flashgrid

Oracle RAC in the Cloud: Options, Challenges, Solutions

NoCOUG Conference Spring 2017

Alex Miroshnichenko
CEO, FlashGrid
May 18, 2017

ORACLE Gold Partner Cloud Standard

About FlashGrid

- Specialists in hyper-converged software-defined storage architecture based on commodity storage and compute resources
- Started in 2015 with an on premise SDS product for Oracle RAC
- Applied the technology to cloud environments in 2016
- Oracle Gold Partner / Cloud Standard
- HQ in Sunnyvale, CA

Growing Demand for HA databases in the Cloud

Advantages of Running Oracle RAC in the Cloud

flashgrid

- All infrastructure moving to the cloud ⇒ Databases must move to the cloud too
- Already using Oracle RAC? Minimize risks, keep using it
- Not using Oracle RAC yet? Maximize database HA in the cloud with Oracle RAC

Until recently, Oracle RAC was the last major component of enterprise IT environment without a clear cloud migration methodology.

Public Cloud Infrastructure Challenges

- No shared block storage
- The fastest storage (local SSD) typically is not persistent
- No network multicast
- Limited network bandwidth between VMs
- Single network pipe for all traffic types
- No Virtual IP support

RAC in Cloud: The Shared Storage Challenge

FlashGrid Storage Fabric

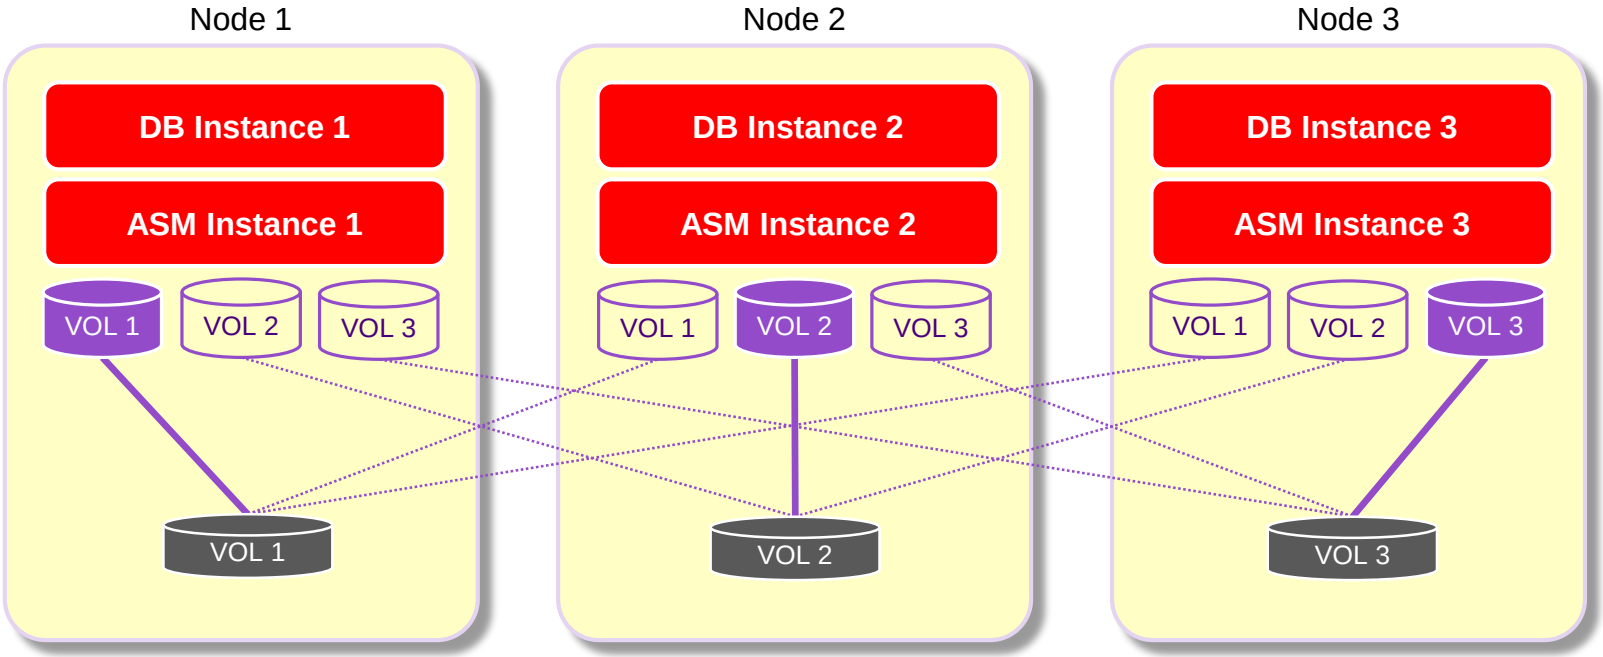

- Oracle ASM manages data, volumes, mirroring, snapshots
- FlashGrid manages storage devices and connections

