

# Lizenzkosten optimieren mit Flash Technologie

August 2016



## Data Warehouse System

- Oracle Enterprise Edition 47.5K
- RAC One Node 10.0K
- Partitioning 11.5K
- Active DataGuard 11.5K
- Advanced Compression 11.5K
- Diagnostic Pack 7.5K
- Tuning Pack 5.0K
- Total **104.5K**

## OLTP System

- Oracle Enterprise Edition 47.5K
- Real Application Cluster 23.0K
- Partitioning 11.5K
- Real Application Testing 11.5K
- Database In-Memory 23.0K
- Diagnostic Pack 7.5K
- Tuning Pack 5.0K
- Total **129.0K**

[1] <http://www.oracle.com/us/corporate/pricing/technology-price-list-070617.pdf>; 21<sup>st</sup> of July 2016



## Cloud Service

• Oracle Enterprise Edition	47.5K
• Real Application Cluster	23.0K
• Multitenant Option	17.5K
• Real Application Testing	11.5K
• Diagnostic Pack	7.5K
• Tuning Pack	5.0K
• Lifecycle Pack	12.0K
• Cloud Pack	7.5K
• Total	<b>131.5K</b>

Ältere ULA Verträge schliessen neue Optionen aus, sodass die pro-core Lizenzierung wichtig bleibt.



Durchschnittliche Lizenzkosten bei 50% Discount  
ca. 60k CHF pro Oracle Prozessor

x86 Prozessor Architektur hat vorteilhaften Core  
Faktor von 0.5, d.h.

- zwei x86 Cores = ein Oracle Prozessor

Oracle Lizenzkosten bei 50% Discount für typische  
x86 Server

rack units	#sockets	#cores	RAM [TB]	~ 60K per Oracle Processor
1	2	36	0.8	1.08M
3	2	36	1.5	1.08M
3	4	72	3	2.16M
5	4	72	3	2.16M
5	8	144	6	4.32M



## Reduktion der Anzahl Cores

- Prozessoren mit hoher Taktrate auswählen
- Wartezustände reduzieren
  - » Möglichst grosse Hauptspeicherkapazitäten
  - » Flash Technologie einsetzen



Server System	Kapazität	Zugriffszeit
CPU register	1 Kbyte	< 1 ns
▪ Level 1 Cache	128 kByte	1 ns
▪ Level 2 Cache	8 Mbyte	5 ns
▪ Level 3 Cache	32 MByte	15 ns
Hauptspeicher	4 TByte	100 ns
PCI attached flash	< 20 TByte	100 $\mu$ s = 100'000 ns
Storage System	Kapazität	Zugriffszeit
Cache		100 ns
SSD oder Flash Module		250'000 ns
Hard drive disks (HDD), 15k rpm		6'000'000 ns
Hard drive disks (HDD), 10k rpm		10'000'000 ns
Hard drive disks (HDD), 7.2k rpm		12'000'000 ns



Server System	Kapazität	Zugriffszeit
CPU register	1 Kbyte	< 1 ns
▪ Level 1 Cache	128 kByte	1 ns
▪ Level 2 Cache	8 Mbyte	5 ns
▪ Level 3 Cache	32 MByte	15 ns
<b>Hauptspeicher</b>	4 TByte	<b>100 ns</b>
PCI attached flash	< 20 TByte	100 µs = 100'000 ns
Storage System	Kapazität	Zugriffszeit
Cache		100 ns
<b>SSD oder Flash Module</b>		<b>250'000 ns</b>
Hard drive disks (HDD), 15k rpm		6'000'000 ns
Hard drive disks (HDD), 10k rpm		10'000'000 ns
Hard drive disks (HDD), 7.2k rpm		12'000'000 ns

A red arrow points upwards from the 'Cache' row to the 'Hauptspeicher' row, with the text 'x 2'500' written below it, indicating a 2500x increase in access time.

Server System	Kapazität	Zugriffszeit
CPU register	1 Kbyte	< 1 ns
▪ Level 1 Cache	128 kByte	1 ns
▪ Level 2 Cache	8 Mbyte	5 ns
▪ Level 3 Cache	32 MByte	15 ns
Hauptspeicher	4 TByte	100 ns
PCI attached flash	< 20 TByte	100 $\mu$ s = 100'000 ns
Storage System	Kapazität	Zugriffszeit
Cache		100 ns
SSD oder Flash Module		250 $\mu$ s = 250'000 ns
Hard drive disks (HDD), 15k rpm		6'000'000 ns
Hard drive disks (HDD), 10k rpm		8 ms = 8'000 $\mu$ s = 8'000'000 ns
Hard drive disks (HDD), 7.2k rpm		12'000'000 ns



