Pivot3

Pivot3 Acuity with Splunk Enterprise

Reference Architecture

vmware READY

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*

Table of Contents

Introduction

The nature of business analytics has undergone dramatic transformation over the last decade. With proliferation of connected devices and digitization of organizational workflows, traditional approaches to business intelligence have become inadequate. New architectures that take fundamentally different approaches utilizing Big Data Analytics have evolved to address the complexities of a modern data-centric world to provide critical business insights. Splunk is a leading Big Data Analytics software solution that allows organizations to flexibly address their business analytics needs for today's highly digitized, highly-connected enterprise.

This paper details a reference implementation for deploying a clustered Splunk solution on a Pivot3 Acuity Hyperconverged Infrastructure (HCI). This reference architecture was tested and validated in one of Pivot3's Center of Excellence labs by Pivot3 engineers in close partnership with a Splunk Elite Services partner.

This reference implementation demonstrates that the Pivot3 Acuity HCI system can deliver industry leading Splunk performance, while simplifying management with intelligent policy-based operation. The architecture delivers consistent and predictable search and indexing performance while lowering the hardware footprint and VM sprawl.

This document is intended for an audience with a modest familiarity with VMware virtual environments, Splunk Enterprise (Splunk) and Pivot3 Acuity HCI.

Splunk Enterprise Overview

Splunk is an enterprise-grade scalable Big Data Analytics platform for collecting, searching, monitoring and analyzing machine data. Splunk can collect and analyze data from a wide array of sources including logs from various systems and applications, machine generated data, social media platforms, application APIs and traditional databases. Splunk's architecture is designed to scale from the needs of small businesses that may ingest a few MBs of data for analysis to large enterprises that may ingest TBs of data per day with a total data retention of many PBs. Initially designed to run on commodity X86 server nodes, many medium to large scale deployments of Splunk today are virtualized. Hyperconverged platforms, utilizing standards-based X86 hardware with right architectural attributes, can help optimize and simplify Splunk deployments.

Splunk Components

Splunk deployments are comprised of three types of Splunk components as shown in the figure below.

