



Storage Developer Conference

December 4-5, 2020

BY Developers FOR Developers

Mission Critical Databases - On-premises vs Cloud

Amol Bhoite
Xoriant Solutions Pvt. Ltd.

<https://www.linkedin.com/in/amol-bhoite-54960b5/>



Agenda

- Terminology
- Database On-premises vs Cloud Considerations
 - High Availability and Reliability
 - Scalability
 - Manageability
 - Control and Compliance
 - Security
 - Cost/Performance
- Database Migration
- Case Study
- Conclusion



Terminology

Mission Critical Databases?



Choice of Scalable Databases

Developer



Development

Small Business



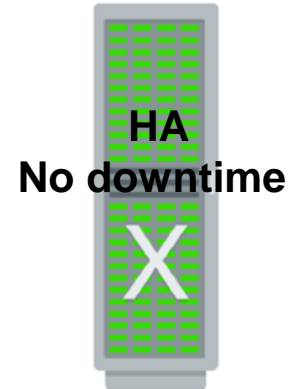
Small business/
departmental
applications

Medium Business



Medium business
applications

Enterprise



Large enterprise
applications/Mission
critical workloads

Data Centers

On-premises



Private Cloud



Customer Data Center

Public Cloud

Managed

24-7 active monitoring and management. Backup, recovery and patching are fully managed with opt-in time windows.

Automated

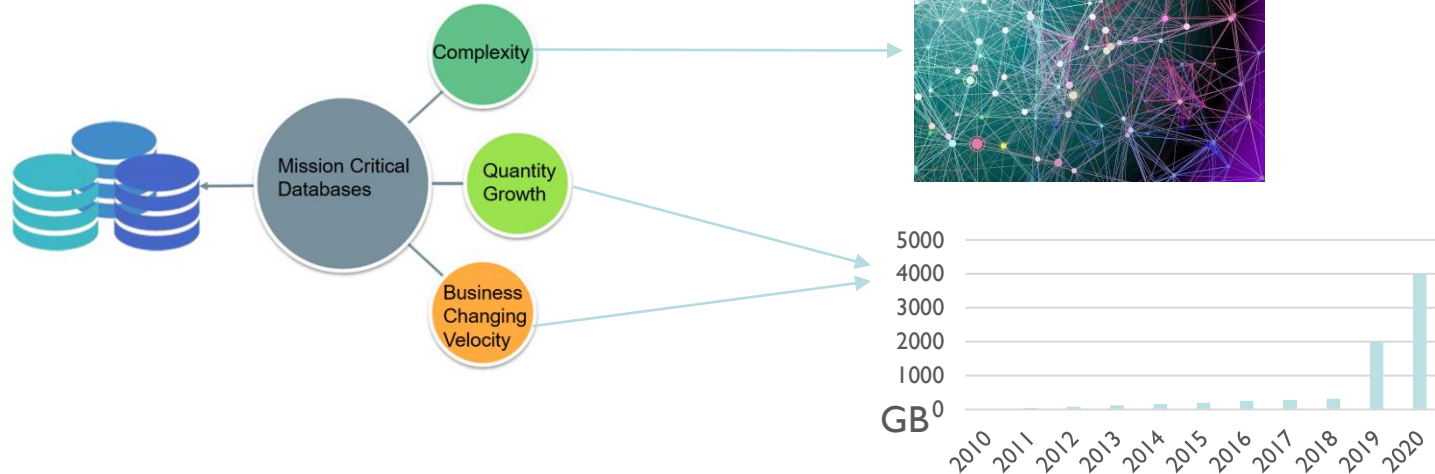
Automated install, patch, upsize / downsize, backup / restore / recovery, configuration & monitoring.

Virtual Image

Database software ready to install. Same software as distributed for on premise use.

IaaS | PaaS | SaaS

Why Database matters?



If we don't well prepare

- Loss of data
- Productivity penalty
- Contractual Liability
- Large scale customer impact if not available
- Traffic Drop
- Affect DB Performance
- Service Unavailable
- Panic!!!

Databases

Database Type	AWS	Azure	GCP	On-premises
Relational/SQL	Amazon Aurora Amazon RDS PostgreSQL, MySQL, MariaDB, Oracle,SQL Server Amazon Redshift	SQL Database - MySQL, PostgreSQL, MariaDB	Cloud SQL – MySQL, PostgreSQL, SQL Server, Oracle, SAP Cloud Spanner	MySQL, PostgreSQL, SQL Server, Oracle, SAP
Key-value	Amazon DynamoDB	Cosmos DB Table Storage	Cloud Bigtable Cloud Firestore Firebase Realtime Database	Redis, Oracle NoSQL, InfinityDB
In-memory	Amazon ElastiCache for Memcached, for Redis	Redis Cache	Cloud Memorystore	Oracle, SAP
Document	Amazon DocumentDB (with MongoDB compatibility)	-	-	MongoDB
Wide-column	Amazon Keyspaces (for Apache Cassandra)	-	-	Cassandra
Graph	Amazon Neptune	-	-	Neo4j
Time-series	Amazon Timestream	-	-	InfluxDB
Ledger	Amazon QLDB	-	-	Hyperledger
Migration service	AWS DMS	Azure DMS	GCP DMS	-

Databases On-premises vs cloud

Considerations

- Reliability and Availability
- Scalability
- Manageability
- Control and Compliance
- Security
- Cost/Performance



Databases High Availability and Reliability

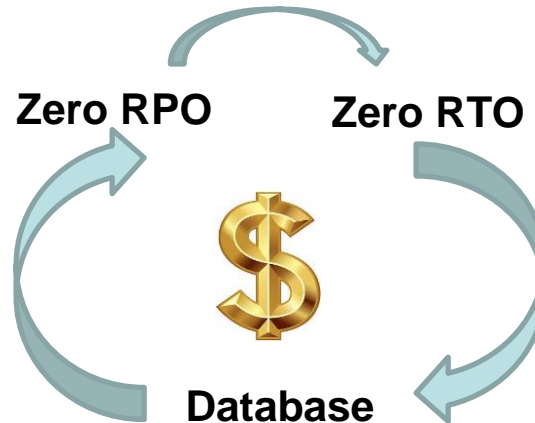
Databases High Availability and Reliability?

