OpenPOWER & IBM Power Systems

The best platform for PostgreSQL workload

René Akeret
Systems Architect
IBM Switzerland

akeret@ch.ibm.com

Version 1.6

PostgresPURE and Power Systems

Higher Performance

PostgreSQL tps

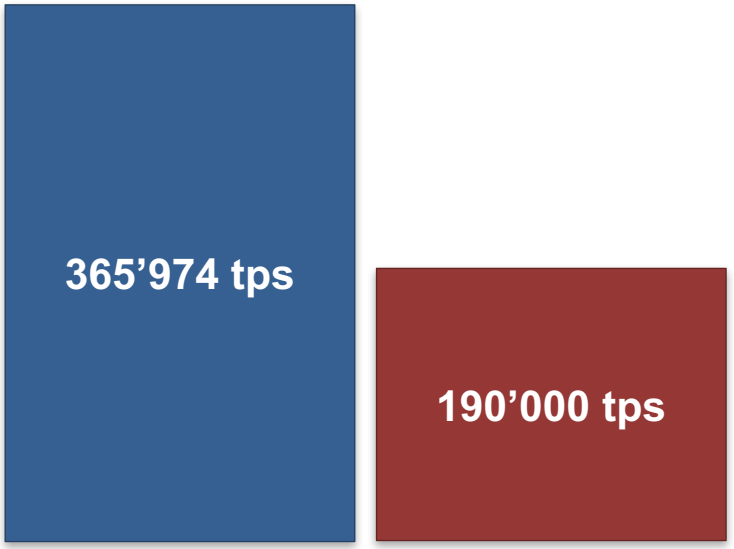

IBM Power S822L
- 20c / 160t

x86 Intel Xeon
- 60c/ 120t

Lower TCO

Price per Transaction

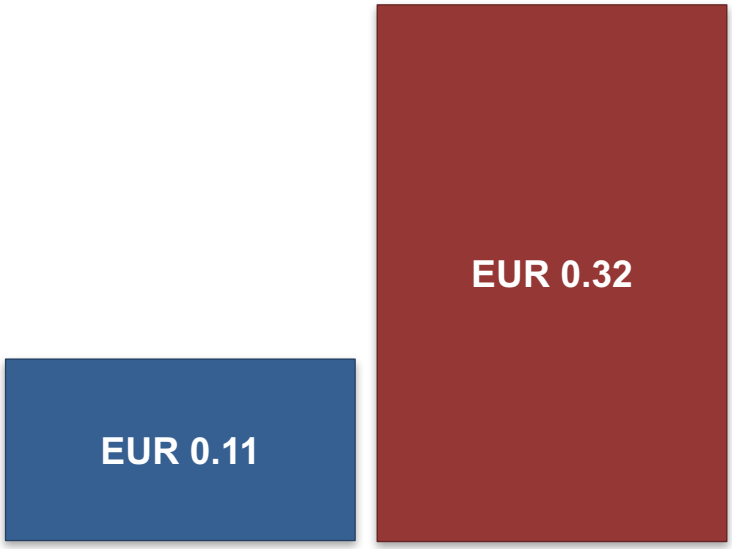

IBM Power S822L
- 20c / 160t

x86 Intel Xeon
- 60c/ 120t

For detailed benchmark results, see Benchmark Report “Splendid Data PostgresPURE on IBM Power System S822L” (http://www.splendiddata.com/benchmark-results-postgrespure-ibm-powerlinux-8/).

Agenda

- POWER8 Processor Technology
- OpenPOWER & IBM Power Systems
- PostgreSQL and POWER - The winning team
- Summary

