



Pivot3 Acuity with Microsoft SQL Server 2014

Reference Architecture



How to Contact Pivot3

Pivot3, Inc.

221 West 6th St., Suite 750

Austin, TX 78701

Tel: +1 512-807-2666

Fax: +1 512-807-2669

General Information: info@pivot3.com

Sales: sales@pivot3.com

Tech Support: support@pivot3.com

Website: www.pivot3.com

Online Support: support.pivot3.com

© Copyright 2018 Pivot3, Inc.

This document is for informational purposes only. Pivot3 reserves the right to make changes without further notice to any products herein. The content provided is as is and without express or implied warranties of any kind.

Table of Contents

SECTION 1	INTRODUCTION.....	4
SECTION 2	MICROSOFT SQL SERVER DATABASE OVERVIEW	5
SECTION 3	VIRTUALIZING SQL SERVER WITH PIVOT3 ACUITY	6
	NVME PCIE FLASH, MULTI-TIER ARCHITECTURE.....	6
	ADVANCED QOS.....	6
	DISTRIBUTED SCALE-OUT ARCHITECTURE.....	6
	PIVOT3’S PATENTED ERASURE CODING.....	7
	INTEGRATION & INTEROPERABILITY.....	7
	RICH DATA SERVICES	7
SECTION 4	REFERENCE SOLUTION ARCHITECTURE	8
	SYSTEM CONFIGURATION.....	8
	DEPLOYMENT & CONFIGURATION CONSIDERATIONS	10
SECTION 5	PERFORMANCE TESTS & RESULTS.....	12
	TEST SYSTEM CONFIGURATION.....	12
	HAMMERDB TEST CONFIGURATION.....	13
	PERFORMANCE RESULTS	13
SECTION 6	CONCLUSION.....	14
	CONSOLIDATE SQL SERVER DATABASES CONFIDENTLY WITH PIVOT3 ACUITY.....	14

Section 1 Introduction

Database infrastructure has played a key role in supporting many organizational workflows for several decades. They are central to many of the applications supporting critical organizational functions. Although databases have been around for a while, the infrastructure supporting them as well as the applications utilizing these databases have undergone a dramatic transformation. Today, many databases are deployed in virtual environments that are shared across multiple applications, creating new challenges around performance management. Pivot3 Acuity Hyperconverged Infrastructure (HCI) platform, with its policy-based Quality of Service (QoS) that leverages NVMe PCIe Flash, distributed scale-out architecture, and rich data services can cost-effectively address the needs of today's performance-sensitive, consolidated database environments while simplifying ongoing management and scaling.

This paper details a reference implementation for deploying Microsoft SQL server database on a Pivot3 Acuity HCI system. The paper also outlines results of performance tests conducted by Pivot3 for a 3-Node Acuity Flash cluster using the industry-standard Online Transaction Processing (OLTP) benchmarking tool HammerDB. It is assumed that the intended audience of this document has modest familiarity with VMware virtual environments, Microsoft SQL Server, and Pivot3 Acuity HCI.

Section 2 Microsoft SQL Server Database Overview

Microsoft SQL Server is one of the most widely used and fastest growing database platforms in use by enterprises today. SQL Server often provides critical backed data infrastructure for many applications that support daily organizational functions, including Customer Relationship Management (CRM), Supply Chain Management (SCM), Accounting and Financial Management, and many other industry and company-specific critical applications

The primary components of a SQL Server database platform are as follows:

Database Engine: A central part of SQL Server database platform that manages the relational databases deployed on the platform

Integration Service: Performs Extract-Transform-Load (ETL) operations to clean and process raw data for database inclusion

Service Broker: Provides messaging and queuing interface for applications that use the database

Reporting Service: Provides comprehensive reporting of the underlying relational databases

Analytics Service: Provides data analytics on underlying database. It can perform OLAP (On-Line Analytics Processing) and data mining.