ASM Disk Group Configuration Example: 3 Nodes

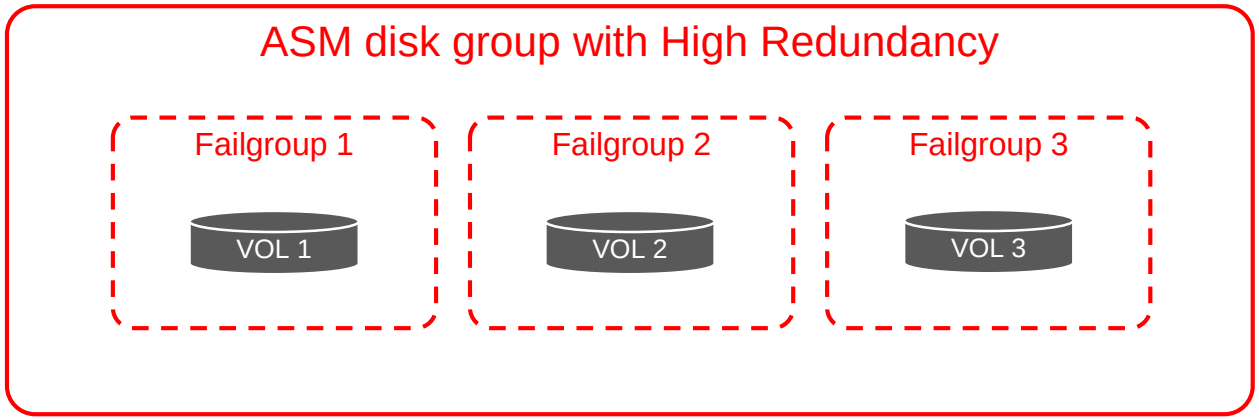

- Volumes from nodes are placed into different ASM failure groups

Multiple Volumes Per Node

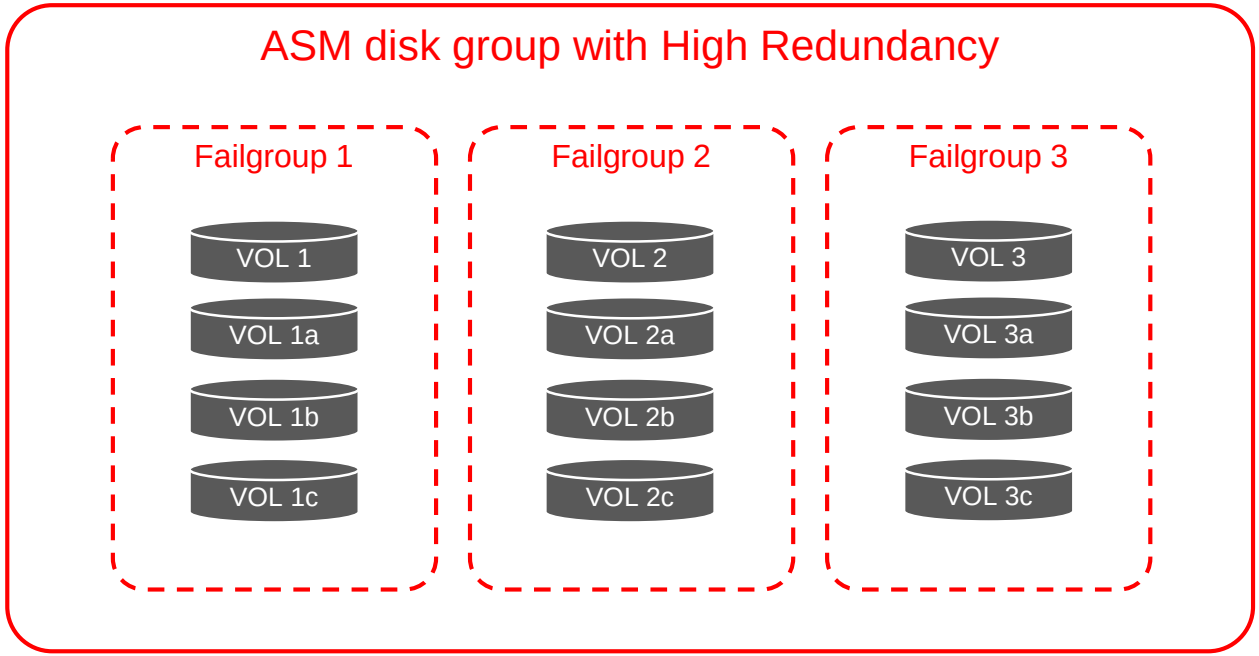

- More than one volume per node is allowed
- Volumes from each node are added into the same ASM failure group