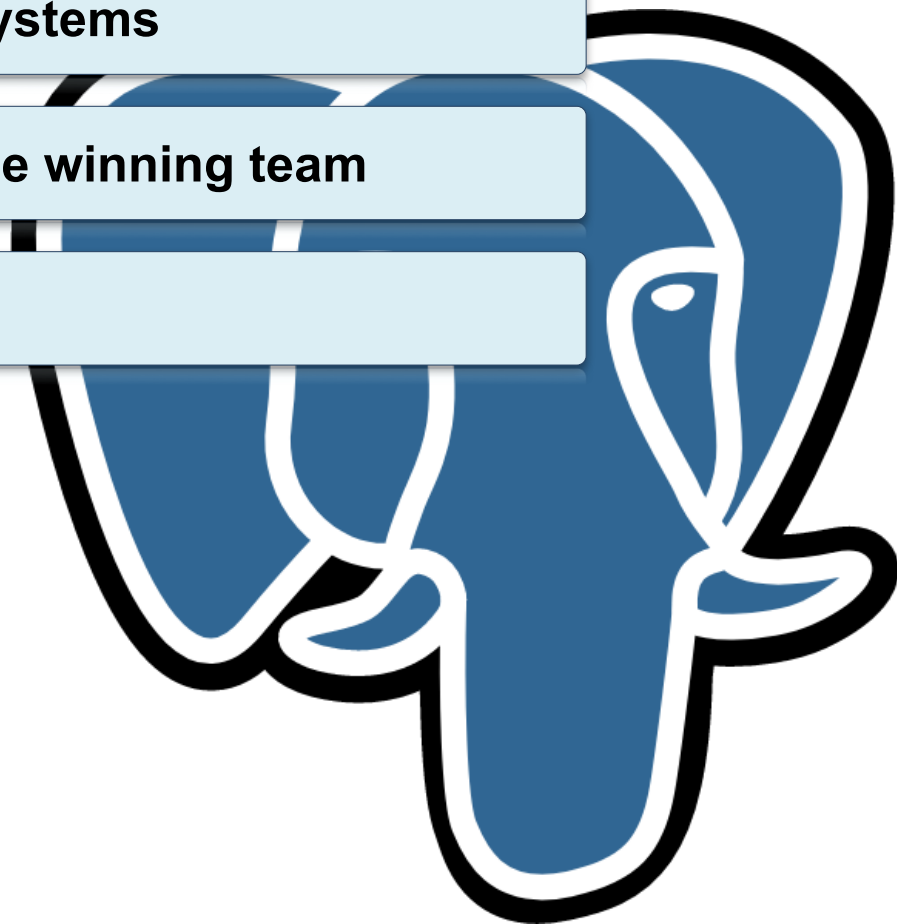

Quarter Century of POWER

Legacy of Leadership Innovation Driving Client Value

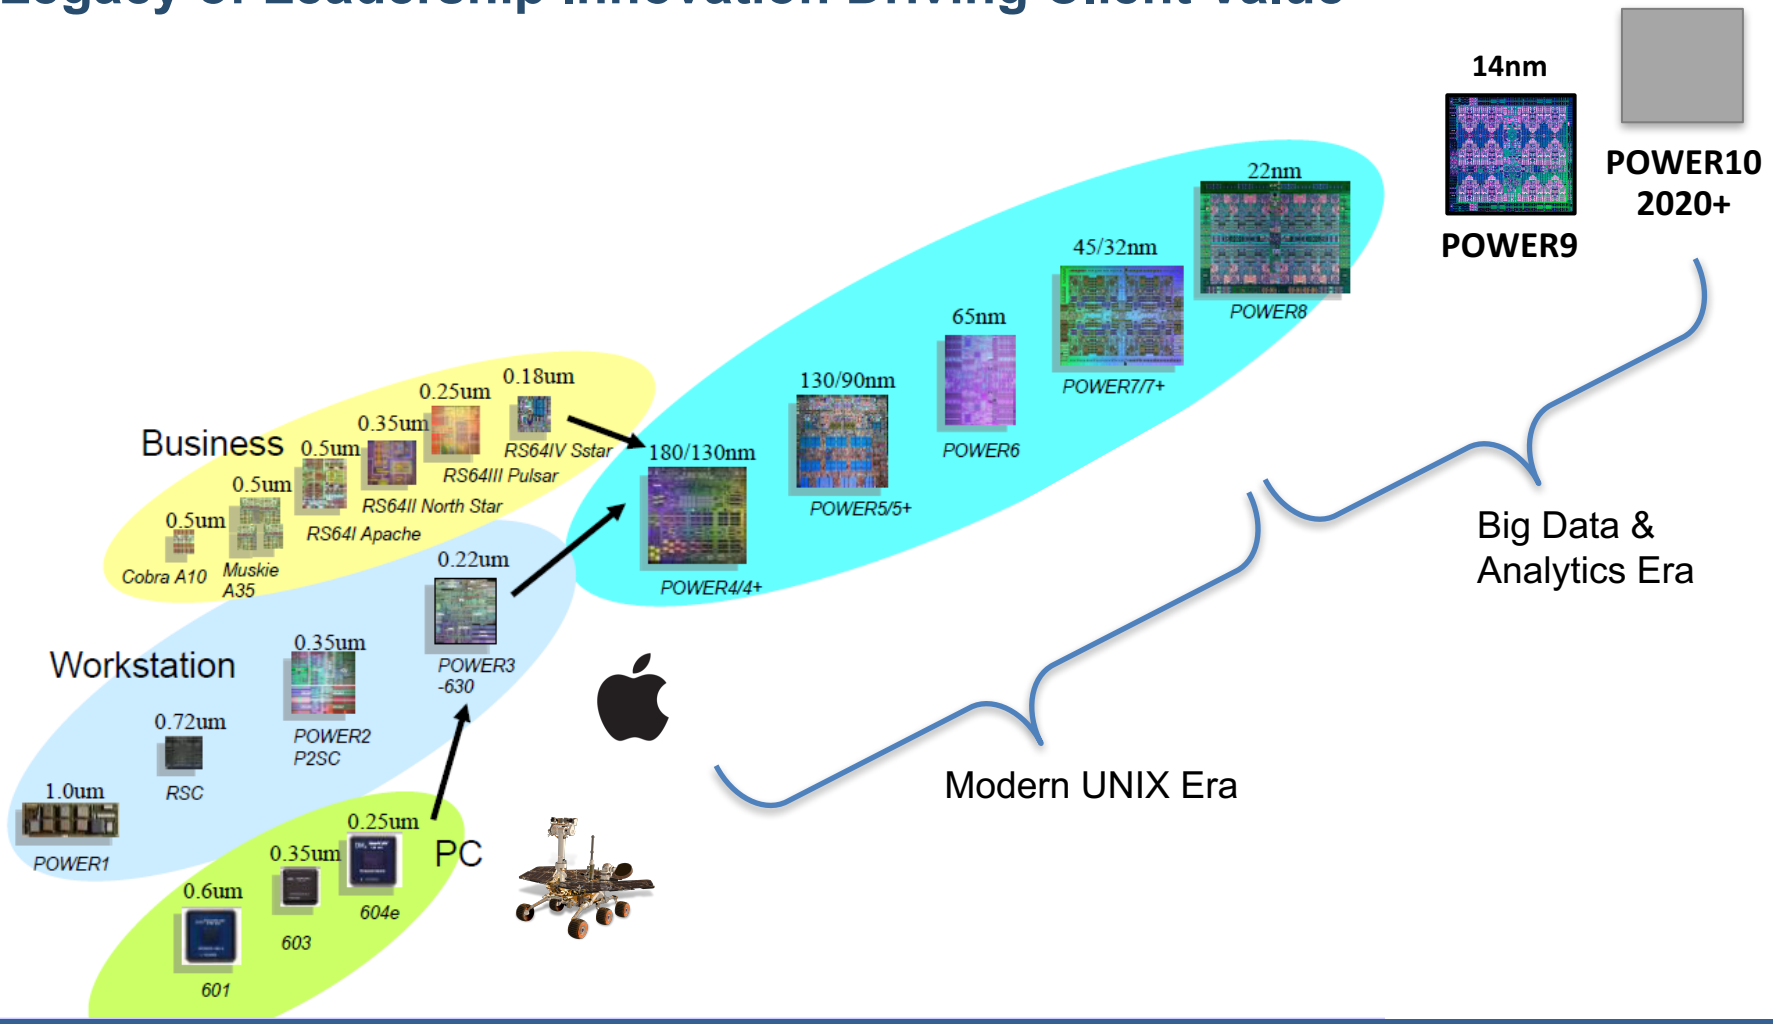

POWER8 Four Level Cache Design

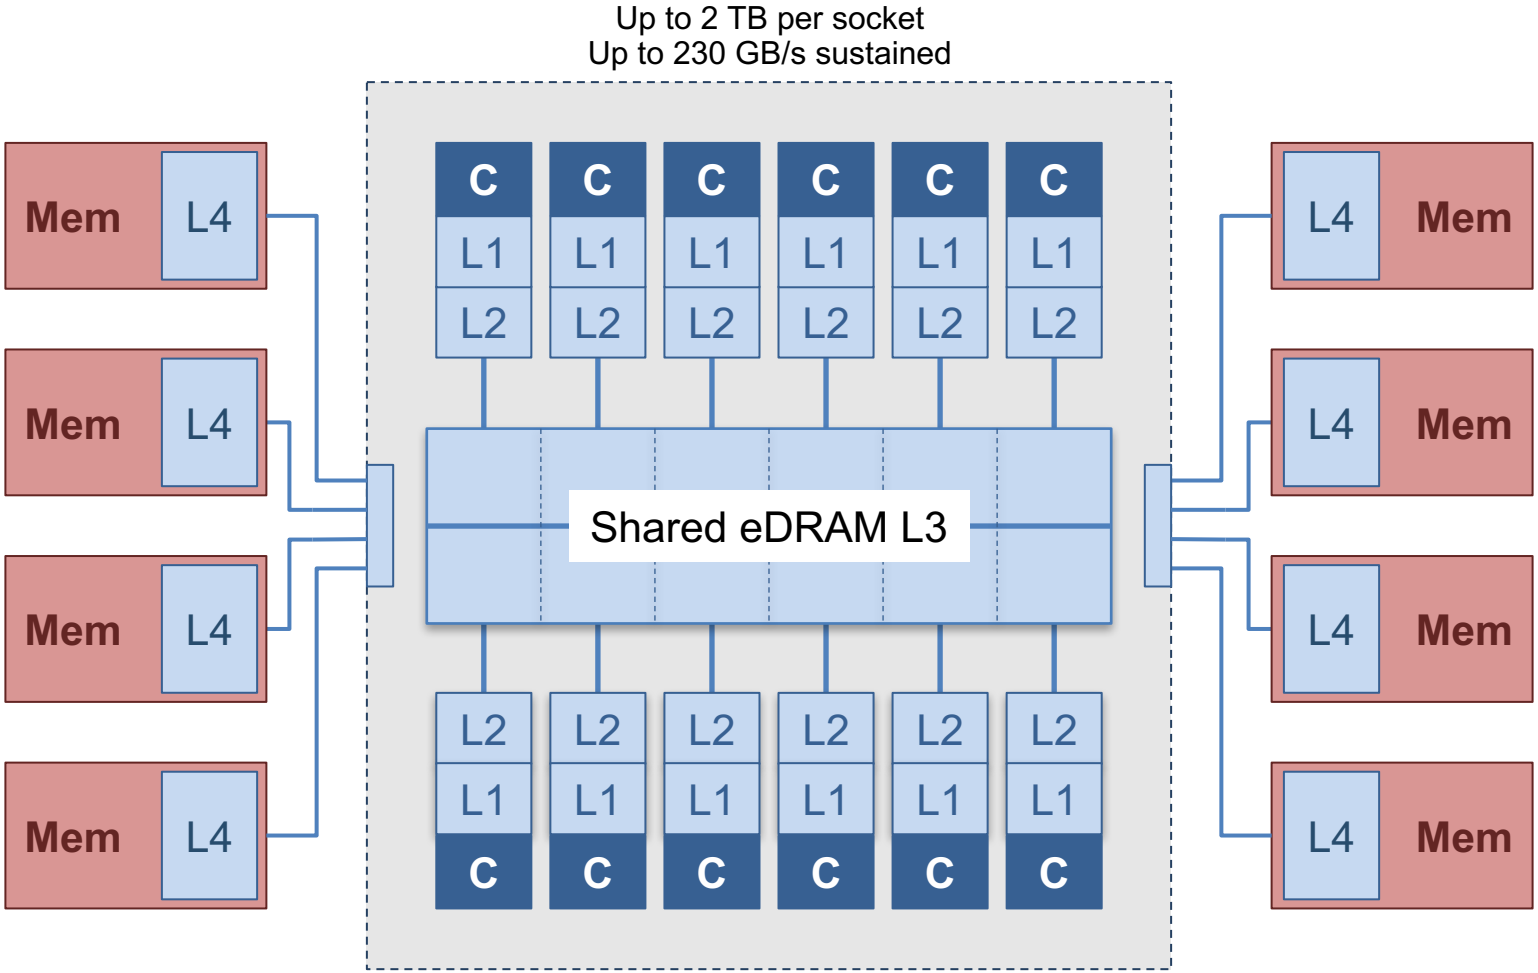

Cache is Critical to Good Performance

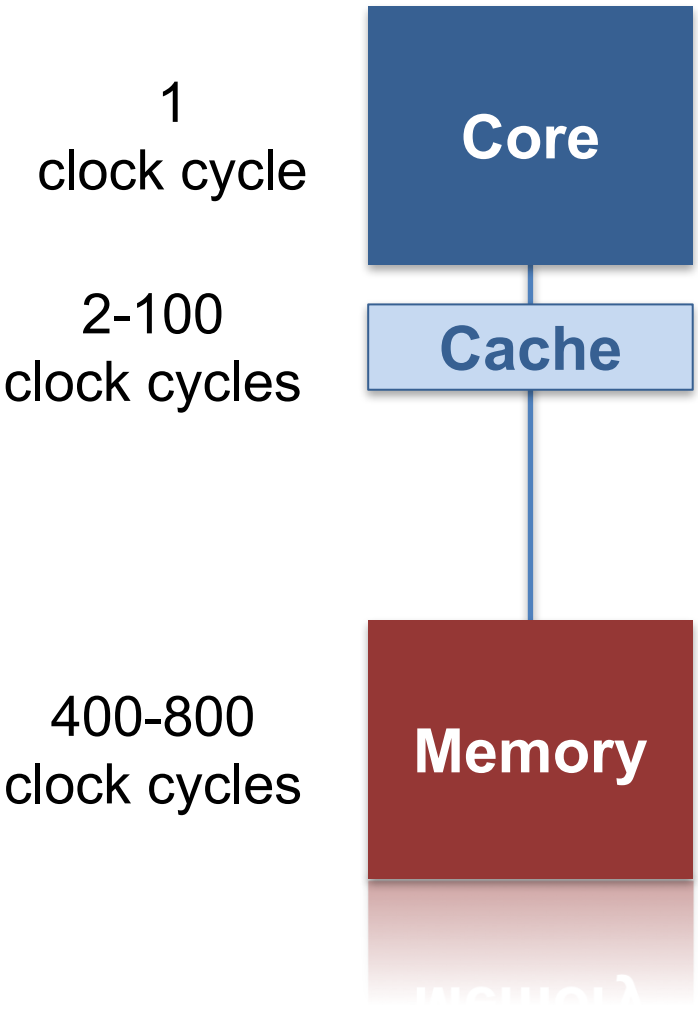

