High Performance Futures

High Performance Embedded Computing Workshop
MIT LL, Bedford, MA

Dr. Mark T. Maybury
Chief Scientist
United States Air Force

21 September 2011

Outline

- Trends in High Performance Computing
- Mission and Application Drivers
- Ensuring a Sustainable High Performance Future

HPC Need

- High Performance computing is foundational to:
  - Weather modeling, financial analysis/forecasting
  - Medical and scientific applications (physics, biology, chemistry)
  - Complex systems design (e.g., Aerospace)
  - Communication
    - Encryption/decryption
  - Multimedia and sensor processing (image, video, acoustic)
  - Search and data mining – pattern detection
  - Analytics
- Promise
  - Speed, discovery, efficiency

High Performance Computing (HPC) Modernization Program

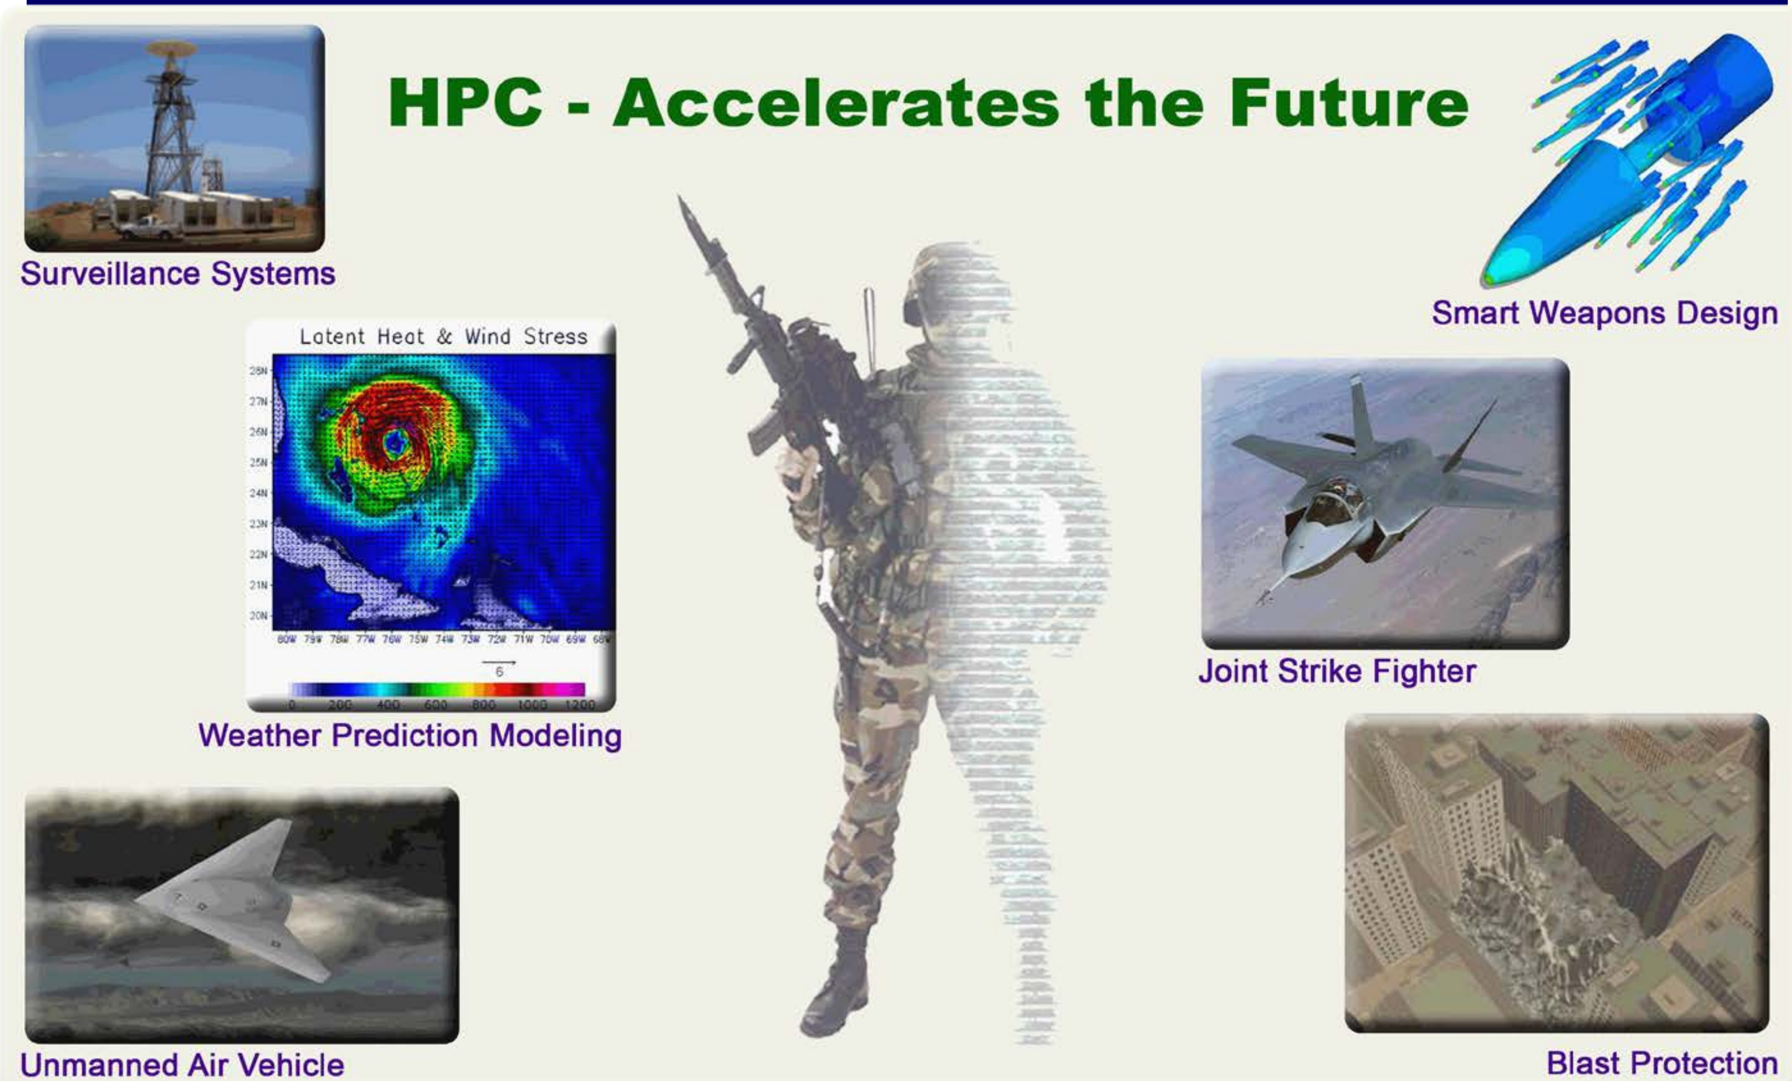

Distribution A. Approved for public release; distribution is unlimited.

High Performance Computing (HPC) Modernization Program

- **FY12 requirements**
  - 515 active projects with 4,661 users at 250 sites
  - 17,324 Habus based on required turnaround time
- **FY11 requirements**
  - 501 projects with 4,408 users
  - 5,098 Habus
- **FY11 usage (through 30 April 2011)**
  - 1,492 Habus

* Requirements and usage measured in *Habus*

Environmental Quality Modeling & Simulation – 189 Users

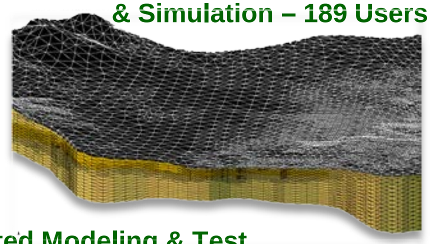

Integrated Modeling & Test Environments – 303 Users

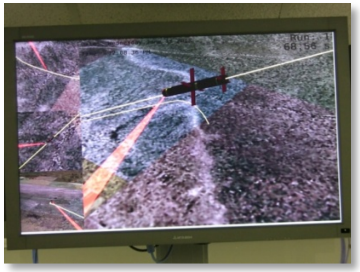

Forces Modeling & Simulation – 109 Users

New CTA added:
Space and Astrophysical Science – 46 Users

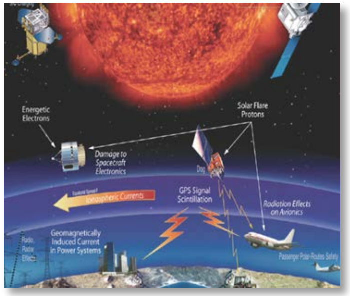

Signal/Image Processing – 545 Users

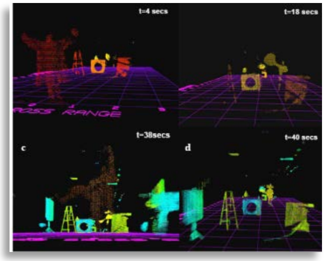

Computational Electromagnetics & Acoustics – 294 Users

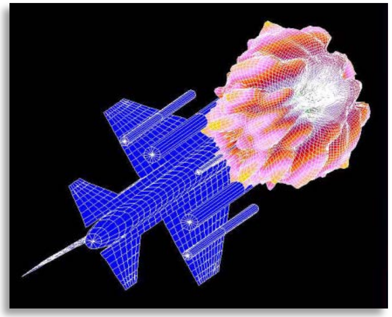

Computational Structural Mechanics – 497Users

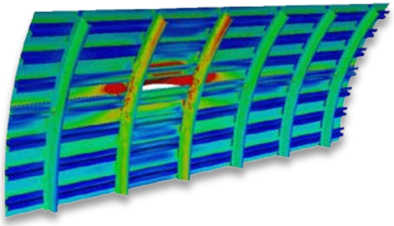

Climate/Weather/Ocean Modeling & Simulation – 326 Users

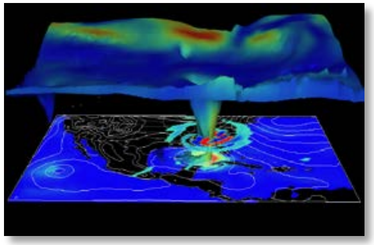

Computational Fluid Dynamics – 1,332 Users

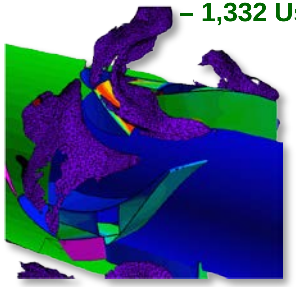

Electronics, Networking, and Systems/C4I – 239 Users

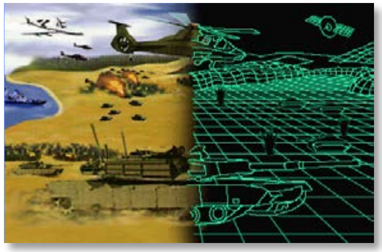

Computational Chemistry, Biology & Materials Science – 683 Users

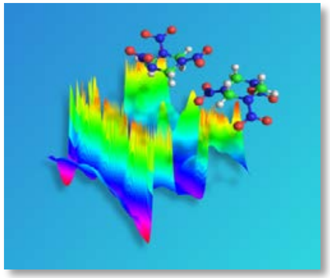

* Source: Portal to the Information Environment – May 2011

98 users are self characterized as "Other"

 MB Revised: 6/9/2011

DoD Supercomputing Resource Centers (DSRCs)

FY11 Capability

Program Structure

Location	System	Processors (Cores)	Memory (GB)	Capacity (Habus)
AFRL	SGI Altix 4700	9,216	20,480	94
	Cray XE6	44,756	90,851	1,037
ARL	Linux Networx Cluster	4,528	8,976	67
	Cray XT5 (Classified)	10,540	40,128	138
	SGI Altix ICE 8200	11,384	34,152	379
	SGI Altix ICE 8200 (Classified)	7,088	21,264	213
ERDC	Cray XT4	8,760	17,696	150
	SGI Altix ICE 8200	16,160	48,480	537
	Cray XE6	22,348	45,473	518
	Cray XE6	13,492	27,617	313
MHPCC	Dell PowerEdge M610	9,216	36,864	293
NAVY	IBM 1600 Power5 Cluster (Classified)	1,920	3,840	17
	Cray XT5	12,872	26,592	222
	IBM Power 6	5,312	8,448	117
	Total	177,592	430,861	4,095

FY10 HPC Systems shown in ORANGE
FY09 HPC Systems shown in GREEN
FY08 HPC Systems shown in RED
FY07 HPC Systems shown in BLUE
Older HPC Systems shown in BLACK

