Using Streaming Analytics to Improve Operating Efficiency

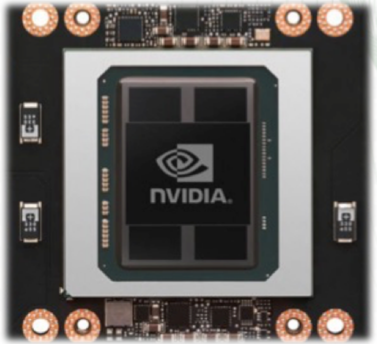

Jim Rogers

Director, Computing and Facilities
National Center for Computational Sciences
Oak Ridge National Laboratory

Using Streaming Analytics to Improve Operating Efficiency

Abstract: ORNL is constructing a new warm-water energy plant that will provide up to 6000 tons of cooling, with supply temperatures that are nearly thirty degrees warmer than the original central energy plant. The initial tenant for this facility is *Summit*, a hybrid computing system from IBM and NVIDIA that could generate demand of up to 20MW. While the control systems for the mechanical plant are well-understood, there is a significant gap between the information available from Summit, and how that information can be used to improve the operational efficiency of the energy plant. To bridge this gap, ORNL is prototyping data collection and real time streaming analytics techniques that can be used to contribute to more efficient plant operation. There are significant challenges related to the data volume, system scale, and the integration with traditional SCADA controls. These challenges, and the corresponding opportunities will be described.

Presenter Overview

Jim Rogers is the Computing and Facilities Director for the National Center for Computational Sciences (NCCS) at Oak Ridge National Laboratory (ORNL). The NCCS provides full facility and operations support for multiple petaFLOP-scale systems including Titan, a 27PF Cray XK7. ORNL has recently signed a subcontract with IBM for their next system, with a baseline description for a 200PF computational resource to be delivered beginning in 2H2017. Jim has a BS in Computer Engineering, and has worked in high performance computing systems acquisition, facilities, integration, and operation for more than 25 years.

Current ORNL Computing Footprint

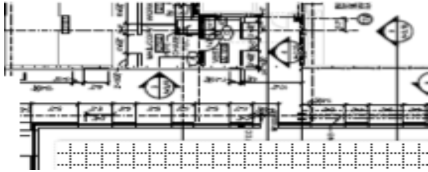

- Infrastructure
  - Floor Space
    - 2x20,000 ft^2 (2x1,859 m^2) raised floor data centers
  - Electrical
    - 13.8kVA medium voltage distribution system
      - Diverse feeds and routes for resiliency
    - CSB Technical Load supported by
      - (6) 1.5/2.0 MVA transformers
      - (8) 2.5/3.3 MVA transformers
      - (1) 3.0/4.0 MVA transformer
  - Mechanical
    - 6,600 tons of cooling (5 separate chillers)
    - 42°F (5.5°C) supply; 55°F (12.8°C) return
    - 4,400-5,000 gpm (16,500-19,000 lpm) flow

- Systems in this environment
  - 200-cabinet Cray XK7 (Titan)
  - 8-cabinet Cray XC-30 (Eos, Darter)
  - 16-cabinet Cray XC-40's (Gaea)
  - 40-cabinets of moderate density clustered computing solutions
  - 80-cabinets of high-density storage
  - 60 cabinets of supporting infrastructure
  - 20 testbed systems
    - ARM
    - Xeon
    - POWER
    - GPU

