RED HAT SUMMIT
10 YEARS and counting
SAN FRANCISCO | APRIL 14-17, 2014

Best Practices for Deploying & Tuning Oracle Database 12c on RHEL6

Roger Lopez, Principal Software Engineer, Red Hat
Sanjay Rao, Principal Performance Engineer, Red Hat
April, 2014

Agenda

Agenda

- Deploying Oracle Database 12c on RHEL6
  - Preinstallation Requirements
  - Installation of Oracle Grid Infrastructure 12c
  - Installation of Oracle 12c R1 Database Software
  - Installation of an Oracle Database using DBCA
- Tuning for Oracle Databases
  - I/O Tuning
  - Memory Tuning

Preinstallation Requirements

Deploying Oracle Database 12c on RHEL6:
Preinstallation Requirements – Network Configuration

- Public Network Configuration
  - Two network interfaces bonded together to provide high availability.
    - If using RHEL 6.2 or below, disable Network Manager and add “NM_CONTROLLED=no” to corresponding slave ifcfg-ethX and bond file ifcfg-bondX
    - If using RHEL 6.3 or higher and disabling Network Manager, besides adding “NM_CONTROLLED=no” to corresponding slave ifcfg-ethX and bond file ifcfg-bondX it is required to add “NM_BOND_VLAN_ENABLED=yes” to /etc/sysconfig/network
    - Starting with RHEL 6.4, nm-connection-editor was introduced to handle the creation of bonding configurations

Deploying Oracle Database 12c on RHEL6:
Preinstallation Requirements – NTP Configuration

- NTP – an operating system daemon which sets and maintains the system time, synchronizing with Internet standard time servers.
- Within the /etc/ntp.conf file, add at least one NTP server (3 recommended)
  - server <IP_ADDRESS> iburst
    - iburst option is used to speed up the time in which it takes to properly sync with the NTP servers

Deploying Oracle Database 12c on RHEL6:
Preinstallation Requirements – Required Packages

- Required packages include the following Group packages:
  - @Base
  - @Core

Required Packages	
cloog-ppl	libXxf86misc
compat-libcap1	libXxf86vm
compat-libstdc++-33	libaio-devel
cpp	libdmx
gcc	libstdc++-devel
gcc-c++	mpfr
glibc-devel	make
glibc-headers	ppl
kernel-headers	xorg-x11-utils
libXmu	xorg-x11-xauth
libXt	libXv
ksh	libXxf86dga

Deploying Oracle Database 12c on RHEL6:
Preinstallation Requirements – SELinux

Why is SELinux important? It enhances security by determining if an action from a particular process should be granted thus protecting against vulnerabilities within a system.

- Does Oracle support SELinux? Yes! As of Oracle version 11.2
- Enabling SELinux does not currently confine Oracle database processes
  - Take advantage of the SELinux policy created by Red Hat that does properly confine Oracle databases. It is currently in tech preview and not officially supported.
    - https://access.redhat.com/site/articles/742603

Deploying Oracle Database 12c on RHEL6:
Preinstallation Requirements – Firewall

Enable the Firewall!

- What ports need to be opened?
  - TNS Listener port, default 1521
  - Enterprise Manager Express, default port 5500
- How do I enable?
  - Best way is to allow access only to the specific database clients
  - Example iptables firewall rule
    - -A INPUT -m state --state NEW -m tcp -p tcp -s <IP_OF_CLIENT> --dport <port> -j ACCEPT

RED HAT SUMMIT
10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

Kernel Parameters

Deploying Oracle Database 12c on RHEL6:
Kernel Parameters

There are many kernel parameters that affect performance of Oracle databases.

Kernel Parameters	
shmmax	shmall
shmmni	Net.ipv4.ip_local_port_range
sem	swappiness
dirty_ratio	dirty_background_ratio
dirty_expire_centisecs	dirty_writeback_centisecs
fs.file-max	fs.aio-max-nr
net.ipv4.ip_local_port_range	net.core.rmem_default
net.core.rmem_max	net.core.wmem_default
net.core.wmem_max	

User Accounts & Groups

User Accounts & Groups

Preinstallation Requirements – grid & oracle user

Recommendation to separate out administrative roles creating users and groups

- What users to create?
  - Oracle user – that handles database related administrative tasks
  - Grid user – that handles grid infrastructure and ASM related tasks
  - Oracle recommends the following groups to be created:
  - dba, asmdba, asmoper, asmadmin, oper, backupdba, dgdba, kmdba
  - Grid contains: dba, asmadmin, asmdba, asmoper
  - Oracle contains: dba, oper, asmdba, asmoper, backupdba, dgdba, kmdba

User Accounts & Groups

Preinstallation Requirements – grid & oracle user shell limits

Setting shell limits to the oracle & grid user is done to prevent Oracle from exhausting all of the OS resources on the system if under a very heavy workload

- Where do I set my shell limits for grid and oracle user?
  - NOT in /etc/security/limits.conf
  - Set in /etc/security/limits.d/99-grid-oracle-limits.conf (file name can be changed but number dictates when file is read.)
- Why not set in /etc/security/limits.conf?
  - After the limits.conf file is read, the files within /etc/security/limits.d are read. If two files contain the same entry, the entry read last takes precedence.

User Accounts & Groups

Preinstallation Requirements – grid & oracle user shell limits

After the shell limits are set within the /etc/security/limits.d/99-grid-oracle-limits.conf file, a file within /etc/profile.d/ needs to be created to ensure shells such as the bash shell are started with the correct ulimit value.

- I thought I already set my ulimit values in the 99-grid-oracle-limits.conf?
  - Without creating an /etc/profile.d/oracle-grid.sh script, when logging into as grid or oracle user, the ulimit value being used would always the soft limit of 1024. Oracle recommends when logging into grid or oracle for the ulimit be the hard limit value of 65536

User Accounts & Groups

Preinstallation Requirements – grid & oracle user shell limits

/etc/security/limits.d/99-grid-oracle-limits.conf

```
oracle soft nproc 16384
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
oracle soft stack 10240
oracle hard stack 32768
grid soft nproc 16384
grid hard nproc 16384
grid soft nofile 1024
grid hard nofile 65536
grid soft stack 10240
grid hard stack 32768
```

User Accounts & Groups

Preinstallation Requirements – grid & oracle user shell limits

/etc/profile.d/oracle-grid.sh

```
#Setting the appropriate ulimits for oracle and grid user
if [ $USER = "oracle" ]; then
    if [ $SHELL = "/bin/ksh" ]; then
       ulimit -u 16384
       ulimit -n 65536
    else
       ulimit -u 16384 -n 65536
    fi
fi

if [ $USER = "grid" ]; then
    if [ $SHELL = "/bin/ksh" ]; then
       ulimit -u 16384
       ulimit -n 65536
    else
       ulimit -u 16384 -n 65536
    fi
fi
```

Storage Configuration

Storage Configuration

Setting up DM-Multipath

Device Mapper multipath is used because it provides:

- High availability
- I/O load balancing
- Persistent Naming

Storage Configuration
Setting up DM-Multipath

Keys to successfully using DM-Multipath

- Blacklist any local disk on the system
- Identify the dm- device, size, and WWID of each device mapper volume for Oracle.
- Create a multipaths section within /etc/multipath.conf that creates an alias for each device mapper volume to achieve persistent naming

Storage Configuration
Setting up DM-Multipath

Keys to successfully using DM-Multipath

- Blacklist any local disk on the system

```
# scsi_id --whitelisted --replace-whitespace --device=/dev/sda
3600508b1001030353434363646301200
```

Blacklist section within /etc/multipath.conf

```
blacklist {
        wwid 3600508b1001030353434363646301200
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
}
```

Storage Configuration
Setting up DM-Multipath

Keys to successfully using DM-Multipath

- Identify the dm- device, size, and WWID of each device mapper volume for Oracle.

```
# multipath -ll
```

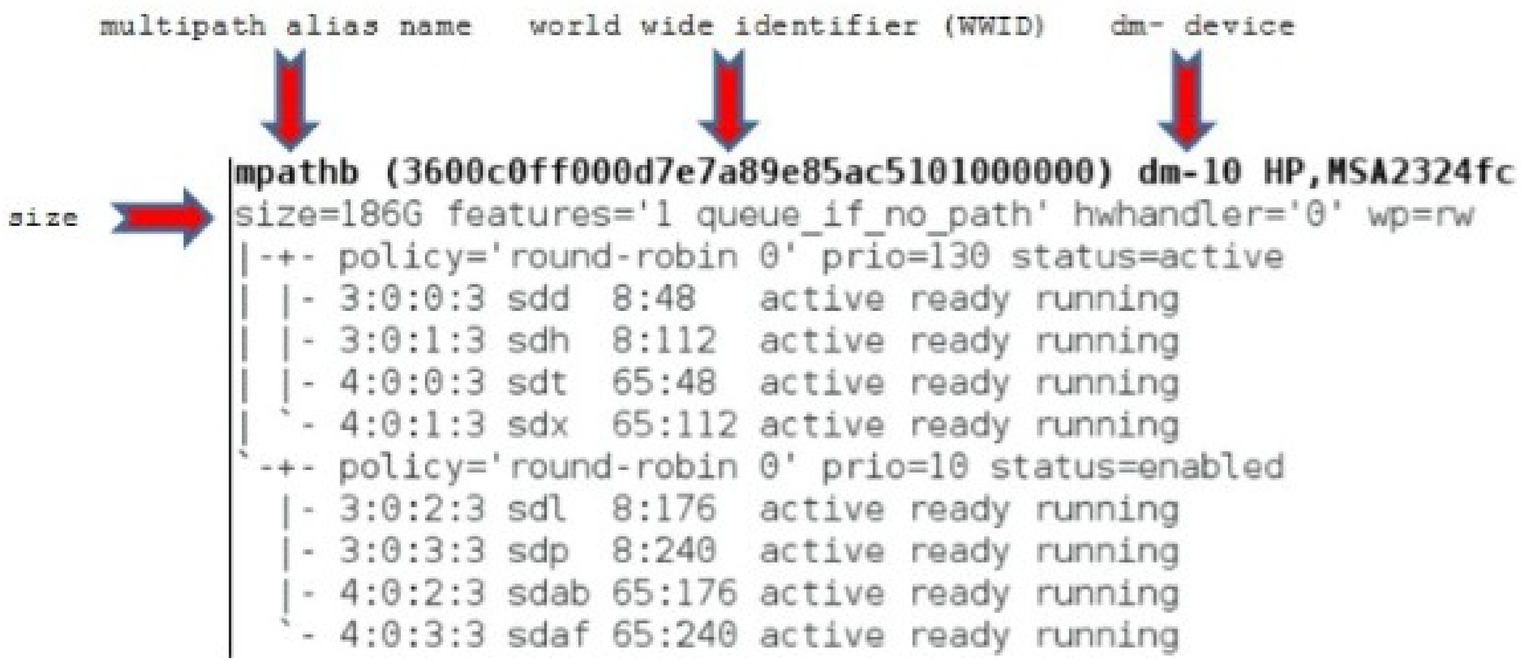

Storage Configuration
Setting up DM-Multipath

Keys to successfully using DM-Multipath

- Create a multipaths section within /etc/multipath.conf that creates an alias for each device mapper volume to achieve persistent naming

```
multipaths {
        multipath {
                wwid                    3600c0ff000d7e7a899d8515101000000
                alias                   db1
        }
        multipath {
                wwid                    3600c0ff000dabfe5a7d8515101000000
                alias                   db2
        }
        multipath {
                wwid                    3600c0ff000d7e7a8dbd8515101000000
                alias                   fra
        }
        multipath {
                wwid                    3600c0ff000dabfe5f4d8515101000000
                alias                   redo
        }
}
```