As of: July 2011

MB Revised: 8/11/2011

Workforce Development

Next-Generation DoD Workforce

Undergraduate STEM Internships

JEOM

HPC Computational Fellows

NDSEG

DoD Institutions of Higher Learning

Cadets & Midshipmen

- **Next-Generation DoD Workforce**
  - JEOM: A program for students from underrepresented groups, who are pursuing studies in the science, technology, engineering and mathematics (STEM) disciplines
  - HPC Computational Fellows: The HPCMP is a partner in the NDSEG program and our participation focuses on helping students interested in the breadth of disciplines related to computational science and engineering
  - Cadets & Midshipmen: The goal is to provide HPC opportunities to DoD Service academies and institutions of higher learning in order to prepare our future military leaders to excel in tomorrow's technologically-complex world
- **Professional Development**
  - Annual conferences and workshops are held that offer HPCMP community, Services, DoD community, peers, scientists and engineers from across the country to present and share ideas
  - These forums encourage collaboration on projects, techniques and methods and the opportunity to hear HPC strategic directions from invited speakers from industry, academia and other government agencies

Professional Development

Oversight & Outreach

Users Group Conference

Software Applications Support

CREATE Developers Review

Resource Management

S/AAA Workshop

Networking & Security

DREN Networking & Security Conference

Trends

- Petascale to exascale
- Hardware diversity
  - Clusters
  - GPUs
  - FPGAs
- Algorithmic complexity
- Mobile devices
  - Driving small energy efficient memory and processors
  - Nanoelectronics for SWAP (e.g., nanowires, memristers)
- Cloud computing
- Green – infrastructure and algorithms


Source: www.dodlive.mil

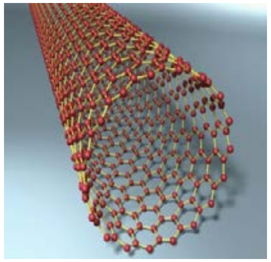

Global Performance

(top500.org)

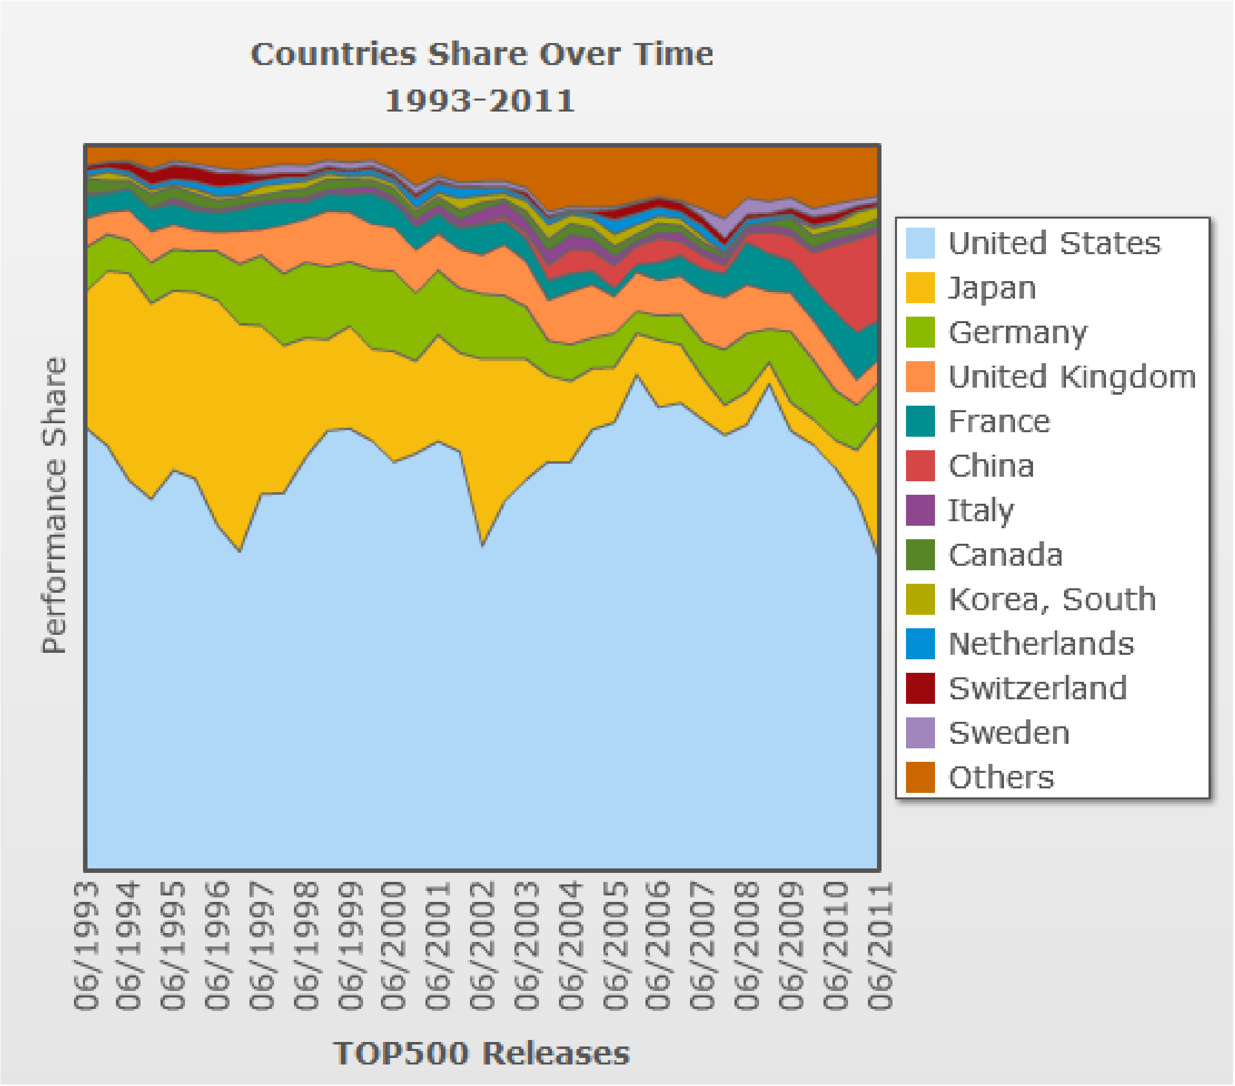

Projected Performance

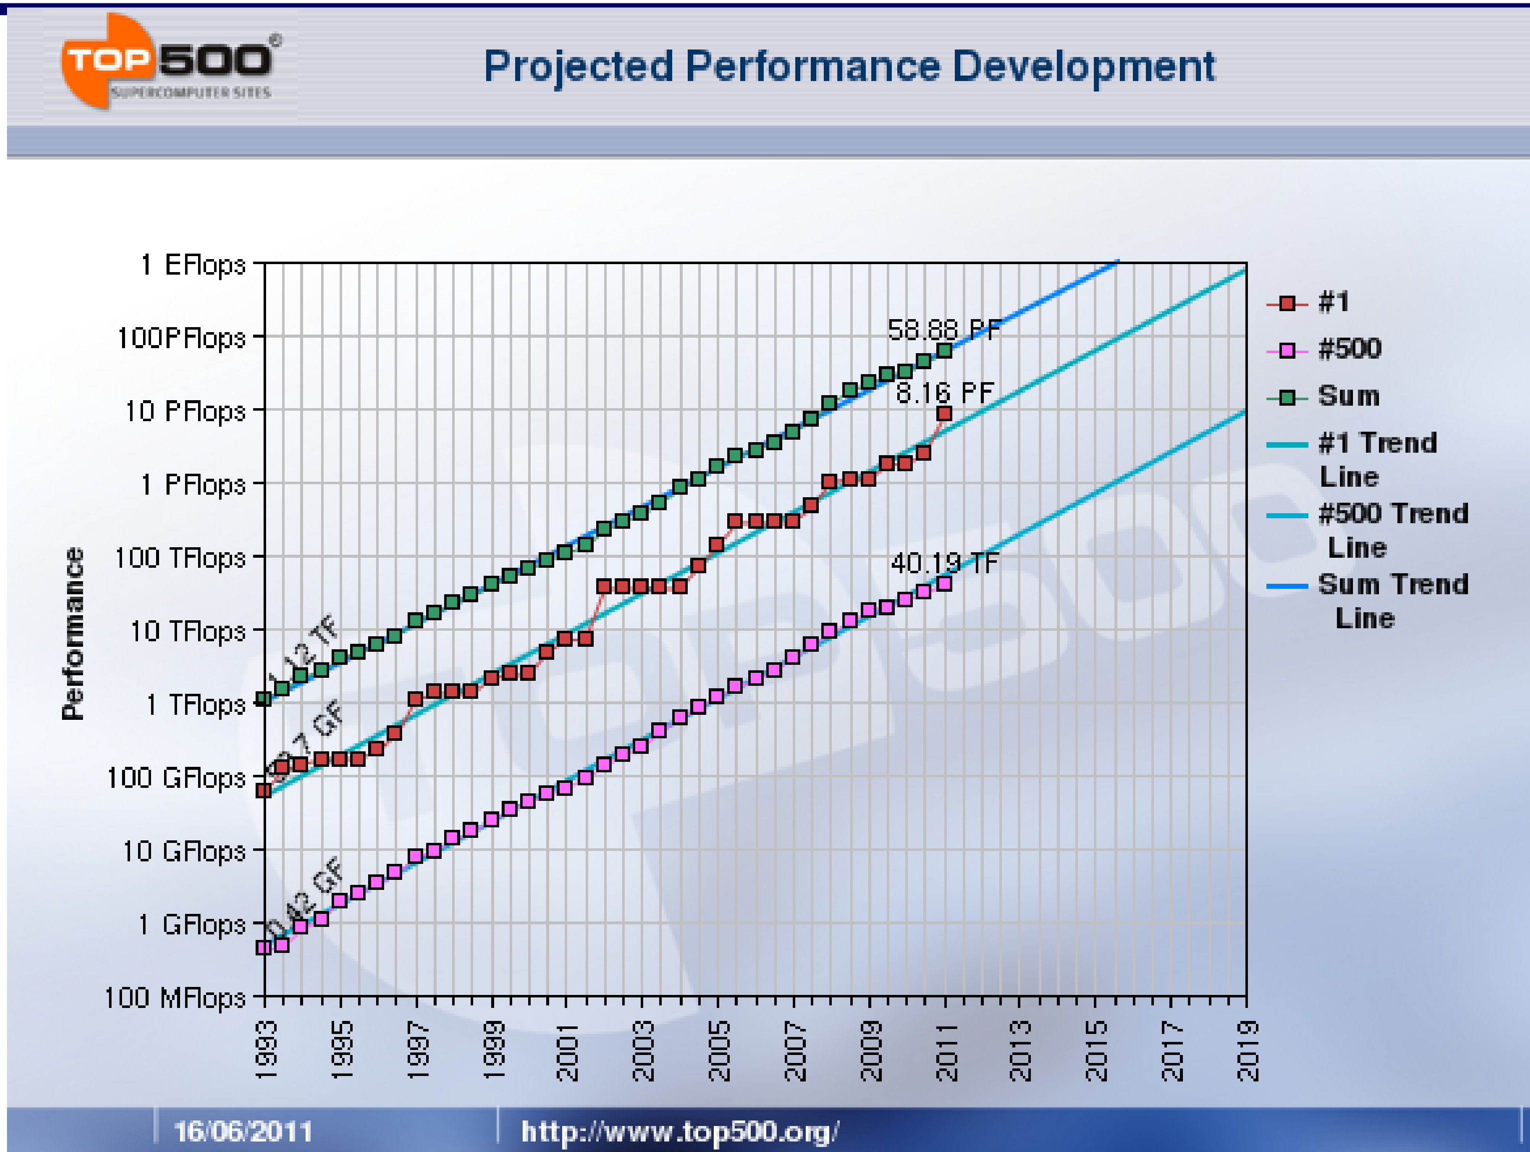

Extreme S&E Discovery Environment (Xsede)

- **Follow on to NSF's 11-year TerraGrid HPC**
- **Xsede: HPC, Data, Collaboration, instruments**
- **Compute: 2720 Peak Teraflops and 1530 TB of disk storage across 13 resources**
- **Visualization: 30 Peak Teraflops and 2709 TB disk storage across 3 resources**
- **Focus: researcher productivity**

University of Illinois at Urbana-Champaign, Carnegie Mellon University/University of Pittsburgh, University of Texas at Austin, University of Tennessee Knoxville, University of Virginia, Shodor Education Foundation, Southeastern Universities Research Association, University of Chicago, University of California San Diego, Indiana University, Jülich Supercomputing Centre, Purdue University, Cornell University, Ohio State University, University of California Berkeley, Rice University, and the National Center for Atmospheric Research. Led by University of Illinois National Center for Supercomputing Applications.