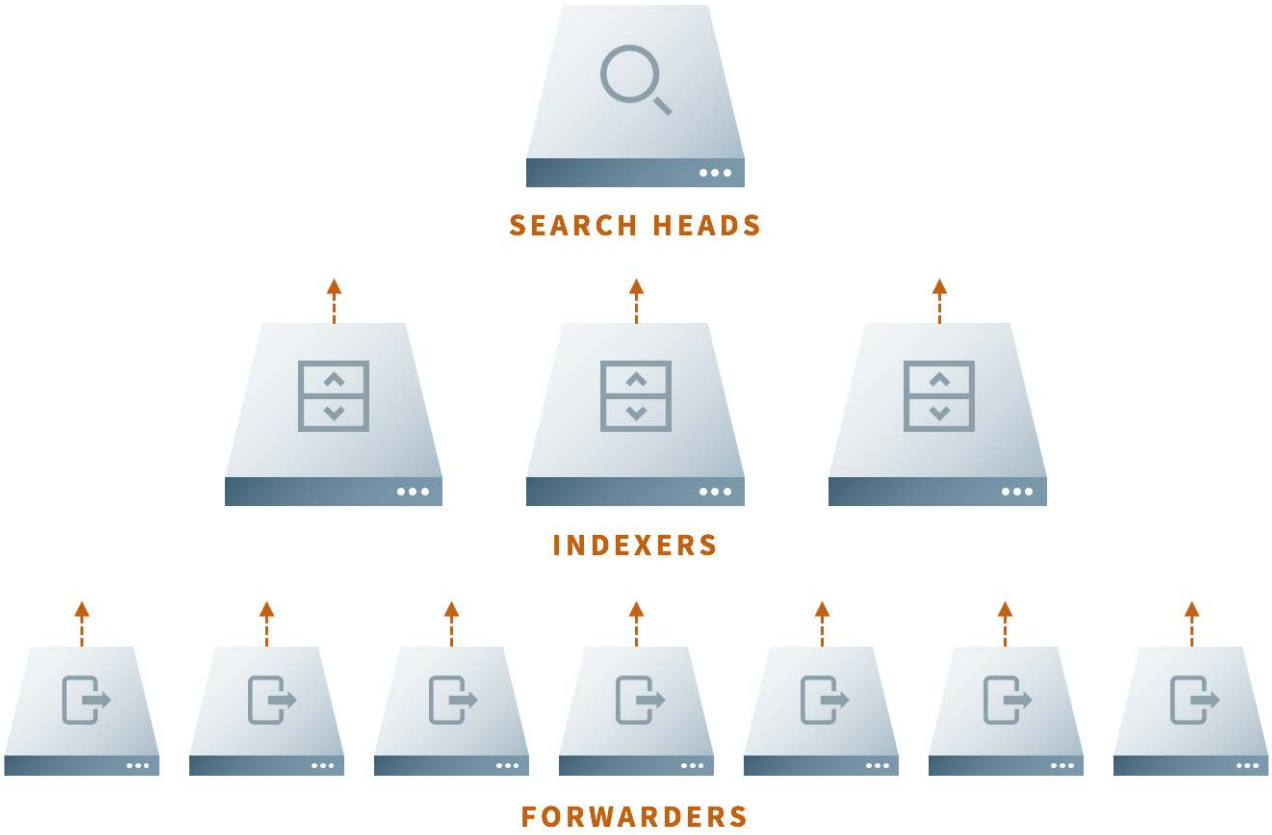

Figure 1: Splunk Components in a Clustered Splunk Deployment

Splunk Forwarders

Splunk Forwarders forward the raw data to Splunk indexers. Splunk Forwarders are deployed on devices and applications where data generates. The data can be log files, syslog, output from a script, REST API inputs, database queries, etc. Typically, many forwarders forward the data to a set of Splunk Indexers deployed on a Splunk cluster. The traffic is usually load-balanced for optimal Indexer utilization.

Splunk Indexers

Splunk Indexers are the central part of a Splunk deployment. They store and index incoming data from the Forwarders and respond to search requests. They use MapReduce to store, process and search large amounts of unstructured data. Typically, a Splunk deployment consists of multiple Indexers. These Indexers are load-balanced for optimal utilization.

Search Heads

Splunk Search Heads are the front-end of the Splunk deployment and are typically accessed via Splunk web interface or APIs to run searches. Search Heads in turn employ Indexers to search and retrieve underlying data.

Splunk Data Buckets

Splunk categorizes and stores data in data buckets. There are three categories of Splunk data buckets: Hot/Warm, Cold and Frozen. As the data is ingested and processed by the Indexers, it is written to Hot/Warm buckets. Based on retention policies, the data is then moved over to Cold data buckets. Hot/Warm data buckets are usually relatively small and the data movement to Cold bucket can be governed by the age of the data or the size of the data bucket. Data is then moved to Frozen buckets from Cold buckets based on retention policies. Typically, data stays in Cold buckets for many weeks or months, before moving to Frozen buckets. The data in Hot/Warm and Cold buckets is searchable by Splunk Search Heads, while the data in Frozen buckets is not.

Splunk Infrastructure Requirements

The Splunk Infrastructure needs to support extreme objectives in terms of performance, capacity and cost-efficiencies. The high-ingest nature of Splunk data requires fast Indexing performance (measured in Mega Bytes per Second [MBPS]) to process and store large amounts of data quickly. The search performance (measured as Events or Buckets per Second) needs to be consistent and high to support large numbers of search queries in a timely manner. On the other hand, a Splunk deployment gathers large amounts of data over time that needs to be stored and made available when required. This requires scalable, cost-effective and high capacity storage on the backend. These extreme requirements for high performance and cost-effective capacity can put a strain on many Splunk deployments that often become too complex, cumbersome and inflexible as they scale.

The ideal infrastructure solution supporting Splunk deployments needs to demonstrate following capabilities:

- Low-latency high-performance storage access and high throughput for fast Indexing and Search performance
- Cost-effective scalable high-capacity storage to support large data retention requirements
- Simple architecture and policy-based infrastructure management that delivers performance and capacity goals while lowering operational overhead.

Optimizing Splunk Deployments with Pivot3 HCI

Pivot3's policy-based, priority-aware distributed scale-out architecture is ideally suited for data and processing intensive Big Data Analytics environments. With its NVMe PCIe Flash Datapath that is dynamically managed by its advanced Quality of Service (QoS), the Acuity HCI platform can deliver superior indexing throughput and search performance. Its scalable multi-tier architecture allows cost-effective high-capacity storage for large sets of data and long-term data retention needs. Pivot3 simplifies and streamlines Splunk deployments by combining the benefits of high-performance NVMe PCIe Flash and cost-effective scalable HDD storage into one easy to manage platform. Its policy-based QoS that can be updated in real-time, helps simplify operational workflows for Splunk administrators and infrastructure teams.

