Demand Response in Data Centers

Feasibility and Proof of Concept Demonstration

Jay Madden, SCE; Mukesh Khattar, EPRI;
Henry Wong, Intel; Chris Battista, Calit2

CalPlug Workshop Series #7
University of California, Irvine
May 12, 2015

Demand Response in Data Centers Project Team Members

Together…Shaping the Future of Electricity

EPRI's Mission

Advancing ***safe***, ***reliable***, ***affordable*** and ***environmentally responsible*** electricity for society through global collaboration, thought leadership and science & technology innovation

Together…Shaping the Future of Electricity

EPRI's Strategic Direction

Innovative solutions that enable the transformation of power systems to be more ***flexible***, ***resilient*** and ***connected*** to provide society with safe, reliable, affordable and environmentally responsible electricity

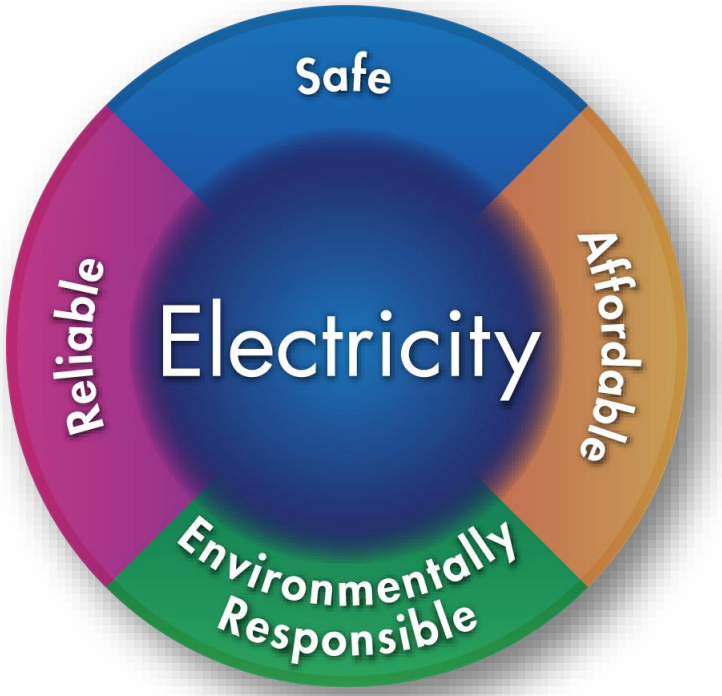

Three Key Aspects of EPRI

Independent
Objective, scientifically based results address reliability, efficiency, affordability, health, safety and the environment

Nonprofit
Chartered to serve the public benefit

Collaborative
Bring together scientists, engineers, academic researchers, industry experts

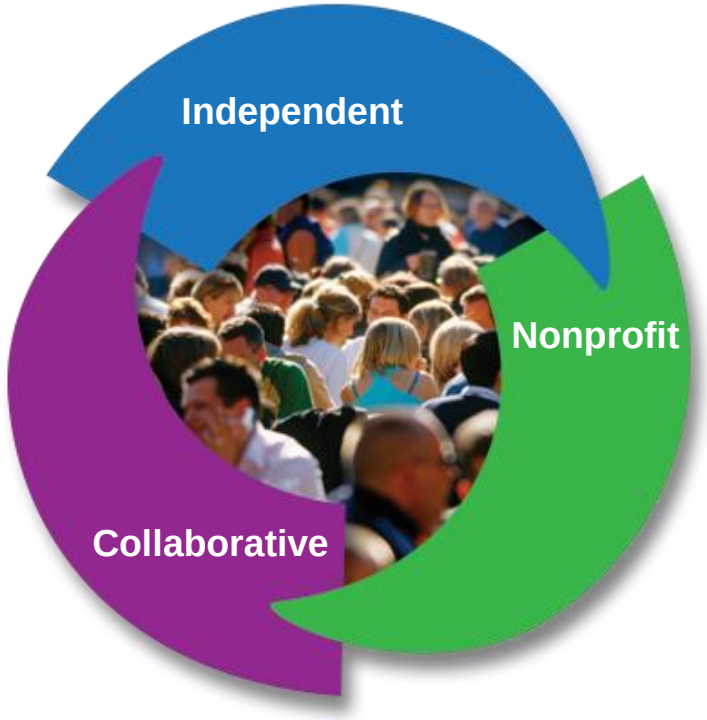

Our Members…

- 450+ participants in more than 30 countries
- EPRI members generate approximately 90% of the electricity in the United States
- International funding of nearly 25% of EPRI's research, development and demonstrations

Global CO_2 Emissions from Information and Telecommunications Technology Systems

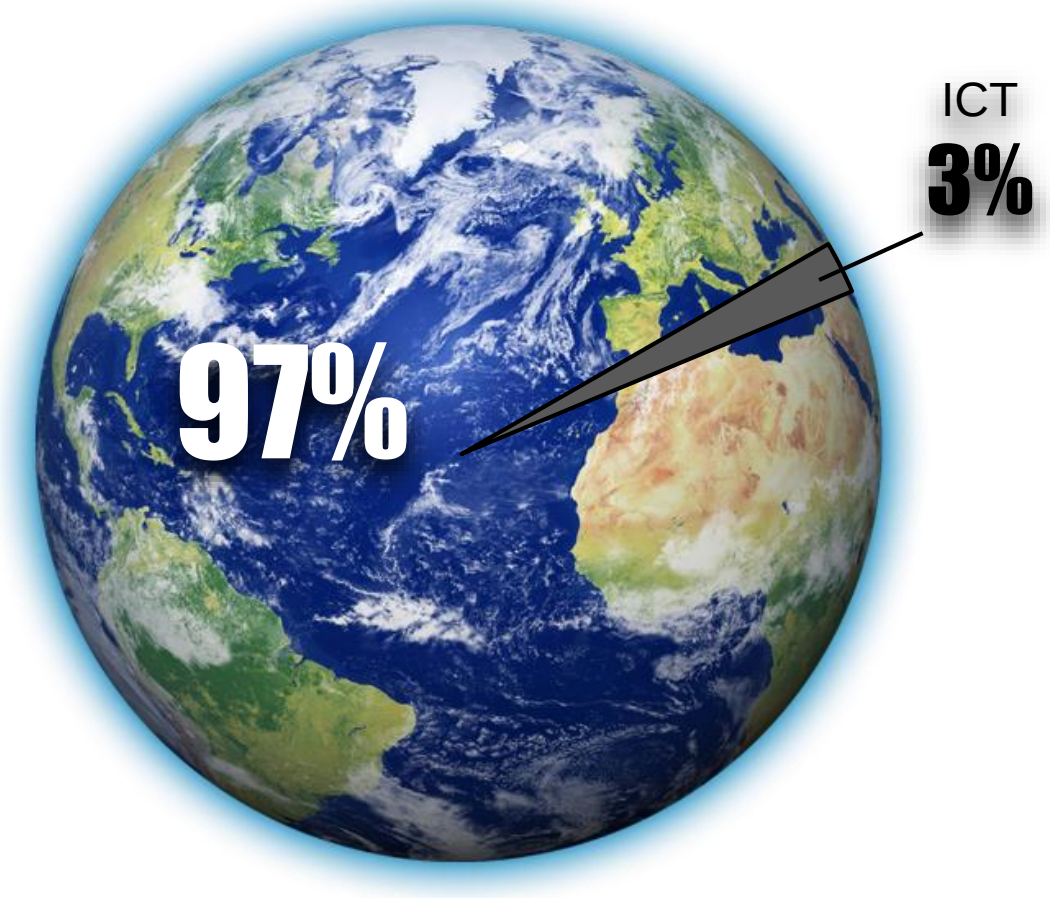

Government and Regulatory Requirements

Carbon Reduction Mandates

- California ~25-30% reduction by 2020 to reach 1990 levels, 33% renewable energy by 2020
- European Union mandates greenhouse gas (GHG) reduction by 40% below the 1990 level. EU-wide binding target for renewable energy of at least 27%

Net Zero Energy Building Mandates & Incentives

- California: "Zero Net Energy" (ZNE) goals for new homes by 2020 and commercial buildings by 2030
- **France**: Le Grenelle de l'environnement- Buildings must be net zero, to positive energy by 2020
- **European Union**: all buildings, including large houses, constructed after 2020 coming close to "nearly zero" energy use.