Storage Configuration
Configuring udev rules

- Why are udev rules required?
  - Ensures that when the system comes up the Oracle volumes have the correct owner, group, and permissions
- How do I set udev rules?
  - Capture the Device Mapper Universally Unique Identifier (DM_UUID) for each DM volume.
  - Create a file labeled 99-oracle-asmdevices.rules within /etc/udev/rules.d
  - Apply and test the rules created within the 99-oracle-asmdevices.rules with udevadm test

Storage Configuration
Configuring udev rules – Cont'd

```
# for i in db1p1 db2p1 frap1 redop1; do printf "%s %s\n" "$i" "$(udevadm
info --query=all --name=/dev/mapper/$i | grep -i dm_uuid)"; done
db1p1 E: DM_UUID=part1-mpath-3600c0ff000d7e7a899d8515101000000
db2p1 E: DM_UUID=part1-mpath-3600c0ff000dabfe5a7d8515101000000
frap1 E: DM_UUID=part1-mpath-3600c0ff000d7e7a8dbd8515101000000
redop1 E: DM_UUID=part1-mpath-3600c0ff000dabfe5f4d8515101000000
```

- /etc/udev/rules.d/99-oracle-asmdevices.rules

```
KERNEL=="dm-*",ENV{DM_UUID}=="part1-mpath-
3600c0ff000dabfe5f4d8515101000000",OWNER="grid",GROUP="asmadmin",MODE="06
60"
```

Storage Configuration

Configuring udev rules – Cont'd

```
# for i in db1p1 db2p1 frap1 redop1; do printf "%s %s\n" "$i" "$(ls
-ll /dev/mapper/$i)"; done
db1p1 lrwxrwxrwx. 1 root root 8 May 20 20:39 /dev/mapper/db1p1 -> ../dm-11
db2p1 lrwxrwxrwx. 1 root root 8 May 20 20:39 /dev/mapper/db1p1 -> ../dm-12
frap1 lrwxrwxrwx. 1 root root 8 May 20 20:39 /dev/mapper/db1p1 -> ../dm-13
redop1 lrwxrwxrwx. 1 root root 8 May 20 20:39 /dev/mapper/db1p1 -> ../dm-14
```

```
 # udevadm test /sys/block/dm-11
[ ... Output Abbreviated ... ]
udevadm_test: DM_NAME=db1p1
udevadm_test: DM_UUID=part1-mpath-3600c0ff000d7e7a86485ac5101000000
udevadm_test: DM_SUSPENDED=0
udevadm_test: DEVLINKS=/dev/mapper/db1p1 /dev/disk/by-id/dm-name-db1p1
/dev/disk/by-id/dm-uuid-part1-mpath-3600c0ff000d7e7a86485ac5101000000
/dev/block/253:11
udevadm_test: ID_FS_TYPE=oracleasm
```

Storage Configuration

Configuring udev rules – Cont'd

```
# ls -lh /dev/dm-11
brw-rw----. 1 grid asmadmin 253, 11 Jun  6 20:59 /dev/dm-11
```

RED HAT SUMMIT
10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

Oracle Grid Infrastructure 12c

Oracle Grid Infrastructure 12c

Key Points to a Successful Installation

- Why install Oracle Grid Infrastructure 12c?
  - It is a required component for the use of Oracle ASM diskgroups.
- Key Points when installing Oracle Grid Infrastructure 12c
  - Redundancy Level (External, Normal, High)
  - Allocation Unit (1MB, 4MB, 8MB, 16MB, 32MB)
    - Larger AU Size is used to decrease amount of extents Oracle needs to manage. Less extents reduces CPU utilization and memory consumption.
    - No “best size” for AU, completely depends on Oracle workload, I/O size per transaction, and overall diskgroup size.

RED HAT SUMMIT
10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

ASM Configuration Assistant (ASMCA)

ASM Configuration Assistant (ASMCA)

Key Points to creating ASM Diskgroups

- Key Points when creating an ASM Diskgroup
  - Redundancy Level (External, Normal, High)
  - Allocation Unit Size (1MB, 4MB, 8MB, 16MB, 32MB) located in Advanced Options
    - Larger AU Size is used to decrease amount of extents Oracle needs to manage. Less extents reduces CPU utilization and memory consumption.
    - No “best size” for AU, completely depends on Oracle workload, I/O size per transaction, and overall diskgroup size.

RED HAT SUMMIT
10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

Oracle Database 12c

Oracle Database 12c Installer

Key Points to a Successful Installation

- At the Installation window, select 'Install database software only'
- At the Grid Installations Options window, ensure to select the appropriate database installation i.e. 'single instance installation'

Oracle Database 12c Database Configuration Assistant (DBCA)

Key Points to a Successful Installation

- At the Database Identification window, allows for the creation of a container database (CDB) to hold X number of pluggable databases (PDB). The ability to create a traditional database is still available.
- Storage Locations window, critical for setting Database Files location, Redo Log file locations within the button labeled “Multiplex Redo Logs and Control Files”, and Fast Recovery Area (Flash Recovery Area)
- Initialization Parameters window, use 'Custom Settings' to set to Automatic Shared Memory Management (ASMM) for best performance. Set SGA and PGA size appropriately.

RED HAT SUMMIT
10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

Oracle Database Tuning

Redo Logs

Checkpoints, Redo Log Size, Multiplexing

- Redo Log – stores all the changes made to an Oracle database data files and control files. Generally consists of two or more files for the storing these changes
- Multiplexing - used to protect the redo logs from a failure. Multiplexed redo logs allows for the redo log to create multiple identical copies of itself to be maintained ideally in separate locations within a storage array.
- Checkpoint – marks a point in time where if there was a failure, the instance recovery must begin at this point.

