Using Samba with a Commercial Clustered File System

Agenda

- Isilon Clustered Storage Overview
- Developing OneFS
- Samba and OneFS
- Clustering Samba
- Future Requirements

Isilon OneFS Cluster

- Single file system (1.6 PB)
- Fully symmetric peers
  - 3 to 96+ nodes
- Fast intra-cluster network
  - InfiniBand
- Multi-protocol access
  - CIFS
  - NFS
  - FTP
  - HTTP/WebDAV
  - Local userspace

What Makes Isilon OneFS Special?

- Single pool of storage
- Granular data protection
  - » Down to the file level
  - » 8x mirroring or +4 recovery
- Easy to manage and grow
  - » Add additional nodes in 60 seconds
  - » Automated data balancing
- Extreme performance for concurrent access
- POSIX and Windows semantics supported

Developer Perspective

Full Control

- File system
- Kernel
- Node's userspace applications

No Control

- Customer environment
  - » Domain topology
  - » Domain policy
- CIFS client
- NFS client

Samba in OneFS

Used over 4 years – **2.2** through **3.0.24**

Isilon Kernel Supported Samba Requirements

- Native Windows ACL storage & enforcement
- Native createfile() syscall implementation
- Alternate Data Streams
- Snapshots with Volume Shadow Copy Service (VSS)
- Zero-copy writes and other performance improvements

Samba in OneFS Continued…

- Cluster coherent share-mode locking
- Cluster coherent oplocks
- Cluster coherent byte-range locks (in development)
- Per-share case-sensitivity
- Unicode normalization insensitive
- Site locator support
- Change notify
- And more…