Demand Response- financial *incentive*

- Load abatement in response to signals
- Higher incentives for faster load-shed and for automated load control

Impact of Renewable Generation

- Increasing renewable generation in utility portfolio
- Renewable generation can be intermittent
- Shift in peak demand profile

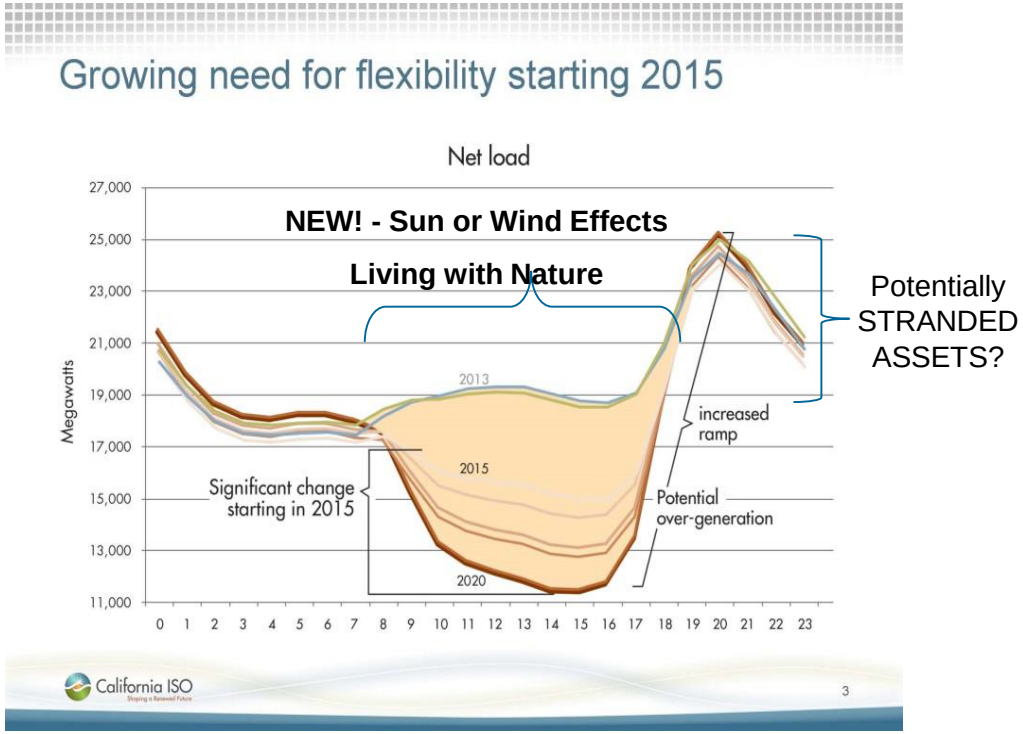

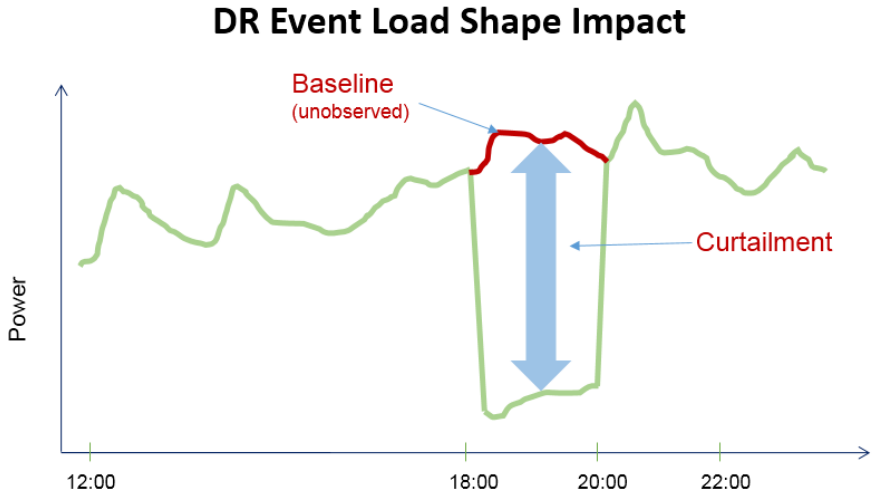

A Quick Overview of Demand Response

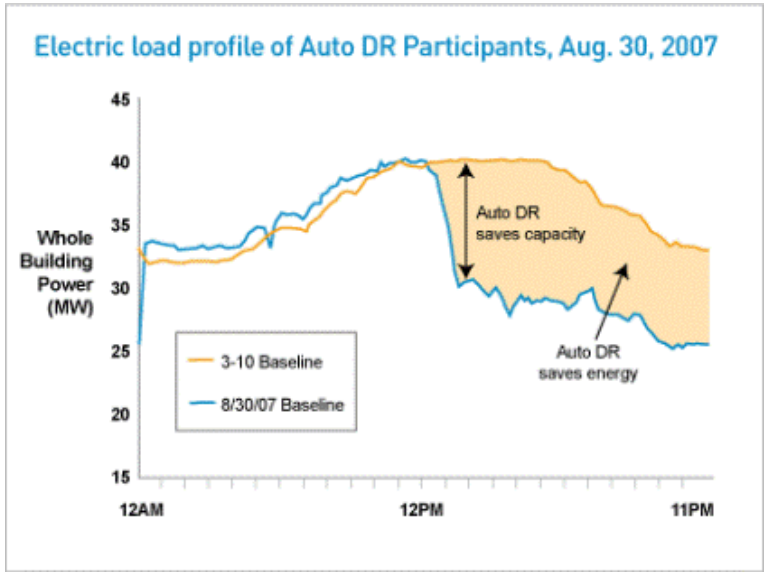

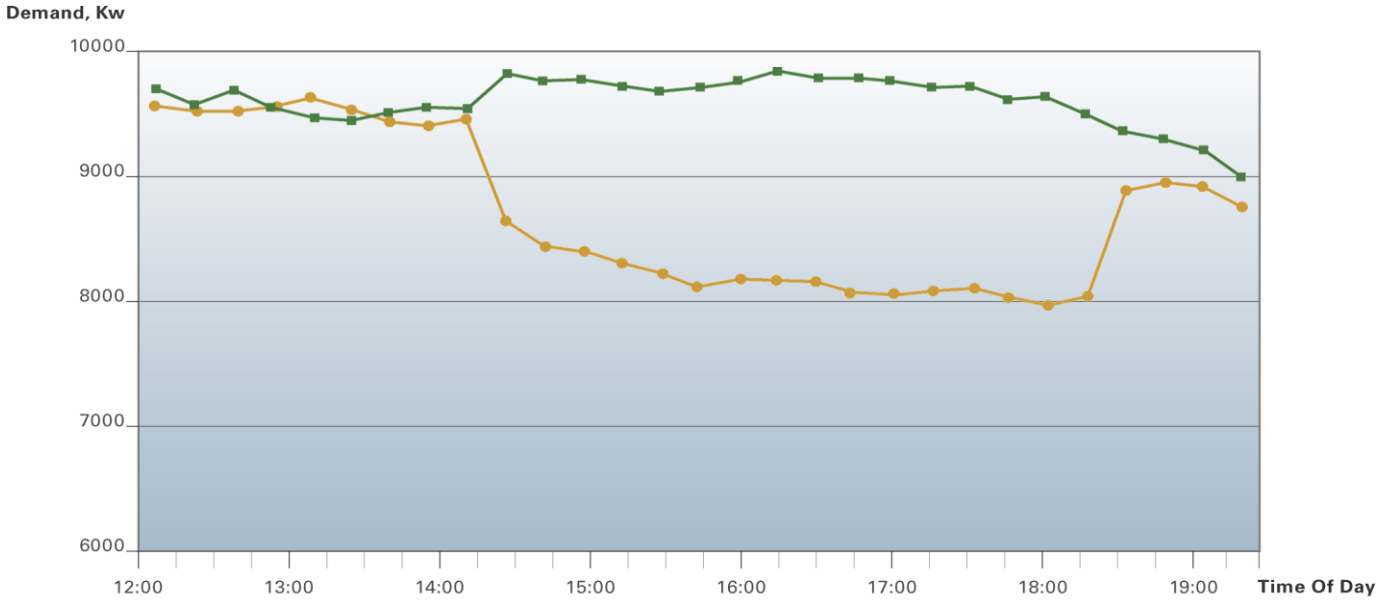

Solar Irradiance Data

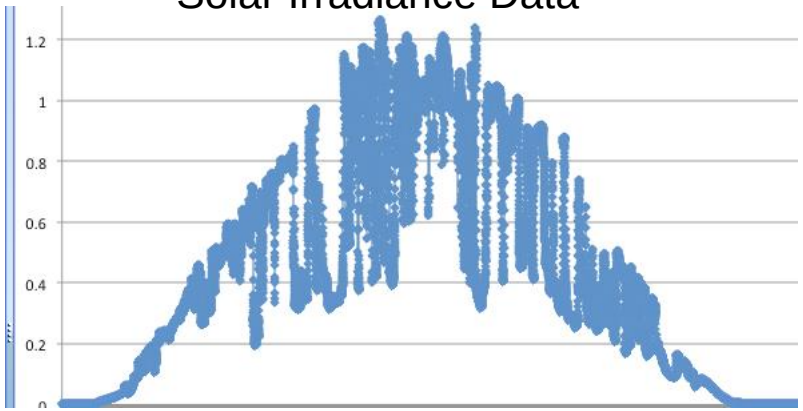

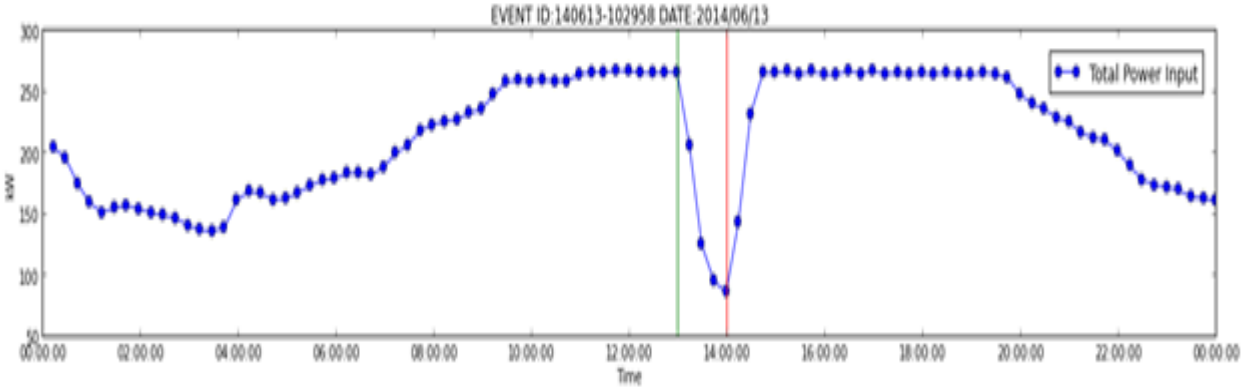

Growing Energy Use in Data Centers in the U.S.

