

The Path to Lower-Cost, Scalable, Highly Available Windows® File Serving

Higher Performance, Modular Expansion, Fault Tolerance at a Lower Cost

File servers are ubiquitous in today's Windows-based data center for sharing data (files, media, etc.) among many clients over a TCP/IP network. However, as the number of file servers and associated storage has grown rapidly, IT organizations have faced multiple challenges:

- Manageability
- Scalability
- Availability
- Utilization
- Cost

What is needed is a way to provide scalable, highly available CIFS file services using inexpensive, industry-standard servers that can be added to incrementally as demands require – while retaining the management simplicity of a single server and a single pool of exported file systems. PolyServe's Matrix Server software, with the File Serving Option, provides breakthrough scalability and failover support for the CIFS protocol.

The Challenges of Cost Effective, Scalable File Services for Windows 3

Traditional, Expensive Upgrade or Growth Paths..... 3

Enabling a Lower-Cost, Flexible Infrastructure for File Servers 4

Applicability to Server Consolidation, NT4 Migration and NAS Replacement Projects 5

Matrix Server Technology – Key Components..... 6

CIFS Manageability - Simplifying File Serving Deployments..... 7

Cost-Effective Scalability..... 7

Enabling Shared Data 9

Affordable High Availability..... 10

Simplify Management and Reduce the Cost of High Availability 10

Reduced Total Cost of Ownership (TCO)..... 11

 Scenario 1 (Small Environment) 11

 Scenario 2 (Large Environment) 12

The File Serving Utility Features and Benefits for CIFS File Serving 13

Summary 14

The Challenges of Cost Effective, Scalable File Services for Windows

File servers are ubiquitous in today's Windows-based data center for sharing data (files, media, etc.) among many clients over a TCP/IP network. However, as the number of file servers and associated storage has grown rapidly, IT organizations have faced multiple challenges:

- **Manageability.** Ever-larger numbers of clients have driven many IT organizations to deploy multiple Windows or NAS file servers, introducing management complexity and overhead in maintaining redundant copies of data.
- **Scalability.** Some have turned to very high-end servers and NAS appliances to simplify the provisioning of Common Internet File System (CIFS) services to a large number of clients or to support heavy workloads. However, this comes at a steep cost and, in some cases, this high-end hardware still represents a performance bottleneck between clients and data.
- **Availability.** As CIFS has increasingly been used for mission-critical tasks, it has been necessary to purchase hardware in pairs to provide high availability. This introduces yet more expense.
- **Utilization.** With the proliferation of departmental standalone servers with dedicated storage, it is impossible to know whether these investments are being well utilized. Many file servers and their storage are less than 20% utilized.
- **Cost.** As the number of islands of storage that require backup and administration increase, there is pressure to reduce costs in infrastructure and its maintenance.

What is needed is a way to provide *scalable, highly available* CIFS file services using inexpensive, industry-standard servers that can be added to incrementally as demands requires – while retaining the *management simplicity* of a single server and a single pool of exported file systems. HP's PolyServe File Serving Utility™, based on Matrix Server technology, provides breakthrough scalability and failover support for the CIFS protocol.

Traditional, Expensive Upgrade or Growth Paths

As the capacity of CIFS servers is exceeded for particular applications, customers have traditionally been forced to choose between upgrading to larger file servers or NAS appliances or adding more of them, perhaps replicating the same data among them. These choices have serious drawbacks:

- Larger servers/NAS appliances are more expensive per increment of serving capacity, and in some cases still remain bottlenecks on overall performance
- More servers create additional points of server and storage management
- Fragmentation of workloads across more servers and storage decreases capacity utilization
- Replicated data wastes storage capacity on copies of data and generates new data management challenges associated with keeping data copies consistent
- High availability becomes increasingly expensive, requiring duplication of large amounts of expensive hardware

PolyServe File Serving Utility running on industry-standard servers and SAN storage addresses all of these problems with a new approach to the problem – shared data clustering.

Enabling a Lower-Cost, Flexible Infrastructure for File Servers

The File Serving Utility and Intel Architecture-based servers running Windows enable a lower cost, modular approach to CIFS file serving. As opposed to larger SMP servers (4, 8, and 8+ CPU), a cluster of inexpensive servers (1, 2 or 4 CPU) can now function as a single, easily managed, incrementally scalable (“pay as you grow”), highly available CIFS server, solving the problems posed by conventional approaches.

