Hewlett Packard Enterprise

Get the Best of Your Database

Red Hat and HPE combined

Yann Allandit – EMEA HP/Oracle Knowledge Center
Yann.Allandit@hpe.com

What is your challenge?

On your journey to SAP HANA?

Migrating your core business solutions from Unix?

Dreaming of mixing OLTP and OLAP in the same machine?

Need more uptime?

Worried about Oracle licensing costs?

Unable to meet your SLAs for SQL Server workloads?

Think you can´t virtualize mission-critical applications?

Modernizing your SAP system?

Todays Topics

- Red Hat / HPE Alliance Overview
- Economic Database Scaling
- Advanced Virtualization Flexibility (Demo)
- SAP HANA on Red Hat / HPE
- Selected Tuning Tools
- Take Away

Red Hat / HPE Alliance Overview

Hewlett Packard
Enterprise

HPE and Red Hat – Proven partnership for success

Market development

Superior results
–
Most servers and storage certified
Leading benchmark results

Certification, integration, and support

Superior experience
–
More than 4,000 Linux Service Professionals
x86 server Linux market share leader

Our alliance

Hewlett Packard Enterprise

redhat.

Our customers

16 years

Open source and open standards innovation

Strategic development

Superior alignment
–
Partnering to deliver transformational solutions for the New Style of Business

Superior commitment
–
More customers run RHEL on HPE servers than any other platform

Example Linux kernel development to support MC x86

Done in community partnership by Hewlett Packard Enterprise and Red Hat developers

- Identify areas to improve and problems to fix
- If no upstream patches already exist:
  - Develop and validate fix
  - Submit the patches upstream for acceptance
- Back-port patches and validate on RHEL kernels
- Submit back-ported patches along with performance justification to Red Hat for inclusion in RHEL
- Changes put into the next RHEL release for all RHEL customers to use.

HPE and Red Hat actively collaborate, with more than 125 patches submitted to advance Linux for mission-critical x86 implementations.

- **Improve scalability and performance** – Very large workload Intel Xeon® ProLiant support
- **Improve security and reliability** – Secure boot support
- **Improve reliability** – PCI LER/AER enabling hot plug for memory and CPU
- **Improve virtualization scalability and performance** – KVM hypervisor becomes NUMA-aware to improve performance to near-bare-metal speed

Economic Database Scaling

Do you want your tpcc-like workload to scale like this?

