Key Factors For a Successful ODA Deployment

Caleb Small
Caleb@Caleb.com

The Goal

To Provide blueprint for successful deployment of ODA devices in a common industry configuration

Background

- Primary database replicating to standby at remote site
- QA and Production environments
- Four ODAs total
- RMAN backup for both primary and standby
- Physical DataGuard between
- Migration & upgrade of existing 10g database

Critical Planning Issues

- Physical installation (briefly)
- Network configuration (most critical)
- Database migration strategy (start early)
- Backup and recovery strategy (where, how)
- Disk group allocation (size, redundancy)
- Disaster recovery strategy
- Initial connection and deployment
- Monitoring and maintenance
- Critical patching & updates

Planning The Physical Installation

- Rack space, power, network cabling, switch ports
- Shipping & delivery logistics
- ODA is *very heavy* – four people to lift
- Initial power up sequence
  - *Follow the instructions*
- Physical access to KVM or Mgmt Port
- Making the initial connection
  - *Especially the remote unit* !
- Power cycle testing

References:

- *Oracle Database Appliance Setup Poster*
- *Oracle Database Appliance Getting Started Guide*
- *Oracle Database Appliance Owner's Guide*
- *Oracle Database Appliance Service Manual*
- *Oracle Database Appliance, Safety and Compliance Guide*

Planning Networks – Required Networks

- Public Interface (dual port)
- ILOM port (out of band)
- Open necessary firewall ports
- DNS and NTP *required*
- SCAN configuration *required*
- NIC bonding

Planning Networks – SCAN Configuration

- One hostname that resolves to two IPs in round-robin configuration
- Set Time To Live (TTL) very short
  - Especially for pre 11gR2 clients
- Beware of
  - Routers w/caching DNS
  - Windows DNS client
- Test with repeated `nslookup` / `dig` commands (or ping)

References:

•11gR2 Grid Infrastructure Single Client Access Name (SCAN) Explained [ID 887522.1]
•How to Configure the DNS Server for 11gR2 SCAN [ID 1107295.1]
•How To Configure SCAN-OTN 129069

Planning Networks – SCAN Example

```
[oracle@beta1 ~]$ nslookup beta-scan.caleb.com
Server:         64.59.160.13
Address:        64.59.160.13#53

Non-authoritative answer:
Name:   beta-scan.caleb.com
Address: 192.168.1.26
Name:   beta-scan.caleb.com
Address: 192.168.1.25

[oracle@beta1 ~]$ nslookup beta-scan.caleb.com
Server:         64.59.160.13
Address:        64.59.160.13#53

Non-authoritative answer:
Name:   beta-scan.caleb.com
Address: 192.168.1.25
Name:   beta-scan.caleb.com
Address: 192.168.1.26
```

Planning Networks – SCAN Example

- A single *hostname* to access the cluster
- Cluster changes are invisible to clients
- Works best with 11gR2 client

```
racdb_taf =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = beta-scan)(PORT = 1521))
    (LOAD_BALANCE = YES)
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = racdb_taf)
      (FAILOVER_MODE =
        (TYPE = SELECT)(METHOD = BASIC)(RETRIES = 180)(DELAY = 5)
      )
    )
  )
```

Planning Networks – Firewall Ports

SP Network Ports		
5120	TCP	Oracle ILOM Remote Console: CD
5121	TCP	Oracle ILOM Remote Console: Keyboard and Mouse
5123	TCP	Oracle ILOM Remote Console: Diskette
5555	TCP	Oracle ILOM Remote Console: Encryption
5556	TCP	Oracle ILOM Remote Console: Authentication
6481	TCP	Oracle ILOM Remote Console: Servicetag Daemon
7578	TCP	Oracle ILOM Remote Console: Video
7579	TCP	Oracle ILOM Remote Console: Serial
CMM Network Ports		
8000 - 8023	HTTP over TCP	Oracle ILOM drill-down to server modules (blades)
8400 - 8423	HTTPS over TCP	Oracle ILOM drill-down to server modules (blades)
8200 - 8219	HTTP over TCP	Oracle ILOM drill-own to NEMs
8600 - 8619	HTTPS over TCP	Oracle ILOM drill-down to NEMs

