UPDATED

Improving Java Applica Availability with UCP and the GridLink Data Source

Simon Haslam

Veriton Limited

UKOUG RAC & HA SIG

31 March 2011

Veriton
Business Systems Integration

Simon Haslam / Veriton

Specialised consultant & established for 15 years

Middleware & SOA !

Demanding **web** & **call-centre** applications

Architecture & development strategy; health-checks; disaster recovery; tuning

Oracle Fusion Middleware
(JEE, SSO, OAM, OID, clustering)
ADF Applications (esp. strategy & admin)

Database & related technologies
(Solaris/Linux, load balancers, firewalls, …)

Veriton
Business Systems Integration

Agenda

1. Why & What? Architecture, FAN, FCF
2. UCP Features (+ compared with ICC)
3. UCP Configuration, Application Changes
4. Exalogic Active GridLink for RAC

Veriton
Business Systems Integration

1 Why… what's the problem ? !

- Lots of connections to database…
- Depending on type of failure, and what the app is doing, it make take a while for app to realise
- But in £££ HA £££ environment we expect:
  1. our application to 'carry on' after node failure
  2. to be able to do controlled shutdown of a node without affecting app

⇨ Otherwise… unresponsive app, users pressing refresh button, calls to helpdesk, etc

Veriton
Business Systems Integration

How do we solve this?

1) Database-tier needs to manage sessions so that app-tier doesn't necessarily notice anything wrong, <caveat> <caveat> <caveat>

or,

2) Database-tier needs to tell app-tier to sort itself out. <world of java>

Veriton
Business Systems Integration

1

- (diagram TAF & FCF – flip chart)

Veriton
Business Systems Integration

Fast Application Notification (FAN)

