

Oracle Database Appliance Technical Architecture

ORACLE

February 2023

Copyright © 2022, 2023 Oracle and/or its affiliates

Copyright © 2022, 2023 Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

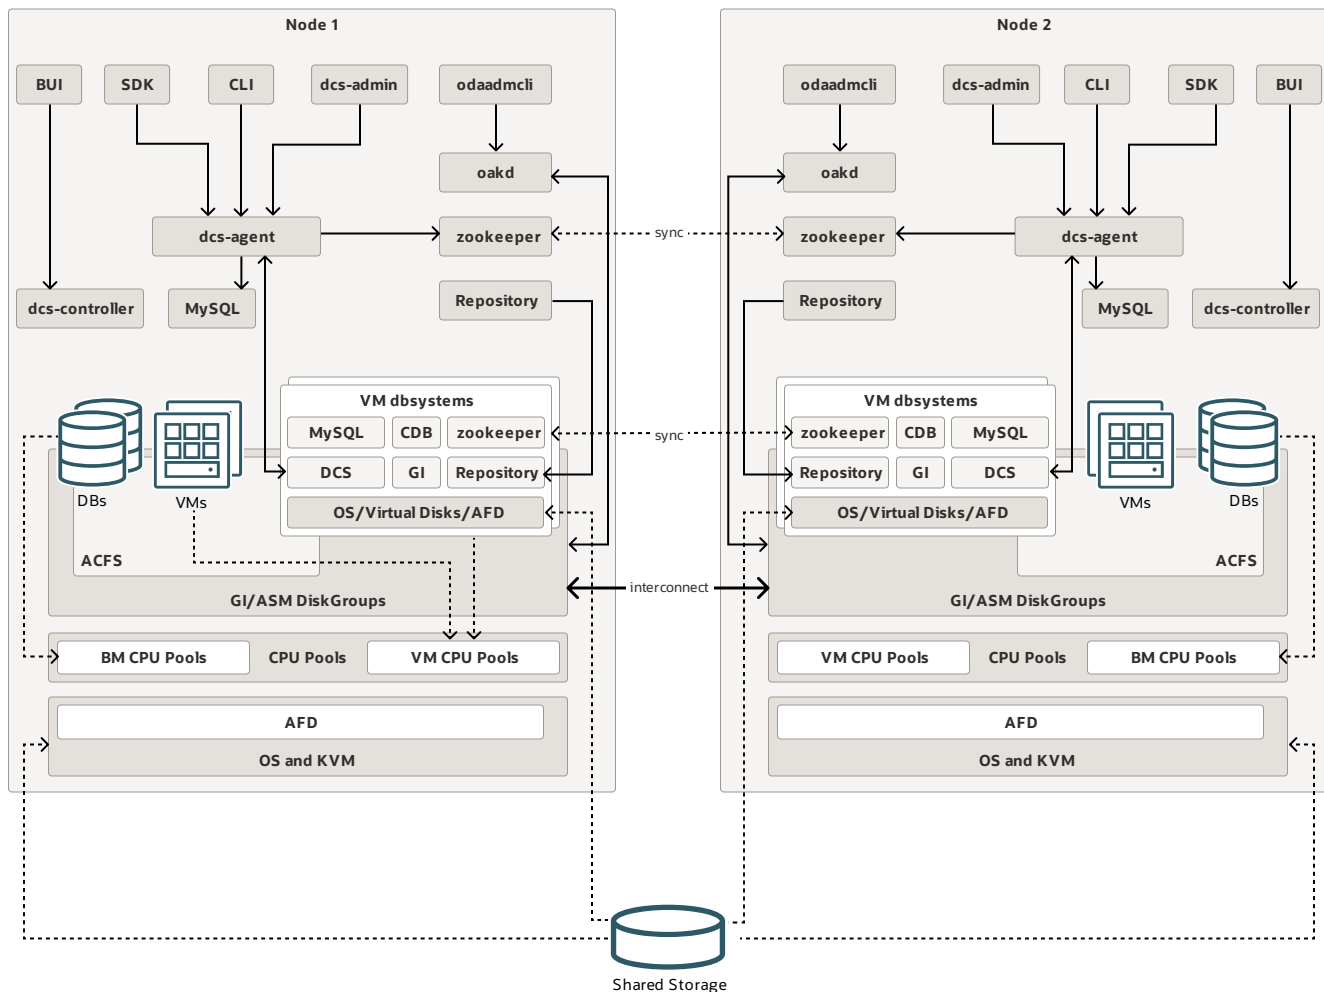
Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Table of Contents

Table of Contents	2
Oracle Database Appliance High-Availability Bare Metal System	3
Notes	3
Related Resources	5
Oracle Database Appliance ODA CLI	6
Notes	6
Related Resources	7
Oracle Database Appliance Browser User Interface (BUI)	8
Notes	8
Related Resources	9
Oracle Database Appliance KVM and DB System	10
Notes	10
Related Resources	11
Oracle Database Appliance Backup and Recovery	12
Notes	12
Related Resources	13
Oracle Database Appliance Oracle Data Guard	14
Notes	14
Related Resources	15
Oracle Database Appliance Network	16
Notes	16
Related Resources	17
Oracle Database Appliance High Availability Storage	18
Notes	18
Related Resources	19
Oracle Database Appliance Disk Partnership on High-Availability Storage	20
Notes	20
Related Resources	21
Oracle Database Appliance Single-Node Storage	22
Notes	22
Related Resources	22
Oracle Database Appliance Disk Partnership on Single-Node Storage	24
Notes	24
Related Resources	25
Oracle Database Appliance Multi-User Access Architecture	26
Notes	26
Related Resources	27

Oracle Database Appliance High-Availability Bare Metal System



Notes

Oracle Database Appliance saves time and money by simplifying deployment, maintenance, and support of database solutions for organizations of every size. Oracle Database Appliance hardware models are optimized to run Oracle Database Standard Edition and Enterprise Edition.

Following are the components of an Oracle Database Appliance deployment:

UI (User Interface): A web-based application to administer and manage an appliance. It interacts with the DCS agent application to perform tasks on the appliance.

CLI: Command Line Interface to manage Oracle Database Appliance. It interacts with the DCS agent application to perform tasks on the appliance.

SDK (Software Development Kit): Oracle Database Appliance Java SDK works with the Oracle Database Appliance software to interact with the DCS agent. You can develop custom client application using REST APIs to manage the appliance.

DCS agent: The DCS agent is a REST application that performs various operations on the appliance, such as provisioning, database lifecycle management, backup and recovery, and storage management.

DCS admin: This application performs lifecycle management operations on the DCS agent, DCS controller, ODA CLI, and other components.

oakd: This process manages the storage components of Oracle Database Appliance.

odaadmcli: The CLI that manages the storage components.

dcs-controller: The application on which the Browser User Interface (BUI) is deployed.

