# The World Beyond x86

Restoring Trust and Owner Control to General Purpose Computing

Timothy Pearson
tpearson@raptorengineering.com

Raptor Engineering, LLC
https://www.raptorengineering.com

# The World Beyond x86

- Purpose
  - Highlight critical problems inherent in x86
  - Generate interest in owner controlled computing
  - Introduce viable x86 alternatives
  - Provide information about Talos™

# The World Beyond x86

- What is x86?
  - A family of ISAs
    - Backward-compatible
    - Derived from the Intel® 8086
    - Anything past the 80486 is Intel®-controlled

# The World Beyond x86

- x86 and Owner Control
  - Machine owners license firmware
  - Licensed firmware implements antifeatures
    - Hardware and software DRM
    - PAVP
  - Licensed firmware cannot be replaced
    - Boot Guard
    - Management Engine
    - Platform Security Processor

# The World Beyond x86

- x86 and Owner Control (cont.)
  - Licensed firmware has full system access
    - Essentially executes in Ring -3
    - High security and privacy risk
  - Licensed firmware is network enabled
    - May communicate with manufacturer
    - May communicate with third parties
    - Becomes a remote hacking risk

# The World Beyond x86

- x86 and Owner Control (cont.)
  - Licensed firmware is typically restricted
    - Unable to independently audit
  - Bottom Line
    - Owner control is effectively absent on x86
    - Restoring owner control is not possible
    - x86 ISA is never licensed by Intel®
    - Lack of competition means no alternative, performant x86 platforms will ever exist!

# The World Beyond x86

- Why is Owner Control important?
  - Freedom
  - Flexibility
  - Reliability
  - Privacy
  - Security
  - Continuance
  - Competition / multi-source capability

# The World Beyond x86

- What about other closed firmware components?
  - Host CPU control is critical
    - Must also control any CPUs with equal or higher privilege level
  - Various host isolation mechanisms available
    - IOMMU / SMMU / TCE / DVMA
    - DMA-incapable busses (USB 1.0 / 2.0)
  - CPU and isolation silicon must be trustworthy!

# The World Beyond x86

- CPU microcode
  - Contentious topic
  - Refers to two distinct firmware components
    - Traditional horizontal / vertical microcode
    - Firmware for embedded on-die “parasite” CPU(s)
  - Traditional microcode
    - Shuttles data to / from internal logic blocks
    - Controls logic sequencing
    - Limited risk

# The World Beyond x86

- CPU microcode (cont.)
  - On-die CPU firmware
    - No different than any other highly privileged CPU
    - Firmware typically has full system access
      - RAM
      - Persistent storage
      - External peripherals (network, keyboard, display)
    - Sufficient hardware resources to execute malware
    - Firmware typically written in a high level language
    - Potentially high risk with closed firmware source!

# Closed Firmware Risk, Relative

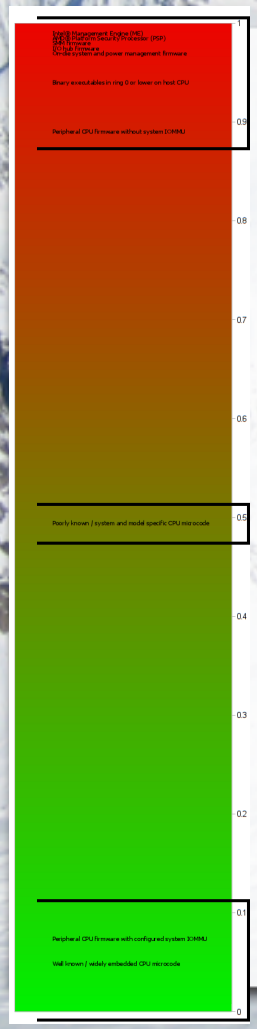


## Low Risk

Peripheral CPU firmware with configured system IOMMU

Well known / widely embedded CPU microcode

0.1

0

## Moderate Risk

Poorly known / system and model specific CPU microcode

0.5

## High Risk

Intel® Management Engine (ME)

AMD® Platform Security Processor (PSP)

SMM firmware

I/O hub firmware

On-die system and power management firmware

Binary executables in ring 0 or lower on host CPU

Peripheral CPU firmware without system IOMMU

1

0.9

RAPTOR ENGINEERING®

# The World Beyond x86

- What can be done to move away from x86?
  - Use a different ISA!
    - This is feasible for most libre software users
    - x86 inertia largely stems from proprietary lock-in
  - Many non-x86 ISAs to choose from
    - ARM
    - MIPS
    - POWER
    - RISC

# The World Beyond x86

- ISA overview
  - ISA alone rarely dictates performance limits
    - Each ISA is targeted at specific market(s)
    - Reference designs typically set performance
  - Developing processors costs money
    - The ISA creator needs to fund further R&D
    - The CPU manufacturer also needs to fund R&D
    - x86 and the SPARC T2 are outliers in openness
    - Long-term manufacturer support is vital

# The World Beyond x86

- Processor market segmentation
  - Consumer x86 positioned in unprofitable region
    - Requires extreme sales volume
    - Not only selling silicon; also licensing software
    - Software lock-in / rental model driving repeat sales
  - Laptop market most affected
    - No challengers apparent on mid to high end
    - ARM only available on low end / crippled laptops
  - Main competitors positioned differently