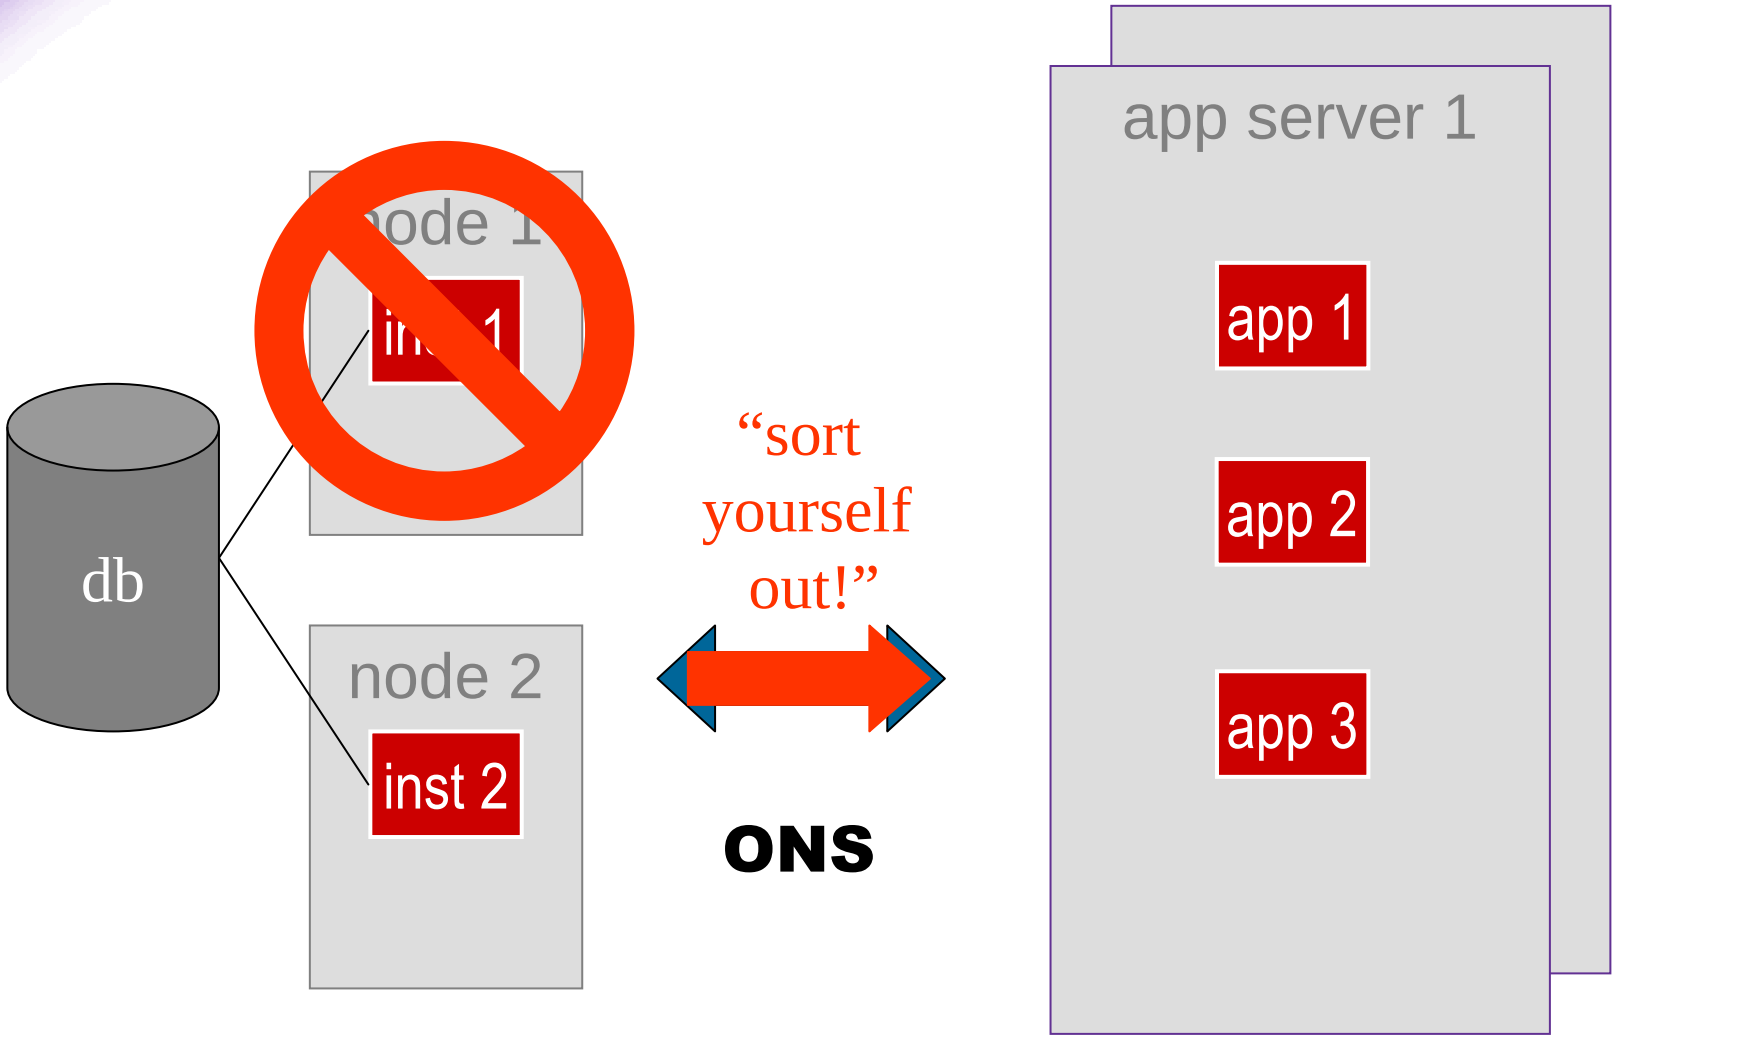

1 Fast Application Notification (FAN)

- FAN is a subset of messages sent out by Oracle Notification Service (ONS)

Veriton
Business Systems Integration

Fast Connection Failover (FCF)