x 32

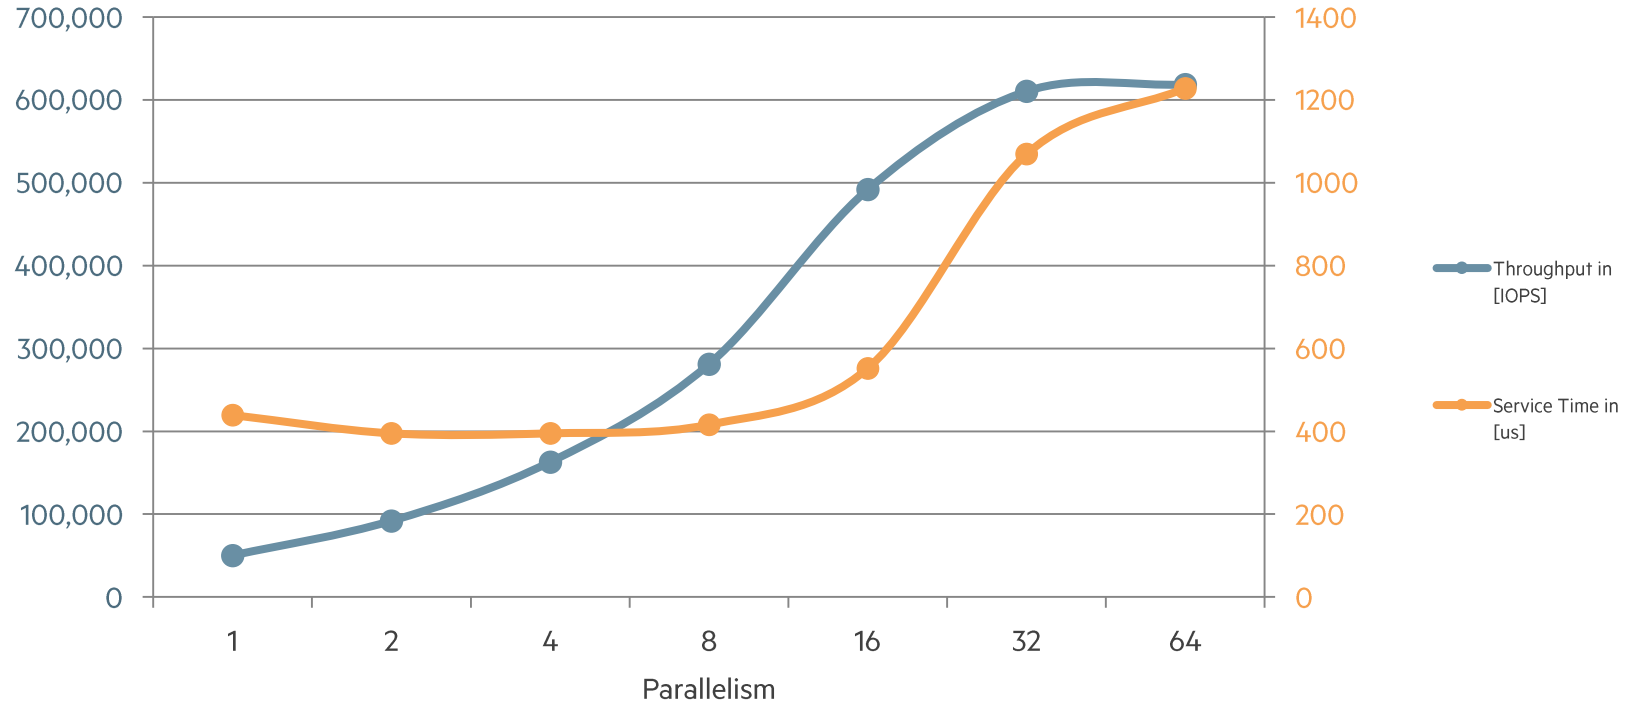




## peakmarks Performance Report

Run	Test	Workload	Nodes	Jobs	DOP	CPU busy [%]	CPU user [%]	CPU sys [%]	CPU iow [%]	CPU idle [%]	Physical read [dbps]	Physical read [us]	Physical read [iops]	Physical read [MBps]	Hitrate db flash [%]	Hitrate exa flash [%]	Elap time [s]
1	24	STO-RR	1	1	1	1	0	0	0	99	5.003E+04	4.393E+02	5.003E+04	3.910E+02	0	0	298
	25	STO-RR	1	2	1	1	1	0	0	99	9.189E+04	3.947E+02	9.189E+04	7.180E+02	0	0	300
	26	STO-RR	1	4	1	2	1	1	0	98	1.629E+05	3.944E+02	1.629E+05	1.272E+03	0	0	305
	27	STO-RR	1	8	1	5	2	2	0	95	2.808E+05	4.162E+02	2.808E+05	2.194E+03	0	0	306
	28	STO-RR	1	16	1	8	4	3	1	92	4.917E+05	5.517E+02	4.917E+05	3.841E+03	0	0	303
	29	STO-RR	1	32	1	12	6	5	3	88	6.102E+05	1.069E+03	6.102E+05	4.768E+03	0	0	300
	30	STO-RR	1	64	1	14	7	6	13	86	6.184E+05	1.228E+03	6.184E+05	4.831E+03	0	0	307
	31	STO-RR	1	128	1	14	7	5	12	86	5.783E+05	1.049E+03	5.783E+05	4.518E+03	0	0	596
	32	STO-RR	1	256	1	17	7	8	34	83	6.015E+05	1.537E+03	6.015E+05	4.700E+03	0	0	330

# Monitoring von I/O wait time




# Monitoring von I/O wait time



## Schlecht konfiguriertes System

### nmon Monitor



## Optimiertes System

### PCI attached flash NVMe

- 1.1 Mio IOPS Oracle
- Datenbank Blockgrösse 8 KByte
- CPU komplett ausgelastet
- kernel mode 49%
- iowait less < 1%





peakmarks

simple. representative. fast.