Although many databases were initially deployed before the advent of virtualization, most organizations today have virtualized and consolidated their database infrastructure or are in the process of doing so. While virtualizing databases can enable economies of scale by improving resource utilization and reducing infrastructure footprint, it can add substantially to performance management complexities. As a result, many databases are virtualized only on the compute layer. Many organizations still manage the storage supporting these databases in traditional storage silos with dedicated SAN (Storage Area Network) or NAS (Network Attached Storage) devices.

With Pivot3's policy-based Acuity HCI platform, it is now possible to predictably virtualize and consolidate multiple SQL Server database instances and other enterprise applications on a shared software-defined platform while ensuring SLA compliance, simple operations, and predictable scalability.

Section 3 Virtualizing SQL Server with Pivot3 Acuity

Pivot3's policy-based, priority-aware, distributed scale-out architecture is ideally suited for virtualizing database infrastructure. With its NVMe PCIe Flash data path that is dynamically managed by its advanced Quality of Service (QoS), the Acuity HCI platform is able to deliver superior, consistent and predictable performance with vastly more Transactions Per Minute (TPM) and lower latency. Pivot3's advanced QoS allows customers to consolidate databases and applications of varying criticality on an Acuity HCI cluster, while ensuring priority and performance for workloads that are most critical to business. Its distributed scale-out architecture simplifies scaling of the infrastructure by delivering predictable and non-disruptive scalability.

With Acuity, IT professionals can consolidate, manage, and scale database and application environments in a simple manner while ensuring stringent SLA compliance.

NVMe PCIe Flash, Multi-Tier Architecture

Pivot3 incorporates a multi-tier architecture that delivers a better combination of performance, density, and economics, enabling a broader range of applications to run on HCI. It consists of RAM, NVME PCIe Flash, and either SSDs or HDDs. Of these tiers, NVMe PCIe Flash delivers ultra-low latency performance to databases and business critical applications in the following ways:

First, it handles every write in the system for faster response.

Secondly, it is used as an intelligent read cache that ensures mission-critical workloads always get priority. This capability is customer defined via Pivot3's advanced QoS policies.

Advanced QoS

Meet application SLAs with QoS policies that set minimum performance levels and prioritize workloads over each other during times of resource contention. Advanced QoS gives business a simple way to manage performance with five preset policies that govern minimum performance in terms of minimum IOPs, Throughput, and Response Time. Additionally, each policy has an associated service level policy that automates resource prioritization, ensuring critical applications don't impact the business during business times. Advanced QoS enables organizations to confidently consolidate multiple, mixed-application workloads and ensure the ones most important to the business deliver guaranteed, predictable performance and response time. Additionally, the policy-based QoS simplifies ongoing performance management by allowing users to schedule, or manually make, real-time QoS policy changes.

Distributed Scale-Out Architecture

Pivot3 Acuity platform is based on distributed scale-out architecture. Acuity aggregates all available resources in all nodes in a cluster to build a global pool of resources. All volumes, virtual machines, and data sets are uniformly distributed across all drives in the cluster. In a fully populated Acuity cluster, there are 144 SATA-controlled drives. All these drives participate in all I/O activities, resulting in inherently superior I/O performance for all data sets. This architecture mitigates I/O bottleneck out of the box, without specialized storage tuning or optimization, thus ensuring a consistent database performance under strenuous load conditions.

Additionally, the cluster is expanded by adding more nodes. The distributed scale-out architecture ensures predictable scalability of capacity, compute, I/O performance, and available bandwidth. Each node adds 20Gbps

Network and 20Gbps Application bandwidth to the cluster. As a result, predicting incremental hardware requirements for anticipated growth is easy and straightforward with Pivot3.

