

Bringing Exadata into an Oracle 11g organisation with VM Hosting

Joseph Rose,

Principal Consultant, Innovative Technical Architecture Ltd.

April 2017



For the Complete Technology & Database Professional



SAFE HARBOUR STATEMENT

- Every organisation is different and every first-time Exadata delivery experience will be different – these are my own views
- The experiences recounted are merely illustrative of challenges that any Exadata delivery project might face
- Learning achieved with the benefit of hindsight are very useful but are not necessarily the 'right answer' and there will often be multiple solutions
- Information provided purely for the purposes of learning and is disseminated in that spirit.
- Exadata technical delivery would not have been possible without the help, dedication and support of a hard-working internal project delivery team, timely support and backing from the management team, and a talented team of consultants from Oracle themselves.... I thank you all.

- Intro and mini-biography
- What these slides are about
- Initial state – how Oracle was delivered before
- Known issues – what were we trying to solve ?
- Typical architecture stack
- Technical Transition – too much change ?
- Physical delivery – OEDA, connecting it all up
- State of training, proposed solution and challenges
- State of operations, proposed solution to overloading
- Security – Exadata consolidation with segregation – gaining consensus
- Constraints – Key applications, supplier support ‘politics’
- Manageability / Monitoring – introducing Oracle Enterprise Manager
- Compliance – Finding a needle in a haystack
- Patching – Downtime and how to avoid it on a shared service
- Capacity planning – ADO policies, compression, maximising value
- The ideal world ? – Some thoughts for Product Management

INTRODUCTION

MINI-BIOGRAPHY

JOSEPH ROSE

<https://www.linkedin.com/in/josephpatrickrose/>

- Delivered Oracle Exadata into my previous client, starting initial discovery in July 2015, moving to procurement in May 2016 – working with Oracle UK
- Supported technical delivery , reaching production workloads in Feb 2017
- Lead Infrastructure Architect background, with 12 years Architecture experience, including Domain Architecture, setting standards, Infrastructure Architecture, Security- low-level Technical Architecture and E2E Solutions Architecture
- 13 years ago, I was an Oracle DBA, supporting a nationally-used application
- 25 years' experience in Infrastructure, Networking, Datacentres / DC Migration and all aspects of IT Architecture
- 20 years working for and delivering projects for blue-chip / Fortune 500 clients.

WHAT THESE SLIDES ARE ABOUT

- A walk-through of a sample delivery of Oracle Exadata into an organisation with no experience of this technology
- No prior experience of Exadata on my part, but 2 years' 8i DBA expertise and Infrastructure Architect experience, designing for Oracle RAC 9i, 10g, 11g
- A look at some of the challenges where the organisation itself is not working with shared services, and where the 'unit of delivery' is the 'VM server'
- A look at the constraints imposed by some major applications, including SAP, on how Exadata is delivered, should be patched and should be supported
- Some steps on the journey towards more maturity, in how Exadata and Oracle databases, are managed and maintained
- Some notes on the Security Controls selected, and some details on desired Security Controls that cannot yet be implemented

INITIAL STATE – HOW ORACLE WAS DELIVERED BEFORE

- Server-centric delivery of 2 or 3 tiered architectures :
 - Redhat Linux 6 or 7 used, pre-dominantly
 - Configuration Management at OS level resulted in an installed Oracle DB with a default database instance created
- Command-line configuration of additional Oracle products such as Database Vault or Advanced Security
- Per-instance local DBA accounts for DBA-level access
- No centralised management of Database instances
- No key management tools or procedures
- Little or no automation of repetitive DBA tasks
- Capacity planning or forecasting usually not attempted
- Additional space requested at VM level , often without analysis or reasoning
- Legacy Oracle versions not addressed, from a compliance perspective

KNOWN ISSUES

WHAT WERE WE TRYING TO SOLVE ?

- Per Oracle suggestion, any organisation with > 100TB of Oracle Databases should realise reduced costs and improve performance by utilising Exadata
- A factory-built RAC cluster would provide greater levels of resilience and availability – needed by increasingly business-critical workloads that were being delivered
- Scalable architecture supported large-scale consolidation
- Transaction peaks handled by Storage Server query offload
- Active Dataguard would provide a method of redirecting read-only requirements such as queries or BI ingestion, towards an alternate synchronous database replica, facilitating more regular ingestion or backup
- ZFS Storage Appliance would provide high-speed backup capability, fronted by clustered Oracle Recovery Manager, as well as a tablespace target for very old data.
- Oracle Enterprise Manager would provide visibility of the whole estate, with additional plug-ins and options providing compliance and delegated administration capabilities.
- Oracle Databases needed to move off VMWare, for licensing reasons.