Zookeeper: A centralized service that maintains configuration information, and provides distributed synchronization and group services.

repository: The file system that holds the images and clone files used by both bare metal Oracle Database Appliance system and VM DB system.

MySQL: The metadata store used by DCS agent.

DBs: Oracle Databases managed by Oracle Database Appliance software.

VMs: The KVM guest machines, also known as application VMs that are managed partially by Oracle Database Appliance. The Oracle Database Appliance software manages basic operation of the VM, such as start, stop, create, delete, and failover. Oracle Database Appliance does not manage the applications in the VM.

ACFS: Oracle ACFS file system used for VM storage, Oracle Database homes, and Oracle Database storage.

GI/ASM disk groups: Oracle Grid Infrastructure and ASM disk groups, such as DATA and RECO.

BM CPU Pools: Use the BM CPU pool to ensure that the database on a bare metal host runs exclusively on the CPU lists specified in the CPU pool. All Oracle processes are bound to the CPUs in this CPU pool and run only on these CPUs. The bare metal CPU pool is assigned to the database by updating the init.ora parameter `processor_group_name`.

VM CPU Pools: Use the VM CPU pools to cage and manage CPU resource allocations to workloads for virtual machines. Workloads are isolated by creating CPU pools and assigning virtual machines

to a specific CPU pool. When you pin a virtual machine to a CPU pool, you ensure that the virtual machine uses CPUs in only that CPU pool.

AFD: Oracle ASM Filter Driver, a kernel module that resides in the I/O path of the Oracle ASM disks. Oracle ASM uses the filter driver to validate write I/O requests to Oracle ASM disks.

OS and KVM: This represents the operating system and KVM Hypervisor.


VM dbsystems: The fully-managed KVM DB System. The DB system contains its own DCS agent, MySQL metadata, zookeeper, Oracle Grid Infrastructure, and Oracle Database. There is no Oracle ASM instance on DB system. The DB system repository is an NFS mount from the repository on the bare metal system. Oracle Database software fully automates the creation of the DB system, including the VM, Oracle Grid Infrastructure, and database creation, patching of the system and database, and deletion of the system.

Shared Storage: The HDD and SSD storage are shared by both nodes on an Oracle Database Appliance high-availability model. On an Oracle Database Appliance single-node system, NVME drives provide the storages. These storages are used by Oracle ASM. The shared storages are passed to the VM DB systems using the virt-scsi protocol.

Related Resources

- [Oracle Database Appliance Documentation](#)
- [Oracle Database Appliance X9-2-HA Hardware Model](#)
- [Oracle Database Appliance X9-2S/X9-2L Hardware Model](#)

Oracle Database Appliance ODACLI



Notes

Use the Oracle Database Appliance command line interface (ODA CLI) to perform operations on the DCS agent.

Oracle Database Appliance establishes communication to the DCS agent is established using REST over HTTPS with the local DCS agent. The credentials and tokens are stored on an Oracle PKI wallet and the public certificates are stored on a Java (PKCS12) TrustStore. Once the secured communication is established correctly, the operation is invoked and the results are printed to the standard output.


Use the `--json/-j` option for any command to print the output in JSON format, which is in a machine-readable form. For advanced interaction with the DCS agent, use Oracle Database Appliance REST APIs. In case of any issue during the operation, either on server or client side, use the respective logs to view the issues.

To query available operations, use `odacli -h`. To view available operations on certain resource types, use `odacli resource -h`.

Related Resources

- [Oracle Database Appliance Command-Line Interface](#)
- [Oracle Database Appliance Command-Line Interface for KVM](#)

Oracle Database Appliance Browser User Interface (BUI)



Notes

The Oracle Database Appliance Browser User Interface (BUI) is a web-based application to administer and manage the appliance. The BUI uses Oracle Javascript Extension Toolkit (JET) framework (<http://oraclejet.org>) which supports the Model-View-ViewModel (MVVM) architectural design pattern.

The Oracle JET framework provides a set of UI components with built-in accessibility support. To support the MVVM design, Oracle JET is built upon a modular framework that includes a collection of third-party libraries and Oracle-provided files, scripts, and libraries.

DCS Controller: When you access the BUI URL, the browser sends an HTTP request to the DCS controller, to send a copy of the web application. The request consists of a line indicating a method such as GET or POST, a Uniform Resource Identifier (URI) indicating which resource is being requested, and an HTTP protocol version.

The index.html page is the first page to download and becomes the blueprint by which the web application is constructed. As the browser reads the HTML code, it sends out more requests to the DCS controller to send more pieces of the page. The browser caches static assets, that is, the parts of a website that do not change from visit to visit, such as images, HTML, CSS, and JavaScript. Once the request is processed successfully, the login page is rendered on the browser.


The BUI code that runs on the client side makes a REST API call to the DCS agent for all requests and processes the results that are sent back. The DCS agent is managed outside of the DCS controller and runs on a separate web server. Each request to the DCS agent is made as an asynchronous call. The DCS agent generates a response based on the REST API call and includes the HTTP protocol version and the HTTP status code. The agent sends back a response to the browser for the request made. The client-side code then processes the result sent back by the DCS agent and presents it on the BUI. The status code and reason phrase tell the client how the server handled the request.

Generally, the status code 200 is returned, indicating that the request was handled successfully and the body data contains the requested item. Other result codes indicate HTTP errors, such as 404 Not Found.

Related Resources

- [Managing Oracle Databases](#)
- [Patching Oracle Database Appliance](#)
- [About the Browser User Interface](#)

Oracle Database Appliance KVM and DB System



Notes

Following are the components of an Oracle Database Appliance KVM and DB system:

VM Storage: A VM storage is a central location for storing resources that are essential to create and manage virtual machines. These resources include ISO files, virtual machine configuration files, and virtual disks. Virtual storage is configured on Oracle ACFS. Oracle Database Appliance KVM uses virtual storage to optimize available disk space usage in the environment, and for easy reallocation of virtual machines if a physical server fails. Create the VM storage first before creating the VM.

vdisk: Virtual disks provide additional storage options for virtual machines by enabling you to attach

additional block storage to your virtual machines. Similarly, you can detach the disk if you no longer need the additional space.

VM CPU Pool: VM CPU Pools include VM and DB system CPU pools. VM CPU pools are for application VMs, and DB system CPU pools are for DB systems. Use the VM CPU pools to cage and manage CPU resource for virtual machines. Workloads are isolated by creating CPU pools and assigning virtual machines to a specific CPU pool. When you pin a virtual machine to a CPU pool, you ensure that the virtual machine uses CPUs from only that CPU pool. CPU pools cannot share CPUs. But one CPU pool can be assigned to multiple VMs or DB systems.

VMs: The KVM guest machines or application VMs on Oracle Database Appliance. The Oracle Database Appliance software manages basic operations of the VM such as start, stop, create, delete, and failover, but it does not manage applications in the VM. Oracle Database Appliance supports all VM operating systems flavors that Oracle KVM supports.

