Linux, Energy, and Networks:

"Saving Large Amounts of Energy With Network Connectivity Proxying"

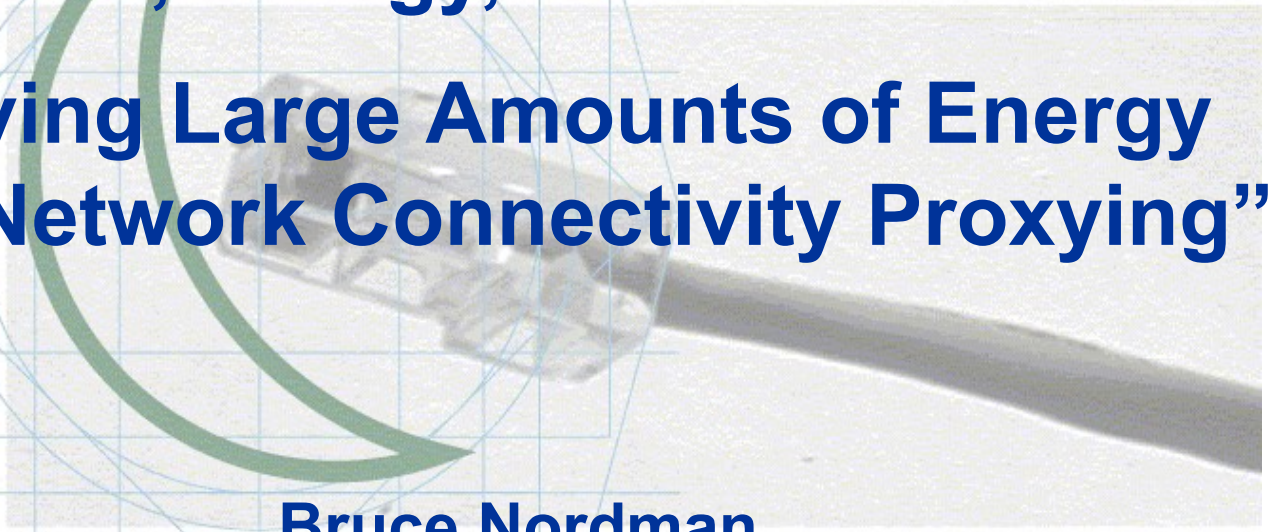

Bruce Nordman

Lawrence Berkeley National Laboratory

April 9, 2009

BNordman@LBL.gov — efficientnetworks.LBL.gov

Linux and Energy Efficiency*

Five basic dimensions of topic

- Efficiency of computing - a lot
- Efficiency of computing - a little
- **Efficiency of doing no computing**
- Effectiveness of communicating with user
- Efficiency imposed on other devices (via network)

**Efficiency not in traditional physics terms*

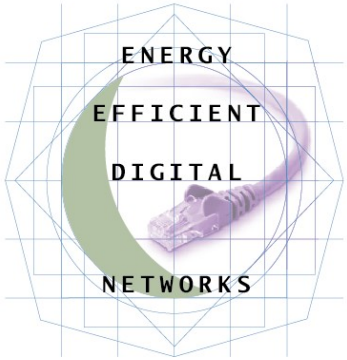

Need to Think Broadly about Networks

While some integrators are skeptical about the prewired, preprogrammed NHS rack from Sony, others embrace the solution for its simplicity.

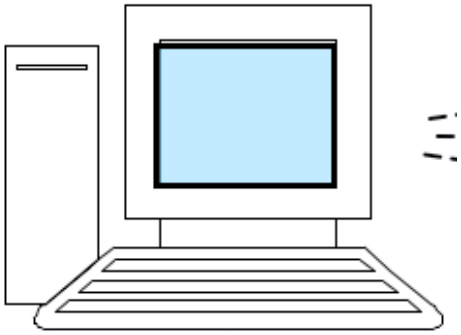

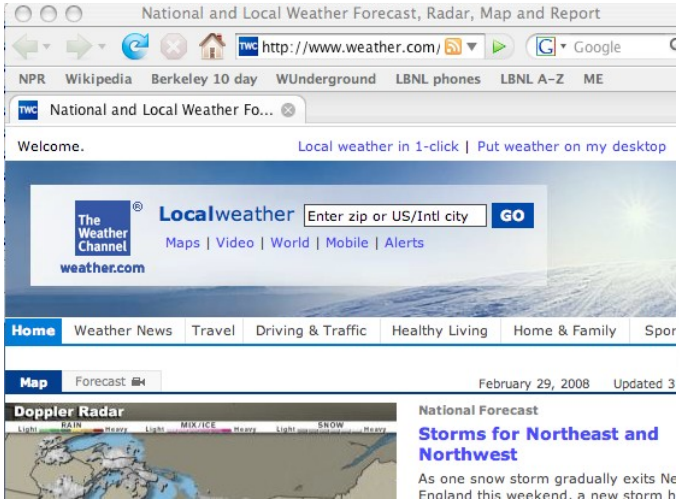

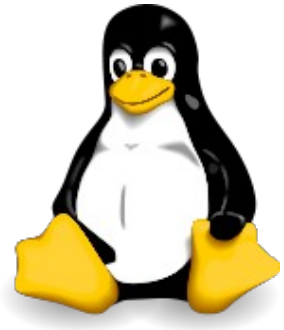

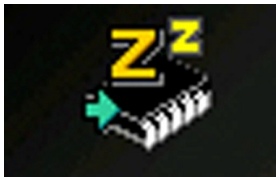

ENERGY EFFICIENT DIGITAL NETWORKS

How much energy does The Internet use?


Forbes.com
U.S. EUROPE ASIA
Search: Forbes.com Quotes Video
HOME PAGE FOR THE WORLD'S BUSINESS
HOME BUSINESS TECH MARKETS ENTREPRENEURS LEADERSHIP
Video Blogs E-mail Newsletters Org Chart Wiki People Tracker Portfolio
E-mail | Print | Comments | Request Reprints | E-Mail Newsletters | My Yahoo! | RSS

Dig more coal -- the PCs are coming

Peter W. Huber and Mark P. Mills, 05.31.99

Southern California Edison, meet Amazon.com. Somewhere in America, a lump of coal is burned every time a book is ordered on-line.

The current fuel-economy rating: about a pound of coal to create, package, store and move 2 megabytes of data. The digital age, it turns out, is very energy-intensive. The Internet may someday save us bricks, mortar and catalog paper, but it is burning up an awful lot of fossil fuel in the process.

1999

"At least 100 million nodes on the Internet, … add up to … **8% of total U.S. demand**. … It's now reasonable to project that **half of the electric grid** will be powering the digital- Internet economy within the next decade."

emphasis added


Slashdot
THE NUTS AND VOLTS OF NEWS FOR NERDS.
Log In | Create Account | Subscribe | Firehose

Sections
Main

Internet Uses 9.4% of Electricity In the US

2007


Uclue
got questions? (beta)
Home | Ask a Question | Categories | Browse Questions
★★★★★ ANSWERED on Fri 17 Aug 2007 - 6:17 pm MDT by davidsarokin
Question: Energy Use of Internet

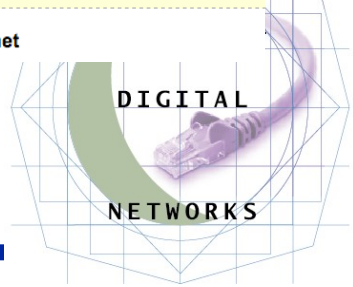

How much energy does The Internet use?

1999

"At least 100 million nodes on the Internet, … add up to … **8% of total U.S. demand**. … It's now reasonable to project that **half of the electric grid** will be powering the digital Internet economy within the next decade."