With Pivot3, IT can effortlessly and cost-effectively deploy and manage demanding Splunk deployments to generate real-time insights into their businesses.

NVMe PCIe Flash, Multi-Tier Architecture

Pivot3 incorporates a multi-tier architecture that delivers a better combination of performance, density and economics to meet objectives of large Splunk deployments for high performance and large data storage. The architecture consists of RAM, NVME PCIe Flash and SATA-controlled drives. These tiers are intelligently managed by Acuity QoS. The architecture delivers ultra-fast performance for indexing and search activities, while providing cost-effective high-capacity storage for data that may not have high performance needs.

Additionally, with the elimination of storage I/O bottleneck by efficiently utilizing NVME PCIe Flash resources, Pivot3 helps improve Splunk Indexer and Search Head densities while reducing the total infrastructure footprint and VM sprawl.

Advanced QoS

Pivot3's advanced QoS gives businesses a simple way to manage performance for data sets or applications with five preset policies that govern minimum performance in terms of IOPs, and Throughput and maximum Response Time. Each policy has an associated service level that automates resource prioritization to ensure that critical and performance intensive operations meet their SLAs.

Pivot3's QoS policies align well with how Splunk organizes Hot/Warm, Cold and Frozen data buckets, allowing administrators to apply appropriate policies to each of the data buckets. This simplifies operational management and ensures consistent search and indexing outcomes without ongoing performance optimization. The QoS policies can be dynamically changed in real-time. This vastly simplifies operational workflows involved in data management by providing on-demand performance for specific data sets when needed and in real-time.

Distributed Scale-Out Architecture

The 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 the I/O bottleneck out of the box, without specialized storage tuning or optimization.

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 additional copies are required to achieve desired availability goals, slashing usable capacity to 33%. In a fully populated cluster, Pivot3 EC delivers 82% usable capacity. Pivot3 EC provides far better efficiencies while providing protection from node and drive failures. Additionally, EC eliminates write I/O duplication to protect against failures, further boosting I/O performance for better indexing and search outcomes.

Integrations and Interoperability

Pivot3 Acuity nodes are preloaded with vSphere, so deployment and setup are 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. 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.

Reference Solution Architecture

This section details a reference solution architecture utilizing a 3 Node Pivot3 Acuity HCI Hybrid cluster for a clustered Splunk deployment. Pivot3 recommends its Acuity Hybrid systems for Splunk deployments. Acuity Hybrid systems provide a combination of low-latency NVMe PCIe Flash Datapath and scalable, cost-effective HDD storage for long term data retention. The architecture and test results presented in this and subsequent sections are based on testing conducted in the Pivot3 Center of Excellence lab by Pivot3 engineers partnered with a Splunk Elite Services partner.

System Configuration

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 hard disk drives (HDD), 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 non-disruptively by deploying additional Acuity X5-2000 Hybrid 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.*

Table 1: Pivot3 Acuity HCI Hybrid Cluster System Configuration

System Configuration	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
Disks	**Per Node:** 12× 1-4TB SATA HDD
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)
Splunk VM OS	Linux 64 bit
Splunk Software	Splunk® Enterprise 7.0.2

