September 24, 2015

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

National Petascale Computing Facility

HPC Utility Metrics and Facility Controls

THOMAS DURBIN, P.E., LEED AP

GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

NPCF

- LEED GOLD CERTIFIED
- 2010 MARCH
- BLUE WATERS BUILD 2012
- 24 MW
- 5,400 TONS
- COOLING TOWERS
- 20,000 SQ. FT. RAF (1.858 M2)
- 87,915 SQ. FT. TTL (8.167 M2)
- 4.5 ACRES SITE (18.210 M2)
- F3 TORNADO PROOF
- DRY AND WET SPRINKLERS
- PUE 1.08 TO 1.215

BLUE WATERS

- 288 CRAY CABINETS
- 26,864 COMPUTE NODES
- >49,000 AMD CPUS
- 405,248 CPU CORES
- 4,224 NVIDIA KEPLER GPUS
- 1.5 PETABYTES RAM
- 25 PETABYTES STORAGE
- 13.34 PETAFLOPS
- 72 XDP COOLING UNITS
- >95% LIQUID COOLING
- PEAK >3,500 TONS, 13 MW
- TAPE LIBRARY 300 PETAB
- 100K MBS EXT NET SPEED

Metrics

- Metrics measure performance and activities
- Utilities and Building Systems
- Compute Systems
- Short Term and Long Term Objectives

Real Time Metrics

- https://bluewaters.ncsa.illinois.edu
- http://isce.ncsa.illinois.edu/PwrConsole.php
- http://isce.ncsa.illinois.edu/powerSummary.php


Transformer 1A
Voltage L-L Voltage L-N Power
482 278 788
Transformer 1B
Voltage L-L Voltage L-N Power
482 279 988
Transformer 1C
Voltage L-L Voltage L-N Power
482 278 846
Transformer 1D
Voltage L-L Voltage L-N Power
482 278 874
Transformer 2A
Voltage L-L Voltage L-N Power
483 279 789
Transformer 2B
Voltage L-L Voltage L-N Power
483 279 893
Transformer 2C
Voltage L-L Voltage L-N Power
482 278 898
Transformer 2D
Voltage L-L Voltage L-N Power
482 279 848
Transformer 3A
Voltage L-L Voltage L-N Power
483 279 694
Transformer 3B
Voltage L-L Voltage L-N Power
483 279 794
Transformer 3C
Voltage L-L Voltage L-N Power
482 278 748
Transformer 3D
Voltage L-L Voltage L-N Power
482 279 854
Transformer 4A
Voltage In Voltage Out Power
Transformer 4B
Voltage L-L Voltage L-N Power
483 279 48
Transformer 4C
Voltage L-L Voltage L-N Power
482 278 316
Transformer 4D
Voltage In Voltage Out Power
Transformer 5A
Voltage L-L Voltage L-N Power
483 279 125
Transformer 5B
Voltage L-L Voltage L-N Power
482 279 88
Transformer 5C
Voltage L-L Voltage L-N Power
481 278 92
Transformer 5D
Voltage L-L Voltage L-N Power
482 278 85
Power draw: 10,592 kW
Total Energy (est.)
Past 24 hours: 227.6 MWh
Past 7 days: 1,687.3 MWh
Past 30 days: 7,232.4 MWh
Hourly use by category