File Serving Utility clusters deliver the following benefits:

- **Manageability.** The cluster can be treated as a single CIFS server with a single pool of storage and a single set of exported file systems. There is no need to create redundant copies of data or to divide clients manually among different servers.
- **Scalability.** Additional servers and storage can be added to the cluster on the fly to meet growing workloads. Each additional server represents incremental performance in terms of I/O throughput.
- **Availability.** In the event of failure, servers in the cluster automatically take over each others’ clients immediately, with no interruption of service. Dynamic on-line addition of servers and storage means that no scheduled downtime is required for capacity upgrades.

- **Utilization.** When servers in the cluster start to reach critical levels new servers can be added dynamically to take up the additional load. By sharing the data across all the servers, there is no redundant duplication of data. Both of these benefits provided by PolyServe result in improved ROI due to more effective utilization of server and storage resources and, hence, spending.
- **Cost.** Centrally managing clusters of low-cost (\$5,000) 2-CPU servers, each providing failover support for one another, creates a substantially lower-cost infrastructure with greater availability than use of costly (\$25,000 and up) 4- to 8-CPU servers or expensive NAS filers.

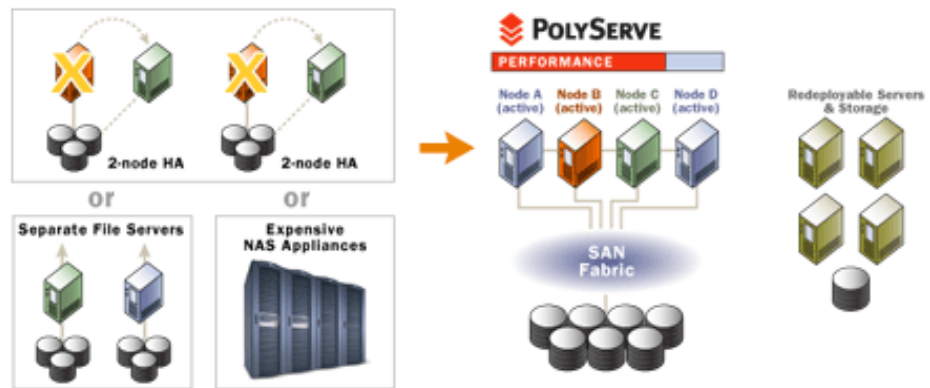


Figure 1. Consolidate & scale file servers with highly available, shared data clustering.

Applicability to Server Consolidation, NT4 Migration and NAS Replacement Projects

The ability to consolidate existing (often standalone) file servers and associated (over-allocated) storage allows existing enterprises to rethink how they are deploying and managing CIFS file serving deployments. The ability for up to 16 servers in a PolyServe cluster to be “active” and provide failover for one another reduces the cost of the hardware footprint and eliminates data silos that have been independently managed. More notably, building clusters out of 1-, 2- or 4-CPU servers eliminates the need for 8-CPU or larger servers and “passive” servers typically found in legacy “high-availability, active-passive” deployments. Three common scenarios that can benefit from a highly available, shared data clustering solution are as follows:

- **Server Consolidation.** Many Windows server environments are composed of many under-utilized file servers. File servers may also be deployed in 2-node active-passive pairs for higher availability. PolyServe allows IT to sweep multiple file servers up into a single ultra-scalable, centrally managed file serving cluster. PolyServe virtualizes these servers and makes them

appear as a single file server with the aggregate performance of many file servers. All servers in the PolyServe cluster can be “active” and can concurrently read and write to shared storage, thereby eliminating the need for “passive” under-utilized servers.

- **NT 4 Migration.** The File Serving Utility allows customers to migrate existing NT4 file servers to a highly available, scalable file serving backbone powered by Windows 2000 or Windows Server 2003. For organizations looking to decrease application and infrastructure downtime, and perhaps consolidate both servers and storage, PolyServe offers a simplified way to group and manage these previously over allocated, un-coordinated server and storage resources through a centralized console that manages highly available, shared file systems across 16 nodes.
- **NAS Replacement.** PolyServe File Serving Utility uses shared data clustering and integrated high-availability services to deliver a truly scalable, low-cost scale-out Windows and Intel alternative to proprietary, expensive NAS filers. The ability to harness industry-standard server building blocks and aggregate their compute capabilities allows customers to deliver a file services backbone that surpasses the performance of NAS, while delivering fault tolerance at a fraction of the cost of typical NAS failover clusters.

