ATLANTIS™ The New Order of Storage Freedom

Atlantis USX Hyper-Converged Solution for Microsoft SQL 2014

ATLANTIS USX™

Microsoft® SQL Server®

atlantiscomputing.com

Table of Contents

Executive Summary

Enterprises rely heavily on databases to deliver pivotal information, providing the analytics that organizations need to make critical business decisions. With the introduction of Atlantis USX, businesses can make better decisions faster while simplifying datacenters with hyper-converged infrastructure. Atlantis USX enables customers to cut the cost per transaction by up to 90%, delivering 5X more transactions per minute than traditional servers with SAN at less than the half the cost. In addition, the increased number of transactions per server socket dramatically reduces database software license costs.

Atlantis USX is a software-defined storage solution, delivering performance that exceeds all-flash arrays for half the cost of a traditional SAN or NAS. Organizations can deploy Atlantis USX by leveraging existing storage, SAN, NAS, DAS, or Atlantis HyperScale hyper-converged appliances to accelerate virtualized Microsoft SQL database transaction processing. Atlantis USX also saves significant storage costs by reducing the capacity required for SQL databases, logs, etc. with deduplication averaging up to 80%. Organizations can also expect additional savings in operational costs with increased productivity as batch jobs are completed over five times faster than traditional storage technologies.

From an IT perspective, organizations can deploy the Atlantis USX platform is as little as one hour. Microsoft SQL infrastructures can easily be deployed or moved to Atlantis USX and immediately deliver value. Atlantis USX not only accelerates the performance and scalability of SQL environments, but also provides additional resilience through high availability and data protection. Atlantis USX is a proven storage platform that gives organizations the confidence to scale and meet the growing demands on databases today.

In order to demonstrate the value the Atlantis USX platform brings to database infrastructures, a series of tests were conducted using Iometer, HammerDB, and SQLIO. These tests validate the ability for Atlantis USX to optimize Microsoft SQL 2014 database environments for performance, capacity and throughput. Using the Atlantis USX hyper-converged architecture, Atlantis was able to achieve compelling results using the architecture and testing methodology defined in this document.

Altantis USX with Microsoft SQL Server 2014	
Characteristic	Results (Up to)
Cost	Transactions per minute as low as $0.18
Capacity	70-80% average reduction in required SQL storage capacity
Performance	-5X more transactions per minute -80% less time for batch processing jobs

Solution Overview

There are three primary goals for most SQL administrators to consider: cost, performance and resiliency. Cost and performance share a close relationship in SQL environments. To lower cost and increase performance, an enterprise needs to increase the amount of transactions performed by each SQL server. This enables organizations to spend less on licensing, storage, and hardware while achieving greater density as a result of faster transactions. Transactions per minute (TPM) is not the only contributing factor for evaluating storage performance, but also the performance of the underlying storage subsystem as well. To guarantee database performance isn't limited by the underlying storage, the storage subsystem needs to be extremely fast and provide very low latency.

Atlantis USX accelerates the performance needed for enterprise SQL by using a combination of inline data deduplication, RAM, and IO acceleration. This ensures that the storage is no longer a bottleneck for TPM. Atlantis USX compliments the new Microsoft SQL 2014 feature 'in memory OLTP'. In order to accelerate the performance of database tables held in RAM by the "In Memory OLTP" feature, transaction log performance also must be accelerated. Without accommodating the increased performance demand for the transaction logs, the full benefits of "In Memory OLTP" will not be achieved. Atlantis USX delivers the performance and low latency to accommodate these needs, allowing the enterprise to deliver the transaction processing speeds required and achieve results faster.

Best practices for architecting Microsoft SQL storage may be intimidating to the uninitiated. Splitting databases, transaction logs, and the "temp.db" across different datastores while accommodating both block and file based storage volumes can be challenging in many environments.

Atlantis USX can easily cater to these diverse requirements. This allows storage volumes to be created for any size, resiliency level, and protocol needed to meet the demands of the organization. Stretching from multiple NFS volumes for logs and data to small iSCSI volumes for Windows clustering, Atlantis USX can deliver any storage scenario required. This empowers administrators to easily create any type of volume necessary to accommodate the requirements of a desired SQL architecture.

Atlantis USX and Microsoft SQL Architecture

Microsoft SQL Architecture

There are a wide variety of architectures that can be leveraged to deliver a Microsoft SQL 2014 environment running on the Atlantis USX platform. Microsoft best practices should always be considered when architecting and deploying a SQL environment. Atlantis USX will accelerate and optimize the various architectures and methodologies that Microsoft recommends for any organizations desired SQL environment.

Stand-alone vs. Always-On Architectures

A stand-alone configuration in this document is defined as an environment where the database is not replicated or mirrored to another location. Only traditional backup and restore methods are used to preserve data and Microsoft Windows Clustering is not used. As a result, NFS volumes can be leveraged for more accurate capacity reporting with stand-alone architectures. In scenarios where application level resiliency is required, the 'Always on' architecture and feature set for

synchronous/asynchronous replication has a pre –requisite for a windows failover cluster, this in turn requires block level storage. Atlantis USX fully supports the mix of file and block storage which would be ideal for an 'always on' infrastructure