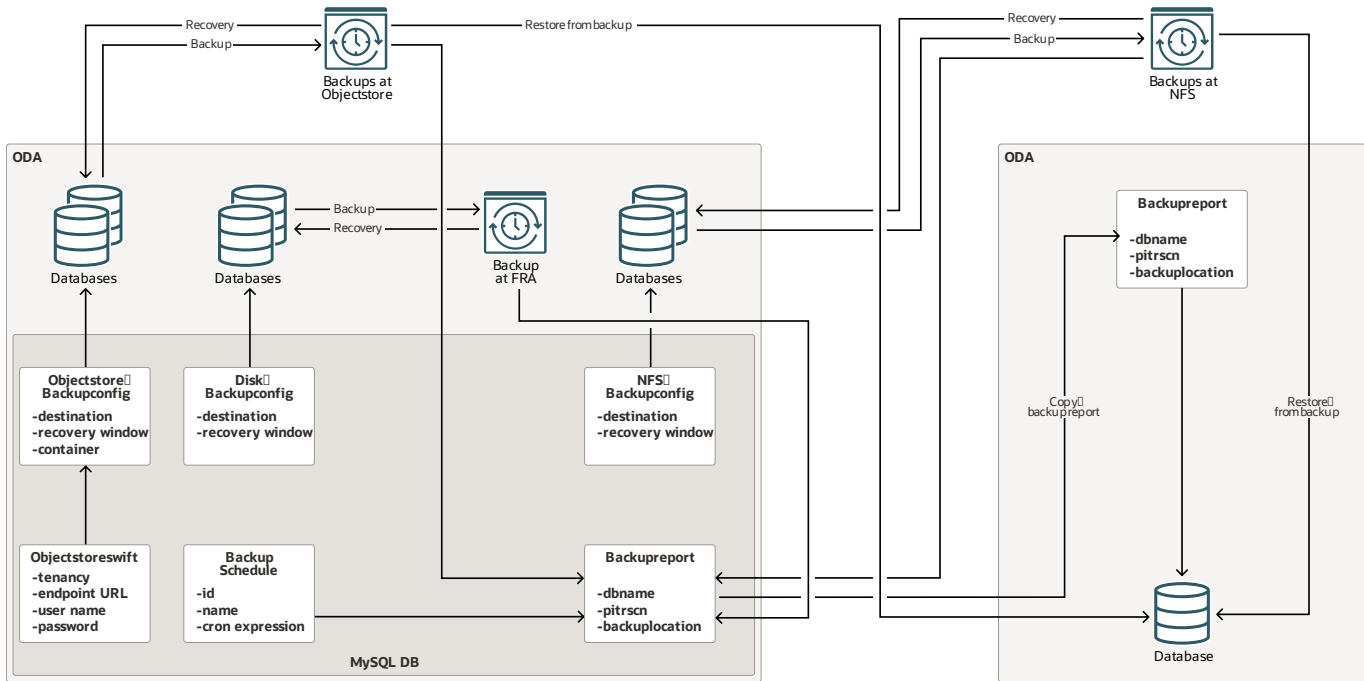
VM DB system: The fully-managed KVM DB System. Oracle Database Appliance software fully automates the creation of the DB system including the VM, Oracle Grid Infrastructure and Oracle Database creation, patching of the system and database, and deletion of the system. KVM database systems enable hard partitioning for Oracle Database licensing, where each KVM database system has its own CPU pool that is automatically assigned during KVM database system creation.

vnetwork: Oracle Database Appliance KVM virtual network supports two type of networks, bridged and bridged-vlan. In a bridged network, a Linux bridge is created and the network interface or bond interface, is attached to the bridge. In a bridged-vlan network, a Linux bridge is created and a VLAN is attached to the bridge. VLAN can be created on all available public interfaces, including the interface where public network is already configured.

Related Resources

- [Managing an Oracle Database Appliance KVM Deployment](#)
- [Migrating from Virtualized Platform to KVM on Oracle Database Appliance](#)
- [Oracle Database Appliance Command-Line Interface for KVM](#)
- [Managing CPU Pools on Oracle Database Appliance](#)

Oracle Database Appliance Backup and Recovery



Notes

Oracle Database Appliance supports three types of backup configuration:

- Backup to an Oracle Fast Recovery Area (FRA) disk (Internal FRA)
- Backup to Oracle Cloud Infrastructure Object Storage (Oracle Object Storage)
- Backup to a Network File System (NFS) location (External FRA)

You can use the backup, restore, and recovery features with databases that have backup policy

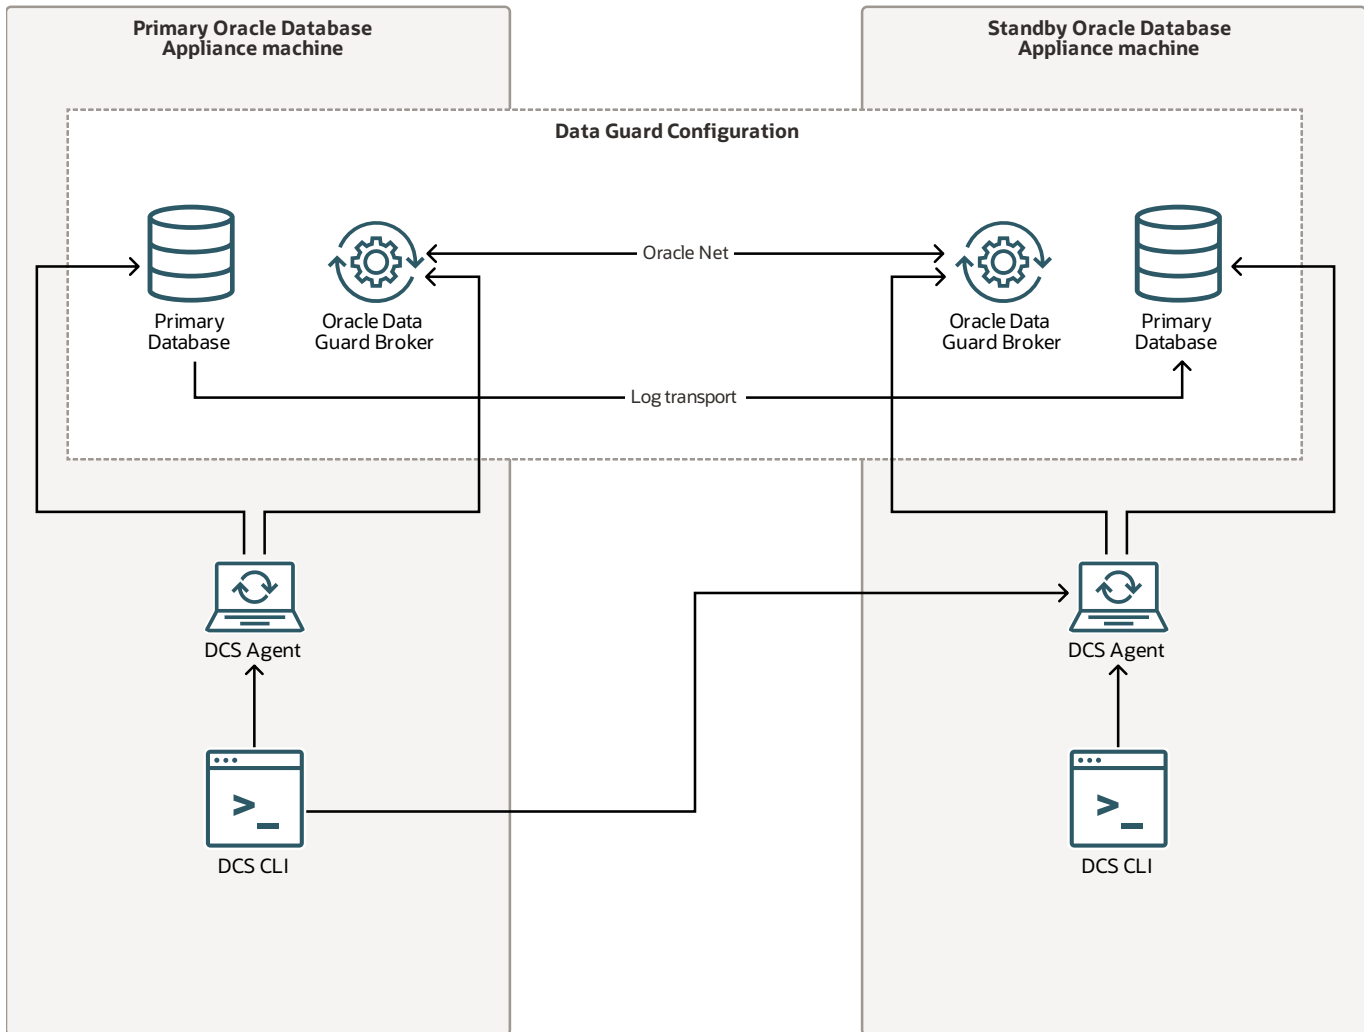
attached to them in the system. You can choose to manage your backups, and set up your backup, restore, and recovery options.

