

Bringing the Benefits of Caringo Scale-Out Object Storage to NFS-Based Applications & Workflows

Benefits

- Improve productivity with data portability
- Reduce risk of data loss with no single point of failure & secure access
- Lower TCO of data distribution at scale with the benefits of Swarm object storage

Caringo® SwarmNFS is the first lightweight file protocol converter to bring the benefits of scale-out object storage—including built-in data protection, high-availability, and powerful metadata—to NFSv4. Unlike cumbersome file gateways, SwarmNFS is a stateless Linux® process that integrates directly with Caringo Swarm and is deployable on any Linux server or virtual machine (VM) running CentOS™ 7.x. SwarmNFS turns any Swarm object storage environment into a fully distributed NFSv4 solution. With SwarmNFS, a single mount point can be accessed across campus, across country or across the world. Our patent-pending technology delivers a truly global namespace across NFSv4, S3, HDFS, and SCSP/HTTP.

FULLY DISTRIBUTED NFS ARCHITECTURE

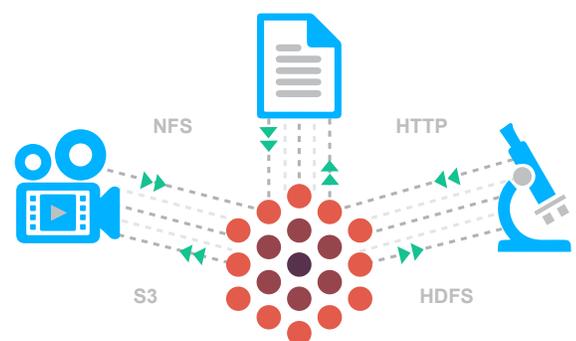
Easily Centralize, Distribute, and Manage Data

SwarmNFS leverages the fully distributed, massively parallel architecture of Swarm to create a distributed NFS solution. Files can be ingested at any site and readily available across the globe via a single mount point. Swarm provides multi-tenancy, versioning and enhanced security, including integration with AD, LDAP and PAM. Once on Swarm, data is automatically protected, eliminating the need for additional backup solutions. Swarm provides automated policy-based data retention and lifecycle management from creation to expiration.

GLOBAL ADDRESS SPACE

Eliminates Storage and Protocol Silos with Universal Access

SwarmNFS is part of the Caringo ecosystem which enables you to eliminate storage silos and access data from any application, device or location. Data is written to a global address space making it portable with multi-protocol read and write through any combination of NFSv4, S3, HDFS or SCSP/HTTP. Data from traditional file and web-based applications can now be streamed directly in and out of a single pool of storage—Caringo Swarm—scaling with no degradation of performance. Once on Swarm, metadata search is available through a web-based UI or programmatically via an API. In addition to Domains and Buckets, saved searches (called Collections) can be preserved and mounted via NFS—ensuring that every time you view a mount point, the data you are looking for is instantly accessible, without the overhead of needing to rerun a search query.

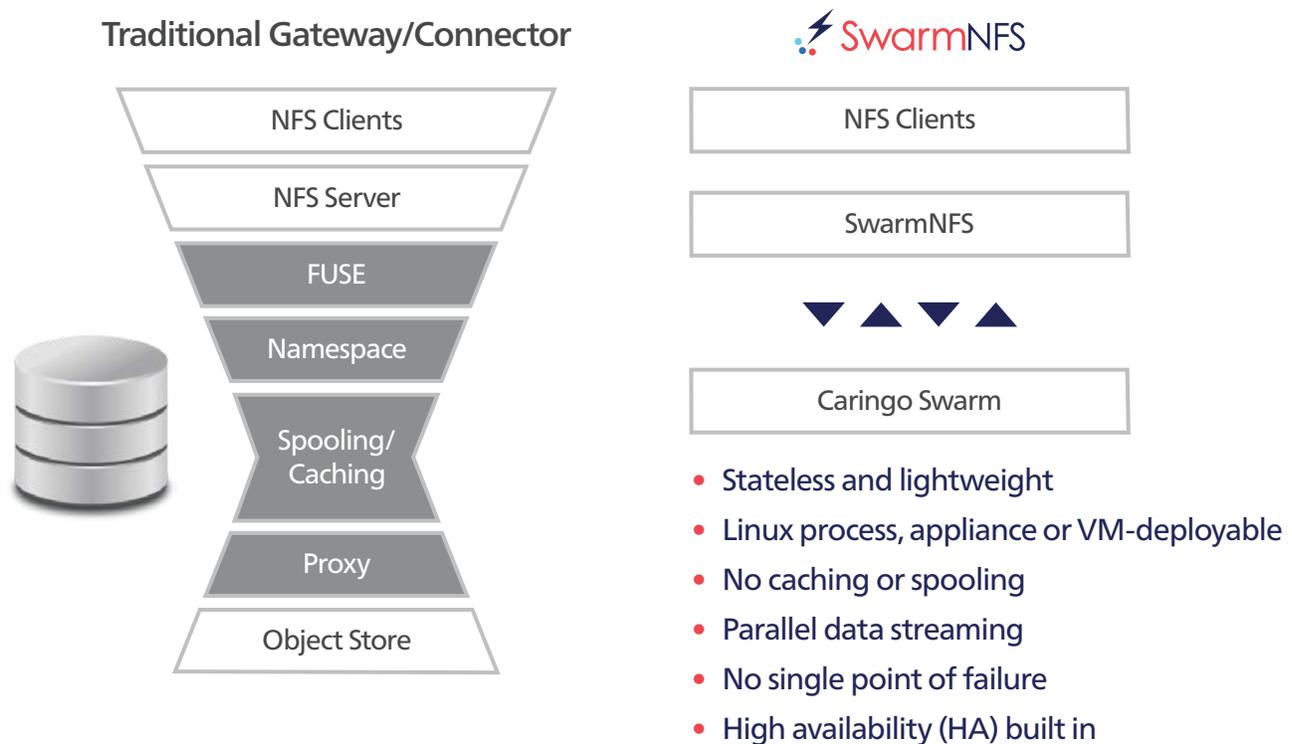


STATELESS AND LIGHTWEIGHT

Uses 80% Less Resources than Traditional Filesystem Gateways and Connectors

The power behind SwarmNFS lies in the innovative protocol conversion approach that streams data from NFSv4 clients directly to and from Caringo Swarm scale-out storage. This eliminates the need for spooling and caching, significantly reducing risk of data loss. Additionally this approach reduces the amount of disk, CPU and RAM that is traditionally required for file-to-object gateways—resulting in limitless scale of throughput while eliminating bottlenecks.

With SwarmNFS, you reap the benefits of out-of-the-box high availability (HA) that requires no local cache and no clustering with easy failover and quick restart. Authentication and authorization settings in Swarm propagate through all protocols, ensuring your data is secure—no matter how it is accessed.



BRINGS THE POWER OF METADATA TO FILES

Eliminates Metadata Databases

SwarmNFS brings the power of metadata, inherent in Swarm object storage, to all files. With SwarmNFS, metadata can be added to files directly from the client. Annotating files with metadata advances the way files can be searched, organized and analyzed at scale. Once on Swarm, data can be profiled with big data analysis tools such as Kibana or Hadoop™. Collections of files can be mounted based on the result of a metadata search.

For More Information