Redo Logs

Checkpoints, Redo Log Size, Multiplexing

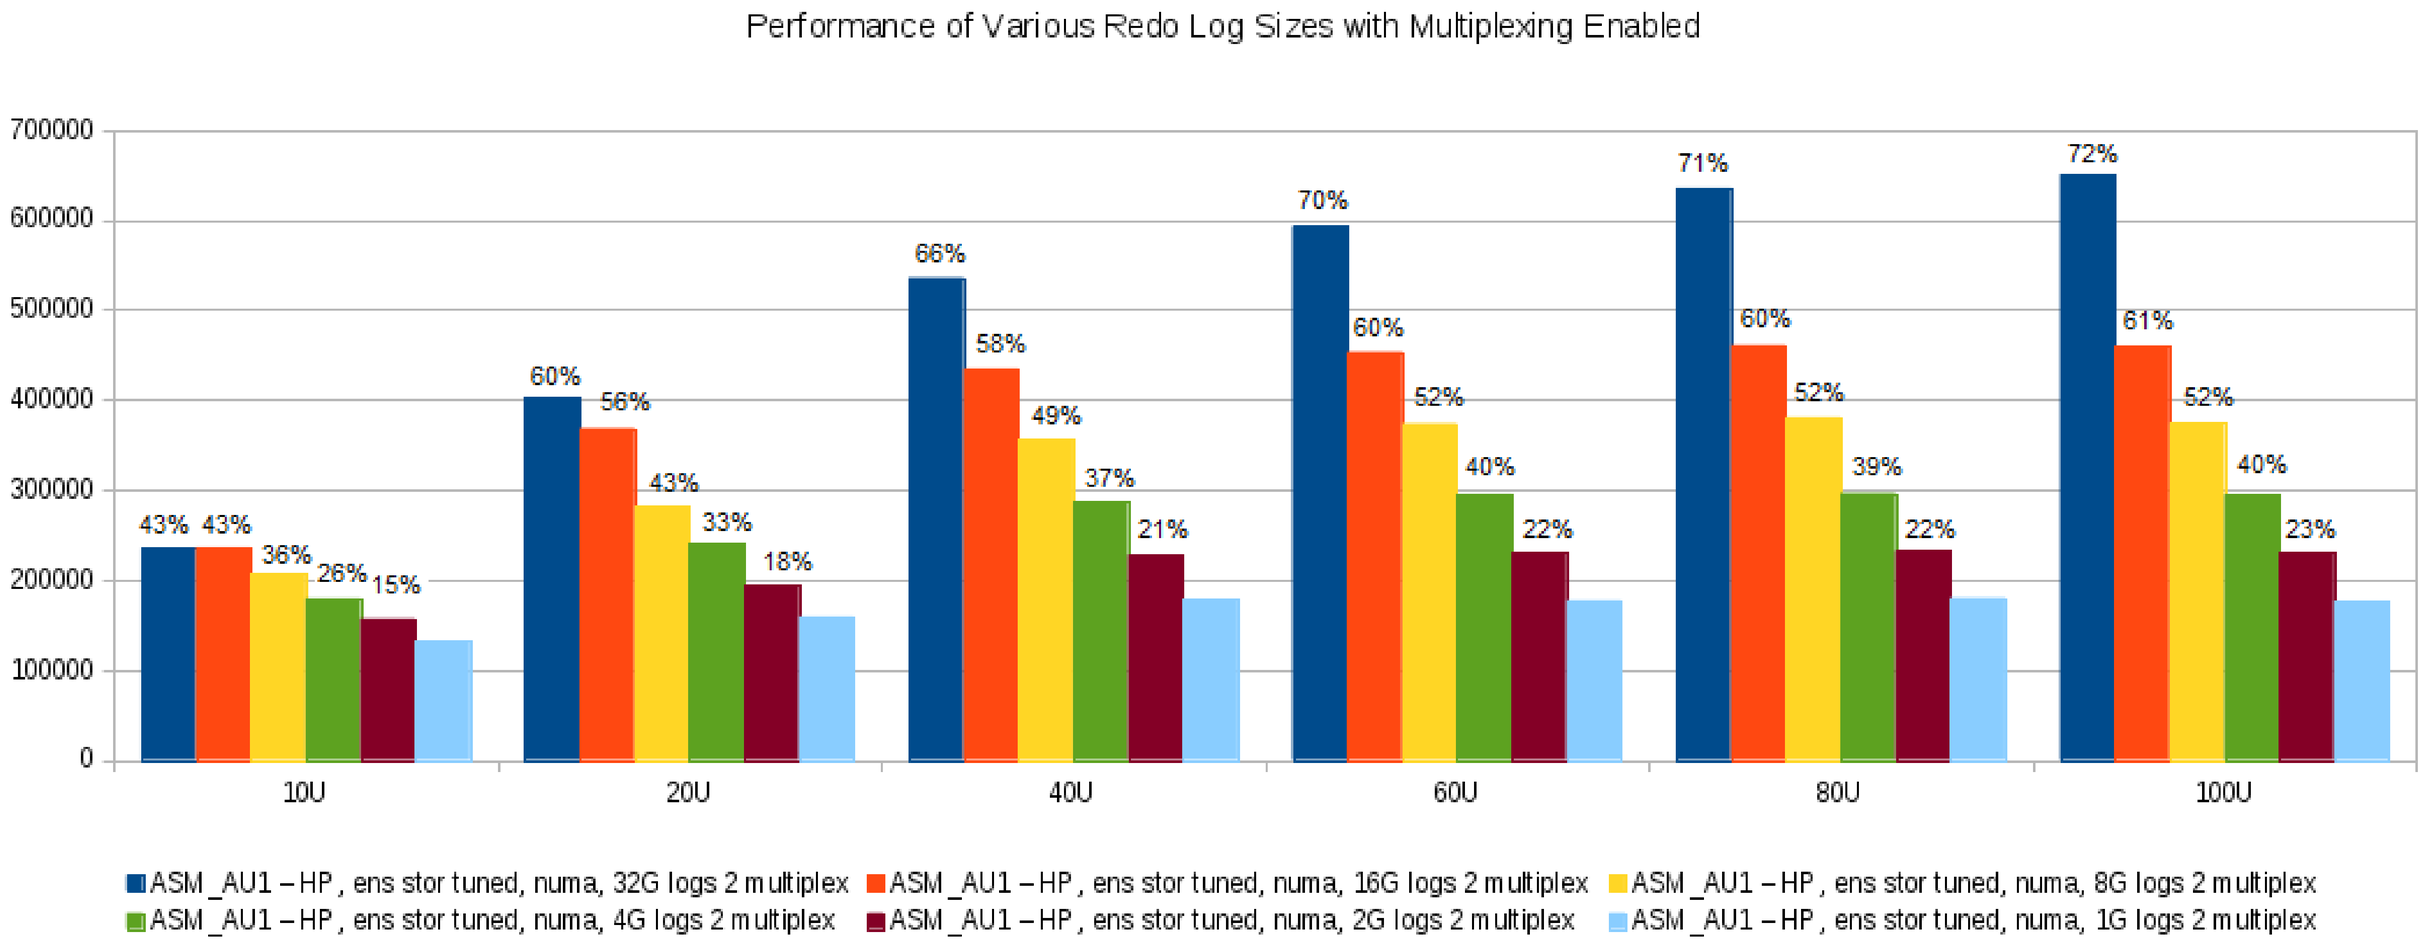

Redo Log size vs Recovery time

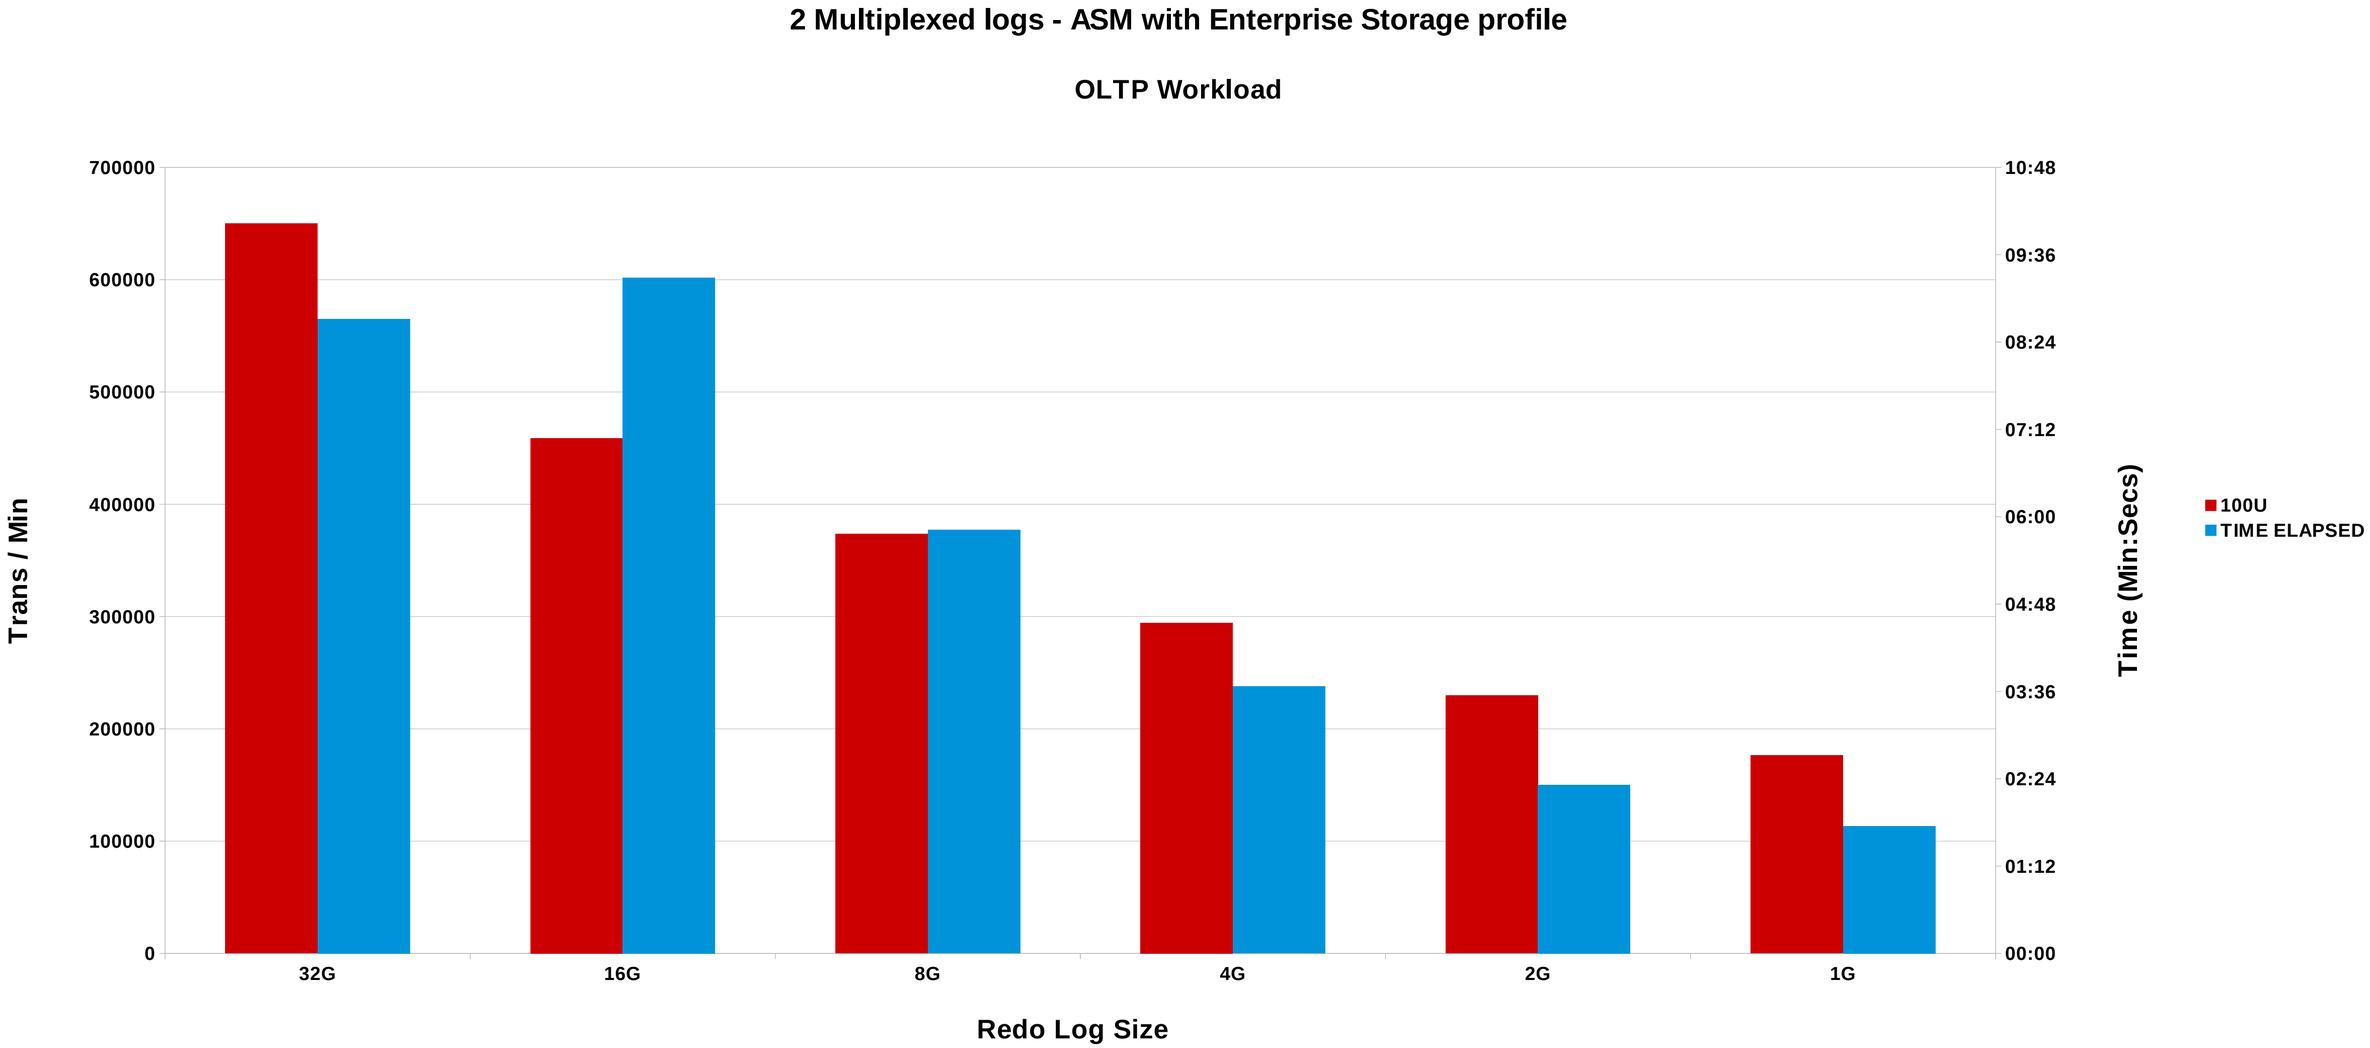

Redo Logs

Checkpoints, Redo Log Size, Multiplexing

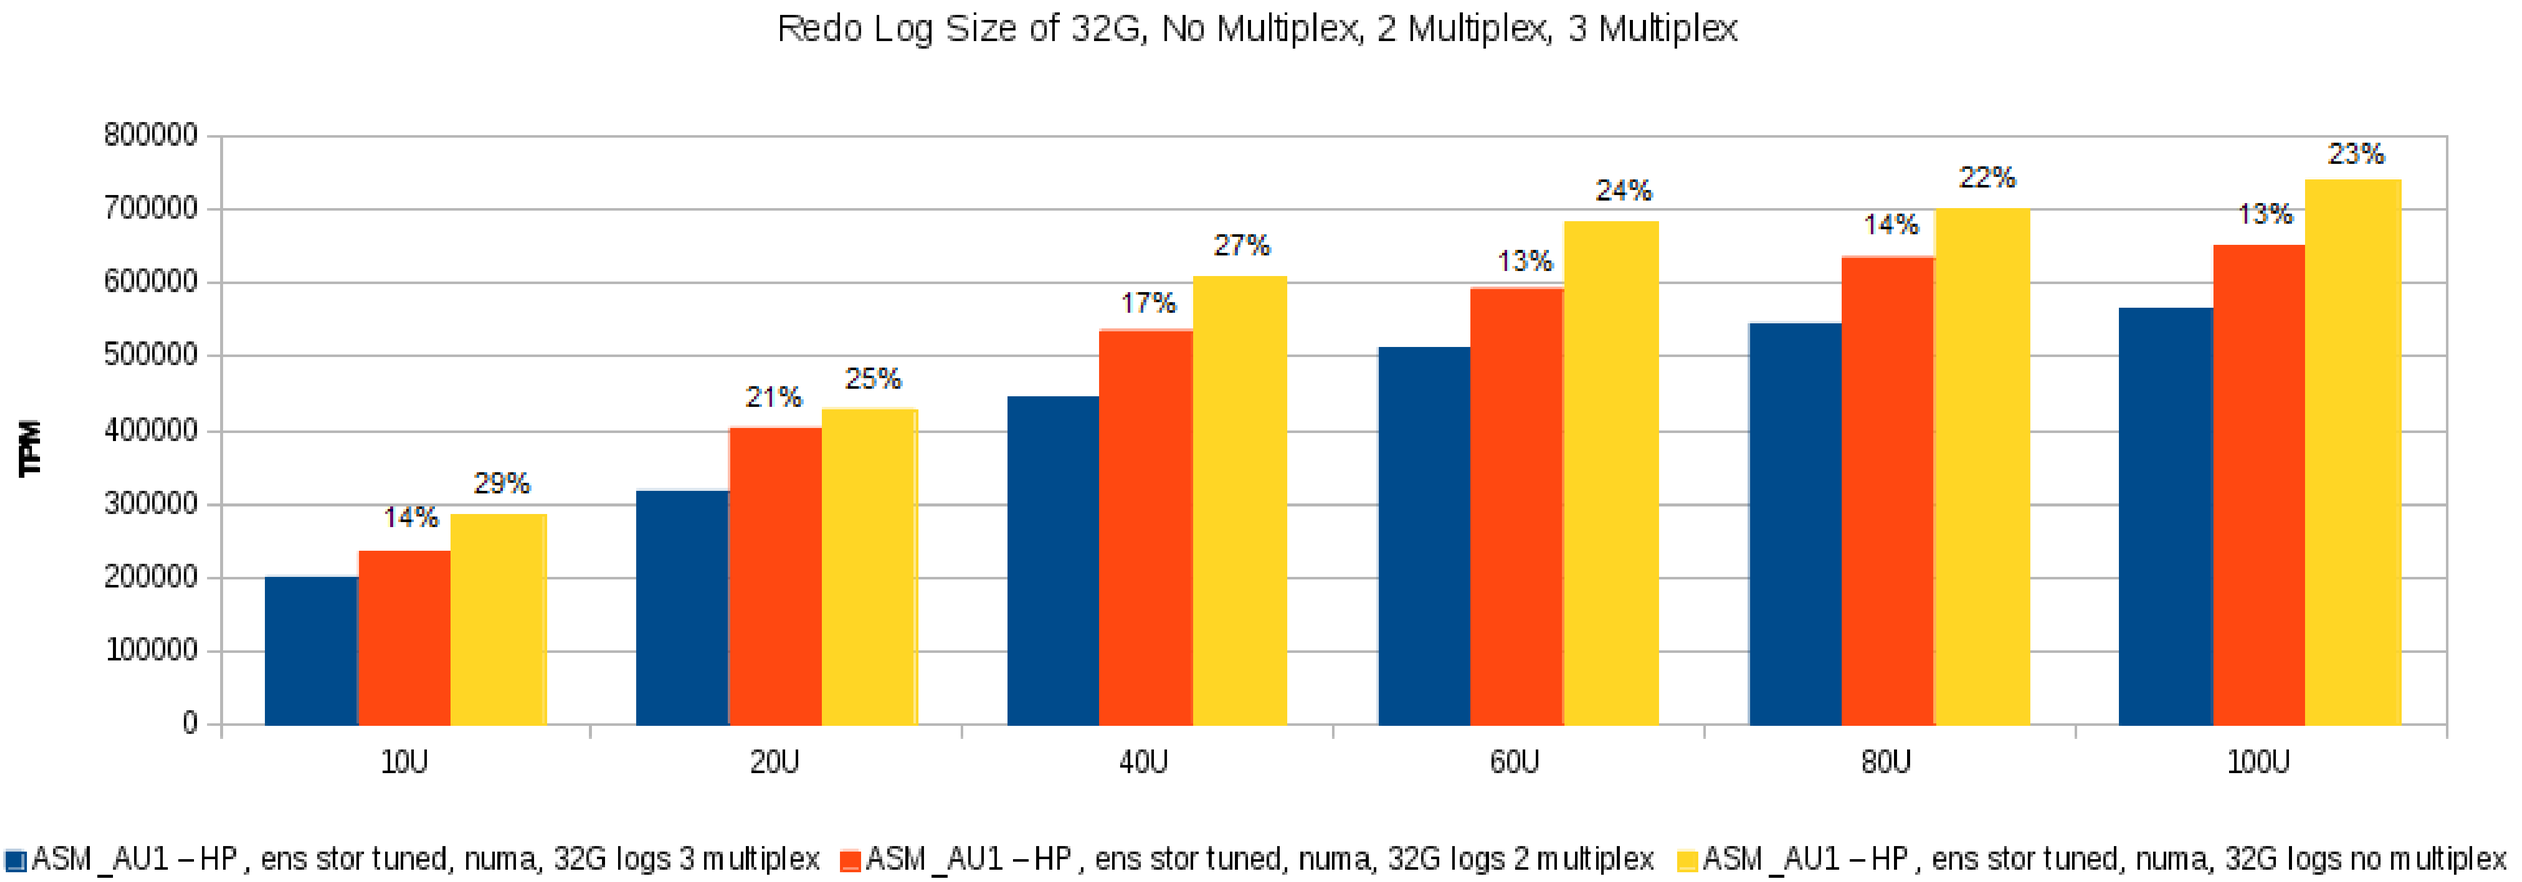

HugePages

- Memory is managed in blocks known as pages. A default page size is 4KB.
- When a system has a lot of memory in it, the use of large pages (2MB in size) keeps the OS page table smaller than if it had to keep every page at the size of 4KB.
- Why do I want to implement HugePages on my Oracle environment?
  - Reduces the OS maintenance of page states
  - Increases Translation Lookaside Buffer (TLB) hit ratios
  - HugePages never get swapped out
- How do I properly configure Huge Pages on my system?
  - Require knowing the total size of all the Oracle SGA instances combined
  - Set the memlock parameter that allows Oracle to use HugePages

HugePages

- Example: Oracle SGA of 19GB, HugePage size in RHEL6: 2MB
- HugePages calculation:
- Convert 19GB to KB: 19922944KB
- Convert 2MB to KB: 2048KB
- Calculate the number of pages required for an SGA of 19GB:
  - 19922944KB / 2048KB = 9728 pages
- For good measure, we will add a few additional pages to ensure we can fit all of the Oracle SGA into HugePages, i.e. 9740
- Within the kernel boot command line within the /etc/grub.conf ensure to include 'hugepages=9740' after the options 'rhgb quiet'

HugePages

- Set the amount of memory Oracle can lock from physical RAM within /etc/security/limits.d/99-grid-oracle-limits.conf file
  - oracle soft memlock <slightly large size than SGA in KB>
  - oracle hard memlock <slightly large size than SGA in KB>
  - 19GB in KB is 19922944, due to this I'd set the value slightly larger, 20000000
- After kernel boot command line & 99-grid-oracle-limits.conf values are set:
  - Reboot the system to allocate the number of huge pages as this is the most reliable method due to memory not yet becoming fragmented.

