SynapSense™ Wireless Environmental Monitoring System of the RHIC & ATLAS Computing Facility at BNL

Alexandr Zaytsev
alezayt@bnl.gov

BNL, USA
RHIC & ATLAS Computing Facility

Computing in High Energy and Nuclear Physics 2013 (CHEP2013), Amsterdam, Netherlands, Oct 14-18, 2013

Outline of the Talk

- RACF facility overview
- Environmental monitoring systems for RACF
  - Historical perspective
  - Requirements and design considerations for a new unified environmental monitoring system
- SynapSense deployment and status
  - Deployment timeline
  - Architecture of the data gathering mechanisms
  - Experience gathered with the system so far
- Summary & future plans
- Q & A

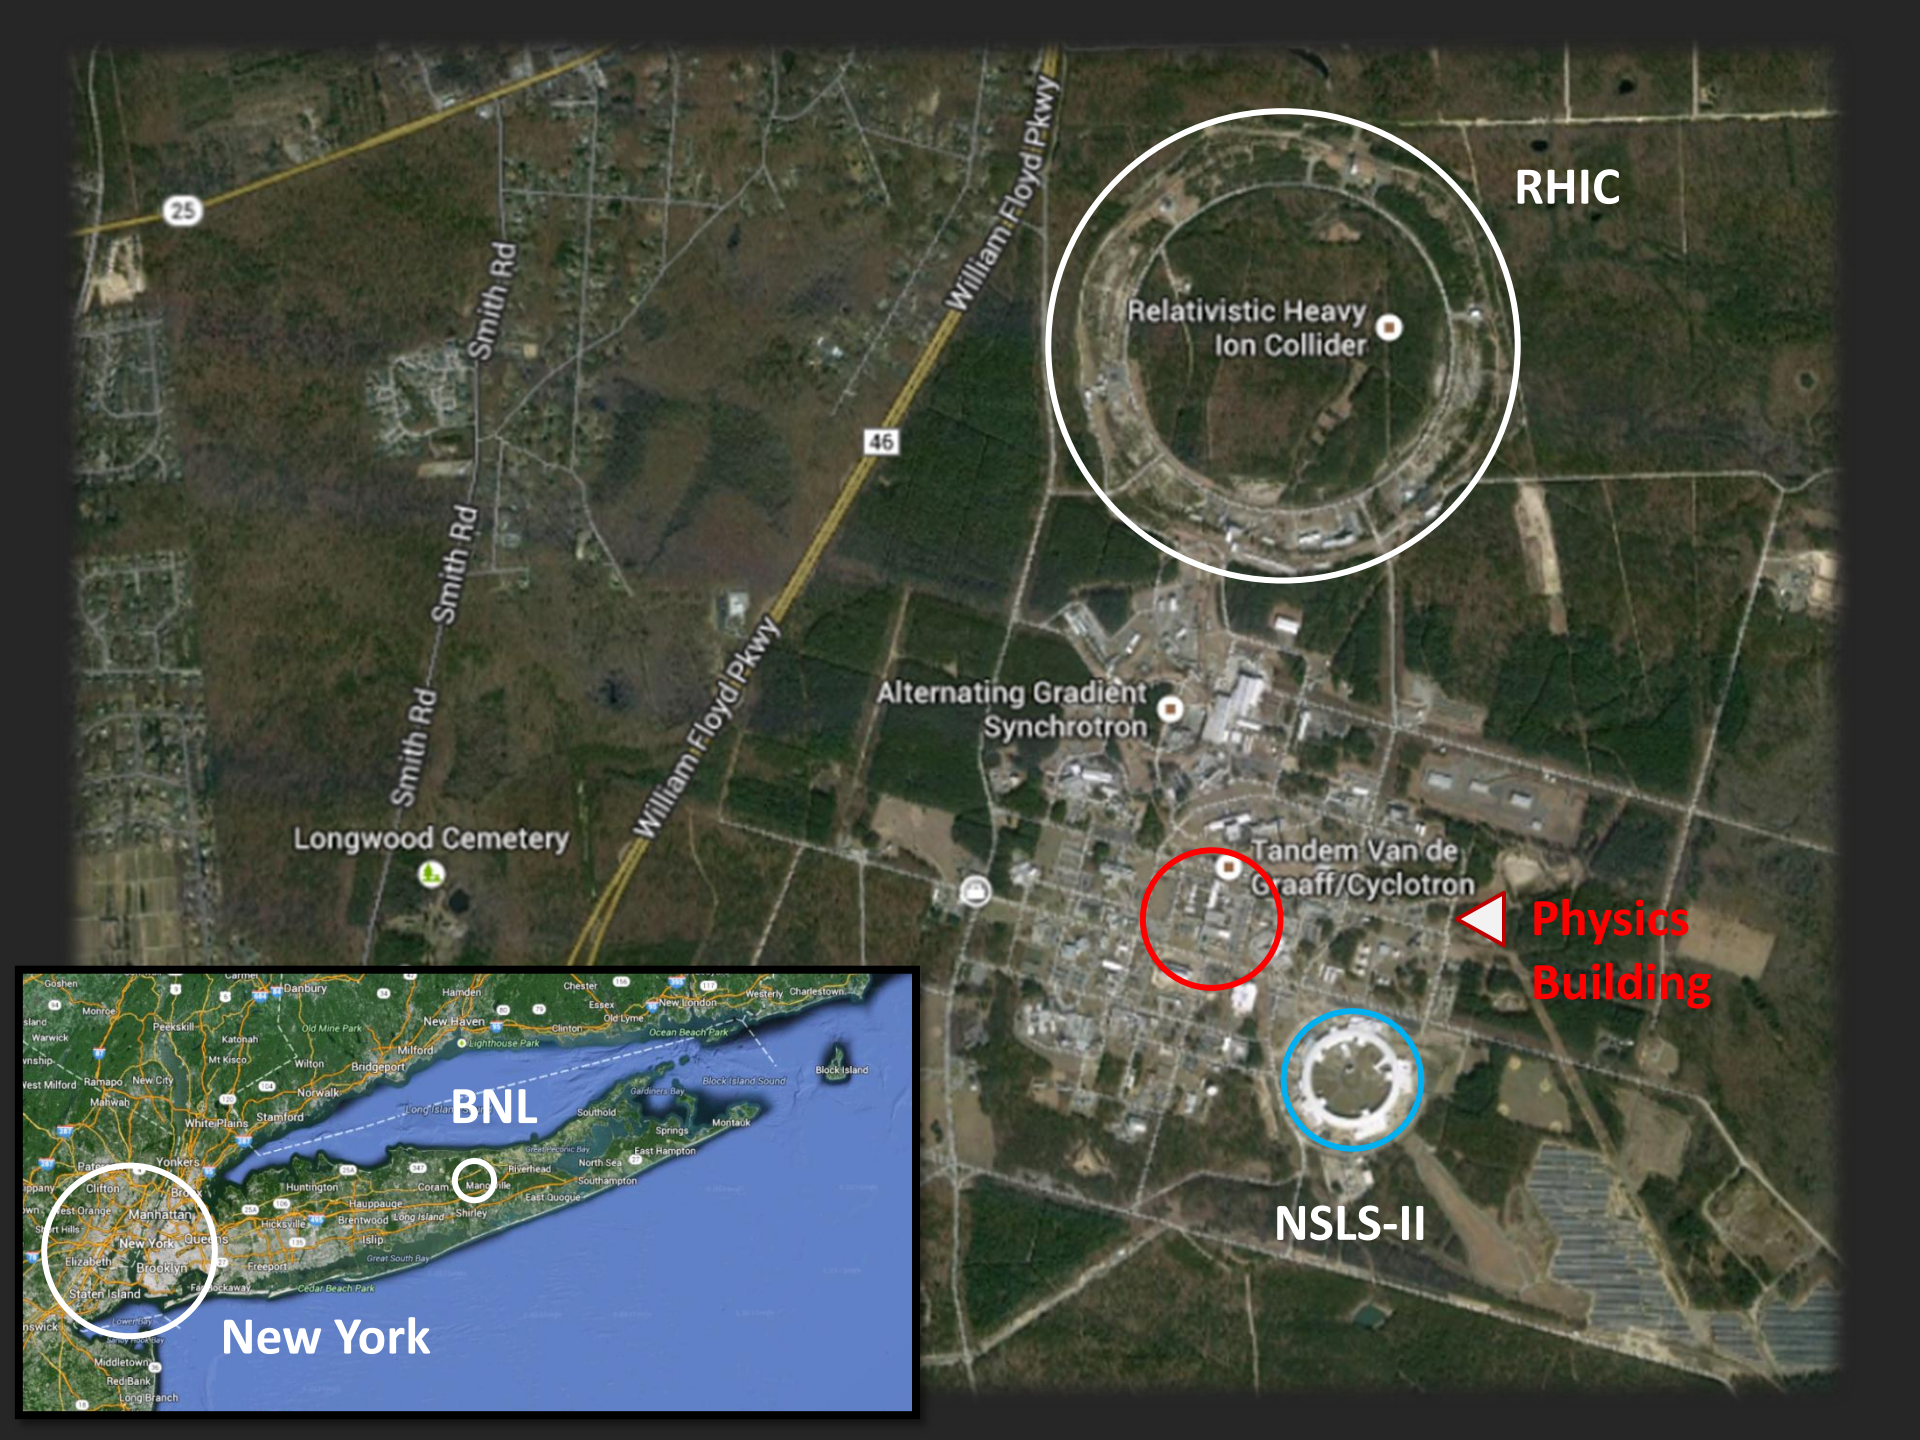
RHIC
Relativistic Heavy Ion Collider
Alternating Gradient Synchrotron
Tandem Van de Graaff/Cyclotron
Physics Building
NSLS-II
Longwood Cemetery
Smith Rd
Smith Rd
Smith Rd
William Floyd Pkwy
William Floyd Pkwy
25
46
BNL
New York