Don’t forget 1521 (Oracle Net), 1158, 5500… (OEM), etc

Planning Networks – Firewall Ports

Common Network Ports		
22	SSH over TCP	SSH - Secure Shell
69	TFTP over UDP	TFTP - Trivial File Transfer Protocol (outgoing)
80	HTTP over TCP	Web (user-configurable)
123	NTP over UDP	NTP - Network Time Protocol (outgoing)
161	SNMP over UDP	SNMP - Simple Network Management Protocol (user-configurable)
162	IPMI over UDP	IPMI - Platform Event Trap (PET) (outgoing)
389	LDAP over UDP/TCP	LDAP - Lightweight Directory Access Protocol (outgoing; user-configurable)
443	HTTPS over TCP	Web (user-configurable))
514	Syslog over UDP	Syslog - (outgoing)
623	IPMI over UDP	IPMI - Intelligent Platform Management Interface
546	DHCP over UDP	DHCP - Dynamic Host Configuration Protocol (client)
1812	RADIUS over UDP	RADIUS - Remote Authentication Dial In User Service (outgoing; user-configurable)

References:

- *Oracle Integrated Lights Out Manager (ILOM) 3.0 HTML Documentation Collection http://docs.oracle.com/cd/E19860-01/E21549/z40001861019988.html*
- *Oracle Integrated Lights Out Manager (ILOM) 3.0 Daily Management — Concepts Guide*
- *Oracle Default Port List, red database security http://www.red-database-security.com/whitepaper/oracle_default_ports.html*

Planning Networks – Optional Networks

- 2 addn'l bonded 1 gig ether
- 1 addn'l bonded 10 gig ether
- Separate DataGuard network
- Separate Backup network
- NFS mounts
- ZFS storage
- etc

References:

•*Data Guard Physical Standby 11.2 RAC Primary to RAC Standby using a second network [ID 1349977.1]*

•*Expanding the Storage Capabilities of the ODA, Oracle White Paper, November 2012*

Planning Database Migration

- Start this first, it may take a while
- Migration vs. Upgrade (or both!)
- Common to find existing 10g (or older) migration to new ODA
- Physical movement vs. shared storage
- Various options available, many technical and white papers
- Not all options support an upgrade
- *Key factor* will be risk & complexity (cost) vs. downtime

Planning Database Migration

Possible options include:

- DataPump (simple, low risk, high down time)
- RMAN & DB Upgrade (min 1hr outage)
- DataGuard (logical for 10g → 11g but not ODA)
- Transportable Tablespace (data only)
- Streams / AQ (much setup)
- Golden Gate (extra license, great for non-Oracle)

All options have some limitation

(data type support, not all objects migrated, etc)

Planning Database Migration

Migration Technique	Complexity (Risk)	Skill Level	Outage Window	Selectivity	Extra Storage	Prep Work	Post Work
Transportable Tablespaces	Medium	Medium	Medium	Low	Yes	Medium	Low
Data Pump (Export/Import)	Low	Low	Long	Medium	Yes	Low	Medium
Recovery Manager	Medium	Medium	Short	Low	No	Low	Low
Procedural Approaches	High	Medium	Long	High	No	High	Medium

References:

- *Oracle Database Appliance: Migration Strategies, An Oracle White Paper, June 2012*
- *Database Rolling Upgrade Using Transient Logical Standby, Oracle MAA White Paper, May 2012*
- *Database Upgrade Using Transportable Tablespaces, Oracle MAA White Paper, February 2009*
- *Complete Checklist for Manual Upgrades to 11gR2 [ID 837570.1]*
- *Parallel Capabilities of Oracle Data Pump, An Oracle White Paper, July 2011*

Planning Backup Strategy

- Disk, tape, de-duplicator, etc
- Local (ASM) diskgroup vs. external storage
- Backup DB to FRA
- Backup FRA to tape, de-dup, snap mirror, etc
- Retention policy, off-site backup, archival backup, etc
- Archive log management
- Compression, encryption
- RMAN best practices (many opinions)

