China Cloud Computing Conference

2010年5月21日–22日　中国·北京

第二届中国云计算大会

One Stone, Three Birds: Performance, Power and Space

Ihab Bishara
Director, Cloud Computing Products
May 21st 2010

What is cloud?

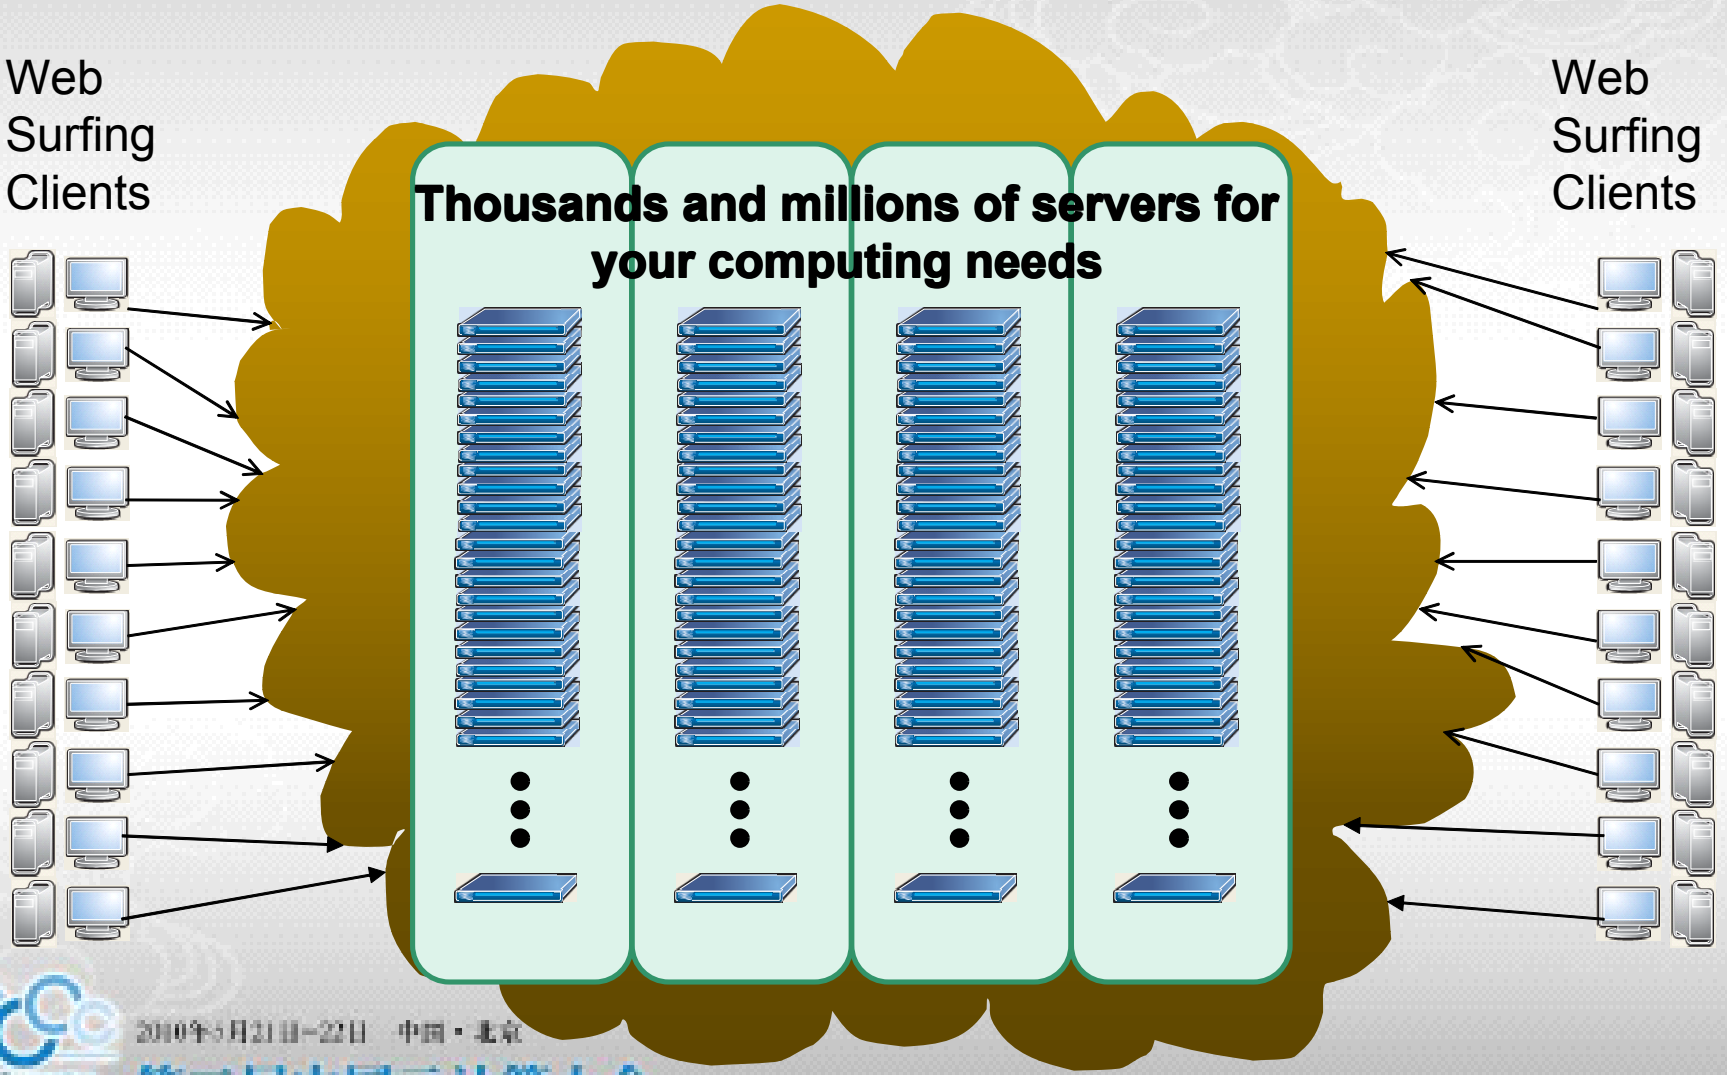

Cloud is needs bigger, more cost efficient, and more power efficient datacenters

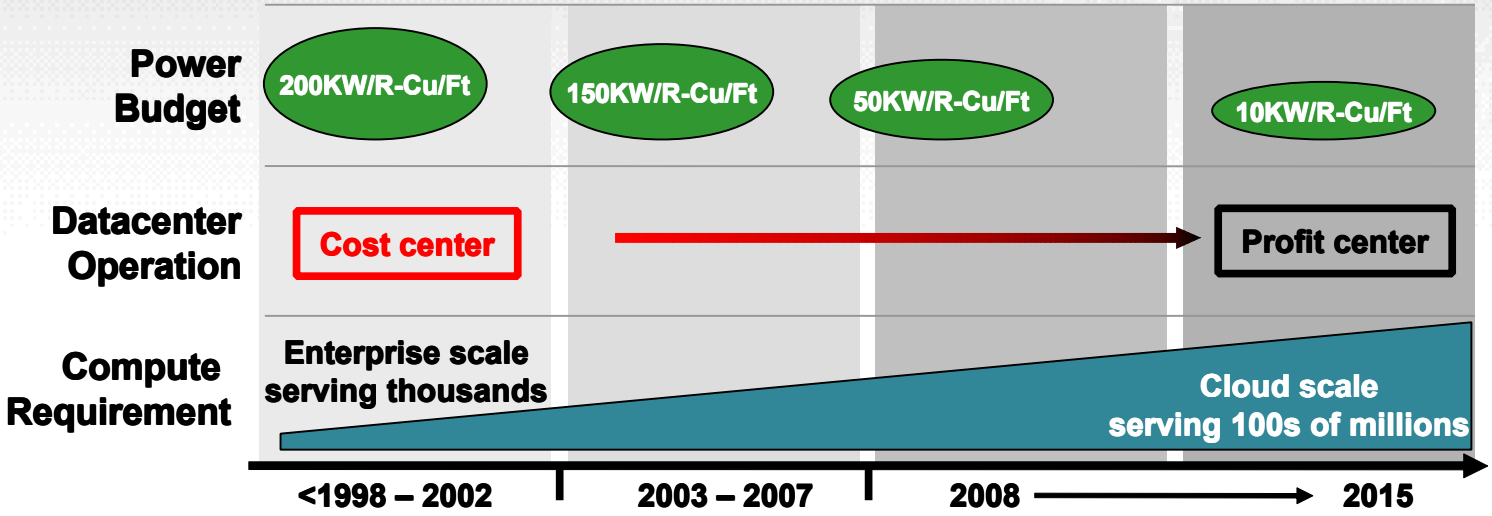

- Compute requirements increasing to serve millions of users
- Focus on maximizing the bottom line by reducing datacenter cost
- Power allowances decreasing due to size and delivery complexities

Power and cooling are top cost issues

- Power and cooling cost is growing faster than new server investment
- Only 3% of the power is used for computing
- IT organization will spend almost $1 on P&C for every $1they spend on new servers
- Developing a power efficient datacenter is more important than ever

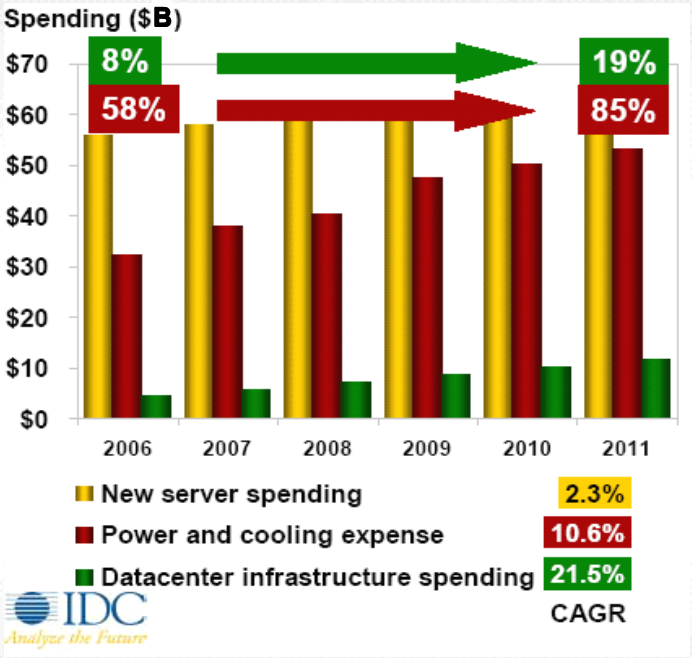

Current solutions are not solving cloud datacenters issues

