



## File Storage Made Whole — NeoPath's File Director FD-220

Analyst: Michael Fisch

### Management Summary

Imagine a factory floor with various machines forming an assembly line. Each machine whirs, spins, or grinds and performs some vital task in the production of, for instance, bicycles or wristwatches. Workers move assemblies from one machine to the next, forming and adding components, until finished products emerge. They finally take products to a warehouse for storage, packaging, and eventual shipment.

Now, imagine installing one more special machine called NFM, which has a profound impact on the factory. First, NFM makes the other machines robotic and easy to operate, so fewer workers are needed and fewer mistakes made. It even moves assemblies seamlessly between stations. Second, NFM makes the other machines more versatile, so they can create products with different options, like gold or leather wristbands. Finally, it makes the machines faster and more efficient, so the factory produces 25% more bicycles or wristwatches each day. The NFM machine takes the factory as it is and dramatically improves it.

Is there such a machine in the world of manufacturing? Unfortunately, there is not, though perhaps a series of new machines, computers, and processes might accomplish all of that – at considerable expense. However, such a machine does exist in the world of file storage. **The vendor NeoPath launched a product called *File Director* in December 2004 that is analogous to this special machine. It fits into an existing file storage environment (i.e., file servers and NAS platforms) and makes everything more functional and economical.** Since business operations depend on data access, and since file storage is where much or most data resides, File Director can provide substantial benefits:

- **Lower hardware costs** through higher utilization and intelligent data placement (ILM),
- **Lower management costs** through a global namespace and central, policy-based control,
- **Higher performance** through load balancing, and
- **Higher availability** through non-disruptive adds, moves, and changes and optimized backup and data protection.

This product falls into an important, nascent category called *network file management (NFM)* or *file virtualization*. NeoPath's solution is solid and has unique features, and it just announced an enhanced version, called the *File Director FD-220*. Read on for details.

### IN THIS ISSUE

➤ File Storage Challenges .....	2
➤ NeoPath File Director .....	2
➤ Conclusion .....	3

## File Storage Challenges

To set the stage, let's look at the formidable file storage challenges that enterprises face.

### *Data Growth*

Continuous data growth (including unstructured data or files) is the backdrop on the stage of enterprise storage. **All enterprises will need more capacity next year than they do today – the only question is how much more.** Therefore, all must plan to add capacity periodically.

### *Wasted Storage Capacity*

**As enterprises try to keep up with this data growth treadmill, most overprovision their capacity.** Traditional NAS platforms and file servers scale in discrete units, so enterprises with requirements of any scale have multiple platforms. Each NAS platform likely contains multiple, discrete file systems (i.e., logical data containers) due to size limitations. This fragmented environment has poor utilization because capacity is not easily transferred between the storage “islands”. To make matters worse, enterprises tend to purchase capacity in large chunks due to the effort and disruption associated with each addition. This excess capacity ties up capital, consumes electricity, space, and administrative resources, and does not help further business objectives.

### *Labor-intensive Management*

Storage hardware is not the only wasted resource in a fragmented environment. **IT administrators also spend too much time and effort managing the storage islands: tracking resources, making changes, manual load balancing, upgrades, migrating data, and consolidation.** For each change or addition, administrators must update the share mappings or mount points on all affected clients. This is like completely rewiring a house every time a lamp is moved or new light bulb installed. There must be a better way. These excess IT staff hours have a major budget impact, over time.

### *File Access Disruption*

**Another problem is downtime.** During these changes and additions, clients may be denied access to data for periods of time that affect their productivity or, at minimum, are an inconvenience.

### *Over-Reliance on Top-tier Storage*

**Yes, most enterprises have long relied on deploying their top tier of storage broadly, probably too much so, because it was easier to deploy and manage.** There is an opportunity to lower file

storage costs yet more by implementing information lifecycle management (ILM)<sup>1</sup>. Enterprises can classify files according to their service level requirements and migrate many (if not most) files from top-tier storage to low-cost media like SATA drives and tape. By streamlining the top tier, its performance also improves and activities like data replication, backup, and restore are faster.

### *Data Retention and Regulatory Compliance*

**A big, ticking time bomb for many enterprises are new regulatory requirements or best practices that will extend how long data must be retained and made accessible.** In general, enterprises need long-term data retention for corporate governance and operational requirements. Low-cost media are preferable for this infrequently-accessed archive data.

### *Business Continuity*

Enterprises need to ensure data access and business continuity through data protection and recovery solutions. **The standard for acceptable recovery time objectives (RTO) and recovery point objectives (RPO) are rising, increasing the challenge and the required investments in time, hardware, and software.**

Network file management (NFM), also called *file virtualization*, tackles these file storage challenges through powerful simplification, data migration, management capabilities. NeoPath is an early and ambitious player in this movement.

