

Peer Software and Scality - A Distributed File System Approach to Scale-out Storage

Contents

Introduction - What We All Want	2
Why Cloud Gateways	2
The Problem With Caching	2
More Sites More Problems	3
Multi-platform Support	3
Resiliency	3
PeerGFS with Scality RING	4
Highly Available	4
Object Store Resilience and Backup	5
Multi-site Architecture	5
Easy to Implement	5
Key Benefits	5
Conclusion	6

This document is designed to help the reader understand how Peer Global File Service (PeerGFS) can be combined with the super scale-out capabilities of the Scality RING to deliver the benefits promised by legacy cloud gateway solutions without having to face the challenges and disadvantages typically encountered in this category.

Introduction - What We All Want

The explosive growth of unstructured data has driven organizations to seek solutions to efficiently store, manage and protect this ever growing universe of data while enabling fast and easy access for their users. In multi-site enterprises, slow WAN-based access to data has proven to be a constant bottleneck leaving users feeling “punished” for having to do their jobs utilizing slow, high latency connections to their work files and applications. Additionally, the need to collaborate across geographies so that two or more users in different offices do not overwrite each other’s work while sharing the same file is a key requirement. Finally, the need for a bulletproof backup option that leverages cloud-based storage and minimizes operational overhead is an imperative.

It is tempting to solve this conundrum by storing everything centrally, typically in a private or public cloud that utilizes scalable cloud-based storage, and then caching frequently accessed data at branch offices on local storage. In fact, Microsoft was on the forefront of this concept when it introduced BranchCache to support this strategy, and shortly after multiple cloud gateway offerings began to emerge with a seemingly familiar approach once you have gotten past the marketing rhetoric.

Why Cloud Gateways

Organizations deploy cloud gateway solutions to solve challenges that traditional storage infrastructure does not address such as data synchronization across multiple locations and collaboration capabilities. Older generation legacy cloud gateway products typically rely on an appliance utilizing proprietary cache storage that is deployed on location, and connects to public or private cloud storage providers.

Appliances are usually sized by anticipated storage utilization (local cache), number of users, require a high throughput/low latency network connection, and in many cases deploy a proprietary file and namespace system.

The Problem with Caching

When you choose to cache “just the data you need”, you are making the assumption that the data needed is known. To help with this, many legacy cloud gateway solutions utilize algorithms to predict and maintain frequently accessed data on local appliances. Granted, older data may not be read often, but it is needed, and when it is the caching approach can fall short. This is especially bothersome in use cases where large project files are in play that may contribute to exceeding local cache resources. For example, many CAD applications will not support the long load times a cloud gateway will require when it is asked to serve up a file that is not in the local cache.

Be aware of local performance and access concerns related to file sharing and collaboration. For example, many CAD applications will not support the long load times a cloud gateway will require when it is asked to serve up a large file that is not in the local cache.

More Sites, More Problems

Many products in the cloud gateway space work well in implementations with a limited number of sites, but are tested when scalability is required to support additional locations. This can stem from high latency connectivity into and out of the cloud, limited resiliency in the event of an appliance failure and the inability to optimize storage systems forcing organizations to deal with all data in the same way.

Multi-platform Support

Organizations have made significant investments in storage platforms they would like to continue utilizing while enhancing their storage strategy with the cloud. This is understandable as each platform has certain capabilities that make it valuable. Additionally, these same organizations benefit from the competition around each tech refresh of their storage platforms. With a legacy cloud gateway, these advantages evaporate due to vendor lock-in not only at the hardware layer, but also with where and how data is stored.