Memory is **slow** relative to cache

PostgreSQL database, server point of view

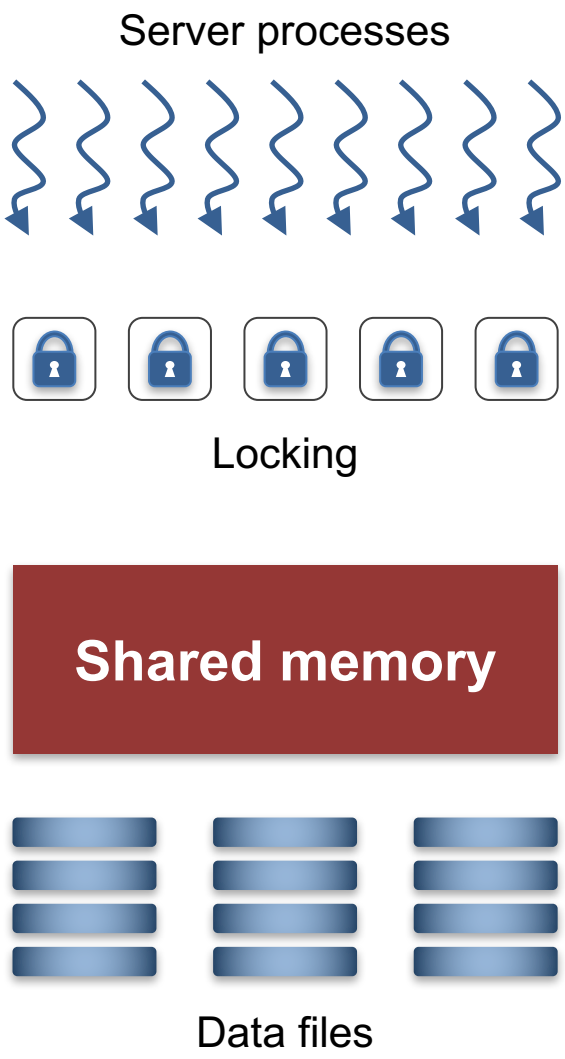

Core performance
- Single thread
- Multithread
- Throughput

SMP interconnect
- Bandwidth
- Latency

Memory performance
- Bandwidth
- Latency

I/O performance
- Bandwidth
