

Virtually possible

Leveraging network intelligence as a pathway to efficient storage

RICK VILLARS is vice president of storage systems at International Data Corp. and a senior member of IDC's storage research team. He recently spoke with writer Paul Gillin.

Q How does IDC define virtualization?

A Virtualization can be defined as many different things, but for most of us there are three basic types. Server virtualization takes a physical server and creates multiple logical servers on top of it. Storage virtualization—block or file—combines several storage devices into a single, logical asset. File virtualization puts a global namespace above independent namespaces; directories and folders look the same to users no matter where the files are located. This makes file server or NAS migration easier and gives you more flexibility when managing backups and long-term archives.

Q How does file-based storage management differ from block-level storage management?

A There's no context for block-level storage, just volumes. Therefore, IT managers know very little about what's being stored; they just know what applications use block-

level storage. That's why when you back up block storage, for example, you need to back up the whole volume rather than parts of it. File-based storage contains metadata, such as the name, the file type, who created it, and when it was created. With the need to archive documents for regulatory reasons, file-based storage is growing quickly. File virtualization enables you to achieve a more detailed level of management. For example, you can consolidate file servers and create different backup schedules so only data that was created in the last 30 days is backed up every few hours, while older files can be moved to cheaper disk drives and backed up monthly.

Q How does file virtualization relate to application-aware networks?

A You can leverage file virtualization to improve the performance of IT operations. You can leverage intelligence in the network to make the file environment cheaper to deploy and more self-sufficient. In many ways, this is similar to load balancing, in that you're managing data more intelligently.


Q What new capabilities does Acopia bring to F5 customers?

A F5 is an important provider

of server load balancing. It has a lot of experience building large systems reliably and quickly moving information between the data center and all the people and systems that need access to that information. Today, enterprises are looking for a partner that can help them adopt server virtualization, new multicore server platforms, and large, rich content applications while implementing more robust data protection, retention, and security practices. Acopia brings products and business relationships in the storage portion of the data center, and can help F5 expand its presence in enterprises as they focus on these growing challenges.

Q How does file virtualization impact SAN and NAS deployment?

A While it doesn't really impact SANs directly, it is a boost to NAS. People increasingly want to consolidate existing file servers and NAS systems, and that can take years if they have a lot of



IDC's Rick Villars expects to see substantial growth in the file virtualization market.

direct-attached storage. This technology makes that move faster and easier. Also, I can have NAS for my high-performance data and a nearline second tier for my archive, while gaining the ability to balance the load across them.

Q What's IDC's forecast for the file virtualization market over the next few years?

A We see the block-level virtualization market growing from \$300 million to more than \$1 billion in the next four years. We haven't done a formal forecast for file virtualization, but we expect to see comparable demand going forward. *