How-to write a xtpmd_plugin for your Cray® XC™ system

Steven J. Martin (stevem@cray.com)

CUG 2017 CAFFEINATED COMPUTING

Redmond, Washington May 7-11, 2017

Cray XC Telemetry Plugin Introduction

- Enabling sites to get telemetry data off the Cray
- Plugin interface enables site specific customization
- This presentation and paper dive into details

> - I plan to move through the slides quickly
> - And have time for question and discussion at the end

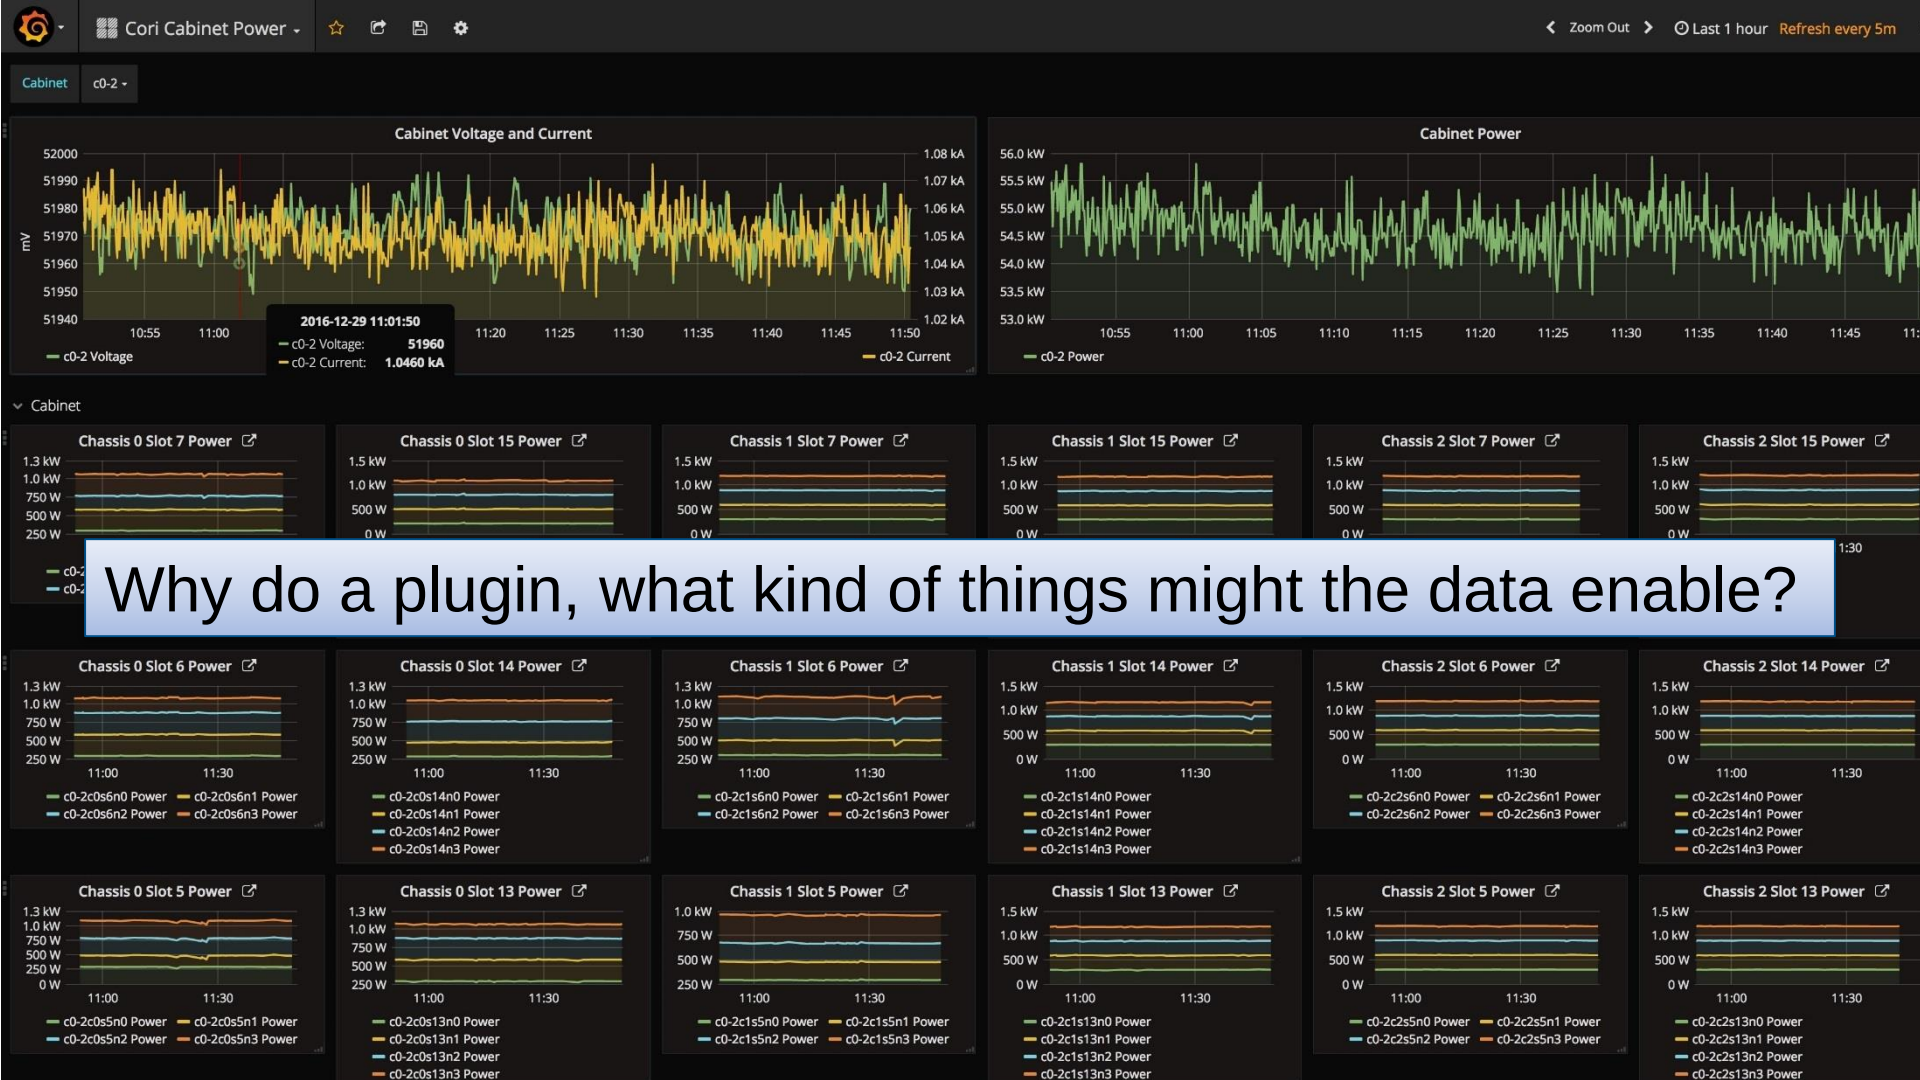
Cori Cabinet Power
Zoom Out
Last 1 hour
Refresh every 5m
Cabinet
c0-2
Cabinet Voltage and Current
mV
1.08 kA
1.07 kA
1.06 kA
1.05 kA
1.04 kA
1.03 kA
1.02 kA
2016-12-29 11:01:50
c0-2 Voltage: 51960
c0-2 Current: 1.0460 kA
c0-2 Voltage
c0-2 Current
Cabinet Power
c0-2 Power
Cabinet
Chassis 0 Slot 7 Power
Chassis 0 Slot 15 Power
Chassis 1 Slot 7 Power
Chassis 1 Slot 15 Power
Chassis 2 Slot 7 Power
Chassis 2 Slot 15 Power
Why do a plugin, what kind of things might the data enable?
Chassis 0 Slot 6 Power
Chassis 0 Slot 14 Power
Chassis 1 Slot 6 Power
Chassis 1 Slot 14 Power
Chassis 2 Slot 6 Power
Chassis 2 Slot 14 Power
Chassis 0 Slot 5 Power
Chassis 0 Slot 13 Power
Chassis 1 Slot 5 Power
Chassis 1 Slot 13 Power
Chassis 2 Slot 5 Power
Chassis 2 Slot 13 Power

