

Calxeda™: RACK TRUMPS THE CHIP

IT Services datacenters have diverged from traditional enterprise datacenters; they have fundamentally different needs driven by massive infrastructure scale. Calxeda has designed an innovative rack-level network fabric architecture for these new IT Services datacenters. Calxeda’s architecture today connects dozens of densely packed, independent server nodes, and will scale in the future to deliver greater operational efficiencies across these new service-oriented mega-datacenters.

State of IT Datacenters Today

For most of the history of enterprise datacenters, a workload image ran within a single server chassis or on an individual blade. Because of this inherent modularity, a modern consolidated and converged enterprise datacenter’s operational model assumes no predictability in scheduling which workloads run on specific servers – any server can run any workload image. Due to Moore’s Law, core enterprise workload performance is now well past “good enough,” the goal now is to boost server utilization through virtualization, and to increase IT flexibility by consolidating those workload instances onto heavily virtualized servers.

Virtualization and consolidation have turned x86 servers into completely interchangeable components; there is no opportunity to leverage significant feature differentiation between servers. Enterprise IT customer demands have pushed x86 servers into an undifferentiated, low margin business.

However, over the last few years, the IT industry has quietly entered a new phase of datacenter build-outs that are unrelated to traditional IT datacenters. These datacenters are not being built to improve business efficiency; the services they enable are directly monetizable.

There are three broad classes of services encompassed in this new “IT Services” model:

Service	Business Model	Workloads *
Outsource IT	Shift non-differentiating IT services to at-scale provider: better, faster, and cheaper	Virtual private cloud, public cloud, native web services
Consumer-to-Consumer (C2C)	Enable consumer collaboration to gain insight into consumer behaviors	New file systems, new databases, media storage and delivery, Big Data analytics
Machine-to-Machine (M2M)	Enable intelligent sensors to collect real-world context for new/enhanced services	

* Web Tier workload is common to all services

Mega-Datacenters: The Difference is Scale

Many companies in the IT Services space have reached Fortune 500 status in their own right. Some have globe-spanning networks of datacenters. By comparison, only a few multi-national companies and governments have the need to build datacenters on this scale. IT Services vendors must competitively lower costs while raising their quality of service. They must drive infrastructure evolution as a business imperative.

Traditional market measurement methods show that HP, Dell and IBM now account for over 75% of server system revenue, however [Wired reported](#) that eight server makers now account for 75% of Intel's server chip revenue. Given the eight now include contract design manufacturers and direct purchasing by mega-datacenters, this is already a significant, high-growth market.

New Architecture Enables Greater Scale

Mega-datacenter pain points are much different from enterprise IT. They have a new model for attaining operational efficiency – they know which clusters of racks will run specific workloads, many times *in advance of building their datacenter*. Optimizing rack architecture for specific workloads directly impacts their bottom line.

These new at-scale datacenter infrastructure customers:

- Control their own destiny through system software source code ownership
- Concentrate their purchasing power on specific workload solutions
- And so attract innovative new hardware and software infrastructure investment

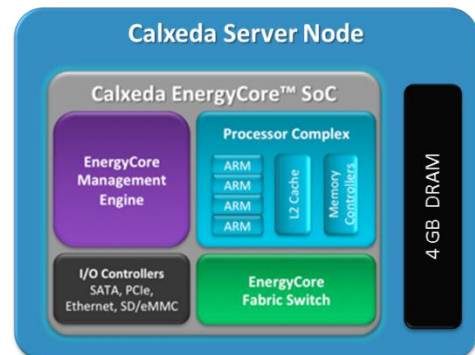
Calxeda's is aiming directly at IT Services with their new rack-level fabric architecture.

The Calxeda Architecture

Network topologies created from connected switches are loosely called “fabrics.” Calxeda's architecture has three major components: rack level interconnect fabric, management engine, and all of the logic for a complete server node integrated into one chip.

Rack-level network interconnect: Calxeda has created a design-time configurable fabric architecture. EnergyCore™ fabric switches are embedded in every server node – they are the foundation of this new fabric architecture. As customers become more sophisticated in analyzing the performance of workloads running on their proprietary runtime software frameworks, they can optimize their fabric topology for a given workload. Configurable fabrics have profound implications for future performance tuning at-scale. It's a fundamentally more economical way to design workload-optimized servers. This kind of optimization is implemented today only in expensive purpose-designed High Performance Computing (HPC) servers.