IFS Clustering Protocols

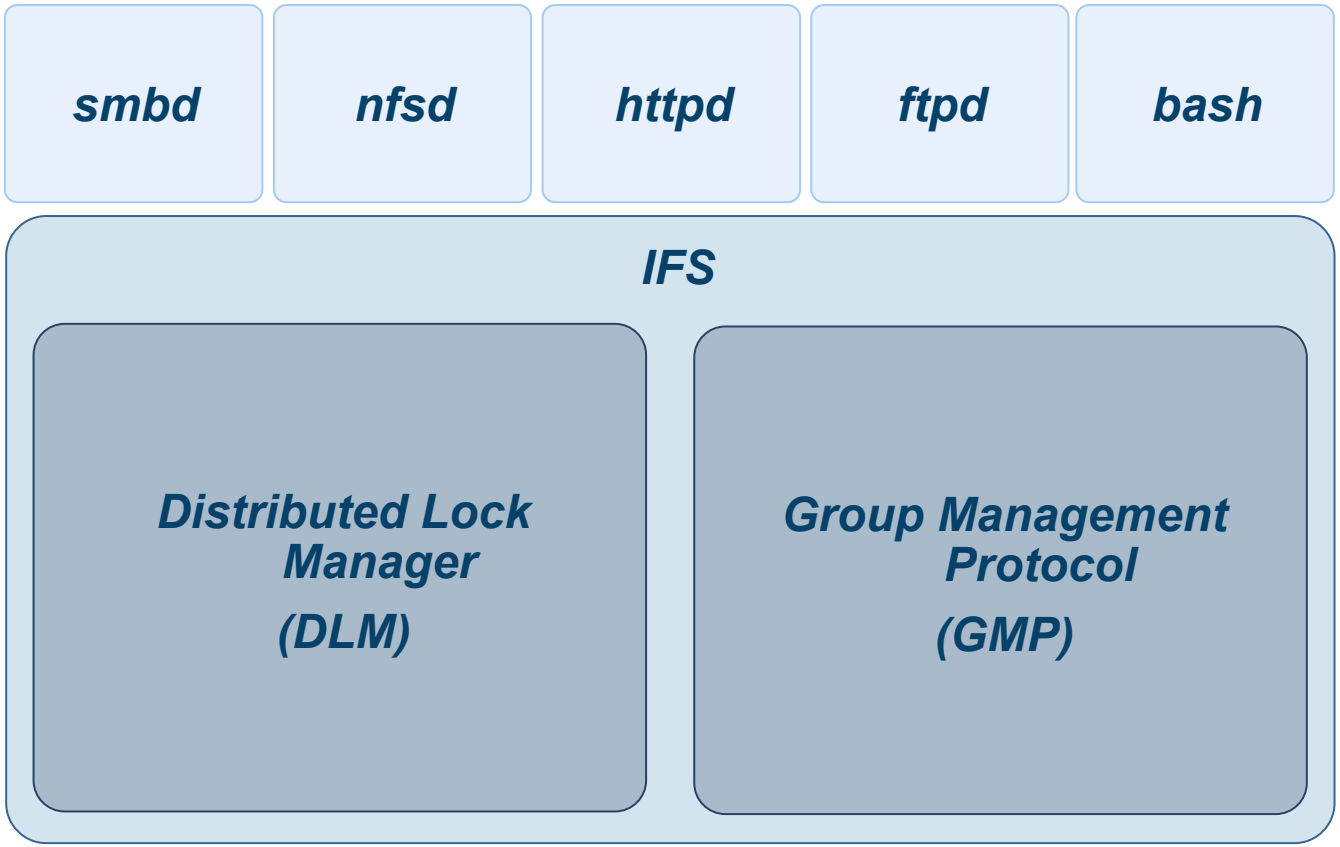

OneFS Distributed Locking [under Samba]

Distributed Lock Manager

- Expressive
  - Arbitrary contention tables
- Two-tier model
  - Each node responsible for subset of cluster locks
  - Local lock caching
- Properly Models
  - Share mode locks
  - Oplocks
- Multi-protocol support without customized applications

OneFS Group Membership [under Samba]

Group Management Protocol

- Rich predicates
  - More than just node “available” / “unavailable”
  - Operations based off health of node
    - Individual I/O directions: Readable / Writeable
  - Predicates dependent on full cluster state
    - Degraded read as last resort
- Group changes integrate with other modules
  - I/O automatically routed to available nodes
  - DLM can renegotiate

CTDB on Clustered File Systems

Leverage Existing Technology

- CTDB is a promising solution for POSIX clustered file systems
- Some file systems extend capabilities
  - » Native distributed locking mechanism
  - » Native cluster membership tracking
  - » Native ACLs and other feature obviate need for specific TDBs
- Native solutions benefit from
  - » Performance tuning
  - » Build on existing testing and deployment
- Samba should benefit from these capabilities when available

Samba 3 Wish List

Moving forward…

- Abstract interface for TDB based solutions
  - » Kernel based solutions and CTDB coexisting
  - » Existing examples
    - iNotify for Linux
    - NTFS VFS layer in Samba 4
  - » Beneficial to other vendors

Samba 3 Wish List

Moving forward…

- Increase communication
  - » Share our features and bug fixes
  - » Minimize parallel development
  - » Vendor specific branches?
  - » Ease merge burden

What We Give Back to the Community

Enterprise Testing and Deployment

- Dedicated Windows enterprise QA team
  - » Hundreds of bugs found
  - » Integration in complex domain topologies
- Real world deployment in large production environments
  - » 50,000+ Users
  - » 300+ Domain Controllers
- Feature development
  - » Windows 2008 Server Authentication

Isilon Customers and Leadership

Select Customers

myspace.com a place for friends
NBC
sony pictures imageworks
Pratt & Whitney A United Technologies Company
Cedars-Sinai
LexisNexis
abc
Sports Illustrated
KTI Kelman Technologies
ClearChannel Worldwide
Paramount
Brigham
E! Entertainment Television
USGS science for a changing world
GlobeXplorer
American Title Inc.
Kodak EasyShare Gallery
Movielink
Technicolor Video • CD • DVD Services

Recognition

2005 Products of the Year — SearchStorage.com
Computerworld Storage Networking World — Best Practices in Storage Awards Program
Best of Show Award — Networld + Interop Tokyo
Bio-IT 50 — Bio-IT World
eWeek
Byte and Switch — Top 10 Private Companies | 2005
Red Herring 100 N. America Winner

> *"Isilon's clustered storage products will play a critical role in meeting the capacity and throughput requirements of leading applications."*
>
> *Richard Villars, Vice President of Storage Systems Research, IDC*

Future

2008 and Beyond

- CIFS/Samba improvements
  - » Many interesting challenges ahead
  - » Clustered performance
    - Closer integration with world-leading clustered storage solution
- Customer Focus
  - » Our customers continue to push the envelope of Samba development
- Growth
  - » Continue to get more involved with Samba community
  - » We need more protocols developers! (Come join us!)

Questions?

Contact

- Steven Danneman
  - steven.danneman@isilon.com
- Tim Prouty
  - tim.prouty@isilon.com