Matrix Server Technology – Key Components

The PolyServe File Serving Utility software is based on the Matrix Server cluster file system, and incorporates several key technologies:

- **A high-performance, symmetric cluster file system.** With Matrix Server, a cluster of industry-standard servers can share concurrent read/write access to file systems resident in a SAN, as shown in Figure 1. Matrix Server provides very high throughput for even the most demanding workloads, with direct disk access speeds and complete cache consistency among servers in the cluster.
- **High-availability infrastructure.** Matrix Server protects against server, network and storage failures. It provides automatic restart of servers and services and immediate fail-over of clients to surviving nodes.
- **Integrated cluster-wide management.** The management infrastructure enables a single console, Matrix Manager, to provide an overview of cluster status. Cluster operations can be reconfigured – either from the console or by using a command line tool – from any member of the cluster.




Figure 2. With Matrix Server, all servers have direct, simultaneous access to file systems and shared data in the SAN.

CIFS Manageability - Simplifying File Serving Deployments

Traditionally, multiple CIFS servers have required administrators to divide data and clients among them, perhaps replicate data among servers, and monitor and provision free space separately on each server. In a File Serving Utility environment, none of this is required:

All servers in the cluster can export a single set of file systems simultaneously; there is no need to divide data among separate servers, nor to replicate data among servers

Similarly, since any server can service any client's requests, there is no need to assign clients manually to individual servers. Clients can be balanced among servers automatically using DFS, DNS round robin or a hardware load-balancer. In either case, no manual intervention is required

Free space is shared among all servers, simplifying storage provisioning

Cost-Effective Scalability

A file server or NAS appliance represents a potential performance bottleneck for mission-critical, data-intensive applications. In CIFS, file system operations performed by clients generate network requests that must be received by the server, interpreted, executed by the server on the client's behalf — perhaps requiring disk I/O — and then replied to. Thus, a large number of CIFS clients, or even a small number of clients with heavy demands, can saturate the CIFS server's network interfaces, CPU or storage connections and cause performance to degrade to an unacceptable level.

In a File Serving Utility cluster providing CIFS services, each individual server contributes incremental resources to the cluster, including network bandwidth for client requests and replies, CPU capacity for file system processing, and I/O bandwidth to storage. No individual server's capacity — in any of these dimensions — imposes a bottleneck on system performance. As needs grow, a cluster can be scaled from as little as a pair of dual-processor servers to a set of 16 servers each with four or eight CPUs, multiple Gigabit Ethernet and Fibre Channel connections — all without interrupting client service and without requiring any migration of data.

Unlike some cluster file systems, Matrix Server is implemented with a completely symmetric architecture in which there is no bottleneck caused by a "master node" on which the other cluster members depend. In a File Serving Utility cluster, each cluster member is equally able to perform all file system operations on behalf of its clients, and every incremental server added to the cluster enhances the bandwidth available to clients.

Because a File Serving Utility cluster can be built out of the most cost-effective industry-standard servers (2-CPU typically), scalability comes at a very attractive price point when compared to traditional, high-end 4- and 8-CPU file servers or NAS appliances.

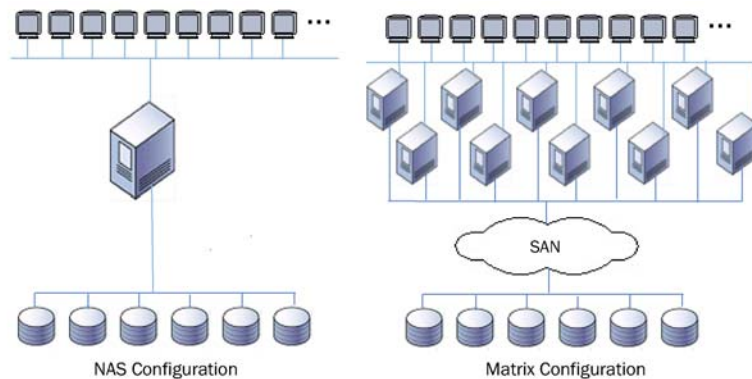
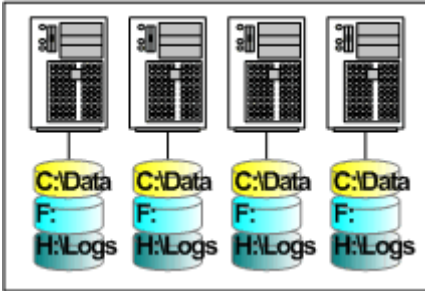
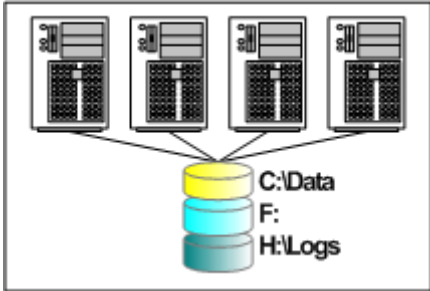