Meters with Wave Form Capture and Send Alerts

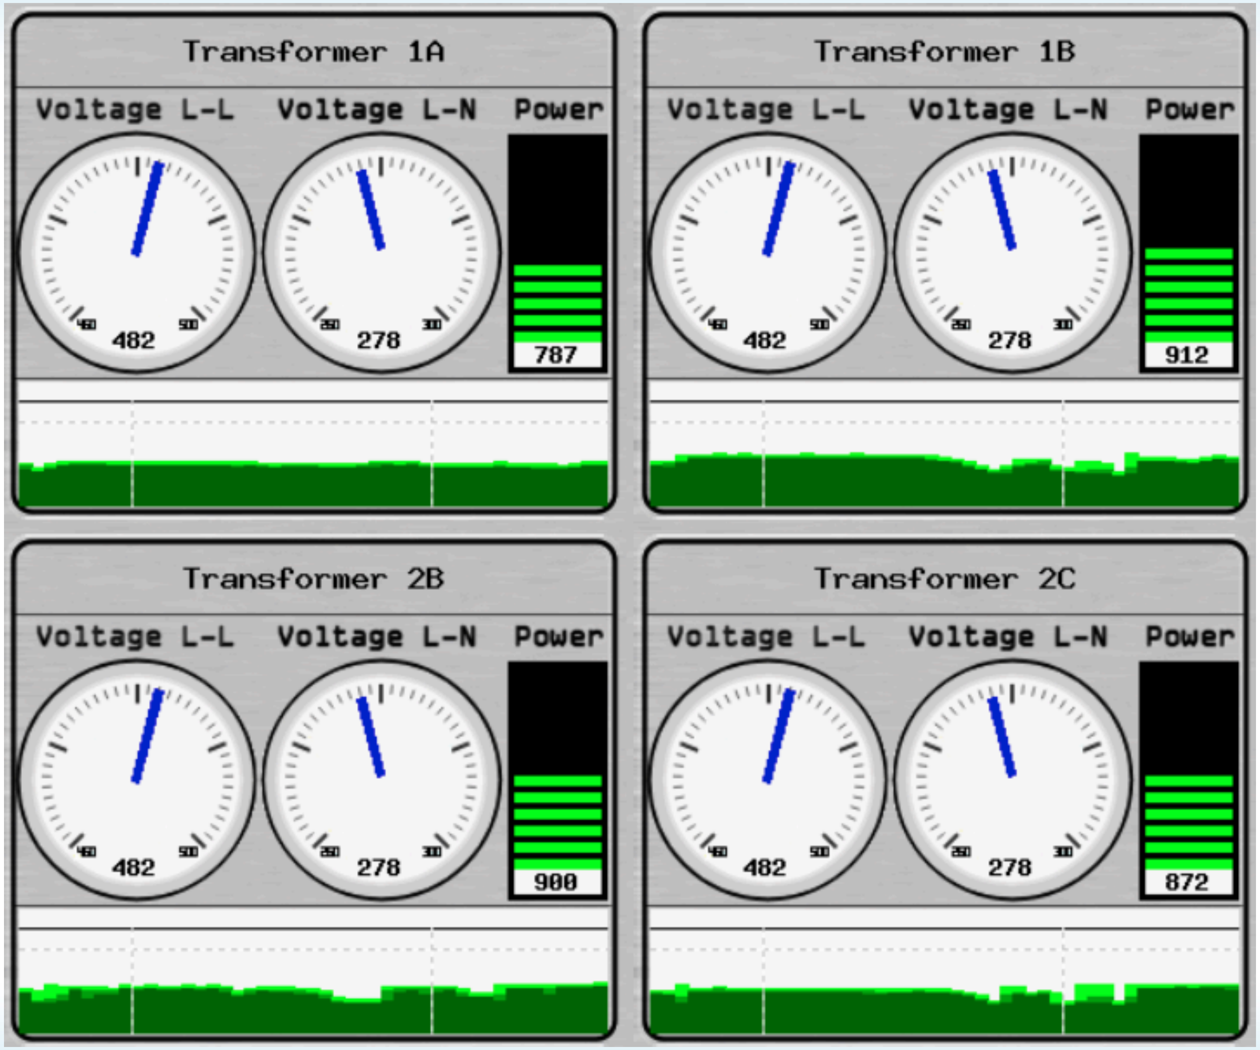

POWER

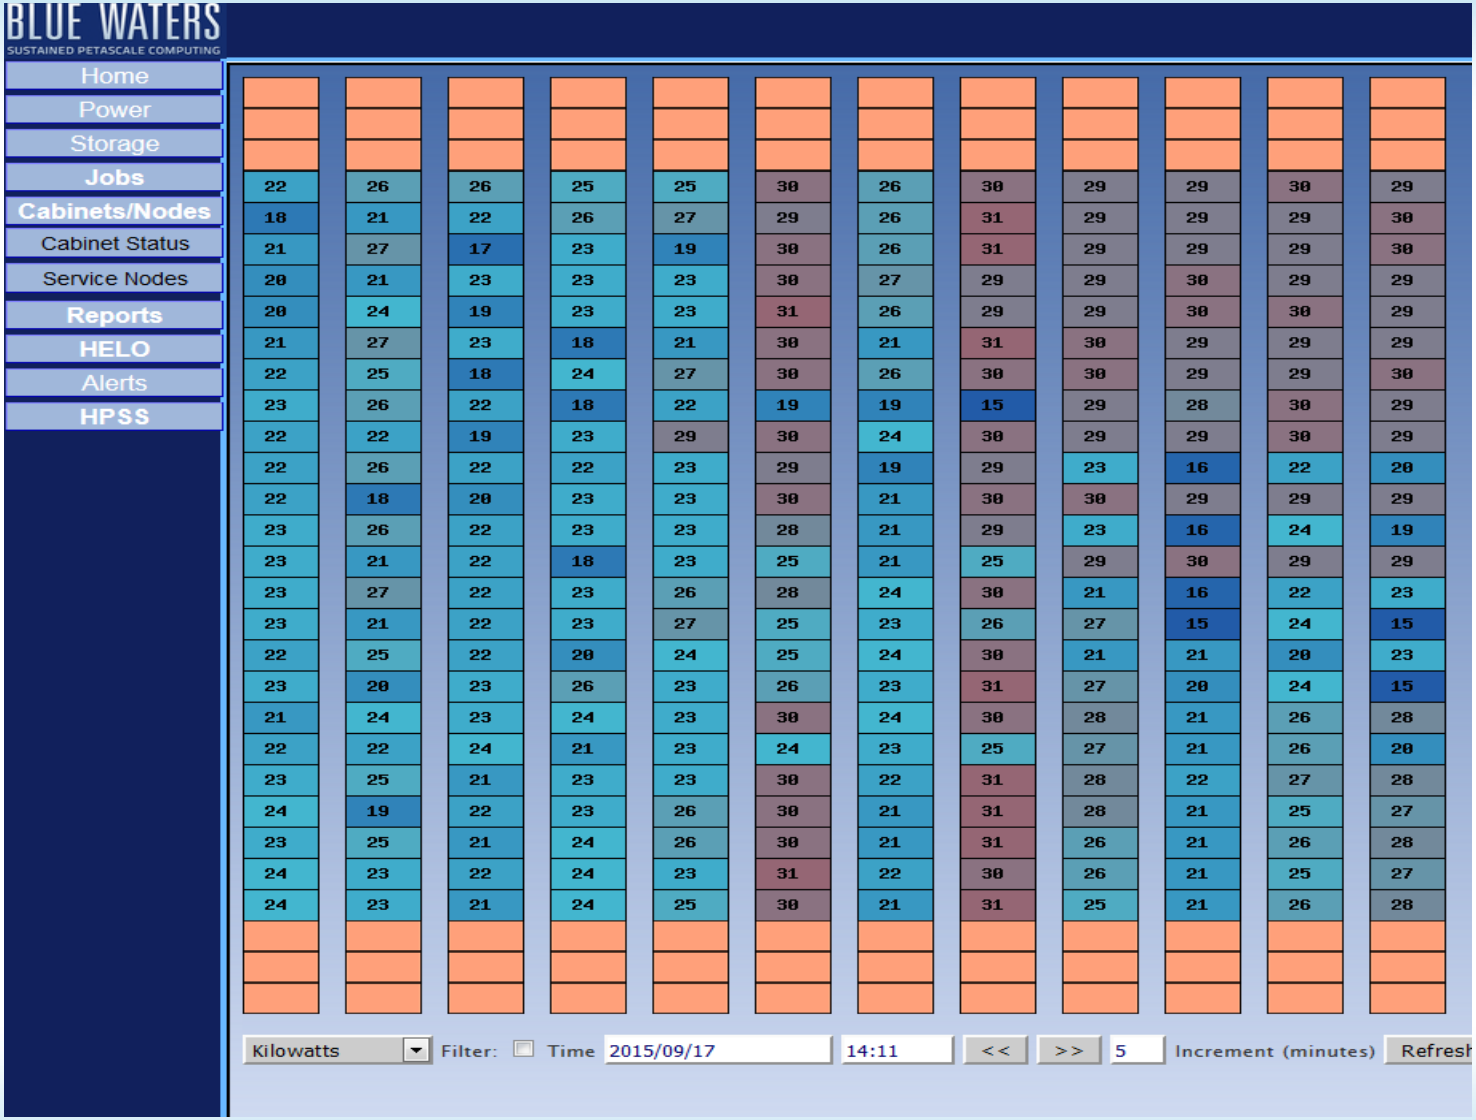

22	26	26	25	25	30	26	30	29	29	30	29
18	21	22	26	27	29	26	31	29	29	29	30
21	27	17	23	19	30	26	31	29	29	29	30
20	21	23	23	23	30	27	29	29	30	29	29
20	24	19	23	23	31	26	29	29	30	30	29
21	27	23	18	21	30	21	31	30	29	29	29
22	25	18	24	27	30	26	30	30	29	29	30
23	26	22	18	22	19	19	15	29	28	30	29
22	22	19	23	29	30	24	30	29	29	30	29
22	26	22	22	23	29	19	29	23	16	22	20
22	18	20	23	23	30	21	30	30	29	29	29
23	26	22	23	23	28	21	29	23	16	24	19
23	21	22	18	23	25	21	25	29	30	29	29
23	27	22	23	26	28	24	30	21	16	22	23
23	21	22	23	27	25	23	26	27	15	24	15
22	25	22	20	24	25	24	30	21	21	20	23
23	20	23	26	23	26	23	31	27	20	24	15
21	24	23	24	23	30	24	30	28	21	26	28
22	22	24	21	23	24	23	25	27	21	26	20
23	25	21	23	23	30	22	31	28	22	27	28
24	19	22	23	26	30	21	31	28	21	25	27
23	25	21	24	26	30	21	31	26	21	26	28
24	23	22	24	23	31	22	30	26	21	25	27
24	23	21	24	25	30	21	31	25	21	26	28

