Improving Efficiency with Dynamic Controls

HPC4HPC

High Performance Cooling for High Performance Computing

David Grant, PE, CEM, DCEP

HPC Mechanical Systems

Modernization Project Office

Oak Ridge National Laboratory

+

Infrastructure Co-Lead for the EEHPCWG

Corresponding Member of ASHRAE TC9.9

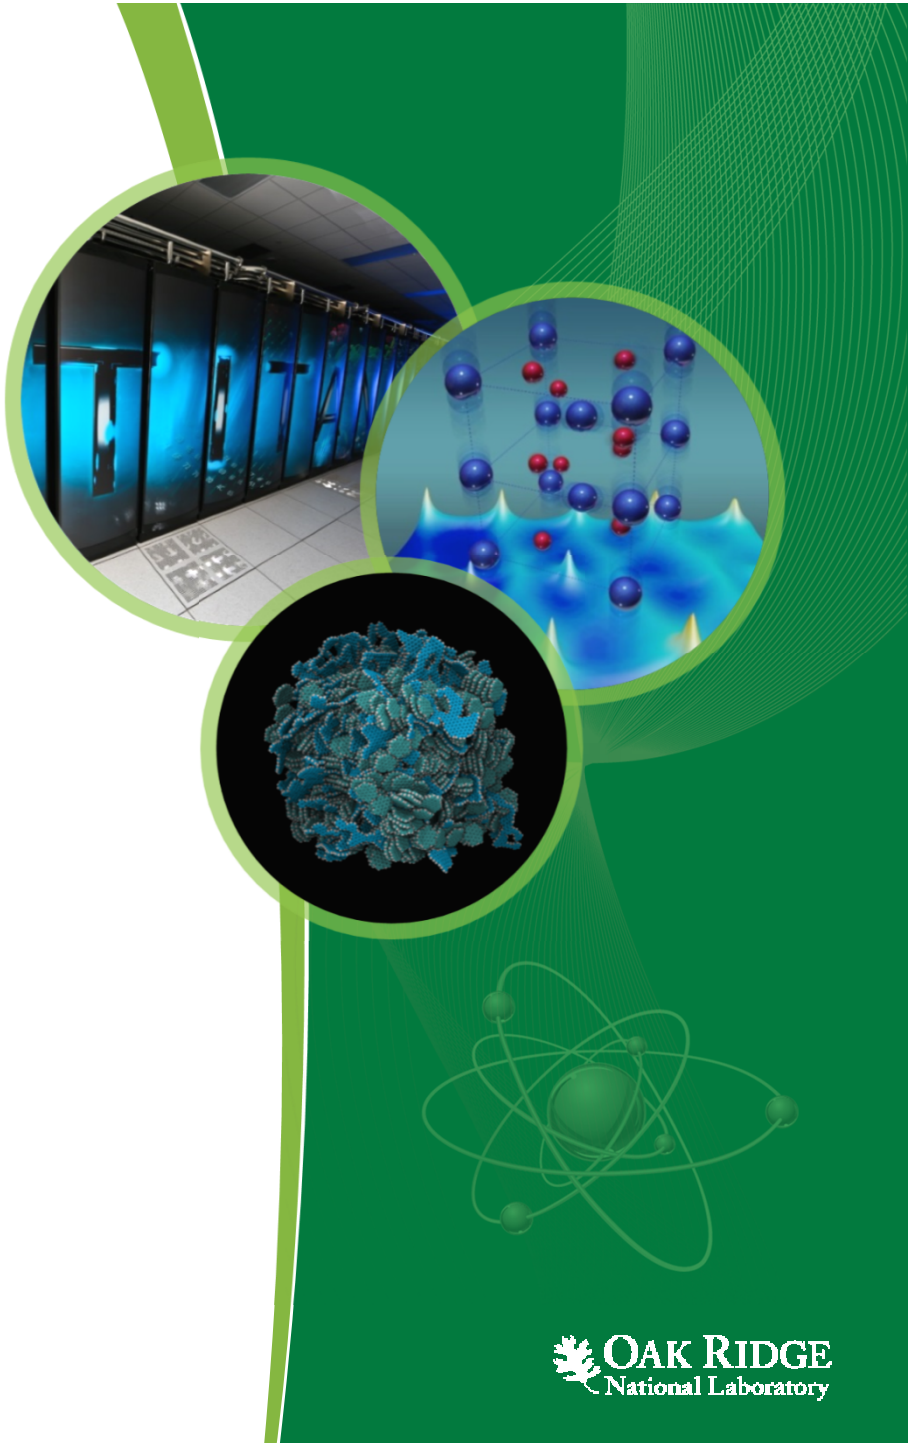

ORNL is managed by UT-Battelle
for the US Department of Energy

OAK RIDGE National Laboratory

Today's Discussion

- Start with the characterization of HPC system loads
  - What are the questions for-
    - The System Integrator
    - The Owner/Customer

- What are we going to do for Summit?
  - Striving for a largely a "hands off" operation.
  - We want the Operators know, how they interact with the system

- What should we do for exascale computing?

System Design - Upfront Design Specifications

- System Integrator Questions
  - Cabinet Count?
  - Max/Idle power draw per cabinet?
  - Min/Max air flow per cabinet?
  - What is the maximum allowed liquid flow rate through the cabinet?
  - What is the heat removal method for air and water?
  - % to water/% to air at worst case vs idle?
  - Have “parasitic” loads been considered?
  - Water flow and inlet temperature curves for 100% heat removal per cabinets at idle, normal, and max power draw? What’s “normal”?
  - What does the water side pressure drop vs flow look like between connection points? Cv
  - What are the allowable cooling fluid flow and temperature excursion magnitudes and durations?
  - Does the system self-protect or does it get too hot?
  - Do individual cabinets have flow control?
  - Water quality specifications?
  - What are the wetted materials within the cabinet?
  - What telemetry is available per node?

System Design – Upfront Design Specifications

- User/Owner/Customer Questions
  - Cabinet Count
  - What is expected system resource utilization?
    - Will power management be utilized?
    - Will the system be partitioned to run multiple jobs?
    - Will the jobs consist of “capacity jobs” and/or “commodity jobs”?
  - Will any power limiting measures be utilized?
    - Demand limiting?
    - What is the rate schedule of the utility provider?
    - How can the power usage of the cooling equipment be leveraged?
  - Can the rate of ramping from idle to normal/max be controlled?
  - Can a system be partitioned such that an outage on a cooling branch line or CDU level only impacts certain cabinets?
  - Will any data analytics be used to gather power usage profile data based on schedule and model types?
  - Will there be a method of communication between the HPC system and the facilities cooling system(s)?