- Latency

Exploits POWER8 Hardware - Speed

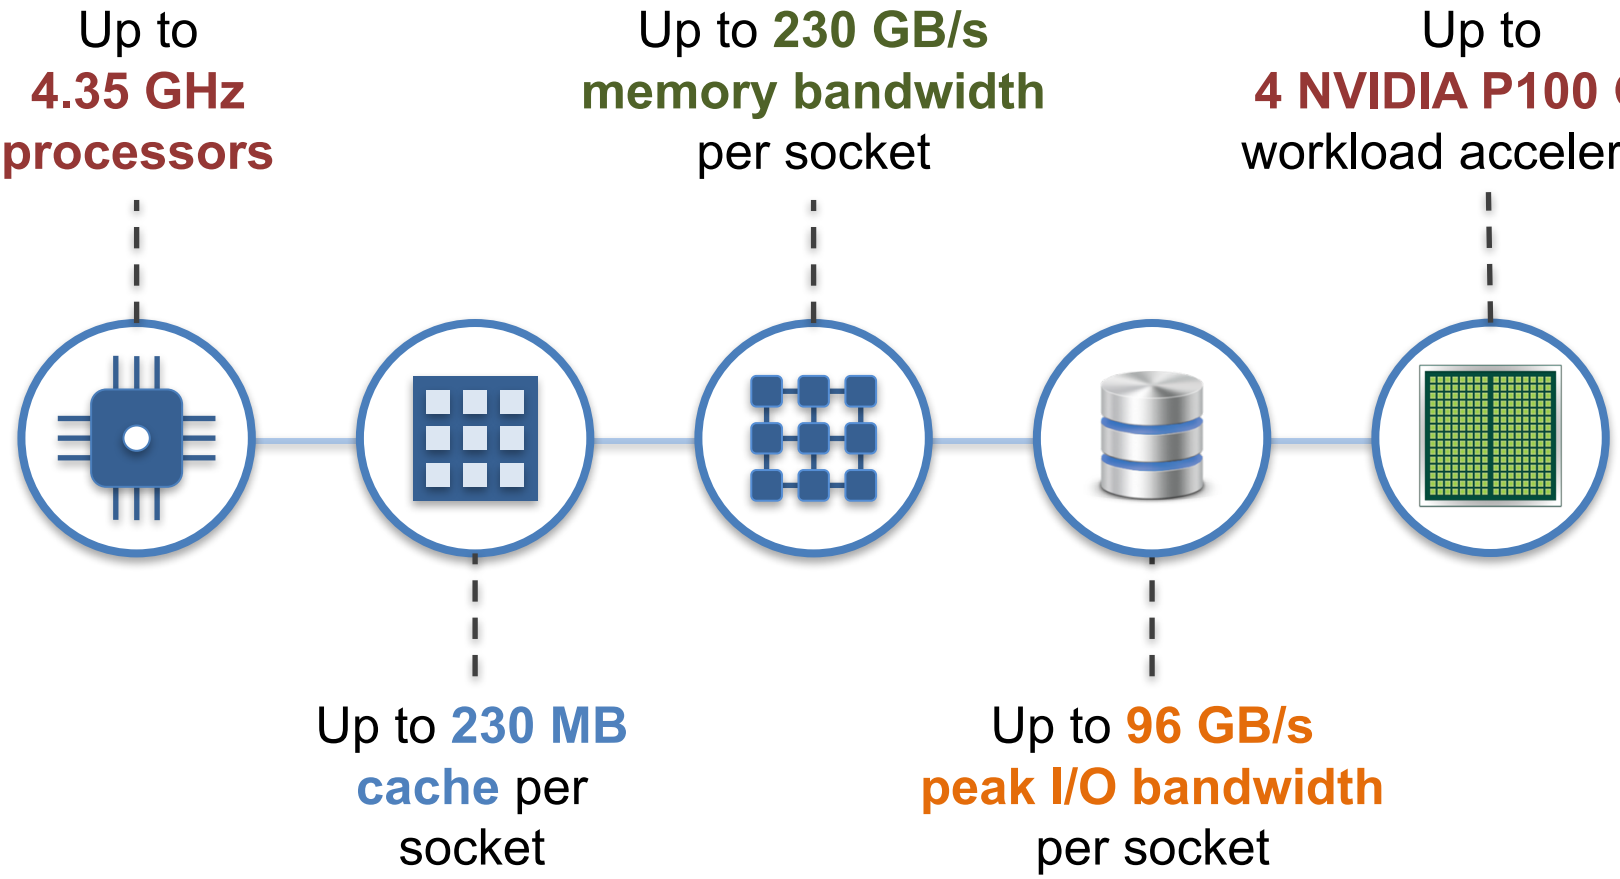

POWER8 developed for Big Data & Analytics

4X

threads per core vs Intel
(up to 1536 threads per system)

4X

memory bandwidth vs Intel
(up to 32TB of memory)

5X

more cache vs Intel
(up to 230MB cache per socket)

