

**performing  
databases**

# Project Report

# Oracle OLTP Benchmarks

against

# Dell EMC RecoverPoint

Martin Klier



Performing Databases GmbH  
Mitterteich / Germany

# Speaker

- Martin Klier
- Solution Architect and Database Expert
- My focus:
  - Performance + Tuning
  - Highly available systems
  - Cluster and Replication
- Linux since 1997
- Oracle Database since 2003

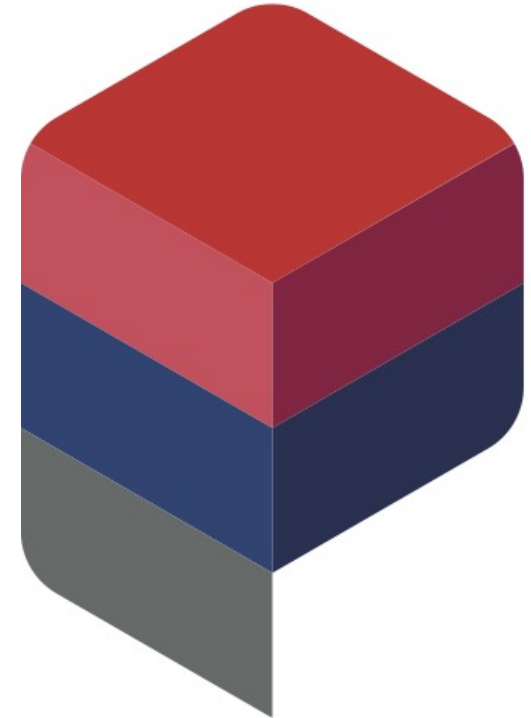


**ORACLE®**  
ACE Director



# Performing Databases

- Three Experts for Database technology
  - Concepts and Project Competence
  - Architecture- and System planning
  - Licensing
  - Implementation and Troubleshooting
- Get in touch
  - Performing Databases GmbH  
Wiesauer Strasse 27  
95666 Mitterteich // Germany
  - <http://www.performing-databases.com>
  - Twitter: @PerformingDB



# TGW Logistics Group

- Systems Integrator for automated warehouse solutions
- End-to-end fulfillment
  - Design
  - Manufacturing
  - Implementation
  - Maintenance
- Locations:  
Austria, Germany, UK, USA  
and many more
- 3.800+ Employees
- 800+ Mio € revenue
- Foundation-owned, former Family-operated



[www.tgw-group.com](http://www.tgw-group.com)

# Project









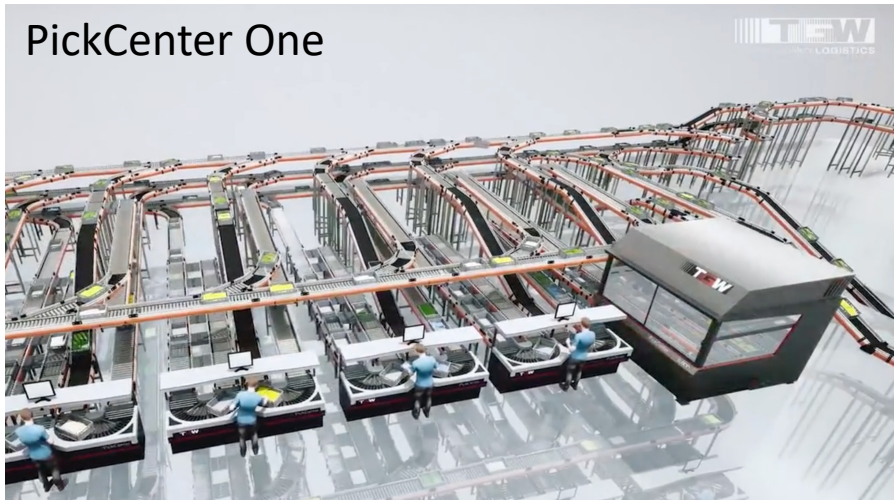
Stingray Shuttle



Stingray Shuttle Block



PickCenter One



Picking  
Robot



# Project Facts

- Greenfield Logistics Hub on 50.000m2
- Highly automated FlashPick® System
- >180.000 Storage Locations
- 200 energy-efficient Stingray shuttles in 10 aisles
- Automatic carton erection and closing
- Automatic crate stacking/lifting



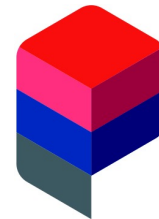
- Two Oracle Databases 19c
  - Material Flow (“Routing”)
  - Warehouse Management (“Business”)
- Zero Data-loss concept w/ Data Guard
- High-class OLTP workload profile
- >2.000 transactions/second
- Oracle Linux on VMware vSphere ESXi



**performing  
databases**

# Project Flow (2-3 years)

- Logistics Requirements
- Logistics Solutions Engineering
- Warehouse Layout Design
- Simulated Layout Validation
- Building / Construction
- Logistics + Software Realization
- IT Infrastructure Sizing based on Layout + Software
- **IT Infra Solution Engineering**
- **Benchmarked IT Infra Validation**
- IT Infra Realization
- Software Stress Testing / Emulation
- Commissioning Phase
- Production Ramp-up Phase
- Acceptance