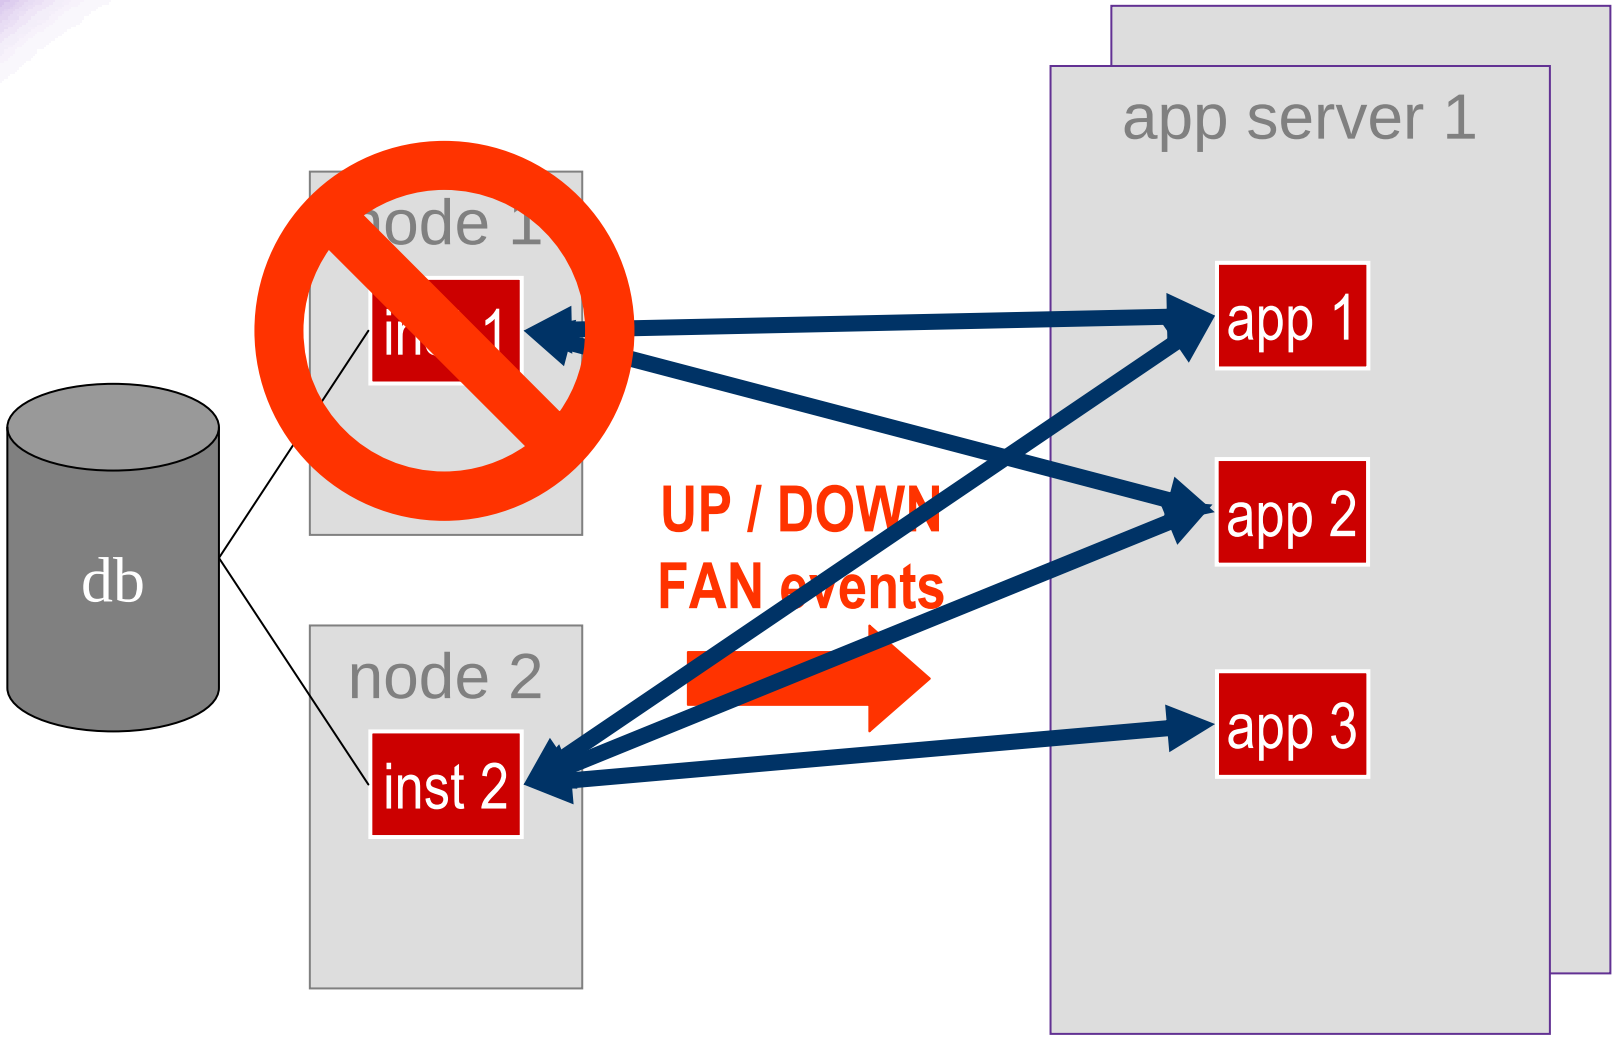

Veriton
Business Systems Integration

Connection Pool

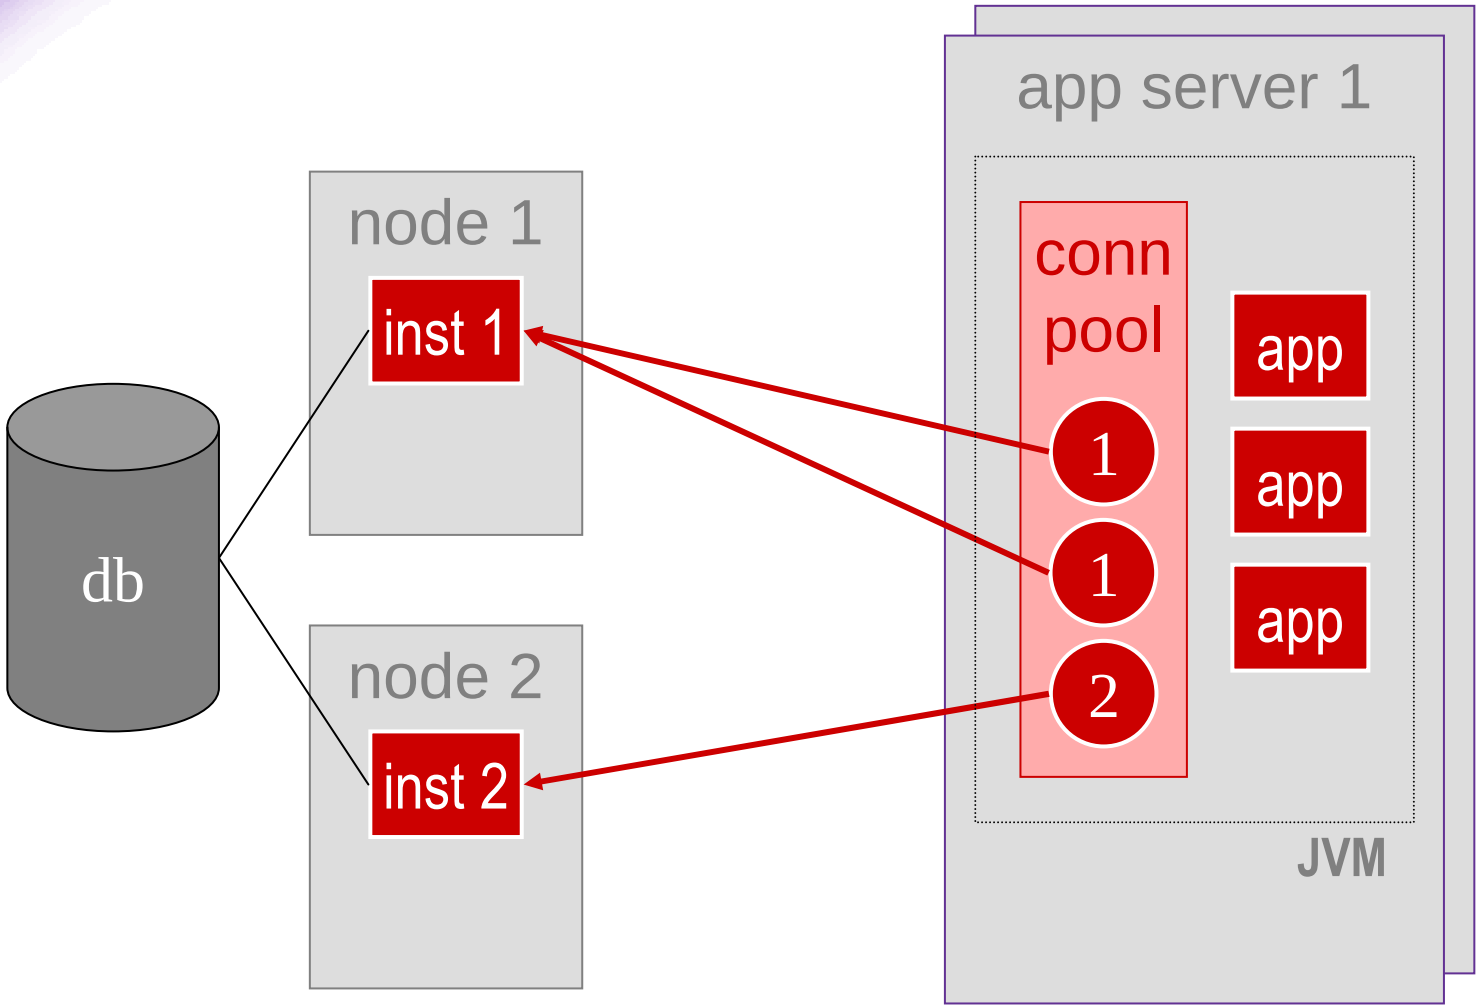

1 ONS is not just about Failures

- UP as well as down events
  - Re-start instance/node in cluster
  - Add node
- Runtime connection load balancing messages (RCLB)
  - Change of nodes)
  - Change of node load

Veriton
Business Systems Integration

Summary

- ONS = a publish/subscribe event notification mechanism
- FAN = subset of ONS messages for UP/DOWN events
- FCF = an application of FAN events to enable app tiers to respond appropriately
- Pool = a collection of shared database connections designed to reduce the amount of sessions & creation/tear-down operations

Veriton
Business Systems Integration

2 Two implementations of FCF for JDBC thin clients

- Implicit Connection Cache (ICC) in JDBC driver
  - deprecated in 11.2
- Universal Connection Pool (UCP) library
  - introduced in 11.1, strategic

Veriton
Business Systems Integration

2 Enter… UCP – the successor to ICC

- Generic connection pool for:
  - JDBC (non & XA), LDAP, JCA
  - any database that has a JDBC driver
- Used standalone or with an app. server
- With Oracle 11g UCP is integrated with:
  - FCF
  - RCLB
  - session / transaction affinity
  - RAC, RAC One-Node / Restart, Data Guard

…but would you?