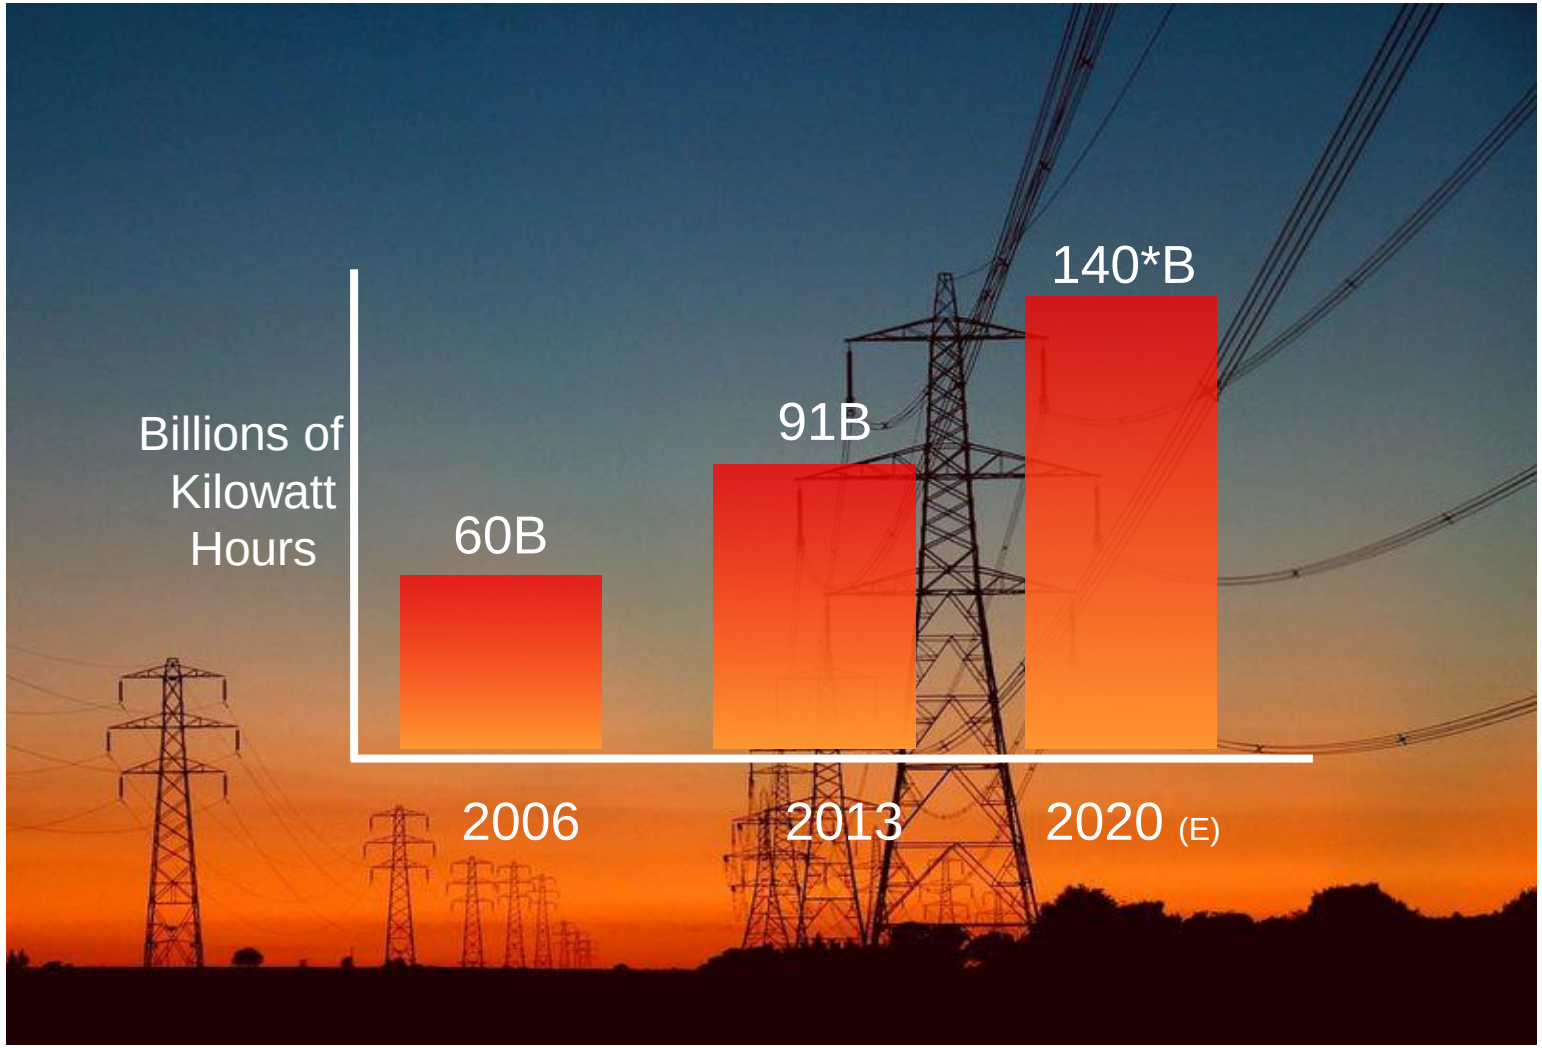

Global Increases in Data Center Power Consumption

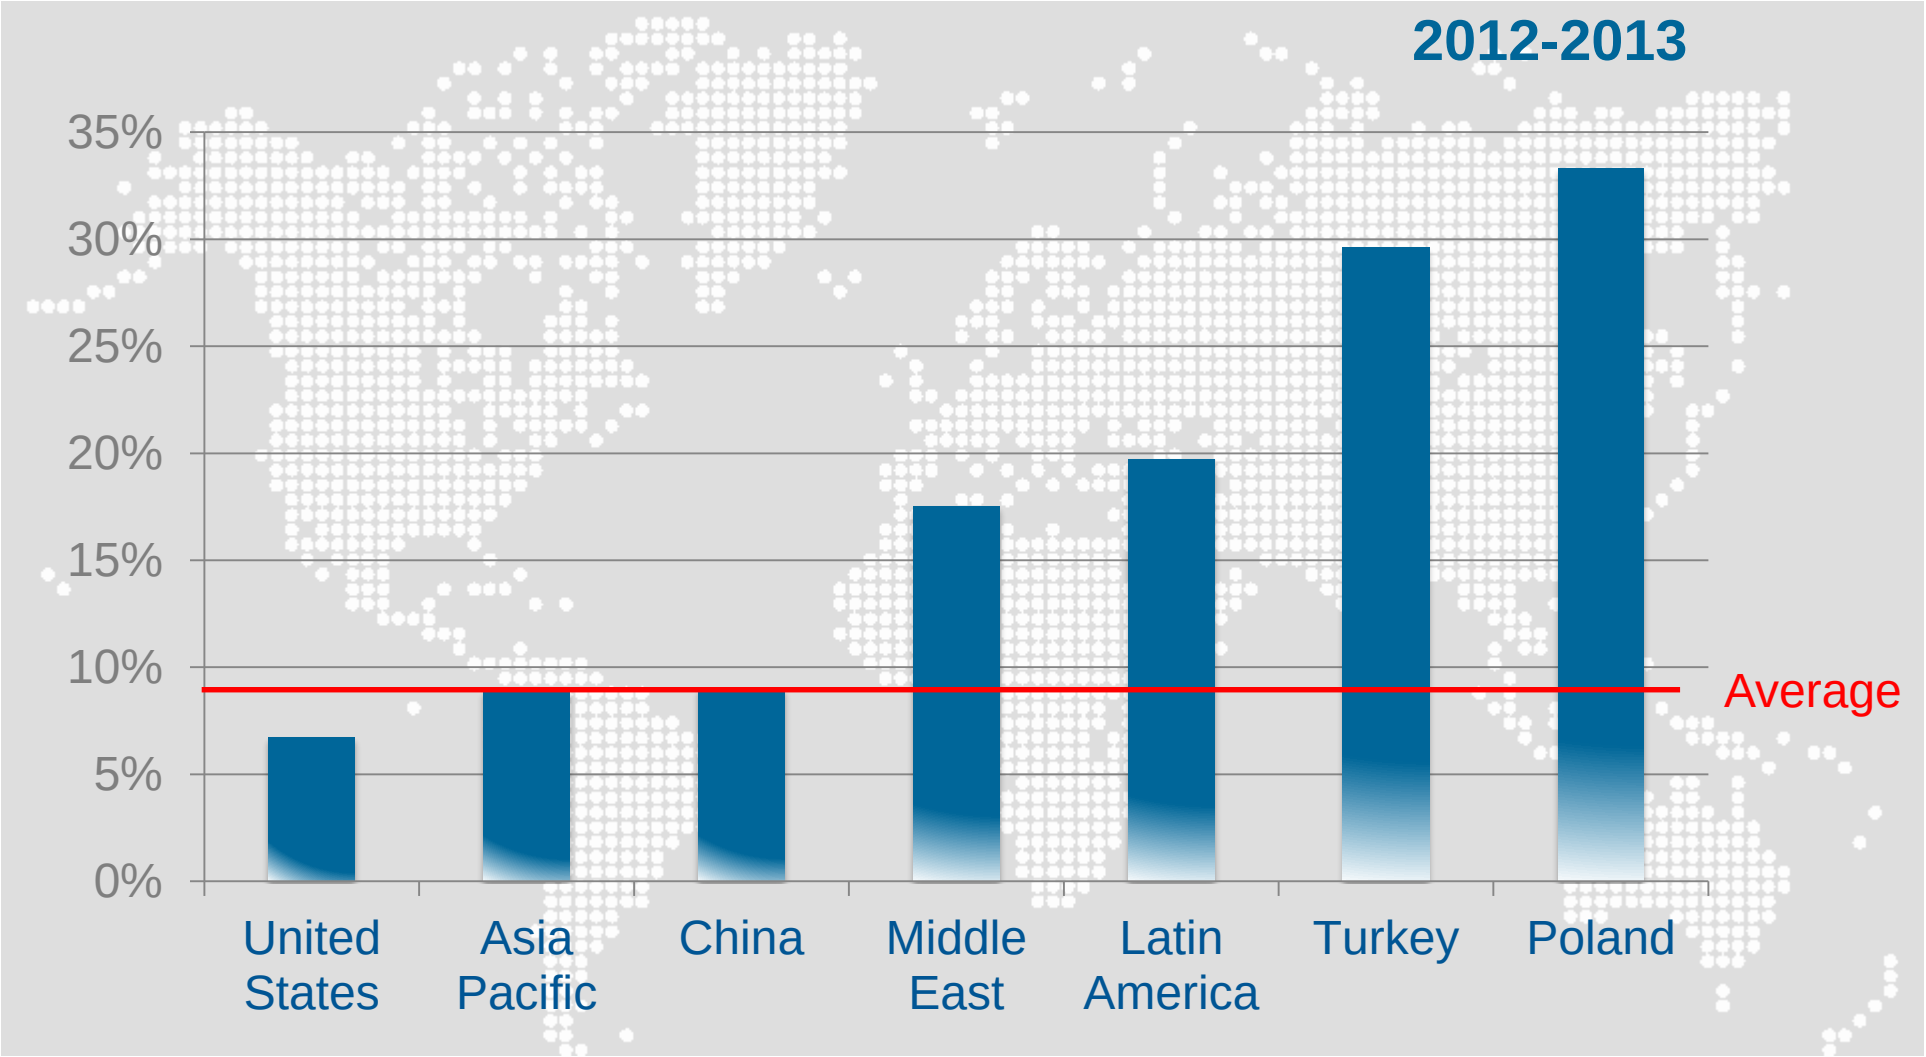

Source: DCD Intelligence 2013 Census Report: *Global Data Center Power 2013*

Data Centers Consume ~2% of Electricity in the U.S.

- Server/network closet
- Server room
- Localized data center
- Mid tier
- Enterprise

		Typical Power Draw		
<200 sq. ft.	<10	kW	8.7*	B kWh
<500	10-100	kW	10.7	
<1,000	100-1,000	kW	12.6	
<5,000	1,000-5,000	kW	11.3	
>5,000	>5,000	kW	25.7	

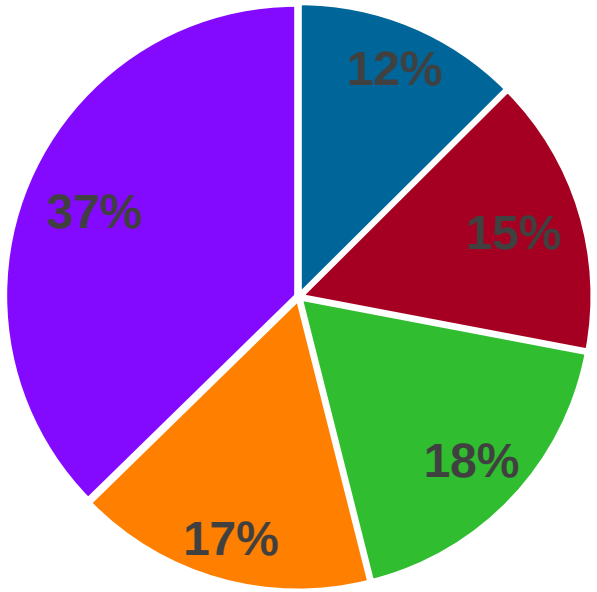

* Ref: 2011 IEEE

Data Center Loads Characteristics

- IT loads are plug loads
- Data Centers operate 7x24x365