emphasis added

Wrong Question

Wrong Answers

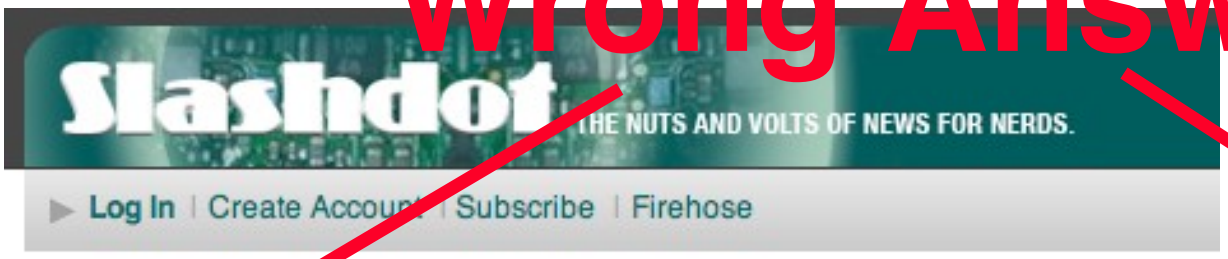

Internet Uses 9.4% of Electricity In the US

2007

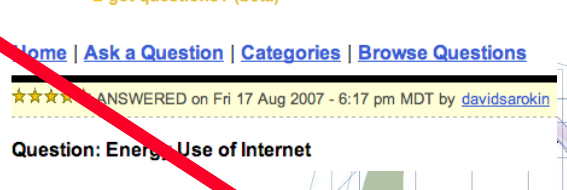

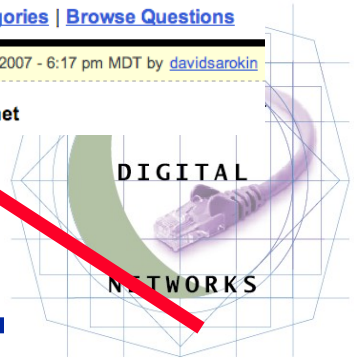

Network Structure and Energy

- Edge devices: PCs, servers - Displays, storage, phones, …

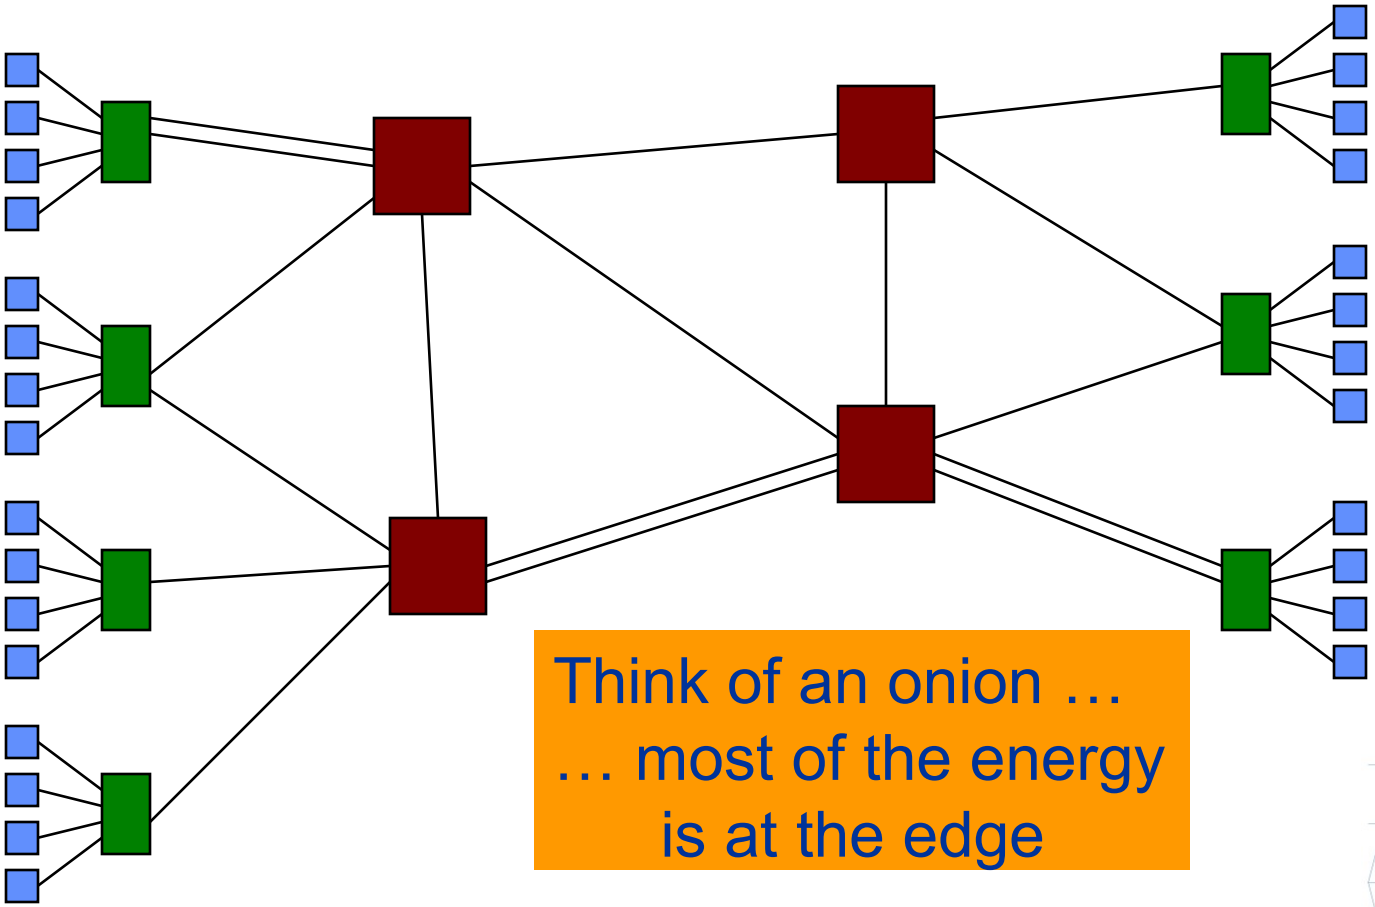

Think of an onion …
… most of the energy
is at the edge

- Network equipment: switches, and routers

Networks and Energy

Network equipment ….

Routers, switches, modems, wireless APs, …

… vs **networked** equipment

PCs, printers, set-top boxes, …

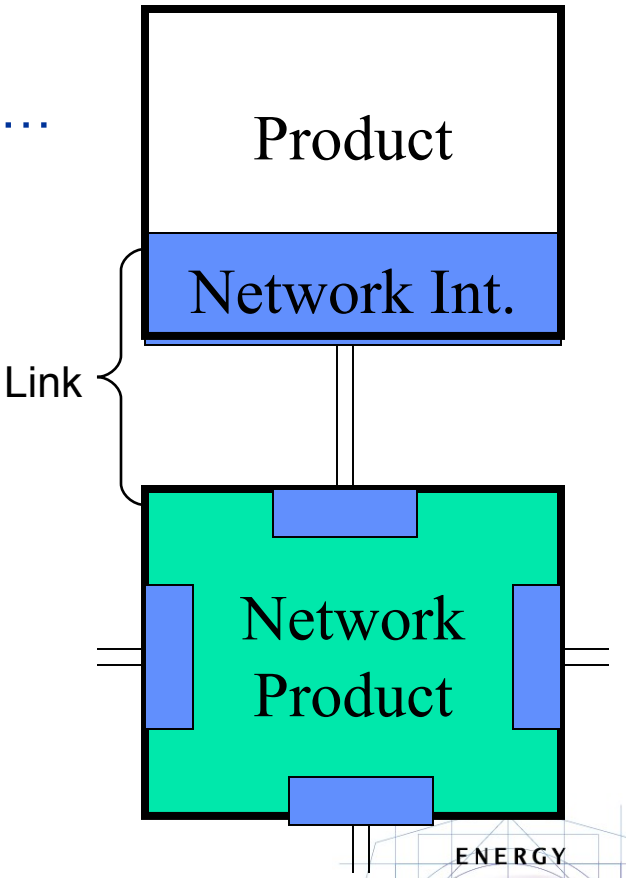

How networks drive energy use

- **Direct**
  - Network interfaces (NICs)
  - Network products
- **Induced** in Networked products
  - Increased power levels
  - Increased time in higher power modes (to maintain network presence)

