



Resilio Connect: Sizing Guide

This guide is intended to help partner companies size their Resilio deployment to achieve their desired performance goals. Sizing requirements vary by use case, desired transfer speed, and payload (i.e., number of files, file size, and file change rate). Other factors like job type, OS, storage, and network also play a role in sizing.

In general, 1 Resilio agent will run on each endpoint. In cases where larger numbers of files need to be transferred, multiple agents can be used to distribute the load of scanning and indexing files. See the Sizing by Payload section.

Resilio Connect Overview

Resilio Connect is a [centrally managed, agent-based solution](#). agents run on a range of devices (physical or virtual) and operating systems (Mac, Windows, Linux, Android, FreeBSD, and some [NAS systems](#)). The system is peer-to-peer (P2P). All agents are centrally controlled through the Resilio Connect management console.

Sizing by Performance Tier

The following table represents examples based on throughput and number of files.

Speed	Number of Files	CPU	Rec'd System Memory ¹	Maximum Capacity	agents per File System
100 Mbps	3 million	4 core	4GB	5TB	1
1 Gbps	10 million	4 core	12GB	50TB	1
2 Gbps	50 million	8 core	32GB	250TB	1 or more ²
10+ Gbps	Unlimited ³	8+ core	128GB	Unlimited	1 or more ⁴

Synchronizing 1 million files of varying file sizes across a few servers with a target throughput of 500Mbps can be achieved on most commercial off-the-shelf servers. By contrast, synchronizing

¹ Resilio consumes about 256 bytes of memory per file or entry for RAM-optimized jobs.

² For jobs requiring periodic scans (indexing), it's recommended that at least 1 agent per 50 million files be deployed. Resilio does not provide automatic load balancing between agents.

³ While there is no technical design limit on the number of files, Resilio recommends a soft limit of 400 million files per job. This will vary by file change rate and the performance of each agent.

⁴ For jobs requiring periodic scans (indexing)--i.e., when file change notifications are turned off--it's recommended that 1 additional agent be assigned to every 50 million files transferred or sync'd.

100s of millions of files across hundreds of servers will have more resource-intensive sizing recommendations. Resilio recommends deploying multiple agents for larger payloads over 50 million files or as needed for a given payload and target throughput.

Sizing Resilio Connect Agents

Resilio Connect agents do the work of scanning, indexing, and replicating files. For sync jobs, file changes are tracked on the OS-level by each agent. When synchronizing large numbers of files across many peers, the faster that each agent picks up and reconciles file changes across all peers, the better the end-to-end performance. To eliminate bottlenecks, each agent can be properly sized based on the target payload, job type, desired speed, and use case.

File Change Notifications

Most modern operating systems and file sharing protocols (such as SMB) publish file change notifications. Most NFS-based NAS systems do not.

File Change Notifications Enabled

When a Resilio Connect agent is installed, Resilio automatically detects and replicates file changes. Resilio provides a fast and efficient mechanism to capture, replicate, and synchronize file changes when installed directly on an OS. Resilio also tracks file changes over SMB.

Scheduled Scans over NFS

Alas, some NAS systems and NFS don't support file changes. In these situations, when file change notifications are not available, period scans may be performed.

Depending on the number of files being replicated, each Resilio system can scan a portion of the files on the NAS. Resilio recommends dedicating a single Resilio agent for every 25 million files being scanned. For example, if your goal is synchronizing 100 million files stored on an NFS server, you would partition the job across (4) NFS client systems running Resilio Connect. And so forth.

Other Performance Considerations

Job Type

The best practice is to spread large numbers of files across multiple directories to spread CPU and disk IO consumption over time. Distribution and consolidation jobs also consume disk IO and require folder scanning more resources than distribution or consolidation jobs.

Number of Files

Resilio has validated up to 400 million files per sync job. It's recommended that the load be spread out across multiple agents, each tasked with replicating a maximum of 50 million files each.

File Sizes

Resilio Connect can transfer and synchronize files of any size and type. Yet file sizes do impact the performance of the job at hand. In general, large numbers of smaller files (under 8MBs) require faster storage, with more IOPS. A best practice is to use faster storage and spread the file count out across multiple directories or nested sub-directories.

Directory Structures

The general best practice is to avoid flat directory structures. Multiple agents can be used to spread out the load scanning and indexing file systems by pointing each agent to its own directory within the file system.

When possible, nest directories in a tree-like structure of files and folders. For smaller files under 8MB in size, limit the number of files to 100K per directory, when possible. If not possible, consider using faster storage such as SSDs to provide more IOPS per agent.

