Domain Services Clusters

Centralized Management & Storage for an Oracle Cluster Environment

Markus Flechtner

BASLE ▪ BERN ▪ BRUGG ▪ DÜSSELDORF ▪ FRANKFURT A.M. ▪ FREIBURG I.BR. ▪ GENEVA
HAMBURG ▪ COPENHAGEN ▪ LAUSANNE ▪ MUNICH ▪ STUTTGART ▪ VIENNA ▪ ZURICH

trivadis
makes IT easier.

Our company.

Trivadis is a market leader in IT consulting, system integration, solution engineering and the provision of IT services focusing on ORACLE and Microsoft technologies in Switzerland, Germany, Austria and Denmark. We offer our services in the following strategic business fields:

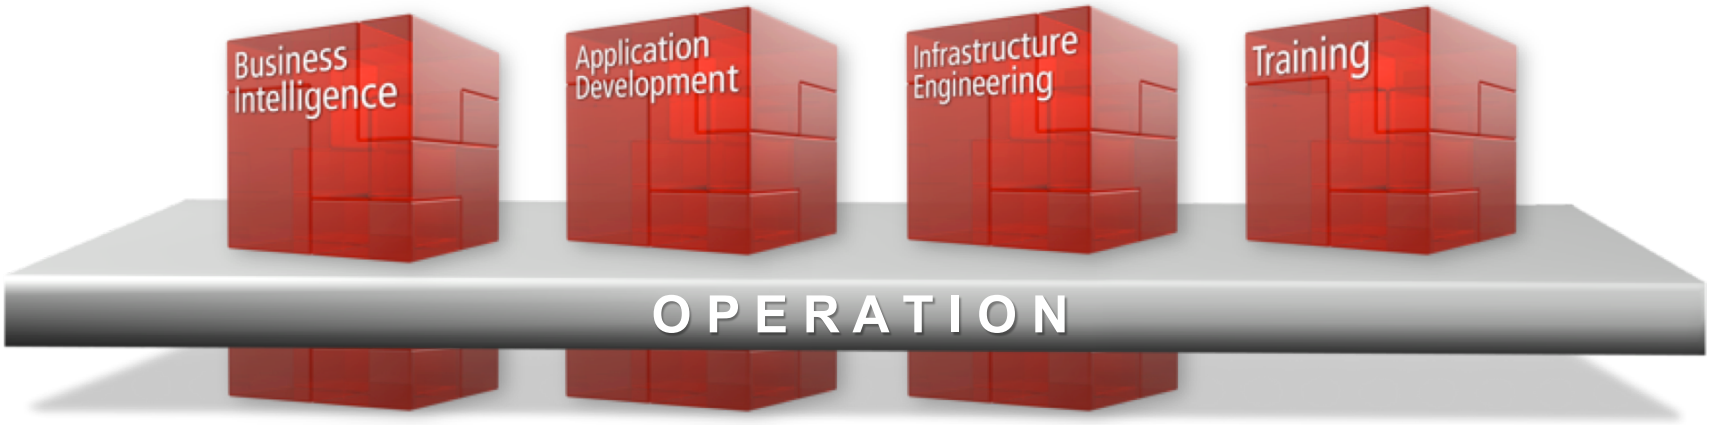

Trivadis Services takes over the interactive operation of your IT systems.

trivadis makes IT easier.

With over 600 specialists and IT experts in your region.

- 14 Trivadis branches and more than 600 employees
- 200 Service Level Agreements
- Over 4,000 training participants
- Research and development budget: CHF 5.0 / EUR 4 million
- Financially self-supporting and sustainably profitable
- Experience from more than 1,900 projects per year at over 800 customers

trivadis
makes IT easier.

About Markus Flechtner

- Principal Consultant, Trivadis, Duesseldorf/Germany, since April 2008
- Discipline Manager Infrastructure Database @Trivadis
- Working with Oracle since the 1990's
  - Development (Forms, Reports, PL/SQL)
  - Support
  - Database Administration
- Focus
  - Oracle Real Application Clusters
  - Database Upgrade & Migration Projects
- Teacher
  - O-RAC – Oracle Real Application Clusters
  - O-NF12CDBA – Oracle 12c New Features for the DBA

Blog:
https://markusdba.net/

@markusdba

DOAG

RACSIG
Real Application Clusters
Special Interest Group

trivadis
makes IT easier.