"To build servers for companies like Facebook, and Amazon, and other people who are operating fairly homogeneous applications, the servers have to be cheap, and they have to be super power-efficient. The latest generations of server processors from Intel and AMD don't deliver the performance"

Jonathan Heiliger,
Facebook's VP of technical operations

"Problems cannot be solved by the same level of thinking that created them." Albert Einstein

- **Current few cores technologies fail to deliver for cloud**
  - Power too high
  - Performance not increasing fast enough
  - Integration too low
  - Cores are inefficient and continue to bloat

- **Manycore is the way to a new horizon**
  - Higher performance at much lower power
  - SoC integration to reduce cost and real estate
  - Standard programming models and improving

Manycore provides performance, low power and low cost

Current Technology

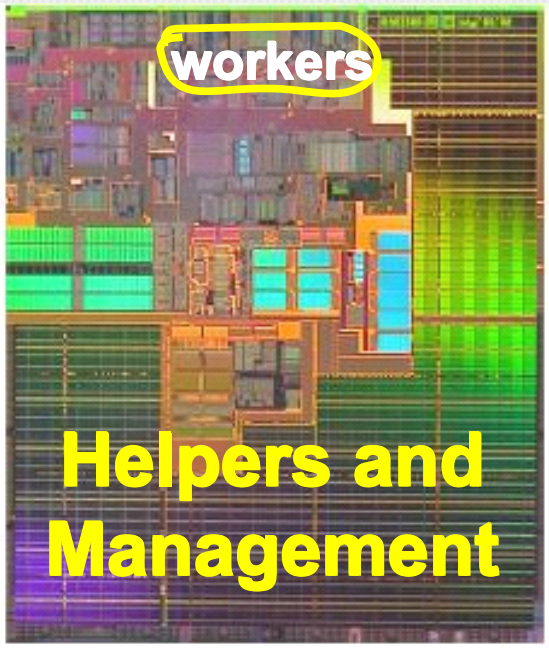

Madison Itanium2

- Less than 4% of area to ALUs
- High frequency, high power

Tilera Manycore

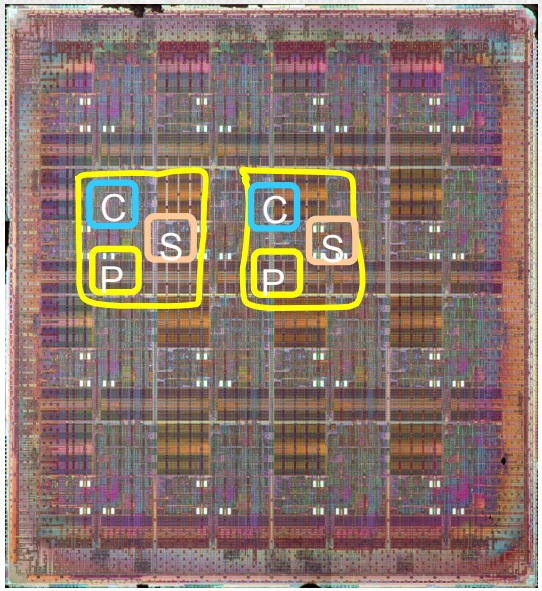

- More cores → more BOPS
- Lower frequency → low power Chandrakasan effect
- Simple, integrated → Lower cost