Up to 16 GB/s of Storage Bandwidth with Local NVMe SSD

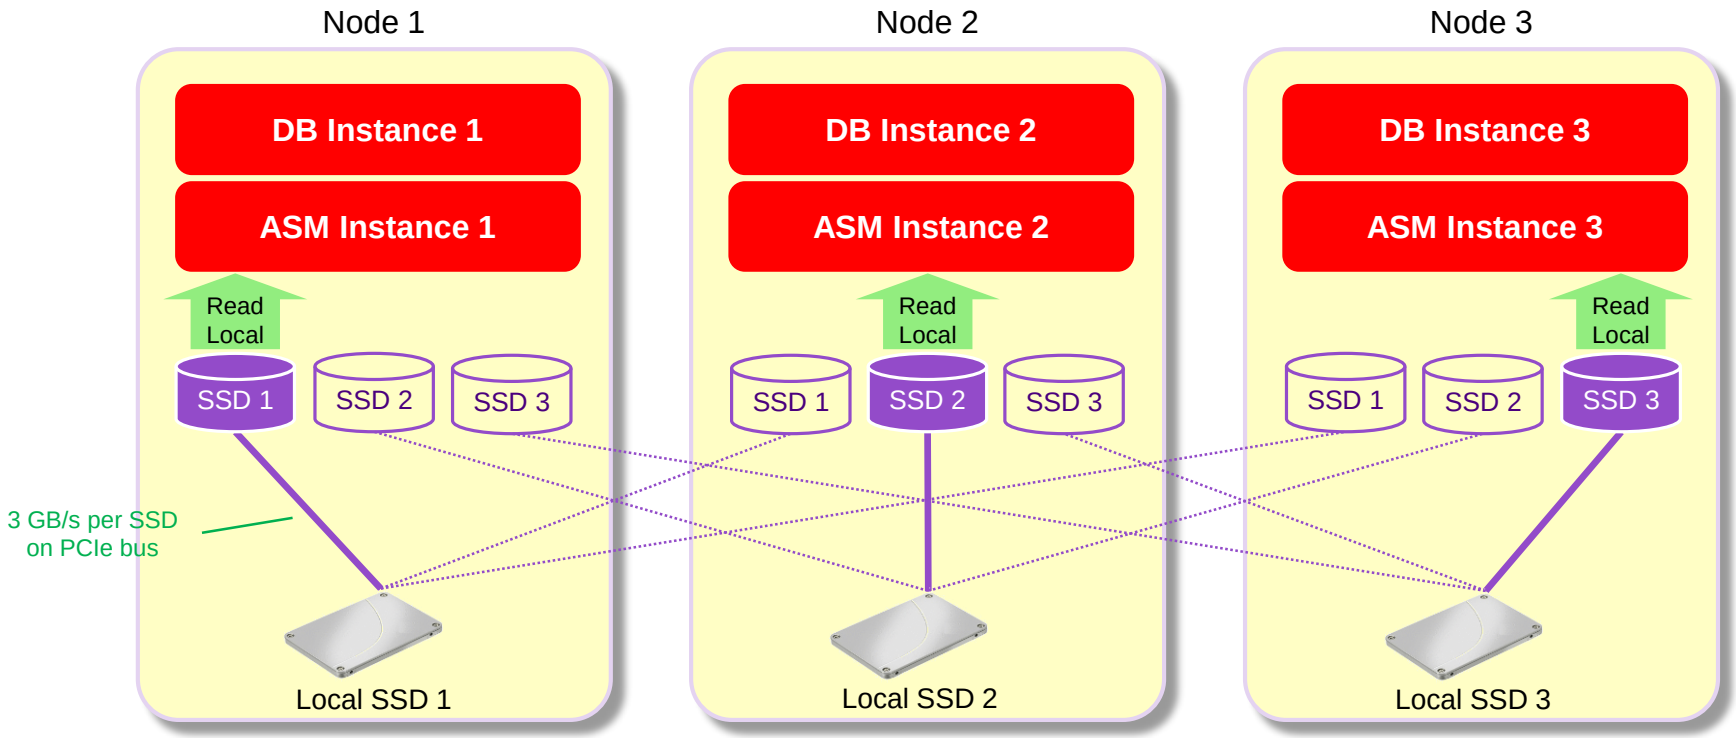

- Local NVMe SSDs available in AWS, GCP, Oracle Bare Metal, Oracle Compute Cloud
- Up to 8 SSDs per node

RAC in Cloud: The Network Challenges

Virtual Network Overlay for Multicast

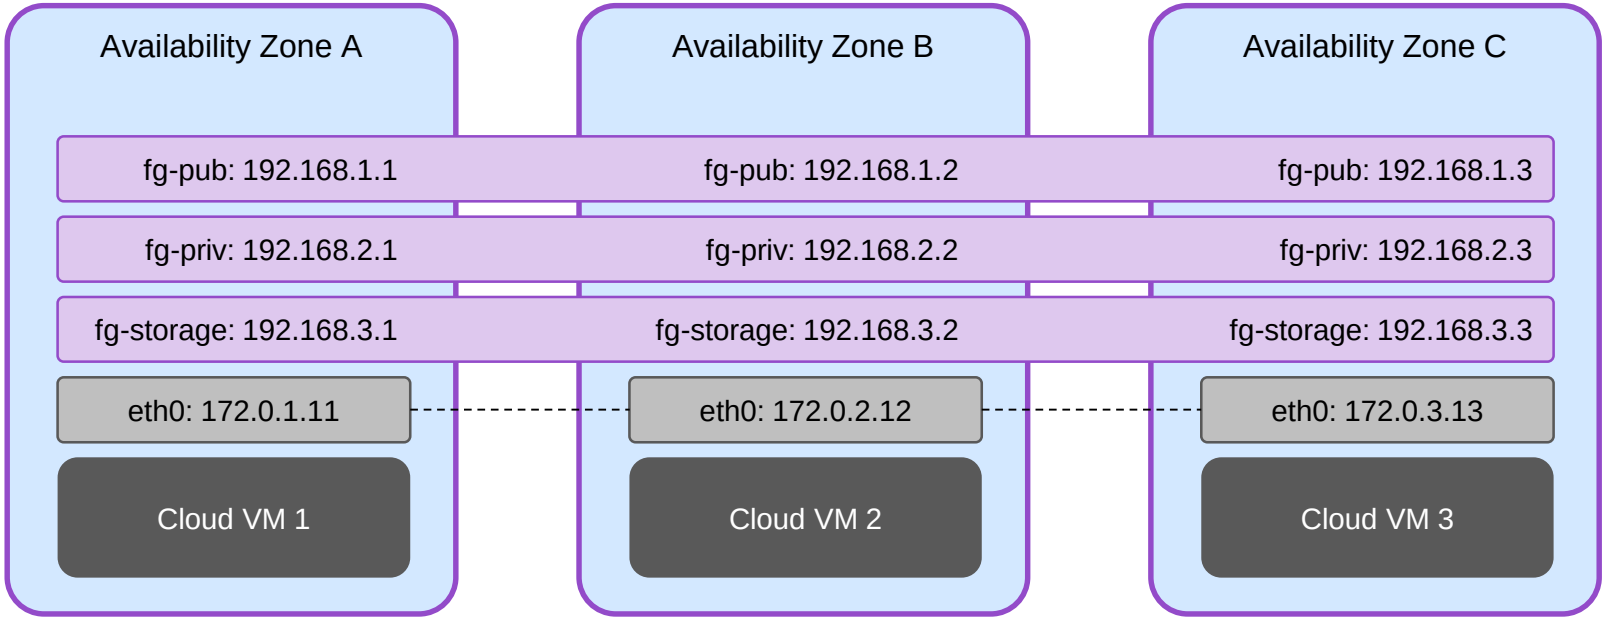

- High-speed network overlay with multicast and QoS
- Separate CLAN subnets for each type of traffic
- Transparent connectivity across availability zones