- Concentrated loads
- High intensity loads
- Load profile is fairly flat
- Average server utilization is quite small

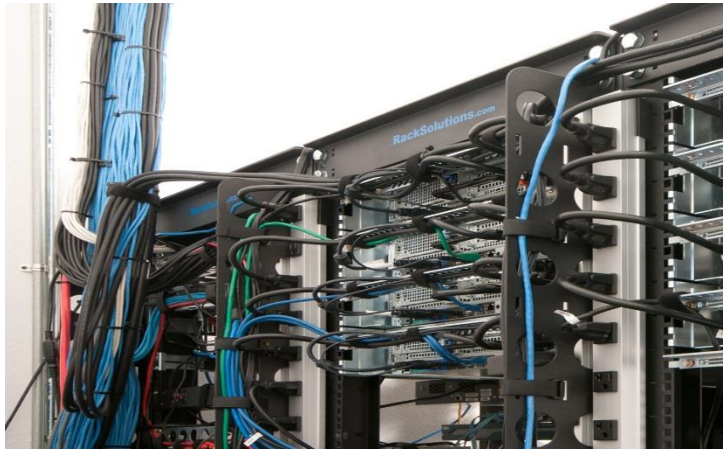

Facebook, Forest City, NC

Project Objectives/Opportunity

- Demonstrate feasibility of automatic IT loads reduction without disruption and without/or minimal user impact
  - Prepare for the future grid with greater renewables
  - Adjust data center power needs to electricity availability
  - Coordinate operations with utilities to avoid power interruptions
  - What kind of info is needed for such coordination?
  - How could this info be communicated?
  - What levels of reductions are feasible?

