ORACLE

Maximum Availability Architecture for Middleware

High Availability and Disaster Recovery best practices

MAA team

2024

Safe harbor statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Agenda

Maximum Availability Architecture for Middleware

1

Maximum Availability Architecture

Middleware MAA introduction and paradigms

2

High Availability

Middleware HA in a single datacenter for

- On premise
- Cloud

3

Disaster Recovery: Active-Active

Middleware Stretched clusters for

- On premise
- Cloud

4

Disaster Recovery: Active-Passive

Middleware A-P for

- On premise
- Cloud
- Hybrid
- Kubernetes

5

Summary

Q&A

Agenda

Maximum Availability Architecture for Middleware

1

Maximum Availability Architecture

Middleware MAA introduction and paradigms

2

High Availability

Middleware HA in a single datacenter for

- On premise
- Cloud

3

Disaster Recovery: Active-Active

Middleware Stretched clusters for

- On premise
- Cloud

4

Disaster Recovery: Active-Passive

Middleware A-P for

- On premise
- Cloud
- Hybrid
- Kubernetes

5

Summary Q&A

Maximum Availability Architecture for Middleware

Introduction

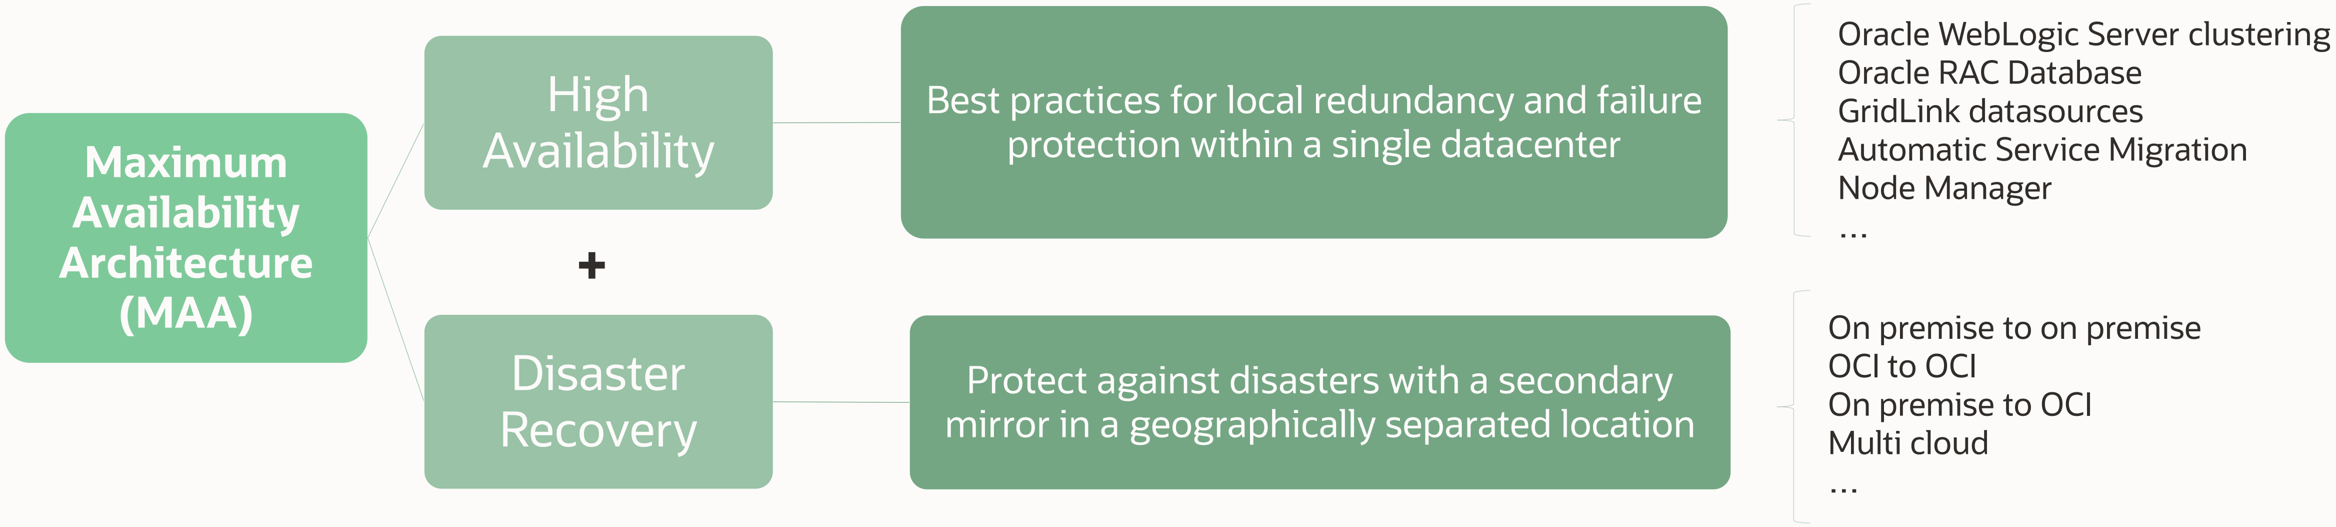

Maximum Availability Architecture (MAA)

High Availability

+

Disaster Recovery

Best practices for local redundancy and failure protection within a single datacenter

- Oracle WebLogic Server clustering
- Oracle RAC Database
- GridLink datasources
- Automatic Service Migration
- Node Manager
- …

Protect against disasters with a secondary mirror in a geographically separated location

- On premise to on premise
- OCI to OCI
- On premise to OCI
- Multi cloud
- …

Oracle MAA team: under Database Dev organization to develop, design, validate and document the MAA best practices for all the products in the Oracle stack.

Maximum Availability Architecture for Middleware

Concepts

RTO

- Recovery Time Objective
- The **time** between the event of a failure and the point where operations resume that an organization can tolerate

RPO

- Recovery Point Objective
- Time-based measurement of the maximum amount of **data loss** that an organization can tolerate

Maximum Availability Architecture for Middleware

Common features

Backup and restore

Oracle Database Data Guard

Storage replication

Load Balancer

Virtual Frontend name and Listen Addresses

WebLogic Server clusters
GridLink datasources
Automatic Service Migration
Node Manager

Fusion Middleware HA features

Maximum Availability Architecture for Middleware

For each topology

- Description (general, on prem, OCI)
- Benefits
- Typical variations not validated for MAA
- Where to find more information

Agenda

Maximum Availability Architecture for Middleware

1

Maximum Availability Architecture

Middleware MAA introduction and paradigms

2

High Availability

Middleware HA in a single datacenter for

- On premise
- Cloud

3

Disaster Recovery: Active-Active

Middleware Stretched clusters for

- On premise
- Cloud

4

Disaster Recovery: Active-Passive

Middleware A-P for

- On premise
- Cloud
- Hybrid
- Kubernetes

5

Summary Q&A

High Availability

Description

- Protection within the scope of a single data center
- MAA best practices for MW HA:

Load balancer in front	• With HA protection (local standby)
Redundancy in all the layers	• At least 2 Oracle HTTP Servers • At least 2 WebLogic nodes (WLS cluster) • Oracle RAC Database • Highly Available Storage
Local recovery features and best practices	• Failover for Admin Server based on Virtual Hostname with floating IP and Shared storage • WebLogic Service Migration for JMS and JTA • JDBC Persistent Stores • WebLogic Node Manager (server's auto restart) • Weblogic Gridlink Datasources • Backup

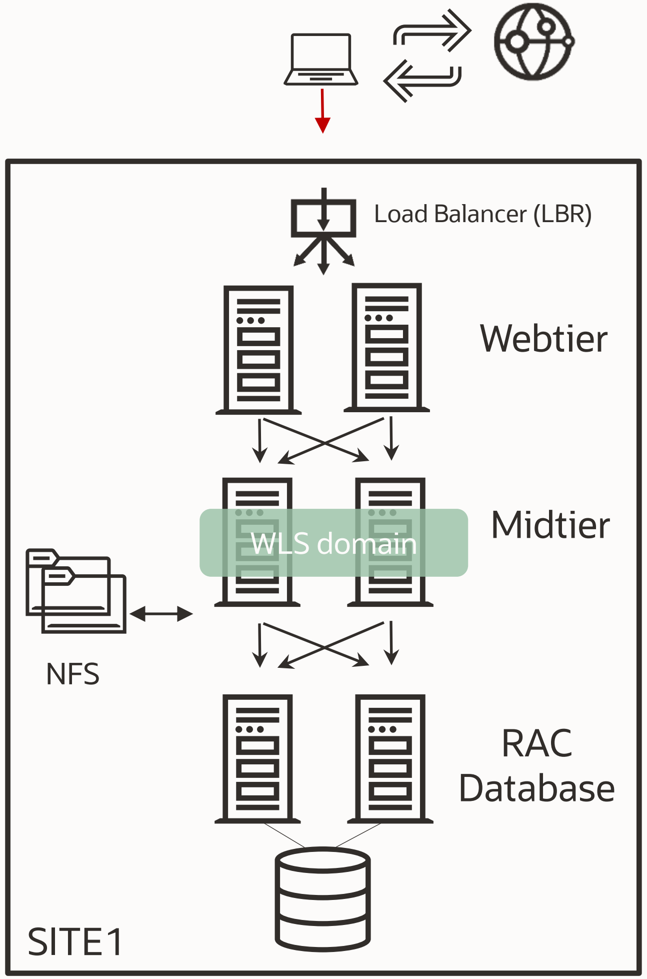

High Availability

Description – On Premise

FMW Enterprise Deployment Guides (EDG):

- Since 10g version (14.1.2 in progress)
- Step by step document
- Implements all the HA best practices for the FMW products

HW requirements:

- Load Balancer
- Storage
  - Shared (NFS)
  - Private (block volumes)
- VHN/VIP for the Admin Server
- (+ computes, networking, FW, backups)

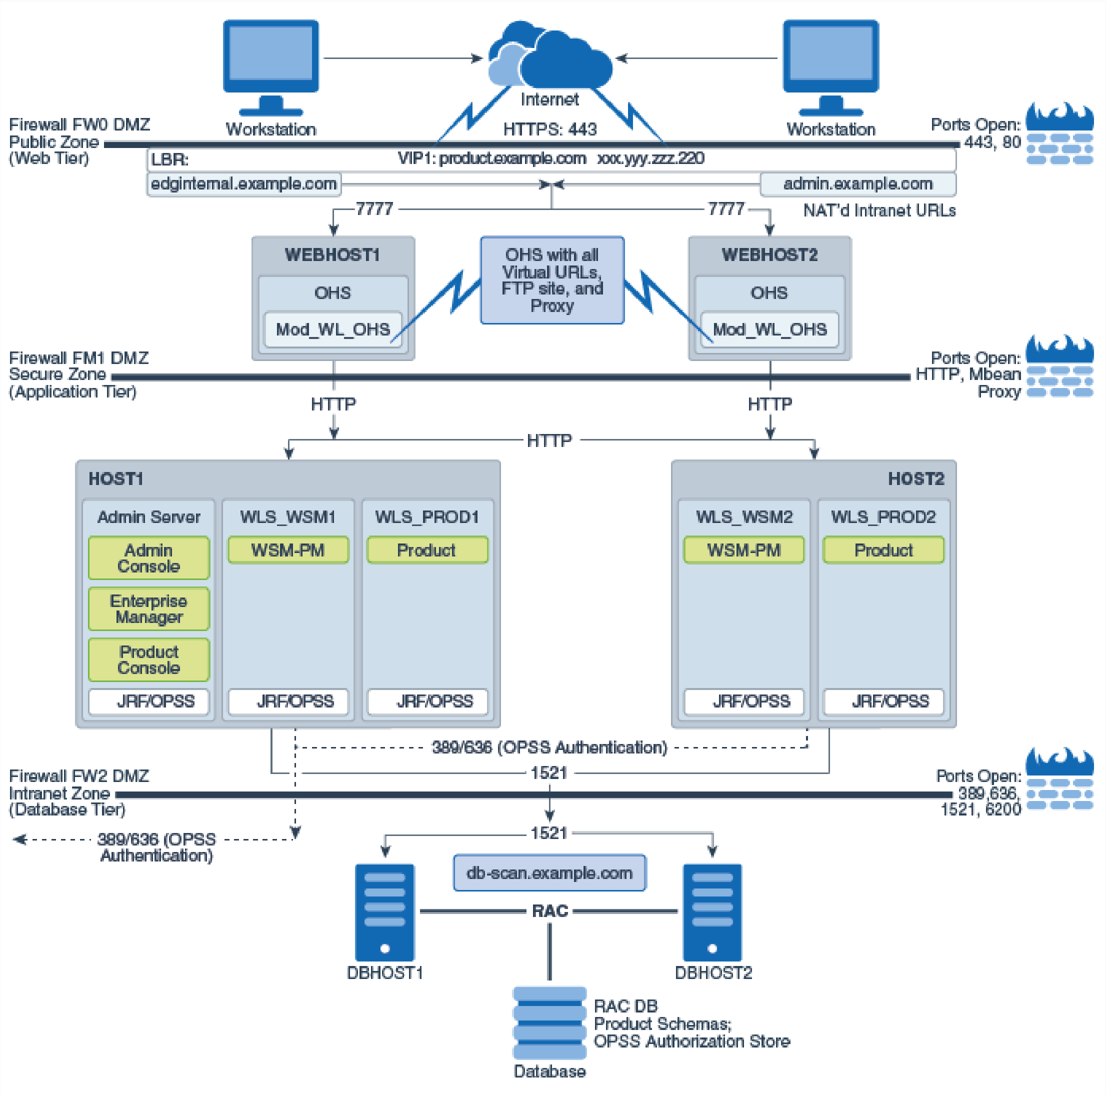

Topology Diagram from EDG document

High Availability

Description – On Cloud

OCI PaaS

- Oracle WebLogic Server for OCI
- Oracle SOA Suite on Marketplace

Not using OCI PaaS

- Manual installation and configuration of FMW products on IaaS

High Availability

Description – On Cloud

When using OCI PaaS services:

- WLS for OCI
- SOA Suite on Marketplace

HA best practices implemented out-of-the-box:

- Load Balancer in front
  - With HA protection (transparent local LBR standby)
- Redundancy in layers
  - WLS Clustering, nodes in different FD and ADs when available
  - RAC with OCI DB services
- Local recovery features:
  - Node Manager (for crash recovery)
  - GridLink datasources when RAC is used
  - Virtual Machine HA features
  - JDBC persistent stores

HA best practices not implemented out-of-the-box :

- WebLogic Service Migration for JMS and JTA
  - Can be configured manually in a post step
- Admin Server Failover
  - No shared storage for Admin config, no VHN/VIP
  - It relies on the virtual machine HA features (live migration)

High Availability

Description – On Cloud

WLS for OCI stack in a single AD region (similar for SOA Marketplace)

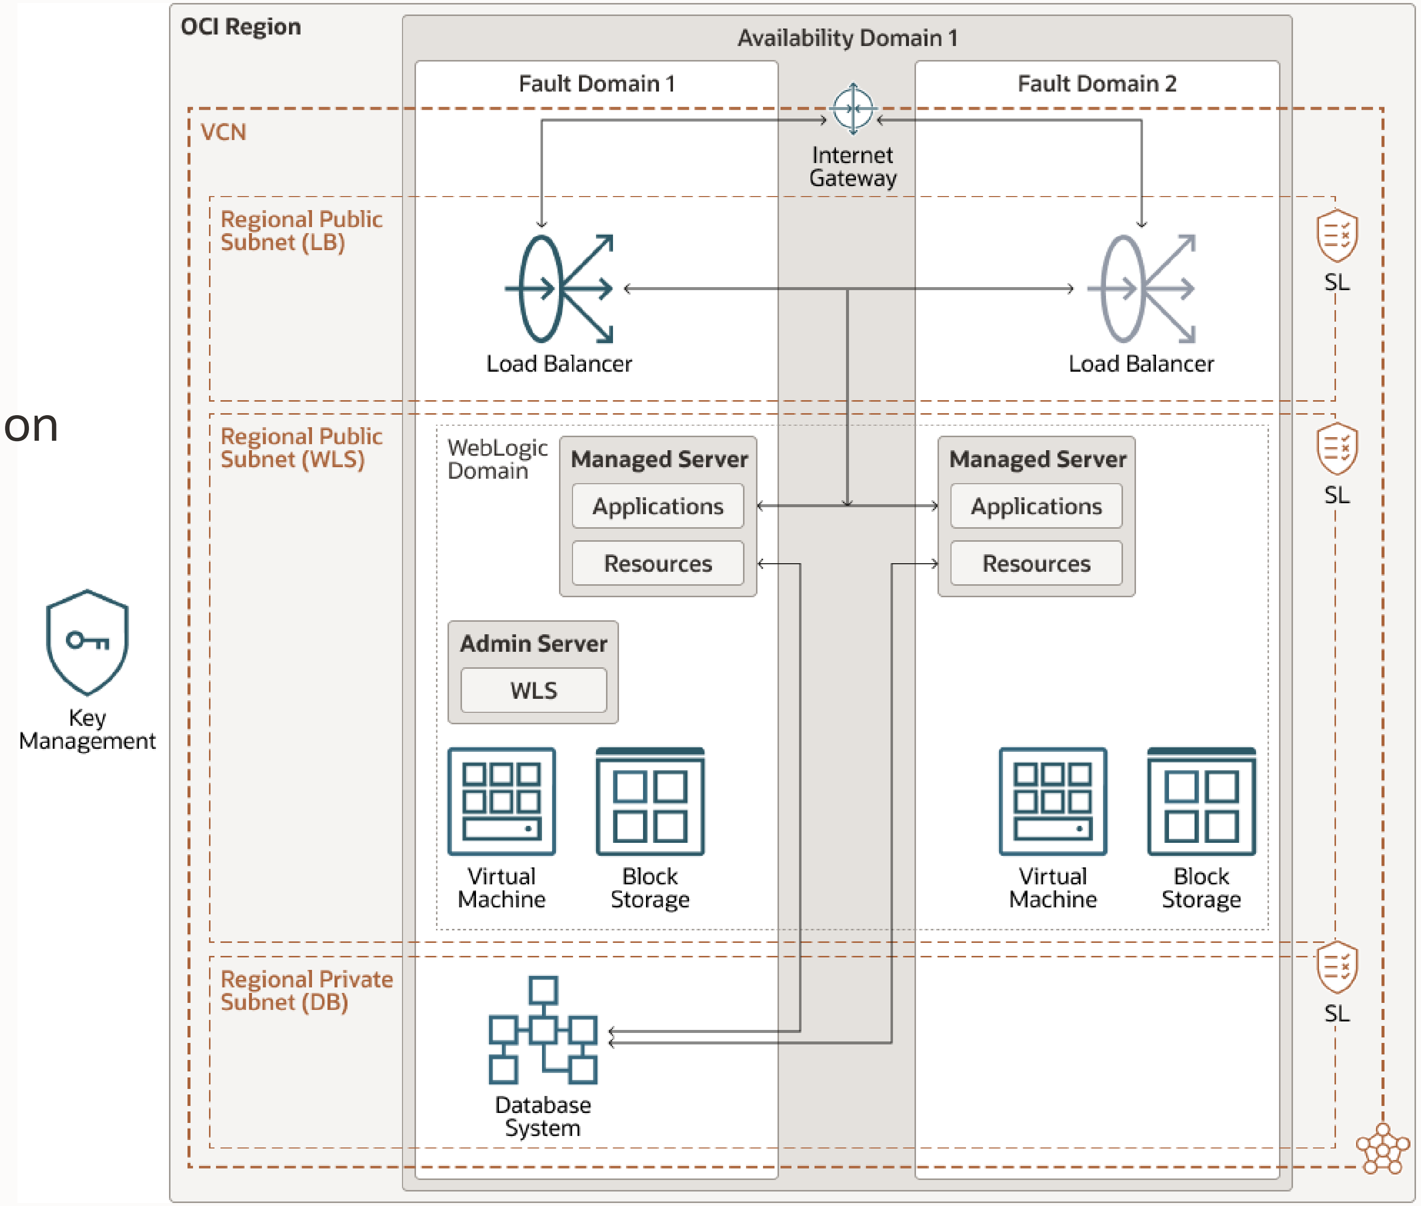

Topology Diagram from PaaS docs

High Availability

Description – On Cloud

- When using manual installation on IaaS:
  - EDG best practices can be implemented using OCI services and features:

OCI Load balancer	• With HA protection (transparent local standby)
Redundancy	• WLS compute nodes in different FD/AD* • RAC Database in OCI, nodes in different FD
OCI Storage	• File Storage services (for shared) • Block storage services
OCI Networking	• VIP for Admin Server • OCI DNS services
Local recovery features	• WLS (NM, GridLink, Service Migration) • Virtual Machine HA (live migration)

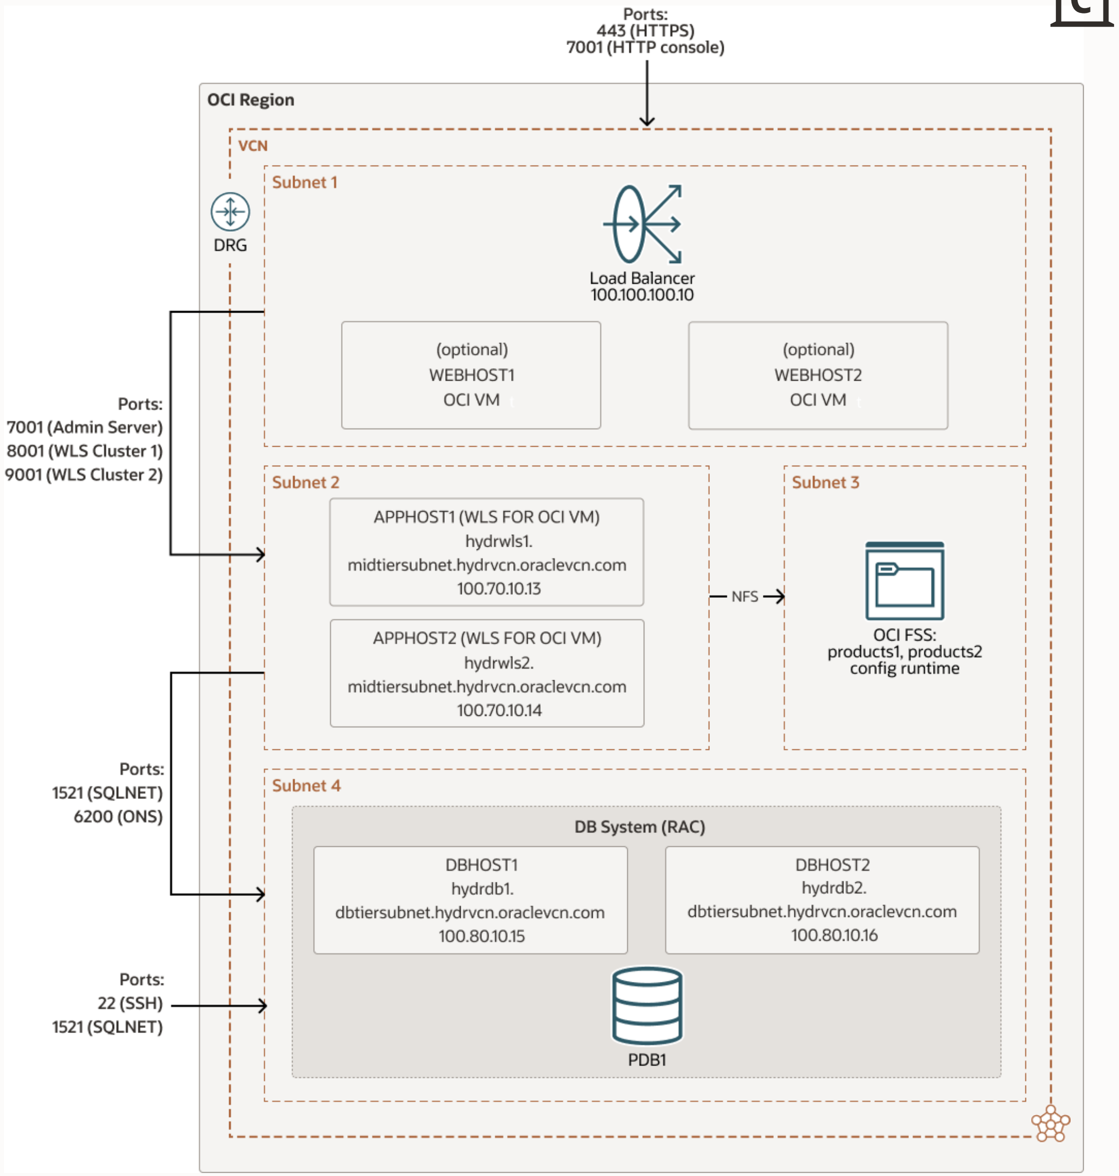

High Availability

Benefits

Benefits

- HA best practices protect the system from the local failures.
- Sometimes the failover is transparent, or with very low RTO.
- For on premise, the EDGs are completely reviewed, updated and validated (HA, DR and functional tests) in every FMW release.
- For OCI, many OCI features provide HA out-of-the-box (PaaS services, regional networks, Fault Domains, OCI LBR, etc.).

Limitations

- HA protection does NOT protect against outages that affect the entire datacenter.