Availability

Percentage of time database remains operations under normal circumstances in order to server its indented purpose

Reliability

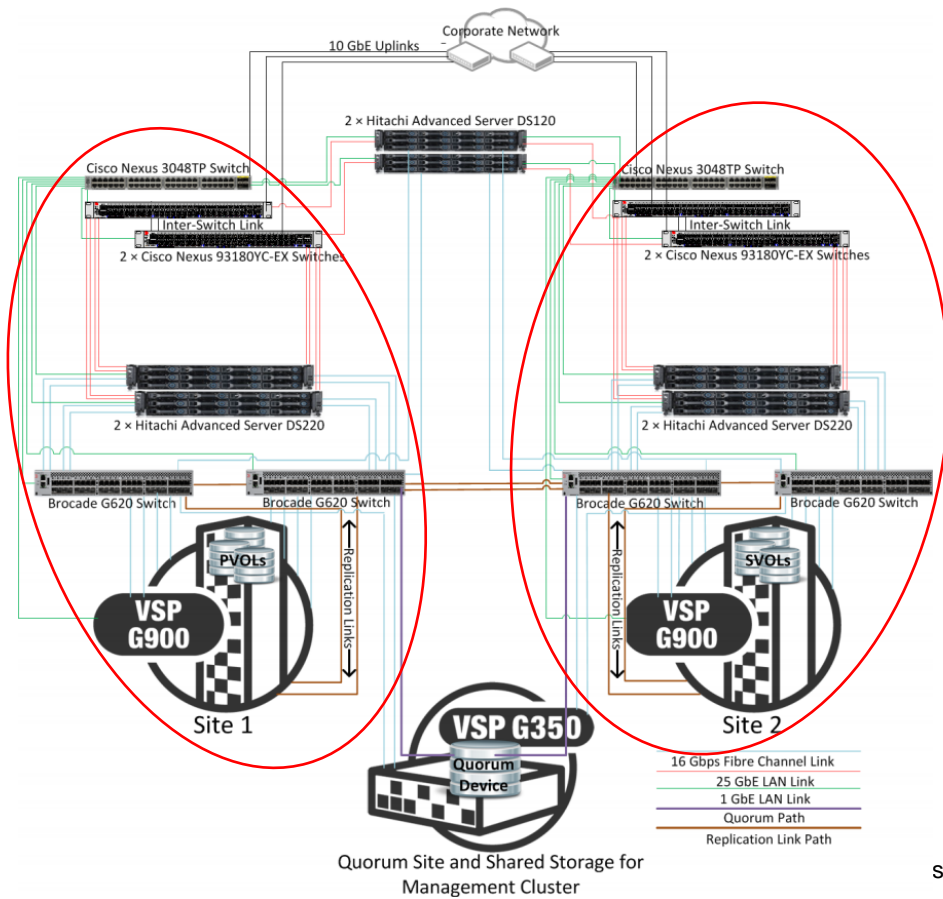
Probability that database will meet certain performance standard when yielding correct output for a set of period time



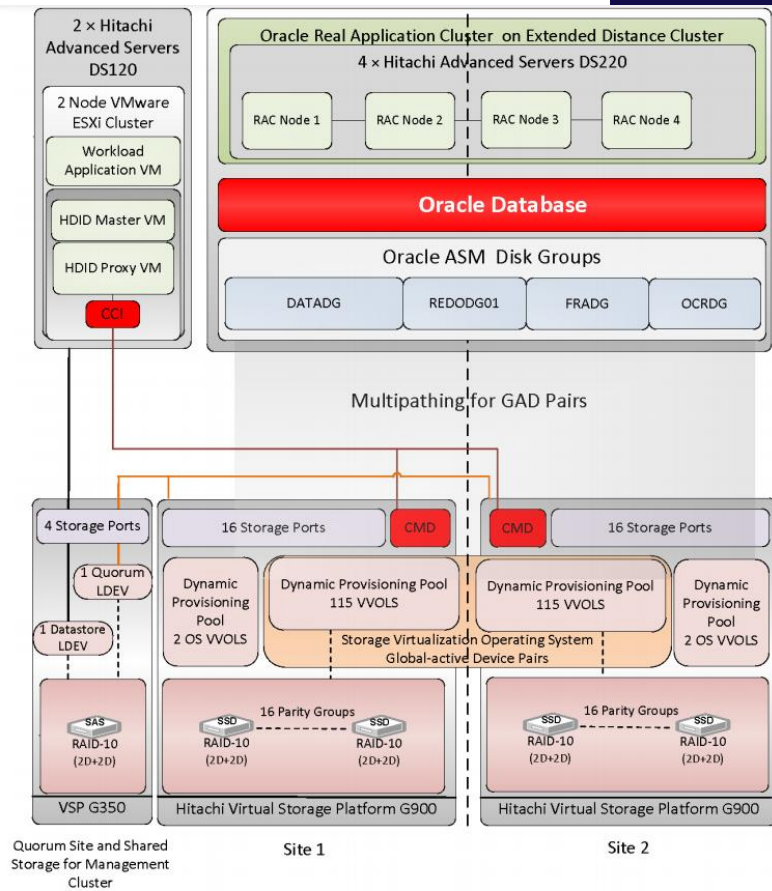
On-premises Database Solution

Hitachi Solution for Databases - Oracle RAC with Global-Active Device using Ops Center Protector