Network electricity use in context

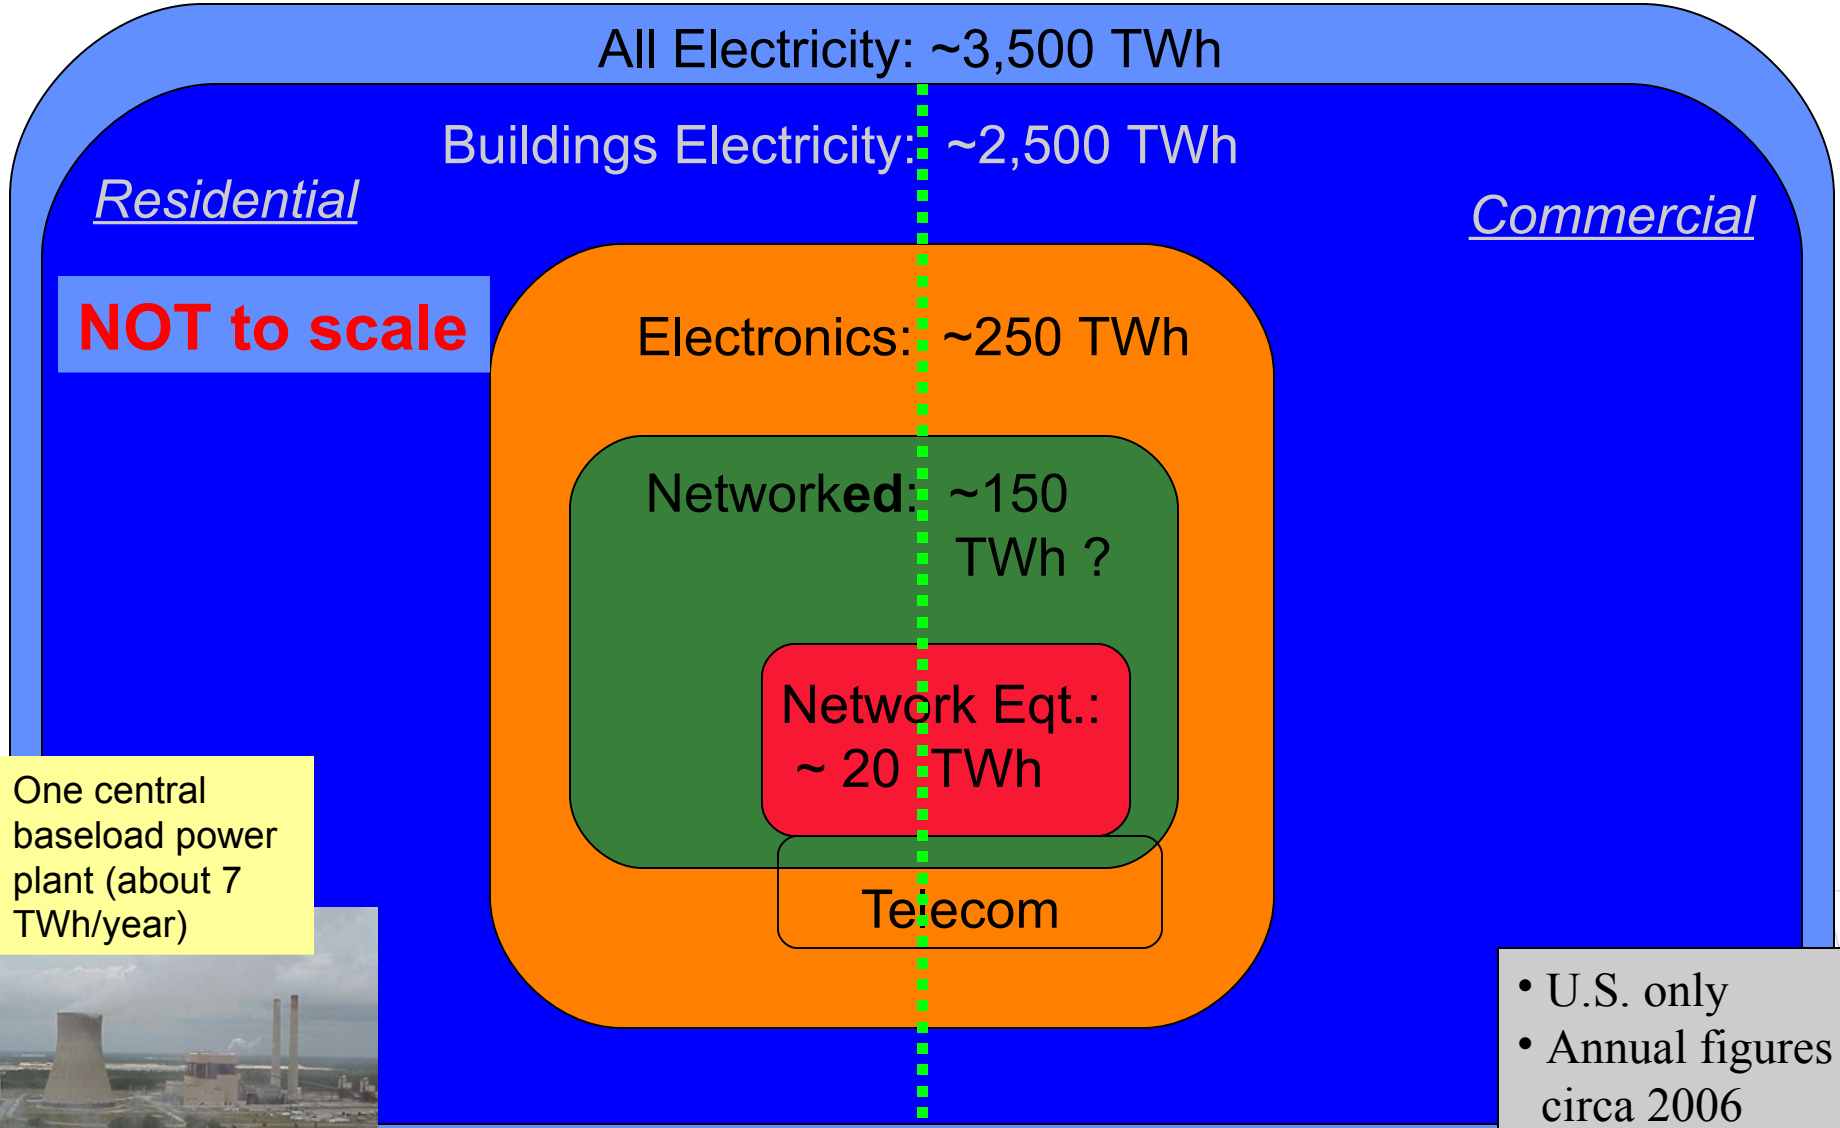

One central baseload power plant (about 7 TWh/year)

- U.S. only
- Annual figures circa 2006
- All approximate

Network electricity use in context, cont.

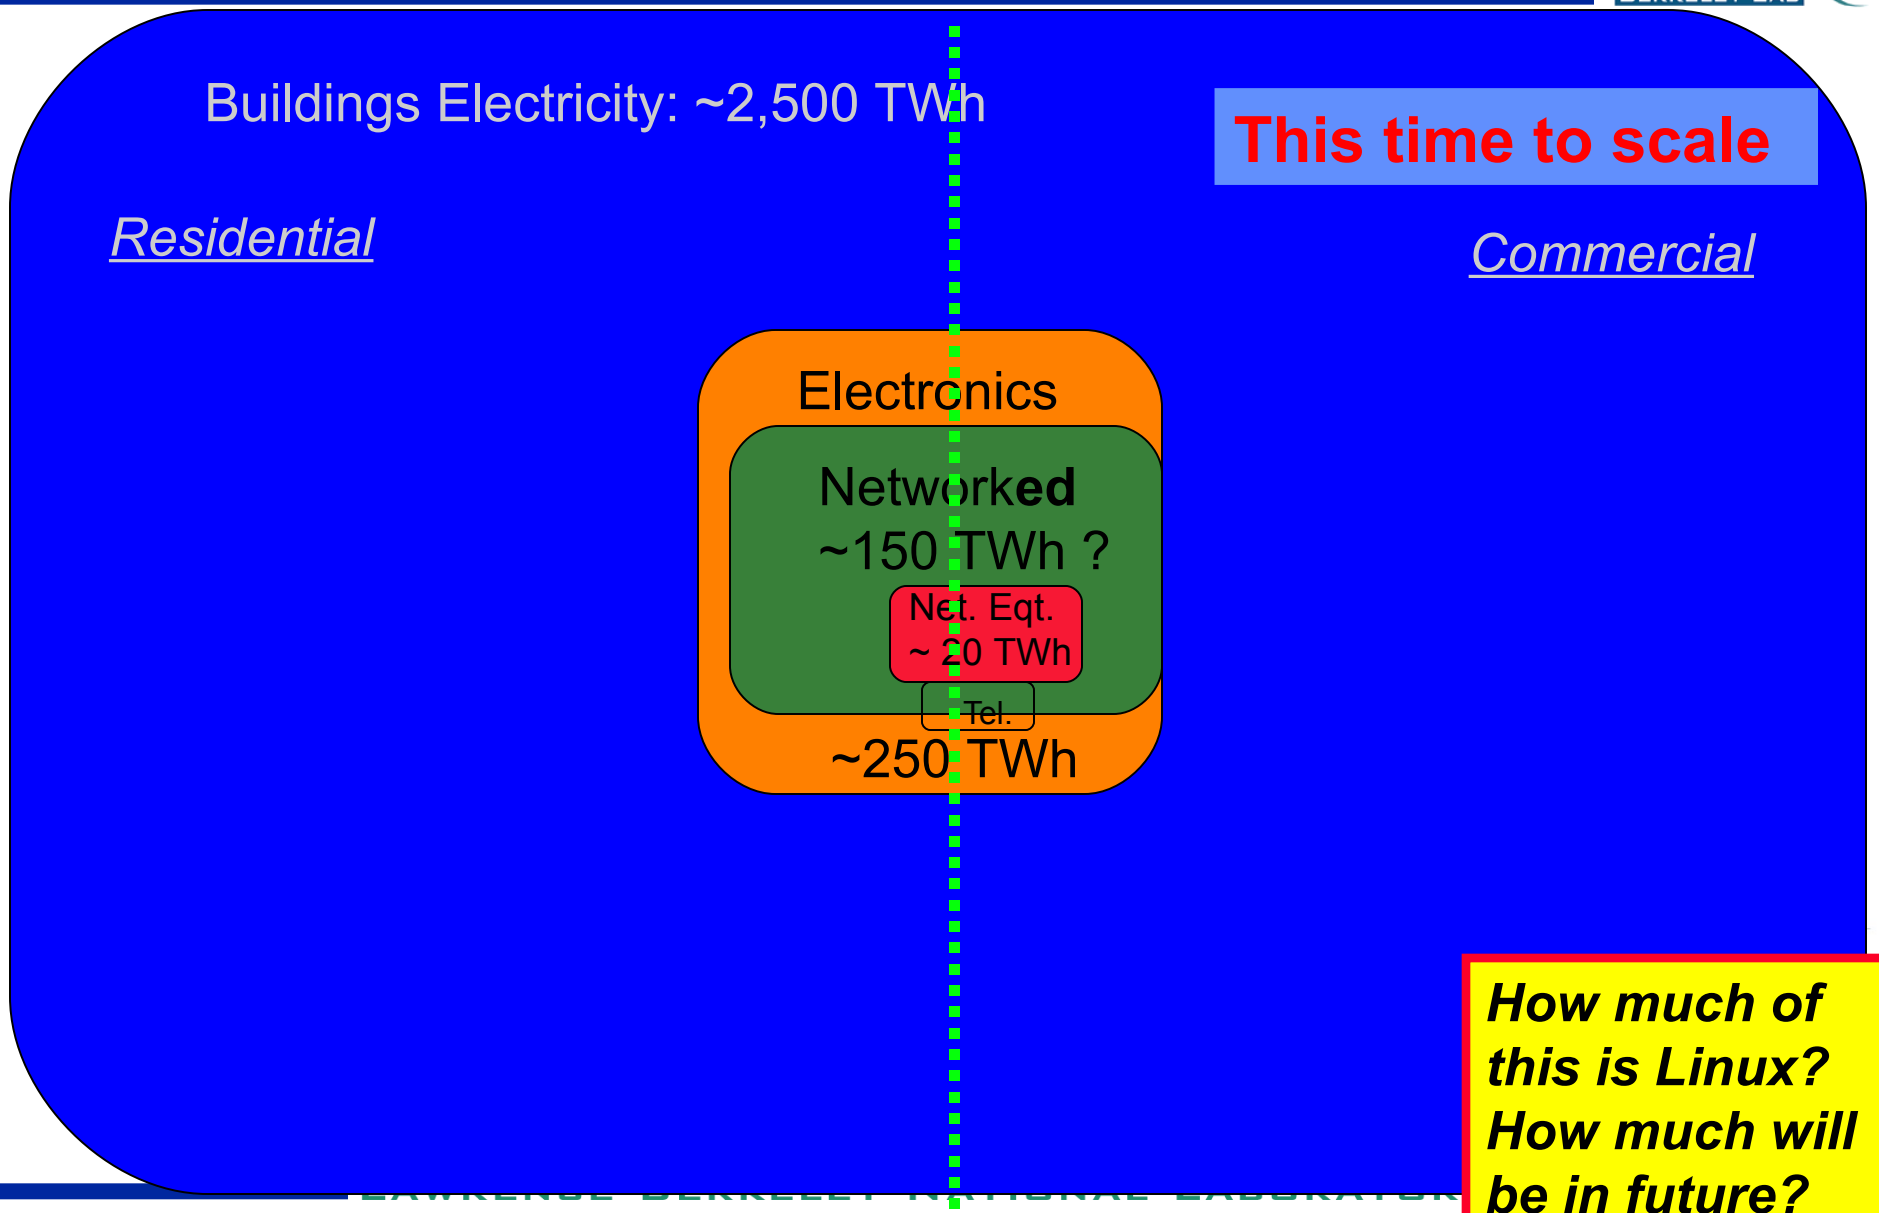

How to think about energy quantities

Our needs only require approximations

1 year = 8,760 hours	~ 10,000 hours
1 kWh costs $0.09	~ $0.10
1 W for 1 year	~ $1
1 TWh = 1 billion kWh	~ $100 million
U.S. annual consumption	~ 3,500 TWh
… buildings portion	~ 2,500 TWh