Cloud Computing

- **Benefits**
  - **Cost efficiency**
  - **Manufacturing agility**
  - **Reliability**
  - **Maintainability**
  - **Democratization**
- **Challenges**
  - **Efficiency**
  - **Security**
  - **Reliability**
  - **Performance**
  - **Verifiability**

Source: https://www.fbo.gov/index?s=opportunity&mode=form&id=8d70683b99829ca8226f4af6163a80c6&tab=core&_cview=0

Dr. Charles Holland, Program Manager, DARPA, MTO

Larger, Smaller, Faster

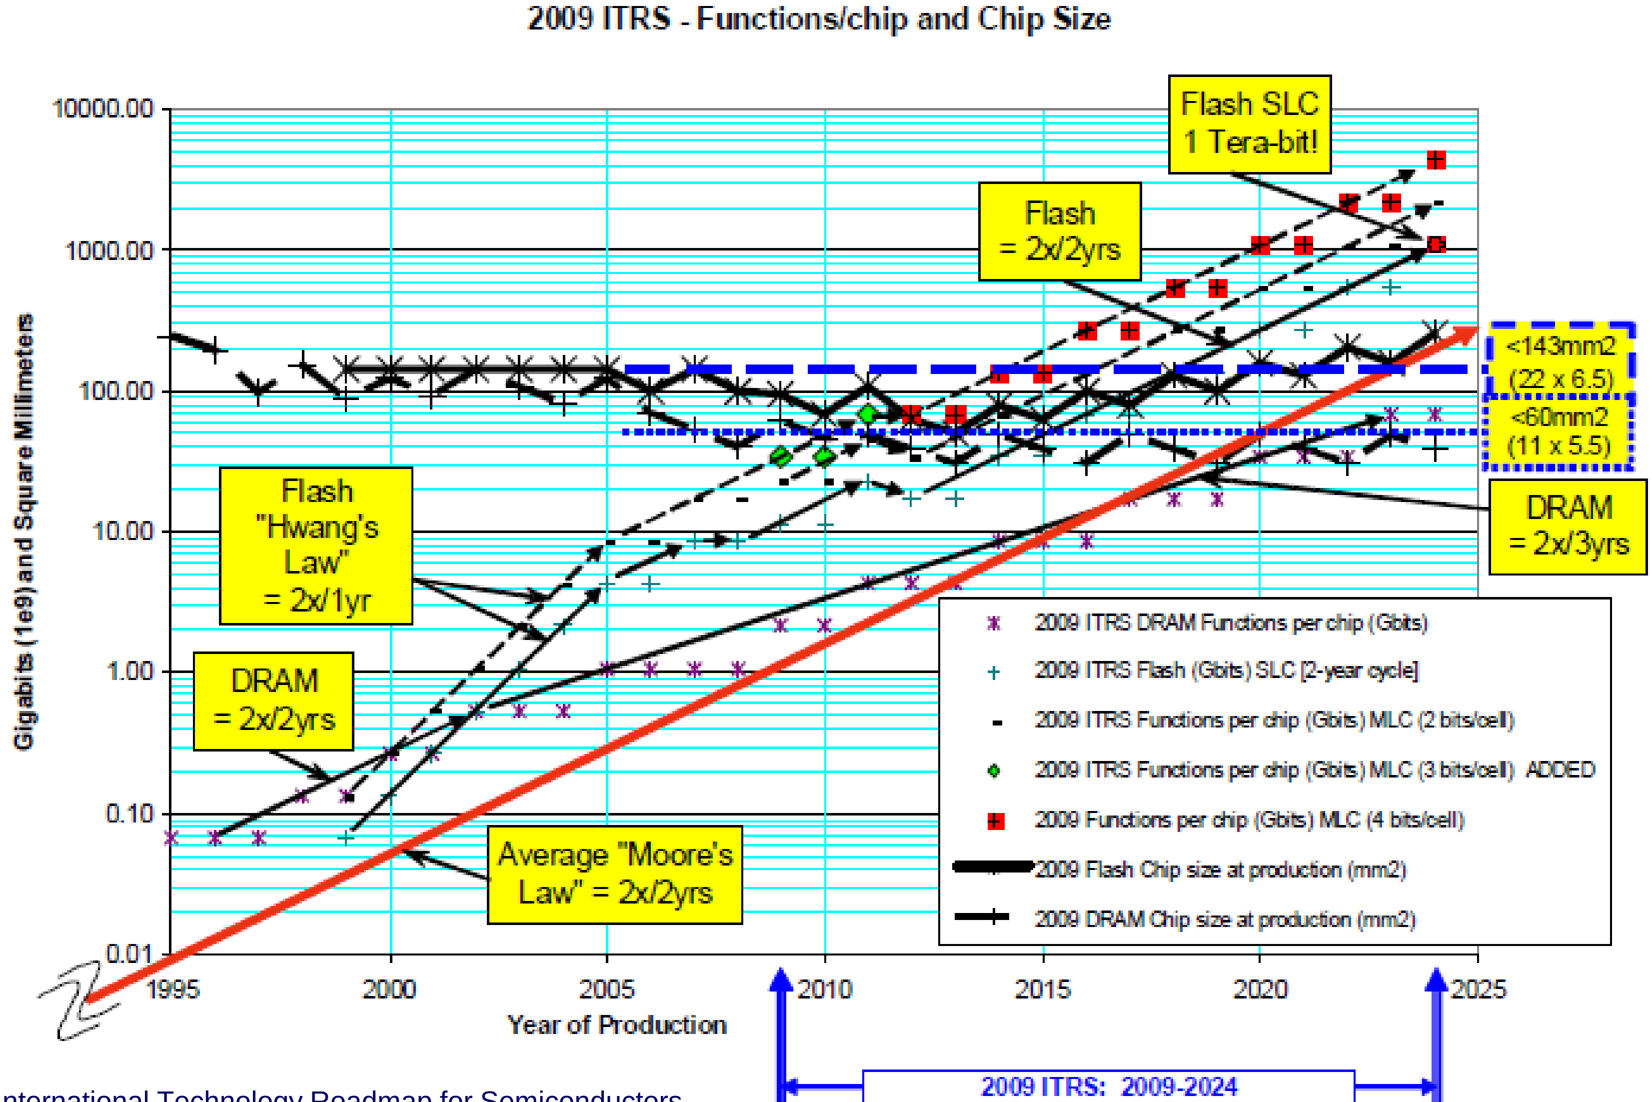

Source: International Technology Roadmap for Semiconductors

More Moore and More than More (MtM)

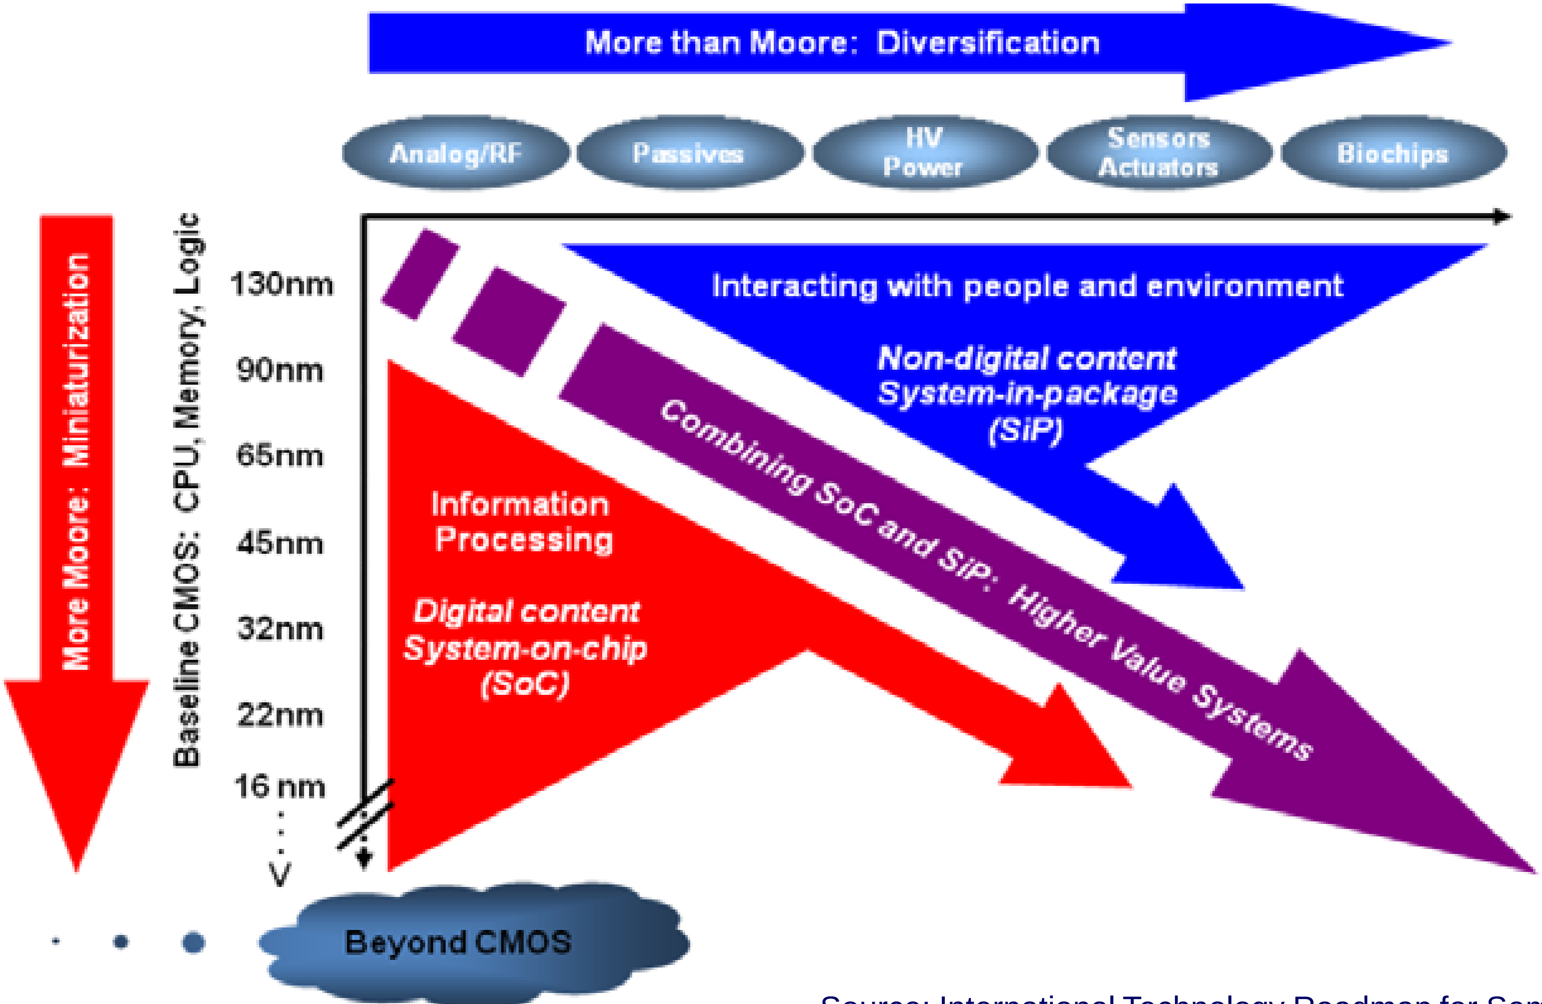

Source: International Technology Roadmap for Semiconductors
Complementary metal–oxide–semiconductor (CMOS)

National Energy Security

"Our military leaders recognize the security imperative of increasing the use of alternative fuels, decreasing energy use, reducing our reliance on imported oil, making ourselves more energy-efficient." **President Barack Obama, 31 Mar 2010**

"The Air Force is pushing forward, focusing on three goals of reducing demand, increasing supply through renewable and alternative sources, and changing the culture. For the last several years, from my perspective, the Air Force has led the way in this area."

Adm Mike Mullen, Chairman of the Joint Chiefs of Staff, 13 Oct 2010

"For the Air Force's part, we must embrace the notion that energy efficiency is not a stand-alone priority because it binds together and enables every dimension of our mission; and the idea that energy efficiency affords us greater resiliency, which translates to greater capability and versatility."

Gen. Norton Schwartz, Chief of Staff of the Air Force, 28 May 2010

"Changing the culture means that all of us, from the Air Staff to Airmen at home or deployed, must learn to think of energy as part of maximizing mission effectiveness."

