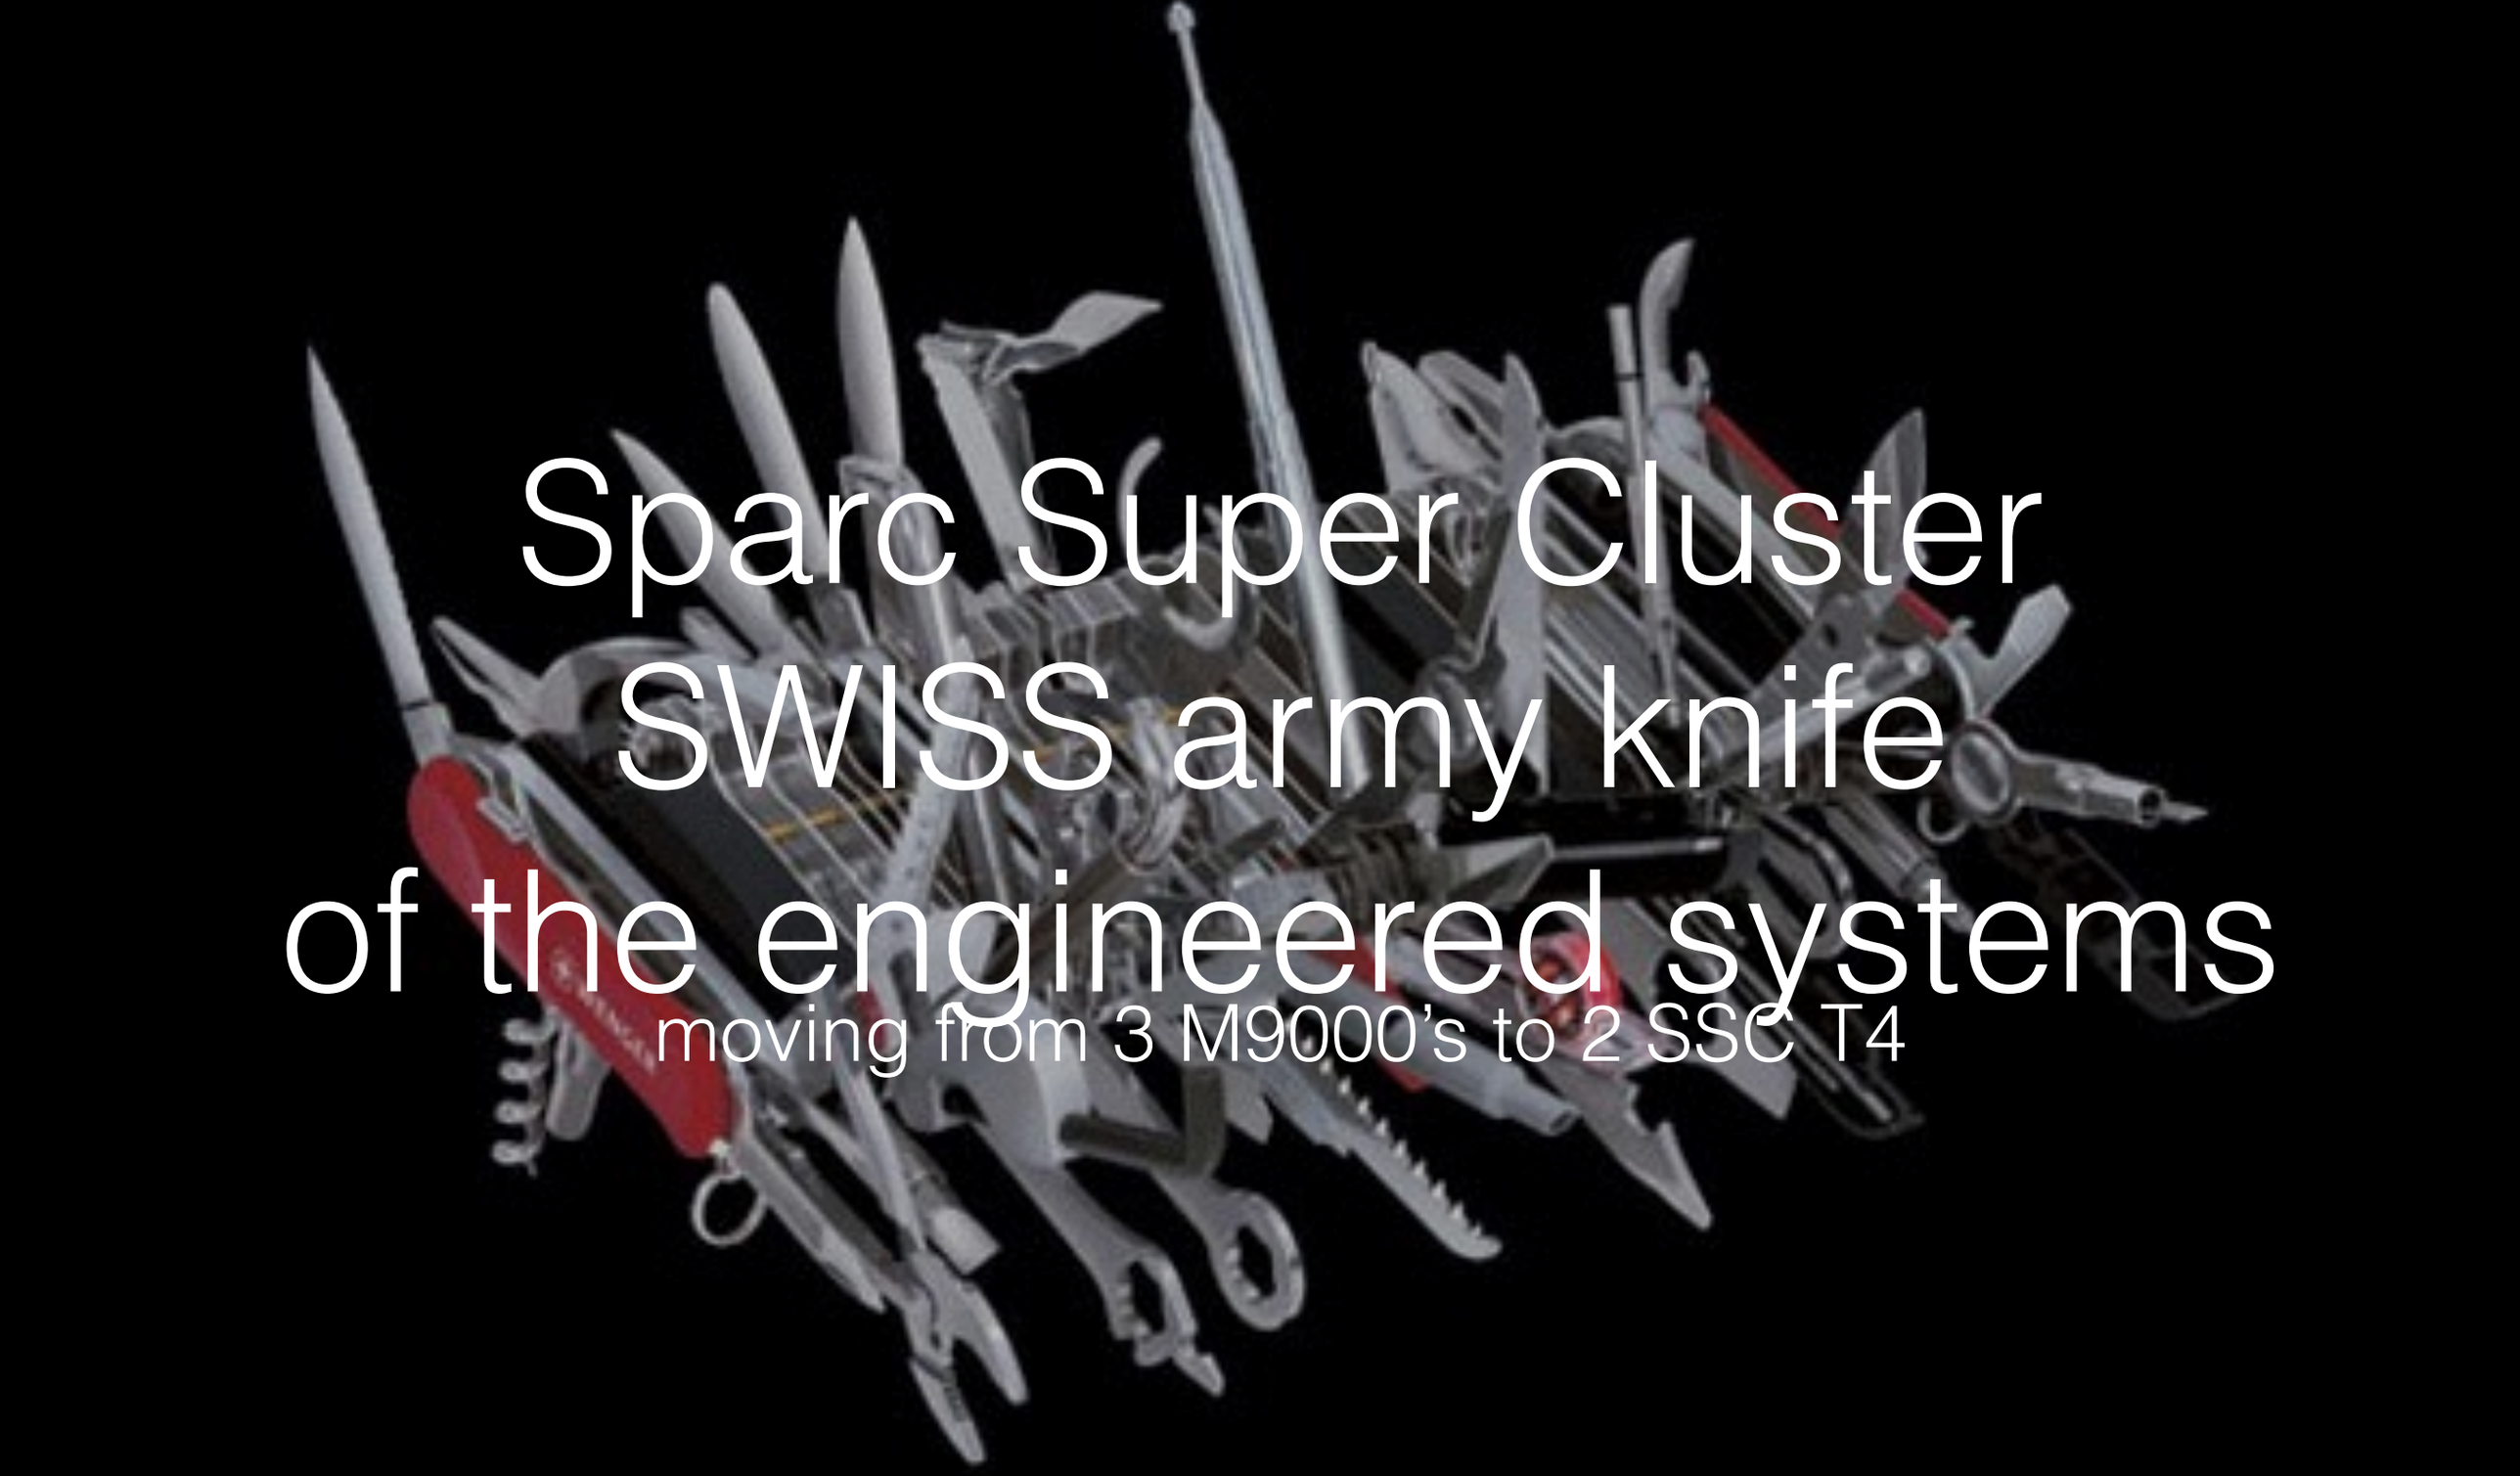

Sparc Super Cluster SWISS army knife of the engineered systems

moving from 3 M9000's to 2 SSC T4

20. hroug
godišnja konferencija

whoami

Philippe Fierens
Independent
Oracle DBA (not solaris admin)
Architect
and
Oracle ACE Associate
and most of all
DAD & husband ;-)
AND
...

BBL


THE BELGIAN BEER GUIDE
WHAT ARE YOU EATING?
LIGHT & LEAN
CHEESE
BARNYARDY
UNSWEETENED
APERITIF
CHICKEN/FISH
SAISON
BELGIAN PALE
LAMBIC
MEAT SNACKS
I'M WATCHING MY FIGURE
(LEAN RED MEAT)
IT'S COOL, I JUST HIT THE LAPBAND
(FATTY RED MEAT)
BLONDE
TRIPPLE
BELGIAN STRONG
SWEETS
SWEETENED
FEELIN' TART
FEELIN' DARK
QUAD
HOW MUCH ARE YOU DRINKING?
1 OR 2
NEW SCHOOL WILD CARD
BELGIAN IPA
TRADITIONAL
PALE
TRIPPLE
DARK
QUAD
MARATHON
FUNKY
FARM HOUSE
SMOOTH
LIGHT
BELGIAN PALE
BELGIAN PILSNER
DARK
BELGIAN AMBER
BELGIAN DARK

What is this presentation about ?

not about beer

SORRY

but there is time after this presentation ;-)

