PolyServe Matrix Server for Linux

Highly Available, Shared Data Clustering Software

PolyServe Matrix Server for Linux is shared data clustering software that allows customers to replace UNIX SMP servers with clusters of price-performant Linux servers. Matrix Server enables applications to be scaled out across an easy to manage cluster of Linux servers.

Matrix Server software enables multiple Linux-based servers to function as a single, easy-to-use, highly available system. Matrix Server is the only clustering solution to integrate next generation shared file system technology with high availability services (also included in Matrix HA) and Multi-Path I/O functionality into a single, easy-to-use framework for cluster management. PolyServe Matrix Server requires the presences of a Storage Area Network (SAN).

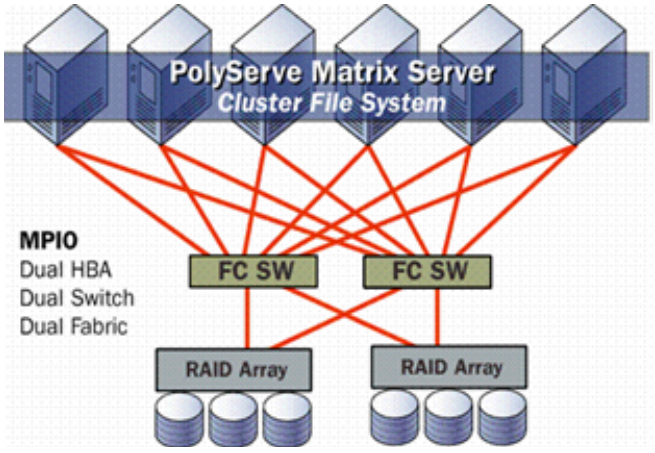

Figure 1. With Matrix Server, all servers in the cluster have direct, simultaneous access to file systems in the SAN

Unlike traditional High Availability (HA) products, PolyServe Matrix Server comprises a fully symmetric cluster file system that enables scalable data sharing, high availability services that eliminate single points of failure to increase system uptime, and cluster and storage management capabilities for managing servers and storage as one. Matrix Server delivers an unparalleled level of scalability, availability and manageability in support of database, file, e-mail, Web and media serving applications.

In addition, PolyServe Matrix Server simplifies application and storage management in the data center. Traditionally, clusters of servers required an application to be installed and

Product Highlights

- High-performance, symmetric cluster file system
- High-availability infrastructure
- Integrated cluster-wide management

Key Benefits

- Replace expensive UNIX servers with scalable Linux clusters
- Increase system availability across applications, data, networks, and servers
- Consolidate data and servers to improve productivity and reduce cost

patched on each server in the cluster.

PolyServe Matrix Server facilitates a single, cluster wide shared file system image for application binaries, configuration information, and data. A customer can install an application into the shared file system once, and have that application run on all nodes in the cluster without additional installation or modification. Consequently, the shared file system allows administrator to patch, change configuration in one location and have those changes be reflected cluster wide. One single point of management.

Also, since the application data is accessible to all nodes in the cluster, there is no need to replicate data from server to server. Replication results in multiple copies of the same data and wasted storage space. With PolyServe, there is a single pool of application data that can be update once and the changes propagated to all servers in the cluster. A shared file system image greatly reduces the time it takes to administer a cluster of Linux servers.

When considering clusters of Linux servers for mission-critical applications, data center managers require a computing infrastructure that is both recoverable and highly available. Matrix Server is the first clustering product to fully deliver on these challenging requirements. The product's fully symmetric design, where all servers in a cluster are peers, coupled with its multi-path I/O functionality guarantees there will be no server bottlenecks or single points of failure cluster-wide.

Benefit Highlights

- **On Demand Scalability** - Enables administrators to aggregate multiple Linux servers (up to 16 nodes) and manage them as a single coherent entity reducing administrative steps and costs
- **Enables Shared Data** - Unlike "shared nothing" clustering, servers within the cluster concurrently read and write data across the SAN resulting in high bandwidth I/O throughput for demanding applications.
- **Simplified Management** - applications can be installed and patched once cluster wide, and applications data
- **Flexible Monitoring and Failover** - Monitors application, network, cluster and hardware health and provides automated failover and fail back eliminating administrative burden and service downtime
- **Broad Application Support** - Supports databases (IBM DB2, Oracle8/8*i*/9*i*, and Oracle9*i* RAC), applications (ERP, CRM, custom), web servers (Apache), application servers (BEA WebLogic, IBM Websphere, and Oracle9*i* Application Server), e-mail, and scalable file serving with NFS.
- **Multi-Platform Support** - Supports all major Intel based server hardware platforms and the most common SAN implementations to ensure easy evaluation and seamless integration into existing infrastructures