INLET TEMP

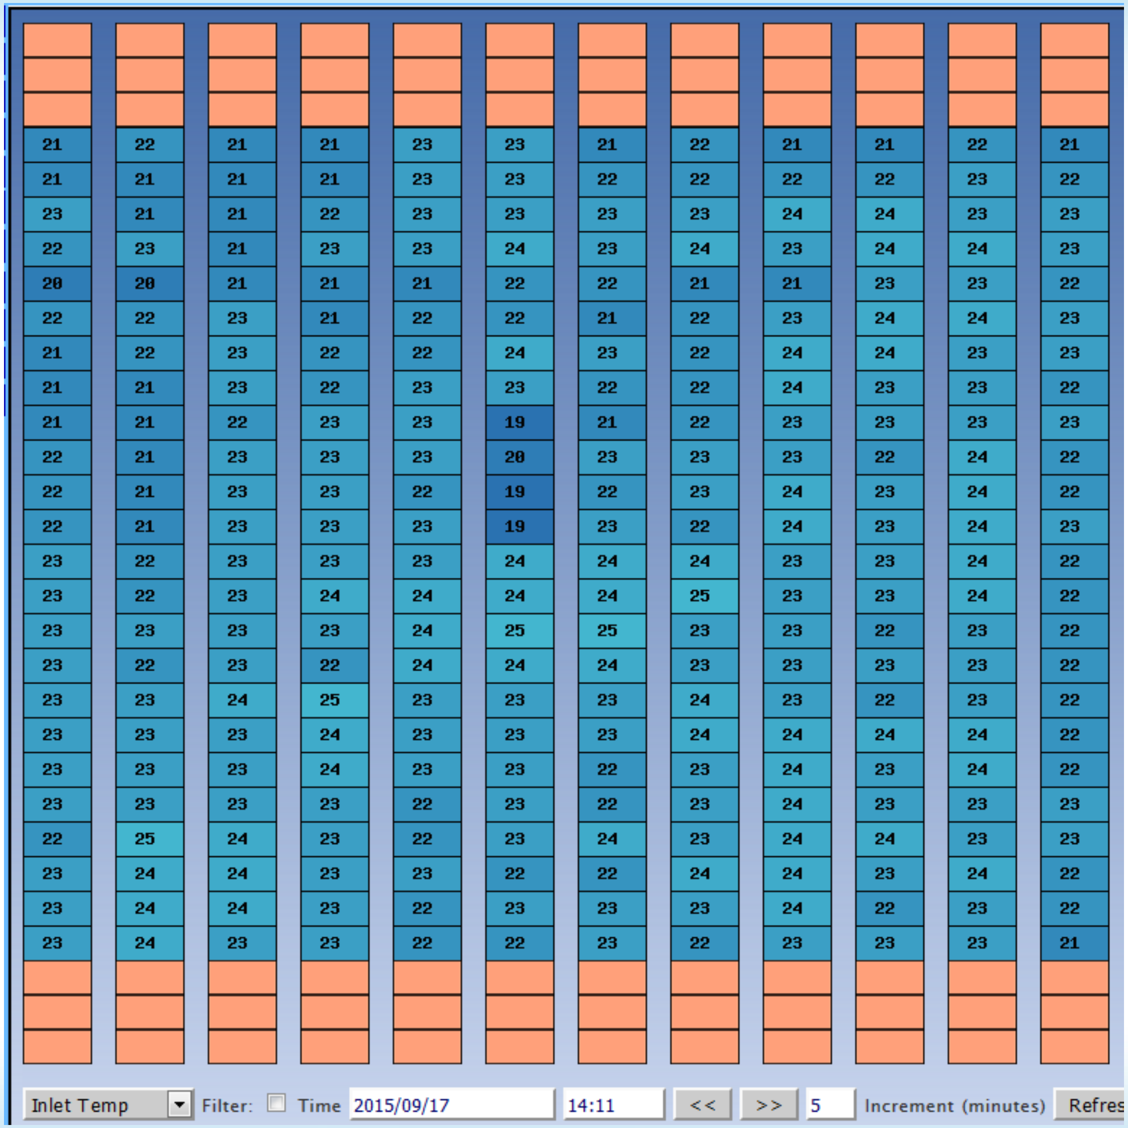
21 22 21 21 23 23 21 22 21 21 22 21
21 21 21 21 23 23 22 22 22 22 23 22
23 21 21 22 23 23 23 23 24 24 23 23
22 23 21 23 23 24 23 24 23 24 24 23
20 20 21 21 21 22 22 21 21 23 23 22
22 22 23 21 22 22 21 22 23 24 24 23
21 22 23 22 22 24 23 22 24 24 23 23
21 21 23 22 23 23 22 22 24 23 23 22
21 21 22 23 23 19 21 22 23 23 23 23
22 21 23 23 23 20 23 23 23 22 24 22
22 21 23 23 22 19 22 23 24 23 24 22
22 21 23 23 23 19 23 22 24 23 24 23
23 22 23 23 23 24 24 24 23 23 24 22
23 22 23 24 24 24 24 25 23 23 24 22
23 23 23 23 24 25 25 23 23 22 23 22
23 22 23 22 24 24 24 23 23 23 23 22
23 23 24 25 23 23 23 24 23 22 23 22
23 23 23 24 23 23 23 24 24 24 24 22
23 23 23 24 23 23 22 23 24 23 24 22
23 23 23 23 22 23 22 23 24 23 23 23
22 25 24 23 22 23 24 23 24 24 23 23
23 24 24 23 23 22 22 24 24 23 24 22
23 24 24 23 22 23 23 23 24 22 23 22
23 24 23 23 22 22 23 22 23 23 23 21
Inlet Temp
Filter:
Time
2015/09/17
14:11
<<
>>
5
Increment (minutes)
Refres

FAN SPEED

64	66	63	63	68	64	63	64	66	64	64	63
64	63	63	63	68	68	68	64	68	64	68	68
70	63	64	64	68	68	68	66	68	71	70	68
66	64	64	66	68	70	66	70	68	71	70	68
61	64	64	64	63	68	64	64	68	70	70	64
64	66	64	66	64	64	64	64	66	70	66	64
63	66	66	68	64	70	66	64	70	68	66	66
64	64	66	66	70	70	64	64	66	70	66	64
64	66	64	70	64	61	64	66	66	64	70	68
64	64	68	66	70	61	66	66	66	64	70	66
64	64	68	68	64	61	64	64	70	66	70	64
66	64	68	68	66	60	66	66	70	66	70	66
68	66	66	68	66	68	66	70	66	66	68	66
68	66	66	68	68	66	70	71	68	70	66	66
68	66	66	68	68	71	71	70	70	64	68	66
70	68	66	66	68	66	70	66	70	70	68	66
70	68	66	71	68	66	70	70	68	66	66	64
66	68	66	66	66	66	66	71	70	68	66	64
66	66	66	66	66	66	66	70	70	66	66	64
66	66	66	66	66	66	66	68	68	70	66	66
64	68	68	70	66	66	66	68	68	71	66	68
70	68	70	66	70	66	66	70	68	66	68	68
66	68	68	66	66	66	64	70	68	64	66	66
66	70	66	70	66	66	64	64	68	66	66	63

