A Data Centric Approach to HPC Energy Reliability and Optimization

Ghaleb Abdulla

HPM 2016

August 24, 2016

Lawrence Livermore National Laboratory

LLNL-PRES-699373
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Acknowledgement

- Anna Maria Bailey
- Aniruddha Marathe
- Alfredo Gimenez
- John Weaver
- Barry Rountree
- Tapasya Patki

Human engineered systems are complex

- HPC Facility
  - HPC have combined load > 20 MW
    - Theoretical peak power requirements > 45 MW
    - Recurring intra-hour variability can exceed 8 MW
  - Electric Grid
    - Peak load during the summer can exceed planned generation
    - Introduction of renewables contributes to grid intermittency
- NIF
  - Very complicated control system with tens of thousands of sensors
  - Need to understand the interaction between the different system parts and relate to experimental results
  - Real time feedback and control

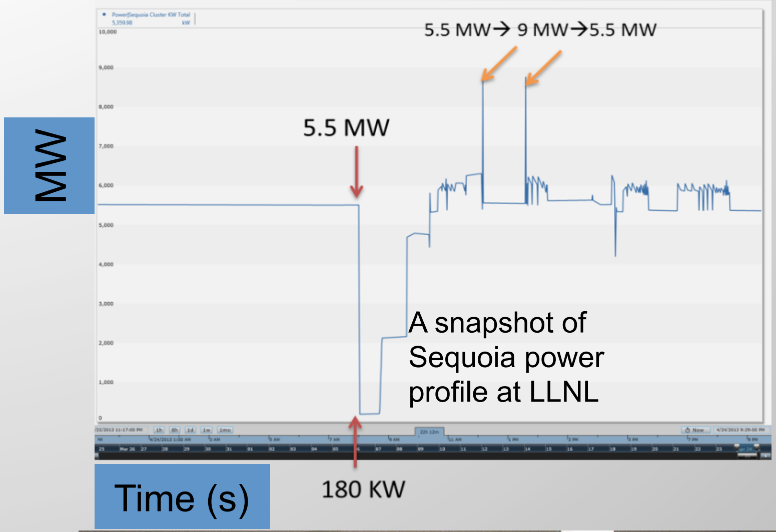

Relate component, system, and facility level data to define optimization strategies