Current Facility Power Use

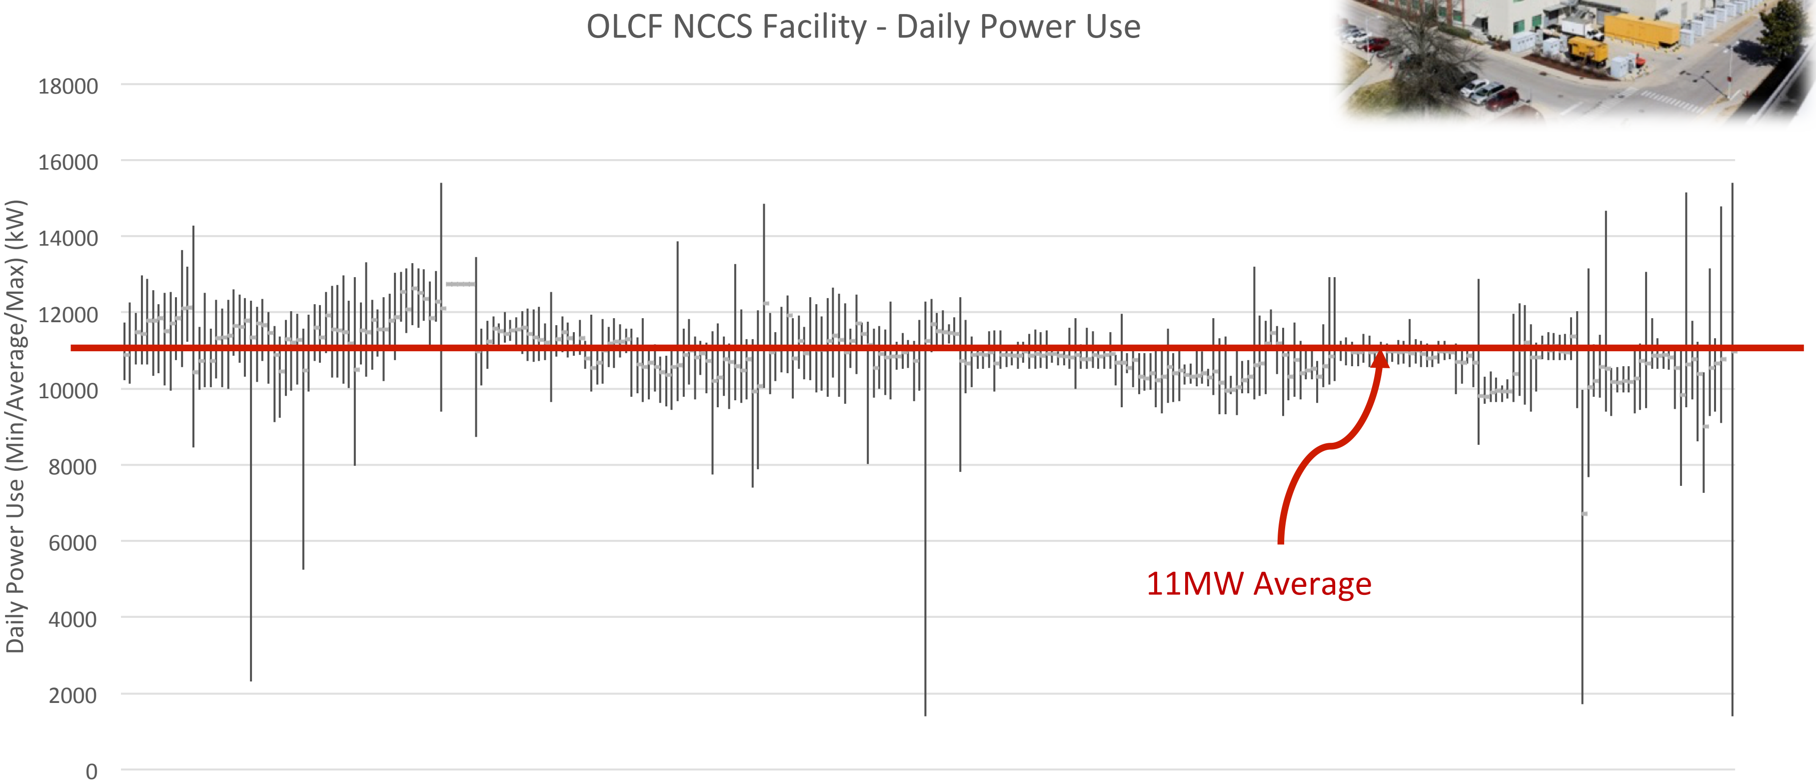

Current Single Largest System

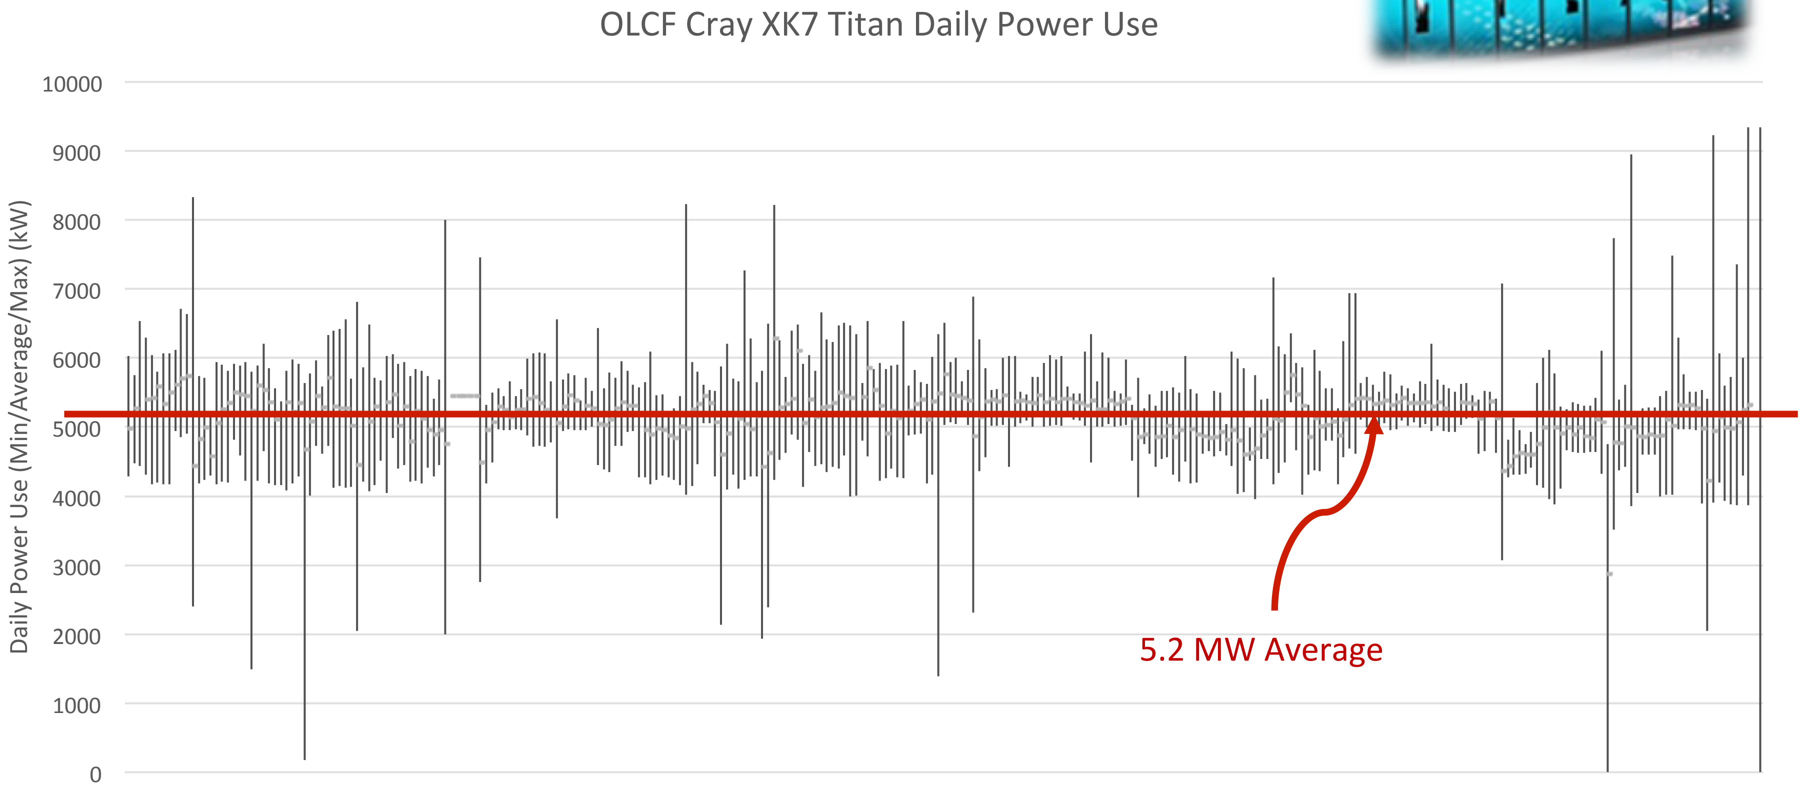

Oak Ridge National Laboratory
OAK RIDGE LEADERSHIP COMPUTING FACILITY

OLCF-4 Leadership System

Collaboration with IBM, NVIDIA, and Mellanox

Hybrid CPU/GPU architecture

Performance Target: ~7x Titan's Application Performance (200PF peak)

Approximately 4,600 compute nodes, each with:

- Multiple IBM POWER9 CPUs and multiple NVIDIA Volta™ GPUs
- CPUs and GPUs connected with high speed NVLink™
- Large coherent memory: over 512 GB (HBM + DDR4)
- An additional 800 GB of NVRAM, which can be configured as either a burst buffer or as extended memory

Dual-rail Mellanox® EDR-IB full, non-blocking fat-tree interconnect

IBM Elastic Storage (GPFS™) – 2.5TB/s I/O and 250 PB disk capacity.

DOE ASCR Computing Portfolio

System attributes	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrade	
