OpenPOWER™

OpenPOWER and the Roadmap Ahead

Brad McCredie
IBM Fellow, VP Power Development

Revolutionizing the Datacenter

Join the Conversation #OpenPOWERSummit

Quarter Century of POWER

Legacy of Leadership Innovation
Driving Client Value

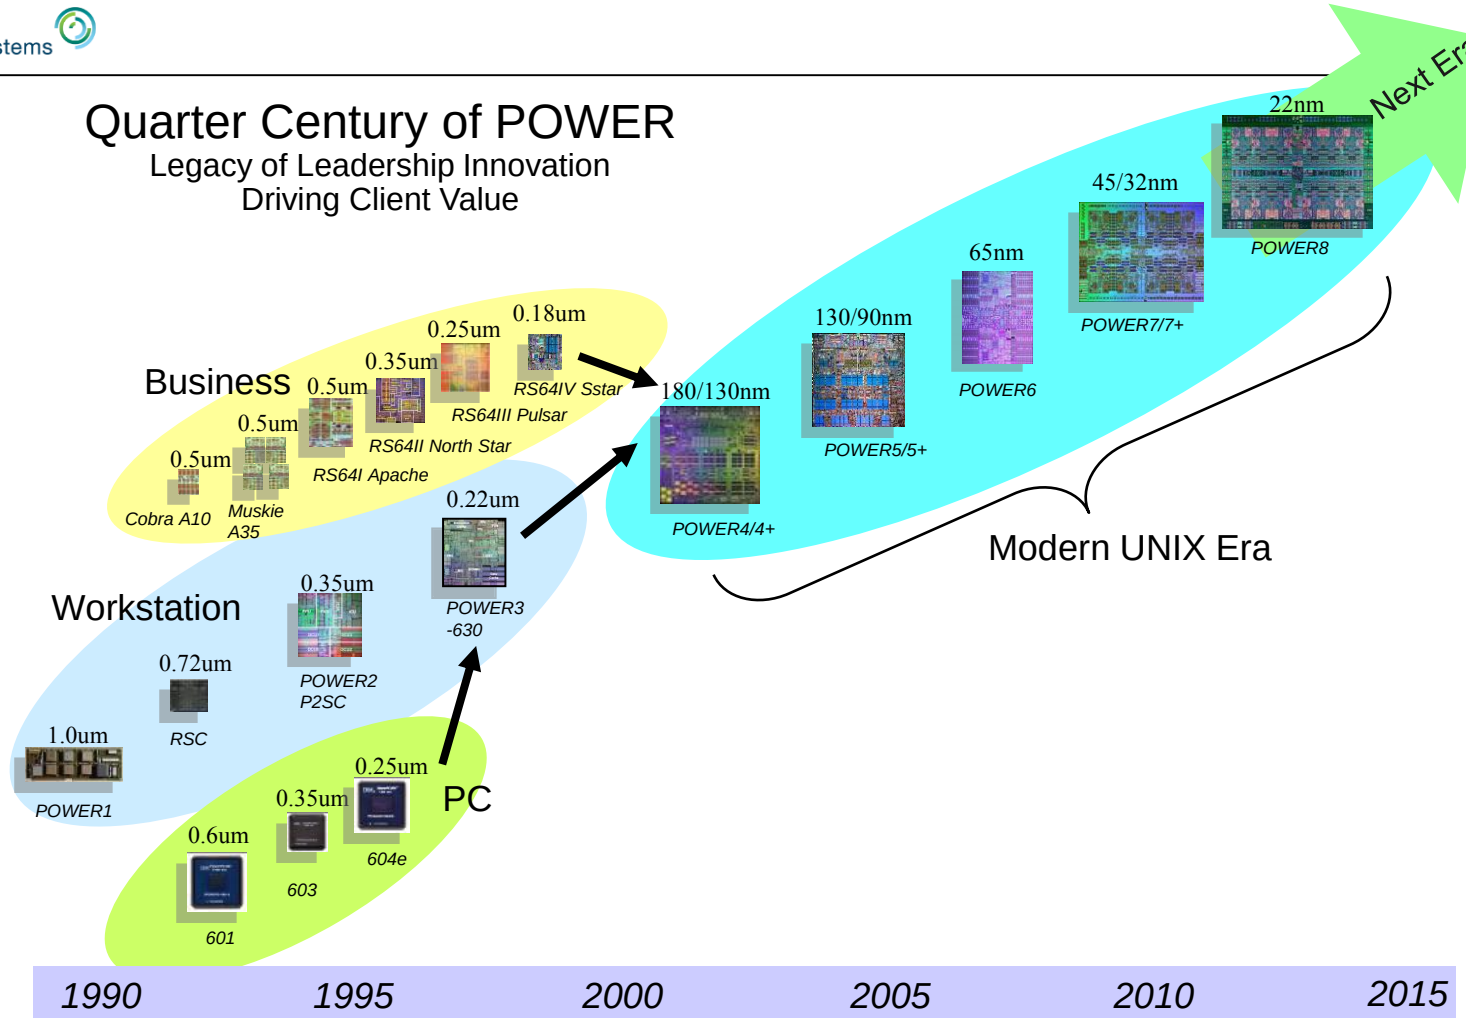

Power Systems

POWER Processor Roadmap

IBM

Focus on Enterprise Technology and Performance Driven				Focus on Scale-Out and Enterprise Cost and Acceleration Driven						Future
POWER6 Architecture		POWER7 Architecture		POWER8 Architecture		POWER9 Architecture			Partner Chip POWER8/9	POWER10
2007 POWER6 2 cores 65nm New Micro-Architecture New Process Technology	2008 POWER6+ 2 cores 65nm+ Enhanced Micro-Architecture Enhanced Process Technology	2010 POWER7 8 cores 45nm New Micro-Architecture New Process Technology	2012 POWER7+ 8 cores 32nm Enhanced Micro-Architecture New Process Technology	2014 POWER8 12 cores 22nm New Micro-Architecture New Process Technology	2016 POWER8 w/ NVLink 12 cores 22nm