Ms. Erin Conaton, Under Secretary of the Air Force, 27 May 2010

Estimated U.S. Energy Use in 2009: ~94.6 Quads

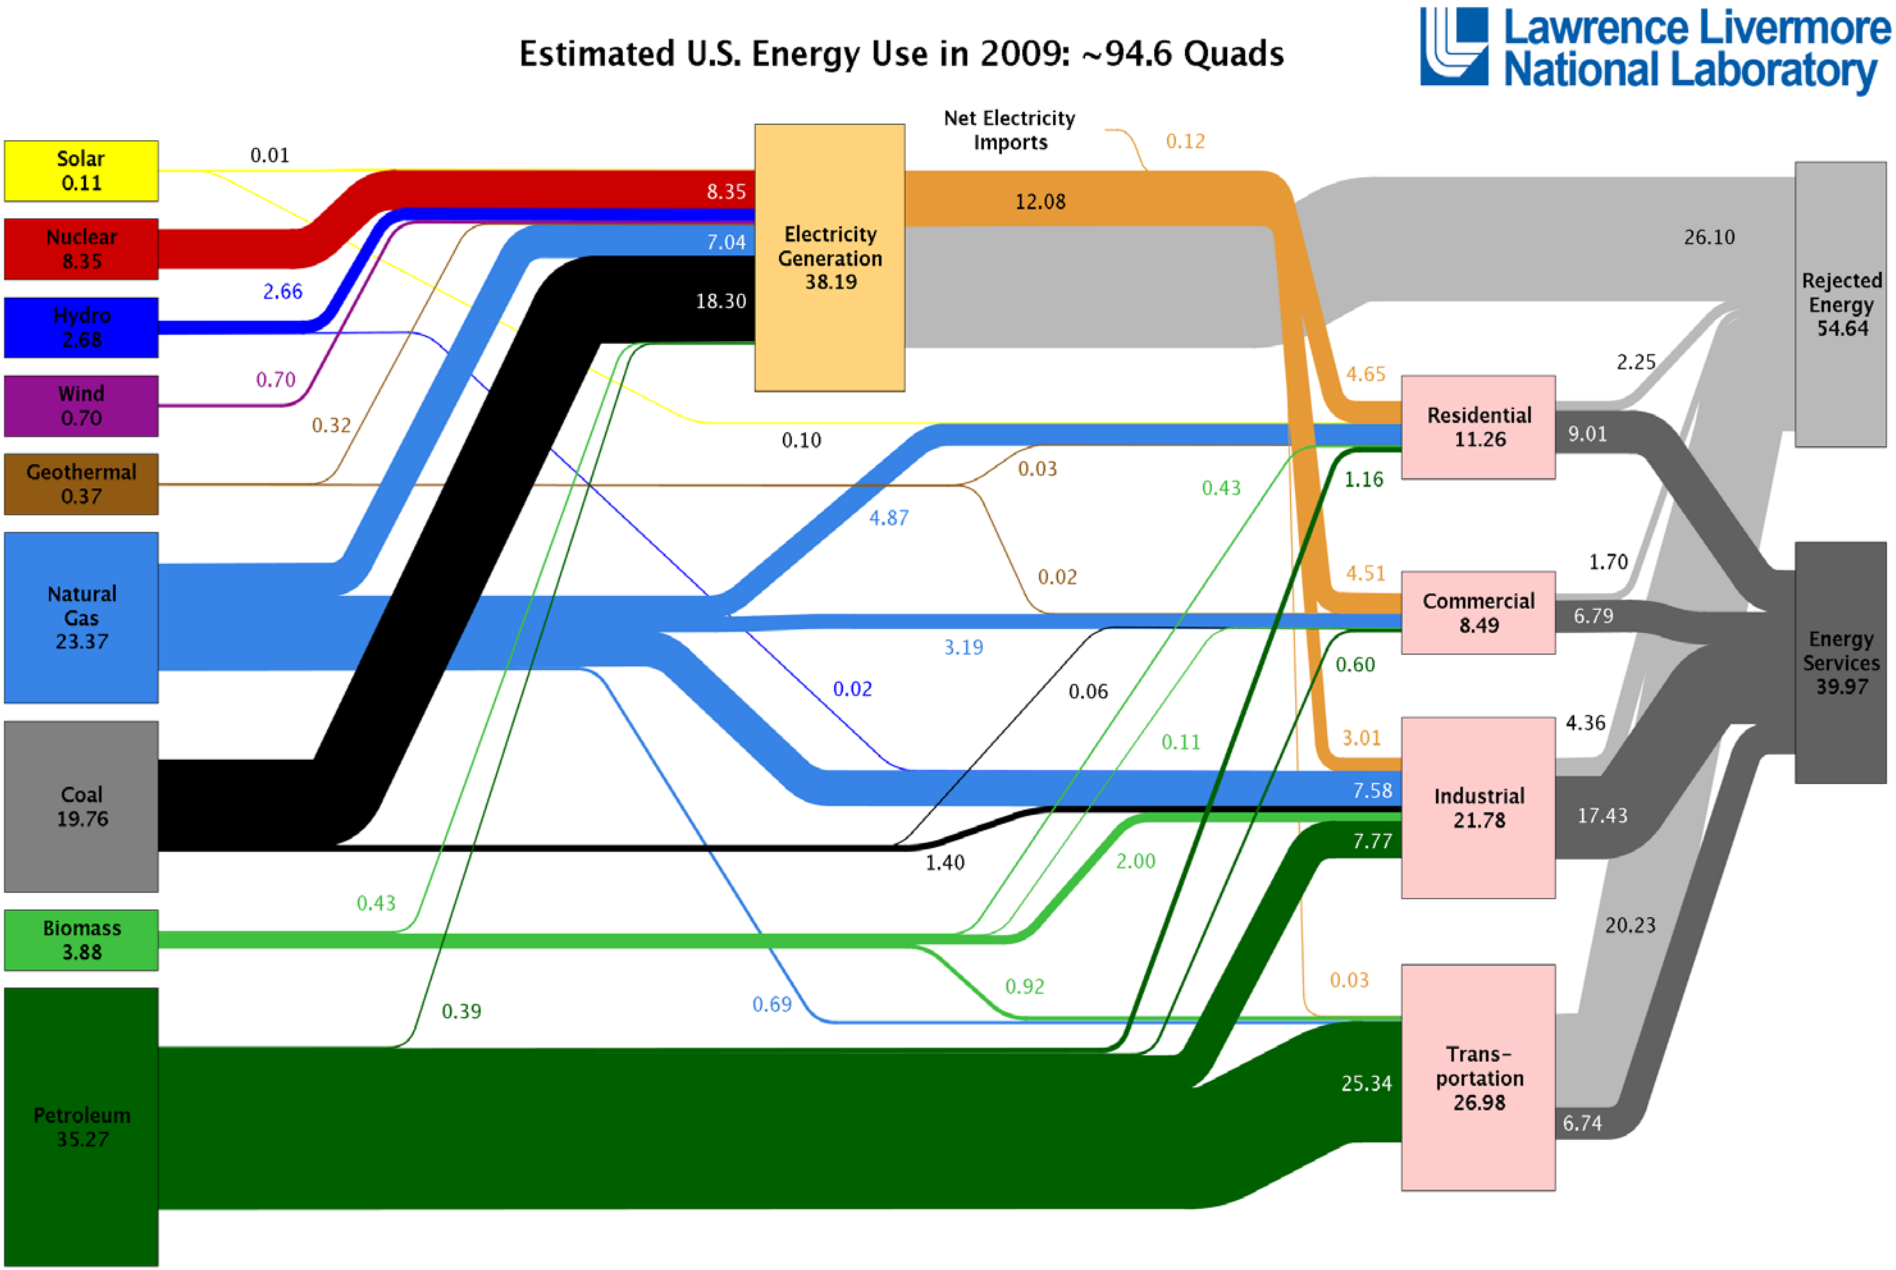

Source: LLNL 2010. Data is based on DOE/EIA-0384(2009), August 2010. If this information or a reproduction of it is used, credit must be given to the Lawrence Livermore National Laboratory and the Department of Energy, under whose auspices the work was performed. Distributed electricity represents only retail electricity sales and does not include self-generation. EIA reports flows for non-thermal resources (i.e., hydro, wind and solar) in BTU-equivalent values by assuming a typical fossil fuel plant "heat rate." The efficiency of electricity production is calculated as the total retail electricity delivered divided by the primary energy input into electricity generation. End use efficiency is estimated as 80% for the residential, commercial and industrial sectors, and as 25% for the transportation sector. Totals may not equal sum of components due to independent rounding. LLNL-MI-410527

RPA Growth

MQ-1, MQ-9 & RQ-4 = Class IV/V RPA

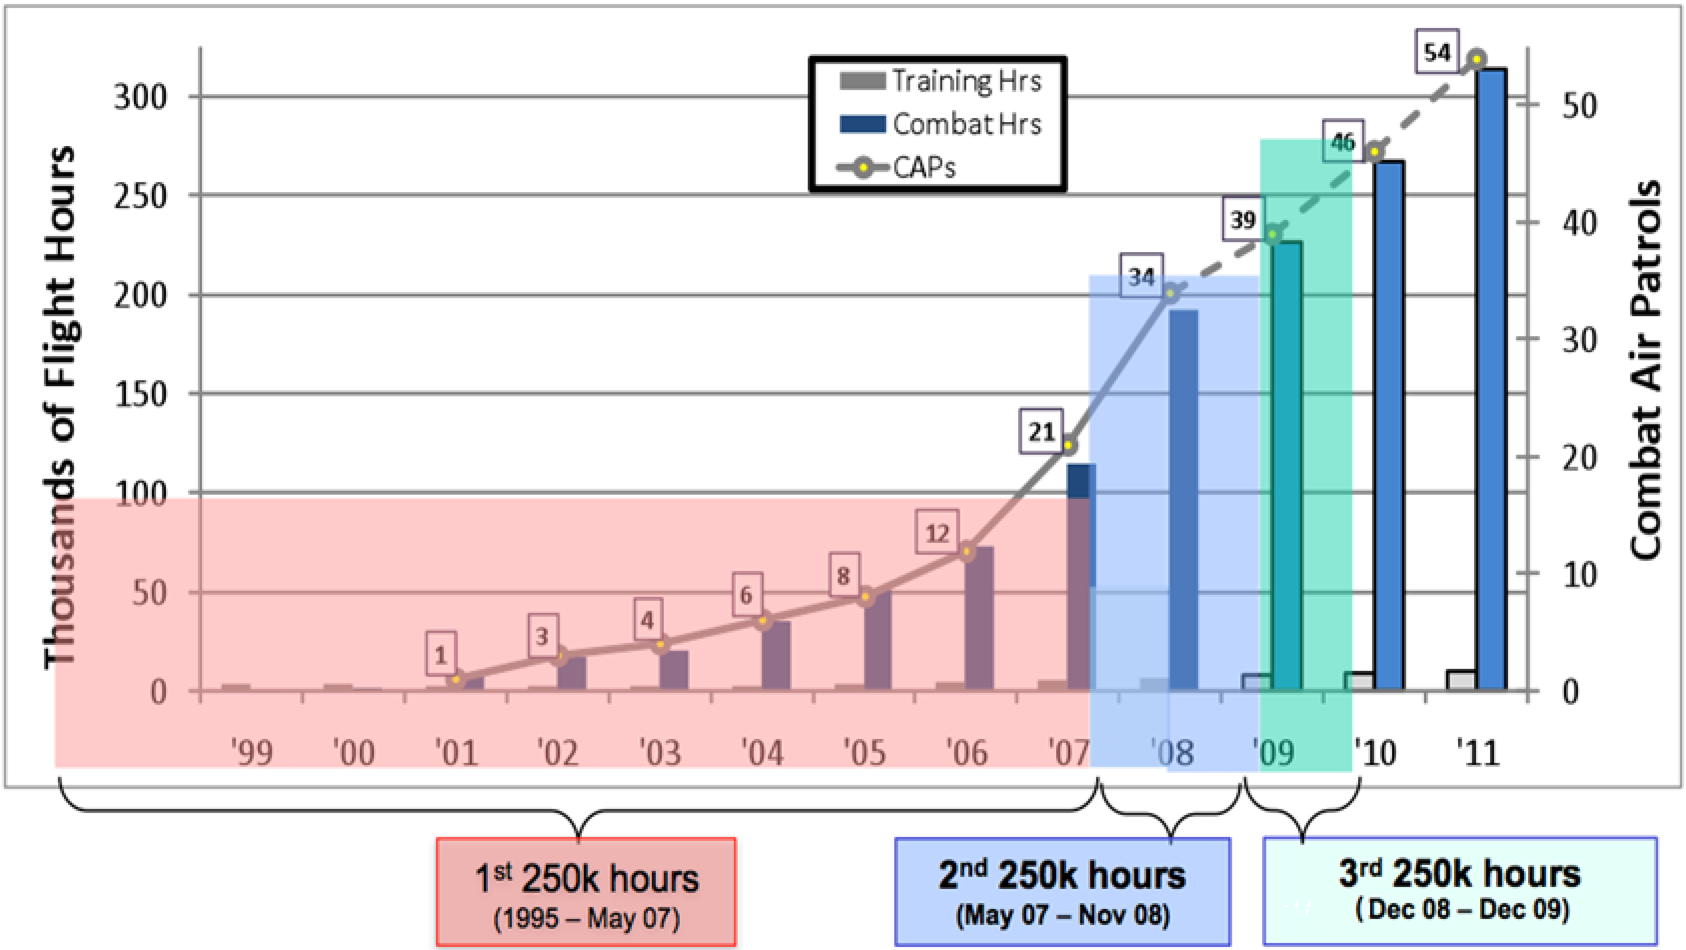

Multi-Spot EO/IR Sensors

- Multi-spot EO/IR cameras allow individually steered low frame rate spots; augment FMV
- Gorgon Stare now; ARGUS-IS will allow 65 spots using a 1.8 giga-pixel sensor at 15 Hz
- Individually controllable spot coverage goes directly to ROVER terminals on ground
- Autonomous Real-Time Ground Ubiquitous Surveillance - Imaging System (ARGUS-IS)