Things we know:
Energy consumption is at edge

- Network equipment < 10% of all electronics
- Most electronics already networked
- More electronic — and non-electronic — devices getting networked
- Network *induced* consumption > all direct
- Network equipment energy will grow …
  … but other electronics will grow faster

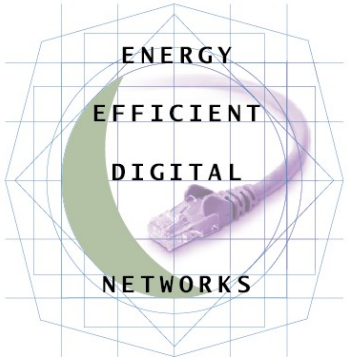

Things we know:
Utilization is low

Data networks are lightly utilized, and will stay that way, A. M. Odlyzko, *Review of Network Economics*, 2003

Network	Utilization
AT&T switched voice	33%
Internet backbones	15%
Private line networks	3~5%
LANs	1%

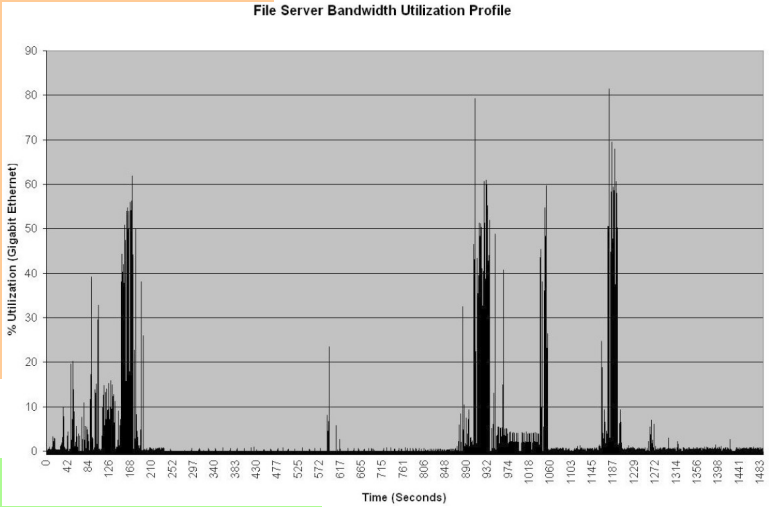

Low utilization is norm in life — e.g. cars

- Average U.S. car ~12,000 miles/year = 1.5 miles/hour
- If capacity is 75 mph, [illegible]

PC and server utilization also low

ENERGY EFFICIENT DIGITAL NETWORKS

Things we know:
Edge device energy is mostly idle

Core Fact: Most PC energy use occurs when no one present

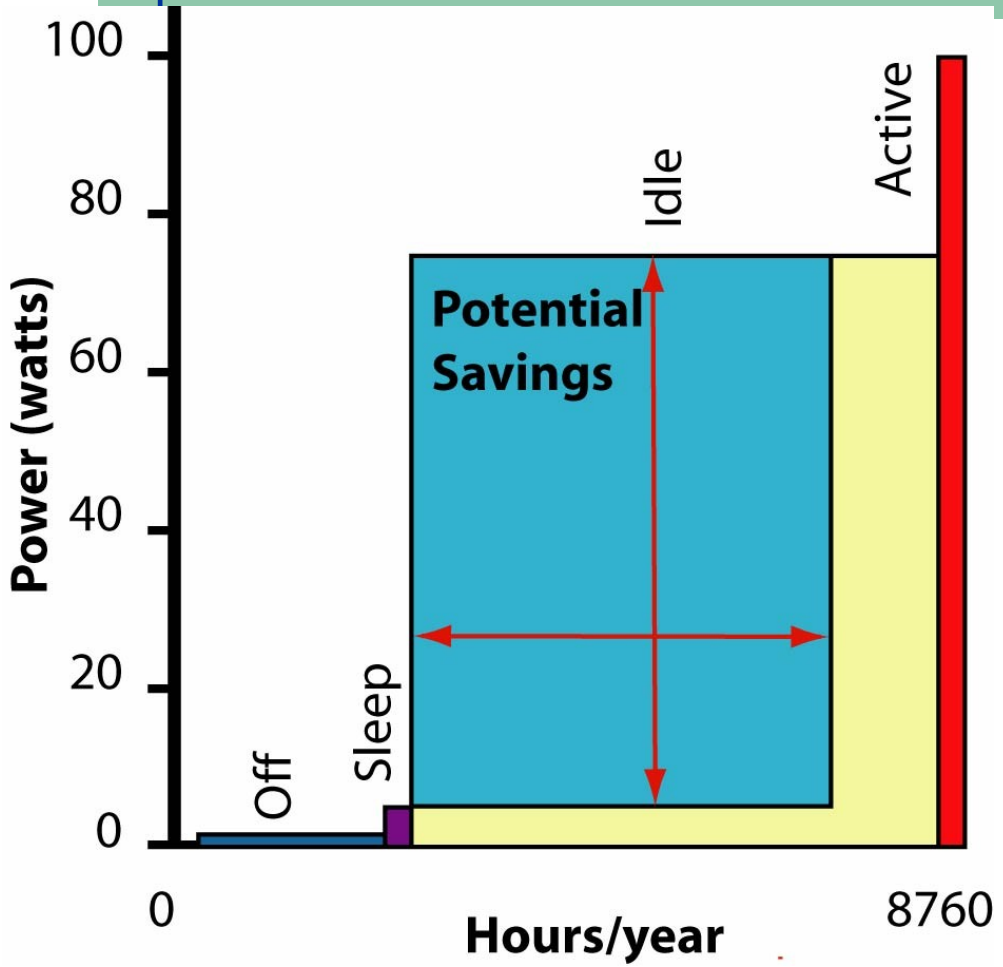

All time for year sorted by power level

Most of time when idle, could be asleep

PC savings potential is **most** of current consumption

Similar patterns apply to set-top boxes, printer, game consoles, …

Proxying addresses this; more later

Things we know:
Speed costs energy / power

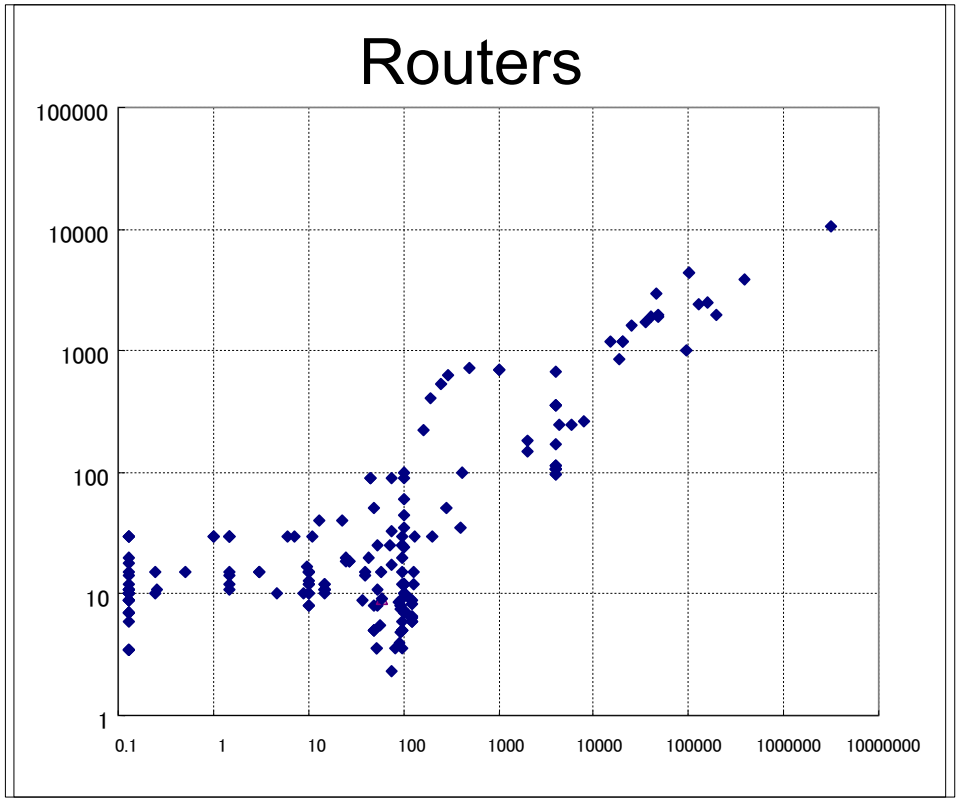

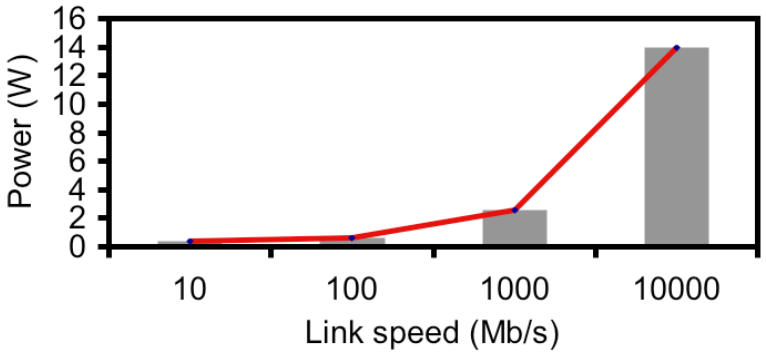

Measured power of various computer NICs (averaged)