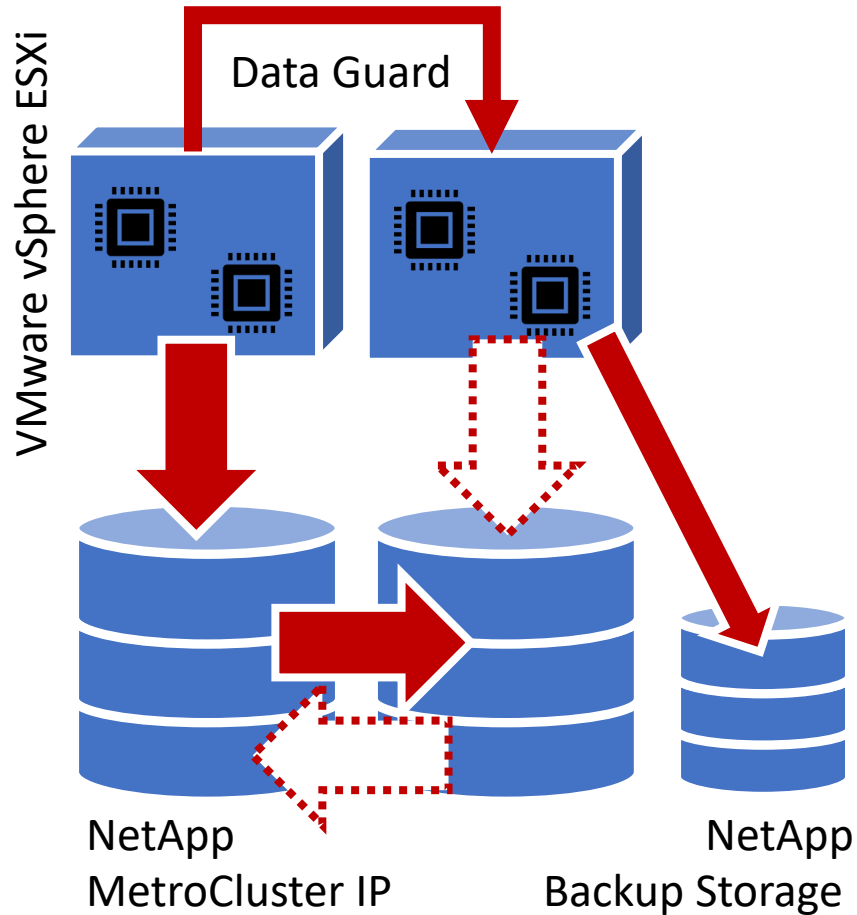
**performing  
databases**

# Solution-Jackstraws [Solution-Mikado]

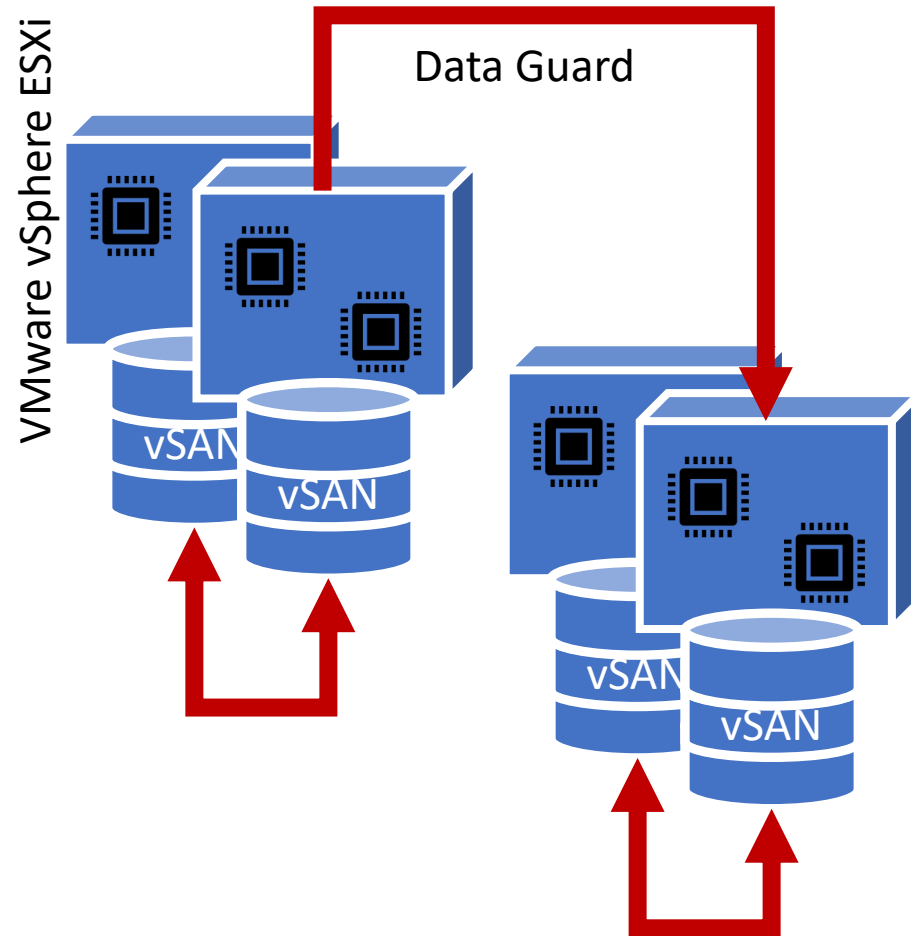


# Solution vs. Solution

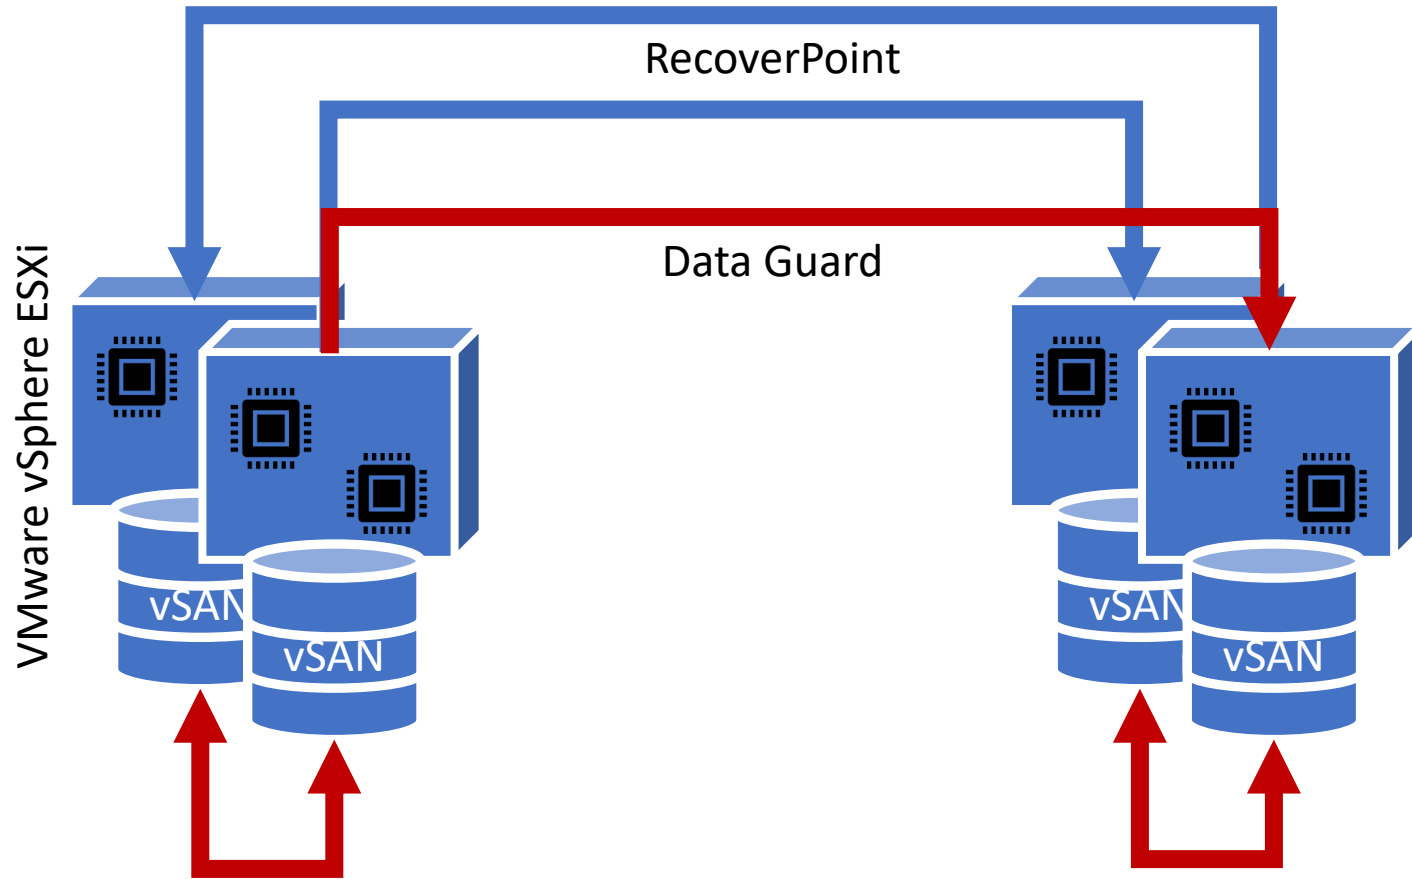
Proposed:



Wanted:



# Compromise



# Per-for-mance

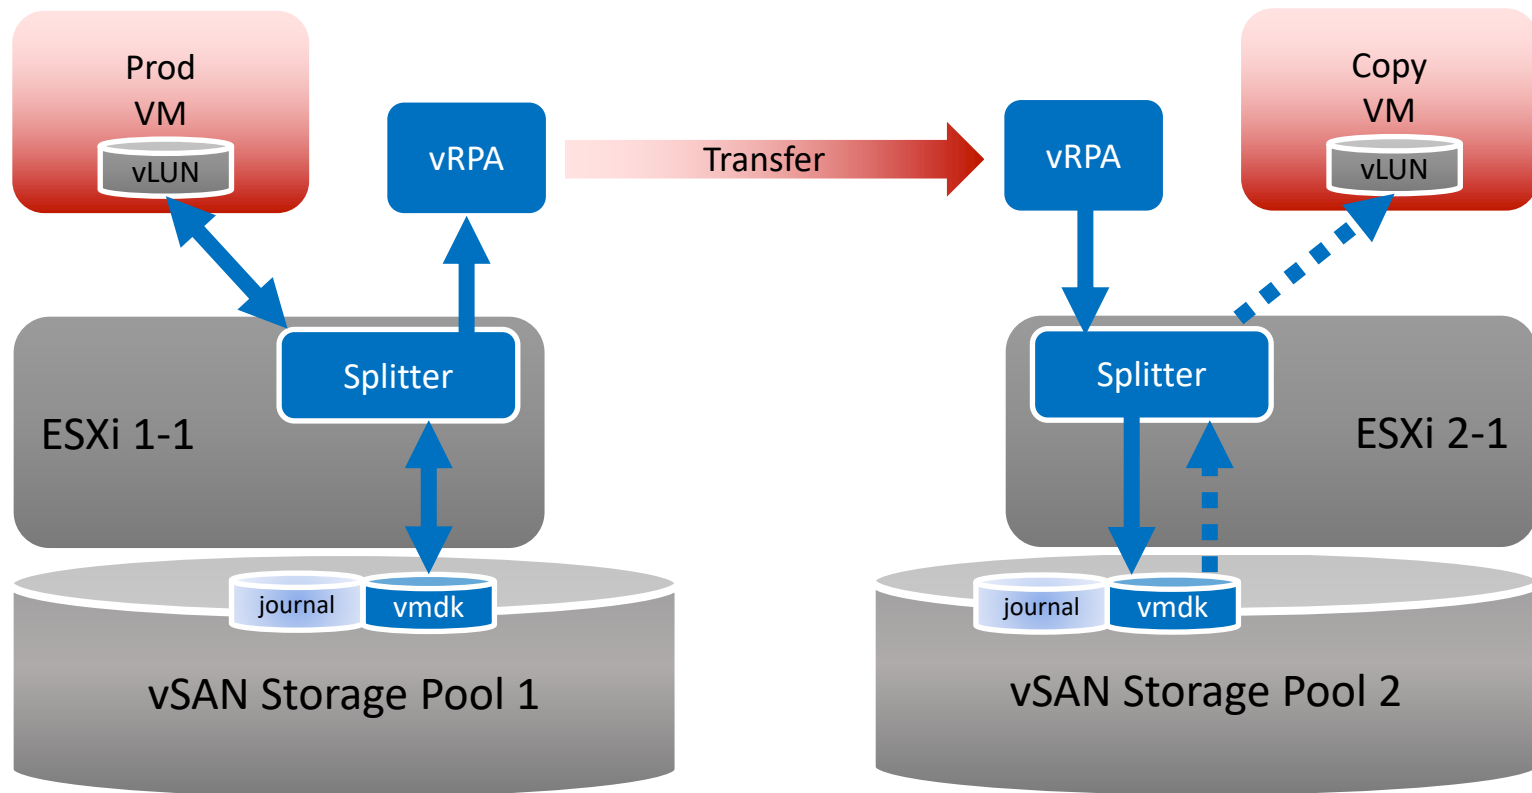
*[pə. 'fɔ:məns]*



# DellEMC RecoverPoint

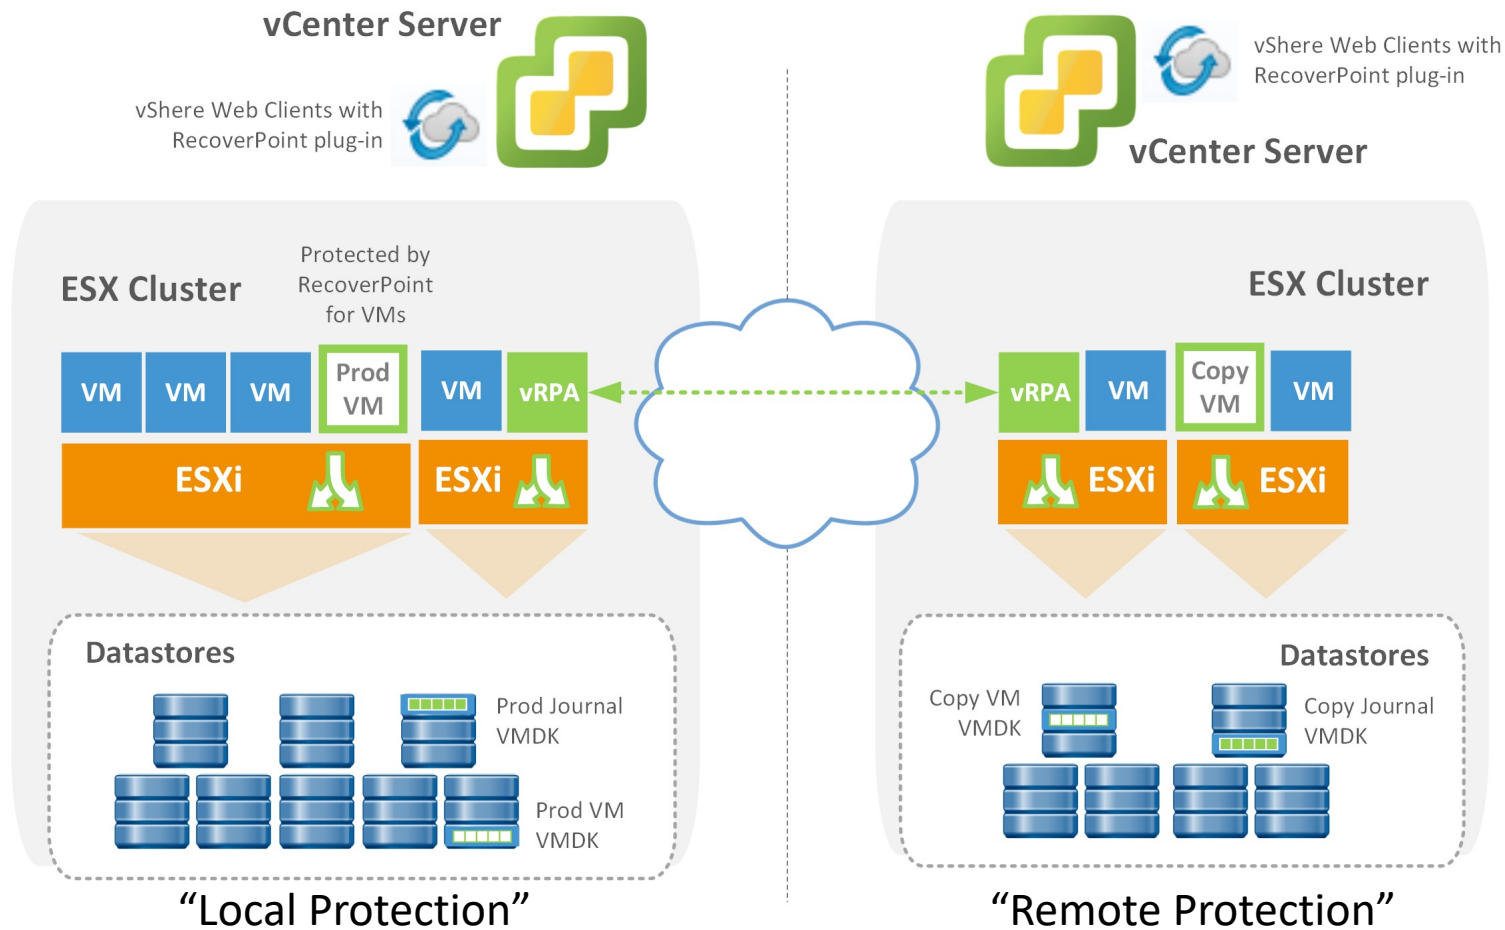
# Del EMC RecoverPoint

How it works, simplified



# Site A

# Site B



Source: Dell EMC RecoverPoint for VMs – Product Guide 5.3

## Legend



RecoverPoint for VMs Splitter



RecoverPoint for VMs Plug-in for vCenter Server



# Benchmarking

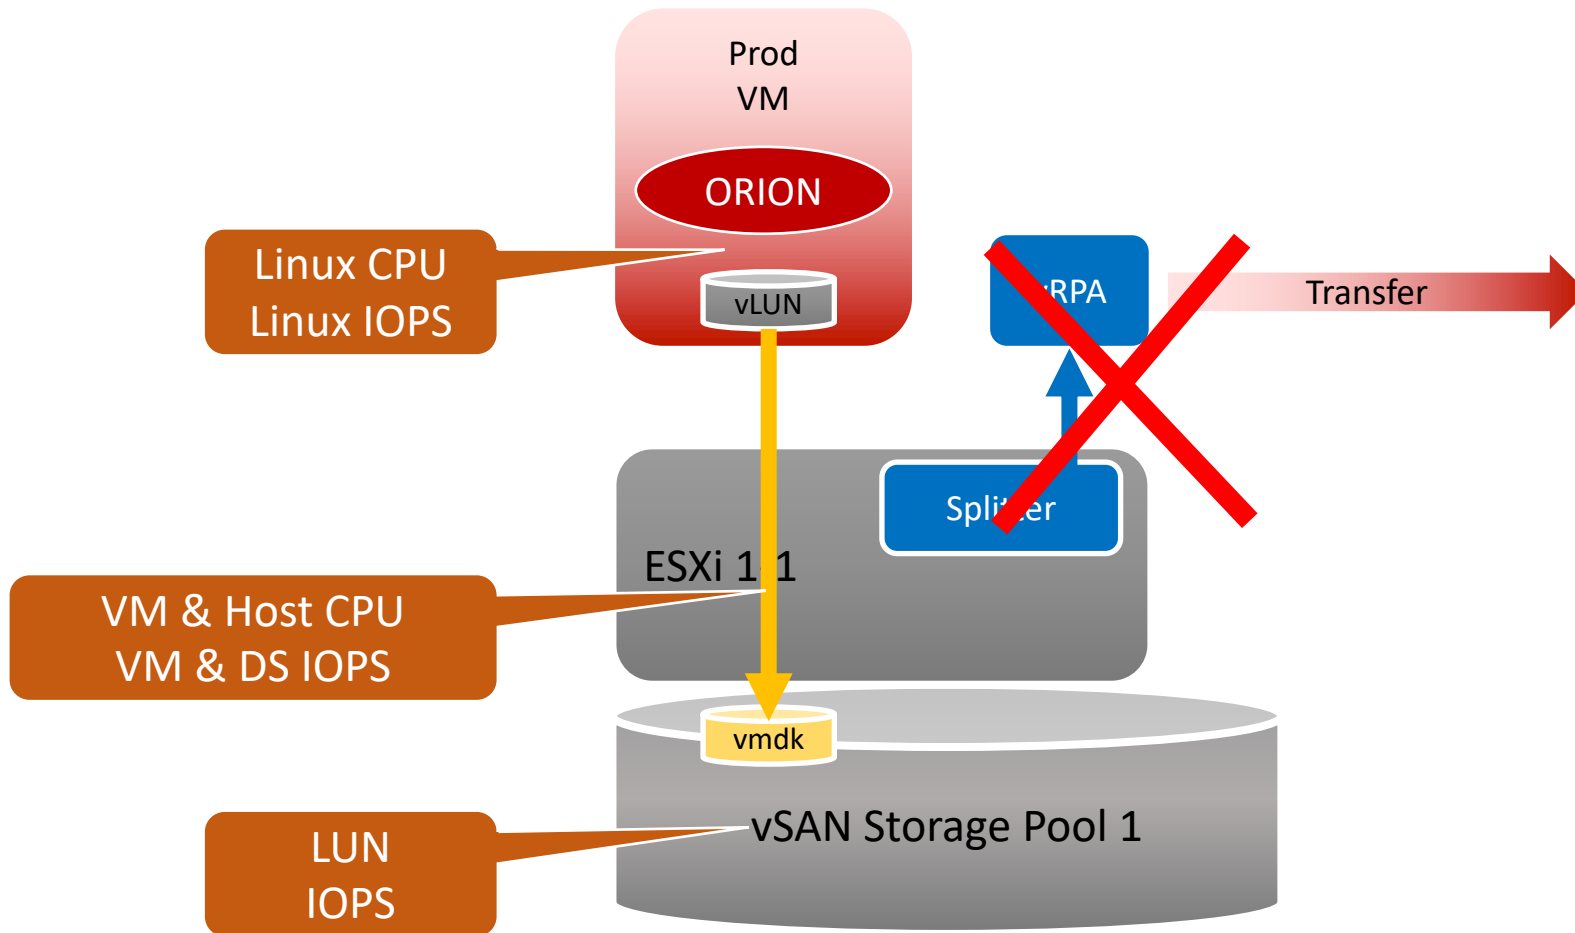
# Limit your Benchmark Scope!