Customer Driven Plugin History

- **Xtreme System Monitoring Collaboration meeting, Jan 28, 2016**
  - Requested access to SEDC and Cray high-speed power and energy data streams before data is injected into the Cray PMDB
- **Cray previewed the plugin feature at CUG 2016 Slides 59-74**
- **Plugin using Redis Pub/Sub transport (Hiredis C-library) demonstrated at Trinity phase-2 factory trial**
- **Example plugin code ships with (SMW) 8.0 and newer software**
- **NERSC now using plugin in production on Cori system**

Streaming Data Available From Five Sources

- **Cabinet SEDC data**
  - xtgetsedcvalues -l -t cc -c c0-0
- **Blade SEDC data**
  - xtgetsedcvalues -l -t bc -c c1-0c1s3
- **Cabinet power and energy data**
  - xtgetsedcvalues -l -t cc_power -c c0-0
- **Blade power and energy data**
  - xtgetsedcvalues -l -t bc_power | grep c0-0c0s8
- **Job and application data**

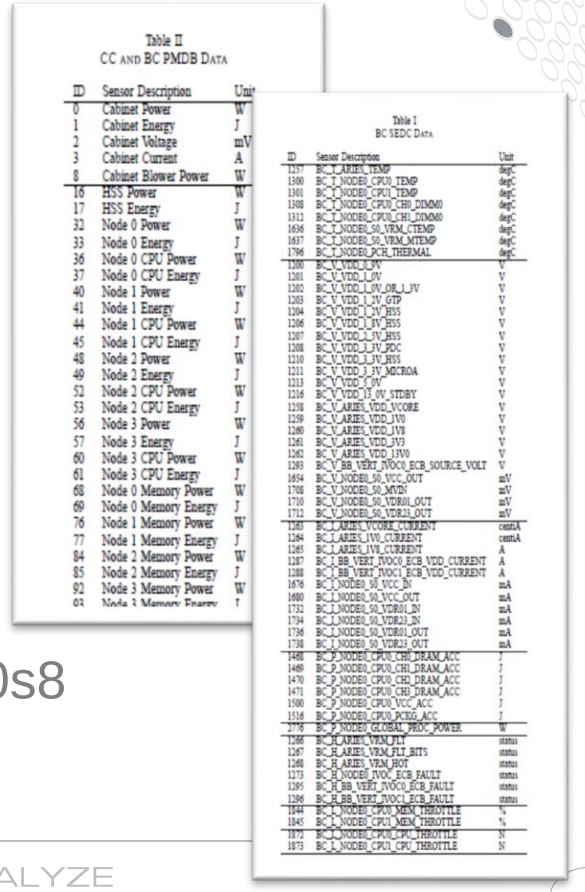

Table II
CC AND BC PMDB DATA

ID	Sensor Description	Unit
0	Cabinet Power	W
1	Cabinet Energy	J
2	Cabinet Voltage	mV
3	Cabinet Current	A
8	Cabinet Blower Power	W
16	HSS Power	W
17	HSS Energy	J
32	Node 0 Power	W
33	Node 0 Energy	J
36	Node 0 CPU Power	W
37	Node 0 CPU Energy	J
40	Node 1 Power	W
41	Node 1 Energy	J
44	Node 1 CPU Power	W
45	Node 1 CPU Energy	J
48	Node 2 Power	W
49	Node 2 Energy	J
52	Node 2 CPU Power	W
53	Node 2 CPU Energy	J
56	Node 3 Power	W
57	Node 3 Energy	J
60	Node 3 CPU Power	W
61	Node 3 CPU Energy	J
68	Node 0 Memory Power	W
69	Node 0 Memory Energy	J
76	Node 1 Memory Power	W
77	Node 1 Memory Energy	J
84	Node 2 Memory Power	W
85	Node 2 Memory Energy	J
92	Node 3 Memory Power	W
93	Node 3 Memory Energy	J

Table I
BC SEDC DATA

ID	Sensor Description	Unit
1257	BC_I_ARIES_TEMP	degC
1300	BC_I_NODE0_CPU0_TEMP	degC
1301	BC_I_NODE0_CPU1_TEMP	degC
1308	BC_I_NODE0_CPU0_CH0_DIMM0	degC
1312	BC_I_NODE0_CPU0_CH1_DIMM0	degC
1636	BC_I_NODE0_S0_VRM_CTEMP	degC
1637	BC_I_NODE0_S0_VRM_MTEMP	degC
1796	BC_I_NODE0_PCH_THERMAL	degC
1200	BC_V_VDD_1_0V	V
1201	BC_V_VDD_1_0V	V
1202	BC_V_VDD_1_0V_OR_1_1V	V
1203	BC_V_VDD_1_2V_GTP	V
1204	BC_V_VDD_1_2V_HSS	V
1206	BC_V_VDD_1_8V_HSS	V
1207	BC_V_VDD_2_5V_HSS	V
1208	BC_V_VDD_3_3V_PDC	V
1210	BC_V_VDD_3_3V_HSS	V
1211	BC_V_VDD_3_3V_MICROA	V
1213	BC_V_VDD_5_0V	V
1216	BC_V_VDD_13_0V_STDBY	V
1258	BC_V_ARIES_VDD_VCORE	V
1259	BC_V_ARIES_VDD_1V0	V
1260	BC_V_ARIES_VDD_1V8	V
1261	BC_V_ARIES_VDD_3V3	V
1262	BC_V_ARIES_VDD_13V0	V
1263	BC_V_BB_VERT_IVOC0_ECB_SOURCE_VOLT	V
1654	BC_V_NODE0_S0_VCC_OUT	mV
1708	BC_V_NODE0_S0_MVIN	mV
1710	BC_V_NODE0_S0_VDR01_OUT	mV
1712	BC_V_NODE0_S0_VDR23_OUT	mV
1263	BC_I_ARIES_VCORE_CURRENT	centiA
1264	BC_I_ARIES_1V0_CURRENT	centiA
1265	BC_I_ARIES_1V8_CURRENT	A
1267	BC_I_BB_VERT_IVOC0_ECB_VDD_CURRENT	A
1268	BC_I_BB_VERT_IVOC1_ECB_VDD_CURRENT	A
1678	BC_I_NODE0_S0_VCC_IN	mA
1680	BC_I_NODE0_S0_VCC_OUT	mA
1732	BC_I_NODE0_S0_VDR01_IN	mA
1734	BC_I_NODE0_S0_VDR23_IN	mA
1736	BC_I_NODE0_S0_VDR01_OUT	mA
1738	BC_I_NODE0_S0_VDR23_OUT	mA
1468	BC_P_NODE0_CPU0_CH0_DRAM_ACC	J
1469	BC_P_NODE0_CPU0_CH1_DRAM_ACC	J
1470	BC_P_NODE0_CPU0_CH2_DRAM_ACC	J
1471	BC_P_NODE0_CPU0_CH3_DRAM_ACC	J
1500	BC_P_NODE0_CPU0_VCC_ACC	J
1516	BC_P_NODE0_CPU0_PCKG_ACC	J
2776	BC_P_NODE0_GLOBAL_PROC_POWER	W
1266	BC_H_ARIES_VRM_FLT	status
1267	BC_H_ARIES_VRM_FLT_BITS	status
1268	BC_H_ARIES_VRM_HOT	status
1273	BC_H_NODE0_IVOC_ECB_FAULT	status
1295	BC_H_BB_VERT_IVOC0_ECB_FAULT	status
1296	BC_H_BB_VERT_IVOC1_ECB_FAULT	status
1844	BC_I_NODE0_CPU0_MEM_THROTTLE	%
1845	BC_I_NODE0_CPU1_MEM_THROTTLE	%
1872	BC_I_NODE0_CPU0_CPU_THROTTLE	N
1873	BC_I_NODE0_CPU1_CPU_THROTTLE	N