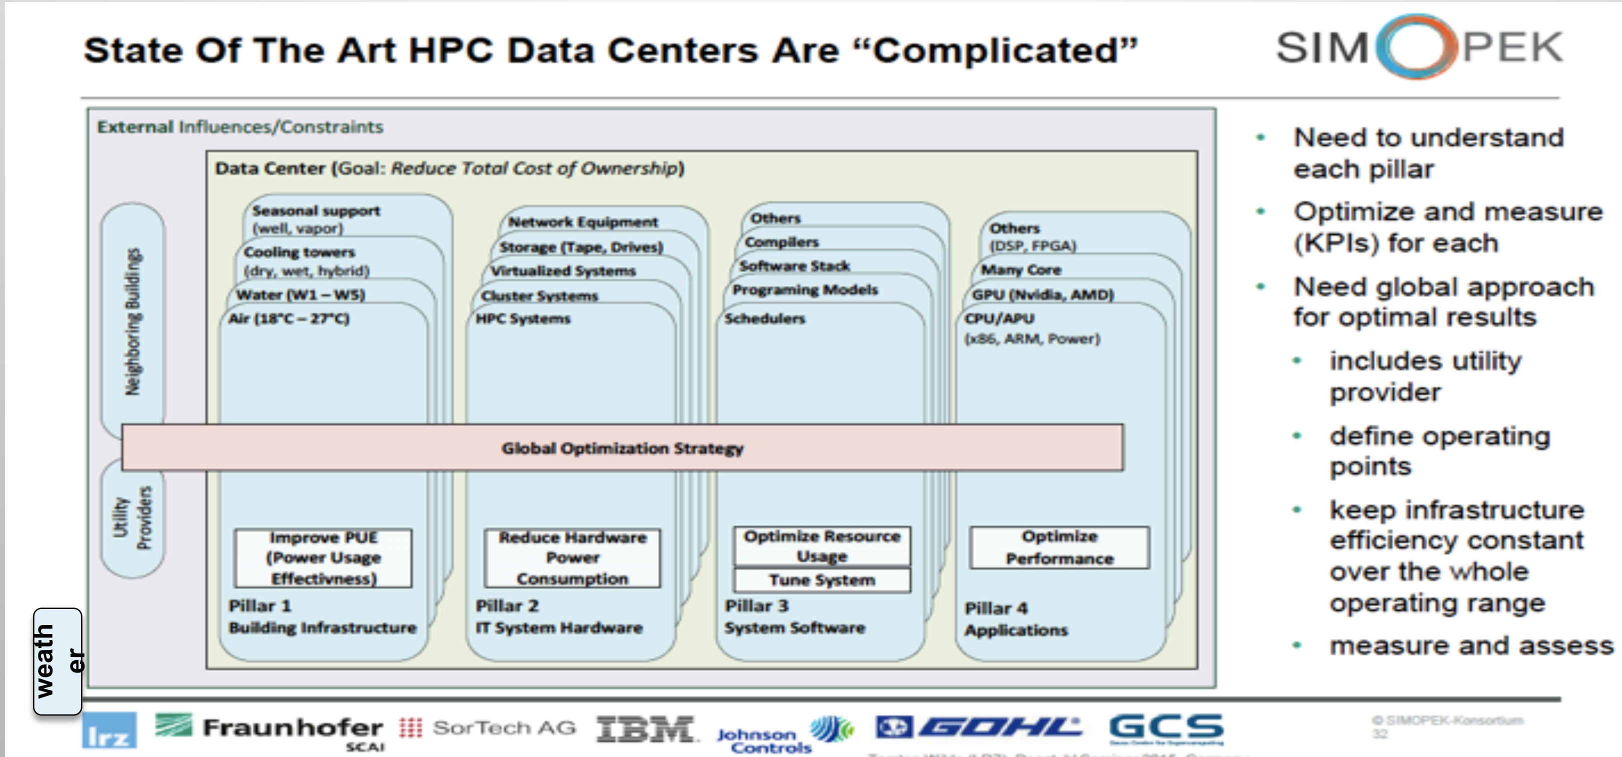

Slide curtesy of Torsten Wilde (LRZ) [www.simopek.de]; HPC Data Center Infrastructure Challenges Under A Power Bound, Dagstuhl Seminar 2015, Germany
Open Access 4 Pillar Framework Paper: http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s00450-013-0244-6

You cannot optimize what you cannot measure

- Collect data
  - Diverse interfaces and different data rates (µsec-15m)
  - All sources are equally important
- 70k tags
  - 20k HPC
  - 50k NIF