High level Infrastructure

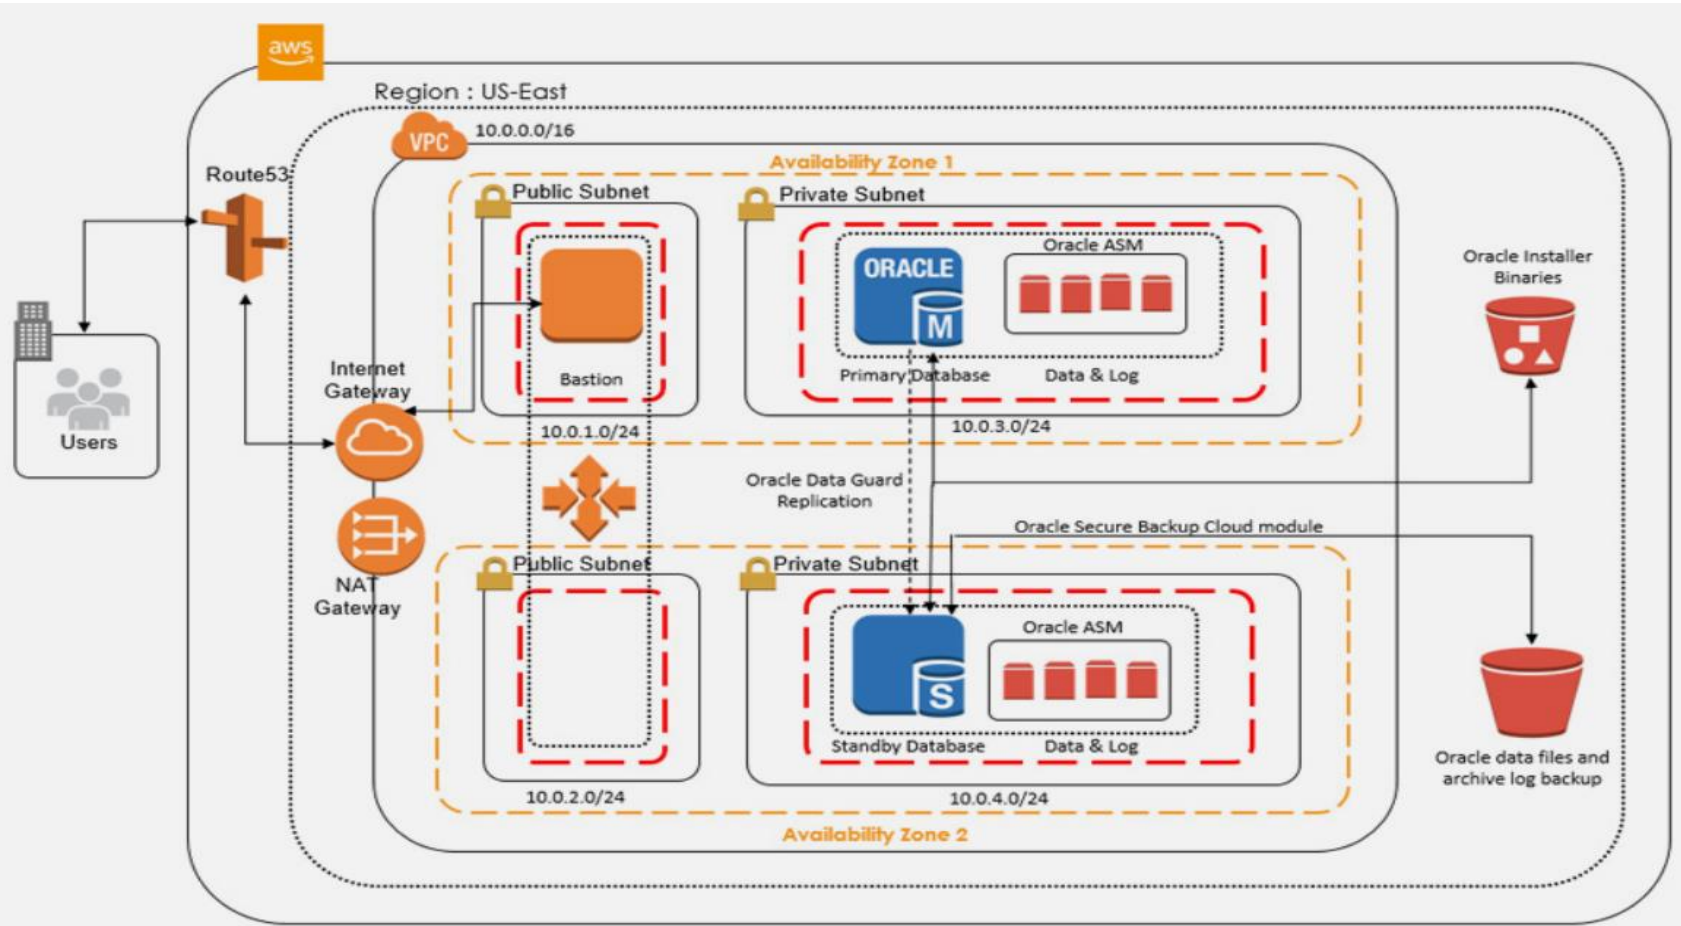


Storage Architecture



AWS cloud solution for Oracle database

Source: <https://www.ctepl.com/oracle-database-high-availability-solution-on-aws-platform/>