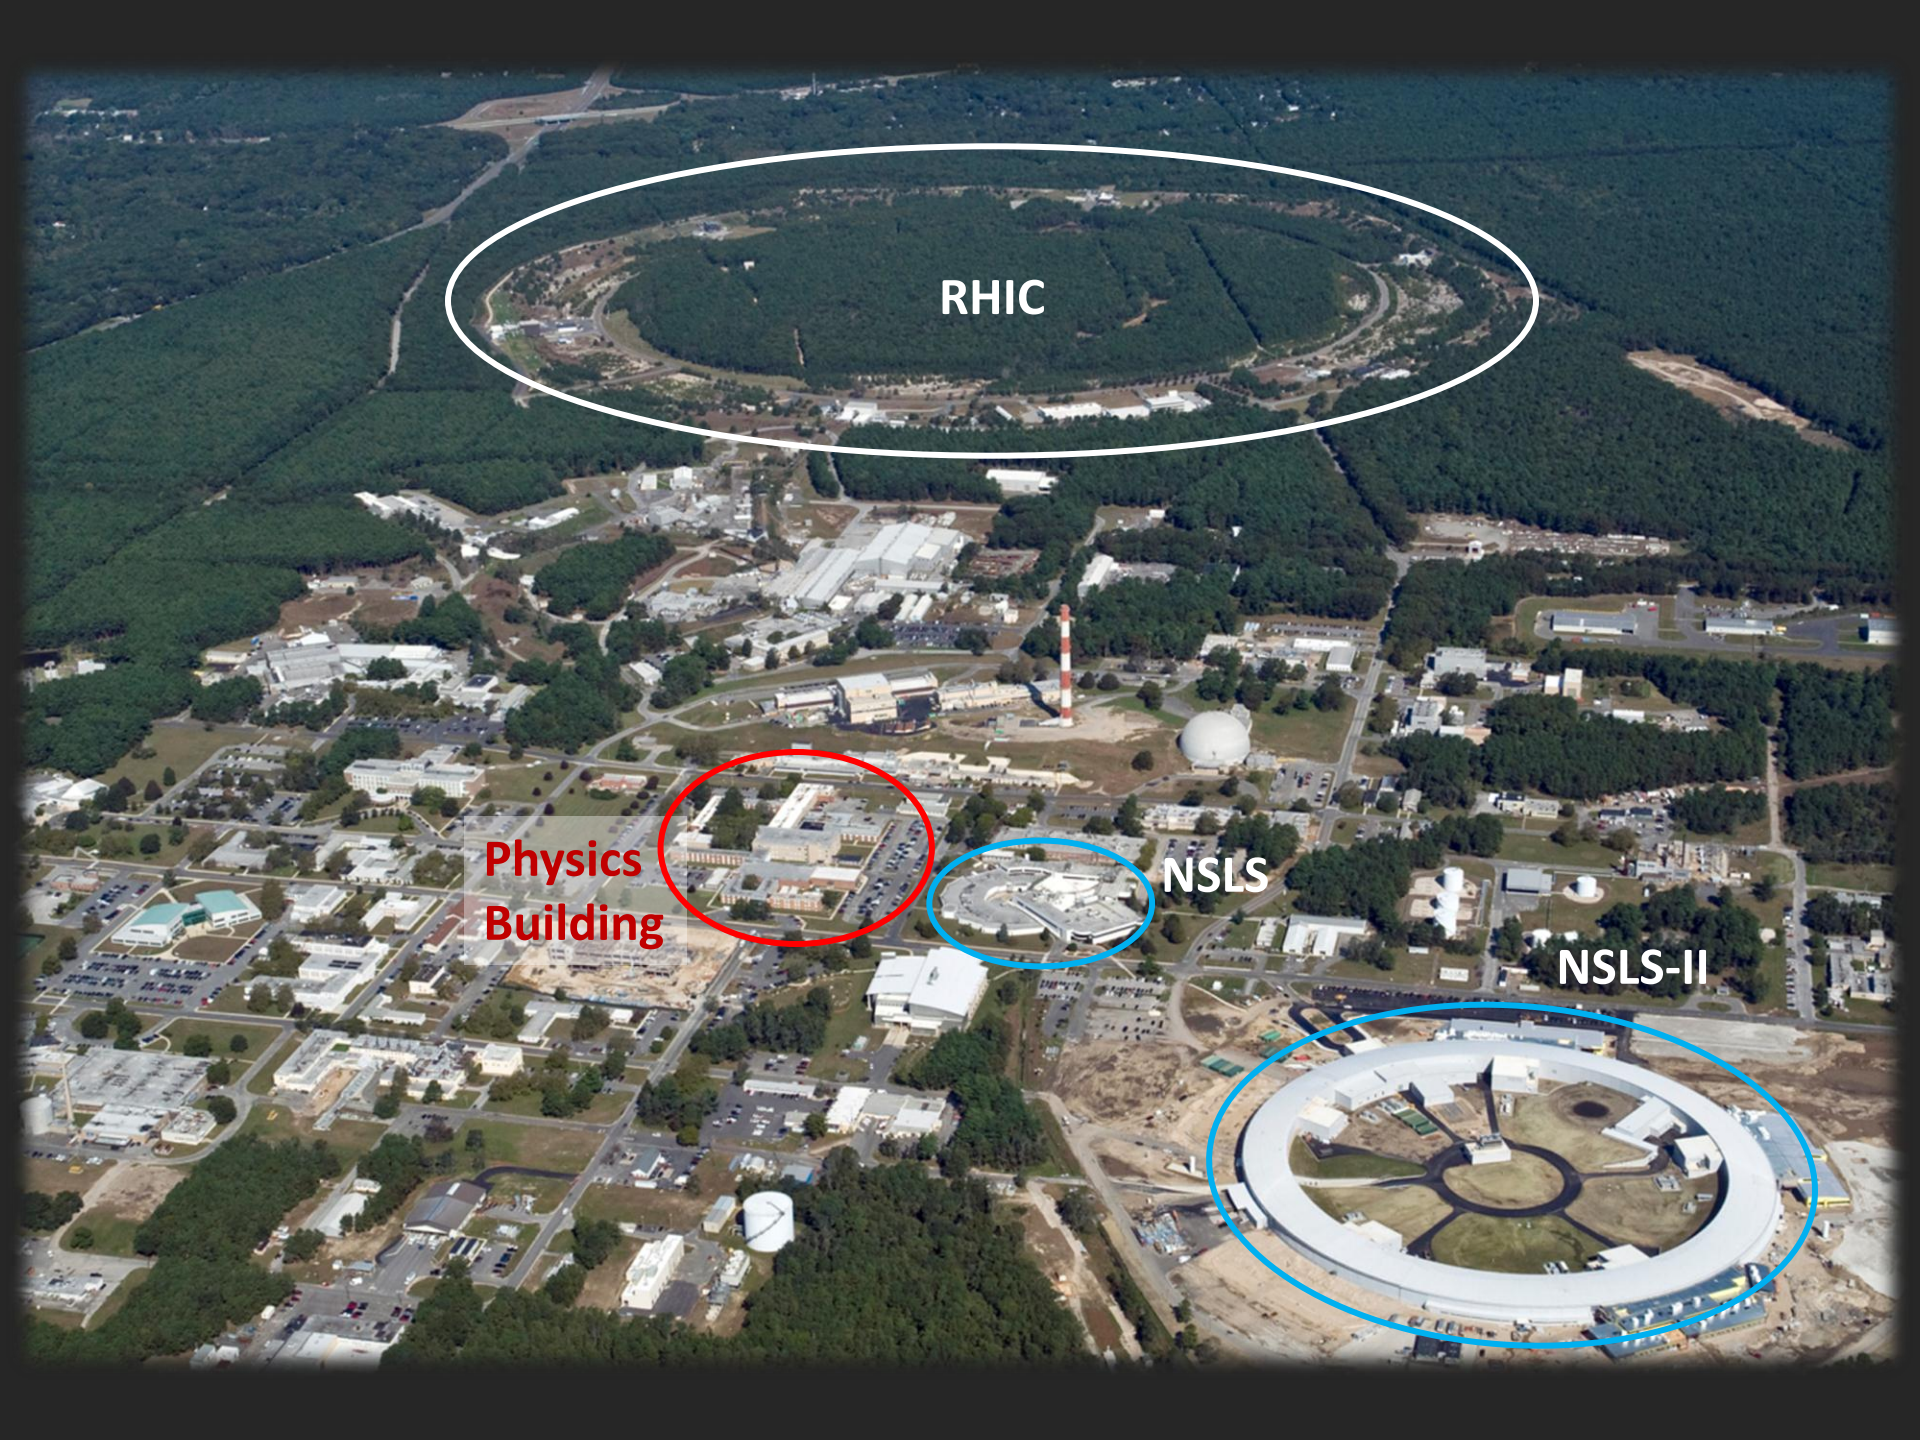
RHIC
Physics Building
NSLS
NSLS-II

RHIC & ATLAS Computing Facility (RACF)

- RACF was created in mid 1990s and currently hosts a variety of computing farms and storage systems used in production:
  - 15000 sq. ft. of combined area / 1000 max. rack capacity (with air cooling)
  - 400 racks at the moment / 1.5 MW of power consumption
  - 34k CPU cores / 72 TB RAM in Linux farms
  - 1 PB in centralised NFS storage (BlueArc/HDS)
  - 23 PB of distributed storage (dCache/XRootD)
  - 66k slots in robotic tape silos: 35 PB in HPSS storage system
- RACF data center consists of multiple isolated and semi-isolated areas that contain equipment grouped by functionality and purpose:
  - BCF (originally the Brookhaven Computing Facility area; the oldest area in the data center)
  - RCF (RHIC Computing Facility serving computing and storage needs of STAR and PHENIX experiments at RHIC collider)
  - Sigma-7 (primarily devoted to storage systems)
  - CDCE (Computing Data Center Expansion devoted to ATLAS Tier-1 site; the newest area in the data center)
  - Other areas (NetLabs, BlueGene supercomputer areas, etc.)

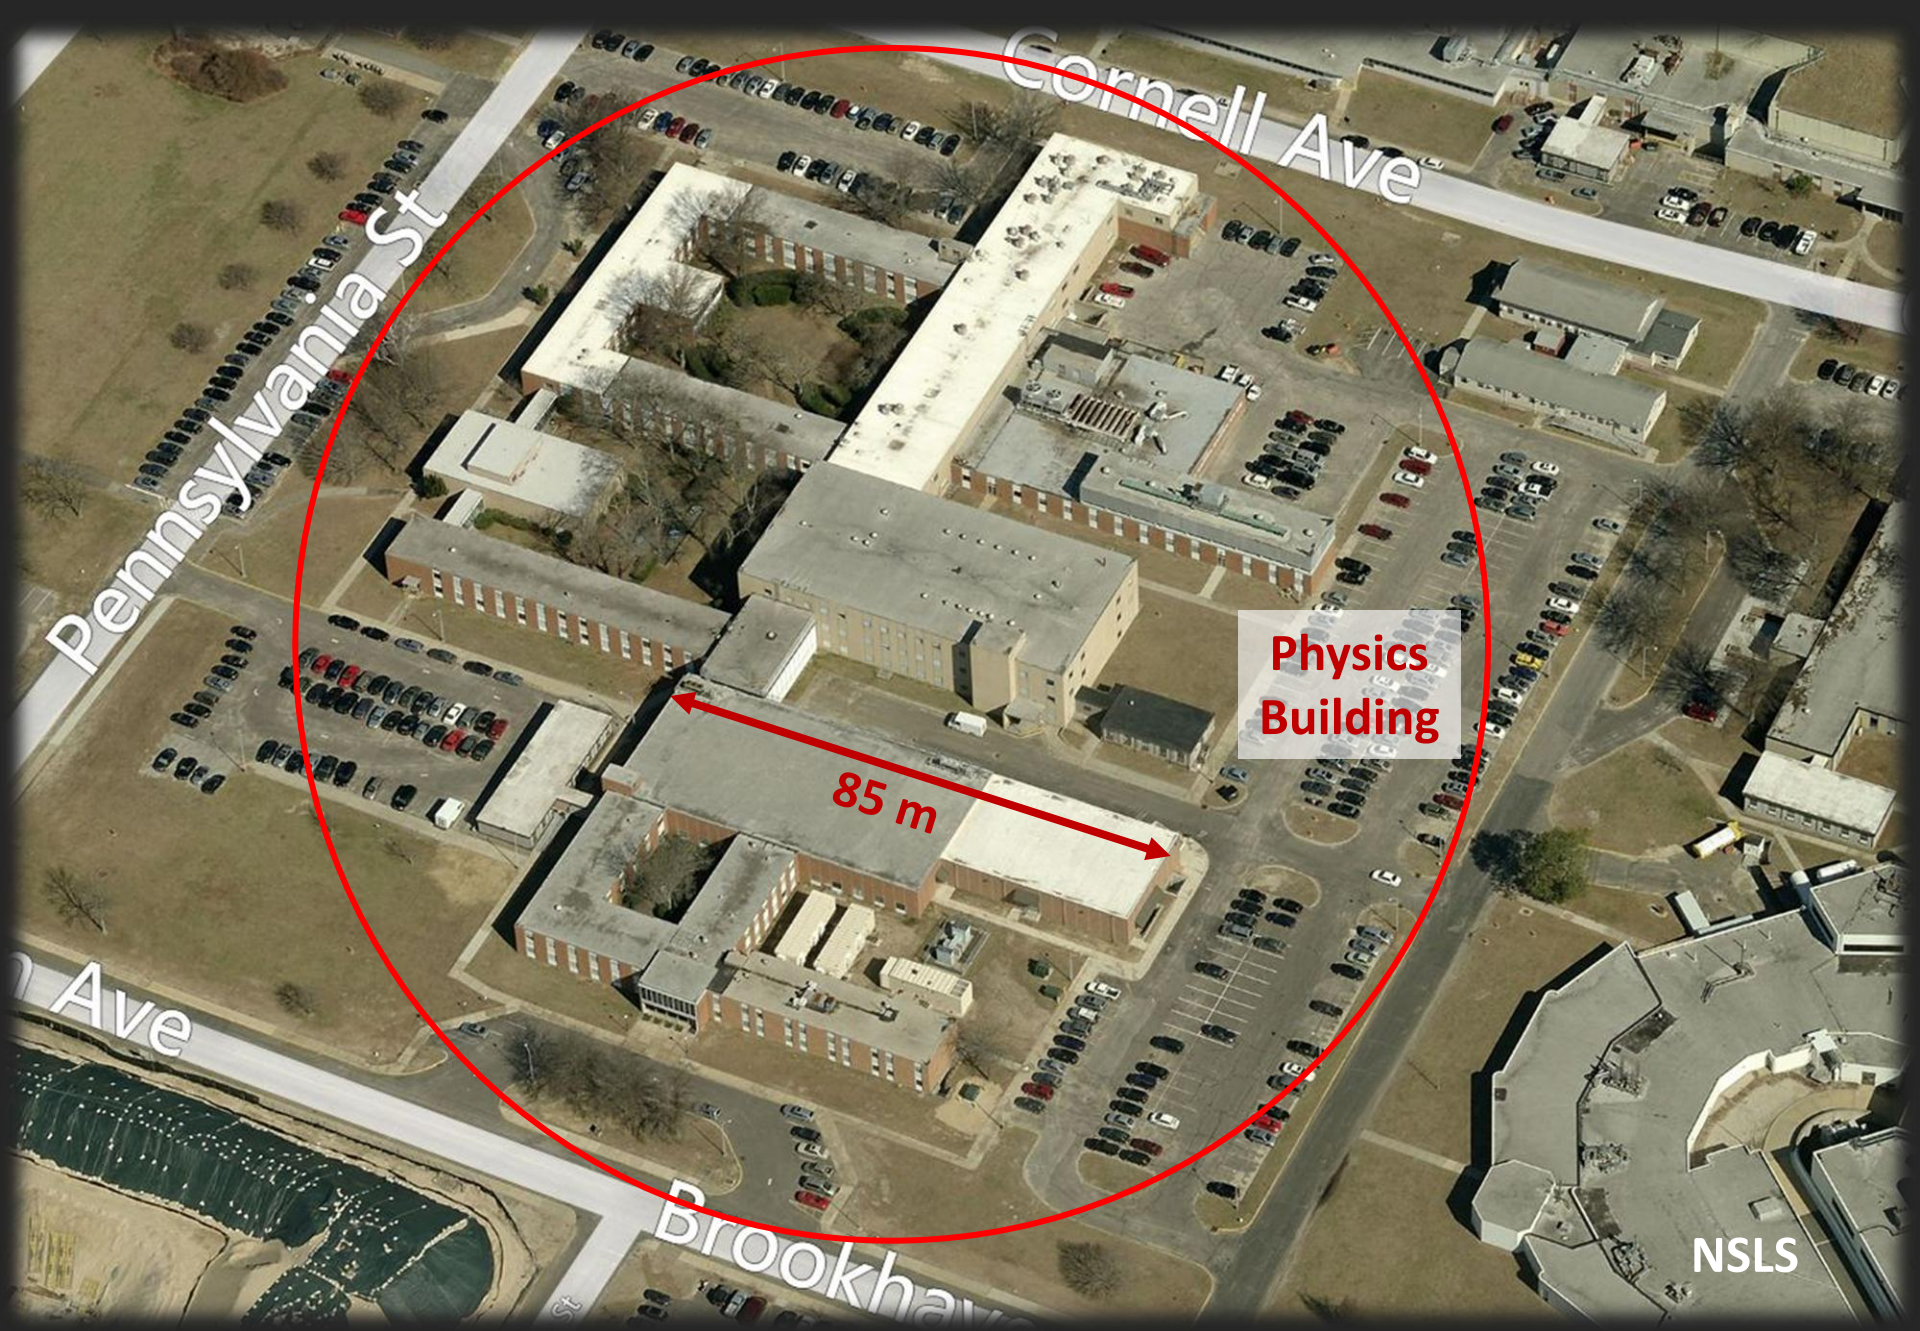
Cornell Ave
Pennsylvania St
Physics Building
85 m
NSLS

