

Scaling the Storage and Data "Wall" with Scale Out File Services

Speakers: **Tim Willeford, Media Relations Lead - GTS Service Product Lines**
 Scott Fadden, GPFS Technical Marketing

Tim:

Welcome to this IBM podcast series focused on how to optimize your technology infrastructure. I am Tim Willeford from IBM. Through this series we will cover topics of how lower IT complexity and improve operating efficiency along with tips and advice on using technology and services to help you and your company succeed.

Today I am joined by Scott Fadden, one of our resident experts on storage and data services to talk about ways to simplify and save on storage and data infrastructure and an exciting approach clients are taking to scale-out their storage and data resources.

So Scott, we hear a lot about the exponential growth of data and information in digital content. How are clients wrestling with this?

Scott:

Well thanks Tim. We've seen this across all industries, right? We've heard it many times -- wow, a 270MB disk I don't know what I'm going to do with all that space, right? But experience has proven that whatever space is available will be filled. This means that we need to find a new approach, enterprise file serving. We need to get away from the past where you buy a box and then in a few years, you buy another box and then move the data. We have to find a way to grow the storage infrastructure as data volumes increase and as technologies change, without downtime to move the data, without "forklift hardware" upgrades which can be really disruptive and expensive.

Tim:

So storage costs are decreasing overall, right? We are seeing it in anything from the price of hard drives for our own PCs to iPods. Is that also reflected in the data center?

Scott:

Absolutely Tim. Sure, hardware cost are decreasing and that's a technology moves. But, the real cost to keep in mind is the cost of maintaining and supporting a whole file service infrastructure and these costs, in general, are not going down with most solutions. With a lot of the solutions, as capacity is added, you are adding another box, another management node, another place where you have to manage, you're increasing your management and you're increasing your power consumption with each edition of space, for example.

The solution that we created at IBM to address this issue is called the Scale Out File Services. Scale Out File Services is built on industry-standard components that keeps costs down and more

importantly improves flexibility and allows us the ability to meet rising demands and quickly adapt to new technologies as they come out. The key to Scale Out File Services or SOFS, as we call it, is it is a mix of hardware and software, it's a complete solution that provides scalable file management services that can be tailored directly to meet your needs and grow as data volumes increase.

Tim:

So Scott, Scale Out File Services sounds like a storage grid or cloud computing. Is that a fair representation?

Scott:

That's exactly right Tim. I mean, storage administrators have been waiting for a product like this. The idea that you have a grid or cloud of storage, that's a dynamic mixture of a number of servers and a number of storage arrays with a backend that can be flexible. So, you can add servers and as you need more processing power, you can add storage if you need more storage. This allows you to add new servers as new machines come out, as time goes on and you retire old ones, and all of the data is online. This gives you flexibility to get exactly what you need out of storage infrastructure. You can have two servers and some storage; you can have 50 servers and some storage. You can have a lot of storage in and a few servers; however you want to design it.

When you go into larger solutions like this, you need to have scalable technologies to support it. With Scale Out File Services it has information lifecycle management tools that can scale to billions of files. Now you can do online migration of data from one disk sort pool to a different one. Let's say an old one or a newer one or a fast one or slow one. You can even move that data off the tape and dynamically retrieve this as needed; all this with a single management infrastructure. It's a one management infrastructure for the whole solution and you get all of these technologies, even multiple sites. So you can have a site in the US. You can have a site in Europe and you can have a site in Asia. All these sites are managed through one interface and one structure.

Within a site our cross sites, Scale Out File Services provides replication. It provides synchronous replication or asynchronous replication. This helps you provide quality of services let's say to each location or each area. This allows you to create a flexible enterprise-wide, scalable infrastructure that's easy to manage.

Tim:

What types of companies, what sizes and what industry might be interested in this solution and what type of benefits are they realizing?

Scott:

Companies that require file services are all types of companies from all different industries. From pharmaceutical companies to companies that use it for computer-aided engineering, for

leading edge Web 2.0 digital media companies. At first we thought this was an enterprise product for large companies, but we're finding small to medium businesses are finding this very useful. No more is it a small business has a little bit of data. There is a company out there with a hundred employees that has billions of files of data. This kind of infrastructure can support that.

For example, we look at a specific major electronics manufacturer. This customer needed a highly scalable, cost-effective environment providing for online media, they are a Web service to their customers. The key to this service is they were looking for a solution to share their videos, audios, on this Website, but the key is they wanted unlimited media storage for their clients. Unlimited uploads and unlimited media storage. They needed a solution that could grow -- unbounded -- something that could grow to support this, because they don't know what the total capacity is going to be.

The initial solution consisted of 50 servers and about 80 terabytes of storage. File Service Storage is expected to grow to about 1 petabyte in the near future. So, this is a very dynamic growing environment and they saw SOFS as the key to solving that problem. It gave them consistent cost advantage, they were able to grow it out knowing their costs and being able to manage it in one context. They get all the services from IBM. We provide hardware, storage, services to implement the solution and it gave them quality of service. We'll be implementing this over multiple locations. They're rolling it out within their European hub and they will be going to Asia-Pacific and the U.S.

Tim:

That's interesting. Clients all over the world are already picking up on this solution. Where do I go for more information about Scale Out File Services?

Scott:

The best place to go to get information is our Website. At www.ibm.com/services/storage and look for storage optimization and integration services.

Tim:

OK, Scott Fadden, thanks for your time and your insights on this really exciting topic.