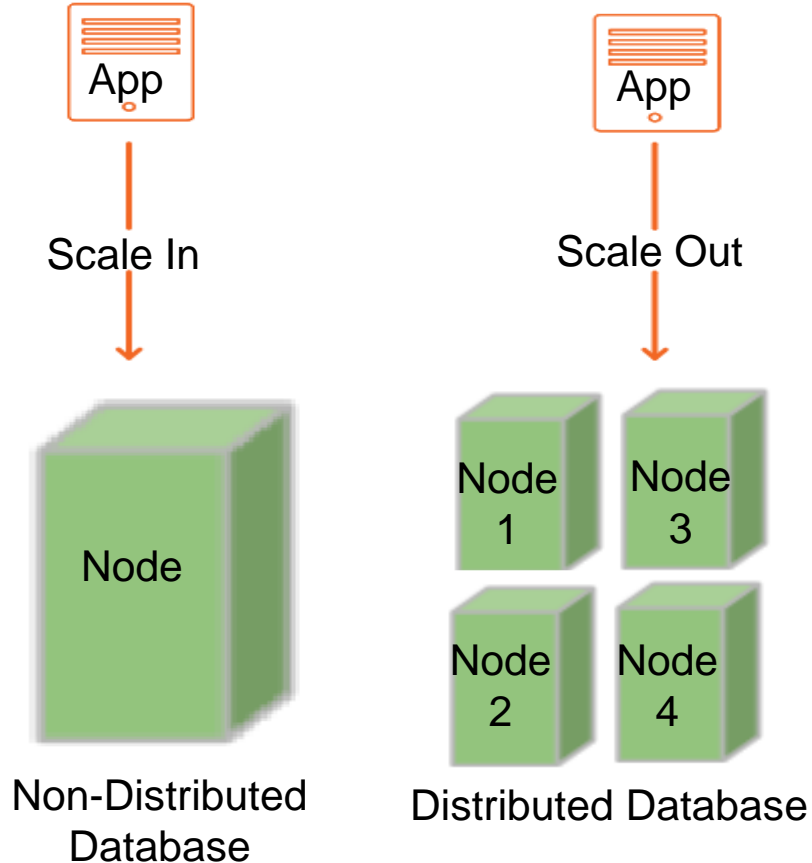
Databases High Availability and Reliability

Database factors	On-premises	Cloud
Database High Availability	✓	✓
Database Reliability	✓	✓
Database Replication	✓ (Hitachi GAD - Two way)	✓ (Standby - One way)
Database Protection	✓	✓
Fault tolerance	✓	✓
Failover mechanism	✓	✓



Databases Scalability

Databases Scalability?



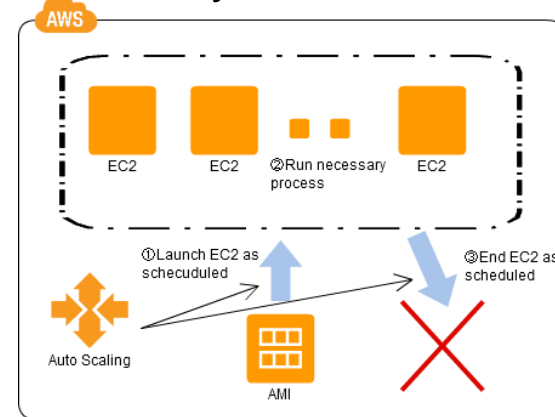
Databases Scalability Comparison

On-premises

- Scaling limitation due to limited resources
- Storage and networking limitations
- Many design and implementation challenges
- Generally scaling horizontally

Cloud

- Automated scaling
- Automatically maintain performance
- Make smart scaling decisions
- Pay only for what you need



Database Scaling Approaches

- There's no One-Size-Fit-All
- Understand your business & database growth
- Attack Top -> Down
- Measure -> Improve



Databases Manageability

On-premises vs Cloud Manageability

On-premises	Cloud
<ul style="list-style-type: none">▪ Self-service - Not available▪ Automation using Ansible, Terraform▪ Licensed tools - Management and orchestration▪ Multiple manual tasks▪ Time consuming▪ Dependability	<ul style="list-style-type: none">▪ Self-service▪ DBaaS provides,<ul style="list-style-type: none">✓ <i>Database management and Monitoring</i>✓ <i>Backup and restore</i>✓ <i>API to enable extendibility and interaction with the database</i>

Manageability - Cloud it better

DBaaS Services



Simple Provisioning



Automated
DBA and Patching



Backup/Recovery
plus HA & DR



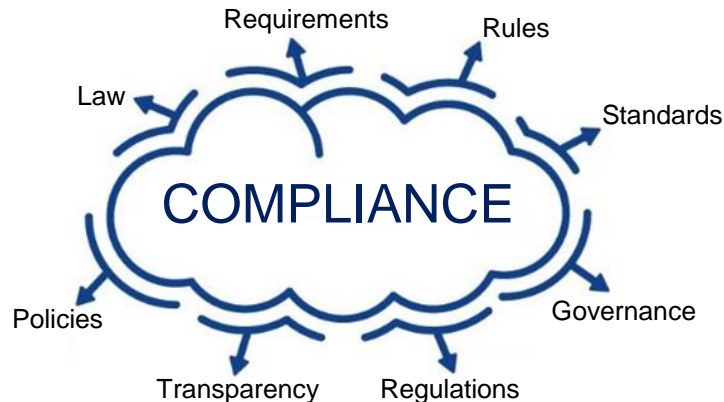


Databases Control and Compliance

Control



Compliance



Parameters	On-Premise	Cloud
What information is stored on a system?	✓	✓
Where is the information stored?	✓	Compliance Challenge and question on compliance
Who can access the system?	✓	
What they can access?	✓	
Is the access appropriate?	✓	