HugePages

- What type of performance improvement can we expect with the use of HugePages?

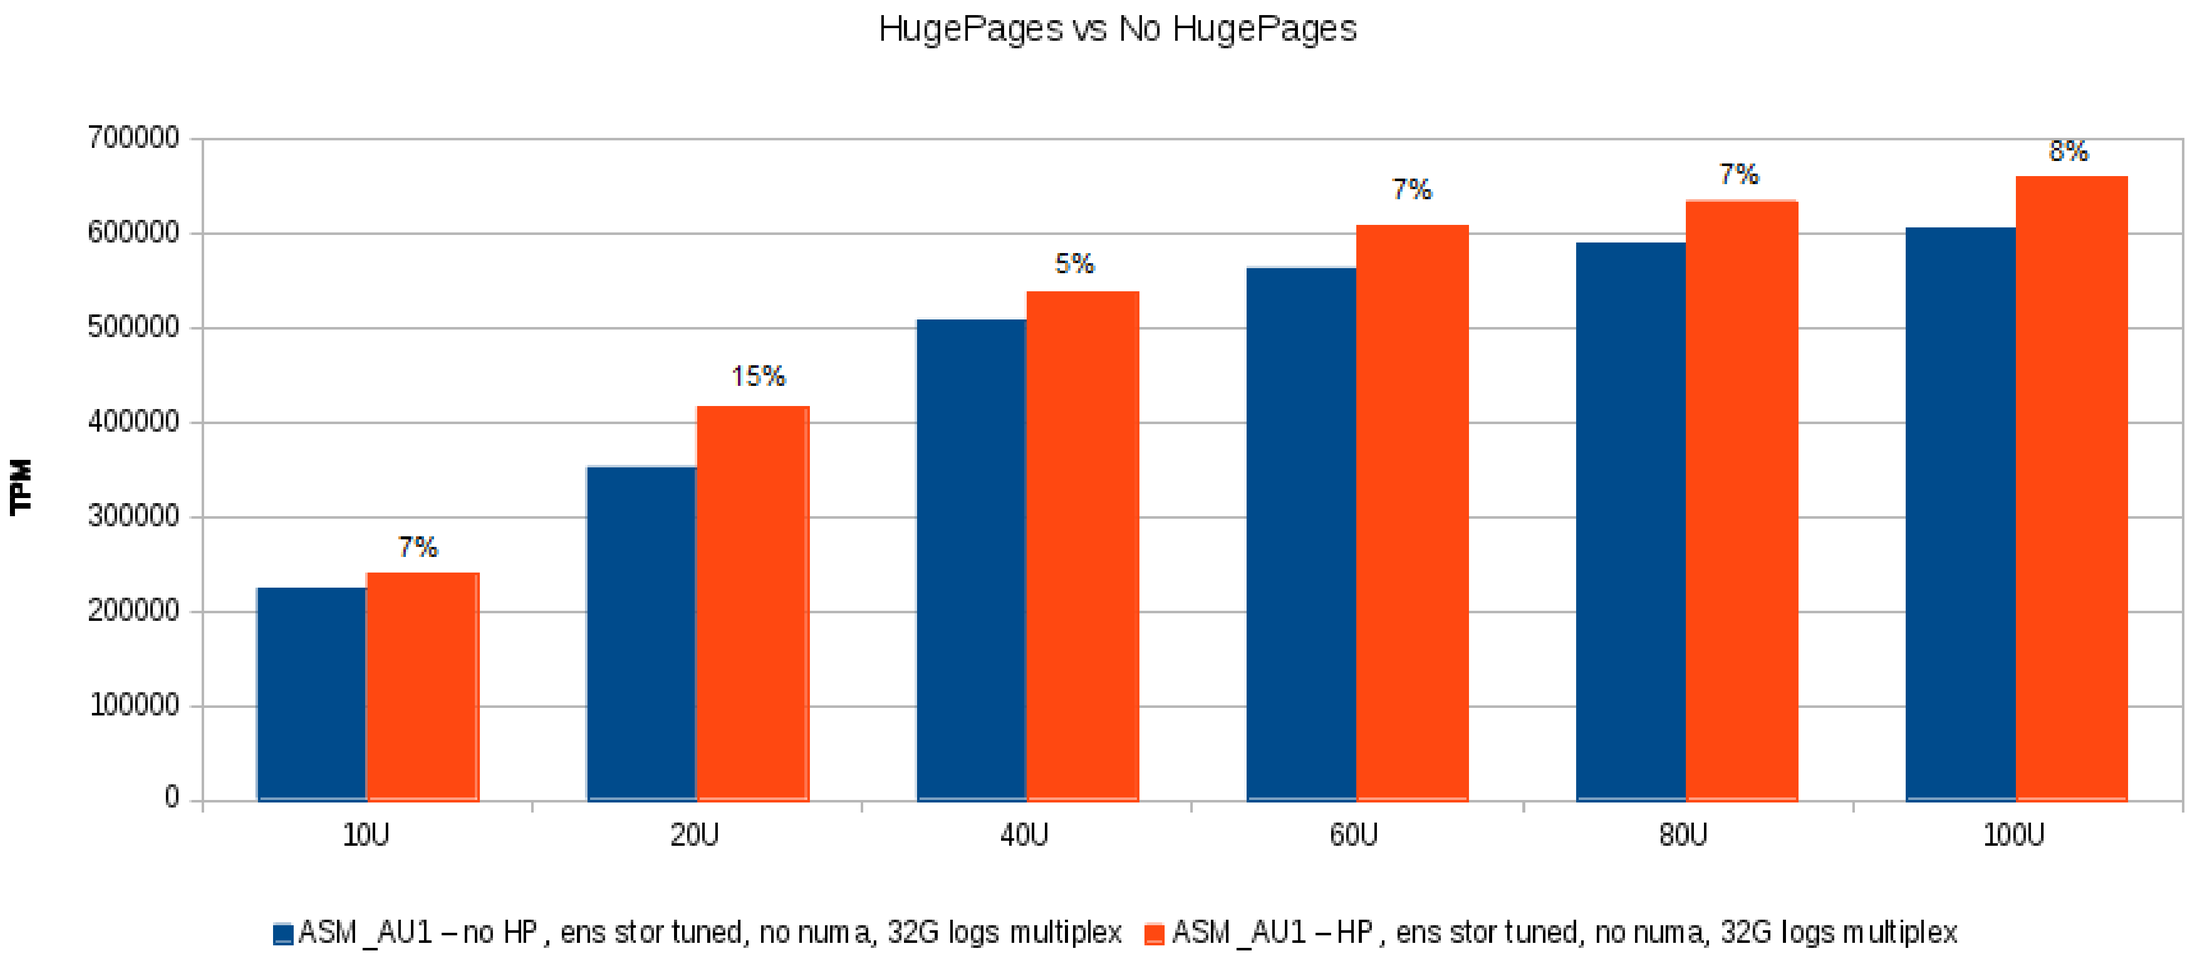

NUMA (Non Uniform Memory Access)

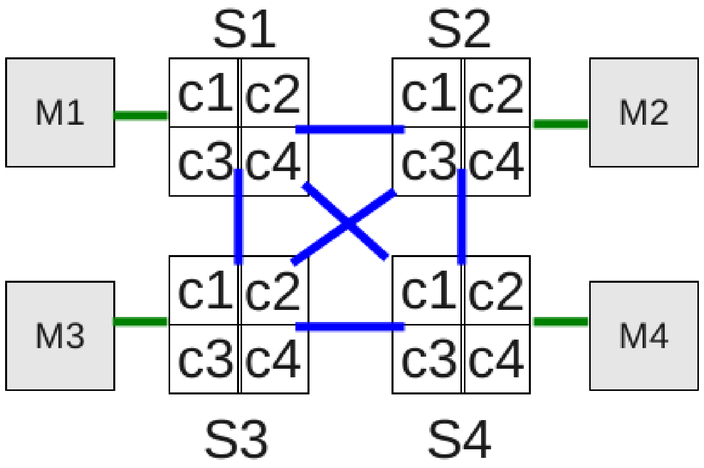

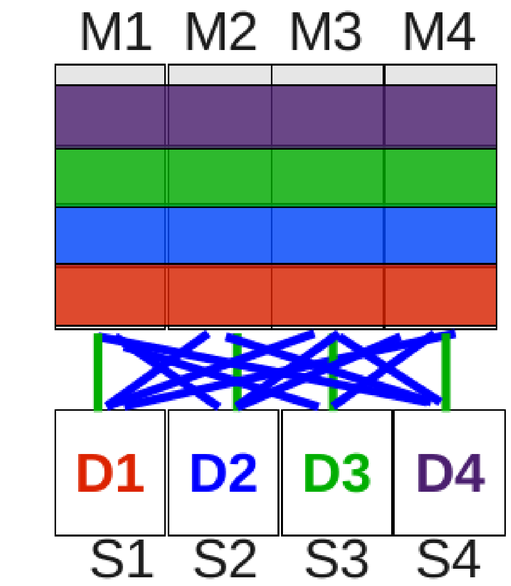

No NUMA optimization

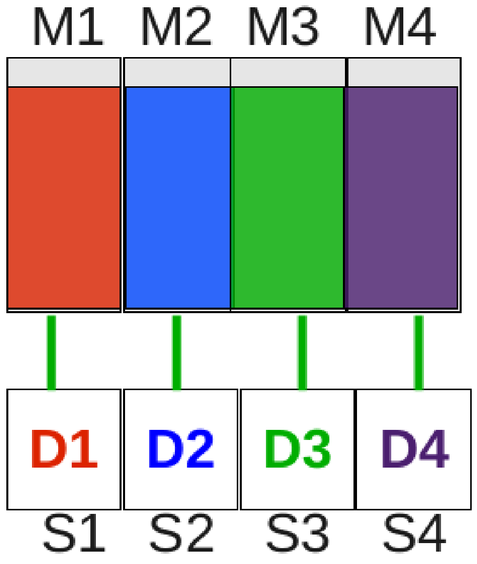

NUMA optimization

- Multi Socket – Multi core architecture
  - NUMA required for scaling
  - RHEL6 completely NUMA aware
  - Performance gains can be achieved when enforcing NUMA

NUMA (Non Uniform Memory Access)

- Easily optimize a system with NUMA via the use of numad
- What is numad?
  - User-level daemon to automatically improve out of the box NUMA system performance
- What does numad do?
  - Monitors available system resources on a per-node basis and assigns significant consumer processes to aligned resources for optimum NUMA performance
  - Rebalances when necessary
  - Provides pre-placement advice for the best initial process placement and resource affinity

NUMA – (Non Uniform Memory Access)

- What type of performance improvement can we expect with the use of NUMA?