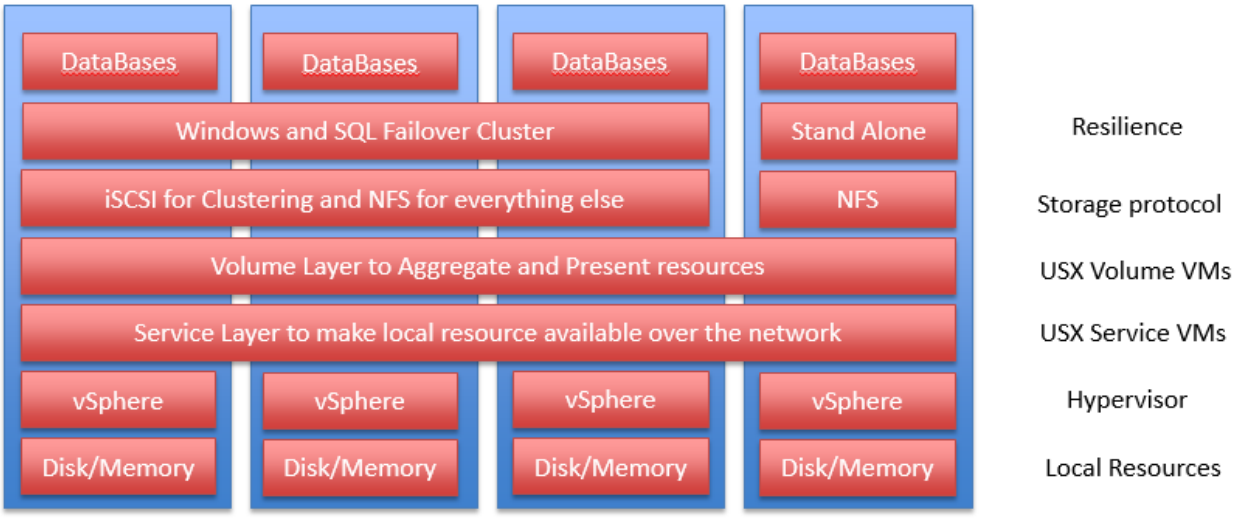

Figure 1: Logical design for both stand-alone and always on SQL architectures

Microsoft SQL 2014 availability groups are fully supported by Atlantis USX:

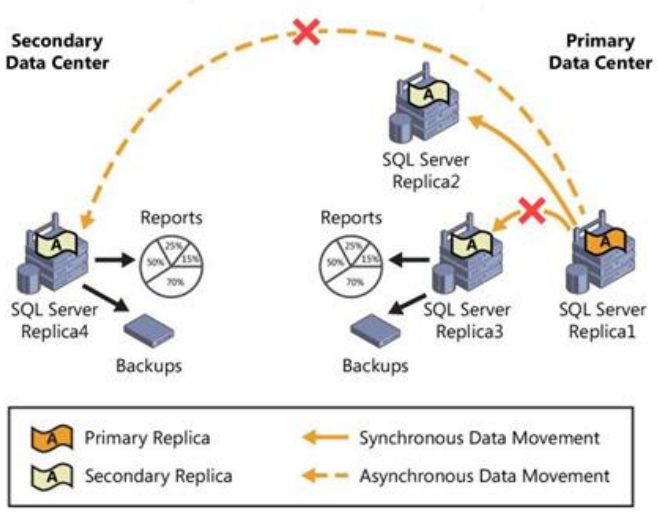

Figure 2: Microsoft SQL availability group architecture

Microsoft SQL 2014 Failover Clusters are fully supported by Atlantis USX:

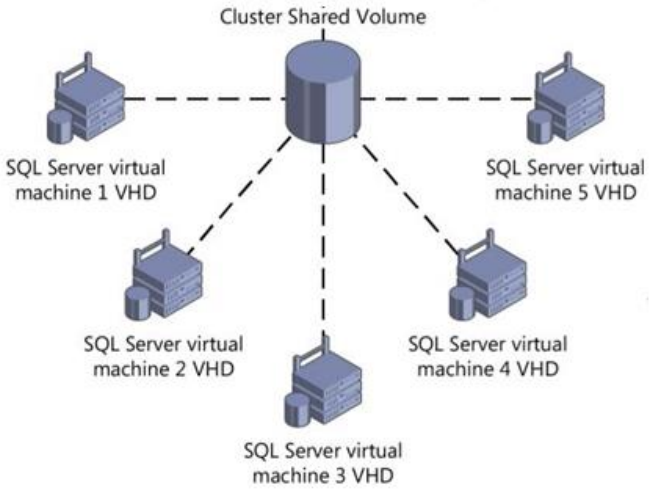

Figure 3: 2014 Microsoft SQL Failover Clusters

In-Memory OLTP

Microsoft describes the use of OLTP as: 'To use In-Memory OLTP, you define a heavily accessed table as memory optimized. Memory-optimized-tables are fully transactional, durable, and are accessed using Transact-SQL in the same way as disk-based tables. A query can reference both memory-optimized tables and disk-based tables. A transaction can update data in memory-optimized tables and disk-based tables. Stored procedures that only reference memory-optimized tables can be natively compiled into machine code for further performance improvements. The In-Memory OLTP engine is designed for extremely high session concurrency for OLTP type of transactions driven from a highly scaled-out middle-tier. To achieve this, it uses latch-free data structures and optimistic, multi-version concurrency control. The result is predictable, sub-millisecond low latency and high throughput with linear scaling for database transactions. The actual performance gain depends on many factors, but 5-to-20 times performance improvements are common.'[1]

An administrator will not be able to achieve the claimed 5-20x performance improvement and benefit from In-Memory OLTP if the storage subsystem for the transaction logs performs poorly. This is primarily due to the fact that Microsoft's In-Memory OLTP technology inheritably only allows for the transaction logs to be placed on disk. Since transaction logs must support of a greater rate of change than a database without OLTP enabled, the storage for these logs needs to be considered carefully. As Microsoft states 'because you now have a higher throughput workload your local I/O device will need to support it'[2] and 'If you use spindles for your log device you won't get a lot of performance benefit from OLTP'[3] . As a result, Atlantis USX and In-Memory OLTP align closely with Microsoft's recommendation for high performance storage for leveraging this feature.

Atlantis USX Architecture

Atlantis USX can be configured in a variety of ways to achieve the demands of numerous applications and infrastructure requirements. While simply accelerating shared storage is perfectly possible, in a new environment using local direct attached storage rather than shared is more cost effective. This allows a choice of Hyper-Converged using a combination of hard disk and SSD or Hyper-Converged with local flash. To conduct the testing in this document, local SSDs were used.