- NOT:
  - ~~Which one is better?~~
  - ~~Comparison to another solution.~~
  - ~~Prove somebody wrong.~~
- Answer this:
  - Is the “Compromise” solution fast enough for proposed warehouse ops?
    - 2000 Transactions/second
    - 10.000 IOPS
    - <5ms Latency
  - ~~Is the “Compromise” solution robust enough? (not covered in this talk)~~
    - ~~“Pull the plug” under full transactional load~~
    - ~~Use the “Remote Protection” VMDKs to re-start the database~~
    - ~~How long does it take to crash-recover? (Or any other recovery necessary?)~~

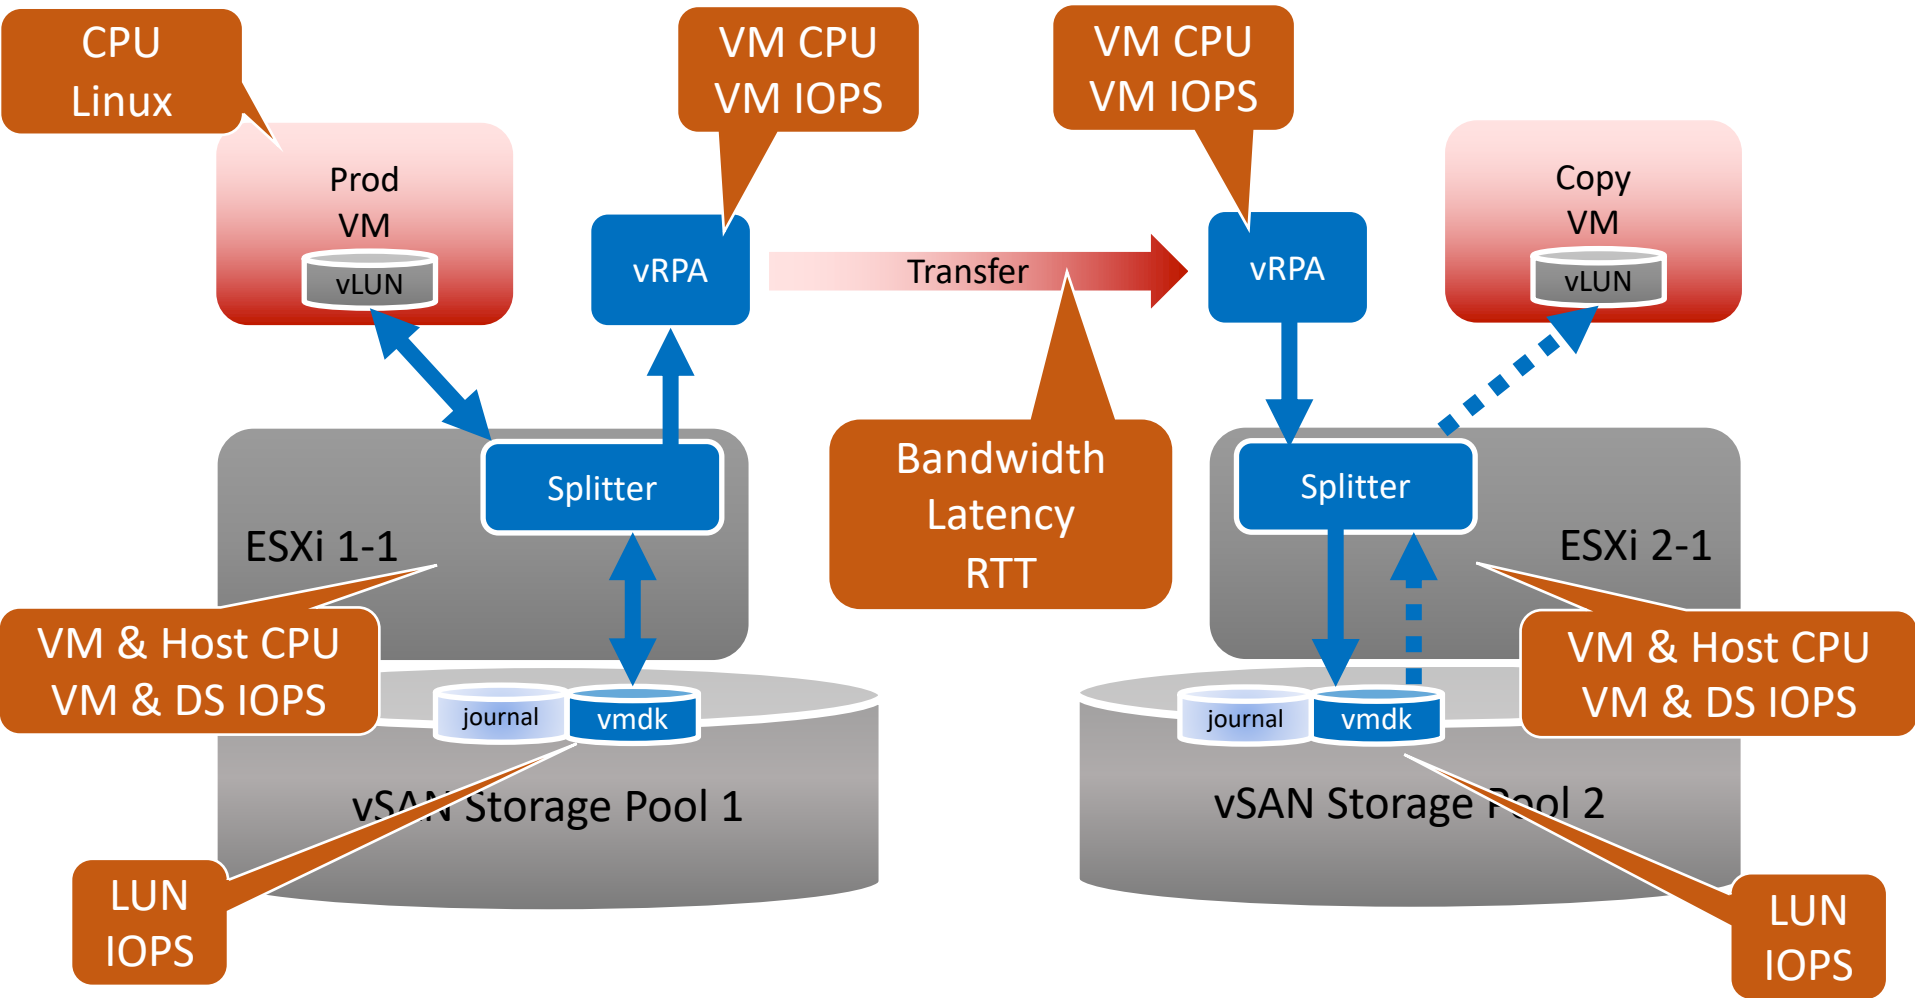
IO per second  
IO Latency



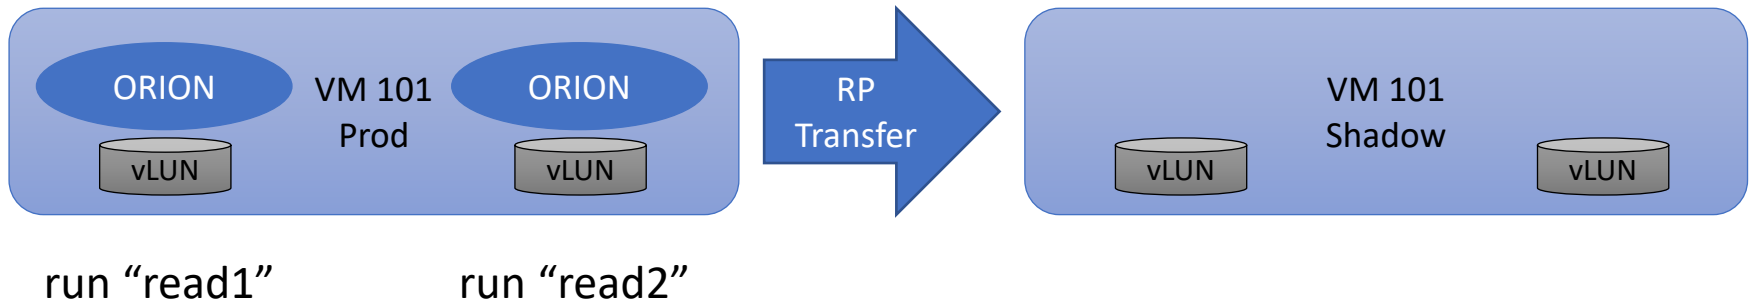
# I/O Stress Test - Baseline Setup



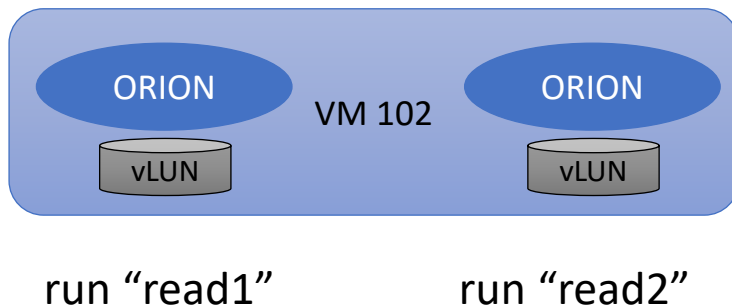
# Bottlenecks United



# ORION Test Setup



“Single-LUN-Performance” .... ☹



- Four ORION executions precisely at the same time
  - Two on a protected VM
  - Two on an unprotected VM
- Four LUNs overall
  - on the same vSAN, to max it out
  - Two protected
  - Two unprotected