Agenda

1. **Cluster Domain Architecture**
2. **Installing a Domain Services Cluster**
3. **Installing Member Clusters**
4. **Cluster Domain Services**
   - Centralized GIMR
   - Centralized TFA
   - Storage Service
   - Rapid Home Provisioning
5. **Summary & Outlook**

Cluster Domain Architecture

Cluster Domain Architecture

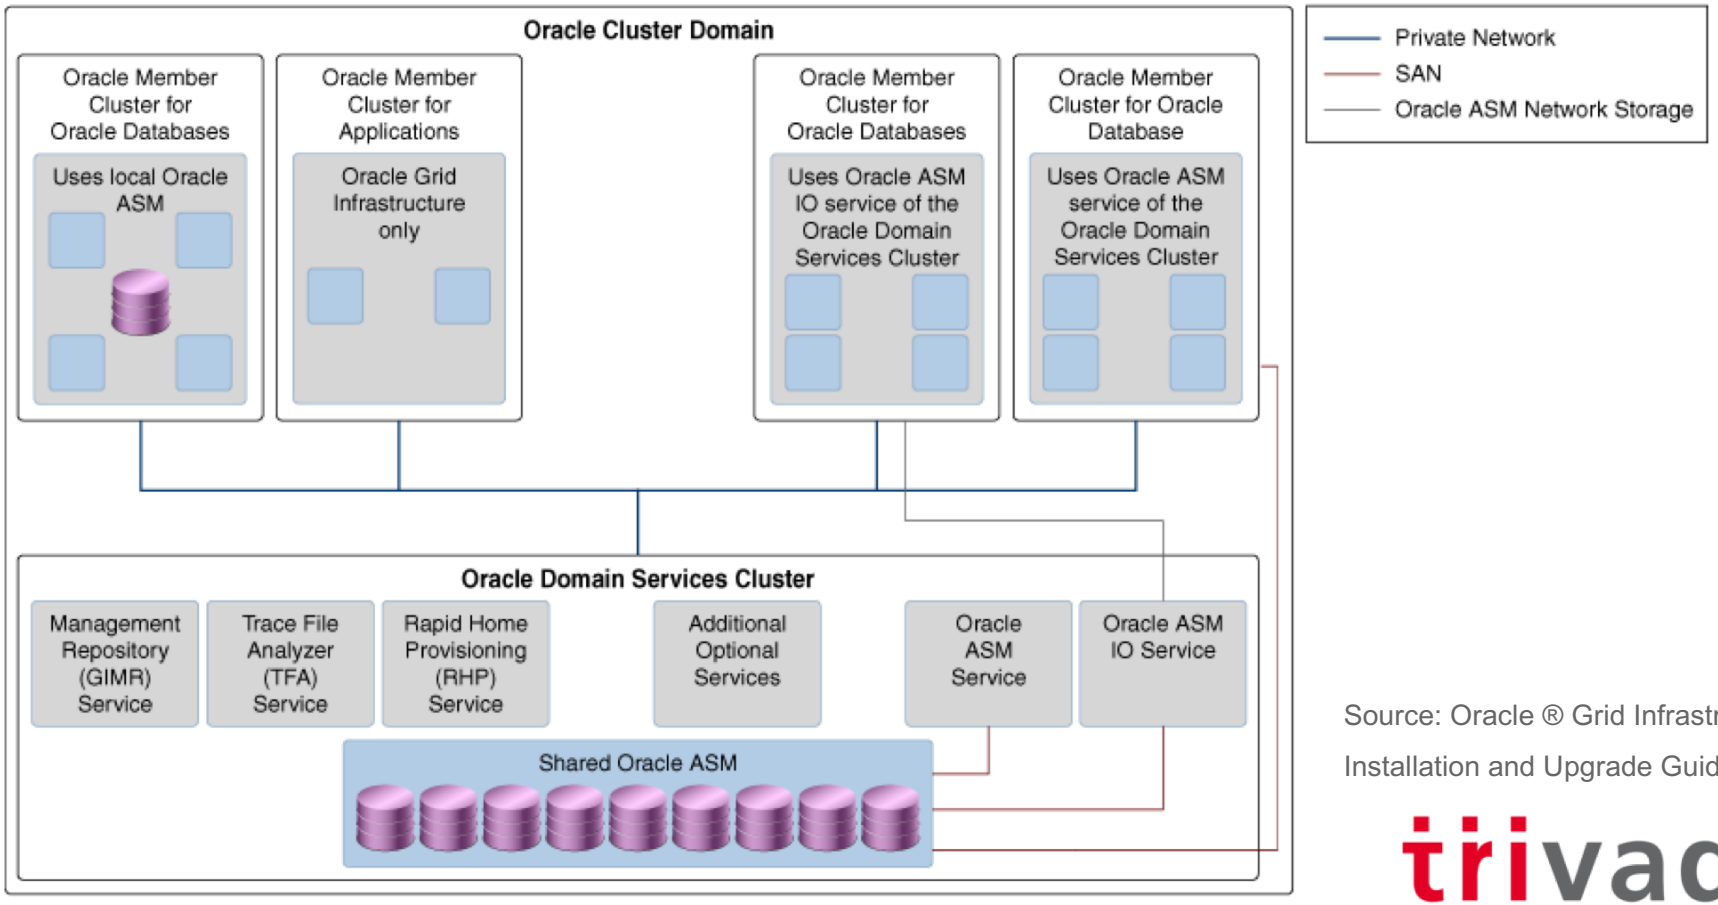

Source: Oracle ® Grid Infrastructure Installation and Upgrade Guide

Domain Services Cluster

- Provides centralized services for the member clusters (Domain Services)
- Standard Services
  - Grid Infrastructure Management Repository (GIMR)
  - Trace File Analyzer Collector (TFA)
- Optional: Rapid Home Provisioning (RHP)
- Optional: Storage Services
  - ASM Service (direct access to centralized storage)
  - IO-Service (indirect access to centralized storage)