Tilera: the only technology delivering on the promise of multicore

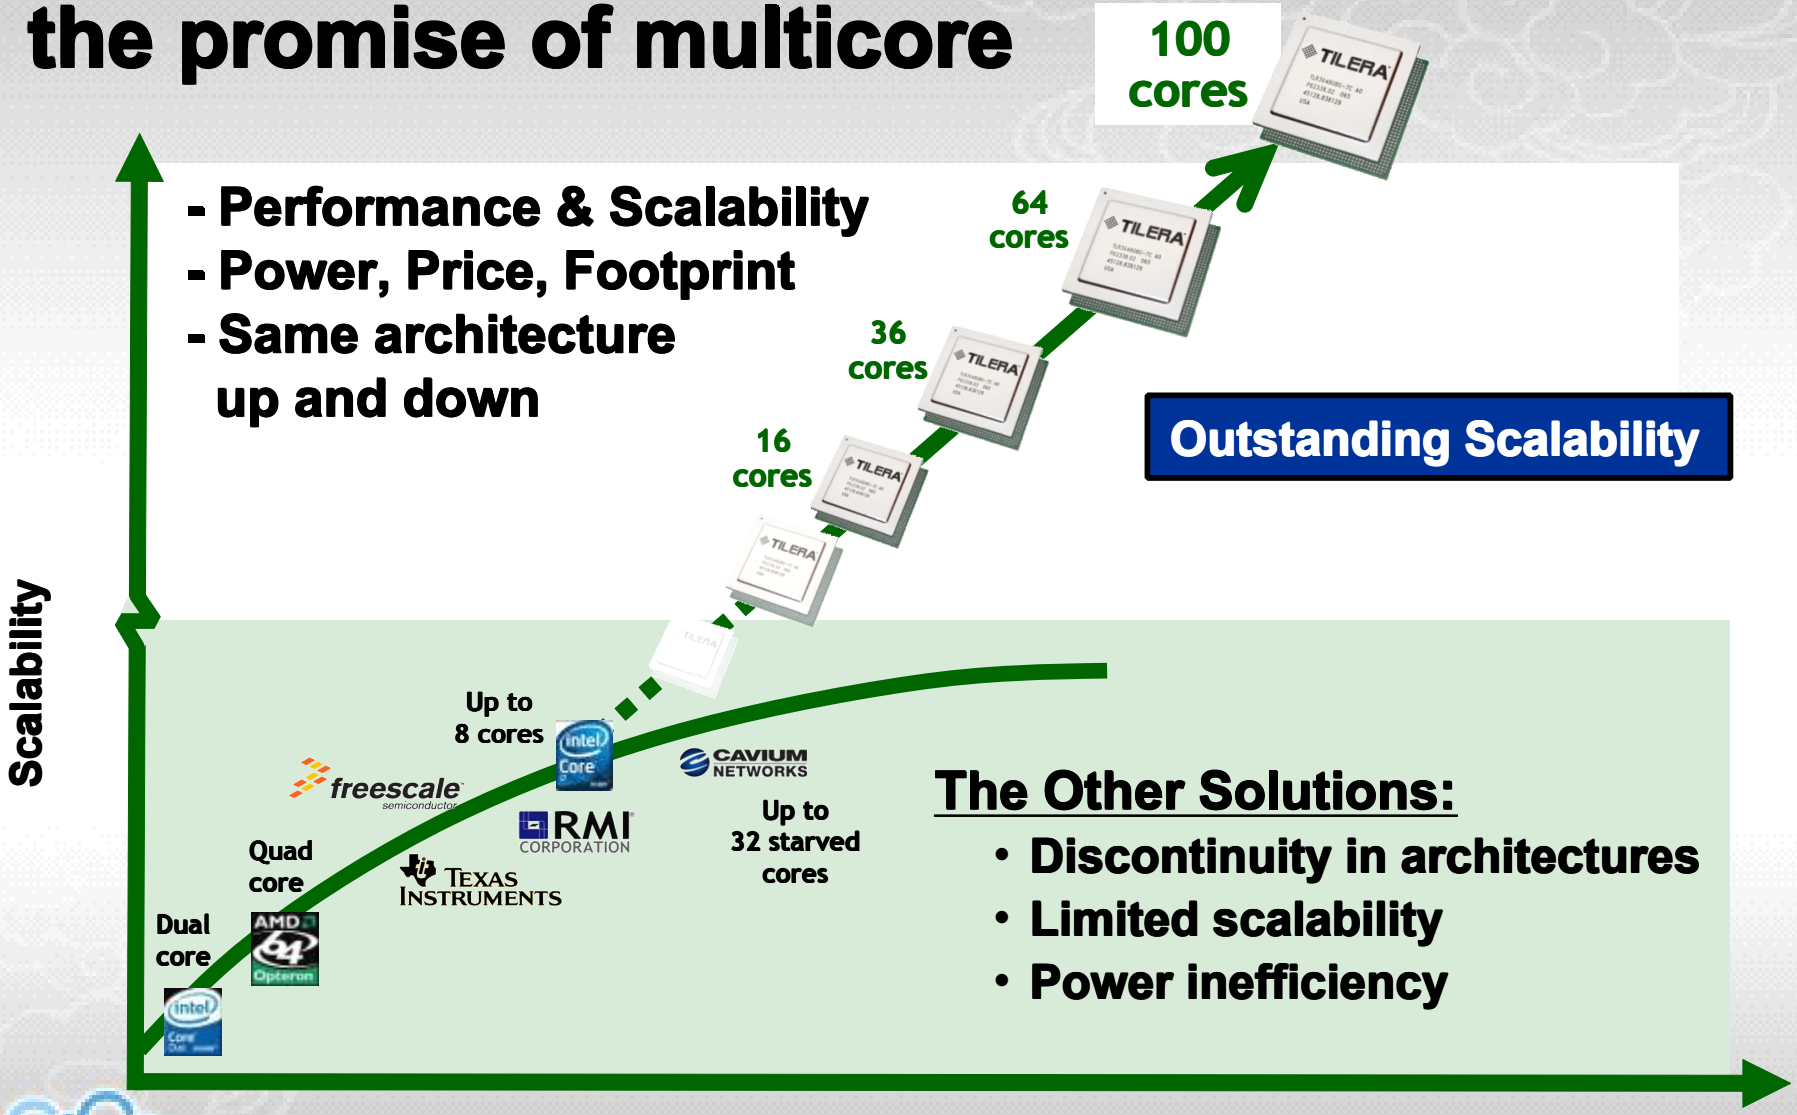

Proven best performance/Watt

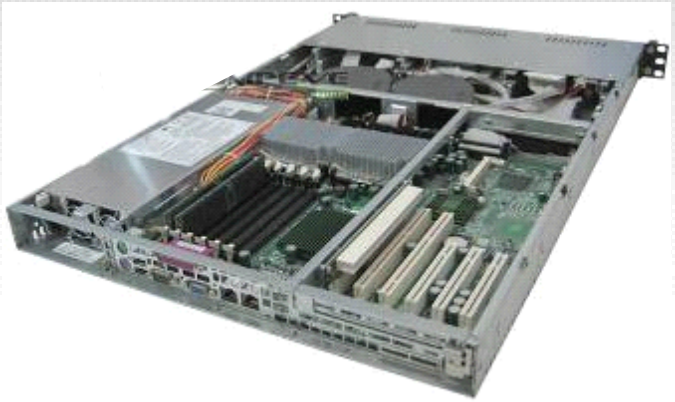

2X Quad XEON x86 Server
300W under load

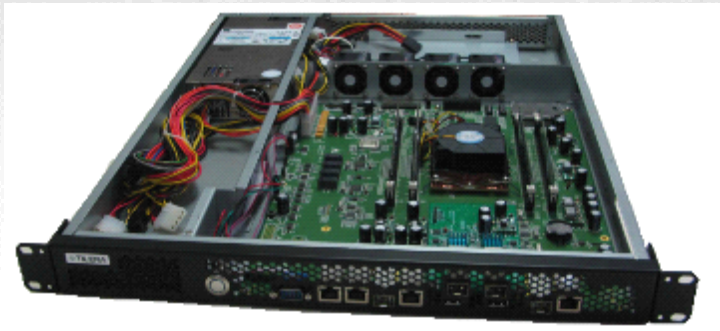

1X TILE*Pro* Server
40W under load
30W target (optimized server)

