

IBM Aspera Desktop Client 4.4.1



Contents

High-Speed Transfer Client Admin Guide for PowerLinux.....	1
Introduction.....	1
Getting started as a transfer client.....	3
Installation and upgrades.....	4
Requirements.....	4
Installing Desktop Client.....	4
Configuring the firewall.....	4
Testing a transfer.....	5
Uninstalling.....	5
Transfer files in the GUI.....	6
Global bandwidth settings.....	6
Enabling a transfer proxy or HTTP proxy.....	6
Adding and editing connections.....	7
Exporting and importing connections.....	14
Creating SSH keys in the GUI.....	14
Transferring content.....	15
Managing transfers.....	16
Scheduling and customizing transfers in advanced mode.....	17
Configuring transfer notifications.....	18
Using transfer notifications.....	22
ascp: Transferring from the command line.....	23
Ascp command reference.....	23
Ascp general examples.....	38
Ascp file manipulation examples.....	41
Using standard I/O as the source or destination.....	42
Using filters to include and exclude files.....	45
Symbolic link handling.....	51
Creating SSH keys	53
Reporting checksums.....	54
Client-Side Encryption-at-Rest (EAR).....	56
Comparison of ascp and ascp4 options.....	57
Ascp FAQs.....	60
ascp4: Transferring from the command line.....	63
Introduction to ascp4.....	63
Ascp4 command reference.....	63
Ascp4 transfers with object storage.....	71
Ascp4 examples.....	72
Built-in I/O provider.....	72
Appendix.....	73
Restarting Aspera services.....	73
Testing and optimizing transfer performance.....	73
Log files.....	75
Preserving IBM Spectrum Scale ACLs of transferred files.....	77
Logging client file system activity on HSTS.....	78
Product limitations.....	79

High-Speed Transfer Client Admin Guide for PowerLinux

Welcome to the High-Speed Transfer Client documentation, where you can find information about how to install, maintain, and use the High-Speed Transfer Client.


Introduction

Thanks for choosing Aspera and welcome to the world of unbelievably fast and secure data transfer.

The basics

Aspera high-speed transfers begin when an Aspera client authenticates to an Aspera server and requests a transfer. If the client user has authorization, then transfer tools are started on the client and server and the transfer proceeds.

For example, an IBM Aspera Desktop Client user connects to an IBM Aspera High-Speed Transfer Server and initiates a transfer:



Depending on the user's transfer request, files and folders can be transferred to the server from the client (uploaded) or transferred to the client from the server (downloaded). The source and destination can be cloud storage, an NFS or CIFS mount, and IBM Spectrum Scale storage, to name a few.

What is the server?

The Aspera server receives transfer requests from Aspera clients, determines whether the user has permission to access the server and authorization to the target area of the file system (source or destination with read or write access), and participates in transfers. The server can be:

- An on-premises installation of HSTS, IBM Aspera High-Speed Transfer Endpoint (which permits one client connection),
- A HSTS installed as part of IBM Aspera Faspex, or
- An HSTS deployed in object storage as an IBM Aspera On-Demand instance, an IBM Aspera on Cloud transfer service node, or an IBM Aspera Transfer Cluster Manager node.

What is the client?

The Aspera client is the program that requests a transfer with the Aspera server. Aspera applications that can act as clients include:

- Desktop Client.
- IBM Aspera Drive.
- IBM Aspera Connect.
- IBM Aspera Command Line Interface (ascli).
- HSTS and HSTE.

What is FASP?

At the heart of your Aspera ecosystem are the FASP transfer engines **Ascp** and **Ascp4**.

Ascp maximizes data transport over any network and is suited to large files. It is a powerful command line tool and also drives transfers that are started in the GUI.

Ascp 4 is another command-line transfer tool that is optimized for both large files and transfers of thousands to millions of small files, handling large amounts of file metadata as part of the high-speed transfer.

Both Ascp and Ascp 4 are installed and enabled with your installation of HSTS, HSTE, and Desktop Client.

The Aspera transfer server

Your Aspera transfer server is a powerful, customizable hub for your high-speed transfer activity. Configuration settings allow you to control which clients have access for uploading or downloading data, how much bandwidth their transfers can use, the priority of those transfers, and how data is secured during and after transfer. The transfer queue can be managed in real time, enabling you to adjust as priorities change. You can also monitor transfers and receive email notifications when transfer sessions or individual file transfers start and stop.