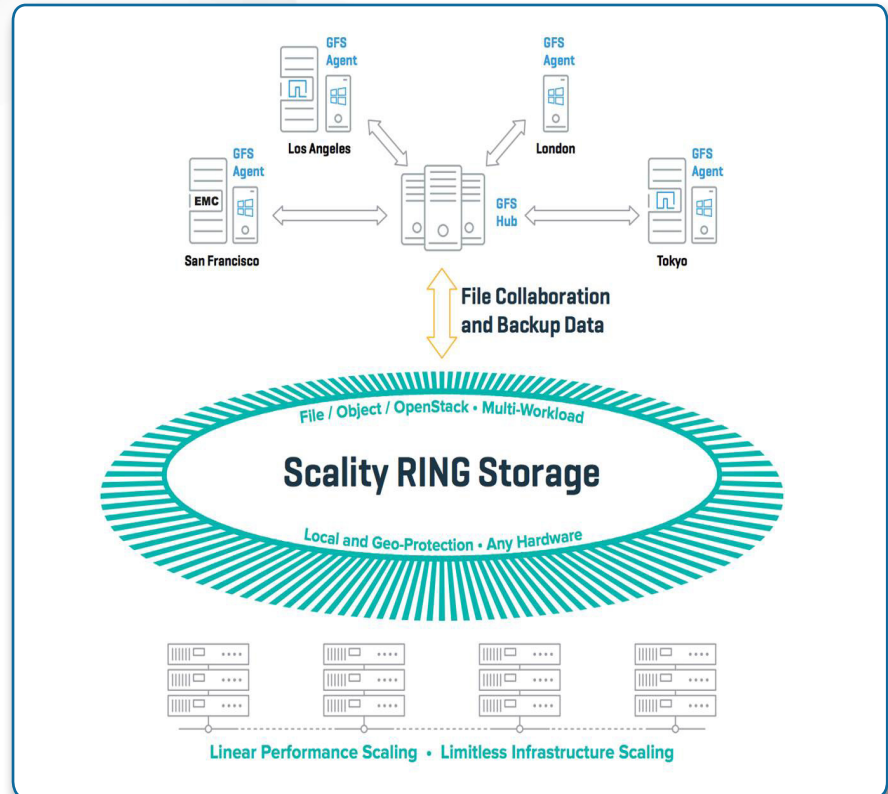
Resiliency

Things go wrong which is why no vendor can promise 100% uptime for any system. The key is to not take the user down with the system. Because legacy cloud gateways tend to replace primary storage on familiar NAS and SAN systems, they make organizations put all their eggs in their basket. Unfortunately, many organizations discover this during a failure event which challenges a solutions resiliency, and can lead to a state that prevents users from accessing their shared work files which hinders productivity and creates frustration.

PeerGFS With Scality RING

By combining Peer Software's distributed global file service (PeerGFS) with the security and scale of the Scality RING we can get all of the advantages promised by legacy cloud gateways, without their drawbacks.

PeerGFS and Scality RING work in concert to seamlessly integrate distributed file and object storage architectures for backup and file collaboration use cases.



Highly Available

PeerGFS allows you to leverage your existing systems and layers on top of them a software-only solution that mitigates the WAN's negative effects by giving users a local copy of their data that remains accessible even if PeerGFS components are shut down. This capability means that productivity does not stop just because a service account expires or a network link goes down.

Object Store Resilience and Backup

Scality RING Storage provides the highly available, yet hardware agnostic backup option that leverages object-based storage technology. PeerGFS ensures that all updates to protected data loads are replicated to Scality so that backups are ongoing all the time and redundant.

Connectivity between PeerGFS and the Scality RING is achieved by using the Scality Scale-Out File System (SOFS) connector. SOFS is a POSIX based virtual file system that provides SMB file access to shared folders. SOFS also features integrated file system load-balancing and failover capabilities that provide enterprises with optimal performance and reliability.

Multi-site Architecture

Common issues faced by file sharing and backup in enterprise environments include file version control, access performance, minimizing backup windows, and scaling data across locations. PeerGFS with Scality RING storage solves all of these without the compromises endemic to legacy cloud gateway solutions.

Easy to Implement

The PeerGFS software solution coupled with Scality's modular design means that implementation is far less complicated than the rip and replace (plus migration) approach.

Key Benefits

- **Multi-site multi-platform solution** - Integrates next generation object-based storage, real-time backup, and file sharing with version management across multiple physical locations and disparate storage platforms
- **Software-based** - Leverages existing investments in namespace, file systems and storage hardware without having to acquire proprietary appliances
- **Scalable** - Grow without imposed limits on number of users or available storage
- **Resilient** - Deploys highly-available storage at both the Windows file sharing level and with the Scality RING

The Scality RING object store integrated with PeerGFS is a powerful combination for managing the explosive growth of unstructured data across a multi-site enterprise.

Conclusion

Legacy cloud gateways promise to solve some very thorny problems for geographically dispersed teams, but they pose their own new sets of challenges typically resulting in proprietary vendor lock-in. Rather than looking for a revolution that disrupts trusted solutions that have worked on many levels, it is better to implement a solution that works in concert with evolutionary technologies like object-based storage without losing the freedom of choice.

There are solutions that do not require throwing out all of the systems you are currently using. As the hype dies down about all things cloud, the Scality RING object store integrated with PeerGFS presents itself as a powerful combination for solving challenges related to managing the explosive growth of unstructured data across a multi-site enterprise.

Learn more at www.peersoftware.com.