### **NeoPath File Director**

Like the fabled machine in the factory, NeoPath's *File Director* installs directly into an existing file storage environment and makes improvements across the board. **It turns a multiplicity of storage islands – heterogeneous file servers and NAS platforms – into a single virtual pool that is accessible by all clients.** It moves data between file servers without disrupting user access, even during production hours. It tracks file attributes in real time and can automatically migrate files based on business policy, as well as flag files for the administrator to delete. IT administrators can manage all of it centrally from one console, utilizing a collection of salient capabilities.

### *Global (or Unified) Namespace*

The foundational feature of File Director is a

---

<sup>1</sup> See **The Clipper Group Explorer** dated May 11, 2004, entitled *Top Ten Things You Should Know About Information Lifecycle Management*, available at <http://www.clipper.com/research/TCG2004041R.pdf>.

*global namespace*, which is like a phone book for file systems. With a phone book in hand, a person can easily contact any individual or business in town by name. He or she does not need to know the physical address or keep track of where everyone lives. As people move and new businesses spring up, the phone company tracks the changes and revises the directory. In a similar fashion, a global namespace presents a single, unchanging, virtual namespace to the enterprise. It eliminates the direct dependencies between clients and file servers. Administrators are free to move files and directories, add capacity, and perform other changes without interrupting client access or revising their share mappings or mount points. This saves a lot of work, makes it easier to manage access security, and improves utilization of the storage assets.

File Director supports the standard *NFS* and *CIFS* file-handling protocols for *Unix*, *Linux*, and *Windows* clients. Any file server running *NFS* or *CIFS* may participate in the global namespace. Clients may access the same file through either protocol, which is unique. Another unique feature is allowing clients to access files simultaneously through and around File Director. This is useful for incremental deployments and ensuring data is accessible if something happens to the File Director. The product stores this namespace information on the file servers themselves, so the state of the environment is never lost.

### ***Non-disruptive File Migration***

A global namespace makes file migration transparent, but non-disruptive migration is a step beyond. It means clients can read and write files even while File Director is moving them – a tricky technical challenge. This feature improves data availability and lets administrators perform management tasks during business hours – *no more late nights or weekends!* Data migration is a routine part of tasks like equipment upgrades and load balancing for performance tuning and capacity utilization.

### ***Real-Time File Characterization***

File Director keeps a real-time database of file attributes for reporting, policy execution, and running *what-if* scenarios. Attributes include file type, size, and access frequency, all of which are useful for making decisions about placement, archiving, and deletion. File Director initially scans the directories and loads the database. From that point, it tracks activity in real time, so the data is always fresh and ready to use. This stands in contrast to other tools that must periodically scan the directory structure, which is time- and resource-intensive and unable to monitor access frequency. File Director's

combination of real-time file characterization and non-disruptive migration also is unique.

### ***Policy-based Automation and Centralized Management***

The icing on the File Director, if you will, is policy-based automation and centralized management. Both of these amplify administrator efforts and reduce operating costs. Centralized management for a global namespace is a powerful capability. File Director can also execute programmed policies based on file attributes. For instance, it could migrate files of a certain age from a high-performance storage tier to a low-cost tier. This lowers overall storage costs and trims the size of primary storage, so data protection processes like backup and replication run faster. File characterization, migration, and automation are key capabilities for implementing ILM strategies.

### ***Hardware Platform***

While the intellectual property in File Director is the software, the hardware platform is an industry-standard Intel server running Linux with redundant power, cooling, and drives (RAID 5), but the storage is not for enterprise data, but its own use. It is in-band, network-resident, and connects to both clients and storage via IP networks (read and writes travel through it). File Director is available in two-node clusters for high availability.

NeoPath states that the recently enhanced *FD-220* delivers ten times the performance of the previous version, primarily due to software refinements, for a total throughput of 3.5 Gbps or 10 Gbps with jumbo frames. It also states that a pair can support up to 20 back-end file servers and 100 TB of data, depending upon usage. List price of File Director is \$39,995 for a single appliance and \$69,995 for a failover pair.

### **Conclusion**

One day, NFM solutions will be standard in file storage environments of any significant size. The current challenges are so many, and the benefits of NFM so clear and broad, that this makes eminent sense today for many enterprises.

**So, if you need to address your enterprise's file storage challenges sooner rather than later, consider NeoPath's next-generation File Director. It is a solid product with unique features.**



### ***About The Clipper Group, Inc.***

***The Clipper Group, Inc.***, is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at [www.clipper.com](http://www.clipper.com).***

### ***About the Author***

***Michael Fisch*** is Director of Storage and Networking for The Clipper Group. He brings over nine years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at [mike.fisch@clipper.com](mailto:mike.fisch@clipper.com) or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

### ***Regarding Trademarks and Service Marks***

***The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, and "clipper.com"*** are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

### ***Disclosure***

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

### ***Regarding the Information in this Issue***

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.