

IBM Aspera Sync

High-performance multi-directional file synchronization and replication

Key benefits & capabilities

- Designed for extreme scalability, with the ability to support millions of files and multi-terabyte file sets
 - Faster speed for bulk-data synchronization over WANs
 - Efficient change detection and updates
 - 100 percent reliable, high-speed data delivery with Aspera FASP transport technology
 - Flexible transfer and bandwidth management
-

IBM® Aspera® Sync is purpose-built by Aspera for high-performance, scalable, multi-directional asynchronous file replication and synchronization. Designed to overcome the performance and scalability short-comings of conventional synchronization tools, Aspera Sync can scale up and out for maximum speed replication and synchronization over WANs, for today's largest big data file stores – from millions of individual files to the largest file sizes.

The FASP advantage


Built upon Aspera FASP®, Aspera Sync transfers data between peers at full bandwidth capacity, regardless of distance and network conditions. Running on commodity hardware, Aspera Sync reconciles file system changes with remote peers at extremely high speed over global distances, and does not degrade in performance as the numbers of files increase, resulting in speed improvements up to 100x faster than others.

Smarter synchronization

Aspera Sync has ultra-fast snapshot performance for synchronization of giant file stores containing 1 million or more files. Unlike conventional replication tools which copy any new data over the WAN, Aspera Sync intelligently recognizes changes and file operations such as moves and renames, instantaneously propagating these to remote peers, avoiding what can be hours of unnecessary copy times. Aspera Sync offers higher performance, without a steep learning curve.

Flexible deployment options

Unlike unidirectional-only tools, Aspera Sync supports both push and pull mode in bi-directional and multi-directional synchronization topologies where content is changing on multiple nodes. Replication jobs can be configured to run continuously for real-time synchronization, or one-time, on demand. Aspera Sync allows users to safely move, rename, and delete files or entire directory structures, without fear of data loss. Windows ACLs and OS X extended attributes are preserved, across syncs and bi-directional workflows, on Windows will tolerate files opened by users during synchronization.



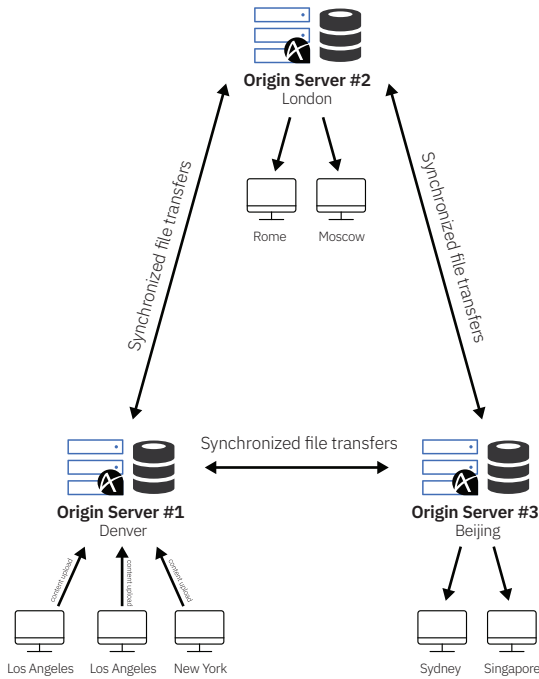


Figure 1: A typical Aspera Sync deployment

Key features

- Multi-directional synchronization of remote files and directories.
- One-to-one, one-to-many, and full-mesh synchronization.
- Open architecture for integration with third-party processes and systems.
- Concurrent synchronization session architecture with support for clustering and multi-gigabit transfer speed.
- Remote sync monitoring through IBM® Aspera® Console web application.
- Locally detects changes and compares them to file system snapshot without having to check with remote systems.
- Replicates file moves and file renames on the source as a file move or rename on the target, avoiding unnecessary data copying.

Supported platforms

Limited-use Internet access

- Linux
- Windows 2000, XP, 2003, Vista, 2008, 7
- OS X
- Solaris and BSD

Typical applications

Disaster recovery and business continuity

With FASP-powered transfers, avoid slow transfer speeds between primary and backup sites that can result in incomplete backups and slow recovery times. Replicate and back up mission-critical data at high speed from a primary site to one or more alternate sites to shrink the recovery point and recovery time, and help ensure systems remain available after an outage or site loss.

Content distribution and collection

Create multi-site, multi-directional synchronization topographies to collect or distribute content, software updates and business data across remote, geographically dispersed locations, regardless of file sizes, locations and distance.

Real-time system mirroring

Replicate on demand or continuously synchronize servers in real time over the WAN to improve data access and service availability.

File archiving and remote storage

Create continuous or scheduled archive of inactive data from high-speed primary storage to remove second-tier storage and allocate a specific amount of bandwidth to be utilized, thereby providing a high level of service for the regular business data network traffic.