Job and Application Data

- **Job and application information**
  - Job ID, app ID, user ID, timestamp
  - Assigned nodes (Cray "nid")

Used by ALPS and Slurm

```
enum {
  APP_CMD_TYPE_START =     1,
  APP_CMD_TYPE_END =       2,
  APP_CMD_TYPE_SYNC =      3,
  APP_CMD_TYPE_SUSPEND =   4,
  APP_CMD_TYPE_RESUME =    5,
  JOB_CMD_TYPE_START =     6,
  JOB_CMD_TYPE_END =       7,
  JOB_CMD_TYPE_SUSPEND =   8,
  JOB_CMD_TYPE_RESUME =    9
};
```

Used by ALPS at reservation time

```
{
  "ts": "2017-03-23T16:04:28.366738Z",
  "event": "APP_START",
  "userid": 27216,
  "job_id": "1723832.sdb",
  "apid": 3801382,
  "nid_count": 2,
  "nid_cname_array": [
    { "nid": 56, "cname": "c0-0c0s14n0" },
    { "nid": 57, "cname": "c0-0c0s14n1" }
  ]
}
```

ALPS Example (Raw Format)

- **Job and application information**
  - Job ID, app ID, user ID, timestamp
  - Assigned nodes (or Cray “nid”)

```
enum {
   APP_CMD_TYPE_START =     1,
   APP_CMD_TYPE_END =       2,
   APP_CMD_TYPE_SYNC =      3,
   APP_CMD_TYPE_SUSPEND =   4,
   APP_CMD_TYPE_RESUME =    5,
   JOB_CMD_TYPE_START =     6,
   JOB_CMD_TYPE_END =       7,
   JOB_CMD_TYPE_SUSPEND =   8,
   JOB_CMD_TYPE_RESUME =    9
};
```

```
ts=1490734709583873,event=6,userid=1205,jobid=3880.sdb,apid=0,nids=',2'
ts=1490734713023993,event=1,userid=1205,jobid=3880.sdb,apid=706691,nids=',2'
ts=1490734718511193,event=2,userid=1205,jobid=3880.sdb,apid=706691
ts=1490734719788394,event=7,userid=1205,jobid=3880.sdb,apid=0
```

ALPS Example (JSON Format)


```
{"ts":"2017-03-31T20:12:32.601051Z","event":"JOB_START",
  "userid":7821,"job_id":"1745267.sdb","apid":0,"nid_count":1,
  "nid_cname_array":[{"nid":76,"cname":"c0-0c1s3n0"}]}

{"ts":"2017-03-31T20:12:39.043667Z","event":"APP_START",
 "userid":7821,"job_id":"1745267.sdb","apid":3860100,"nid_count":1,
 "nid_cname_array":[{"nid":76,"cname":"c0-0c1s3n0"}]}

{"ts":"2017-03-31T20:12:50.091905Z","event":"APP_END",
  "userid":7821,"job_id":"1745267.sdb","apid":3860100}

{"ts":"2017-03-31T20:12:51.102143Z","event":"JOB_END",
  "userid":7821,"job_id":"1745267.sdb","apid":0}
```

Slurm Example

```
system:~> salloc -N 4
salloc: Granted job allocation 369936
salloc: Waiting for resource configuration
salloc: Nodes nid00[122-125] are ready for job
system:~> srun -N 2 hostname
nid00123
nid00122
system:~> srun -N 4 hostname
nid00124
nid00125
nid00123
nid00122
system:~> exit
exit
salloc: Relinquishing job allocation 3
```

1

2

```
{"ts":"2017-04-20T19:56:43.595141Z",
    "event":"APP_START", "userid":26914, "job_id":"369936",
    "apid":6056184802581266704, "nid_count":4,
    "nid_cname_array":[{"nid":122,"cname":"c0-0c1s14n2"},
                       {"nid":123,"cname":"c0-0c1s14n3"},
                       {"nid":124,"cname":"c0-0c1s15n0"},
                       {"nid":125,"cname":"c0-0c1s15n1"}]}
```

```
{"ts":"2017-04-20T19:57:10.623100Z","event":"APP_START",
    "userid":26914, "job_id":"369936",
    "apid":369936, "nid_count":2,
    "nid_cname_array":[{"nid":122,"cname":"c0-0c1s14n2"},
                       {"nid":123,"cname":"c0-0c1s14n3"}]}

{"ts":"2017-04-20T19:57:11.543174Z","event":"APP_END",
    "userid":26914,"job_id":"369936","apid":369936}
```

Slurm Example

```
system:~> salloc -N 4
salloc: Granted job allocation 369936
salloc: Waiting for resource configuration
salloc: Nodes nid00[122-125] are ready for job
system:~> srun -N 2 hostname
nid00123
nid00122
system:~> srun -N 4 hostname
nid00124
nid00125
nid00123
nid00122
system:~> exit
exit
salloc: Relinquishing job allocation 3
```

3

```
{"ts":"2017-04-20T19:57:30.087104Z","event":"APP_START",
    "userid":26914, "job_id":"369936",
    "apid":10000369936, "nid_count":4,
    "nid_cname_array":[{"nid":122,"cname":"c0-0c1s14n2"},
                       {"nid":123,"cname":"c0-0c1s14n3"},
                       {"nid":124,"cname":"c0-0c1s15n0"},
                       {"nid":125,"cname":"c0-0c1s15n1"}]}

{"ts":"2017-04-20T19:57:31.007176Z","event":"APP_END",
    "userid":26914, "job_id":"369936",
    "apid":10000369936}
```

4

```
{"ts":"2017-04-20T19:57:35.783165Z","event":"APP_END",
    "userid":26914, "job_id":"369936",
    "apid":60561848025812667 04}
```

Slurm Example (Closer look at steps 1 and 4)

```
system:~> salloc -N 4
salloc: Granted job allocation 369936
salloc: Waiting for resource configuration
salloc: Nodes nid00[122-125] are ready for job
…
system:~> exit
exit
salloc: Relinquishing job allocation 369936
```

1

```
{"ts":"2017-04-20T19:56:43.595141Z","event":"APP_START",
    "userid":26914, "job_id":"369936","apid":60561848025812667O4,
    "nid_count":4,"nid_cname_array":[
         {"nid":122,"cname":"c0-0c1s14n2"},{"nid":123,"cname":"c0-0c1s14n3"},
         {"nid":124,"cname":"c0-0c1s15n0"},{"nid":125,"cname":"c0-0c1s15n1"}]}
...
{"ts":"2017-04-20T19:57:35.783165Z","event":"APP_END",
    "userid":26914, "job_id":"369936","apid":60561848025812667O4}
```

4

Plugin Configuration File

- **/opt/cray/hss/default/etc/xtpmd_plugins.ini**
  - Configuration file the in smw release supports the plugin_csv example
  - **smw:/opt/cray/hss/default/pm/xtpmd_api**

```
smw:/opt/cray/hss/default/etc> egrep "=|\[" xtpmd_plugins.ini | grep -v _path
[plugins]
# shmsize=4194304
# instances=csv;other
# instances=csv
[plugin_csv]
object=/opt/cray/hss/default/lib64/xtpmd_plugin_csv.so
log_dir=/tmp
pmdb_cc_enabled=yes
…
```