The architecture for SQL stand-alone with Atlantis USX is as demonstrated in this diagram:

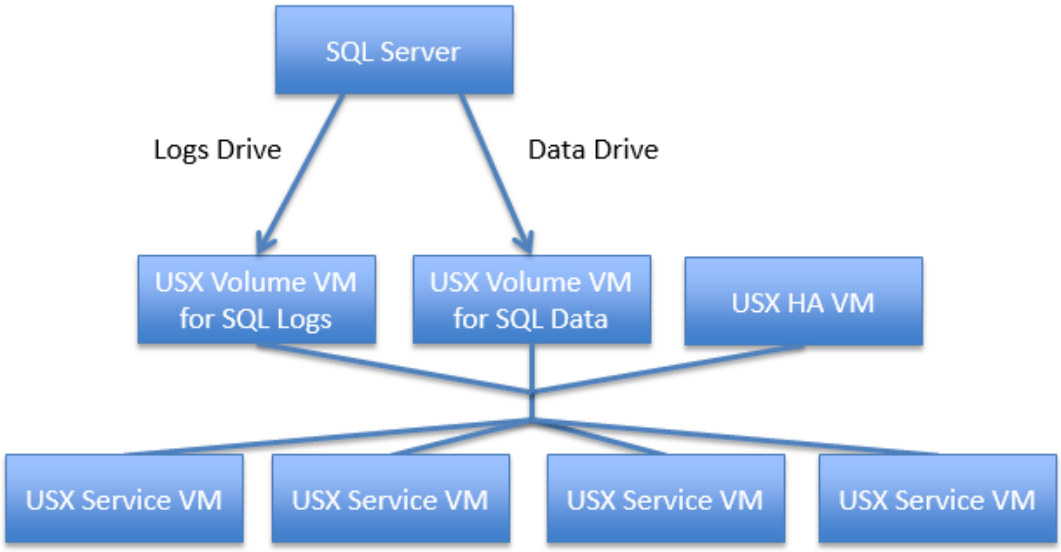

Figure 4: Standalone SQL Architecture with Atlantis USX

The Atlantis USX architecture for "Always-On" is as follows:

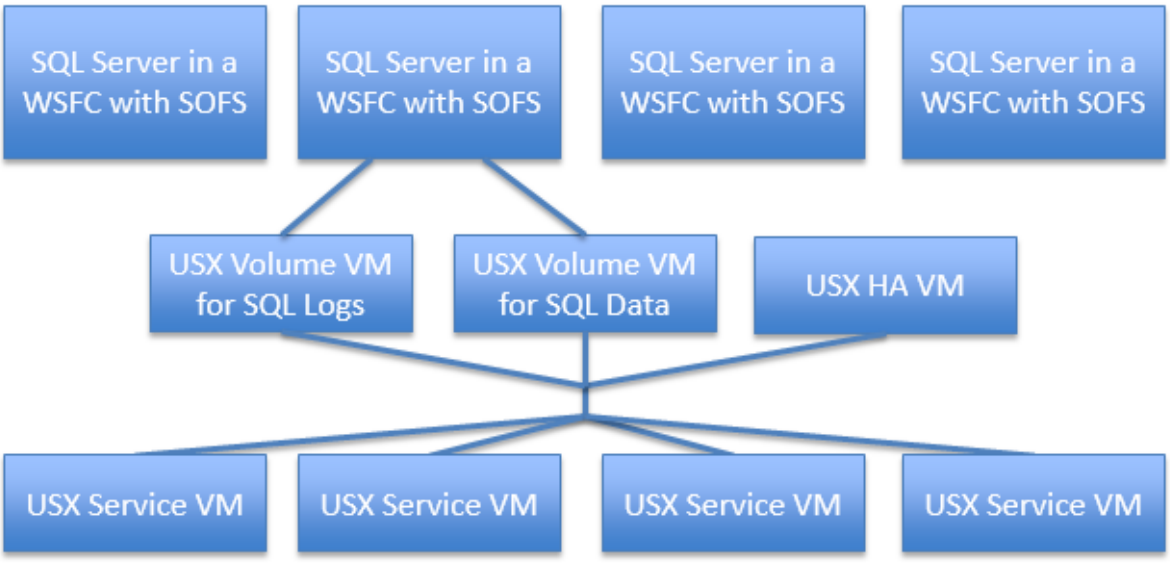

Figure 5: Always-On SQL Architecture with Atlantis USX

Atlantis USX High Availability and Resiliency

This section will only deal with the high availability of the Atlantis USX components, as SQL application level HA has been previously detailed. Atlantis USX best practices with HA include having a single HA volume for all other volumes of the same size. The more volumes that can be sized equally, the fewer HA volumes that are needed.

Atlantis USX HA will ensure that in the event of a failure of any single host, the Atlantis USX environment will continue to function without an impact in storage availability. Specifically, the Service VMs and Volume VMs work together harmoniously to enable the performance and capacity benefits that Atlantis USX provides. In the event of a service VM failure, the Atlantis USX service will be uninterrupted and operate continuously. Volume service VMs however assume the role of the lost volume, providing business continuity to data and eliminating downtime.

Server Hardware Recommendations

Atlantis USX benefits from a higher performing CPUs, therefore it's beneficial to purchase higher clock speed processors. Speed of the RAM bus has been found to have little impact on performance. For local SSD drives, a spec in the middle of current hardware offerings should be sufficient. Atlantis would recommend at least 800GB of data throughput per day for the SSD drives.

Testing Methodology

A series of various tests were conducted to validate the performance, scalability, and reliability of Atlantis USX for SQL workloads. The first phase of testing focused primarily on the raw IO storage performance that Atlantis USX provides, leveraging Iometer and HammerDB configured to simulate SQL workloads. The second phase of the testing leveraged the SQLIO tool from Microsoft. This tool provides the ability to simulate production level Microsoft SQL workloads for scalability and performance validation. SQLIO was configured and tested for a variety of block sizes typically leveraged by enterprise implementations to represent the value Atlantis USX provides in a myriad of scenarios.

Test Environment Overview

The environment to conduct the series of SQL tests was configured as illustrated below.