Veriton
Business Systems Integration

UCP cf ICC

Looks like a `PoolDataSource`

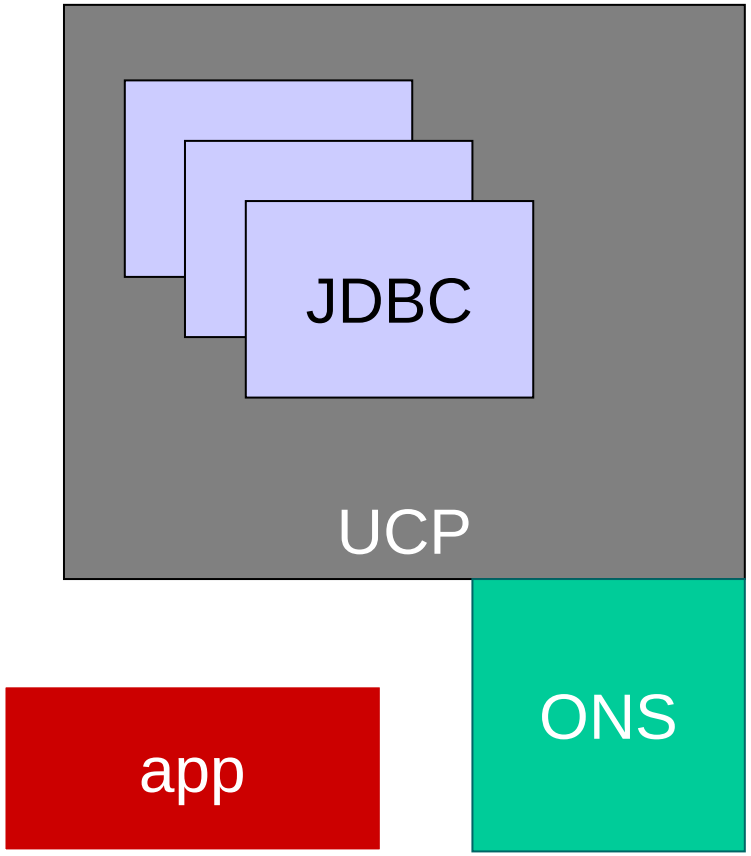

Looks like a `PoolDataSource`

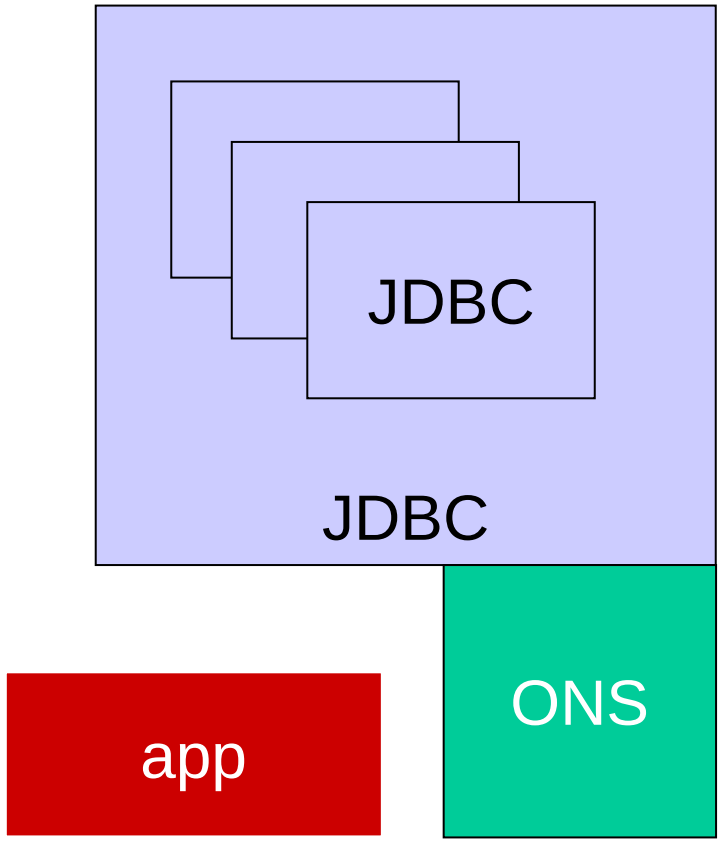

Veriton
Business Systems Integration

2 UCP has lots of nice features…

- RCLB (see Alex Gorbachev's presentations)
- Planned outages – allows in-flight txns to finish (cf ICC)
- Web-session affinity, e.g. shopping basket

- Management interface for pool status (& exposed as JMX too!)
- Pool configuration changes allowed at run-time

Note: most of these features require Oracle database

Veriton
Business Systems Integration

2

…and there's more

- XA transaction support
- Connection labelling
- Advanced timeouts
- Advanced diagnostics
- Application callbacks
- Can use SCAN addresses
- Dynamic ONS

Veriton
Business Systems Integration

Agenda

1. Why & What? Architecture, FAN, FCF
2. UCP Features (+ compared with ICC)
3. UCP Configuration, Application Changes
4. Exalogic Active GridLink for RAC

Veriton
Business Systems Integration

Remote ONS access

- Client can connect directly to ONS in database tier
- Clients need to know about 'all' database nodes…
  - 11.2.0.2: dynamic ONS
- No longer need to run ONS daemon on client
  - Oracle recommends remote now