Demand Response for Data Centers: Challenges

- What info does it a signal need to communicate?
  - OpenADR 2 profiles
  - Opportunity to create additional info/profile for data centers
- How is it communicated?
  - Transmission across firewalls
- How is signal received?
  - Which software program element receives it?
  - Acknowledgement/confirmation/handshaking
  - Can handshaking confirm the amount of DR will the data center provide?
- What are decision making tree/rules for demand response?
  - What are the methodologies for power capping?
  - Is it simply an absolute amount of power to be reduced?
    - Is the power reduction from nominal power capacity or from current power draw?
    - Is it the same for each IT equipment or for the overall data center?

Automated Demand Response

Utility

Facility Equipment

End User

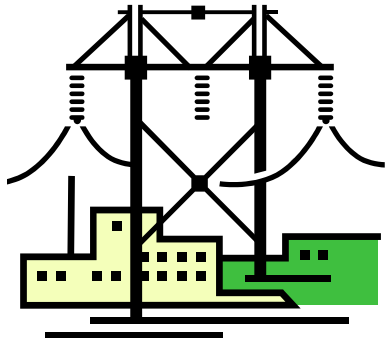

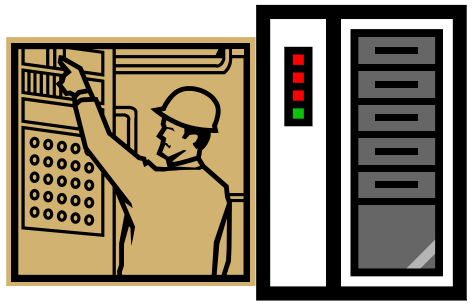

Power delivered to meet needs → Drives University facility demand → University workload

Field Setup and Monitoring

Chris Battista

Calit2

University of California, Irvine

Test Set Up

- 5 servers are set up in a rack
- Different type of activity in servers

2x Dell R520 (2U)

3x Dell R320 (1U)

Monitoring Equipment

Metering and Data Monitoring

- Sensor data downloaded remotely by EPRI
- Performance data collected using Intel IPMI and DCIM

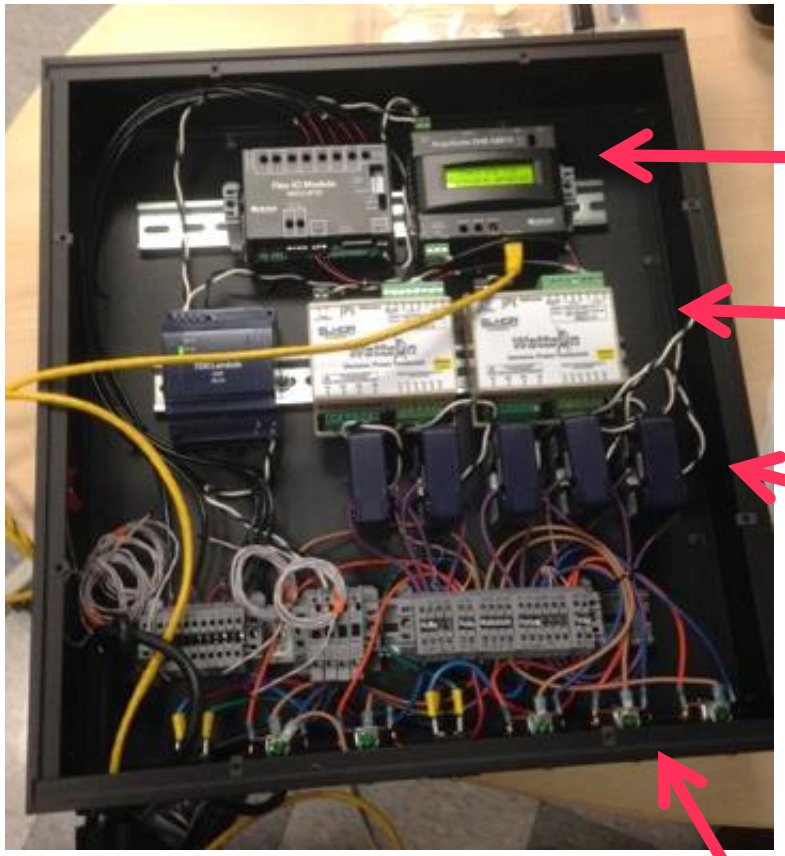

Baseline Data

- Data collection begun 10/29/2014
- Analysis in progress

Server	Average Power (W)	Min. Power (W)	Max. Power (W)	Average Exhaust Temp (°F)
1	73	64	110	95
2	77	76	100	99
3	46	45	58	104
4	57	56	63	103
5	46	45	56	100