Measured by Tier 1 Server OEM running MemcacheD on its own Tilera and x86-based servers

One TILE*Pro*64 performance = Dual Quad XEON
Much Lower Power
7X Compute/Watt advantage

An order of magnitude better processing in a standard 5Kwatt rack

High efficiency x86 Server

- **12** 2U servers per rack
- **48** processors
- **196** cores
- **1,944** BOPS
- **96** Gbps I/O
- **5K** watt

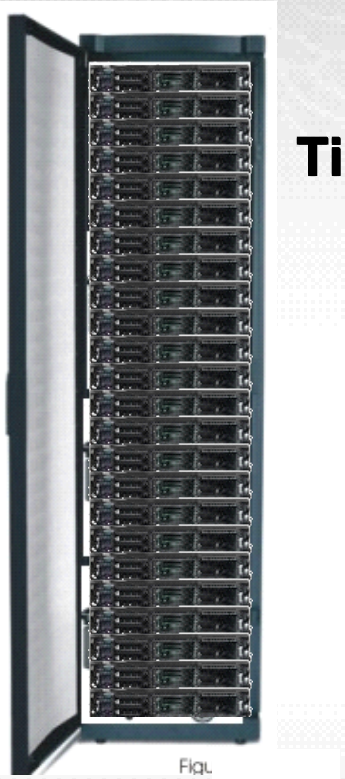

Tilera-based high density 2U production server

- **20** 2U servers per rack
- **160** processors
- **10,240** cores
- **20,000+** BOPS
- **3,500** Gbps I/O
- **5K** watts

New Tilera server ideal for cloud throughput applications
Best performance, I/O, power, and density
Complete utilization of power and space of a standard rack

Slashing Total Cost of Ownership for a given performance

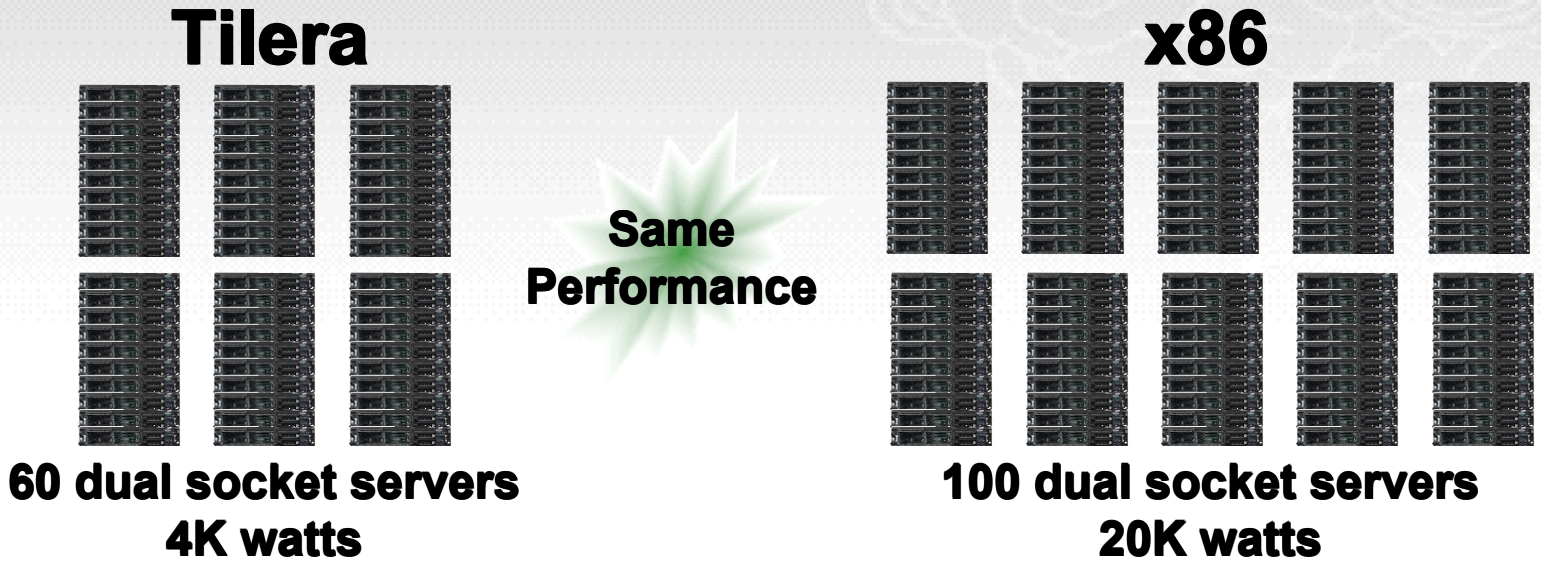

- Up to 40% CAPEX savings
- Up to 80% OPEX savings
- Slashing the TCO by up to 50%

Standard SW stack on Tier 1 servers

Markets	Front end Web; Database; Data-Mining
Infrastructure Apps	Lighttpd; APACHE; MySQL; Memcached; hadoop
Language Support	php; python; Ruby (A Programmer's Best Friend); Perl; Java
Compiler, OS Hypervisor	GNU gcc & g++; Xen; KVM; Commercial Linux Distribution
Tilera-based Servers	Low Power Prototype Server (now)  ~40 Watts; Tier 1 ODM/OEM; Production Server starting Q3 10