Comparison: accepted ratio is 1.6

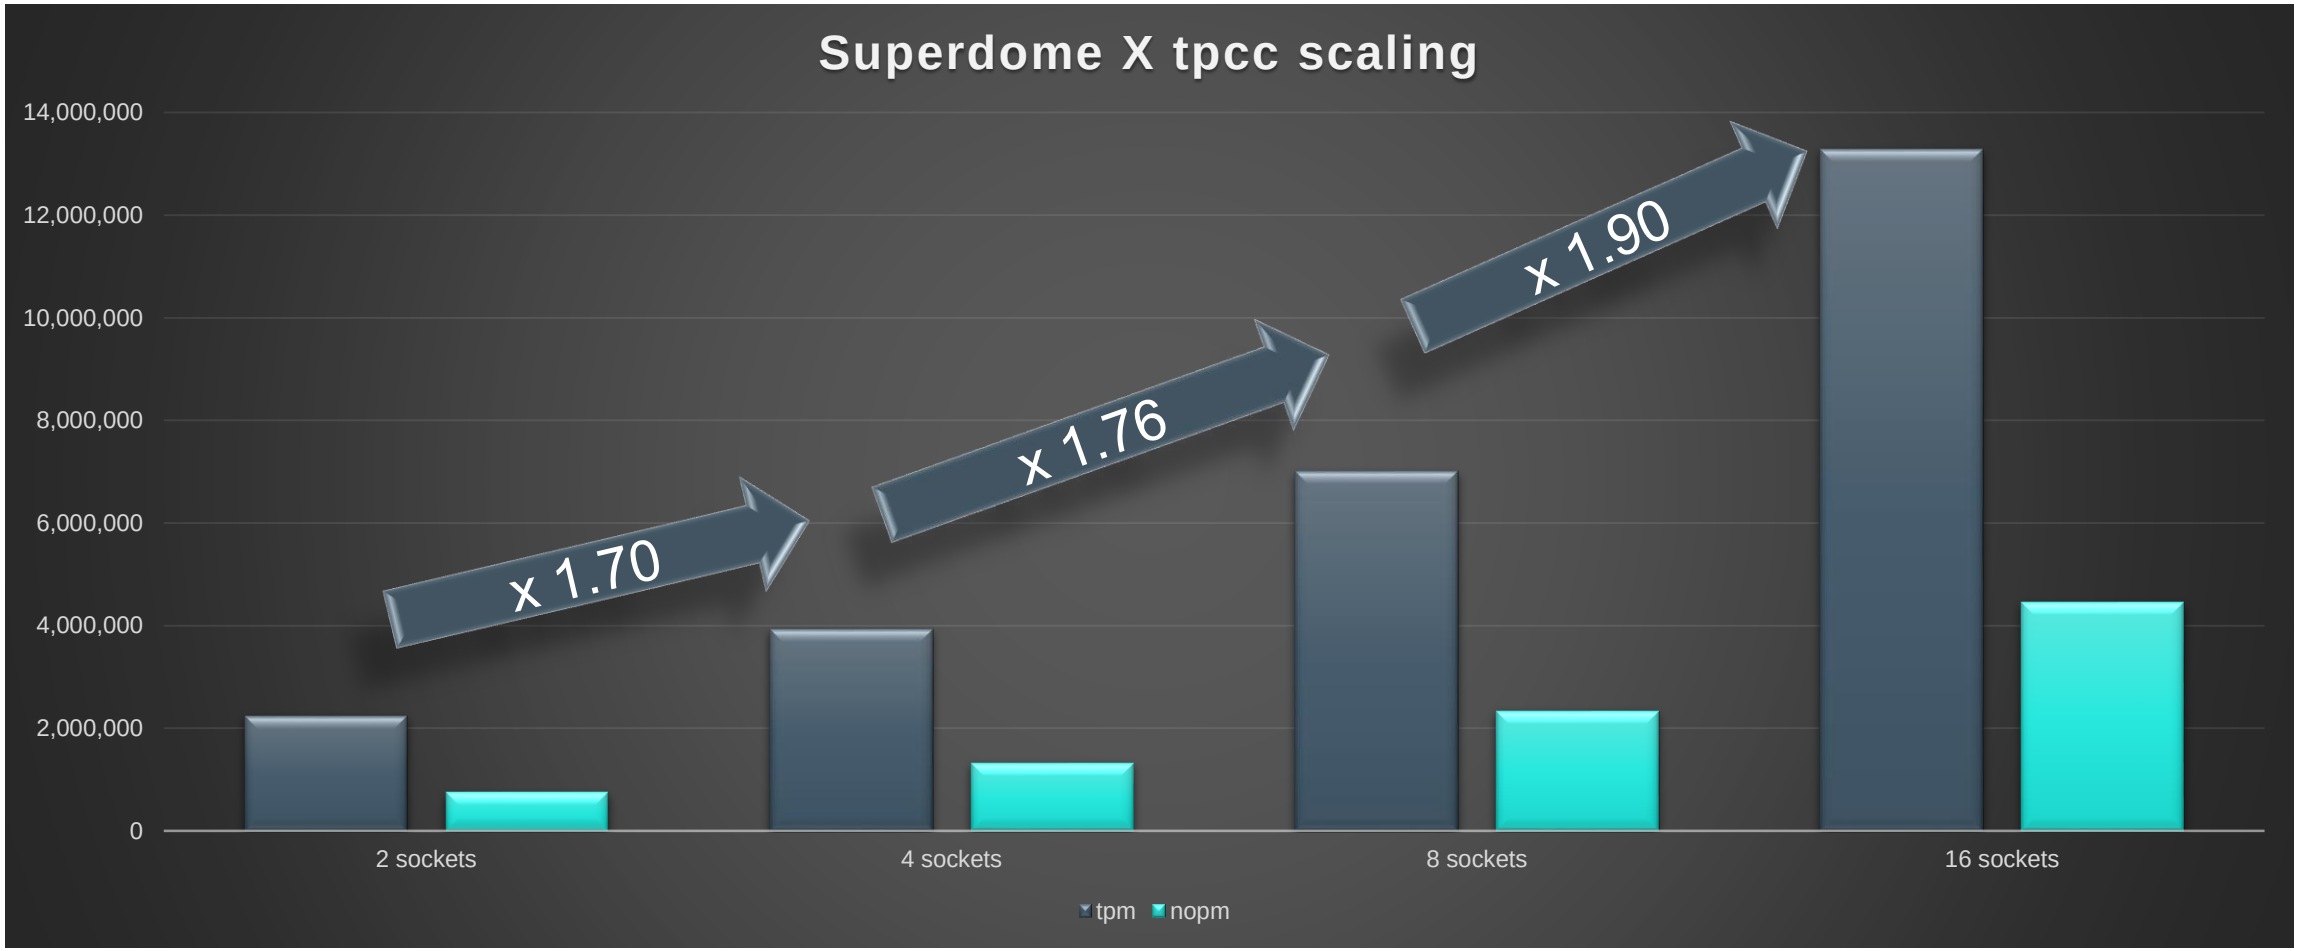

Hewlett Packard Enterprise

Scale on your terms: **Start small, grow seamlessly when needed**

1-8
blades

2-16
sockets

8-384
core count

24 TB
memory

Superdome X at a glance

Front view

Height: 18U

Rear view

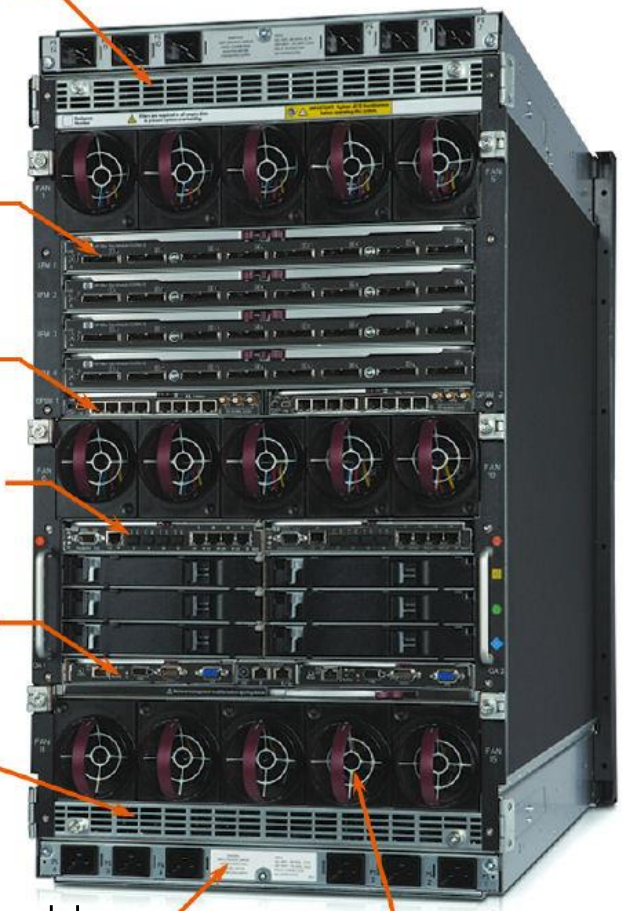

The unique value of HPE nPars

Hard partitions add flexibility and cost efficiencies

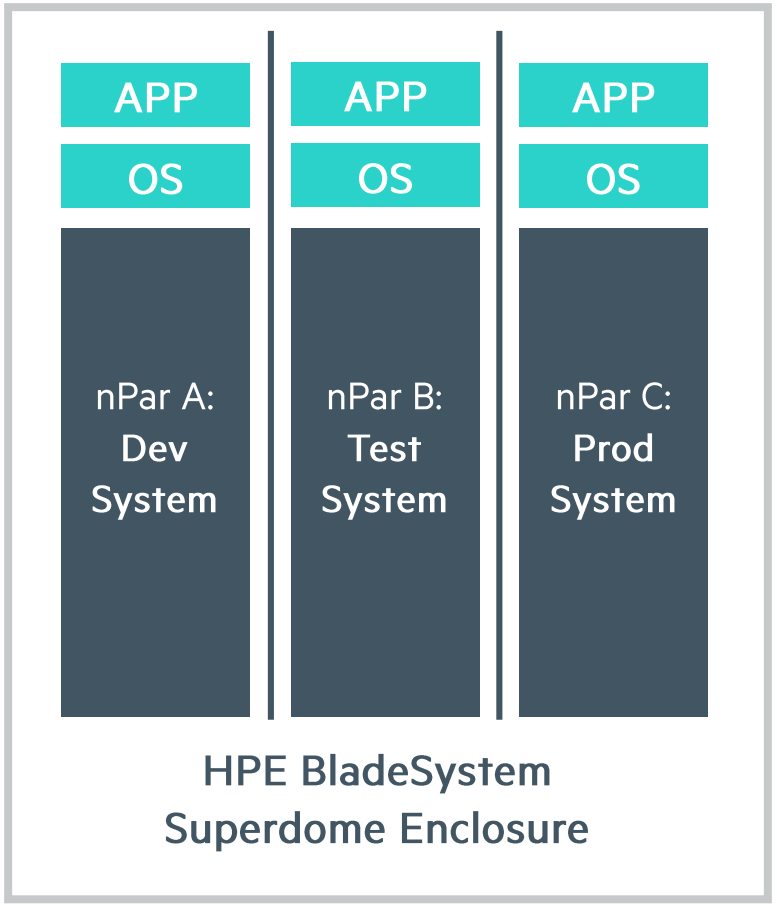

Lower your TCO	Optimize software costs by using HPE nPars
Maximize resource utilization	Create different development, test, and production environments within a single enclosure
Minimize downtime	Take one partition offline, perform maintenance and/or reconfiguration, while other partitions continue to run undisturbed 20x greater reliability
Protect your data	Electronic isolation provides a high degree of security between partitions

Hewlett Packard Enterprise

Configuration Example

HammerDB performance in tpm		
Configuration	4 BL920 120 cores	2 BL920 60 cores
TPM	7,000,507	3,931,072