High Availability

Solutions **not** validated for MAA

Unique domain folder shared by all the hosts	• Contention (logs) and potential conflicts (NodeManager) • EDG shared folder only for Admin Server config and a few shared items, but each host has its mount for the manager servers' domain folder.
Multi Datasources	• Use GridLink to access RAC databases
Non clustered servers	• Use WebLogic clusters and load balance in front
File persistent Stores for JMS and JTA	• Use JDBC persistent stores, advantages for service migration and consistency in DR

High Availability

More information

Enterprise Deployment Guides:

- EDG for Oracle SOA Suite and BPM
- EDG for Oracle Business Intelligence
- EDG for Oracle Webcenter Content
- EDG for Oracle Webcenter Portal
- EDG for Oracle Identity and Access Management

PaaS documentation:

- WLS for OCI docs
- SOA Suite on Marketplace docs

Agenda

Maximum Availability Architecture for Middleware

1

Maximum Availability Architecture

Middleware MAA introduction and paradigms

2

High Availability

Middleware HA in a single datacenter for

- On premise
- Cloud

3

Disaster Recovery: Active-Active

Middleware Stretched clusters for

- On premise
- Cloud

4

Disaster Recovery: Active-Passive

Middleware A-P for

- On premise
- Cloud
- Hybrid
- Kubernetes

5

Summary Q&A

Disaster Recovery: Active-Active

Description

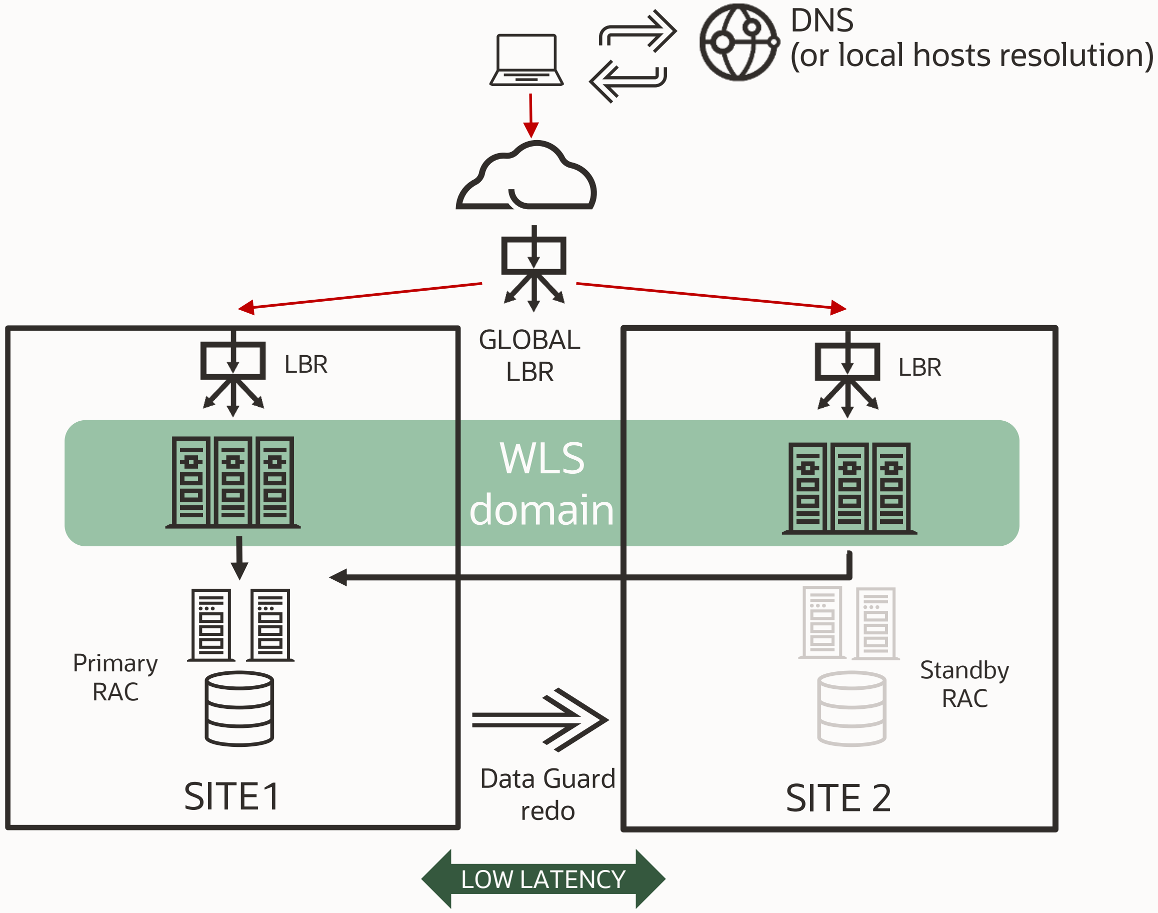

Stretched cluster topology

- Midtier A-A
  - Same WebLogic domain: some nodes in site1, other nodes in site2
  - Best practices to minimize the traffic between sites
  - No file system replica needed . The config is replicated between nodes with WebLogic mechanisms (it is a standard WLS cluster!).
- Database A-P
  - Data Guard for the database
- Global Load Balancer
  - To balance traffic between the sites' load balancers.

Supported only when latency between datacenters is low (<10ms RTT)

Disaster Recovery: Active-Active

Description

After a database only switchover

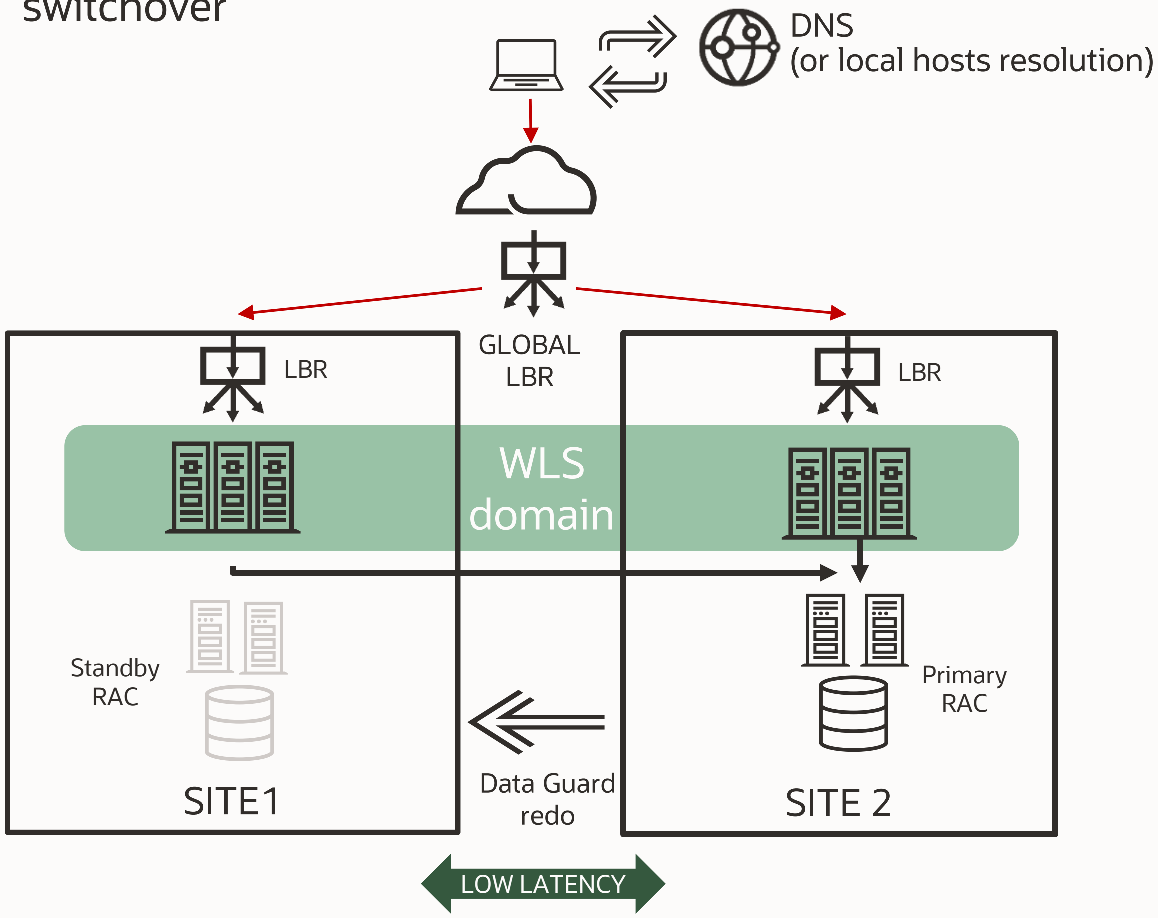

Stretched cluster topology

- Midtier A-A
  - Same WebLogic domain: some nodes in site1, other nodes in site2
  - Best practices to minimize the traffic between sites
  - No file system replica needed . The config is replicated between nodes with WebLogic mechanisms (it is a standard WLS cluster!).
- Database A-P
  - Data Guard for the database
- Global Load Balancer
  - To balance traffic between the sites' load balancers.

Supported only when latency between datacenters is low (<10ms RTT)

Disaster Recovery: Active-Active

Description

Stretched cluster - Complete site outages

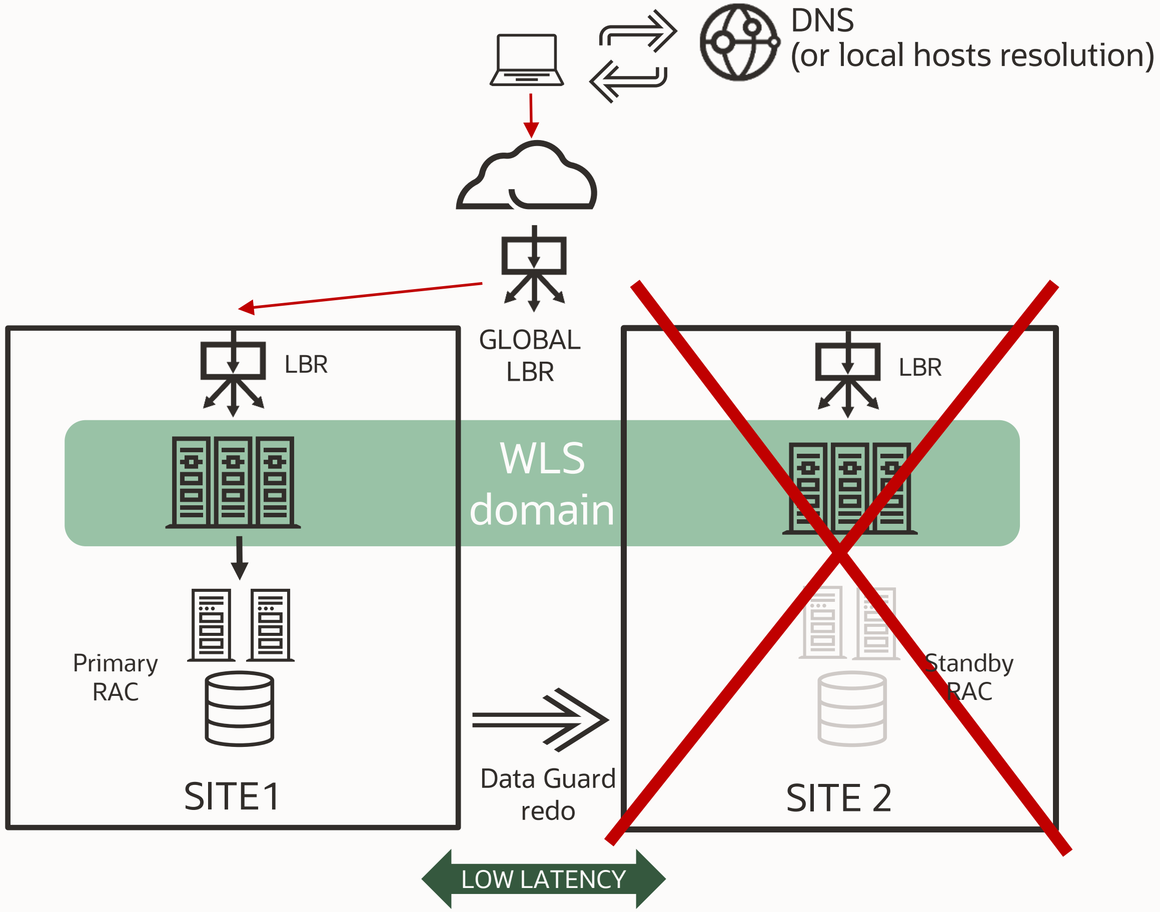

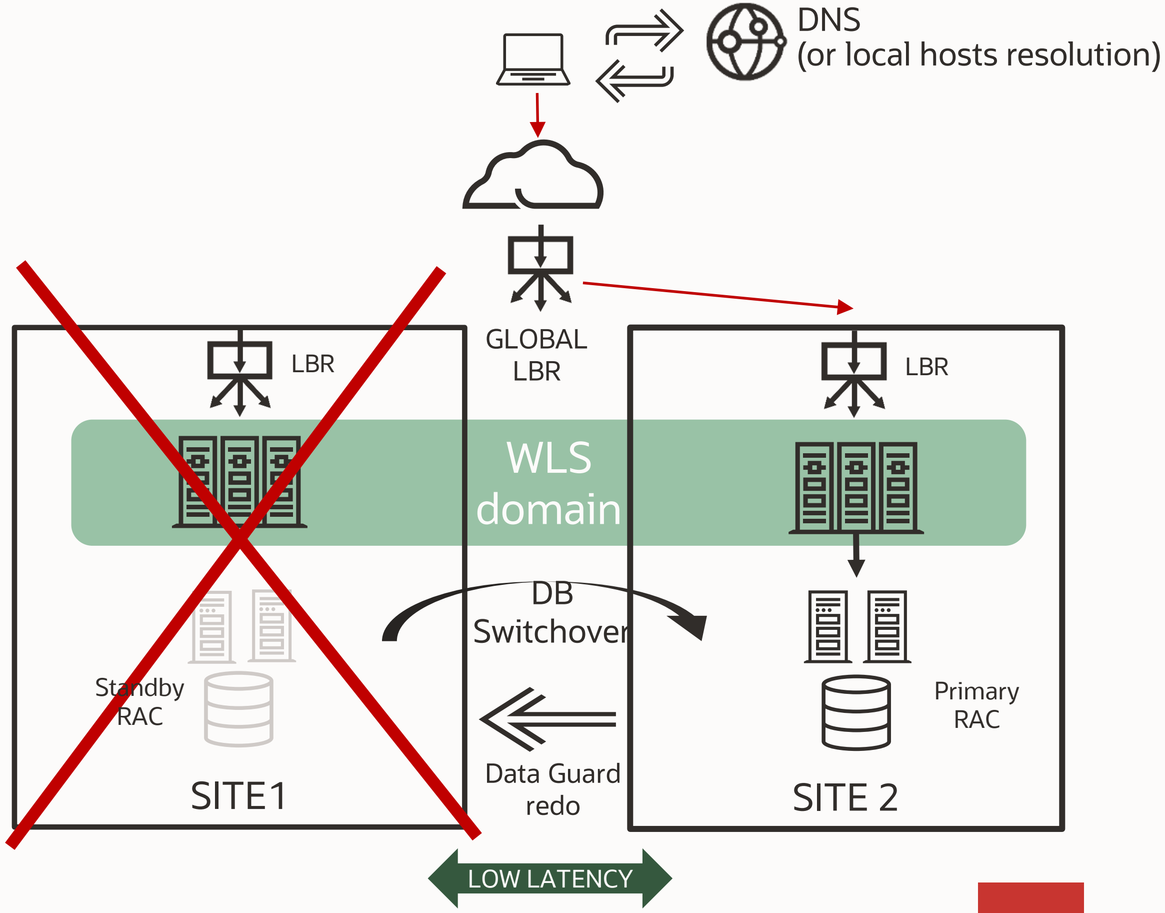

Disaster Recovery: Active-Active

Description – On Premise

Best practices document since 11g

HW requirements:

- Global Load Balancer
- 2 Local Load Balancers
- RAC DB with Data Guard
- Storage:
  - Shared: NFS (restricted within each site)
  - Private: Block Volumes
  - Storage replication between sites NOT needed
- Network:
  - Good bandwidth and low latency between sites (<10 ms RTT)
- (+compute, etc.)