Member Cluster Types

- Member Cluster for Applications
  - Lightweight cluster stack (DB specific services were removed)
- Member Cluster for Databases
  - With local storage (= local ASM instance)
  - Without local storage / access to storage via DSC ASM service
  - Without local storage / access to storage via DSC IO service

- In Oracle 12.2 it is not possible to convert an existing cluster into a member cluster

Benefits

- Centralized Management
- Less overhead on the member clusters
- Sharing storage between clusters

Installing a Domain Services Cluster

Installing a Domain Services Cluster (1)

- It's an option when installing a Cluster (gridSetup.sh)

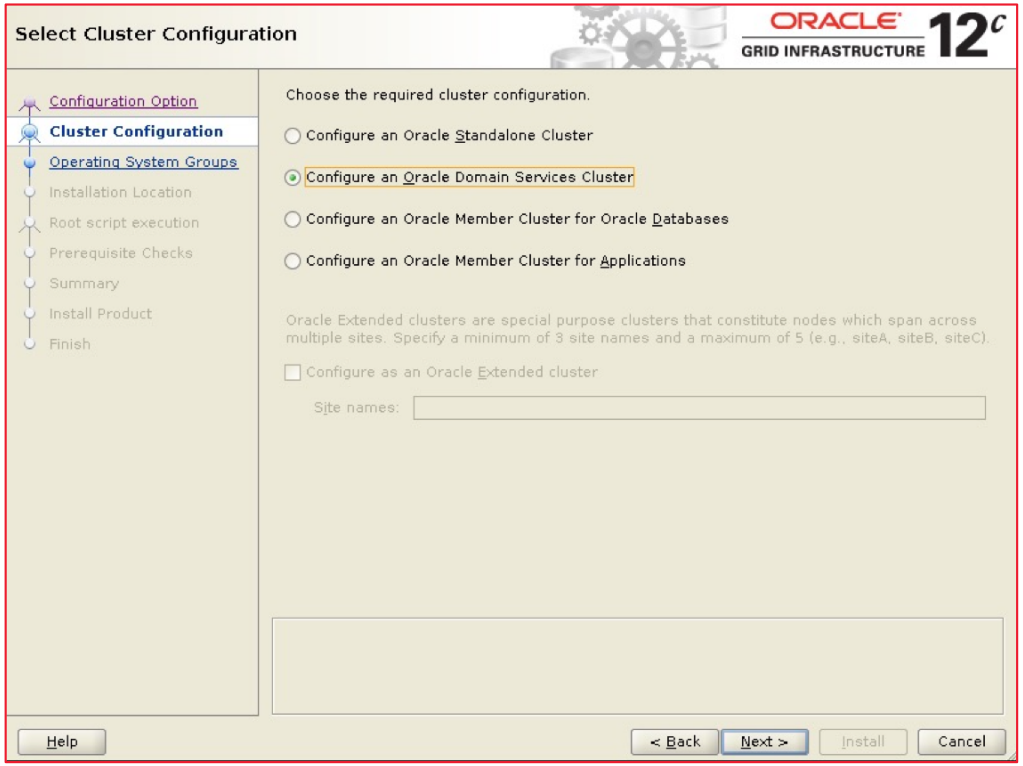

Installing a Domain Services Cluster (2)

- No major difference to installing a "standard standalone cluster"
- Diskspace requirements
  - Ca. 300 GB ASM storage (with RHP)
  - Ca. 100 GB in Oracle-Base
- RHP as an optional service has to selected during installation

Installing a Domain Services Cluster (3)

- "Additional" non-standard services after installation (with RHP)

```
[grid@dsc-node1 ~]$ crsctl stat res -t
Local Resources
--------------------------------------------------------------------------------
ora.helper
               ONLINE  ONLINE       dsc-node1      STABLE
               ONLINE  ONLINE       dsc-node2      IDLE,STABLE
ora.mgmt.ghchkpt.acfs
               ONLINE  ONLINE       dsc-node1      mounted on /mnt/oracle/rhpimages/chkbase,STABLE
               ONLINE  ONLINE       dsc-node2      mounted on /mnt/oracle/rhpimages/chkbase,STABLE
Cluster Resources
--------------------------------------------------------------------------------
ora.ioserver
      1        OFFLINE OFFLINE                     STABLE
      2        OFFLINE OFFLINE                     STABLE
      3        OFFLINE OFFLINE                     STABLE
ora.rhpserver
      1        ONLINE  ONLINE       dsc-node2      STABLE
```

