Standards for Server Architectures

Intersect360 RESEARCH

Michael Feldman
michael@intersect360.com

Overview

- The Case for Standards
- Application Drivers
- Server Architectures
- A Special Case: Microservers
- Components
- Component Integration
- Vendor Dynamics
- Conclusions
- Recommendations for SSI

The Case for Standards

- Hardware ecosystems broaden as components become interchangeable
- Standardization of components (e.g., x86 processors, DDR memory, SATA disks, etc.) encourage commoditization, lowers costs
- Standardization of interfaces enables interoperability, avoids vendor lock-in, and lowers cost of manufacturing (e.g., PCI, Ethernet, InfiniBand, USB, SATA, Switch SERDES)
- Solution for non-standard hardware was software – operating systems, compilers, and other software components that abstracted hardware differences

Intersect360 RESEARCH

Application Drivers

- Cloud/Web-based computing – Internet of things, virtual desktop computing, cloud storage, online retail, search, online publishing, web hosting, streaming media
- Big Data – Financial risk management, social networking, business intelligence, customer analytics, complex event processing, security/intelligence, scientific analytics and visualization
- High Performance Computing (HPC) – Scientific research, manufacturing R&D, financial risk management/trading, oil and gas exploration, medical research/drug discovery
- All segments growing much faster than generic enterprise computing

Intersect360 RESEARCH

Application Drivers (cont'd)

- With some exceptions, data movement for these categories of applications is more important than compute performance (most are memory- or I/O-bound)
- Requires computing at scale, which drives leaner server designs of higher density and greater energy efficiency
- Diversity of applications encourages purpose-built servers targeted to specific software requirements
- All have profound ramifications for server design

Intersect360 RESEARCH

Server Architectures

- Rack servers – general-purpose standalone computers with integrated power supplies, cooling, but external network switching; advantage is flexibility in deployment and operation; 1U rack server the de facto standard
- Blade servers – also general-purpose computers but with shared infrastructure: power supply, networking, and often cooling integrated into enclosure, no standard form factors; advantage is efficiency in power and space
- Microservers – low-power alternative based on single-socket, low-power processors with shared infrastructure as in blades; advantage is lower cost with higher efficiencies in power and space than blades

Intersect360 RESEARCH

A Special Case: Microservers

- Progression of rack servers to blades to microservers reflects the scalability and efficiency needs of the application drivers outlined above
- Microservers, in particular, reimagines the server design, encapsulating many cluster components into a single enclosure, reducing network switching, cabling, and peripheral storage
- Note that microservers are not just applicable for lightweight applications like web hosting, but also heavy-duty applications like big data analytics and HPC codes that can be decomposed into many lightweight tasks.

Components

- Processors – Today: x86, POWER, SPARC, GPUs; Future: ARM64, integrated CPU-GPU like AMD APU, manycore CPUs like Intel Xeon Phi; industry is moving from CPUs plus discrete devices to SoCs
- Network Fabrics – Ethernet, InfiniBand, Fibre Channel, custom fabrics; Ethernet will remain the de facto standard for the foreseeable future
- Interconnects and Buses – PCI Express for external devices, vendor-specific buses (QPI, HyperTransport, etc.) for inter-processor and I/O hub communication
- Power Supplies
- I/O Controllers

Component Integration

- Server processors becoming more specialized to serve a widening array of applications and deliver better energy efficiency
- SoC model emerging; integration of network controllers, I/O controllers and compute accelerators is underway by all major server processor manufacturers: Intel, AMD, NVIDIA, and IBM
- Microservers will take advantage of SoC efficiencies, but will be used generally across server architectures
- Integration will reduce motherboard complexity, but not necessarily in a standard manner as is evident by initial microserver designs

Intersect360 RESEARCH

Power and Cooling

- Power costs are driving designs: more energy-efficient components, stripped down servers, shared infrastructure within the rack, integrated components; microservers incorporate all of these
- Air cooling remains the default, but water cooling becoming more prevalent as designs become denser and heat loads can’t be dispersed air
- Server and rack-level cooling; standard rack (42U) suggests opportunity for standard cooling solution

Intersect360 RESEARCH

Broader Trends

- The 1U single-node server giving way to the cluster in a box with integrated power, cooling, and I/O
- The new “server” is a multi-node cluster in standard (42U) rack or chassis sub-rack (2U, 4U, 8U, etc.)
- The rack unit (U) endures as the standard measurement of server real estate, but internals are being revamped and specialized to serve new demands
- Largest server users (Facebook, Amazon, and Google) now make their own servers from standard parts

Approaches to Implementation of Standards (OCP)