ObjectStoreSwift object must be created before creating the Object Store backup configuration. When a database is modified with either Object Store or Disk or NFS backup configuration, Database backup and archive log backup schedules are created for that database. The database can be restored only from Object Store and NFS backups. Corresponding backup report must be copied manually to the system where the database needs to be restored.

Related Resources

- [Backup, Restore and Recover Databases](#)
- [odacli Backup and Recovery Commands](#)

Oracle Database Appliance Oracle Data Guard



Notes

Oracle Database Appliance provides client interface through ODA CLI commands for easy configuration and management of Oracle Data Guard for high availability, data protection, and disaster recovery. Oracle Data Guard Broker is a distributed management framework that automates the creation, maintenance, and monitoring of Oracle Data Guard configurations.

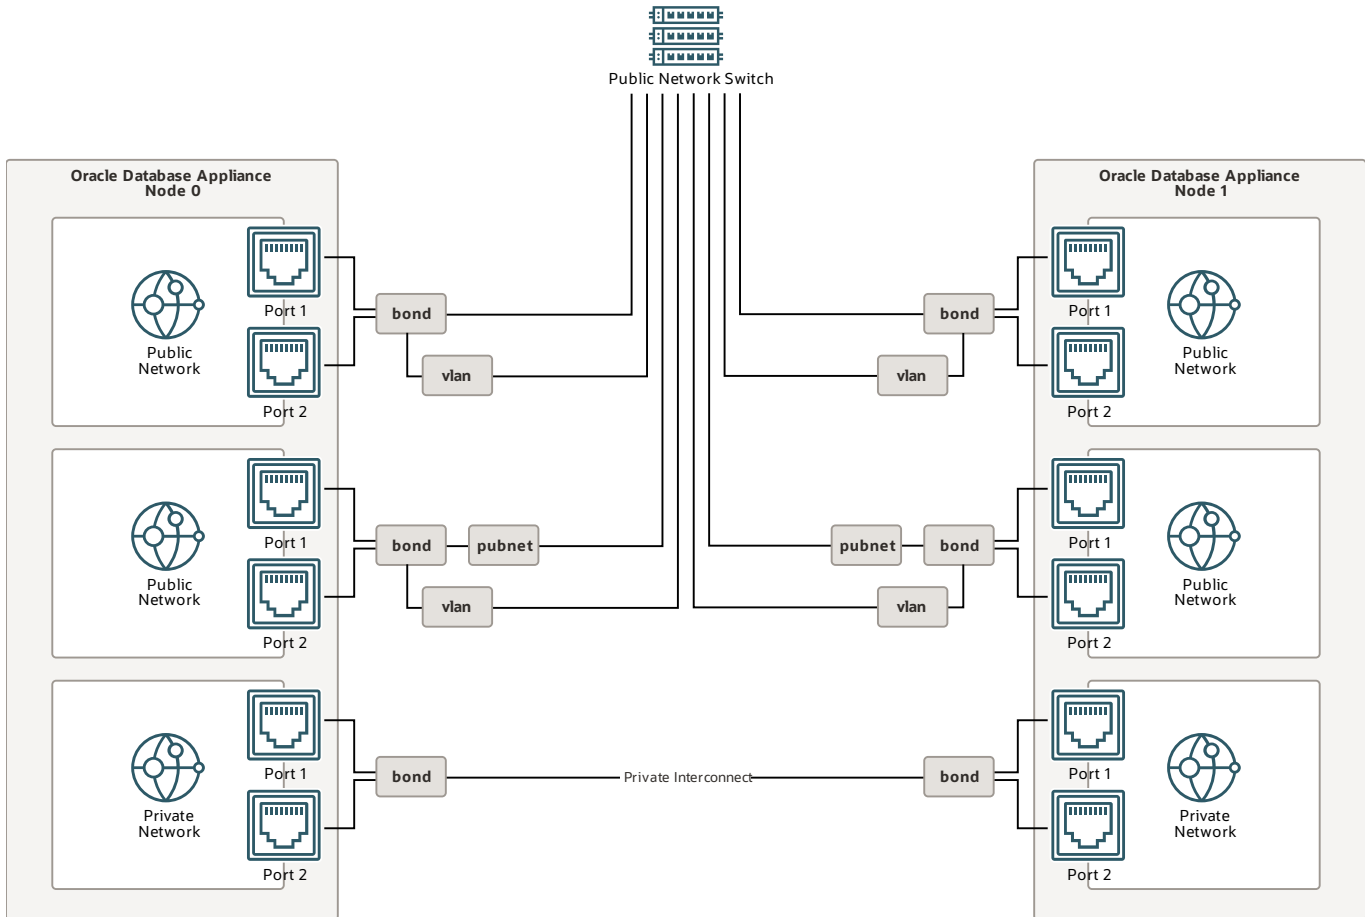
An Oracle Data Guard configuration set up with ODA CLI commands consists of one primary database and one standby database. The primary system is the Oracle Database Appliance system that hosts the primary database. The standby system is the Oracle Database Appliance system that hosts the standby database.