References:

•Oracle10g / 11g - Getting Started with Recovery Manager (RMAN) (Doc ID 360416.1)

•Top 10 Backup and Recovery best practices. (Doc ID 388422.1)

•Rman Don't Forget the Basics, Michael S. Abbey, NYOUG Webinar, February 3, 2012
http://www.nyoug.org/info/pod/2012-02-03-UKOUG-rman-best-practices.pdf

•10 Problems With Your RMAN Backup Script, Yury Velikanov
http://www.slideshare.net/yvelikanov/10-problems-with-your-rman-backup-script-whitepaper

Planning Disk Group Allocation

Redundancy

- Normal (2-way) vs High (3-way, default)
- 6Tb vs 4Tb divided for +DATA and +RECO
- Fault tolerance vs storage capacity
- +REDO (on SSD) is unaffected

Planning Disk Group Allocation

Allocation

- Both disk groups are striped across all disks
- Disks are partitioned during initial deployment
- “Internal” vs “External” backup option determines (raw) partition sizes
- /cloudfs is deducted from +RECO after split

diskgroup	internal backup	external backup
+DATA	1.6 TB	3.2 TB
+RECO	2.4 TB	0.8 TB
+REDO	97.3 GB	97.3 GB

Planning Disk Group Allocation

- Beware of REQUIRED_MIRROR_FREE_MB and USABLE_FILE_MB
- ASM “reserves” space to restore redundancy in the event of a disk failure
- You can not ***safely*** use all of the free space in a diskgroup

References:

- *Oracle ASM Administrator’s Guide 11g Release 2, ch 4 Administering Oracle ASM Disk Groups*
- *Demystifying ASM REQUIRED_MIRROR_FREE_MB and USABLE_FILE_MB, Harald van Breederode*
- *Oracle Database Appliance – Safely usable ASM diskgroup size, Marcel Lambrechts*

Planning Disaster Recovery (DataGuard)

- Physical vs. Logical
- Synchronous vs. Asynchronous
- Active Standby (extra license)
- Consider separate network & Listener
- Consider Advanced Compression option
- Method of (re)creating standby DB
- Numerous documents exist, plus so called “one button standby”

Continued….

Planning Disaster Recovery (DataGuard)

- Use DG Broker
- Consider carefully DG Observer
- Enable Flashback Database
- Implement DG auto health check
  - ***DO NOT*** rely on DG Broker health check !
  - Check SCN shipped ***and*** SCN applied
- Testing: Switchover vs. Failover

References:

- *Deploying Oracle Data Guard with Oracle Database Appliance, An Oracle White Paper, April 2012*
- *Data Guard 11g Installation and Configuration On Oracle RAC Systems, White Paper, October 2008*
- *Creating a Physical Standby Database on Oracle 11.2.0.x [ID 1475344.1]*
- *Step by Step How to Create Dataguard Broker Configuration [ID 984622.1]*
- *Redo Transport Compression in a Data Guard Environment [ID 729551.1]*

Next Steps

I’ve just unwrapped my brand new ODA what next?

Making the Initial Connection

Out of the box

- No network interfaces are configured
- ILOM is DHCP

Two best options:

- KVM to console (not ILOM)
- Laptop and console cable (RS-232) to Management Port (ILOM command line)

ILOM Access

- Configure static IP
- Required firewall ports
- Required browser version & Java plug-in

References:

- *ILOM configuration via Serial port [ID 1395445.1]*

Initial Network Configuration

for KVM

logon to OS and run firstnet to configure public and ILOM

```
$ /opt/oracle/oak/bin/oakcli configure firstnet
```

for console cable

logon to ILOM cli and configure ILOM IP

```
-> set /SP/network pendingipdiscovery=static
pendingipaddress=<IPv4_address> pendingipgateway=<gateway_address>
pendingipnetmask=<netmask_address>
-> cd /SP/network
-> set commitpending=true
```

- ILOM is now accessible with known IP
- Connect with crossover cable or management network
- Load in web browser and open console window to OS (requires Java plug-in !)

