



ARDENT PERFORMANCE COMPUTING

# Oracle Services on RAC

*How They Work and What  
You Should Use Them For*



**Five Things You Might Not Know**

# JEREMY SCHNEIDER

[jeremy.schneider@ardentperf.com](mailto:jeremy.schneider@ardentperf.com)

Chicago  
ENFJ

Music  
Theology  
Swing Dancing  
Motorcycles




# Services

- What are services?
- Do you use RAC?
- Do you use services? (With or without RAC?)
- What do you hope to learn? What questions do you hope to have answered?



# Memorize or Understand?

The more they understand,  
the less they need to memorize




\*Thanks to Jared Still and Alex Gorbachev. Originally from:

[http://headrush.typepad.com/creating\\_passionate\\_users/2006/09/how\\_to\\_get\\_user.html](http://headrush.typepad.com/creating_passionate_users/2006/09/how_to_get_user.html)



# Memorize or Understand?



# Five things you might not know


- **Services are a lot more than just a connect method**
- There are three different ways to setup services
- Service configuration is stored in four different places
- There are multiple layers of load balancing and connection failover
- Services don't always do what you would expect



# What is a SERVICE?

```
ARDENT-SALES.LAB.ARDENTPERF.COM =  
  (DESCRIPTION =  
    (ADDRESS = (PROTOCOL = TCP) (HOST = rh4lab10) (PORT = 1521))  
    (CONNECT_DATA = (SID = ORCL))  
  )
```

```
ARDENT-ACCT.LAB.ARDENTPERF.COM =  
  (DESCRIPTION =  
    (ADDRESS_LIST=  
      (FAILOVER = on)  
      (LOAD_BALANCE = on)  
      (ADDRESS= (PROTOCOL = TCP) (HOST = rh4lab12) (PORT = 1521))  
      (ADDRESS= (PROTOCOL = TCP) (HOST = rh4lab13) (PORT = 1521))  
      (ADDRESS= (PROTOCOL = TCP) (HOST = rh4lab14) (PORT = 1521))  
    )  
    (CONNECT_DATA= (SERVICE_NAME = orcl.lab.ardentperf.com))  
  )
```



# What is a SERVICE?

```
rh4lab12:/home/oracle[orcl]$ lsnrctl stat
```

```
[...]
```

## **Services Summary...**

```
Service "PLSExtProc" has 1 instance(s).
```

```
Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
```

```
Service "orcl.lab.ardentperf.com" has 1 instance(s).
```

```
Instance "orcl", status READY, has 1 handler(s) for this service...
```

```
Service "orclXDB.lab.ardentperf.com" has 1 instance(s).
```

```
Instance "orcl", status READY, has 1 handler(s) for this service...
```

```
Service "orcl_XPT.lab.ardentperf.com" has 1 instance(s).
```

```
Instance "orcl", status READY, has 1 handler(s) for this service...
```

```
Service "test1.lab.ardentperf.com" has 1 instance(s).
```

```
Instance "orcl", status READY, has 1 handler(s) for this service...
```


```
Service "test2.lab.ardentperf.com" has 1 instance(s).
```

```
Instance "orcl", status READY, has 1 handler(s) for this service...
```

```
Service "test3.lab.ardentperf.com" has 1 instance(s).
```