Pivot3's Patented Erasure Coding

Pivot3 HCI solutions utilize patented erasure coding (EC) to achieve high availability and fault tolerance. Conventional methods used to protect against component failures in HCI environments rely upon replicating data sets across the nodes to ensure availability of data in case of a node failure. This method proves to be inefficient, offering maximum usable capacity of 50%. Often, two copies are required in order to achieve desired availability goals, slashing usable capacity to 33%. In a fully populated cluster, Pivot3 EC delivers 82% usable capacity. Storage efficiency is further enhanced by Acuity's in-line data reduction. Pivot3 EC provides much better efficiencies while providing protection from node and drive failures. Additionally, EC eliminates write IO duplication in order to protect against failures, further boosting I/O performance of databases and applications.

Integration & Interoperability

Pivot3 Acuity nodes are preloaded with vSphere, so deployment and setup are very simple and straightforward processes. The Acuity vCenter Plugin further simplifies management of Infrastructure through integrated management, helping to provision, manage, and monitor Acuity HCI environments from native VMware interfaces. Acuity supports a comprehensive set of vStorage APIs for Array Integration (VAAI) functions that optimize many storage-intensive operations. Further, by leveraging DirectPath (vSphere PCIe Passthrough), Acuity bypasses the hypervisor for any storage-related I/O activities, removing a potential bottleneck in the I/O path and improving latency by up to 40%. Acuity also supports a broad set of native VMware capabilities including High Availability (HA), Distributed Resource Scheduling (DRS), and NSX network virtualization. This allows administrative flexibility.

Additionally, through its support for vStorage APIs for Data Protection (VADP) and Microsoft Volume Shadow Service (VSS), Acuity provides application-consistent, space-efficient snapshots and clones that can simplify management and protection of database environments.

Rich Data Services

Pivot3 Acuity Hyperconverged platform offers a rich set of data services with its all-inclusive licensing model. These data services include Thin Provisioning, space-efficient application-aware snapshots, clones, replication, and data reduction. This combined with Pivot3's ability to deliver higher densities by eliminating I/O bottleneck allows Pivot3 Acuity to deliver up to 3X density with up to 6X the overall performance.

Section 4 Reference Solution Architecture

System Configuration

This section details examples of both Pivot3 Acuity Hybrid and Flash system configurations for deploying Microsoft SQL Server workloads. Pivot3 conducted exercises to validating SQL Server reference solution architecture in December 2017. This exercise used SQL Server 2014 as the database software and Windows Server 2012 R2 as the guest operating systems for SQL Server Virtual Machines (VMs). The configuration recommendations provided in this section are for a SQL Server OLTP (On-Line Transaction Processing) database workload.

Flash Systems

Pivot3 Acuity Flash minimum system configuration includes two Pivot3 X5-6500 Flash Accelerator nodes and one X5-6000 Flash Standard node. The Accelerator nodes each have 1.6TB NVMe PCIe Flash storage. Each node has 8 SATA-controlled SSD drives, which are available in sizes ranging from 480GB to 3.84TB. Each node is available with RAM ranging from 256GB to 1.5TB. Clusters can be scaled by adding more Acuity X5-6000 Flash Standard nodes with identical configurations.

NOTE: The cluster can be scaled to up to a total of 16 nodes (2X X5-6500 Accelerator nodes + 14X X5-6000 Standard nodes). Multiple Acuity clusters can be managed as a single domain from the management interface within VMware vCenter.