Focused on internet datacenters running LAMP stack

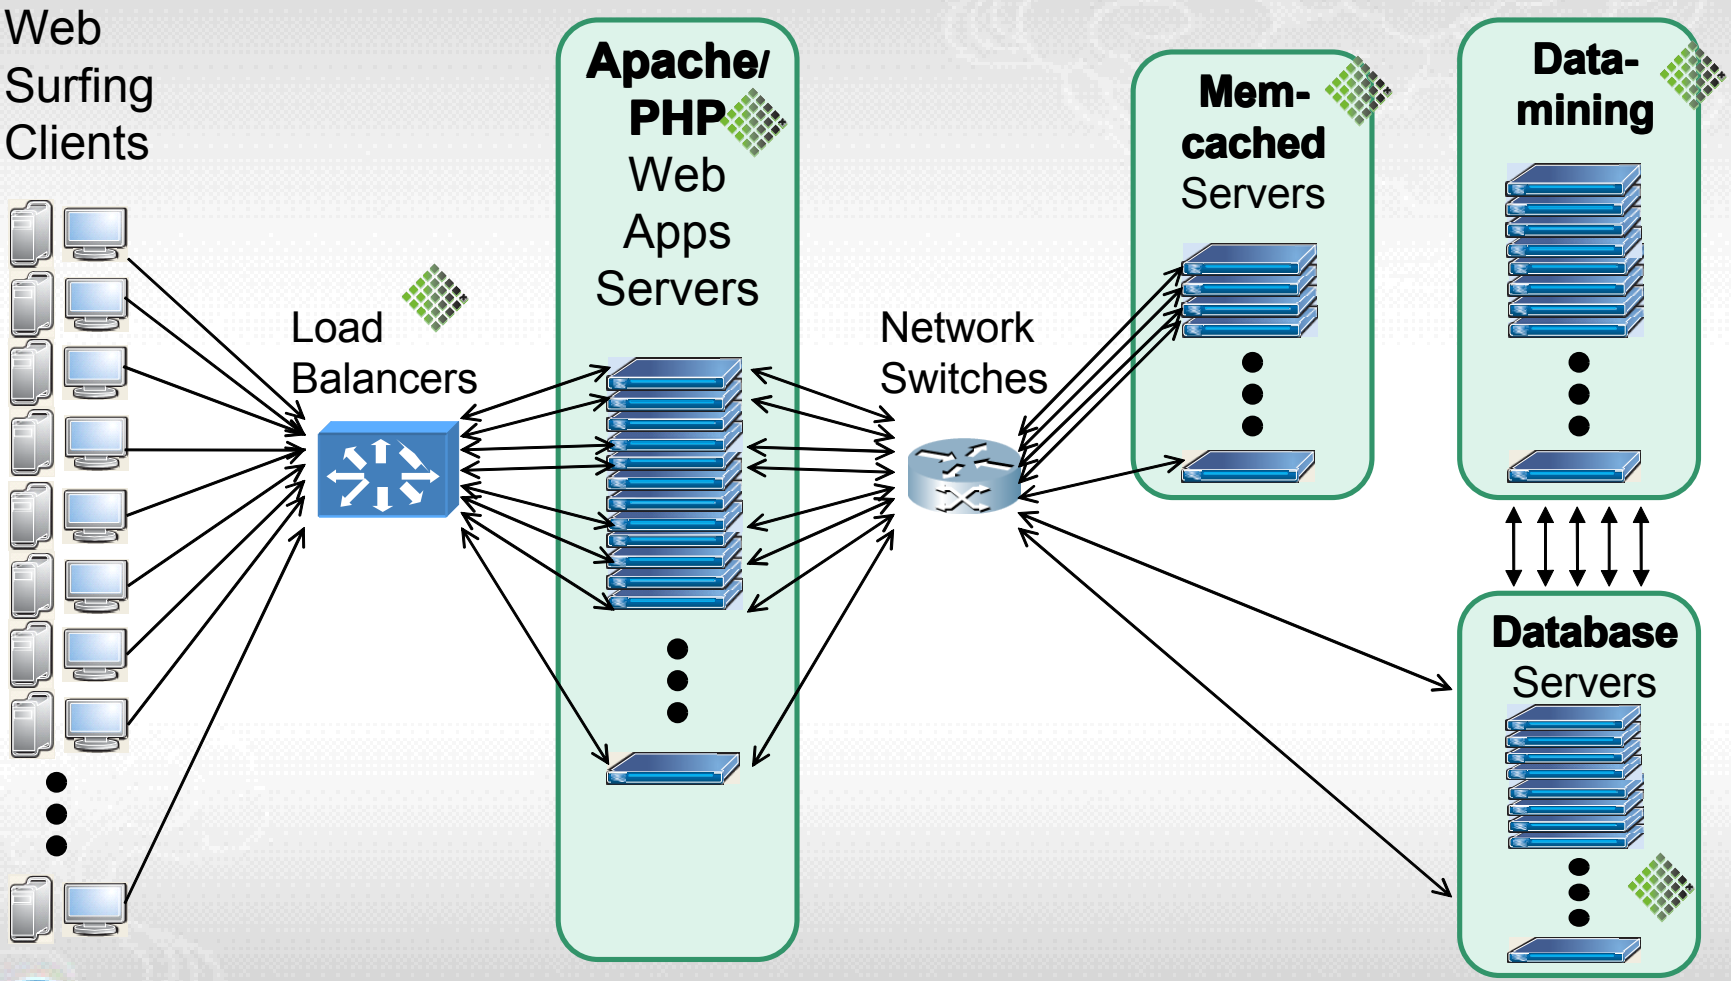

Scalable performance with complete cores and on-chip iMesh network

TILE Architecture

Complete cores with integrated cache

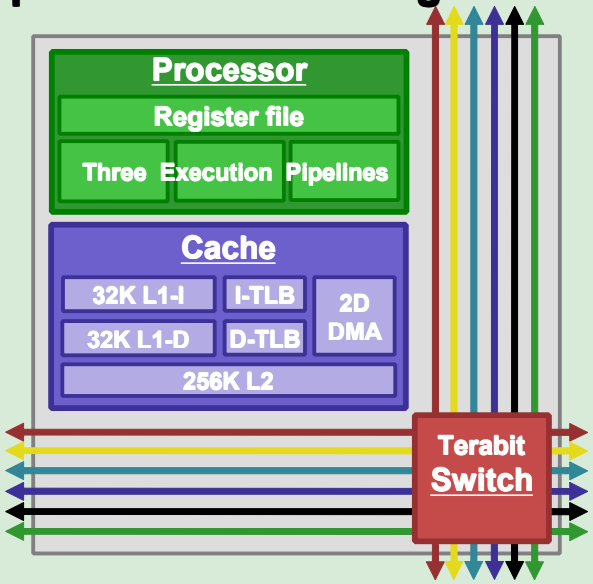

Tile = Core + Switch

2 Dimensional on-chip mesh network

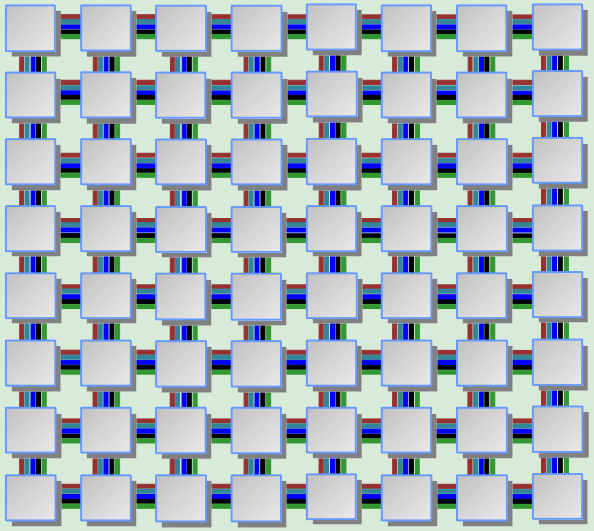

iMesh = 200 Tbps* on-chip bandwidth

*Based on TILE-Gx100

Tilera single core performance comparable to Atom & ARM Cortex-A9 cores

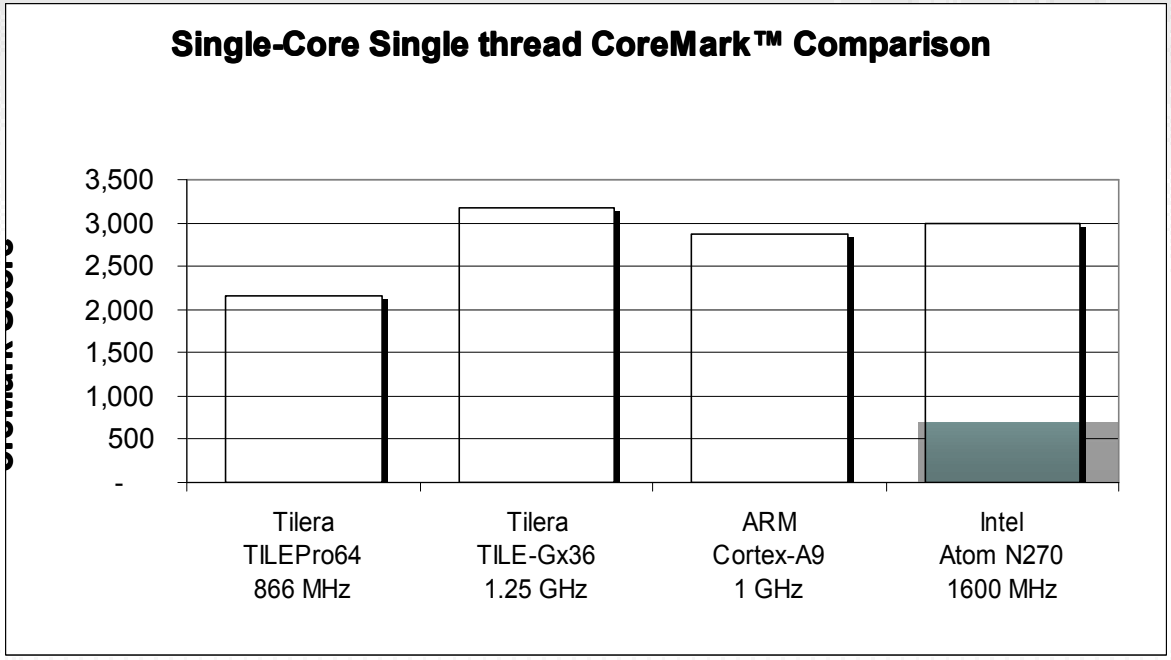

- Data for TILEPro, ARM Cortex-A9, Atom N270 is available on the CoreMark website http://coremark.org/home.php
- Telex and single thread Atom results were measured in Tilera labs
- Single core, single thread result for ARM is calculated based on chip scores

Tilera offers standards-based tools and software stack

Multicore Development Environment

Standards-based tools

Standard programming
- SMP Linux 2.6.26
- Java
- ANSI C/C++

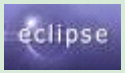

Integrated tools
- SGI or GCC compiler
- Standard gdb gprof
- Eclipse IDE