# Architecture Market Segmentation

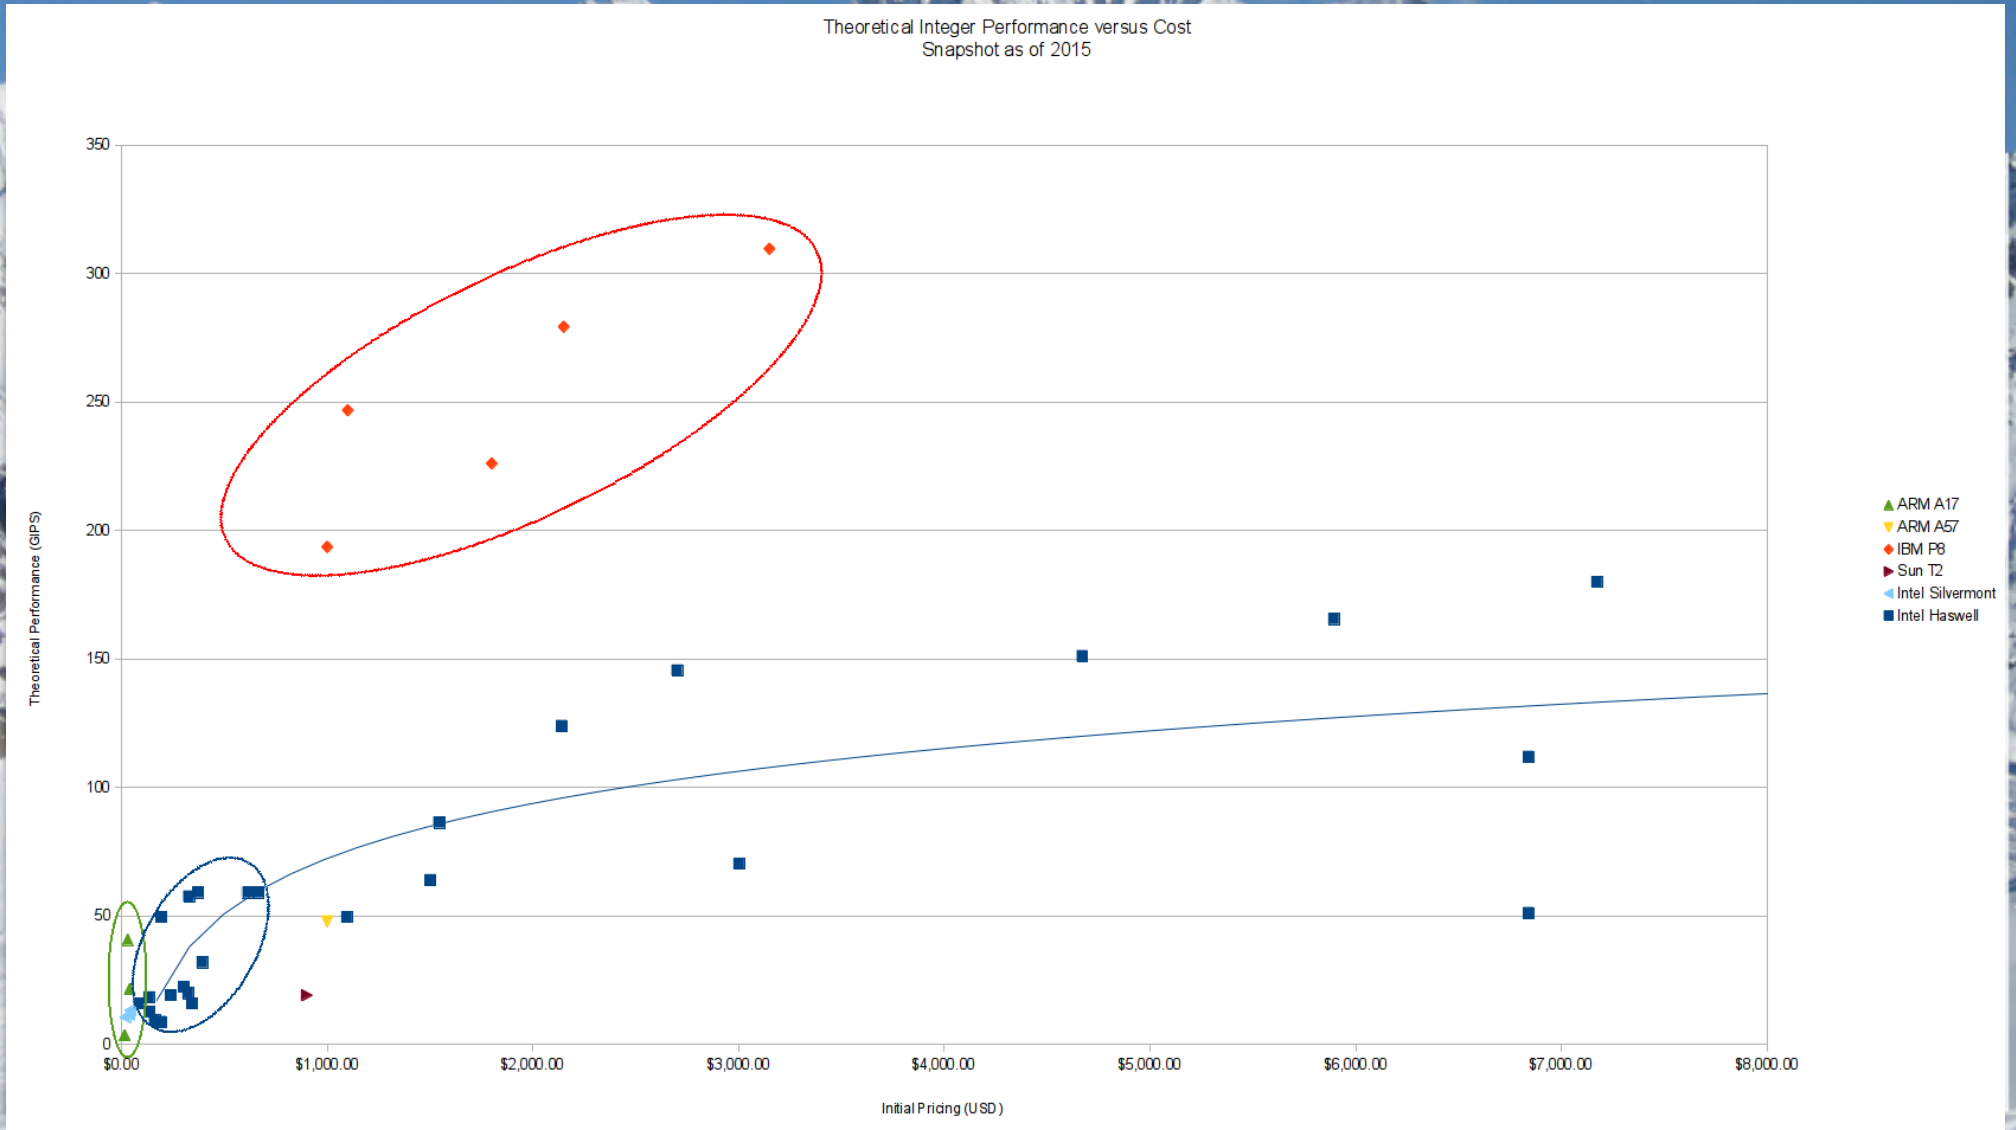


RAPTOR ENGINEERING®

# The World Beyond x86

- Choosing an ISA
  - Avoid vendor lock-in
    - Choose a licensable ISA
    - Make sure the license allows modification
  - Verify suitability for general purpose computing
    - ISAs linked to cheap CPUs can perform poorly
    - Many cheap CPUs use limited cache hierarchy
    - Many cheap CPUs trim features needed for GPC
    - Some CPUs still require signed firmware