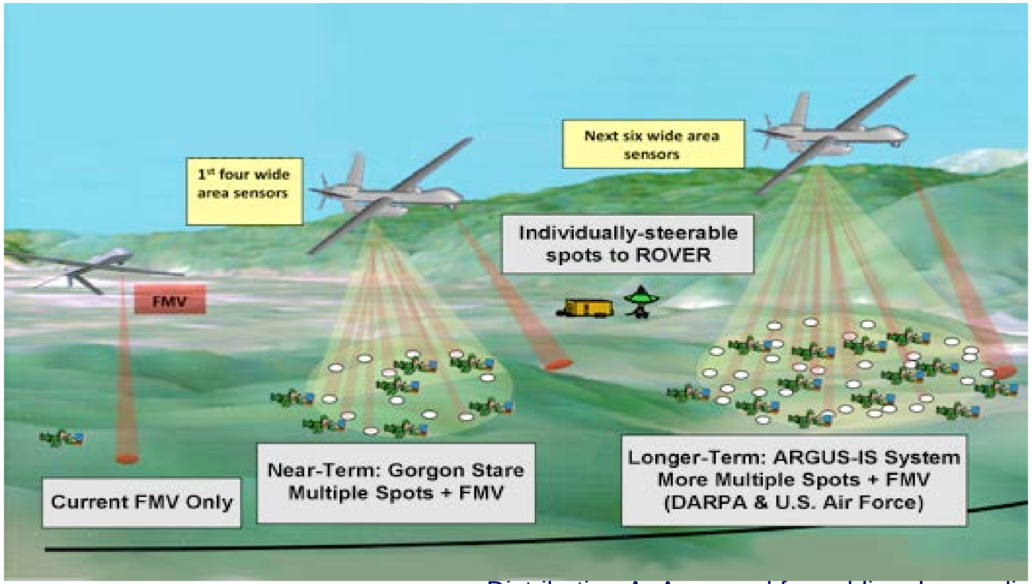

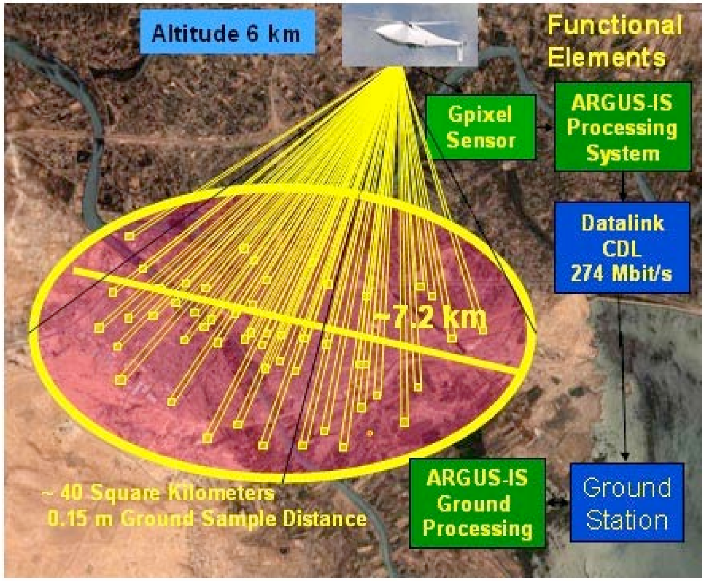

AF Energy Horizons S&T Vision

"Assured energy advantage across air, space, cyberspace and infrastructure"

- ***Assured*** **– We need to ensure operations in spite of vulnerabilities and (physically, economically, politically) contested environments**
- ***Energy*** **– its acquisition, storage, distribution, use**
- ***Advantage*** **– we seek an effectiveness and efficiency edge over our adversaries to ensure operational supremacy/dominance**
- ***Across*** **– we require supremacy within and across**
- ***Air, space, cyber, infrastructure*** **– we require full spectrum energy solutions**

Space Operational Energy Dominated by Ground Facilities

FY10 AFSPC Energy

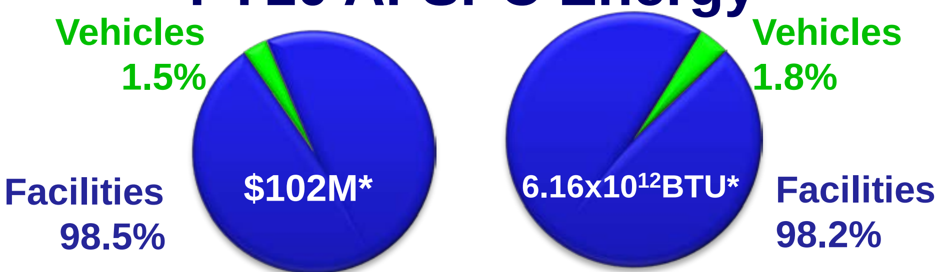

*Figures do not account for rocket fuel or satellite energy

Since FY03

- **Energy intensity decreased by 11.2%**
- **Energy consumption decreased 6.8%**
- **Total energy cost increased 22.8%**
- **Cost/MBTU increased by 28.4%**
- **Avoided energy cost in FY10 was $9.01M**

Energy

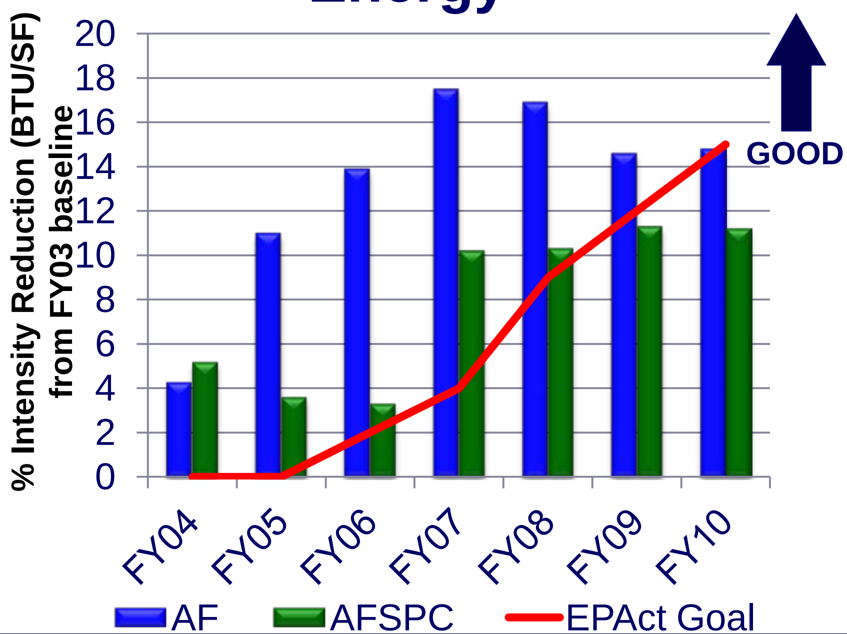

Water

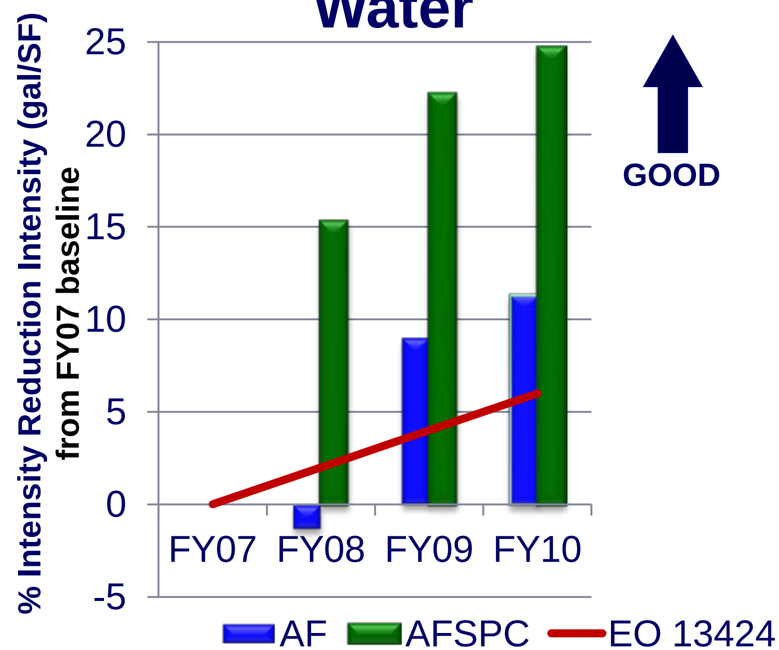

S&T (e.g., solar cell efficiency, on-orbit refueling) can revolutionize space power

DoD and AFRL Cyberspace Strategy Aligned

DEPARTMENT OF DEFENSE STRATEGY FOR OPERATING IN CYBERSPACE

Strategic Initiative 1: Treat cyberspace as an operational domain to organize, train, and equip so that DoD can take full advantage of cyberspace's potential

Strategic Initiative 2: Employ new defense operating concepts to protect DoD networks and systems

Strategic Initiative 3: Partner with other U.S. government departments and agencies and the private sector to enable a whole-of-government cybersecurity strategy

Strategic Initiative 4: Build robust relationships with U.S. allies and international partners to strengthen collective cybersecurity

Strategic Initiative 5: Leverage the nation's ingenuity through an exceptional cyber workforce and rapid technological innovation

Align with others and grow AF expertise

AFRL Strategic Cyber Thrusts

Assure and Empower the Mission

Optimize Human-Machine Systems

Enhance Agility and Resilience

Invent Foundations of Trust and Assurance

Partner with acquisition and operational communities

Cyber Energy

- Considerations include:
  - Reduce energy footprint, Performance/watt/$
  - Analytics, Automated processes
  - Focus on Facility Efficiency vs. Compute Efficiencies
  - Impact of Policy, acquisition strategy

- 5 May 2011: Cyber Energy S&T Summit: Industry/Academia Participants from Intel, HP, Microsoft, IBM, Virginia Tech, HPCMO, University at Albany
  - Preliminary Observations
    - 3D chip stack and memristor nanotech breakthroughs allow 10X density and 2X efficiency within 5 years
    - Cloud computing offers multiplicative efficiencies from scale, sharing, low cost energy supply, and architecture tailoring
    - 100X efficiency improvements possible by optimizing application to architecture mapping
    - Many cyber energy ideas benefit cyber security

Excessive Energy Required for Cyber Servers and Facilities

EXAMPLE: NNSA's IBM Sequoia "energy efficient" supercomputer does 3,050 calculations per watt of energy, requiring >6 MW of power (not including facility).
This equates to 6,000 gal/hour using diesel generators.