Setting it Up: Turning Cloud VMs and Storage into a Working Oracle Real Application Cluster

Oracle RAC in AWS with FlashGrid

Oracle RAC in AWS with FlashGrid

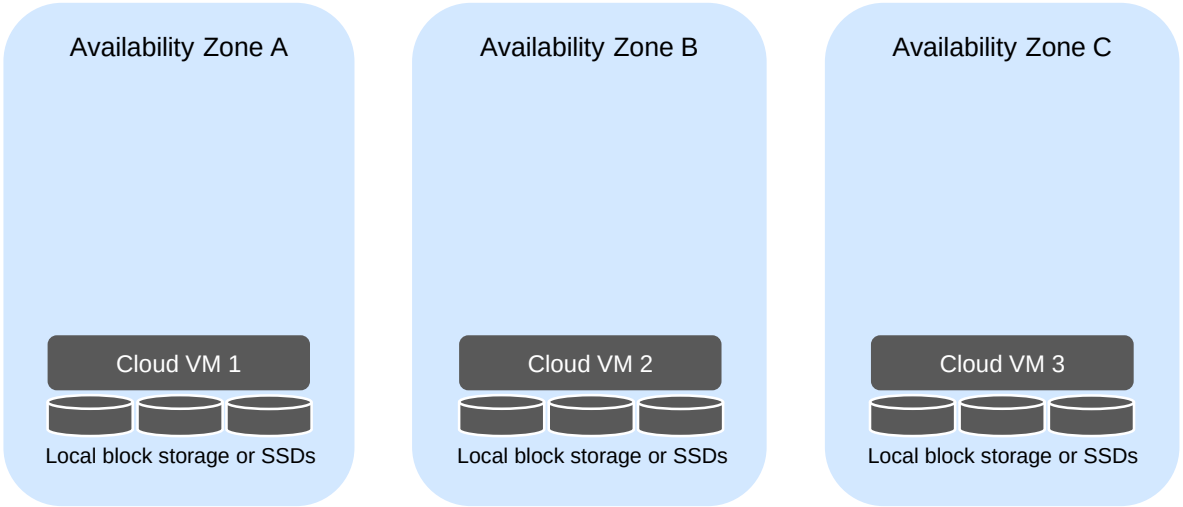

Oracle RAC in AWS with FlashGrid

FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS

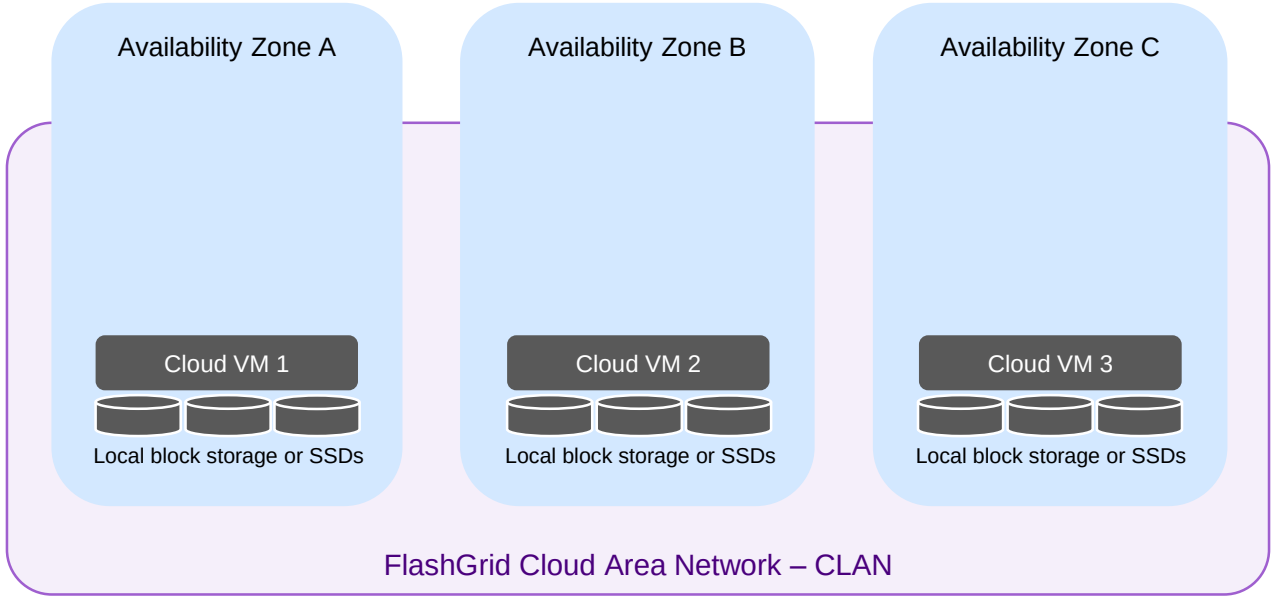

Oracle RAC in AWS with FlashGrid

FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS

FlashGrid Storage Fabric creates shared storage from local drives (elastic or local SSDs)