Tonight

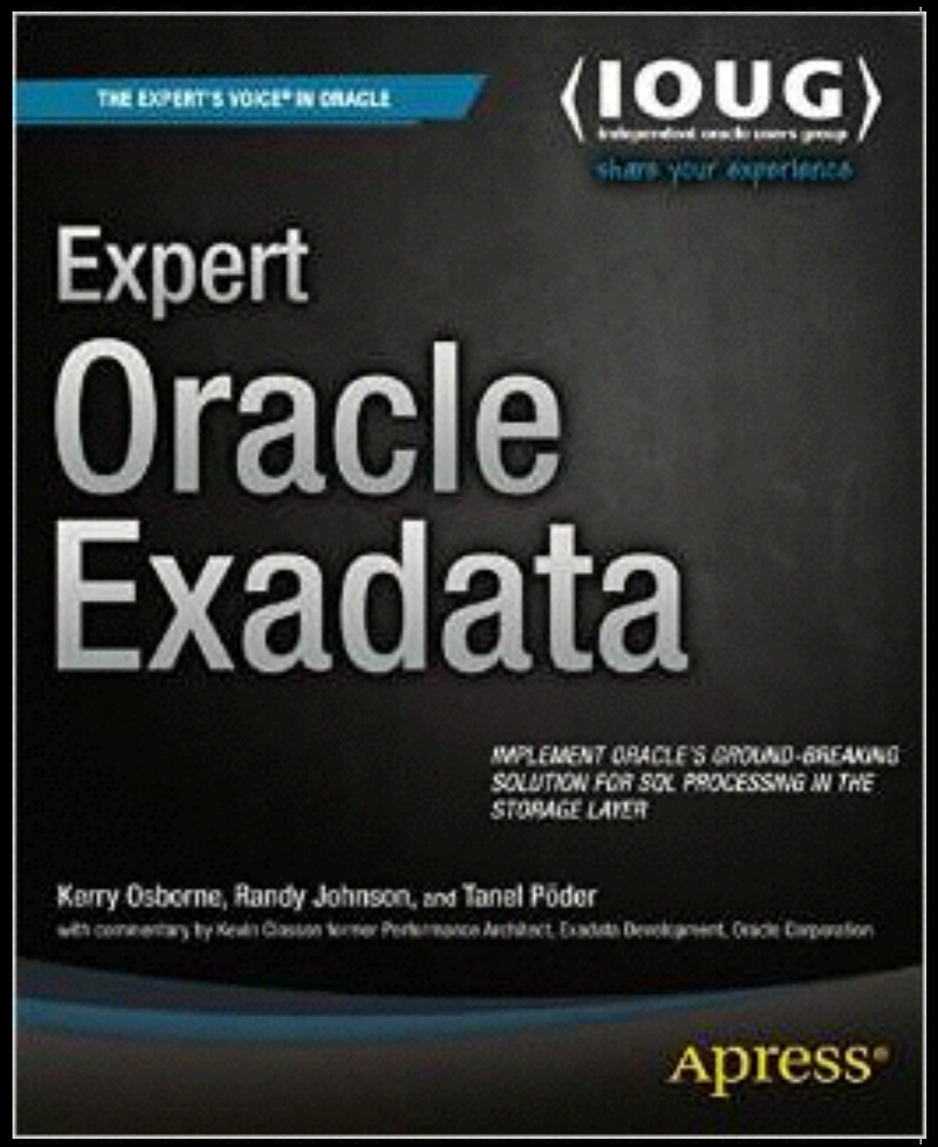
THE EXPERT'S VOICE® IN ORACLE
IOUG
share your experience
Expert
Oracle
Exadata
IMPLEMENT ORACLE'S GROUND-BREAKING SOLUTION FOR SQL PROCESSING IN THE STORAGE LAYER
Kerry Osborne, Randy Johnson, and Tanel Põder
with commentary by Kevin Closson former Performance Architect, Exadata Development, Oracle Corporation
Apress®

Agenda

Sparc **S**uper **C**luster what is it ?

Project

Overview : phases

Technologies used

Issues encountered

Sparc Super Cluster

SPARC based flexible EXADATA

ZFSSA integrated

Several versions exist

SSC Versions

SSC	T4-4	T5-8	M6-32
processor	T4	T5	M6
cores/proc	8	16	12
nr nodes	4	2	M9000 style
RAM/RACK	4TB	4TB	**32TB**
cores/RACK	128	256	384
storage cells	6	8	9


Sun
ORACLE
S

T4-4


SUPERCLUSTER
Sun
ORACLE
EXTREME PERFORMANCE
S

T5-8


EXADATA
Sun
ORACLE
STORAGE EXPANSION
X
M6 32 TERABYTES
BIG MEMORY MACHINE
Sun
ORACLE

M6-32

SSC T4-4

4 times T4-4 => 32 cores * 4

6 X3-2 storage cells

3 x IB Switch (1 spine 2 leafs)

1 TB RAM per compute node

ZFSSA 7320 with 20 x 3 TB raw capacity

Management Switch 1Gbe

Solaris 11

Customer config

3 High Capacity X4-2 CELLS
~~5~~ 3 High Performance X3-2 CELLS

15k vs 7,2 k disks
600 GB vs 4TB disks
4x400Gb F40 vs 3,2 TB Flash

So SSC = normal Solaris?