Verify and Upgrade System Software

- Two layers of software
  1. OS including OAK and oakcli
  2. Oracle software including GI and RDBMS
- Each is updated separately
- ODA is shipped with OS/OAK installed but *not* Oracle software
- Verify factory image version – may not be current
- Download and upgrade system software *BEFORE* deploying Oracle software (or may not get all features)

Install End User Bundle

- Download End User Bundle (Oracle Support)
- Copy to ODA Node 0 and ***verify checksum***
- Unpack with **`oakcli`**
- Run **`oakcli`** to deploy
- Be prepared to provide all required information
- Required network configuration ***must be in place***
- Must create a database, suggest throw-away, only limited options available
- Run dbca later with full options available
  - *Beware default DATA and RECO are reversed*

Recovering From a Failed Install – Wipe Clean

1. Run **`CleanDeploy.sh`**
   - Removes Oracle software and network configuration
   - ILOM configuration remains intact
   - Restore /etc/hosts if using X-Windows
2. Bare metal restore – factory image
   - Start at square one
   - Requires ILOM "*CD ROM Image*"
   - Restores factory OS image & wipes disks
   - Does not update device drivers, firmware, etc

References:

•*Oracle Database Appliance Bare Metal Restore Procedure. [ID 1373599.1]*

Monitoring and Maintenance

- OEM fans can use Grid Control/DB Console
- Daily alert log scanner (DB & ASM)
- Daily RMAN backup report
- Monthly log file rotation (archive & purge)
- Hourly DataGuard health check
- ILOM alerts
- Configure Automatic Service Requests (ASR), if possible

Monitoring & Maintenance

6:49 am	SUCCESS Oracle Alert Log Scanner [redacted]S2
6:46 am	SUCCESS Oracle Alert Log Scanner +ASM1 on [redacted]2a
6:45 am	SUCCESS Oracle Alert Log Scanner +ASM2 on [redacted]2b
6:44 am	SUCCESS Oracle Alert Log Scanner [redacted]S1
6:43 am	SUCCESS Oracle Alert Log Scanner [redacted]S2
6:42 am	SUCCESS Oracle Alert Log Scanner +ASM1 on [redacted]2a
6:41 am	SUCCESS Oracle Alert Log Scanner +ASM2 on [redacted]2b
6:40 am	SUCCESS Oracle Alert Log Scanner [redacted]S1
6:39 am	SUCCESS Oracle Alert Log Scanner [redacted]S2
6:34 am	SUCCESS Oracle Alert Log Scanner [redacted]RC1
6:33 am	SUCCESS Oracle Alert Log Scanner [redacted]RC2
6:30 am	SUCCESS Oracle Alert Log Scanner [redacted]RC1
6:29 am	SUCCESS Oracle Alert Log Scanner [redacted]RC2
6:25 am	SUCCESS: Oracle DataGuard health check on [redacted]S2
6:25 am	SUCCESS: Oracle DataGuard health check on [redacted]RC2
6:24 am	SUCCESS: Oracle DataGuard health check on [redacted]CS2

Critical Patching and Updates

- ODA patches are separate from usual GI / DB patches
- Still (roughly) follow quarterly PSU schedule
- Two separate patches (three layers):
  1. OS/OAK (Infrastructure)
  2. GI & DB patch (one patch, two layers)
- May require a short outage
- Will run DB upgrade scripts *IF* DB is open
- *Always* checksum patches after downloading

References:

•*Oracle Database Appliance - 2.X Supported Versions & Known Issues [ID 888888.1]*

Script Tips & Traps

- Install/update scripts run as root
- Docs expect root login – may (*should*) be a security issue!
- su from oracle or grid may fail if these users are re-created by the script
- Session may crash if default directory is under /cloudfs and CRS is re-started
- Custom configuration (eg SSH trust) of oracle or grid may be wiped out by patch / upgrade
- Consider creating alternate / wheel user(s)
- /cloudfs is very useful for common scripts, log file locations, backup job, dg_config, etc