Makefile Example

```
# Makefile for xtpmd_xjson plugin
OBJ = xtpmd_plugin_xjson.o
LIB = xtpmd_plugin_xjson.so

CFLAGS  += -O3 -fPIC
CFLAGS  += $(shell pkg-config --cflags glib-2.0 gthread-2.0 libpq)
LDFLAGS += $(shell pkg-config --libs glib-2.0 gthread-2.0 libpq)
LDFLAGS += -lm -lz

ALL: xtpmd_plugin_xjson.so

xtpmd_plugin_xjson.o: xtpmd_plugin_xjson.c xtpmd_plugin.h
        $(CC) $(CPPFLAGS) $(CFLAGS) -c -o $@ $<

xtpmd_plugin_xjson.so: xtpmd_plugin_xjson.o
        $(CC) -shared $< -o $@ $(LDFLAGS)
```

Starting a Test Plugin Unsupervised

```
smw:~> ipcs | grep -v post
------ Message Queues --------
key        msqid   owner     perms   used-bytes    messages
------ Shared Memory Segments --------
key        shmid   owner     perms   bytes     nattch      status
0x7a060809 32769   crayadm   600     4194304   1
------ Semaphore Arrays --------
key        semid   owner     perms   nsems

smw:~> xtpmd_plugd 0x7a060809 4194304 plugin_xjson ./xtpmd_plugins.ini
```

Starting a Test Plugin Unsupervised

```
smw:~> ipcs | grep -v post
------ Message Queues --------
key        msqid   owner     perms   used-byt
------ Shared Memory Segments --
key        shmid   owner                    nattch     status
0x7a060809 32769                            194304   1
------ Se                    ----
key                owner     perms   nsems

smw:~> xtpmd_plugd 0x7a060809 4194304 plugin_csv ./xtpmd_plugins.ini
```

Note: Use 'Ctrl-D' to kill the test plugin, or kill -9 from another window…

- 'Ctrl-C' is blocked.

Design Considerations Covered in the Paper

- Limiting the plugin's impact
- Library usage
- Time stamp formatting
- Translating binary fields
- Getting data off node
- Other formatting considerations

NERSC Plugins

- **Plugin data sent into NERSC Center-wide Data Collect**
  - https://cug.org/proceedings/cug2016_proceedings/includes/files/pap101.pdf
- **Implemented 5 plugins one for each data source**
  - instances = bp ; bc ; cp ; cc ; j o b
- **RabbitMQ is used to send data off the SMW**
  - Connection details in the configuration file
  - Changes can be made without recompiling code
- **Data written in JSON format**
  - Significantly increases Elastic ingest rate

Cori Problem Information Gathering

The next 5 slides are graphics from the paper

Collected looking into conditions at the time of a reported thermal throttling event