System Design

System Level
Load Diversity

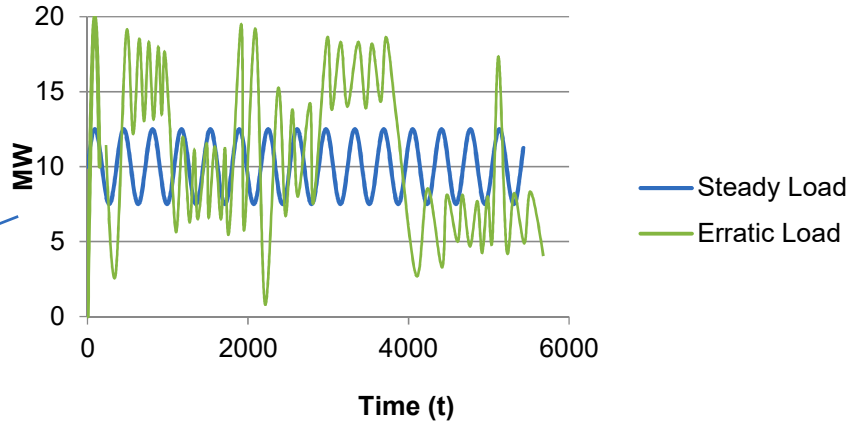

Rack Level
Load Diversity

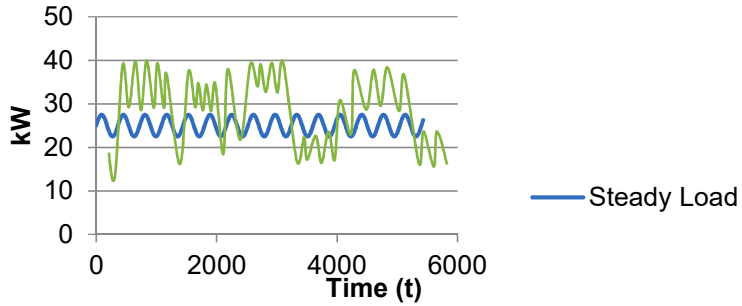

System Partitioned High
Rack Load Diversity

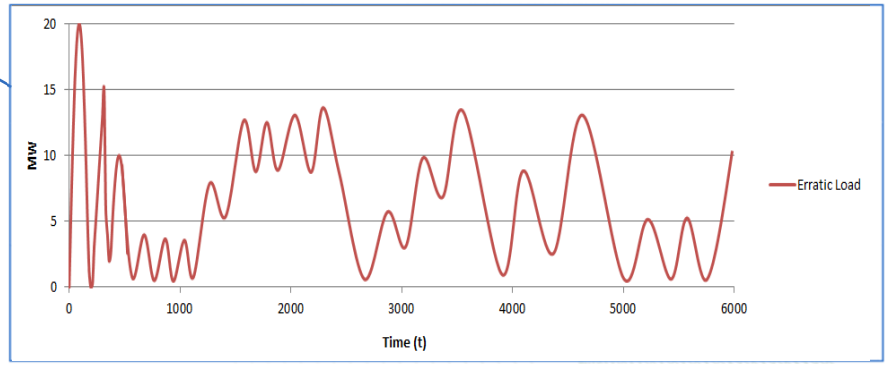

OAK RIDGE National Laboratory

Chiller Plant System Design

System Level
Load Diversity

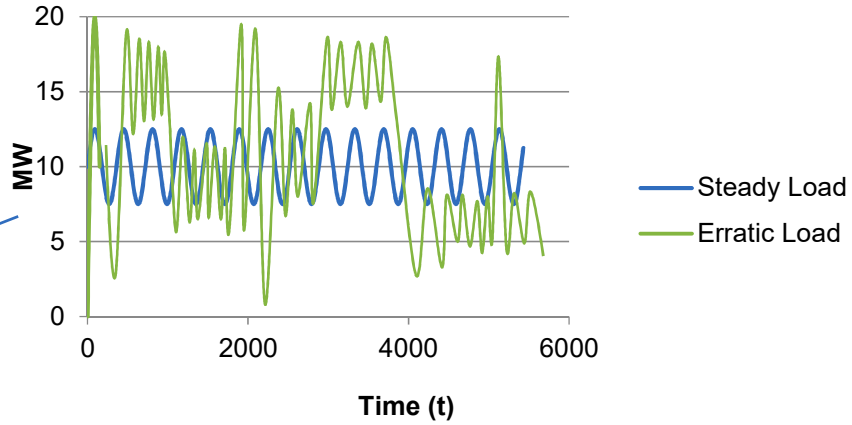

Rack Level
Load Diversity

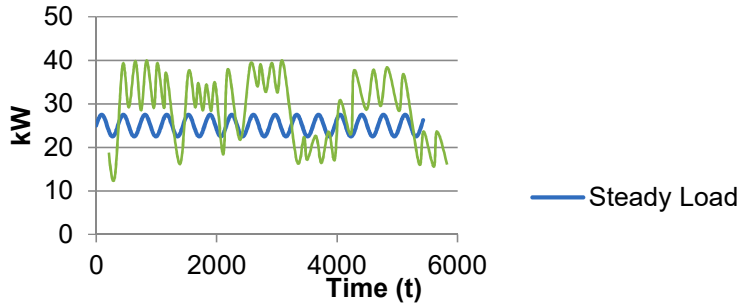

System Partitioned High
Rack Load Diversity

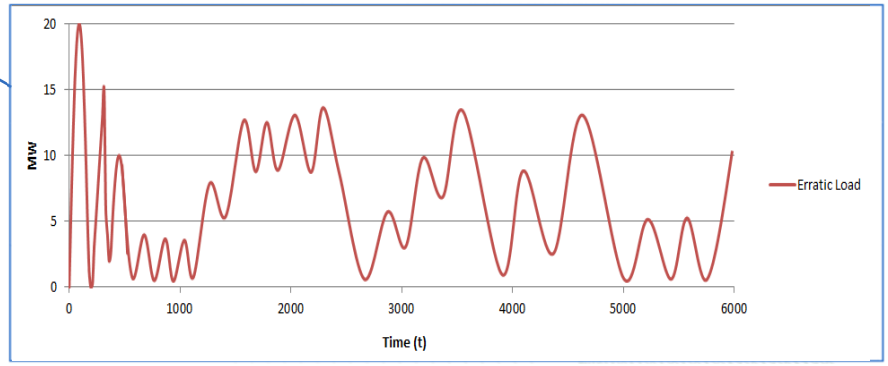

OAK RIDGE National Laboratory

Chilled Water Plant Load Diversity

- Each cooling system must be analyzed for vulnerabilities
  - Rate of load change
  - Outdoor conditions
  - System volume “turn-over” rate
  - What are the uptime requirements (wet bulb maximums?)
  - If chillers – more complicated – low dTs, high wet bulb “extreme-events”, rate of changes, etc.