AF Cyber Energy must focus on Power Usage Effectiveness (PUE) and alternative energy sources

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

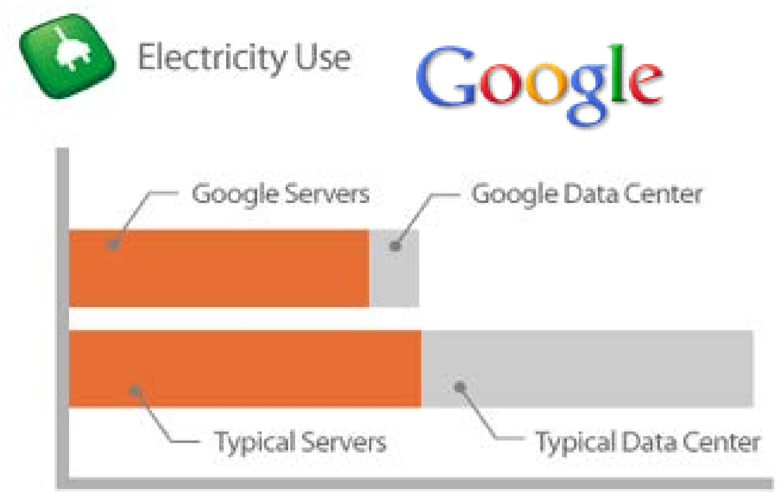

IBM, Google, others, follow PUE closely

Green Computing saves energy

Why build a GREEN Super Computer?

- **Meet AF demand for large computing power while reducing the environmental impact to our natural resources**
- **US EPA created the Energy Star program to create guidelines for energy-efficient devices**
  - **Reduce the amount of energy consumed by a product by automatically switching it into sleep mode or reducing power consumption when in standby mode.**
- **Reduce computer power requirements**
- **Reduce the stress on power-supply infrastructure**
- **Reduce room cooling requirements**
- **Reduce sustainment costs**
- **Computing Data Centers consume ~1.57% of the total electricity used on the planet**

Energy Star Logo

Save the Environment; Save Dollars

AFRL Using Sony PS3s to Build Green Super Computer (Condor Cluster)

- AFRLs Advanced Computing Architecture group is designing, integrating and building the largest interactive computer in the US Air Force.
- Using Sony Playstation3 (PS3) game consoles and other off-the-shelf commercially available components.
- 1716 PS3s and 156 high end graphics cards, known as General Purpose Graphic Processor Units (GPGPU).
- Costs to build and operate the Condor super computer are significantly lower than traditional super computers.
- Leverage built-in power saving sleep feature to lower the operational cost.
  - Intelligent Power Management (based on users needs, machine reservation)
- Power of super computers is measured in Floating Point Operations per Second (FLOPS)
- A typical household laptop can achieve approximately 10 billion FLOPS
- AFRLs Condor super computer will achieve 500 trillion FLOPS of processing power, equivalent to about 50,000 laptops.
  - 260KW Peak Power Consumption 1.92 GFLOPS/W and 1.05 GFLOPS/W sustained

Introducing Condor
DoD's Largest Interactive Supercomputer

- **Cost-effective**
  - **Low-cost commercial parts**
  - **10X price-performance advantage**
- **A "green" supercomputer**
  - **Reduced power consumption**
  - **15X power-performance advantage**
- **Military applications - examples**
  - **Real-time radar processing for high resolution images**
  - **Embedded HPC experiments**
  - **Computational intelligence**
  - **Space situational awareness**

A DoD High Performance Computing Modernization Program (HPCMP) Affiliated Research Center

AFRL
Joint Capability Technology Demonstration
Energy Efficient Computing (E2C)

Escalating energy costs and rapidly expanding demands on computational capacity and power are straining the ability of military data centers to deliver cost effective support to the warfighter in an environmentally sound manner.

- **Exorbitant cost associated with DoD computing; $2.5M per year alone for Maui High Performance Computing Center (MHPCC)**
  - **Inefficient electrical system - switching from DC to AC multiple times**
  - **Inefficient facility – cooling, lighting, etc.**
  - **Inefficient hardware – needs upgrade to state-of- the-art Energy Star**
  - **Inefficient software – no attention to energy in scheduling jobs**
- **Unable to expand existing computing capability under fixed energy budget and increasing computing requirements**
  - **Retrofits required for most DoD data centers**
- **Unable to manage data center power holistically**
  - **Facility energy management system not linked to IT infrastructure**
  - **Facility infrastructure unable to keep up with IT tech advances**
  - **Inability to actively manage power consumption and performance output**

"Data centers consume up to 100 times more energy than a standard office building." U.S. Department of Energy, May 2009

Reference: **Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431,** U.S. Environmental Protection Agency, August 2, 2007

E2C JCTD Operational View

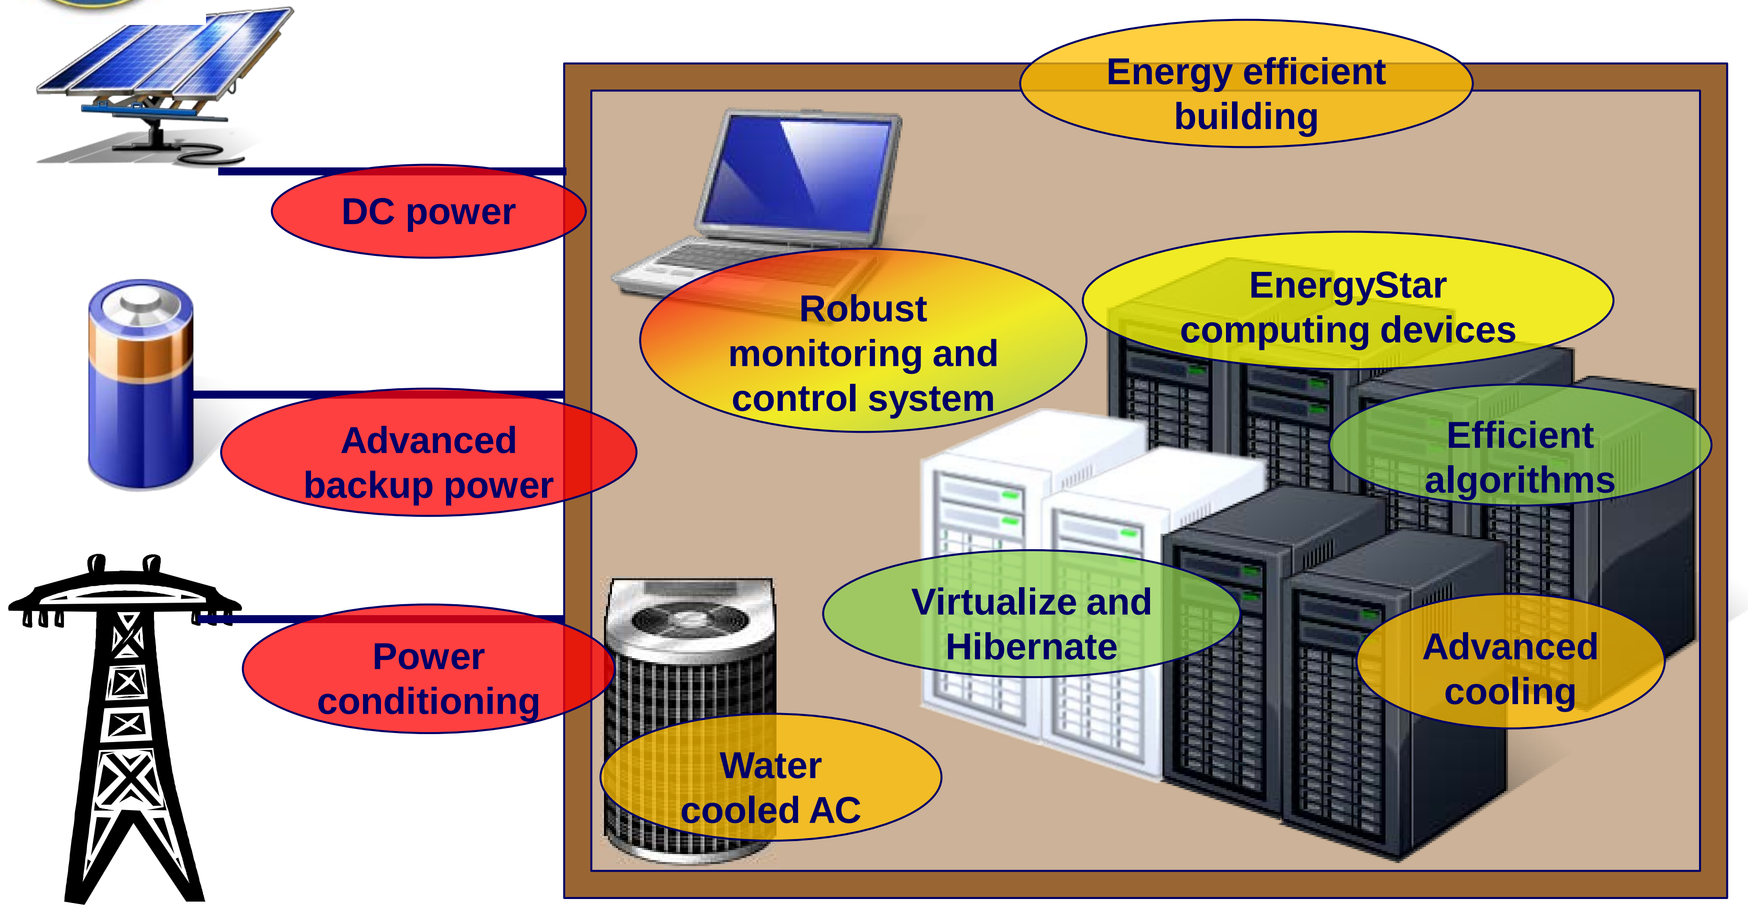

Electrical System	Facility	Hardware	Software
Efficient power transmission	Minimize non-computing energy demand	Efficient computing devices	Energy-aware software