Installing Member Clusters

Installing a Member Cluster (1) – Manifest File

- Before you install a member cluster, you have to create a Member Cluster Manifest File

```
[grid@dsc-node1 ~]$ crsctl create member_cluster_configuration -help
  crsctl create member_cluster_configuration <member_cluster_name> -file
<cluster_manifest_file>  -member_type <database|application>  [-version
<member_cluster_version>] [-domain_services [asm_storage <local|direct|indirect>][<rhp>]]
  where
     member_cluster_name    name of the new Member Cluster
     -file                  path of the Cluster Manifest File
     -member_type           type of member cluster to be created (database|application)
     -version               5 digit version of GI, if different from DSC
     -domain_services       services to be configured (asm_storage with local, direct, or
                            indirect access paths, and rhp)
     asm_storage            indicates the storage access path for database member clusters
                            local : storage is local to the cluster
                            direct or indirect : direct or indirect access to storage
     rhp                    generate credentials and configuration for an RHP client
```

Installing a Member Cluster (2) – Manifest File

- Example

```
[grid@dsc-node1 ~]$ crsctl create member_cluster_configuration dbcluster\
-file /home/grid/dbcluster.xml \
-member_type database -domain_services asm_storage local
--------------------------------------------------------------------------------
ASM GIMR TFA ACFS RHP GNS
================================================================================
 NO   YES   NO    NO   NO   NO
================================================================================

[grid@dsc-node1 ~]$ crsctl query  member_cluster_configuration dbcluster
      dbcluster      12.2.0.1.0 4a34243b9a4f7fdeff42e66c983b79dd GIMR
```

- Copy the manifest file to the new cluster afterwards

Installing a Database Member Cluster

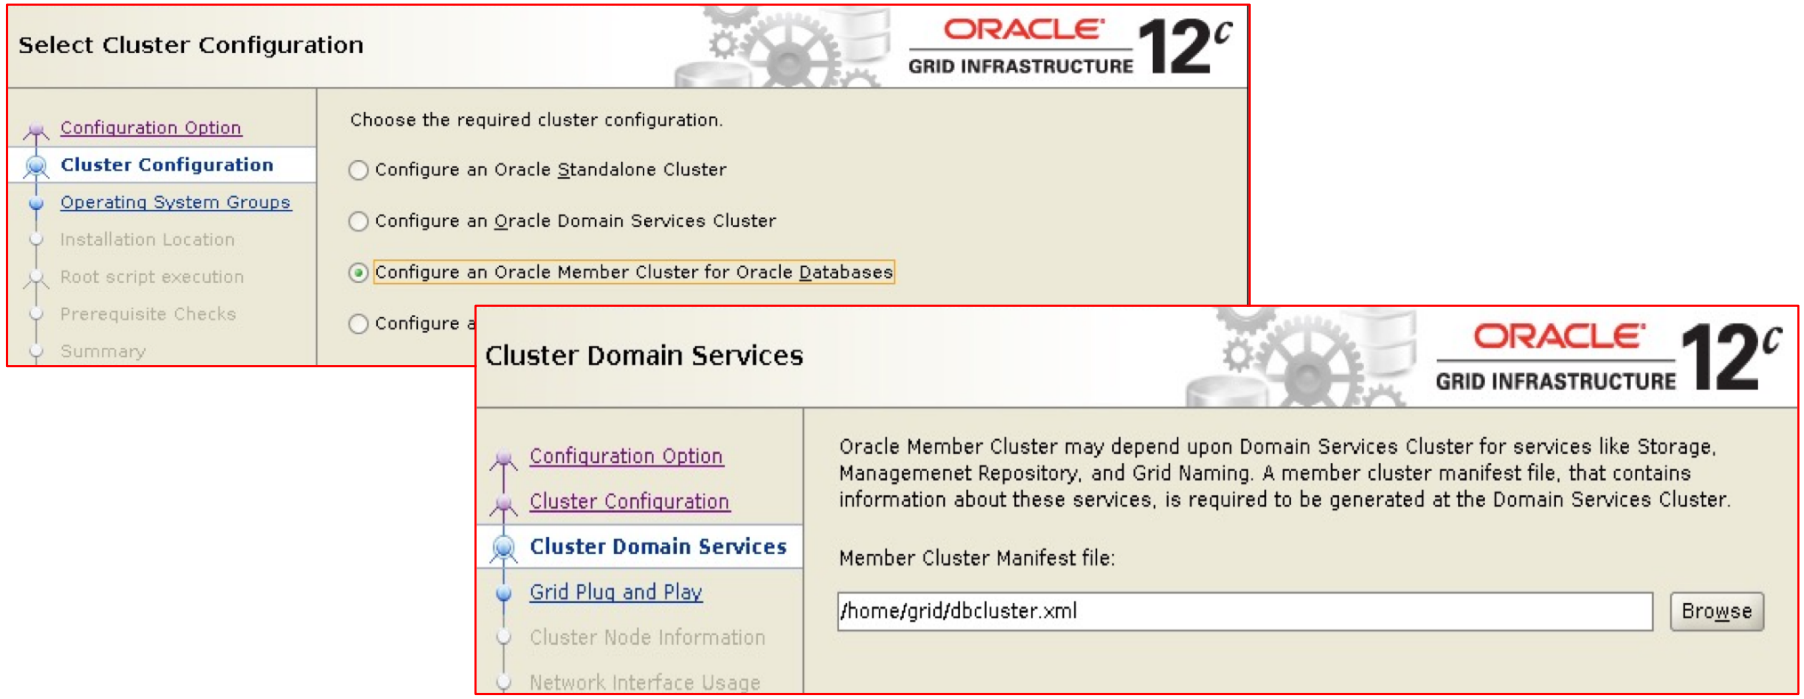

- In this case the database member cluster uses local storage (local ASM instance)