Hardware Components – 1 Node		
Component	Quantity	Description
Processor	2	Intel Xeon 12C E5-2680v3 120W 2.5GHz/2133MHz/30MB (12 cores/socket)
Networking	1	Dual Port - Emulex VFA5 ML2 10GbE SFP+ Adapter
Networking	1	Quad Port - BCM5719 Broadcom RJ-45 Gigabit Ethernet 1000BASE-T

RAM	256GB	16GB TruDDR4 (2 Rx4, 1.2V) PC4-17000 CL15 2133MHz (x24 modules)
Drive	1.2TB	Intel S3700 400GB SATA 2.5" MLC G3HS Enterprise SSD (x3 drives)

Software and Benchmarking Tools

Iometer 1.1

Iometer is an I/O subsystem measurement and characterization tool for single and clustered systems. It is used as a benchmark and troubleshooting tool and is easily configured to replicate the behaviour of applications. In this test, Iometer was configured to simulate a SQL workload and determine the maximum available IOPS of Atlantis USX environment. The SQL workload tested was configured to be 50% write, 50% read and 80% random with a 4K block size using a two 20GB test files on two Atlantis USX volumes. A virtual machine with Iometer was run with this test profile to determine the maximum available IOPS for Atlantis USX Solution.
http://www.iometer.org/

SQLIO

SQLIO is a tool provided by Microsoft which can also be used to determine the I/O capacity of a given configuration. SQLIO is also an I/O subsystem measurement and characterization tool for single systems. In this test SQLIO was used to determine the underlying storage performance at varying metrics, including block size and queue depth.

http://www.microsoft.com/en-gb/download/details.aspx?id=20163

HammerDB

HammerDB is an open source database load testing and benchmarking tool for Oracle, SQL Server, TimesTen, PostgreSQL, Greenplum, Postgres Plus Advanced Server, MySQL, Redis and Trafodion SQL on Hadoop. It is used to test the transactional performance of databases. It was used for two tests in this architecture, TPC-C while creating the database and TPC-C scaling after database creation.

http://www.hammerdb.com/about.htm

VMware vSphere

VMware vSphere 5.5 Update 2 was installed on the host. VMware tools were installed into the Microsoft SQL VM, and no additional performance optimizations were performed than the default configurations for vSphere and the VM.

Microsoft SQL 2014 SP1

Microsoft SQL 2014 SP1 was used to provide the SQL role function in this test. The database and transaction logs were placed on different Atlantis USX volumes to optimize SQL performance accordingly.

Test Results

Test Results Summary

The tables below provide a summary of the results based on the various testing methods used to demonstrate the value Atlantis USX provides with Microsoft SQL workloads. A more in depth analysis is provided in the testing section to explain how these results were achieved.

Iometer/HammerDB Test Results	
Test	Results
Iometer (IOPS) – 2 Volumes	76,650
HammerDB TPC-C Auto-Pilot (TPM)	319,036

SQLIO Test Results	
4K Block Size	
Latency (in milliseconds, R/W)	0/0
IOPS (Read+Write)	122,137
Throughput (MB/s)	447
8K Block Size	
Latency (in milliseconds, R/W)	0/0
IOPS (Read+Write)	115,084
Throughput (MB/s)	899
64K Block Size	
Latency (in milliseconds, R/W)	0.5/0.5
IOPS (Read+Write)	33,437
Throughput (MB/s)	2089

Solution Price by Performance

Based on the performance results above, the cost of the solution by metric can be evaluated when considering:

- Server Hardware (with 3 year support)
- Atlantis Software (with 3 year support)
- Microsoft Windows Server 2012 R2 Standard
- Microsoft SQL 2014 Standard Edition (8 vCPUs)
- VMware vSphere 5.5 U2 (2 socket licensing)

Solution Pricing	
Solution Pricing by Metric	
TPM (Auto-Pilot) as low as:	$0.18
SQLIO (4K Block)/IOPS as low as:	$0.20
SQLIO (8K Block)/IOPS as low as:	$0.21
SQLIO (64K Block)/IOPS as low as:	$0.94

Iometer Test

Iometer is a tool that assists in determining the amount of IOPS a subset of hardware is capable of producing, leveraging both reads and writes to the drive. In this scenario, since the system is writing to the Atlantis USX SSD volumes, the IOPS capabilities are significantly higher than traditional storage architectures. The test was conducted leveraging the parameters described in the benchmarking tools section.

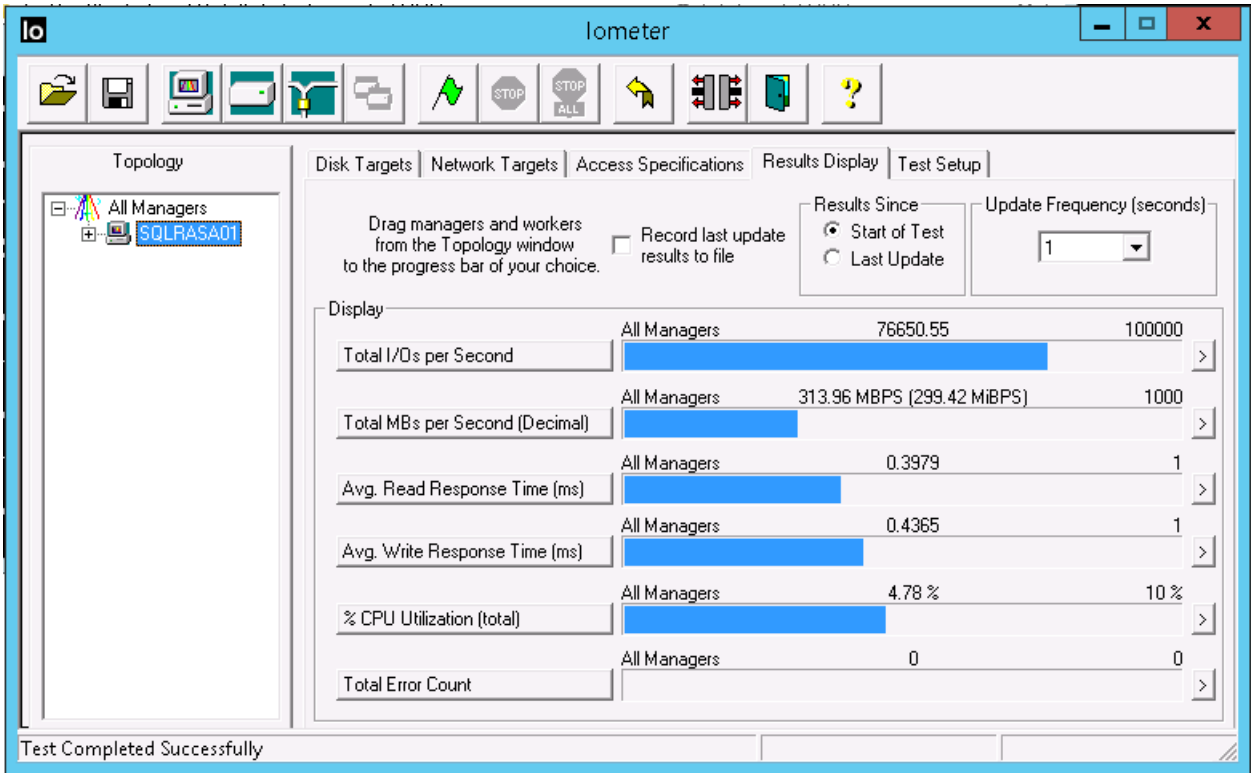

Figure 6: Iometer Test Results

IOPS

Iometer was run against two volumes simultaneously generating 76,650 IOPS. Microsoft best practice dictates splitting log and data across two different storage locations, however there is also the ability to split a database across several locations. By placing the VM's VMDK files on different Atlantis USX