# The World Beyond x86

- Libre Computing ISA Candidates
  - Two architectures most viable
    - ARM
      - Extremely popular
      - Diverse; many vendors own core licenses
      - Suffers from poor design decisions
      - Cheap, but relatively slow
    - POWER
      - Extremely fast, competitive with high-end x86
      - Feature rich but expensive
      - OpenPOWER not yet as diverse as ARM

# The World Beyond x86

- ARM
  - Midrange locked, libre capable on high/low end
    - High end
      - Mainly offered by NXP (formerly Freescale)
      - Has PCIe, ECC, SATA, USB3, hypervisor, SMP
      - Cache structure limits GPC performance
      - Expensive
    - Low end
      - Offered by RockChip, Allwinner, etc.
      - No PCIe, no ECC, no SATA, no hypervisor
      - SoC support spotty in upstream u-boot / Linux

# The World Beyond x86

- ARM (cont.)
  - First non-x86 architecture to gain traction
  - Leverages typical x86 weaknesses
    - Simpler design (less power, silicon)
    - Highly customizable SoC designs
    - Multi-sourced w/ similar features across vendors
    - Licensable
  - Primarily targets mobile and network markets
  - Hardwired microsequencer

# The World Beyond x86

- Example owner-controllable ARM64 system
- NXP LS2085A NPU
  - 8 A57 cores (4 SMP clusters w/ L2 cache)
  - Standard DDR4 DIMMs, ECC, SMMU, PCIe
  - No L3 cache
  - RDB uses odd form factor – BTX?
  - DPAA2 requires proprietary firmware
    - Removing firmware disables integrated NICs

# NXP LS2085ARDB NPU


RAPTOR ENGINEERING®

# The World Beyond x86

- POWER
  - Main focus of presentation
  - ONLY libre-capable core on par with x86 cores
  - Low- to mid-range libre capable (OpenPOWER)
  - High end locked
    - Mainly aimed at ultra-high-reliability use
    - x86 servers do not offer similar capabilities
    - Raw performance comparable to OpenPOWER

# The World Beyond x86

- POWER (cont.)
  - Traces roots back to RS/6000
  - Multiple firmware components, on-die CPUs
    - Source available / modifiable on OpenPOWER
  - Traditional microcode
    - Unlike x86, can be read back out and verified
    - Microcode not lost on power off / reset
    - Can be locked in hardware to prevent updates

# The World Beyond x86

- POWER (cont.)
  - Historically big endian (BE)
  - Hardware little endian support as of P8
    - Allows non-endian-aware code to function
    - Eases porting burden
  - Bi-endian capable
    - Booted kernel sets endianness
    - Bare metal firmware can boot BE or LE
    - KVM VM kernels can select either endianness

# The World Beyond x86

- POWER (cont.)
  - OpenPOWER “Turismo”
    - Single Chip Module (SCM)
    - Up to 12 “chiplets” (core, L2, and sharable L3)
  - Expansive cache hierarchy
    - L3 and L4 cache
    - eDRAM allows large caches
  - Large die size, high TDP
    - TDP in line with similarly performant CPUs