Source: Christensen, 2005

Source: METI, 2006

Energy cost is a function of capacity, not throughput

Things we know:
IP will go everywhere

- IT equipment - IP already universal

- IP for phone calls (VOIP)
- IP for TV (IPTV)
- IP for consumer electronics generally
- IP for buildings (lighting, climate)
- IP for …..

ENERGY EFFICIENT

How much of this will be Linux?

Efficiency Approaches

Product Focus

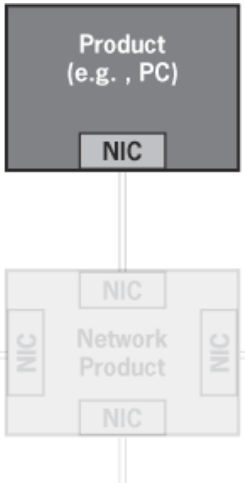

Network Product Focus

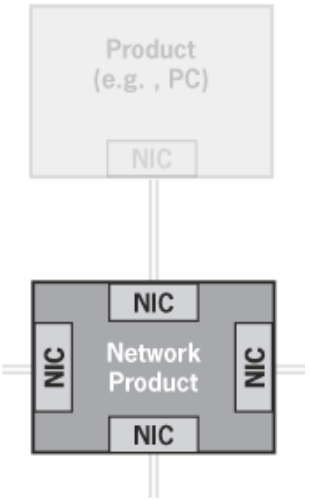

Interface Focus

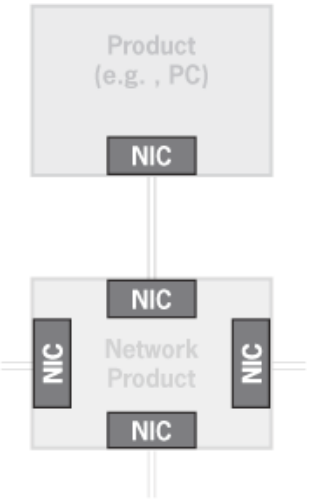

Protocol / Application Focus

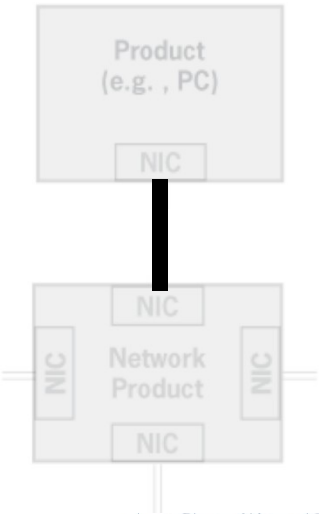

Examples:

Proxying

ENERGY STAR

Energy Efficient Ethernet

CE

ENERGY EFFICIENT DIGITAL NETWORKS

Adaptive Link Rate (ALR)

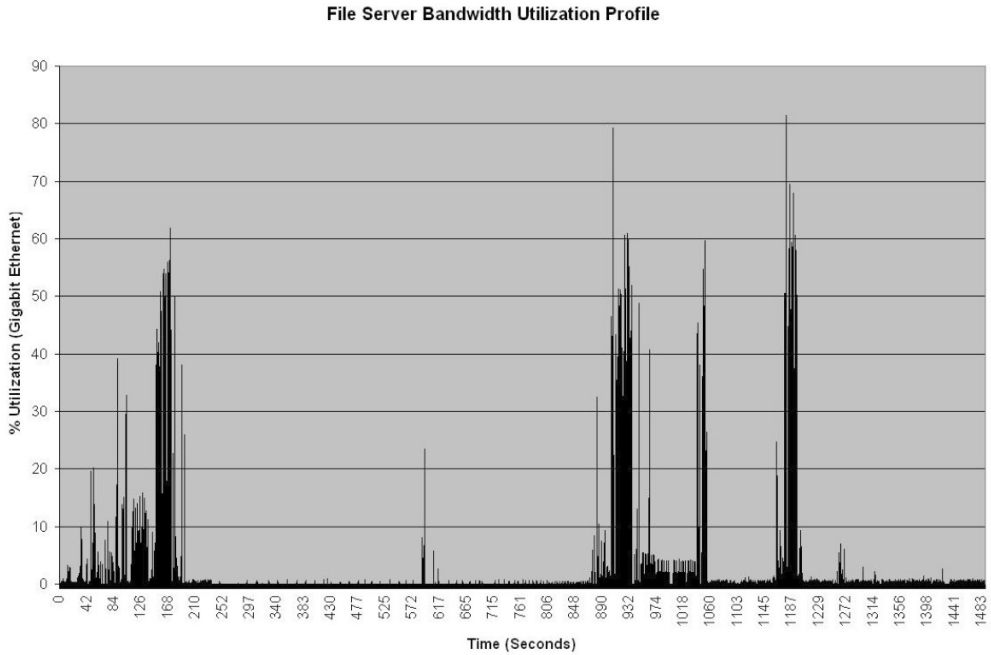

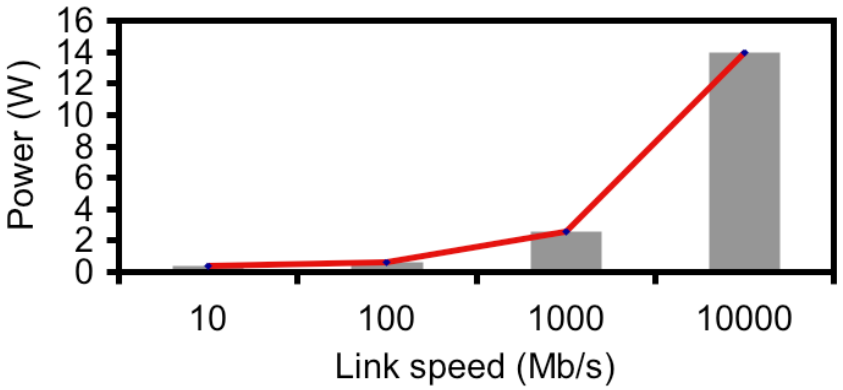

Observations

- Most of time, full link capacity not needed
- Notebooks already dropped link rate in sleep

Proposal (LBNL & USF)

- Enable changing link rate **quickly** in response to traffic levels
  *(**ms** not s)*

(ALR now) Energy Efficient Ethernet

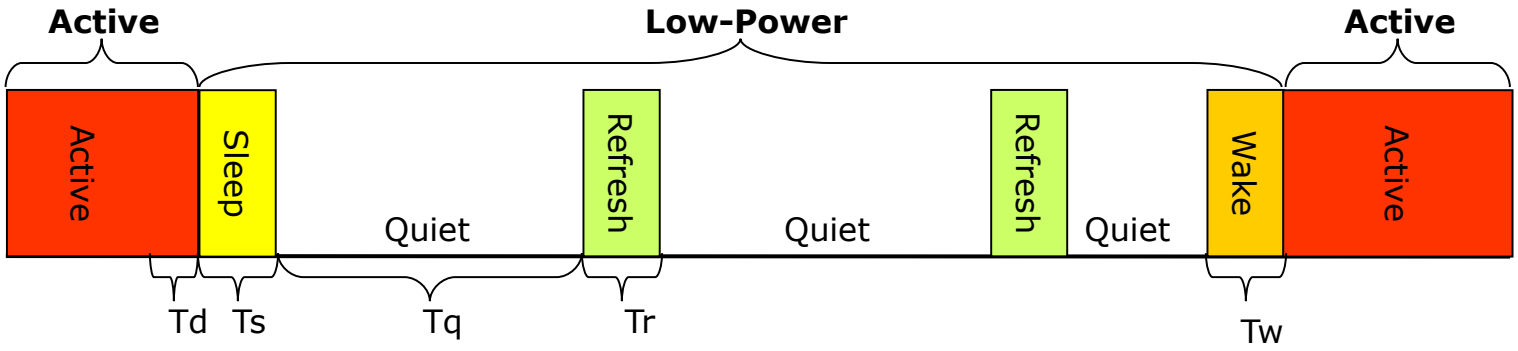

- IEEE 802.3az created to standardize EEE
- Standards process began with ALR; eventually settled on alternate method "Low Power Idle"
  - Stop transmitting between packets
  - Switch now takes ***micro***seconds
- Standards process needs about 1 more year
  - Goal to get EEE technology into ALL Ethernet network hardware globally over next few years

Energy Efficient Ethernet

IEEE

IEEE 802

ENERGY STAR

Linux needs to be aware of EEE but not much else


CEPro.com
Consumer Electronics
This the CE equipment in a real house

Our CE Future ?

While some integrators are skeptical about the prewired, preprogrammed NHS rack from Sony, others embrace the solution for its simplicity.