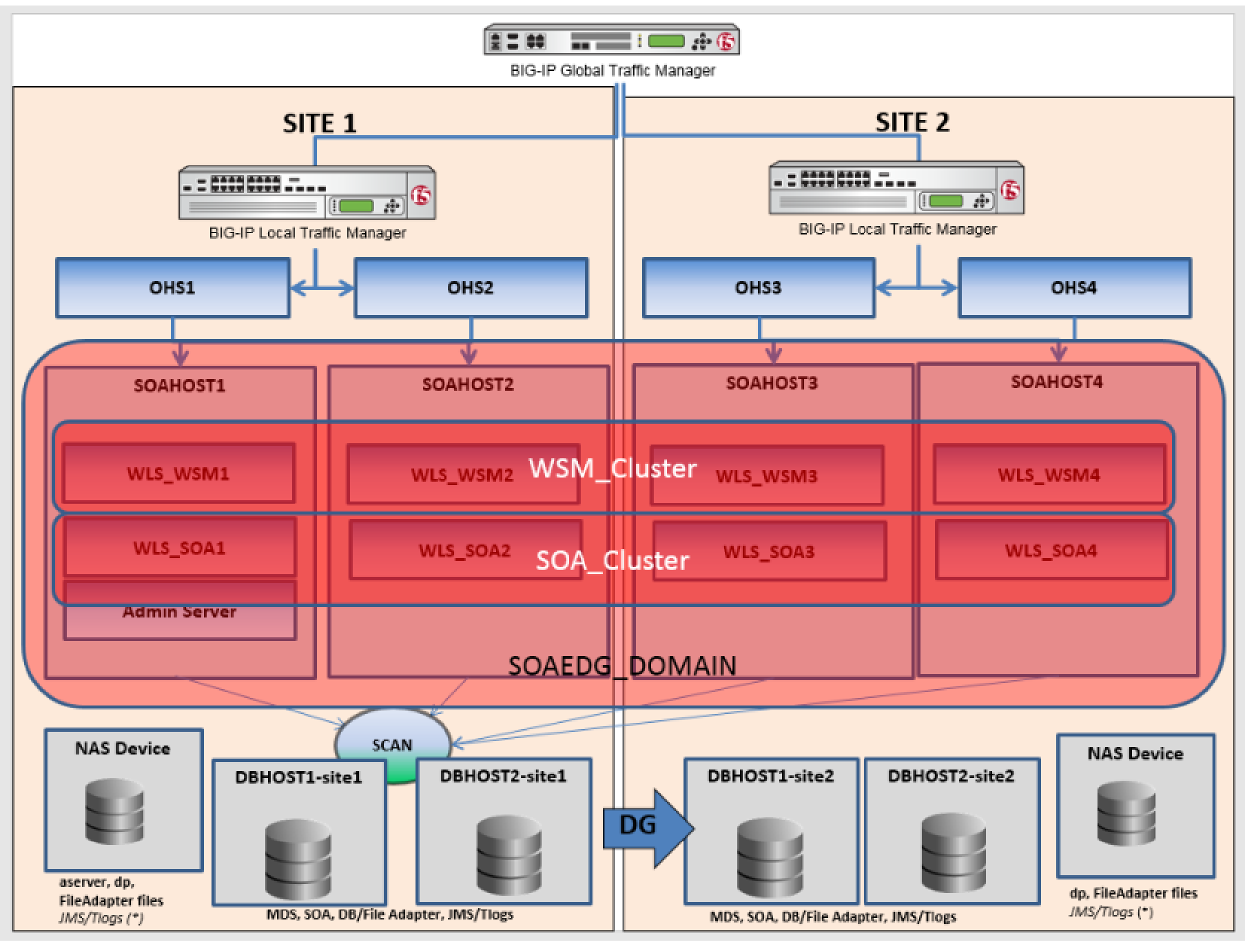

Topology Diagram from active-active document

Disaster Recovery: Active-Active

Description – On Cloud

- Latency between OCI regions is normally high.

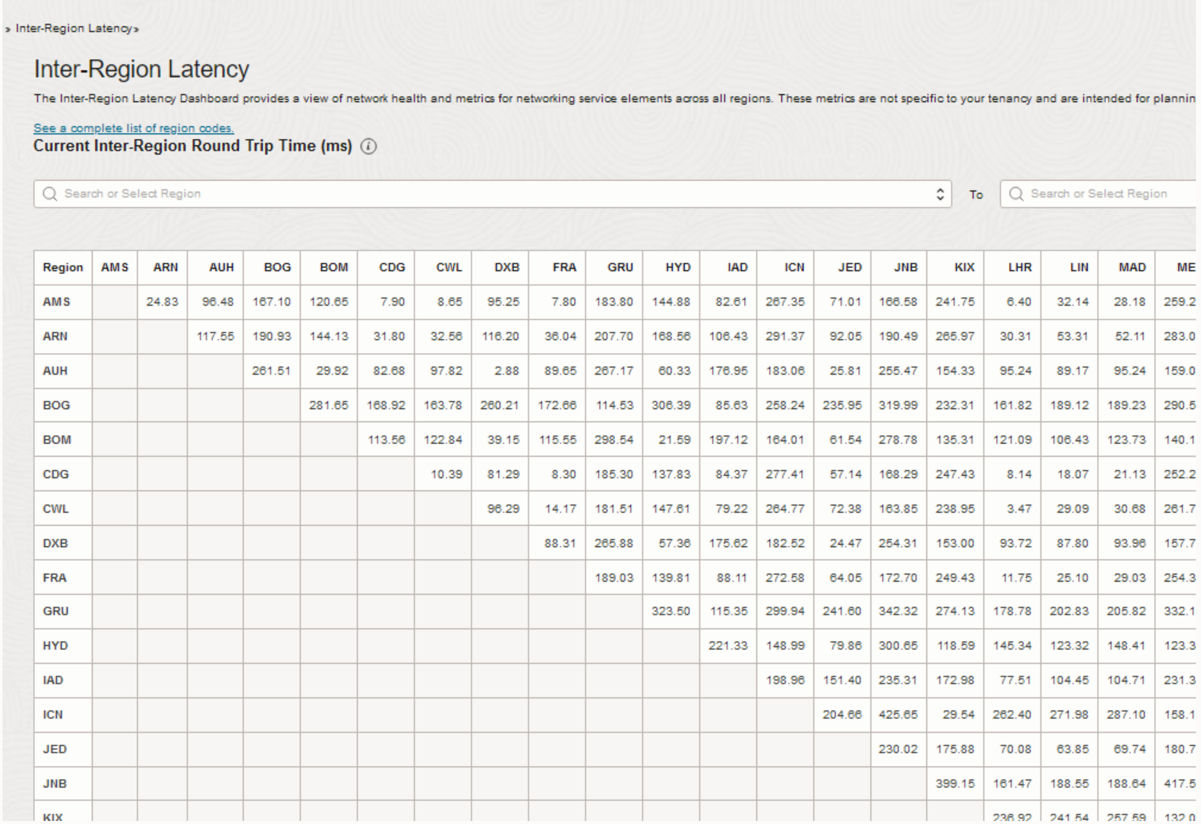

› Inter-Region Latency›

Inter-Region Latency

The Inter-Region Latency Dashboard provides a view of network health and metrics for networking service elements across all regions. These metrics are not specific to your tenancy and are intended for plannin

See a complete list of region codes.

Current Inter-Region Round Trip Time (ms) ⓘ

Region	AMS	ARN	AUH	BOG	BOM	CDG	CWL	DXB	FRA	GRU	HYD	IAD	ICN	JED	JNB	KIX	LHR	LIN	MAD	ME
AMS		24.83	96.48	167.10	120.65	7.90	8.65	95.25	7.80	183.80	144.88	82.61	267.35	71.01	166.58	241.75	6.40	32.14	28.18	259.2
ARN			117.55	190.93	144.13	31.80	32.56	116.20	36.04	207.70	168.56	106.43	291.37	92.05	190.49	265.97	30.31	53.31	52.11	283.0
AUH				261.51	29.92	82.68	97.82	2.88	89.65	267.17	60.33	176.95	183.06	25.81	255.47	154.33	95.24	89.17	95.24	159.0
BOG					281.65	168.92	163.78	260.21	172.66	114.53	306.39	85.63	258.24	235.95	319.99	232.31	161.82	189.12	189.23	290.5
BOM						113.56	122.84	39.15	115.55	298.54	21.59	197.12	164.01	61.54	278.78	135.31	121.09	106.43	123.73	140.1
CDG							10.39	81.29	8.30	185.30	137.83	84.37	277.41	57.14	168.29	247.43	8.14	18.07	21.13	252.2
CWL								96.29	14.17	181.51	147.61	79.22	264.77	72.38	163.85	238.95	3.47	29.09	30.68	261.7
DXB									88.31	265.88	57.36	175.62	182.52	24.47	254.31	153.00	93.72	87.80	93.96	157.7
FRA										189.03	139.81	88.11	272.58	64.05	172.70	249.43	11.75	25.10	29.03	254.3
GRU											323.50	115.35	299.94	241.60	342.32	274.13	178.78	202.83	205.82	332.1
HYD												221.33	148.99	79.86	300.65	118.59	145.34	123.32	148.41	123.3
IAD													198.96	151.40	235.31	172.98	77.51	104.45	104.71	231.3
ICN														204.66	425.65	29.54	262.40	271.98	287.10	158.1
JED															230.02	175.88	70.08	63.85	69.74	180.7
JNB																399.15	161.47	188.55	188.64	417.5
KIX																	236.92	241.54	257.59	132.0

- Stretched clusters possible only between Availability Domains within a multi-AD region.
  - Availability domain: a data center within a region
  - ADs within the same region are connected by a low latency, high bandwidth network

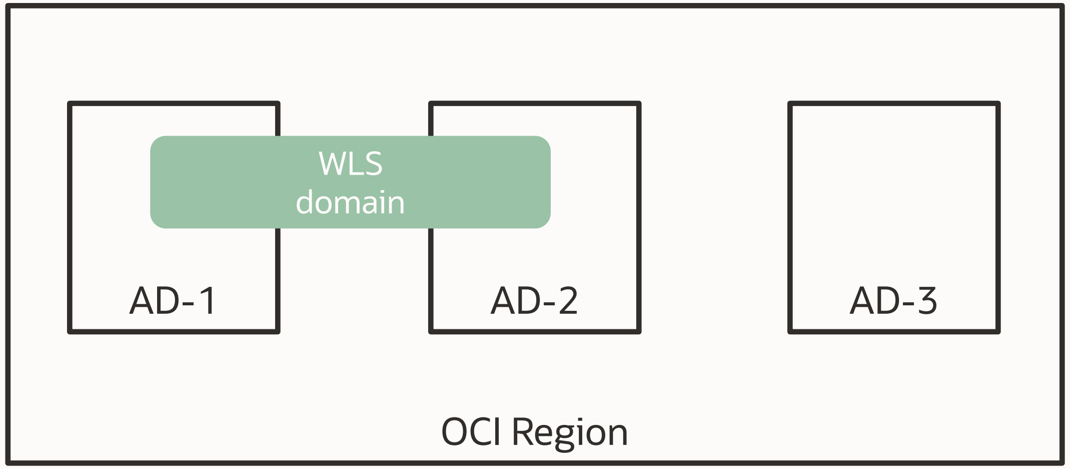

Disaster Recovery: Active-Active

Description – On Cloud

PaaS: some stretched cluster features implicitly provided between ADs:

- With regional subnets, midtier nodes distributed across ADs
- OCI LBR implicit cross-AD HA (standby in other AD)
- Data Guard (standby in other AD)

If not using PaaS (manual install on IaaS):

- Manually distribute computes in different ADs
- Take advantage of the OCI services and features:
  - Regional subnets
  - Data Guard in other AD
  - OCI Load Balancer