Cname c4-2

Chassis 0

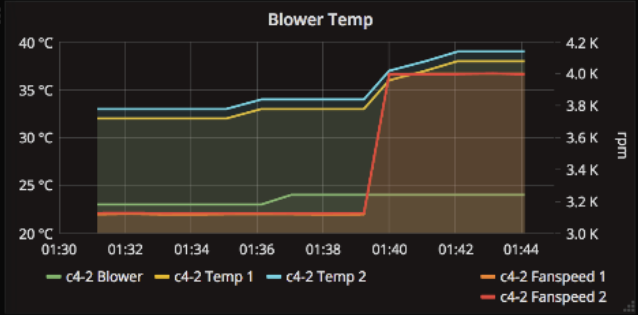

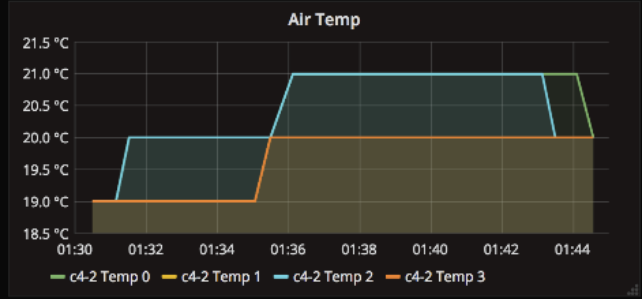

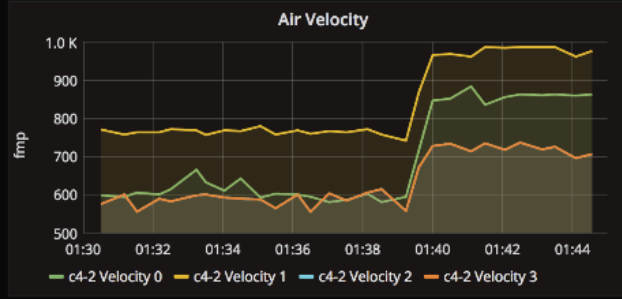

Chassis 1

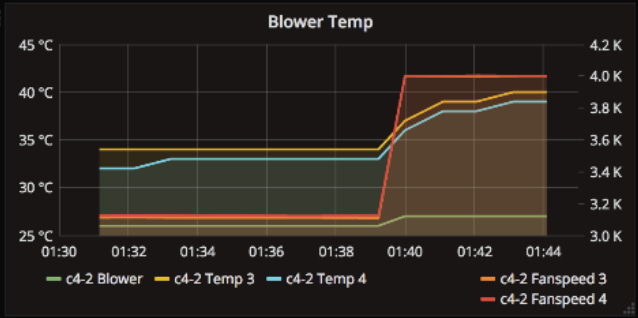

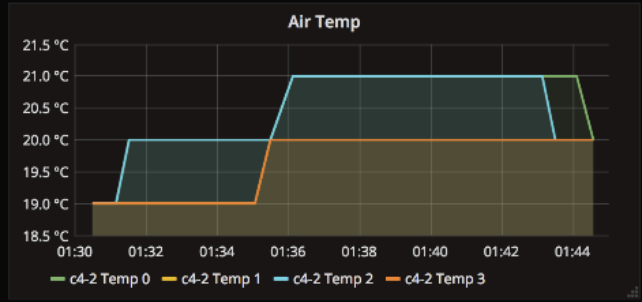

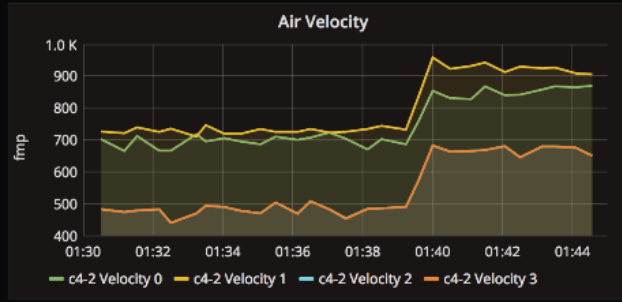

Chassis 2

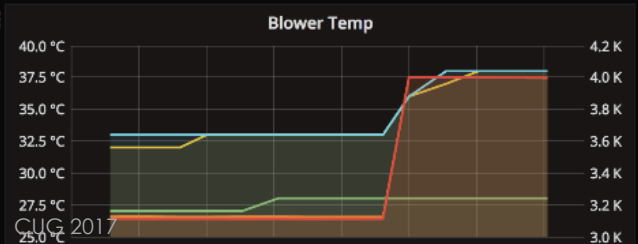

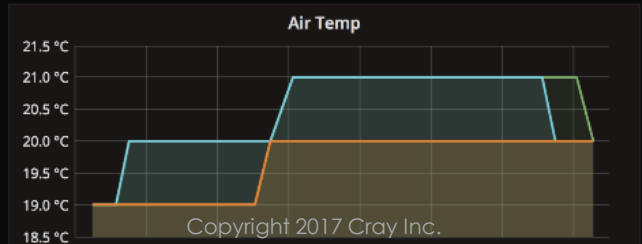

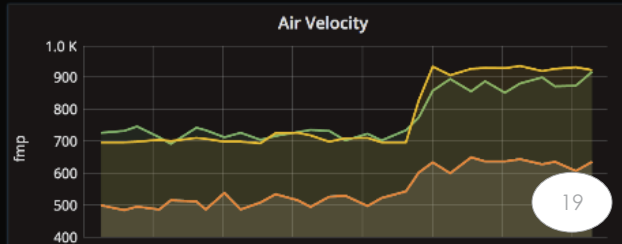

CName c4-2

Cabinet Water Temp

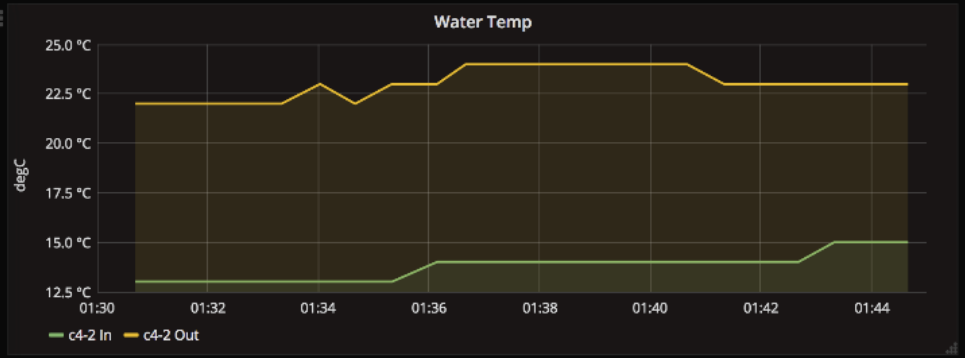

Cabinet Water Pressure

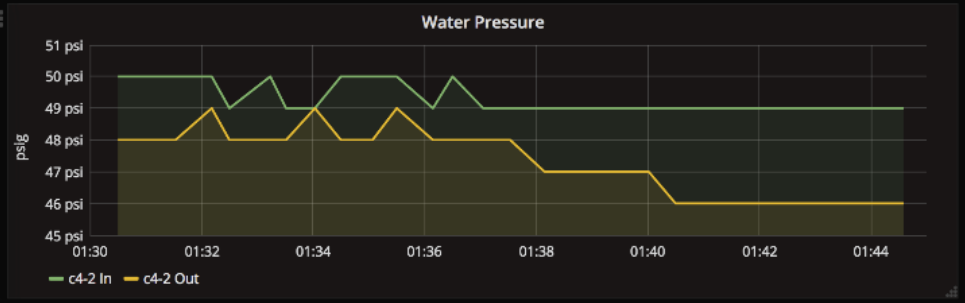

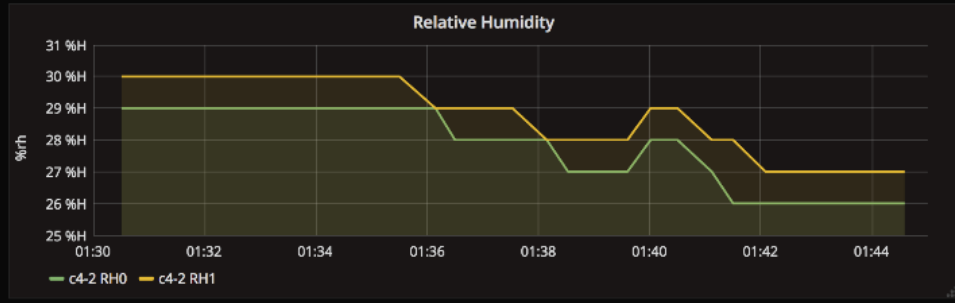

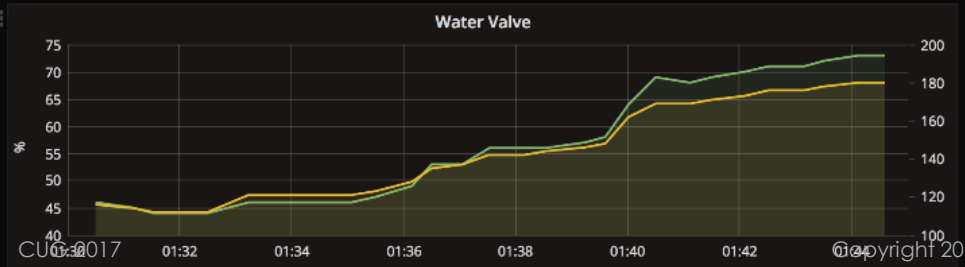

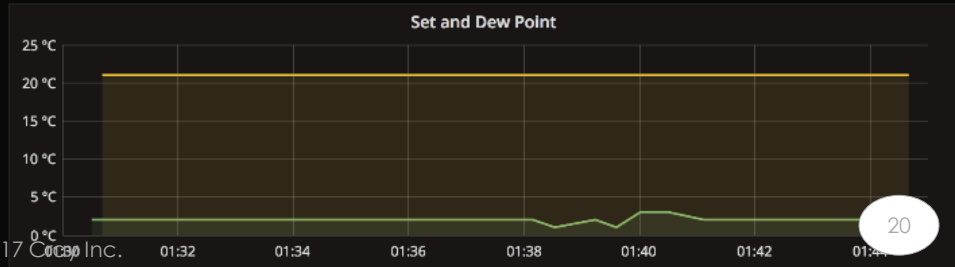

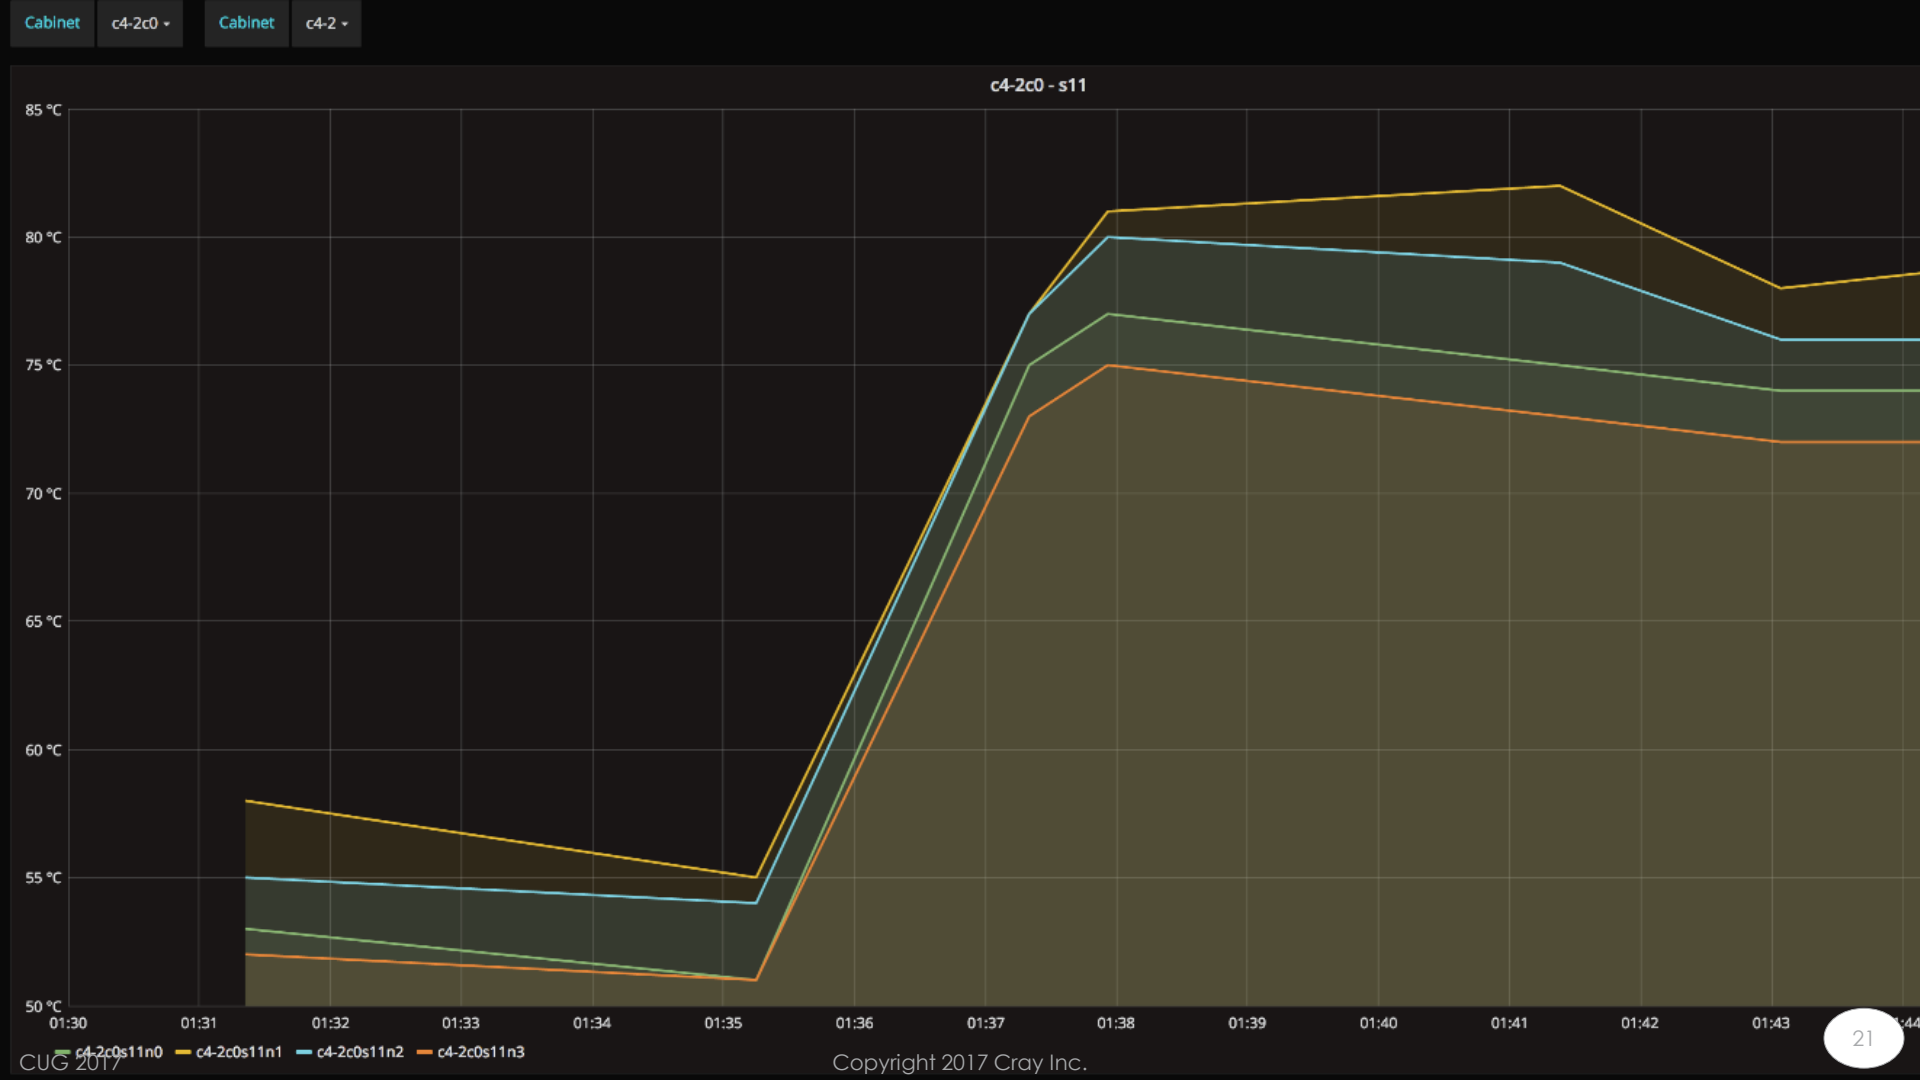

Cabinet c4-2 Chassis c0 Slot s11

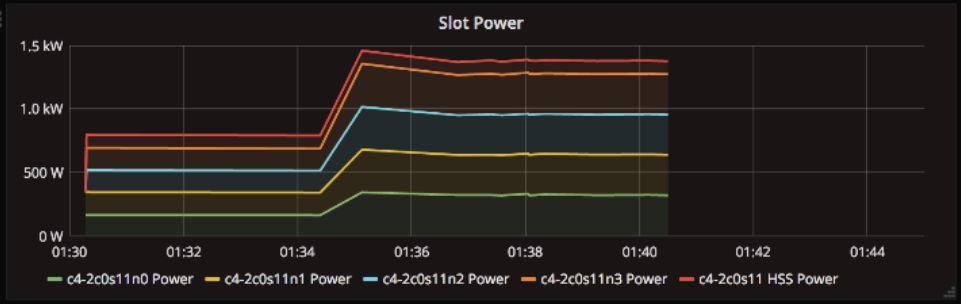