Veriton
Business Systems Integration

11.2 RAC Default ONS Ports

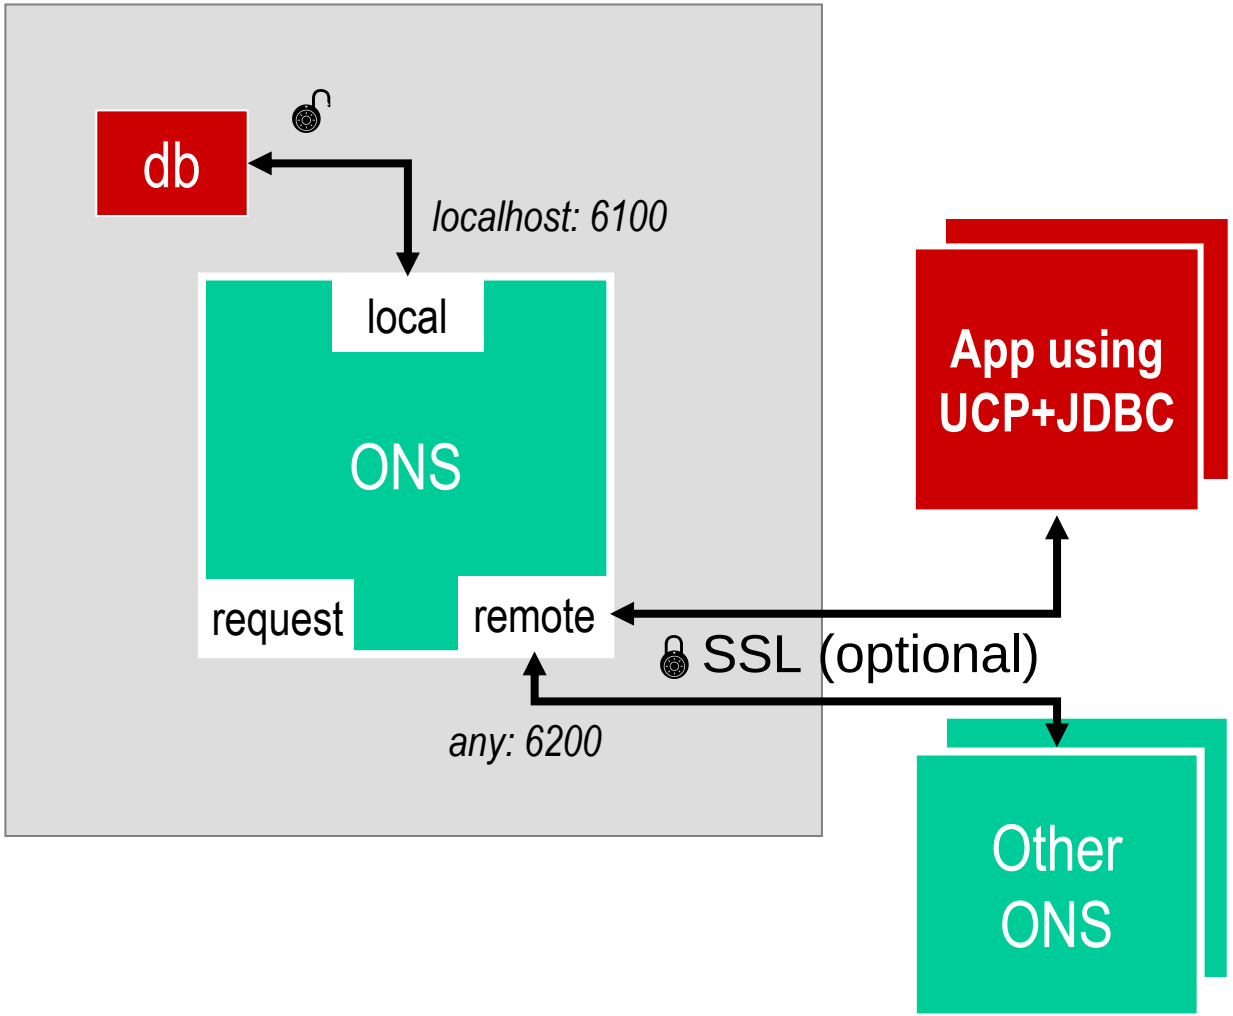

Note: Example host:port shown in italics

3 How to set up UCP + FCF (1)

```
pds = PoolDataSourceFactory.getPoolDataSource();
pds.setConnectionFactoryClassName(
      "oracle.jdbc.pool.OracleDataSource");
pds.setUser("scott");
pds.setPassword("tiger");
pds.setURL(“"jdbc:oracle:thin:@"
   + "(DESCRIPTION=(ADDRESS_LIST=(LOAD_BALANCE=ON) "
   + "(ADDRESS=(PROTOCOL=TCP)“
   + "(HOST=node1-vip.example.com)(PORT=1521)) "
   + "(ADDRESS=(PROTOCOL=TCP)“
   + "(HOST=node2-vip.example.com)(PORT=1521)))“
   + "(CONNECT_DATA=(SERVICE_NAME=appservice)))");
```

- or

```
pds.setURL("jdbc:oracle:thin:@"+
   "cluster-scan.example.com:1521/appservice");
```

3 How to set up UCP + FCF (2)

- Sensible:

```
pds.setInitialPoolSize(<number>);// def 0
pds.setMinPoolSize(<number>);    // def 0
pds.setMaxPoolSize(<number>);
                          // def Integer.MAX_VALUE
```

- To switch on FCF & remote ONS

```
pds.setONSConfiguration(
  “nodes=node1.example.com:6200,
  node2.example.com:6200{+wallet}”);
                          // Wallet is optional
pds.setFastConnectionFailoverEnabled(true);
```

- Refinements:

```
pds.setConnectionPoolName(…);
pds.setValidateConnectionOnBorrow(true);
  // validation over SQL*net
```

Veriton
Business Systems Integration

3

Application Changes

- App will still get failures if connected so needs to retry SQL
  - Request a new connection
  - Replay the transaction

- The application should **not** try to rollback the transaction as it has already been rolled back

Veriton
Business Systems Integration

3 Error Check with UCP

- Use the `isValid()` method when handling SQL exceptions:

```
if (conn == null || !((ValidConnection)
  conn).isValid())
{
  // retry what you were doing
  // since last commit
…
}
```

where conn is the Connection

Veriton
Business Systems Integration

3 Comparison with 'old' ICC

- Method on `OracleConnectionCacheManager` class to determine if the exception was fatal:
  `boolean isFatalConnectionError(SQLException e)`
- …basic (checks error code) and so not as effective as UCP's `isValid()`

Veriton
Business Systems Integration

3 Using UCP at Higher Level

- Data Source in an Application Server
  - E.g. Tomcat - I recommend Martin's blog post:
  - **http://martincarstenbach.wordpress.com/2011/02/03/getting-up-and-running-with-universal-connection-pool/**

- Are Java frameworks (e.g. Hibernate) UCP-aware???

Veriton
Business Systems Integration

Remember

- Logging:
  - if you have separate Oracle Clusterware user (e.g. `grid`) you need to look at ONS in `$GRID_HOME` not `$ORACLE_HOME`

3

Beware! ons.jar version

- If you choose not to install the full
- Oracle client (>1GB!) you need an `ons.jar` file*
- There are 3 `ons.jar` versions installed in Oracle server home (11.2):
  - `$OH/ons/lib` <= messages ignores, API changes
  - `$OH/oc4j/opmn/lib` <= messages ignored
  - `$OH/opmn/lib` <= you must use this one!!!
- (currently one post on OTN about it)

* Not *currently* available with UCP & JDBC downloads

Veriton Business Systems Integration

11.2 ONS Logging

- ons config files are suffixed with hostname
- Logging
  - loglevel=9 doesn't work in 11.2.0.1
  - Error:

```
"Unknown configuration parameter --
  loglevel=9, In file … "
```

- Instead you need to set:

  debugcomp=ons

- Doc (11.2 JDBC Dev Guide) was/is wrong
- Not published on web/MOS yet?

Veriton
Business Systems Integration

3 Note: 986065.1 Bug: 8779597

- Fast Connection Failover (FCF) Does not Work in 11.2 RAC Cluster Due to Incorrect ONS Events
- Fast Connection Failover (FCF) does not work correctly using Implicit Connection Cache (ICC) or Universal Connection Pool (UCP) when connected to a 11.2 RAC database cluster.
- When a database instance is shut down, all connections associated with that instance are not cleared up from the ICC or UCP.
- An ONS event is sent from the RAC cluster to the JDBC client when the database instance is shut down, but the connections associated with the shut down instance are not cleared as expected.