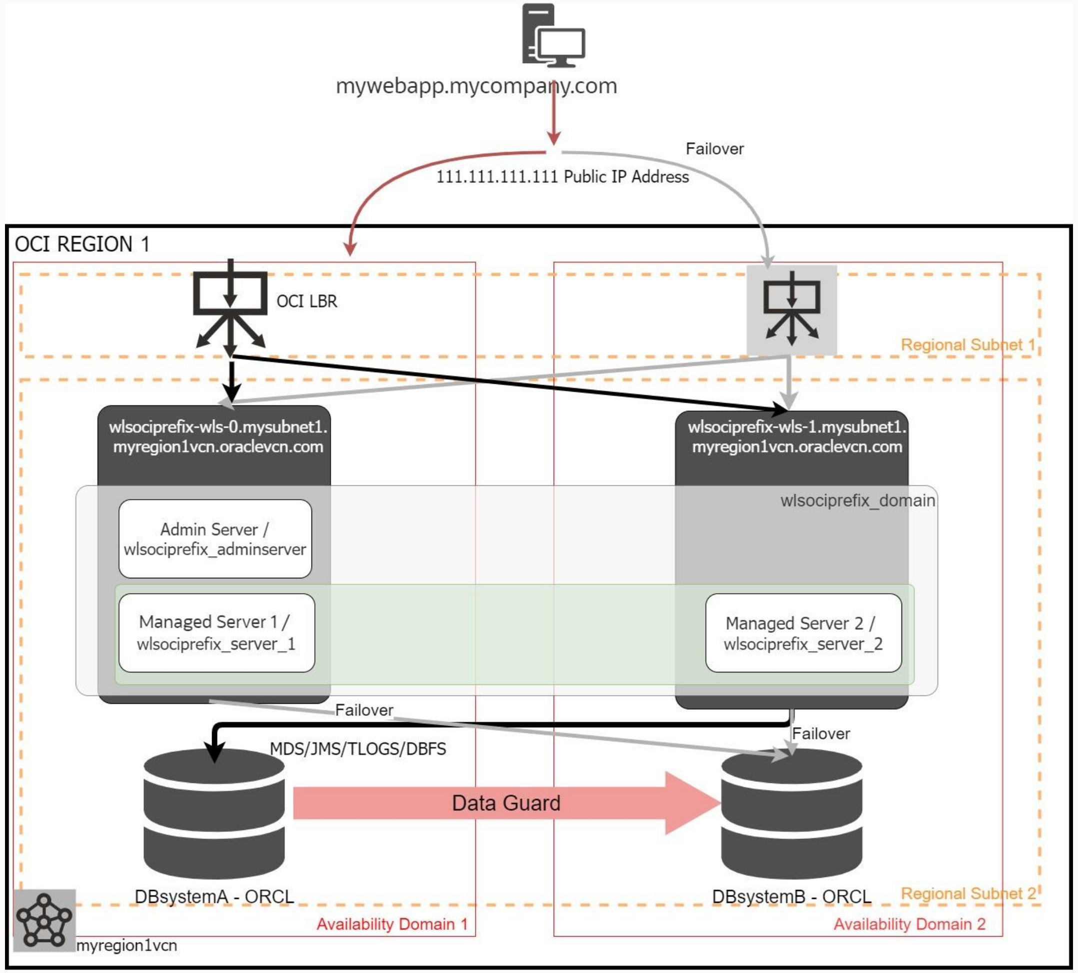
mywebapp.mycompany.com
Failover
111.111.111.111 Public IP Address
OCI REGION 1
OCI LBR
Regional Subnet 1
wlsociprefix-wls-0.mysubnet1.myregion1vcn.oraclevcn.com
wlsociprefix-wls-1.mysubnet1.myregion1vcn.oraclevcn.com
wlsociprefix_domain
Admin Server / wlsociprefix_adminserver
Managed Server 1 / wlsociprefix_server_1
Managed Server 2 / wlsociprefix_server_2
Failover
MDS/JMS/TLOGS/DBFS
Failover
Data Guard
DBsystemA - ORCL
DBsystemB - ORCL
Regional Subnet 2
Availability Domain 1
Availability Domain 2
myregion1vcn

Stretched cluster in a multi-AD OCI region

Disaster Recovery: Active-Active

Benefits

Benefits

- Stretched cluster is easy to manage
  - Single WLS domain
- Replication for file systems not required
  - The config is replicated between nodes with WebLogic mechanisms.
- Minimal RTO in a complete site outage
  - All midtier nodes are already up, just a DB failover
  - Midtier nodes automatically reconnect to DB (dual connect string)
- This topology provides HA and DR* at the same time

Limitations

- Supported only when latency between datacenters is low
  - <10ms RTT
- The midtier in each site need to be able to sustain the combined load
  - Appropriate capacity planning, sites designed with exceeding power under normal business.
- Shared storage between sites not possible
- Admin server failover to the other site requires manual actions (to copy the AS config)
- No standby system for test/validations like in an A-P DR

Disaster Recovery: Active-Active

Solutions **not** validated for MAA

Active-Active topologies different than stretched clusters

- Other A-A topologies not suitable for most FMW products because of the FMW metadata schemas, JMS and tlogs consistency

Database replica based on Golden Gate

- GG not applicable to all data types.
- GG is more oriented to data integration than to DR (TX consistency issues).

SOAMP and WLS for OCI (PaaS) stretched cluster across regions

- Can't provision stacks across regions
- Can be done with “manual installation” between some regions if latency is low (see Inter-region latency)

Disaster Recovery: Active-Active

More information

On premise:

- Best Practices for Oracle Fusion Middleware SOA 12c Multi Data Center Active-Active Deployment

PaaS documentation:

- WLS for OCI docs
- SOA Suite on Marketplace docs

Agenda

Maximum Availability Architecture for Middleware

1

Maximum Availability Architecture

Middleware MAA introduction and paradigms

2

High Availability

Middleware HA in a single datacenter for

- On premise
- Cloud

3

Disaster Recovery: Active-Active

Middleware Stretched clusters for

- On premise
- Cloud

4

Disaster Recovery: Active-Passive

Middleware A-P for

- On premise
- Cloud
- Hybrid
- Kubernetes

5

Summary Q&A

Disaster Recovery: Active-Passive

Description

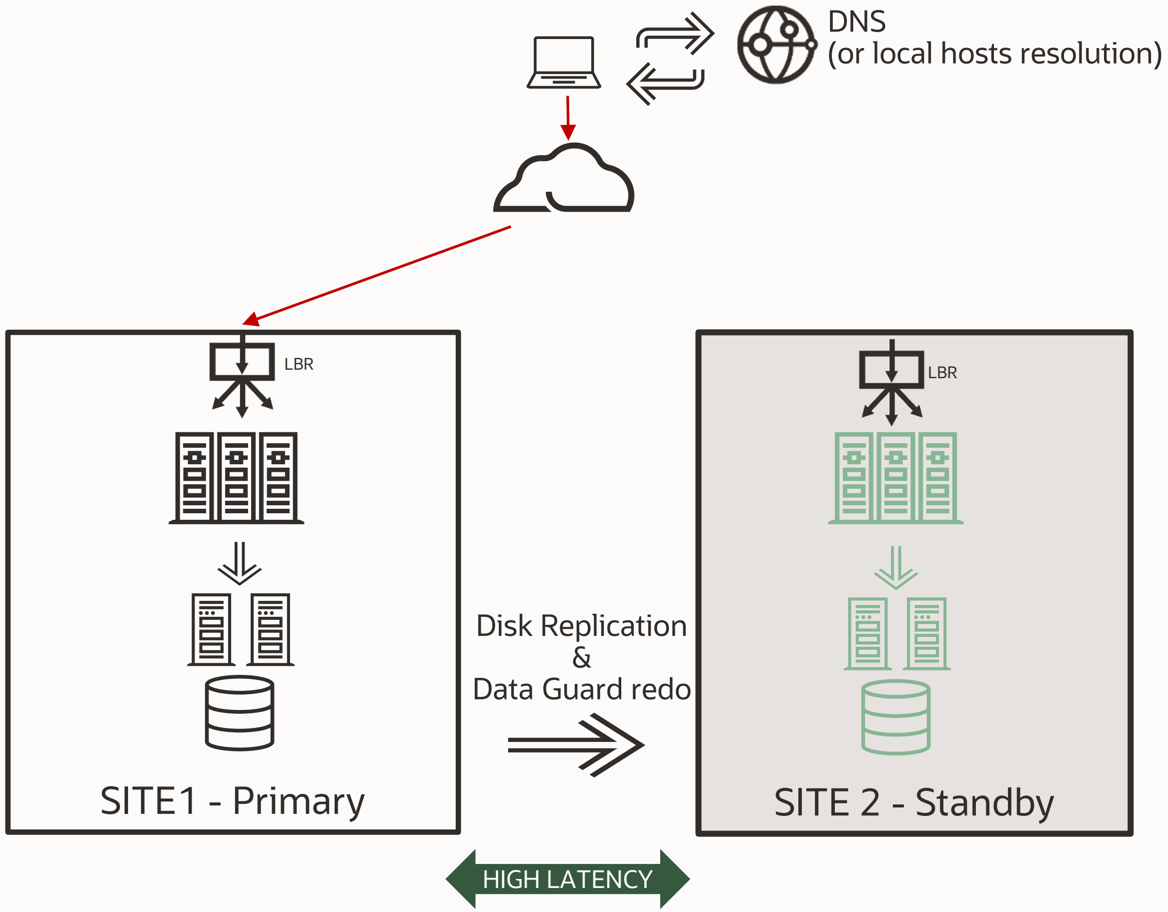

- Primary is Active
  - DB in primary role, sending redo to standby
- Standby is Passive
  - DB in standby role, applying redo from primary
  - Midtier hosts exists and can be up, but processes are stopped
- Frontend name in DNS
  - a.k.a. “vanity url”, “virtual frontend”, not IP!
- Listen addresses for MW components
  - Names, not IPs!
  - “Virtual names”: resolved in each site to the site’s IPs
- TNS alias for Data Sources
- A way to replicate
  - File systems (Disk replication, rsync*, DBFS)
  - Database (Data Guard)
- Symmetric resources in standby
  - Same number of nodes and capacity

Disaster Recovery: Active-Passive

Description

Lifecycle operations

- Switchover
  - Planned operation
- Failover
  - Unplanned operation
- Open standby for validations
  - This not possible in A-A model
  - Convert standby db into snapshot standby and start WLS processes in standby site for testing.
  - Manipulate frontend hostname resolution in the scope of the test.
  - Production continues in primary.
  - Any change in stby snapshot is lost

About TNS alias

Database connect string in WebLogic datasources
jdbc:oracle:thin:@mydatabase

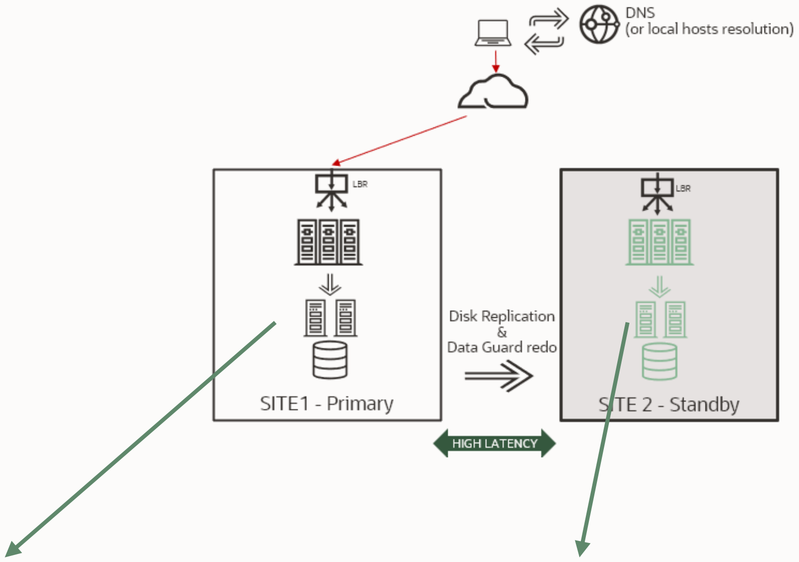

```
mydatabase =
(DESCRIPTION=
(ADDRESS_LIST= (LOAD_BALANCE=ON)
(ADDRESS=(PROTOCOL=TCP)(HOST=prmy-scan)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=mydb.example.com)) )
```

```
mydatabase =
(DESCRIPTION=
(ADDRESS_LIST= (LOAD_BALANCE=ON)
(ADDRESS=(PROTOCOL=TCP)(HOST=standby-scan)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=mydb.example.com)) )
```

tnsnames.ora in midtier nodes

Disaster Recovery: Active-Passive

Description – On Premise

- FMW Disaster Recovery guide
- Step by step guide to configure the DR in another site
- Assumes HA best practices implemented in each site (EDG)
- Procedure described "universally":
  - Not for any specific underlying technology
  - You should be able to implement with your technology (any LBR, any storage solution, etc.)