Name Planned Installation	Edison	TITAN	MIRA	Cori 2016	Summit 2017-2018	Theta	Aurora 2018-2019
System peak (PF)	2.6	27	10	> 30	200	2016	180
Peak Power (MW)	2	9	4.8	< 3.7	13	1.7	13
Node perf (TF)	0.460	1.452	0.204	> 3	> 40	>3	> 15 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron + NVIDIA Kepler	64-bit PowerPC A2	Intel Knights Landing (KNL)	IBM Power9 CPUs + NVIDIA Volta GPUs	Intel Knights Landing (KNL)	Intel Knights Hill (KNH)
System size (nodes)	5,600 nodes	18,688 nodes	49,152	9,300 nodes; 1,900 nodes in data partition	~4,600 nodes	>2,500 nodes	>50,000 nodes

Summit's Electrical Distribution Design

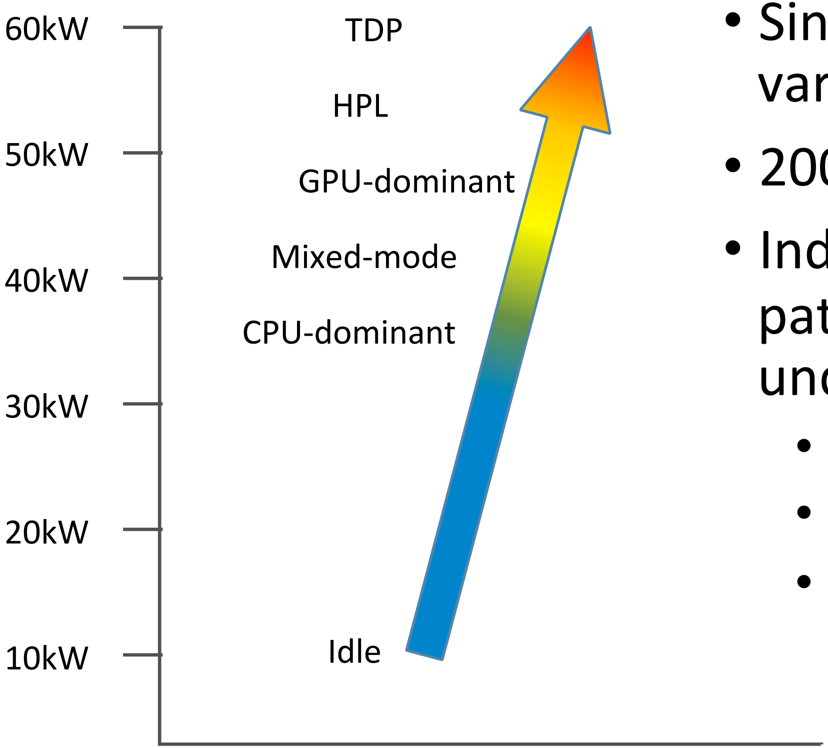

- Single cabinet design point: ~50kW across a variety of workloads (10kW idle)
- 200PF system design point: ~13MW
- Individual applications will exhibit different patterns of consumption. This works fine under most conditions
  - Steady state operations system-wide
  - A mix of jobs terminating and others starting
  - Draining of queues for system maintenance