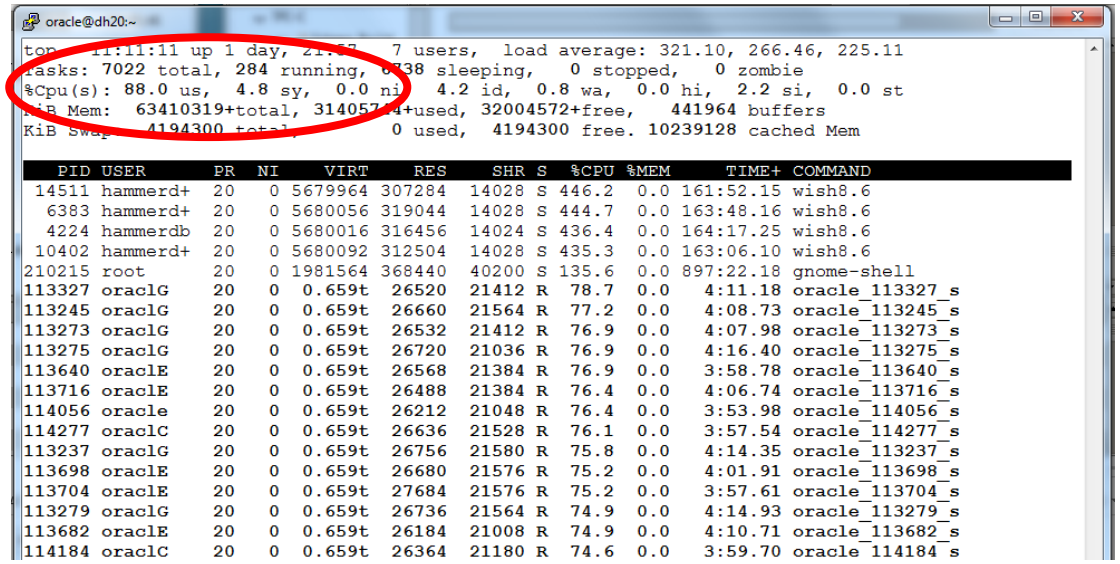

```
oracle@dh20:~
top - 11:11:11 up 1 day, 21:57,  7 users,  load average: 321.10, 266.46, 225.11
Tasks: 7022 total, 284 running, 6738 sleeping,   0 stopped,   0 zombie
%Cpu(s): 88.0 us,  4.8 sy,  0.0 ni,  4.2 id,  0.8 wa,  0.0 hi,  2.2 si,  0.0 st
KiB Mem:  63410319+total, 31405764+used, 32004572+free,   441964 buffers
KiB Swap:  4194300 total,        0 used,  4194300 free. 10239128 cached Mem

   PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
 14511 hammerd+  20   0 5679964 307284  14028 S 446.2  0.0 161:52.15 wish8.6
  6383 hammerd+  20   0 5680056 319044  14028 S 444.7  0.0 163:48.16 wish8.6
  4224 hammerdb  20   0 5680016 316456  14024 S 436.4  0.0 164:17.25 wish8.6
 10402 hammerd+  20   0 5680092 312504  14028 S 435.3  0.0 163:06.10 wish8.6
210215 root      20   0 1981564 368440  40200 S 135.6  0.0 897:22.18 gnome-shell
113327 oraclG    20   0  0.659t  26520  21412 R  78.7  0.0   4:11.18 oracle_113327_s
113245 oraclG    20   0  0.659t  26660  21564 R  77.2  0.0   4:08.73 oracle_113245_s
113273 oraclG    20   0  0.659t  26532  21412 R  76.9  0.0   4:07.98 oracle_113273_s
113275 oraclG    20   0  0.659t  26720  21036 R  76.9  0.0   4:16.40 oracle_113275_s
113640 oraclE    20   0  0.659t  26568  21384 R  76.9  0.0   3:58.78 oracle_113640_s
113716 oraclE    20   0  0.659t  26488  21384 R  76.4  0.0   4:06.74 oracle_113716_s
114056 oracle    20   0  0.659t  26212  21048 R  76.4  0.0   3:53.98 oracle_114056_s
114277 oraclC    20   0  0.659t  26636  21528 R  76.1  0.0   3:57.54 oracle_114277_s
113237 oraclG    20   0  0.659t  26756  21580 R  75.8  0.0   4:14.35 oracle_113237_s
113698 oraclE    20   0  0.659t  26680  21576 R  75.2  0.0   4:01.91 oracle_113698_s
113704 oraclE    20   0  0.659t  27684  21576 R  75.2  0.0   3:57.61 oracle_113704_s
113279 oraclG    20   0  0.659t  26736  21564 R  74.9  0.0   4:14.93 oracle_113279_s
113682 oraclE    20   0  0.659t  26184  21008 R  74.9  0.0   4:10.71 oracle_113682_s
114184 oraclC    20   0  0.659t  26364  21180 R  74.6  0.0   3:59.70 oracle_114184_s
```