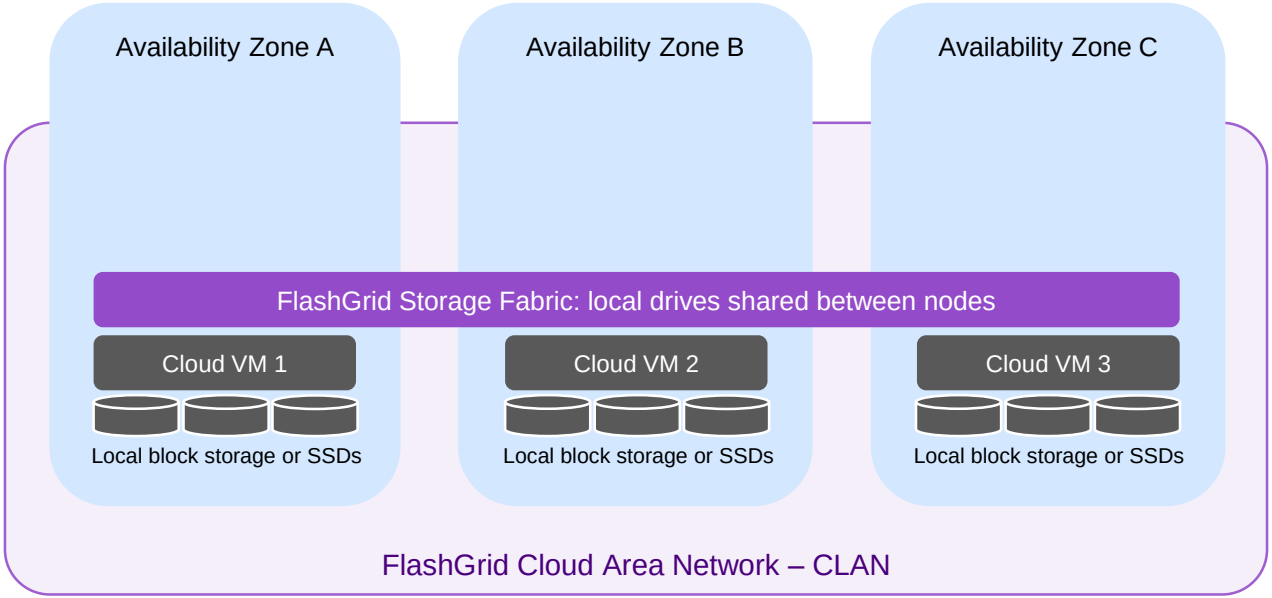

Oracle RAC in AWS with FlashGrid

FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS

FlashGrid Storage Fabric creates shared storage from local drives (elastic or local SSDs)

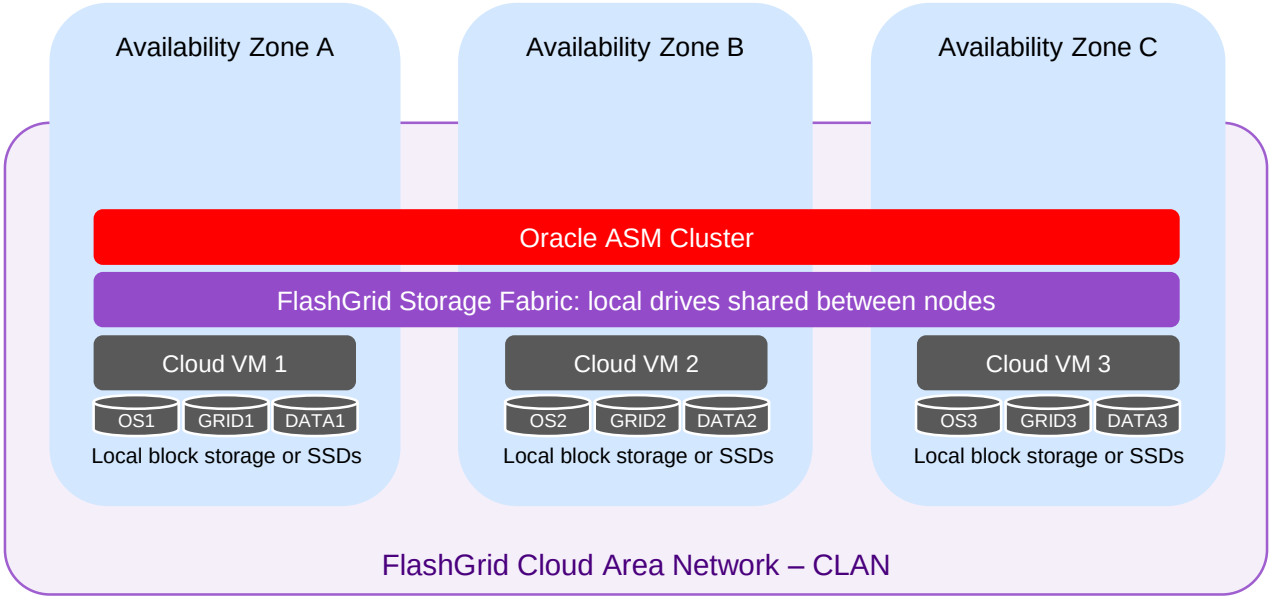

Oracle RAC in AWS with FlashGrid

FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS

FlashGrid Storage Fabric creates shared storage from local drives (elastic or local SSDs)