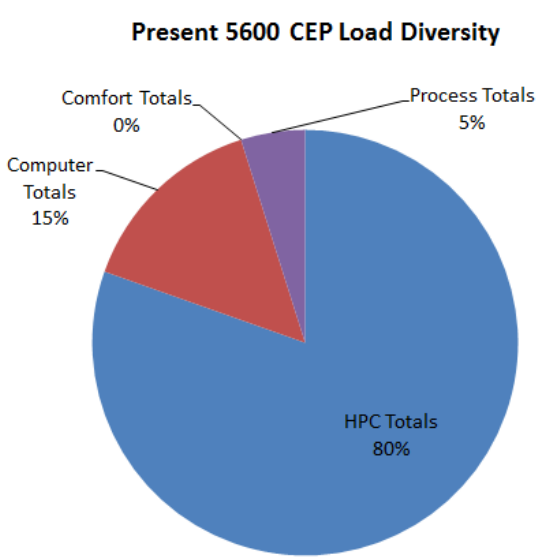

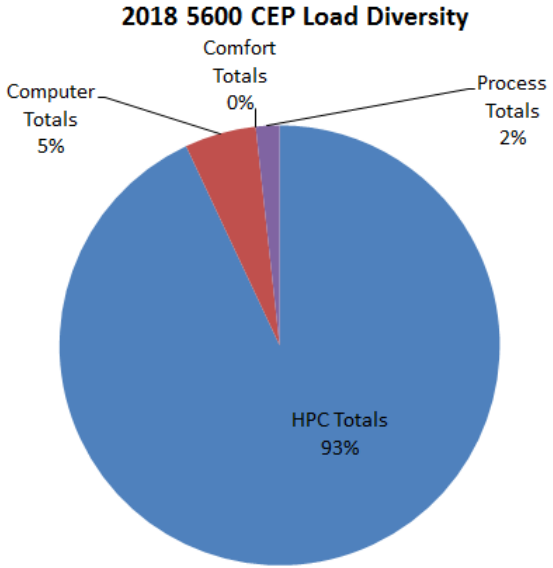

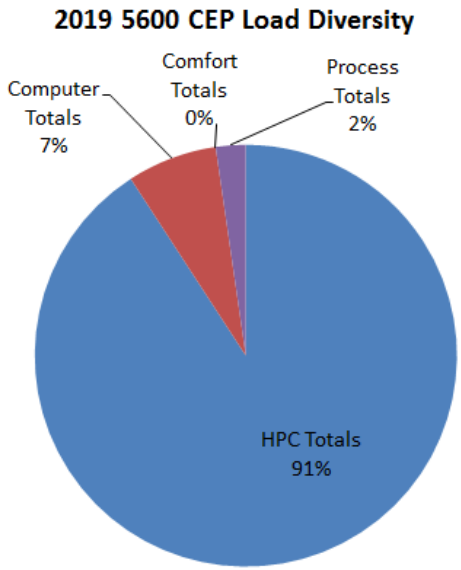

Chilled Water Plant Load Diversity

Possible Scenarios:

- Dedicated Chiller System for a single HPC System
- Campus chiller system w/ a single HPC system
- Campus chiller system w/ single HPC + Other + Comfort + Customer/Process Load
- Multiple HPC Systems
- Split between air cooled vs. water cooled

Transitions of Cooling Systems

- Opportunity exists where:
  - Variable speed pumping exists
    - Pump level
    - Cabinet or node level
  - Transitions between cooling systems with different COPs
  - Leveraging heat recovery

- Must always keep owner's desires for uptime and expected component life in mind.

What Cooling System Is Summit Fitting Into?