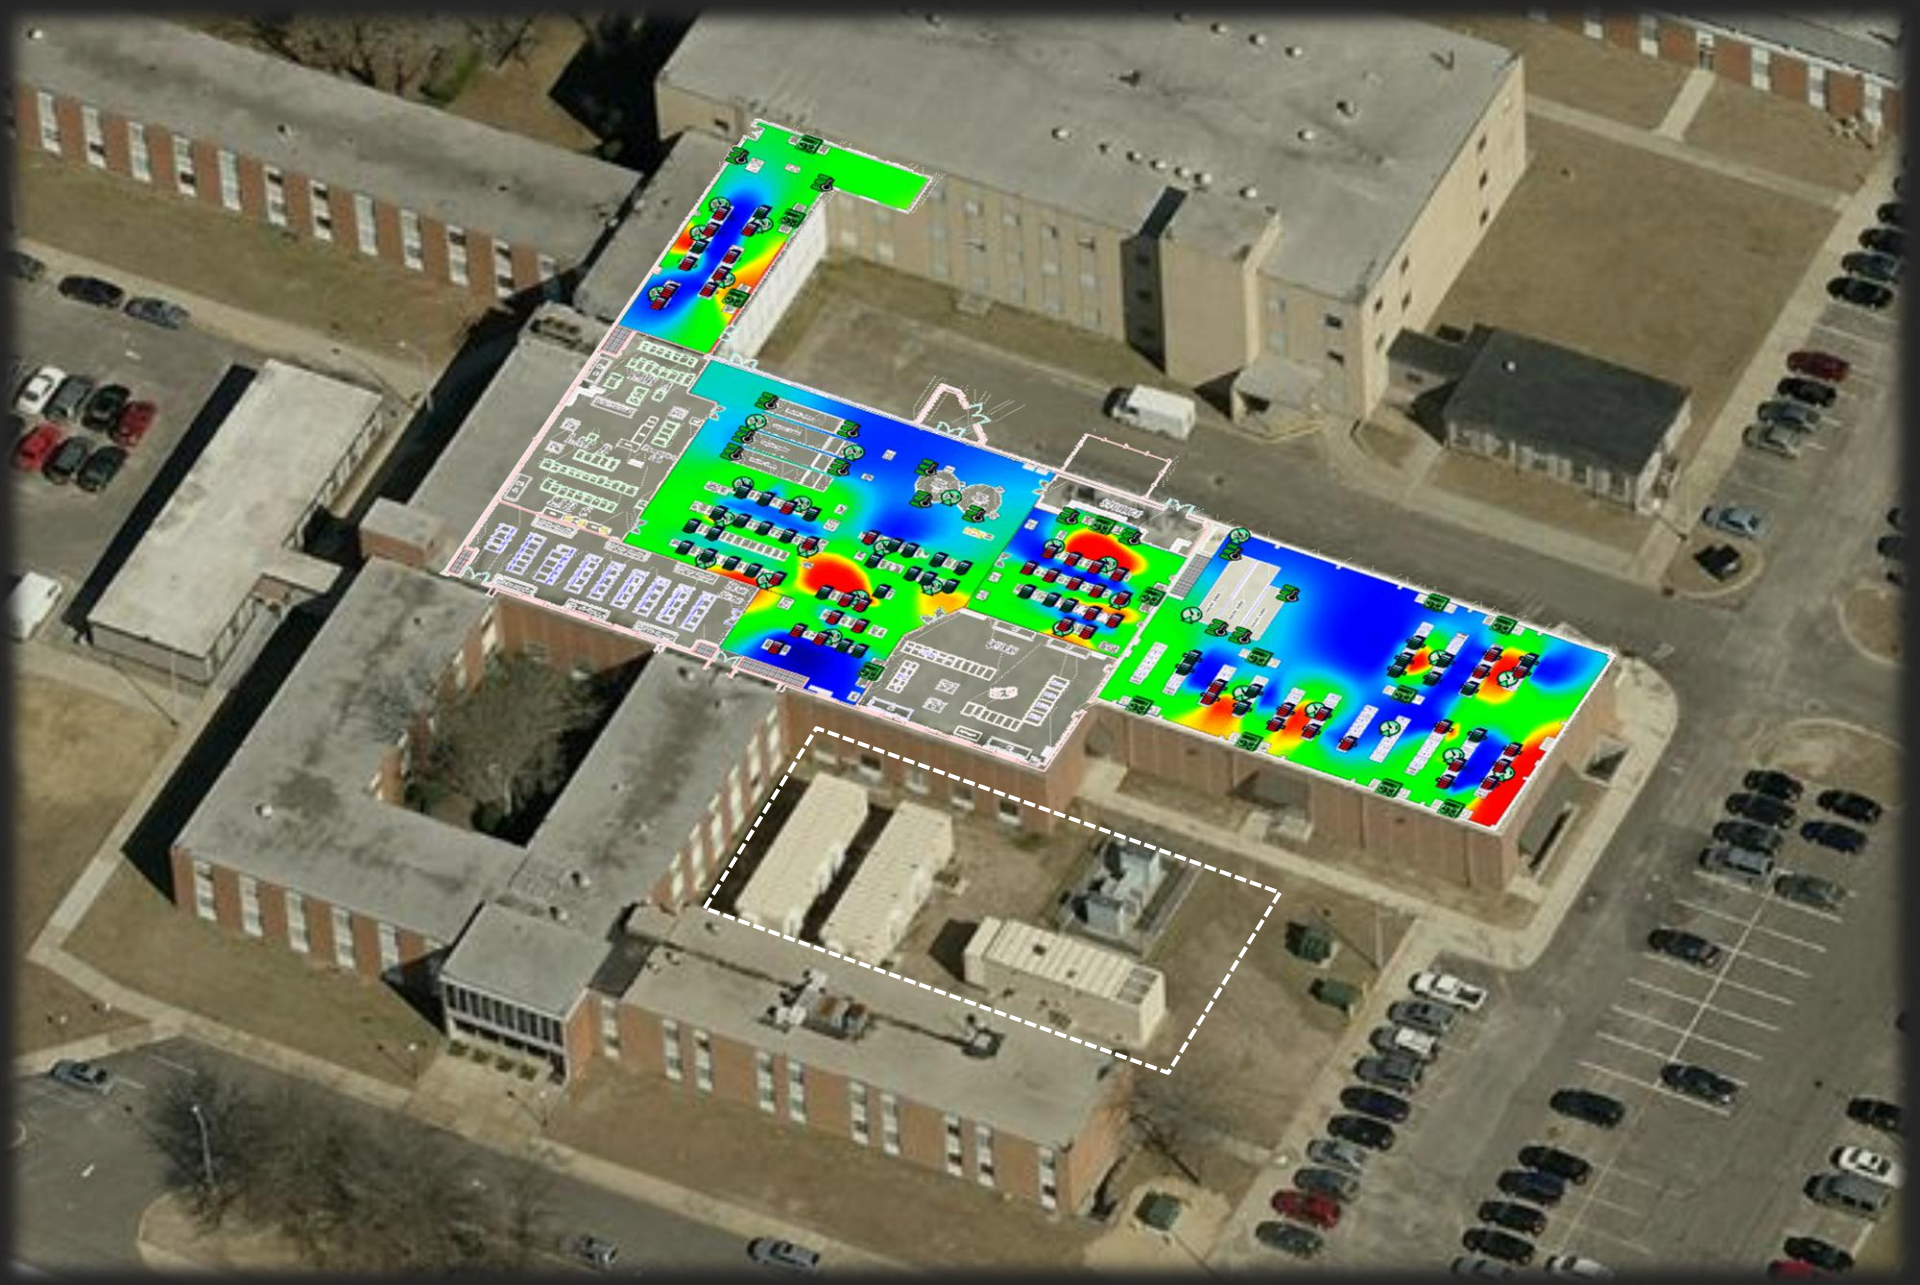

RHIC & ATLAS Computing Facility (RACF)

- All the areas of the data center are making use of the room level air cooling systems provided with humidity control supplied by:
  - CRAC units distributed among the area of every IT room (central facility-wide chilled water supply is used; under the raised floor exhaust on all the CRAC unit)
  - Central building-level AC units in the basement
  - Liebert racktop cooling units (based on pumped refrigerant technology) are used in certain areas
- Pressurized raised floor layout is used for the cooled air distribution in all the areas:
  - Maintaining a uniform pressure differential on the entire area of the facility is important for ensuring the high cooling efficiency
  - Cooling solutions based on row and individual rack isolated may be used in the future in high density areas, thus the ability to represent the compartmentalization layout of the facility is important

RHIC & ATLAS Computing Facility (RACF)

- Impact of high availability / infrastructure resilience features of the RACF facility on the monitoring system requirements:
  - Certain areas with relative small air volume and/or large portion of IT equipment on UPS are particularly sensitive to cooling system failures: thus fast and fully automated mechanism for IT load reduction is needed
  - Local overheating of the racks containing storage equipment caused by a racktop cooling unit failure or temporary raised floor pressure differential misbalance (for instance, caused by removal of a section of the raised floor in the vicinity) may trigger disk failures in the RAID systems, thus reducing their performance for the period of RAID group recovery: thus the high spatial resolution of the monitoring system is required, monitoring ambient temperature in the room is not good enough
- Important requirements for the monitoring system:
  - Ability to analyse historical data is needed for retrospective analysis of maintenance / disaster recovery operations with respect to their impact on environmental conditions in the data center
  - Ease of reconfiguration / reinstallation of the sensors – RACF is a rather dynamic environment where rack movements and row reconfigurations do happen frequently

RACF Environmental Monitoring Systems

Historical Perspective

- The following systems were in use within RACF computing facilities before 2012:
  - Ambient temperature / humidity sensors with local displays deployed in every room (no central data gathering, no alarm system)
  - Embedded CRAC unit sensors with local alarms (no central information gathering implemented so far)
  - DigiTemp ambient rack temperature sensors (provided with Ethernet based connectivity and central data gathering mechanism, but no high spatial resolution data)
  - IPMI based temperature information: (very) high spatial granularity, but poor absolute calibration / precision, and reduced sensitivity to the room-level environmental changes (IPMI data is sufficient for implementing host level automatic emergency shutdown mechanisms, but those are normally used only as a last resort)

RACF Environmental Monitoring Systems

Considerations for the New Unified System Design

- The following architectural possibilities were available for the new unified & centralized environmental monitoring system:
  - Extension of the existing sensor groups provided with the Ethernet connectivity (factor of 10 increase of sensor count is needed): on the scale of RACF that would require deploying too much of the dedicated cable infrastructure
  - Deploying new groups of sensors provided with wired loop topology interlinks and central controllers (e.g. APC / NetBotz solutions): require a dedicated cable infrastructure intersecting multiple racks limiting the flexibility of rack movement in the data center
  - Fiber optic based temperature monitoring solutions (Raman / Brillouin distributed fiber based thermometers): require a dedicated fiber optics infrastructure intersecting multiple racks, requires precision placement of the fiber in the racks, limited number of junctions allowed on a fiber line due to signal degradation, systems with high spatial resolution (better than 30 cm) are expensive
  - Wireless sensor network based solutions such as one provided by SynapSense Corporation: found to be the most price / performance efficient solution for our environment (the same decision is made by several other DOE and DOD labs)

SynapSense Monitoring System

Building Blocks

Ethernet Gateway

(Ext. AC PSU)

Approx. 50m range

Max 0.25k sense points per unit

Pressure Differential Sensor Base Station

(Local AA Batteries)

Rack / CRAC Unit LiveImaging™ Unit Sensor Base Station

(Local AA Batteries)