Installing an Application Member Cluster

- Small differences compared to installing a database member cluster
  - Virtual hostname can be defined for client access
  - Application Member Clusters store the OCR on the DSC (via ASM Service)

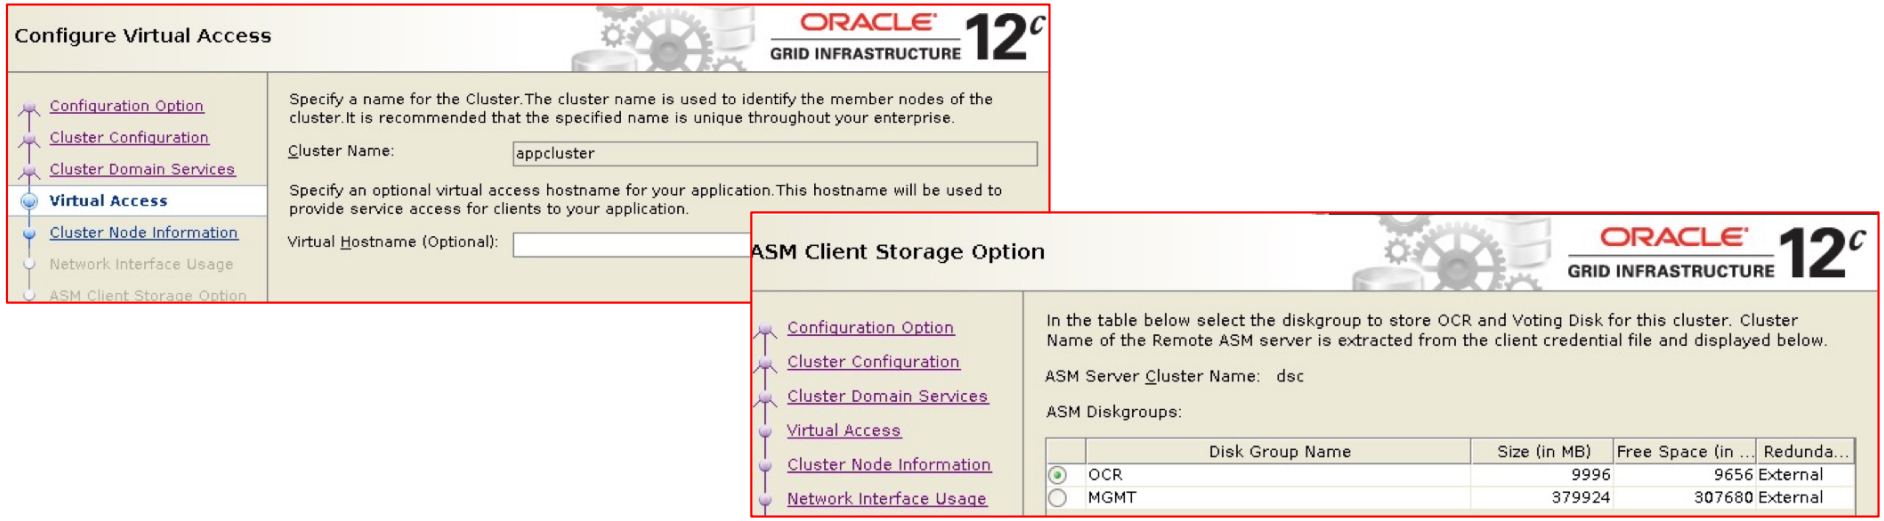

Cluster Domain Services

Available Cluster Domain Services

- Centralized GIMR
- Centralized TFA
- Storage Service
- Rapid Home Provisioning

Centralized GIMR

- There's a PDB for every member cluster in the GIMR (-MGMTDB):

```
[grid@dsc-node1 ~]$ sqlplus / as sysdba
SQL*Plus: Release 12.2.0.1.0 Production on Mon Mar 5 21:19:07 2018
Copyright (c) 1982, 2016, Oracle.  All rights reserved.
Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production
SQL> show pdbs
CON_ID    CON_NAME             OPEN MODE            RESTRICTED
--------- -------------        -------------------- ----------
2         PDB$SEED             READ ONLY            NO
3         GIMR_DSCREP_10       READ WRITE           NO
4         GIMR_CLUREP_20       READ WRITE           NO
5         GIMR_CLUREP_30       READ WRITE           NO
SQL> exit
```

- If the DSC is not available, GIMR data is temporarily stored locally and transferred later

Centralized Trace File Analyzer Collector (TFA)

- TFA
  - Daemon, independent on the Grid Infrastructure
  - Collects trace and log files and system information from all nodes into a cluster with a single command initiated on one cluster node

Why isn't that writen somewhere in the documentation?