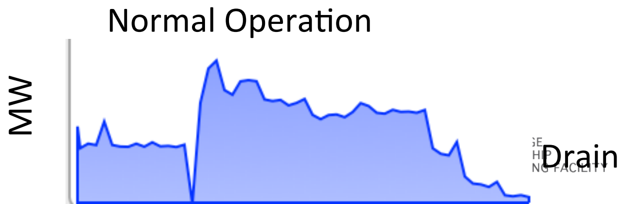

Summit's Mechanical Cooling Design – Single Cabinet

- Water supplied via a cabinet-level manifold
- CPU and GPU heat loads are transferred directly to medium-temperature water (on-chip cooling)
- Residual heat (DIMMs, power supplies, HCAs) is extracted with fans, and moved to water via RDHX.
- Serial design eliminates second loop

❗ No flow control

ORNL Medium Water Temperature (MWT) Facility

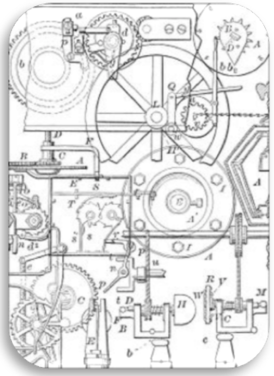

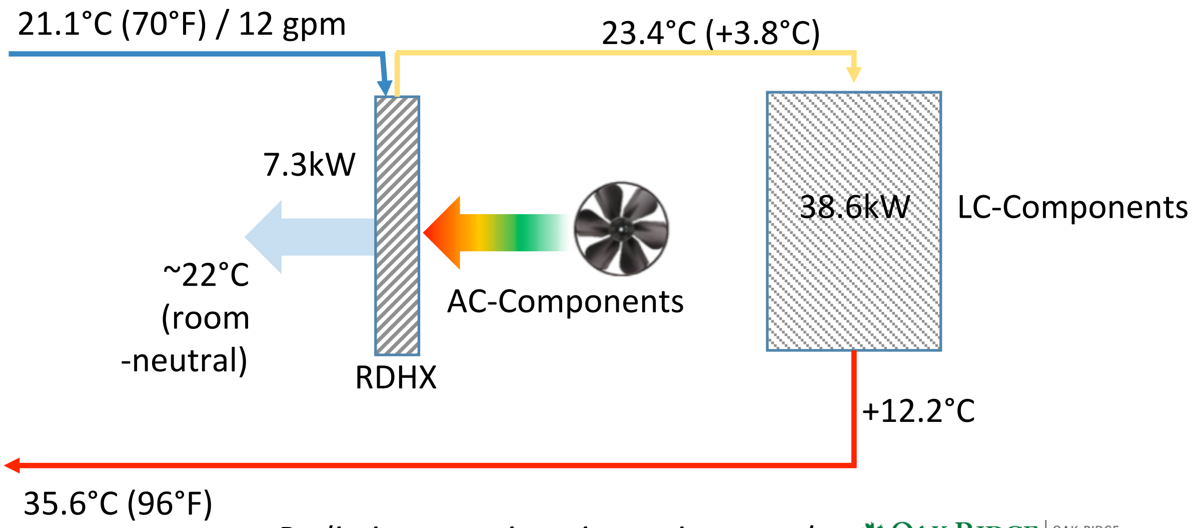

Preliminary engineering estimate only.

Necessity

Electrical

- Per-node power consumption forecast is 10kW (idle) to ~55kW (absolute worst case).

! Fast rate of change (application launch or crash) could produce

$$4{,}600 * 45\text{kW} == 11.5\text{MW}$$

load swings (both directions).

Mechanical

- Mechanical
  - Each node has multiple water-cooled components, each with an operating envelope, and a limited tolerance (time/temp) to excursions

  ! No flow control at the cabinet level means *protect the worst case (at the node-level)*

  - Fast rate of change (load swings may create significant stress on the mechanical plant.

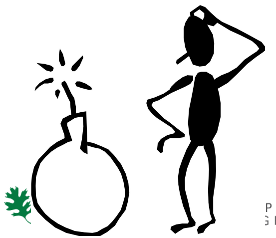

P
FACILITY

Improving Operations

❗ ***Caution- make sure that the tail does not wag the dog…***

Things I Know (and Control System Does Not)

- A node is reporting over-temp
  - ✓Request increased flow
- A new eligible job is pending
  - ✓Report pending increase in aggregate heat load and provide adequate time to adjust plant state
- Demand will decrease (reservation, no eligible job, etc)
  - ✓Report pending decrease (and potentially the time interval)

Things I Need from the Control System

- How much capacity is available?
- Are there issues that would preclude normal operation?
  - Weather, chiller, maintenance…
- (correlation) plant state, as an input to long-term analysis of applications, power consumption, detailed node data, and corresponding plant info