EnergyCore Management Engine: This logic block includes standard server baseboard management controller (BMC) functions for node reliability and availability. Because each Calxeda server node contains an embedded switch, the EnergyCore Management Engine adds real-time fabric traffic optimization and fabric power optimization for on-the-fly performance tuning, while a system is fully operational. These are very sophisticated features not found in volume-manufactured servers. The EnergyCore Management Engine runs on a separate, dedicated ARM Cortex-M3 core, and supports standard IPMI and DCMI interfaces to enable control and monitoring of on-chip functions, such as power on, boot sequence control, temperature and power readings, etc. This reduces cost and power consumption compared to standalone BMCs needed for most server environments.



EnergyCore Server-on-Chip:

Calxeda’s novel fabric architecture and on-the-fly performance tuning are enabled by integrating an EnergyCore Fabric Switch and an EnergyCore Management Engine onto the same die as a quad-core server System-on-Chip (SoC).

Integration enables all of these components to operate together at chip-speed and without burning the additional power needed for chip-to-chip communications in a less integrated design, such as a traditional Intel processor plus southbridge server node. The EnergyCore Management Engine manages on-chip power consumption, and also includes additional hardware acceleration to enable its real-time fabric traffic and power optimization.

Choosing Sides – ARM vs. x86

Processor integration enables Calxeda’s unique architecture. Enterprise servers today are overwhelmingly based on the x86 instruction set. Until recently that posed a problem – Intel, AMD are the only two processor vendors who have server-class x86 processors, and they are not licensing their designs for other processor vendors to use. The expensive and increasingly boutique RISC processor vendors are also unlikely to grant a start-up a good licensing deal.

ARM has the most competitive merchant processor architecture on the market. Silicon design and systems modeling tools are readily available, ARM has a vibrant ecosystem of licensable third-party functional blocks, and software runtime environments and development tools are even more widely available than their design tools. ARM is also focused on the server market and has a credible roadmap focused on the same set of IT Services customers as Calxeda.

Calxeda’s EnergyCore ECX-1000 processor is now in full production and shipping to customers. It is available in the EnergyCard reference design through HP, Dell, Boston Ltd, System Fabric Works, and Penguin Computing.

Calxeda Architecture Performance

The first iteration of Calxeda’s integrated server SoC is based on ARM’s Cortex-A9 32-bit core – by all definitions the Cortex-A9 design is a “small” core. While much has been written comparing [big cores and small cores](#), the important facts to remember about Calxeda’s choice are:

- The ARM Cortex-A9 core is very power efficient.
- It has adequate performance for several I/O intensive IT Services workloads, where system performance is limited by network and storage access.
- The Cortex-A9 core’s performance per watt, augmented with Calxeda’s power management features, is more than enough to permit Calxeda entry into the low-end of the IT Services market.

Power of:	Calxeda EnergyCore ECX-1000 (4-core)	Intel E3-1240 (4-Core Sandy Bridge)
Processor	2W*	80W*
SoC or add chipset	3.8W*	86.7W
Node (SoC + 4GB DRAM)	5W*	~91W

* Manufacturer’s stated TDP

Calxeda and Boston Limited have published power measurements that seem to bear out the company’s claims (below). Their measurements are all under 200 watts, even under various

stress loads, with the exception of the notorious STREAM benchmark. However, the more interesting performance metrics will be to measure racks and clusters of racks dedicated to single workloads. I am looking forward to seeing those results in coming months.

Workload*	Total System Power**	Power per ECX-1000 Node***
Linux at Rest	130W	5.4W
phpbench	155W	6.5W
Coremark (4 threads per SoC)	169W	7.0W
Website at 70% utilization	172W	7.2W
LINPACK	191W	7.9W
STREAM	205W	8.5W

* 24 nodes, 24 SSDs, 96GB DRAM cluster in a 2U chassis

** Measured with disks at the wall

*** Approximate: calculated

Calxeda's R&D Roadmap

Calxeda says that they plan to spend much of their recently announced funding round on R&D, to ensure that they will be a leading competitor in the IT Services datacenter market by mid-decade.

ARM's design and manufacturing ecosystem enables Calxeda to ship Cortex-A9 cores today and then substitute new cores as ARM releases them to production.

	Calxeda Silicon	Calxeda Platform	OSV	Workloads
Q4 2012	Cortex-A9 32-bit volume production availability	Seed OSVs with ARM systems using production Gen 1 fabric	Deploy with Ubuntu for ARMv7; enable development for Fedora and others	C2C, M2M, and telecom data plane; Static/small web pages; Some analytics (Hadoop, NoSQL databases, etc.)
2013	Cortex-A9 32-bit virtual/40-bit physical production ramp; Bring up A15 32-bit silicon	Initial 32-bit deployments; Support Linaro Enterprise Work Group standardization	Linux 3.7 kernel with AArch64 picked up by distros; Linaro Enterprise WG accelerates ARM ecosystem	
2014	Cortex-A15 32-bit virtual / 40-bit physical volume production ramp; Bring up 64-bits	Production systems with Gen 2 fabric	AArch64-based distro release candidates	All Web tier; General cloud hosting
2014-2015	64-bit production ramp	Production Systems with Gen 3 fabric	AArch64 enterprise OS volume ramp	Virtualized cloud; Analytics; etc.