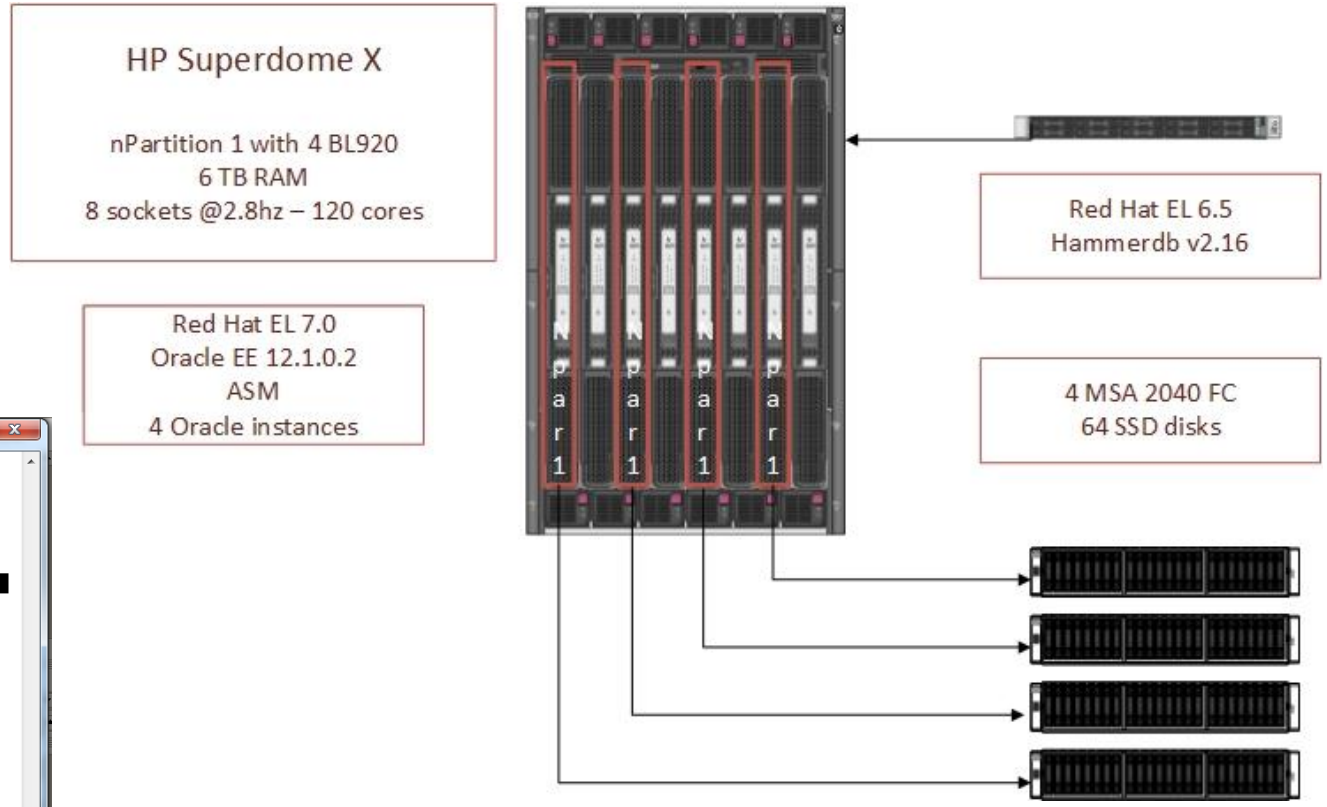

AutoNUMA with SPECjbb workload

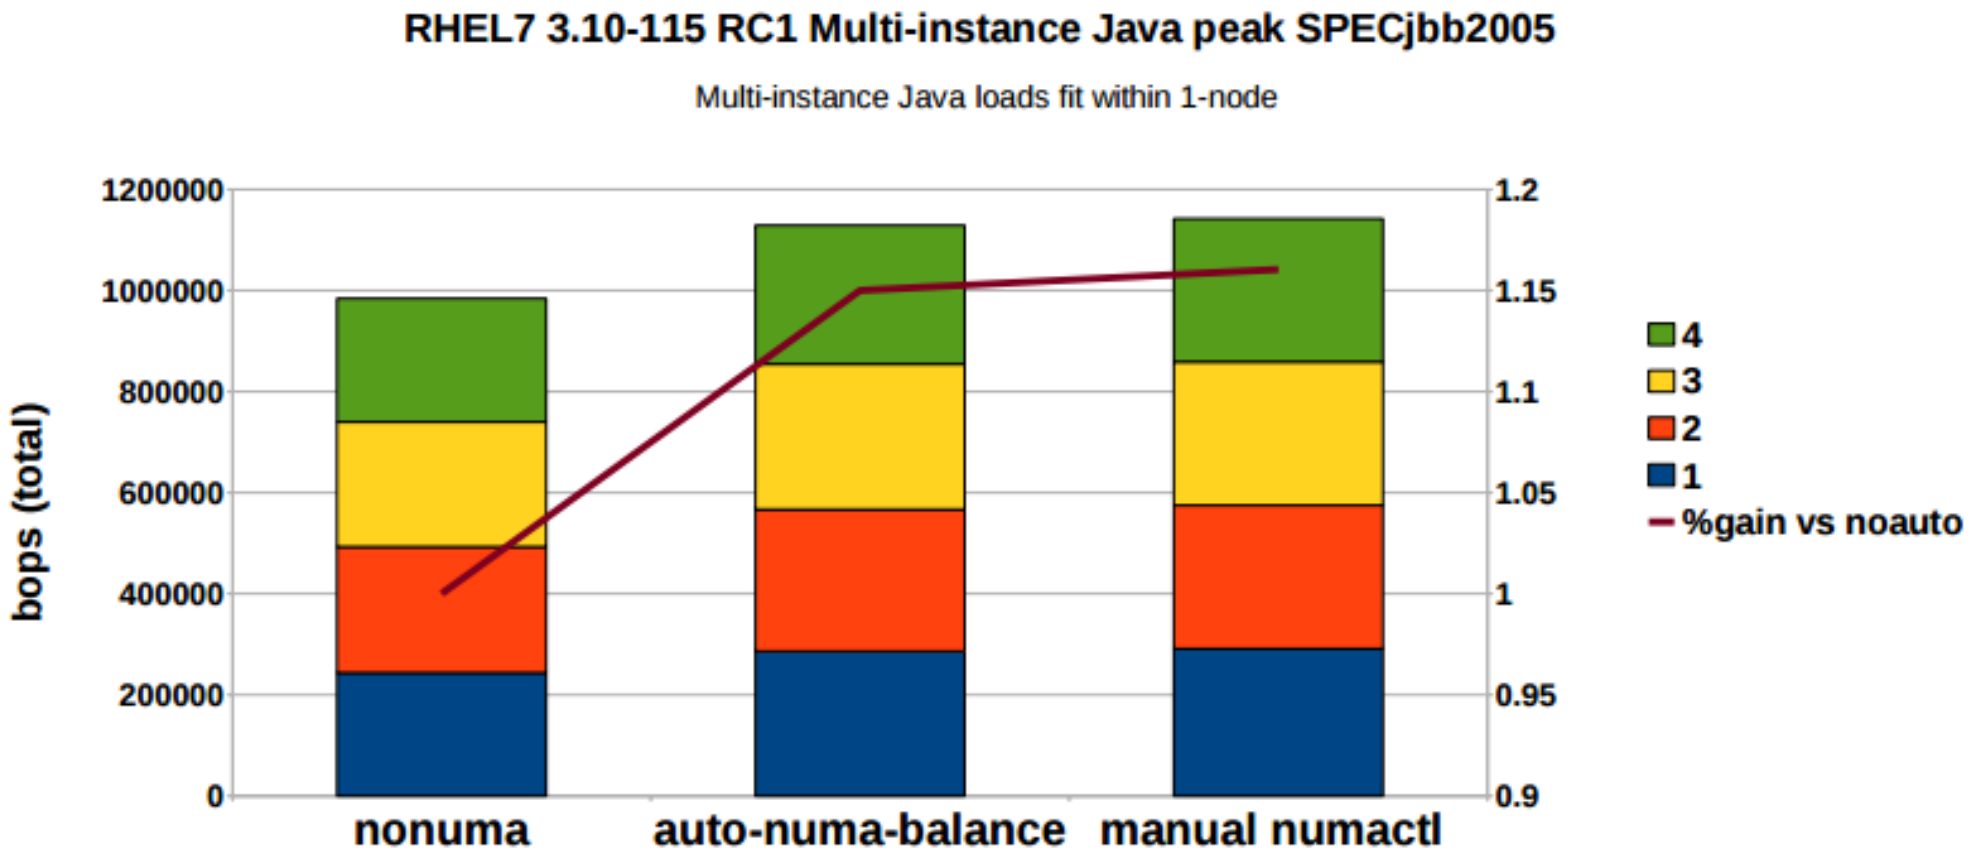

This graph comes from a Red Hat presentation delivered during the RH Summit 2014

NUMA tuning with Oracle database workload

No NUMA vs. Auto-NUMA vs. Auto-NUMA + HPE Optimization

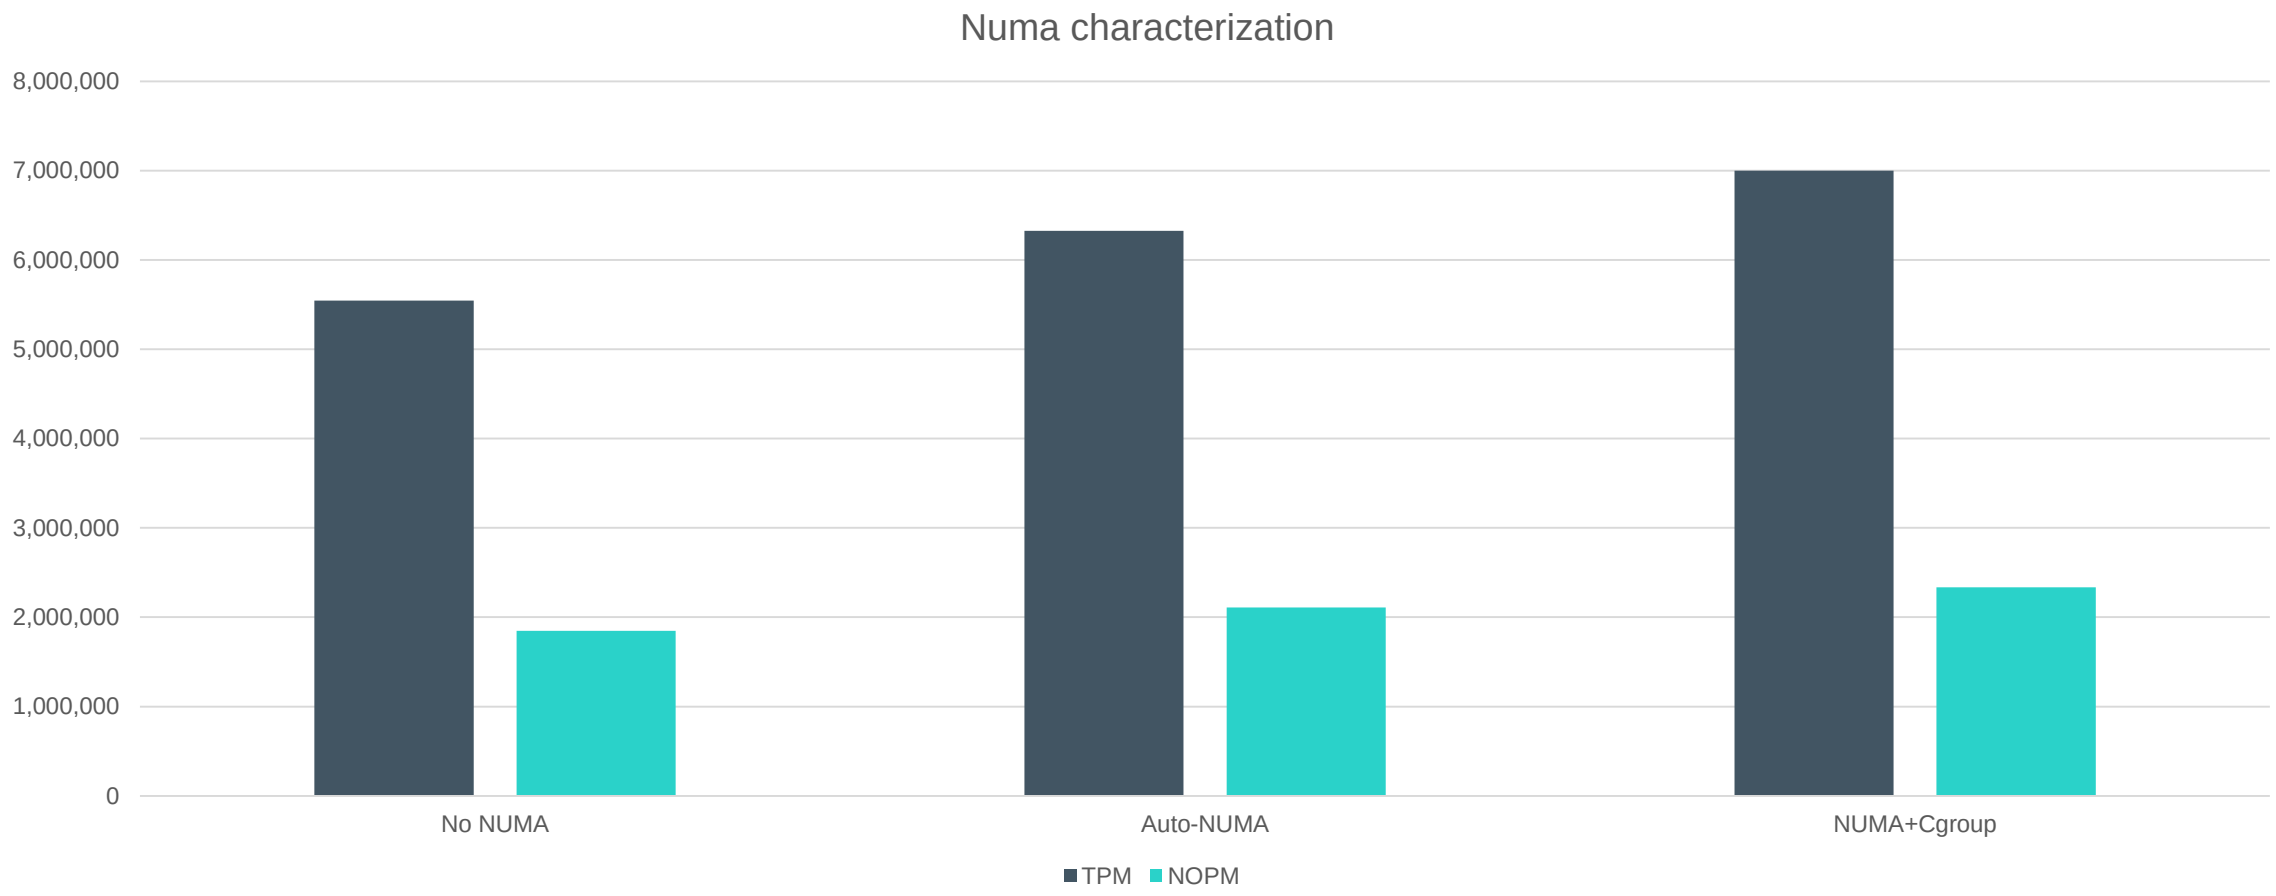

Hewlett Packard Enterprise

Tuned-adm

To be definitly considered

```
[root@dh20 ~]# tuned-adm list
Available profiles:
- balanced
- desktop
- latency-performance
- network-latency
- network-throughput
- powersave
- sap
- throughput-performance
- throughput-performance-no-thp
- virtual-guest
- virtual-host
Current active profile: throughput-performance-no-thp
```

These 3 profiles were tested
throughput-performance-no-thp
shows slightly higher performances

```
