VERITAS

Providing High Availability, Resiliency, and Performance in Amazon Web Services Cloud Using InfoScale Enterprise

Veritas InfoScale™ Enterprise 7.3

Last updated: 2017-07-12

Summary

Veritas InfoScale™ Enterprise 7.3 builds on the strengths of the Amazon Web Services (AWS) platform and extends it to provide scale-out performance, resiliency and high availability (HA) that protects your databases against the failure of Availability Zones (AZs). InfoScale Enterprise comprises the Veritas InfoScale™ Storage and Veritas InfoScale™ Availability components, among others, which let you manage storage and clustering for your data centers.

- InfoScale Storage delivered 1,54,000 Oracle transactions per minute (TPM) and provided HA using a two-node cluster in AWS, both within and across AZs.
- The Flexible Storage Sharing (FSS) feature of InfoScale Storage lets you overcome the limitations of AWS shared-nothing storage by providing shared storage and a single namespace across servers and even across AZs.

Overview

Enterprises are moving their workloads into cloud environments as they offer agility, infrastructure cost reduction, and ease of manageability. AWS is a leading cloud service provider. AWS provides a range of servers as Elastic Compute Cloud (Amazon EC2) nodes and a range of storage devices as Elastic Block Store (Amazon EBS) volumes. However, lack of HA, storage media failures and protection against failure at a rack level, or failure of AZs, are some common concerns encountered in AWS.

Further, AWS storage cannot scale-out or cannot be shared amongst multiple compute nodes easily, which requires either migration of data or redeployment of applications on scaled-up compute nodes.

As these issues are critical to larger enterprises, they have been reluctant to adopt the AWS platform. InfoScale Storage addresses these concerns by providing HA for applications and scaling-out of storage and compute resources in AWS.

InfoScale Storage

The Veritas Cluster File System (CFS) component provides a single namespace across multiple servers allowing applications to share files as per their workloads and workflows. CFS has a robust transaction model, which is critical to online transaction processing (OLTP) systems such as banking, stock exchanges and travel bookings. Several large enterprises use CFS for large on-premises clusters that run business-critical applications.

The Flexible Storage Sharing (FSS) feature of InfoScale Storage allows shared-nothing storage to be available across multiple servers of a CFS cluster as locally available storage. FSS can combine multiple EBS volumes, which individually have an IOPS limit, to provide increased IOPS capacity. FSS can also combine EBS volumes across multiple AZs. InfoScale Storage delivers scale-out of the following components across multiple EC2 instances and EBS volumes provided by Amazon, as per the business needs:

- Storage capacity
- Compute power
- Resiliency

InfoScale Availability (formerly Veritas Cluster Server – VCS) provides HA for applications against server or storage failures. InfoScale Availability can be configured for servers within an AZ or across AZs, thereby ensuring the continuity of business applications in case a rack or an entire AZ fails.

Configuration

A two-node InfoScale Storage cluster was set up within a single Amazon Virtual Private Cloud (VPC). The cluster was configured using 2 Amazon EC2 M4.16xlarge instances and 6 EBS gp2 volumes of 3.3 TB each attached to each EC2 node. The maximum budget for this configuration was 60K IOPS. Of the 6 gp2 volumes, 5 were used to form a striped-mirrored layout across 2 EC2 instances.

Table 1: AWS configuration

Instances	M4.16xlarge: 2 instances RAM: 256 GB CPU: 64 core
Storage	6 gp2 volumes of 3.3 TB each
Region	Mumbai
AZs	ap-south-1a ap-south-1b
Network tuning	Elastic Network Adapter (ENA) enabled to go up to 20 Gbps
Storage tuning	EBS-optimized storage used to separate storage and network traffic

Table 2: InfoScale cluster configuration

InfoScale Enterprise	Version 7.3
InfoScale cluster	2 nodes
FSS volume	2 TB RAID-10 across 5 gp2 volumes DCO log on 1 gp2 volume
Network tuning	Veritas Low Latency Transport (LLT) for TCP
Storage tuning	vxtune vol_use_rq=0 BIO mode set for faster IO path for Veritas volumes

Striping provided increased performance, while mirroring across the nodes provided data-protection. The remaining gp2 volume was used for the Oracle Redo log. The network was configured using AWS-recommended enhanced networking to provide 20 Gbps.

Figure 1: Configuration across AWS AZs

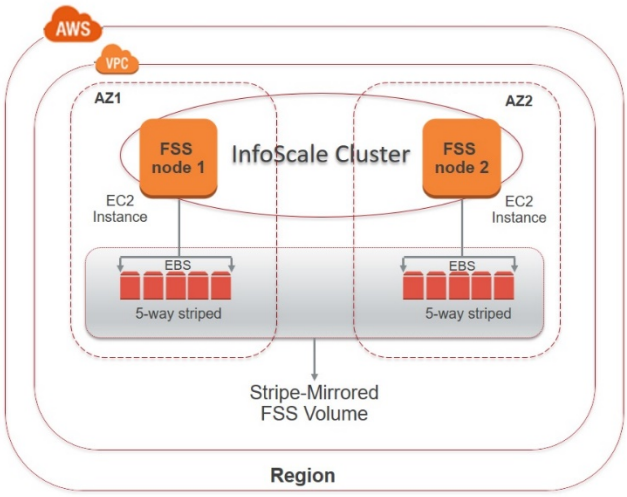

InfoScale Cluster across Azs with FSS

For the first experiment, 2 EC2 instances with EBS volumes were configured in single AZ (ap-south-1a in the Mumbai region). For the second experiment, the 2 EC2 instances were configured in 2 different AZs (ap-south-1a and ap-south-1b in the Mumbai region) and the 20 Gbps enhanced network was setup across the AZs.