The Aspera Server GUI

The Aspera desktop GUI is primarily a client transfer tool, but it also offers a user-friendly interface for managing users and configuring your server on supported platforms (Windows, Linux, macOS). Security settings, bandwidth use policies, and file handling rules can all be set in the GUI. Configurations can be applied to all users (globally), to groups, or to individual users.

HSTS Web Portal

Your HSTS can be made even more accessible to clients by hosting a web-based storage directory. Authorized clients can browse files by using any modern web browser, and transfer using the free, automatically installed Connect.

Asconfigurator: The Aspera Configuration Tool

If you are unfamiliar with the XML formatting that is required for your Aspera server's configuration file, you can edit your configuration with confidence by using **asconfigurator**. These commands ensure that the XML structure is correctly maintained when you add or change settings.

Tap into the Aspera Ecosystem

If you have various data storage systems and internal and external customers who need access to the content in that storage, HSTS can be incorporated into a scalable Aspera data transfer ecosystem that meets your needs. Your Aspera server can be monitored and managed by IBM Aspera Console, and added as a node to IBM Aspera Faspex, IBM Aspera Shares, IBM Aspera on Cloud, and IBM Aspera Application for Microsoft SharePoint.

The Aspera Client transfer tools

Your installation includes the following transfer tools, some of which require an additional license for activation.

The FASP transfer engines: ascp and ascp4

These command line tools enable you to run transfers to any server to which you have access, and to customize the transfers (within the parameters set by the server). They are scriptable, supporting unattended data transfer and custom pre- and post-transfer file processing.

Hot Folders: Automatic data transfer in the GUI

Sending or receiving files can be even easier and faster by using Hot Folders. Available only on Windows, you can set up a Hot Folder to watch for and automatically transfer any new files that are added to that folder. Automatically send files to a server as they are added to a folder on your own desktop, or receive files as they are added to a folder on the server. Transfers use Ascp and are easily managed from the GUI.

Watch Folders: Automatic content delivery at any scale

Using `asperawatchd` and Watch Folders creates a powerful, efficient file system monitoring and automatic transfer tool that can comfortably handle millions of files and "growing" sources. Automatically transfer files as they are added to a source folder. With a REST API interface, you have full programmatic control for custom, automatic transfer processing.

Watch Folders offer the same transfer and bandwidth management options as **ascp**, and can be monitored and managed through Console. Watch Folders are enabled in your HSTS or HSTE.

IBM Aspera Sync: Directory synchronization at the speed of FASP

When everyone needs to see the same files or you need to be sure that every file is replicated, Aspera Sync provides a high-speed tool to do it. Unique among Aspera transfer tools, Aspera Sync supports bidirectional synchronization for optimum collaboration and consistency between computers.

Aspera Sync uses efficient file system monitoring and change detection to minimize redundant data transfer and to reduce database storage requirements. Aspera Sync offers the same transfer and bandwidth management options as **ascp**, and can be monitored and managed through Console.

Aspera Sync is installed with your HSTS and HSTE, but both the client and server require a Aspera Sync-enabled license.

Getting started as a transfer client

Aspera transfer clients connect to a remote Aspera transfer server and request a transfer with that server. Your Aspera application can be used as a client to initiate transfers with Aspera servers, as described in the following steps.

1. Review the system requirements and install Desktop Client.
See [“Requirements” on page 4](#) and [“Installing Desktop Client” on page 4](#).
2. Configure the firewall, if it is not already configured.
See [“Configuring the firewall” on page 4](#).
3. Test a locally initiated transfer to the Aspera demonstration server to confirm that your installation and firewall configuration are operational.
For instructions, see [“Testing a transfer” on page 5](#). This provides a simple walk-through of how to set up a connection with a server and transfer.
4. If you need to authenticate to the remote server with an SSH key, create an SSH key and send the public key to the server admin.
For instructions on creating an SSH key, see [“Creating SSH keys ” on page 53](#).
5. To run transfers from the command line, review the instructions for the Aspera command line client. Your Aspera product comes with two command line clients: **ascp** and A4. They are similar but have different capabilities. For a comparison, see [“Comparison of ascp and ascp4 options” on page 57](#).
 - For more information about **ascp**, see [“Ascp command reference” on page 23](#) and [“Ascp general examples” on page 38](#).
 - For more information about A4, see [“Ascp4 command reference” on page 63](#) and [“Ascp4 examples” on page 72](#).