Matrix Server Functionality

PolyServe Matrix Server is based on several distributed computing breakthroughs. At the center of Matrix Server is a fully distributed, fully journaled cluster file system (CFS) that supports online additions and deletions of nodes and concurrent multi-node access to shared data. It also includes a completely symmetric, distributed lock manager, no master/slave relationships among servers, and innovative, distributed metadata management, which avoids the single server bottleneck on file operations.

Reduce Costs and Simplify Clustering

With Matrix Server, companies reduce costs and complexity associated with managing Linux based server and blade clusters. Responding to enterprises' need for simplified cluster management and zero tolerance for downtime, Matrix Server uses simple and flexible failover and fail back policies for M:N clusters and supports multiple simultaneous failures of network, storage and servers. The result is a hands-off, self-healing cluster, lower administrative costs, and increased freedom for IT personnel.

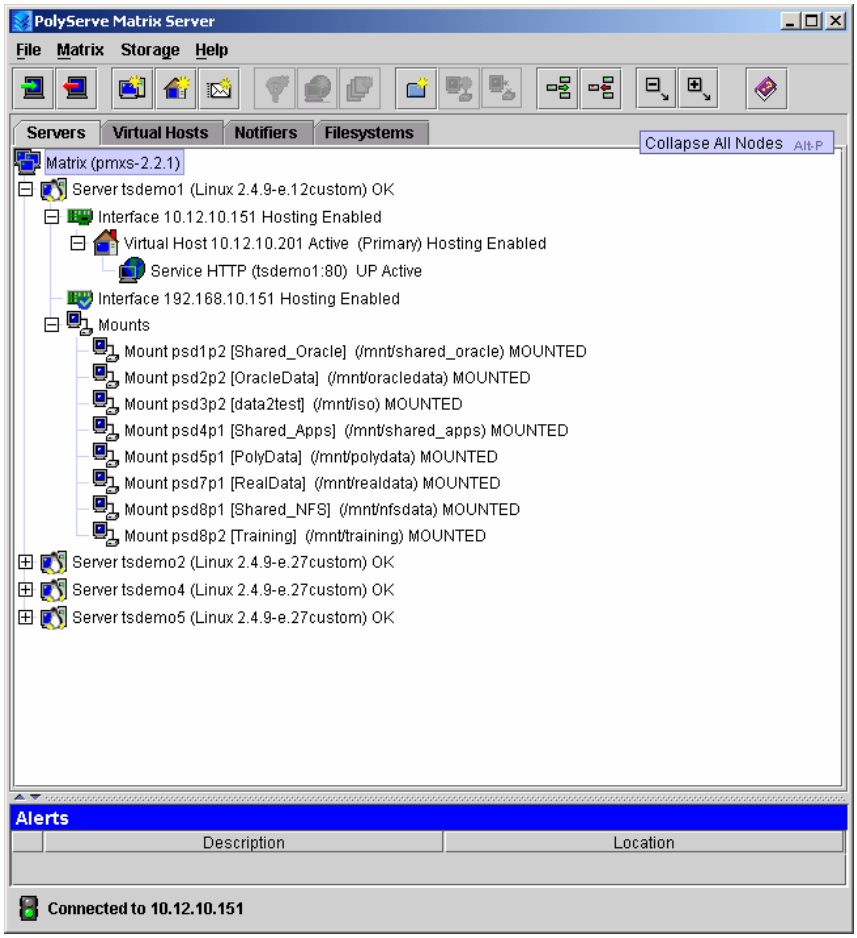

Figure 2. PolyServe Matrix Server Cluster Management Console

Offers a New Alternative for Hardware Purchases

Enabling more than just enhanced availability, Matrix Server allows concurrent read and write data access across a shared file system for up to 16 server cluster nodes. This new level of concurrent access to business critical data along with the low cost nature of blade or volume servers makes Matrix Server clusters a superior choice over large expensive SMP UNIX servers.

Specifically, the horizontal scalability of the product also allows IT budget stakeholders to consider "scaling-out" using multiple two (2) CPU servers as opposed to single 8 or 16 CPU "scale-up" servers resulting in a hardware cost savings often exceeding 85% d far exceeding that on proprietary SMP UNIX servers.