TYPICAL ARCHITECTURE STACK

- While in some areas of the business, midrange platforms were still in use, the typical architecture stack would employ Redhat Linux as the OS, with:
 - Apache Web Server
 - Application Server such as Tomcat or JBoss
 - Oracle Database 11g or 11g R2:
 - ITIL Bronze/Silver services would typically use a single server, with VMWare resilience
 - ITIL Gold/Platinum services would use Redhat Linux clustering and Active/Passive Oracle instances
 - RAC was very rarely used, owing to reliability problems with historical deployments

TECHNICAL TRANSITION TOO MUCH CHANGE ?

- Although the project was delivering Exadata as new technology, other things were changing too:
 - Oracle 12c had not been available as a Configuration Management package, so 11g R2 was invariably used
 - Individual Oracle instances on VM meant no ASM storage
 - RAC clusters had not been employed, with OS clusters used instead
 - DBAs used to command-line operation, not OEM 13c
 - VMWare Site Recovery Manager and Storage replication used to protect VMs, not Active Dataguard
 - ZFS was also new technology and hadn't been deployed elsewhere in the business
 - Lifecycle Management Pack and Cloud Management Pack - totally new options

PHYSICAL DELIVERY OEDA / CONNECTING IT ALL UP

- As with any brand-new equipment, Exadata anticipates up-to-date connectivity options, when commissioning takes place
- On the networking side:
 - 10Gb networking was only available via fibre-presentation
 - Dataguard networking (copper) was only provided with 1Gb or bonded connections
 - During 2017, support for additional 2-port Fibre Ethernet card in slot 1 of Database Server announced, with Exadata 12c R2 software
- Physical placement of ZFS tied to IB connection:
 - Infiniband basically a 'backplane in a cable'
 - As a result, IB cables are very expensive
 - ZFS rack location should be left of Exadata rack 1, supporting future growth
- OEDA process designed to support collection of customer-specifics such as Exadata / Server names, IP addressing and options:
 - Version 111.5 used
 - IP subnets allocated for various purposes
 - Not possible to reserve IP addresses for specific purposes (VIPs, DB server addresses etc.) without editing the XML file produced by OEDA

STATE OF TRAINING

PROPOSED SOLUTION AND CHALLENGES

- One of the challenges of making any Exadata delivery successful is good DBAs, sufficient training, time to learn and put new features into practice
- The amount of technical transition taking place necessitated good training provision:
 - Classroom training was proposed (to ensure it happened)
 - Operations team could not release all DBAs for training, so CBT was proposed instead
 - Unlimited CBT access for 12 months, to all Oracle training was negotiated
 - Challenges continued, owing to the amount of project delivery activity that required DBA contribution
 - Build team had more Oracle background ; Run team more skilled with DB2
 - Day-release for training on new features - Oracle 12c, RAC, Exadata, ASM, ZFS – very challenging on an ongoing basis
 - Oracle University provided a feedback loop on progress ; ensuring that DBAs complete the relevant training is important when there is a lot of Technical Transition and this needs to be managed.

STATE OF OPERATIONS

PROPOSED SOLUTION TO OVERLOADING

- Very apparent, during Exadata delivery and delivery work on key projects that DBA team appeared to be overloaded
- Proposed a list of daily / weekly / monthly checks:
 - Time to complete each check to be recorded
 - Additional checks required for critical workloads
 - Further checks and due diligence required for SAP workloads
 - Time taken to complete checks could be totalled and used to estimate the number of DBAs required to meet Project Delivery and Operate – known as the ‘FTE Impact’.
 - Standard Operating Procedures used to formalise how each check should be done
 - Figures produced could be used to justify and drive further recruitment, to meet the longer-term requirement for more DBAs
 - ...DBA costs still had to be competitive with other parts of the business, where a service provider delivered to a fixed cost

SECURITY

CONSOLIDATION WITH SEGREGATION

GAINING CONSENSUS

- Exadata was new to the organisation, including the Security team
- Security team working globally on new policies, while Exadata delivery was progressing
- Shared services represent challenges when data needs to be segregated
- Approach to gain design consensus centred around presenting:
 - A series of technical options, with increased Security Controls applied
 - Obtaining consensus on the interim and final technical approach
 - Security team focussed on separate hosting for most sensitive workloads as an objective
 - Interim approach involved segregation of network traffic and ASM storage groups
- ...Networks team concerned about single routing table, although policy-based routing technique was successfully implemented.

CONSTRAINTS

KEY APPLICATIONS, SUPPLIER SUPPORT 'POLITICS'

- Early on in the Exadata delivery process, it was clear that SAP was going to feature in a number of business-critical projects
- SAP provided a number of constraints:
 - Oracle VM database servers could not be used
 - Various Oracle database options could not be used or had restrictions placed on their use
- SAP provided a large body of 'SAP notes' describing various issues or constraints:
 - A very thorough body of technical documentation
 - A large additional dimension to build or run team DBA workload ... increased due diligence to perform patches, configuration changes
- Other major suppliers offered 'no support' or 'conditional support' for Exadata:
 - Sometimes, the issue is RAC usage, where the s/w house also offers its own DB engine
 - Inherent issues with RAC, such as DND, addressed with Exadata improvements