# tuned-adm profile throughput-performance-no-thp
```

Database Consolidation: Customer case

Two (2) databases per core – 80 DBs – 8,000 concurrent users in a partition

Initial result with configuration define by the customer team:
280,000 tpm

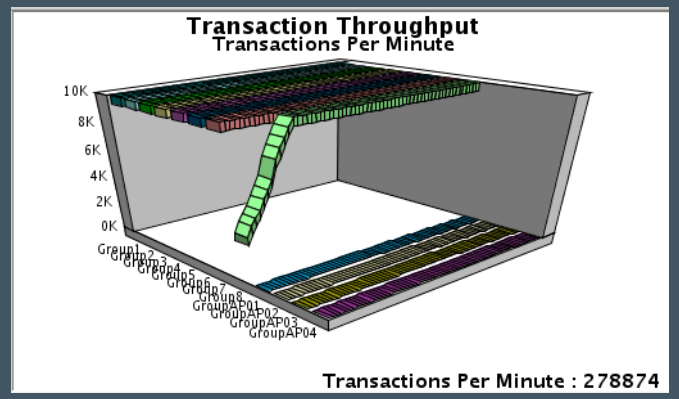

HPE Superdome X test with similar configuration:
900,000 tpm

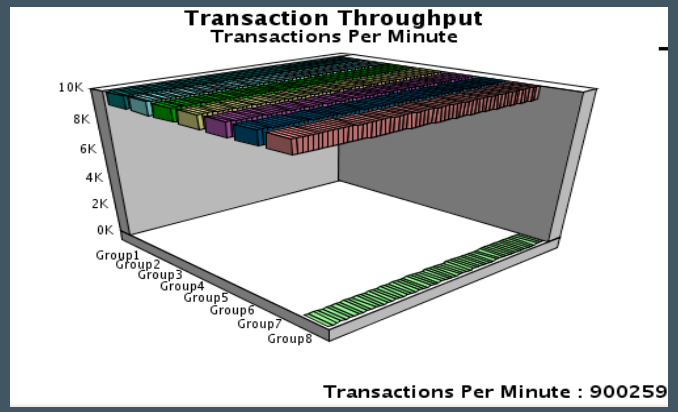

HPE top performance measured with up to 20 PDBs
1,600,000 tpm (at peak)

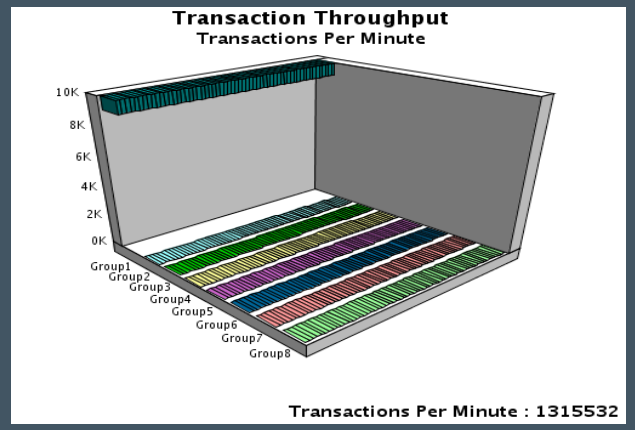

```
top - 12:01:38 up 14 days, 23:23,  8 users,  load average: 40.
Tasks: 9058 total, 142 running, 8916 sleeping,   0 stopped,
Cpu(s): 77.5%us,  6.9%sy,  0.0%ni, 12.3%id,  1.4%wa,  0.0%hi,
Mem:  3176572064k total, 2069124056k used, 1107448008k free,
Swap:  4182012k total,        0k used,  4182012k free, 1592998
```