Max. 10 sense points per unit

10 years (max.) battery life

SynapSense Monitoring System

Building Blocks

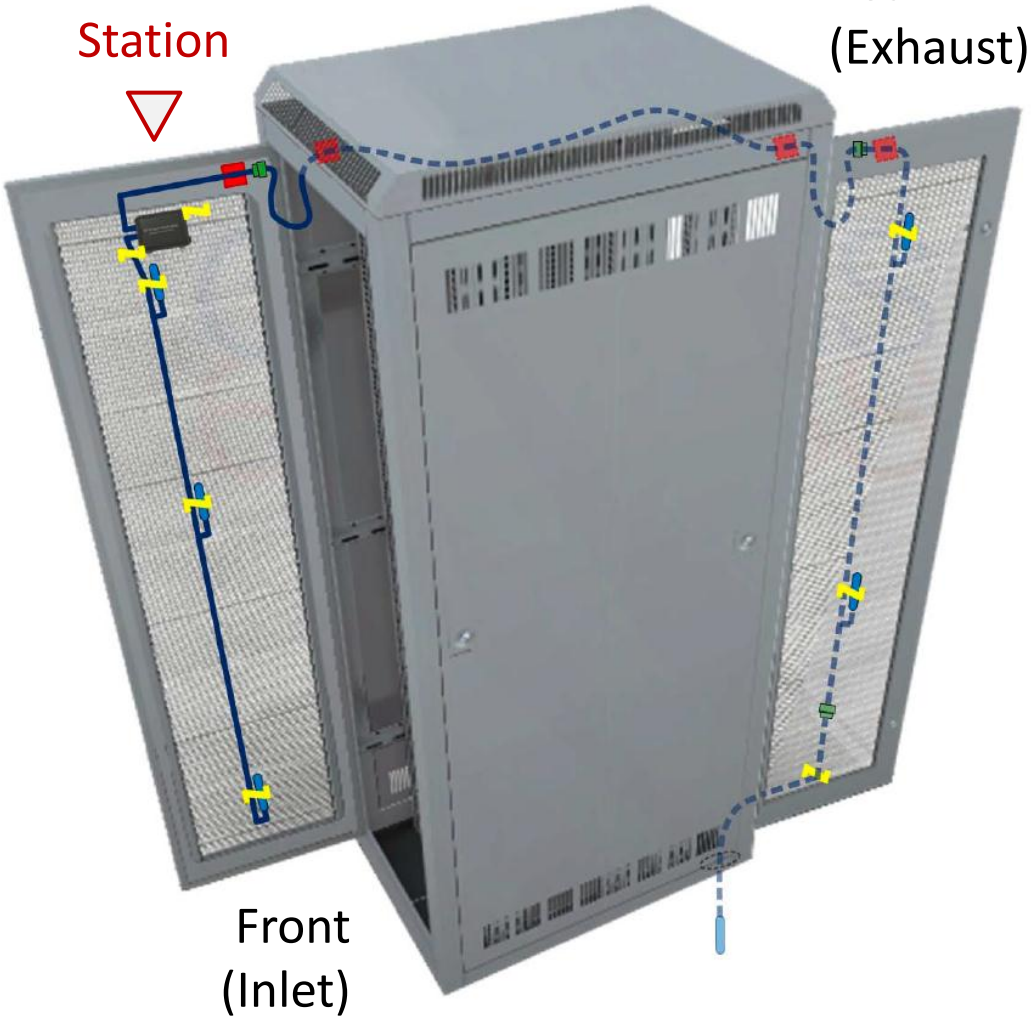

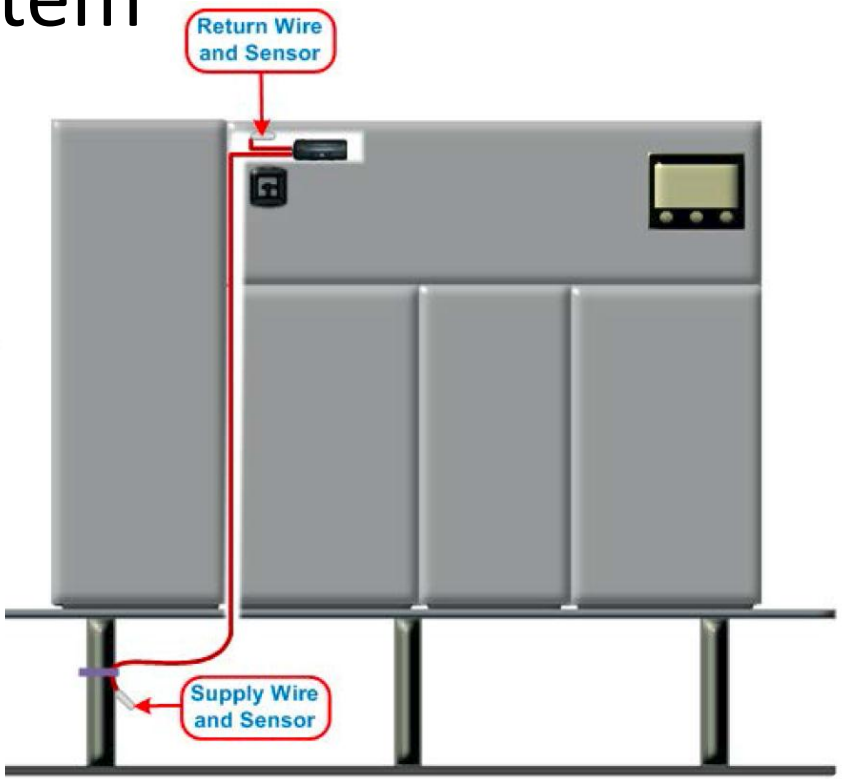

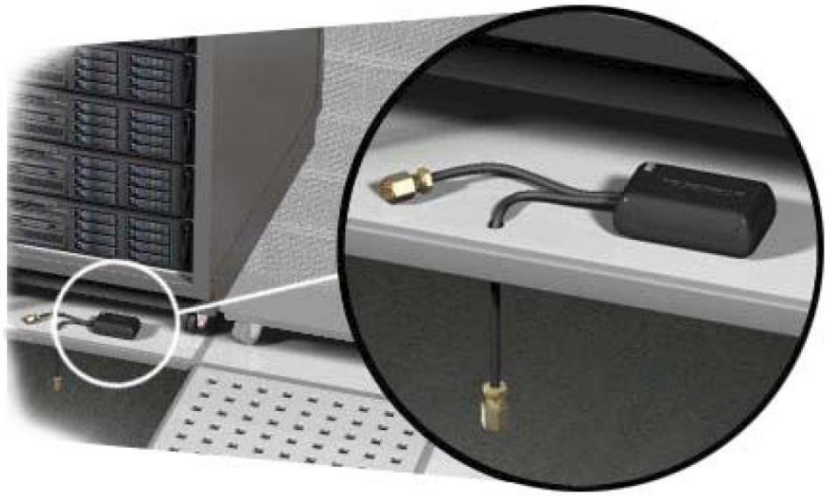

SynapSense Monitoring System

RACF Data Gathering Layout

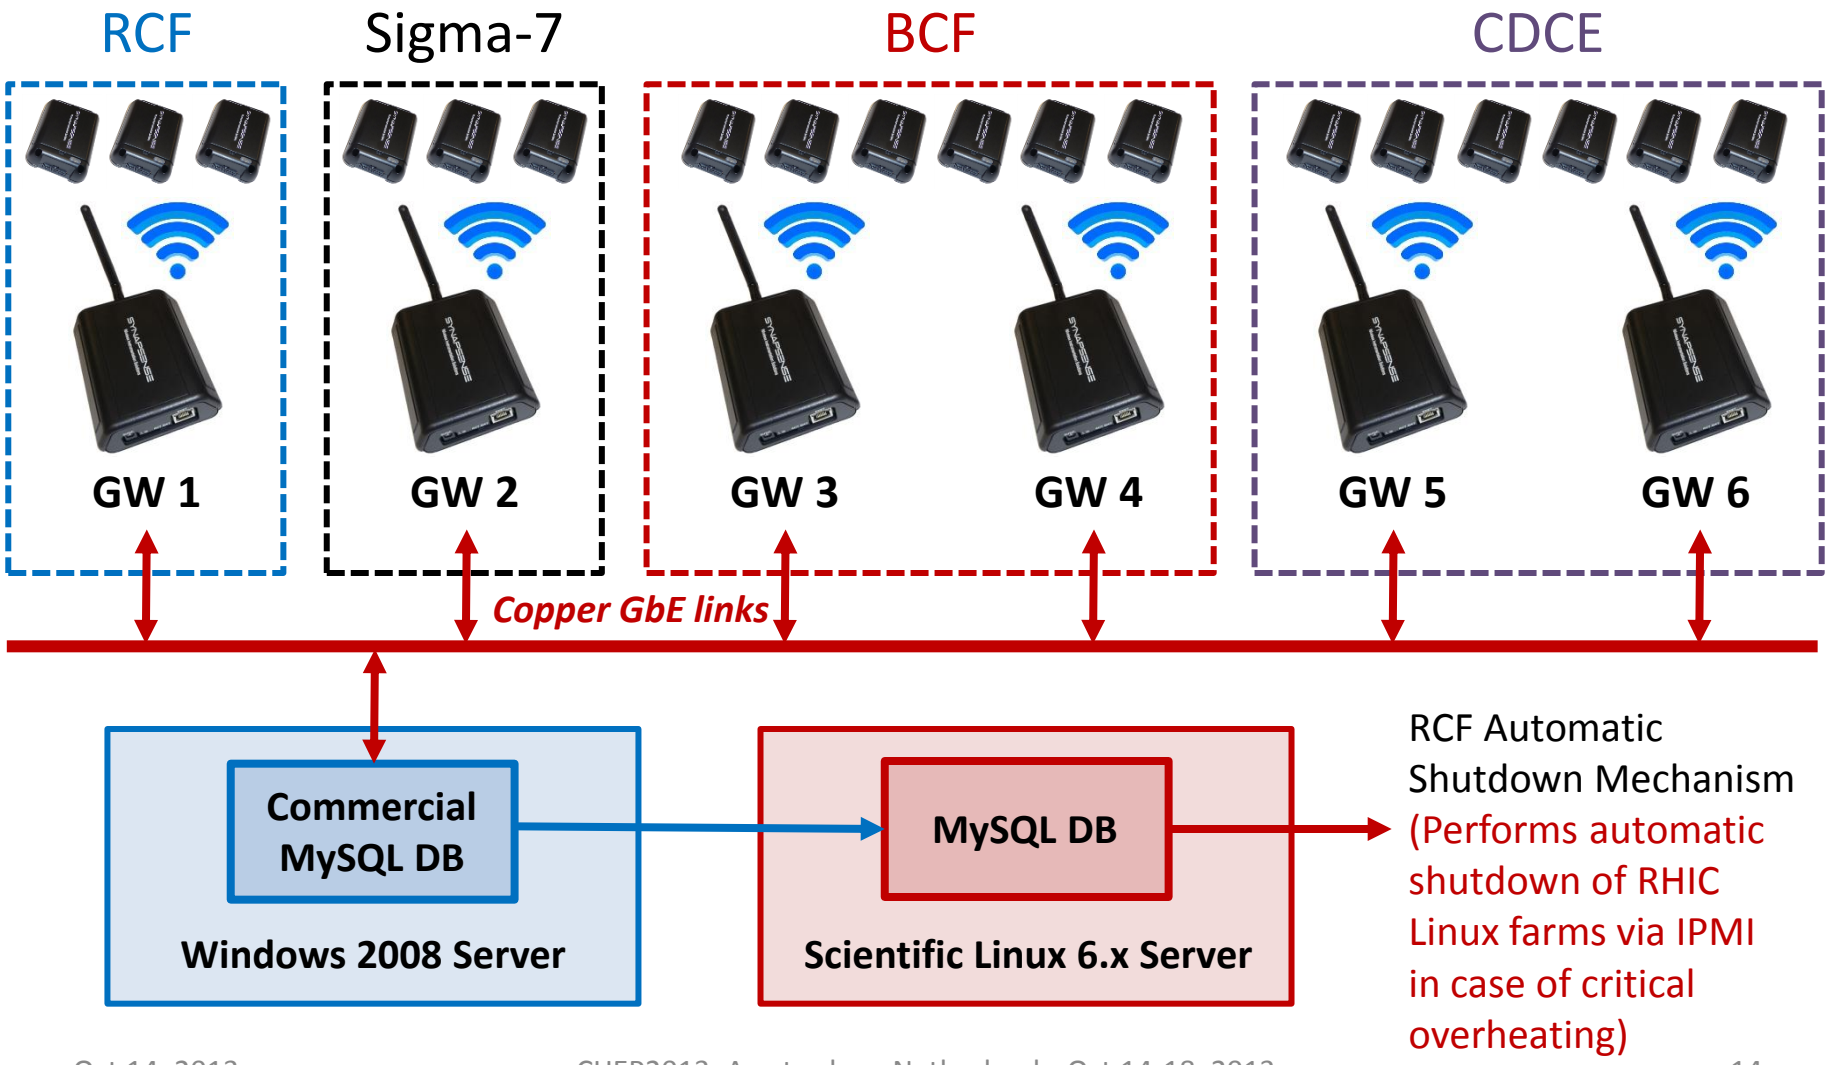

SynapSense Monitoring Systems in RACF

Timeline for Deployment

- June 2012: RCF area
  - Central server infrastructure is built
  - First SynapSense sensor group deployed in RACF
  - Deployment performed by SynapSense engineers
- Dec 2012: Sigma-7 area
  - Due to the relatively small scale of Sigma-7 area deployment performed by RACF staff
- Aug 2013: BCF + CDCE area
  - Finalizing the coverage of all main areas of the data center
  - Deployment performed by group of SynapSense engineers
  - Bringing the total count of base stations up to 152 and the total number of sensors ("sense points") – up to 926