The diagram below illustrates the ability of Matrix Server to aggregate computing resources and simplify "scale-out" server clustering.

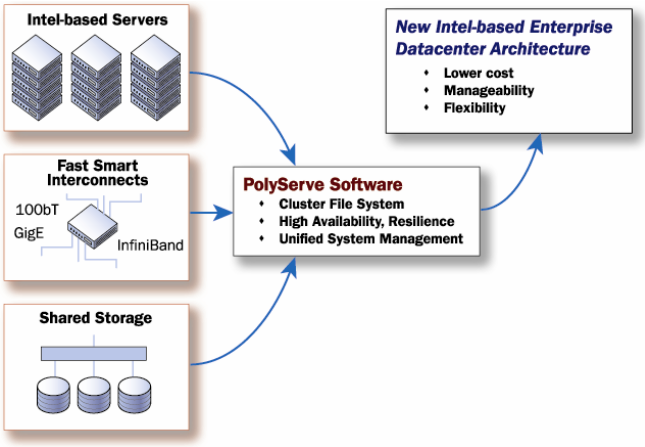

Figure 3. Matrix Server simplifies and aggregates "scale-out" server clusters

Application and Solution Support

Matrix Server augments the scalability, availability, and manageability of databases, application servers, web servers and file servers. To an application, Matrix Server behaves like a single-node file system so applications run unchanged. On Matrix Server for Linux, third party applications that use the POSIX APIs will run without recompilation or modification to their executables.

Solutions Scenarios Enabled by Matrix Server

Solution	Customer Needs	Description and Benefits
Oracle9*i* RAC	• Enables migration off large Unix systems to more easily managed and supported Linux servers • Lower TCO and spend less time developing, integrating, and managing applications	• Improves the manageability, availability and recoverability in Oracle9i RAC environments • Build larger, more flexible clusters than is possible with current technology • Provide single point of installation, configuration, and patch updates through a cluster wide shared Oracle home. • Provides the manageability of a file system with the same or better performance than raw partitions • Improves performance of Oracle9i Data Warehousing functions such as ETL
Scalable File Serving (NFS)	• Obtain a higher throughput, higher availability file serving environment than deployed today • Utilize cost-effective industry standard building blocks rather than expensive, slow to upgrade NAS filers or UNIX servers	• Aggregate file serving performance of up to 16 nodes on industry standard hardware while eliminating bottlenecks and single points of failure • Adds failover infrastructure to achieve a highly available enterprise wide namespace • Allows servers and storage to be scaled on demand as need dictates
Server Consolidation for Oracle 8, 8i, 9i Databases	• Consolidate large numbers of database servers to increase utilization on active servers and redeploy passive ones • Improve uptime for databases • Achieve a more unified view of company's servers, storage, and databases	• Eliminates active-passive pairs increasing utilization and reducing hardware, software and management costs • Provides inherent high availability where every server node becomes a backup to others for all Oracle instances • Enables very fast application and Oracle failover and recovery • Build larger, more flexible clusters than is possible with current technology • Provide single point of installation, configuration, and patch updates through a cluster wide shared Oracle home. • Provides the manageability of a file system with the same or better performance than raw partitions
Server Consolidation for IBM DB2 Databases	• Consolidate large numbers of database servers to increase utilization on active servers and redeploy passive ones • Improve uptime for databases • Achieve a more unified view of company's servers, storage, and databases	• Eliminates active-passive pairs increasing utilization and reducing hardware, software and management costs • Provides inherent high availability where every server node becomes a backup to others for all DB2 instances • Enables very fast application and DB2 failover and recovery • Build larger, more flexible clusters than is possible with current technology • Provide single point of installation, configuration, and patch updates through a cluster wide shared DB2 home. • Provides the manageability of a file system with the same or better performance than raw partitions
Web Content Serving (Apache)	• Simplify application or web site deployment, maintenance and change management • Centralize data to increase storage utilization and intra-company data accessibility	• Simplifies application or web content publishing and maintenance by enabling multiple servers to concurrently share a single copy of data (or content) eliminating redundant server-to-server file replications and synchronization challenges • Enables consolidation of isolated pools of direct attached storage (DAS) onto shared SAN infrastructure, resulting in reduced management costs and increased availability • Enables data or log file consolidation for easy post-processing, backup, or archiving for compliance audits
Application Serving (IBM WebSphere, BEA	• Simplify application or web site deployment, maintenance and change management • Centralize data to increase	• Share one common set of application binaries, Java Virtual Machine (JVM) files, configuration information, and web content • Single point of administration for installation, patches, and