Not really: own patches

Own patch cycle

Lag behind a bit on normal solaris

Quarterly **F**ull **S**tack **D**ownload **P**atch

Patches firmware, PDU, Storage Cells, OS, …

JIDR …

Flexibility ++

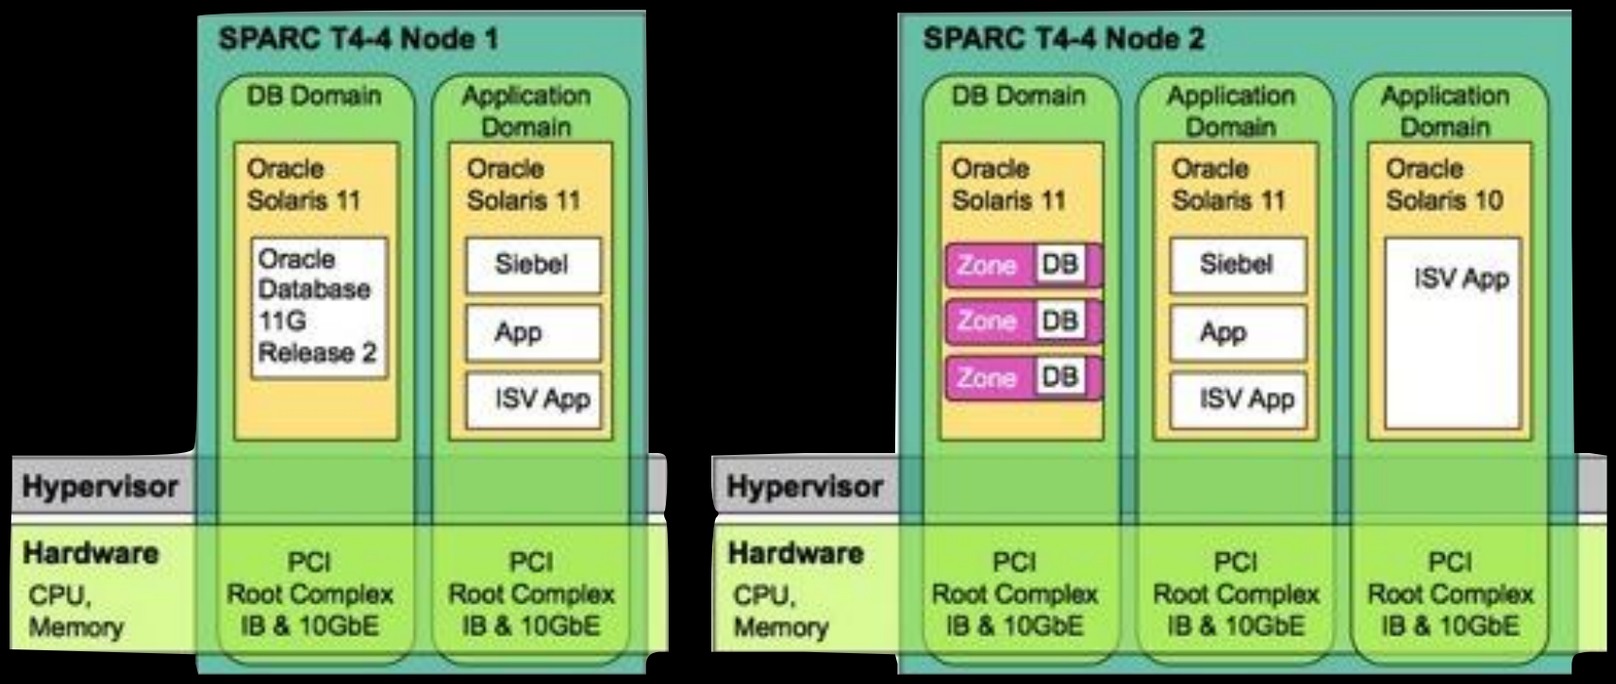

2 Types :

DB Domains zones for more env

App Domains

Integrated IB connected Storage

LDOMS / Zones

IB

2 partitions

- interconnect and cell communication FFFF
- ZFSSA 8503

shared in a similar way as with normal network

IB

On the global zone

```
dladm show-phys

LINK              MEDIA                STATE      SPEED  DUPLEX    DEVICE
net24             Infiniband           up         32000  unknown   ibp1
net0              Ethernet             up         1000   full      igb0
net5              Ethernet             unknown    0      unknown   ixgbe1
net7              Ethernet             up         10000  full      ixgbe3
net23             Infiniband           up         32000  unknown   ibp0
net1              Ethernet             up         1000   full      igb1
net21             Infiniband           up         32000  unknown   ibp2
net4              Ethernet             up         10000  full      ixgbe0
net22             Infiniband           up         32000  unknown   ibp3
net6              Ethernet             unknown    0      unknown   ixgbe2
net38             Ethernet             unknown    0      unknown   e1000g0
net39             Ethernet             unknown    0      unknown   e1000g1
net12             Ethernet             up         10     full      usbecm2
```

IB partitions

```
dladm show-part
LINK         PKEY  OVER         STATE    FLAGS
ipmpapp0_0   8503  net22        up       f---
ipmpapp0_1   8503  net23        up       f---
bondib0_0    FFFF  net21        up       f---
bondib0_1    FFFF  net22        up       f---
bondib1_0    FFFF  net23        up       f---
bondib1_1    FFFF  net24        up       f---  --> in global
n21_1019_bondib0 FFFF net21     up       ----  --> in local zone
n22_1019_bondib0 FFFF net22     up       ----
n23_1020_bondib1 FFFF net23     up       ----
n24_1020_bondib1 FFFF net24     up       ----
ipmpapp0_z01_0 8503 net22       up       ----
ipmpapp0_z01_1 8503 net23       up       ----
```

IB

```
nodeadm03 [ ~ ] $ dladm
LINK                CLASS     MTU    STATE    OVER
net24               phys      65520  up       --
net23               phys      65520  up       --
net21               phys      65520  up       --
net4                phys      1500   up       --
net22               phys      65520  up       --
net12               phys      1500   up       --
net24               phys      65520  up       --
net0                phys      1500   up       --
…
ipmpapp0_0          part      65520  up       net22
ipmpapp0_1          part      65520  up       net23
bondib0_0           part      65520  up       net21
bondib0_1           part      65520  up       net22
bondib1_0           part      65520  up       net23
bondib1_1           part      65520  up       net24
n21_1019_bondib0    part      65520  up       net21
nodedat03z01/n21_1019_bondib0 part 65520 up net21
n22_1019_bondib0    part      65520  up       net22
nodedat03z01/n22_1019_bondib0 part 65520 up net22

n23_1020_bondib1    part      65520  up       net23
nodedat03z01/n23_1020_bondib1 part 65520 up net23
n24_1020_bondib1    part      65520  up       net24
nodedat03z01/n24_1020_bondib1 part 65520 up net24
ipmpapp0_z01_0      part      65520  up       net22
nodedat03z01/ipmpapp0_z01_0 part 65520 up  net22
ipmpapp0_z01_1      part      65520  up       net23
nodedat03z01/ipmpapp0_z01_1 part 65520 up  net23
nodedat03z01/net0 vnic  1500   up       net0
```

net24 phys 65520 up --

bondib1_1 part 65520 up net24

n24_1020_bondib1 part 65520 up net24

nodedat03z01/n24_1020_bondib1 part 65520 up net24

IB

```
ipadm
NAME              CLASS/TYPE STATE        UNDER      ADDR
bondeth0          ipmp       ok           --         --
   bondeth0/v4    static     ok           --         10.191.99.71/24
bondib0           ipmp       ok           --         --
   bondib0/v4     static     ok           --         192.168.10.5/22
bondib1           ipmp       ok           --         --
   bondib1/v4     static     ok           --         192.168.10.6/22
bondib0_0         ip         ok           bondib0    --
bondib0_1         ip         ok           bondib0    --
bondib1_0         ip         ok           bondib1    --
bondib1_1         ip         ok           bondib1    --
bondmgt0          ipmp       ok           --         --
   bondmgt0/v4    static     ok           --         10.191.174.11/24
ipmpapp0          ipmp       ok           --         --
   ipmpapp0/v4    static     ok           --         192.168.30.4/22
ipmpapp0_0        ip         ok           ipmpapp0   --
ipmpapp0_1        ip         ok           ipmpapp0   --
lo0               loopback   ok           --         --
   lo0/v4         static     ok           --         127.0.0.1/8
   lo0/v6         static     ok           --         ::1/128
net0              ip         ok           bondmgt0   --
net1              ip         ok           bondmgt0   --
net4_iud1_734     ip         ok           bondeth0   --
net7_iud1_734     ip         ok           bondeth0   --
```

IPMP

```
ipmpstat -g
GROUP        GROUPNAME    STATE      FDT        INTERFACES
bondeth0     bondeth0     ok         --         net7_iud1_734 net4_iud1_734
bondmgt0     bondmgt0     ok         --         net0 (net1)
bondib1      bondib1      ok         --         bondib1_0 (bondib1_1)
bondib0      bondib0      ok         --         bondib0_0 (bondib0_1)
ipmpapp0     ipmpapp0     ok         --         ipmpapp0_0 (ipmpapp0_1)
```

IB presentation Local Zone

using a Vnic ?