Veriton
Business Systems Integration

3

…continued

- Unpublished bug 8779597 - [11GR2-LNX-090731] A NUMBER OF JDBC CONNECTION IS INCORRECT WITH FAN UP EVENT

  11.2 RAC Database sends incorrect ONS events when the DB_DOMAIN database parameter is set.
  The Database Name field in the ONS event includes the domain, however the connections in ICC and UCP are created without the domain. For example , if the Database Name is *ORCL* and the db_domain is *uk.oracle.com*, all connections are created with ORCL as Database Name. The notification event from ONS, however, incorrectly sets the Database Name with domain as *ORCL.uk.oracle.com*. Since, the Database Name in the event does not match the Database Name in the connections, the event is not recognized by UCP or ICC and the connections do not get cleared.

Upgrade database to 11.2.0.2 or Apply patch for bug 8779597

Veriton
Business Systems Integration

Agenda

1. Why & What? Architecture, FAN, FCF
2. UCP Features (+ compared with ICC)
3. FCF Configuration, Application Changes
4. Exalogic Active GridLink for RAC

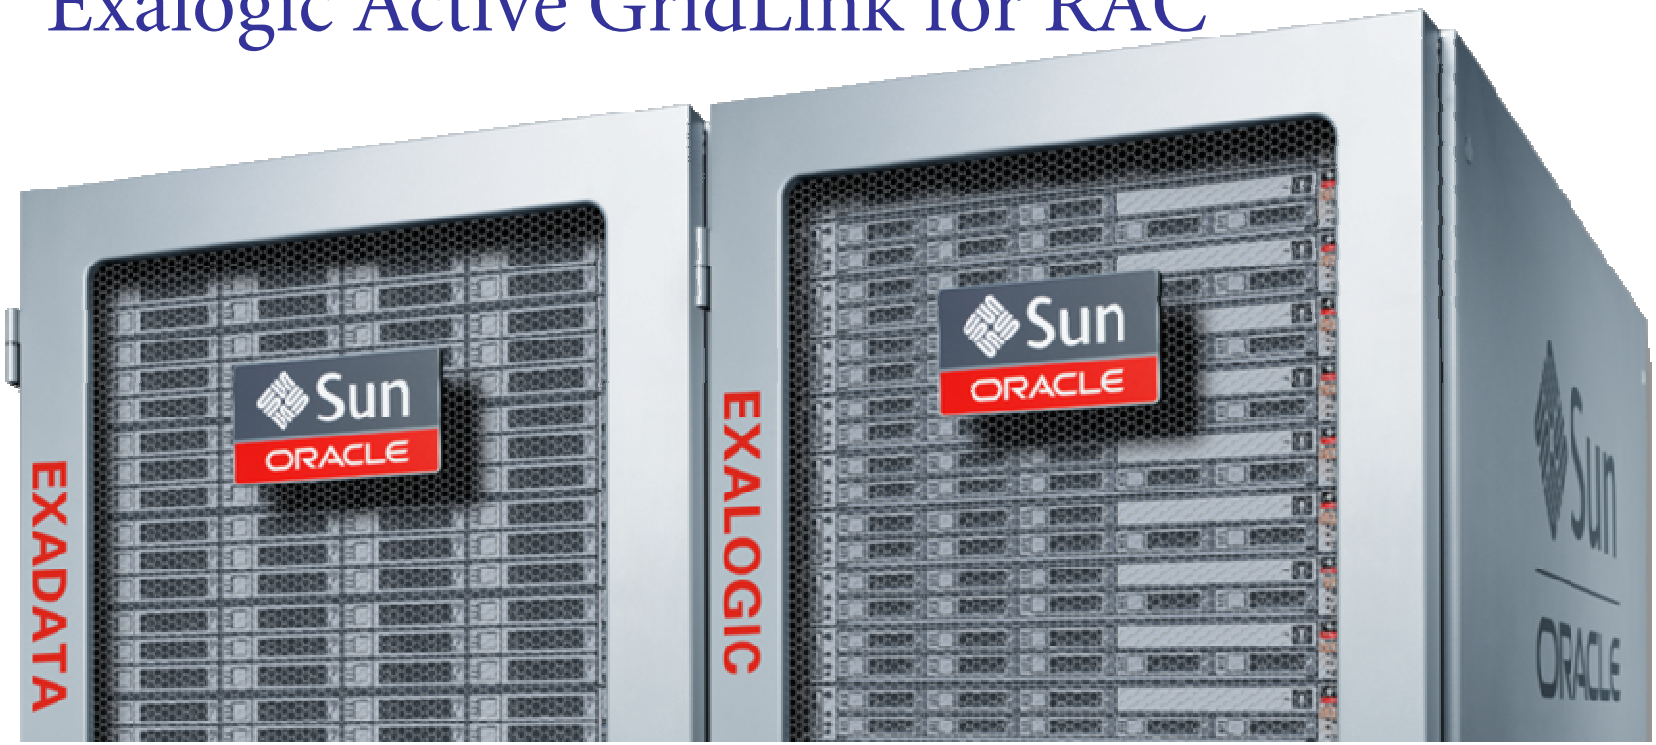

4

Active GridLink for RAC

- Oracle's post-BEA engineering work to improve integration between WebLogic and (Oracle) Database
- = GridLink Data Source
- Part of the Exalogic Elastic Cloud software (US$20k per Oracle Processor)
- Available from WebLogic Server 10.3.4 = 11g PS 3