Workload

The benchmarking of InfoScale Enterprise on AWS was tested with a large Oracle database. The OLTP workload was generated on the Oracle 12c database with the Transaction Processing Committee's (TPC) benchmark C (TPC-C) in its unaudited form.

Table 3: Oracle database configuration

Oracle database version	12c
Database size	1.5 TB
Oracle System Global Area (SGA) size	32 GB
Redo log size	100 GB
Database tuning	Hard and soft limits for Oracle users: soft_nproc=4096, hard_nproc=16384 soft_nofile=65536, hard_nofile=65536 soft_stack=10240, hard_stack=32768

The TPC-C benchmark simulates a wholesale supply company that comprises multiple regional warehouses, which supply 10 sales districts per warehouse. Each district serves 3,000 customers. Each warehouse stocks 1,00,000 items that are sold by the company.

Table 4: Benchmark and workload details

Benchmark	TPC-C-like OLTP benchmark
Random Read:Write ratio	70:30 percentage ratio
SGA	32 GB
Users	200
Number of warehouses	5000
Database size	1.5 TB

The company's OLTP application performs a set of read-only, update, insert, and delete transactions to process customer orders, payments, and delivery of goods. The TPC-C database consists of 9 tables with a logical entity-relationship defined between the tables. Oracle 12c was used to create a database of 1.5 TB with 5000 warehouses. The Oracle SGA size was set to 32 GB. The workload was generated using 200 users.

Results

The performance of InfoScale Enterprise with the OLTP workload on the Oracle database was measured by the TPM that the configuration delivered. System resource metrics such as the CPU usage and storage IOPS were also measured.

Table 5: Database performance across 2 AZs compared to performance within an AZ

	SGA	Users	TPM	IOPS	CPU
Within AZ	32 GB	200	154.2 K	54.4 K	25%
Across AZ	32 GB	200	154.6 K	54.4 K	25%

Performance within and across AWS availability zones

The benchmark results show that InfoScale Enterprise delivered 1,54,000 OLTP TPM on a large Oracle database with M4.16xlarge instances and 6 gp2 volumes. The CPU usage was 25% and the SGA size was 32 GB, which means that the remaining CPU and memory was spared for other application processing.

Figure 2: Same database performance within and across AZs

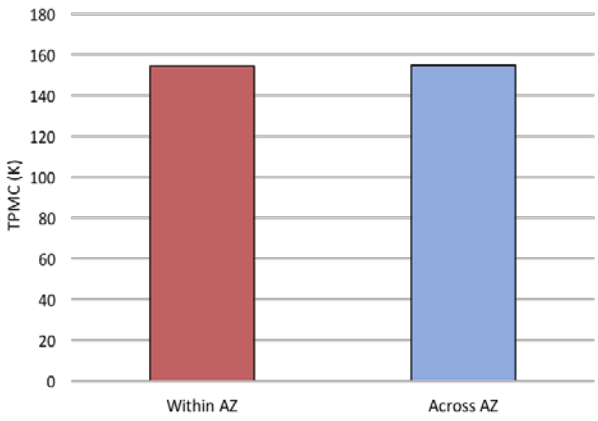

When the instances were deployed in 2 different AZs, there was no drop in OLTP performance. This shows that InfoScale Enterprise can deliver matching database performance within as well as across AZs.

Failover within and across AWS availability zones

This experiment demonstrated the fast failover of an Oracle database instance across 2 EC2 instances located in different AZs using an InfoScale Availability cluster.

Of the 2 EC2 instance in the InfoScale Enterprise configuration, the EC2 server in the ap-south-1a AZ was intentionally made to crash during the benchmark run. This triggered the migration of services from the failed EC2 instance (which formed a cluster node) to the EC2 instance present in the ap-south-1b AZ. The failover process first reconstructed the mirrored Oracle database to a consistent state and restarted the database. The failover process lasted for 2 minutes and 34 seconds.

The application process, which was the benchmark workload in this case, was then restarted with no loss of performance. This demonstrates that InfoScale Enterprise can provide fast failover of Oracle dbase instances even across AWS AZs with no loss of data or performance.

References

You can find more information about InfoScale Enterprise at:

- Veritas InfoScale Enterprise
- Enabling application and database availability

About Veritas Technologies LLC

Veritas Technologies LLC enables organizations to harness the power of their information, with solutions designed to serve the world's largest and most complex heterogeneous environments. Veritas works with 86 percent of Fortune 500 companies today, improving data availability and revealing insights to drive competitive advantage.

For specific country offices and contact numbers, please visit our website.

Veritas World Headquarters

500 East Middlefield Road

Mountain View, CA 94043

+1 (650) 933 1000

www.veritas.com