NOPE

local zone definition

```
zonecfg -z nodedat03z01 export
create -b
set brand=solaris
set zonepath=/zoneHome/nodedat03z01
set autoboot=true
set pool=pool_nodedat03z01_17347
set limitpriv=default,proc_priocntl,proc_clock_highres
set ip-type=exclusive
add fs
…
add net
set configure-allowed-address=true
set physical=n21_1019_bondib0
end
add net
set configure-allowed-address=true
set physical=n22_1019_bondib0
end
add anet
set linkname=net4
set lower-link=net4
set configure-allowed-address=true
set link-protection=mac-nospoof
set mac-address=random
set vlan-id=734
..
add dataset
set name=nodedat03z01/nodedat03z01DB
```

On the local zone

```
ipmpstat -g
GROUP        GROUPNAME    STATE      FDT        INTERFACES
bondib1      bondib1      ok         --         n24_1020_bondib1 (n23_1020_bondib1)
ipmpapp0     ipmpapp0     ok         --         ipmpapp0_z01_1 (ipmpapp0_z01_0)
bonddg0      bonddg0      ok         --         net71 (net41)
bondib0      bondib0      ok         --         n21_1019_bondib0 (n22_1019_bondib0)
bondeth0     bondeth0     ok         --         net4 (net7)
bondmgt0     bondmgt0     ok         --         net0 (net1)
```

Project

M9000 state of the art in it's time

REPLACE

before maintenance is over

Upgrade EOL software to current version

10g -> 11.2.0.3 later this was 11.2.0.4

Prepare SSC and move to Exadata storage

Past

SITE A

Mirroring implemented via ZFS

SITE B

Past

Local zones share cpu's in pool with sybase

ZFS fs is used to host the datafiles

ZFS does the cross site mirroring not the storage

Static Database entries aka SID

Each database SAME name

Different port per DB -> 1522 - 16xx

Present / Future

RAC One node

DB resource manager

Consolidation

RAT

Solaris 11.1

parallelism

EM 12c

Data Guard

DOP

ASM

Instance caging

SCAN

Data Guard Broker

IORM

Services

VIP

Clusterware

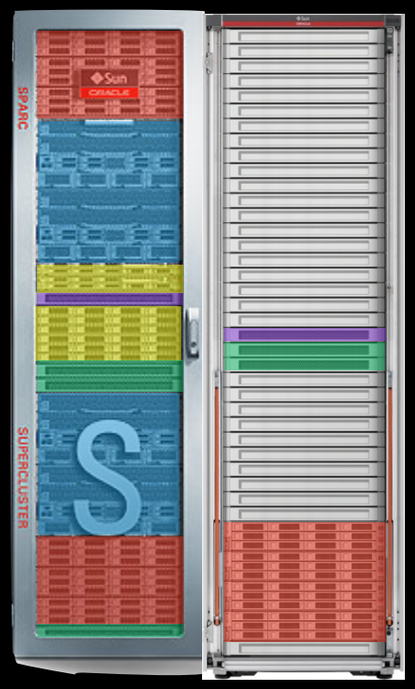

SSC01

Data Guard

SSC02

SSC design

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

UAT EXA
PRS EXA

APP DOMAIN

UAT EXA

APP DOMAIN

ZFS SA

PRO EXA

APP DOMAIN

PRO EXA

APP DOMAIN

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

Data Guard UAT /PRO

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

UAT EXA
INT EXA

DEV
EXA

APP
DOMAIN

UAT EXA

APP DOMAIN

ZFS SA

PRO EXA

APP DOMAIN

PRO EXA

APP DOMAIN

Exadata Storage Server

Exadata Storage Server

Exadata Storage Server

Project Overview

EOL

phase 1

LIFT & SHIFT

Move zones from M9000 to APP domain on Super Cluster

phase 2

Parallel : Upgrade from 9i and 10g to 11g

MAKE USE OF STORAGE CELLS

Simple Right ?

Not really

Technically not too complicated

But logistically

300 databases

even more applications

Limited team :

One part-time Unix Admin

One part - time DBA

ME

Lift and Shift

From native Solaris 10 to “Branded” zones

LDOM Solaris 10

Sun4u to Sun4v conversion

Detach zone

Attach ZONES target

QFSDP OCTOBER 2013

Post Lift & Shift

Post Lift and Shift Issues (1)

Performance

Full TBS take much longer from <1m to > 3m

Sybase Dumps take much longer

Sybase 15.7 not supported on Branded Zones

Post Lift and Shift Issues (2)

Lot’s of Conf Calls further

ZFS tuning : prefetch off

DISM not supported in 11.2 on SSC Doc ID 1468297.1

Slow segments in Data Dictionary fixed in 11.2.0.4

Multipathing to Logical Block instead of RR

Increased buffer cache

Increase priority of LGWR to FX-60

Dedicated resource pools on Solaris Level

Performance Call Lessons Learned

Route ticket to right team

Escalate !!!

Confs call help

Support focusing on AWR it helps to put retention > 1 w

Re read recommendations and interact with support

Summary post L&S actions

ZFS tuning

OS Tuning

Upgrade from 11203 to 11204 PSU 2

PART II

phase 2

RAC One node

DB resource manager

Consolidation

RAT

Solaris 11.1

Data Guard

EM 12c

parallelism

ASM

Instance caging

DOP

SCAN

Data Guard Broker

IORM

Services

VIP

Clusterware

Initial “temporary” configuration done by ACS in 2013

Some stuff missing :

- grid slicing dicing
- zones
- cell config X4 cells added
- QFSDP january
- RAC One install
- node listener running on admin nw

potential issues

About 300 db’s

Some with data dictionary bigger then data

RAC One in conjunction with >50 db’s
with just 16 cores

Oracle Client version …. SCAN

Consolidate

Consolidate the 11g way
Characterset
Criticality
Public synonyms …?
Access Rights …