Server and VM migration, replication, back-up and recovery

Automate replication and leverage high-speed transfers to minimize setup time, duplication and re-installation of new and existing systems and to maintain accurate copies for development, sandbox, quality assurance and standby.

Features and benefits

Industry-leading Aspera performance

- Built on Aspera FASP technology for maximum transfer speed regardless of file size, transfer distance and network conditions.
- Precise bandwidth control ensures that the entire allocated bandwidth is utilized to help achieve maximum transfer speeds, while being fair to other critical network traffic.
- 100 percent reliable data delivery, real-time reporting of transfer progress and performance.

First Run

Performance comparison synchronizing many small files (average size 100 KB) over WAN of 100 ms/1%

SMALL FILES	NUMBER OF FILES	DATA SET SIZE	SYNC TIME	THROUGHPUT
Aspera Sync	978,944	93.3 GB	9,986 sec (2.8 hours)	80.4 Mbps
Rsync	978,944	93.3 GB	814,500 sec (9.4 days)	0.99 Mbps
Speed up difference				1x

Performance comparison synchronizing many small files (average size 100 KB) over WAN of 100 ms/1%

LARGE FILES	NUMBER OF FILES	DATA SET SIZE	SYNC TIME	THROUGHPUT
Aspera Sync	5,194	500.1 GB	4,664 sec (1.3 hours)	921 Mbps
Rsync	5,194	500.1 GB	4,320,000 sec (50 days)	0.98 Mbps
Speed up difference				940x

Second Run

Synchronization time after adding 31,056 files to 1 million small files (100 KB each) over WAN of 100 ms/1%

CHNAGE FILES	INITIAL FILES	ADDED FILES	TOTAL SIZE	SYNC TIME	THROUGHPUT
Aspera Sync	978,944	31,056	2.97 GB	947 sec (16 min)	26.9 Mbps
Rsync	978,944	31,056	2.97 GB	37,076 sec (10.3 hrs)	0.68 Mbps
Speed up difference					39x

Synchronization time after adding 31,056 files to 1 million small files (100 KB each) over WAN of 100 ms/1%

CHNAGE FILES	INITIAL FILES	ADDED FILES	TOTAL SIZE	SYNC TIME	THROUGHPUT
Aspera Sync	5,194	54	5.49 GB	54 sec	871 Mbps
Rsync	5,194	54	5.49 GB	54,573 sec (15 hours)	0.86 Mbps
Speed up difference					1000x

Figure 2: IBM Aspera Sync performance compared with rsync

Flexible options for replication and sync of remote files and directories

- One-to-one, one-to-many, and full-mesh synchronization.
- Supports both uni- and bi-directional synchronization in both push and pull mode between peers over WAN/LAN.
- Windows bidirectional synchronization tolerates open files during sync.
- One-time replication as well as continuous synchronization.
- Remote sync monitoring through Aspera Console web application.

Designed for extreme scalability

- Optimized performance (up to 10x faster) with highly scalable architecture supports millions of files and multi-terabyte file sets.
- Concurrent synchronization session architecture with support for clustering and multi-gigabit transfer speed.

Efficient change detection and updates

- Locally detects changes and compares them to file system snapshot without having to check with remote systems.
- Built-in deduplication (in bidirectional as well as unidirectional mode) detects multiple copies of a file at the source and creates links to a single copy of the file at destination, saving transfer and storage capacity.
- Super fast scan to detect changes in scan mode for large incremental data sets.
- Advanced features for preservation of file timestamps, secure docroots, and token authorization.
- Preview of Aspera Sync to S3 for unidirectional sync of file stores with cloud storage
- Replicates file moves and file renames on the source as moves or renames on the target.
- Distributed event collection system enables the fastest possible capture of file system changes on clusters of ingest servers and high availability deployments.

About IBM Aspera

Aspera, an IBM company, is the creator of next-generation transport technologies that move the world's data at maximum speed regardless of file size, transfer distance and network conditions. Based on its patented, Emmy® award-winning FASP® protocol, Aspera software fully utilizes existing infrastructures to deliver the fastest, most predictable file-transfer experience. Aspera's core technology delivers unprecedented control over bandwidth, complete security and uncompromising reliability. Organizations across a variety of industries on six continents rely on Aspera software for the business-critical transport of their digital assets.

For more information

On IBM Aspera solutions, please visit us at <https://www.ibm.com/cloud/high-speed-data-transfer> or contact aspera-sales@ibm.com.



© Copyright IBM Corporation 2018

IBM Corporation
Route 100
Somers, NY 10589

Produced in the United States of America
November 2018

IBM, the IBM logo, ibm.com and Aspera are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: ibm.com/legal/us/en/copytrade.shtml

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other product, company or service names may be trademarks or service marks of others.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on the specific configurations and operating conditions. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM product and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.



Please Recycle