Power Trends in Past Month

3/31/15 to 4/28/15

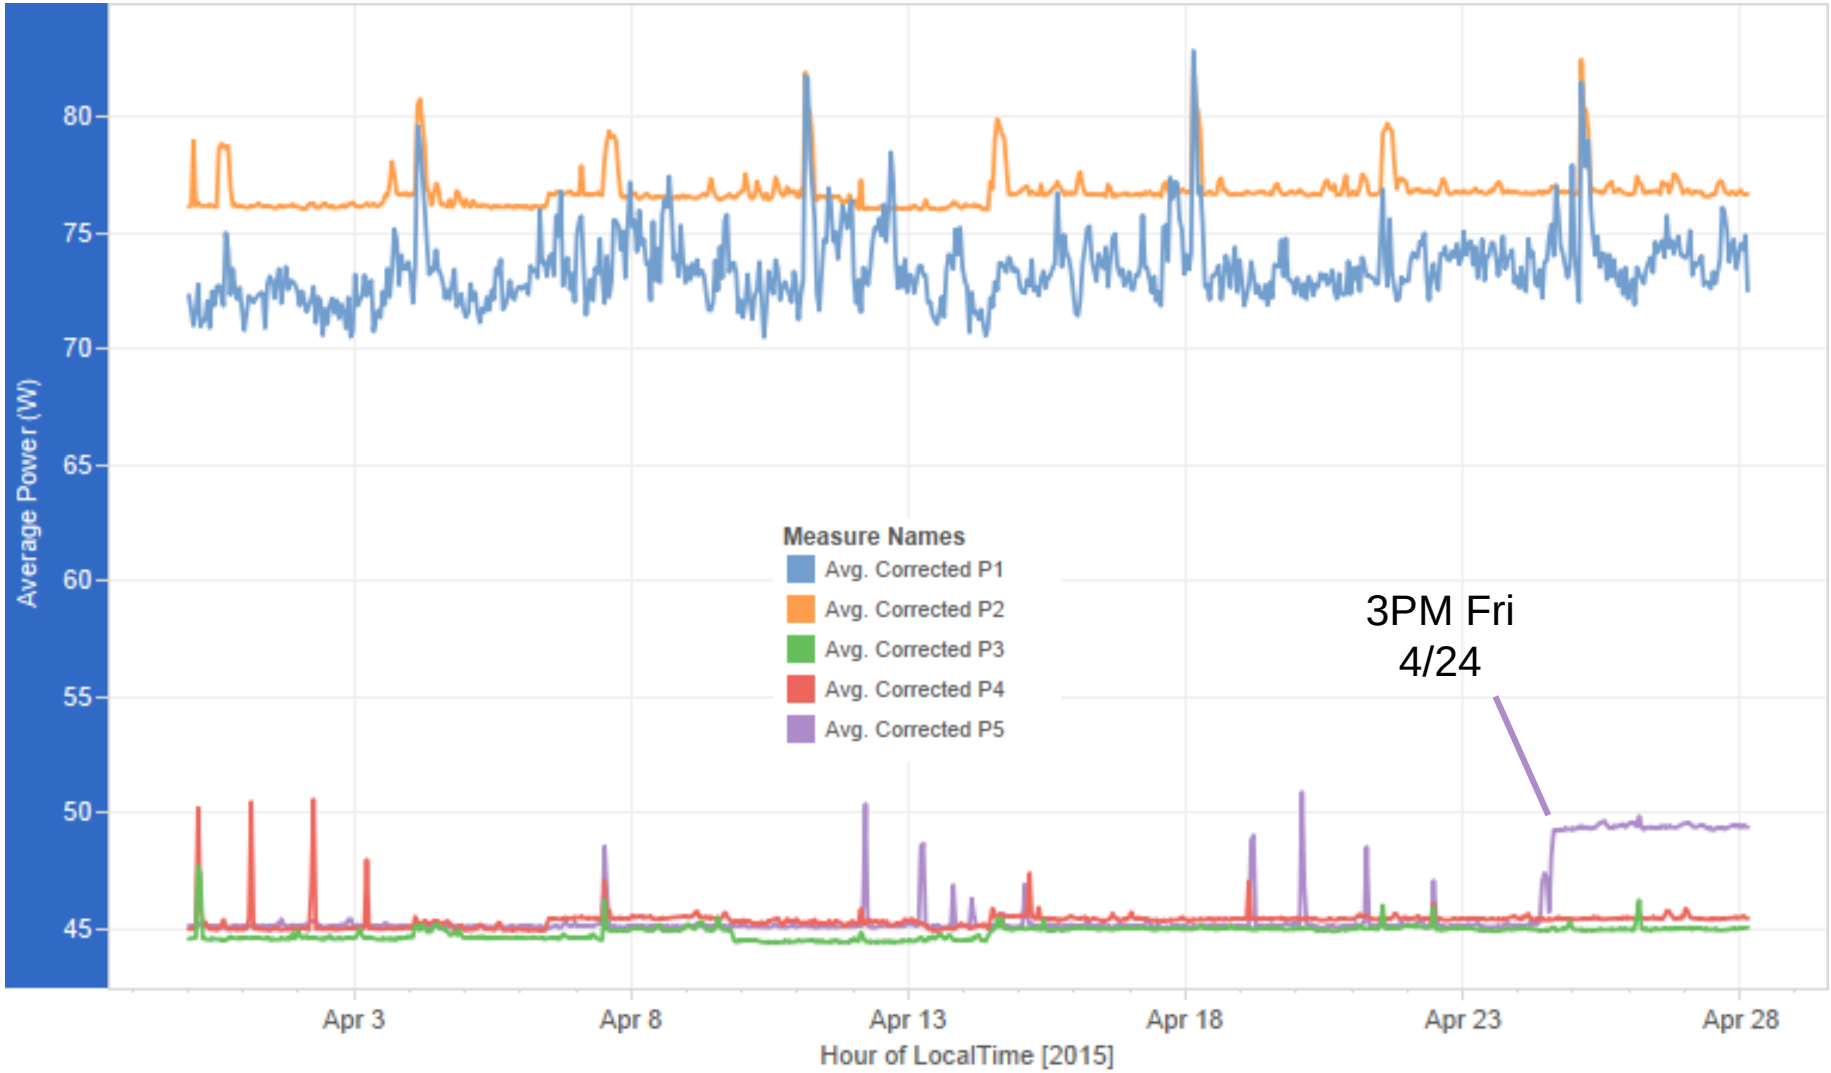

IT Level Power Provisioning

Business Continuity and ADR

Henry M.L. Wong
Sr. Power Technologist

EPRI ELECTRIC POWER RESEARCH INSTITUTE

Economic Efficiency and Technology Growth

>3.6 Billion

Connected Users by 2017 [1]

More Users

>2.2 Billion

Connected Devices shipping in 2017 [2]

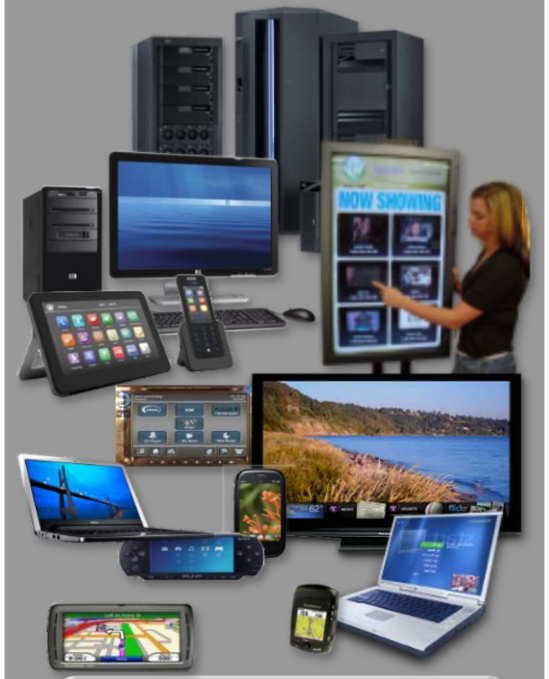

More Computing Devices

Data Centre Data Traffic 2017: 7.0

Zetabytes/annum [1]

More Data

Data demands continue to grow!

[1] Cisco Global Cloud Index Report, 2014

[2] IDC Tracker 2013

EPRI ELECTRIC POWER RESEARCH INSTITUTE

Critical: Availability of Services

BUSINESS

Regulatory compliance alone not protecting practices agair

A report notes that resources are being concentrated on government rules and i at guarding against information theft or loss.

By PAMELA LEWIS DOLAN, amednews staff. *Posted April 26, 2012.*

Regulatory compliance insufficient to safeguard and secure personal data

Availability and outages continues to jeopardize commerce

Top 13 Internet outages of 2007

Dec. 31, 2007 (8:46 am) By: Matthew Humphries

The uptime monitoring service company Pingdom has compiled a list of the 13 most notable Internet outages from 2007. Below is a summarized version with some of the big names in technology today featuring:

Gmail Down: Google Email Suffers Service Disruption (UPDATE)

April 2012

The Huffington Post | By Craig Karaliey
Posted: 04/17/2012 1:34 pm Updated: 04/17/2012 2:14 pm

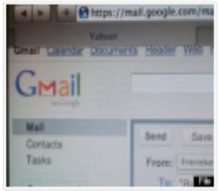

A number of Gmail users are finding error messages this afternoon, as the popular email service is suffering some rare downtime.

UPDATE: Google says the issue has been resolved:

"The problem with Google Mail should be resolved. We apologize for the inconvenience and thank you for your patience and continued support. Please rest assured that system reliability is a top priority at Google, and we are making continuous improvements to make our systems better."
EARLIER:

Google acknowledges the problem and says it is "working on it." It also released a statement that the outage is affecting less than 2 percent of its users worldwide:

Twitter crashes hard, Internet freaks out

By Julianne Pepitone @CNNMoneyTech June 21, 2012: 3:34 PM ET

June 2012

NEW YORK (CNNMoney) -- Cue the collective Internet freakout! Twitter went down for several hours on Thursday afternoon, depriving users of a place to complain that Twitter was down.

The Twitter outage began at 11:59 a.m. ET, according to Twitter's page on tracking site Pingdom. Service returned intermittently around 1 p.m., but less than an hour later, Twitter crashed again.

Balance Security, Availability, and Data Compliance in addition to Productivity and Energy Efficiency

Power Provisioning & Energy Efficiency

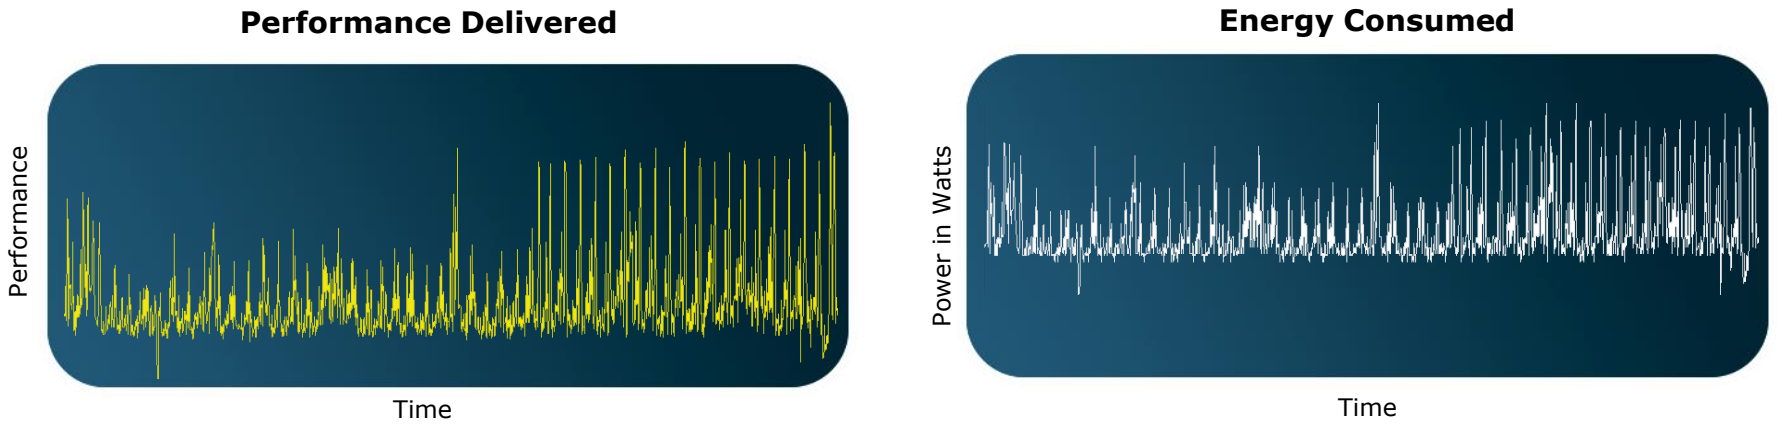

Power Levels and Energy Provisioning

- Static provisioning defines upper power bands
- Dynamic provisioning could map to an ideal energy profile
  - Standardized power controls to facilitate static and autonomous dynamic energy provisioning