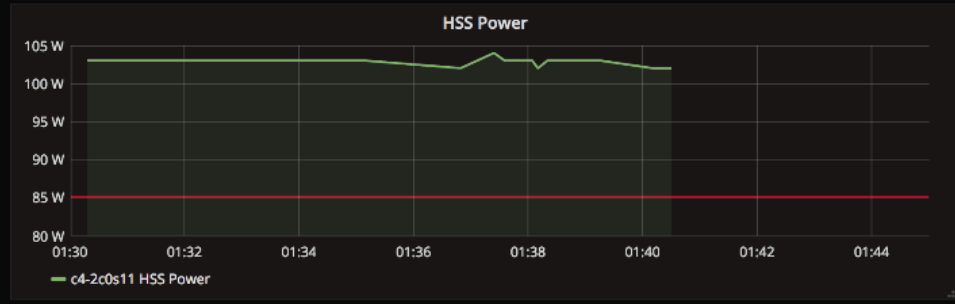

Slot

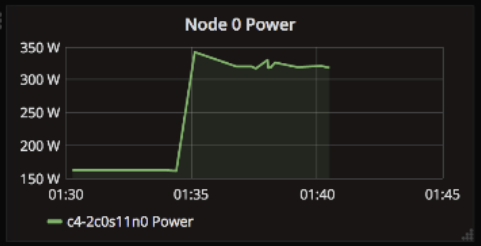

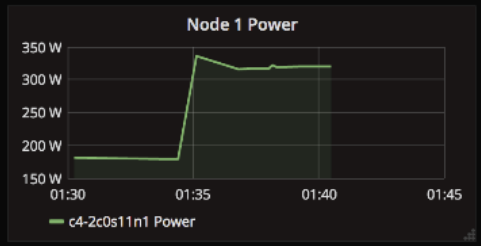

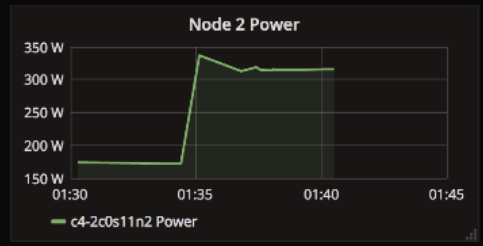

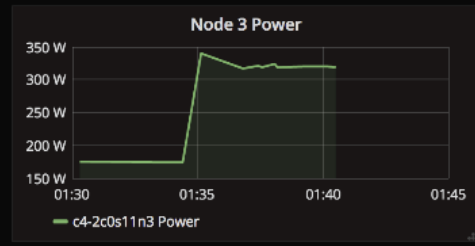

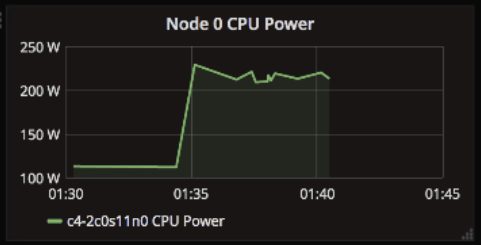

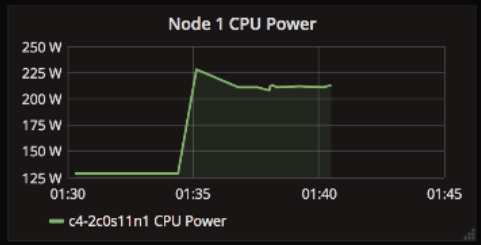

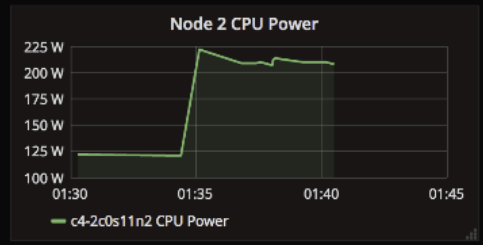

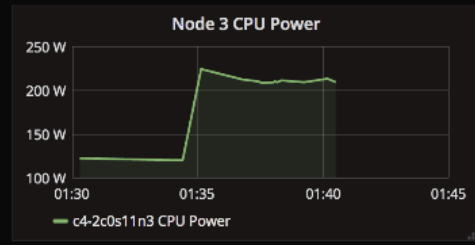

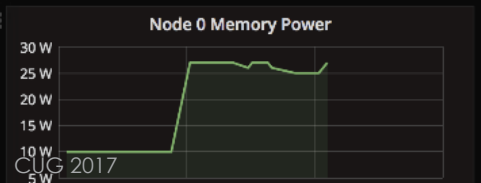

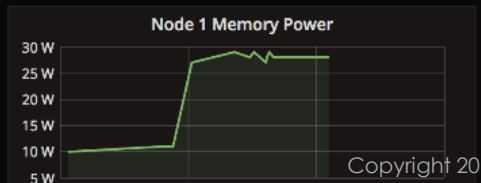

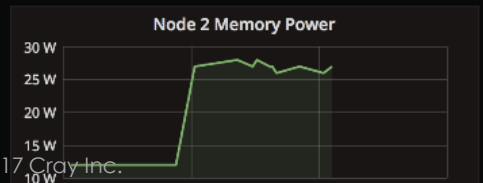

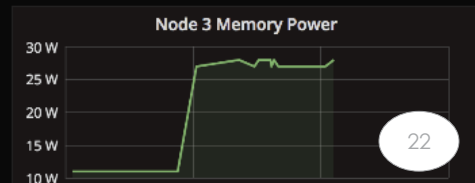

Row 0

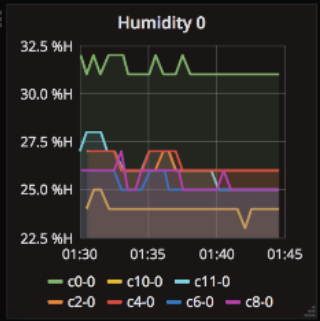

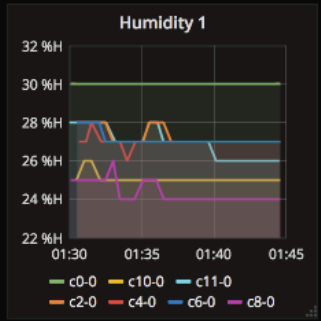

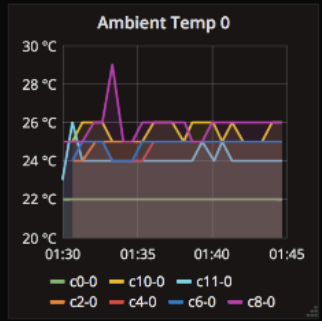

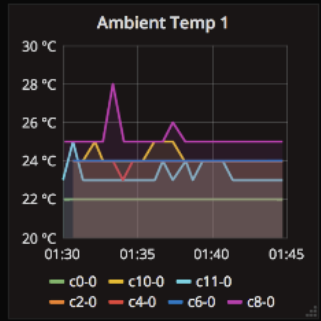

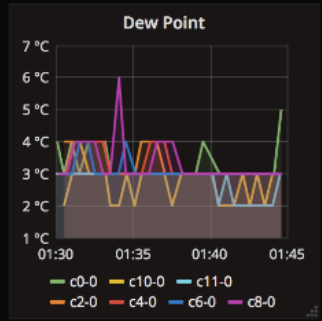

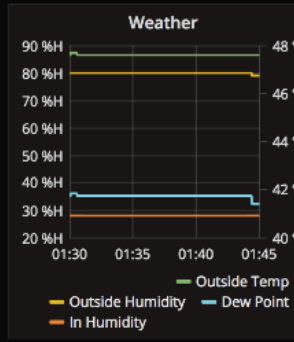

Row 1

Row 2