Fan Speed Filter: Time 2015/09/17 14:11 << >> 5 Increment (minutes) Refre

Node Usage and Electricity Demand

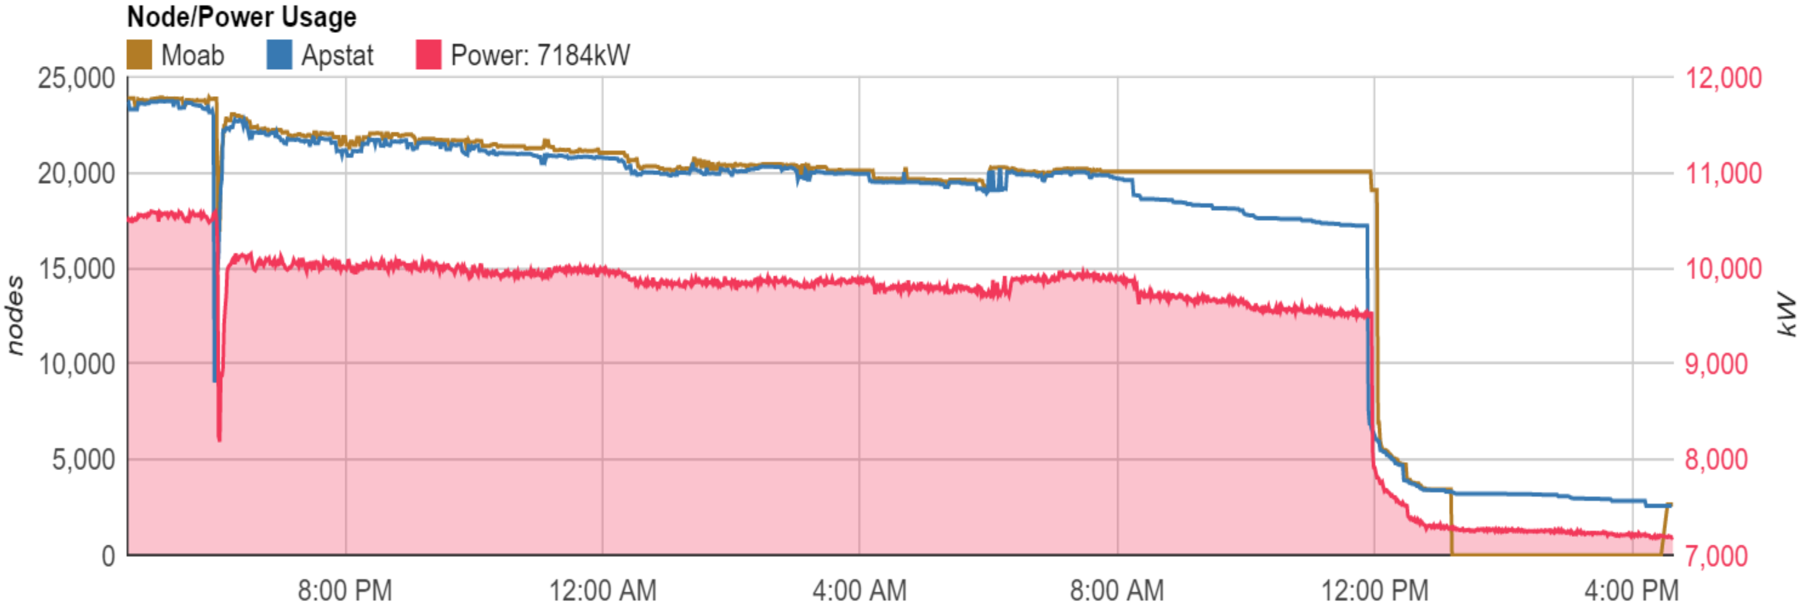

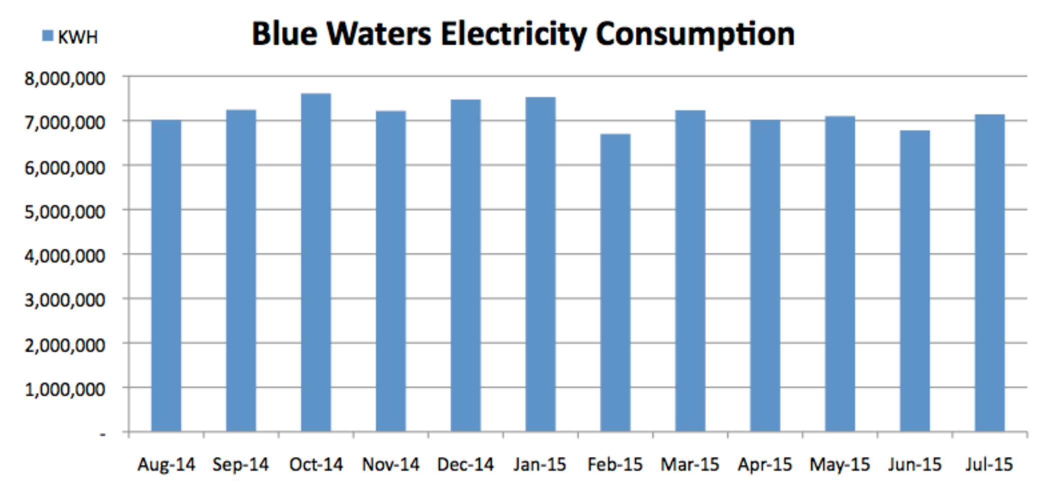
Blue Waters Electricity Consumption
KWH
8,000,000
7,000,000
6,000,000
5,000,000
4,000,000
3,000,000
2,000,000
1,000,000
-
Aug-14 Sep-14 Oct-14 Nov-14 Dec-14 Jan-15 Feb-15 Mar-15 Apr-15 May-15 Jun-15 Jul-15

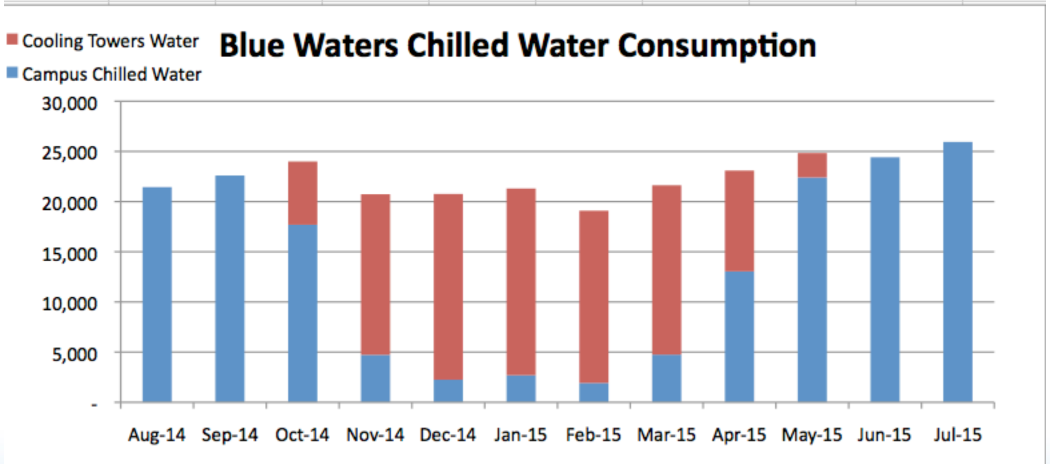
Blue Waters Chilled Water Consumption
Cooling Towers Water
Campus Chilled Water
30,000
25,000
20,000
15,000
10,000
5,000
-
Aug-14 Sep-14 Oct-14 Nov-14 Dec-14 Jan-15 Feb-15 Mar-15 Apr-15 May-15 Jun-15 Jul-15