volumes, an organization can scale effectively to the desired level of performance and meet the needs of the business. If greater performance is needed, additional volumes can be created to increase IOPS and throughput. As a result, the incremental cost per IOPS is kept relatively low with the Atlantis USX platform.

Latency

Latency during Iometer testing was very low, with 0.3979ms for read and 0.4365ms for write operations during 100% system utilization and saturation. Under normal operation, response times are further minimalized, achieving the performance needs to produce high SQL transactions per minute. Comparatively, traditional spindle based systems generally produce 40ms of latency in heavily utilized scenarios. This test demonstrates the ability of Atlantis USX to gracefully maintain high performance even in the most extreme scenarios, outperforming both traditional spindle and flash based storage systems.

Throughput

Throughput is also excellent, exceeding more than 300 MBPS during this test. In some scenarios, throughput can scale even higher and Atlantis USX can easily scale to meet these demands.

HammerDB

HammerDB was used to measure the effectiveness of the infrastructure by evaluating the number of transactions per minute the system can provide. When leveraging this tool, a consistent high TPM is the most optimal scenario.

TPC-C Creation Phase

During the creation phase, Atlantis USX enabled a large amount of transactions per minute and maintained a high level of performance at a consistent rate.

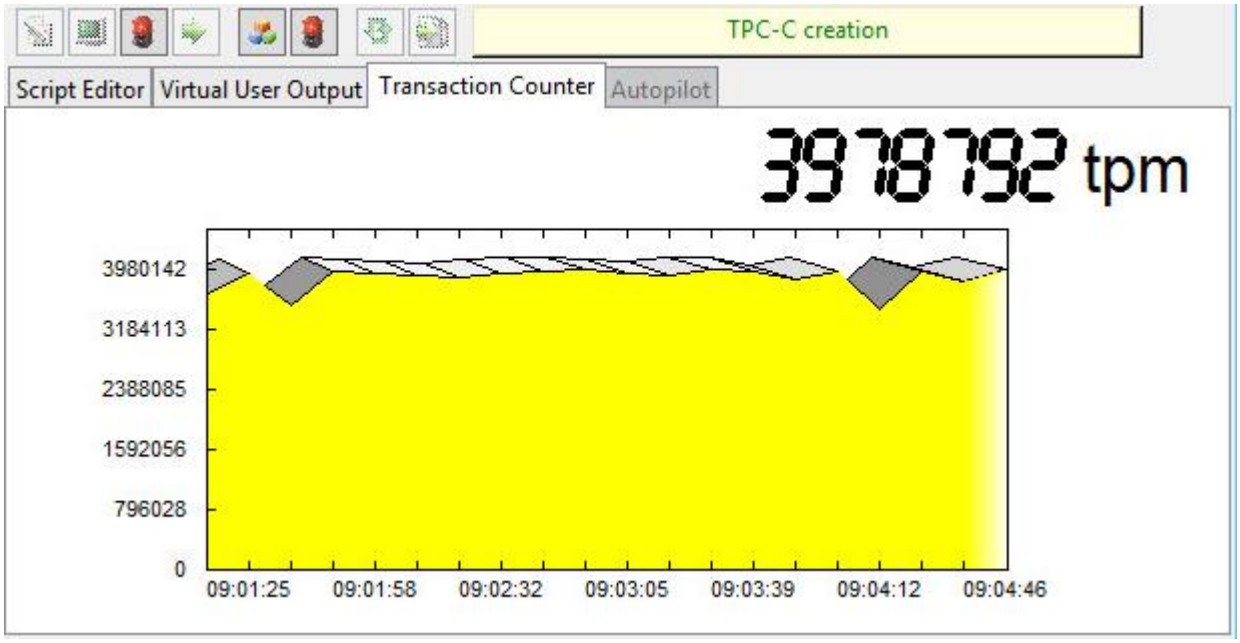

Figure 7: HammerDB TPC-C Creation Results

TPM Auto-Pilot

The HammerDB TPM Auto-Pilot test demonstrates the ability for Atlantis USX to provide reliable scalability as a user workload demand is progressively increased.