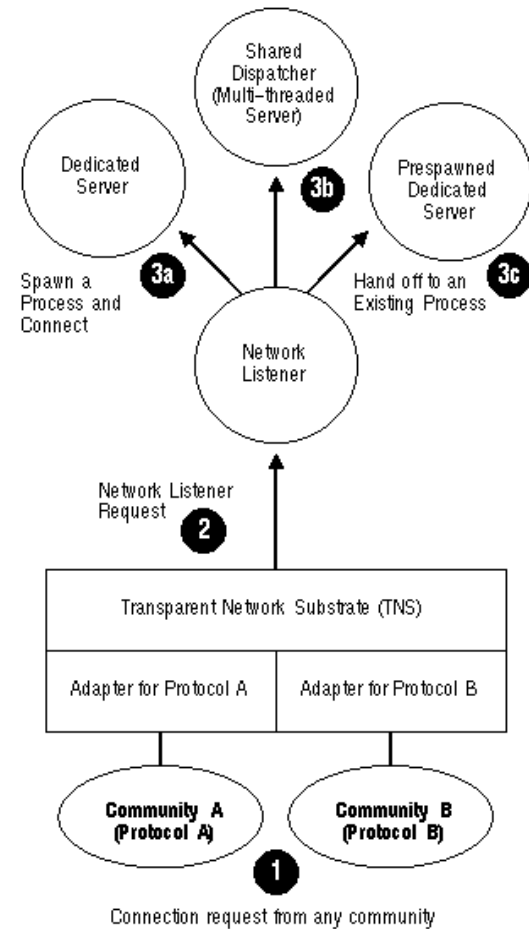

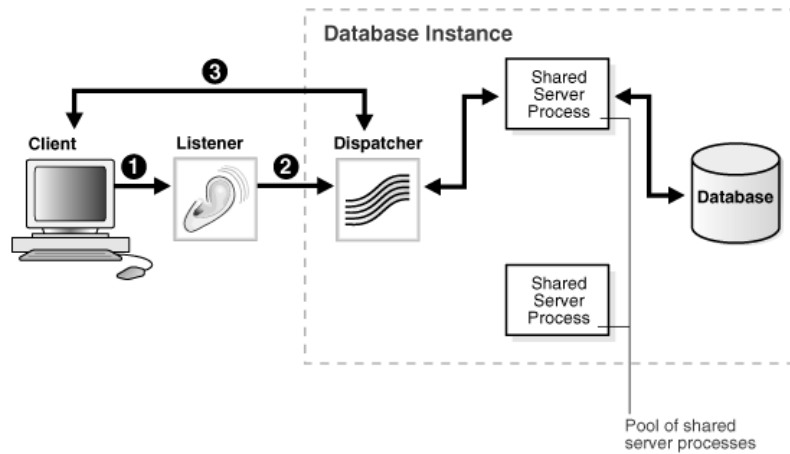
```
Instance "orcl", status READY, has 1 handler(s) for this service...
```

```
The command completed successfully
```






# What is a SERVICE?



# What is a SERVICE?



# More than just a connect method

## Tuning and Monitoring Aid

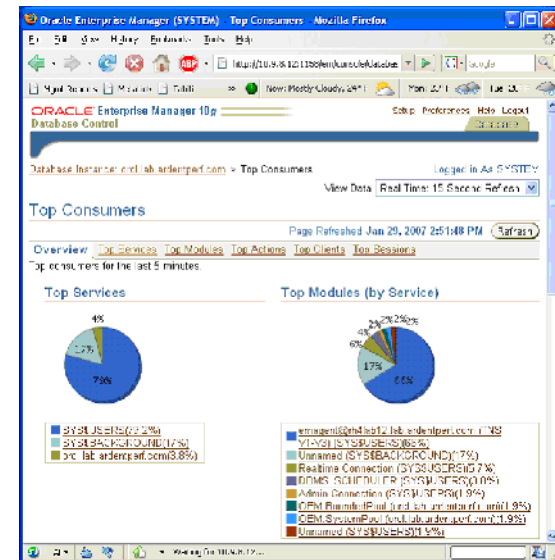
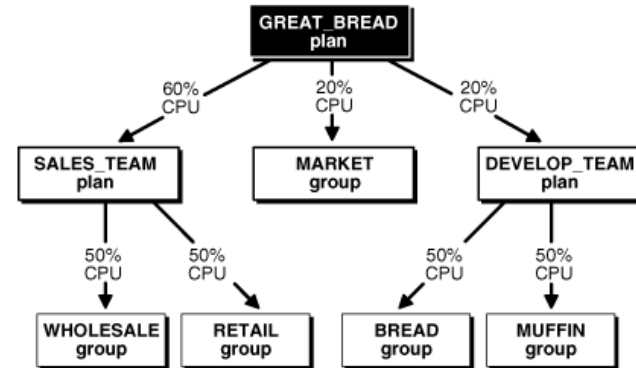
- Integration with dynamic views
- All stats automatically collected and aggregated by service
- AWR integration
- Alerts for response time thresholds