# Theoretical Performance Versus TDP

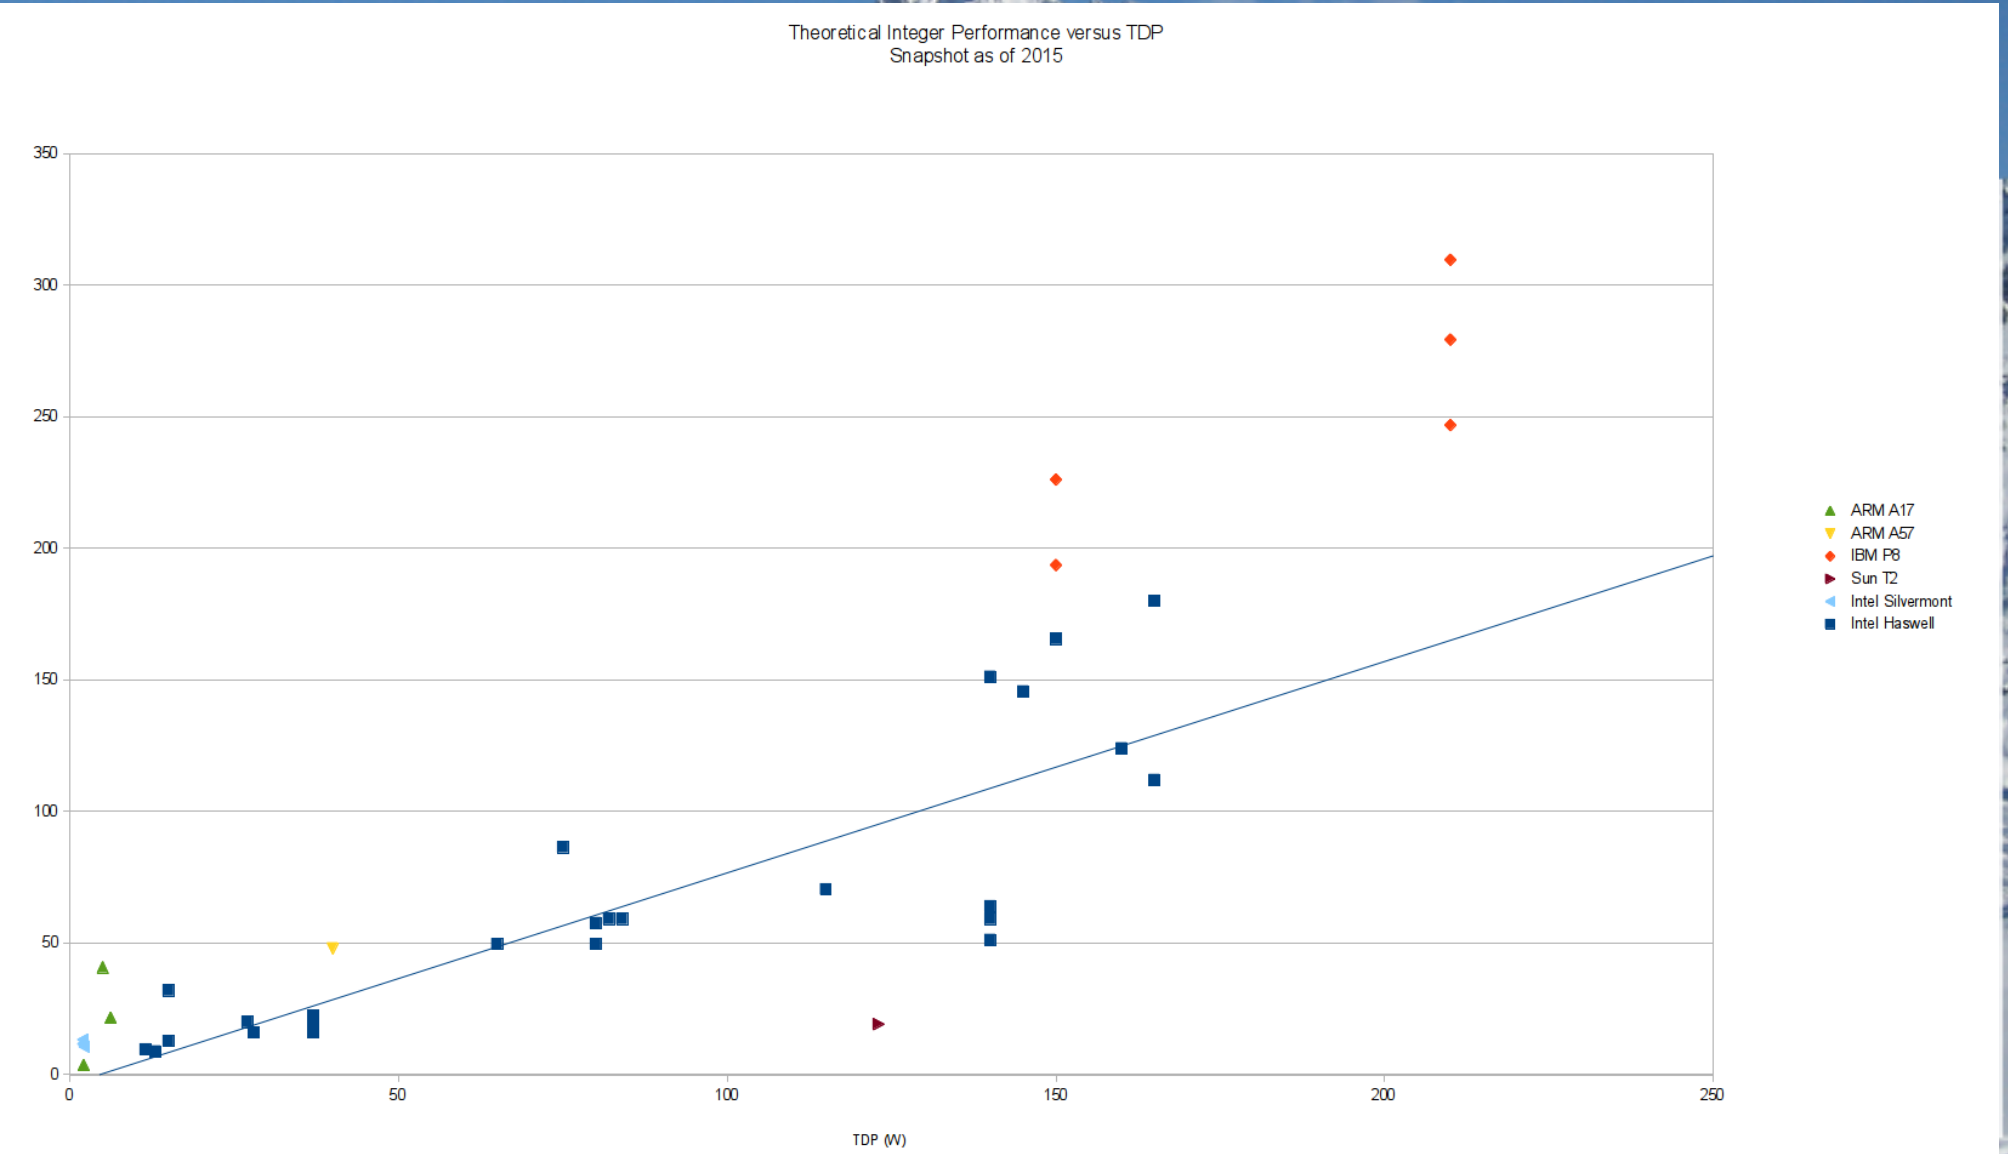


RAPTOR ENGINEERING®

# The World Beyond x86

- POWER (cont.)
  - High Bandwidth
    - Higher memory read / write bandwidth vs. x86
    - Higher I/O bandwidth than most other offerings
    - Low-latency CAPI interface for accelerators
    - Outperforms most (all?) competing, owner-controllable machines in bandwidth capacity

