

White Paper

The Inevitable Move Toward a Modern Global File System

A Detailed Look at the Future of File Storage
in an Era of Distributed Digital Business

By Steve Duplessie, ESG Founder and Senior Analyst
and Scott Sinclair, ESG Senior Analyst

December 2019

This ESG White Paper was commissioned by Nasuni
and is distributed under license from ESG.



Contents

Abstract.....	3
What a File System Is, and Why It Needs to Change	3
The Failure of Traditional File Systems to Keep Pace	4
It Gets Worse	5
Welcome to a World of Cloud-native Global File Systems	6
Why It Is Inevitable?.....	7
Modern Applications Require Modern File Infrastructure	7
The Bigger Truth.....	7

Abstract

In the age of on-demand, instant-on, cloud-scale applications, using file architectures developed for single, local, stove-piped applications simply cannot work. It’s time to modernize, globalize, and simplify our unstructured data architectures. Just as SaaS and cloud-native architectures have almost completely replaced legacy software applications, so will cloud-based global file systems replace legacy file architectures.

What a File System Is, and Why It Needs to Change

In simple terms, a file system is an index that keeps track of all the data we store. In the good old days, we had a file system to keep track of what we put on our disks. Then we bought more disks. We had to spread out our file system to cover them all. When we didn’t have much data, file systems were fairly easy to maintain. Those were the days. Those days are gone.

The main issues with classic file systems tended to relate to performance (the bigger it gets, the slower it goes) and management (the ability to organize the file system so we can find things). Those issues were initially limited to *individual* machines, and originally, to individual users.


As our world moved from monolithic mainframe computing to distributed computing, we moved the processing and applications out to more users, and we centralized the data back at the data center’s server. When we started distributing servers all over the place and created a management nightmare, it was clear we needed a better way. Enter the likes of the NetApp filer in the mid-1990s, which consolidated a million small file systems onto one mega-file system, making it significantly easier to organize, protect, and manage that data.

But, like everything in IT, data growth eventually breaks everything we know and love.

File infrastructure is no different. It’s not a surprise then that two-thirds (66%) of IT decision makers surveyed by ESG report their IT environments are more complex than they were just two years ago, with 30% of surveyed IT decision makers seeing increases in complexity and citing higher data volumes as the cause (see Figure 1).¹

Figure 1. Top Six Drivers of IT Complexity

What do you believe are the biggest reasons your organization’s IT environment has become more complex? (Percent of respondents, N=400, three responses accepted)



Source: Enterprise Strategy Group

¹ Source: ESG Master Survey Results, [2019 Technology Spending Intentions Survey](#), February 2019.

Data growth caused even the biggest monolithic file servers to choke, back in the day. It forced the IT industry to come up with new architectures to try to keep up, namely:

- **Scale-up.** This architecture let monolithic file servers add capacity and grow the file system dynamically, so users didn't have to migrate from one file system to a bigger one. It made life much easier for admins ... if it were only about physical capacity. Alas, that's not how the real world operates. Performance became the next bottleneck.
- **Scale-out.** This architecture allowed performance to be scaled independent of capacity by adding compute/network nodes regardless of capacity. It forced us to create the "clustered file system," that is, one file manager executing between geographically confined compute nodes.

Every file system advancement developed because data growth was affecting file infrastructure performance or management has been a "one step forward, two steps back" scenario.

Those advancements helped centralize computing/data implementations in the workplace. But they also created significant issues with existing data operations.

And that's why file systems need to change. Every file system advancement developed because data growth was affecting file infrastructure performance or management has been a "one step forward, two steps back" scenario.

The Failure of Traditional File Systems to Keep Pace

The haphazard evolution of file systems has led us to an untenable position. No one option does everything we need. So, we leverage multiple products. The result is an unsustainable hodgepodge of disparate systems.

For example, if you already had your backup and recovery systems in place to handle one architecture, chances are, you had to go back to the drawing board to implement protection for your newer file architecture. The same holds true for your data operations functions: disaster recovery, business continuity, test/dev, analytics, etc.

We have too many architectures. Consider that 83% of IT professionals we surveyed who are responsible for data protection technology reported that they expected to use more than one data protection solution in the following two years.²

The rise of the data-driven economy is poised to make all this worse for IT admins dealing with data growth. In an ESG research study, the top three most often cited workloads driving storage growth were digital media (named by 31% of respondents), collaboration (30%), and business intelligence (29%).³ All those workloads are fueled by companies' desire to get more value from data, all those workloads leverage file data, and all those workloads are showing no signs of slowing down.

The fact is that the modern digital business that leverages file data is encountering the same hurdles presented by old file storage technologies. Notably, not only do 26% of surveyed IT decision makers who consider IT more complex report that digital transformation is a reason for the complexity,⁴ but also technology issues, along with collaboration challenges, are the top reported inhibitors of digital transformation (see Figure 2).

² Source: ESG Master Survey Results, [2018 Data Protection Landscape Survey](#), November 2018.

³ Source: ESG Master Survey Results, [2017 General Storage Trends](#), November 2017.

⁴ Source: ESG Master Survey Results, [2019 Technology Spending Intentions Survey](#), February 2019.

Figure 2. Biggest Challenges with Digital Transformation Initiatives



Source: Enterprise Strategy Group

It Gets Worse

Users are no longer just sitting in the cube farm upstairs. They are everywhere, around the globe. Even if they have a formal office somewhere out there, historically, IT would have to give them their own set of infrastructure, including file systems to manage and support. No one wanted to do that, but there was no choice.