## Resource Management Aid

- Consumer groups can be mapped to services, modules and actions

## Job Control Aid

- Tuning, resource management
- Enable/disable jobs




# Single System Image (SSI)

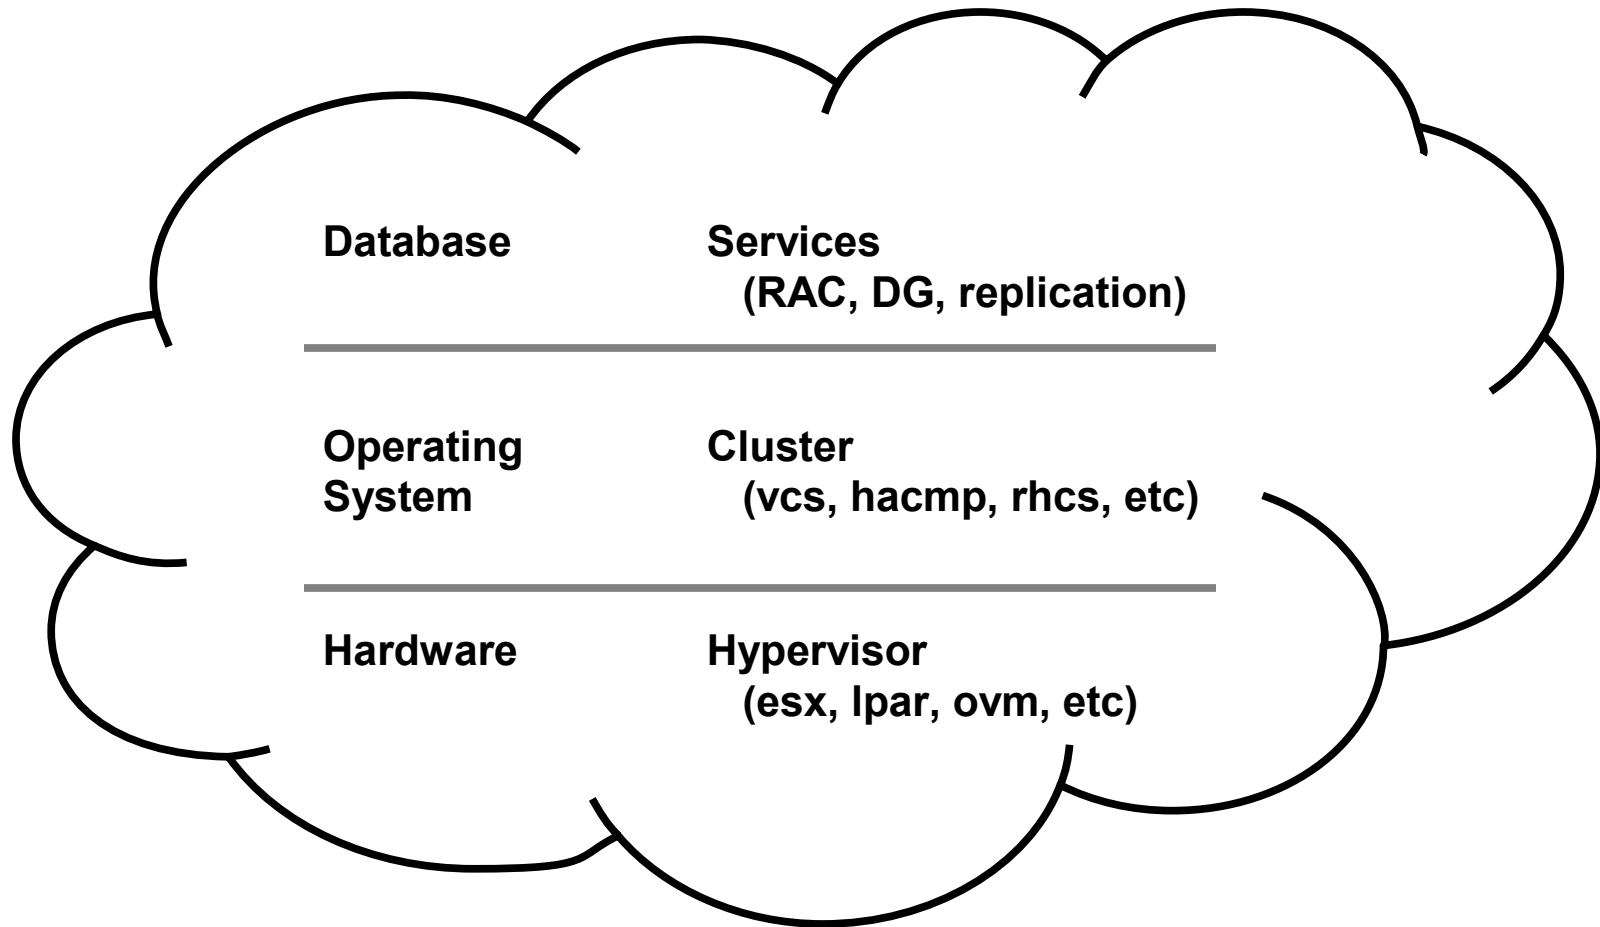
**A single system image is the illusion, created by software or hardware, that a collection of computing elements is a single computing resource.**