# Opteron™ 6328 Memory Bandwidth

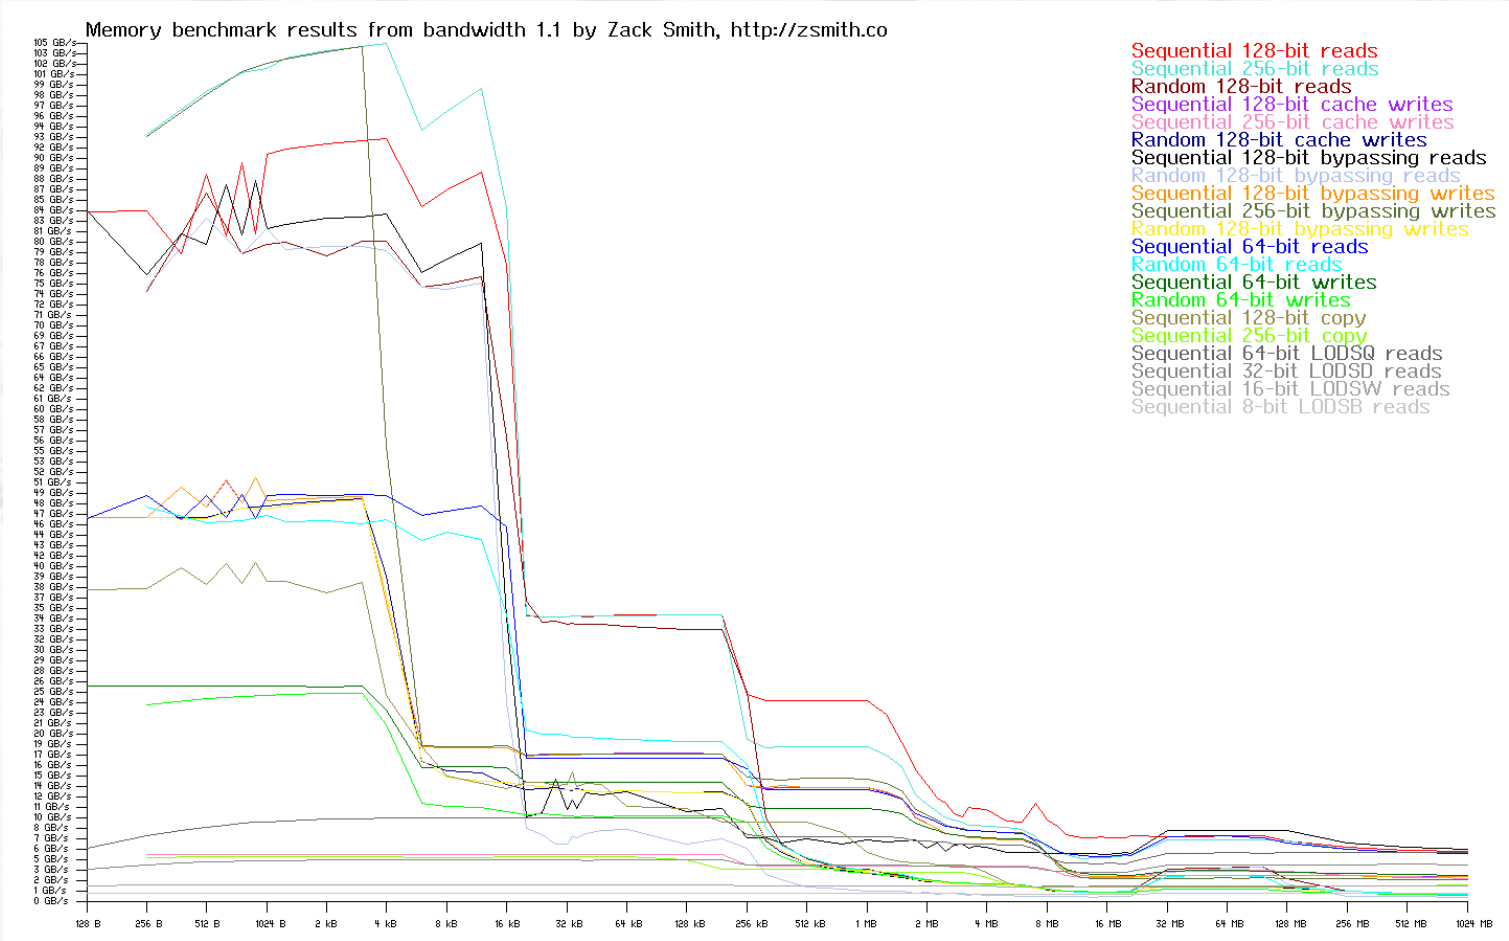


RAPTOR ENGINEERING®

# POWER8 Memory Bandwidth

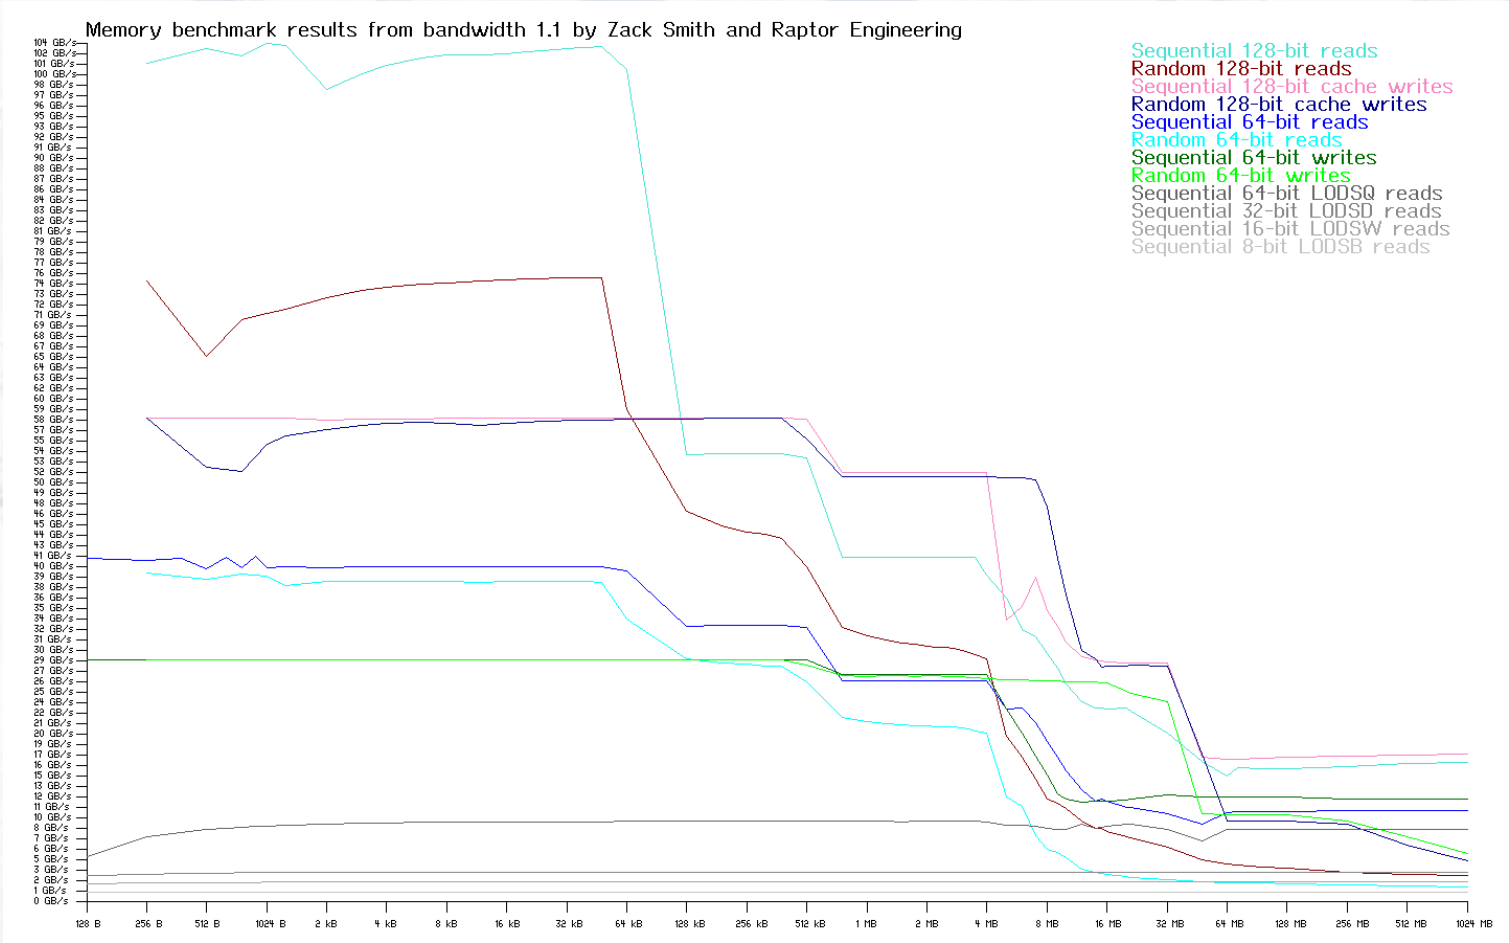


RAPTOR ENGINEERING®

# The World Beyond x86

- POWER (cont.)
  - Significant investment in Linux ecosystem
    - Debian archive coverage > 95%!
    - Full KVM support including TCE
  - Unprecendented level of firmware access
    - No regions cryptographically locked by hardware
    - OCC firmware is open source
  - Hardware capable of ChromeOS security model
    - Keeps full control in hands of machine owner

# The World Beyond x86

- POWER (cont.)
  - Hardware Support
    - LE mode means most open drivers work OOB
    - Closed drivers must be recompiled by vendor
    - Industry-standard high-speed interfaces
      - PCIe
      - USB 3.0
    - Special-purpose interfaces
      - CAPI
    - In practice, most hardware works without hassle

# The World Beyond x86

- POWER (cont.)
  - Software Support
    - Extensive support for open Linux applications
    - Some support for closed Linux applications
      - IBM Java
      - NVIDIA CUDA
    - Notable Exceptions
      - Chromium (lack of Google interest circa 2014)
      - Flash (becoming largely irrelevant)
      - VirtualBox (x86 only, use KVM / QEMU instead)

# The World Beyond x86

- POWER (cont.)
  - Software Support (cont.)
    - 3D functionality verified w/ Radeon 290X
      - Engineering Tools
        - Kicad
        - FreeCAD
      - 3D Modelling
        - Blender
      - Games
        - Xonotic
    - No perceivable difference from x86!

# The World Beyond x86

- POWER Future Direction
  - Next major revision will be POWER9
    - Due late 2017
    - New socket, different RAM architecture
    - Even more open than POWER8!
      - Fully open microcode and toolchain expected
    - Adopting open P8 systems helps ensure P9 later