Enhanced Micro-Architecture With NVLink	2017 P9 SO 24 cores 14nm New Micro-Architecture Direct attach memory New Process Technology	TBD P9 SU TBD cores 14nm Enhanced Micro-Architecture Buffered Memory	TBD	2018 - 20 P8/9 SO 10nm - 7nm Existing Micro-Architecture Foundry Technology	2020+ New Micro-Architecture New Technology
High Frequency Enhanced RAS Dynamic Energy Management		Large eDRAM L3 Cache Optimized VSX Enhanced Memory Subsystem		Optimized for Data-Centric Workloads Integrated PCIe CAPI Acceleration / I/O		Scale-Out Datacenter TCO Optimization Scale-up performance Optimization Acceleration Enhancements to CAPI and NVLINK Modularity for OpenPOWER			OpenPOWER Ecosystem Design Targeting Partner Markets & Systems Leveraging Modulatrity	New Features and Functions

Price, performance, feature and ecosystem innovation →

POWER8 Processor

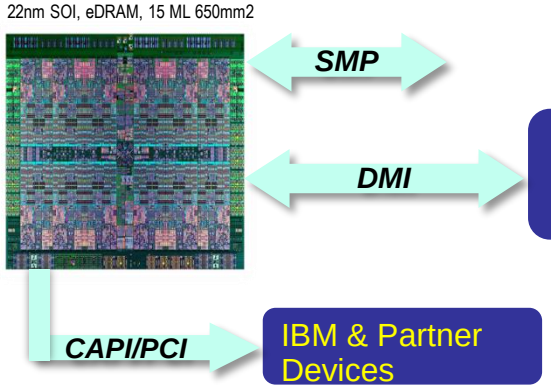

Cores
- 12 cores / 8 threads per core
- TDP: 130W and 190W
- 64K data cache, 32K instruction cache