File Change Rates

For capabilities such as real-time sync, file change notifications are tracked on the host OS. When file change notifications are not available, each Resilio agent performs a periodic scan of a given pathname, share, or drive mapping. The more often that files change, the higher the performance burden on each system in the job. In general, faster storage and more memory helps to keep up with frequently changing files.

Running Multiple agents

For larger payloads, or in situations where file change notifications are not enabled, multiple agents can be run to spread out the load of scanning, indexing, and replicating files across multiple peers.

There are two modes for running multiple agents:

- Run agents side-by-side on the same OS instance. Resilio supports running multiple agents on the same system as long as each agent is replicating a separate directory. Each agent can be used to scan its own set of files and directories.
- Mount the source file system via SMB or NFS. Multiple agents can mount and scan a source file system across high-speed connections. This can be used to partition the scan and index process of large file systems.

For frequently changing files, Resilio suggests deploying an agent for every 25 million files. For less active files, each Resilio agent could be sized to index 50 million files, on average. For example, if the source file system (that does not support file change notifications) contains 150 million files, 3 or more agents are recommended to mount and index the file system. agent 1 would scan and index /fs/dir1, agent 2 /fs/dir2, agent 3 /fs/dir3, and so forth.

Storage Performance - Throughput and IOPS

The faster the storage, the better the performance. In general, synchronizing many smaller files requires more IOPS than larger files. As the payload size and throughput increases, faster storage — using SSD or All-Flash Arrays (AFA) — will provide the highest level of throughput and IOPS to eliminate bottlenecks.

Bandwidth & Network

Achieving faster end-to-end throughput across multiple peers requires that each peer contain sufficient CPU, memory, storage throughput, and IOPS. For example, if your target speed is 10 Gbps, you'll need to ensure all endpoints meet the following system requirements at a minimum:

CPU	Xeon E5 2600 v4 series or i7-7000 series, with 4 physical cores or more (double that on a VM). Minimal frequency 2GHz.
Memory	16 GB or more (in addition to files' based requirements above); minimal memory bandwidth throughput of 10 GBps is required.
NIC	For expansion cards, the expansion slot must support 4 lanes for PCIe 2.0. Link aggregation should be avoided. For NUMA machines with some CPU slots empty, the network card PCIe slot should connect to the occupied CPU slot.
OS	The performance requirements of each operating system (Windows, Linux, macOS, etc.) should be taken into account when sizing the deployment.

Disk

Read/write speed more than 1.2 GB/s for speed of 10 Gbps.

Other Considerations

For most use cases, there will be a 1-to-1 relationship between the number of physical systems and Resilio agents. Each computer system participating in a file transfer or sync job runs an agent. In most cases, there will be 1 Resilio agent installed on each system. But depending on the storage architecture and factors such as file notifications (on or off), your environment may need more agents to transfer or sync your target payload.

In general, higher performance (10+ Gbps) with larger numbers of files depends on faster system performance. Faster storage (with direct- or SAN-attached SSDs) and more memory will improve performance.

Sizing by Payload

The total number of agents deployed depends on the total number of files, file change rate, and scan intervals, and if file change notifications (on/off) are supported. A simple rule to follow is:

- File change notifications: 50 million files per agent (per job), depending on file change rate
- Period scans: 25 million files per agent for each job, depending on file change rate.

Number of Files per Job	50M files			150M files			300M files		
	Weekly	Daily	Hourly	Weekly	Daily	Hourly	Weekly	Daily	Hourly
File Change Rate ⁵									
Change Notifications - ON	1	1	1	2	2	3	4	5	6
Change Notifications OFF ⁶	1	1	2	3	3	6	6	6	12
Recommended Number of agents per File System									

⁵ Some customers may run frequent scans. Others may limit scans to 1 or 2 per day, depending on their file system size and system performance.

⁶ When the OS does not have file change notifications, Resilio relies on a periodic scan. Multiple Resilio agents can scale-out the indexing process.

Sizing the Management Console

The Resilio Connect management console runs on Linux or Windows. Each management console instance can be hosted on a physical or virtual server. Each instance supports managing up to approximately 10,000⁷ Resilio agents.

That said, most customers will not need to support that many agents per console instance. The following chart provides some sizing guidelines for each console instance.

Each of the following estimates varies by workload. More jobs, more files, and more frequently changing files drives up the requirements. The following approximations provide a baseline starting point.