- Does not work in (my) DSC configuration
- TFA service not "included" in member cluster manifest file
- No information in documentation how to enable it
- .. I created a Service Request …

Centralized Trace File Analyzer Collector (TFA)

- Result of the Service Request

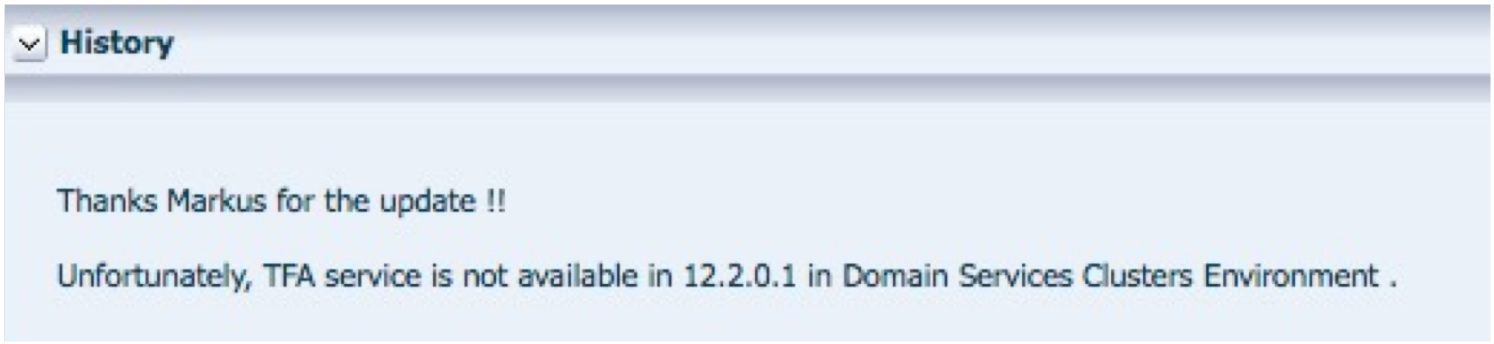

Why isn't that written somewhere in the documentation?

Storage Service (ASM service)

- Member cluster has a connection to the storage
- The ASM service manages the disk groups
- Member cluster retrieves storage information from DSC (via ASM network) and retrieves data from the ASM disks

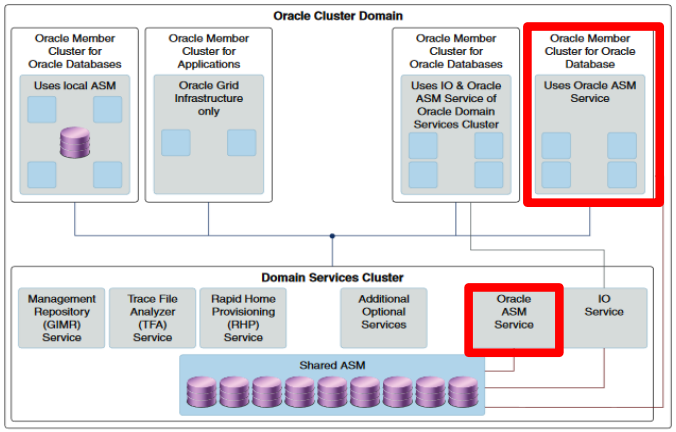

From Oracle ® Grid Infrastructure Installation and Upgrade Guide

```
[grid@dsc-node1 ~]$ srvctl status asm -verbose -detail
ASM is running on dsc-node2,dsc-node1
ASM is enabled.
ASM instance +ASM1 is running on node dsc-node1
Number of connected clients: 4
Client names: +APX1:+APX:dsc +IOS1:+IOS:dsc -MGMTDB:_mgmtdb:dsc dsc-
node1.markusdba.net:_OCR:dsc
[..]
Detailed state on node dsc-node1: Started
[..]
```

Storage Service (IO Service)

- The I/O server provides network-based connectivity to ASM file for database instances on nodes that do not have direct access to ASM managed disks
- On the storage cluster, an I/O server instance opens up network ports to which clients send their I/O
- The I/O server instance receives data packets from the client and performs the appropriate IO to ASM disks

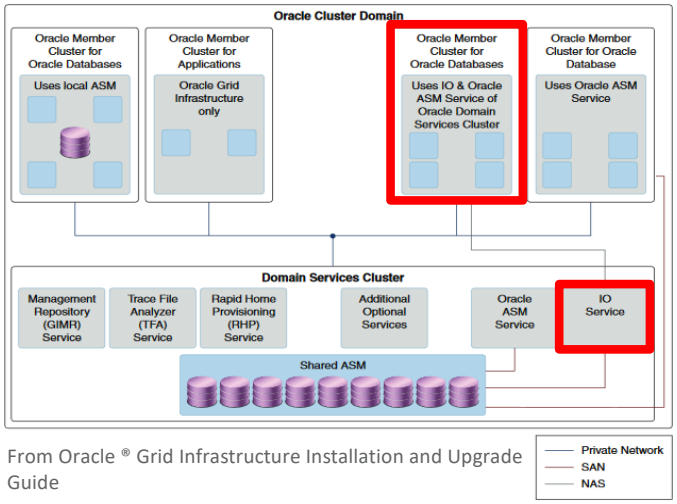

From Oracle ® Grid Infrastructure Installation and Upgrade Guide

```
[grid@dsc-node1 ~]$ srvctl status ioserver -verbose -detail
ASM I/O Server is running on dsc-node2,dsc-node1
ASM I/O Server is enabled.
ASM I/O server instance +IOS1 running on node dsc-node1 is connected to ASM instance
Number of connected clients: 1
Client names: RAC1:RAC_SITE1:a38f8b271ff9efbcbf367068f36c0c0e
ASM I/O server instance +IOS2 running on node dsc-node2 is connected to ASM instance
Number of connected clients: 1
Client names: RAC2:RAC_SITE1:a38f8b271ff9efbcbf367068f36c0c0e
```