Once you confirm that you can transfer with your server, your basic setup is complete.

Installation and upgrades

Before you install the current release, review the following information about hardware and software requirements, system preparation for upgrades or downgrades, installation instructions, and product security configuration.

Requirements

System requirements for Desktop Client.

- Product-specific Aspera license file .
- Ubuntu 16.04.2 LTS; Kernel: Linux 4.4.0-116-generic; architecture: ppc64-le.

Note: Your OS version must support Little Endian (LE) and it must run on IBM Power hardware that supports LE.

- SSH Server. Make sure that your operating system is up to date for the latest version of SSH Server.

Installing Desktop Client

To install Desktop Client, log in to your computer with root permissions.

1. Download the HSTS installer from [Fix Central](#).
2. For product upgrades, ensure that you prepared your system to upgrade to a newer version.
Although the installer performs your upgrade automatically, you must complete the tasks that are described in
3. Run the installer.

Run the following commands with the admin permissions. Replace the product version with that of your package.

OS	Commands
Red Hat, zLinux, CentOS	<pre>\$ rpm -Uvh /path_to_installer/aspera-desktopclient-version.rpm</pre> <p>Note: If your Linux OS is a minimal clean system, ensure that all the required dependencies are installed with your Aspera application by installing the product with a yum install:</p> <pre>\$ yum --nogpgcheck install /path_to_installer/aspera-desktopclient-version.rpm</pre>

Configuring the firewall

Desktop Client requires access through specific ports. If you cannot establish the connection, review your local corporate firewall settings and remove the port restrictions.

The following is basic information for configuring your firewall to allow Aspera file transfers. The outbound TCP port for SSH might differ depending on your organization's unique network settings. Although TCP/33001 is the default setting, refer to your IT Department for questions that are related to which SSH ports are open for file transfer. Consult your operating system's documentation for instructions on configuring your firewall. If your client host is behind a firewall that does not allow outbound connections, you need to allow the following ports:

- **Outbound TCP/33001:** Allow outbound connections from the Aspera client on the TCP port (TCP/33001 by default, when connecting to a Windows server, or on another nondefault port for other server operating systems).

- **Outbound UDP/33001:** Allow outbound connections from the Aspera client on the FASP UDP port (33001, by default).
- **Local firewall:** If you have a local firewall on the client (such as iptables), verify that it is not blocking your SSH and FASP transfer ports (such as TCP/UDP 33001).

Testing a transfer

To make sure that the software is working properly, set up a connection with a server and test downloads and uploads.

1. Download test files from the server.

Use the `ascp` command to download, press `y` to accept the server's key, and enter your password when prompted.

For example,

```
# ascp -T xeno@my_demo.example.com:test-dir-large/100MB /tmp/
```

The transfer command is based on the following settings:

Item	Value
Server address	my_demo.example.com
Login account	xeno
Test file	/test-dir-large/100MB
Download location	/tmp/
Transfer settings	Fair transfer policy, target rate 10 M, minimum rate 1 M, encryption disabled.

You must see a message similar to the following:

```
100MB          28%  28MB  2.2Gb/s  01:02 ETA
```

This message provides the following information:

Item	Description
100 MB	The name of the file that is being transferred.
28%	The percentage completed.
28 MB	The amount transferred.
2.2 Gbps	The current transfer rate.
01:02 ETA	The estimated time the transfer completes.

2. Upload test files to the demo server.

Run the command to upload the same file (100MB) back to the demo server, to its `/Upload` directory. Enter your password when prompted.

For example,

```
# ascp -T /tmp/100MB xeno@my_demo.example.com:Upload/
```

Uninstalling

Desktop Client can be uninstalled without removing existing configuration files.

1. Close or stop the following applications and services:

- FASP transfers.
 - SSH connections.
2. Uninstall by running the following command:

Platform	Command
Red Hat, CentOS	<code># rpm -e aspera-scp-client</code>
Debian	<code># dpkg -P aspera-scp-client</code>

Note: This process does not remove Aspera configuration files. If you reinstall an Aspera product, these configuration files are applied to the new installation.