MANAGEABILITY / MONITORING

INTRODUCING ORACLE ENTERPRISE MANAGER

- Historically, there was a siloed way of working with Database servers, and OEM was not used
- With an increased number of 'Oracle targets' , a redundant OEM implementation was a firm requirement to manage the estate
- OEM plug-ins provided additional insight into operation of the Exadata, ZFS and other aspects
- Monitoring and ticketing required additional steps:
 - SNMP-type monitoring type used to collect Linux and OS-level alerts for 24 hour monitoring
 - OEM link to Ticketing system a requirement for formal Service Introduction, as Service Desk did not want to learn how to use another tool.
 - SMTP / Email alerts to team mailboxes of questionable practical use.
 - OEM Integration proved valuable to ensure any Exadata or Oracle alerts were immediately actioned and properly routed, for rapid resolution.
- MOS Training for support team was quickly completed and this helped with escalation of any very complex technical issues.

COMPLIANCE

FINDING A NEEDLE IN A HAYSTACK

- Exadata delivery was driven by and funded by a small number of very large projects
- Nevertheless, a large number of legacy Oracle databases exist, and in-time these need to be remediated:
 - Legacy Oracle DBs might contain security vulnerabilities and require patching or upgrade
 - Versions may be out of extended support
 - Critical or Security patches may not be applied
 - Oracle configuration files such as Sqlnet.ora need to be compliant
- Discovery of ALL the Oracle estate using OEM is a firm requirement
- Pushing a Lifecycle Management Pack agent to each Oracle execution environment allows us to capture metadata about running workloads
- Lifecycle Management Pack reports access should be given to Security team and/or Auditors, so that overall compliance can be managed and improved.

PATCHING

DOWNTIME AND HOW TO AVOID IT ON A SHARED SERVICE

- A long time was spent debating the correct way to do Exadata patching
- Serious roadblocks existed when it appeared Exadata downtime would be required to complete full-stack patching – various business-critical services had negotiated specific downtime arrangements with the business units and were unwilling to repeat this process
- Proposals were made to delay Database Server patching, maybe by 12 months – it appeared to be ‘too hard’
 - With further consideration, a sensible approach is to :
 - For each DB node:
 - Evacuate the RAC node
 - Complete OS/Clusterware patching, per database node
 - Allow workloads to run on the newly patched node
 - (repeat over successive nights)
 - Complete Oracle Database Home cloning for various Oracle home types (SAP, non-SAP, others)
 - Complete Oracle Database Home patching, at times to suit each critical business service, Data-guarding the active workload to an alternate Exadata machine, as required

CAPACITY PLANNING

ADO POLICIES, COMPRESSION – MAXIMISING VALUE

- Oracle 12c introduces Automated Database Optimization (ADO), which provides a mechanism for managing data movements between tablespaces, according to business rules
- If database delivery onto Exadata exceeds capacity forecasts, the platform can quickly be filled up
 - It's important to identify opportunities to use compression and to set this up on various tablespaces
 - Advanced Row compression, HCC query and HCC Archive compression can be used at different times in the Data Lifecycle
 - ADO rules can manage automatic movement of data between tablespaces
 - Business use of the data and how this changes must be carefully understood
 - Projections for use of each type of tablespace is critical for forecasting capacity usage
 - Effect of all workloads over all Exadata machines can be totalled – good management information
 - Early warning of capacity shortages is important to allow Exadata Storage servers to be ordered:
 - 2-4 weeks to get an approved purchase order in a typical large organisation
 - 6 weeks actual lead-time for the equipment
 - 1-2 weeks to arrange the Oracle engineer, worst case
 - Actual compression ratios and space consumed should be recorded at regular intervals and compared to projected values
 - Value of 'actuals' for estimation purposes must take into account business onboarding of users and business processes for new solutions – when is each solution 'fully live' ?

THE IDEAL WORLD

SOME THOUGHTS FOR PRODUCT MANAGEMENT

- Every organisation is different ☺
- Networking :
 - Switch included to provide Management port reduction
 - What about 10GbE Copper ?
 - Should X7-2 contain 10/25GbE Copper switches, with QSFP uplink cables ?
 - Should X7-2 contain the option for all-fibre NICs ?
 - Exadata sometimes needs to connect into legacy core networks...
- OEDA
 - Make first-time commissioning easier, avoiding editing the XML file
 - Infrastructure teams want more control over naming and use of IP address ranges
 - Allow naming policies to be configured more easily, ensuring consistency across multiple OEDA files
- Platinum vs ACS managed patching
 - One team managing all patching ?
 - SAP-patching expertise to be included in standard offering ?
 - Downtime challenges to be explained to organisations at the start of the process – in full ...
- Workload discovery
 - Oracle DB discovery tools to be provided free-of-charge to customers
 - Reports required to explain the required capacity and Exadata resources needed – customer's understanding of what they need is often inaccurate – very difficult to revise estimates and purchase more resources, after procurement stages are complete



THANKS FOR LISTENING ...

Any questions ?