Accelerators
- Crypto & memory expansion
- Transactional Memory

Caches
- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3

Memory Subsystem
- Memory buffers with 128MB Cache
- ~70ns latency to memory

Bus Interfaces
- Durable Memory attach Interface (DMI)
- Integrated PCIe Gen3
- SMP Interconnect for up to 4 sockets

Server Class Memories (SCM)
- First Functioning Demo of SCM in an Enterprise system
- 15x better than SSD
- Natively non-volatile ST-MRAM DIMMs (Everspin)
- Avoids NVDIMM complications (DRAM to FLASH and back)
- No Supercaps with umbilicals required

Coherent Accelerator Processor Interface (CAPI)

Virtual Addressing
- Accelerator can work with same memory addresses that the processors use
- Pointers de-referenced same as the host application
- Removes OS & device driver overhead

Hardware Managed Cache Coherence
- Enables the accelerator to participate in “Locks” as a normal thread
- Lowers Latency over IO communication model

6 Hardware Partners developing with CAPI

Over 20 CAPI Solutions
- All listed here http://ibm.biz/powercapi

Examples of Available CAPI Solutions
- IBM Data Engine for NoSQL
- DRC Graphfind analytics
- Erasure Code Acceleration for Hadoop

OpenPOWER
DEVELOPER CHALLENGE

Two tracks to challenge and win:

1. The Open Road Test
   - Port and optimize for OpenPOWER
   - Go faster with accelerators *(optional)*
2. The Spark Rally
   - Train an accelerated DNN and recognize objects with greater accuracy
   - Show you can scale with Spark

SuperVessel Cloud

OpenPOWER™

Key Dates

Register today
openpower.devpost.com

Sun May 1st:
Submission periods opens

Tue Aug 2nd:
Submission period closes

Grand prizes include a trip to **Supercomputing 2016**
Other prizes include **iPads, Apple Watches**

Join the conversation at #OpenPOWERSummit

POWER8 with NVLink

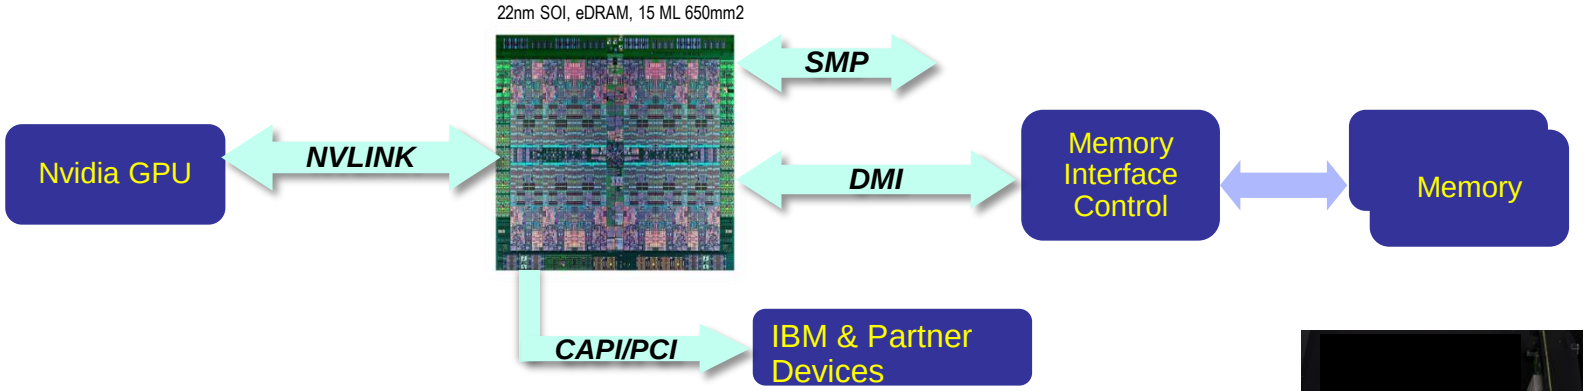

- NVLink High Speed CPU <-> GPU Interconnect
- 160+ GigaBytes per second bi-directional
- 5-12x faster than PCIe Gen3 x16