- Network / Data connectivity a Mess
- Number of CE devices is LARGE
- For energy use, digital networking could easily:
  - cause large increases, or
  - enable significant reductions
- We cannot rely on manual power control

ENERGY EFFICIENT DIGITAL NETWORKS

Consumer Electronics – What to do

- Move to 3-state power model
- Address link power consumption
- Provide for persistent network presence
- Expose power state to network
- Standardize some user interface elements
  - Displays
- Create a model for standard behaviors / expectations for CE devices

Many of these devices will run Linux. Any implications for OS or related activities?

User Interfaces

- Standard Interface elements common throughout daily life
- Key to safety, ease of use, efficiency
- Many use graphics, color, location, etc. to improve functionality and reduce language-dependence
- Commonality limited to comprehension needs
- Can deviate from standards when there is a good reason

User Interface Standards

- Consistent across:
  - Manufacturers
  - Products
  - Countries
- Simple
- Accessible
- Portable

Key Elements
- *Terms*
- *Symbols*
- *Colors*
- *Metaphors*
- *…*

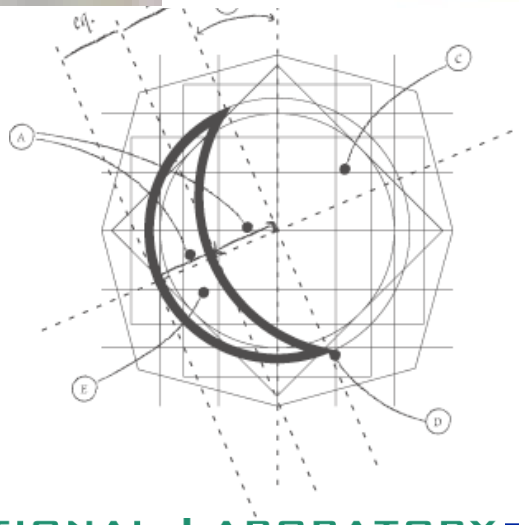

ENERGY EFFICIENT DIGITAL NETWORKS

Non-Interoperability
w/ devices or w/ people

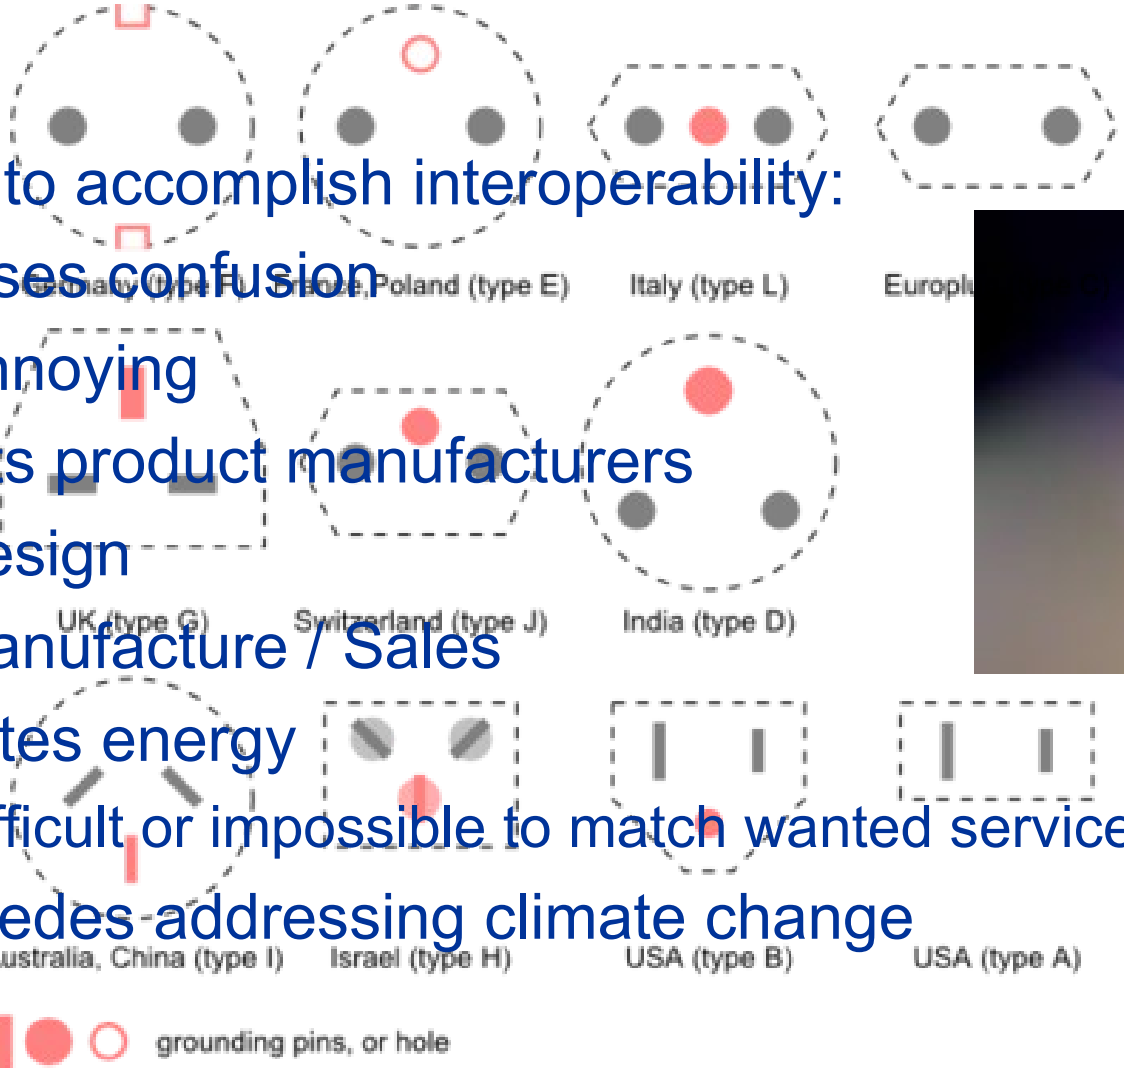

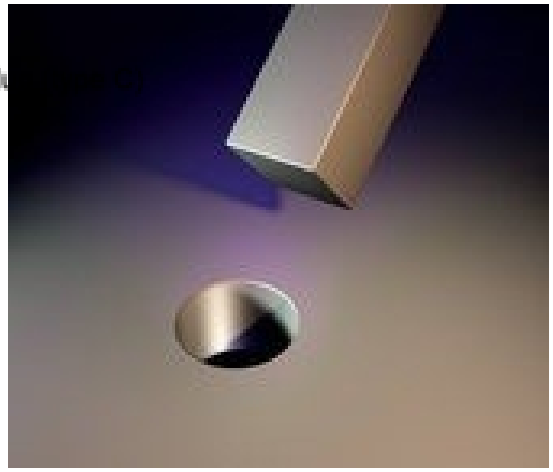

- Failure to accomplish interoperability:
  - Causes confusion
  - Is annoying
  - Costs product manufacturers
    - Design
    - Manufacture / Sales
  - Wastes energy
    - Difficult or impossible to match wanted service to delivered
  - Impedes addressing climate change

ENERGY EFFICIENT DIGITAL NETWORKS

User Interfaces

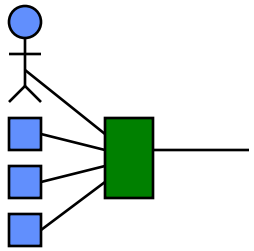

People:

- … are best understood as nodes on the network
  - Even more than portable electronics, they move
- … are often absent from design, presentation of networks
- … need standard interfaces, just like devices do
  - Nature of interface different, but principle same

Past LBNL work: “Power Control User Interface Standard”, IEEE 1621 - terms, symbols, colors, metaphor

IEEE

ENERGY EFFICIENT

Linux community should adopt IEEE 1621

"Network Connectivity Proxying"

A low-power entity that maintains "full" network connectivity for a sleeping high-power device

- Addresses energy use by devices persistently network-connected, but often doing little or nothing
- Key goal: hide host's sleep state from rest of network
- Need standard definition of proxy behavior
- Need cooperation of operating system

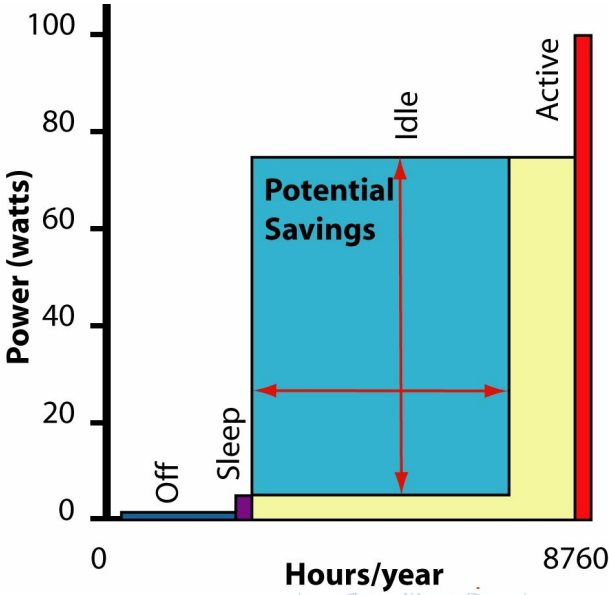

Key collaborator: Ken Christensen, University of South Florida

Linux community should engage proxying topic

Networks and Energy

Network equipment ….