  - Secure Boot
    - Owner control central to IBM Secure Boot model

# The World Beyond x86

- OpenPOWER Firmware
  - OCC
    - On-chip power management / thermal control
  - Hostboot
    - Low-level startup
  - Skiboot
    - OPAL runtime services
  - Petitboot
    - User interaction / final bootloader

# The World Beyond x86

- OpenPOWER Firmware (cont.)
  - GRUB (on disk / distro provided)
    - IEEE 1275-1994
  - Coreboot port in progress
    - Initial target: QEMU
    - Initial toolchain and bootblock ready
    - Needs attention / more developers
    - P8 starts in BE mode, complicating port

# The World Beyond x86

- OpenPOWER Firmware (cont.)
  - Closed components
    - Microcode, including Self Boot Engine
      - Roughly equivalent to FPGA bitstream
      - No direct equivalent in x86 or ARM worlds
      - Not cryptographically checked by hardware
      - Stored on processor module
      - Written with external programmer; not loaded by host
      - IBM may open these components for audit on P8
      - P8 toolchain unreleased due to size / complexity
      - Will be fully open for P9

# The World Beyond x86

- POWER Availability
  - Existing products targeted at server market
    - 1 and 2 CPU systems available for rack mount
    - Bare ATX boards not available
    - Low-end machines tend to be proprietary
      - Tyan
      - Certain IBM servers designed by third parties
      - No firmware source, no schematics, no support
      - Recycled proprietary firmware used for BMC, etc.
    - High-end machines are open, but expensive

# The World Beyond x86

- Example owner-controllable POWER8 system
- IBM S822LC
  - 2 POWER8 190W CPUs
  - Unique RAM backplane system w/ Centaurs
  - Numerous PCIe slots
  - Ships with proprietary BMC firmware
    - Open firmware under development
    - Basic functionality available under open firmware