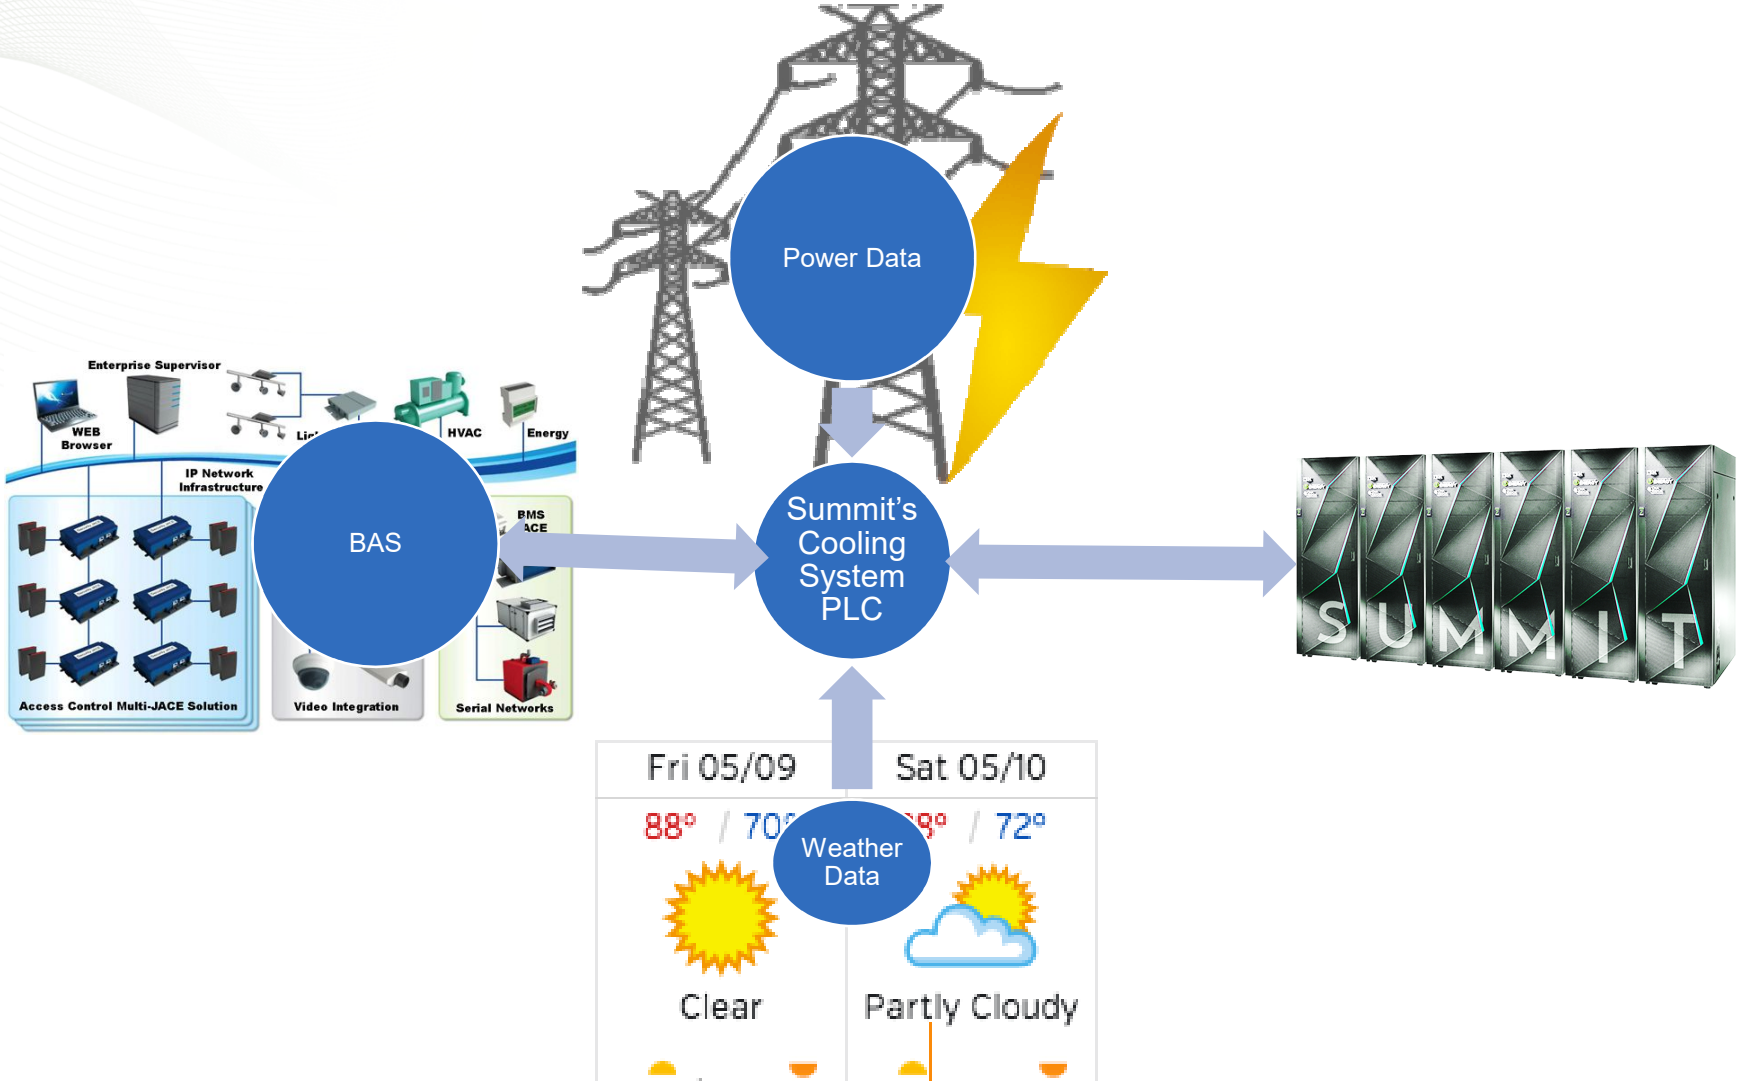

What will Summit need?

- Computer loads from 0-100% for this geographic area.
- Maintaining the temperature set point of the supply water to the computer system.
- Maintaining the minimum flow requirements of the computer system at the given supply water temperature.
- After satisfying minimum system requirements, the controls system stages equipment and operates at the highest efficient point possible.
- Monitoring of all system parameters, providing warnings and alarms. The thresholds, dead-bands, delays, and communication types for these warnings and alarms will be determined by the Design Team.
- Monitoring redundancy requirements and alert Operators should it fall below minimum requirements.
- A sudden loss of a cooling tower, MTW pump, chiller, CHW pump, or CW pump due to a power event or maintenance activity.
- A sudden load swing when the load on the chiller plant could put the chiller plant at risk of a shutdown.
- Control of the shifting of load between the economizing HXs and the trim HX's, including the required staging of chillers. This may require anticipatory controls.
- Design for efficiency, reliability, maintainability, scalabilty…