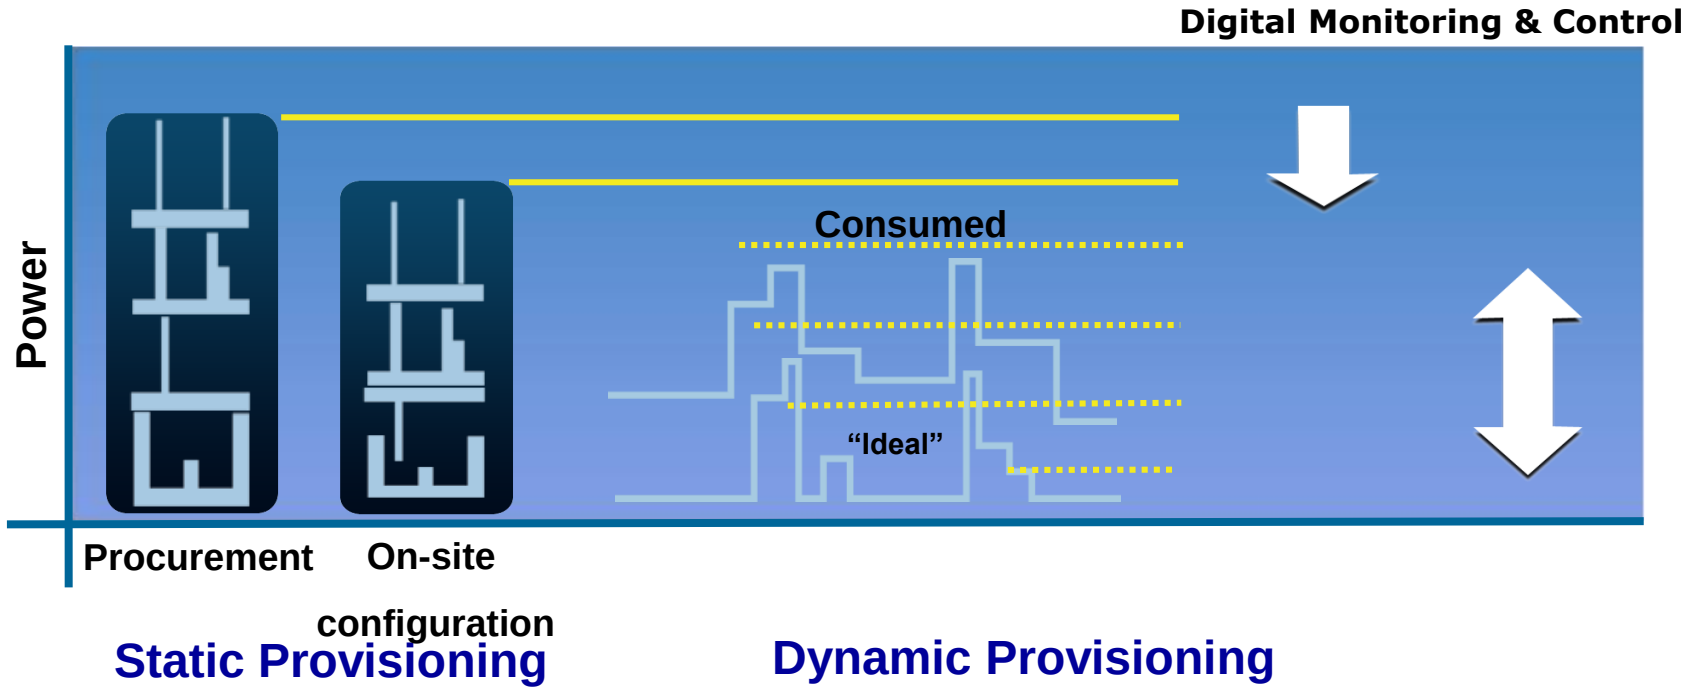

Digital Monitoring & Control Enables Static and Dynamic Provisioning

Data Center Power Provisioning

- PUE* measures the efficiency of data center infrastructure (UPS, CRACs, etc.)

$$PUE = \frac{Energy\ Consumed\ by\ the\ Data\ Center}{Energy\ Consumed\ by\ IT\ equipment} = \frac{IT\ Energy + nonIT\ Energy}{IT\ Energy}$$

Efficient Data Center Energy Usage (PUE)