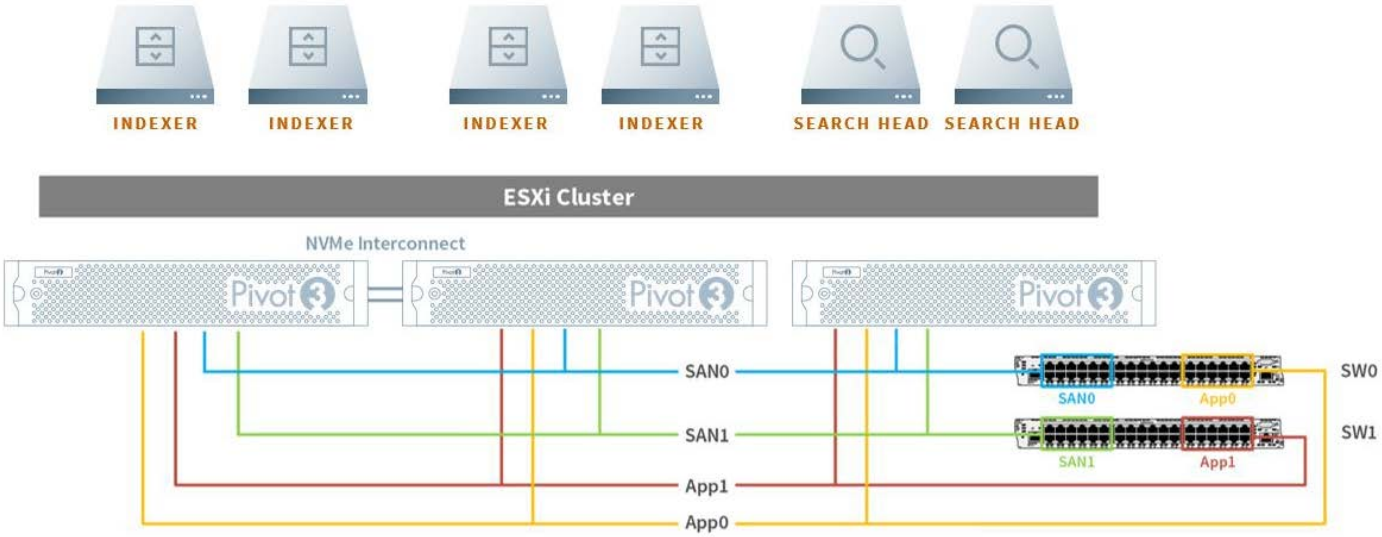

Figure 2: 3 Node Pivot3 Acuity Solution Architecture for Clustered Splunk Deployment

Deployment and Optimization Considerations

Pivot3 Acuity and VMware ESXi Deployment

Standard Pivot3 Acuity deployment procedures found in the *Acuity Setup & User Guide* should be followed for setting up the initial cluster for Splunk deployment.

Splunk Deployment and Data Bucket Optimization

This section outlines considerations for optimizing performance of the clustered Splunk deployment on Pivot3 Acuity. Standard procedure for installing and deploying Splunk components should be followed. More information can be found at *https://docs.splunk.com/Documentation/Splunk/7.0.2/Installation/*

Splunk VM performance optimizations

To optimize Splunk VM performance, Pivot 3 recommends the following:

- All Splunk Indexer and Search Head node VMs should be provisioned as Eager-Zero Thick provisioned VMs. There should not be any over-provisioning of vCPU allocated to these VMs.
- Splunk Indexer node VMs should be evenly distributed on two of the Acuity Accelerator nodes in the cluster. This ensures that the Indexers have the lowest latency access to Acuity's NVMe PCIe Flash Datapath. If the two accelerator nodes do not have enough RAM and CPU resources necessary to host all the Indexer nodes, the remaining Indexer nodes should be distributed evenly on the Standard nodes in the cluster after exhausting Accelerator node resources.
- Search Head node VMs should be evenly distributed over the cluster to balance the overall Indexer and Search Head workloads on each of the Acuity nodes.

Splunk Data Bucket and Acuity QoS Optimization

To optimizes the Splunk Data Bucket and Acuity QoS policies, Pivot3 recommends the following:

- All OS environments for Splunk VMs should be hosted on a separate data volume with the volume assigned an Acuity QoS policy of 2 or 3 (Business Critical).
- There should be an even number of volumes hosting Hot/Warm Splunk data buckets. These volumes should be evenly assigned to Acuity Controller 1 and Acuity Controller 2. This will ensure balanced utilization of Acuity Data path and NVMe PCIe Flash storage resources. These volumes should be assigned an Acuity QoS policy of 1 (Mission Critical).
- There should be an even number of volumes hosting Cold data buckets. These volumes should be evenly assigned to Acuity Controller 1 and Acuity Controller 2, with an Acuity QoS policy of 2 or 3 (Business Critical).
- Data volumes containing Frozen data buckets should be assigned an Acuity QoS policy of 4 or 5 (Non-Critical). Frozen data has the lowest performance requirements. Assigning a lower QoS policy will ensure that Frozen data volumes are not consuming higher-performance storage resources. With Acuity's policy-based QoS, Frozen data can be easily thawed and made available for analysis at high performance by simply switching QoS policy to a higher priority in real-time.

NOTE: *Sizing of the volumes hosting Splunk data buckets will be influenced by expected data ingest rates, retention policies and NVMe PCIe Flash storage available in the system. For specific requirements, please contact a Pivot3 representative or channel partner.*

Performance Tests and Results

This section summarizes the performance results for a clustered Splunk deployment on a 3 node Pivot3 Acuity Hybrid system.