Summit - Cooling Sources

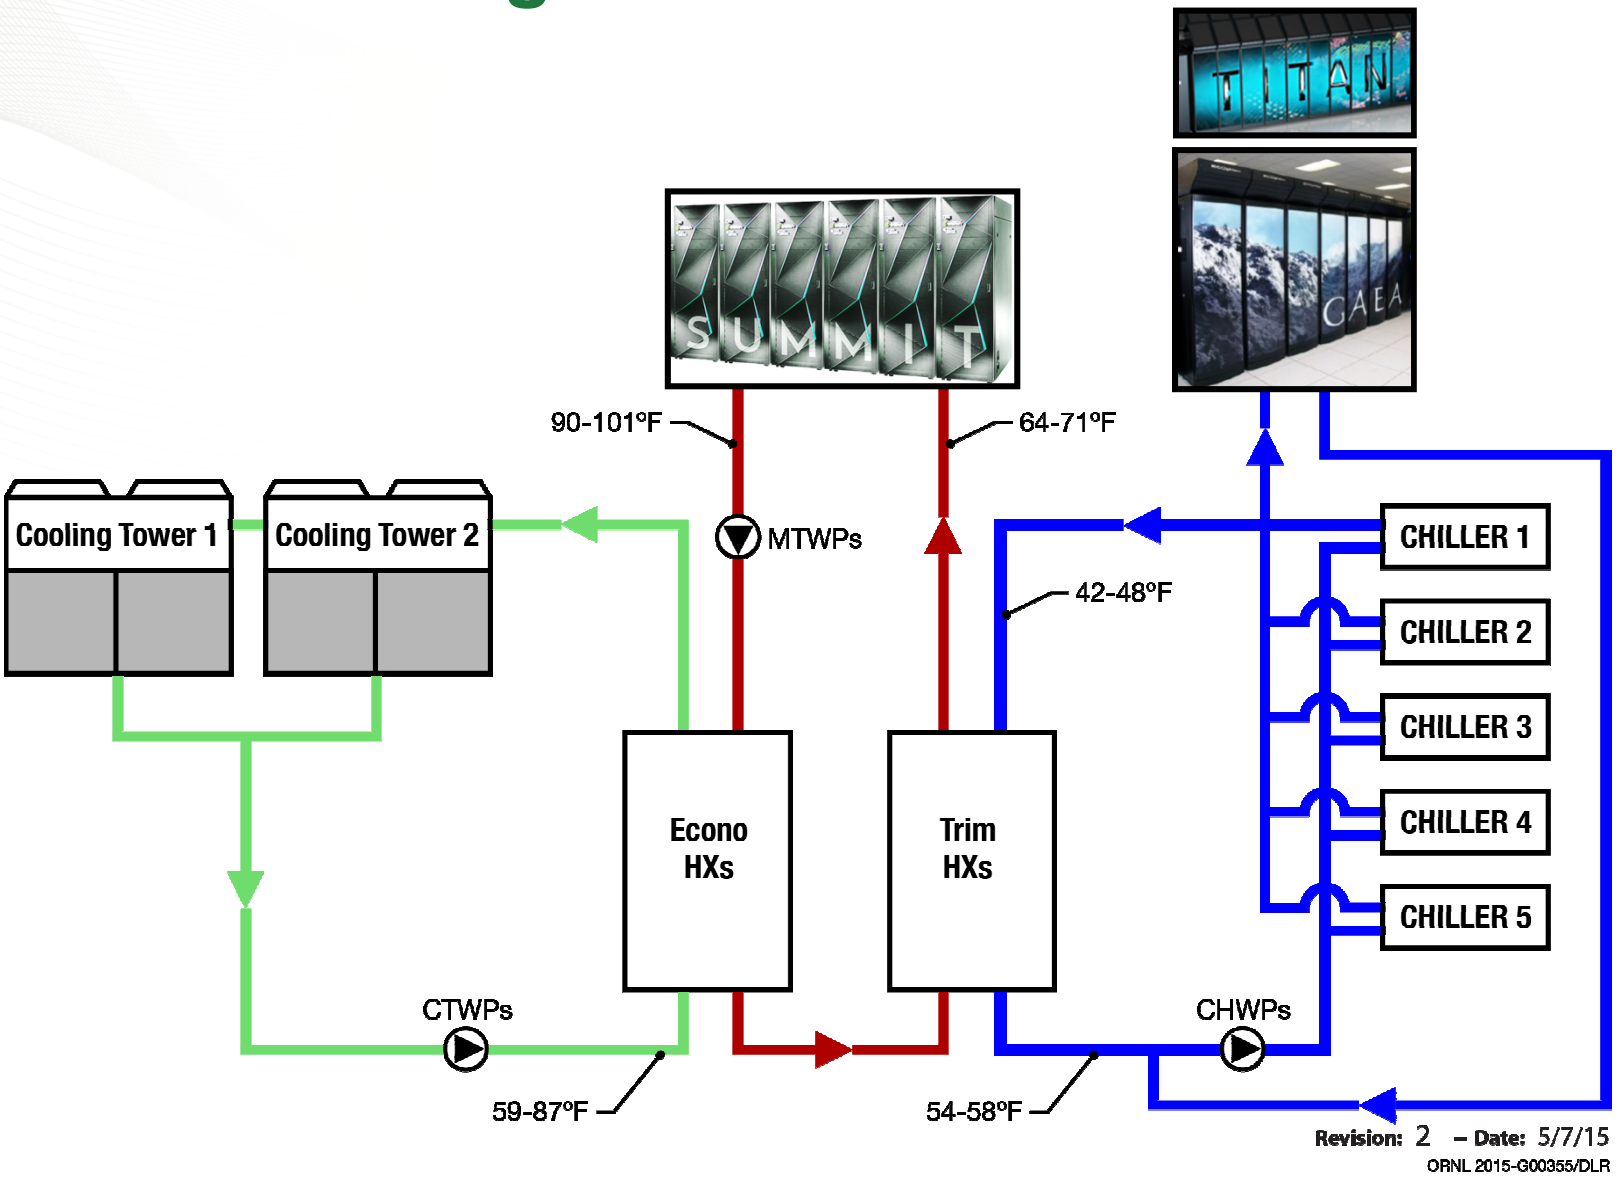

OAK RIDGE National Laboratory

Summit - Expected Cooling Source

- Variables
  - System Size (changes the slope of CHW demand)
  - Cabinet Load (changes the overall load)
  - Cabinet Cooling Supply Temperature (shifts the demand line left/right)
  - Approach Temperatures of HXs and Cooling Towers (shifts the demand line left/right)
  - Outside Air Wet Bulb Temperature