-- Gregory Pfister, In Search of Clusters

- Two key properties of SSI:
  - Every SSI has a boundary.
  - SSI can exist at different levels.
- In cluster discussions SSI usually means OS (perspective of admin).
- Services create SSI for database client.




# Single System Image (SSI)



# Five things you might not know

- Services are a lot more than just a connect method
- **There are ~~four~~ three different ways to setup services**
- Service configuration is stored in four different places
- There are multiple layers of load balancing and connection failover
- Services don't always do what you would expect



# Three ways to setup services

## 1. Manual for Non-RAC Databases:

- SERVICE\_NAMES parameter
  - Text editor (init file)
  - SQLPlus (spfile)
- DBMS\_SERVICES

## 2. Manual for RAC Databases:

- `srvctl` and `DBMS_SERVICES`
  - Best practice to always perform both steps in 11.1 and older.
  - Never set SERVICE\_NAMES!


```
[oracle@rh4lab15 ~]$ srvctl add service -d db2rac1 \  
> -s reports -r db2rac11,db2rac12 -P basic
```

```
SQL> begin  
2   dbms_service.create_service(  
3     service_name=>'reports',  
4     network_name=>'reports',  
5     goal=>dbms_service.goal_none,  
6     dtp=>FALSE,  
7     aq_ha_notifications=>FALSE,  
8     clb_goal=>dbms_service.clb_goal_long  
9   );  
10 end;  
11 /
```

PL/SQL procedure successfully completed.

```
[oracle@rh4lab15 ~]$ srvctl start service -d db2rac1 \  
> -s reports
```

```
[oracle@rh4lab15 ~]$ lsnrctl services  
Service "reports.lab.ardentperf.com" has 2 instance(s)  
Instance "db2rac11", status READY, has 2 handler(s)  
Instance "db2rac12", status READY, has 1 handler(s)
```



# Three ways to setup services

## 3. Database Console/Grid Control

- Only point/click method in 11g
- Only available for RAC databases
- Many available options
  - Service level thresholds
  - Consumer groups, job classes
  - Notification properties
  - Load balancing goal
  - DTP option

### High Availability Configuration

| Instance Name | Service Policy                             |
|---------------|--|
| db2rac11      | Preferred <input type="button" value="v"/> |
| db2rac12      | Preferred <input type="button" value="v"/> |

**TIP** Must select at least one preferred instance.

### Service Properties

Transparent Application Failover (TAF) Policy

Enable Distributed Transaction Processing

Choose this option for all Distributed transactions including XA, JTA. Services with exactly one preferred instance can enable this.

Connection Load Balancing Goal  Short  Long

Load balance connections based on elapsed time (Short) or number of sessions (Long).

### Notification Properties

Enable Load Balancing Advisory

Service Time  Throughput

Enable advisory for load balancing based on service quality.