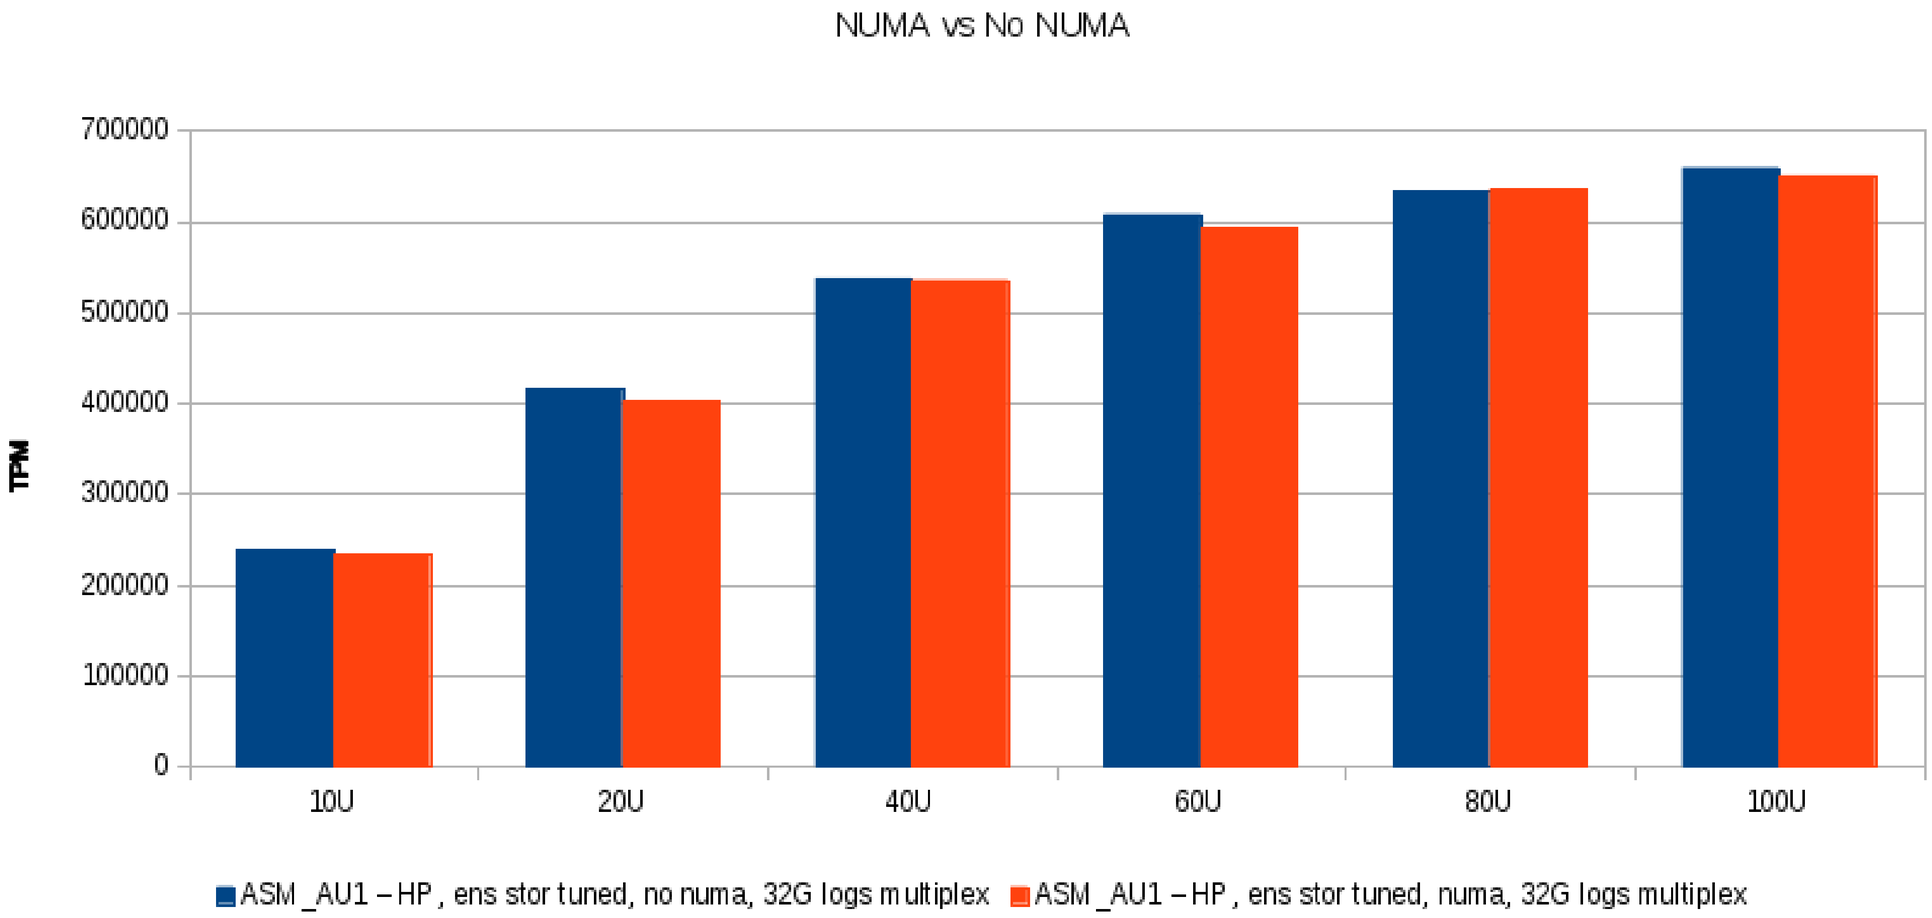

Memory Tuning – Effect of NUMA Tuning

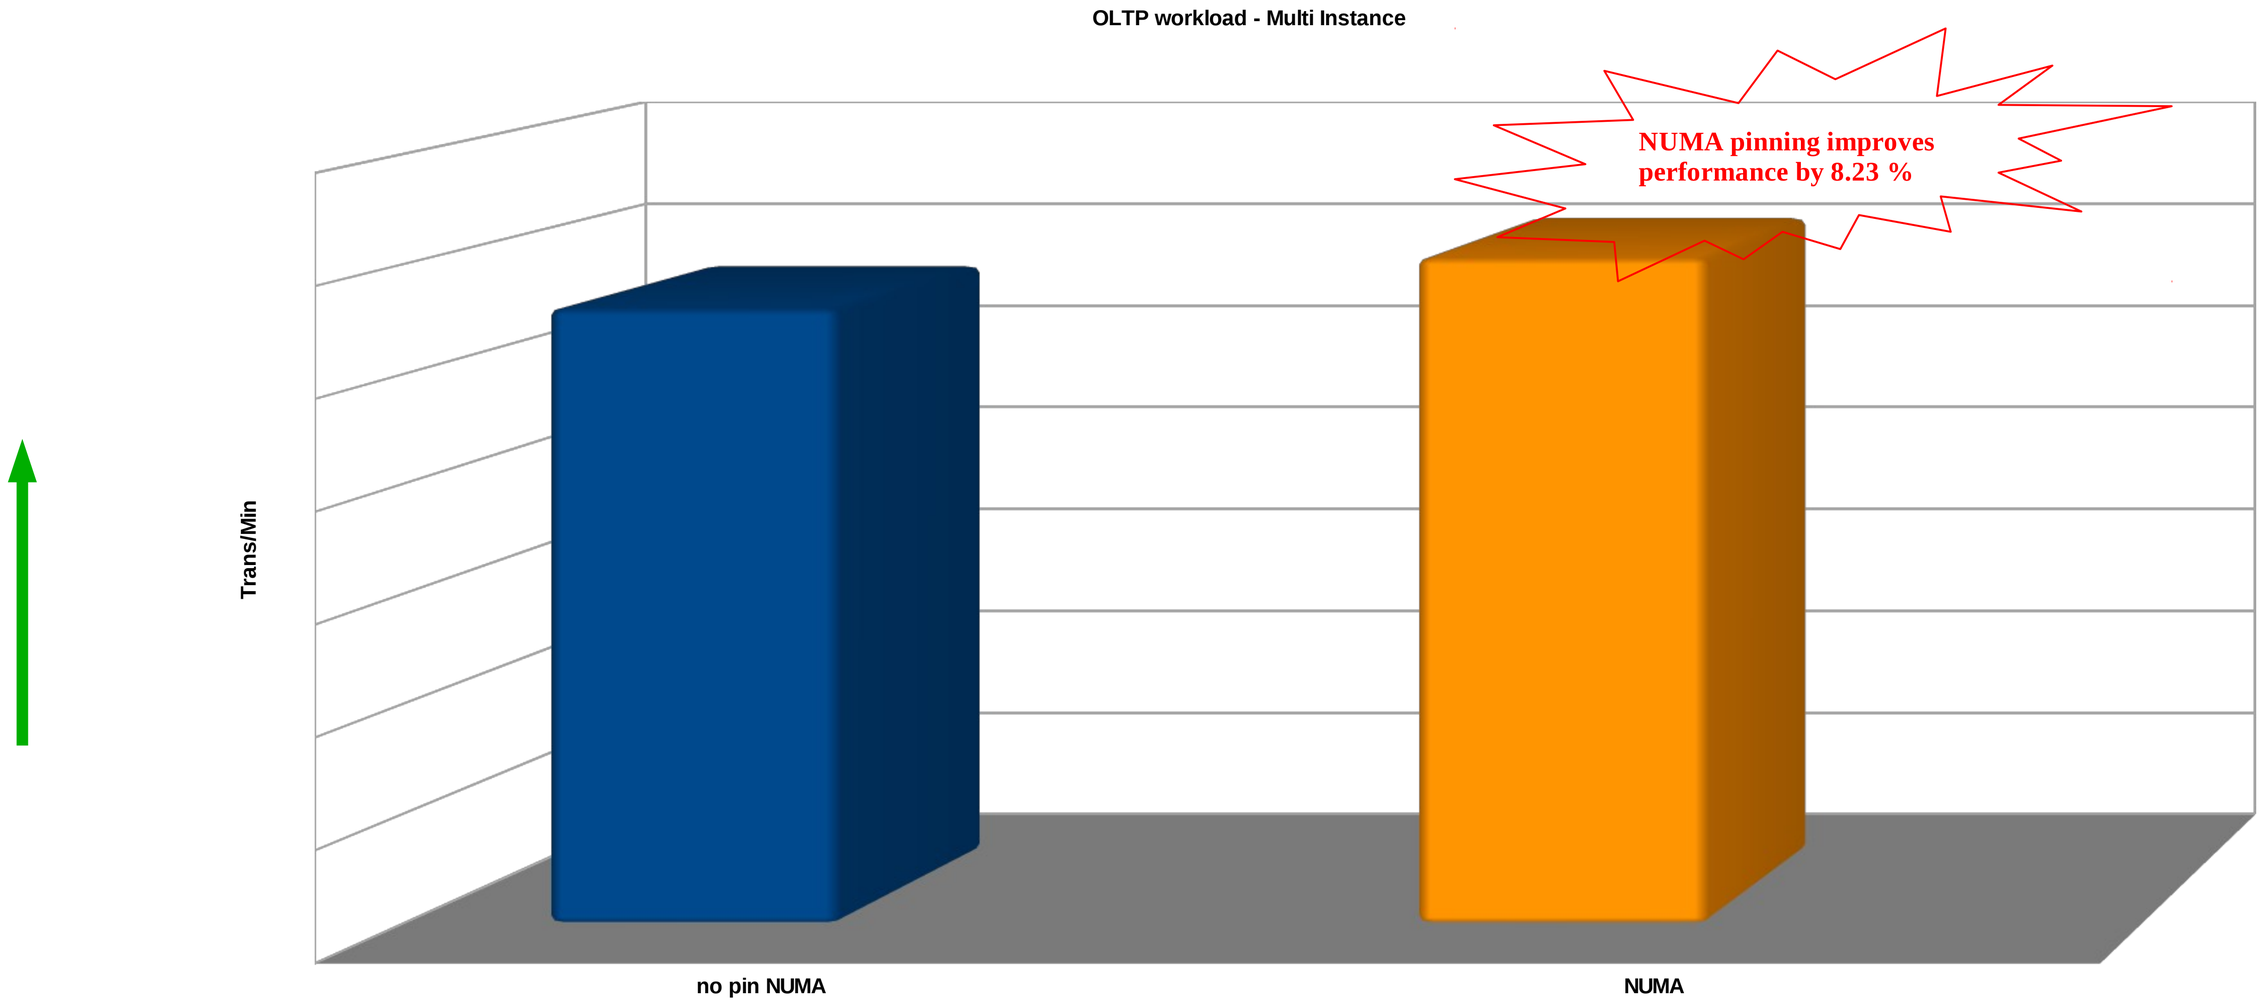

Tuned

- Tuned – dynamic adaptive system tuning daemon that tweaks the system for performance based on different performance profiles
- Benefits
  - Easily installed via yum
  - Over 10 profiles available
  - Can create own custom profiles
  - Can be rolled back

Tuned Profiles	
laptop-battery-powersave	virtual-host
enterprise-storage	default
server-powersave	latency-performance
laptop-ac-powersave	throughput-performance
desktop-powersave	virtual-guest
spindown-disk	

Tuned Profile Summary

Tunable	default	latency-performance	throughput-performance	enterprise-storage	virtual-host	virtual-guest
kernel.sched_min_granularity_ns	4ms		10ms	10ms	10ms	10ms
kernel.sched_wakeup_granularity_ns	4ms		15ms	15ms	15ms	15ms
vm.dirty_ratio	20% RAM		40%	40%	10%	40%
vm.dirty_background_ratio	10% RAM				5%	
vm.swappiness	60				10	30
I/O Scheduler (Elevator)	CFQ	deadline	deadline	deadline	deadline	deadline
Filesystem Barriers	On			Off	Off	Off
CPU Governor	ondemand	performance	performance	performance	performance	performance
Disk Read-ahead				4x	4x	4x