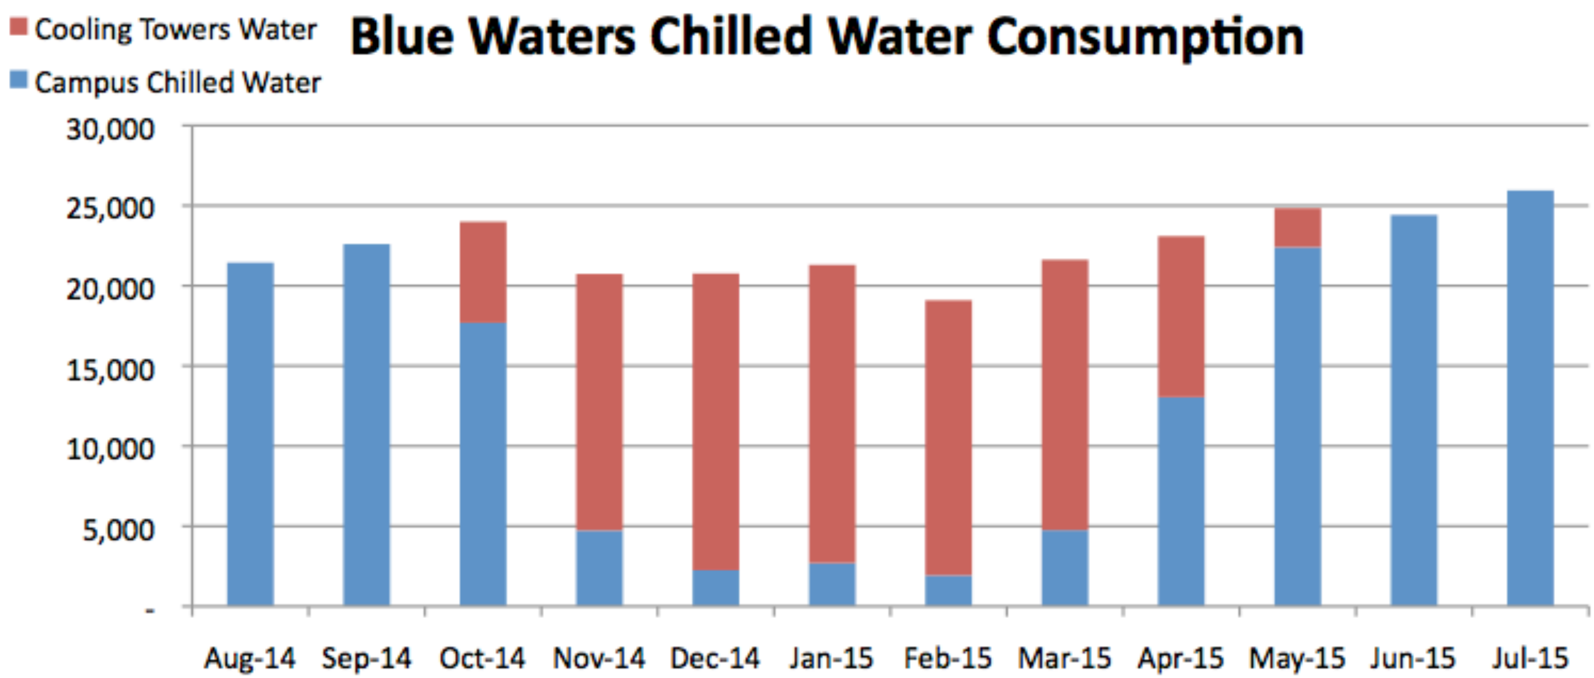
Blue Waters Chilled Water Consumption
Cooling Towers Water
Campus Chilled Water
30,000
25,000
20,000
15,000
10,000
5,000
-
Aug-14
Sep-14
Oct-14
Nov-14
Dec-14
Jan-15
Feb-15
Mar-15
Apr-15
May-15
Jun-15
Jul-15