Figure 3. NAS Configuration vs. Matrix Configuration

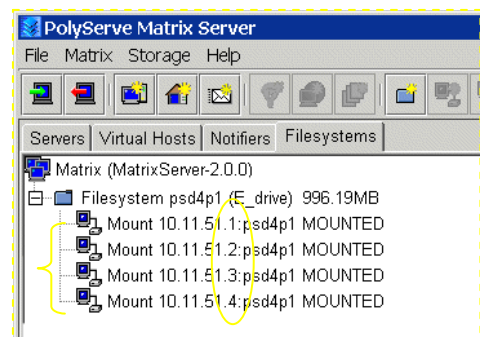
Enabling Shared Data

A cluster file system (CFS) for Windows drastically simplifies how easily one can keep drive letters consistent across the cluster. For example, when drive letter F: or G: is created in the CFS, all server nodes in the cluster immediately have access to this drive. With alternative HA software, creating and maintaining consistent drive letters is very difficult and often labor intensive.

Without CFS	With CFS
	
<p>Drive letters and shares must be maintained on each node independently (manage on each server in the cluster)</p>	<p>Drive letters and shares automatically consistent across cluster (manage coherently across the cluster)</p>

The screen shot below shows how multiple servers share the same file system (designated by the “E_drive”)

E_drive shared by 4 servers



Affordable High Availability

The File Serving Utility includes built-in high-availability functionality that permits the cluster to detect and respond to failures of network interfaces, storage interfaces, servers and operating systems within the cluster, including transparent failover of CIFS sessions.

- The File Serving Utility transparently fails over clients' CIFS sessions and therefore delivers the following benefits:
- Clients reconnect automatically to the CIFS share
- No change in Windows / CIFS clients required
- Works with the DFS/MFS solution or with external load balancers

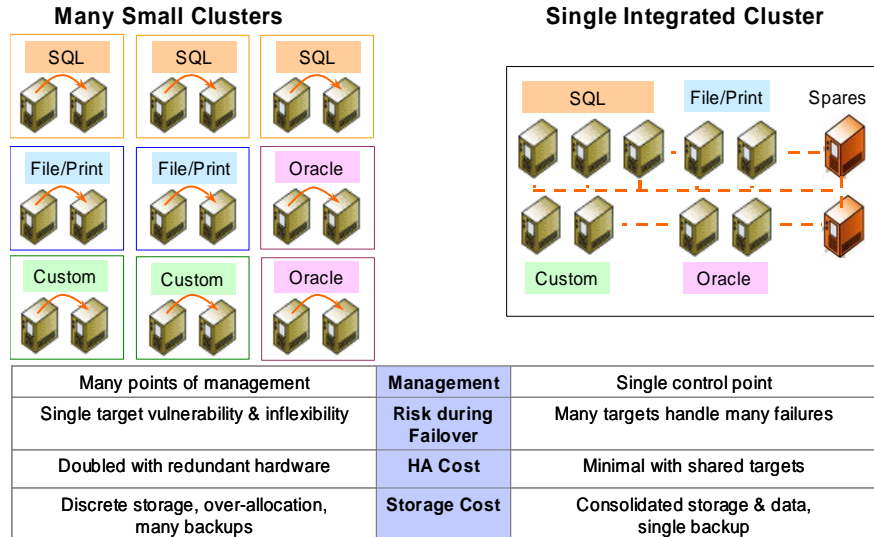
When configured to use virtual IP addresses assigned to virtual file servers, a single server can provide failover protection for the entire cluster. If a server fails, the File Serving Utility will automatically reassign the failed server's IP address to the standby server, allowing affected clients to continue making CIFS requests without any interruption of service. In fact, typically a client will have minimal or no indication that a server failure has occurred.