Benchmark testing on the Pivot3 Acuity Hybrid Cluster was performed to establish architecture design and recommended configuration baselines for common Splunk data workloads. The emphasis for the testing was to measure performance and QoS metrics that showcase Splunk.

Test scenarios matched typical Splunk installations. Tests were performed to find peak indexing performance on the Acuity platform. Additional tests were performed to find search performance under moderate indexing load.

Test System Configuration

For this exercise, a 3-node Pivot3 Acuity Hybrid cluster with 2× X5-2500 Accelerator nodes and 1× X5-2000 Standard node was used. The Accelerator nodes had 1.6TB NVMe PCIe Flash storage each. Each node had 12× 2TB SATA-controlled Hard Disk Drives, dual 18-core Intel Xeon E5-2695v4 processors and 512GB RAM.

Table 2: Pivot3 Acuity Hybrid Cluster System Configuration

System Configuration	Details
System Cluster	2× Acuity X5-2500 + 1 Acuity X5-2000
CPU	**Per Node:** 2× Intel® Xeon® E5-2695v4 (36 physical [72 logical] cores)
RAM	**Per Node:** 512GB
NVMe PCIe Flash	2× 1.6TB
Disks	**Per Node:** 12× 2TB SATA HDD
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
Networking Switches	2× 10GbE 10BaseT
Splunk VM OS	Linux 64 bit
Splunk Software	Splunk® Enterprise 7.0.2

NOTE: *For this exercise there were two volumes to host Hot/Warm data buckets and two volumes to host Cold data buckets.*

Single Splunk Instance Indexing Performance Results

This test utilized a Splunk Light instance that bundles an Indexer and a Search Head. More information on Splunk Light can be found at *https://www.splunk.com/en_us/products/splunk-light.html*. The Indexer was configured with two indexing pipelines. **Table 3** shows the performance comparison of the Acuity configuration with Splunk Reference Hardware recommendations. More information on Splunk Reference Hardware can be found at *http://docs.splunk.com/Documentation/Splunk/7.0.3/Capacity/Referencehardware*

Table 3: Single Splunk Instance Indexing Performance Results

	Max Index Rate
Splunk Reference	20,000.00 KB/Sec
Acuity Hybrid	81,136.48 KB/Sec

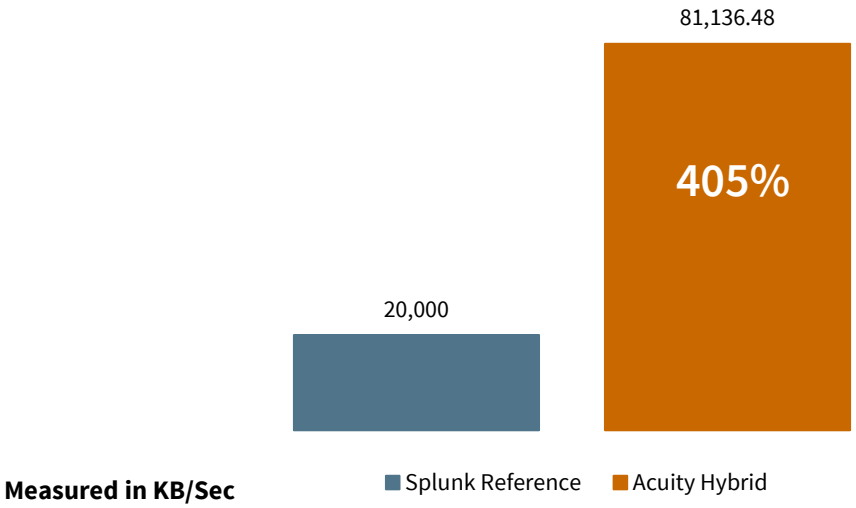

Figure 3: Indexing Performance for Single Instance Splunk Deployment

Clustered Indexing Performance Results

This test was designed to highlight the Indexing performance of Splunk on the Pivot3 Acuity Hybrid system. The test utilized 3 Indexers with 2 indexing pipelines each. **Table 4** shows the performance comparison of the Acuity configuration with Splunk Reference Hardware recommendations.

Table 4: Clustered Indexing Performance Results

Reference	Max Index Rate
Splunk Reference	60,000.00 KB/Sec
Acuity Hybrid	153,967.00 KB/Sec

Search Performance Under Concurrent Indexing Load Results

This test was designed to highlight the search performance of Splunk on the Pivot3 Acuity Hybrid system under a moderate indexing load. The test utilized 1 Indexer and 1 Search Head. **Table 5** shows the performance comparison of the Acuity configuration with Splunk Reference Hardware recommendations.