Databases Compliance Requirements

- GDPR Database Compliance
- Sarbanes-Oxley Act (SOX)
- HIPAA Database Compliance
- GLB Act or GLBA
- PCI DDS Database Compliance



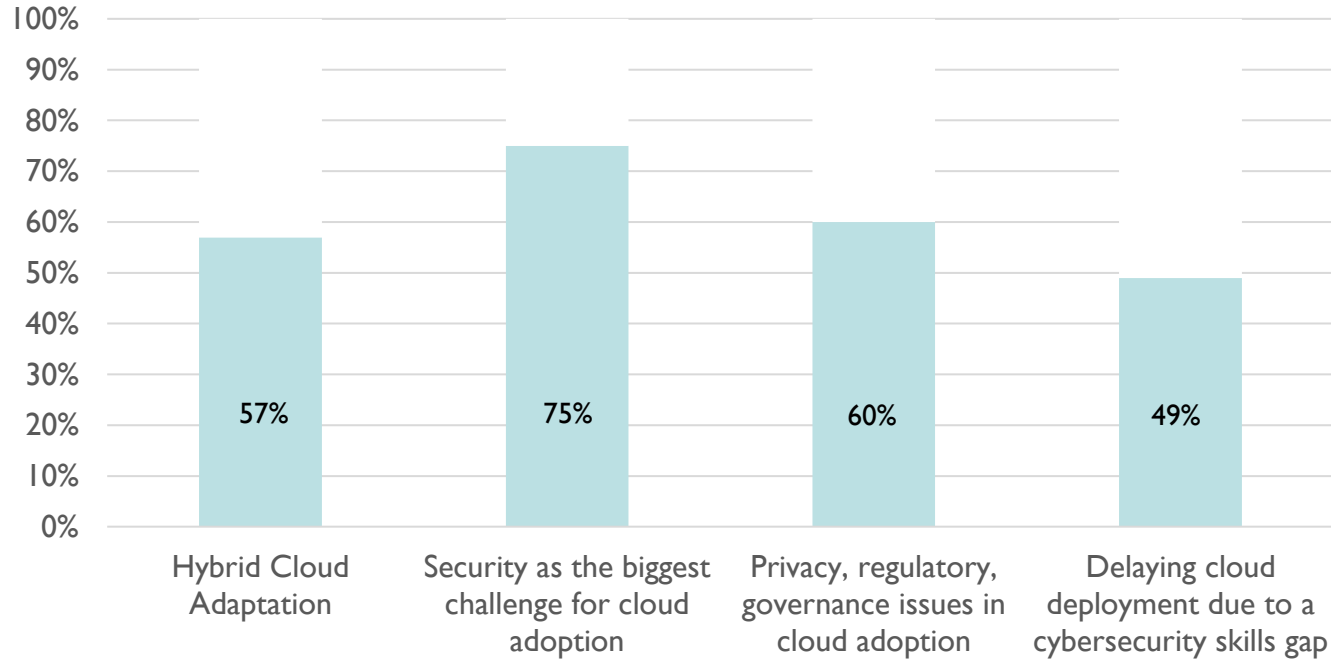
Databases Security

Database Security Checklist

- Encrypt
- Database Audit
- Install firewall
- Perform updates and backups



Data Security in the Cloud



Which One Is Better?

- Not one-size-fits-all
- Hybrid options?
- The short answer is that, based on what priorities you have, either type of data security could be a good fit for your business