# IBM S822LC Commercial Server

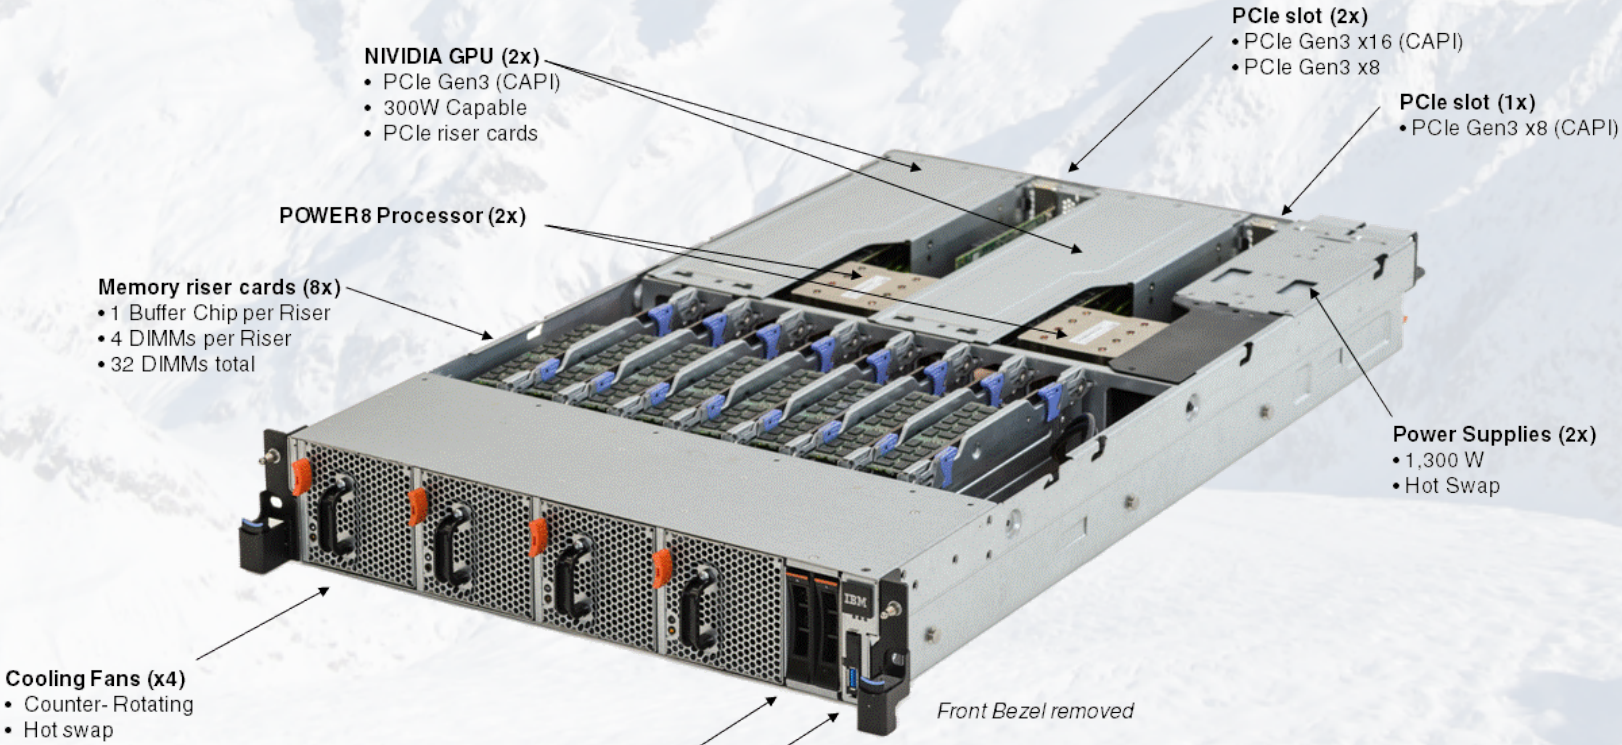


Courtesy of IBM Redbooks

RAPTOR ENGINEERING®

# The World Beyond x86

- Raptor Engineering's Talos™ system
  - Standard ATX-compatible mainboard
  - Specifically targeted at workstation market
    - 1 POWER8 CPU – 130W stock, 190W optional
    - Up to 256GB DDR3 w/ ECC
    - High PCIe lane count
    - SATA, USB 3.0, etc.
  - Fully open firmware
  - Includes schematics for most hardware!