Calxeda is planning to upgrade from the 32-bit Cortex-A9 core to the 32-bit Cortex-A15 core for increased performance. This will not obsolete their Cortex-A9-based processors – the Cortex-A15 core offers better performance per watt, and the new processors will offer higher performance with more memory. This new processor will be pin-compatible within the EnergyCard implementation with the EnergyCore ECX-1000 processor (customers can use their existing EnergyCard system board layouts), it will not be backward compatible for power draw or thermal solutions.

As Calxeda launches their Cortex-A15 based processor, they will also upgrade their EnergyCore Management Engine to implement a second-generation fabric. The fabric upgrade is aimed at better automating power management across the fabric and within the processor, optimizing routing, and increasing reliability across thousands of server nodes. It will be a

software-only upgrade that will also apply to their existing Cortex-A9 based EnergyCore ECX-1000 processors.

Calxeda's customers will then have a choice of best-in-class power consumption based on the Cortex-A9 vs. a highly power-optimized, higher performance, larger memory processor based upon the Cortex-A15. Those who do not need the performance uplift can continue to use EnergyCore ECX-1000 processors, which will continue to be the best low-power server choice.

Given their current product investments in ARM silicon, software and systems, Calxeda stands a very good chance of leading the field in delivering Cortex-A15 based production servers.

In 2014, Calxeda is planning to bring-up an ARMv8 based processor to support mainstream 64-bit workload stacks with light virtualization and large memory profiles. They will also upgrade their Management Engine to Fleet Services™ with a third-generation fabric aimed at datacenter scale, tens of thousands of server nodes, and enterprise network Quality-of-Service implemented across clusters of racks. While it is likely that several other ARM licensees are implementing their own ARMv8 microarchitectures, Calxeda will stay focused on maintaining their differentiation via fabric integration and will license a standard ARMv8 implementation.

There are no software changes needed for Calxeda to upgrade from the 32-bit Cortex-A9 core to the 32-bit Cortex-A15 core. Upgrading from 32-bits to the ARMv8 64-bit architecture will require a little work, but it is well understood work as the x86 server market took exactly the same path a decade ago.

Calxeda's initial IT Services workload targets do not require 64-bit hardware or OS support. Canonical's Ubuntu Server with long-term support (LTS) is already available on Calxeda's customers' systems. I suspect that other 32-bit Linux distributions will follow with ARM releases shortly, given the amount of pent-up interest in ARM-based server solutions.

64-bit OS and hypervisor qualification cycles will most likely wait until early ARMv8 silicon and systems are available in mid-2014. Best efforts aside, there is a reasonable chance that 64-bit OSes and hypervisors will not see production release until early 2015.

As Calxeda upgrades their fabric and cores, more IT Services workloads will become available to them. Over the next few years, Calxeda's delivered performance will increase much faster than the performance demands for many workloads while maintaining low-power advantages.

Call to Action

The mega-datacenter infrastructure landscape is changing rapidly. Sun was right, the network is the computer – Sun was simply a couple of decades too early.

If you are an IT Services vendor in C2C or M2M businesses, or if you are building [Big Data](#) analytic systems, then you should include Calxeda in your evaluation matrix. Calxeda's novel architecture is well suited to new IT Services C2C and M2M build-out. Their first implementation, based on the ARM Cortex-A9 core, is worthy of production testing for Web Tier servers and a few other low-end mega-datacenter workloads.



Author

[Paul Teich](#), Contributing Analyst, [Moor Insights & Strategy](#)

Editor

[Patrick Moorhead](#), Principal Analyst, [Moor Insights & Strategy](#)

Inquiries

Please contact us at the email address above if you would like to discuss this report, and Moor Insights & Strategy will promptly respond.

Licensing

Creative Commons Attribution: Licensees may copy, distribute, display and perform the work and make derivative works based on this paper only if *Paul Teich and Moor Insights & Strategy* are credited.

Disclosures

Moor Insights & Strategy has a consulting relationship with Calxeda, Inc. No employees at the firm hold any equity positions with Calxeda, Inc.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

©2012 Moor Insights & Strategy

Calxeda, EnergyCore, and combinations thereof are registered trademarks of Calxeda, Inc.

Other names are used for informational purposes only and may be trademarks of their respective owners.