Databases Cost/Performance

Factors affects Database Performance

- Storage
- Network
- Server – CPU, Memory
- OS
- Database



Benchmarking

Storage

- Vdbench
- Orion
- Iometer
- FIO
- HClbench

Network

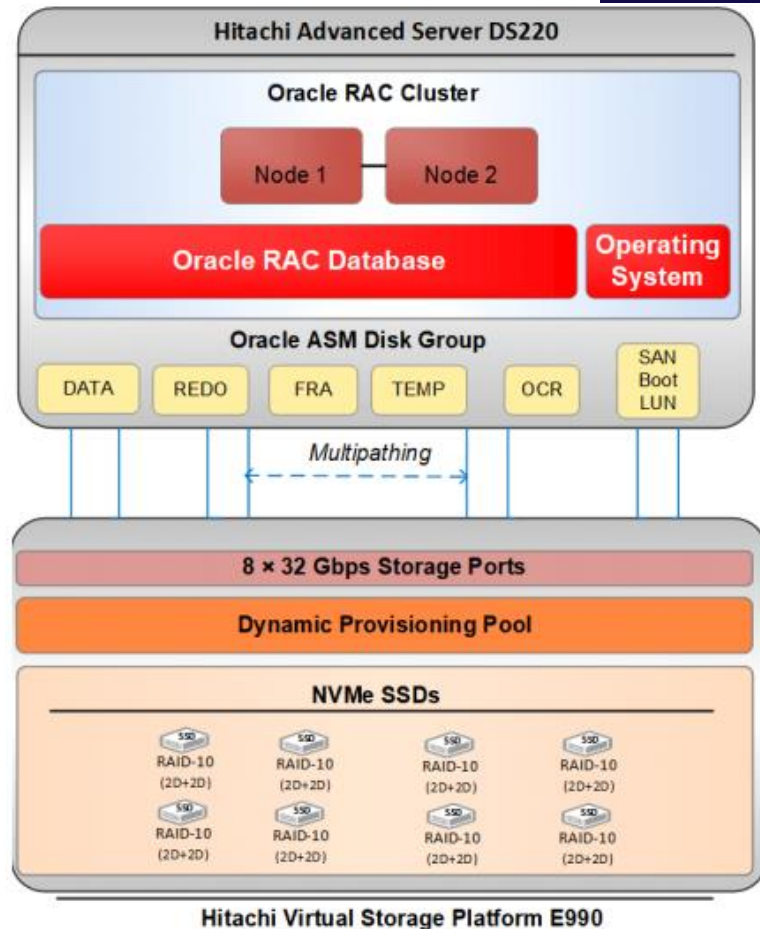
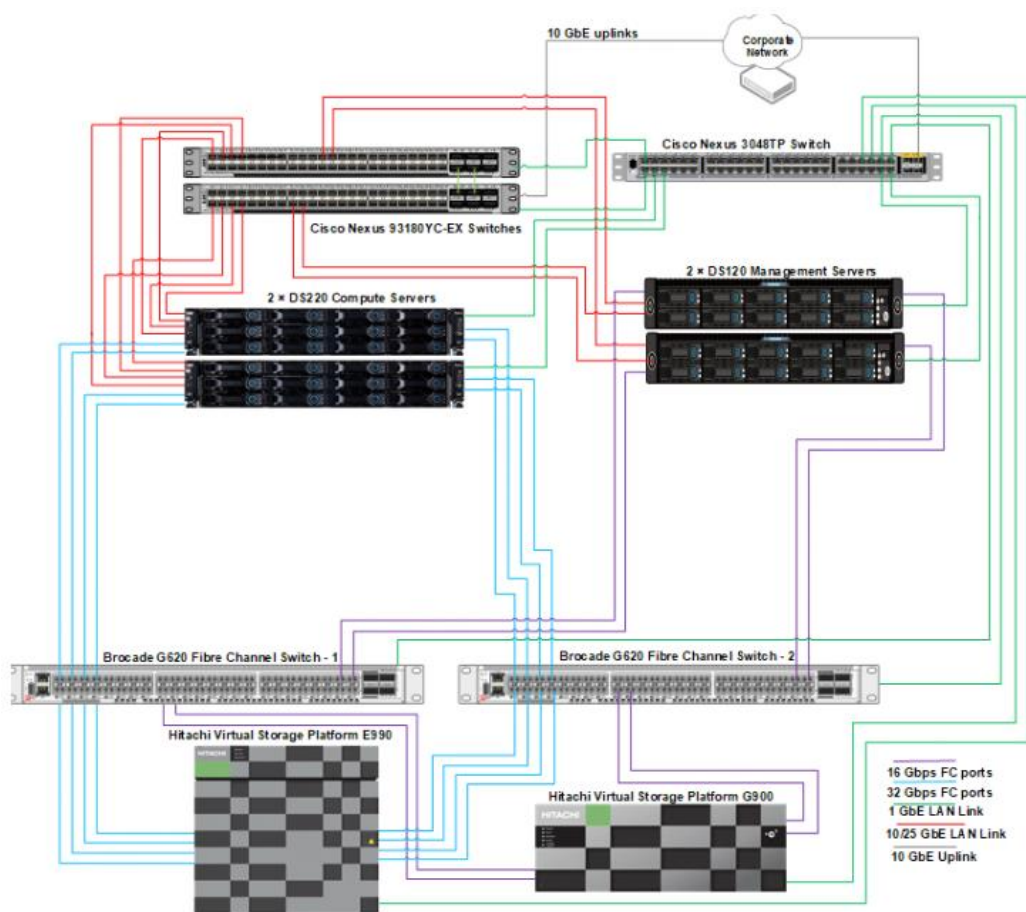
- Acunetix
- Solarwinds
- Datadog
- Obkio

Database

- HammerDB
- Benchmark Factory
- YCSB
- Peakmarks
- Database Benchmark

On-premise solution - Performance

Reference- [Hitachi Solution for Databases - Oracle Composable Platform with Hitachi VSP E990 and Hitachi Advanced Server DS220](#)



Database Cost vs Performance

On-Premises

Storage Performance

Test Category	Test Case	Throughput	RT (ms)
Orion	8KB Random Read	1,021,750 IOPS	0.47
	8KB Random Write	394,600 IOPS	0.93
	1MB Sequential Read	24.56 GB/s	N/A
	1MB Sequential Write	11.29 GB/s	N/A

16TB Oracle Database Benchmarking

Test Category	Test Case	Encryption Disabled		Encryption Enabled	
		Throughput	RT (ms)	Throughput	RT (ms)
Peakmarks	Random 100% read	948,916 IOPS	0.92	822,027 IOPS	0.96
	Sequential Read	17.63 GB/s	1.7	18.53 GB/s	1.26
	Full Table scan	91.5 GB/s	N/A	100.94 GB/s	N/A
	Database Writer – Buffer Cache	3.32 GB/s	N/A	3.004 GB/s	N/A

Cloud

Cloud cost calculator

<https://calculator.aws/#/createCalculator>
<https://azure.microsoft.com/en-in/pricing/calculator/>
<https://cloud.google.com/products/calculator>

- Pay-per-use
- Elasticity
- No over provisioning
- Sharing resources to share cost
- Measuring exact usage