Tuned – OLTP workload

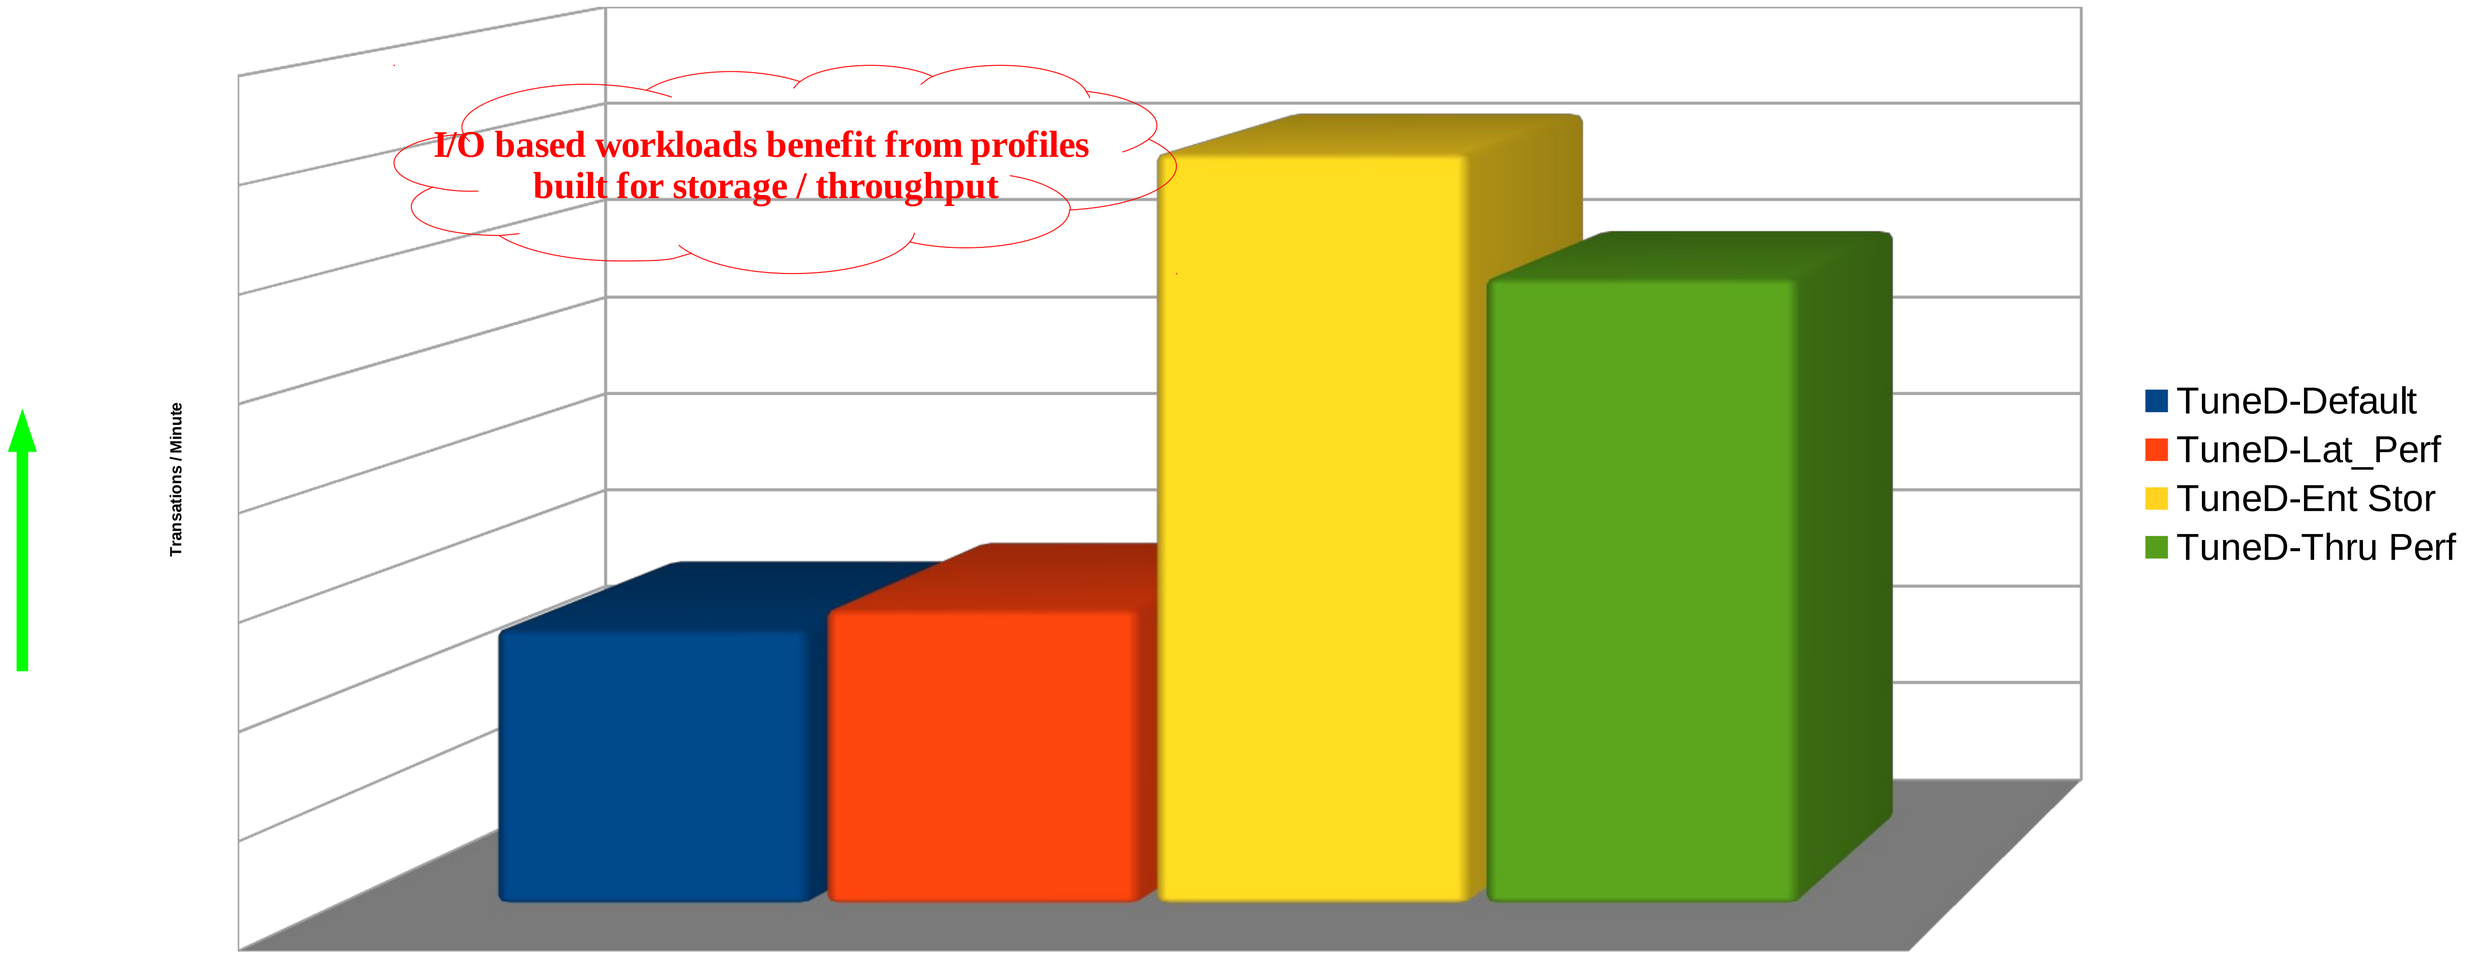

Virtual Memory Settings

- vm.dirty_background_ratio - % of system memory that can be filled with dirty pages.
  - What are dirty pages?
    - Dirty pages are memory pages that still need to be written to disk.
  - What does that mean?
    - If I had a system with 48GB of RAM, and I allowed 3% that means only 1.44GB of data can be in RAM before flushing to disk is required

- vm.dirty_ratio – absolute maximum amount of system memory that can be filled with dirty pages before everything must be committed to disk

Virtual Memory Settings

- vm.dirty_ratio – absolute maximum amount of system memory that can be filled with dirty pages before everything must be committed to disk
  - What does that mean?
    - For example, let us assume we have that same 48GB amount of RAM on a system and I set the vm.dirty_ratio to 15%. At 15%, I can have 7.2GB of data sitting in RAM before the system must write all the dirty pages to disk. The side of effect of this is that any new I/O that comes in is blocked until all the pages have been written to disk.
    - Should I increase this value?
      - Depends, increasing the value allows for more data to sit in RAM but it does increase the possibility of too much data being cached in memory.

Virtual Memory Settings

- vm.dirty_expire_centisecs – how long a dirty page should sit in RAM prior to being written to disk
  - How does it work?
    - When the flush processes looks to see what dirty pages to write to disk, it first checks how old the pages are. If the dirty page is older than the value specified, it is written to disk asynchronously. Oracle recommends a value of 500 which is expressed in hundredths of a second (5 seconds)
  - vm_dirty_writeback_centisecs – how often the flush process wakes up to see if it needs to flush any dirty pages

Performance Monitoring Tools

- Monitoring tools
  - Top, vmstat, ps, iostat, netstat, sar, perf
- Kernel tools
  - /proc, sysctl, AltSysRq
- Profiling
  - oprofile,strace,ltrace,systemtap,perf

Performance Monitoring Tool – sar

Output of "sar -N DEV 3"

For a DSS workload running on iSCSI storage using different MTUs

1500 MTU

```
01:40:08 PM     IFACE   rxpck/s   txpck/s    rxkB/s    txkB/s   rxcmp/s   txcmp/s  rxmcst/s
01:40:11 PM      eth0      0.34      0.34      0.02      0.02      0.00      0.00      0.00
01:40:11 PM      eth5 135016.78  19107.72 199178.19   1338.53      0.00      0.00      0.34

01:40:14 PM      eth0      0.66      0.00      0.05      0.00      0.00      0.00      0.66
01:40:14 PM      eth5 133676.74  18911.30 197199.84   1310.25      0.00      0.00      0.66

01:40:17 PM      eth0      0.67      0.00      0.05      0.00      0.00      0.00      0.67
01:40:17 PM      eth5 134555.85  19045.15 198502.27   1334.19      0.00      0.00      0.33

01:40:20 PM      eth0      1.00      0.00      0.07      0.00      0.00      0.00      0.67
01:40:20 PM      eth5 134116.33  18972.33 197849.55   1325.03      0.00      0.00      1.00
```

9000 MTU

```
06:58:43 PM     IFACE   rxpck/s   txpck/s    rxkB/s    txkB/s   rxcmp/s   txcmp/s  rxmcst/s
06:58:46 PM      eth0      0.91      0.00      0.07      0.00      0.00      0.00      0.00
06:58:46 PM      eth5 104816.36  48617.27 900444.38   3431.15      0.00      0.00      0.91

06:58:49 PM      eth0      0.00      0.00      0.00      0.00      0.00      0.00      0.00
06:58:49 PM      eth5 118269.80  54965.84 1016151.64  3867.91      0.00      0.00      0.50

06:58:52 PM      eth0      0.00      0.00      0.00      0.00      0.00      0.00      0.00
06:58:52 PM      eth5 118470.73  54382.44 1017676.21  3818.35      0.00      0.00      0.98