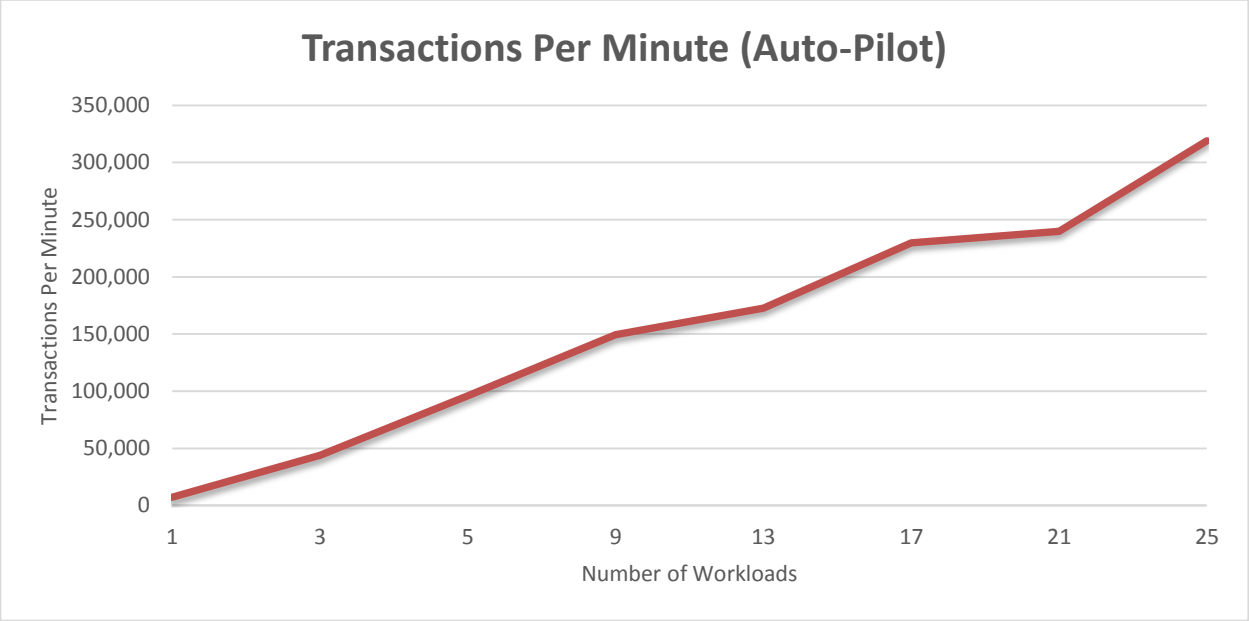

Figure 8: HammerDB TPM Auto-Pilot Results

SQL IO

SQLIO is a tool developed by Microsoft that assists with determining the I/O capacity and capabilities of an infrastructure. In this example, SQLIO was used to demonstrate the extreme performance Atlantis USX provides in a SQL environment.

IOPS and Throughput by Block Size

Traditionally, many SQL workloads operate in the block size range below. USX accelerates SQL workloads in all block ranges, providing the performance needed to deliver enterprise scale implementations and increase data center efficiency.

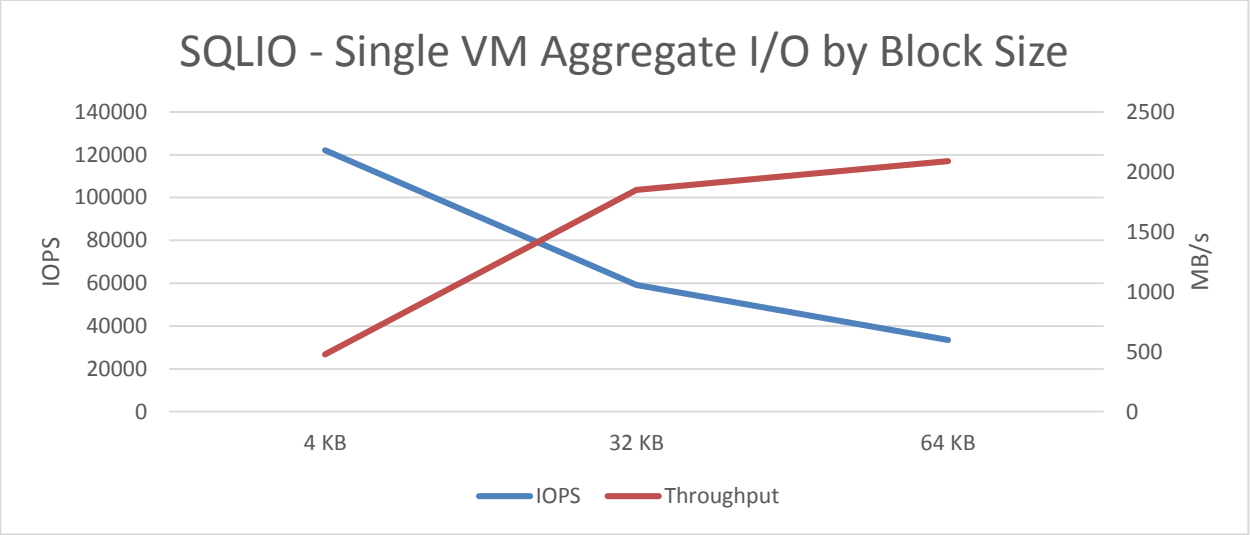

Figure 9: SQLIO - Single VM Aggregate I/O by Block Size

To further demonstrate USX's ability to accelerate SQL performance, additional testing was performed with extremely large block size levels. Even in the most extreme scenarios, USX continues to provide increased efficiency to these workloads.

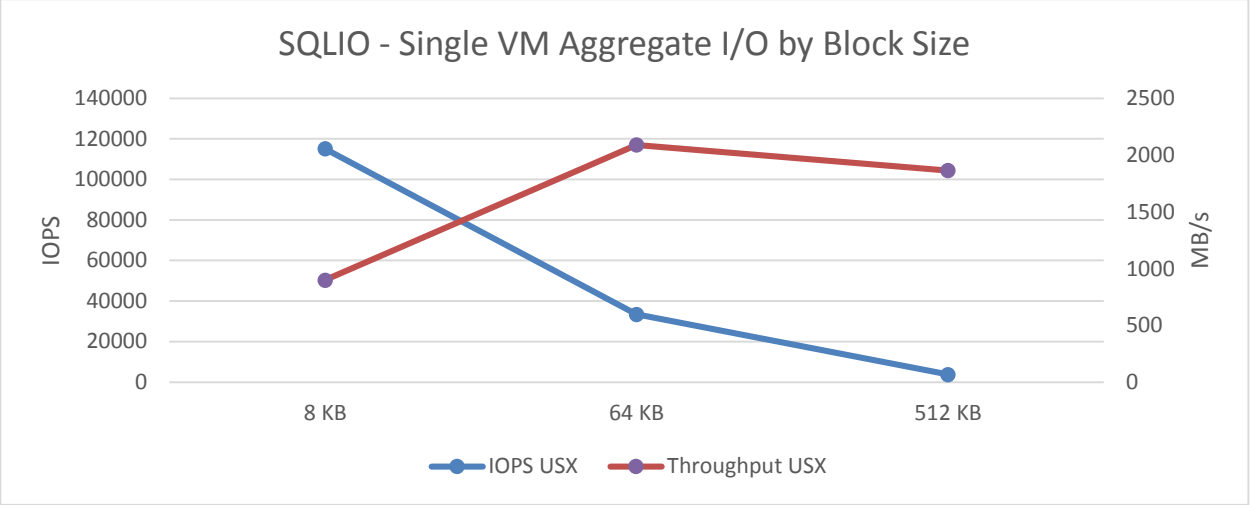

Figure 10: SQLIO - Single VM Aggregate I/O by Block Size

Latency by Block Size

While IOPS are always an important factor to consider, it's more important to deliver them at low latency. The data below provides visibility into how Atlantis USX maintains the high level of IOPS without sacrificing performance. Latency below 0ms is not measured by the tool.