Because a single machine can provide failover protection for the entire cluster, it is not necessary to buy hardware in pairs to achieve high availability. For example, in Figure 3, one server could provide fail-over protection for the other nine, giving just a 10% overhead for availability, as opposed to 50% in a typical, paired configuration.

In addition to protecting clients from failures, the File Serving Utility also enhances availability by reducing the need for scheduled downtime. Online addition of servers and storage is supported, making it possible to grow from, say, a three-node cluster supporting just a few hundred gigabytes of storage to a 10-node cluster hosting many terabytes, with no interruption of service for system upgrades and with no need to reconfigure or restart any clients.

Simplify Management and Reduce the Cost of High Availability

The ability for many servers, applications and their data stores to be centrally managed in one consolidated cluster reduces the number of entities to manage and administer. It also offers greater availability (more failover targets) while reducing the necessary number of servers to create and support a highly available environment. The following illustration shows how many small clusters can be consolidated into a more usable, lower-cost infrastructure.



Reduced Total Cost of Ownership (TCO)

The File Serving Utility reduces Total Cost of Ownership (TCO) substantially over typical file-serving infrastructure deployments today. File Serving Utility-based configurations reduce the necessary hardware footprint across both servers and storage, reduces downtime through configurable, high-availability services and simplifies file-system and data-store management. The examples below demonstrate how both small and large file serving environments can reduce TCO approximately 50% and realize a payback in less than 1 year.

Scenario 1 (Small Environment)

This scenario demonstrates how the File Serving Utility can transform 10 dual-processor servers – either clustered in two-node active-passive pairs or as uncoordinated, under-utilized standalone servers – into a consolidated group of six active dual-processor servers, where each server provides failover services to every other server.

The operational benefits are as follows. With increased failover capability, downtime is reduced. With the ability to access shared file systems across all active servers in the cluster, a cluster of six nodes can adequately handle the load of the existing deployment. Duplicated, unnecessary and orphaned files can be eliminated when the storage is consolidated from the existing environment, translating to moderate storage-consolidation gains. Providing the ability to now manage all servers and storage as one computing entity simplifies manageability over the 10 under-utilized servers and un-coordinated storage.

Table 1 and Figure 4 summarize the operational simplification, as well as over 60% savings on TCO with a payback of less than a half-year.

Assumptions	Scenario 1 (Small)	Business Results	
Server Count	10 servers reduced to 6	Net Pres Value Savings (\$M)	\$1.93
Avg CPUs/Server	2	Annualized ROI (%)	51%
Server Avg Purchase Cost	\$ 6,000	Payback (Yrs)	0.35
Downtime	0.22 hr reduced to 0.07	3 yr TCO Reduction (%)	61%
Storage	4 TB reduced 30%		

Table 1. Savings in Small Environments

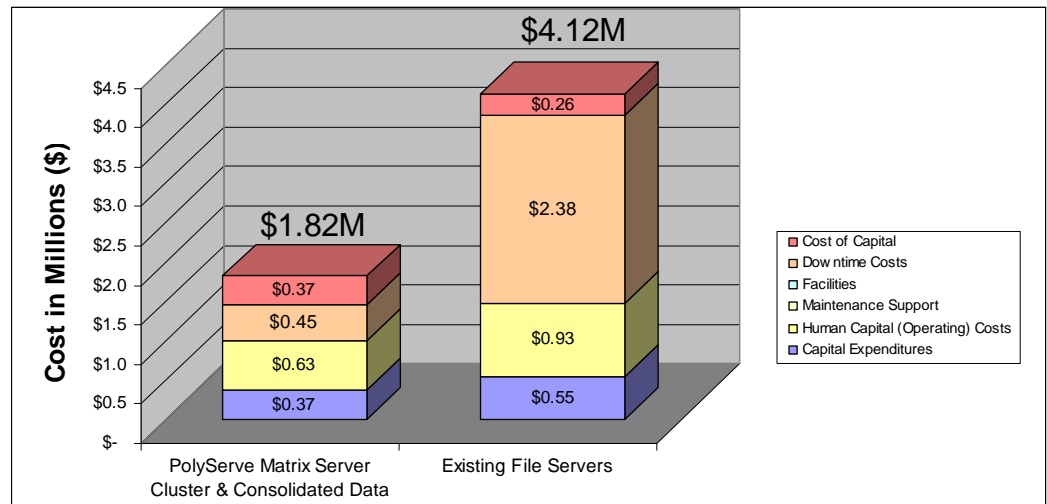


Figure 4. TCO Savings Comparison of PolyServe to Existing File Servers

Scenario 2 (Large Environment)

Similarly with large environments, servers and storage can be more effectively consolidated and therefore managed. Consolidating a large environment of 40 file servers with 10 TB of data provides a savings opportunity of nearly \$8 million as shown below.