Using a traditional file infrastructure to support an on-demand, instant-on, collaborative world is impractical. The world needs a new architecture to handle infinite capacity, infinite performance, and infinite access—from anywhere, at any time—all without breaking or even requiring all of our traditional data operations functions. In traditional solutions, distance means needing to have multiple systems at multiple locations, dramatically increasing capital expense, operating inefficiencies, and expensive synchronization, which means slower performance and costly networking. There is simply nothing good about having to replicate systems, but that has been the only way to solve the problem until this point.

In a data-driven economy, just storing data isn’t good enough. Businesses must use their data to get smarter and more efficient, unlock new markets, drive revenue, and keep shareholders satisfied. ESG research shows that in regard to these requirements, the typical IT organization is not doing a good enough job.


One quarter (25%) of line-of-business (LoB) executives view IT as a business inhibitor, and that is four times the percentage who view IT as a differentiator. When the LoB executives were asked what technology-related challenges impede their jobs (see Figure 3), the most common response was difficulty accessing the data they need (30%). Among those executives who viewed IT as a business inhibitor, that number increased to 43%.⁵

So, even with all the technology being thrown around inside IT organizations, users are not getting the data they need to perform their jobs.

⁵ Source: ESG Master Survey Results, [2019 Technology Spending Intentions Survey](#), February 2019.

Figure 3. Technology Challenges with the Biggest Negative Impact on Productivity

Which of the following technology challenges have the biggest negative impact on the productivity of you and your colleagues? (Percent of respondents, N=210, multiple responses accepted)



Source: Enterprise Strategy Group

Welcome to a World of Cloud-native Global File Systems

A global file system is just that—global. A single namespace that lives in the cloud, that scales in every dimension in real time dynamically, and that provides data access to, and from, anywhere. Imagine a single “thing” to manage, no matter where the people who utilize it are located—a single thing that one part-time person can manage regardless of the size and complexity of the file system.

And it keeps getting better. A global file system enables:

- **Access.** Everyone can see and access the same file system from anywhere. It doesn't matter where on the planet your users and collaborators are located; if they can get to the cloud, they can get to the file system.
- **Availability.** The entire file system is automatically protected, altogether eliminating the need for backup. An organization can keep multiple copies of the file system in multiple clouds concurrently, and each cloud copy would have essentially infinite point-in-time copies—including persistent WORM versions.
- **Scale.** Imagine the ability to scale to any level. Need more performance? Simply add more front-end cache. Need to go from a 10PB file system to a 100PB? It happens dynamically. Scale up and out dynamically with zero management. Capacity is unlimited.
- **Portability.** Move enterprise-scale volumes of file data between clouds seamlessly and easily.
- **Collaboration.** Enable global file locking to provide collaboration on an unprecedented level—again, no matter where collaborators are located.

Why It Is Inevitable?

Let's be honest. If we were starting with a clean slate, none of us would architect our unstructured data infrastructure the way most enterprises have. No one anticipated the staggering volumes of data we would have to manage, which has created the need for a brand-new technique.

Originally, we bought a big file server. When it got too slow or ran out of capacity, we bought another. And so on and so on. Once we realized that not every one of our users was in the same building, we had more problems: We needed to send replicate file servers all over the world. Then we needed to coordinate their file systems, adding complex tasks such as replication. The management burden of distributed architectures became massive. How do we ensure the protection and availability of those systems? The answer—buy more. “Buy more” has been industry's response to this problem. But if you were starting out today, there is no way you would burden yourself with this type of complexity and expense.

Modern Applications Require Modern File Infrastructure

Just like we no longer write legacy applications that are monolithic stovepipes, we must stop deploying infrastructure in the same manner. Modern SaaS providers and hyperscalers do *not* build monolithic systems. They build highly scalable elastic systems that use the resources they need when they need them, relinquishing those resources back into the resource pool after they are done. They scale in real time in every dimension.

Amazon and Salesforce aren't running their businesses on scale-up NAS systems. All hyperscalers are built on true global file infrastructures. The key to hyperscale infrastructure is not its ability to get really big. Scale alone is not enough. Hyperscalers succeed not only because their deployed infrastructure is massive, but also because they leverage the right level of automation to manage that scale with as few people as possible. Personnel are the most valuable resource. A modern global file system offers the scale and flexibility to minimize the IT staff required to manage and support it. Simultaneously, the always-available-everywhere approach maximizes the productivity of end-user knowledge workers. It is a win-win. Hyperscalers don't worry about scale. They worry about making sure that users are able to deploy easily and seamlessly even while the back-end might be in flux. They worry about user experience, which means it's easy, it's available, and it's always accessible.

The Bigger Truth

Don't be dumb. Stop doing things the old way. Customer expectations and business in general are drastically different than they were 10-15 years ago. Why would we expect the technologies of 10-15 years ago to still apply? The fundamental truth of the data-driven economy is that data is not just necessary for business any more, it also unlocks new opportunities, new efficiencies, and new revenue growth. Before, we were all just “storing.” Now, we're accessing that data and using it a lot more to extract its value.

And this is just the beginning. The data deluge is not slowing down. Stop cobbling together traditional systems to solve a modern problem. Leverage a modern, global file system to make your life easier and your business better.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.