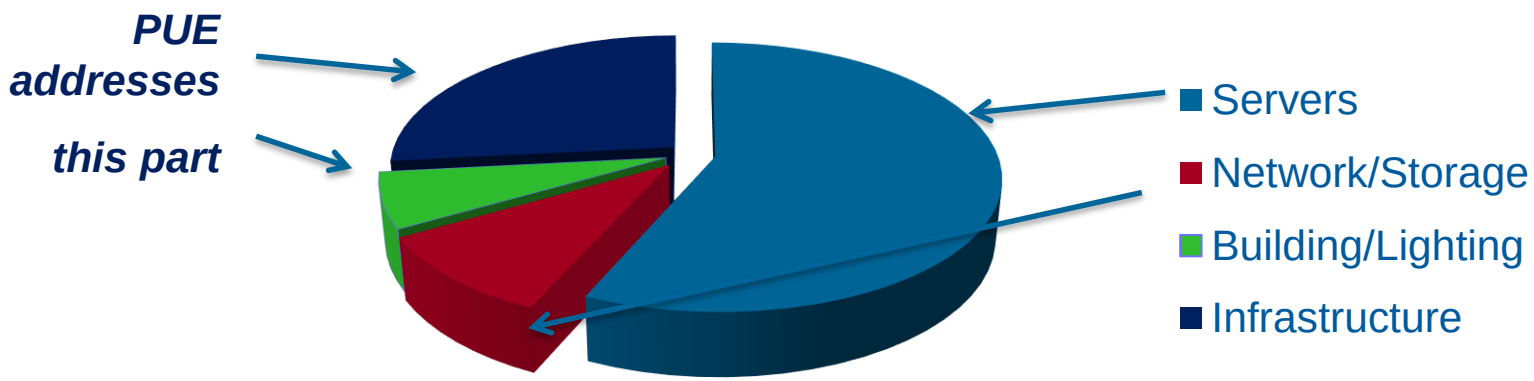

Server Energy Efficiency

DRIVE PERFORMANCE IMPROVEMENTS

CONTAIN GROWTH IN PEAK POWER

REDUCE POWER AT LOW UTILIZATION

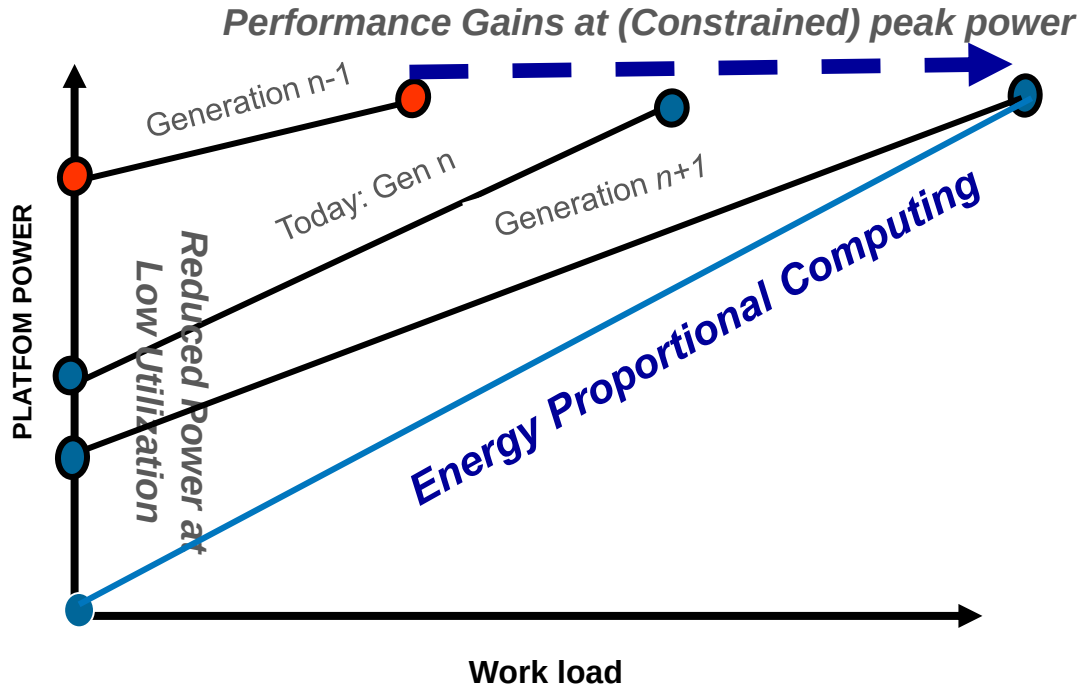

Server Energy Efficiency Across the Work Load

Energy Proportional Computing

Sustained Efficiency Improvement

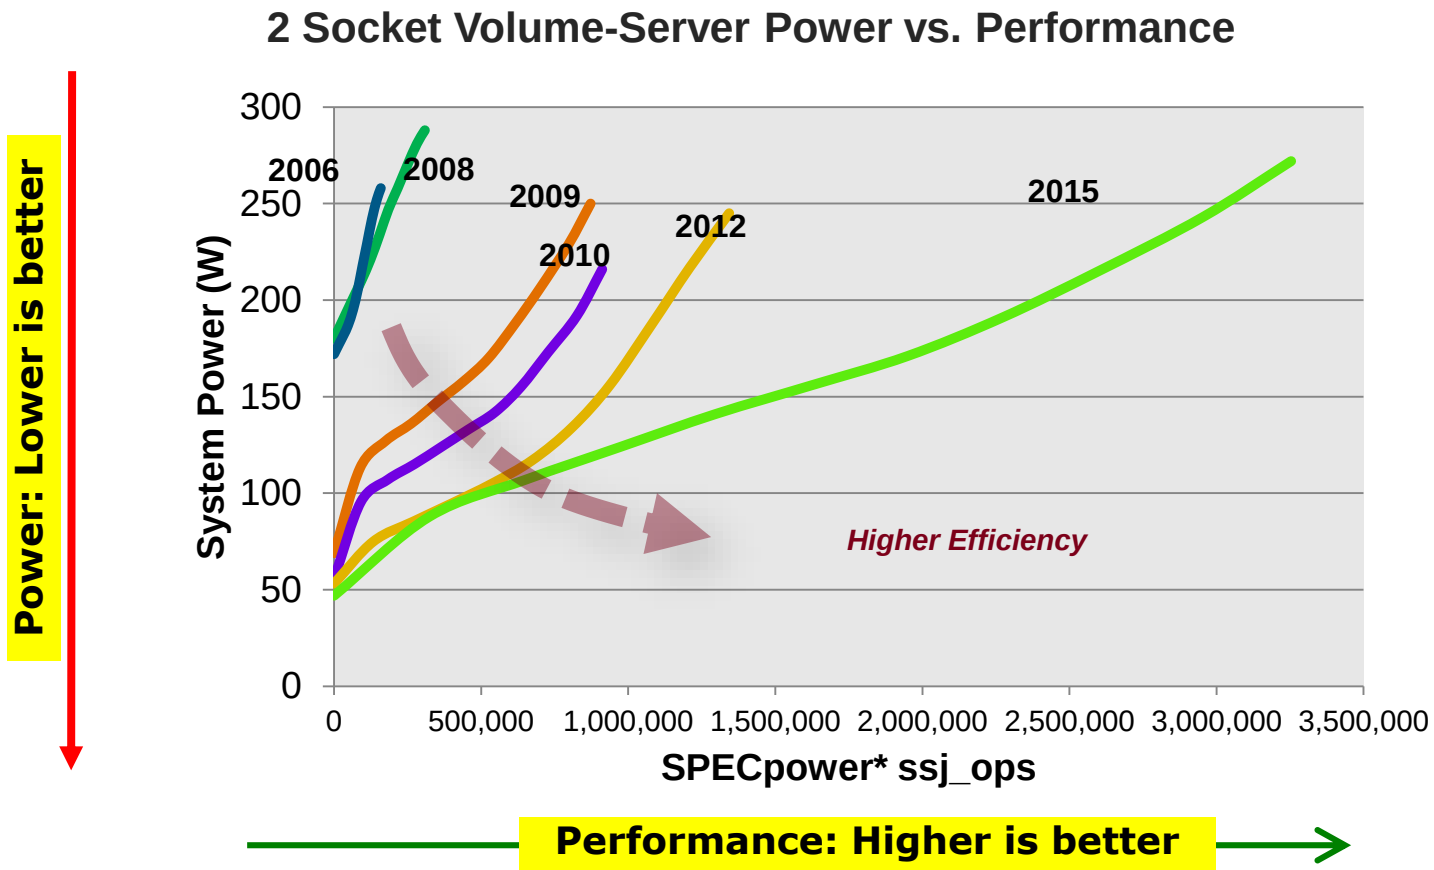

Source: SPECpower_ssj2008* 2 socket results from SPEC.org as of April 2015

Generational improvement in Energy Efficiency

Performance and power consumption results are based on certain tests measured on specific computer systems. Any difference in system hardware, software or configuration will affect actual performance. Configurations: Two-socket Systems, Test Results for SPECpower_ssj2008, **For more information go to** http://www.intel.com/performance

Materializing Energy Efficiency and Business Continuity

Nice in *theory*. But…

How does this keep my business running and save energy?

Extending Data Centre Operations
(2012/2013)

In response to Japan's power grid situation after the 2011 Earthquake, NTT Data Center operations were able to:

- Extend data centre run time capacity during backup generation from 36hrs to 65hrs
- Balanced extended operating time vs. service response level
- Realized up to 17% energy savings in the data center.

Generator Operating Time

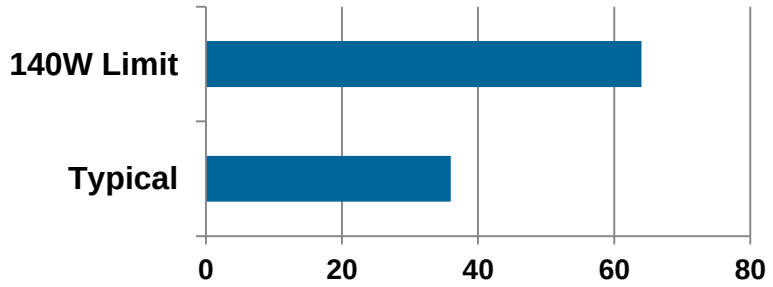

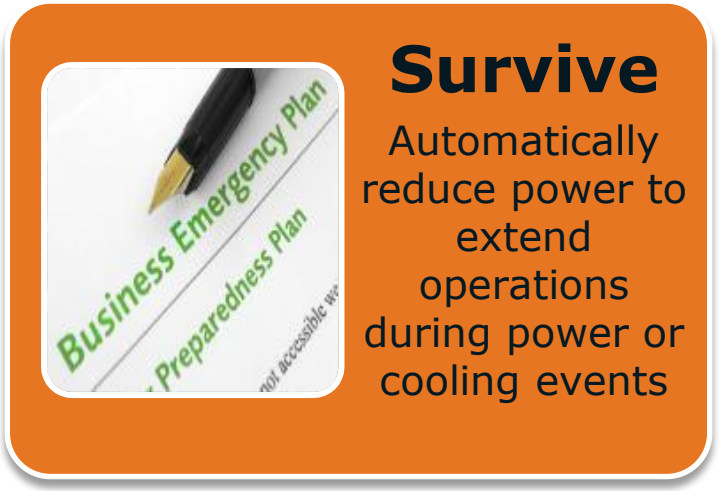

[Whitepaper available at http://software.intel.com/sites/datacentermanager/US_NTTDATA_CaseStudy_041912.pdf](http://software.intel.com/sites/datacentermanager/US_NTTDATA_CaseStudy_041912.pdf)

Power Capping and Automated Demand Response (ADR)

- ✓ Recognize and set power caps based on ADR signaling
- ✓ Determine user impact and latencies if any
- ✓ Determine communication requirements and negotiation parameters
- ✓ Provide recommendations for ADR and data center operations

Together…Shaping the Future of Electricity