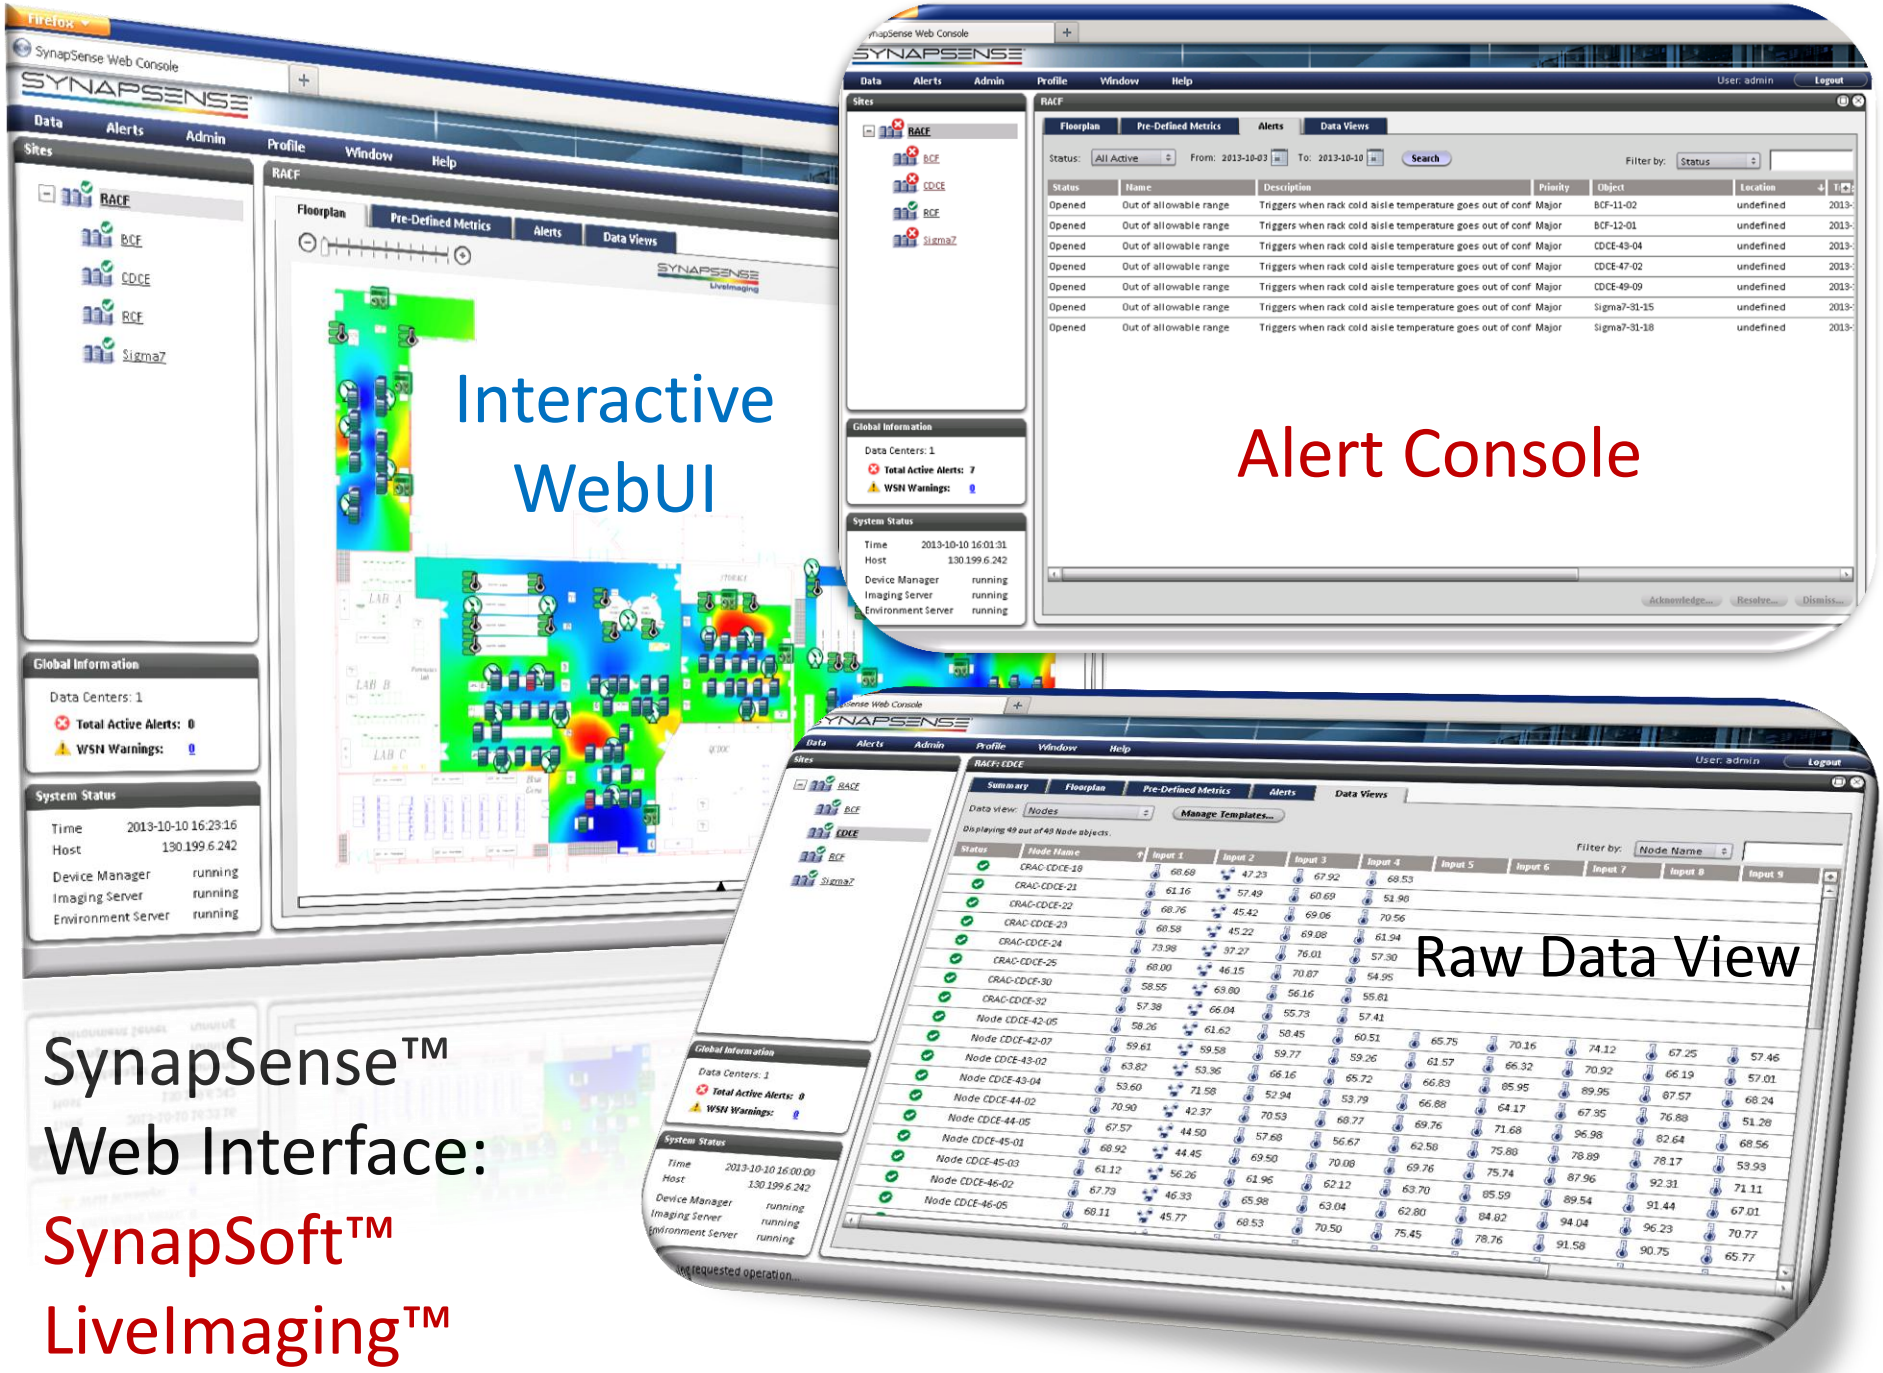

SynapSense™ Web Interface:

SynapSoft™ LiveImaging™

SynapSense™ Configuration GUI: MapSense™

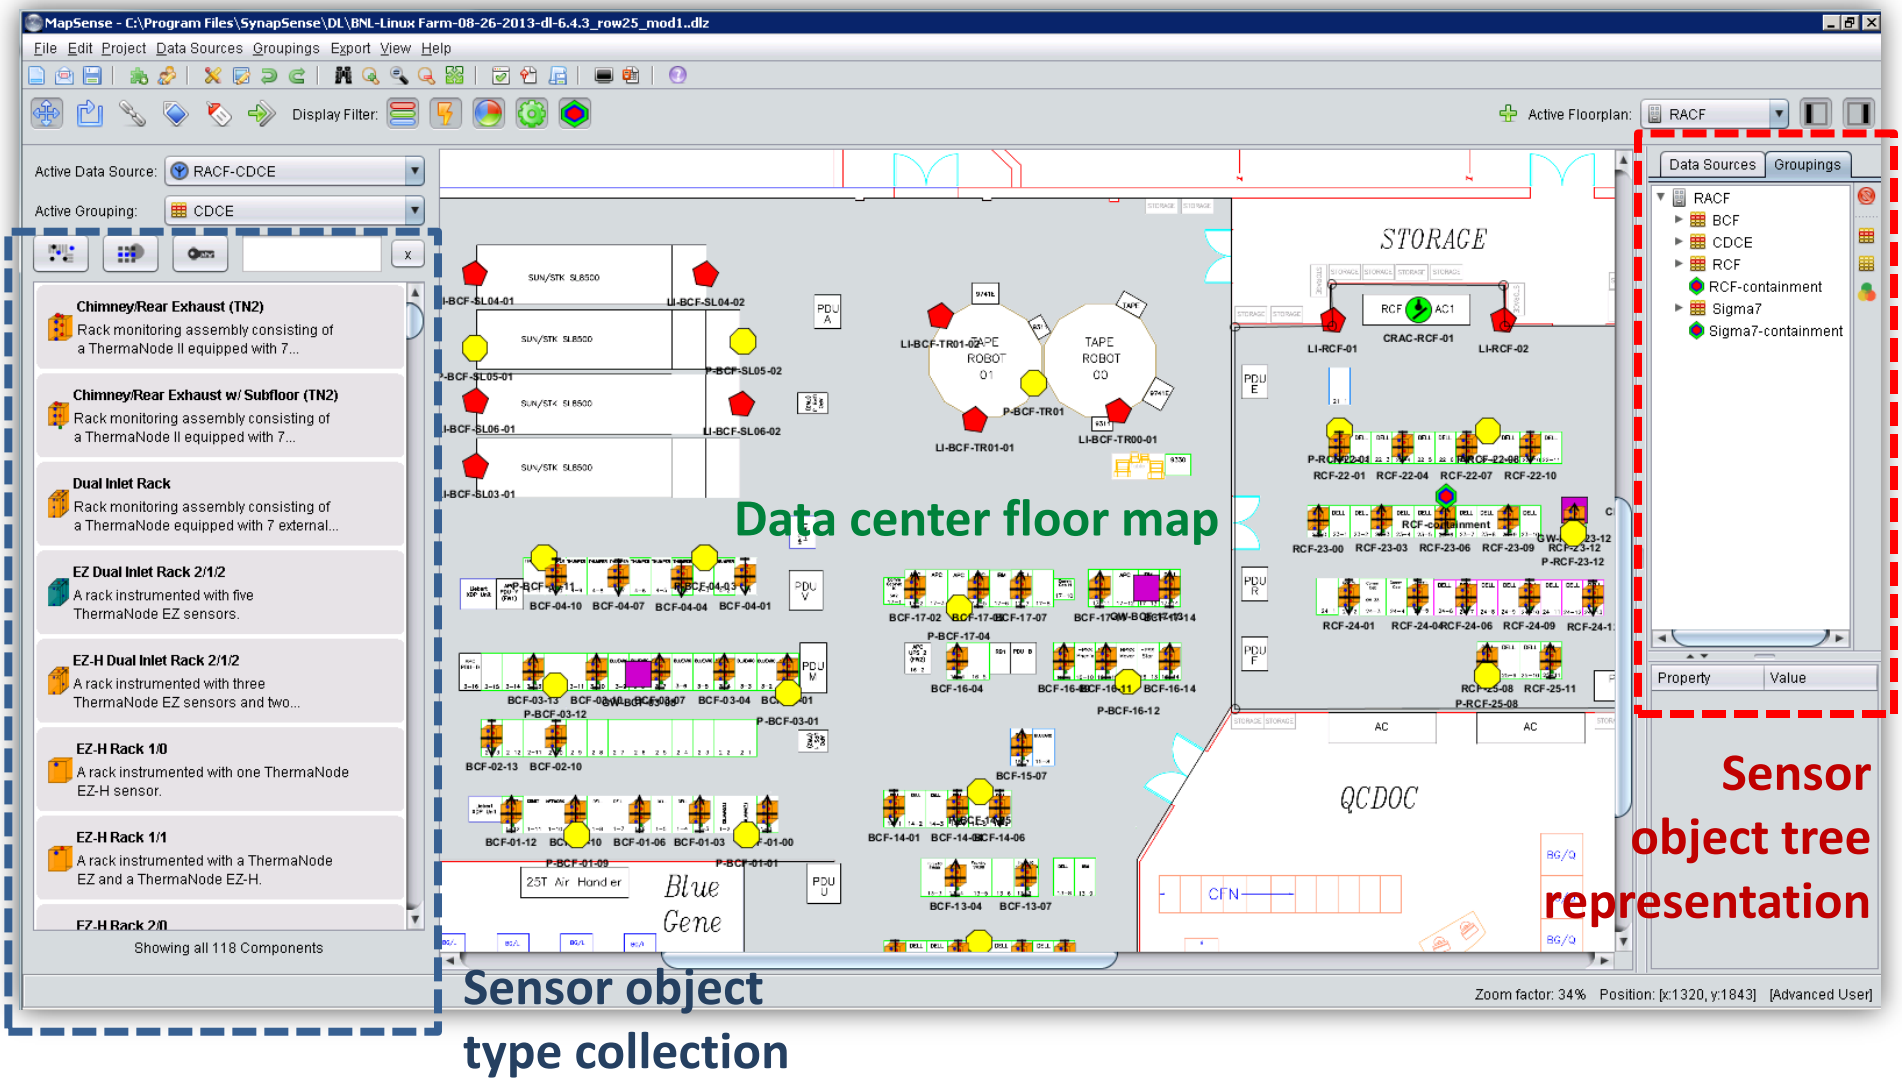

RACF: Top of the Rack Temperatures, F

Oct 10, 2013 / 4pm EDT

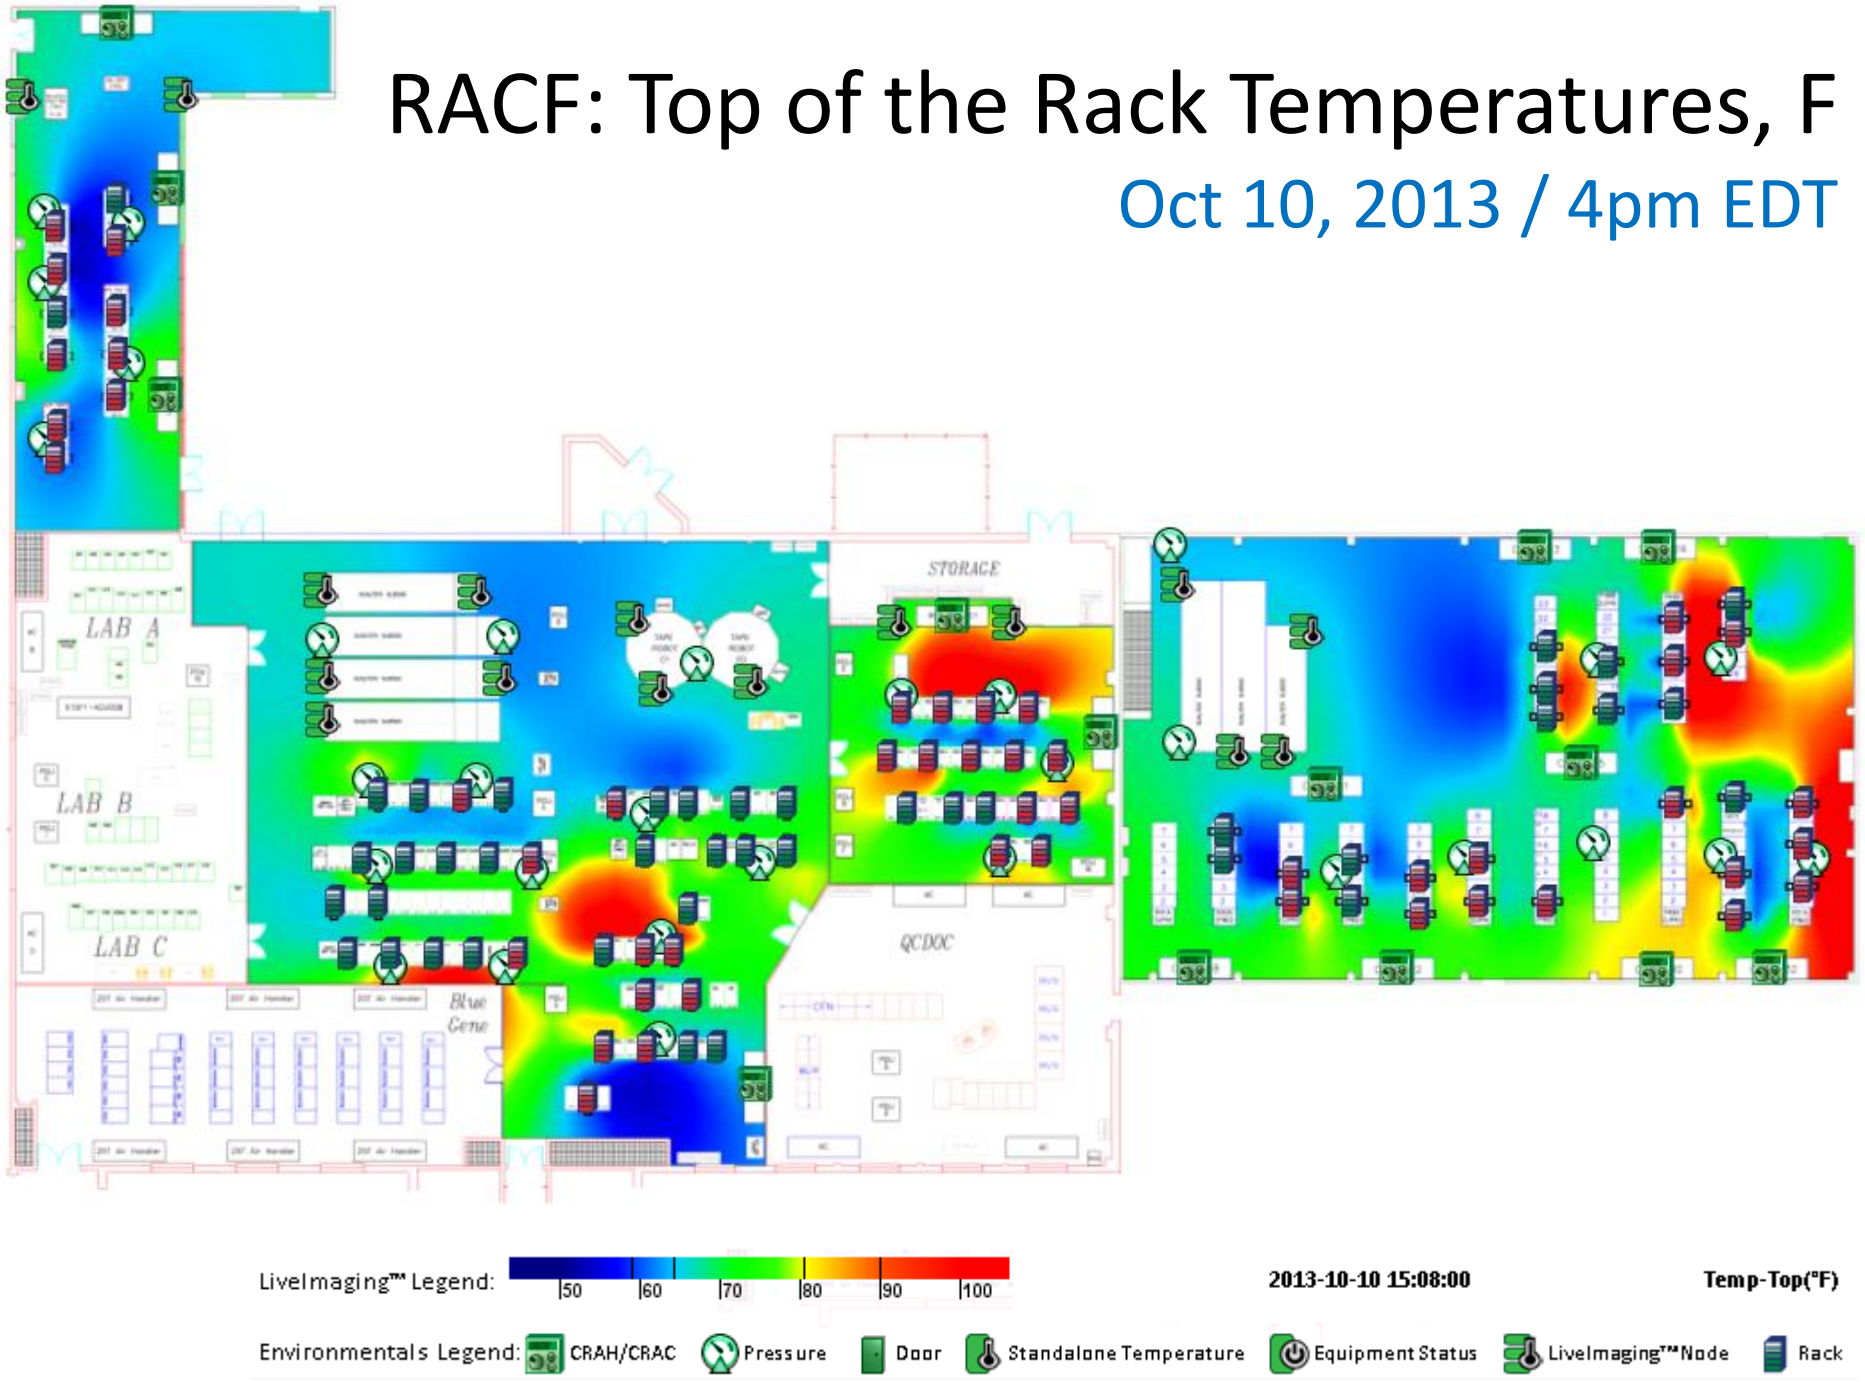

RACF: Top of the Rack Temperatures, F

Oct 10, 2013 / 4pm EDT

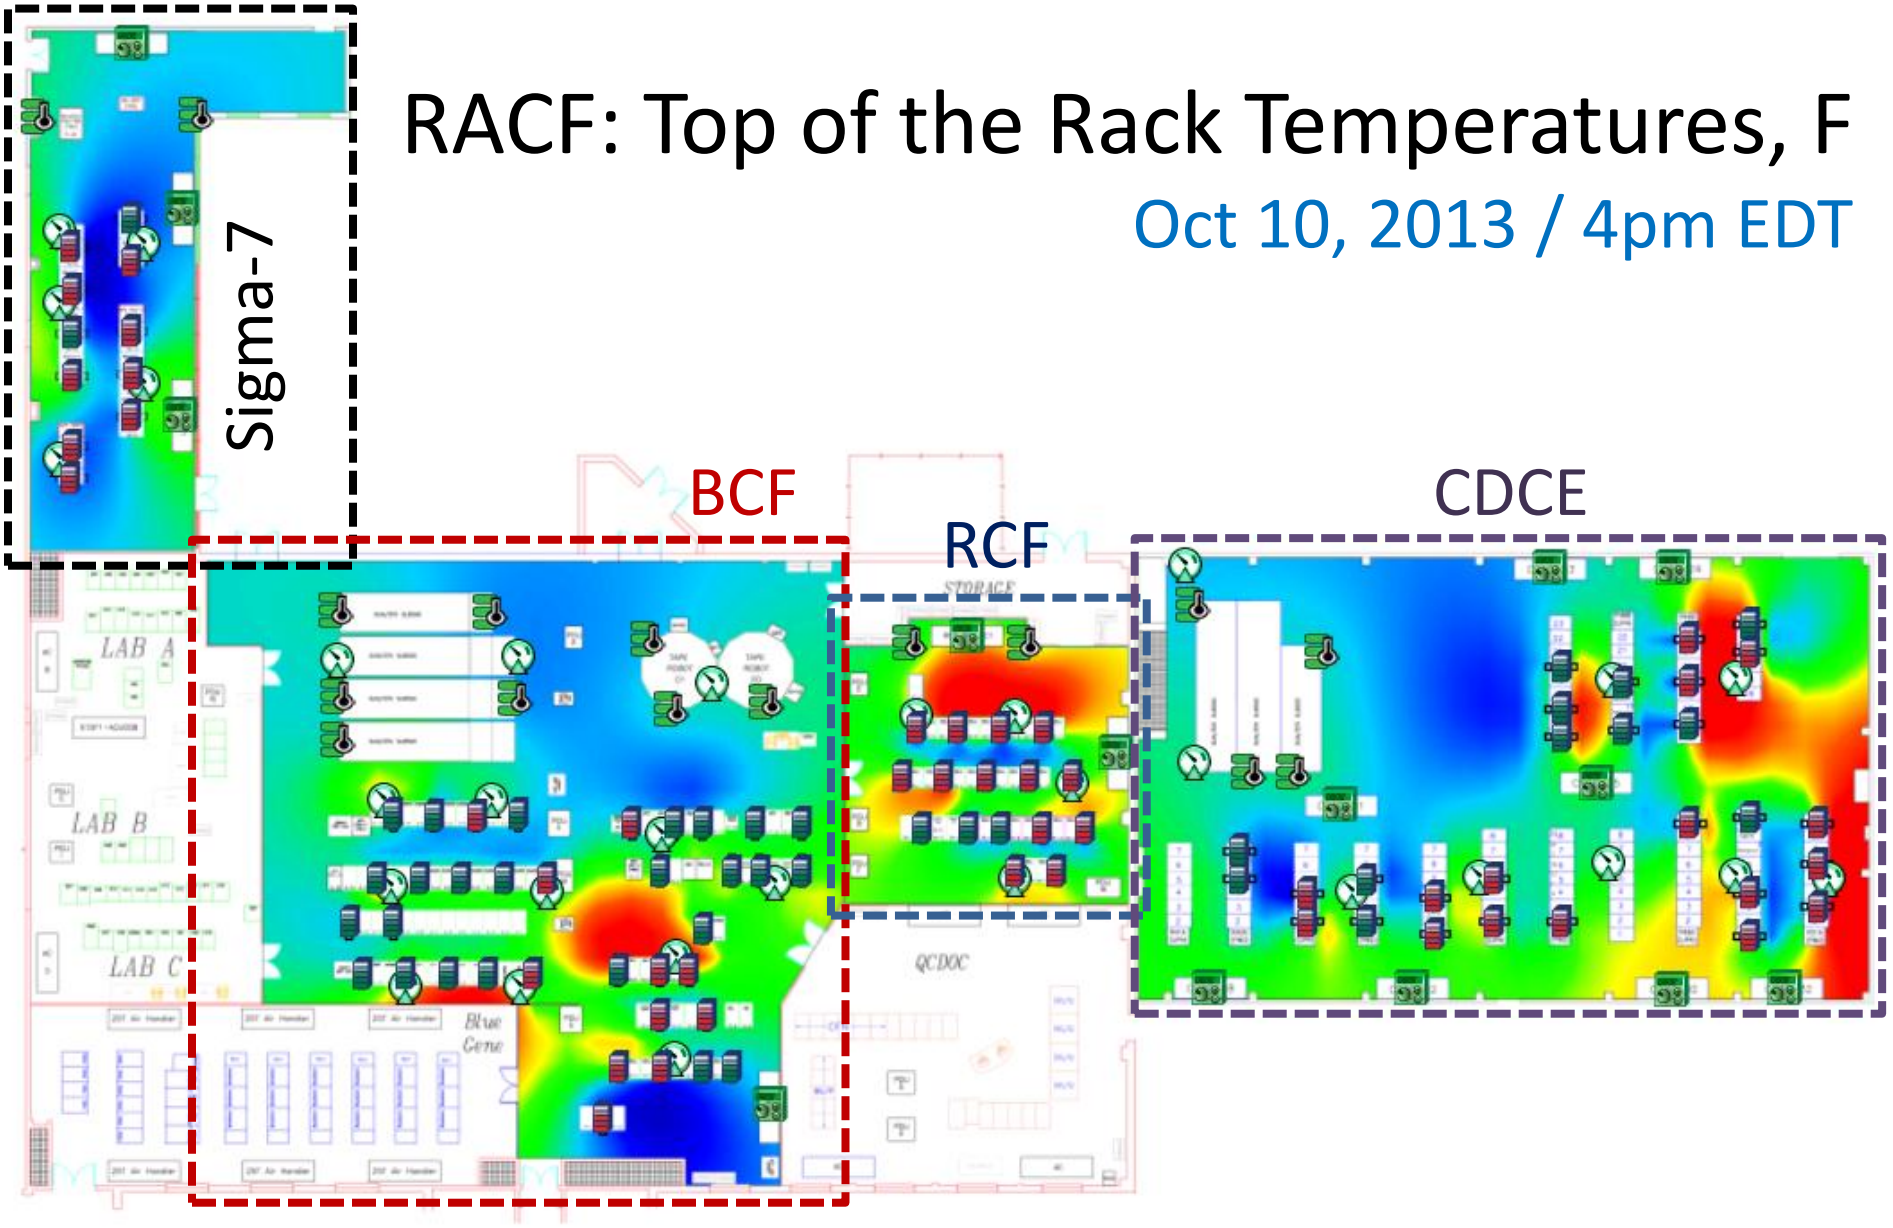