Reduce from >50 db’s per ENV to +- 20 db’s

issues encountered

Upgrade from QFSDP OCT 2013 to QFSDP Jan 2014

10.2 db and 11.1 db fail to start

work around ipcrm …

IDR delivered

QFSDP APR 2014 again regression

IDR delivered

Griddisk config

PRO DAT

PRO FRA

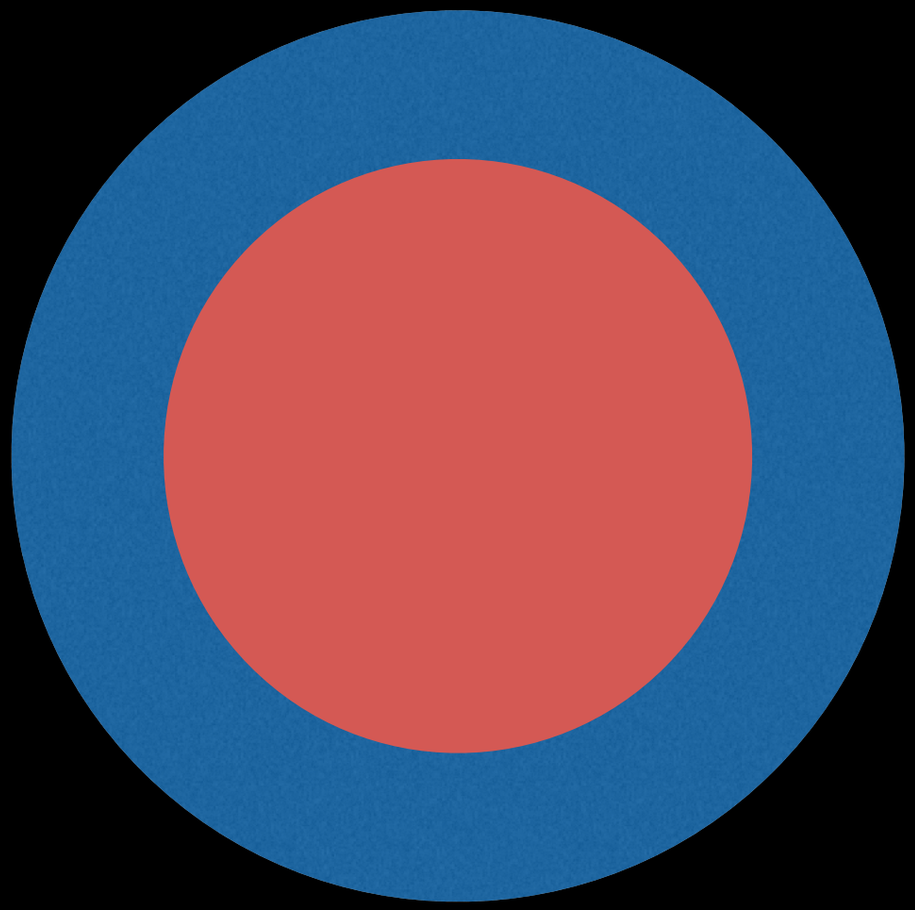

High Perf
X3

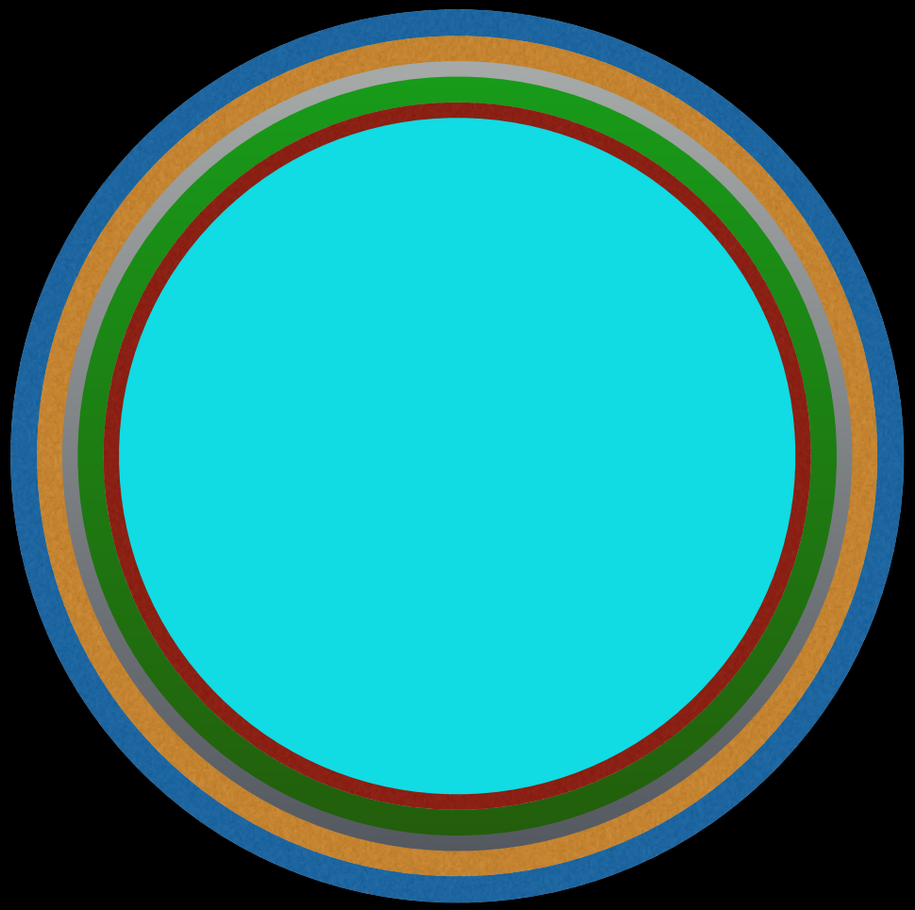

High Cap
X4

UAT DAT

INT DAT

DEV DAT

UAT FRA

INT FRA

DEV FRA

Rac One

Differences with Regular RAC

Controlled Active Passive solution

Always one instance active

Except during relocate this to allow failover

No 1 to 1 mapping node - instance nr

Rac One Relocate

```
srvctl relocate database -d SWINGUAT_SSC01 -n nodedat04z01 -v
Configuration updated to two instances
Instance SWINGUAT_2 started
Services relocated
Waiting for up to 30 minutes for instance SWINGUAT_1 to stop ...
Instance SWINGUAT_1 stopped
Configuration updated to one instance
```

Instance naming different

Rac => instance_name1
Rac One => instance_name_1

Rac One Binding

```
srvctl status database -d DGTEST_SSC02
Instance DGTEST_1 is running on node ssc02node02z01
Online relocation: INACTIVE

srvctl relocate database  -d DGTEST_SSC02 -n ssc02node01z01 -v
Configuration updated to two instances
Instance DGTEST_2 started
Services relocated
Waiting for up to 10 minutes for instance DGTEST_1 to stop ...
Instance DGTEST_1 stopped
Configuration updated to one instance

srvctl status database -d DGTEST_SSC02
Instance DGTEST_2 is running on node ssc02node01z01
Online relocation: INACTIVE

srvctl stop database -d DGTEST_SSC02
srvctl start database -d DGTEST_SSC02 -n ssc02node02z01

srvctl status database -d DGTEST_SSC02
Instance DGTEST_2 is running on node ssc02node02z01
Online relocation: INACTIVE
```

Rac One + Dg

requirements :

separate logically dg redo traffic from regular
use data guard broker

broker requires a static listener entry

Data Guard Physical Standby 11.2 RAC Primary to RAC Standby using a second network (Doc ID 1349977.1)
How to setup LOCAL_LISTENER for RAC One Node Database (Doc ID 1497619.1)
Data Guard: Redo Transport Services – How to use a separate network in a RAC environment. (Doc ID 1210153.1)
Martin Carsten Bach's blog

DG

create an extra network vips etc…

create a listener

static entry DGMGRL

like this one

```
SID_LIST_LISTENER =
        (SID_LIST =(
          (SID_DESC =
                        (GLOBAL_DBNAME = DGTEST_DGMGRL)
                        (ORACLE_HOME = /u01/app/oracle/product/11.2.0.4/dbhome_1)
                        (SID_NAME = DGTEST_1)
                    )
                (SID_DESC =
                        (GLOBAL_DBNAME = DGTEST_DGMGRL)
                        (ORACLE_HOME = /u01/app/oracle/product/11.2.0.4/dbhome_1)
                        (SID_NAME = DGTEST_2)
                    )
        )
```