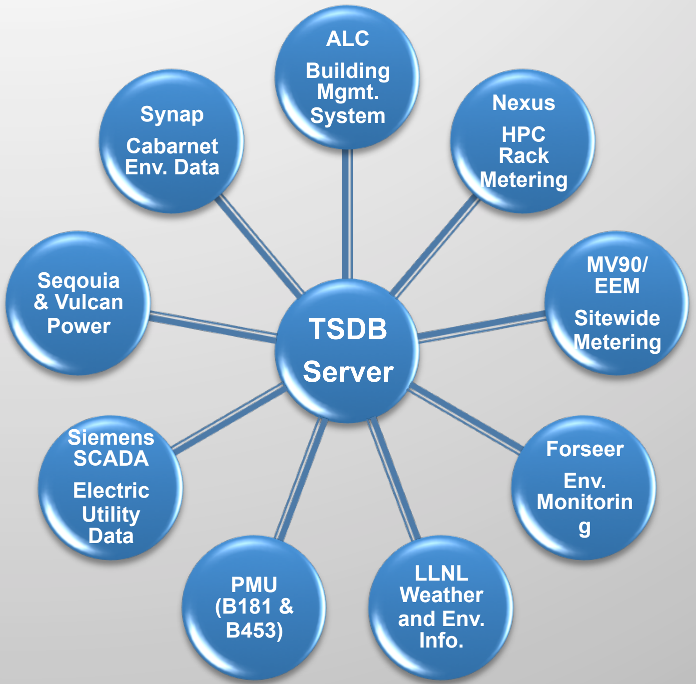

Data acquisition and management is necessary but time consuming

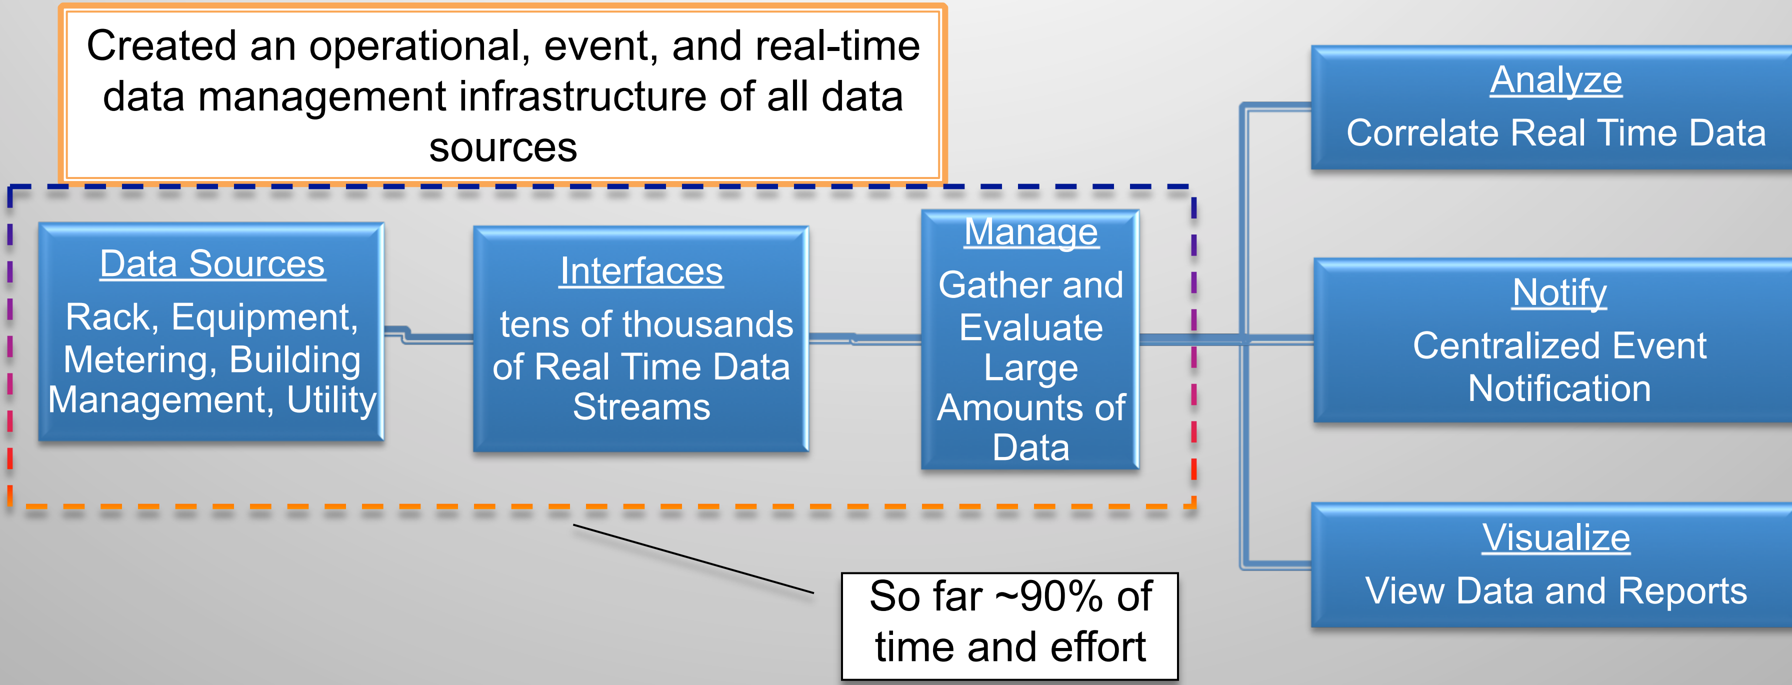

PI Server architecture

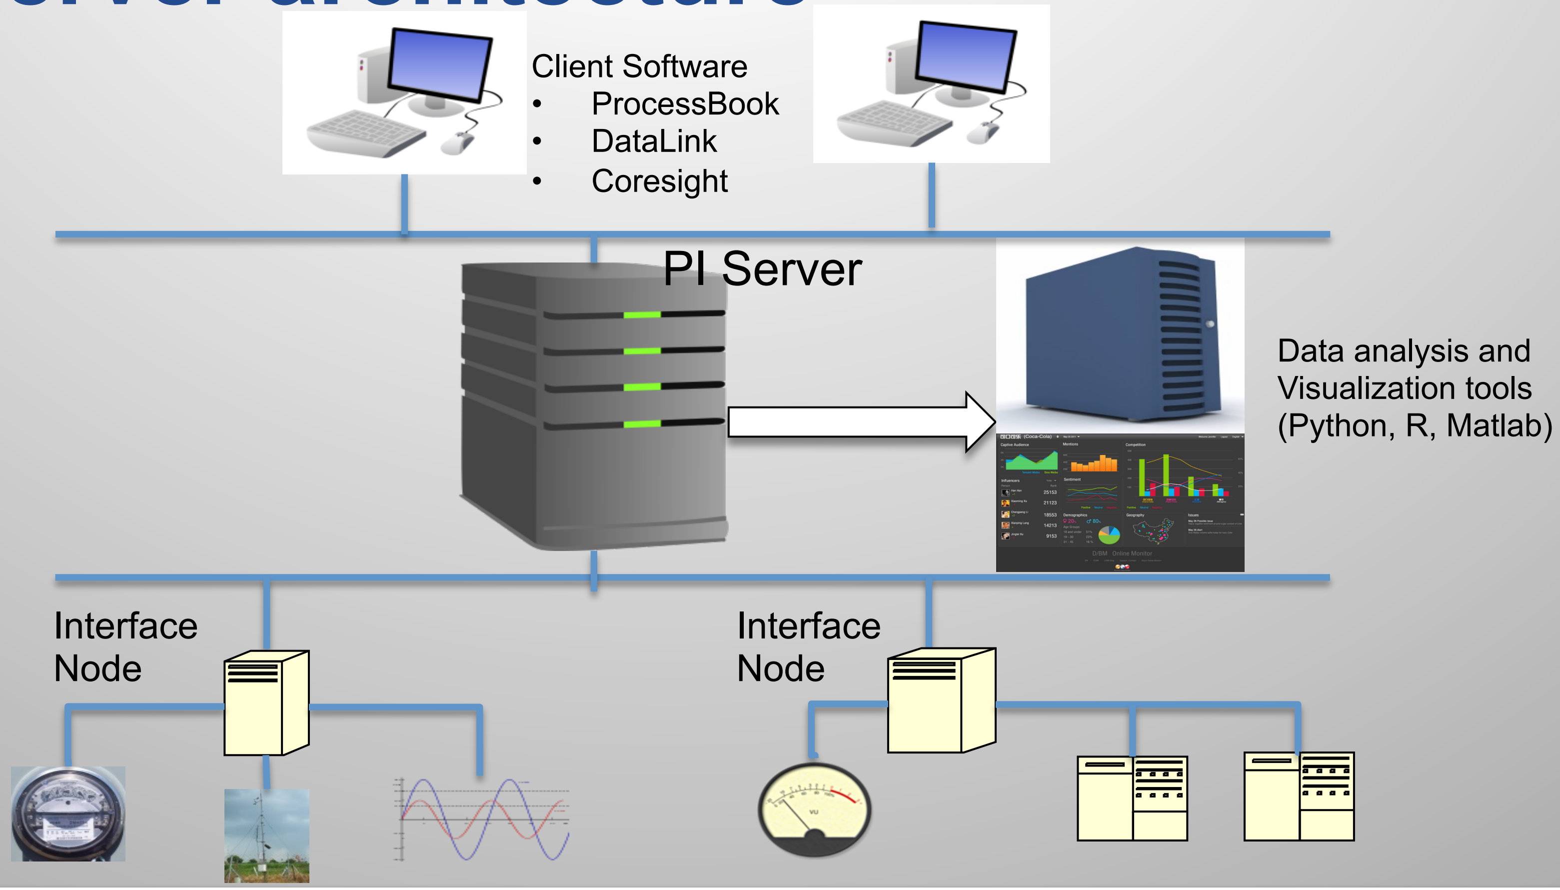

Efficient and user friendly data discovery interface

- Hierarchical data model reflects the physical structure of the monitored systems
- Multiple paths to the same data point
- Related assets