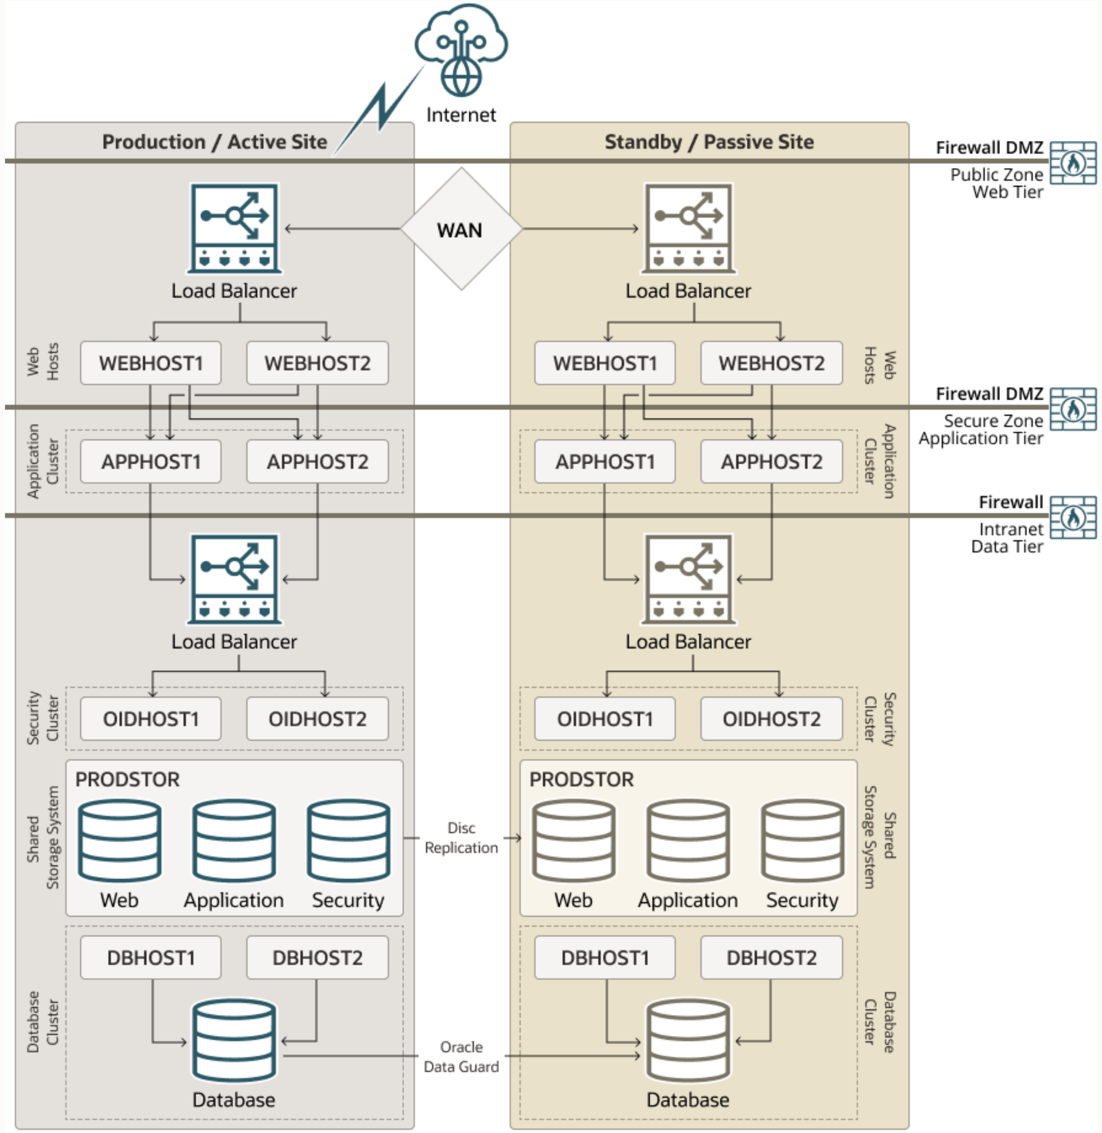

Topology Diagram from FMW DR guide

Disaster Recovery: Active-Passive

Description – On Cloud

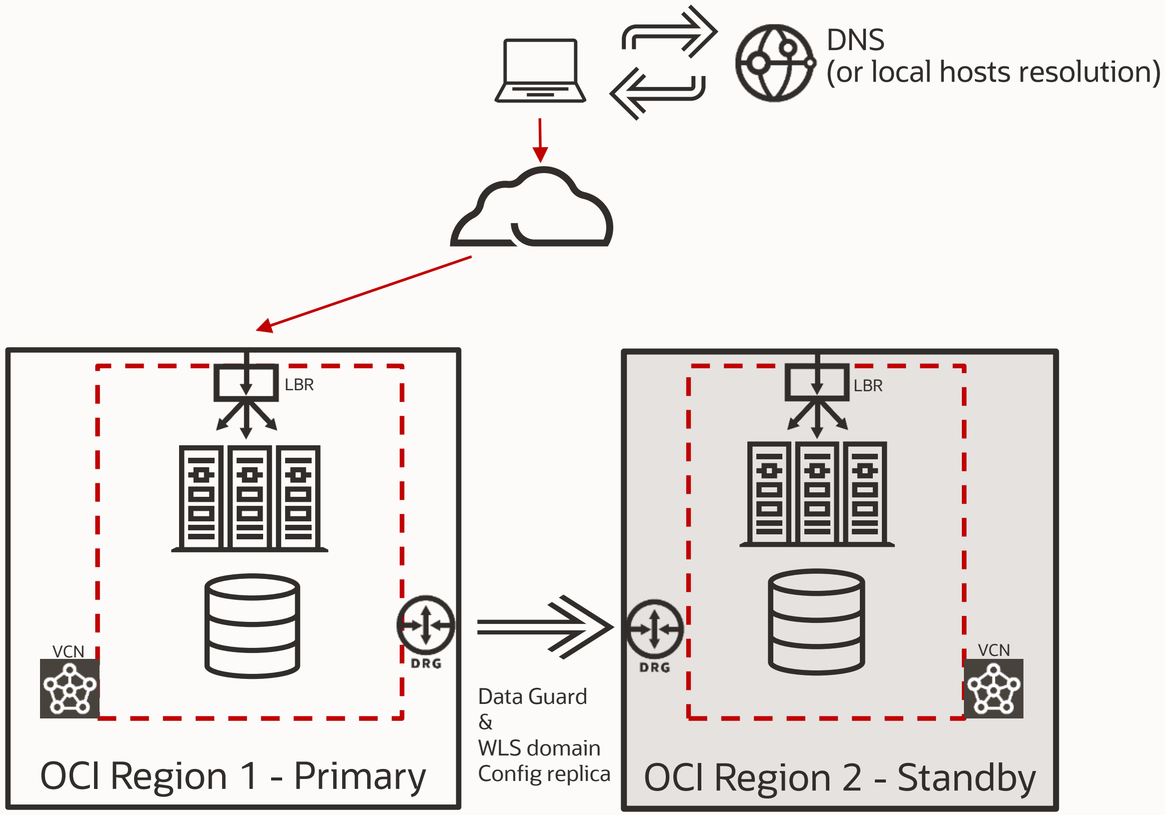

- SOAMP DR and WLS for OCI DR documents
- Database flavors supported:
  - Oracle Base Database Service
  - Oracle Exadata Database Service
  - Oracle Autonomous Database

- 3 methods for the config replication:
  - Block Volume replica (recommended)
  - rsync
  - DBFS

Disaster Recovery: Active-Passive

Description – On Cloud

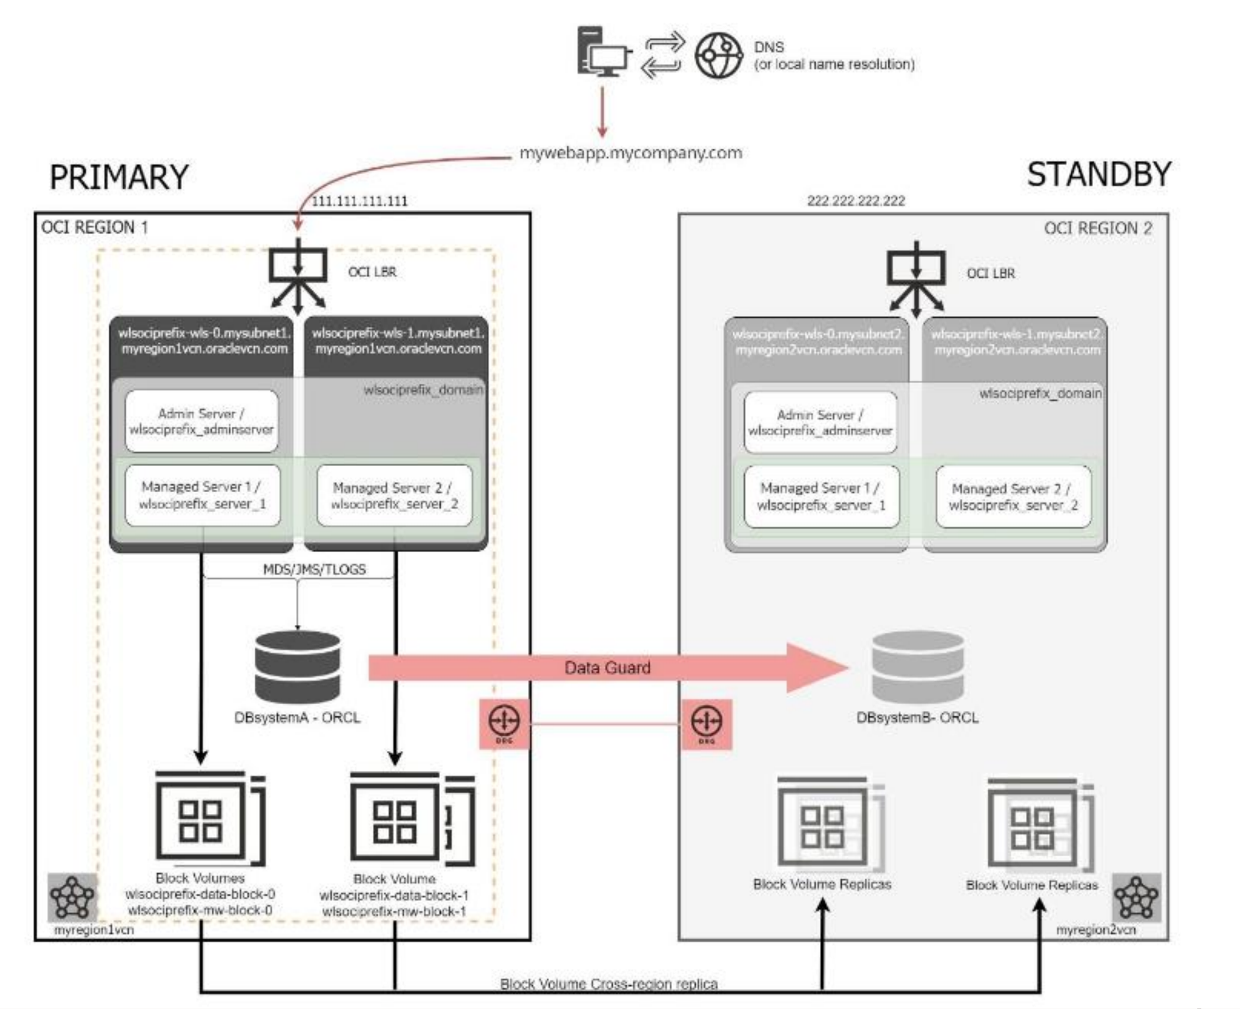
DNS
(or local name resolution)
mywebapp.mycompany.com
PRIMARY
111.111.111.111
OCI REGION 1
OCI LBR
wlsociprefix-wls-0.mysubnet1.myregion1vcn.oraclevcn.com
wlsociprefix-wls-1.mysubnet1.myregion1vcn.oraclevcn.com
wlsociprefix_domain
Admin Server / wlsociprefix_adminserver
Managed Server 1 / wlsociprefix_server_1
Managed Server 2 / wlsociprefix_server_2
MDS/JMS/TLOGS
DBsystemA - ORCL
Data Guard
DRG
Block Volumes wlsociprefix-data-block-0 wlsociprefix-mw-block-0
Block Volume wlsociprefix-data-block-1 wlsociprefix-mw-block-1
myregion1vcn
STANDBY
222.222.222.222
OCI REGION 2
OCI LBR
wlsociprefix-wls-0.mysubnet2.myregion2vcn.oraclevcn.com
wlsociprefix-wls-1.mysubnet2.myregion2vcn.oraclevcn.com
wlsociprefix_domain
Admin Server / wlsociprefix_adminserver
Managed Server 1 / wlsociprefix_server_1
Managed Server 2 / wlsociprefix_server_2
DBsystemB- ORCL
DRG
Block Volume Replicas
Block Volume Replicas
myregion2vcn
Block Volume Cross-region replica

Topology Diagram from PaaS DR docs

- Requirements for DR
  - OCI cross-region communication (Dynamic Routing Gateways)
  - DNS
  - OCI LBR
  - OCI Data Guard
  - OCI Block Volume replication

- Full Stack Disaster Recovery Service (FSDR):
  - OCI service to orchestrate the switchover and failover steps (not the setup).
  - Provides built-in steps for many OCI components (Data Guard, File System, Block volumes)
  - User defined steps
  - Create plans and execute with one click

Disaster Recovery: Active-Passive

Description – Hybrid

- Between On premise and OCI
- Symmetric system in OCI, implementing HA best practices:
  - Compute instances for midtier, webtier (WLS images)
  - OCI LBR
  - OCI RAC Database
  - Block Storage
  - File Storage Service for shared storage
- Fast Connect
  - For connectivity between on prem and OCI
- Data Guard
  - For the Database replication
- File system replication
  - Disk replica not possible
  - Bastion in OCI to perform rsync copies
- New automation framework from MAA available!