Innovative tools
- Multicore debug
- Multicore profile

Standard application stack

Application layer
- Open source apps
- Standard C/C++ libs

Operating System layer
- 64-way SMP Linux
- Zero Overhead Linux
- Bare metal environment

Hypervisor layer
- Virtualizes hardware
- I/O device drivers

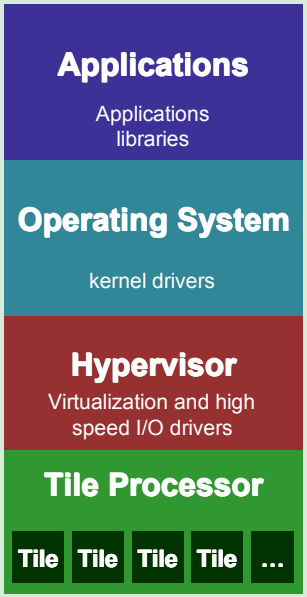

Example implementation on a Manycore-based server

- Complete functionality
  - Networking
  - Load balancer
  - Application
- Manycore values
  - Granularity: Complete server on one chip
  - Integration: I/O and memory controllers
  - Low power: 20watts
  - High performance
  - highly efficient communication on-chip

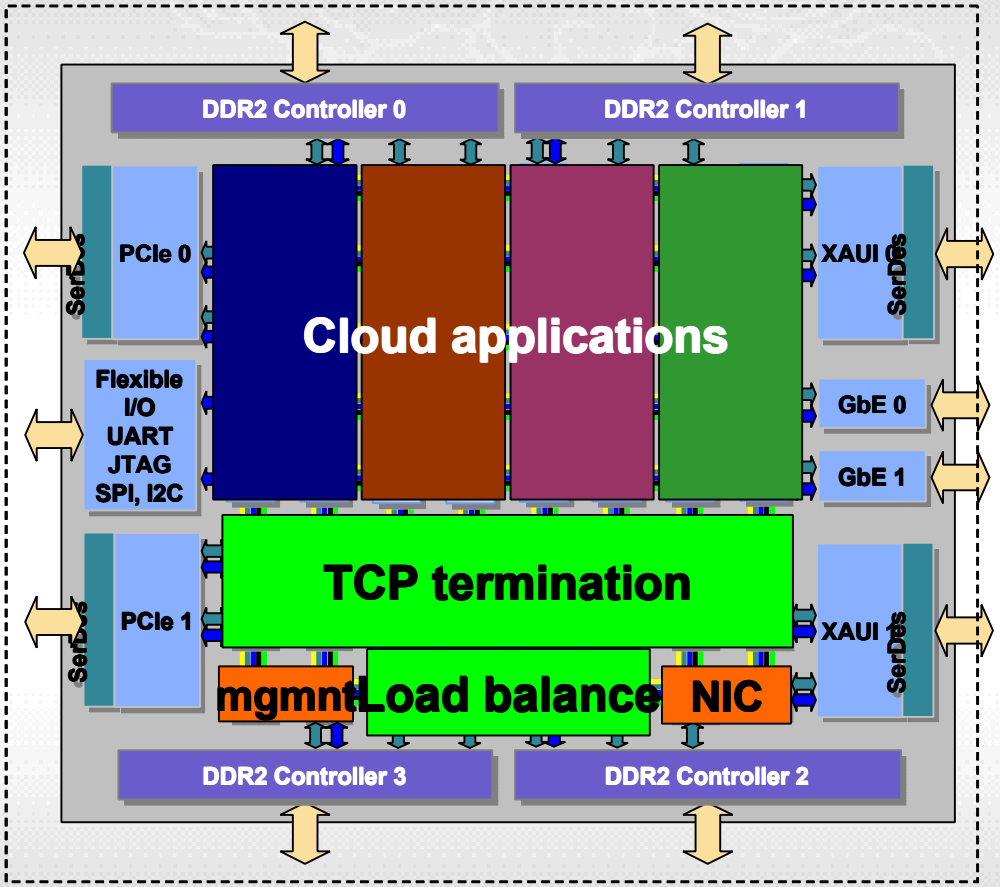

One SMP OS

A complete LAMP server at 40watts

- **Complete LAMP stack running on server**
  - Linux, Apache, PHP, MySQL
- **Serving standard cloud applications**
  - SugarCRM for enterprise
  - Gallery2 for photo sharing
- **Using standard Server management protocols**
  - Standard SNMP
  - MRTG
- **At 40 watts**
  - 40 watts power draw
  - Target 30 watts for optimized platform

