Pivot3 RAIGE™ Architecture

TECHNICAL OVERVIEW

Breaking the Limits of Physical RAID

> Flexibility and investment protection with totally virtualized environment

> Boundless performance as capacity scales

> New standard of availability with virtual sparing and high-bandwidth recovery

> Simple provisioning and management

The Storage Performance Challenge

Data intensive applications continue to demand additional storage capacity and performance at an unprecedented rate. Direct attach storage cannot scale to meet the requirements. Traditional networked storage is capacity bound by the choice of controller. Clustered storage products offer the promise of scaling both capacity and performance, but clustered solutions built on traditional hardware RAID architectures are still bound by box-level performance. The ground-breaking distributed RAID architecture from Pivot3 removes traditional barriers of clustered storage and provides unparalleled scalability in performance and capacity.

Distributed Virtualized Environment

Pivot3 RAIGE™ (RAID Across Independent Gigabit Ethernet) supports the most flexible and extensible system possible without the designed-in physical barriers or limitations found in legacy implementations.

> RAID is distributed and executed in parallel across multiple Databanks connected via common gigabit Ethernet as opposed to single, performance-bound controllers or sets of controllers.

> Capacity and performance scales flexibly with dynamically changeable RAID levels and volume definitions. Databanks are simply added to scale capacity to hundreds of Terabytes.

> Virtualized environment protects your investment. Future expansion and feature set is not limited by your current hardware investment choices.

Performance Plus

IOP and bandwidth performance scales linearly as Databanks are added to a RAIGE cluster. No controller-based limits or bounds. No upfront costs. Just pay as you grow.

> Every Databank includes processing, caching and network ports, increasing system performance as capacity is added.

> Every client has direct access to every Databank—no node-hopping required—providing lower latencies and highly parallel throughput.

> Data is automatically distributed across the entire system providing maximum performance. As the system scales, data is rebalanced and performance increases.

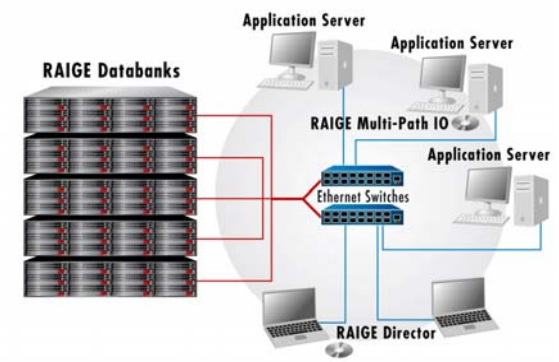

Pivot3 RAIGE™ Industry Firsts

Distributed RAID: RAIGE is a distributed RAID architecture that eliminates the need to replicate data across Databanks while still protecting against a drive or Databank failure. This break-through technology substantially decreases the overhead needed for fault tolerance and provides maximum efficiency.

Parallel I/O Throughput: RAIGE Dynamic Multi-path IO running on client systems directly distributes data transfers across all the Databanks maximizing overall efficiency and throughput while protecting against a switch or network failure.

Scalable Performance: Each Databank adds processing power, cache and network ports which contribute to overall performance and bandwidth. RAIGE harnesses and delivers this parallel processing power even to a single workload, guaranteeing scalable performance. The result is enterprise-class scaling and throughput from affordable processors and SATA drives.

Virtual Sparing: RAIGE virtual sparing distributes spare capacity throughout all components in the system. This allows all system assets (disks and controllers) to contribute to the performance of the running system rather than sit idle waiting for a failure.

Rapid Recovery: In a drive or Databank recovery scenario, every drive and Databank contributes to the process. This parallelization, only possible with a distributed RAID design, drastically reduces the time necessary to recover from a failure.

Data Availability

Pivot3 RAIGE turns a standards-based hardware investment into enterprise-class protection.

> RAIGE insures data access in case of network failures, drive failures and even Databank failures for availability generally found only in the most costly systems.

> With virtualized sparing, there are no wasted or idle assets. All disk drives and Databanks are utilized to deliver maximum performance.

> Data recovery times are up to ten times faster with system-wide parallel processing and a proprietary algorithm which optimizes the rebuild process.

> RAIGE data protection has low overhead so your data is protected without substantially increasing your storage capacity needs.

Simple Management

RAIGE Director provides an easy-to-use top down view of your storage hierarchy and all its important characteristics. You can dynamically increase Logical Volume capacity, change Logical Volume RAID Levels, and view key performance metrics of your RAIGE Storage without interrupting data access. You no longer have to be concerned with the physical grouping of your storage. RAIGE Director does it for you and simplifies storage management to an easy-to-understand logical level. RAIGE Director CLI provides for batch mode scripting operation.

RAIGE Architecture

Standard server hardware and industry-standard protocols ensure the lowest possible acquisition and infrastructure management costs. The Pivot3 RAIGE architecture minimizes the investment needed to achieve scalable, high-performing, always-on storage.

Contact Pivot3 for an in-depth architecture paper and to learn how we can help ease your storage challenges.

Contact Pivot3:

Pivot3, Inc
6605 Cypresswood Drive
Spring, TX 77379
www.pivot3.com
Tel: 1.877. 574.8683
Fax: 281.516.6099

DS RA V2.3 March 2008