Enable Fast Application Notification (FAN) for OCI and ODP.NET Applications

### Service Threshold Levels

If thresholds are specified, alerts will be published when the service elapsed response time and/or CPU time exceed the threshold.

|                                       | Warning              | Critical             |
|---------------------------------------|----------------------|----------------------|
| Elapsed Time Threshold (milliseconds) | <input type="text"/> | <input type="text"/> |
| CPU Time Threshold (milliseconds)     | <input type="text"/> | <input type="text"/> |

### Resource Management Properties

**TIP** To edit service associations with one or more consumer groups: [Click here.](#)

**TIP** To edit service associations with one or more job classes: [Click here.](#)







# Three ways to setup services

## 4. DBCA

- Final steps of database creation (number 12 of 15)
- Services management page from DBCA main menu
- Only available in 10g
- Maintains server's TNSNAMES.ORA
- Does not call DBMS\_SERVICES



# Five things you might not know

- Services are a lot more than just a connect method
- There are three different ways to setup services
- **Service configuration is stored in four different places**
- There are multiple layers of load balancing and connection failover
- Services don't always do what you would expect



# Where's the configuration stored?

cluster registry (ocr)

voting disk/quorum

ASM disk headers

control file

redo log

temporary tablespace

rollback/undo

datafile headers

data dictionary

workload repository (awr)

OS startup scripts (init.d)

oratab

oraInst.loc

inventory

init.ora or spfile

sqlnet.ora

listener.ora

tnsnames.ora



# Where's the configuration stored?

cluster registry (ocr)

ASM disk headers

control file

datafile headers

data dictionary

workload repository (awr)

oratab

init.ora or spfile

sqlnet.ora

listener.ora

tnsnames.ora



# Where's the configuration stored?

**cluster registry (ocr)**

ASM disk headers

control file

datafile headers

**data dictionary**

workload repository (awr)


oratab

**init.ora or spfile**

sqlnet.ora


listener.ora

**tnsnames.ora**



# Where's the configuration stored?

- Data Dictionary  
(Bulk of service configuration)
  - DBA\_SERVICES
  - DBA\_RSRC\_GROUP\_MAPPINGS
  - DBA\_THRESHOLDS
  - DBA\_SCHEDULER\_JOB\_CLASSES
- TNSNAMES.ORA
  - Server-side
    - Entries for LOCAL\_LISTENER and REMOTE\_LISTENER
  - Client-side
    - CONNECT\_DATA
    - Special entries for PRECONNECT services
- Cluster Registry
  - Resource
    - Start/stop script
    - Dependencies
    - Restart policy
  - Stringpairs
    - Instance list
    - Preferred/available
    - Enabled/disabled
    - TAF policy
- Initialization Parameters
  - LOCAL\_LISTENER
  - REMOTE\_LISTENERS
  - DB\_DOMAIN
  - DISPATCHERS
  - STATISTICS\_LEVEL




# Five things you might not know

- Services are a lot more than just a connect method
- There are three different ways to setup services
- Service configuration is stored in four different places
- **There are multiple layers of load balancing and connection failover**
- Services don't always do what you would expect




# Three layers of load balancing






# Three layers of connection failover



# Three layers of TNSNAMES

```
ARDENTACCT.LAB.ARDENTPERF.COM =  
  (DESCRIPTION_LIST =  
    (LOAD_BALANCE = ON)  
    (FAILOVER = ON)  
    (DESCRIPTION = ... )  
    (DESCRIPTION =  
      (LOAD_BALANCE = ON)  
      (FAILOVER = ON)  
      (ADDRESS_LIST = ... )  
      (ADDRESS_LIST =  
        (LOAD_BALANCE = ON)  
        (FAILOVER = ON)  
        (ADDRESS= (PROTOCOL = TCP) (HOST = rh4lab12) (PORT = 1521))  
        (ADDRESS= (PROTOCOL = TCP) (HOST = rh4lab13) (PORT = 1521))  
      )  
      (CONNECT_DATA= (SERVICE_NAME = orcl.lab.ardentperf.com))  
    )  
  )  
)
```

*“Client TNSNAMES” load balancing and “Connect-Time” failover.*



# Three layers of runtime failover

## 1. Client-Side TAF


- The original runtime failover – existed since Oracle8
- Use `FAILOVER_MODE` block in `TNSNAMES` file

## 2. Server-Side TAF

- New in 10g
- Attribute of service
- Overrides `TNSNAMES`

## 3. Fast Connection Failover

- New in 10g
- JDBC Connection Pools
- Utilizes FAN events



# Five things you might not know


- Services are a lot more than just a connect method
- There are three different ways to setup services
- Service configuration is stored in four different places
- There are multiple layers of load balancing and connection failover
- **Services don't always do what you would expect**