8” Strainers Serving Blue Waters XDP HXs

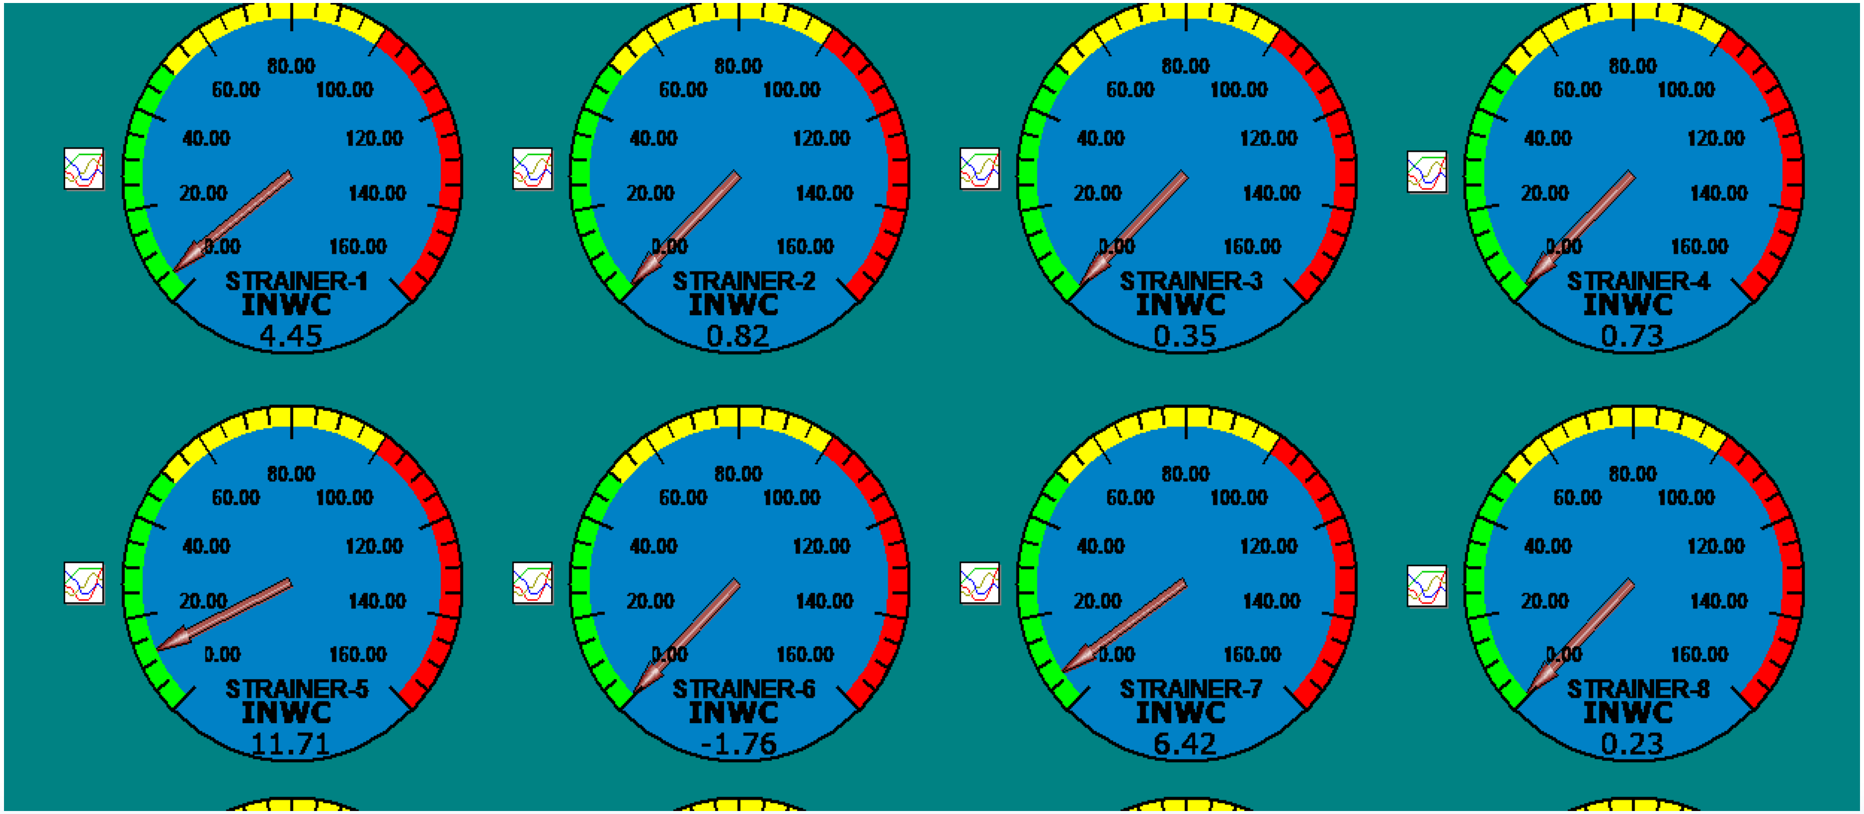

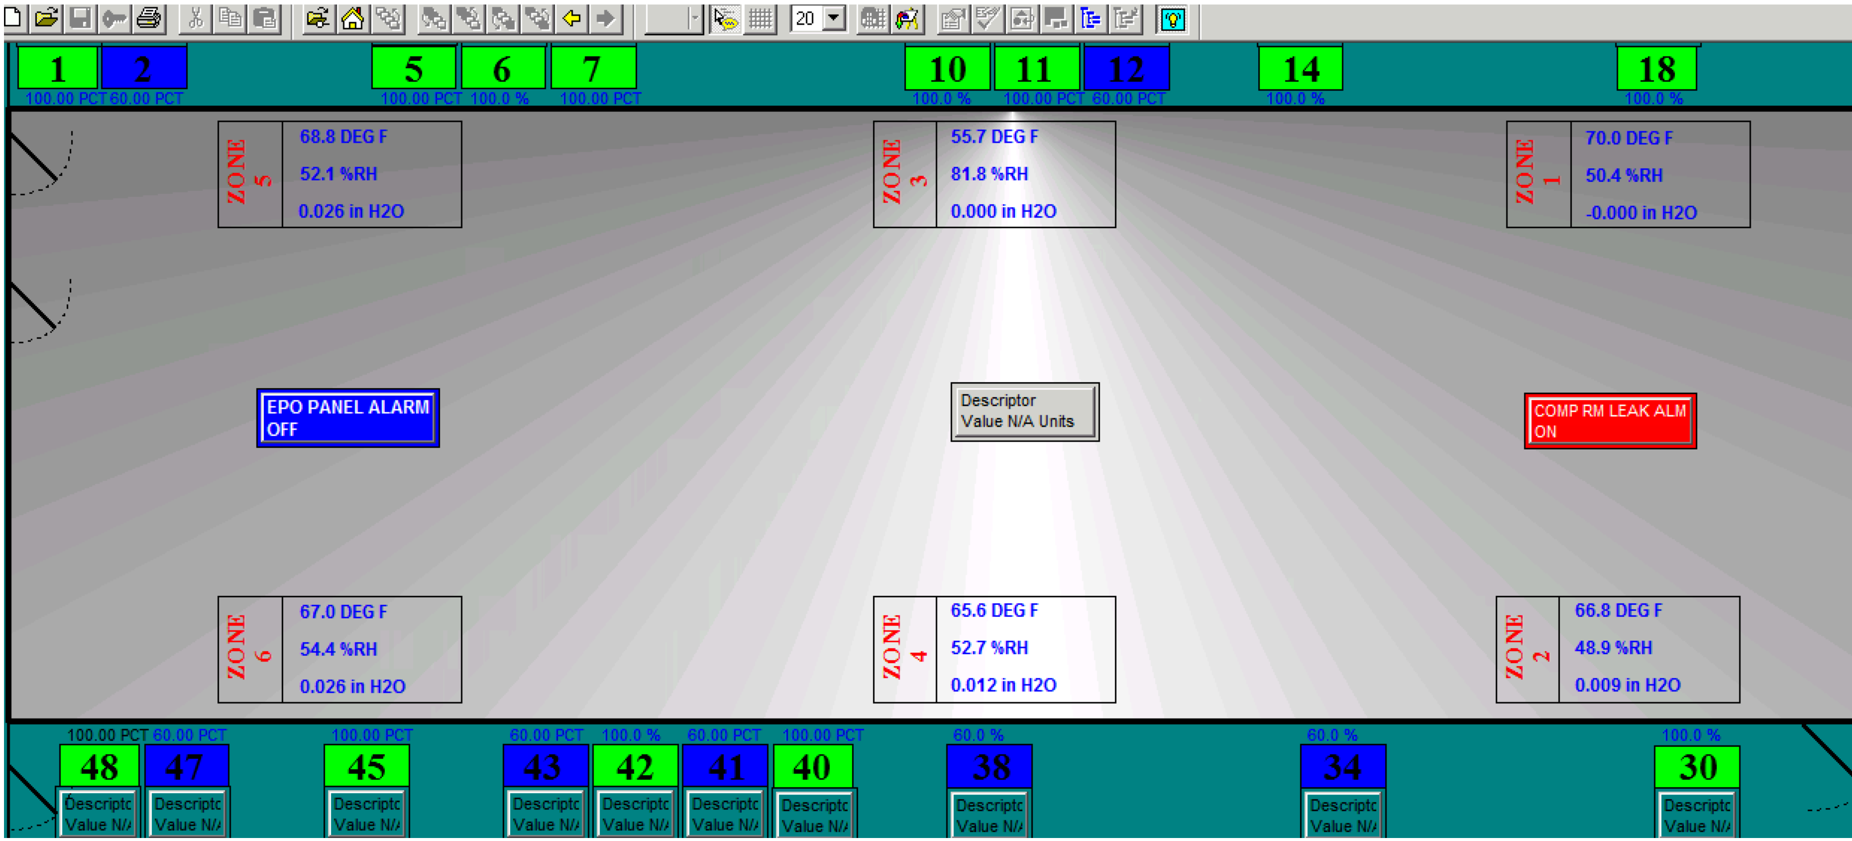
1
2
100.00 PCT 60.00 PCT
5
6
7
100.00 PCT 100.0 %
100.00 PCT
10
11
12
100.0 %
100.00 PCT 60.00 PCT
14
100.0 %
18
100.0 %
ZONE 5
68.8 DEG F
52.1 %RH
0.026 in H2O
ZONE 3
55.7 DEG F
81.8 %RH
0.000 in H2O
ZONE 1
70.0 DEG F
50.4 %RH
-0.000 in H2O
EPO PANEL ALARM
OFF
Descriptor
Value N/A Units
COMP RM LEAK ALM
ON
ZONE 6
67.0 DEG F
54.4 %RH
0.026 in H2O
ZONE 4
65.6 DEG F
52.7 %RH
0.012 in H2O
ZONE 2
66.8 DEG F
48.9 %RH
0.009 in H2O
100.00 PCT 60.00 PCT
48
47
100.00 PCT
45
60.00 PCT
43
100.0 %
42
60.00 PCT
41
100.00 PCT
40
60.0 %
38
60.0 %
34
100.0 %
30