Routers, switches, modems, wireless APs, …

… vs **networked** equipment

PCs, printers, set-top boxes, …

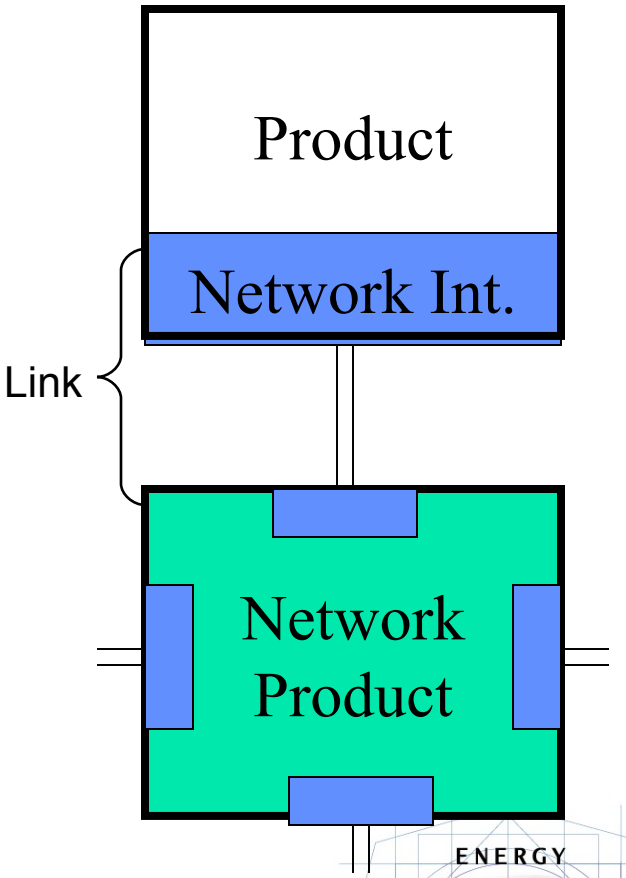

How networks drive energy use

- **Direct**
  - Network interfaces (NICs)
  - Network products
- **Induced** in Networked products
  - Increased power levels
  - Increased time in higher power modes (to maintain network presence)

ENERGY EFFICIENT DIGITAL NETWORKS

Proxying: Origins

3.4 PC Power Management with Networks

Networks pose special challenges for power management. Depending on the systems (hardware and software), the network can partially or entirely defeat power management, or may require extra configuration changes for it to function.

LBNL Report: 1997

INTERNATIONAL JOURNAL OF NETWORK MANAGEMENT
Int. J. Network Mgmt., 8, 120–130 (1998)

Enabling Power Management for Network-attached Computers

Power management is an emerging area of interest for network management. This article reviews current developments and describes methods for enabling power management in network-attached computers. © 1998 John Wiley & Sons, Ltd.

By Kenneth J. Christensen and Franklin 'Bo' Gulledge*

USF paper: 1998

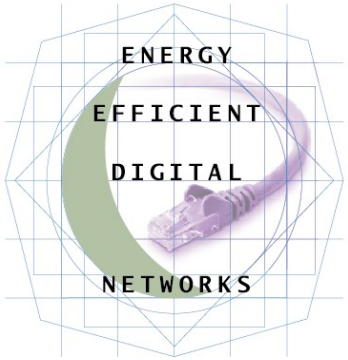

PC energy is mostly idle

Core Fact: Most PC energy use occurs when no one present

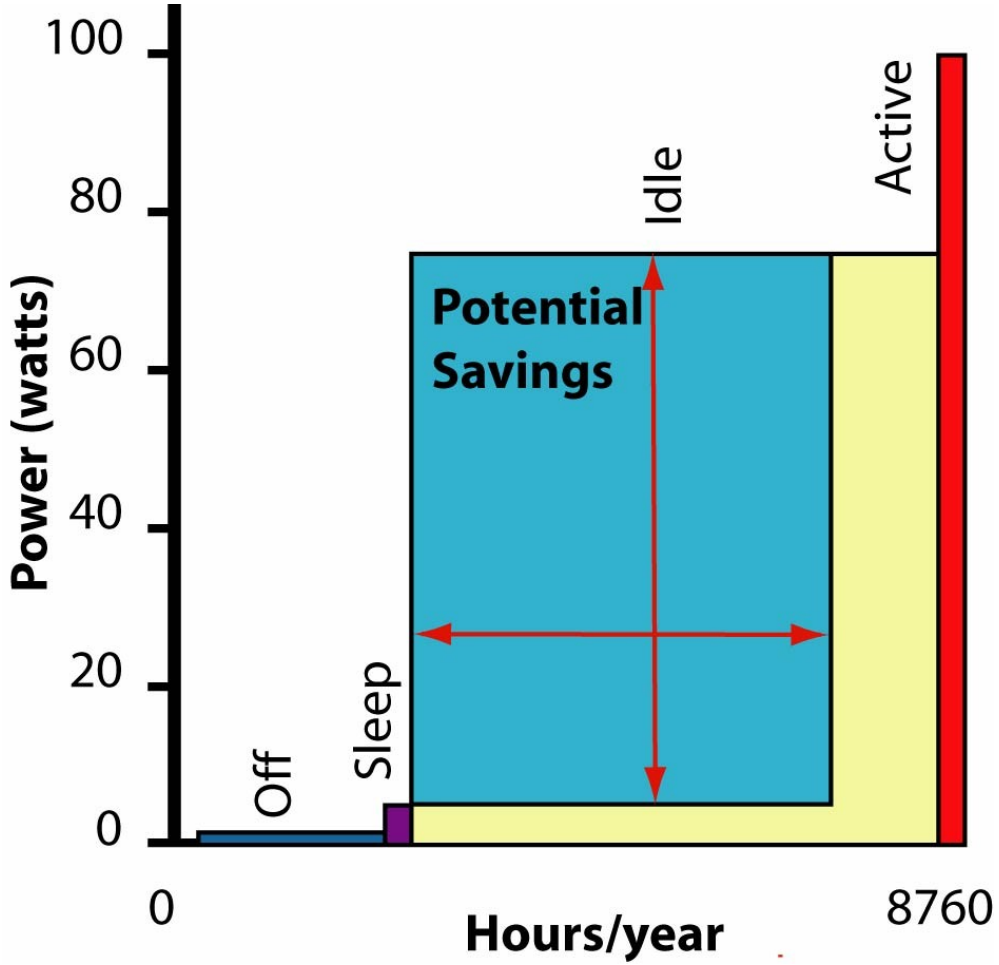

All time for year sorted by power level

Most of time when idle, could be asleep

PC savings potential is **most** of current consumption

Similar patterns apply to set-top boxes, printer, game consoles, …

ENERGY EFFICIENT DIGITAL NETWORKS

Proxying: Operation

Proxy operation

1. **PC awake; becomes idle**
2. **PC transfers network presence to proxy on going to sleep**
3. **Proxy responds to routine network traffic for sleeping PC**
4. **Proxy wakes up PC as needed**

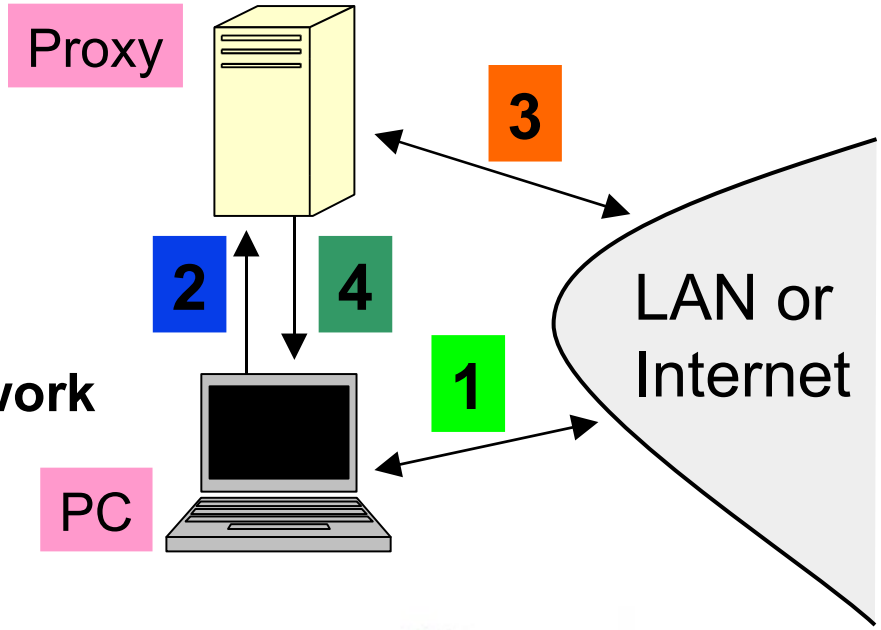

Proxy can be internal (NIC), immediately adjacent switch, or "third-party" device elsewhere on network

Proxy does: ARP, DHCP, TCP, ICMP, SNMP, SIP, ….

Proxying: Relevant Protocols

- **What is network presence?**
  - –Host-level reachability
    - ARP, IGMP
  - –Application-level reachability
    - TCP SYN, SIP invitations
  - –Addressability
    - DHCP
  - –Manageability
    - ICMP, SNMP
  - –Liveliness
    - TCP connections, application heartbeats

General Goals

- Enable large majority of PC users to use sleep without breaking their own or IT admin applications
  - At least 80%. > 90% better. > 95% or > 98% even better.
- Enable both current and emerging common applications
- Enable standard to directly (or easily adapted) for use in printers, set-top boxes, game consoles, etc.
- Describe behavior of “green applications” that don’t break proxying
  - Create *de facto* guide for new applications

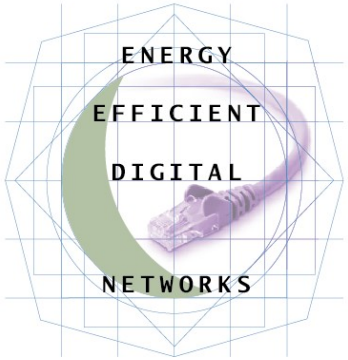

General Goals, restated

PC (or other edge device):

- Is always available
- Doesn't wake up if doesn't need to
- Does wake up when does need to
- Provides good user experience
- Provides consistent user experience
- Hides sleep status from rest of network
  - Except when explicitly tells

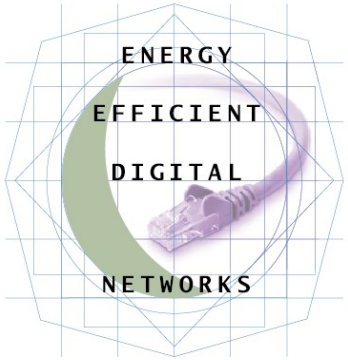

Ecma Context – TC38-TG2 (May, 2007)

Similar graphic in May 2006 meeting

Which one is more energy efficient?