Transfer files in the GUI

Use the Desktop Client GUI to create connections to Aspera servers, configure transfer settings, set up transfer notifications, and start, stop, pause, and schedule transfers.

Global bandwidth settings

Aspera FASP transport has no theoretical throughput limit. In addition to network capacity, transfer rates can be limited by user-configured rate settings and the resources of the local and remote machines. You can configure bandwidth usage limits and the number of concurrent FASP transfers that are allowed by Desktop Client.

1. Start the application with administrator privileges and click **Tools > Global Preferences**.
2. Click **Transfers**.
3. To limit total bandwidth usage by FASP uploads and downloads, edit **System-Wide Settings**.
System-wide settings set the aggregated bandwidth cap for all FASP transfers on this computer.
To override the default bandwidth limits, under **System-Wide Settings** select the boxes next to **Limit Download Bandwidth** and **Limit Upload Bandwidth** and enter new values in the fields. The global settings for download and upload bandwidth limits cannot be reset by nonadmin users. However, users can view the global limit from the **Preferences > Transfers** dialog.
4. To set default target rates for all users, edit **Default Target Rate**.
Nonadmin users can adjust their personal default target rates above or below the global default value.
5. To limit the number of active transfers, edit **Maximum Active Transfers**.
Nonadmin users can adjust their personal maximum active transfers above or below the global default value.
6. Click **OK** to activate your changes.

Enabling a transfer proxy or HTTP proxy

If for network security reasons, you are behind a transfer proxy server, you can enable the proxy for Aspera file transfers. If you have admin privileges, you can enable transfer proxies for all users by setting global preferences. If you are a nonadmin user, you can override global transfer-proxy settings for your own account, including enabling or disabling the feature. By default, proxy is disabled.

Open the proxy configuration dialog by clicking **Preferences > Proxy**.

Clicking **Preferences** opens the user-account proxy settings. If you have permission, you can click **Global Preferences** to access those settings.

Configuring global transfer and HTTP proxy settings

You must have admin privileges to set global preferences.

To enable a transfer proxy:

1. Go to **Global Preferences > Proxy**.
2. Select **Enable transfer proxy**.
3. Enter the proxy server's hostname or IP address and port number.
4. Select **Secure** if your proxy server allows secure connections.
5. Enter your username and password to authenticate with your proxy server.

To enable HTTP proxy:

1. Go to **Global Preferences > Proxy**.
2. Select **Enable HTTP proxy**.
3. Enter the HTTP proxy's hostname or IP address and port number.
4. If your HTTP proxy requires authentication, select **Authenticated** and enter the username and password for your HTTP proxy.

To clear all settings, click **Restore System Defaults**.

User Proxy Settings

To override the global settings, edit the proxy settings for your account. Click **Preferences > Proxy**. The values are those that you inherited from the global proxy settings.

To configure user transfer proxy settings:

1. Select or clear **Enable transfer proxy** to enable or disable transfer proxy.
2. Enter the proxy server's hostname or IP address and port number.
3. Select **Secure** if your proxy server allows secure connections.
4. Enter your username and password to authenticate with your proxy server.

To configure user HTTP proxy settings:

1. Select or clear **Enable HTTP proxy**.
2. Enter the HTTP proxy's hostname or IP address and port number.
3. If your HTTP proxy requires authentication, select **Authenticated** and enter the username and password for your HTTP proxy.

To revert all user settings to the global values, click **Restore Defaults**.

Adding and editing connections

To transfer with HSTS, HSTE, IBM Aspera Shares, or IBM Aspera on Cloud transfer service (AoC), add it as a connection in the **Connection Manager**. The following instructions describe how to create and configure a connection and edit or delete connections.

To connect with cloud storage, you must meet the following prerequisites:

- You have permissions to access the cloud storage and the necessary authentication information.
- To transfer files with an S3 storage device that uses an S3 direct connection, the user must have a restriction rather than a docroot set on the server.


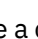
Once you create connections, you can export and import connection lists. For instructions, see [“Exporting and importing connections”](#) on page 14.