Each system runs its own ODA CLI which can call operations on the DCS agent. The primary DCS CLI can call operations on primary and standby DCS agent for the `odacli configure-dataguard` command. It orchestrates a sequence of jobs on two hosts to complete Oracle Data Guard configuration on Oracle Database Appliance. The DCS agent connects with Oracle Data Guard Broker and databases to manage Oracle Data Guard configuration on Oracle Database Appliance.

Related Resources

- [Using Oracle Data Guard for Disaster Management and Recovery on Oracle Database Appliance](#)
- [odacli Oracle Data Guard Commands](#)

Oracle Database Appliance Network



Notes

Oracle Database Appliance has a fixed network configuration on Oracle Database Appliance X7-2 and earlier. The public network ports are from the onboard network card. All network bonding setup are configured during imaging.

On Oracle Database Appliance X8-2 and later, the network set up is not fixed. The public network ports are from PCI network card. Each PCI slot can be installed with either two ports 25G or four ports 10G network cards. Additional network cards can be added into the system after the system deployment.

Public network: Public bonding interfaces are created during imaging of the system.

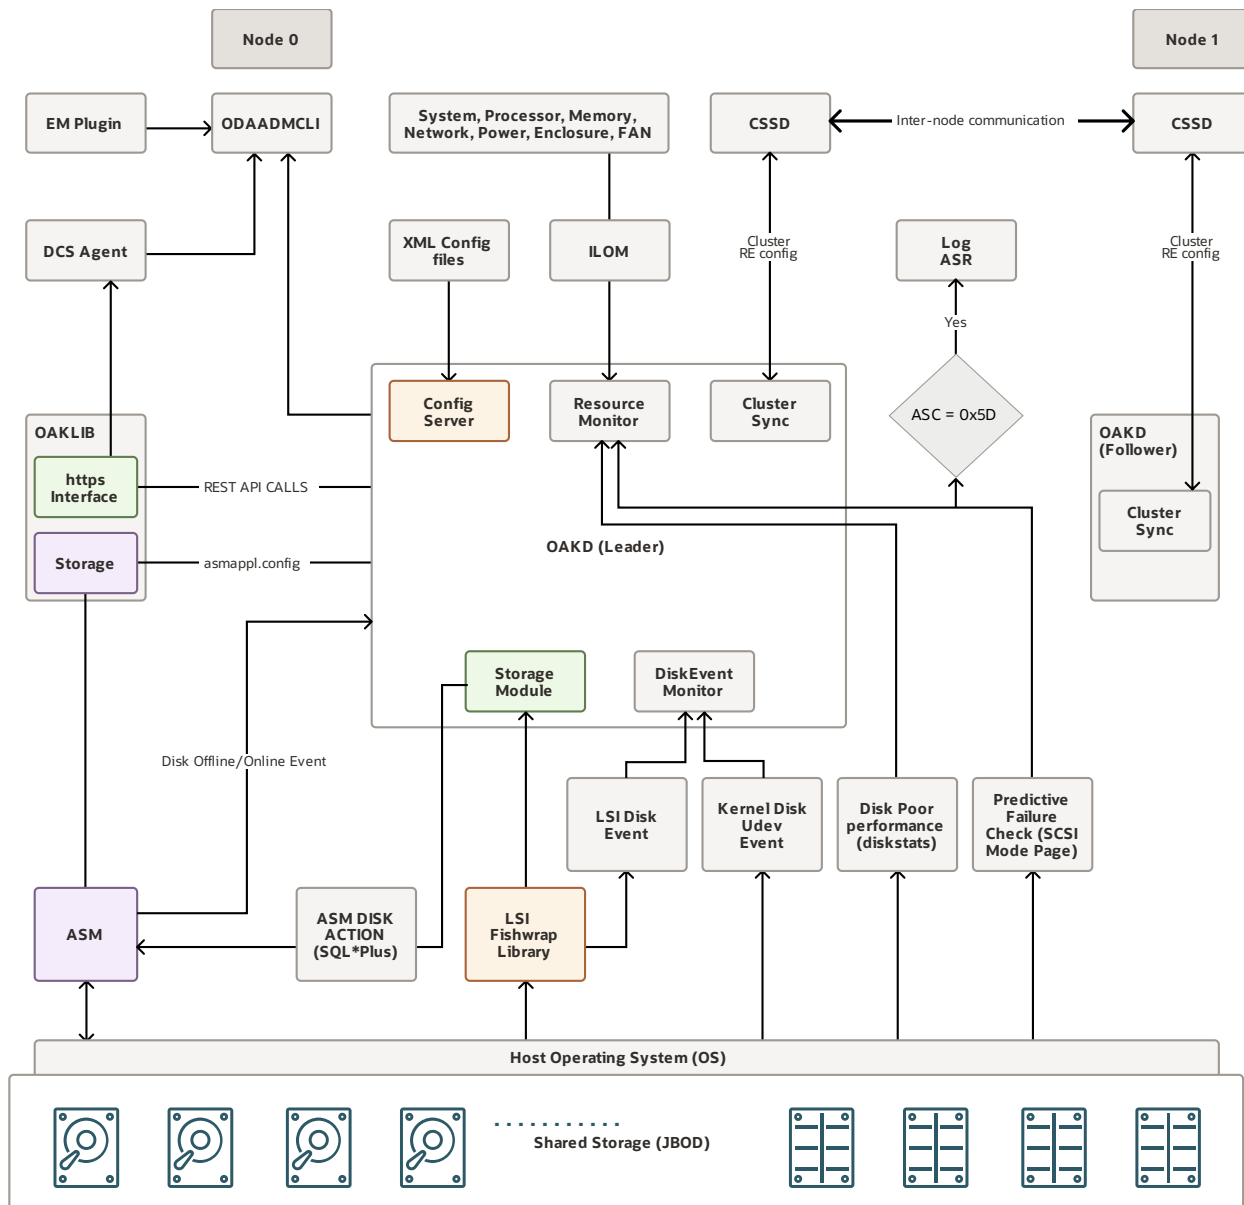
Private network: For Oracle Database Appliance high-availability systems, there is a private network card for the interconnection between two nodes. This network card can be Ethernet ports or InfiniBand ports on Oracle Database Appliance X6-2 and X5-2. A bonding interface is created over the two ports for communication between the two nodes. When the node number for the system is confirmed, the private IP address for the node is assigned automatically.