DG REDO

```
# this is the tnsnames for the dg config which is SPECIFIC for DG and DIFFERS PER NODE
#

RACONE,NodeDAT01Z01_LOCAL_NET1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedat01z01)(PORT = 1521))
  )
DG_VIP,NodeDAT01Z01_LOCAL_NET2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg01z01-vip)(PORT = 1522))
  )
NodeDAT02Z01_LOCAL_NET1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedat02z01)(PORT = 1521))
  )
NodeDAT02Z01_LOCAL_NET2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg02z01-vip)(PORT = 1522))
  )
#REMOTE LISTENERS
NodeDATPRO_REMOTE_NET2 =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
     (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg01z01-vip)(PORT = 1522)))
    (DESCRIPTION =
     (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg02z01-vip)(PORT = 1522)))
  )
```

DG REDO

```
# this is the tnsnames for the dg config which is SPECIFIC for DG and DIFFERS PER NODE
#

NodeDAT01Z01_LOCAL_NET1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedat01z01)(PORT = 1521))
  )
NodeDAT01Z01_LOCAL_NET2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg01z01-vip)(PORT = 1522))
  )
RACONE,NodeDAT02Z01_LOCAL_NET1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedat02z01)(PORT = 1521))
  )
DG_VIP,NodeDAT02Z01_LOCAL_NET2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg02z01-vip)(PORT = 1522))
  )
#REMOTE LISTENERS
NodeDATPRO_REMOTE_NET2 =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
     (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg01z01-vip)(PORT = 1522)))
    (DESCRIPTION =
     (ADDRESS = (PROTOCOL = TCP)(HOST = nodedg02z01-vip)(PORT = 1522)))
  )
```

DG

configure listener networks

Data Guard Physical Standby 11.2 RAC Primary to RAC Standby using a second network (Doc ID 1349977.1)
How to setup LOCAL_LISTENER for RAC One Node Database (Doc ID 1497619.1)
Data Guard: Redo Transport Services – How to use a separate network in a RAC environment. (Doc ID 1210153.1)

```
listener_networks='((NAME=network1)(LOCAL_LISTENER=RACONE)
(REMOTE_LISTENER=ssc01scan01-pro:1521))'
,'((NAME=network_dg)(LOCAL_LISTENER=RAC_ONE,DG_VIP)
(REMOTE_LISTENER=NodeDATPRO_REMOTE_NET2))'
```