Database Migrations

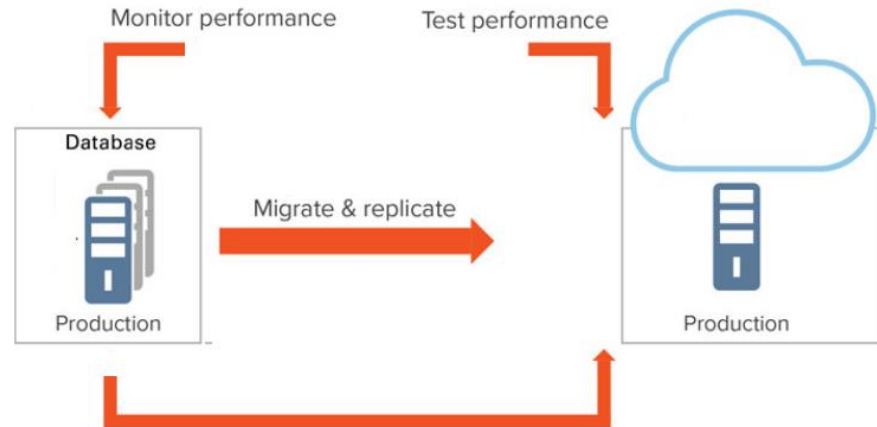
Database Migration

Steps:

- Assess your current data environment
- Migrate access privileges & security settings
- Assess the skills & expertise
- Select the right migration tools
- Deploy a POC
- Review your data
- Migrate, validate

Tools:

- DMS
- Native





Case Study

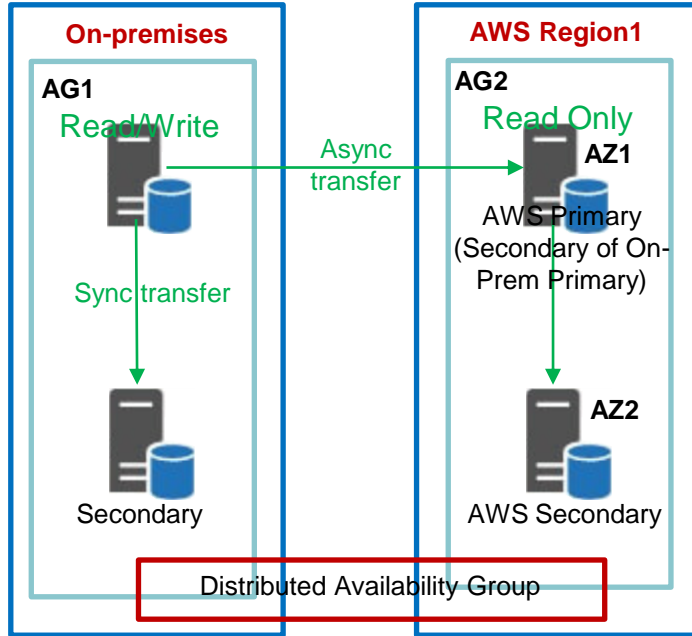
Case Study 1

Hybrid Architecture: Migrate on-premises SQL Server to AWS EC2 SQL Server DB Instance

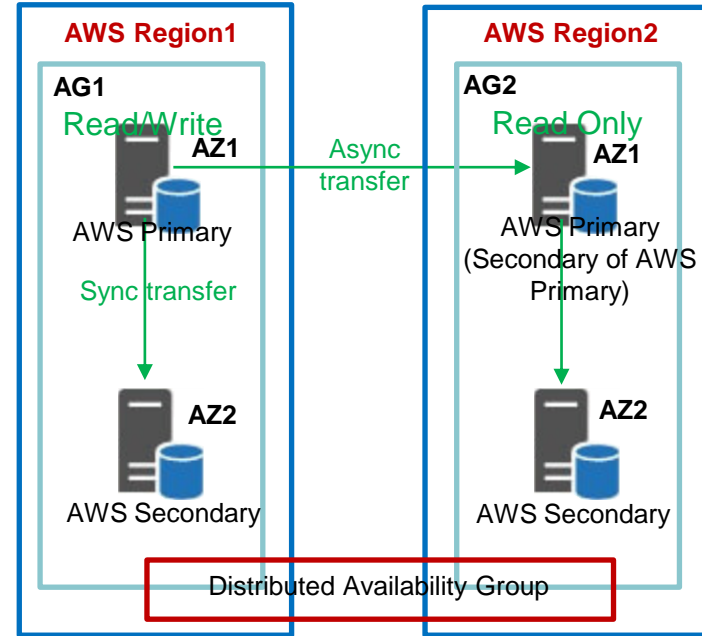
- **Customer:** Advertisement Domain
- **Business Problem Statement:**
 - Application tier moved to cloud and database was still on-premises
 - Facing database latency issues
 - Increased data transfer cost
- **Solution Offered**
 - Database migration to AWS EC2 SQL server Database using distributed Always On availability group with SQL server 2016
 - **Phase 1:** On-premises database nodes (read/writes) + EC2 database nodes (read only)
 - **Phase 2:** AWS-only EC2 DB nodes (read/writes)

Solution Architecture

Phase 1



Phase 2



▪ Solution Benefits:

- Increased application response time (latency)
- Reduced data transfer cost
- Less application down time during migration