For DB system support on the Oracle Database Appliance system, create a network bridge on the public network so that it can be shared by multiple KVM systems. During the deployment of the system, a network bridge (pubnet) is created on the selected public network interface. You can use the pubnet to connect to the appliance.

Related Resources

- [Managing Networks](#)
- [Readying Oracle Database Appliance for Deployment](#)

Oracle Database Appliance High Availability Storage



Notes

OAKD is the storage management system for Oracle Database Appliance. It provides storage discovery, auto configuration of storage, and disk management.

It has three modes: non-cluster, upgrade, and foreground. OAKD starts in non-cluster when Oracle Grid Infrastructure stack is not yet installed. After installation of Oracle Grid Infrastructure, OAKD starts in cluster mode.

Use `odaadmcli` commands to query information from OAKD or provide information to OAKD. It has various modules such as server configuration, storage management, resource monitoring, cluster

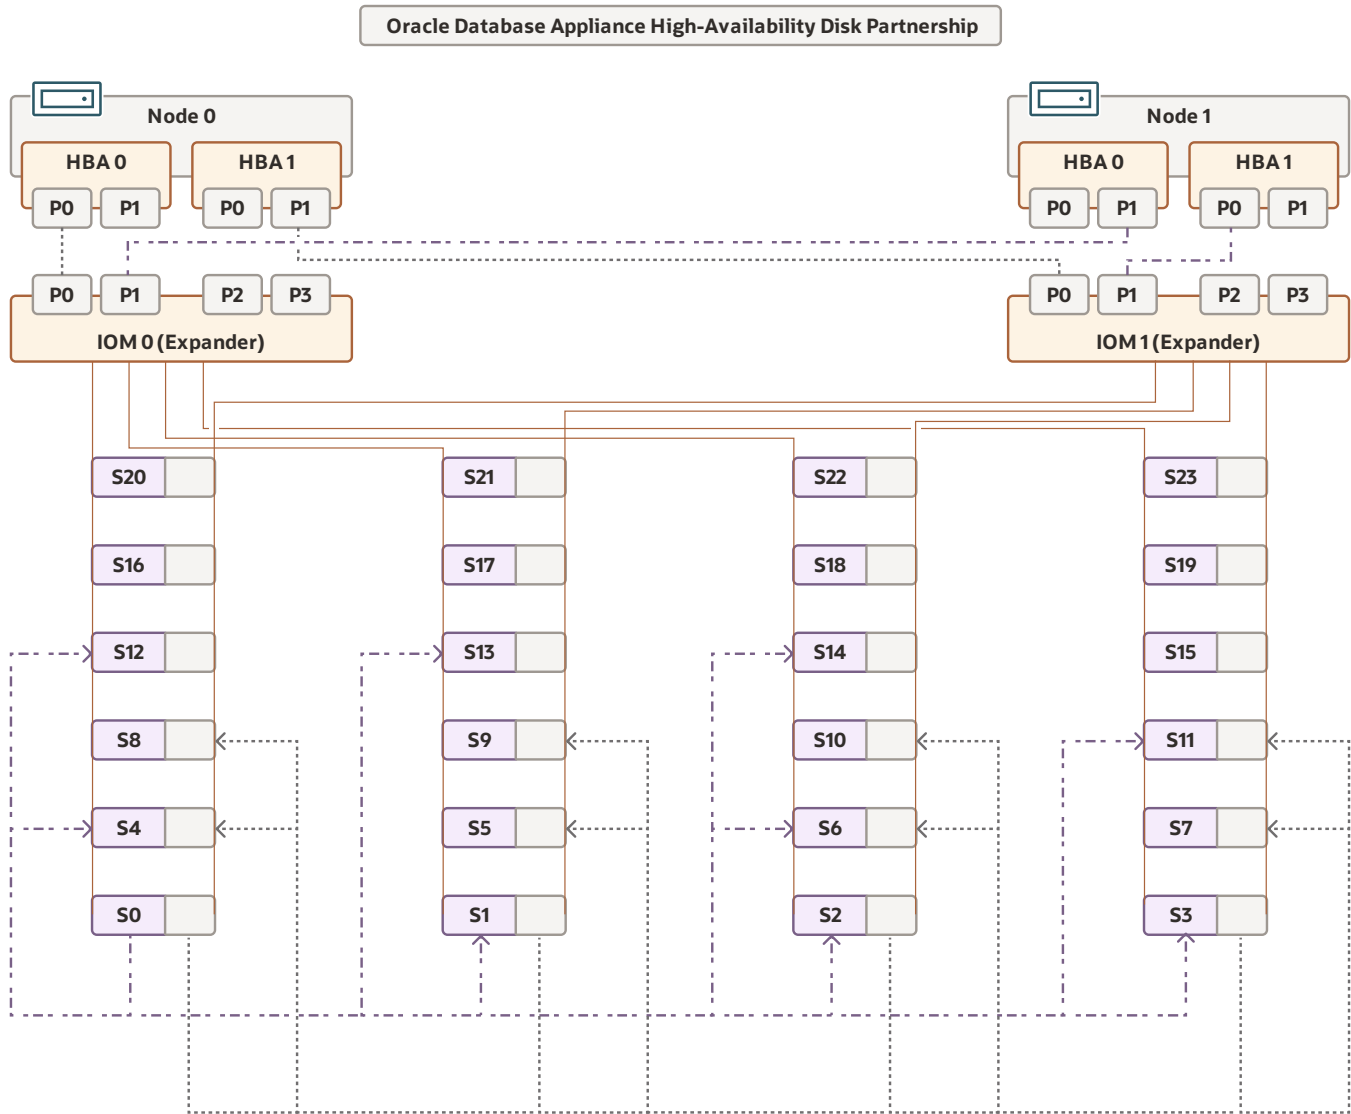
synchronization, and disk events monitoring. When the appliance is started, OAKD uses LSI and Fishwrap library to discover storage. Configuration servers read XML configuration files describing various properties of various OAKD resources. OAKD provides storage configuration information contained in asmappl.config to Oracle ASM through oaklib library. oaklib also has an HTTPS interface used by OAKD to make REST API call to the DCS agent. The resource management module monitors various hardware components at regular intervals. Cluster communication modules waits for cluster reconfiguration events to take appropriate actions. OAKD monitors LSI and UDEV events for disk operations such as insert or removal and completes the action in OAKD and Oracle ASM.

OAKD monitors for predictive failure information on disk. OAKD also reads /proc/diskstats periodically for disk performance matrix. OAKD monitor for Oracle ASM action events such as disk offline or online events. It also performs auto online operation for transient failures which can cause Oracle ASM to offline a disk. Other modules such as DCS and Oracle Enterprise Manager plugin for Oracle Database Appliance use the odaadmcli interface to query storage information.