Processors
flexible, fast execution of analytics algorithms

Memory
large, fast workspace to maximize business insight

Cache
ensure continuous data load for fast responses

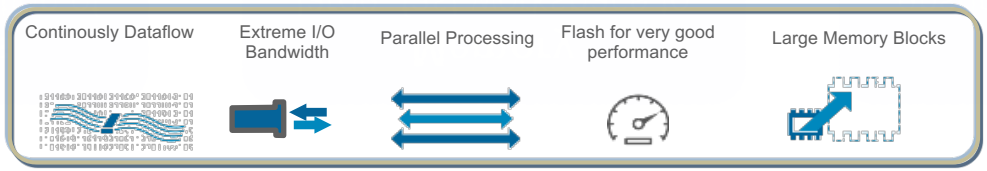

First especially for Big Data developed Processor

POWER8

OpenPOWER and IBM Power Systems

IBM Power Systems portfolio

S822
- Up to 20 Cores
- Up to 1TB Memory

S814
- Up to 8 Cores
- Up to 1TB Memory

S824
- Up to 24 Cores
- Up to 2TB Memory

redhat. SUSE ubuntu

S812L
- Up to 12 Cores
- Up to 512GB Memory

S822L
- Up to 24 Cores
- Up to 1TB Memory

S824L
- Up to 24 Cores
- Up to 2TB Memory

Enterprise Family

Scale - Out Family

E850C
- Up to 48 Cores
- Up to 4TB Memory

E870C
- Up to 80 Cores
- Up to 16TB Memory

E880C
- Up to 192 Cores
- Up to 32TB Memory

AIX i for Business Linux

Power Enterprise Pools & CoD Lower Costs

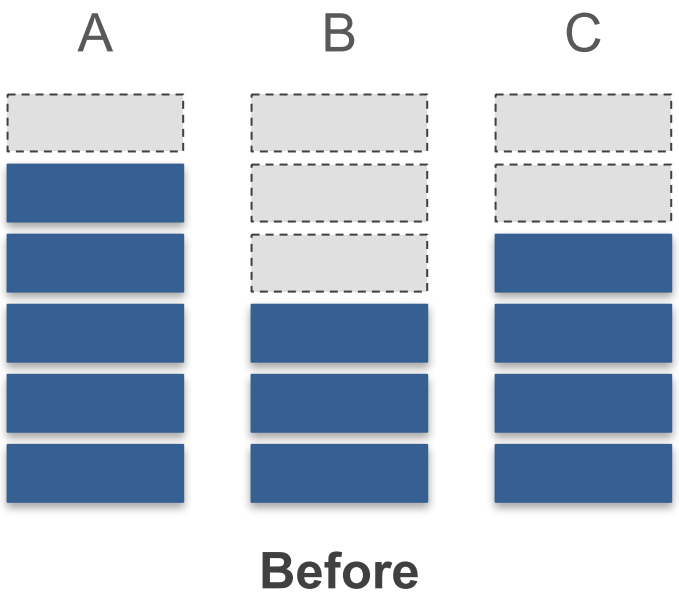

Before

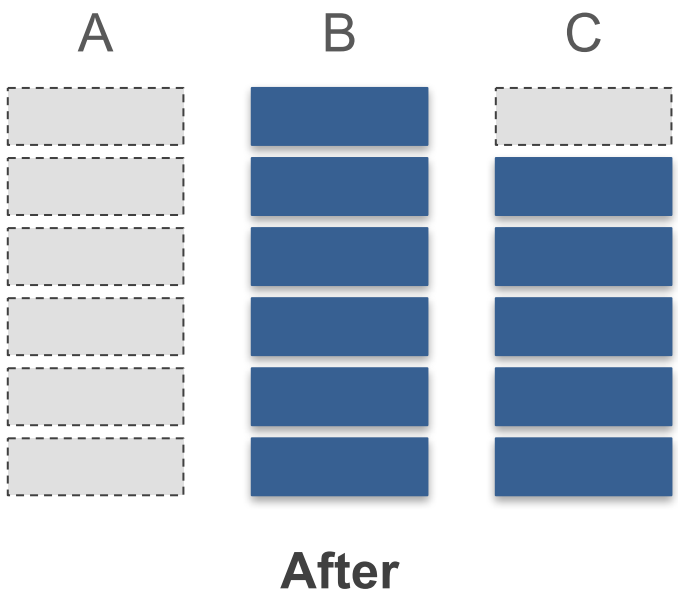

After

- Planned maintenance
- Rebalance capacity
- Failover clusters
- POWER8 migration

Power Reliability and Availability Strategy

Avoid errors by design

Make soft errors harmless

Let hardware take care of hardware

Focus on planned and unplanned outages

Virtualization Options for Power Systems

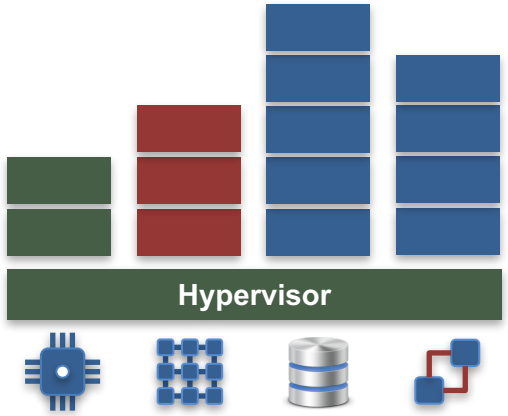

PowerVM is Power Virtualization that will continue to be enhanced to support AIX, IBM i Workloads as well as Linux Workloads.
PowerVM is supported on non-LC Systems (Exxx, Sxxx, SxxxL)

KVM provides an open source choice for Power Virtualization for Linux workloads. Best for clients that aren't familiar with Power and Linux centric admins.
PowerKVM is supported on L and LC Systems

IBM Power Systems run at Higher Utilization

For a single workload, we require a machine capacity **6.0x** the average demand

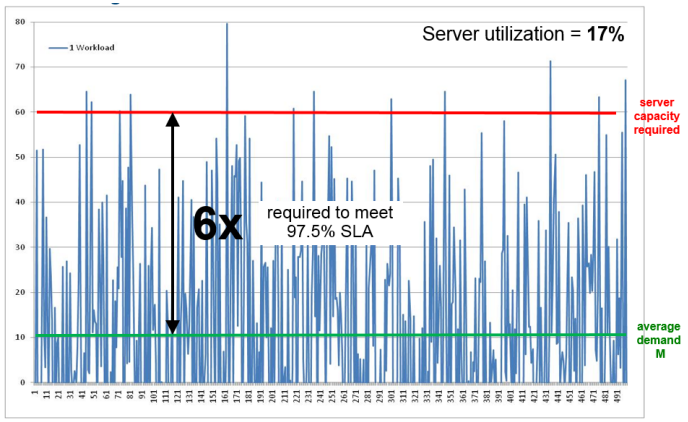

When we consolidate 4 workloads we only require **3.5x** average demand

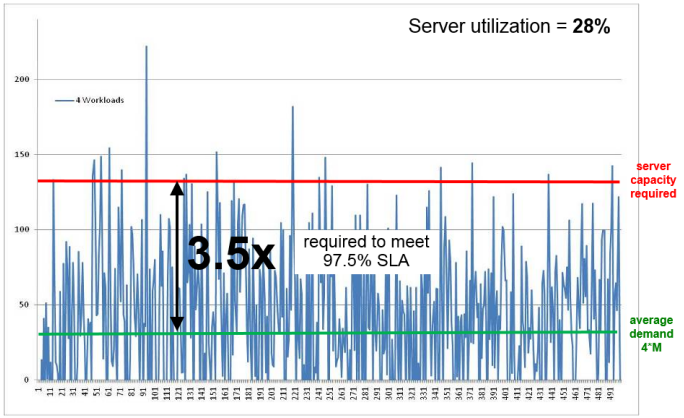

When we consolidate 144 workloads we only require **1.42x** average demand

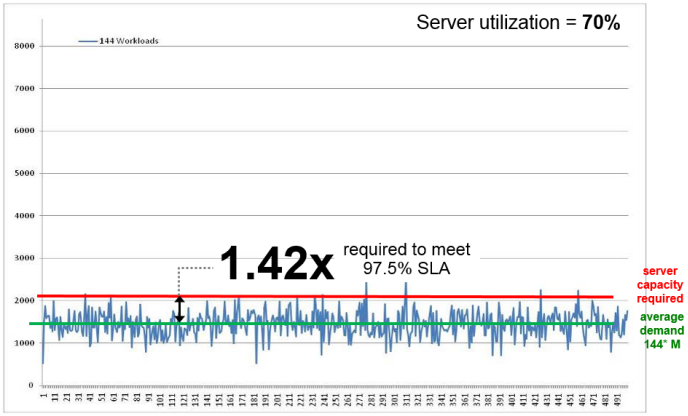

- Fewer servers
- Improved data center management
- Reduced software licensing
- Reduced HW & SW maintenance
- Opportunity for significant growth within a single physical System

☛ **Lower TCO**

Performance Utilization Guarantee

POWER8
Scale-Out

POWER8
E850

POWER8
E870 or E880

Real **TCA** and **TCO** Savings with POWER8

To qualify clients must implement each of the criteria set forth in the "POWER8 Performance Guarantee Requirements". See IBM for full details and remediation steps IBM will take in accordance with the Guarantee.

