

Top Reasons to Deploy Your Big Data Analytics on Pivot3 Hyperconverged Infrastructure



Simple, Unified Architecture that Addresses All Big Data Needs

Pivot3 HCI features a multi-tier storage architecture utilizing NVMe PCIe Flash, RAM cache, and SATA-controlled drives that is managed by Pivot3's Intelligence Engine. By prioritizing performance of data buckets, the Intelligence Engine offers a single platform that can address diverse performance, scalability and cost-efficiency objectives of a Big Data Analytics environment using a simple policy-based framework.



More Search and Indexing with Less Infrastructure

Pivot3's NVMe PCIe Flash datapath eliminates potential IO bottleneck in search and indexing operations found in Big Data Analytics environments. As a result, more searching and indexing is achieved with less infrastructure and fewer VMs. In recent lab testing for Splunk enterprise workload, Pivot3 HCI delivered 6X search and 4X indexing speeds compared to Splunk Reference Hardware recommendations. A single 3 node Pivot3 HCI Hybrid cluster was able to deliver indexing speeds to support over 13TB of indexing per day.



Cost-Effective for Long-Term Retention

Pivot3 Hybrid HCI solutions that utilize high-capacity, cost-effective HDDs as one of its tiers are ideally suited to address the long-term retention needs of many Big Data Analytics environments. Pivot3 further reduces cost of the storage with its patented erasure coding (EC) that delivers up to 82% usable capacity. Additionally, Pivot3 solutions are designed to work with leading cloud service providers, simplifying data movement to the Cloud for retention and archival.



On-demand Scalability

Many Big Data Analytics environments grow at a consistent pace over time. This is because of a predictable amount of new data that gets ingested in the environment on daily basis. Pivot3's distributed scale-out architecture provides a flexible way for enterprises to add capacity to their deployment in small steps and only when required. The architecture aggregates all resources from all nodes in the cluster, and enables new nodes to be added on-demand and non-disruptively.

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