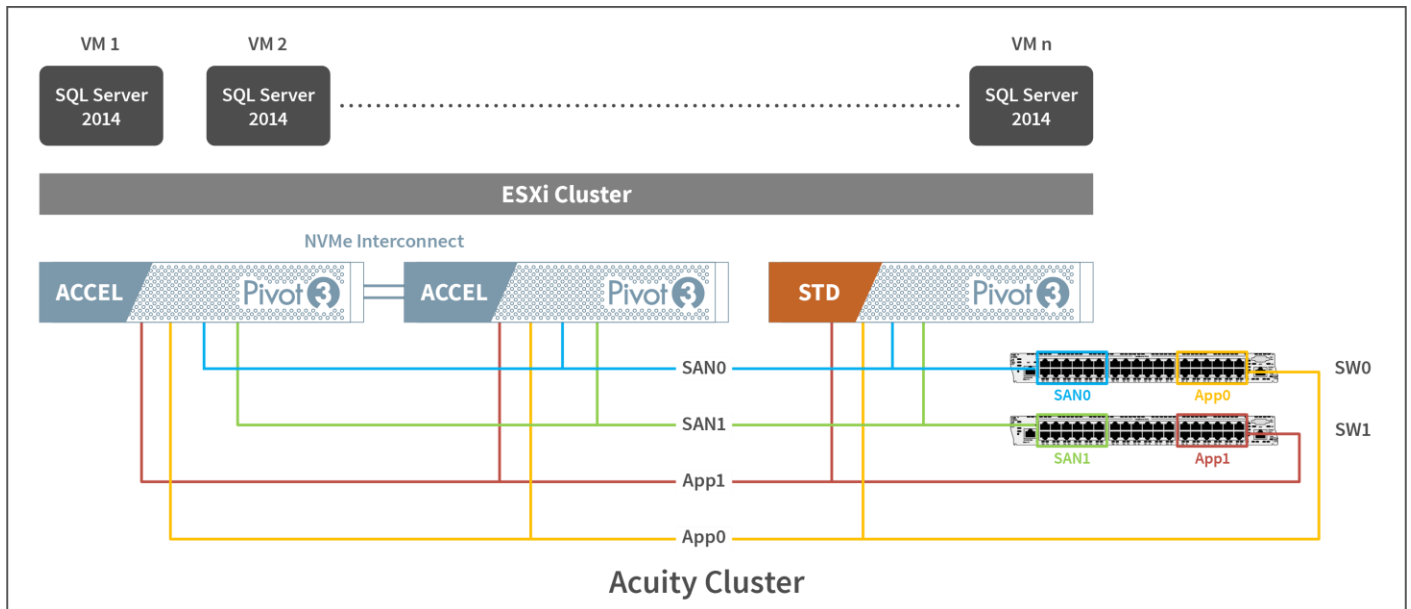
Information	Details
System/Cluster	2× Acuity X5-6500 + 1 or more Acuity X5-6000
CPU	Per Node: 2× Intel® Xeon® E5-2695v4 (36 physical [72 logical] cores)
RAM	Per Node: 256GB to 1.5TB
NVMe PCIe Flash	2× 1.6TB per cluster
Disks	Per Node: 8× 480GB to 3.84TB SATA SSDs
PCIe NIC	Per Node: 2× 10GbE SAN Per Node: 2× 10GbE Application Per Node: 2× 1 or 10GbE Management 2× 10GbE NVMe Interconnect
Hypervisor	ESXi 6.0.0.0-5050593
Pivot3 Acuity Version	2.1
Networking Switches	2× 10GbE 10BaseT (SFP+ Optional)
Database VM Guest OS	Windows Server 2012 R2
Database Software	SQL Server 2014

Hybrid Systems

Pivot3 Acuity Hybrid minimum system configuration includes two X5-2500 Hybrid Accelerator Nodes and one X5-2000 Hybrid Standard Node. The Accelerator nodes each have 1.6, 2, or 3.2TB NVMe PCIe Flash storage. Each node has 12 SATA-controlled SSD drives, which are available in sizes ranging from 1TB to 4TB. Each node is available with RAM ranging from 256GB to 1.5TB. Clusters can be scaled by adding more Acuity X5-2000 Flash Standard nodes with identical configurations.

NOTE: The cluster can be scaled to up to a total of 12 nodes (2× X5-2500 Accelerator nodes + 10× X5-2000 Standard nodes). Multiple Acuity clusters can be managed as a single domain from the management interface within VMware vCenter.