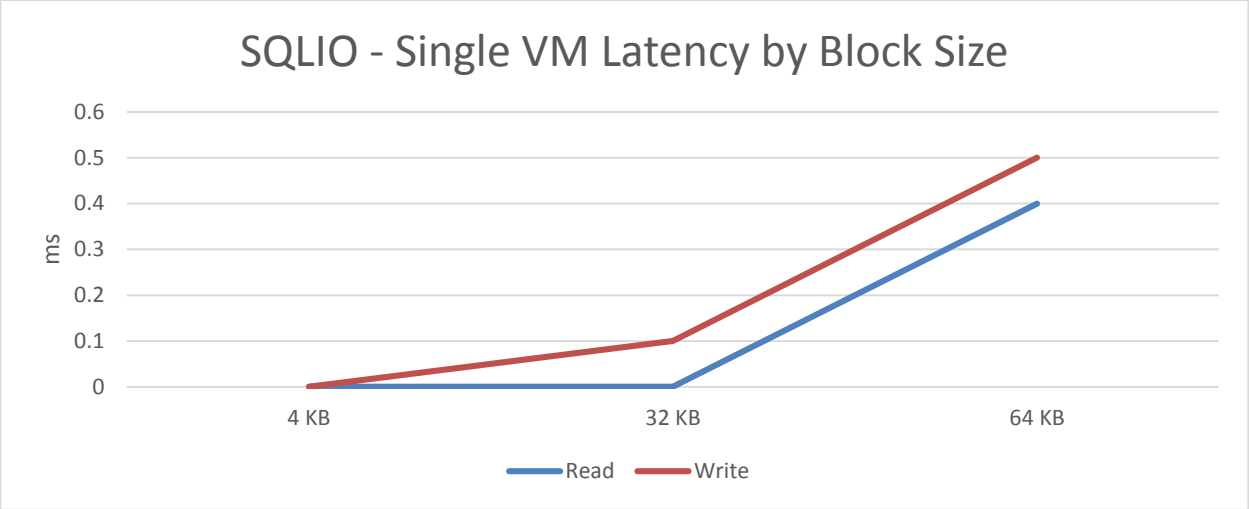

Figure 11: SQLIO – Single VM Latency by Block Size

The read and write operations have essentially no latency up until 32KB block sizes. Even at 64KB block sizes, Atlantis USX still continues to provide a significant number of IOPS at low latencies resulting in an extremely high level of performance.

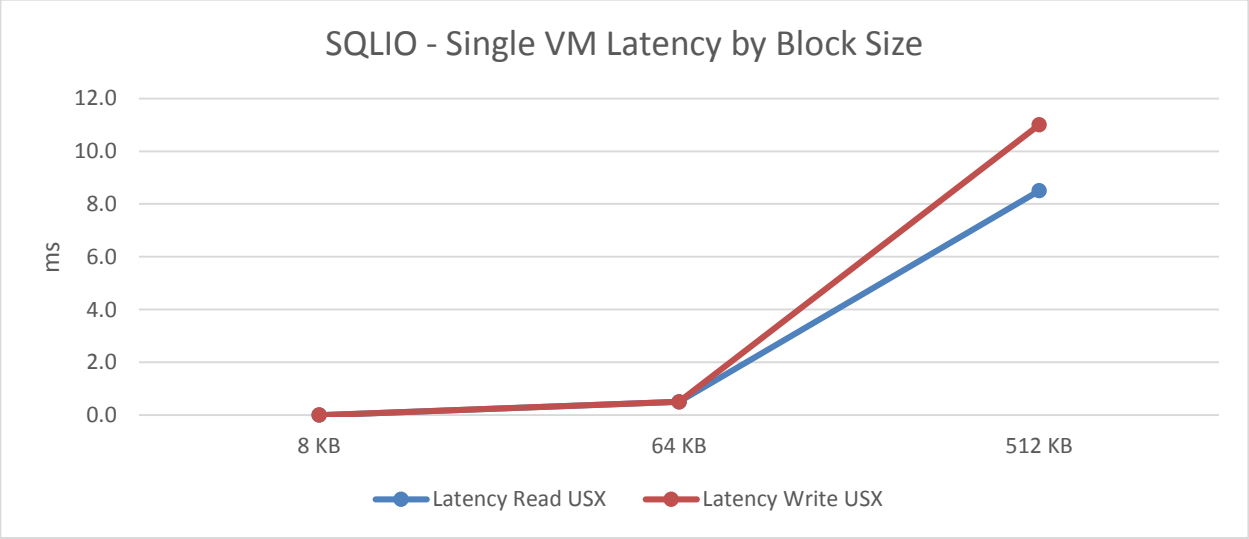

Figure 12: SQLIO – Single VM Latency by Block Size

While testing extreme and unlikely cases of 512KB, latency was manageable for very large block sizes while keeping low for the most common block levels. With this level of performance Atlantis USX provides, organizations can be comfortable scaling SQL efficiently.

Latency vs. Outstanding IO

The Latency vs Outstanding IO test helps demonstrate how efficient the underlying storage is at serving IO when requests are queued. This clarifies how the response time is affected when an application is demanding more IO than the storage system is providing at a given time.