RACF: Subfloor Temperatures, F

Oct 10, 2013 / 4pm EDT

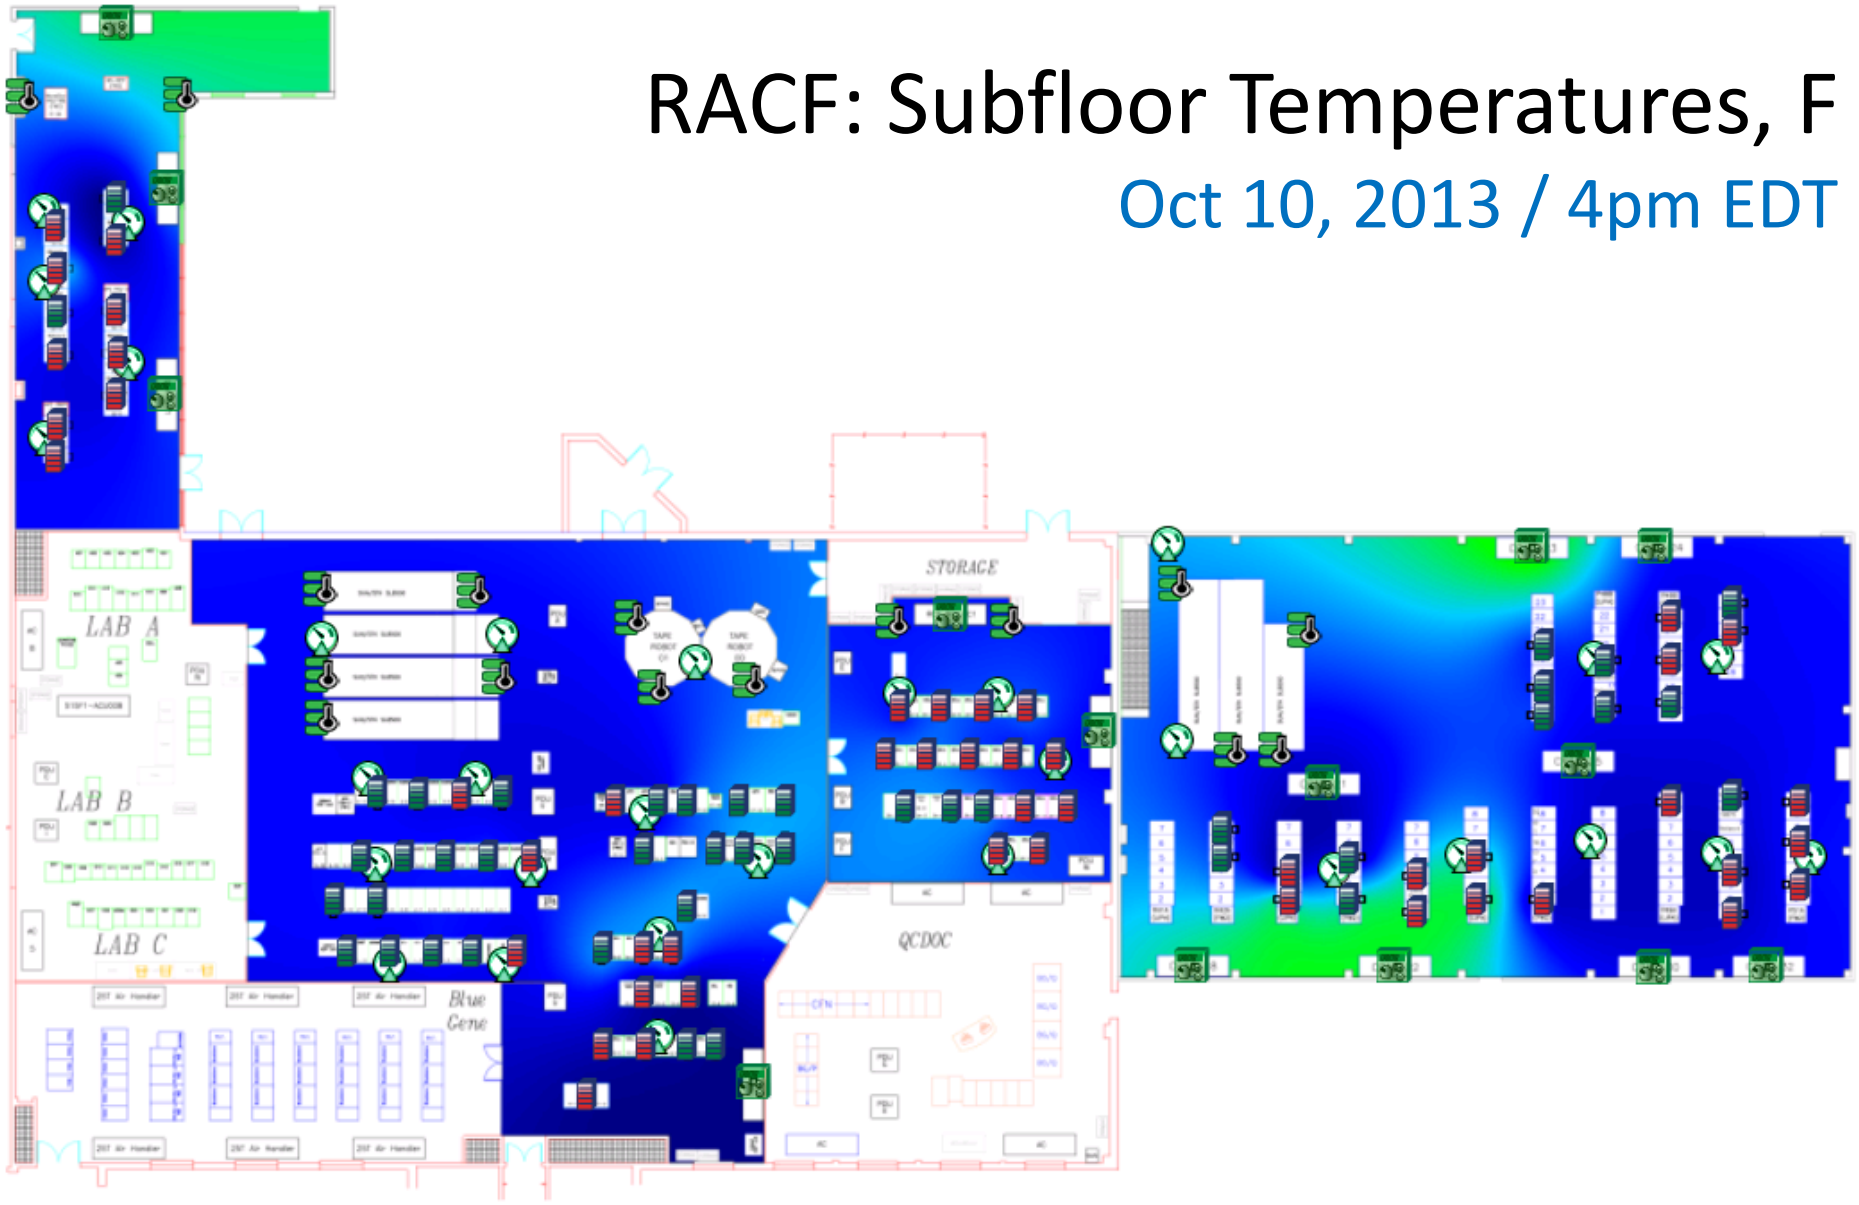

RACF: Subfloor Pressure, inH_20

Oct 10, 2013 / 4pm EDT

inH_20 = 248.8 Pa (60 °F = 15.6 °C)

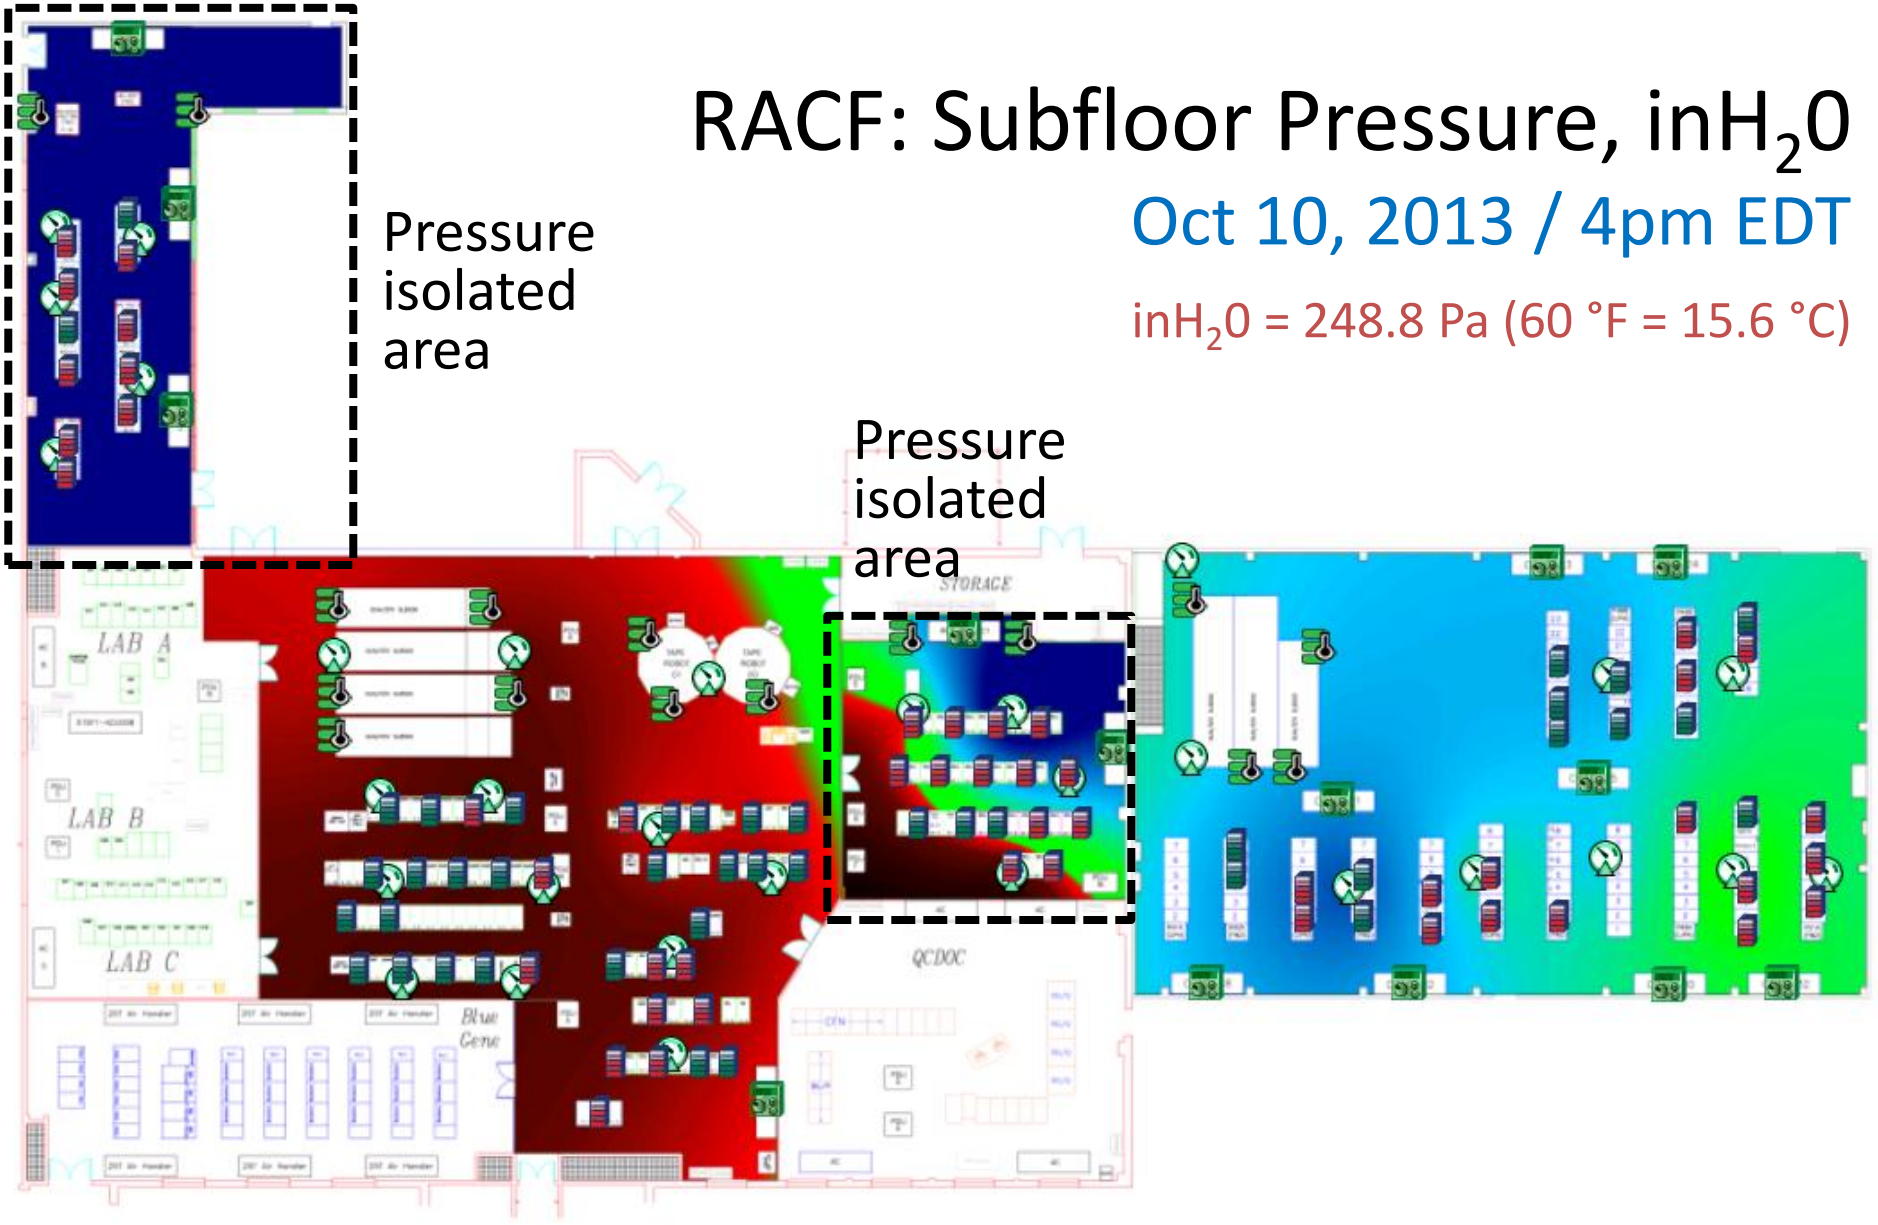

RACF: Subfloor Pressure, inH_20

Oct 10, 2013 / 4pm EDT

inH_20 = 248.8 Pa (60 °F = 15.6 °C)