Deployment Size	Job Runs	Minimum CPU	Minimum RAM	Storage Free Space ⁸
50 agents 10 jobs total 1 million files	Daily	2 Core	4 GB	4 GB
	Hourly	4 Core	8 GB	4 GB
100 agents 50 jobs total 10 million files	Daily	4 Core	16 GB	14 GB
	Hourly	4 Core	16 GB	14 GB
1,000 agents 300 sync jobs total 100 million files	Daily	4-8 Core	32 GB	120 GB
	Hourly			
10,000 agents 750 jobs total 250 million files	Updated Hourly	8+ Core	64 GB	200 GB
	Updated Daily	8+ Core	64 GB	200 GB

For specific OS version support and a more detailed breakdown of storage space requirements, please refer to the [system requirements on Resilio.com](https://resilio.com/system-requirements).

⁷ This is not a hard limit but a recommended limit. Customers with more than 10,000 agents may deploy additional console instances or use the API. Guidance provided by Resilio engineering.

⁸ Does not include backup or archive storage. This number is based on 4MB for each job plus additional space for events and logging.

Solution Sizing by Use Case

Each solution scenario will describe a sample small, medium, and larger configuration to use as a starting point for your deployment.

DFSR Replacement on Windows Server

While there may be caveats to the rule, there will be a 1-to-1 relationship between the number of Windows Servers running DFSR and the number of Resilio agents required. One Resilio agent should be installed on each Windows Server previously running DFSR, joined to DFS.

For example, if you have 10 Windows Servers running DFSR, you'd license 10 Resilio Connect agents.

Number of Files	Change Rate ⁹	CPU ¹⁰	RAM	Storage Space	Sync Optimization
10 million	Daily	2 core	4 GB	40 GB	Built in
	Hourly	2 core	4 GB	40 GB	Built in
50 million	Daily	4 core	12 GB	190 GB	Built in
	Hourly	4 core	12 GB	190 GB	Built in
100 million +	Daily	4-8 core	32 GB	380+ GB	Built in + Map agents to Folders
	Hourly	4-8 core	32 GB	380+ GB	Built in + Map agents to Folders ¹¹

When synchronizing larger numbers of files above 50 million, Resilio can be performance optimized by structuring the file system and directory layout across multiple directory. Many files stored in flat directory structure perform worse than many files spread out (nested) across multiple directories. Also, agents can be mapped to monitor a share or drive mapping (e.g. D:\shareName) to spread out the load synchronizing large file systems.

⁹ For DFSR replacements, Resilio uses file change notifications in Windows Server.

¹⁰ 2Ghz minimum CPU requirement (e.g., Xeon E5 2600 v4 series or i7-7000 series)

¹¹ Sync jobs should be structured to spread the load across multiple file shares or drive mappings (vs. pointing the sync job at the DFS root)

Server Sync

The number of agents deployed will depend on the payload and OS type. Linux, Windows, and some NAS based OS's can run multiple Resilio agents on the same system. Also, multiple Resilio agents can mount a file store via NFS to scale-out the indexing process. When multiple agents are used, the metrics used in the [Sizing by Payload](#) section apply here.

Scenario 1: To sync 50 million files across 5 Linux or Windows systems using 1 Proxy Server in Site 1, and another 5 Linux or Windows Servers and 1 more Proxy Server in Site 2, you'd license 10 Resilio agents (optionally with WAN Optimization) and 2 Proxy Servers.

Scenario 2: Sync 150 million files stored on a central NAS in Site 1 across 25 Linux servers in Site 2. Depending on the agent systems' performance, you might license 4-6 Resilio agents in Site 1 to mount and index directories on the NAS via NFS. And then another 25 Resilio agents running on each target Linux system in the remote site (Site 2).

Number of Files	Change Rate	CPU	RAM	Storage Space	Sync Optimization
10 million	Daily	2 core	4 GB	20 GB	Built in
	Hourly	2 core	4 GB	20 GB	Built in
50 million	Daily	4 core	12 GB	95 GB	Built in
	Hourly	4 core	12 GB	95 GB	Built in
100 million +	Daily	4-8 core	32 GB	190+ GB	Built in + Scale-out
	Hourly	4-8 core	32 GB	190+ GB	Built in + Scale-out ¹²

Virtual Desktop Infrastructure

This 1 agent per server formula usually also holds true in other use cases as well: Any source or destination machine replicating files requires a Resilio agent. In VDI Profile Sync, for example, 1 agent is required per file server used for replicating VDI profiles. If you have 3 file servers in site A and another 3 in site B, then you'd have a total of 6 Resilio agents.