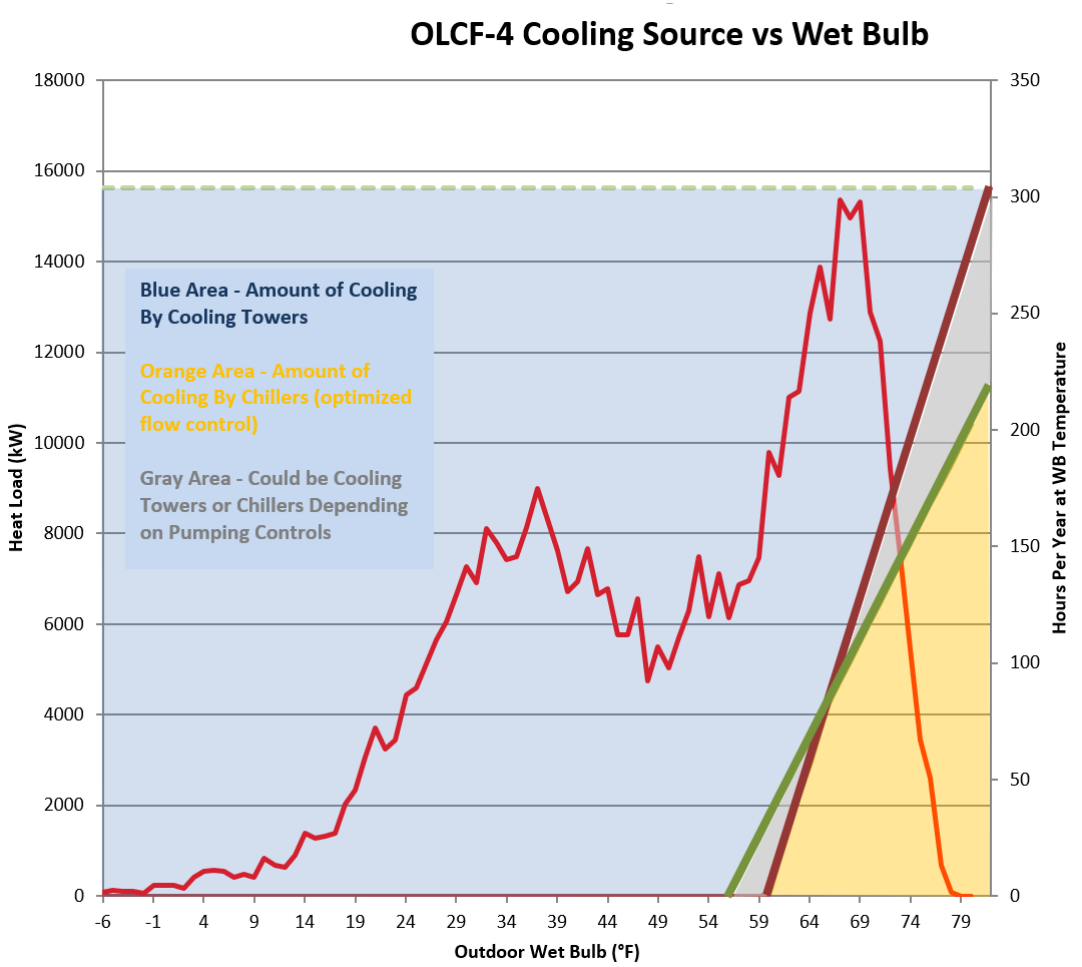

OAK RIDGE National Laboratory

Wet Bulb Data

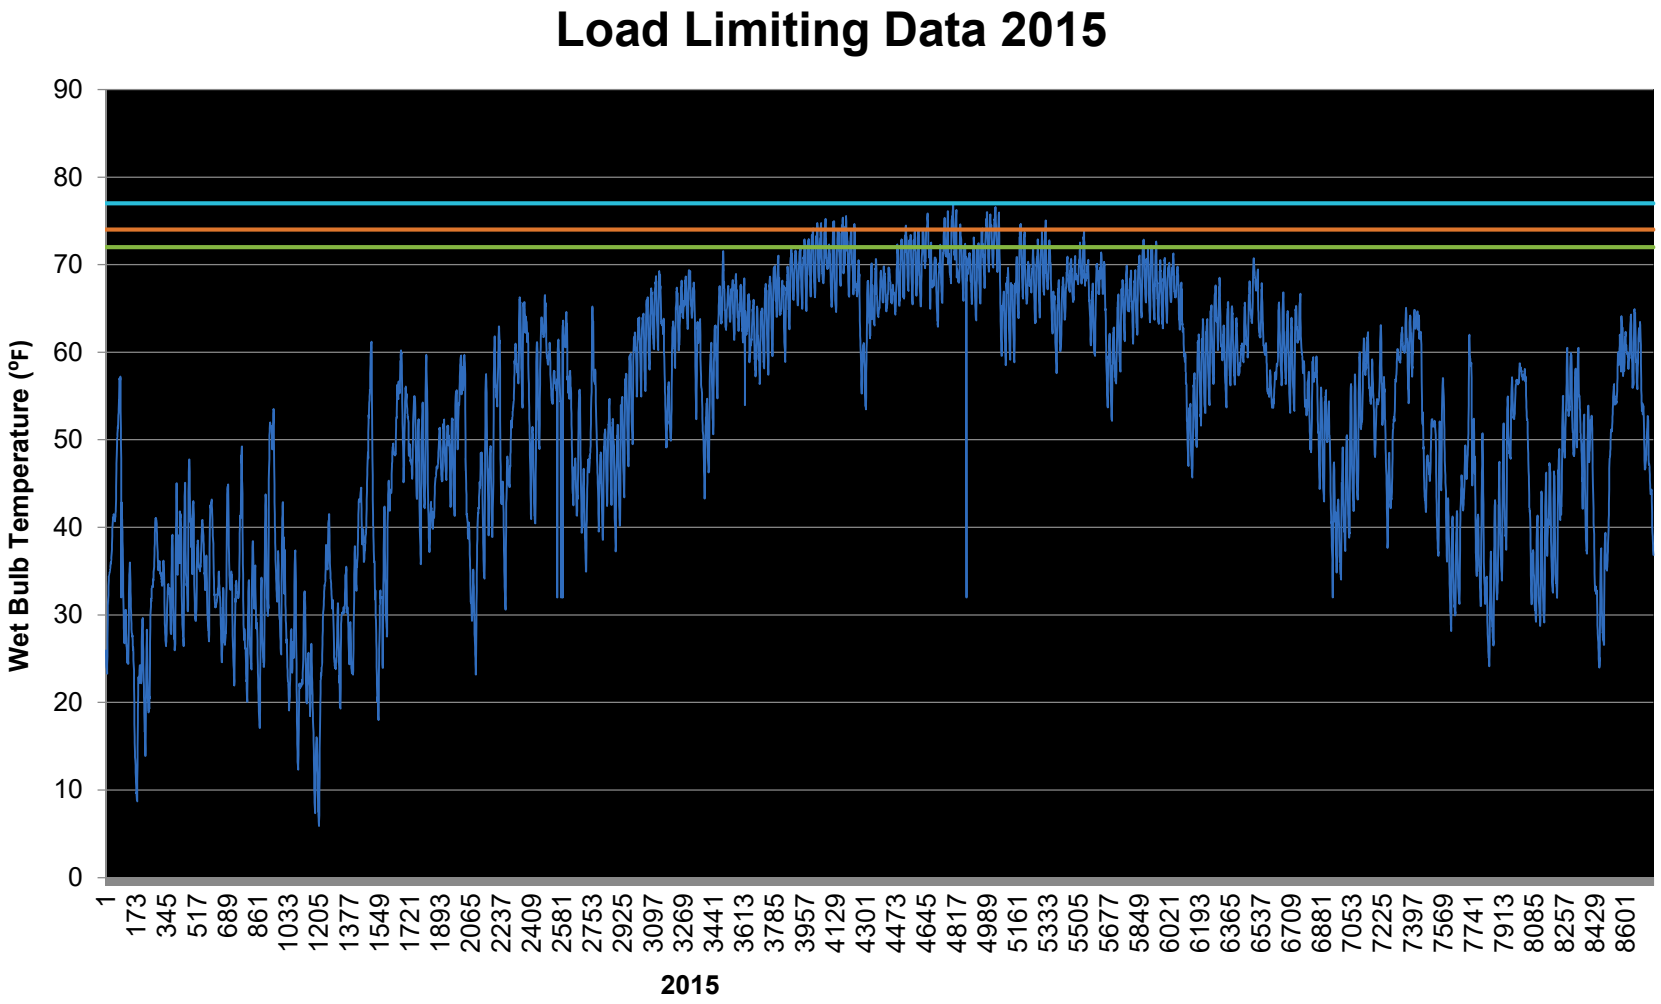

OAK RIDGE National Laboratory