To create a new connection:

1. Start the application.

2. To open the **Connection Manager**, click **Connections**.

3. Click **+** to create a new connection.

Click  to duplicate a selected connection (to copy all information into a new profile) and  to delete a connection profile.

4. Configure the connection name, if wanted.

By default, connections are named **username@host**.

To name or rename a connection, click the connection name and enter the new name in the pop-up window. Click **OK** to save your changes.

5. Configure the required settings for the connection.

On the **Connection** tab, enter the following information. In most cases, only **Host**, **User**, and **Authentication** are required.

Connection Option	Description
Host	The server's address, such as 192.168.1.10 or companyname.com. For Shares, Node API, or AoCts connections, enter the URL and port for asperanoded, such as <code>https://ats-aws-us-west-2.aspera.io:443</code> .
User	The transfer user's username, the Shares user, Node API credentials, or an access key ID.
Authentication	<p>The authentication method. Select Password to authenticate with the transfer user's password, the Shares user's password, the Node API user password, or an access key secret (such as for AoCts or ATC Manager).</p> <p>Select Public Key to authenticate with the transfer user's public SSH key. For more information, see "Creating SSH keys in the GUI" on page 14.</p>
Storage Type	<p>The default option is local storage. Use this option to connect to:</p> <ul style="list-style-type: none">• On-premises servers.• AoCts.• Cloud-based servers when the transfer user has the storage credentials that are configured in their docroot on the server. <p>When the server is in the cloud but the storage credentials are not configured in the transfer user's docroot, use the drop-down menu to select the object storage type and enter credentials.</p> <p>Supported object storages include the following:</p> <ul style="list-style-type: none">• Akamai NetStorage• Amazon S3: Requires your Access ID, Secret Access Key, and bucket path. The local machine must be reasonably time-synchronized in order to communicate with the Amazon servers. You can also select the Advanced button to modify the following settings:<ul style="list-style-type: none">– Host: Amazon S3 hostname (default: s3.amazonaws.com).– Port: Default is port 443.– HTTPS connection for file browsing: Enable for secure browsing.– Server-side file encryption: Enable for AES256 encryption.– Reduced redundancy storage class: Assign objects to the reduced redundancy storage class (durability of 99.99%).• Google Storage: Requires your project number and bucket path.

Connection Option	Description
	<ul style="list-style-type: none"> • Limelight: Requires your account, username, and password. • Windows Azure: Requires your storage account and access key. Azure storage is set to use the Azure block blob REST API by default. To specify the REST API for page blobs, enter your account credentials then click Advanced. Select PAGE from the drop-down menu next to Api and click OK. • Windows Azure SAS: Requires your Shared URL. • Azure Files: Requires your File service endpoint and password.

6. Configure other connection settings, if needed.

On the **Connection** tab, configure nondefault connection settings by editing any of the following settings:

Connection Option	Description
Target Directory (or Bucket Path for most object storage)	The default directory when connecting to this computer. When left blank, browsing the remote host brings up either the user's docroot or the last-visited folder. When a path is set, the connection opens to the exact directory.
Advanced Settings > SSH Port (TCP)	The TCP port for SSH connections. Default: 33001. If the application cannot connect on 33001, it automatically attempts a connection on port 22. If the connection on 22 succeeds, the setting is updated to 22.
Advanced Settings > FASP Port (UDP)	The UDP port for FASP transfers. Default: 33001.
Advanced Settings > Connection Timeout	Time out the connection attempt after the specified time.
Test Connection	Click to test the connection to the remote server with the settings you configured.

7. Configure the connection's transfer settings, if needed.

On the **Transfer** tab, configure nondefault transfer settings by editing any of the following settings:

Transfer Option	Description
Transfer Name	Select from the following options: Automatically generate allows the user interface to generate the transfer name; Automatically generate and add prefix uses auto-generated name with a customizable prefix; Specify uses the user-specified name.
Policy	Select the FASP transfer policy. <ul style="list-style-type: none"> • high - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, the transfer rate is twice as fast as a fair-policy transfer. The high policy requires maximum (target) and minimum transfer rates. • fair - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, bandwidth is shared fairly by transferring at an even rate. The fair policy requires maximum (target) and minimum transfer rates.