# ORION Test Setup

`$ORACLE_HOME/bin/orion`

<code>-run advanced</code>	Test Profile
<code>-testname oltp-read</code>	Identification (LUN file, Output)
<code>-num_disks 20</code>	How much power (=threads, Q)
<code>-cache_size 8192</code>	How much cache to saturate
<code>-size_small 8</code>	Small IO size
<code>-size_large 16</code>	n/a (limited test matrix)
<code>-type rand</code>	IO Profile
<code>-simulate raid0</code>	n/a (here single LUN)
<code>-write 20</code>	Write Percentage (0=100% read)
<code>-duration 30</code>	How long per data point
<code>-matrix basic</code>	Which tests to run

# ORION Test Results

```
$ ls -1 oltp-read1*  
  oltp-read1.lun  
  oltp-read1_20210415_1320_hist.txt  
  oltp-read1_20210415_1320_iops.csv  
  oltp-read1_20210415_1320_lat.csv  
  oltp-read1_20210415_1320_mbps.csv  
  oltp-read1_20210415_1320_summary.txt  
  oltp-read1_20210415_1320_trace.txt
```

## Command line:

```
-run advanced -testname oltp-write -num_disks 16 -cache_size 8192 -size_small 8 -size_large 16 -type rand -simulate raid0
```

These options enable these settings:

Test: oltp-write

Small IO size: 8 KB

Large IO size: 16 KB

IO types: small random IOs, large random IOs

Sequential stream pattern: RAID-0 striping for all streams

Writes: 80%

Cache size: 8192 MB

Duration for each data point: 5 seconds

Small Columns: 0

Large Columns: 0, 1, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22,

Total Data Points: 41

Name: /dev/disk/by-id/BENCHMARK1

Size: 322122547200

1 files found.

Maximum Large MBPS=209.60 @ Small=0 and Large=32

Maximum Small IOPS=19563 @ Small=60 and Large=0

Small Read Latency: avg=1616.936 us, min=153.961 us, max=5673.406 us, std dev=574.261 us @ Small=60 and Large=0

Small Write Latency: avg=3428.564 us, min=1516.842 us, max=7062.202 us, std dev=678.150 us @ Small=60 and Large=0

Minimum Small Latency=852.863 usecs @ Small=1 and Large=0

Small Read Latency: avg=372.455 us, min=190.038 us, max=2469.994 us, std dev=117.754 us @ Small=1 and Large=0

Small Write Latency: avg=976.756 us, min=637.569 us, max=3245.912 us, std dev=206.006 us @ Small=1 and Large=0

Small Read / Write Latency Histogram @ Small=60 and Large=0

Latency:		# of IOs (read)	# of IOs (write)
0 - 128	us:	0 ( 0.00%)	0 ( 0.00%)
128 - 256	us:	177 ( 14.74%)	0 ( 0.00%)
256 - 512	us:	990 ( 97.17%)	0 ( 0.00%)
512 - 1024	us:	31 ( 99.75%)	2927 ( 62.85%)
1024 - 2048	us:	2 ( 99.92%)	1717 ( 99.72%)
2048 - 4096	us:	1 (100.00%)	13 (100.00%)
4096 - 268435456	us:	0 (100.00%)	0 (100.00%)

**oltp-write\_20210414\_1043\_summary.txt (END)**





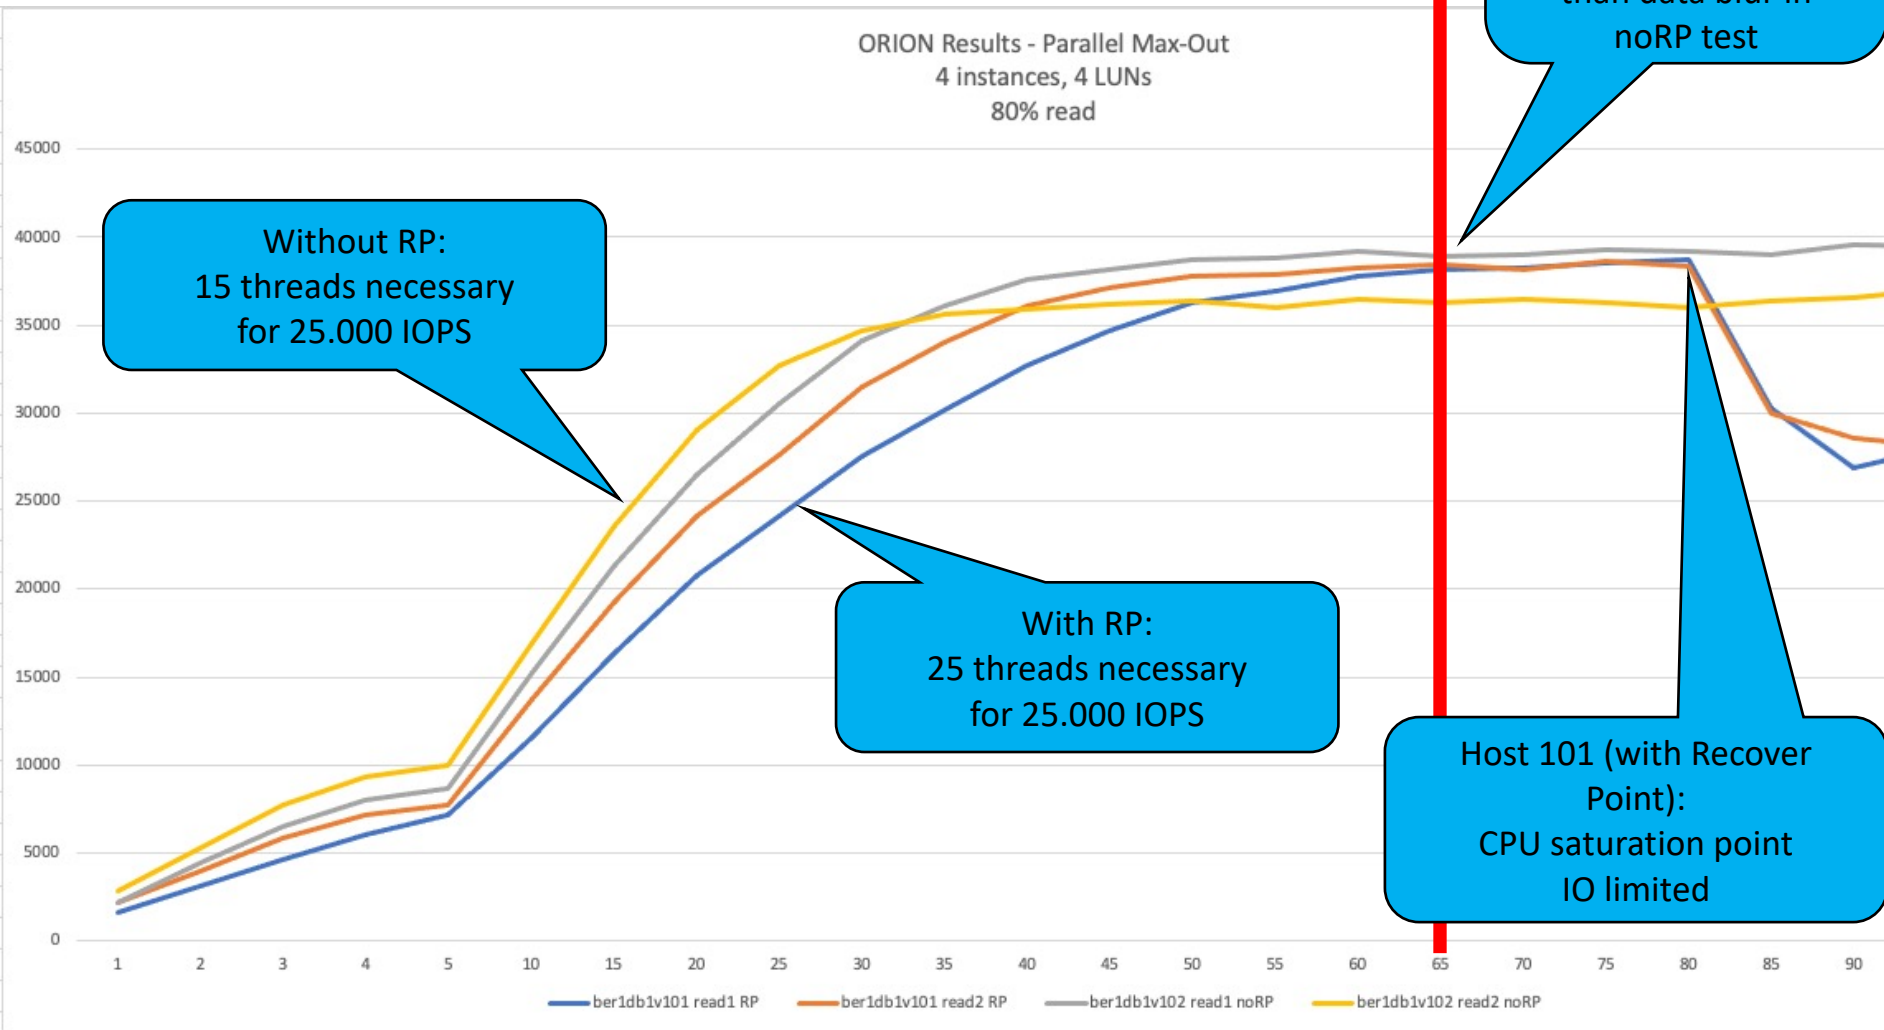
# ORION Test Results

This comma-separated-value file contains the rates sustained by small I/Os in IOPS.  
Each value corresponds to a data point test that used a fixed number of outstanding small and large I/Os.  
The number of outstanding small I/Os for a value is specified by its column header in the first row.  
The number of outstanding large I/Os for a value is specified by its row header in the first column.

