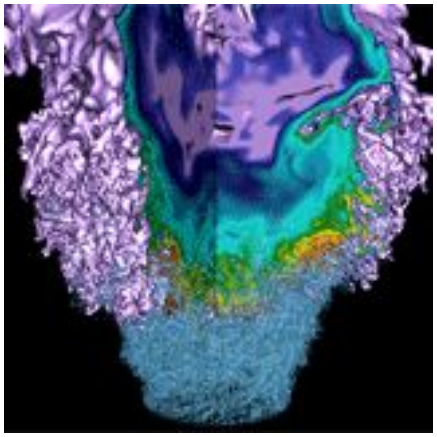

NERSC

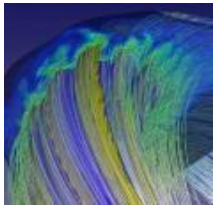
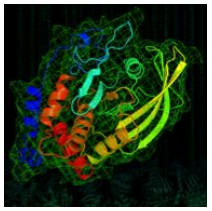
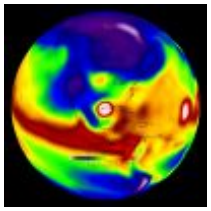
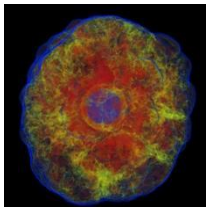
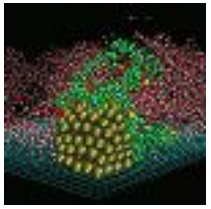

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Centerwide monitoring

Installation of Collectd on Cori/Hopper nodes.

- IB, Lustre client, Lustre MDS, Lustre OSS plugins
  - All written in C instead of using existing Perl plugin for efficiency.
  - Lustre server side modules support jobid which is batch job related stats
- Also used to collect load, memory, interface, disk stats.
- Compute nodes do not get this.
  - Concerns about jitter and memory pressure.

Local logstash instance.

- Syslog forwarding to this instance.
- Also does collectd udp to json/tcp/amqp conversion for us.
- Local grok filters in this instance.
- AMQP/SSL into the rabbitmq cluster.
  - Not everything is HA'd at this point.