MegaFile: Faster Backups and Restores for Large Files

Reduce Downtime and Meet Critical Business SLAs

Key Benefits

- High performance for backups and restores of large, multi-terabyte files
- Reduce downtime in recovering from a disaster or ransomware attack
- Applicable for a broad range of sources including VMware, Hyper-V, Linux, AIX, and relational databases

As a high-performance solution, Cohesity provides quick backups and restores even for large, multi-terabyte files. MegaFile is a unique component of Cohesity's architecture that allows fast backup and recovery for large files, helping reduce backup times as well as downtime when recovering from a disaster.

Differentiated Parallel Streaming

MegaFile is a differentiated approach to data streaming across multiple nodes, implemented using patented technology. MegaFile intelligently distributes files across all nodes in a cluster. An aspect of Cohesity SpanFS®, MegaFile breaks large files into smaller chunks for parallel backup and recovery across nodes. The minimum size of these chunks is optimized to maximize performance.

Unlike other vendors, Cohesity dynamically splits and distributes data, allowing large files to be divided into smaller chunks and streamed across the entire cluster.

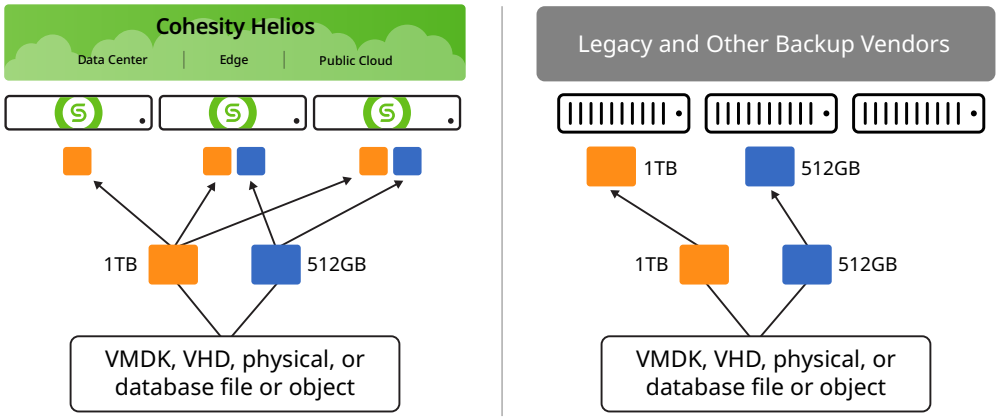

Dramatically Faster Backup and Recovery

When backing up a 2 TB file to an 8-node cluster, for example, MegaFile creates eight segments and ingests each of these data chunks in parallel across the Cohesity cluster. For this file, MegaFile decreases backup times by up to 8x. Adding more nodes in a cluster increases the performance further.

The same benefits of MegaFile are applicable for restores as well. With MegaFile, restores are dramatically faster, even for large, multi-terabyte files. This helps reduce downtime and helps teams meet SLAs.

Cohesity's intelligent multi-node streaming is applicable for VMware, Microsoft Hyper-V, Pure Storage, the leading databases, and physical servers—supporting VMDK, VHD, physical, and database objects.

Higher Performance Regardless of Size

Dramatically Faster Backups and Retores

- Higher backup and recovery performance that scales with more nodes
- Increase the performance of protecting and recovering multi-terabyte VMs, databases, and physical servers

Supporting a Range of Data Sources

- Supports block-based and file-based data
- Supports VMware, Hyper-V, Linux, AIX, Pure Storage and SQL Server

Higher Performance Across Workflows

- Supported for both full and incremental backups
- Meet business SLAs more easily, with shorter backup windows even for large files

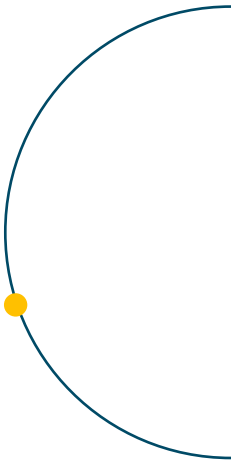

Learn more at www.cohesity.com

COHESITY