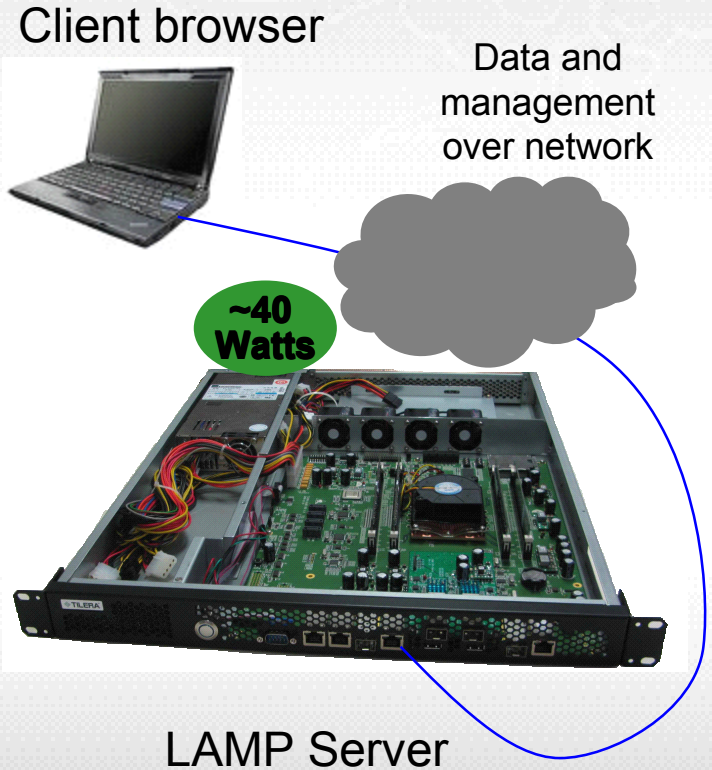

This is just the beginning…

Tilera's portfolio demonstrates the scale of many-core

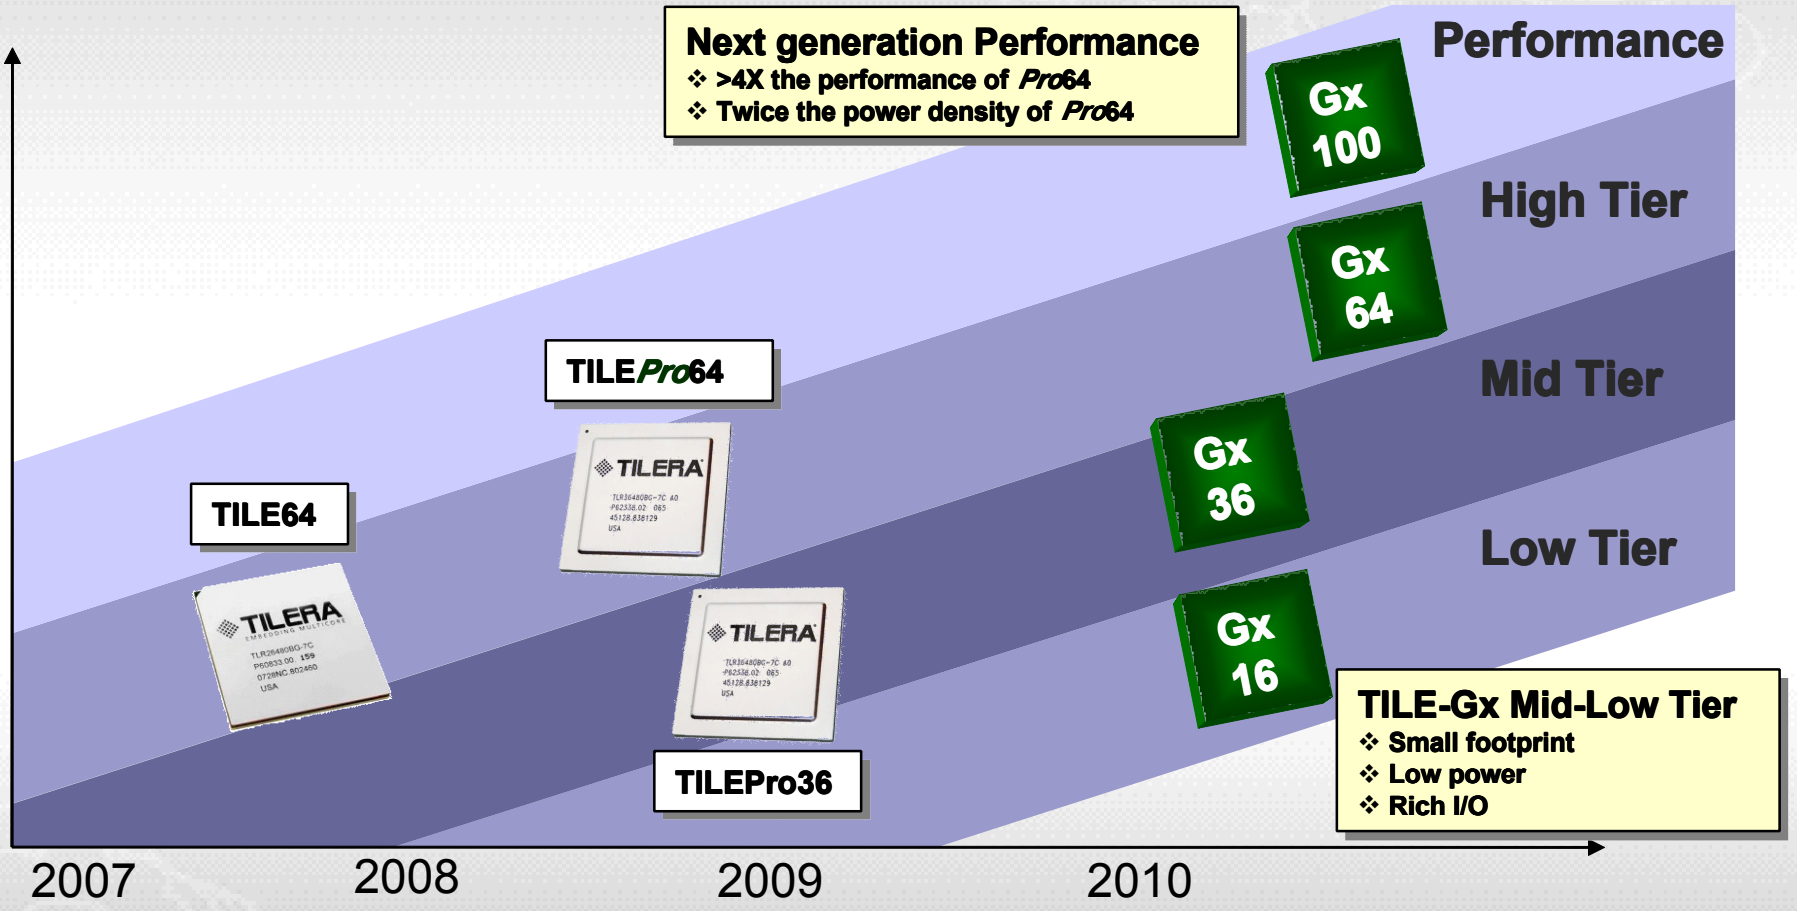

TILE-Gx100: High end system-on-a-chip Under 55 watts!