Wet Bulb Data

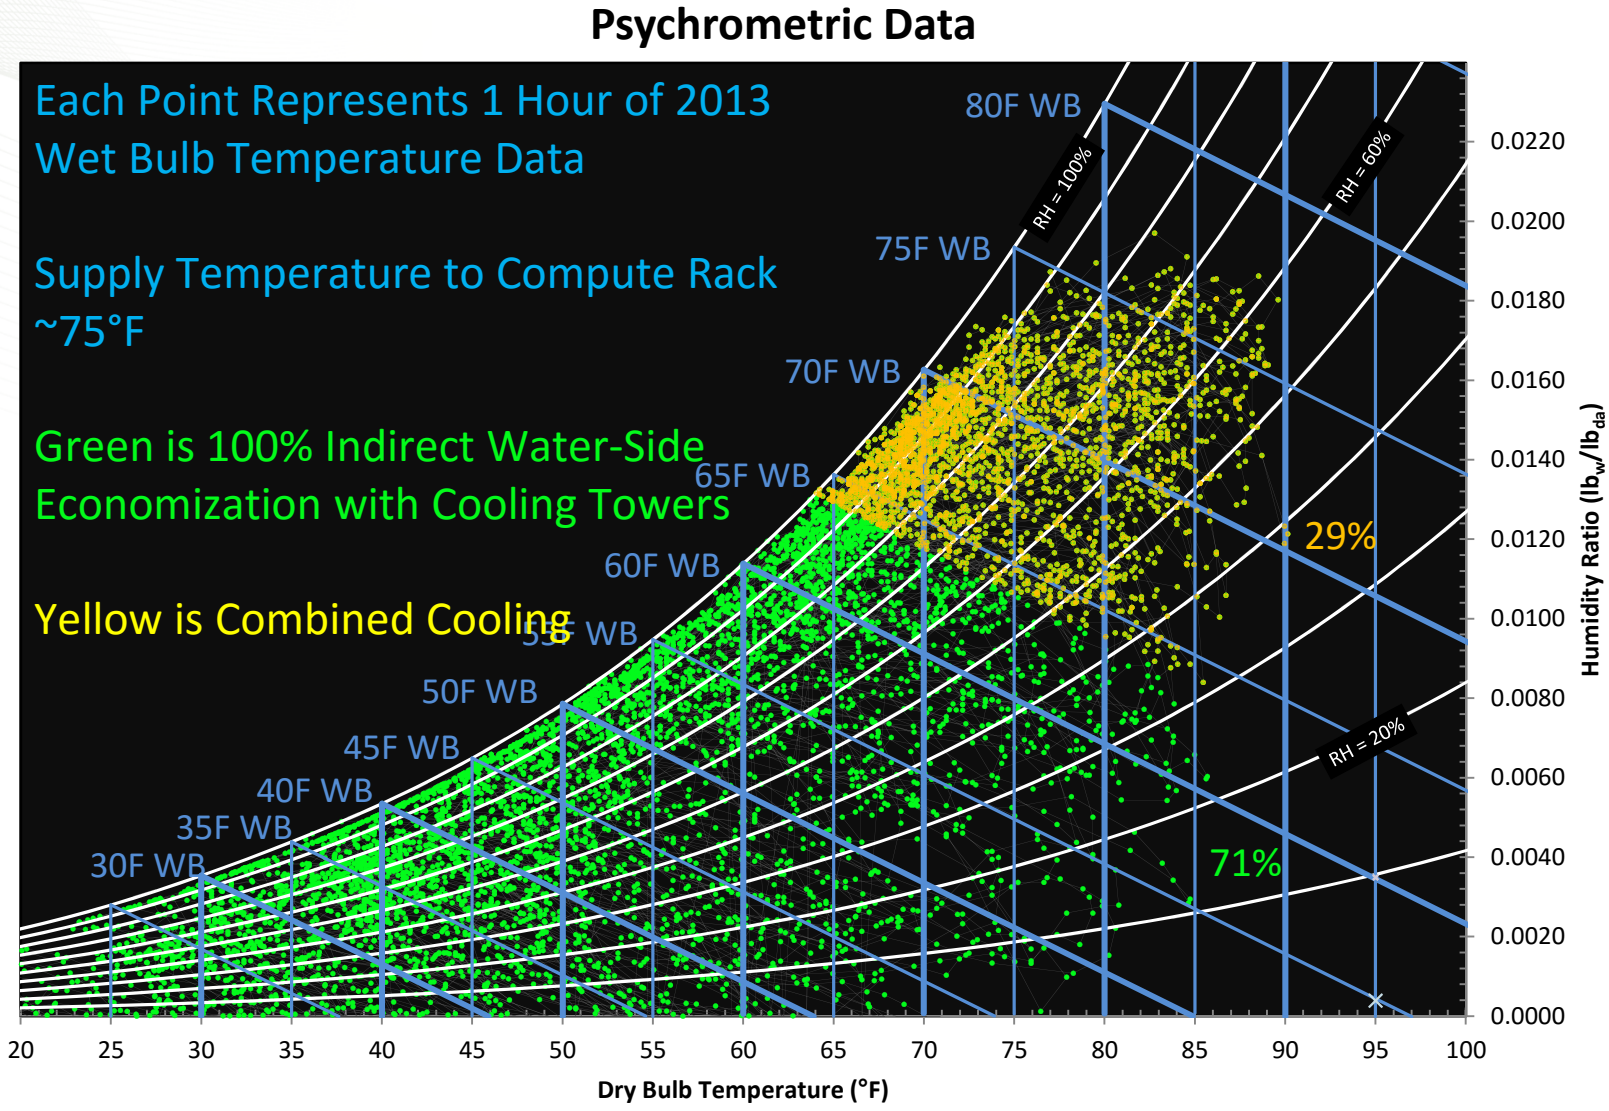

OAK RIDGE National Laboratory

What can change across the system boundary?