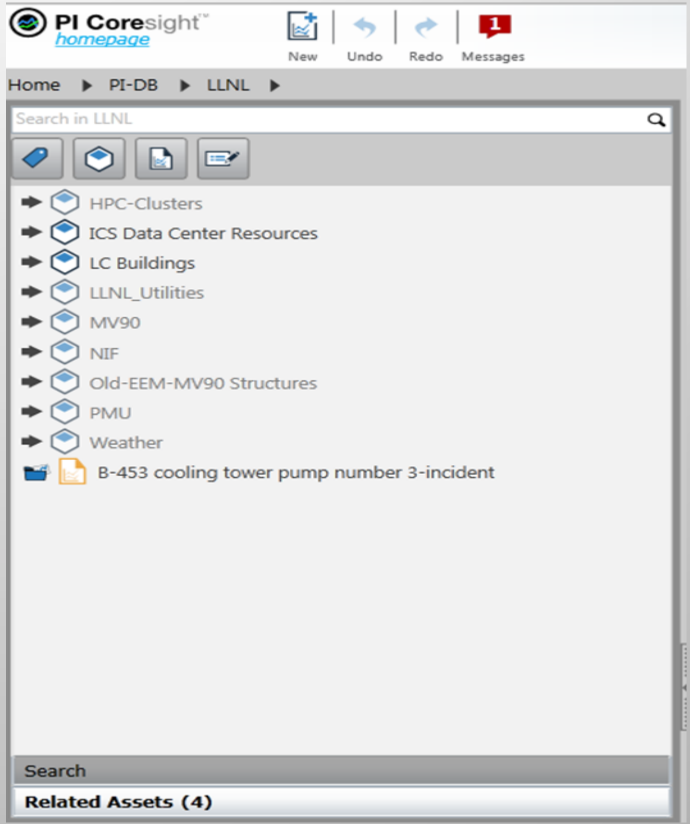

Example 1, Visual Analytics: Chiller shutdown events

- Date and Time: 10/30/2012 9:08:18 PM
- Buildings: 453 chiller #1
- T1883 fed by LGS-37, 3705.

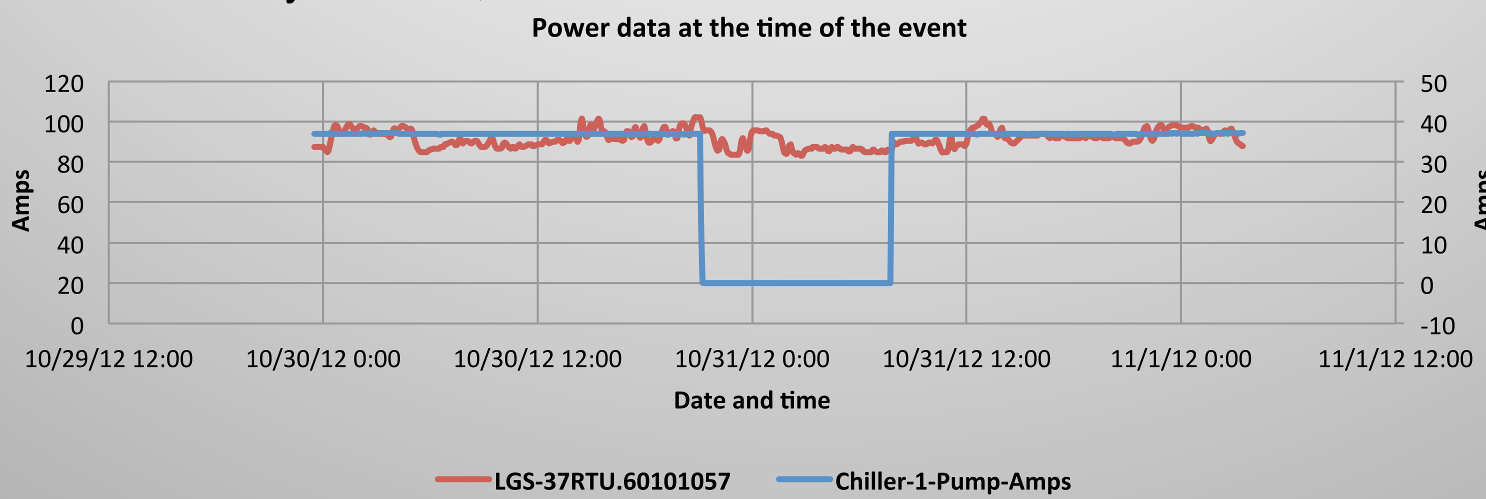

Data over a year explains the chiller behavior

- The pump shuts down regularly and more frequently during the Summer
- No correlation with the electric signal
- Chiller is used for “staging”
    - Solution: redesign the system and replace with a larger chiller