Motivation – Minimize Operating Expense

- Warm-water design needs only evaporative cooling about 70% of the year.
- Some chiller contribution is needed up to 30% of the year.
- PUE for 100% evaporative cooling mode: 1.10
- PUE impact when chiller is required: 1.40 for those hours.
- Blended PUE estimate: ~1.15 (which is a big deal on a 13MW computer)

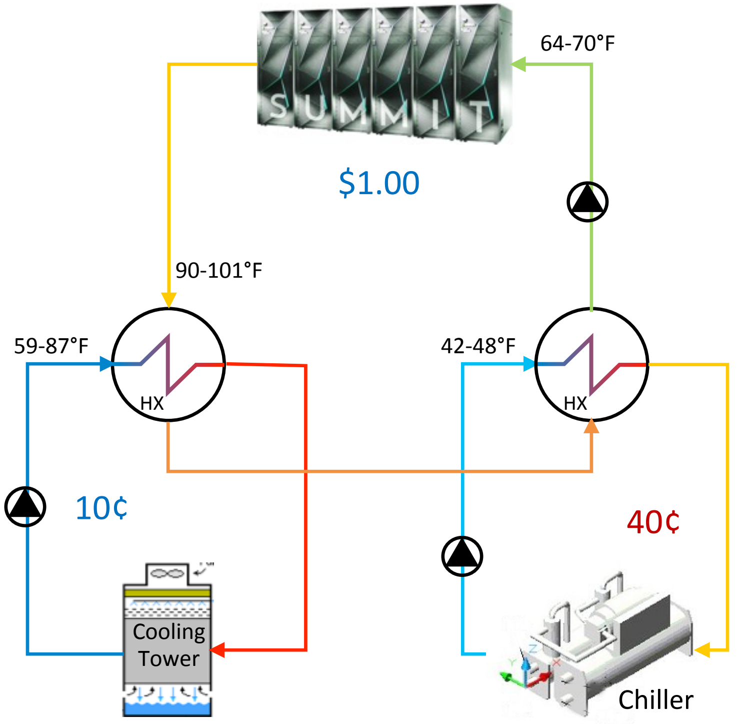

Design Basis

- Monitoring of Individual Nodes via the OCC
- Awareness of individual nodes exceeding operating conditions
- Use streaming analytics to classify range of operation as "normal" as well as "outliers" that must be addressed
- Feedback loop from the PLC to know current delivery metrics
- Integration with scheduler for special cases (fast rate of change)

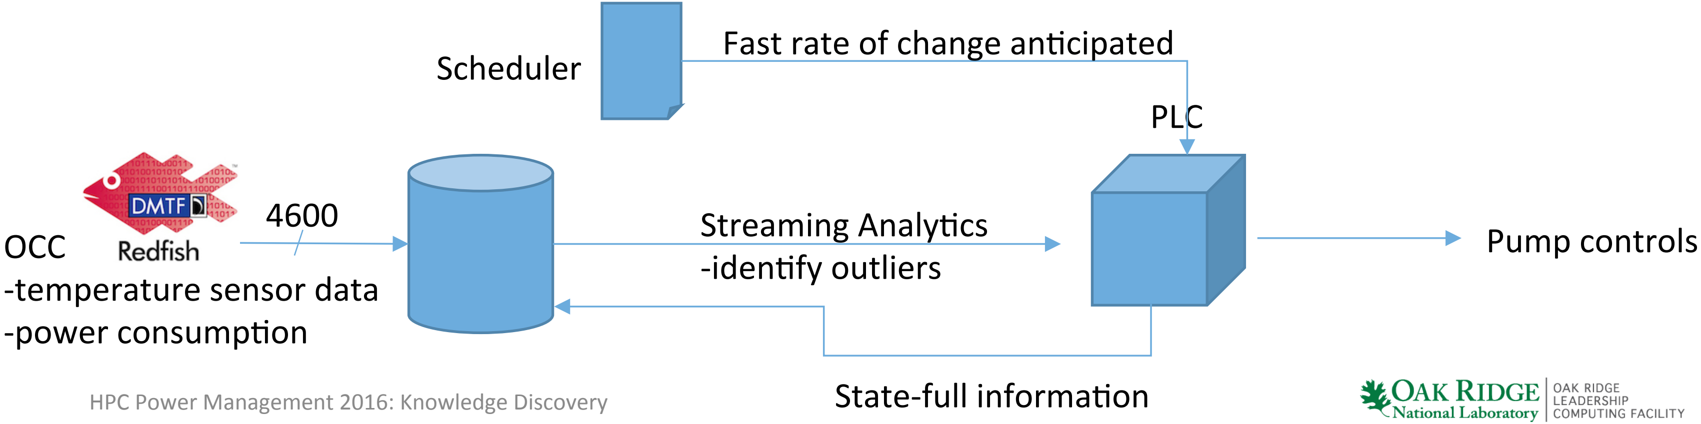

Streaming Analytics to Support Summit Operations

- Clear patterning of data across time points:
  - Across a node (intra)
  - Across nodes (inter)
- Inter-node correlations are high indicators of collective behavior amongst processors
- Can contribute to recommendation to add/remove cooling