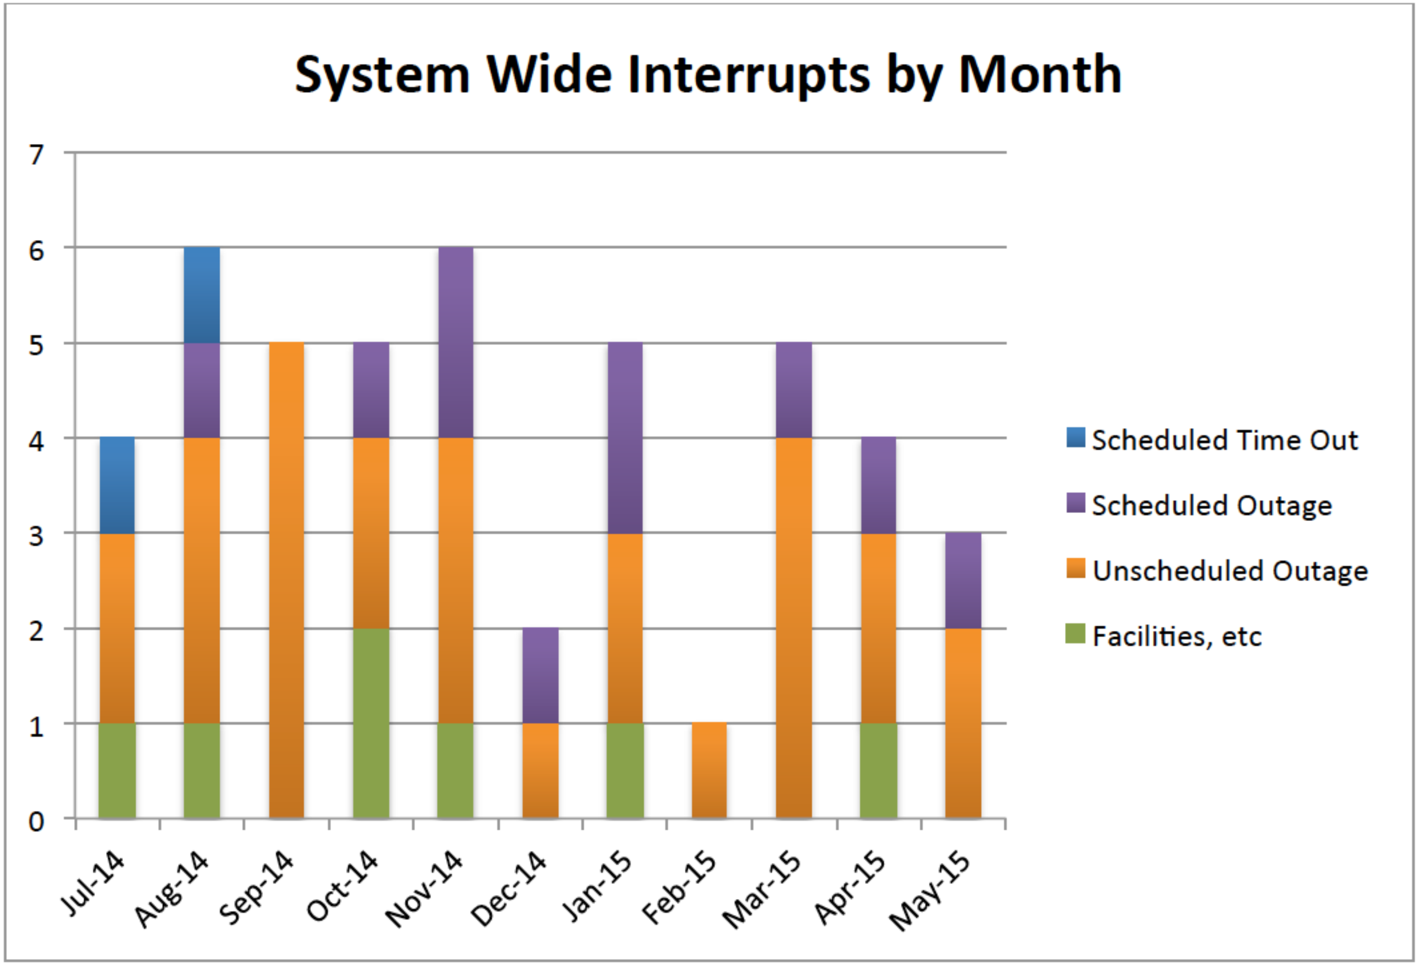

Figure 3 The number and categories of System Wide Interrupts per month

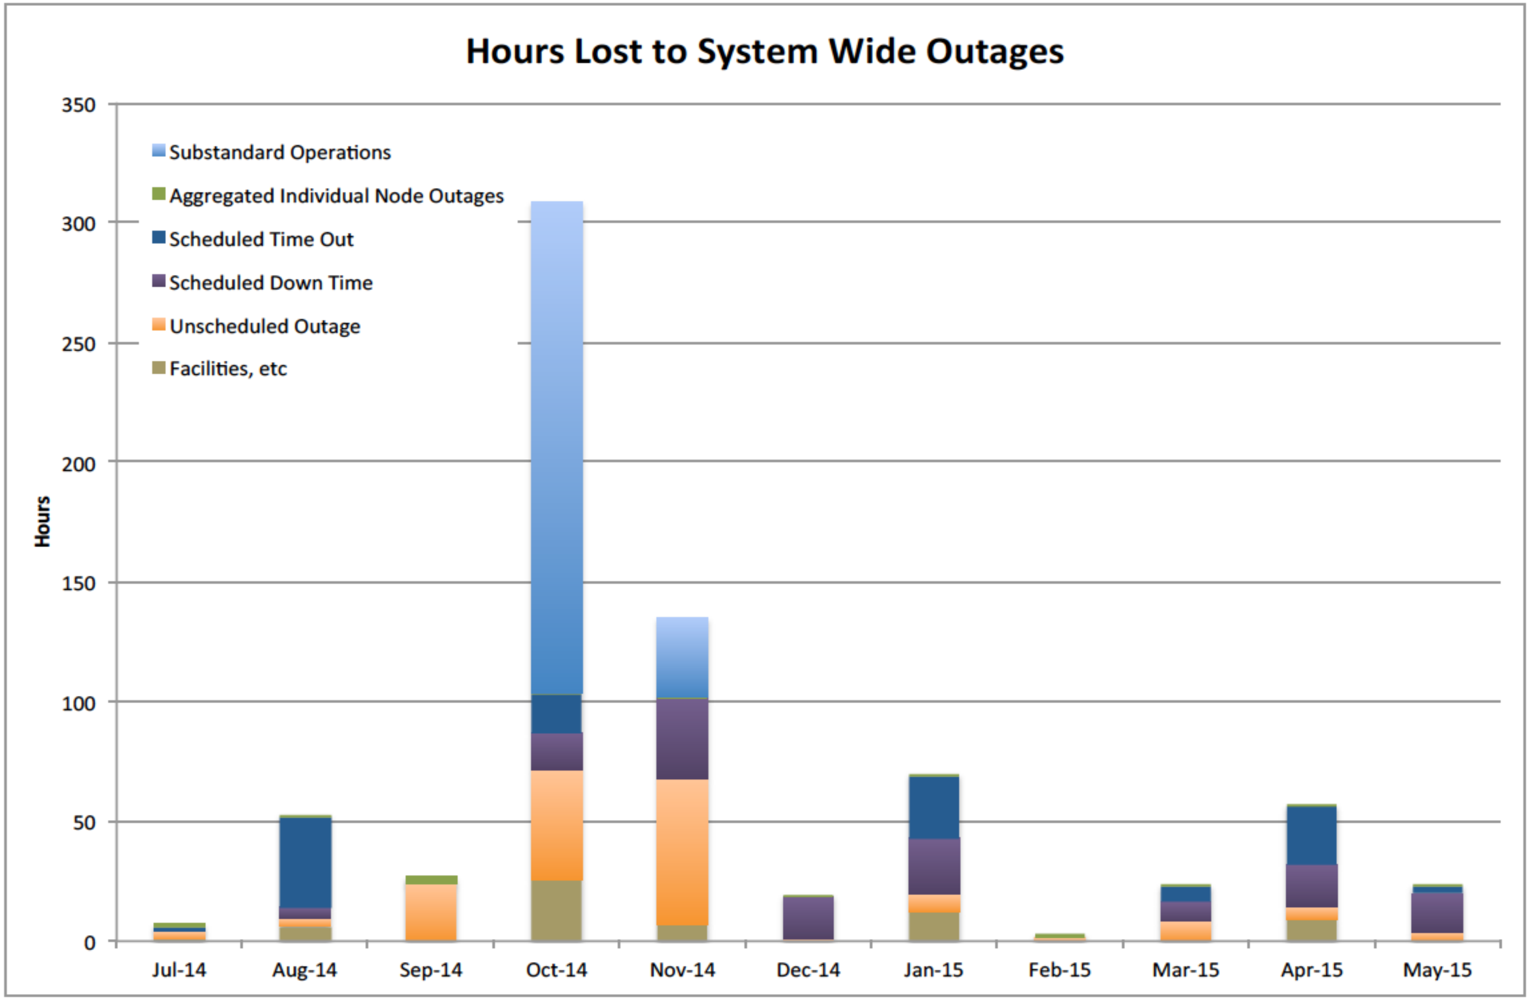

Figure 4: Hours Attributed to System Outages. Colors are as indicated in Figure 3 with the addition of green for the aggregate time for all individual node outages.

Availability

Table 1 System availability metrics

Metric	Target	7/1/2014-8/31/2014	9/1/2014 - 11/30/2014	12/1/2014-2/28/2015	3/1/2015-5/31/2015	PY2 7/1/2014-5/31/2015
Scheduled Available Computing Time	≥ 92%	96.0% **Exceeds Expectations**	78.4% **Below Expectations**	95.7% **Above Expectations**	95.4% **Above Expectations**	91.0% **Below Expectations**
Mean Time Between System Wide Failure	≥ 5 day	8.6 days **Significantly Exceeds Expectations**	7.9 days **Significantly Exceeds Expectations**	17.4 days **Significantly Exceeds Expectations**	9.9 days **Significantly Exceeds Expectations**	11.2 days **Significantly Exceeds Expectations**

Building Components

- Utilities: Electrical, Cooling, Humidification
- Computing Equipment (Machine) Rooms
- Controls, BAS and Compute Management
- Networking
- Security, Physical and Cyber
- Receiving, Processing, Recycling
- Physical Storage
- Personnel Areas

Facility Operation Concerns

- Adaptability – Utilities design