OpenPOWER Foundation - 300+ Members

This is What a Revolution Looks Like

OpenPOWER Linux Cluster (LC) Systems

S812LC

1 socket, 2U, Linux
8 or 10 cores
Up to 1 TB memory
Up to 112 TB Storage
12+2 Disk Bays
4 PCI Slots

KVM / Bare Metal

Hadoop & Spark

S822LC - GCA

2 socket, 2U, Linux
16 or 20 cores
Up to 1 TB memory
2 Disks
5 PCI slots

KVM / Bare Metal

Commercial

S822LC for Big Data

Up to 20 Cores
Up to 512 GB memory
12 HDD/SSD / NVMe
5 PCI Slots
(4 CAPI, 2 K80 GPU)

KVM / Bare Metal

Big Data

S822LC for HPC

Up to 20 Cores
POWER8 with NVLink
Up to 1 TB memory
2 HDD/SDD
3 PCI Slots. 2 CAPI
Up to 4 P100 GPU

Bare Metal

HPC

S821LC

Up to 20 Cores
Up to 512 GB memory
4 HDD/SDD / NVMe
4 PCI Slots
(3 CAPI, 1 K80 GPU)

KVM / Bare Metal

Compute

PostgreSQL and POWER

The winning team

Battle: High-Density consolidation test

Get the maximum number of VMs running *pgbench* while **average response-time is < 1ms**