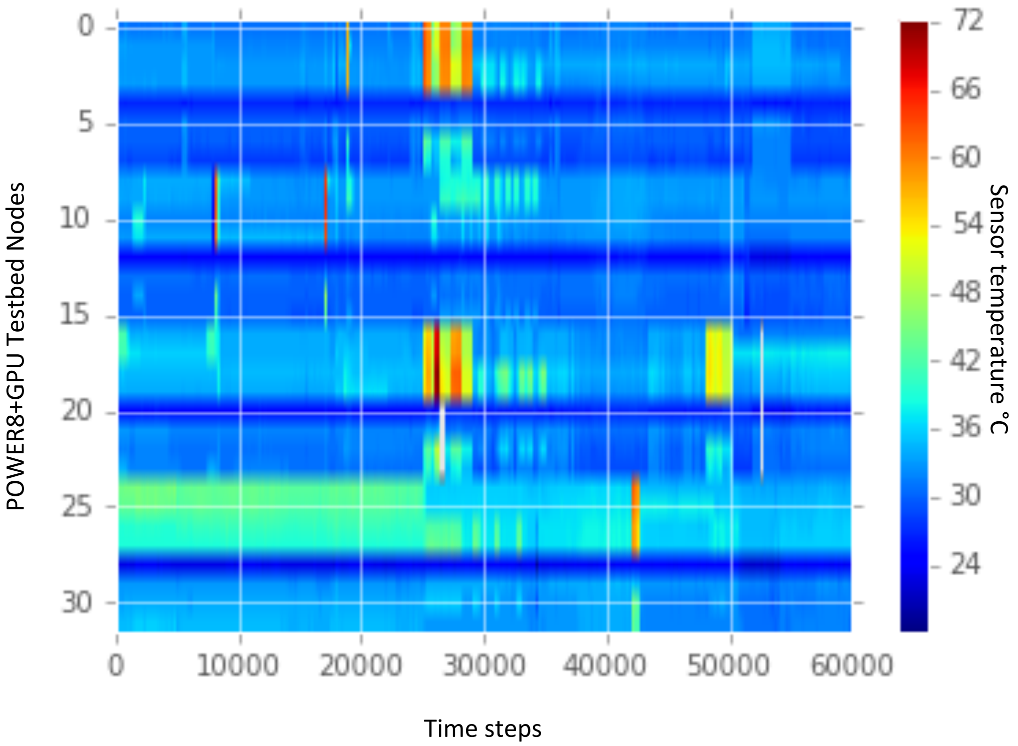

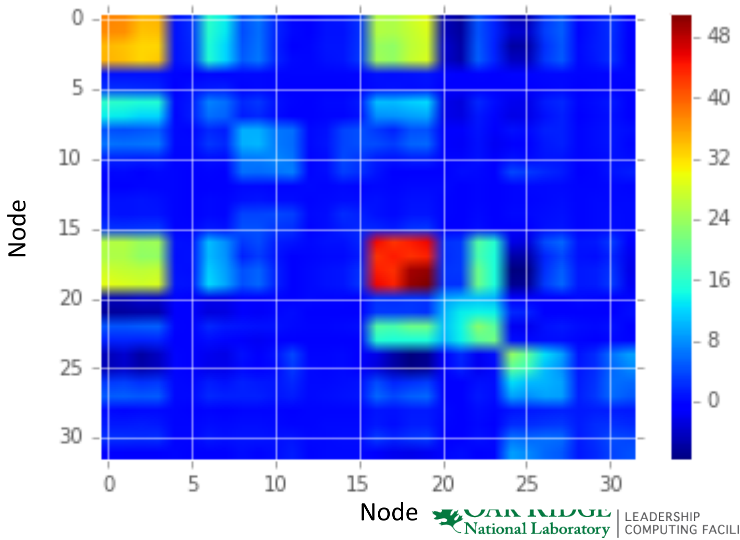

Single-node variation in temperature, POWER8

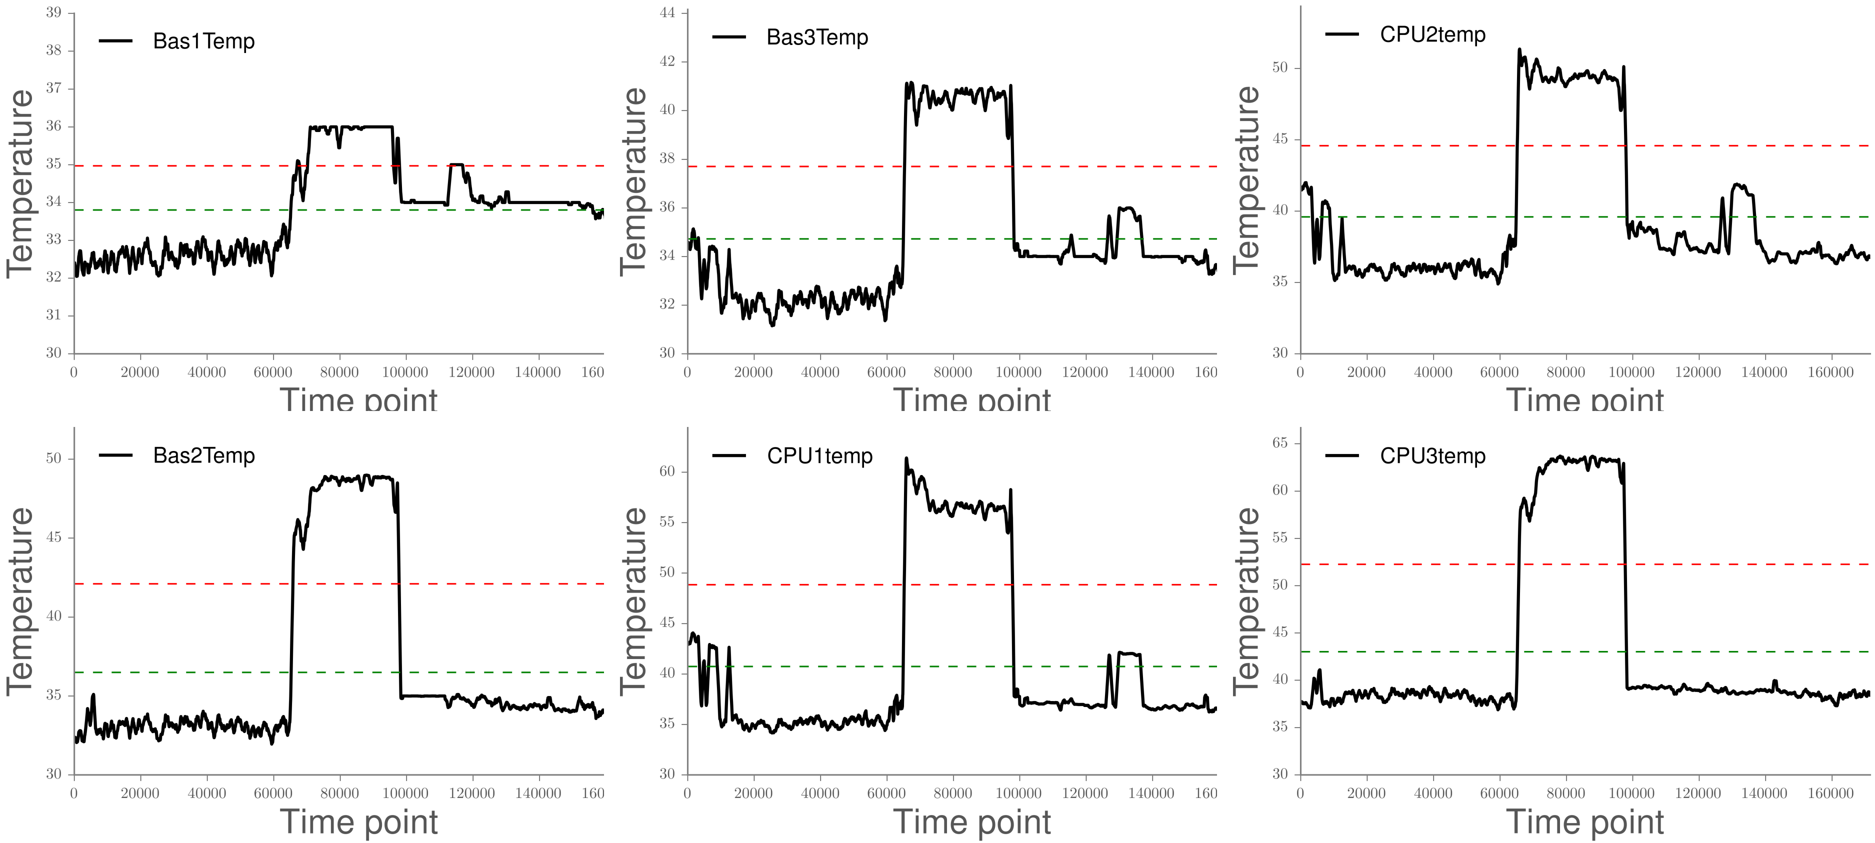

Snapshot

Today

- Maintaining adequate flow to cool compute nodes is customer (ORNL) responsibility
- Balancing flow will reduce operating expense (maximize chiller-less cooling)
- Potential $M savings annually
- Power8+/Pascal arrive in September'16 - Very similar mechanical and OCC design to Summit

Tomorrow

- Complete DB schema for scheduler/analytics/stateful data
- Solve security/communication issues from Summit enclave to SCADA
- Integrate Summit-aware communication in to PLC
- Extensive test and tune
- Correlate application, hardware, power, and mechanical plant data