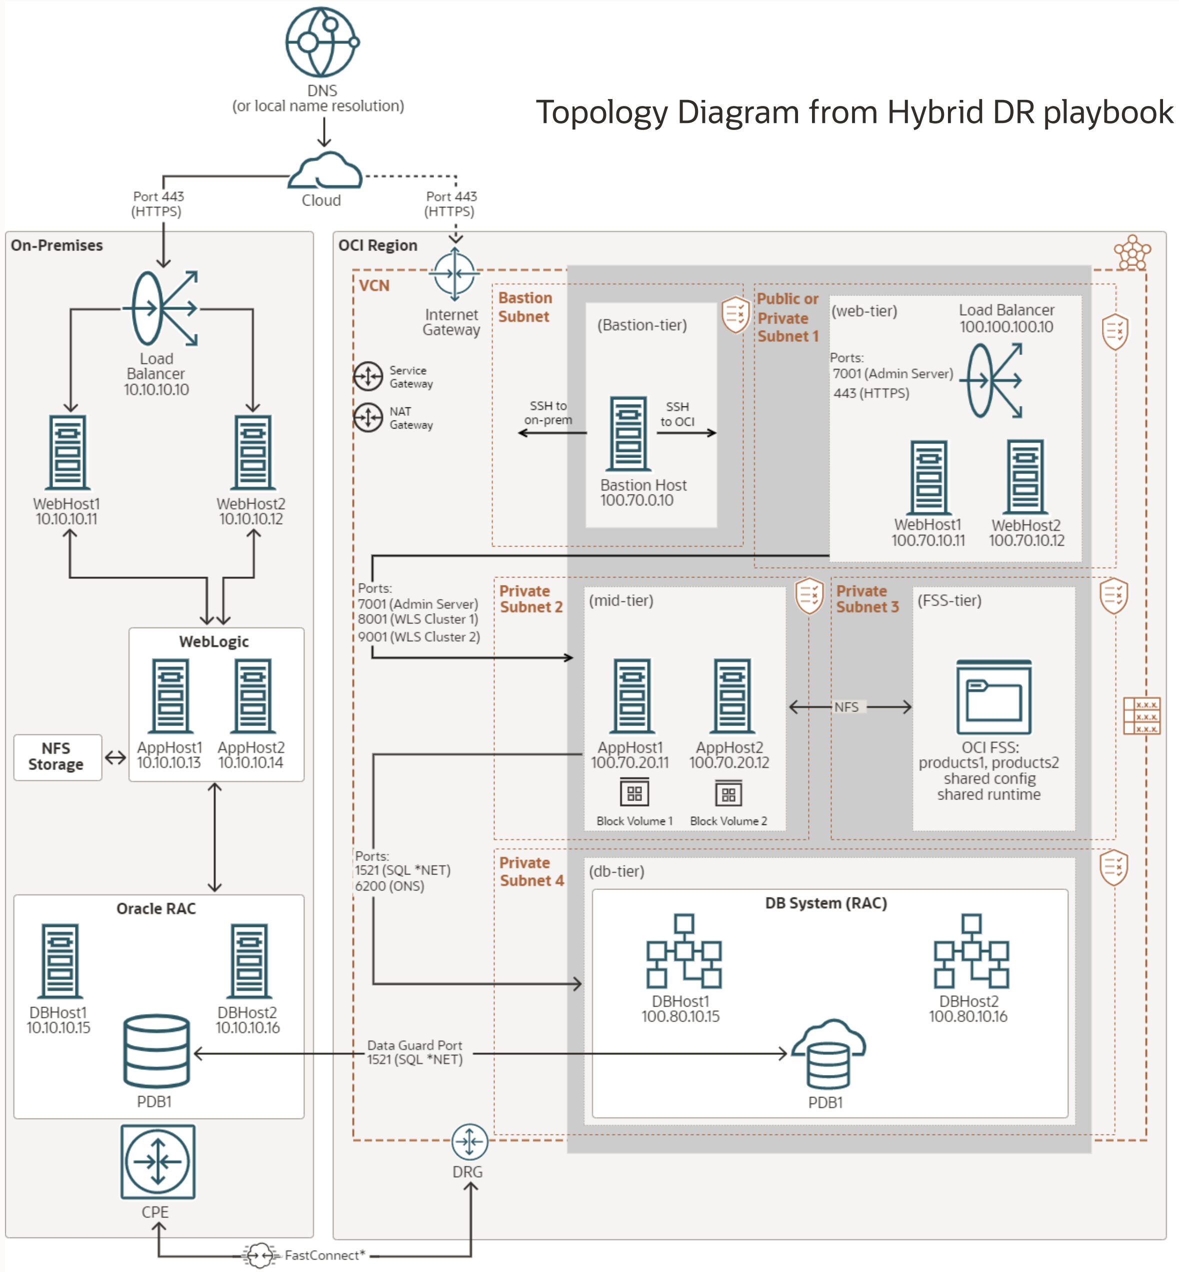

Topology Diagram from Hybrid DR playbook

Disaster Recovery: Active-Passive

Benefits

Benefits

- Protection for failures that affect to the entire site
- Supports high latencies between sites
- Ability for testing the standby without affecting primary (snapshot standby mode)
- Hybrid DR model can be used to migrate on premise environments to OCI (now automated)
- Topologies validated by MAA (DR tests and functional tests)

Limitations

- Replication for the file systems is required (in addition to the Data Guard)
- Higher RTO in a complete site outage than in stretched cluster model (because of the start of midtier processes)

Disaster Recovery: Active-Passive

Solutions **not** validated for MAA

Database replica based on Golden Gate	• GG not applicable to all data types. • GG is more oriented to data integration than to DR.
Complete recreation of the standby hosts from zero (VM replica) in each switchover	• Bad RTO compared with switchover to a pre-existing secondary. • High Risk operational model
Asymmetric standby	• Performance issues in standby • Unexpected application errors due to missing nodes
Completely stopped standby	• No validations, huge backlog of modifications

Disaster Recovery: Active-Passive

More Information

- FMW Disaster Recovery guide

- MW PaaS Disaster Recovery docs:
  - Oracle WebLogic Server for Oracle Cloud Infrastructure Disaster Recovery
  - SOA Suite on Oracle Cloud Infrastructure Marketplace Disaster Recovery
  - Configure FMW DR on OCI with an autonomous database
  - Use OCI Full Stack Disaster Recovery Service with Oracle WebLogic Server domains

- Hybrid DR playbooks
  - Configure a hybrid DR solution for Oracle WebLogic Server
  - Configure a hybrid DR solution for Oracle SOA Suite
- **New!** Framework in Github to automate the Hybrid DR setup process
  - https://github.com/oracle-samples/maa/tree/main/wls-hydr

Agenda

Maximum Availability Architecture for Middleware

1

Maximum Availability Architecture

Middleware MAA introduction and paradigms

2

High Availability

Middleware HA in a single datacenter for

- On premise
- Cloud

3

Disaster Recovery: Active-Active

Middleware Stretched clusters for

- On premise
- Cloud

4

Disaster Recovery: Active-Passive

Middleware A-P for

- On premise
- Cloud
- Hybrid
- Kubernetes

5

Summary Q&A

Disaster Recovery: A-P on Kubernetes

Description

Topology Diagram from K8 DR playbook

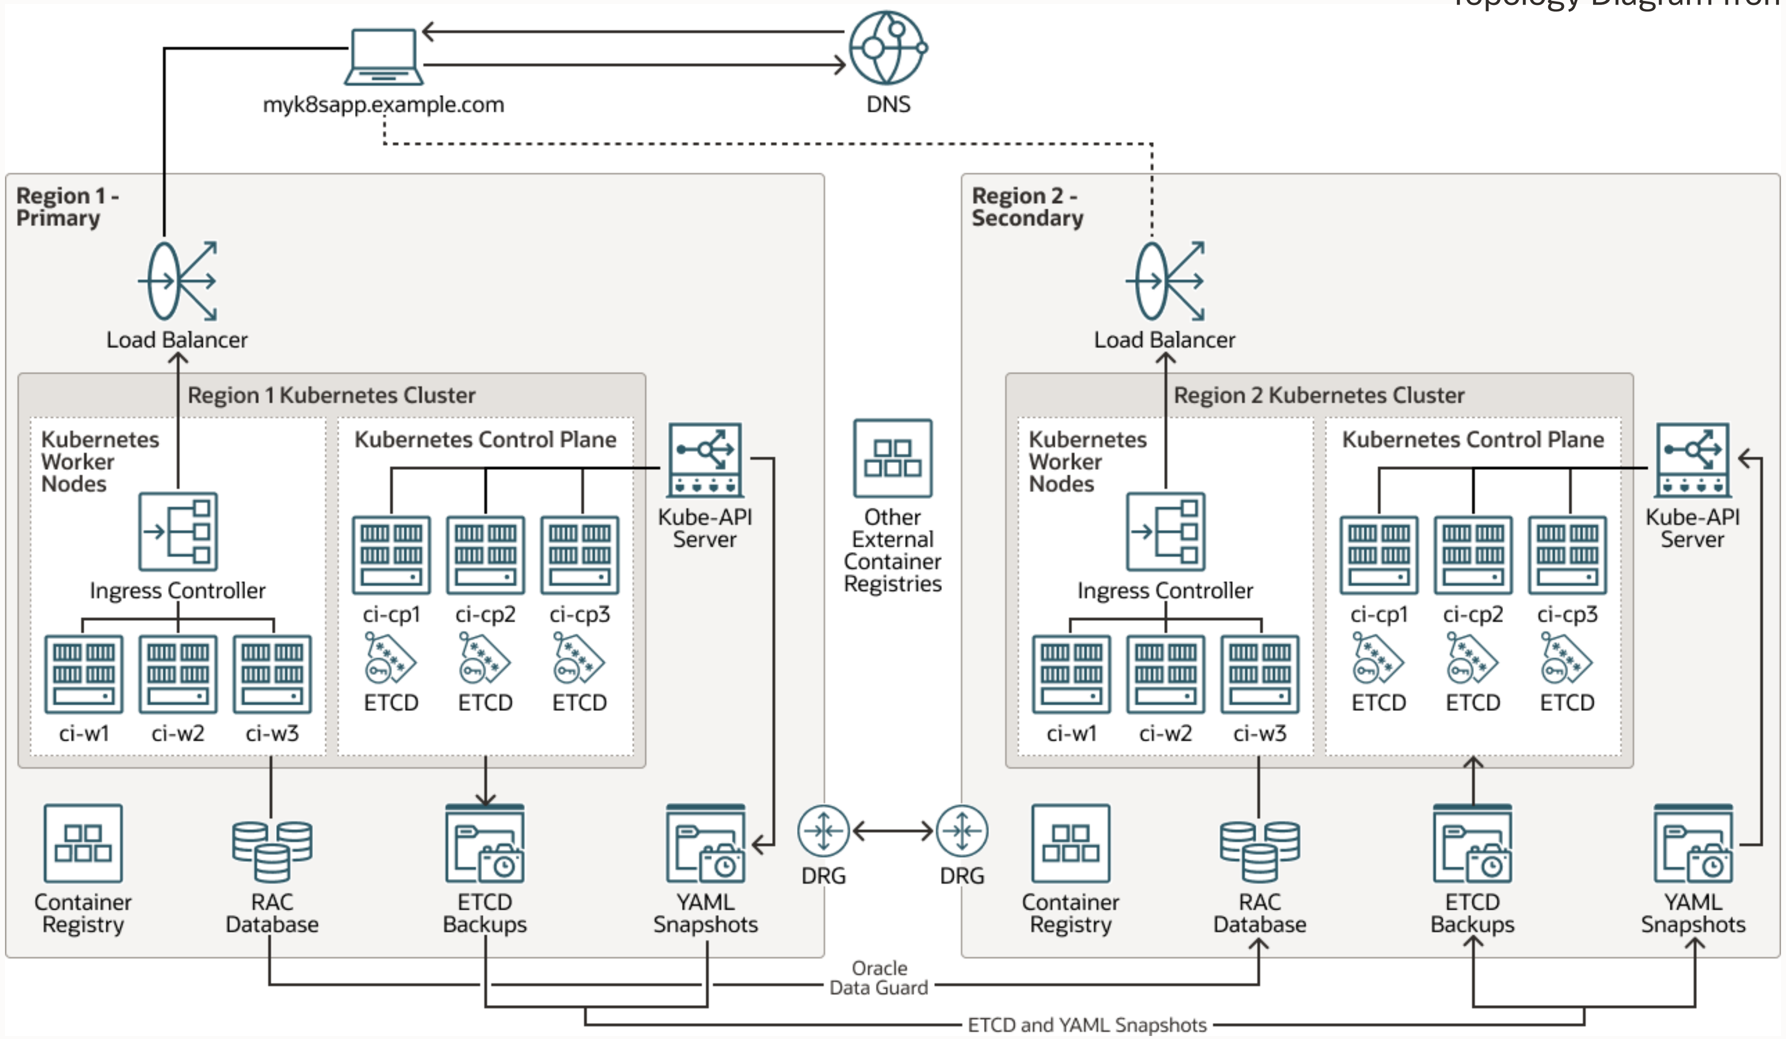

Disaster Recovery: A-P on Kubernetes

Description

- Active-Passive model, there are components that require its own DR protection as usual:
  - Database: Data Guard
  - Content in file systems: disk replication, rsync
  - Load Balancer: configure Load Balancer in each site
- The container images are like “binaries”:
  - Can reside locally in worker nodes
  - Can reside in totally external Container Registry
  - Can reside in Container Registry specific to each region
- For replicating the K8s resources, 3 main approaches:
  - Replicate etcd
  - Dual Apply
  - Extract & Apply
- Avoid polyglot persistence
  - Use Oracle DB (multitenancy if required) but keep all “interdependent” runtime data in a single store

Disaster Recovery: A-P on Kubernetes

Description

3 approaches for replicating the K8s resources:

Replicate etcd

- Backup (snapshot) etcd in primary “as is”
- Replicate snapshot to secondary

Dual apply

- Create and manage 2 separate K8s clusters
- Configure and deploy everything twice (pipelines, CI/CD)

Extract & apply

- Create 2 separate clusters
- Use a DR tool to replicate deployments (Rancher, Rackware, Velero++, Kasten, maak8 framework)

Disaster Recovery: A-P on Kubernetes

Description – etcd replica DR method

- **etcd** is a consistent and HA store for all K8s cluster data
  - PODs, services, secrets, config maps, daemon sets, deployments, etc.
  - Except the control plane components config.