Veriton
Business Systems Integration

4

Active GridLink for RAC

- Provides all the UCP features ‘out of the box’
  - FCF (e.g. graceful shutdowns)
  - RCLB (e.g. graceful startups)
  - Instance Affinity for XA (e.g. state mgmt)
- …plus InfiniBand performance enhancements

Veriton
Business Systems Integration

4

You are only allowed to use Active GridLink for RAC / GridLink Data Sources if you've purchased Exalogic Elastic Cloud software

Veriton
Business Systems Integration

4 BEA's Idea for RAC Connectivity

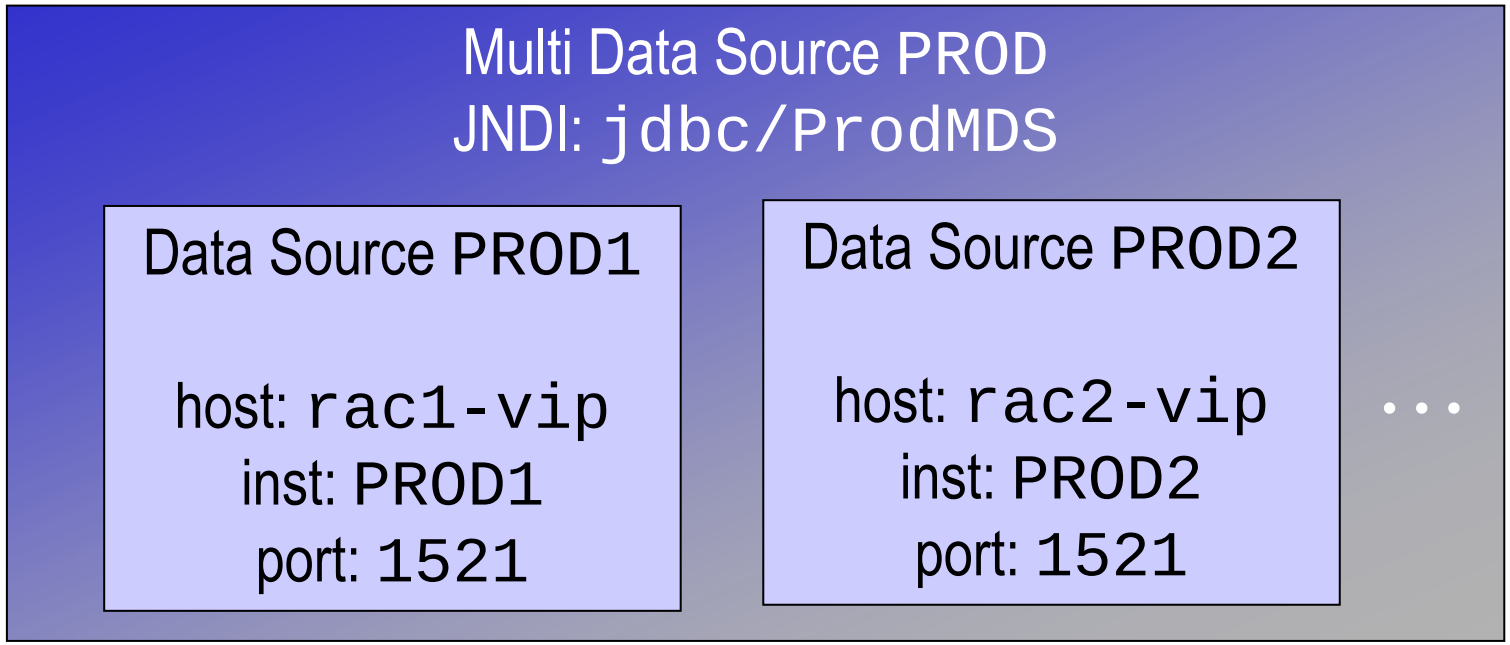

- Duplicated details
- Static configuration
- Failover & LB done by WebLogic

Veriton
Business Systems Integration

Another brainiac idea

- `testConnectionsOnReserve`
  - usually means `SELECT 1 FROM DUAL` prior to every call (e.g. transactional page request)
- enhancement: `secondsToTrustAnIdleConnection`

Veriton
Business Systems Integration

uses GridLink Data Source

Active GridLink for RAC ≠ GridLink for RAC

Very nice!

uses Multi Data Source

Veriton
Business Systems Integration

4 Connection Pools the Oracle way

- 'Database is King'
  - FAN
  - RCLB
- Pedigree: 10g JDBC & Oracle iAS

Veriton
Business Systems Integration

4

Configuration

- From WLS Console or WLST
  - See my January blog post **http://www.veriton.co.uk/roller/fmw/entry/gridlink_for_rac_data_source**
- Versions:
  - WLS 10.3.4
  - DB 11.2.0.2 + BP3* recommended
- Application:
  - Treat the same as UCP, i.e. test `isValid`

* thanks to Martin Bach

Veriton
Business Systems Integration

- ONS is the best thing we've got for pro-active service announcements
- Careful mid-tier configuration can increase application resilience
- App tier still needs code for handling connection failure
- GridLink Data Source is very nice (but comes at a price)… best to use 11.2.0.2 BP3

Veriton
Business Systems Integration

Thanks for listening!

SimonH@veriton.co.uk

Twitter: @simon_haslam

Oracle Fusion Middleware Admin blog:

http://simonhaslam.co.uk

ADF ENTERPRISE METHODOLOGY GROUP

Veriton
Business Systems Integration