Device:	tps	Blk_read/s	Blk_wrtn/s
dm-0	51,33	0,00	773,33
dm-1	0,00	0,00	0,00
dm-2	0,00	0,00	0,00
dm-3	96,67	0,00	773,33
dm-4	12,00	0,00	96,00
dm-5	0,00	0,00	0,00
dm-6	84,67	0,00	677,33
dm-7	5805,67	254154,67	64794,67
dm-8	5287,00	229248,00	64208,00
dm-9	4949,33	223866,67	63458,67
dm-10	5693,00	242608,00	63973,33
dm-11	0,00	0,00	0,00
dm-12	1797,33	10,67	158355,33

Hewlett Packard Enterprise

Superdome X – In Memory workload

Inital run:
101 transactions/hour with 4 PDBs

Total completed transactions	28	20	20	33
Average transactions per second	0,01	0,01	0,01	0,01
Maximum transaction rate	4	4	3	5
Total Failed Transactions	0	0	0	0
Average response time	1795824,43	1974997,00	1583030,15	1341288,30
Standard deviation	787799,09	1183070,45	1332714,83	832673,72
Minimum response time	58 925	2 010	48 706	34 740
Maximum response time	2 851 955	3 650 090	3 707 168	2 932 297
Geometric mean	1445946,09	1158679,38	858535,10	1002175,56

60000 Ios of 80 k block : 4,8 Gbs

In Memory enablement:
956 transactions/hour

Total completed transactions	234	240	233	249
Average transactions per second	0,06	0,06	0,06	0,06
Maximum transaction rate	11	15	9	11
Total Failed Transactions	0	0	0	0
Average response time	302812,40	303198,86	311310,99	292431,06
Standard deviation	272905,61	223966,35	347260,94	268021,86
Minimum response time	32 445	33 497	34 227	9 970
Maximum response time	1 342 491	1 405 363	2 347 809	1 814 952
Geometric mean	210396,92	239090,39	217370,82	215612,46

Oracle In Memory setup

```
*.inmemory_max_populate_servers=120
*.inmemory_size=1099511627776
*.open_cursors=4500
*.parallel_degree_policy='AUTO'
*.parallel_execution_message_size=32768
*.parallel_max_servers=480
*.parallel_min_servers=160
*.pga_aggregate_limit=483183820800
*.pga_aggregate_target=429496729600
*.resource_manager_plan=''
*.sga_max_size=1717986918400
*.sga_target=1610612736000
*.use_large_pages='ONLY'
```

Oracle In Memory usage

- The 4 SH schemas were loaded into the columnar format.
- The compression ratio was 4 on average
- The tables were loaded with « query low » option (means low compression)
- The table were pinned into the memory with a « critical » status. It means they were loaded at the boot time
- The time to upload the 4 entire schema (2TB on disk) into memory was 20 to 25 minutes.
- During the test, the CPU load was aroung 95%
- During the test, the IO usage was low.

Hewlett Packard Enterprise

Advanced Virtualization Flexibility

Solution Architecture

2xBL920s hard partition
60 cores – 120 threads
1TB of RAM

RHEV 3.6.1.3

Virtual Machine 1
Guest RHEL 7.2
80 vCPUs
512GB RAM
Oracle Database 12.1.0.2
In Memory option enabled
Average CPU usage: 85%

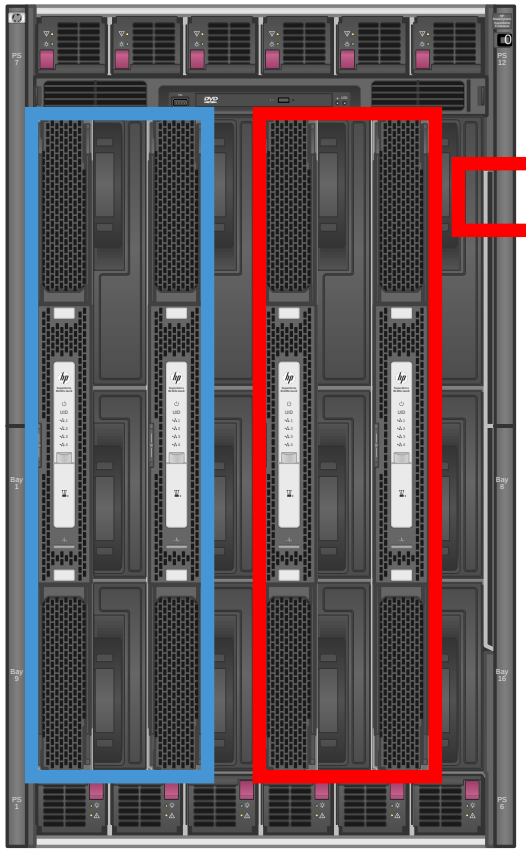

HPE Superdome X

2xBL920s hard partition
60 cores – 120 threads
1TB of RAM

RHEV 3.6.1.3

Virtual Machine 2
Guest RHEL 7.2
4 vCPUs
1GB RAM
Swingbench 2.5
70 users – DSS workload