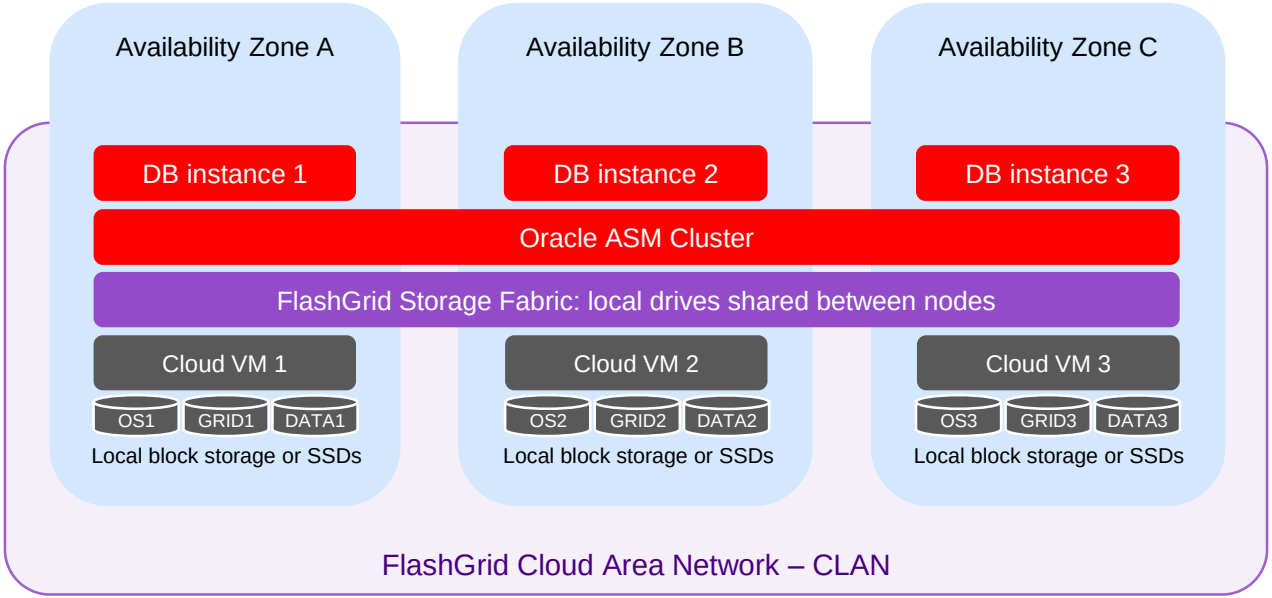

Typical total RAC deployment time in AWS : 90 minutes

Cloud-Agnostic

Microsoft
AZURE

Google
GCP

Examples of 2- and 3-Node Clusters, Multiple and Single Availability Zones

2 RAC Nodes, 3 Availability Zones

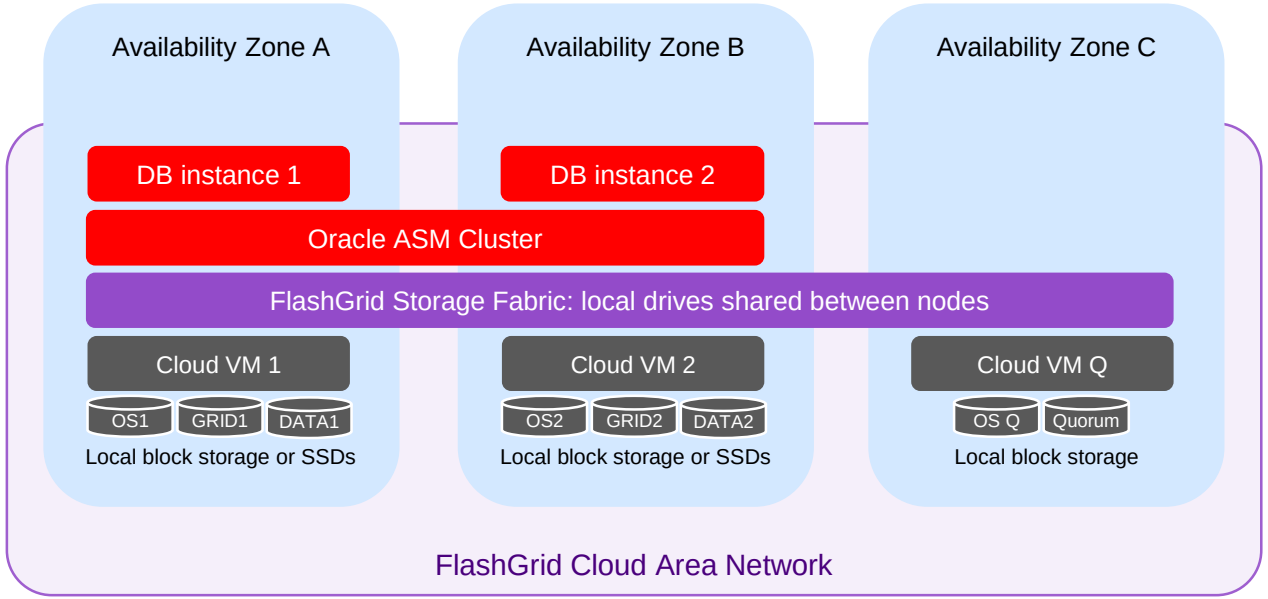

- FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS
- FlashGrid Storage Fabric creates shared storage from local drives (elastic or local SSDs)
- Leverage proven Oracle ASM for high availability and data mirroring
- On any public cloud – virtual or bare metal