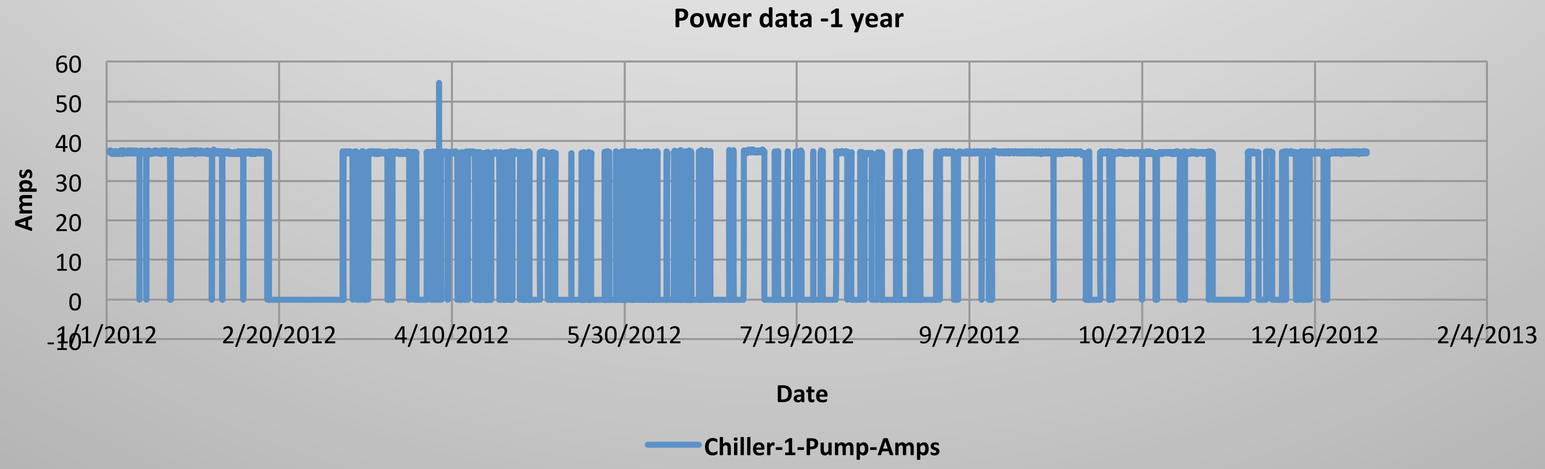

Example 2, Visual analytics: Correlating different data streams

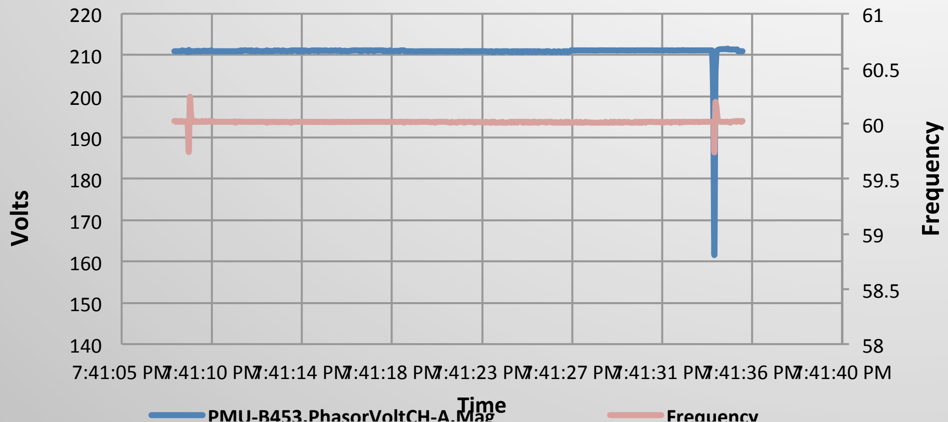

- Voltage dip:
  - 7:41 pm on 10/27/2013
  - 10:51 pm on 02/08/2015

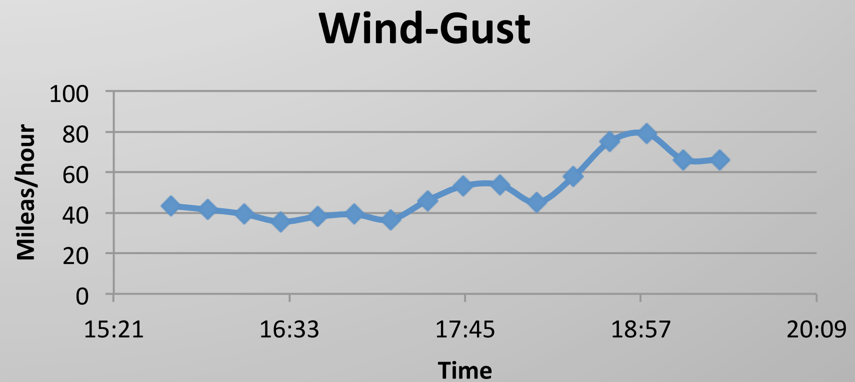

- Wind with gusts over 67 m/h
  - SW and SSW direction
- Winds gusting below 67 m/h
  - Other directions