Table 5: Search Performance Under Concurrent Indexing Load Results

Search Type	Dense	Sparse	Rare
Splunk Reference	<150,000 Events/Sec	<15,000 Events/Sec	<30-150 Buckets/Sec
Acuity Hybrid	**482,415.46 Events/Sec**	**96,890.56 Events/Sec**	**226.30 Buckets/Sec**

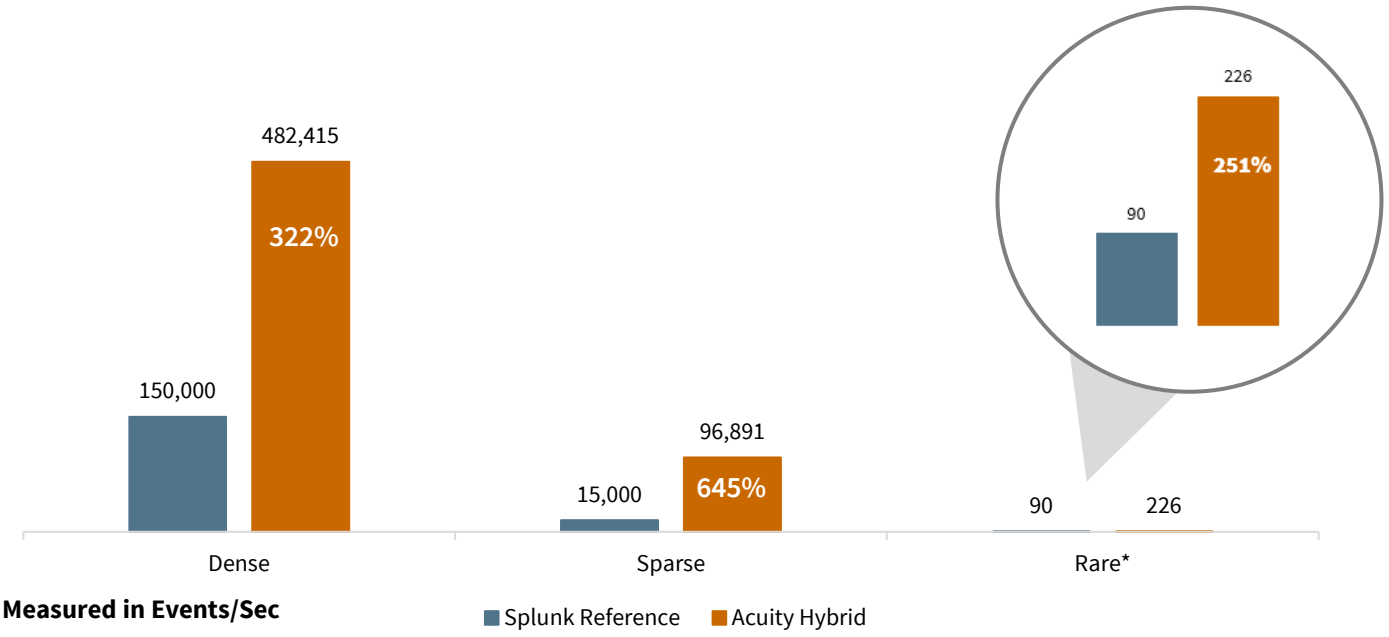

***Rare search results are measured in Buckets/Sec.**

Figure 4: Search Performance under Moderate Indexing Load for Clustered Splunk

NOTE: *Splunk Reference performance calculated by assuming linear scalability of reference hardware. Realistically it should be significantly lower.*

Test Results Summary

The results demonstrate that the Pivot3 Acuity system can deliver industry leading Splunk performance, while simplifying performance management by intelligent use of advanced QoS. The architecture delivers consistent and predictable search and indexing performance while lowering the hardware footprint and VM sprawl.

Conclusion

Pivot3 HCI provides an ideal platform for deploying Splunk analytics workloads as well as other Big Data Analytics solutions. Its multi-tier architecture that combines NVMe PCIe Flash Datapath with advanced QoS delivers superior indexing and search performance while proving cost-effective storage for large data sets. Further, the QoS drastically simplifies ongoing operational workflows by proving a policy-based way for applying, scheduling and modifying performance priorities.

With Pivot3, enterprises can achieve industry-leading performance for their Splunk environments while slashing footprint, reducing VM sprawl and simplifying operations.