3 RAC Nodes, 3 Availability Zones

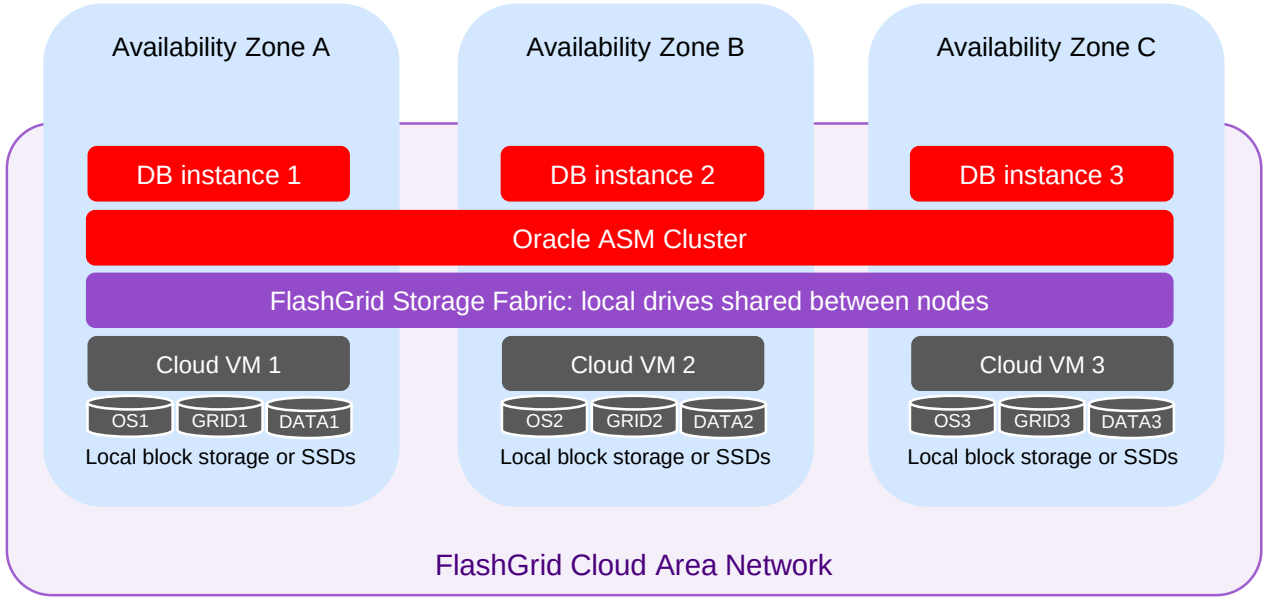

- FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS
- FlashGrid Storage Fabric creates shared storage from local drives (elastic or local SSDs)
- Leverage proven Oracle ASM for high availability and data mirroring
- On any public cloud – virtual or bare metal

2 RAC Nodes, 1 Availability Zone

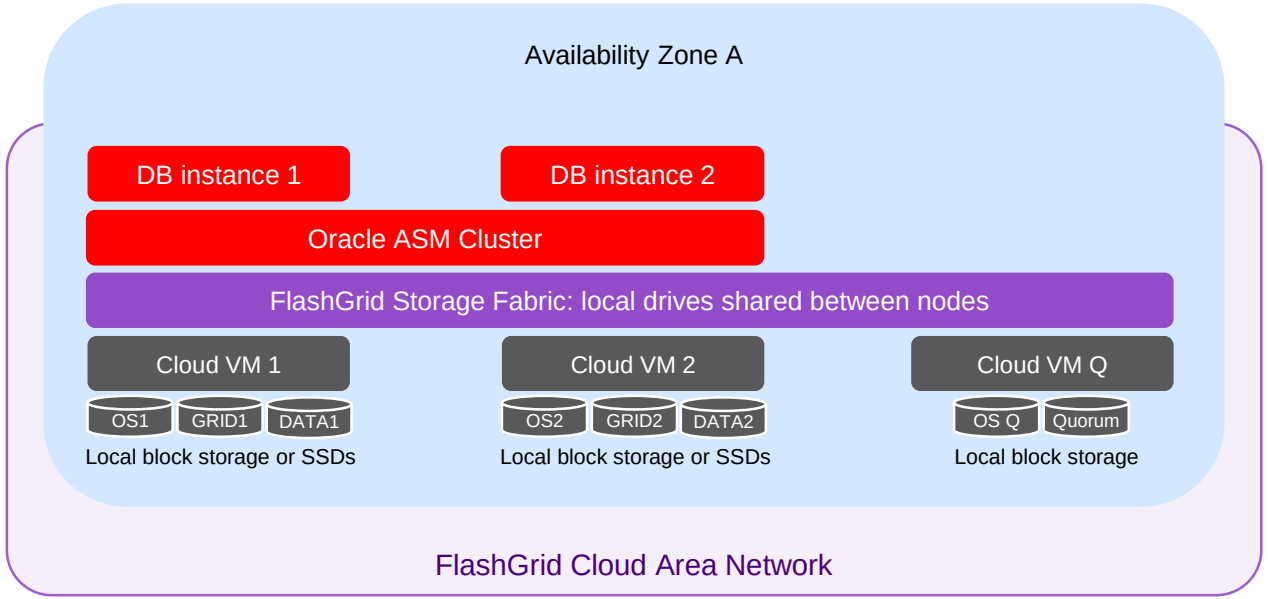

- FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS
- FlashGrid Storage Fabric creates shared storage from local drives (elastic or local SSDs)
- Leverage proven Oracle ASM for high availability and data mirroring
- On any public cloud – virtual or bare metal