Rapid Home Provisioning Service

- The DSC can be configured as a Rapid Home Provisioning (RHP) server
- Features:
  - Provisioning Oracle Software (GI + RDBMS) for member clusters
  - Patching Oracle Software on member clusters
  - Upgrading Oracle databases on member clusters
- Licence:
  "Lifecycle Management Pack" on targets required

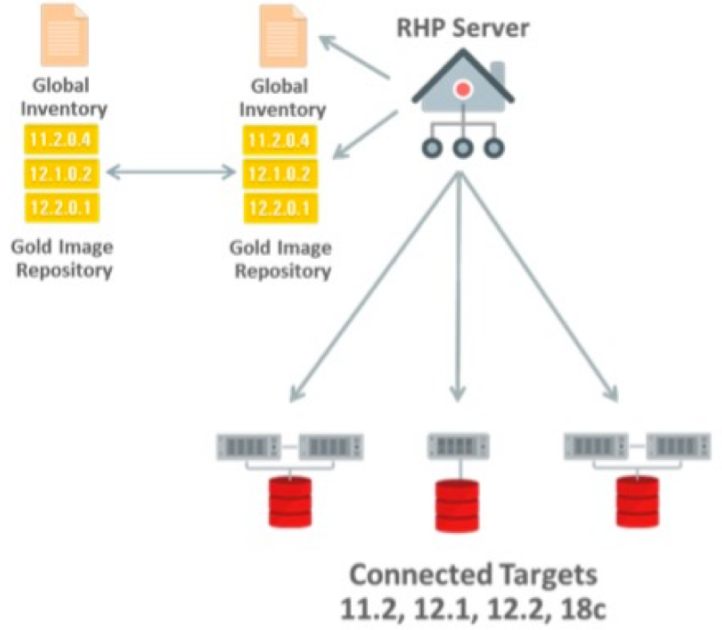

Source: Oracle-Whitepaper on RHP 18c

trivadis makes IT easier.

Summary & Outlook

Summary

- ☺ Centralized Management for a group of clusters (> 8)
- ☺ Storage Optimization when using shared storage
- 😐 Install Enterprise Manager Cloud Control on the DSC – and you'll get the "one cluster to rule them all"
- ☹ TFA not working, poorly documented
- ☹ Not very much documentation available
- ☹ The documentation is partly misleading
- ☹ Oracle 12.2: Existing Clusters cannot be converted into member clusters

Domain Services Cluster – New Features in Oracle 18c

- According to the docs ☺
  - Conversion of a standalone cluster to a member cluster is possible
  - Storage Conversion for member clusters (e.g. from direct ASM to indirect ASM)
  - ACFS remote service

Further Information & References

- Oracle 12.2 - Grid Infrastructure - Installation + Upgrade (Chapter 8.5, 9.3)
  https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwlin/index.html
- Oracle Whitepaper "Oracle Grid Infrastructure – Cluster Domains"
  http://www.oracle.com/ocom/groups/public/@otn/documents/webcontent/3633615.pdf
- http://www.hhutzler.de/blog/install-12-2-oracle-domain-cluster/

trivadis
makes IT easier.

Questions and Answers

Markus Flechtner
Principal Consultant

Phone +49 211 5866 64725
Markus.Flechtner@Trivadis.com
@markusdba https://markusdba.net