Assumptions	Scenario 2 (Large)	Business Results	
Server Count	40 servers reduced to 24	Net Pres Value Savings (\$M)	\$7.71
Avg CPUs/Server	4	Annualized ROI (%)	52%
Server Avg Purchase Cost	\$ 31,000	Payback (Yrs)	0.54
Downtime	0.22 hr reduced to 0.07	3 yr TCO Reduction (%)	62%
Storage	10 TB reduced 40%		

Table 2. Savings in Large Environments

File Serving Utility Features and Benefits for CIFS File Serving

The following table summarizes the capabilities of the File Serving Utility.

Feature	Benefit
Symmetrical Cluster File System with Distributed Lock and Metadata Management	Scalability. Each additional server in a cluster adds more aggregate I/O throughput. Due to PolyServe's unique architecture, the File Serving Utility can deliver near linear scalability of aggregate I/O throughput (as measured by CIFS operations per second) as additional servers are added to the cluster.
Cache Coherency Across All Nodes	Scalability and Manageability. Cache coherency means that clients can connect to any of the servers in the cluster and receive exactly the same responses to CIFS requests. This guarantees data integrity clusterwide.
Automated Failover for CIFS Client	Availability. Higher system availability and minimized downtime. Includes integrated clusterware that monitors the health of the network, server hardware and CIFS sessions. Initiates failover to migrate a CIFS client from a failed server to a healthy server. In the background, the CIFS clients reconnect to the CIFS share(s)
Single Clusterwide CIFS Mount Point (Matrix File Shares)	Scalability, Load-Balancing and Failover. Ensures even distribution of CIFS client load across the cluster as all servers have concurrent access to shared data.
Storage Area Network with Consolidated Data	Data Capacity Utilization and Storage Management Efficiency. Centralization of data makes storage management more effective and increases storage utilization rates. Significant reduction in storage footprint is achieved as storage and data is centralized and consolidated. Number of backups are reduced as data is consolidated
Ability to Use Industry-Standard 2 CPU and Blade Servers Rather than Large UNIX Servers	<p>Lower Total Cost of Server Capacity. Ownership (TCO) and compelling price per performance (measured by \$/CIFS operations). The ability to use industry-standard servers and storage delivers lower hardware, software and maintenance costs while decreasing system downtime.</p> <p>Lower Cost of Providing Availability. Because a cluster is composed of many smaller servers rather than one large server, and because the cost of buying a smaller server is significantly less expensive per increment of processing power, the cost of providing failover capacity is much lower. Additionally, since the File Serving Utility supports many:1 failover capabilities, there is no need to maintain a spare server for every server that requires high-availability services.</p>

Feature	Benefit
Centralized Clusterwide Administration with GUI and Command Line Interface	Lower Systems Management Cost. Reduces the complexity of administering a cluster of CIFS servers. Cluster and file system configurations can be changed from a central point.
Online Addition of Servers and Storage	Lower Ongoing Cost of Capacity/“Pay as You Grow.” No need to over-provision to avoid “fork-lift” upgrades. Grow cluster and storage as business dictates without having to provision hardware forecasted peak load. Ride the curves that cause each increment of compute capacity and each megabyte of storage to get less expensive every six months.

Table 3. Features and Benefits of PolyServe File Serving Utility

Summary

HP’s PolyServe File Serving Utility™ for Windows provides a unique solution that addresses the scalability, availability and manageability of CIFS file serving. The PolyServe solution addresses the scalability and performance bottleneck posed by the most demanding, data-intensive applications. In addition, PolyServe File Serving Utility can be utilized to consolidate numerous CIFS servers onto a single centrally managed cluster that provides inherent high availability and a clear, inexpensive, industry-standard growth path.

Join us for an online demo to learn more about [Windows File Serving Utility](#)

PolyServe, Inc.
 20400 NW Amberwood Drive, Suite 150
 Beaverton, OR 97006
 Toll Free: 877-765-7378
 Tel: 503-617-7574
 Fax: 503-617-7592

© Copyright 2007 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

WinFileServing_UTILITY_Whitepaper_042407.doc