Related Resources

- [Managing Storage](#)
- [odacli Storage Commands](#)
- [Readying Oracle Database Appliance for Deployment](#)

Oracle Database Appliance Disk Partnership on High-Availability Storage



Notes

Oracle ASM supports fixed partnership for Oracle ASM disks. Oracle ASM performs splitting and mirroring of data across all disk partners. Oracle ASM supports either four or eight disk partners for Oracle Database Appliance high-availability. Oracle Database Appliance X6-2 and X5-2 support four disks partnership.

Oracle Database Appliance X7-2 and later models have an eight disk partnership, if there are a minimum of eight disks. If there are not enough disks, then each disk partners with all others. For

example, Oracle Database Appliance X8-2 with six SSDs does not have eight disks to maintain eight disks partnership, hence all six disks partner among themselves.

For two JBOD systems, disk partnerships are across the JBOD boundary. It means that any disk in the first JBOD partners with eight disks from second JBOD and vice versa.

Disk partnership information is a part of the `asmappl.config` file.

The disk partnership on Oracle Database Appliance high-availability system is as follows:

DATA Disk S0 partners with S1, S2, S3, S4, S11, S12, S13, S14

RECO Disk S0 partners with S4, S5, S6, S7, S8, S9, S10, S11


If disk in slot 0 is lost for any reason, then Oracle ASM restores and rebalances data on lost disks among the partners of the lost disk.

If disk in slot 0 is lost for any reason, Oracle ASM restores and rebalances data on lost disks among the partners of lost disk.

Related Resources

- [Managing Storage](#)
- [odacli Storage Commands](#)
- [Readying Oracle Database Appliance for Deployment](#)

Oracle Database Appliance Single-Node Storage



Notes

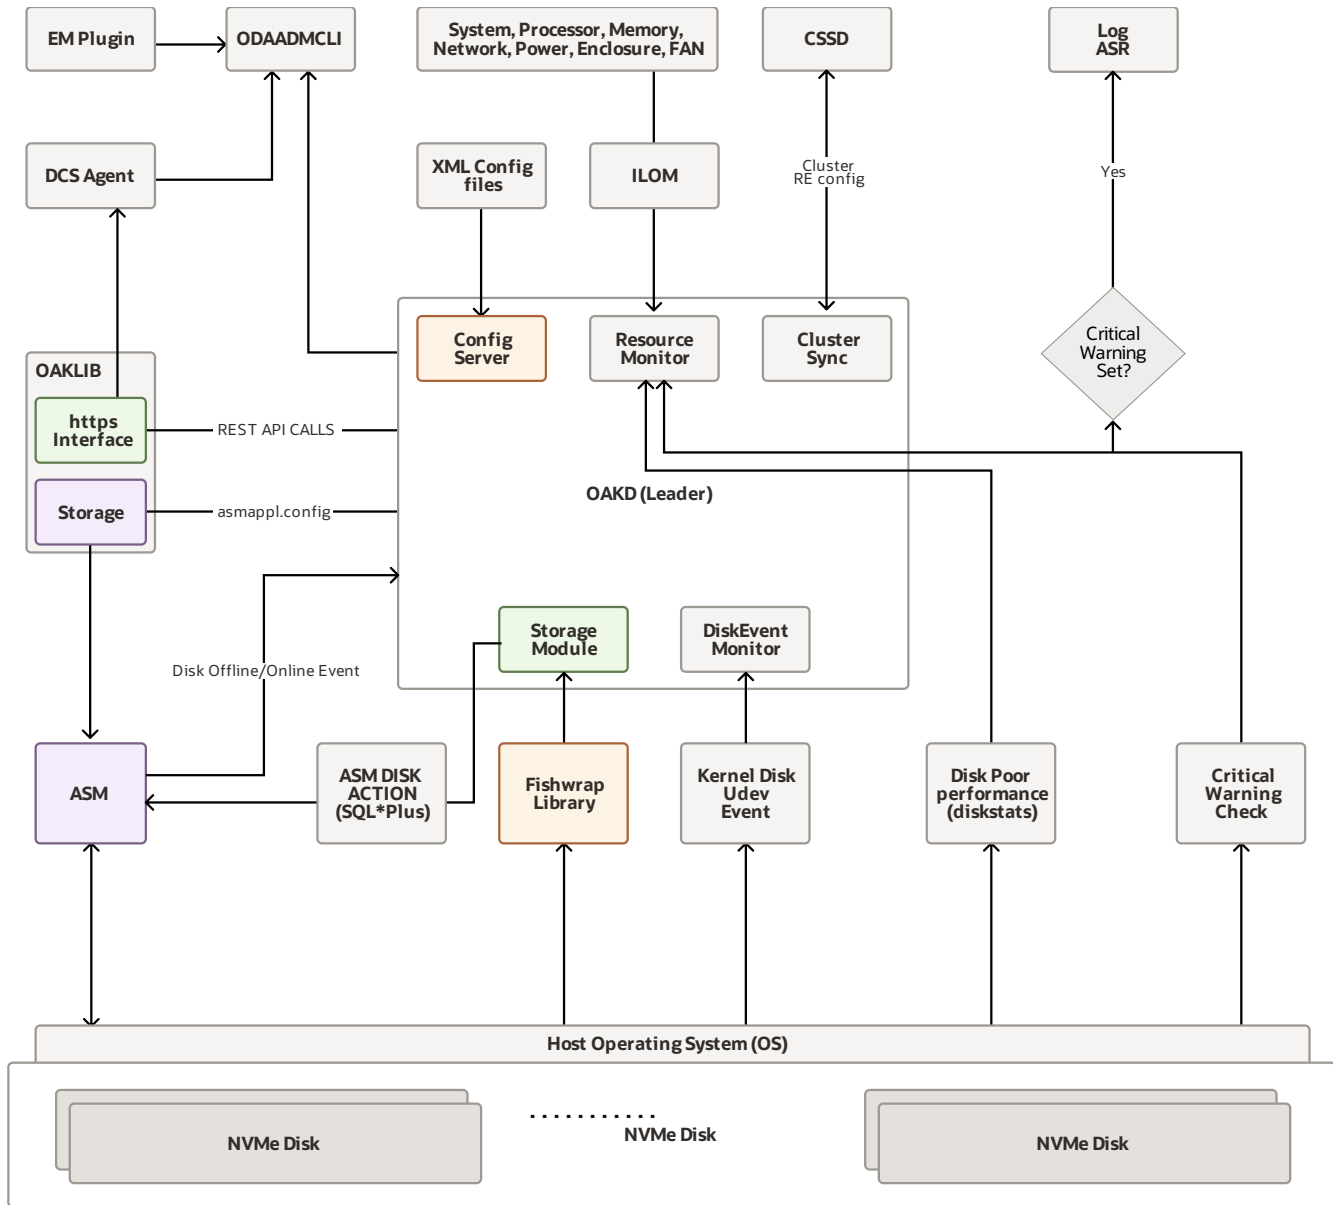
On an Oracle Database Appliance single-node server, OAKD is always in the leader role. Oracle Database Appliance single-node systems do not have SAS storage, hence LSI library are not used. Fishwrap library is used for storage discovery.