- Nvlink Accelerator Lab
  accellab@us.ibm.com

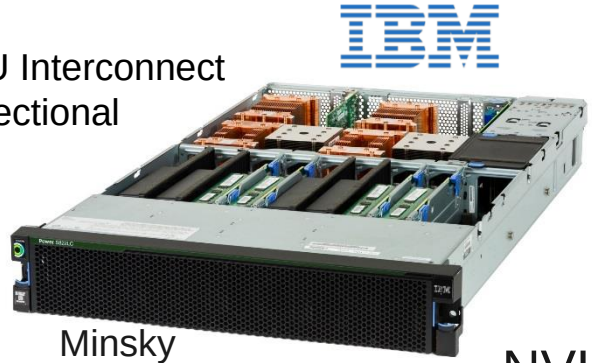

Minsky

Zoom

Inventec

NVLink POWER Systems

POWER9 SO with Advanced Accelerator Attach

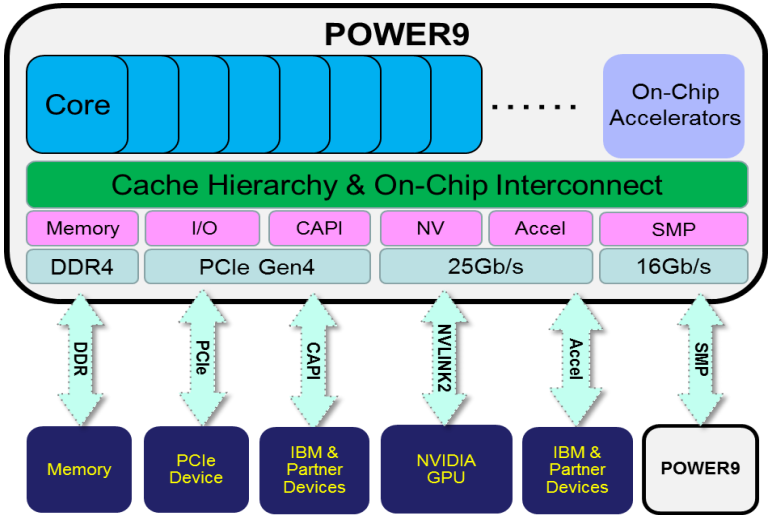

24 newly designed POWER9 cores

- Leveraging execution slices for improved performance on cognitive, analytic, and big-data applications

Large, low-latency, eDRAM cache for big datasets

Global Foundries 14HP finFET technology with eDRAM

Cloud-focused innovation in Energy Efficiency, Security, and Quality of Service

State-of-the-art IO subsystem using PCIe Gen4

Leadership platforms for hardware acceleration

- High bandwidth, GPU interconnect (NV link2.0)
- Next-generation CAPI2.0 interface for coherent accelerator and storage attach
- On-chip compression & cryptography accelerators
- New 25Gb/s advanced accelerator attach bus

1st chip in POWER9 family

- Optimized for 2 socket scale out servers & hyperscale datacenters
- DDR4 direct attach memory channels

Full POWER9 family will address a broad range of scale out & enterprise servers

POWER Processor Roadmap

Denser Roadmap
- Not tied to technology
- More teams developing

Performance and cost/performance leadership

Leadership accelerator attach technologies including
- GPUs
- Advanced Storage
- Advanced Networking
- FPGAs

Expanded focus includes Enterprise and Hyperscale datacenters

All of these technologies will be made available to our partners