Large/Small,	1,	2,	3,	4,	8,	12,	16,	20,	24,	28,	32,	36,	4
0,	1171,	2266,	3271,	4274,	6432,	7580,	9685,	11587,	13236,	14462,	15888,	16880,	1773
1													
2													
4													
6													
8													
10													
12													
14													
16													
18													
20													
22													
24													
26													
28													
30													
32													

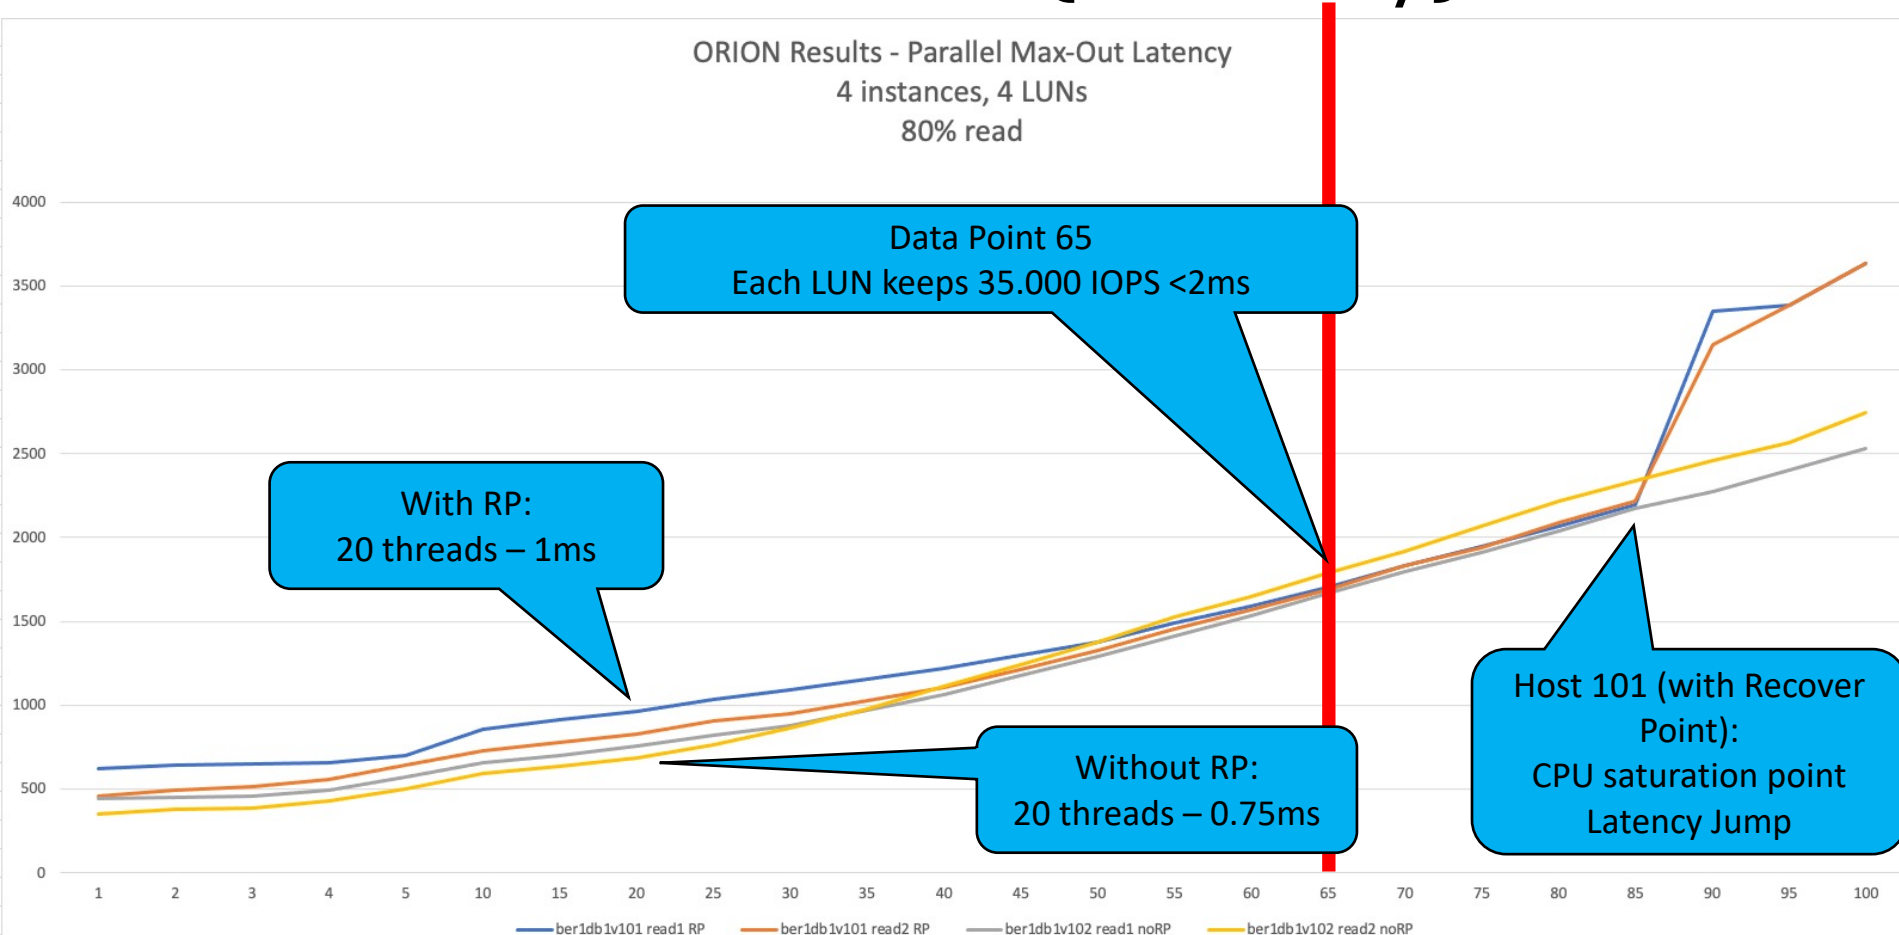
Example

# ORION Results (IOPS)



Remark: This is baseline, RecoverPoint result,  
and Max-Out test in one single graph

# ORION Results (Latency)



Remark: This is baseline, RecoverPoint result,  
and Max-Out test in one single graph

Maxed out: Each LUN delivers >35.000 IOPS  
RP impact smaller than data blur in noRP test

Baseline  
Test Run

no  
RecoverPoint

Max Out

Latency  
800ms?

with  
RecoverPoint

CPU saturation  
breakdown

Max Out

# ORION Results (Latency)

Latency Histogram for small IOs @ Small=65 and Large=0

Latency:		# of IOs (read)	# of IOs (write)
0 - 64	us:	0 ( 0.00%)	0 ( 0.00%)
64 - 128	us:	65 ( 0.01%)	0 ( 0.00%)
128 - 256	us:	115 ( 0.02%)	0 ( 0.00%)
256 - 512	us:	19743 ( 2.18%)	0 ( 0.00%)
512 - 1024	us:	420127 ( 48.10%)	0 ( 0.00%)
1024 - 2048	us:	466006 ( 99.04%)	0 ( 0.00%)
2048 - 4096	us:	8794 (100.00%)	98737 ( 43.08%)
4096 - 8192	us:	17 (100.00%)	129608 ( 99.64%)
8192 - 16384	us:	0 (100.00%)	824 (100.00%)
16384 - 268435456	us:	0 (100.00%)	0 (100.00%)

Latency Histogram for small IOs @ Small=100 and Large=0

Latency:		# of IOs (read)	# of IOs (write)
0 - 64	us:	0 ( 0.00%)	0 ( 0.00%)
64 - 128	us:	2 ( 0.00%)	0 ( 0.00%)
128 - 256	us:	71 ( 0.01%)	0 ( 0.00%)
256 - 512	us:	1131 ( 0.18%)	0 ( 0.00%)
512 - 1024	us:	11666 ( 1.95%)	0 ( 0.00%)
1024 - 2048	us:	410422 ( 64.21%)	0 ( 0.00%)
2048 - 4096	us:	234620 ( 99.80%)	3911 ( 2.36%)
4096 - 8192	us:	1093 ( 99.97%)	157912 ( 97.86%)
8192 - 16384	us:	0 ( 99.97%)	1368 ( 98.68%)
16384 - 32768	us:	0 ( 99.97%)	1 ( 98.68%)
32768 - 65536	us:	10 ( 99.97%)	0 ( 98.68%)
65536 - 131072	us:	0 ( 99.97%)	0 ( 98.71%)
131072 - 262144	us:	19 ( 99.97%)	91 ( 98.79%)
262144 - 524288	us:	185 (100.00%)	1905 ( 99.94%)
524288 - 1048576	us:	8 (100.00%)	92 (100.00%)
1048576 - 268435456	us:	0 (100.00%)	0 (100.00%)