IBM Power System

Intel x86 Xeon System

Number of VMs

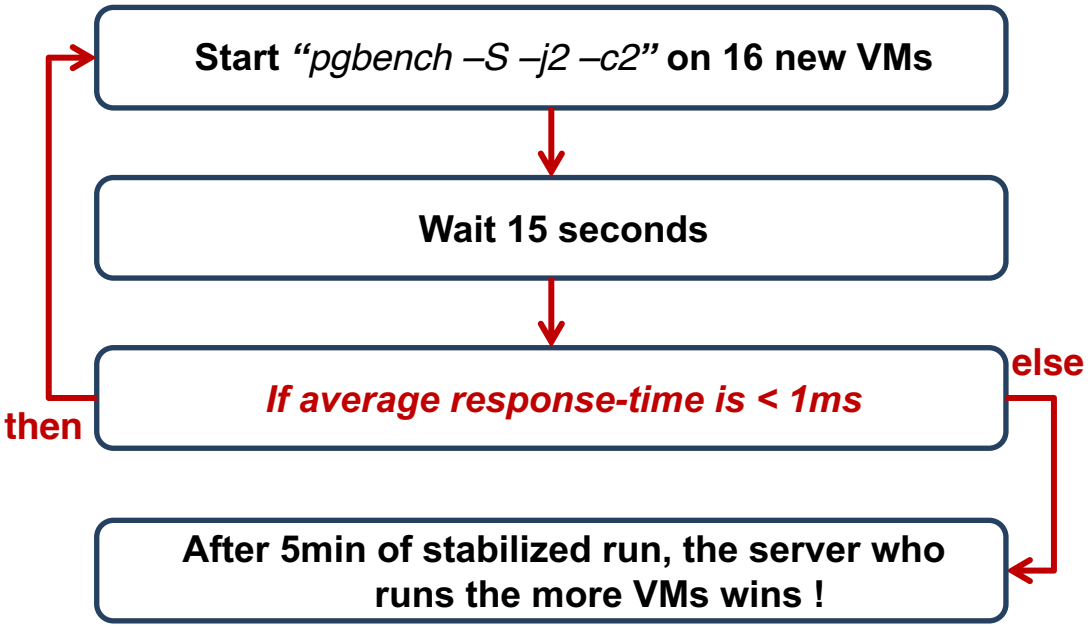

Higher Scalability on Power Systems

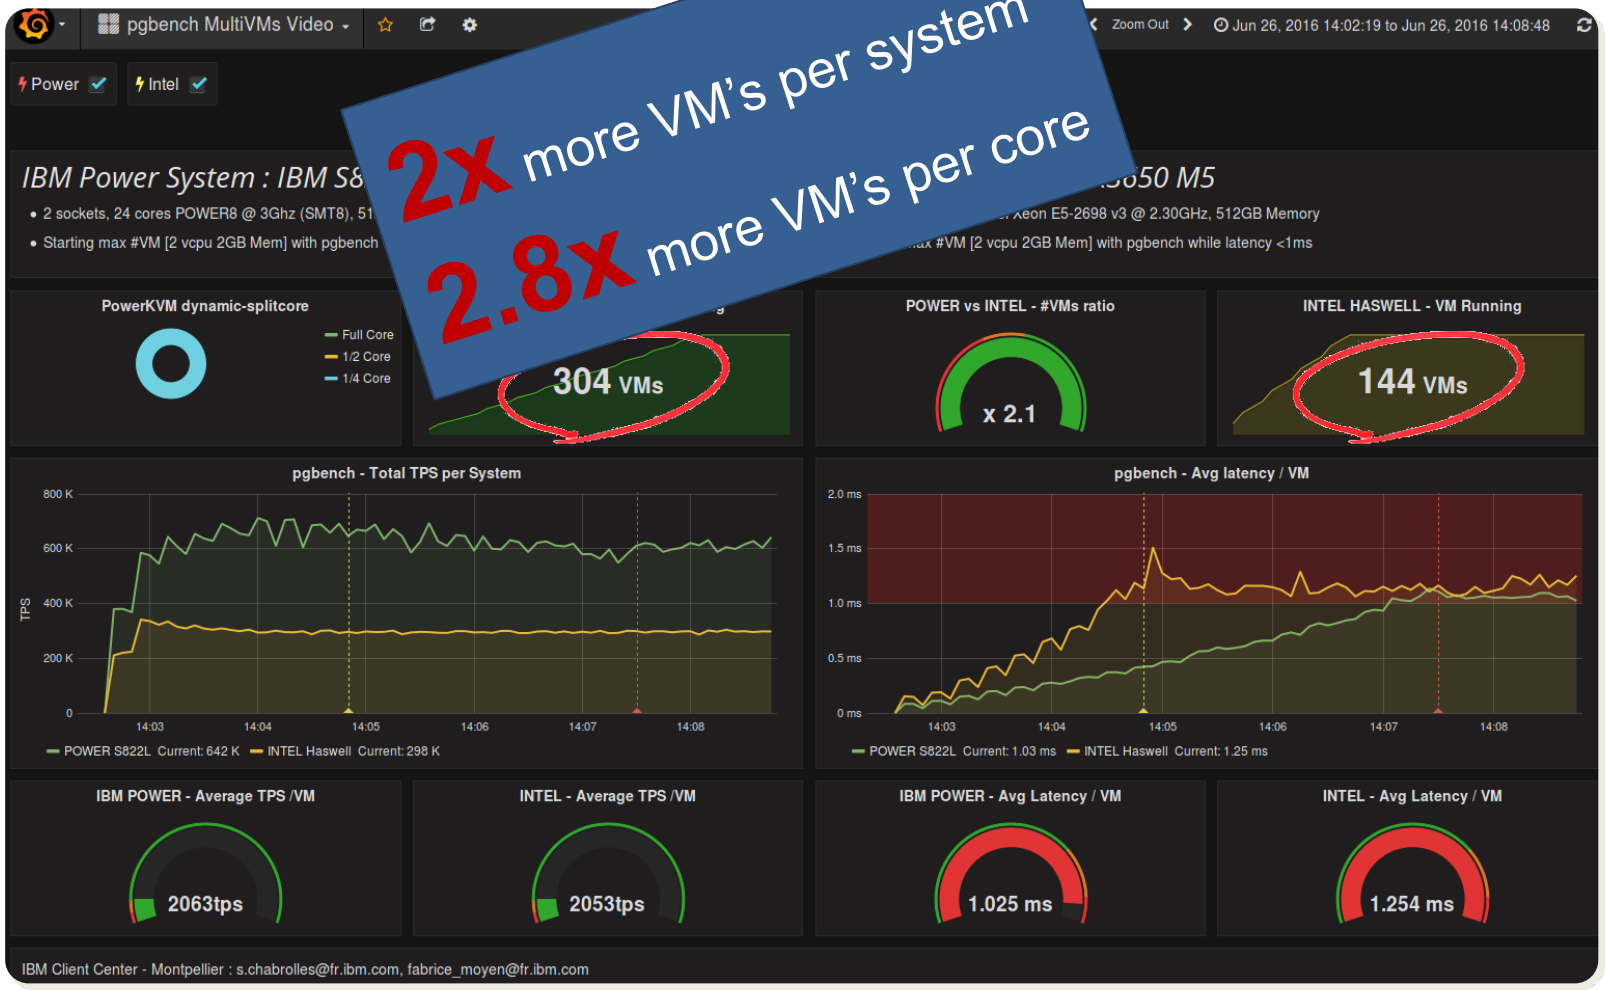

For detailed benchmark info's, see https://www.youtube.com/watch?v=5uKuQ8nzJns&t=120s

Higher Performance with Linux on POWER

IBM Power System S822L

- 20x POWER8 cores, 3.42GHz
- 256 GB memory
- PowerKVM, RHEL 7.1 LE, PostgreSQL 9.4.4

365’974 transactions / second

x86 Intel Xeon

- 4x Intel Xeon 2.80GHz (60 cores)
- 256 GB memory
- RHEL 7.1 / PostgreSQL 9.4.4

190’000 transactions / second

For detailed benchmark results (pgbench), see Benchmark Report “Splendid Data PostgresPURE on IBM Power System S822L”
http://www.splendiddata.com/benchmark-results-postgrespure-ibm-powerlinux-8

Lower costs with Linux on POWER

IBM Power System S822L

- 20x POWER8 cores, 3.42GHz
- 256 GB memory
- PowerKVM, RHEL 7.1 LE, PostgreSQL 9.4.4

0.11 EUR / transaction

x86 Intel Xeon

- 4x Intel Xeon 2.80GHz (60 cores)
- 256 GB memory
- RHEL 7.1 / PostgreSQL 9.4.4

0.32 EUR / transactions

For detailed benchmark results (pgbench), see Benchmark Report “Splendid Data PostgresPURE on IBM Power System S822L”
http://www.splendiddata.com/benchmark-results-postgrespure-ibm-powerlinux-8

Price-performance advantage guarantee

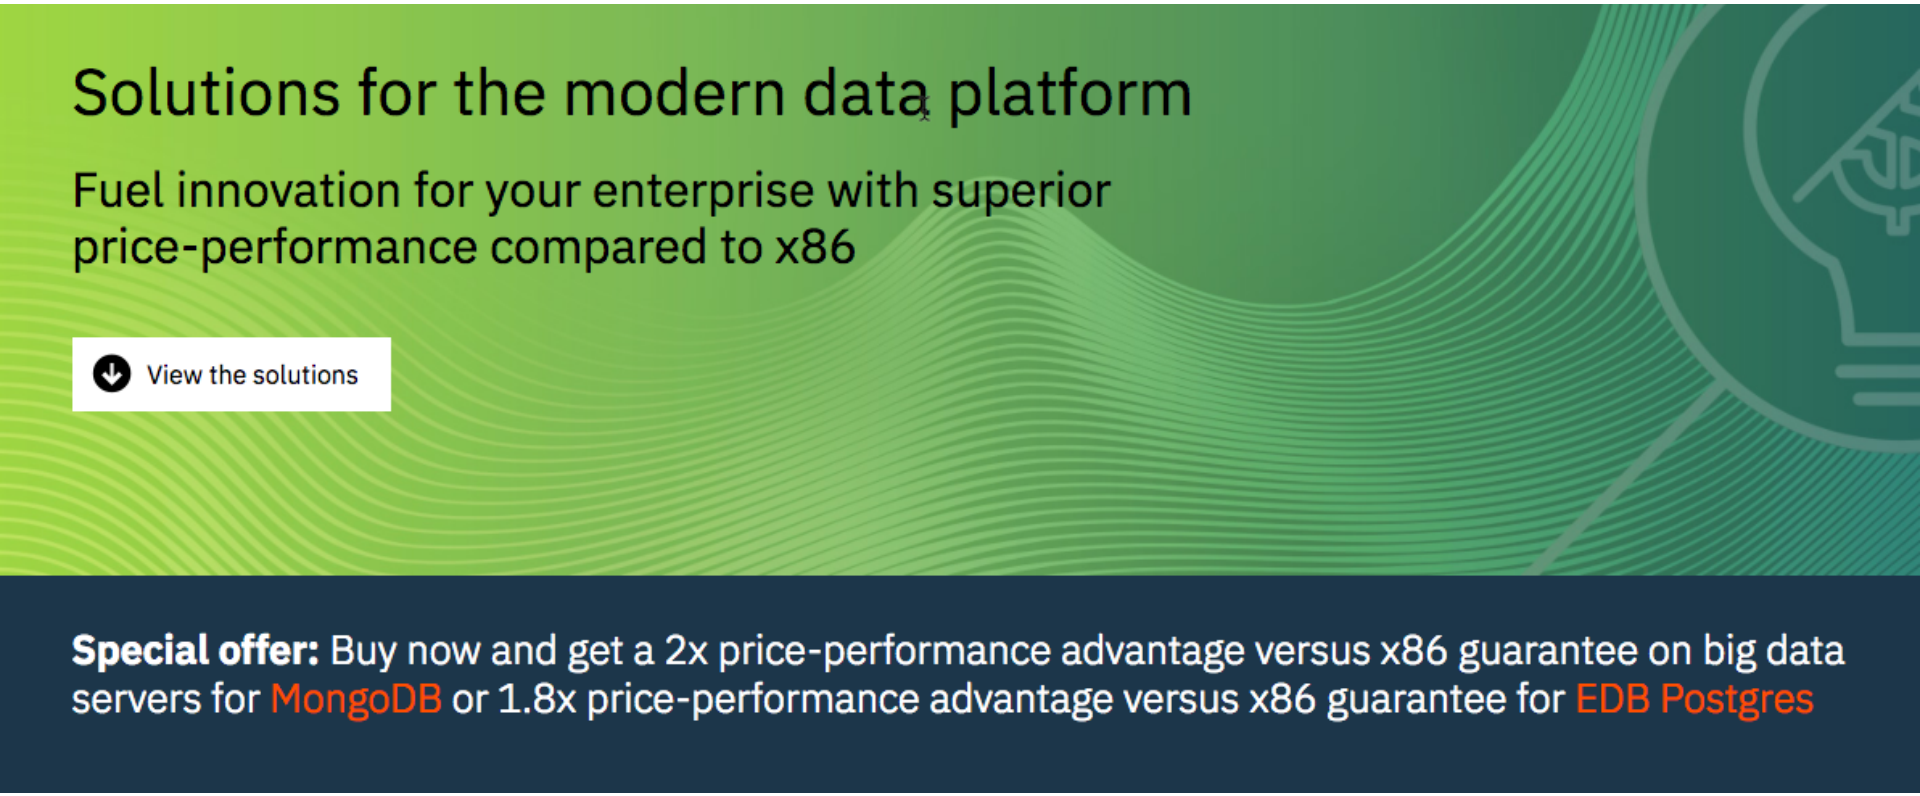

For detailed benchmark info's, see https://www-03.ibm.com/systems/power/solutions/data-platform/