Live-Migrating a Memory-intensive Workload

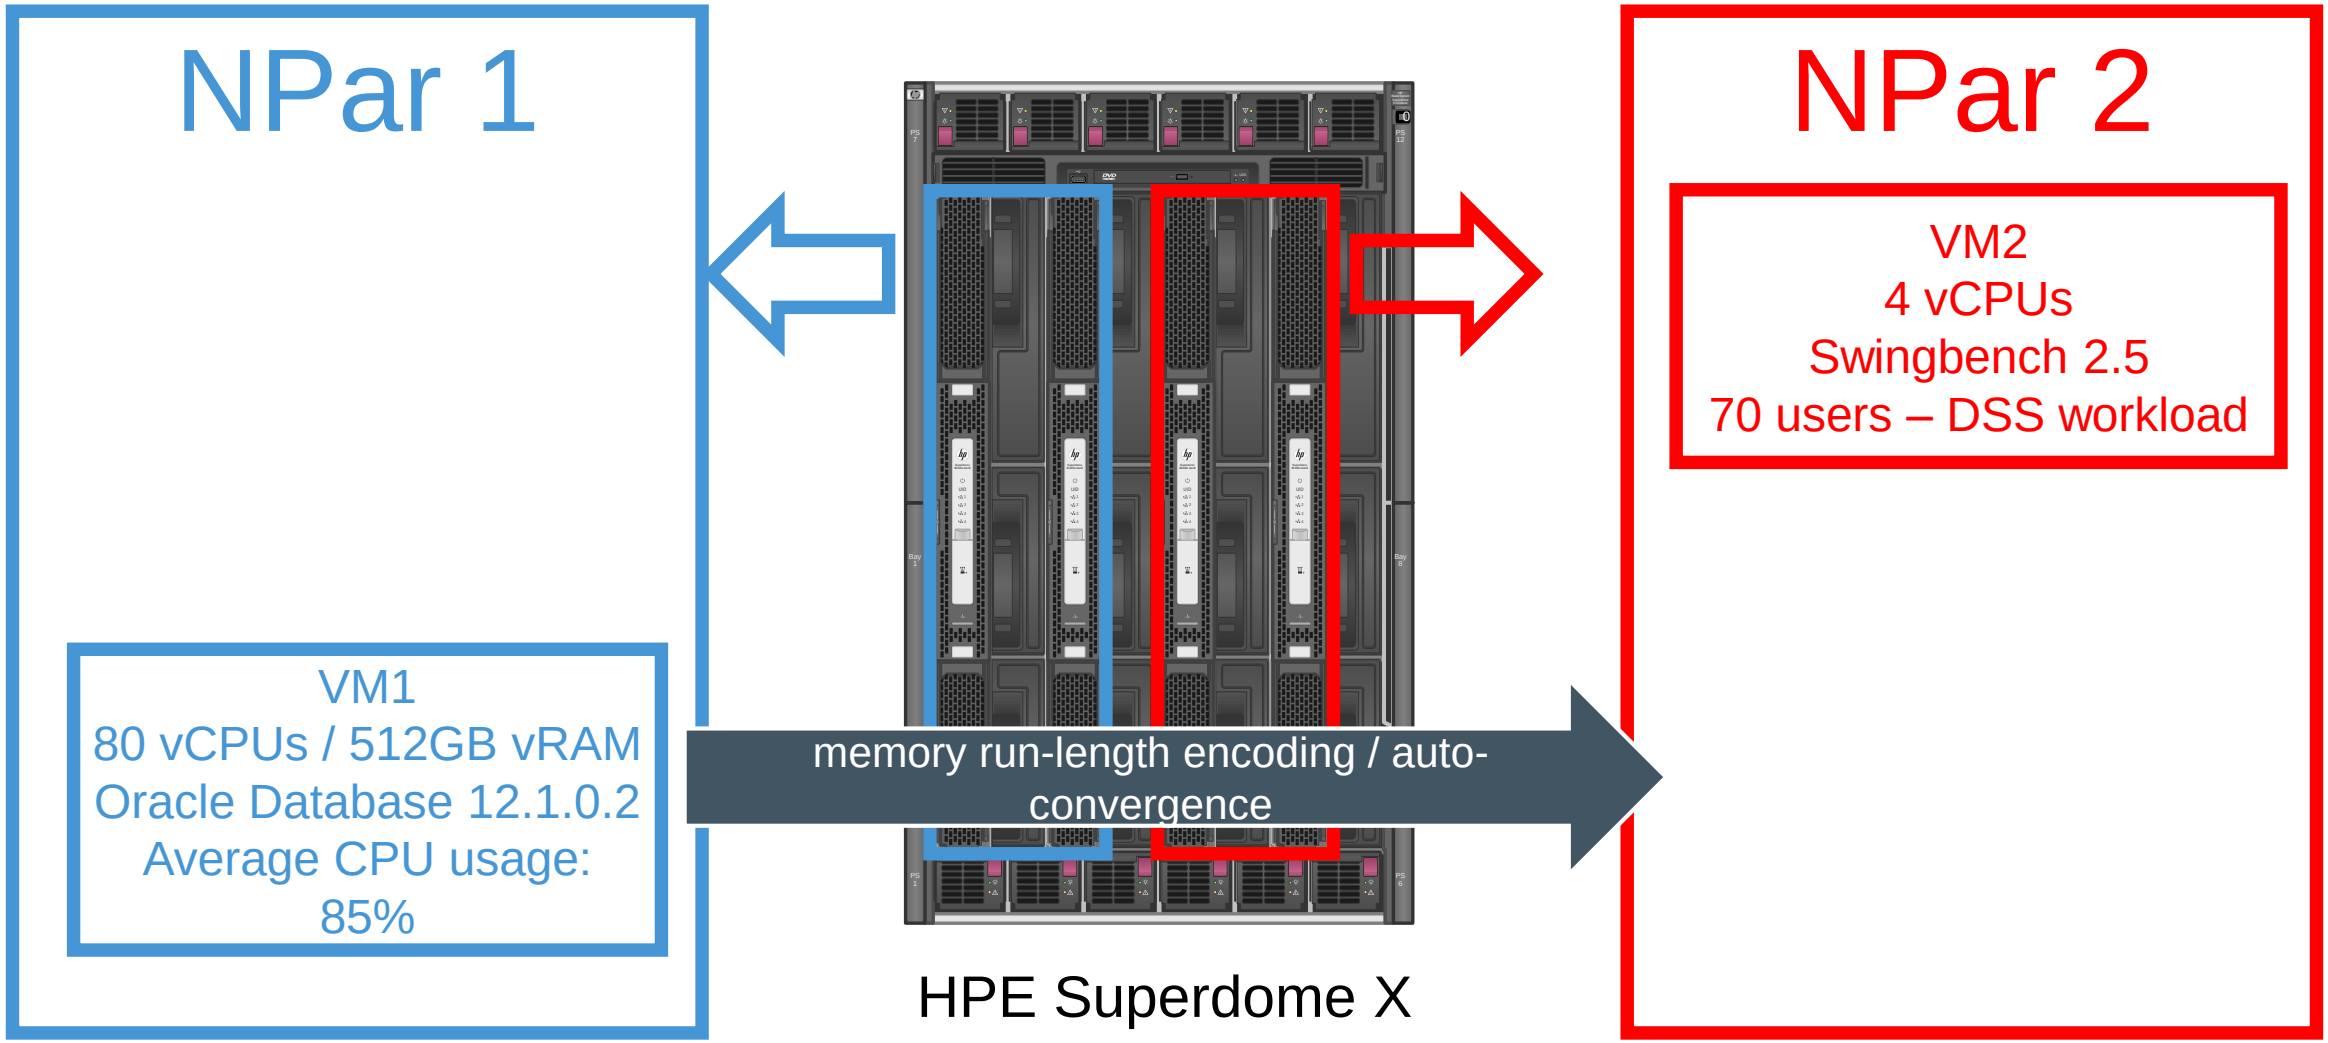

Hewlett Packard Enterprise

Demonstration

SAP HANA on Red Hat / HPE

Hewlett Packard
Enterprise

Red Hat: Key to SAP HANA Success

“The #1 requested thing that we got was not a feature for the product, it was support for Red Hat” (in regards to SAP HANA)

“Our largest financial customers came to us and said that we are not going to buy anything from you unless you support Red Hat.”

"The Red Hat platform is the fastest growing platform for our SAP core product SAP HANA on the market today."

Steve Lucas
President, Platform Solutions SAP
2015 Red Hat Summit Keynote
https://www.youtube.com/watch?v=_UiKdgasgQ0

Hewlett Packard Enterprise

Portfolio of real-time data management systems

Optimized appliances for SAP HANA

HPE ConvergedSystem 500

- **Innovative:** Built on the next generation Intel® Xeon® E7 Haswell Processor
- **Fastest Performance:** #1 SAP BW-EML benchmark results
- **Right sized for you:** Start as small as 128GB, and grow up to 24TB[1]
- **Choice:** Over 1,000+ scale-up and scale-out configuration possibilities

HPE ConvergedSystem 900

- **Bigger business applications:** Large, demanding, complex SAP landscapes; Application aggregation and consolidation
- **Right sized for you:** Highly flexible scale-up systems for SoH and S/4HANA from 1 TB to 12 TB
- **Unmatched Scalability:** Massive scale-out system for BWoH from 2 TB to 144 TB[2]
- **Reliability:** Mission-critical high availability and data protection

Hewlett Packard Enterprise

[1] Larger scale-out systems may be available at HPE and SAP's discretion

[2] Certified up to 32 TB; customers may scale out to 144 TB with a custom SOW

Selected Tuning Tools

Significant OLTP performance gains with HPE-ATX

– Database performance gains with HPE Integrity Superdome X Gen9

HPE Application Tuner Express (ATX): Launches applications on large core count systems in an efficient way, delivering significant performance improvements

Cut your per-core licensing costs almost in half!