# Unexpected behavior




# Unexpected behavior

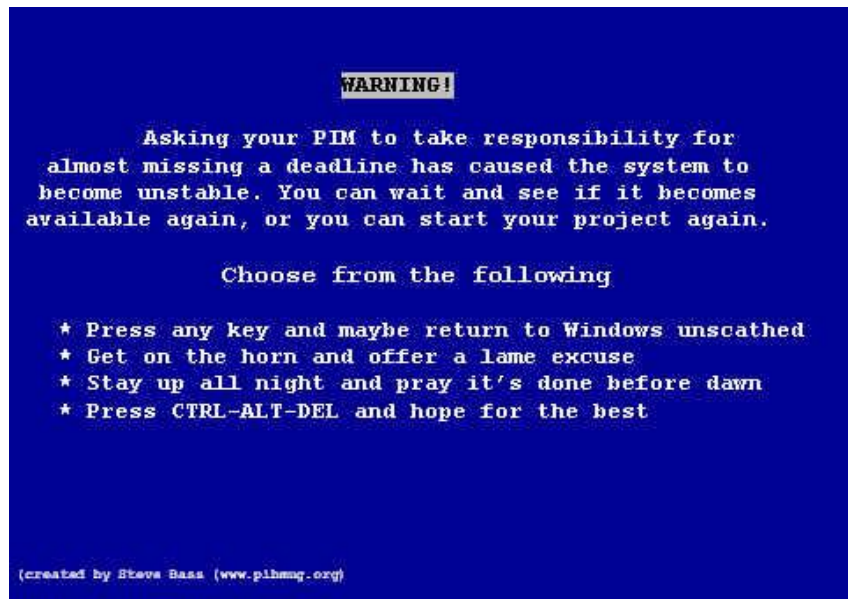


## Preferred-Available $\neq$ Single Node Performance

- What is an active-passive cluster?
  - Note: services obsolete `active_instance_count` for all practical purposes
- Cluster overhead impacts performance
- Caution: Partitioning different kinds of work on the same objects?
- If clustering only for HA, avoiding RAC is often better



# Unexpected behavior



**WARNING!**

Asking your PIM to take responsibility for almost missing a deadline has caused the system to become unstable. You can wait and see if it becomes available again, or you can start your project again.

Choose from the following

- \* Press any key and maybe return to Windows unscathed
- \* Get on the horn and offer a lame excuse
- \* Stay up all night and pray it's done before dawn
- \* Press CTRL-ALT-DEL and hope for the best


(created by Steve Bass (www.plhmg.org))

## No Failback

- After preferred node comes back online, service will NOT automatically move back from available node.

## No Auto-Start


- 10g clusterware will attempt to return services to their last-known state. If you shutdown with srvctl then the service will not start after reboot.



# Unexpected behavior

## Parallel Execution

- 10g and older versions
  - PQ slaves inherit service name
  - Slaves run on all nodes by default, ignoring PREFERRED and AVAILABLE nodes for service
  - Execution controlled through INSTANCE\_GROUPS and PARALLEL\_INSTANCE\_GROUPS
- 11g
  - Slaves run only on nodes where service is active
  - PARALLEL\_INSTANCE\_GROUPS overrides default behavior





# Unexpected behavior

## SCAN and DB Clients


- 11gR1 and older versions
  - Cannot handle multiple IP addresses returned by DNS or GNS
  - No connect-time failover
  - Include all three SCAN VIP addresses into TNSNAMES



# Five things you might not know

1. Services are a lot more than just a connect method
2. There are four different ways to setup services
3. Service configuration is stored in four different places
4. There are multiple layers of load balancing and connection failover
5. Services don't always do what you would expect

## How do I get started?



# What is a SERVICE?

**A service is a logical abstraction of a single group of database clients or a single application.**

Each service represents a workload with common:

- attributes
- service-level thresholds
- priorities

The grouping is based on attributes of work that might include:

- the application function to be used
- the priority of execution for the application function
- the job class to be managed
- the data range used in the application function or job class

## **Example Service Names**


cust-portal-prod

cust-portal-bulk-load

crm-prod

development-wiki

rman-repository



# Q&A

Questions, comments, suggestions?