E2C JCTD

DC Power Efficiency Example

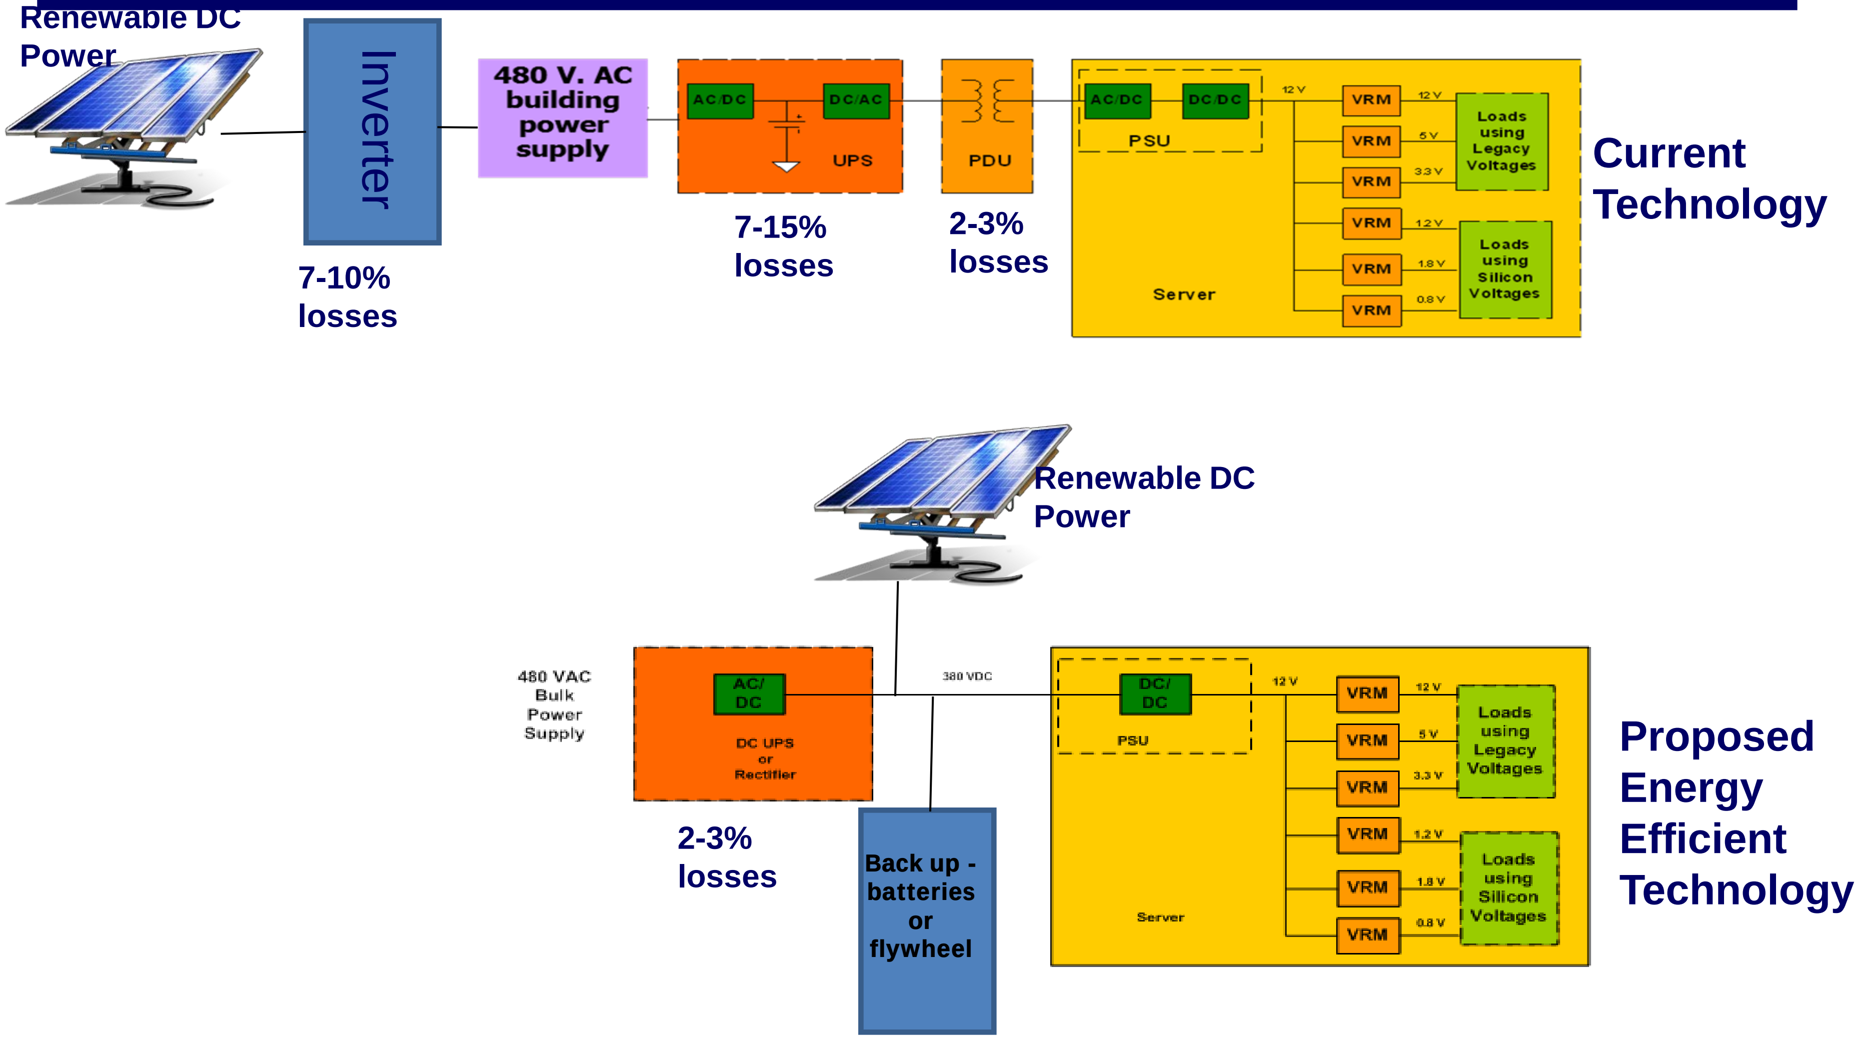

MIT-Lincoln Lab TX-Green High Performance Computer

- Department of Defense High Performance Computing Modernization Program DHPI project
- Exploring revolutionary supercomputing management approaches to dynamically manage energy consumption
- Each compute node contains:
  - Dual 2.2 GHz, 12-core Magny-Cours processors
  - 64 GB RAM
  - 10 GB ethernet interface
  - 2 TB disks
- Software includes:
  - LL Grid interactive cloud computing
  - Cluster provisioning and management
  - Central shared file storage – serves 12 PB of hi-speed storage
- Applications include developing novel sensor processing algorithms

TX-Green System Specification

- 6,000 Processing Cores
- 53 TeraFLOPs
- 16 TB RAM
- 10 Parallel Storage and Management Nodes
- Non-Blocking 10 GbE Switching
- 2.2 PB Storage
  - 1.0 PB Internal
  - 1.2 PB DDN SFA10000 (320 Gbps)
- 20' Portable Data Center (Shipping Container)

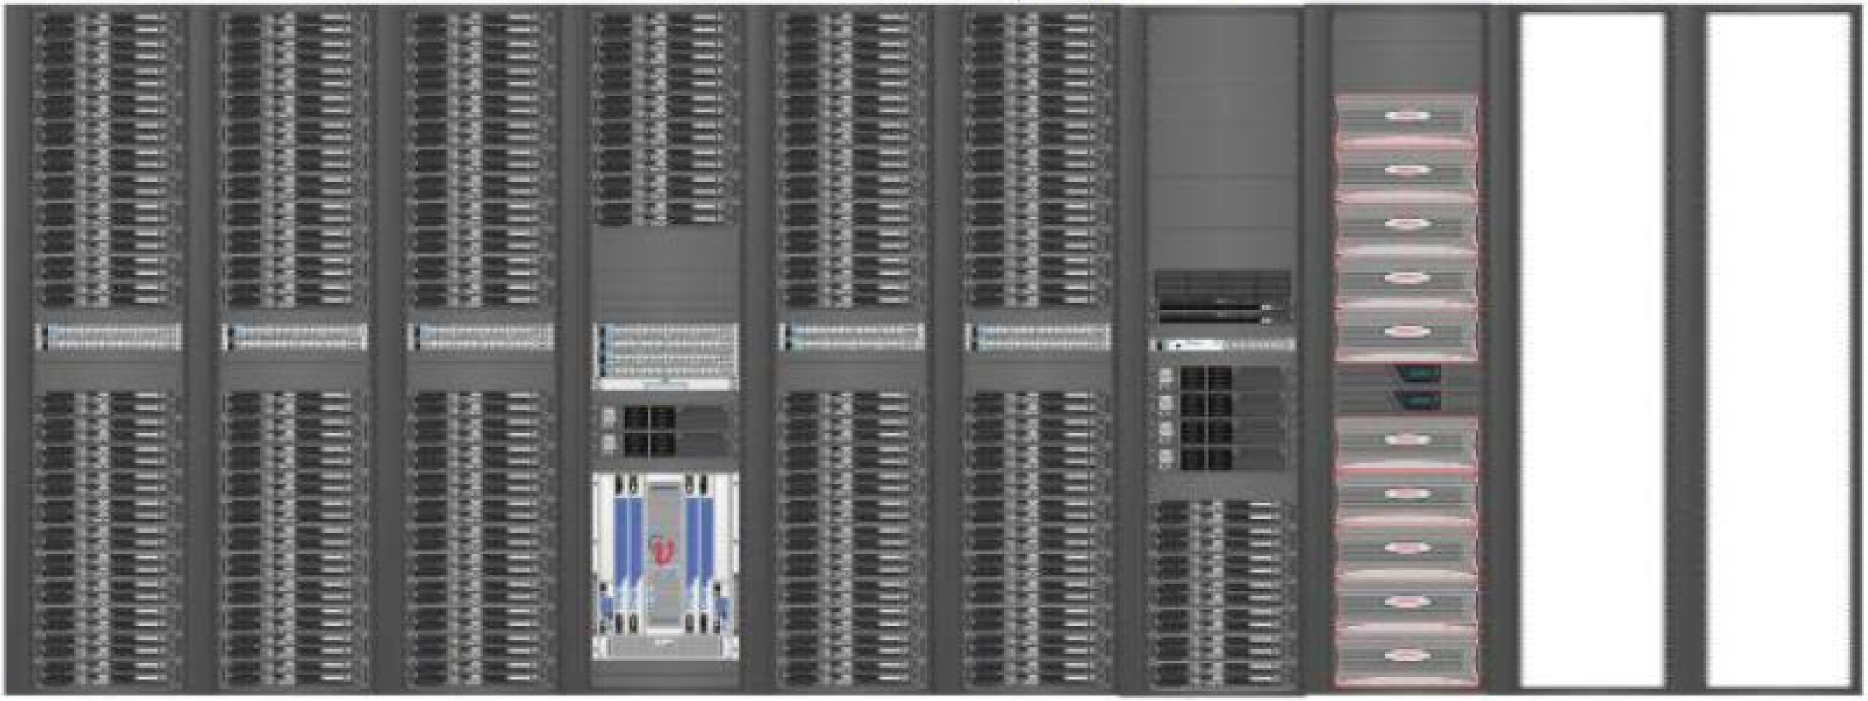

HPC and 3D Background

- **Systems such as Condor and Blue Gene are State of the Art for large clusters**
- **JSTARS uses 15 embedded parallel processors for radar images**
  - **X86 Legacy systems have fundamental heat and latency limitations to Size Weight and Power (SWAP), and parallelism**
  - **3D stacking will solve part of the latency roadblocks but heat is still problematic for further improvements in legacy system SWAP**
  - **Current 3D: Current commercial approach is memory on memory, research is ramping for heterogeneous systems. Latency reducing logic on logic is >3 years out. Intel's is researching 80 Atom-scale-cores on a chip. It uses 62W / chip @ 1TFLOP. Too hot to stack.**

Current AF needs point towards concurrent SWAP trades that the commercial world is not yet motivated to make for Big Iron and Smart Phones.

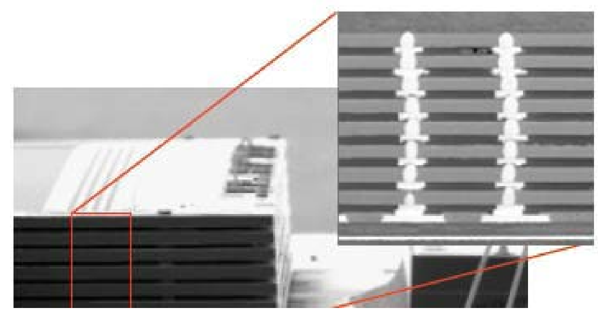

Samsung Stacked DDR3 RDRAM

AFRL Chip Level Energy Savings

Chip Level Energy Savings:
Low power processor
- **ALU @ 50 GFLOPS/W**
- **Gated Clock**
- **Shallow Pipeline**
- **SOC integration: AES, AFPGA…**
- **Instruction set reduces clocks / function**

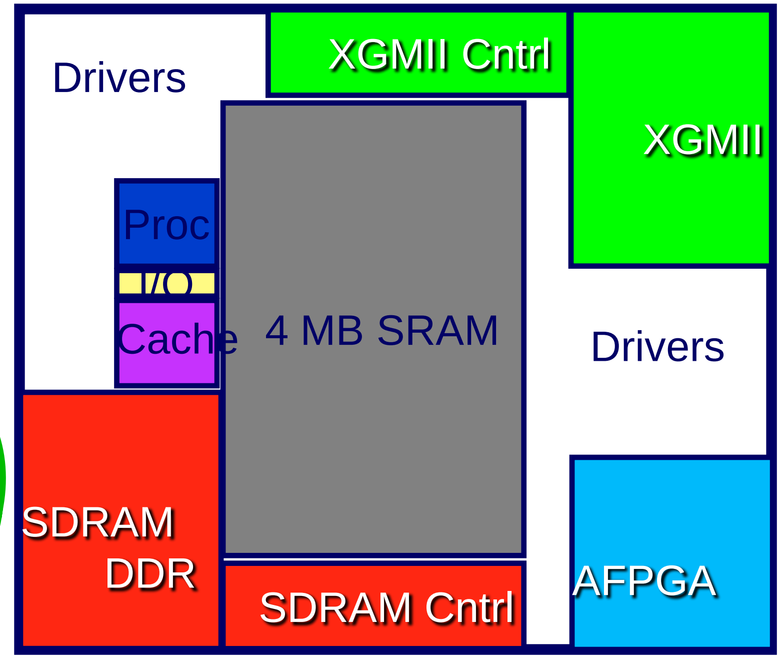

AFRL CyberCog Core V1.0

Low processor energy enables 3D

Cluster Level Savings, 3D stacking:
- **Reduces Pin and Pin Driver count**
- **Reduces Wire length**
- **Reduces order of off cluster information**
- **Reduces chip cooling**

Commercial Example
Tezzaron: DDR3 ~40mW/pin
2G -> 1024 Data Pins = 40 W
= 160W for pins on 8G DDR3 RAM

Die on Wafer interconnect ~24 μW / pin
= 20.9W for 3D 8G DDR3 RAM

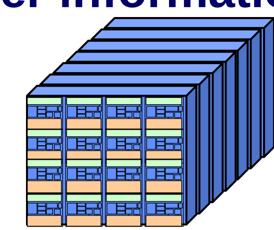

3D Stacked Cluster of CyberCog Chips

System Savings

System Level Energy Savings

Energy savings at the core level enable more savings as we go from processor to cluster to system