WebLogic, and Oracle9*i* Application Server)	storage utilization and intra-company data accessibility	configuration • Improved high availability and flexible failover of the application server farm • Enables consolidation of isolated pools of direct attached storage (DAS) onto shared SAN infrastructure, resulting in reduced management costs and increased availability • Enables data or log file consolidation for easy post-processing, backup, or archiving for compliance audits
Data Processing and Data Sharing	• Parallel process large data sets • Share data from one server to another as part of a business workflow process	• Provide better performance for sharing data than traditional file transfer protocols (e.g. FTP, NFS, or CIFS) • Support the ability to share data and parallel process data at high speeds with full data integrity
Backup (Legato NetWorker, Veritas NetBackup)	• Achieve zero downtime backup window • Gain flexibility in terms of deployment of backup servers and storage	• Enable the ability to initiate backup procedures from any node in the cluster • No system downtime to perform backups • Backups can be processed in parallel to other operations in the cluster

Feature Benefits

Cluster File System	
16 node Scalability	Distributed lock and metadata management ensures there is no single point of failure or single system bottleneck in the cluster, avoiding the typical limitations of competing CFS or other clustering offerings
Shared Data File System	Enables concurrent, fast access for multiple servers to “shared data” as opposed to “shared nothing” clustering approaches
Journaled File System	Allows file system to quickly recover and maintain data coherency among all nodes in the cluster
Single System Semantics	An application running on top of Matrix Server does not need to be cluster aware. Matrix Server for Linux is fully compliant with all POSIX API calls and exhibits the same behavior as an ext3 file system
High Availability	
Online Insertion or Deletion of Nodes	An administrator can add or delete a node from the cluster without pausing or halting the processing of other nodes in the system resulting in no downtime
High Availability Monitoring	Integrated and customizable monitors assess health of applications, operating system, servers, network, and storage throughout the cluster and will initiate and oversee failover and fail back
Flexible Failover Configurations	Supports N:1, N:N, and N:M failover models and allows fine grained control over failover and fail back policies
Cascading NIC Failover Support	Multiple NICs can be configured as standby interfaces and will transparently handle network failover
Multi-Path I/O	Supports seamless failover and multiple data paths to shared storage through multiple HBAs and redundant switches while maintaining data integrity
Fabric Fencing	I/O fencing prevents rogue servers from outside the cluster or a malfunctioning server within the cluster from accessing the SAN
Manageability	
Integrated Fabric Management	Supports Brocade and McData fabrics and manages fabric load balancing within nodes and across the cluster. Supports recovery of failed components within fabric and automatically re-enables data paths when failures are corrected
Context Dependent Symbolic Links (CDSLs)	Allows administrators to create shared application images with node specific configuration files. CDSLs map designated filenames to node specific files, while still maintaining consistent pathnames and directory trees cluster wide. The use of CDSLs results in improved cluster management and simplified installation and configuration.
Consistent Device	Guarantees that each server in the cluster refers to cluster-wide accessible devices using the same

Naming	name and prevents device slippage when new (or rebooted) servers join the cluster
LUN Discovery and Management	Allows administrator to discover LUNs that exist in the shared storage and easily map those LUNs to file systems
Central GUI Management Console and CLI	Powerful GUI and Command Line Interface (CLI) that allows administrator to configure and manage the cluster from a central management station. The GUI provides a single unified view of the cluster its resources (file systems, application services, LUNs, Virtual IP addresses) and allows the administrator to drill down and reconfigure where necessary. In addition, an administrator can administer the cluster from a remote client or any node in the cluster – all changes are instantly propagated clusterwide.

Hardware and Platform Support

Operating System Support	Red Hat, Red Hat Enterprise Server, SuSe Linux Enterprise Server (SLES), and other United Linux distributions
Fibre Channel HBAs Support	QLogic and Emulex HBAs
Fibre Channel Switch Support	Brocade, McData, and QLogic switches
Server Support	Intel IA32 or compatible servers and blades

PolyServe, Inc.
20400 NW Amberwood Drive
Suite 150
Beaverton, OR 97006

(877) 476-5973 (within U.S.)
(503) 617-7592 (fax)
(503) 617-7574 (outside U.S.)
www.polyserve.com

Mxs linux ds 102003