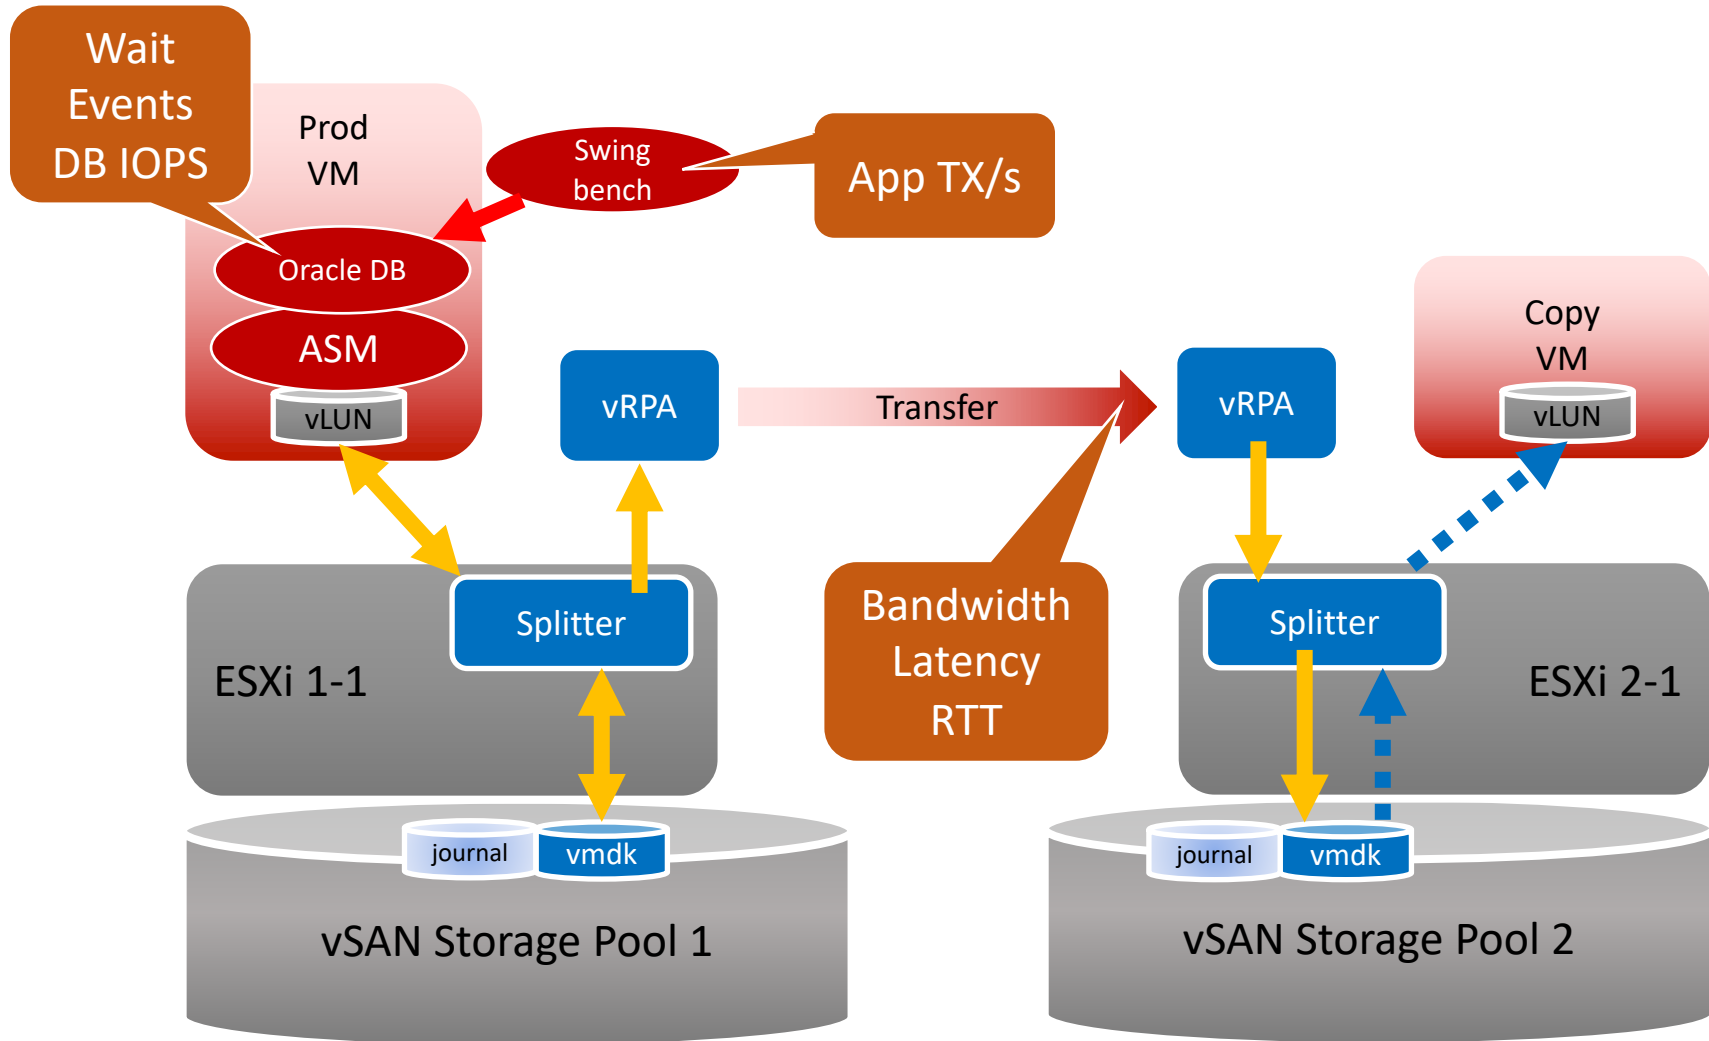
Average sucks ;)

Remark: This is baseline, RecoverPoint result, and Max-Out test in one single graph

# Transactions / Second



# Swingbench Setup

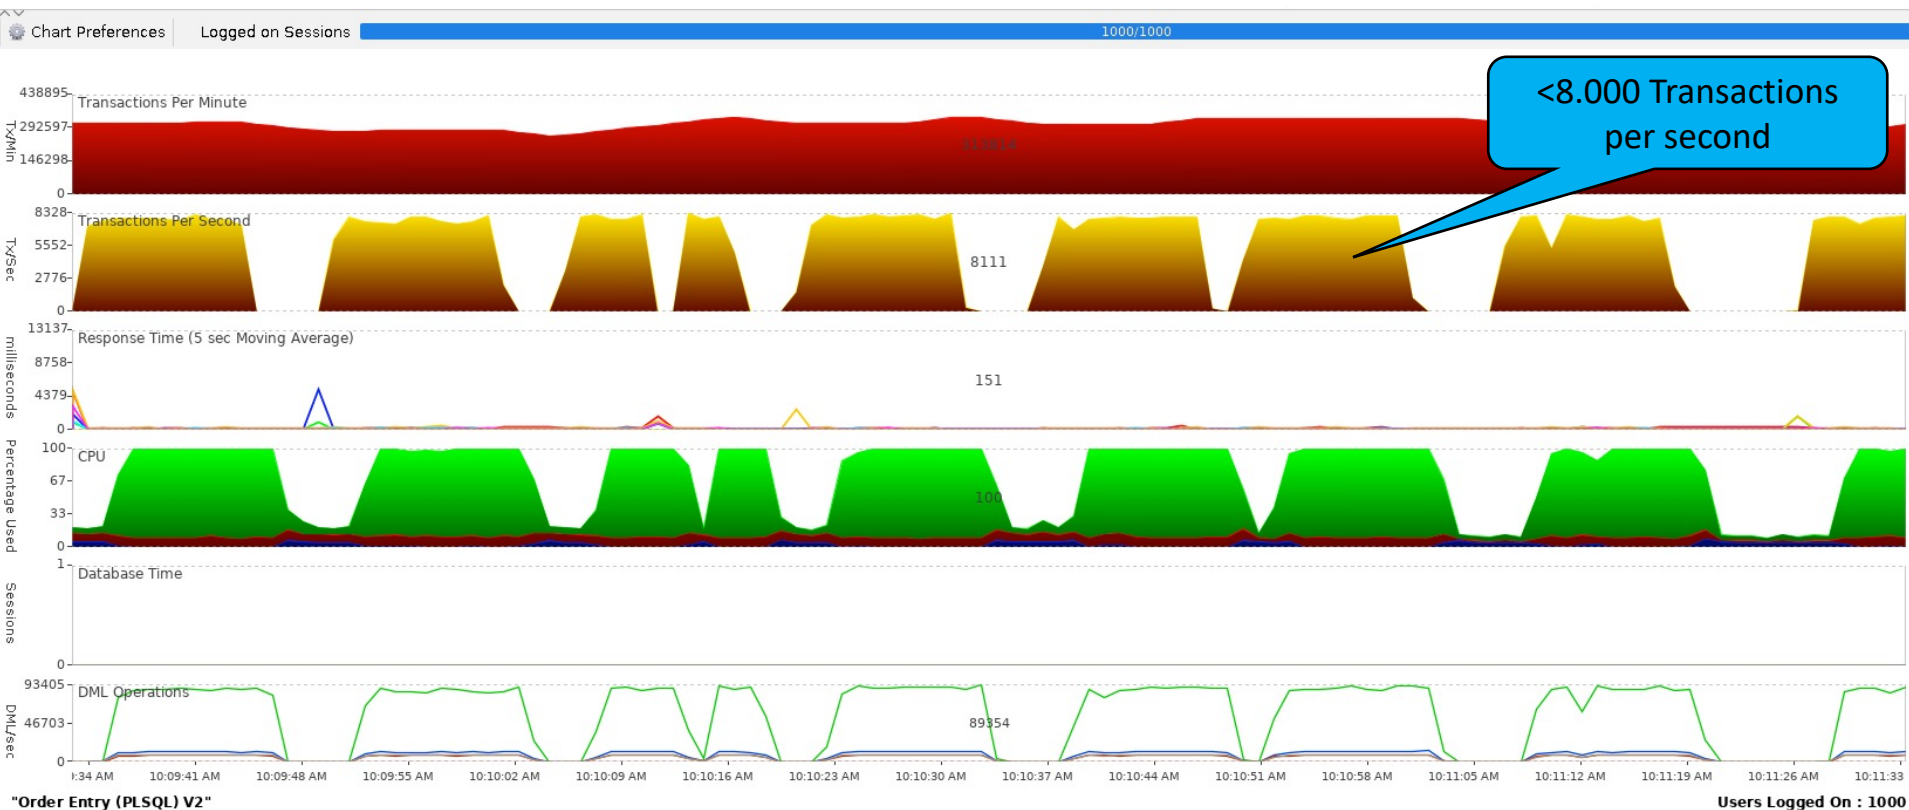


# Swingbench Setup

- Order Entry v 2.0
- IO-bound setup
  - „small“ SGA (10GB)
  - „big“ Tablespace (300GB)
- Increased No. of users until saturation (=1000)