Visual analytics: Correlating different data streams

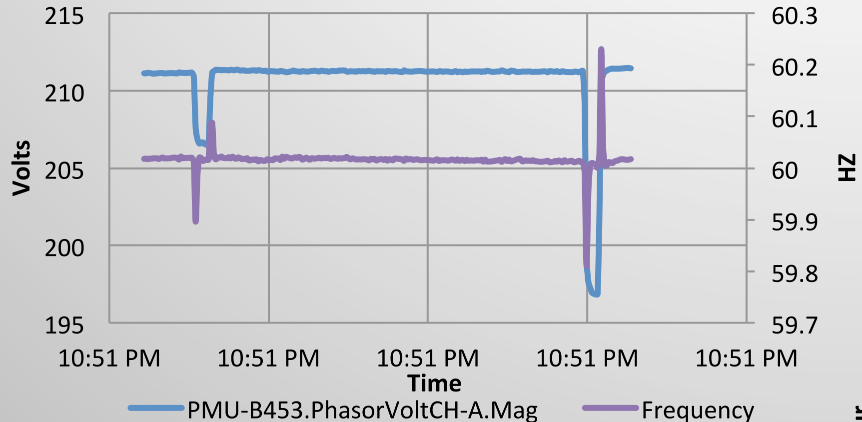

- 2/28/2015 event
- Very similar behavior

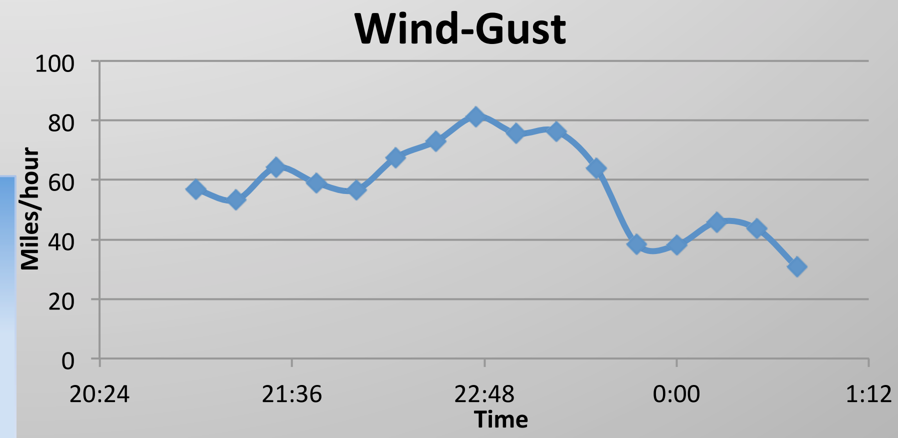

Two important steps:

1. Find similar events
2. Find what caused these events

Can we automate the two steps?

1. Search
2. Cause and effect, correlation

Example 3: Retrofitting the Cab cluster with liquid cooling technology

- **Cabernet (CAB) Overview**
  - Appro System
    - Intel Xeon ES-2670
  - OS – TOSS
  - Interconnect – IB QDR
  - 426 TeraFLOP/s peak
  - Memory 41,472 GB
  - 1296 nodes, 16 cores/node
  - Power – 564kW in 675 ft^2
- #94 on November, 2013 Top 500