Temporary pressure imbalance due new CRAC unit installation

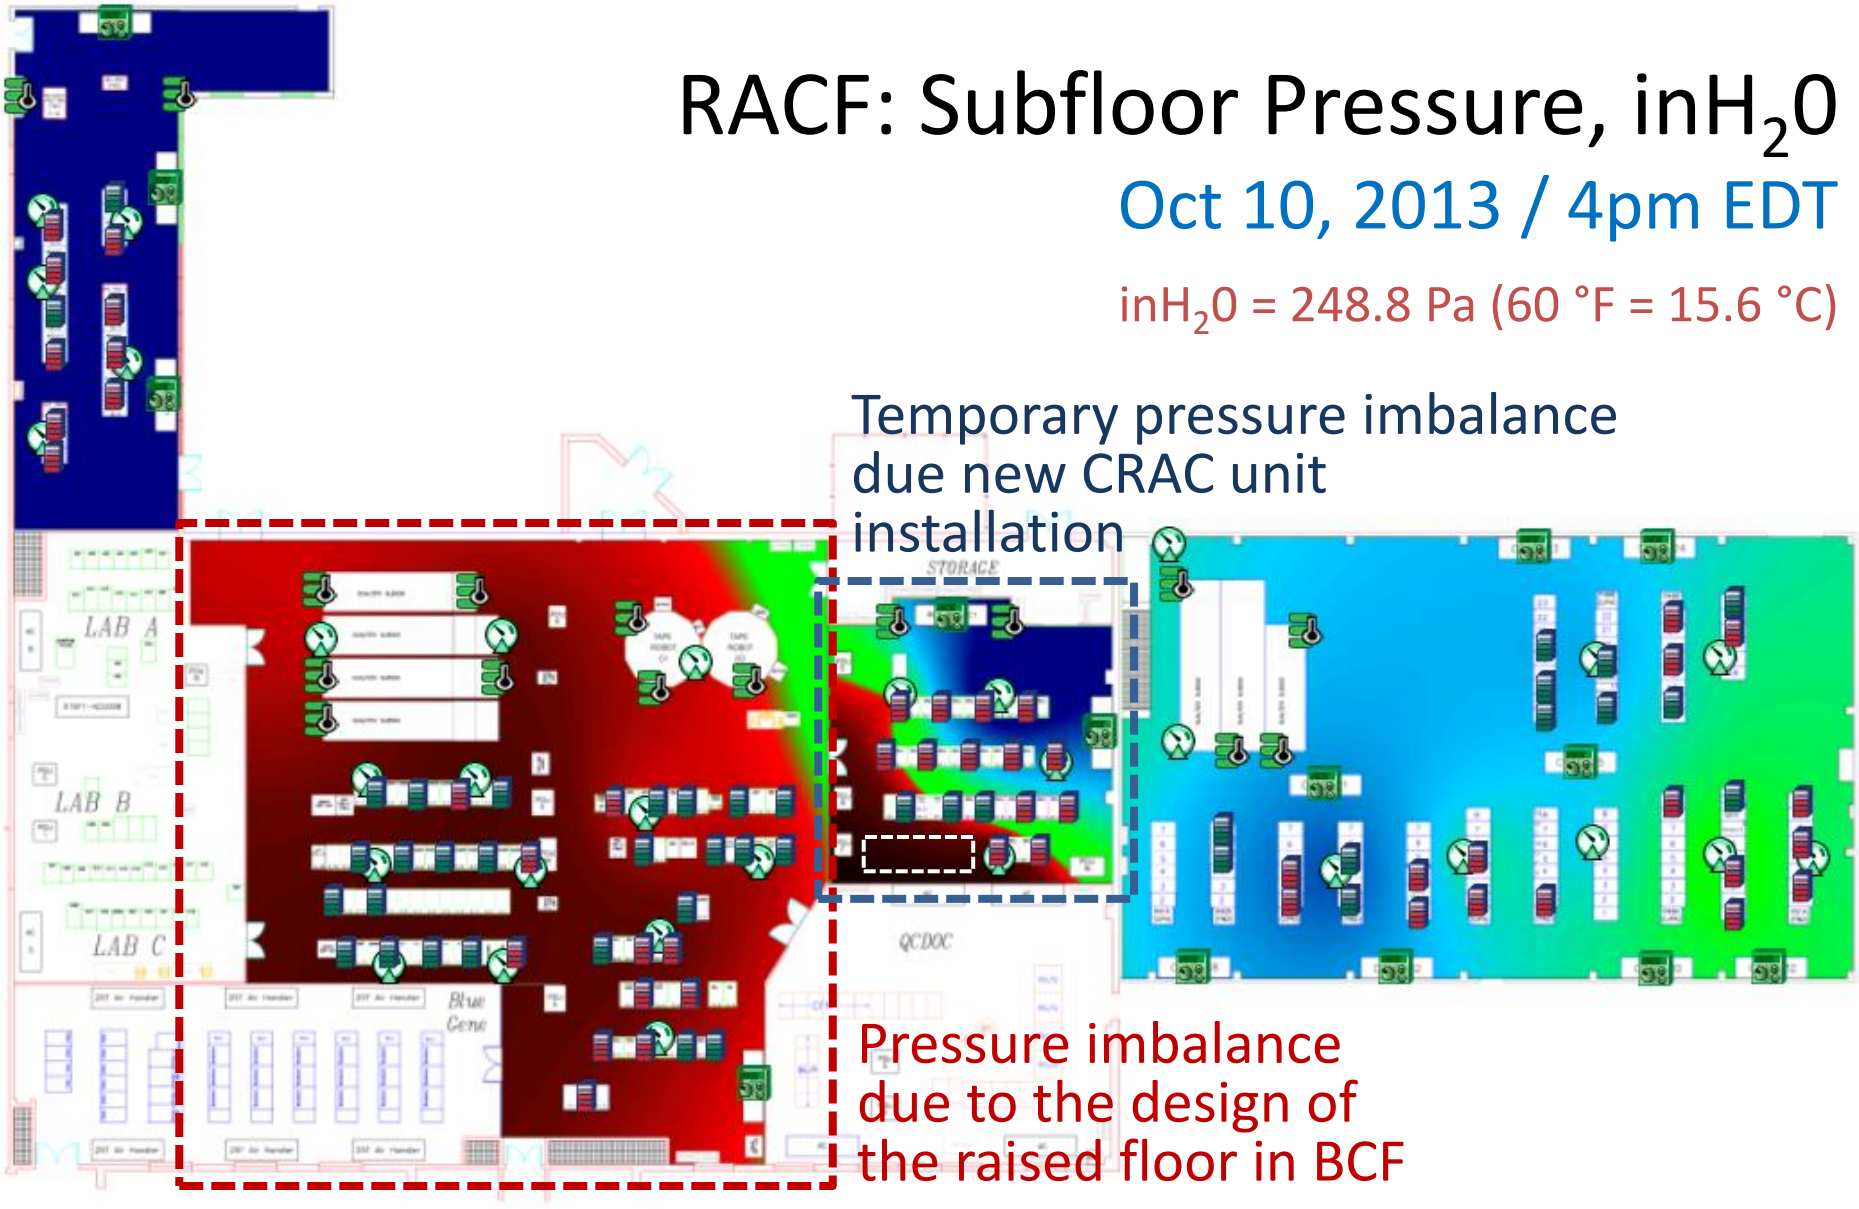

Pressure imbalance due to the design of the raised floor in BCF

RACF: Relative Humidity, %

Oct 10, 2013 / 4pm EDT

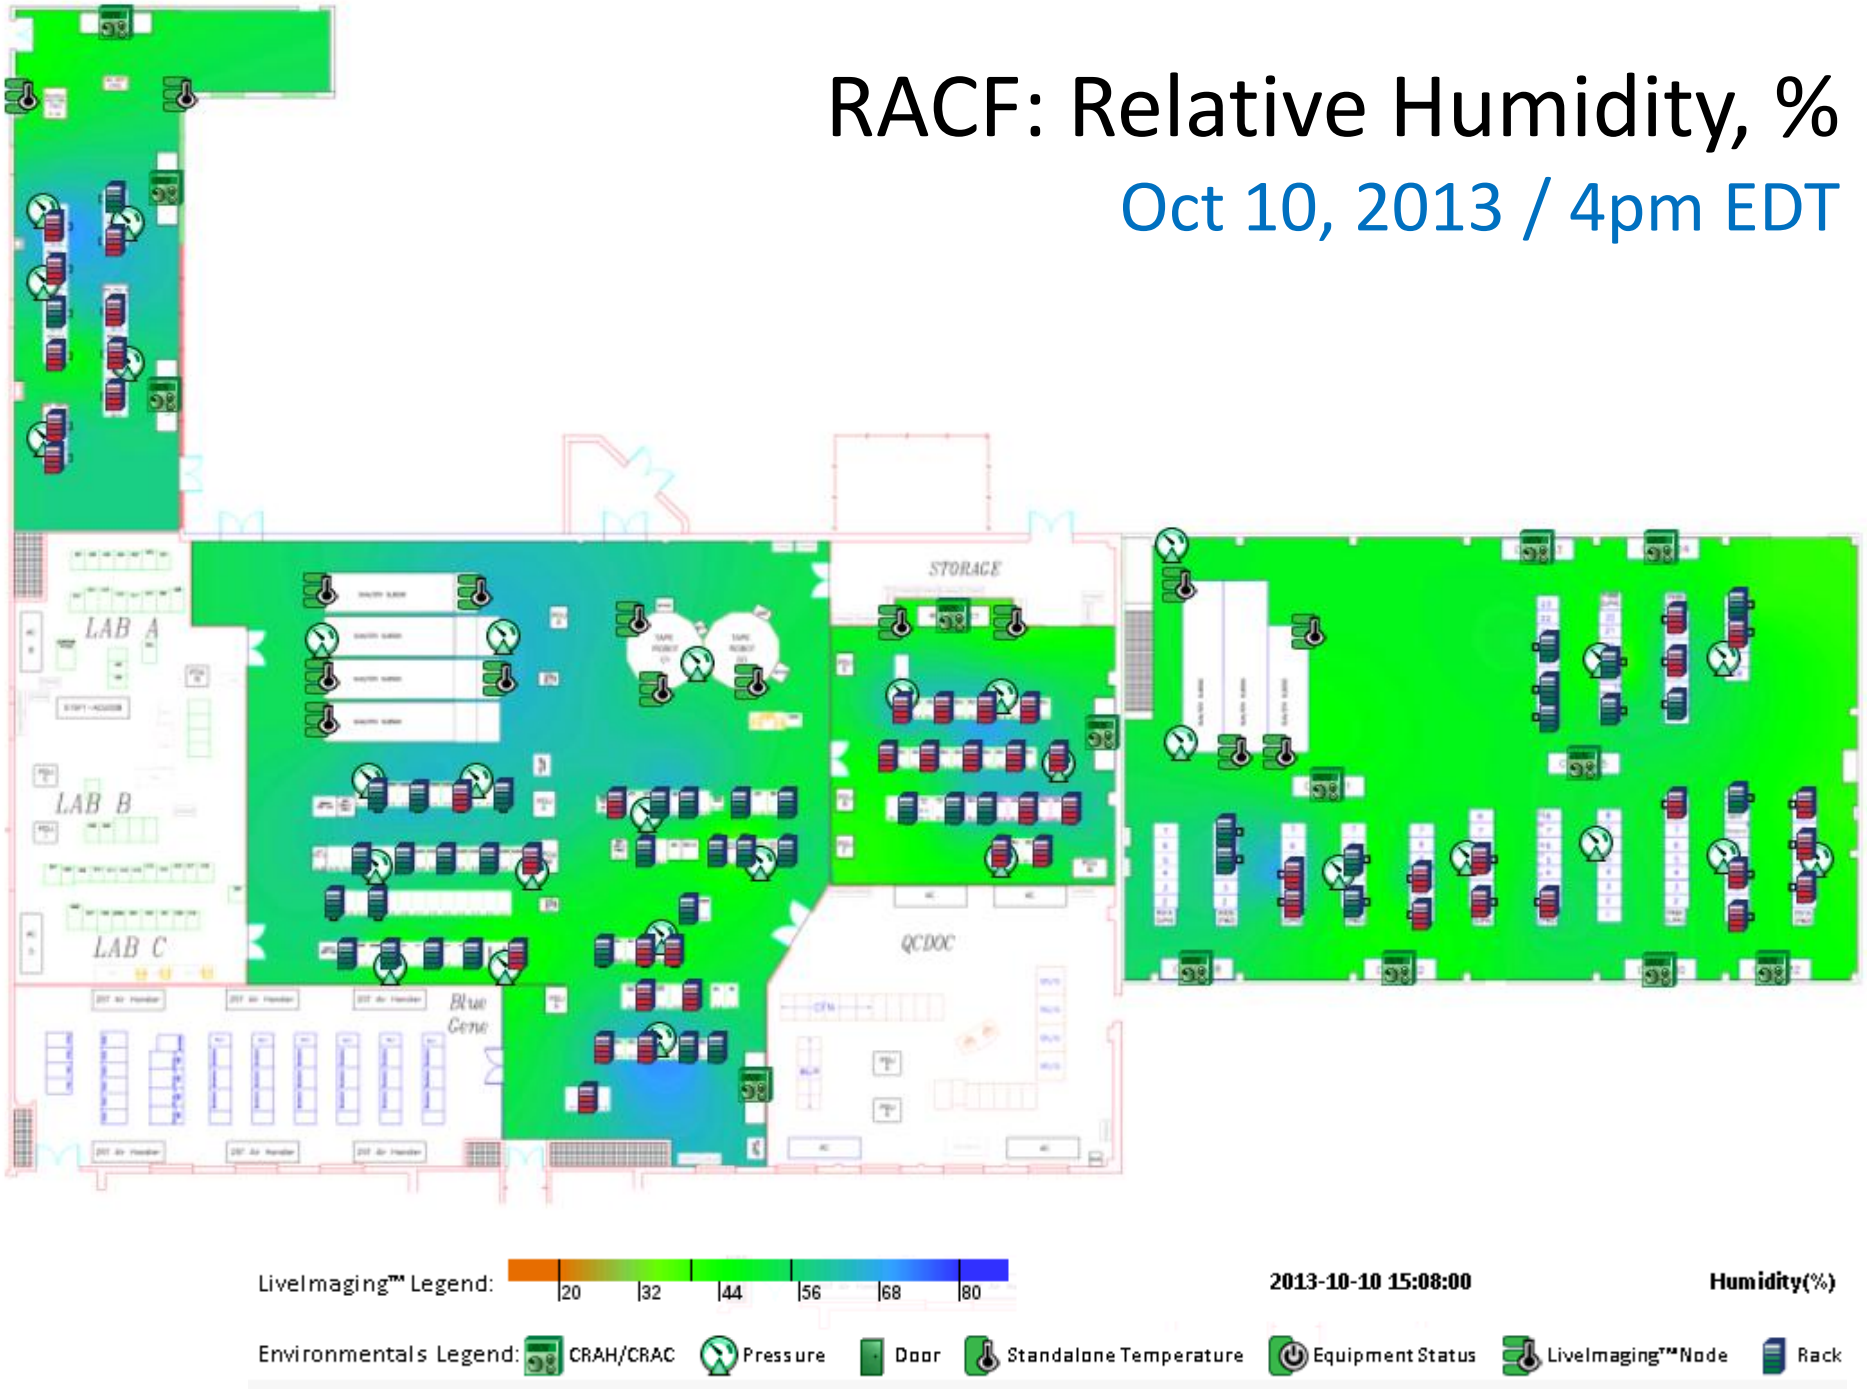

Summary and Future Plans

- Full scale deployment of the SynapSense™ environmental monitoring system was performed in RACF data center during the period of 2012Q2-2013Q3
  - In the present configuration the system has 150+ base stations provided with 920+ low systematic temperature / humidity / pressure sensors reporting to the central servers every 5 minutes (0.27M readings per day)
  - The integral cost of the system is not exceeding the cost of 2 racks of equipment typical for RACF Linux farms
- The experience gathered with the system so far indicate that we achieved the necessary level of time and spatial resolution that allows us to
  - Monitor the rack-by-rack cooling issues such as racktop cooling unit failures
  - IT room level issues such as raised floor pressure differential imbalance, AC unit failures and obstructed airflow issues
  - Observe and perform retrospective analysis of evolution of environmental conditions within the entire RACF data center
- There is a potential of extending the SynapSense™ monitoring system to include power consumption monitoring for all the CRAC and PDU devices in the facility and SynapSense™ ActiveControl™ features, thus providing real time estimates of the PUE of the data center and the means to optimize it

Q & A