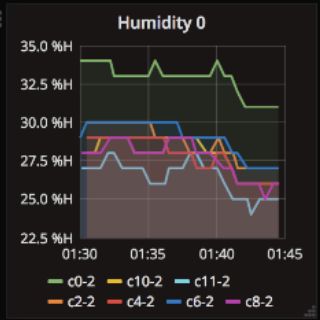

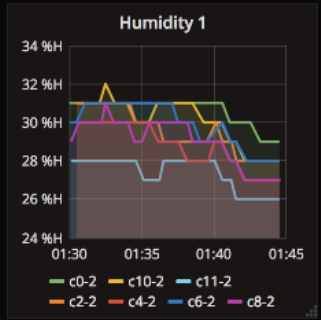

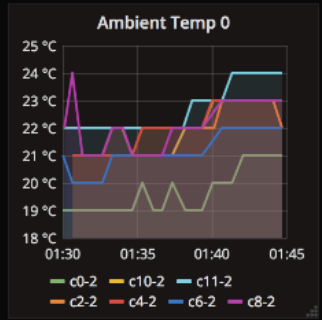

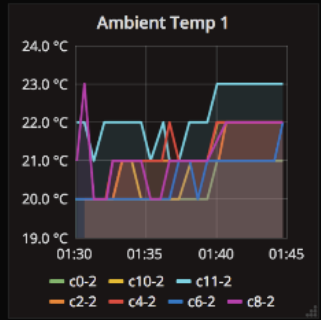

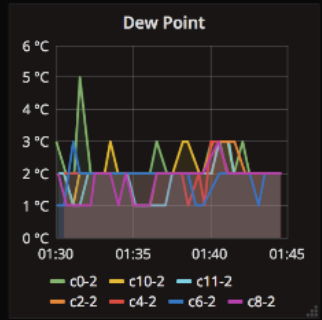

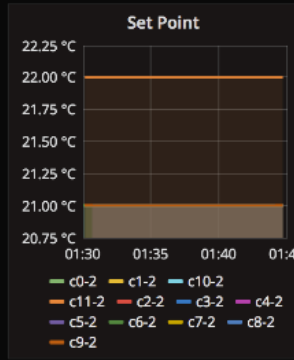

Row 3

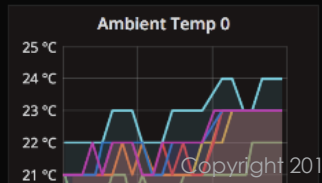

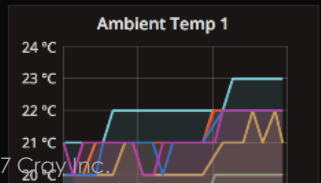

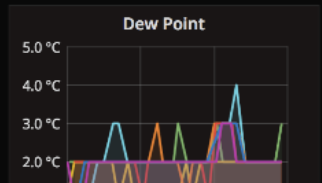

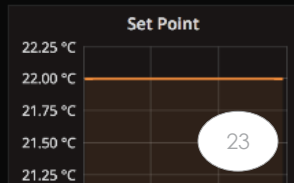

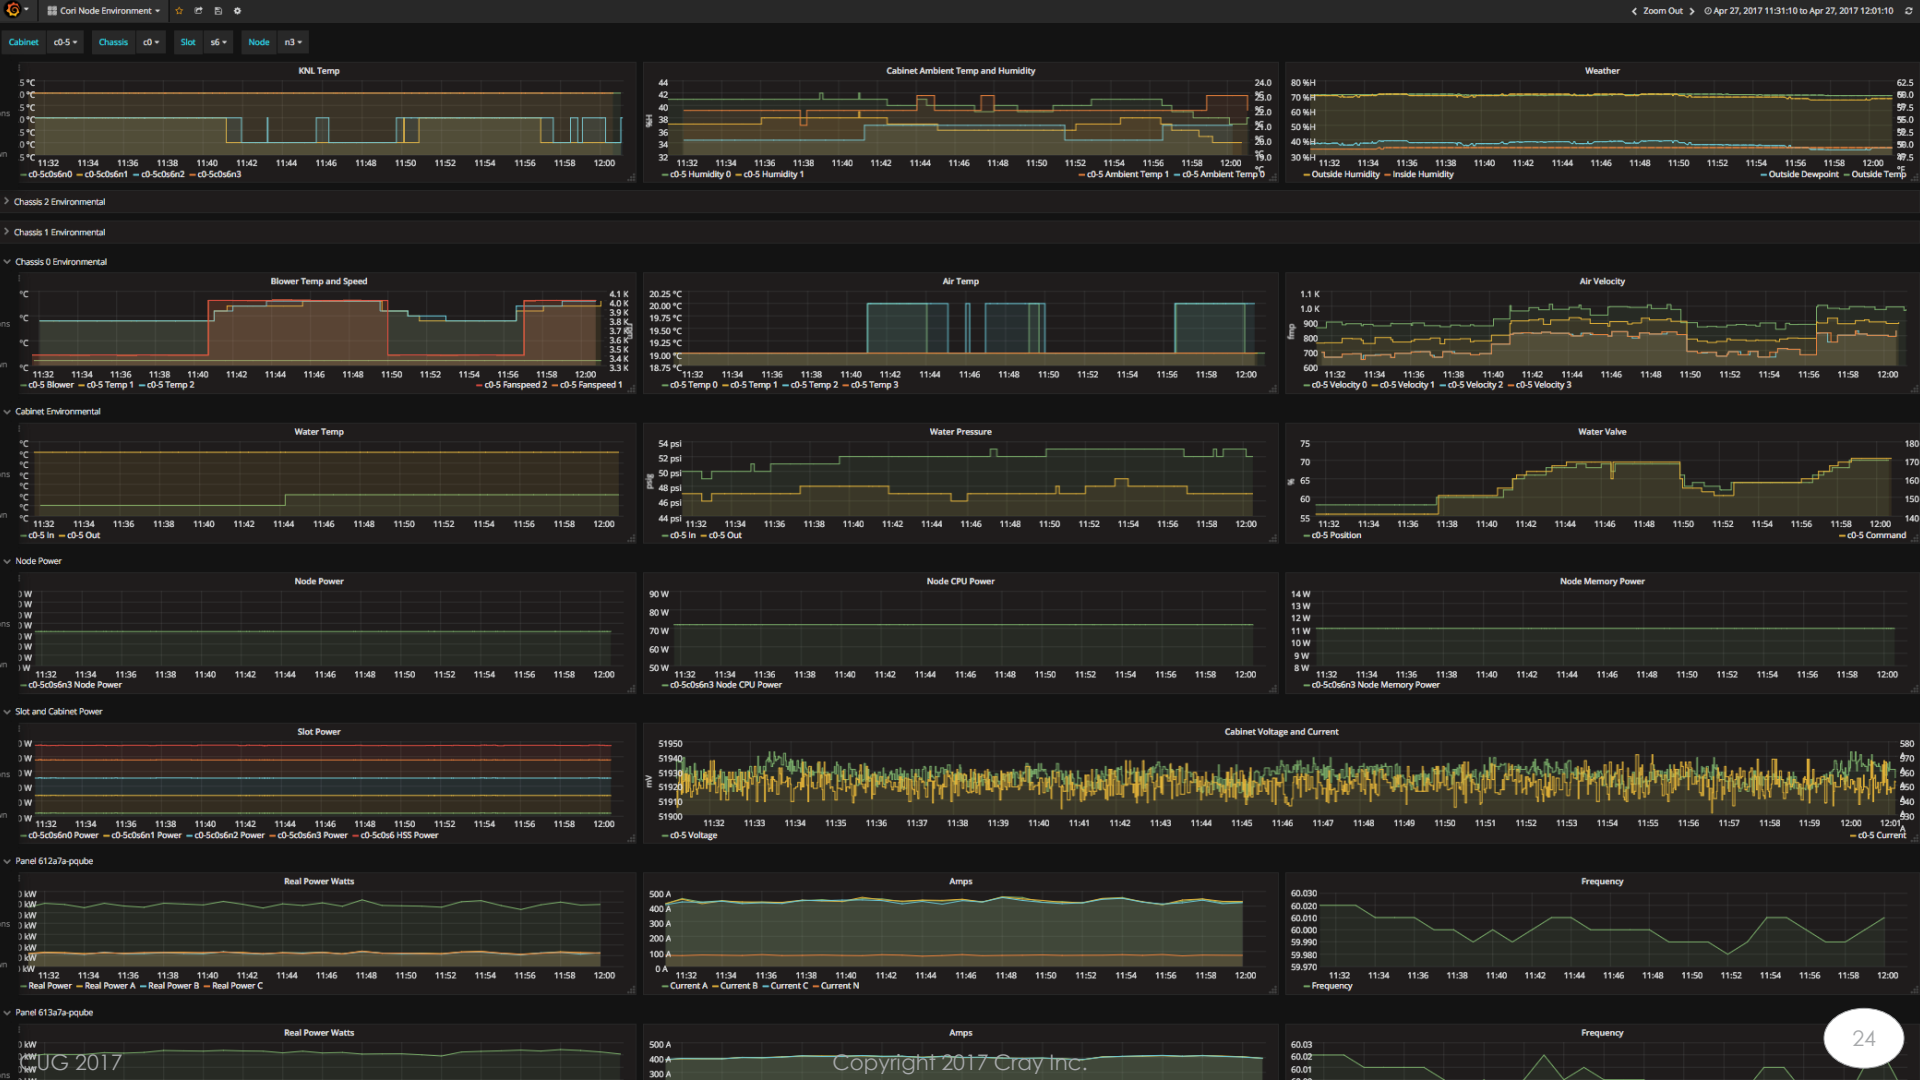

Wrap-up

- **Collaborative effort**
  - Cray, customers, and the user community
- **Allows streaming of**
  - Power, energy, thermal, and application meta-data
- **Making that data available**
  - System administrators
  - Application developers
  - HPC research community
- **As appropriate given site-level policy**

Acknowledgment

- **NERSC**
  - Specifically Cary Whitney for co-authoring this paper
  - NERSC is leading the way in use of this new plugin feature

- **XTreme**
  - Large system customer support is driving changes
  - See also: monitoring_wg@lists.cug.org

- **HPC Community**
  - Broad community push for improved monitoring capabilities
  - See also: EEHPC WG

Legal Disclaimer

Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publically announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, and URIKA. The following are trademarks of Cray Inc.: APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYPAT, CRAYPORT, ECOPHLEX, LIBSCI, NODEKARE, REVEAL,THREADSTORM. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.

Q&A

Steven Martin
stevem@cray.com

CUG.2017.CAFFEINATED COMPUTING

Redmond, Washington May 7-11, 2017