Compare efficiency and reliability of liquid cooling versus air cooling

Run a Dedicated Access Test on Cab before and after the retrofitting

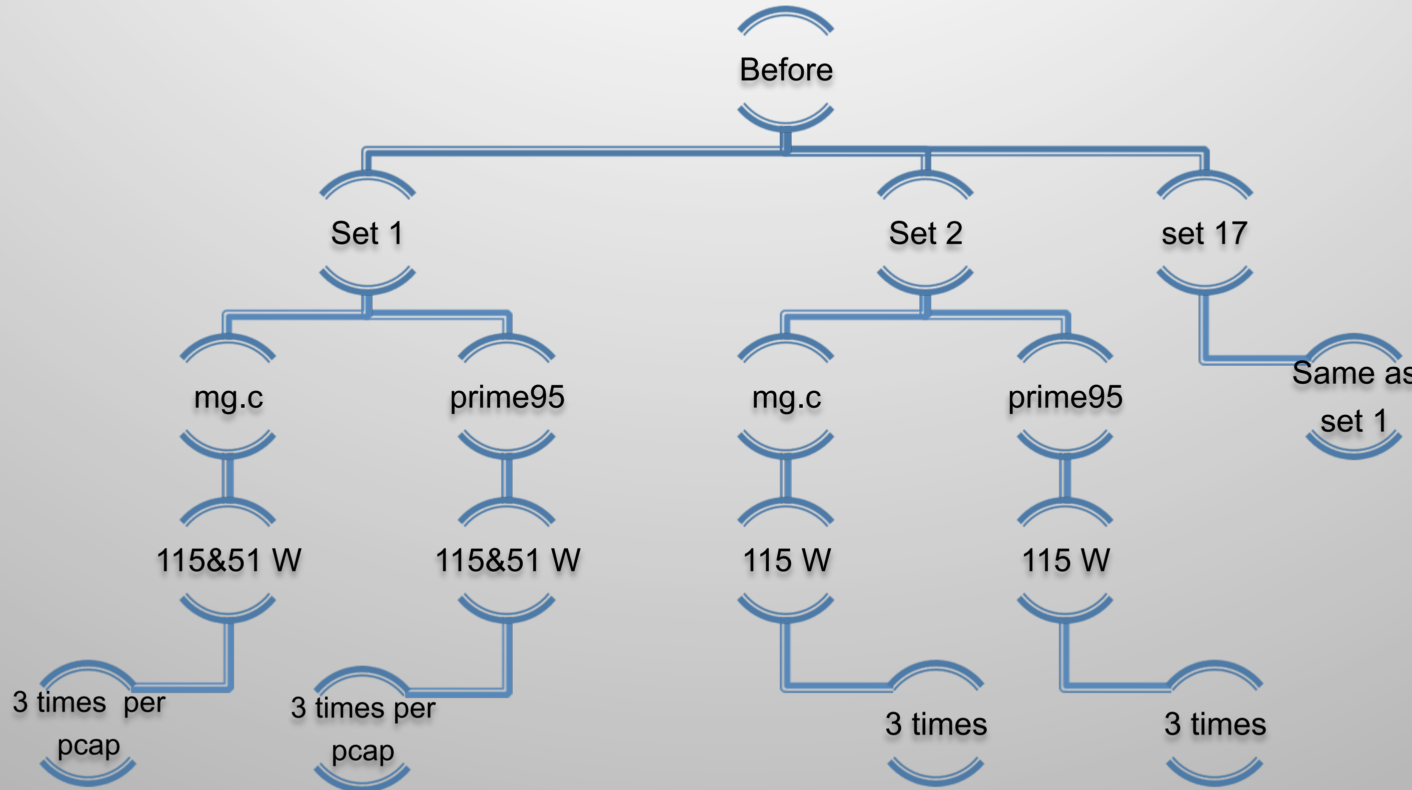

Power profile for Cab while running the DAT

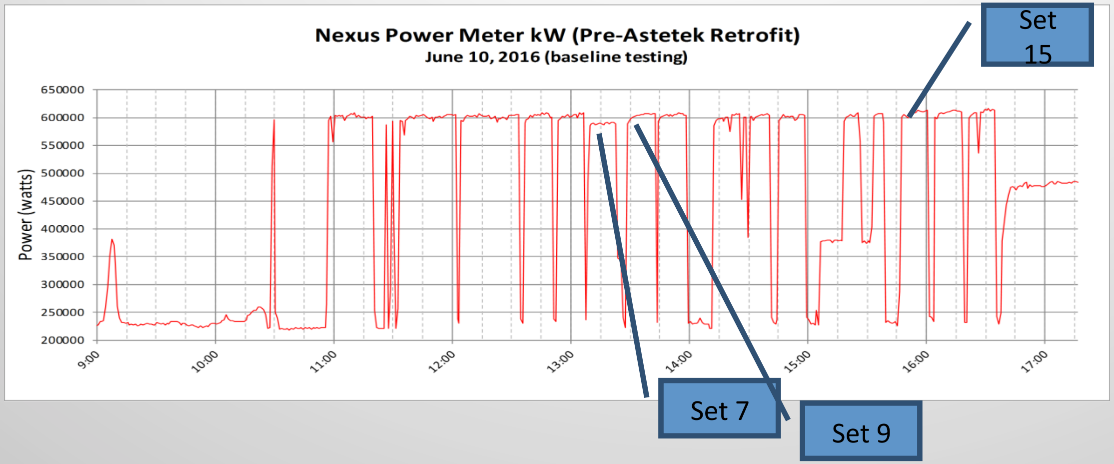

Power profile for set 9

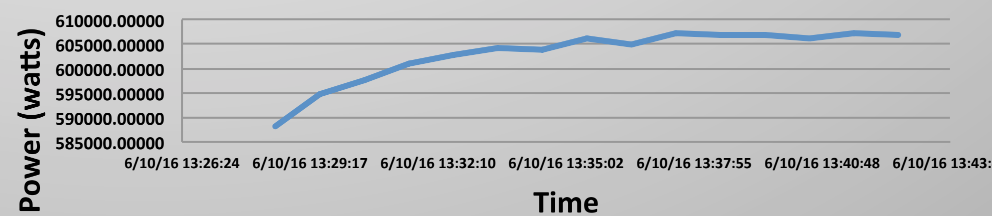

Large data set collected using different tools and requires its own data model

What is next?

- Used PI for:
  - Historian
    - Efficient/ real time data ingest and analysis
  - Event analysis
  - Correlate time series data
  - Enable simple visualizations and dashboards

- PI has limitations:
  - We cannot integrate non time series data
  - Data pre-processing
    - Especially experimental data
  - Scalable data analysis
    - Hard to get large data out