- Open Compute Project (OCP) – Facebook-inspired server specifications targeted to scalable, low-cost, energy-efficient computing
- OCP is mainly concerned with defining standards for stripped down server hardware – as well as corresponding storage, switches, racks, and other data center hardware
- By implicitly limiting the application set to Facebook-like applications and encompassing other data center hardware, OCP is able to focus its standards effort

Approaches to Implementation of Standards (SSI)

- Server System Infrastructure (SSI) Forum – Group focused on defining standards of server infrastructure including motherboard, chassis, mezzanine cards, power supply and management modules; also IBM open switch spec for blades
- For member OEMs, ODMs, and server component makers (processors, power supplies, chassis, etc.)
- Published specs on rack servers, blade systems and microservers available
- Provides Product Development Kits (PDKs) for compliance/interoperability

Vendor Dynamics

- Server OEMs and chip vendors are looking to exploit opportunities created by Cloud, Big Data, and HPC growth
- Commodity-based systems tend to provide narrow margins; purpose-built processors and systems for specific application profiles promise better returns, even if system costs drop
- Using application virtualization for ever more powerful servers and processors is a questionable proposition for OEMs and component providers
- Low-power SoCs and microservers offer opportunities for all scalable computing growth segments

Intersect360 RESEARCH

Vendor Dynamics (Processors)

- Intel: Focused on network controller integration and designing low-power x86 processors for the emerging microserver space; also looking to move up the food chain with native server motherboards and chassis
- AMD: Providing servers (SeaMicro SM15000) for emerging microserver market, developing APUs (CPU-GPU integrated processors) and 64-bit ARM SoCs for servers
- ARM licensees: Companies like Samsung or Qualcomm could enter the server chip business via their ARM expertise developed with mobile market; Texas Instruments already has low-power ARM-DSP offering aimed at specialized server applications

Intersect360 RESEARCH

Vendor Dynamics (Systems)

- IBM: renewing focus on POWER-based servers and has opened up the architecture for third-party providers; also partnered with NVIDIA to co-develop GPU-accelerated servers for big data applications; reduced x86 dependency with sale of System X business to Lenovo
- HP: Heavily invested in microserver paradigm via its “Moonshot” systems; incorporates low-power x86 processors today, but designed to incorporated ARM, GPUs, DSPs or other SoCs in future.
- Dell: Offers microserver variant under PowerEdge line (C5220); developed ARM-based prototypes “Zinc” and “Copper”

Conclusions

- Server architectures are undergoing a transition from general-purpose rack servers to more purpose-built servers (blades and microservers) that use shared infrastructure to lessen power and space requirements
- Cloud/web computing is volume growth market; servers are less demanding of compute than traditional enterprise computing relative to data storage and data movement; simple, scalable designs are key
- Big data and HPC are lower volume growth markets, but often (not always) with more specialized compute, data storage and data movement requirements

Intersect360 RESEARCH

Conclusions (cont'd)

- As a result, market is bifurcating into high-volume, generic servers for the web and cloud, with little product differentiation and lower-volume, more specialized servers for big data and HPC
- Standards are easier to establish for the former; emphasis on very dense, low-power, low-cost, scalable designs; such designs can be satisfied by ODMs and second tier OEMs
- Big data and HPC servers tend to be more demanding, requiring extra architectural expertise, and these are best left to the large OEMs; note that easily decomposable big data and HPC applications fall naturally into generic server model

Intersect360 RESEARCH

Conclusions (cont'd)

- For servers built to optimize data storage and data movement, PCI and Ethernet standards need to keep pace to serve more demanding requirements
- Overall, as more data center and motherboard infrastructure is shrunk down onto processor, server standards will tend to simplify even as certain elements of system design (interprocessor interconnect, memory and storage hierarchies, more diverse processors, for example) becomes more complex
- Standardization will be challenging through transition to server designs that rely on more tightly integrated processors, more diverse processors and shared infrastructure

Intersect360 RESEARCH

Recommendations for SSI

- Continue to split specifications taxonomically across the major architectures: rack, blade, and microserver
- For the time being, focus on x86-based specifications; ARM and third-party POWER-based solutions are not yet established (and may never be), however...
- Incorporate OCP specifications to support emerging technologies, specifically ARM/microserver designs
- Wait at least another year for accelerator-based server and chassis designs; market is still immature and volatile

Intersect360
RESEARCH

Recommendations for SSI (cont'd)

- Consider broader support/cooperation with OCP for other specifications, including Open Rack, storage, and networking
- Open Rack in particular could be common area of agreement given diversity of traditional server forms and architecture
- Finally, consider shifting direction of SSI into an advisory forum, rather than strictly a standards-making body; create a platform for discussion that can guide members

Intersect360 RESEARCH