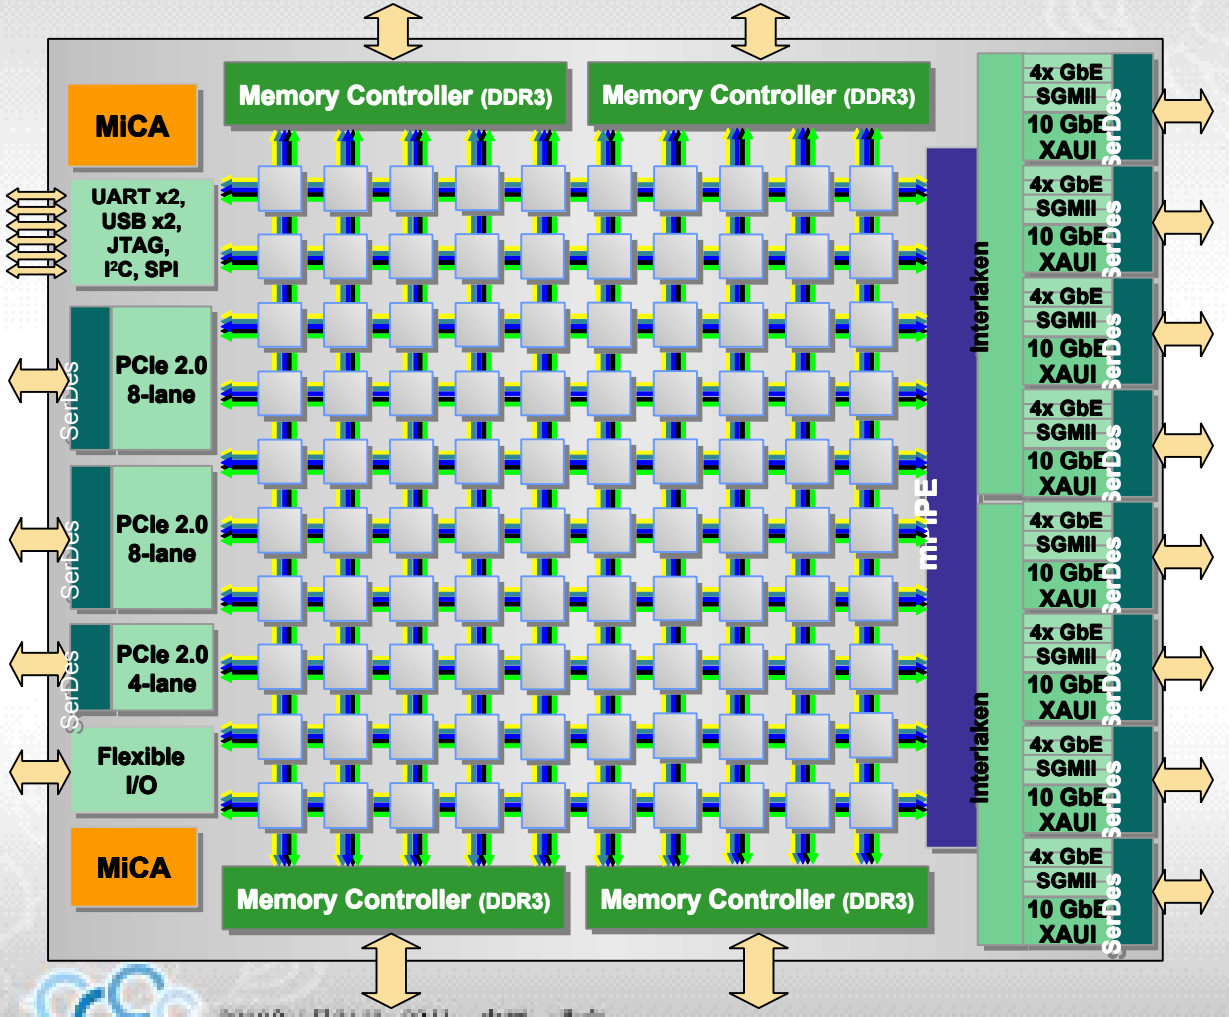

- 1.25GHz – 1.5GHz
- 32 MBytes total cache
- 200 Tbps iMesh BW

- >½ Tb/sec memory BW
- Addresses 1TByte of RAM
- 80-120 Gbps packet I/O
- 80 Gbps PCIe I/O

- Wire-speed packet engine
  - 120Mpps

- MiCA™ engines:
  - 40 Gbps crypto
  - 20 Gbps compress & 20 Gbps decompress

TILE-Gx36™: Midrange system-on-a-chip Under 22 watts!

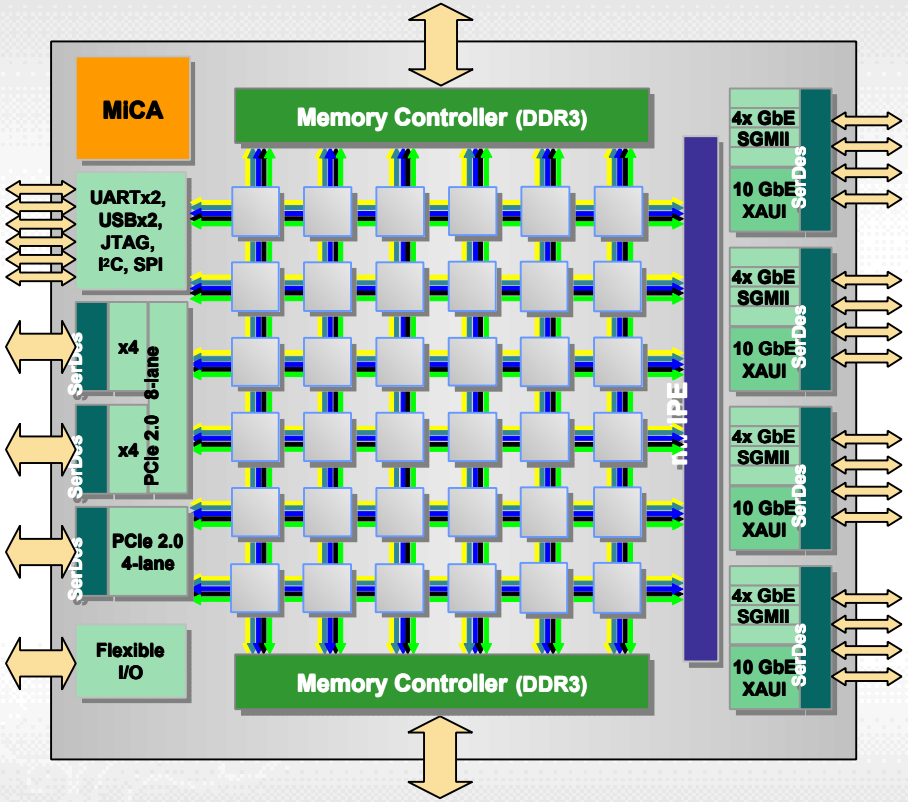

- 1.0, 1.25, 1.5 GHz speeds
- 12 MBytes total cache
- 66 Tbps iMesh BW

- 200 Gbps memory BW
- 40 Gbps total packet I/O
- 48 Gbps PCIe I/O

- Wire-speed packet engine
  - 60Mpps

- MiCA crypto engine
  - 20 Gbps crypto
  - 10 Gbps compress & 10 Gbps decompress

Summary

- Cloud faces critical performance, power and cost issues
- Tilera TILEPro processors deliver a proven solution to all three problems. Slashing TCO
- Tilera roadmap continues to deliver the promise of many core for performance, power and cost

China Cloud Computing Conference

2010年5月21日-22日 中国·北京

第二届中国云计算大会

Thank you