



# Big Data Analytics

Optimize Your Splunk and Big Data Deployments with Pivot3

## KEY ADVANTAGES



### Accelerate Performance

4 to 6X indexing and search performance



### Control Sprawl

Less hardware and fewer VMs to manage



### Simplify Management

Policy-based Quality of Service simplifies performance management



### Scale Efficiently

Distributed architecture and patented erasure coding simplify scaling and deliver storage efficiency

With the proliferation of connected devices and digitization of organizational workflows, businesses are looking at new architectures that take fundamentally different approaches to utilizing Big Data Analytics to address the complexities of a modern data-centric world and provide critical business insights.

Splunk is the market-leading, enterprise-grade Big Data Analytics platform for collecting, searching, monitoring and analyzing machine data. Splunk can analyze data from a wide array of sources including logs, machine-generated data, social media platforms, application APIs and traditional databases. Initially designed to run on commodity X86 nodes, many Splunk deployments today are virtualized for simplicity and efficiency.

## Infrastructure Requirements for Splunk

Infrastructure for Splunk needs to support extreme objectives in terms of performance, capacity and cost-efficiencies. The high-ingest nature of data requires fast indexing performance to process and store large amounts of data quickly. The search performance 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. This requires scalable, cost-effective storage on the backend. These 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.

## A Policy-Based Foundation for Splunk and Big Data Analytics

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) engine, Pivot3 can deliver superior indexing throughput and search performance. Its scalable multi-tier architecture delivers cost-effective high-capacity storage for large datasets and long-term 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. Performance and protection policies that can be updated in real-time, 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.

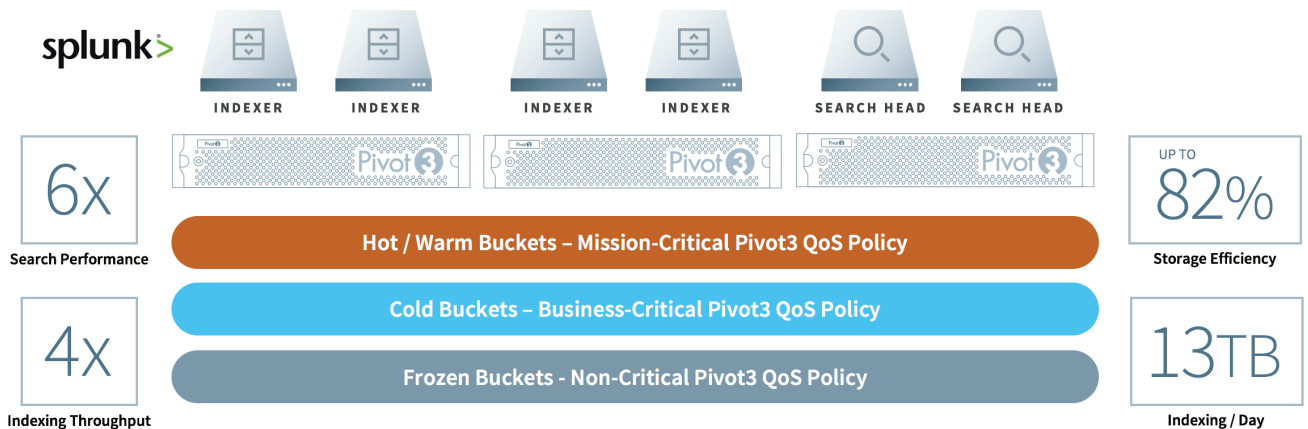
### More Indexing and Searching in Smaller Footprint with NVMe PCIe Flash Datapath

Pivot3's multi-tier architecture that intelligently utilizes NVMe PCIe Flash storage, RAM cache, and SATA-controlled drives, eliminates potential IO bottleneck in indexing and search operations, delivering higher performance with lower hardware footprint and fewer VMs to manage. In recent lab testing, Pivot3 delivered more than 6X the search and 4X the indexing speeds compared Splunk Reference Hardware recommendations. A single 3 node Pivot3 Hybrid cluster was able to deliver indexing rates that would support over 13TB of indexing per day.

### Simple Data Management with Policy-Based QoS

Pivot3's advance QoS that is both policy-based and priority-aware, dramatically simplifies management of Splunk data buckets. Splunk categorizes data into Hot/Warm, Cold, and Frozen buckets depending on its age and retention policies. These buckets have distinct performance requirements. With Pivot3's QoS, Splunk administrators can easily apply right QoS policies to right data buckets to ensure optimal performance under any condition. Hot/Warm data buckets can be applied a Mission-Critical QoS policy, while Cold data buckets can be assigned a Business-Critical policy and Frozen data buckets can be assigned a Non-Critical policy to ensure superior indexing and search performance while eliminating cumbersome performance optimization exercises. QoS policies can be changed in real-time or scheduled at known points in time. This simplifies operational workflows. For example, if there is a compliance requirement that necessitates analysis of older data that is stored in Frozen buckets, Splunk admins can easily switch QoS policies on appropriate data buckets to higher priority and make that data available to search and analysis in no time and at a higher performance level.

#### PIVOT3: POLICY-BASED HIGH-PERFORMANCE PLATFORM FOR SPLUNK AND BIG DATA ANALYTICS



## Scale and Efficiency

Splunk deployments gather large amounts of data over time. In many situations, this data needs to be retained for long periods of time due to regulatory and compliance reasons. Pivot3's scale-out architecture and its patented erasure coding (EC) ensure smooth scalability for Splunk data infrastructure while providing high storage efficiency. Pivot3 architecture aggregates all resources - storage, compute, bandwidth and distributed cache - to form a unified resource pool. As new nodes are added to the existing cluster, the storage capacity, IO capacity and bandwidth pool available to applications scales seamlessly and non-disruptively. Additionally, Pivot3's patented EC enables up to 82% usable capacity while protecting the infrastructure from component and node failures. This compares to less than 50% efficiency observed in most HCI products that rely on replicating data across nodes for fault-tolerance.

## Optimize Splunk and Big Data Deployments with Pivot3

With Pivot3, enterprises can achieve market-leading performance for their Splunk and Big Data environments while slashing footprint, reducing VM sprawl and simplifying administration. Its multi-tier architecture that combines NVMe PCIe Flash Datapath with advanced QoS delivers superior indexing and search performance while providing cost-effective storage for large data sets. Further, the QoS drastically simplifies ongoing operational workflows with a policy-based way for applying, scheduling and modifying performance priorities for Splunk data buckets. Pivot3 offers the next-generation policy-based platform for optimizing Splunk and Big Data deployments in a simple, scalable, cost-effective manner.

---

## About Pivot3

Pivot3 provides easy to use and efficient solutions that deliver intelligence and automation to simplify surveillance, IoT, and multi-cloud environments. Built on hyperconverged infrastructure, Pivot3's advanced intelligence engine combines policy-based management, inference and orchestration capabilities, and industry-leading performance for automated management, resource allocation, workload mobility and data placement so IT can focus on delivering the applications and services that power your business. With more than 2,500 customers in 54 countries and over 18,000 deployments in healthcare, government, transportation, security, entertainment, education, gaming, retail and more, Pivot3 is defining the future of intelligent hybrid cloud computing with smarter infrastructure solutions.

**For more information, visit [Pivot3.com](https://www.pivot3.com)**