3 RAC Nodes, 1 Availability Zone

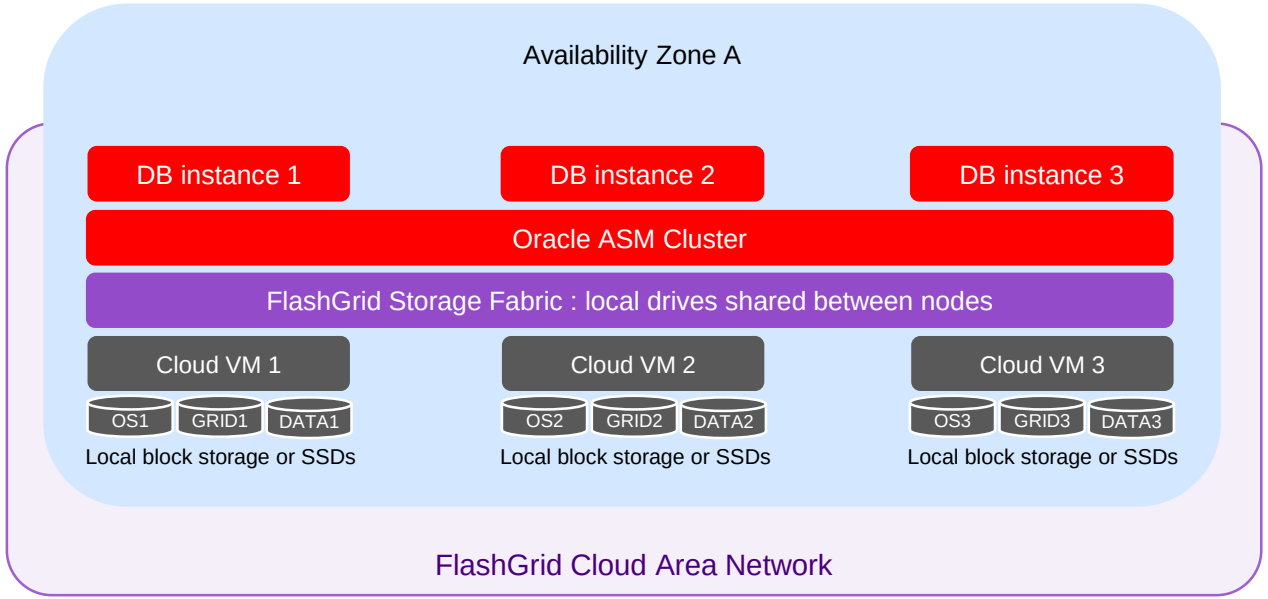

- FlashGrid Cloud Area Network creates a high-speed network overlay with multicast and QoS
- FlashGrid Storage Fabric creates shared storage from local drives (elastic or local SSDs)
- Leverage proven Oracle ASM for high availability and data mirroring
- On any public cloud – virtual or bare metal

Performance

Performance | Azure

- Oracle Real Application Cluster (RAC) with two hyper-converged nodes
- 2 x DS15_V2 VMs with sixteen 513 GB Premium SSD disks each
- Oracle Linux 7.3 with Oracle Grid Infrastructure 12.1 and Oracle Database 12.1

- Calibrate_IO
  - 121,597 IOPS
  - 1.3 GB/s bandwidth
  - Latency below 1ms

- SLOB IOPS

	2-node RAC, both nodes combined	Single-instance
Read+Write Database Requests	53,839	52,990
Read Database Requests	43,020	43,159
Write Database Requests	10,819	9,831

Performance | Amazon Web Services

- Oracle Real Application Cluster (RAC) with two hyper-converged nodes
- 2 x M4.16xlarge instances with four io1 20,000 IOPS 400GB volumes each
- Oracle Linux 7 with Oracle Grid Infrastructure 12.1 and Oracle Database 12.1

- Calibrate_IO
  - 154,864 IOPS
  - 2.2 GB/s bandwidth
  - Latency below 1ms

- SLOB, nodes in different availability zones
  - 92,081 IOPS for physical reads
  - 19,465 IOPS for writes
  - 111,546 IOPS combined

- SLOB, nodes in the same availability zone
  - 121,237 IOPS combined, 8% increase of performance
  - Interconnect Ping latency 2x-3x shorter, average 0.23 ms vs 0.69 ms

Performance | Oracle Bare Metal Cloud Services

flashgrid

- Oracle Real Application Cluster (RAC) with three hyper-converged nodes
- 3 x DenseIO1.36 instances with nine 3.2TB NVMe SSDs each
- Oracle Linux 7 with Oracle Grid Infrastructure 12.1 and Oracle Database 12.1

- Calibrate_IO
  - 6.4 million IOPS
  - 55 GB/s bandwidth
  - Latency below 1ms

- SLOB
  - 537,000 IOPS for physical reads
  - 120,000 IOPS for writes
  - 657,000 IOPS combined

Compare:

Dell EMC XtremIO X2
All-flash SAN
220,000 IOPS
6 GB/s peak bandwidth

StorageReview.com, May 2017

Oracle RAC in The Cloud – Easier than On-Premise

flashgrid

- No need to purchase server or storage hardware
- Create a fully configured cluster with a few mouse clicks
- Easily modify VM sizes or storage capacity when needed

Further Reading

- AWS Products & Solutions Article “Oracle RAC on Amazon EC2”
  https://aws.amazon.com/articles/7455908317389540

- White paper: “Mission-Critical Databases in the Cloud – Oracle RAC in Microsoft Azure”
  https://www.flashgrid.io/wp-content/sideuploads/resources/FlashGrid_OracleRAC_in_Azure.pdf

- www.flashgrid.io

Questions please!

alex at flashgrid dot io