06:58:55 PM      eth0      0.94      0.00      0.06      0.00      0.00      0.00      0.00
06:58:55 PM      eth5 115853.05  53515.49 995087.67   3766.28      0.00      0.00      0.47
```

Performance Monitoring Tool – iostat

Device:	rrqm/s	wrqm/s	r/s	w/s	rMB/s	wMB/s	avgrq-sz	avgqu-sz	await	svctm	%util
dm-5	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	5.50	5.50	0.18
dm-6	0.00	374.13	0.00	357.20	0.00	80.23	459.98	0.64	1.80	1.78	63.50
dm-7	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	5.60	5.60	0.19
dm-8	0.00	373.23	0.00	358.10	0.00	80.23	458.82	0.71	1.98	1.95	69.87
dm-12	0.00	0.17	318.73	0.80	10.24	0.01	65.71	5.16	16.14	2.96	94.60
dm-13	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	5.40	5.40	0.18
dm-15	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	5.50	5.50	0.18

Device:	rrqm/s	wrqm/s	r/s	w/s	rMB/s	wMB/s	avgrq-sz	avgqu-sz	await	svctm	%util
dm-5	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	0.90	0.80	0.03
dm-6	0.00	390.03	0.00	394.37	0.00	76.27	396.07	0.67	1.69	1.67	65.90
dm-7	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	0.70	0.70	0.02
dm-8	0.00	389.67	0.00	394.73	0.00	76.27	395.70	0.73	1.85	1.83	72.30
dm-12	0.00	0.40	266.17	604.80	8.45	34.07	99.98	7.07	8.11	1.10	96.04
dm-13	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	1.10	1.10	0.04
dm-15	0.00	0.00	0.00	0.33	0.00	0.00	8.00	0.00	0.90	0.90	0.03

Device:	rrqm/s	wrqm/s	r/s	w/s	rMB/s	wMB/s	avgrq-sz	avgqu-sz	await	svctm	%util
dm-5	0.00	0.00	0.00	0.30	0.00	0.00	8.00	0.00	2.11	2.11	0.06
dm-6	0.00	611.60	0.00	1036.10	0.00	36.00	71.16	0.61	0.59	0.59	61.19
dm-7	0.00	0.00	0.00	0.30	0.00	0.00	8.00	0.00	2.00	2.00	0.06
dm-8	0.00	610.33	0.00	1037.37	0.00	36.00	71.07	0.69	0.66	0.66	68.43
dm-12	0.00	0.40	112.50	3239.17	3.53	134.76	84.50	24.62	7.28	0.30	100.00
dm-13	0.00	0.00	0.00	0.30	0.00	0.00	8.00	0.00	1.11	1.11	0.03
dm-15	0.00	0.00	0.00	0.30	0.00	0.00	8.00	0.00	1.22	1.22	0.04

Performance Monitoring Tool – vmstat

procs		----------memory-------				-swap-		----io----		--system--		-----cpu-----				
r	b	swpd	free	buff	cache	si	so	bi	bo	in	cs	us	sy	id	wa	st
5	21	0	7021756	126152	3129252	0	0	3649	211810	19946	75994	26	5	29	40	0
22	19	0	7014420	126320	3128412	0	0	5627	194181	22729	96493	49	8	25	17	0
22	11	0	7017680	126516	3128724	0	0	6443	205404	24029	106927	64	10	18	8	0
1	16	0	7010932	128496	3129072	0	0	5169	223335	19931	84992	52	8	32	8	0
33	11	0	7009068	128672	3129740	0	0	5775	202376	23412	105580	66	10	18	6	0
33	11	0	7002660	128920	3130536	0	0	5094	209813	22997	102622	64	10	20	6	0
40	13	0	7000120	129088	3131312	0	0	4727	205395	23208	103987	66	10	19	5	0
20	8	0	6999472	129256	3132268	0	0	4630	204699	23395	105048	67	10	18	4	0
2	15	0	6992544	130492	3133136	0	0	3543	232340	21024	94246	57	9	24	11	0
11	7	0	6983392	130952	3134084	0	0	3830	211579	22739	101220	65	10	20	5	0
60	4	0	6978504	131264	3134208	0	0	2666	241644	19730	83332	44	7	37	13	0
14	6	0	6977012	131436	3134056	0	0	2952	213545	21980	91155	50	8	32	11	0
43	8	0	6976668	131608	3134700	0	0	3615	210458	23811	102238	64	10	22	5	0
26	4	0	6974312	131768	3135632	0	0	3278	203879	24058	104138	68	10	18	3	0
45	3	0	6975760	131936	3136440	0	0	3237	211308	24026	104433	68	10	18	3	0
50	3	0	6976124	132124	3137380	0	0	2912	206853	23865	103986	68	10	19	3	0
43	2	0	6973608	132304	3138184	0	0	3090	208478	23962	103676	69	10	18	3	0

Summary

- I/O
  - Enable tuned and use one of its profiles for Oracle databases, i.e. Enterprise-storage or throughput-performance
- Memory
  - Enable NUMA with the help of numad
  - Enable HugePages
- Oracle Redo Logs
  - Consider the price of downtime vs performance gain

RED HAT SUMMIT

10 YEARS *and counting*

SAN FRANCISCO | APRIL 14-17, 2014

Q&A

Follow me on Twitter @Lopez
For Reference Architectures visit:
http://red.ht/1f7THRa

redhat.