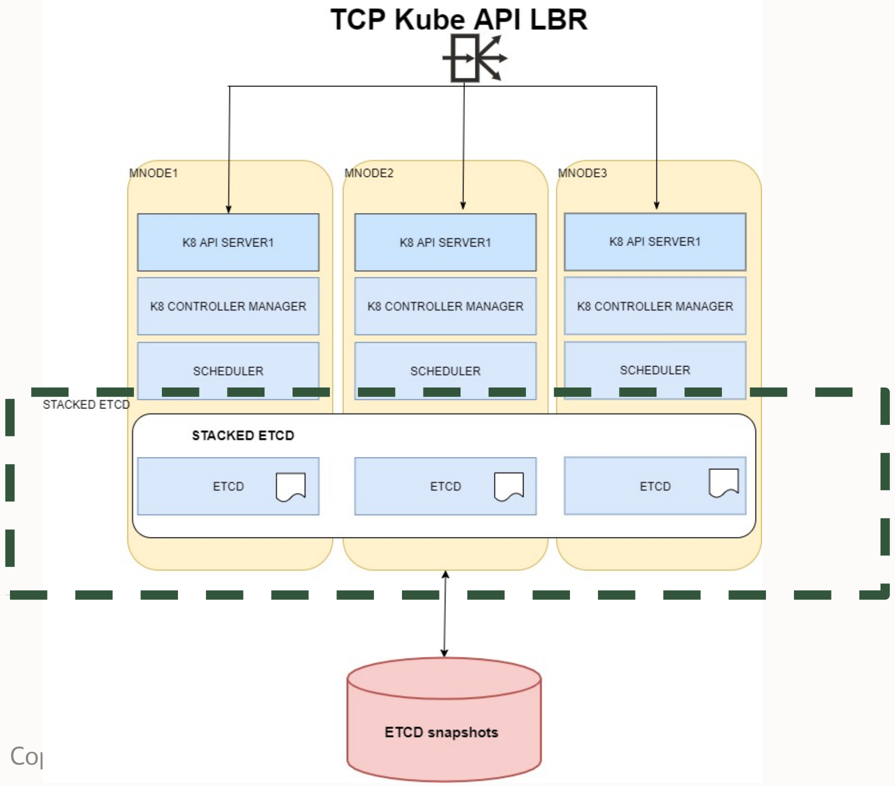

- You can replicate the etcd information to another K8s cluster for DR:
  - The standby K8s must have the same K8s configuration!

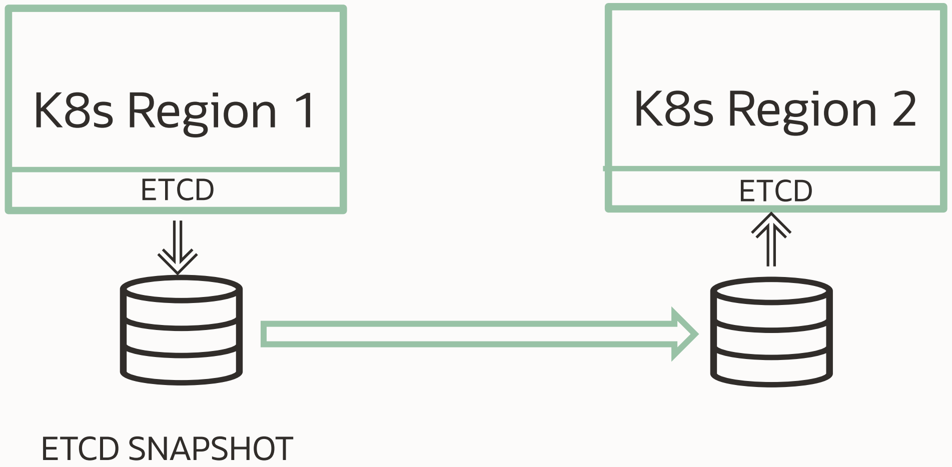

- This method can be used in
  - Generic/custom K8s clusters
- This method can't be used
  - In OKE (no access to etcd)

Disaster Recovery: A-P on Kubernetes

Description – Dual Apply DR method

- Manually configure the K8s clusters to be the same
- Then, apply any config changes (image, pod, label, allocation, etc.) at the same time in both regions.

- This method can be used in
  - Generic/custom K8s clusters
  - In OKE

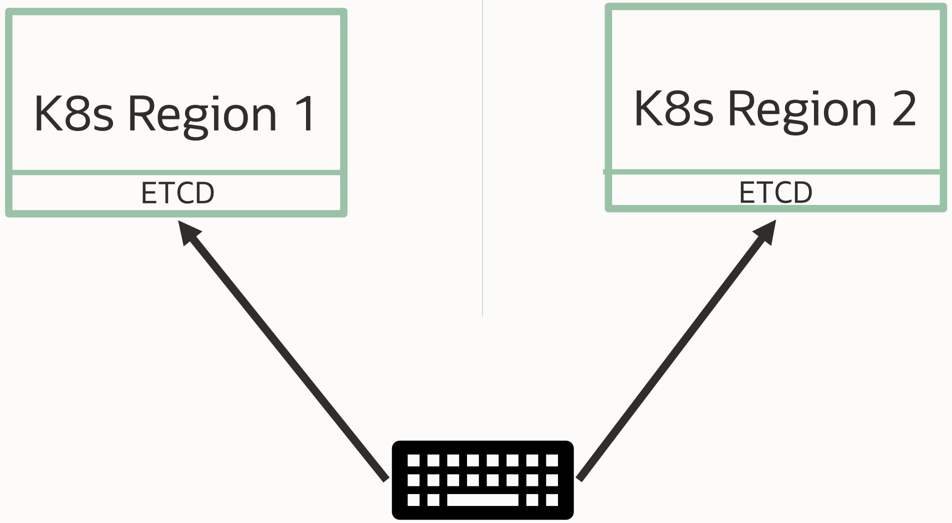

```
kubectl --kubeconfig K8sregion1.conf apply -f myapp.yaml
kubectl --kubeconfig K8sregion2.conf apply -f myapp.yaml
```

Disaster Recovery: A-P on Kubernetes

Description – Extract & Apply DR method

- Manually configure the K8s clusters to be the same
- Use 3rd party tool (Rancher, Rackware, Velero++, Kasten, maak8 framework) for the extracts/apply
- Performs extract and apply artifacts regularly

- This method can be used in
  - Generic/custom K8s clusters
  - In OKE

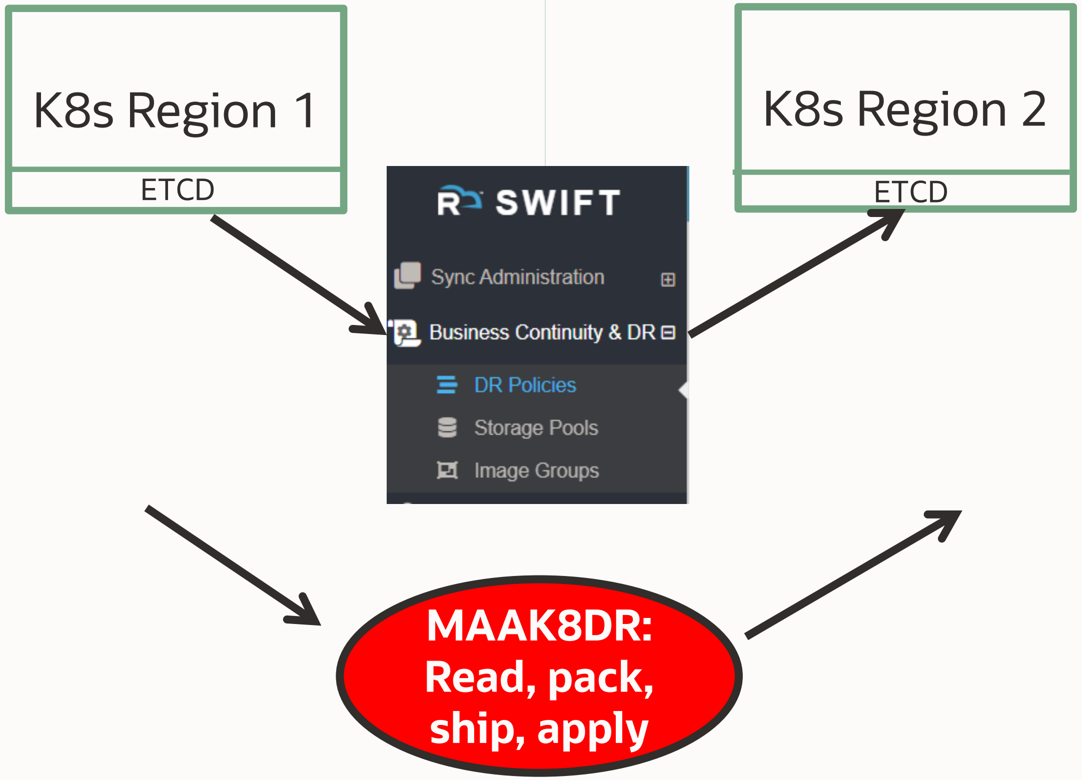

K8s Region 1

ETCD

K8s Region 2

ETCD

MAAK8DR: Read, pack, ship, apply

Oracle MAA team provides a framework for this approach since 2020

Disaster Recovery: A-P on Kubernetes

Benefits

Etcd replication

- Benefits
  - Total consistency (not only K8 config and deployments are mirrored, also settings for node, pod allocations..)
- Limitations
  - This method **can't** be used in OKE (there is no access to etcd)
  - Requires exact mirror: same K8s cluster/resource/hosts name alias, mem/cpu settings, workers, etc.

Dual Apply

- Benefits
  - Flexible: primary and standby K8s cluster may differ
  - Better RTO and RPO that others
  - Can be used in generic/custom K8s and in OKE
- Limitations
  - Consistency maintenance is manual
  - DR creation on day N
  - All the artifacts may not be able to deploy in standby (e.g. because DB is in standby)

Extract & Apply

- Benefits
  - Easy for large systems
  - Flexible: allows differences between primary and standby K8s clusters
  - Can be used in generic/custom K8s and in OKE
- Limitations
  - Possible inconsistencies if not all dependencies or namespaces are replicated exactly

MAA recommends using a combination of etcd replication (for quick local recovery and flashback) and Extract & Apply (for multi-flavor cluster DR)

Disaster Recovery: A-P on Kubernetes

More information

- Use artifact snapshots to protect your Kubernetes Clusters from disaster
- Kubernetes Clusters restore based on etcd snapshots

Agenda

Maximum Availability Architecture for Middleware

1

Maximum Availability Architecture

Middleware MAA introduction and paradigms

2

High Availability

Middleware HA in a single datacenter for

- On premise
- Cloud

3

Disaster Recovery: Active-Active

Middleware Stretched clusters for

- On premise
- Cloud

4

Disaster Recovery: Active-Passive

Middleware A-P for

- On premise
- Cloud
- Hybrid
- Containers

5

Summary Q&A

Summary of documents and resources

	On Premise	Oracle Cloud Infrastructure	Hybrid
High Availability	Enterprise Deployment Guides (SOA, OBI, Webcenter Content, Webcenter Portal, IAM)	PaaS Services documentation (WLS for OCI, SOA Suite on Marketplace) Enterprise Deployment Guides (SOA, OBI, Webcenter Content, Webcenter Portal, IAM)	Enterprise Deployment Guides (SOA, OBI, Webcenter Content, Webcenter Portal, IAM)
Disaster Recovery A-A	Best Practices for Oracle Fusion Middleware SOA 12c Multi Data Center Active-Active Deployment	PaaS Services documentation (WLS for OCI, SOA Suite on Marketplace)	n/a
Disaster Recovery A-P	Oracle Fusion Middleware Disaster Recovery Guide Use artifact snapshots to protect your Kubernetes Clusters from disaster Kubernetes Clusters restore based on etcd snapshots	Oracle WebLogic Server for Oracle Cloud Infrastructure Disaster Recovery SOA Suite on Oracle Cloud Infrastructure Marketplace Disaster Recovery Configure FMW DR on OCI with an autonomous database Use OCI Full Stack Disaster Recovery Service with Oracle WebLogic Server domains Use artifact snapshots to protect your Kubernetes Clusters from disaster Kubernetes Clusters restore based on etcd snapshots	Configure a hybrid DR solution for Oracle WebLogic Server Configure a hybrid DR solution for Oracle SOA Suite https://github.com/oracle-samples/maa/tree/main/wls-hydr

Q&A

ORACLE