Case Study 2

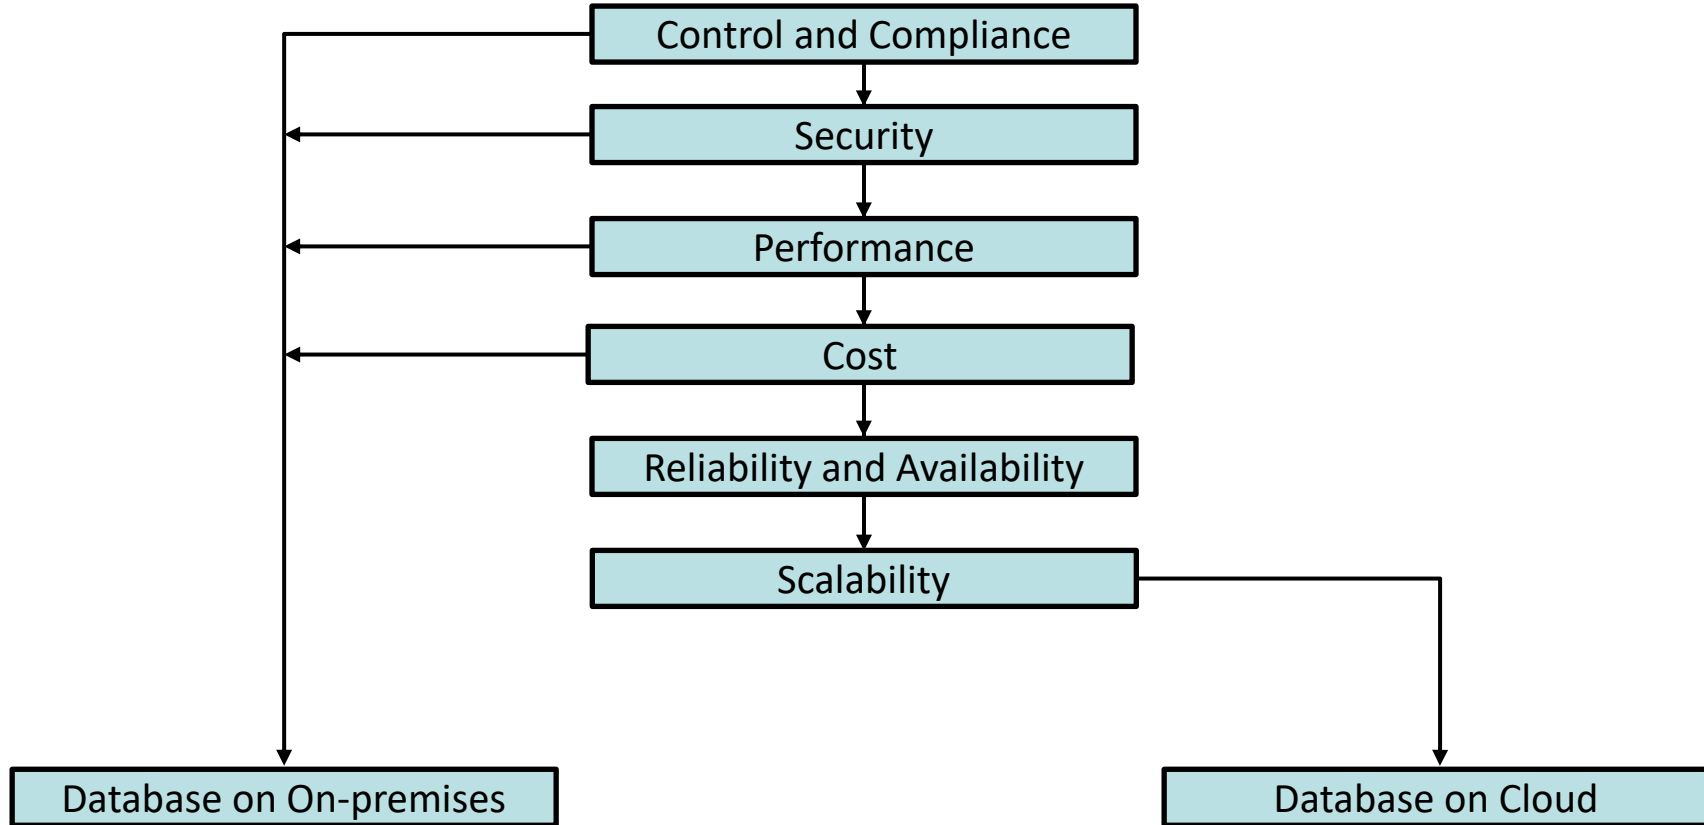
Migrate database from Cloud Database to On-premise Database

- **Customer:** Telecom Domain
- **Business Problem Statement:**
 - Customer moved multiple databases from on-premises to cloud due to cost benefits
 - After 1 year
 - Some databases size grown drastically
 - Performance issues
 - Difficult to troubleshoot easily
 - Database response issues
- **Solution**
 - Customer migrated critical databases back to on-premises
- **Benefits**
 - Complete control over the infrastructure and database
 - Troubleshooting and working with database support team was quicker
 - Resolved performance issues and prepared for near and long-term future performance issues



Conclusion

Conclusion



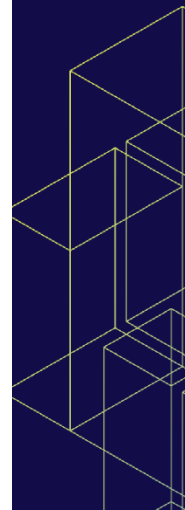
Thank you for your help and support!

- Eduardo Freitas

Director, IDSE – Integrated Data Solutions Engineering & Engineering Operations
Hitachi Vantara

- Gargi Singh

Delivery Director
Xoriant Software Solutions





Questions?

[Email: amol.bhoite@hitachivantara.com](mailto:amol.bhoite@hitachivantara.com) / amol.bhoite@xoriant.com



Thank You!