Information	Details
System/Cluster	2× Acuity X5-2500 + 1 or more Acuity X5-2000
CPU	Per Node: 2× Intel® Xeon® E5-2695v4 (36 physical [72 logical] cores)
RAM	Per Node: 256GB-1.5TB
NVMe PCIe Flash	2× 1.6TB, 2TB, or 3.2TB per cluster
Disks	Per Node: 12× 1-4TB SATA HDDs
PCIe NIC	Per Node: 2× 10GbE SAN Per Node: 2× 10GbE Application Per Node: 2× 1 or 10GbE Management 2× 10GbE NVMe Interconnect
Hypervisor	ESXi 6.0.0.0-5050593
Pivot3 Acuity Version	2.1
Networking Switches	2× 10GbE 10BaseT (SFP+ Optional)
Database VM Guest OS	Windows Server 2012 R2
Database Software	SQL Server 2014



Deployment & Configuration Considerations

Pivot3 Acuity and VMware ESXi Deployment and Configuration

Standard Pivot3 Acuity deployment procedures should be followed for setting up the initial cluster for virtualized SQL Server deployment. The Acuity Setup Guide (https://support.pivot3.com/doc_categories/manuals) provides instructions and considerations for deploying an Acuity cluster.

After the deployment of an Acuity cluster, the following configuration changes should be made to ensure maximum performance for SQL Server workloads.

- In the firmware setting of the physical host, set the UEFI (Universal Extensible Firmware Interface) to **High Performance**.
- Set each ESXi node performance to **Maximum**.

NOTE: Pivot3 Acuity supports VMware DRS and HA capabilities. These should be used to ensure optimal workload distribution and uninterrupted operations in case of a host failure.

Guest VM Configuration Considerations

The guest VMs that host the SQL Server workload should be configured according to the performance requirements of the database workload and the size of the database. Depending on performance requirements, assign more or less vCPU and RAM to the VM. In general, higher RAM and vCPU allocation for the guest VM will support higher OLTP transaction performance and lower response times for the workload.

Each guest VM hosting the SQL Server workload should have:

- 2× vNICs using VMXNET 3 adapters for 2× 10GbE application networks
- 2× SCSI controllers, LSI logic for Boot volume of the VM, and Para-Virtual for DB and Logs volumes
- 3× volumes, one each for Boot, DB, and Logs data
 - The Boot volume will contain the OS data.
 - The DB volume will host the SQL Server database.
 - The Log volume will host the SQL Server logs.

NOTE: *Separating the volumes is important to ensure superior performance and for Acuity Quality of Service (QoS) to effectively manage and prioritize the correct data sets.*

Guest OS Windows Server 2012 R2 Configuration

Windows Server 2012 R2 was used as the guest OS during this exercise. Default Windows Server configurations are recommended. The following changes should be made to the Windows Server Guest OS after deploying:

- Windows Server Power Policy should be set to **Maximum Performance**.
- DBs and Log volumes should be formatted **NTFS** using default block sizes.

SQL Server 2014 Configuration

Default processes for SQL Server installation and configuration should be followed while virtualizing SQL Server on Pivot3 Acuity. VMware's best practices for virtualizing SQL Server on VMware vSphere can be found on VMware's website at <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/sql-server-on-vmware-best-practices-guide.pdf>.

Acuity Quality of Service Considerations

Pivot3 Acuity offers advanced, policy-based QoS that helps prioritize workloads and deliver predictable performance in environments that may host workloads of varying performance requirements and criticality. QoS policy level of 1 indicates highest priority while QoS policy level of 5 indicates lowest priority. When consolidating production and non-production SQL Server databases on Acuity HCI, Pivot3 recommends assigning mission and business critical workloads a QoS policy level of 1 or 2 and lower priority levels to other workloads.

Acuity also supports scheduled QoS policy changes. In instances where a particular workload requires higher performance during predefined windows, this scheduling function can simply performance management by dynamically managing QoS schedules.