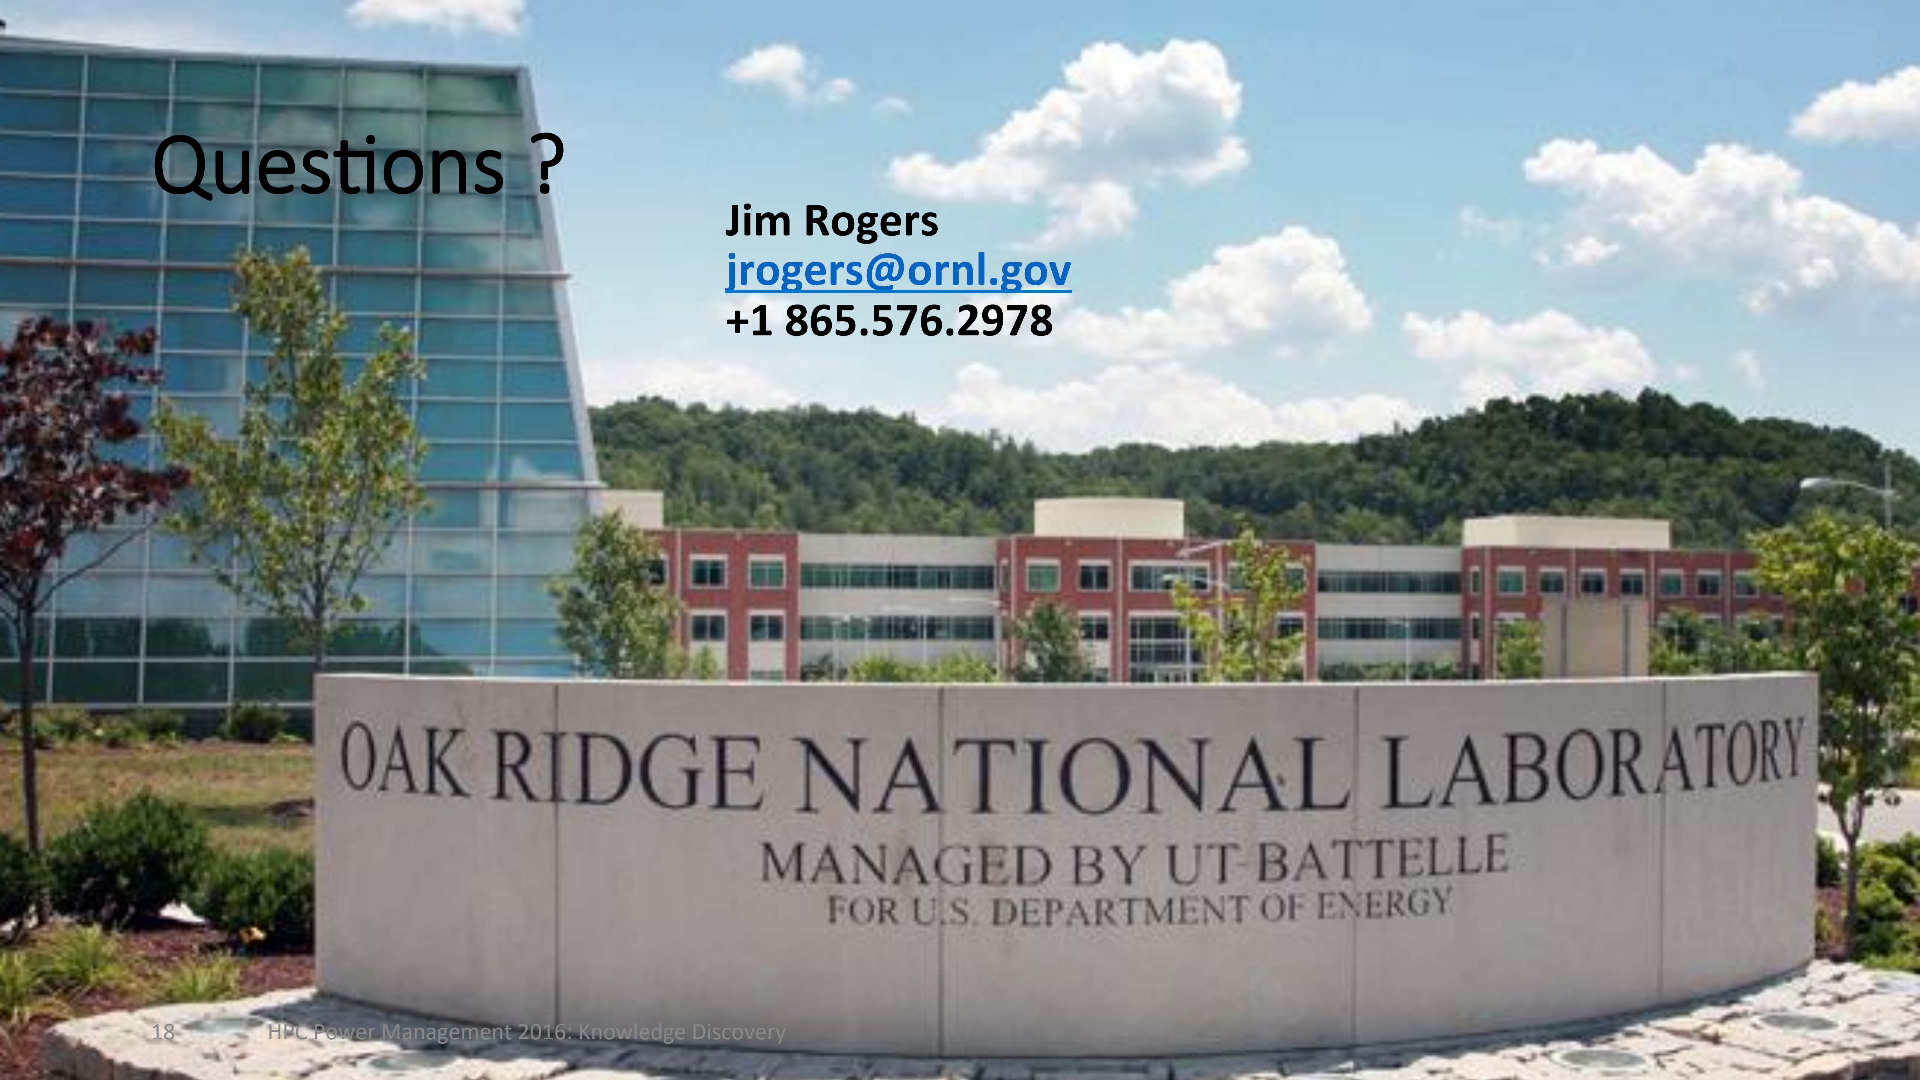

Questions ?

Jim Rogers
jrogers@ornl.gov
+1 865.576.2978