Working on a scalable data management and analysis system

Potential Data Sources from an HPC facility

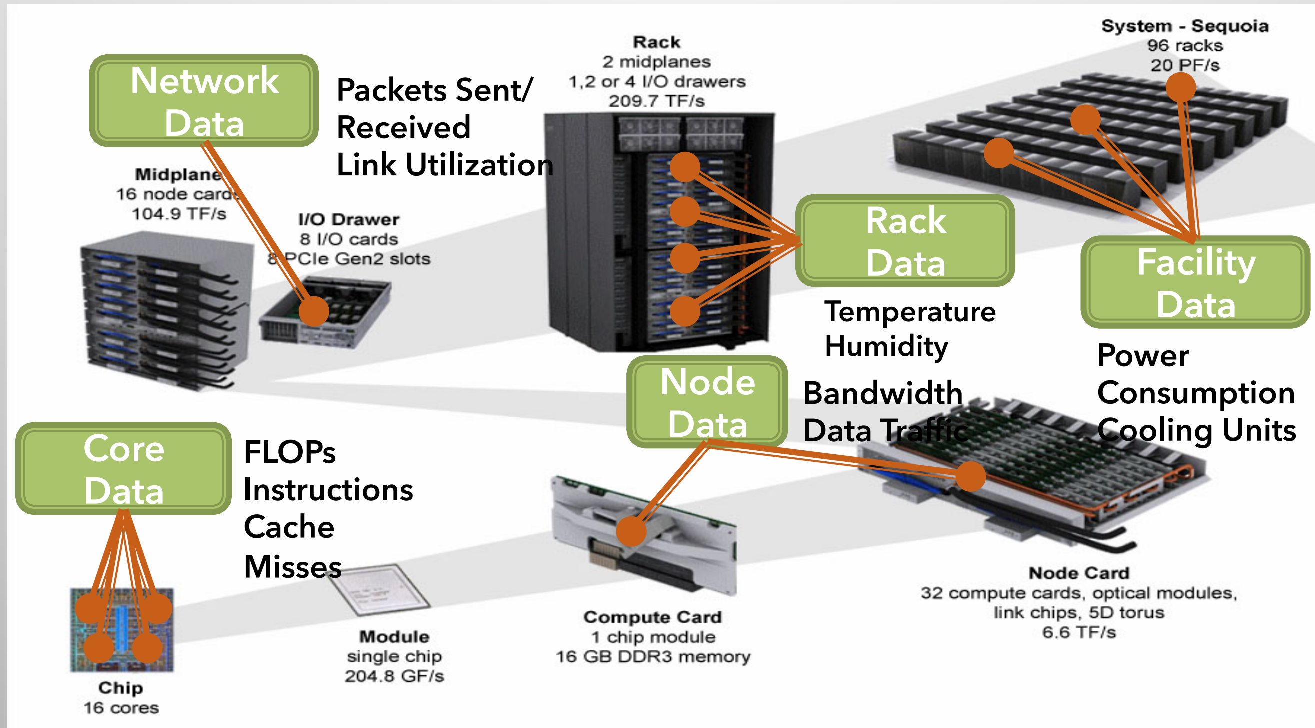

Data driven optimization strategy

- Enable data ingestion with different data rates
- Define data semantics
- Define different data transformations
- Enable querying of data using data sources, transformations, and semantics
- Scalable DM using Spark and Cassandra

Rack X Time X Heat X JobList

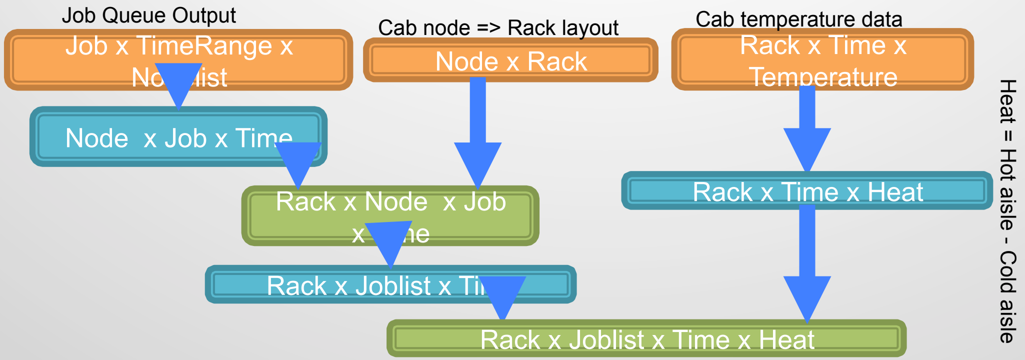

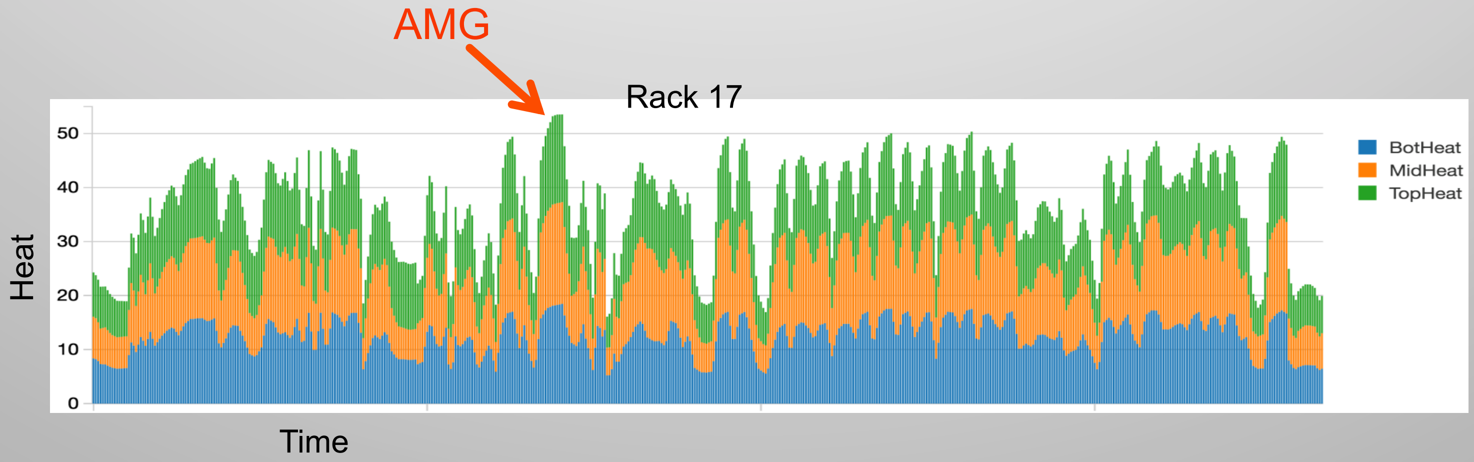

Summary
Wish list

- "You cannot control what you cannot measure"
  - Major effort to collect and manage relevant data
  - Still missing important data
- Visual analytics
  - Easy data access and quick visual displays are very powerful
    - Develop better Vis tools to help Define questions for advanced data analytics
    - Automate event and correlation analysis
- Sensor data are mostly time based
  - TSDB is appropriate for managing this type of data, however we collect and correlate other data types
  - Need to :
    - integrate with other data types
    - Spark based scalable data analysis
    - Real time prediction and control