Hint: Add up the total area of the graphs

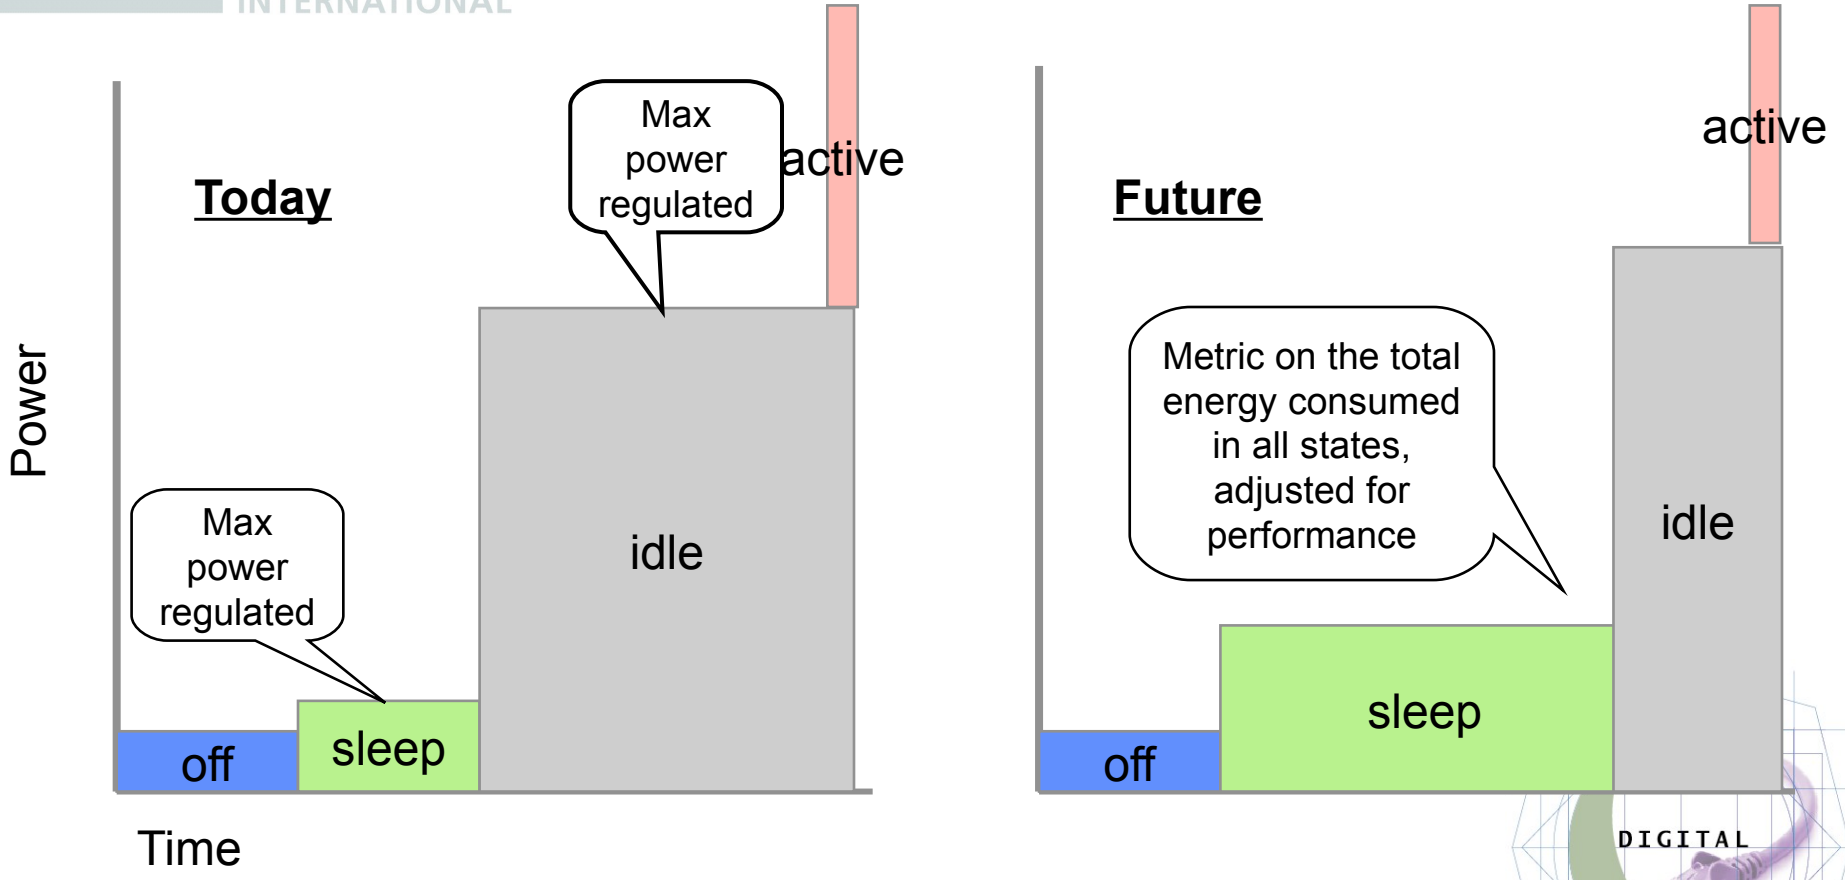

Note: This is simplified conceptual example and represents a sum of activities in each state

**Note: Active rectangle adjusted from original*

Energy Star context

Computer Performance and Energy Assessment Tool Stakeholder Meeting, June 20, 2007, Washington, D.C.

Background

- Most energy used by desktop PCs in U.S. when no one is present
- Enabling power management could save > 50% of desktop PC energy use
- Network connectivity the major impediment to enabling sleep moving forward
- Topic dates back to beginning of Energy Star PC process in 2004
- Intent is to enable sleep without requiring any changes to existing protocols and applications used on great majority of PCs
- Wake On LAN inadequate for general solution for many reasons

Goal

- Drive proxying functionality into all networked electronic products that have significant On / Sleep power difference (printers, consumer electronics, etc.)

EPA Announcement of V4.0 Process, September 2004

Tier 2
I) Fix the “network problem” with power management

In future, Linux community should become more engaged with Energy Star

Energy Star V5.0 Computer Spec (eff. June, 2009)

Definition (emphasis added)

Full Network Connectivity: **The ability of the computer to maintain network presence while in sleep and intelligently wake when further processing is required.** Maintaining network presence may include obtaining and/or defending an assigned interface or network address, responding to requests from other nodes on the network, or sending periodic network presence messages to the network all while in the sleep state. In this fashion, presence of the computer, its network services and applications, is maintained even though the computer is in sleep.

Requirement: None

Incentive

- Reduced idle time in TEC calculation

Proxying: Process

Standard
- Ecma TC32-TG21

Trace Analysis
- Intel Research Berkeley

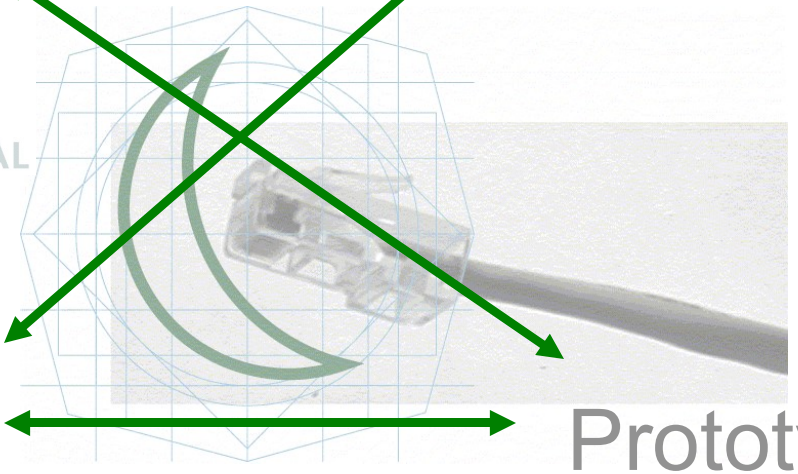

Use Cases
- In development

Prototypes
- Microsoft Research "Somniloquy"
- ???

Standard by end of 2009

ENERGY EFFICIENT DIGITAL NETWORKS

Key Points

- Standard intended for PCs and any “PC-like” device
  - Desire for persistent network connectivity
  - Non-trivial difference between idle/sleep power
  - PCs, printers, set-top boxes, game consoles, …
- Establishes a floor of functionality, not a ceiling
- Not designed with servers in mind
- Avoid any content that limits location of proxy between Internal (NIC), and External (closest switch or router)
  - NOT get distracted by “third-party” proxy location

Need contacts in Linux community to assure this gets implemented in timely fashion

Thank you!

efficientnetworks.LBL.gov

Bruce Nordman
Lawrence Berkeley National Laboratory
BNordman@LBL.gov
510-486-7089

(or google)