IBM Linux Technology Centers

- 600 developers in over 100 open source projects to improve Linux
- 500+ IBM Linux SW products
- Work with industry organizations to expand the use of Linux

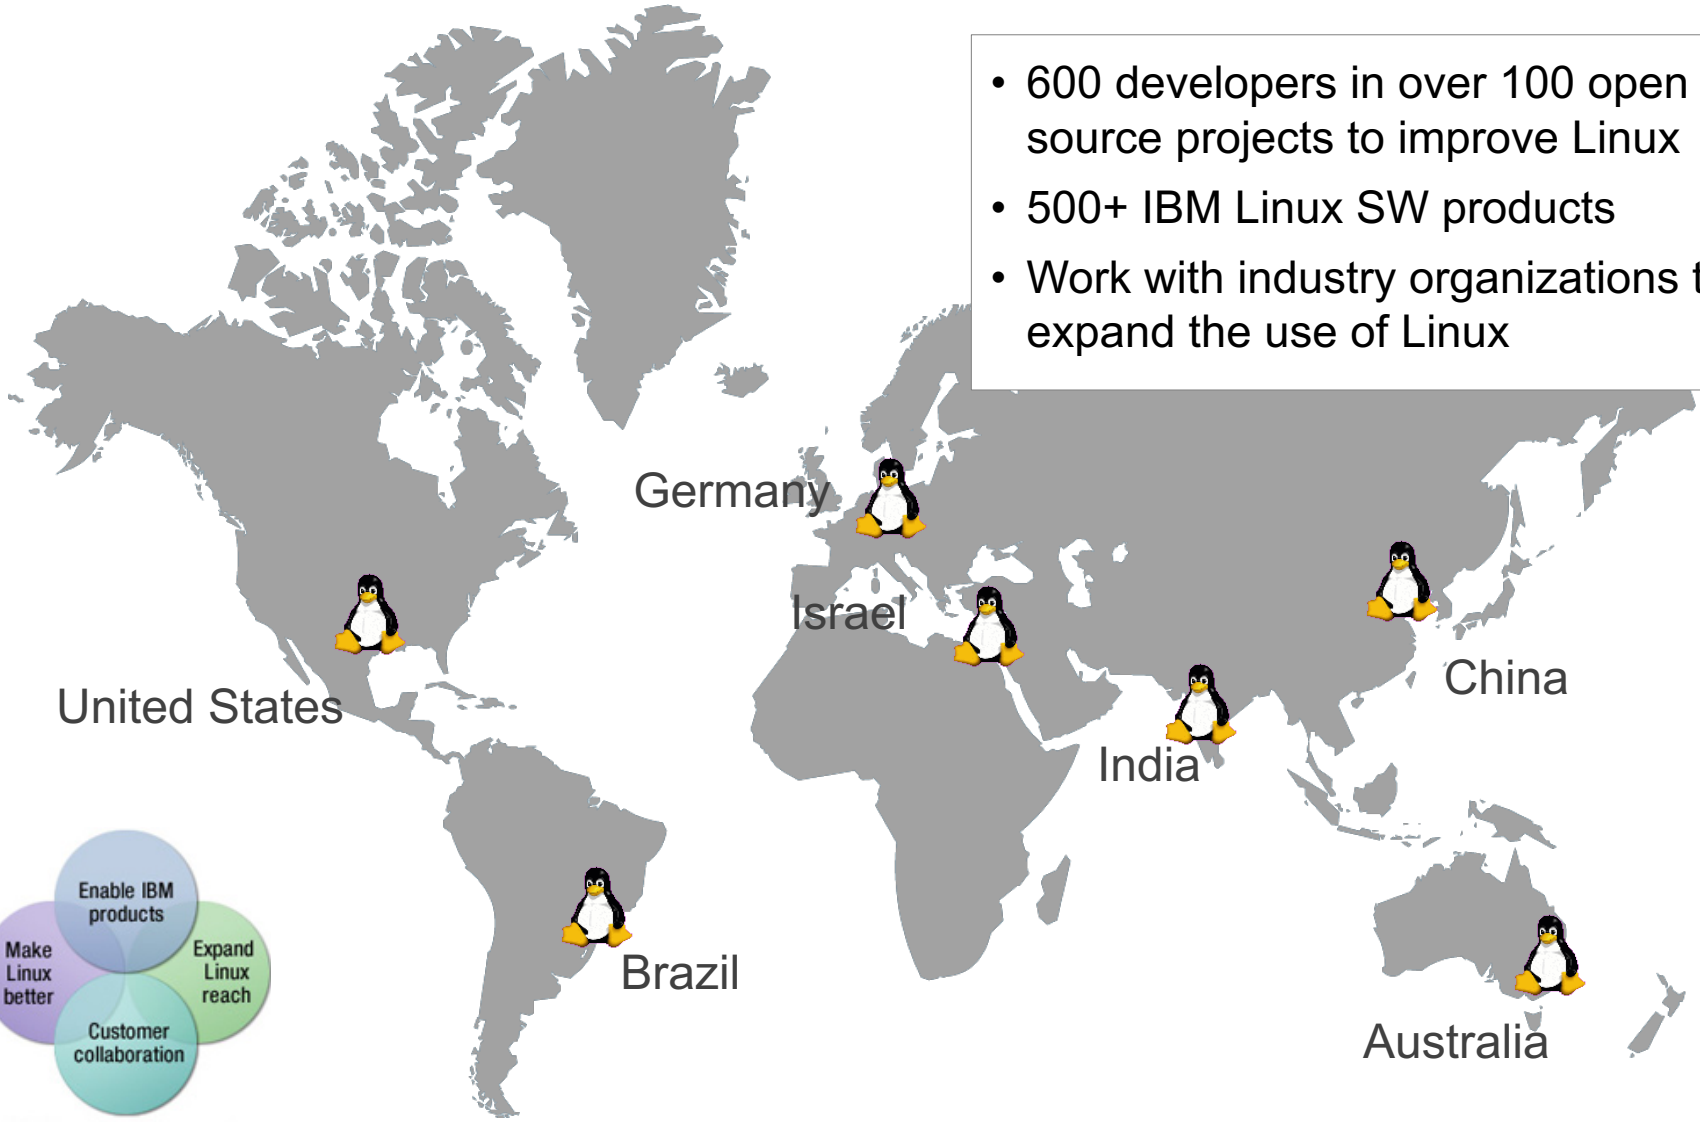

Summary

Power Systems, the best Systems for PostgreSQL workload

Highest throughput per core and core/memory bandwidth to deliver faster business results, up to 2x Intel-based alternatives.

Superior virtualization and management features to afford flexibility and maximum utilization.

Truly open ecosystem and collaboration for extreme accelerator innovation with industry-leading technology partners

Lower total cost of ownership compared to virtualization on x86, including reduced database licensing costs.

Power Systems

Open innovation to put data to work

Thank you!

IBM (Switzerland)
René Akeret - Systems Architect