Section 5 Performance Tests & Results

This section summarizes the performance tests and results for SQL Server performance testing using HammerDB Workload Generator as tested by Pivot3 in December of 2017.

HammerDB is a versatile database load testing tool commonly used with many of the leading database platforms, including SQL Server. HammerDB can be used for testing and benchmarking database performance on a variety of operating systems; by simulating order creation and processing transactions, it allows customers to test and benchmark database performance for various combinations of HammerDB Warehouses and Users.

HammerDB provides metrics on Transactions Per Minute (TPM) and New Orders Per Minute (NOPM) for the chosen combination of Warehouses and Users by generating load for the underlying database.

Test System Configuration

For this exercise, a 3-node Pivot3 Acuity Flash cluster with 2× X5-6500 Accelerator nodes and 1× X5-6000 Standard node was used. The Accelerator nodes had 1.6TB NVMe PCIe Flash storage each. Each node had 8× 800GB SATA-controlled SSD drives, dual 18-core Intel Xeon E5-2695v4 processors, and 512GB RAM.

Information	Details
System/Cluster	2× Acuity X5-6500 + 1 Acuity X5-6000
CPU	Per Node: 2× Intel® Xeon® E5-2695v4 (36 physical [72 logical] cores) Total System: 3× 2x Intel® Xeon® E5-2695v4 (108 physical [216 logical] cores)
RAM	Per Node: 512GB
NVMe PCIe Flash	2× 1.6TB per cluster
Disks	Per Node: 8× 800GB SATA SSDs
PCIe NIC	Per Node: 2× 10GbE SAN Per Node: 2× 10GbE Application Per Node: 2× 10GbE Management 2× 10GbE NVMe Interconnect
Hypervisor	ESXi 6.0.0.0-5050593
Pivot3 Acuity Version	2.1
Other Configuration Details	vCenter Server Appliance 6.5.0 Embedded
Networking Switches	2× 10GbE 10BaseT
Database VM Guest OS	Windows Server 2012 R2
Database Software	SQL Server 2014

HammerDB Test Configuration

Three Windows Server 2012 R2 guest VMs with 32× vCPU and 384GB RAM each were configured to host three SQL Server database instances. Each VM was configured for 5000 HammerDB warehouses. Three HammerDB tests were run for 128, 256, and 512 concurrent users; each test was run for a duration of 5 minutes. Transactions Per Minute (TPMs) and New Orders Per Minute (NOPMs) were recorded for each user count.

Performance Results

The below table summarizes the HammerDB OLTP test results for a 5000-warehouse tests for 128, 256, and 512 concurrent users on a 3-node Acuity Flash systems cluster.

# of Warehouses	# of Users	Transactions Per Minute	New Orders Per Minute
5000	128	1,461,531	317,938
5000	256	2,231,898	484,929
5000	512	2,859,156	617,940

The results above demonstrate that Pivot3 Acuity is able to deliver industry-leading transactional performance when compared with other HCI products with similar hardware configurations for SQL Server OLTP workload. Pivot3 chose this specific hardware configuration for this exercise because we expect most of our customers to deploy Pivot3 Acuity clusters with this configuration for OLTP database workloads. It is possible to achieve higher transactional performance by adding more RAM to each of the nodes.

Section 6 Conclusion

Consolidate SQL Server Databases Confidently with Pivot3 Acuity

Pivot3 Acuity provides an ideal platform for consolidating SQL Server databases as well as other enterprise applications. Its NVMe data path coupled with advanced QoS delivers predictable and high transactional performance, while its distributed scale-out architecture simplifies scaling exercises. Further, its performance-optimized architecture reduces the infrastructure footprint by eliminating I/O bottleneck, delivering up to 3X the density per node with up to 6X overall performance.

With Pivot3, IT organizations can confidently consolidate databases and applications of varying performance characteristics and business criticality while ensuring stringent SLA adherence, simple operations, and predictable scalability.