There are no SAS HBA (Controller) and hence LSI events notification is available. UDEV events notification are used for disk hot plugs. Instead of predictive failure checks, NVMe disks uses critical warning check using OAK NVMe interface.

Related Resources

- [Managing Storage](#)
- [odacli Storage Commands](#)
- [Readying Oracle Database Appliance for Deployment](#)

Oracle Database Appliance Disk Partnership on Single-Node Storage



Notes

Oracle ASM supports fixed partnership for disks. Oracle ASM performs splitting and mirroring of data across all disk partners. All Oracle Database Appliance single-node systems support 8 disks partnership provided there are enough disks to support it. If the number of NVMe disks are less than eight then all disks partner among themselves.

NVMe disks are populated across the PCI switch card for equal distribution of IO bandwidth across IO sub-systems. QRM1 and QRM2 have 1GB partitioned carved out from the operating system disk. They are used as DATA and RECO quorum disk respectively in Oracle ASM. These disk are used if the number of NVMe disks is less than 5. They are used for storing PST and voting files. They are

used when a disk with voting file or PST get lost. Their content helps to decide the majority without which the Oracle ASM disk group cannot be mounted.

The disk partnership information is part of the asmappl.config file.

The disk partnership on Oracle Database Appliance single-node system is as follows:

DATA Disk S0 partners with S2, S3, S4, S5, S8, S9, S10, S11

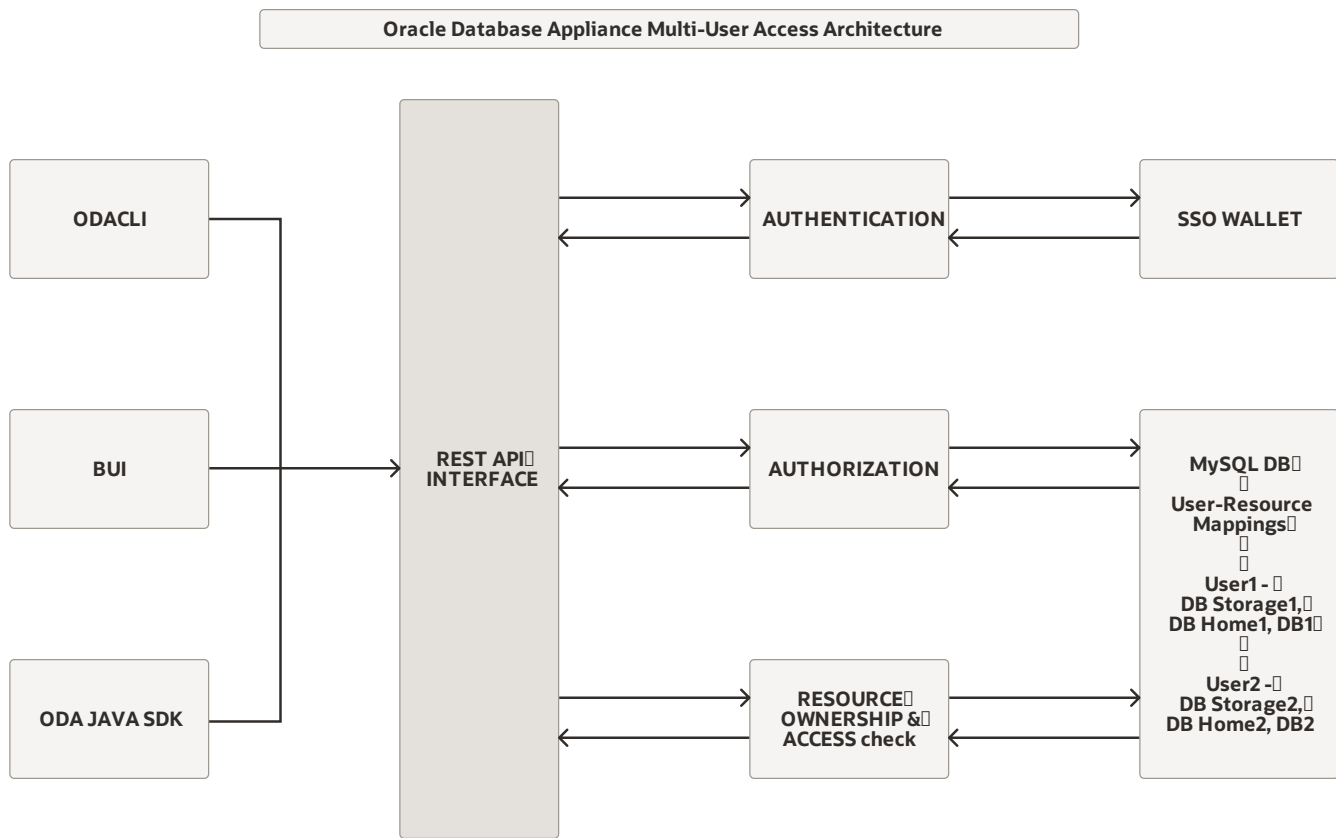
RECO Disk S0 partners with S2, S3, S4, S5, S8, S9, S10, S11

If disk in slot 0 is lost for any reason, Oracle ASM restores and rebalances data on lost disks among the partners of lost disk.

Related Resources

- [Managing Storage](#)
- [odacli Storage Commands](#)
- [Readying Oracle Database Appliance for Deployment](#)

Oracle Database Appliance Multi-User Access Architecture



Notes

On an Oracle Database Appliance multi user access-enabled system, an Oracle Database Appliance user runs a REST call on the DCS agent using an ODA CLI command, or a Browser User Interface (BUI) operation, or an Oracle Database Appliance Java SDK call.

Within the DCS agent, the REST API layer receives this request and makes a call to the authentication module of the Oracle Database Appliance multi-user access framework. The authentication module of the multi-user access framework authenticates the credentials of the REST caller against the SSO wallet.

After successful authentication, the REST API layer makes a call to the authorization module of the

Oracle Database Appliance multi-user access framework. The authorization module of the multi-user access framework checks if the user is authorized to perform this operation by checking against the user role and entitlements-related information stored in the MySQL metadata of the multi-user access framework.

After successful authorization, the REST API layer makes a call to the Resource Ownership and Access module of the multi-user access framework. The Resource Ownership and Access module checks against the resource ownership information stored in the MySQL metadata of the multi-user access framework to determine if the user has the required access privileges to perform lifecycle management operations on the specified resource. If the Resource Ownership and Access check is successful, then the REST API layer runs the REST endpoint and displays the results. Otherwise, an error message to this effect is displayed.

Related Resources

- [Implementing Multi-User Access on Oracle Database Appliance](#)
- [Oracle Database Appliance X9-2-HA Hardware Model](#)
- [Oracle Database Appliance X9-2S/X9-2L Hardware Model](#)