¹² For larger file systems, multiple agents can work together to scale-out the scanning and indexing process. Multiple agents can mount a common file system over NFS or SMB. Each agent should mount a unique folder or directory.

Add resource consumption for:

- WAN optimization
- NTFS permissions enabled

For larger deployments with many thousands of user profiles, you could distribute the replication load by logical business units or user group functions. For example, you might have 1 sync job for the Accounting department, 1 sync job for IT, and 1 job for Engineering.

Another consideration is logon and logoff behavior. If all users logon or logoff at the same time, there's a benefit to spreading the load across more Resilio agents to handle peak load.

Profile Containers	CPU	Available RAM	Storage ¹³	Sync Optimization
< 1000	4 Core	8 GB	300 MB/s	1 agent per file server
< 3000	4 Core	8 GB	500 MB/s	1 agent per file server Map sync jobs to individual user groups
< 8000 ¹⁴	8-16 Core	16 GB	700 MB/s	1 agent per file server Map sync jobs to individual user groups or divisions

Cloud Object Storage

Resilio supports a variety of cloud providers and their file, block, and object storage platforms. This includes most S3-compatible object storage and any block or file storage. Each provider has their own best practices and nuances when it comes to optimizing performance and right-sizing instances based on your specific use case or customer needs.

¹³ For faster speeds such as 10 Gbps, read/write storage speed of more than 1.2 GB/s is recommended.

¹⁴ More than 8000 containers are supported but hardware requirements will need to be determined based on size and use case.

For example, from an API perspective, Azure limits the number of access calls to each instance per hour. A single blob supports up to 500 requests per second. If you have multiple clients that need to read the same blob and you might exceed this limit, then consider using a block blob storage account. A block blob storage account provides a higher request rate, or I/O operations per second (IOPS).

Cloud Platform	Desired Speed	Instance Type	Required agents
AWS EC2	< 1 Gbit	t3.xlarge instance and below	1 per cloud instance
	> 1 Gbit	m5d.12xlarge instance and greater	1 per cloud instance
Azure	< 1 Gbit	D4s v3 instance and below	1 per cloud instance
	> 1 Gbit	D14 v2 instance and greater Standard_F32s_v2 for speed up to 9 Gbps (same region to sync files between storages)	
Google GCP	> 1 Gbit	n1- standard-64 and greater (same region to sync files between stores)	1 per cloud instance

Sizing the Proxy Server

The [Proxy Server](#) is an optional component used to route requests through the DMZ. Both Windows Server and Linux operating systems are supported. It's recommended that a dedicated server (physical or virtual) be used to host the Resilio Connect proxy server.

In general, the more concurrent connections you have, and the faster the speed requirements, the more resources you will need on each Proxy Server; numbers of files and agents is also a consideration. On average, each Proxy instance is capable of handling up to 20 million files and 2 thousand agents without significant performance degradation.

Concurrent Connections	Number of Files	Minimum CPU	Minimal Available RAM ¹⁵	Recommended RAM
512	< 100,000	1 core	512 MB	
1024	< 1,000,000	2 core	1024 MB	
2048	> 10,000,000	4 core	2048 MB	

Other Proxy Server Considerations

Proxy speed depends on factors such as hardware configuration, number of connections, and desired target throughput. With indirect connections through a Proxy, speed deterioration of up to 10% may be observed compared to direct connections between agents.

Here's an example set of best practices to meet a target speed of 5 Gbps:

CPU	Xeon E5 2600 v4 series or i7-7000 series, with 4 cores or more. Minimal frequency 2GHz.
Memory	16 GB or more. Minimal memory bandwidth throughput of 10 GBps is required.
NIC	For extension cards, the extension slot must support 4 lines for PCIe 2.0. Link aggregation should be avoided. For NUMA machines with some CPU slots empty, the network card PCIe slot should connect to the occupied CPU slot.
OS	For Windows, Windows Server 2019 is the recommended OS to meet these speed requirements. For Linux, a variety of OS types could be used to meet the targets above.

For the latest, most up-to-date system requirements on the Proxy Server, Management Console, and Agents, please see the [System Requirements section of the web site](#).

¹⁵ RAM requirements depend on the number of connections. A single agent may instantiate several connections if it uses the Proxy for several jobs. Other variables include the distance between the agents and Proxy Server and the bandwidth between them. The Minimal Value is 1 MB per connection. On 10 Gbps networks it may require 300 MB per connection.

More Information

If you need further assistance, please contact Resilio Sales (sales@resilio.com) or call 415-851-9884 (US only, M-F 9:00am-4:00PM).