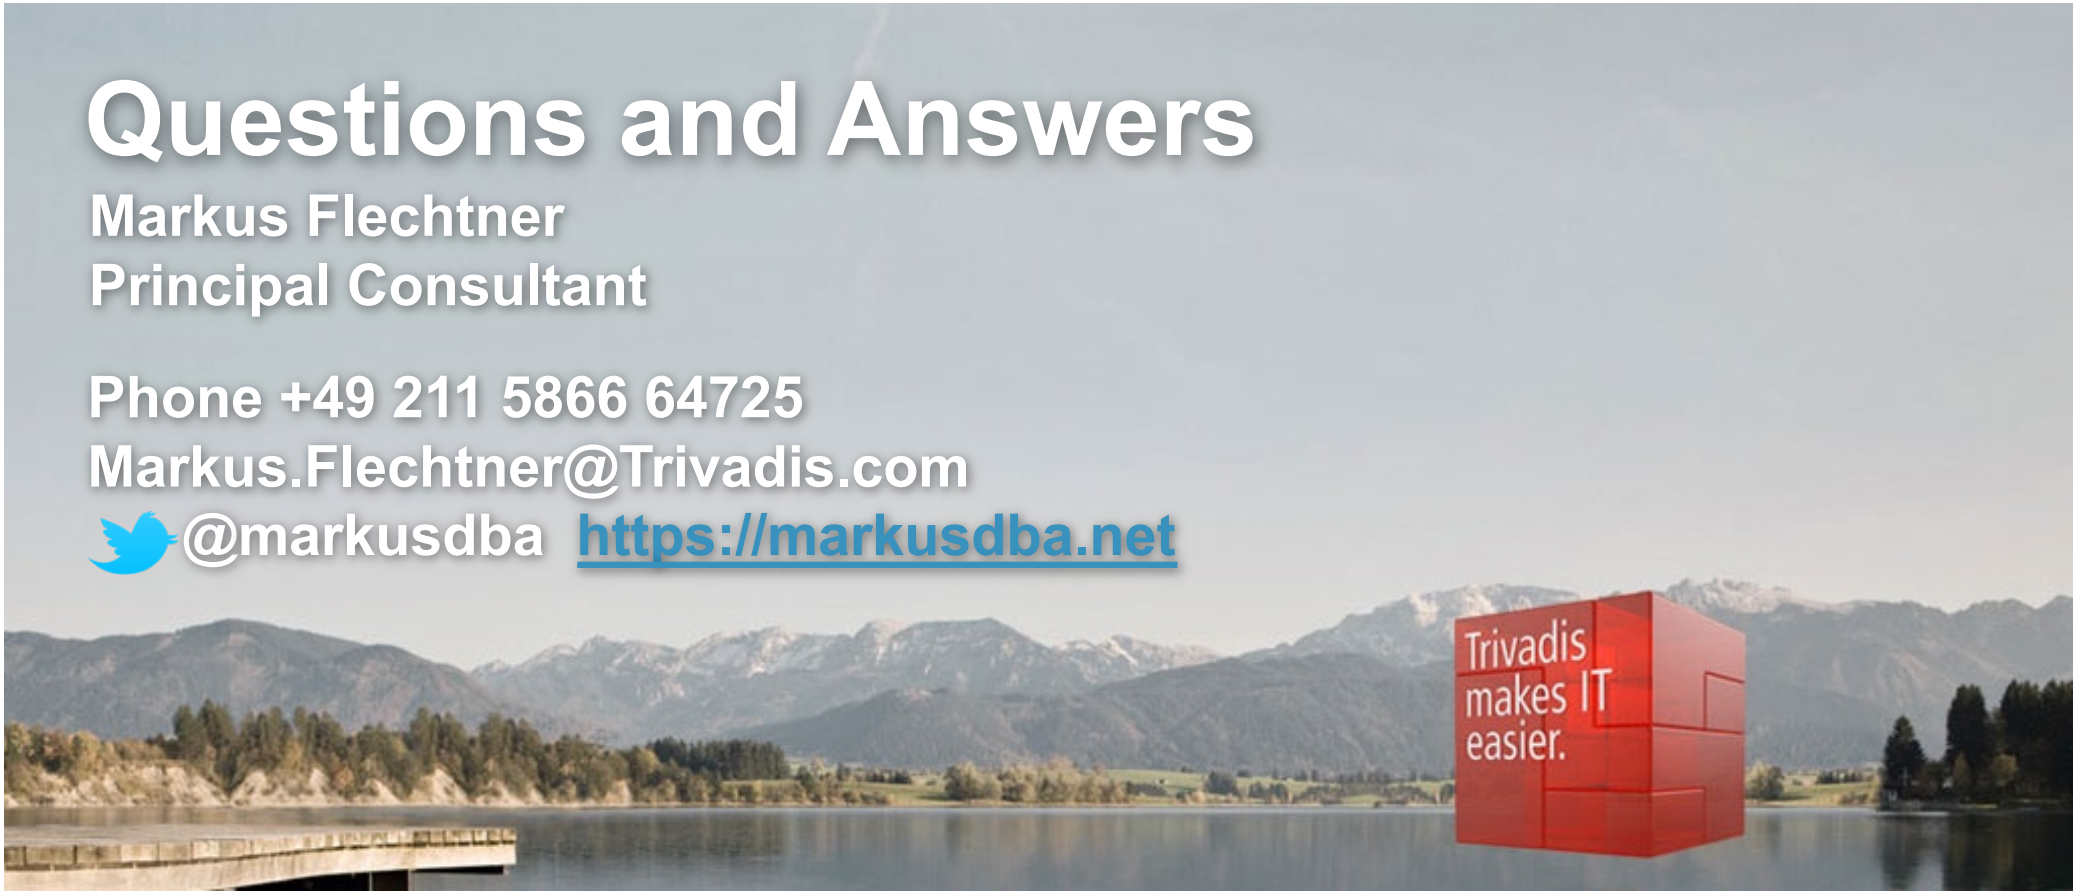

Download the slides from http://www.slideshare.net/markusdba

Please don‘t forget the session evaluation – Thank you!

trivadis
makes IT easier.