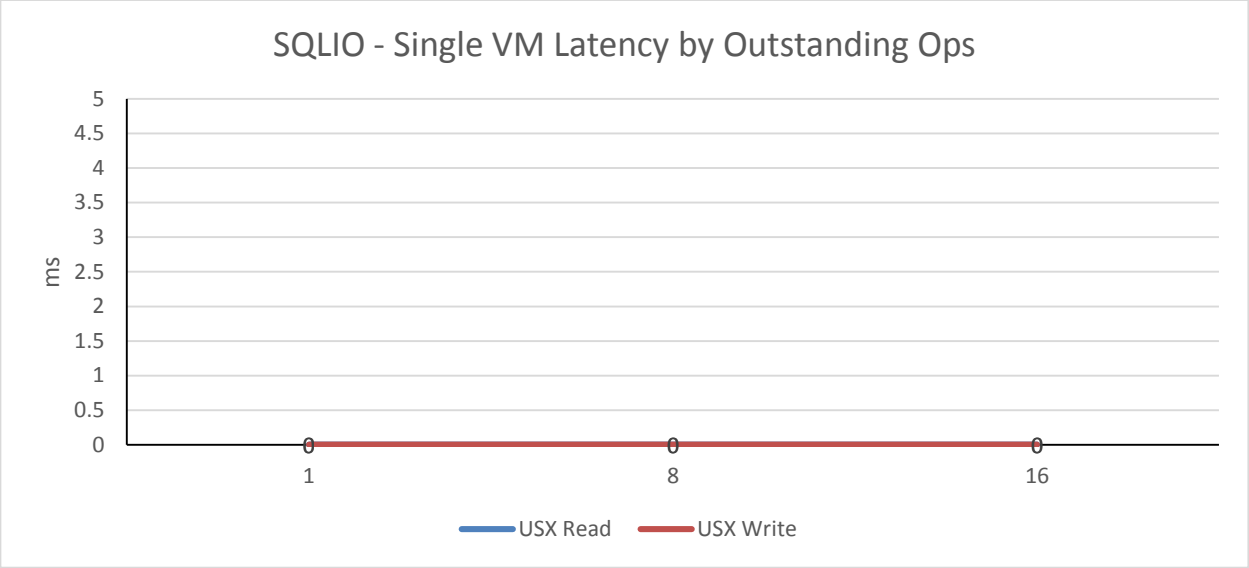

Figure 13: SQLIO – Single VM Latency by Outstanding Ops

In this exercise, the response time was too low for SQLIO to measure. With traditional SAN storage solutions, increasing latency would commonly be observed in this scenario.

Random IOPS by Block Size

Finally in this example, Atlantis USX demonstrates the ability to maintain high performance on both IOPS and throughput with random I/O over the 4-64KB range.

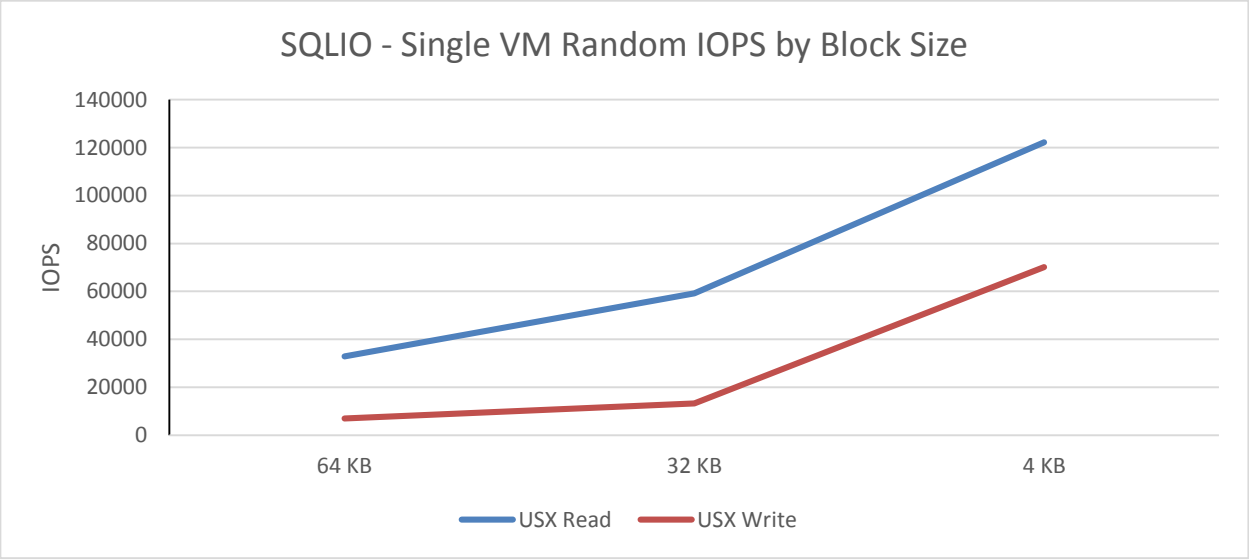

Figure 14: SQLIO – Single VM Random IOPS by Block Size

About the Authors

Patrick Brennan, Sr. Product Marketing Manager, Atlantis Computing

Twitter:@vPatrickB

Patrick leads the development of solutions and marketing campaigns around applications and platforms such as database, analytics, and VDI. Prior to Atlantis, Patrick worked at Citrix as a Sr. Solution Architect for Global Systems Integrator HP where he was responsible for architecting and developing service and product offerings related to virtualization, cloud, mobility, and datacenter transformation.

Jim Moyle, Infrastructure Technologist, Atlantis Computing

Twitter: @JimMoyle

Jim Moyle is an infrastructure technologist with almost two decades of passionate involvement in the IT industry. Over the last decade he has been a specialist in the server based computing, application, and desktop delivery sector. He has worked with many blue chip companies around the UK delivering and architecting some of the largest Citrix farms in Europe. He also speaks at various industry events such as BriForum, E2EVC, and Citrix Synergy.

References

1. Web: https://msdn.microsoft.com/en-us/library/dn133186.aspx

2. Web: https://channel9.msdn.com/Shows/Data-Exposed/SQL-Server-2014-In-memory-OLTP-Logging

3. Web: http://www.nimblestorage.com/blog/technology/storage-performance-benchmarks-are-useful-if-you-read-them-carefully/

4. Web: http://www.nimblestorage.com/blog/technology/storage-performance-benchmarks-are-useful-if-you-read-them-carefully/

atlantiscomputing.com

Atlantis Computing
U.S. Headquarters
2525 E. Charleston Road, Suite 100
Mountain View, CA 94043
PHONE: 650.917.9471
EMAIL: sales@atlantiscomputing.com
TWITTER: @AtlantisSDS

Atlantis Computing
European Headquarters
20 Birchin Lane
London, EC3V 9DU
PHONE: +44 2034052851
EMAIL: sales@atlantiscomputing.com
TWITTER: @AtlantisSDS