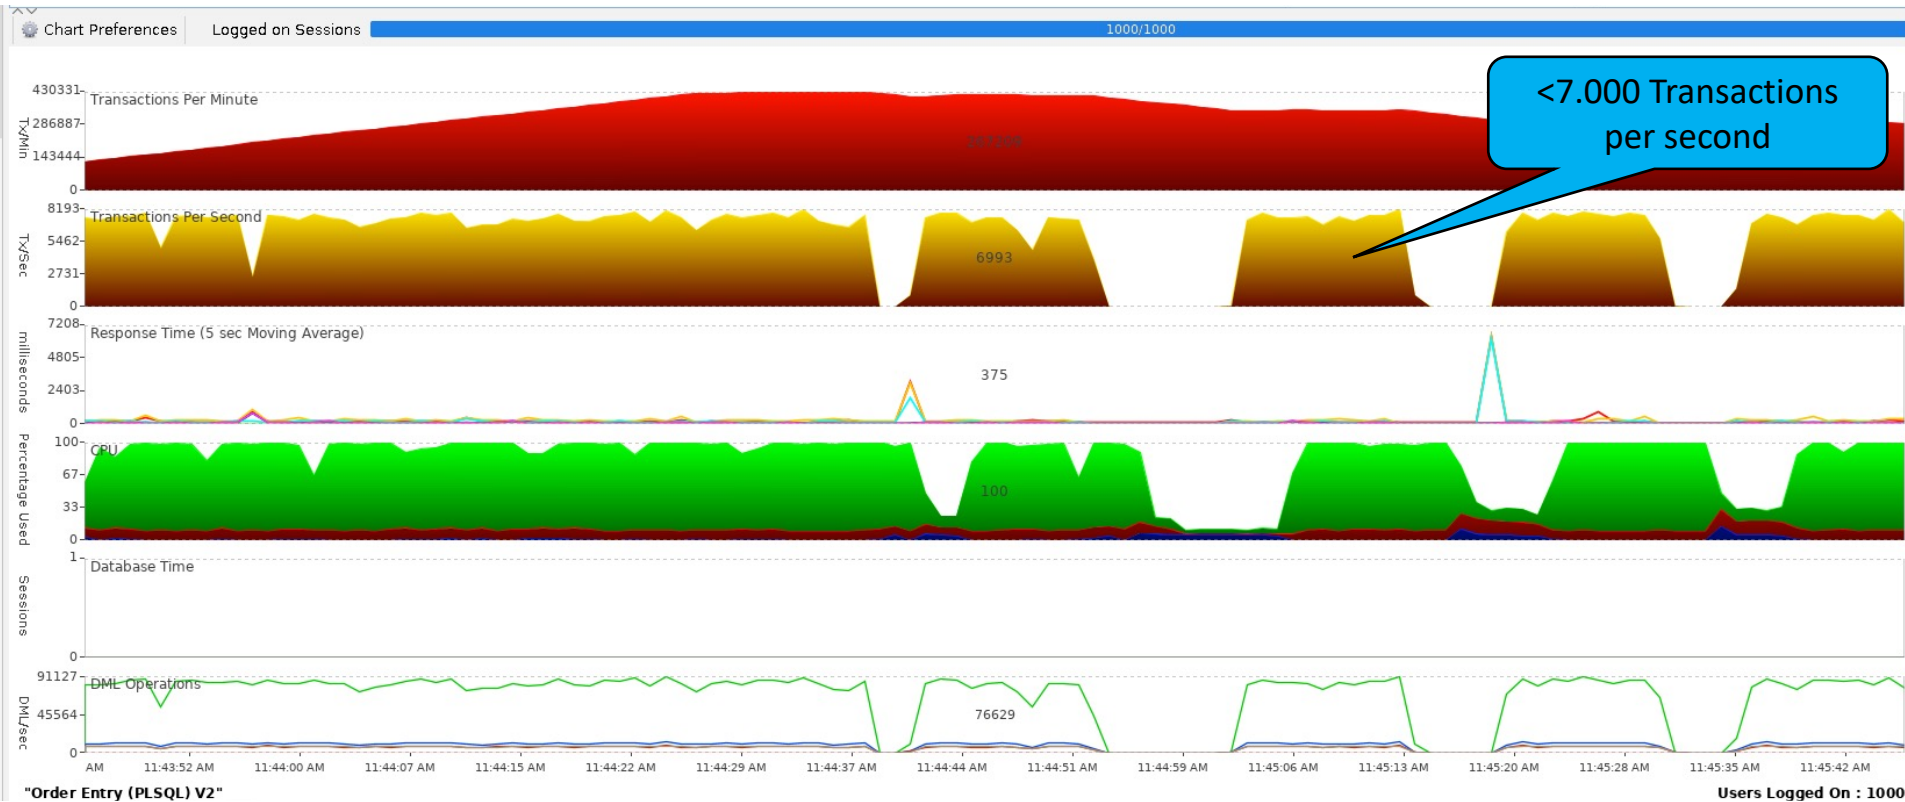


# Swingbench Baseline – no RP



Expected  
2.000+ TX/s

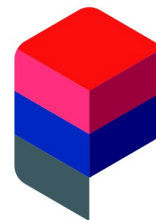
# Swingbench Results – with RP



Expected  
2.000+ TX/s

# Results

Metric	Warehouse Ops Request	w/o RecoverPoint	with Recover Point
IOPS	>10.000	>35.000 steep scaling	>35.000 relatively “flat” scaling
IO Latency	<5ms	avg. <2ms peak 5ms	avg. <2ms peak 800ms
Transactions/s	>2.000	>8.000	>7.000



**performing  
databases**

# Thank you!



**performing  
databases**



# Meet & Greet

[martin.klier@performing-db.com](mailto:martin.klier@performing-db.com)

[www.performing-databases.com](http://www.performing-databases.com)

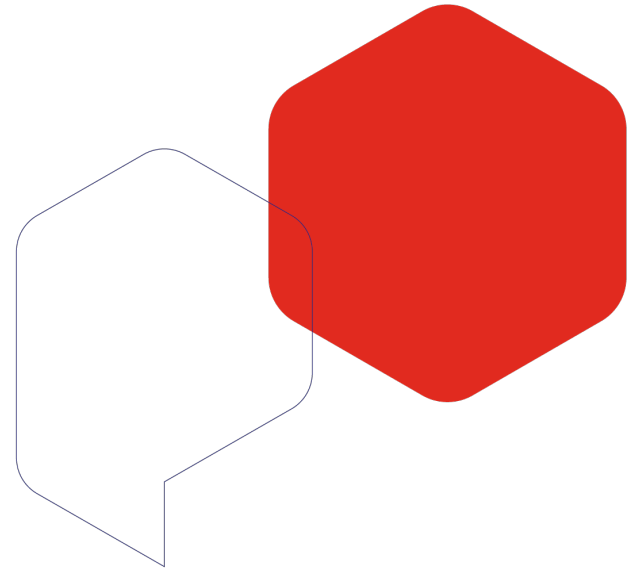
Many national // international events

# DOAG

Deutsche ORACLE -Anwendergruppe e.V.



**Efficiency is**  
to stay in touch!

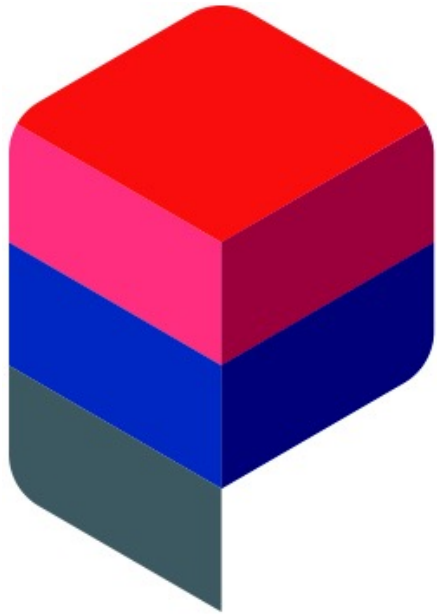


[martin.klier@performing-db.com](mailto:martin.klier@performing-db.com)

<http://www.performing-databases.com>



**performing  
databases**



**performing  
databases**