- Energy Consumption and Efficiency
- Metering and Energy Consumption Data
- Controls Systems Integration
- Preventive Maintenance and Outage Scheduling
- Connectivity, Networking and Speed
- Reliability aka "Up Time"
- Notification Procedures (based on metrics)

Lessons Learned

- Early Identification of problems is critical (Metrics help with this)
  - Valve position data and Strainer DP data indicate condition and status of cooling system
  - Cabinet level data provides early warning
  - Electricity Meters and Dashboard also provide early warnings
  - Visual, Audible, and Olfactory Cues + Experience (the Bonus Metrics)

Lessons Learned

- Automation and Controls provide Metrics Data
  - Differential Pressure sensors: strainer maintenance costs reduced, improved pump utilization and saves energy and $s
- Alternate Cooling Strategies Integration
  - Evaporative Cooling Towers reduce cooling costs by more than 50% and improving PUE
- Electricity Metering and Metrics
  - Metering at panel level rather than circuit level saves time, lower capital cost, but less granularity

Blue Waters Active Jobs Metric

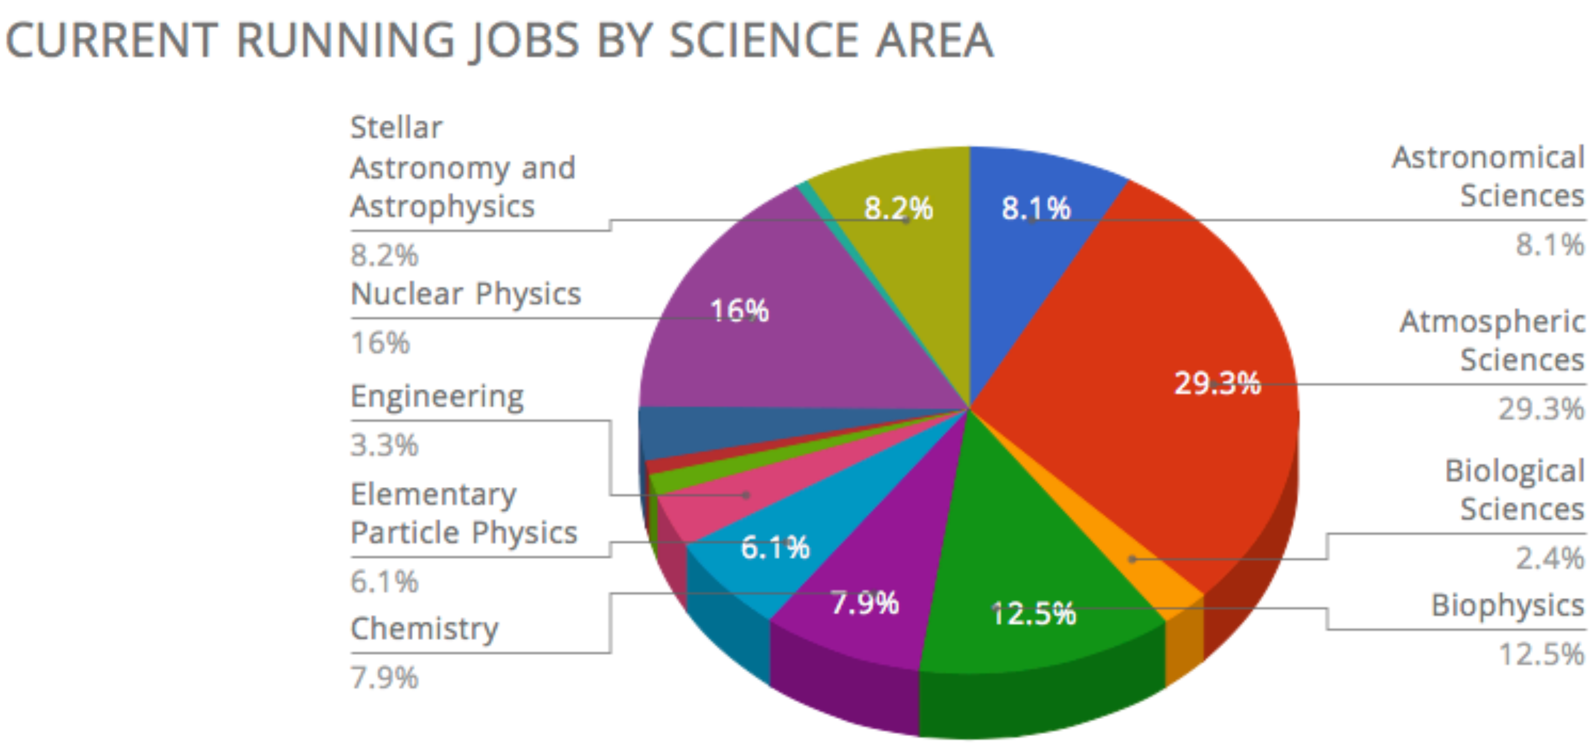

QUESTIONS

http://www.ncsa.illinois.edu

THOMAS DURBIN, P.E., LEED AP
217-333-4024 office
217-260-9559 cell
tedurbin@illinois.edu