IBM Blue Gene/L System Architecture

<3 GFLOPS / W

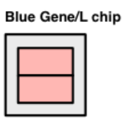

Chip

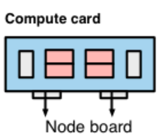

Node

16GFLOPS / W

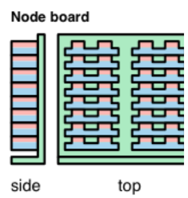

Board

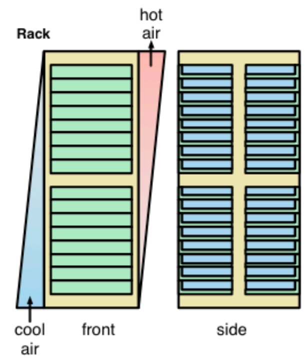

Rack

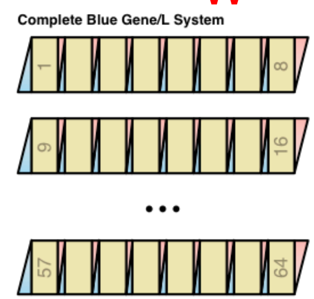

System

3X

2X
Pins
SOC

1.25X
Power
Supplies

1.5X
Network
Overhead

2X
Cooling
Efficiency

Core

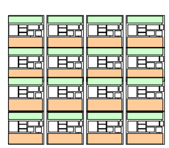

Chip

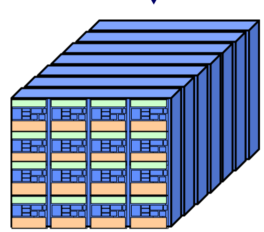

3D Stack / 800 Cores

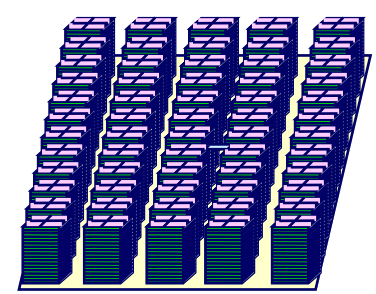

Board / Many Stacks

Effects Level Multiplier

The greatest energy savings to be had are at the Effect level E.g. - processing and resulting autonomy can be a big force multiplier for precision strike systems

- **More autonomous systems => more intelligent effect at the tip of the spear, fewer personnel - with their logistics**
- **Trusted localized intelligence => data to information reduction => smaller communication packages**
- **More capable systems => fewer separate systems to do the job => better utilization, reducing standby power use**
- **Best for Last: Improved planning enabled by more capable decision tools => fewer energy consuming resources needed in the first place**

Conclusion

- **Energy efficient HPC is an operational imperative**
- **Assured cyber energy advantage can**
  - **Conserve resources and protect the environment**
  - **Provide economic efficiencies**
  - **Achieve energy goals and congressional mandates**
  - **Enhance operations (e.g., efficiency, performance, resilience)**
- **Revolutionary advances in Cyber Energy S&T:**
  - **Will demand a focused roadmap and targeted investment**
  - **Promise transformational operations**

Energy Horizons **Vision:**
Assured Energy Advantage
across Air, Space, Cyber and Infrastructure

Questions?

Energy Scaling: BioInspiration

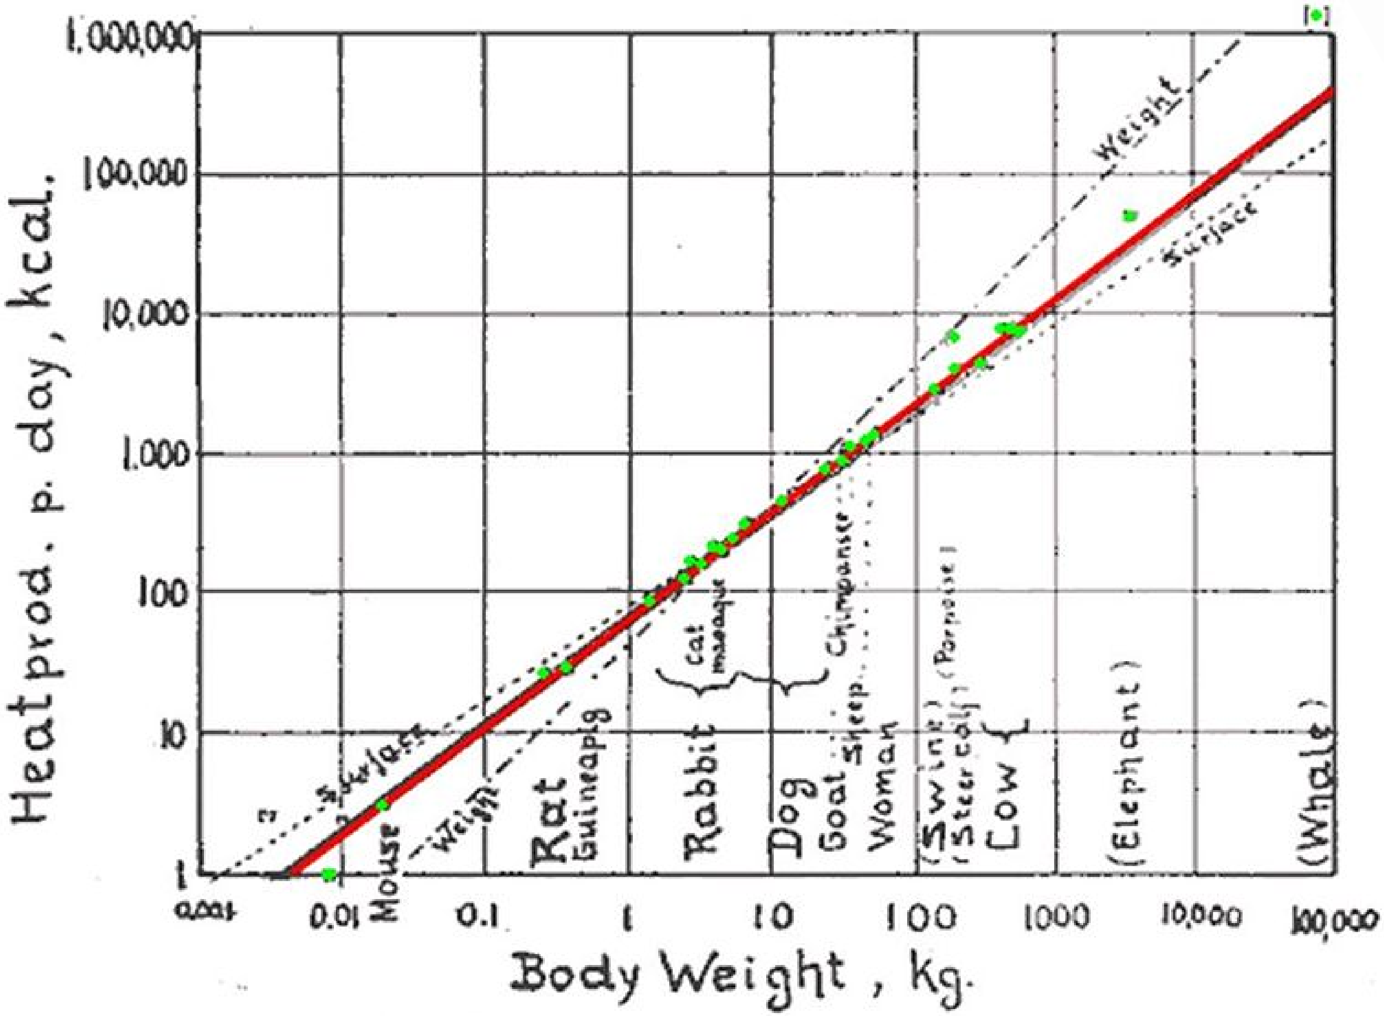

Fig. 1. Log. metabol. rate/log body weight

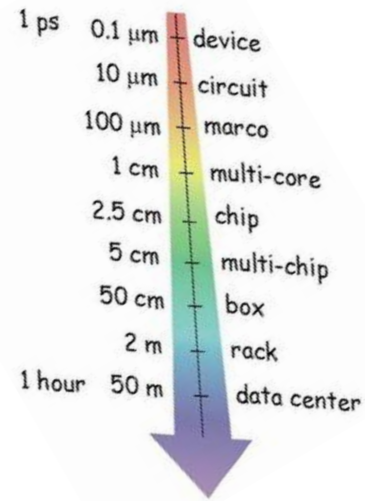

$q_0 \sim M^{3/4}$

Biologist Max Kleiber, early 1930s

en.wikipedia.org/wiki/Kleiber's_law

Energy and Power Density

★ **Human metabolism** ★ **Hummingbird metabolism**

Electric Ray: ***Torpedo Marmorata;***
Total stored energy in EOD: 38 kWhr (135MJ)
Power: >10^5 W

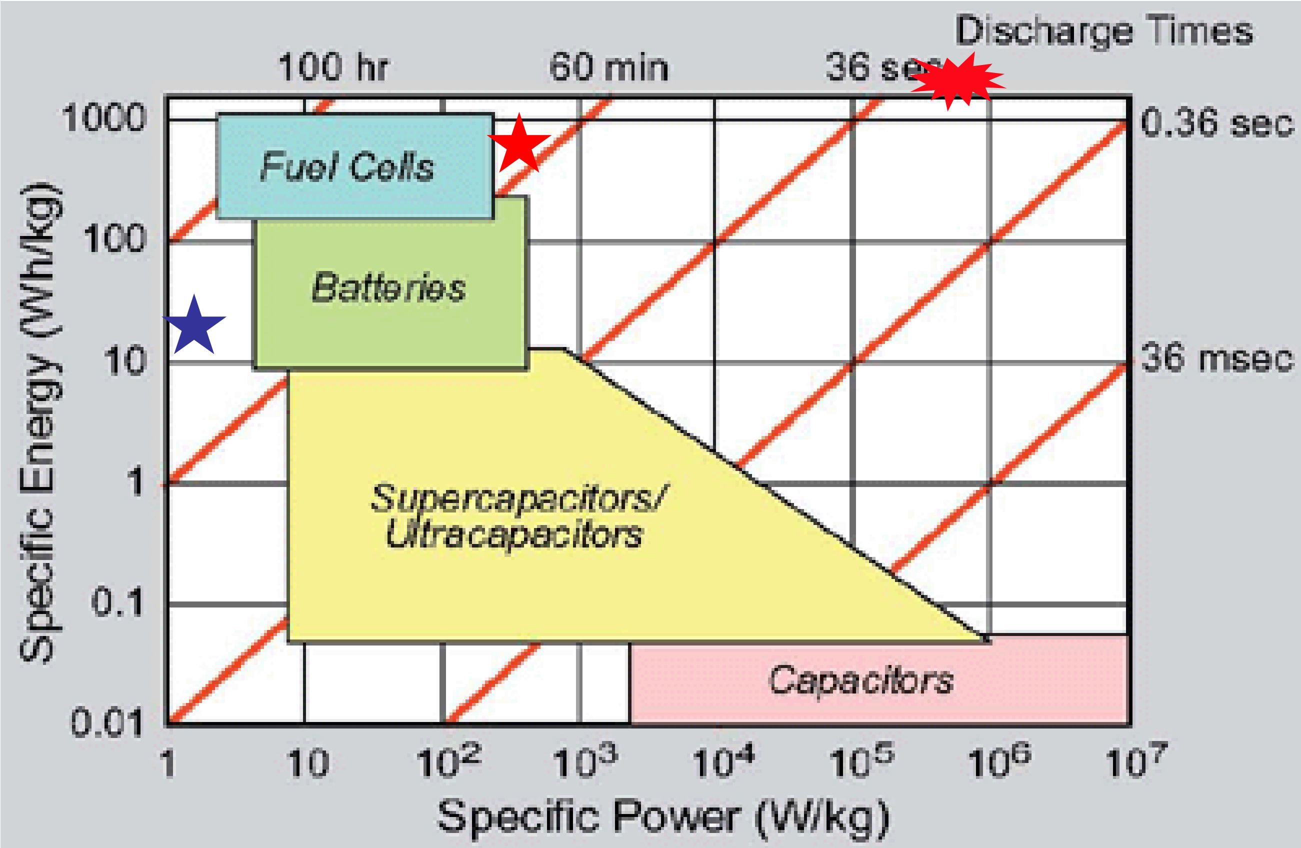

A Ragone plot comparing relative energy storage delivery performance for the state-of-the-art man-made energy-storage devices, relative to the energy-storage performance achieved by Mother Nature with the hummingbird, the torpedo ray, and with human metabolism. Red lines indicate times for complete discharge of stored energy.