Resource Management

Before

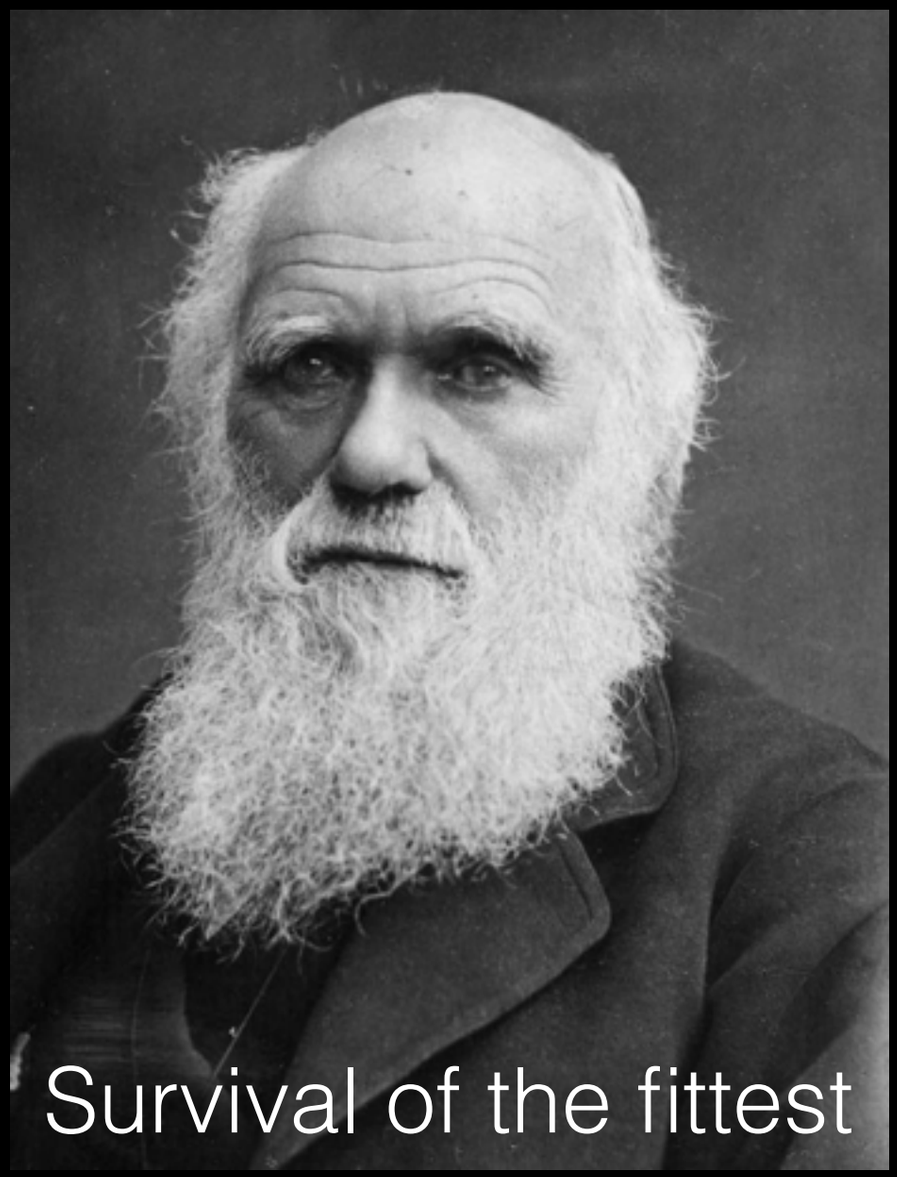

Resource Management

Inter database

Intra database

Category

1st step

Instance Caging

Simple to implement

Enable a resource plan

Specify a cpu_count

In practice

```
CPU count 2 => 12 sessions

PID USERNAME    USR SYS TRP TFL DFL LCK SLP LAT VCX ICX SCL SIG PROCESS/NLWP
  1995 oracle    18 0.0 0.0 0.0 0.0 0.0  82 0.0   9  12  18   0 oracle/1
  2016 oracle    18 0.0 0.0 0.0 0.0 0.0  82 0.0   9  11  18   0 oracle/1
  3782 oracle    17 0.0 0.0 0.0 0.0 0.0  83 0.0   9  10  18   0 oracle/1
  2010 oracle    17 0.0 0.0 0.0 0.0 0.0  83 0.0   9  11  18   0 oracle/1
  3749 oracle    17 0.0 0.0 0.0 0.0 0.0  83 0.0   8  11  16   0 oracle/1
  1999 oracle    16 0.0 0.0 0.0 0.0 0.0  84 0.0   8  10  16   0 oracle/1
  2014 oracle    16 0.0 0.0 0.0 0.0 0.0  84 0.0   8  11  16   0 oracle/1
  3772 oracle    16 0.0 0.0 0.0 0.0 0.0  84 0.0   8  10  16   0 oracle/1
  1750 oracle    16 0.0 0.0 0.0 0.0 0.0  84 0.0   8   9  16   0 oracle/1
  1757 oracle    16 0.0 0.0 0.0 0.0 0.0  84 0.0   8  11  16   0 oracle/1
  1744 oracle    16 0.0 0.0 0.0 0.0 0.0  84 0.0   8  14  16   0 oracle/1
  1721 oracle    16 0.0 0.0 0.0 0.0 0.0  84 0.0   8  11  16   0 oracle/1
Total: 373 processes, 1330 lwps, load averages: 2.99, 3.68, 4.48

CPU count 4 => 12 sessions
  2016 oracle    40 0.0 0.0 0.0 0.0 0.0  60 0.0  19  21  36   0 oracle/1
  1999 oracle    38 0.0 0.0 0.0 0.0 0.0  62 0.0  18  18  33   0 oracle/1
  1750 oracle    38 0.0 0.0 0.0 0.0 0.0  62 0.0  21  21  40   0 oracle/1
  1721 oracle    36 0.0 0.0 0.0 0.0 0.0  64 0.0  20  19  37   0 oracle/1
  2014 oracle    36 0.0 0.0 0.0 0.0 0.0  64 0.0  19  21  35   0 oracle/1
  1995 oracle    36 0.0 0.0 0.0 0.0 0.0  64 0.0  17  23  31   0 oracle/1
  3772 oracle    34 0.0 0.0 0.0 0.0 0.0  66 0.0  19  19  36   0 oracle/1
  1757 oracle    34 0.0 0.0 0.0 0.0 0.0  66 0.0  19  27  35   0 oracle/1
  2010 oracle    34 0.0 0.0 0.0 0.0 0.0  66 0.0  18  22  33   0 oracle/1
  3749 oracle    34 0.0 0.0 0.0 0.0 0.0  66 0.0  16  19  31   0 oracle/1
  3782 oracle    32 0.0 0.0 0.0 0.0 0.0  68 0.0  17  19  30   0 oracle/1
 1744 oracle    32 0.0 0.0 0.0 0.0 0.0  68 0.0  16  26  32   0 oracle/1
Total: 373 processes, 1330 lwps, load averages: 4.64, 4.01, 4.50
```

In practice

```
CPU count 8 => 12 sessions

   PID USERNAME USR SYS TRP TFL DFL LCK SLP LAT VCX ICX SCL SIG PROCESS/NLWP
  3772 oracle    86 0.0 0.0 0.0 0.0 0.0  14 0.0  22  31  41   0 oracle/1
  2014 oracle    82 0.0 0.0 0.0 0.0 0.0  18 0.0  25  36  43   0 oracle/1
  2016 oracle    82 0.0 0.0 0.0 0.0 0.0  18 0.0  23  30  44   0 oracle/1
  2010 oracle    82 0.0 0.0 0.0 0.0 0.0  18 0.0  19  27  32   0 oracle/1
  1999 oracle    80 0.0 0.0 0.0 0.0 0.0  20 0.0  18  31  34   0 oracle/1
  1757 oracle    80 0.0 0.0 0.0 0.0 0.0  20 0.0  21  29  34   0 oracle/1
  3749 oracle    80 0.0 0.0 0.0 0.0 0.0  20 0.0  24  31  40   0 oracle/1
  1995 oracle    78 0.0 0.0 0.0 0.0 0.0  22 0.0  21  41  38   0 oracle/1
  1750 oracle    76 0.0 0.0 0.0 0.0 0.0  24 0.0  21  37  32   0 oracle/1
  1744 oracle    76 0.0 0.0 0.0 0.0 0.0  24 0.0  25  32  38   0 oracle/1
  3782 oracle    72 0.0 0.0 0.0 0.0 0.0  28 0.0  21  29  34   0 oracle/1
  1721 oracle    70 0.0 0.0 0.0 0.0 0.0  30 0.0  24  28  35   0 oracle/1

Total: 378 processes, 1336 lwps, load averages: 8.21, 6.68, 4.71


   PID USERNAME USR SYS TRP TFL DFL LCK SLP LAT VCX ICX SCL SIG PROCESS/NLWP
  3772 oracle   100 0.0 0.0 0.0 0.0 0.0 0.0 0.0   0  25   0   0 oracle/1
  2016 oracle   100 0.0 0.0 0.0 0.0 0.0 0.0 0.0   1  31   2   0 oracle/1
  1757 oracle   100 0.0 0.0 0.0 0.0 0.0 0.0 0.0   3  31   5   0 oracle/1
  3749 oracle   100 0.0 0.0 0.0 0.0 0.0 0.1 0.0   1  27   2   0 oracle/1
  2014 oracle   100 0.0 0.0 0.0 0.0 0.0 0.1 0.0   3  38   3   0 oracle/1
  1995 oracle   100 0.0 0.0 0.0 0.0 0.0 0.1 0.0   4  29   5   0 oracle/1
  2010 oracle   100 0.0 0.0 0.0 0.0 0.0 0.2 0.0   1  30   1   0 oracle/1
  1744 oracle   100 0.0 0.0 0.0 0.0 0.0 0.2 0.0   2  27   4   0 oracle/1
  1750 oracle   100 0.0 0.0 0.0 0.0 0.0 0.2 0.0   3  27   6   0 oracle/1
  1721 oracle   100 0.0 0.0 0.0 0.0 0.0 0.2 0.0   6  27   8   0 oracle/1
  3782 oracle   100 0.0 0.0 0.0 0.0 0.0 0.3 0.0   7  27   9   0 oracle/1
  1999 oracle   100 0.0 0.0 0.0 0.0 0.0 0.4 0.0   5  29   6   0 oracle/1
Total: 372 processes, 1328 lwps, load averages: 12.15, 8.55, 5.65
```

2nd step IORM

topic for a next talk ;-)

for the moment we started easy

on database level

Later : category based IORM

3rd step DBRM

still under evaluation

paradigm shift

PARALLELISM

direct path reads are smart scan triggers

we tried auto DOP

Not a good idea for us

TOO AGGRESSIVE

the starting point

parallel_degree_policy = LIMITED

- no parallel statement queueing
- no in memory parallelism

parallel_degree_limit = 4 instead of CPU

degree on tables…
better on query level but not always possible

Tips & Tricks aka Lessons Learned

Compose a dedicated project team

Solaris on SSC is Solaris but with minor differences

Have a pilot

Make use of your platinum support, route to correct team

Follow QFSDP patches but not too closely ;-)