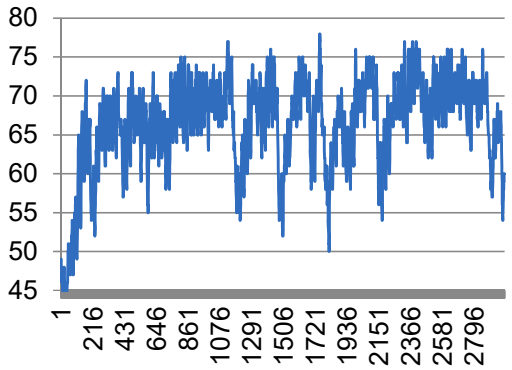

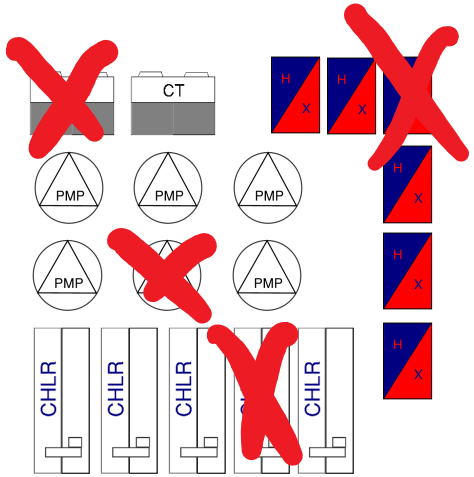

Computer System State	Outdoor Conditions	Equipment Availability and Performance
Off Idle Steady State Load Max Load Going between the states above at various rates of change	Design Day Beyond Design Day Low Wet Bulb Daily weather variations Seasonal weather variations	Design Capacity Plant Staging Maintenance Outages Performance

Control Scenarios – 180 to be Explored

OAK RIDGE National Laboratory

Controls Point List – For Summit 500+ points

Summit to PLC-

Scope	System	Point Type	Description	RTD	VIRTUAL	Connectivity/SCOPE OWNER					
						SYS-1	SYS-2 *Write TO PLC	PLC	Summit/API	Predicted OAWBT	SYSA/SYSB
SP-PH2	DC	AI	Flow Demand From Summit	-	-	Tertiary	Secondary	Secondary	Primary	-	-
SP-PH2	DC	AI	Predicted Total Computer Power From Summit	-	-	-	Secondary	Secondary	Primary	-	-
SP-PH2	DC	DI	Temperature Cooler/Warmer From Summit	-	-	Tertiary	Secondary	Secondary	Primary	-	-

PLC to Summit-

Scope	System	Point Type	Description	RTD	VIRTUAL	Connectivity/SCOPE OWNER					
						SYS-1	SYS-2 *Write TO PLC	PLC	Summit/API	Predicted OAWBT	SYSA/SYSB
CEPX-CTRLS	MTW	AI	Controlling MT Supply Temperature	-	Y	Tertiary	Secondary	Primary	Secondary	-	-
CEPX-CTRLS	MTW	AI	MT Flowrate	-	-	Tertiary	Secondary	Primary	Secondary	-	-
SP-PH1	DC	AI	Average Data Center Air Temperature	-	Y	Tertiary	Secondary	Primary	Secondary	-	-
SP- PH1	DC	AI	Available CEPx Capacity at Hour X	-	Y	Tertiary	Secondary	Primary	Secondary	-	-

Exascale

I hope to see:

1) A closed box with network, power, and cooling water supply/return connections (no direct air exchange with the ambient air in the data center)

2) Cooling water supply temperature no less than the location's 99.4% design wet bulb temperature + HX approach + cooling tower approach + 1

Wish List

Need to know more about IT networking and protocols.

What devices can act as network tied I/O devices: Arduino, beagle boards, raspberry pi, etc?

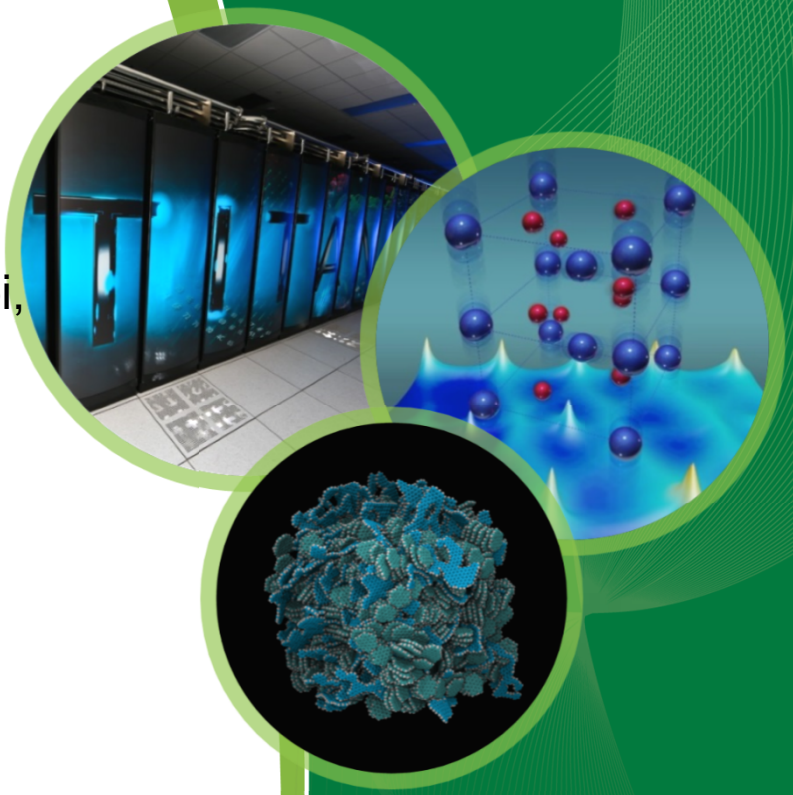

David Grant, PE, CEM, DCEP

HPC Mechanical Systems

Modernization Project Office

Oak Ridge National Laboratory