# TALOS

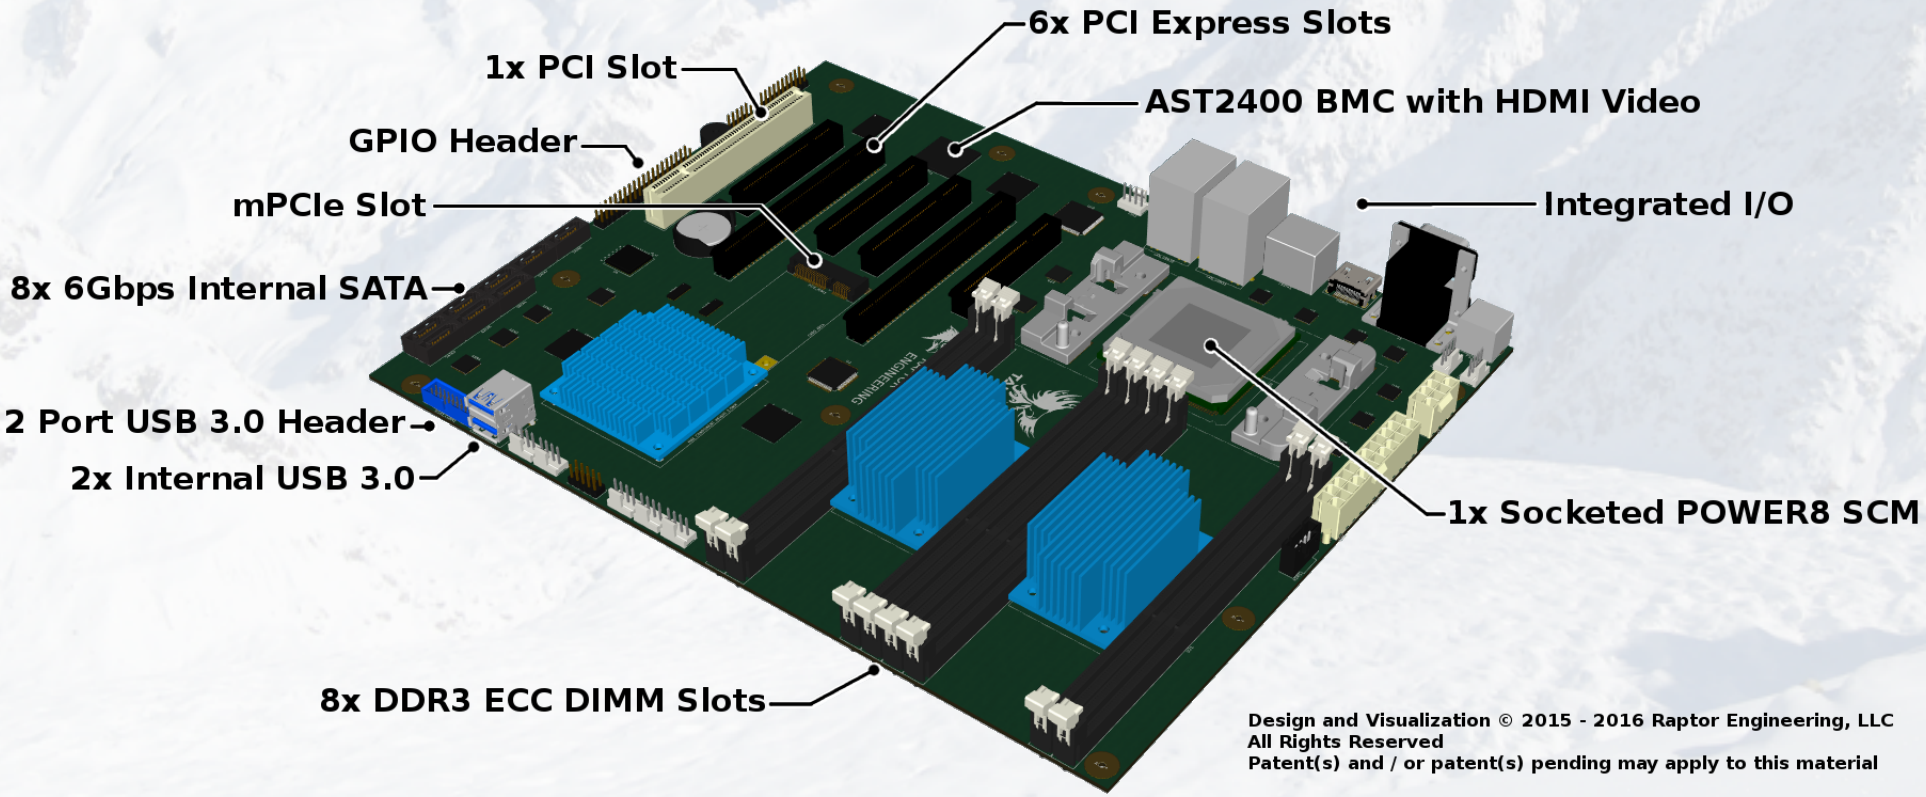




RAPTOR ENGINEERING®

# The World Beyond x86

- Raptor Engineering's Talos™ system (cont.)
  - Unique Features
    - Open toolchain FPGAs for signal routing
    - Programming headers for firmware development
    - Will ship with OpenBMC & OpenPOWER firmware
    - Multiple hardware-enforced security options
      - PNOR write protect switch
      - Key store write protect switch
      - Flash partition write protect switches
      - BMC network disable switch

# The World Beyond x86

- Raptor Engineering's Talos™ system (cont.)
  - Primary goals
    - Allow trusted, complex general purpose computing
    - Encourage development of libre low-level firmware
    - Provide a viable x86 workstation alternative
    - Enable innovation through OpenPOWER access
  - Status
    - System development well underway
    - Gauging overall market demand

# The World Beyond x86

- The GPU Situation
  - NVIDIA
    - NVE (Kepler) and below work well under nouveau
    - All cards above NVE (Kepler) are locked
    - >= NV110 (Maxwell) unusable w/o signed firmware
    - NVIDIA has been very slow to enable even basic nouveau functionality via signed firmware images
    - NVLink (P8+ and above) requires binary drivers

# The World Beyond x86

- The GPU Situation
  - AMD
    - Requires unsigned firmware download to card
    - Radeon driver executes vendor ATOMBIOS
    - Libre radeonsi driver functions well
  - Lack of competition
    - Developing GPUs is expensive
    - NVIDIA has adopted anti-competitive techniques
    - Use standardized OpenCL wherever possible!

# The World Beyond x86

- Summary
  - x86 is familiar, but increasingly locked down
  - Alternative architectures are far more open
  - Switching from x86 is becoming a viable option
  - Libre software has enabled architecture choice
  - ARM primarily targets low end and NPU markets
  - OpenPOWER / Talos™ target workstation and server markets

# The World Beyond x86

- Summary (cont.)
  - Fully libre CPUs are not currently viable
    - Most attempts have ended in failure
    - RISC-V may have chance in embedded systems
    - Major manufacturing technology developments required before GPC viability
  - Balance possible between proprietary and libre
    - Sell hardware, not firmware and software licenses
    - ARM / OpenPOWER seem to have found balance

# The World Beyond x86

- Summary (cont.)
  - Era of cheap, commodity GPC has ended
    - Consumer machines are cryptograpically locked
    - Consumer markets have shifted to rental / cloud
  - Libre software advocates have three choices
    - Use cheap, low-end, reverse engineered machines for as long as they remain unlocked
    - Accept loss of owner control (and libre software)
    - Pay for full featured, owner-controlled machines

# The World Beyond x86

- Summary (cont.)
  - OpenPOWER has re-enabled trusted GPC
    - Still needs more support from libre community
    - Trends are encouraging
    - Projects such as Talos™ need market demand!
  - After battle for control of CPU has concluded...
    - ...battle for control of GPU begins!

# The World Beyond x86

Thank you for your attention!

Check out the Talos™ Secure Workstation
https://www.raptorengineering.com/TALOS/

Follow us on Twitter or GNU Social!
https://twitter.com/RaptorEng
https://social.raptorengineering.io/raptoreng

# The World Beyond x86

- Additional Resources
  - Platform Embedded Security Technology Revealed: Safeguarding the Future of Computing with Intel Embedded Security and Management Engine
    - http://www.apress.com/9781430265719
  - Intel x86 considered harmful - The Invisible Things
    - http://blog.invisiblethings.org/2015/10/27/x86_harmful.html
  - Intel & ME, and why we should get rid of ME
    - https://www.fsf.org/blogs/licensing/intel-me-and-why-we-should-get-rid-of-me
  - Debian archive coverage for supported machine architectures over time
    - https://buildd.debian.org/stats/graph-quarter-big.png
  - OpenPOWER machine firmware
    - https://github.com/open-power