1.77x scaling with HPE-ATX!

Up to 13% improvement on 4P with HPE-ATX!

Up to 24% improvement on 8P with HPE-ATX!

Up to 58% improvement on 16P with HPE-ATX!

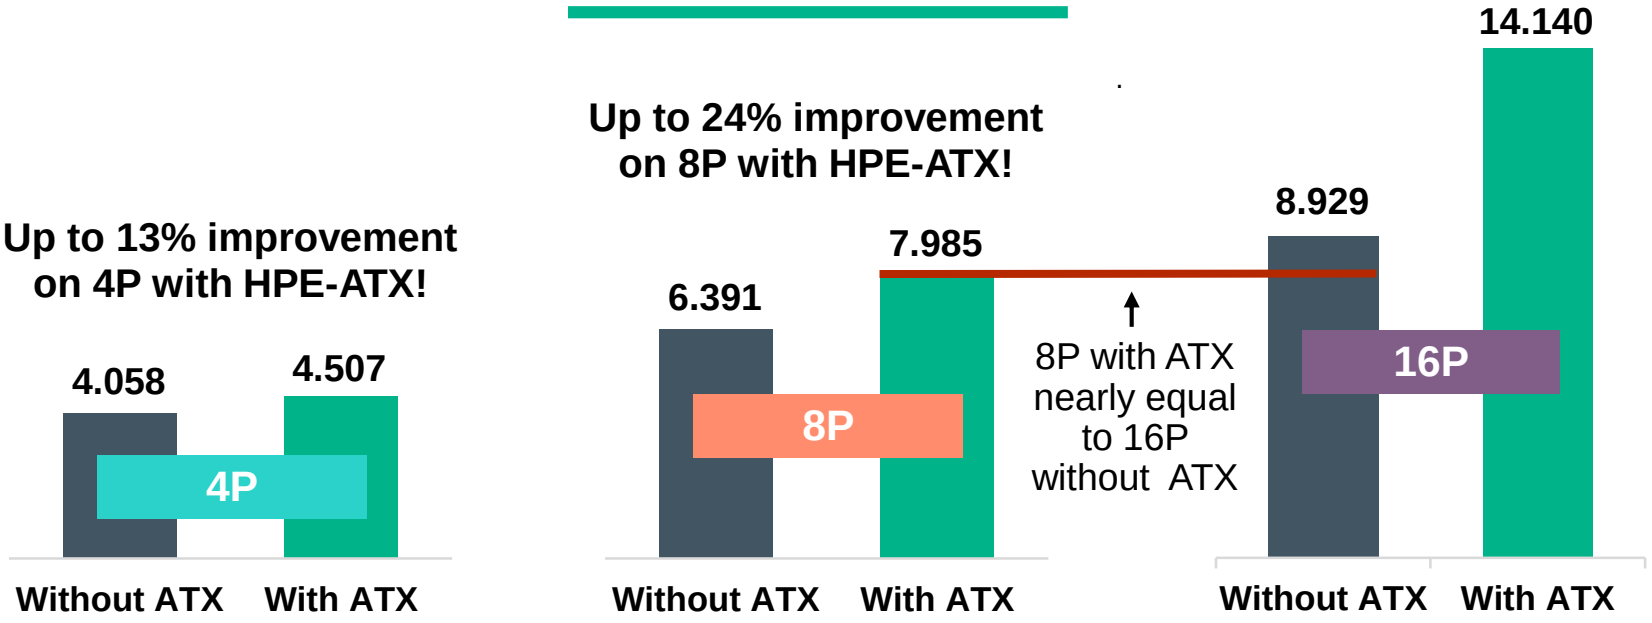

Hammer DB OLTP multi-tenant workload

HPE Superdome X using BL920s Gen9 blades with Intel® Xeon® E7-8890 v3 processors

Hewlett Packard Enterprise

Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. HPE internal testing. Performance results with HPE-ATX (patent pending) on the Superdome X with Gen9 were achieved with one tenant totaling 144 users (4P), two tenants totaling 288 users (8P) and four tenants totaling 576 users (16P) for a 1-to-1 mapping of users to logical CPUs. Configurations that are over-subscribed may show less performance gain. Results as of April 6, 2016.

System Resource Orchestrator (SRO)

SRO is a workload manager for Linux/ Windows for scale-up environments orchestrating memory, CPU, storage and networking, guaranteeing the right amount of resources at the right time for your business critical applications

Orchestration of
- CPU
- Memory
- Networking
- Storage

OS supported
- RHEL 6.5
- RHEL 6.6
- RHEL 7.x
- SUSE 11
- Windows

Hardware Supported
- Superdome X
- Proliant (SL, DL, BL, ML)

Integrates with
- One View
- Service Guard
- Workload Insight
- Docker/ Docker Orchestration Tools

Policies
- Fixed
- Utilization
- Own-Borrow
- Conditional
- Custom

BENEFITS

- Optimal utilization of System Resources
- Optimized workload consolidation
- Application agnostic
- Easy user interface for workload management
- Lower TCO with OS License and Support costs reduction
- Reduced administration costs
- Better performance with same resources

Hewlett Packard Enterprise

Text: *Would be available in future versions*

Take Away

Proof Point #1: **Statens Innkrevingssentral modernizes collections infrastructure**

Statens innkrevingssentral

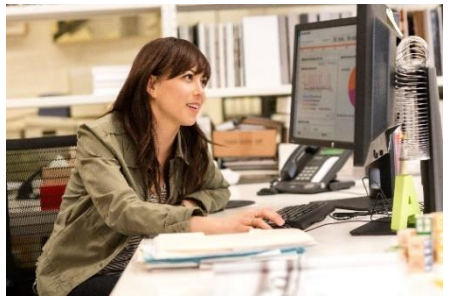

- Profile
  - Agency collecting fines, compensation, court case costs and confiscations, among others.
  - 188 types of claims for 35 clients from 15 departments
- Approach
  - Migrate to x86 reliability and availability on the Superdome X Server, 3PAR and RHEL, including RHEV to drive down operational costs
  - Focus on increased performance and improved availability
- Results
  - Streamlined online debt and fee collection for 35 government agencies
  - Enabled 90% of collections to be processed online
  - Delivers highly reliable infrastructure to support collection activities totaling up to ~2.5 M USD per day

10x performance boost

1/3 TCO vs. previous infrastructure

>30% deduplication of stored data

Proof Point #2: **Cerner chooses Superdome X to deliver hosted services**

- Powering core database for healthcare applications
- Replacing both x86 and older UNIX environments with Superdome X, as they standardize on x86 and RHEL
- Growing client base requires higher levels of infrastructure scalability, over and above what current x86 solution could offer
- With Superdome X, they now have an x86 server powerful enough for their largest environments

Cerner

"As a global supplier of health information technologies, Cerner continues to innovate to meet our clients' needs. We look to HPE as a partner to provide compute power in support of our innovation. The HPE Integrity Superdome X provides the scalability to support our largest clients, the performance to deliver the results our clients have come to expect and the availability that health care demands."
- Kent Scheuler, SVP, Managed Service, Cerner

More information is available

- Operating Environments
  - HPE & Red Hat Alliance
  - HPE Operating System Support Matrices
  - Red Hat Certified Hardware
- Superdome X external portals
  - www.hpe.com/servers/superdomex
  - www.hpe.com/info/superdomex
- Solutions
  - Application Tuner eXpress
  - System Resource Orchestrator
  - Rapid Setting For Oracle

Thank you

Hewlett Packard
Enterprise