Transfer Option	Description
	<ul style="list-style-type: none"> • low - Adjust the transfer rate to use the available bandwidth up to the maximum rate. Similar to fair mode, but less aggressive when the bandwidth is shared with other network traffic. When congestion occurs, the transfer rate is reduced to the minimum rate until other traffic decreases. • fixed - Attempt to transfer at the specified target rate, regardless of network or storage capacity. This can decrease transfer performance and cause problems on the target storage. Use the fixed policy only for specific contexts, such as bandwidth testing, otherwise, avoid the use of this policy. The fixed policy requires a maximum target rate. • aggressiveness - The aggressiveness of transfers that are authorized by this access key in claiming available bandwidth. Value can be 0.00-1.00. For example, these values correspond to the policy option where a policy of high approximates to aggressiveness of 0.75, fair to 0.50 and low to 0.25. Aggressiveness can be used if you need to fine-tune the transfer policy.
Speed	Select Override default transfer rate settings to specify the transfer rate. The target rate is constrained by the global bandwidth settings; for more information, see “Global bandwidth settings” on page 6 .
Retry	<p>Select to automatically retry the transfer after a recoverable failure for a set amount of time, in seconds, minutes, or hours. You might set the initial and maximum retry intervals by clicking the More Options button.</p> <ul style="list-style-type: none"> • Initial interval: The first retry waits for the initial interval. Input in seconds, minutes, or hours. • Maximum interval: After the initial interval, the next interval doubles until the maximum interval is met, and then stops retrying after the retry time is reached. Input in seconds, minutes, or hours. <p>For example, if the retry is set for 180 seconds, the initial interval is 10 seconds, and the maximum interval is 60 seconds, then the transfer will retry at 10, 30, 70, 130, and 180 seconds after the first try, such that the interval progression is 10, 20, 40, 60, 60, and 50 seconds. The last interval is not the maximum because the retry period ends with a final retry.</p> <p>As another example, if the retry is set for 600 seconds, the initial interval is 30 seconds, and the maximum interval is 120 seconds, then the transfer will retry at 30, 90, 210, 330, 450, 570, and 600 seconds after the first try, such that the interval progression is 30, 60, 120, 120, 120, 120, and 30 seconds. Again, the last interval is not the maximum because the retry period ends with a final retry.</p>
Show Advanced Settings	<p>Click Show Advanced Settings to edit the following options:</p> <ul style="list-style-type: none"> • Specify FASP datagram size (MTU): By default, the detected path MTU is used. Select to specify a value in the range 296 - 10000 bytes. • Disable calculation of source files size before transferring: Select to turn off job size calculation on the client side, if allowed by the server.

8. Configure tracking and email notifications, if needed.

On the **Tracking Tab**, configure nondefault transfer settings by editing any of the following settings:

Tracking Option	Description
Generate delivery confirmation receipt	Select to create a delivery receipt file in the specified location.

Tracking Option	Description
Send email notifications	Send email notifications based on specified events (start, complete, and error). For more information, see “Using transfer notifications” on page 22.

9. Configure filters to automatically exclude certain files from transfers with this connection, if needed.

On the **Filters** tab, click **Add** and enter the pattern to exclude files or directories with the specified pattern in the transfer. The exclude pattern is compared with the whole path, not just the file name or directory name. Two special symbols can be used in the setting of patterns:

Filter Symbol	Name	Description
*	Asterisk	Represents zero to many characters in a string, for example *.tmp matches .tmp and abcde.tmp.
?	Question mark	Represents one character, for example t?p matches tmp but not temp.

For more information on filter rule syntax, see [“Using filters to include and exclude files”](#) on page 45.

10. Configure security settings, if needed.

On the **Security** tab, configure nondefault transfer settings by editing any of the following settings:

Security Option	Description						
Encryption	<p>Select the encryption cipher. Aspera supports three sizes of AES cipher keys (128, 192, and 256 bits) and supports two encryption modes, Cipher Feedback mode (CFB) and Galois Counter Mode (GCM). The GCM mode encrypts data faster and increases transfer speeds compared to the CFB mode, but the server must support and permit it.</p> <p>Cipher rules</p> <p>The encryption cipher that you are allowed to use depends on the server configuration and the version of the client and server:</p> <ul style="list-style-type: none"> • When you request a cipher key that is shorter than the cipher key that is configured on the server, the transfer is automatically upgraded to the server configuration. For example, when the server setting is AES-192 and you request AES-128, the server enforces AES-192. • When the server requires GCM, you must use GCM (requires version 3.9.0 or newer) or the transfer fails. • When you request GCM and the server is older than 3.8.1 or explicitly requires CFB, the transfer fails. • When the server setting is any, you can use any encryption cipher. The only exception is when the server is 3.8.1 or older and does not support GCM mode; in this case, you cannot request GCM mode encryption. • When the server setting is none, you must use none. Transfer requests that specify an encryption cipher are refused by the server. <p>Cipher Values</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> <th>Support</th> </tr> </thead> <tbody> <tr> <td>AES-128 AES-192 AES-256</td> <td>Use the GCM or CFB encryption mode, depending on the server configuration</td> <td>All client and server versions.</td> </tr> </tbody> </table>	Value	Description	Support	AES-128 AES-192 AES-256	Use the GCM or CFB encryption mode, depending on the server configuration	All client and server versions.
Value	Description	Support					
AES-128 AES-192 AES-256	Use the GCM or CFB encryption mode, depending on the server configuration	All client and server versions.					

Security Option	Description				
	Value	Description			Support
		and version (see cipher negotiation matrix).			
	AES-128-CFB AES-192-CFB AES-256-CFB	Use the CFB encryption mode.			Clients version 3.9.0 and newer, all server versions.
	AES-128-GCM AES-192-GCM AES-256-GCM	Use the GCM encryption mode.			Clients and servers version 3.9.0 and newer.
	NONE	Do not encrypt data in transit. Aspera strongly recommends against using this setting.			All client and server versions.
	Client/Server Cipher negotiation The following table shows which encryption mode is used depending on the server and client versions and settings:				
		Server AES-XXX-GCM	Server AES-XXX-CFB	Server AES-XXX	Server AES-XXX
	Client AES-XXX-GCM	GCM	Server refuses transfer	GCM	Server refuses transfer
	Client AES-XXX-CFB	Server refuses transfer	CFB	CFB	CFB
	Client AES-XXX	GCM	CFB	CFB	CFB
	Client AES-XXX	Server refuses transfer	CFB	CFB	CFB
Content Protection	<p>Select Encrypt uploaded files with a password to encrypt the uploaded files with the specified password (client-side encryption at rest). The protected file has the extension <code>.aspera-env</code> appended to the file name. Anyone downloading the file must have the password to decrypt it.</p> <p>Select Decrypt password-protected files downloaded to prompt for the decryption password when downloading encrypted files.</p> <p>For more information about client-side encryption at rest, see “Client-Side Encryption-at-Rest (EAR)” on page 56.</p>				

11. Configure file handling, if needed.

On the **File Handling** tab, configure non-default transfer settings by editing any of the following settings:

File Handling Option	Description
Resume	<p>Select Resume incomplete files to enable the resume feature. Select the file comparison method from the When checking files for differences drop-down menu:</p> <ul style="list-style-type: none"> • Compare file attributes - Compares the sizes of the existing and original file. If they are the same, then the transfer resumes, otherwise the original file is transferred again. • Compare sparse file checksums - Performs a sparse checksum on the existing file and resumes the transfer if the file matches the original, otherwise the original file is transferred again. (Default) • Compare full file checksums - Performs a full checksum on the existing file and resumes the transfer if the file matches the original, otherwise the original file is transferred again. <p>Under When a complete file already exists at the destination, select an overwrite rule when the same file exists at the destination. By default, files on the destination are overwritten if different from the source.</p>
File Attributes	<ul style="list-style-type: none"> • Select Preserve Access Time to set the access time of the destination file to the same value as that of the source file. • Select Preserve Modification Time to set the modification time of the destination file to the same value as that of the source file. • Select Preserve Source Access Time to keep the access time of the source file the same as its value before the transfer. <p>Note: Access, modification, and source access times cannot be preserved for node and Shares connections that are using cloud storage.</p>
Source Handling	<p>Select Automatically delete source files after transfer to delete the files that transferred successfully from the source.</p> <p>Select Automatically move uploaded source files to a directory after transfer and specify the location on the source machine to which they must be moved. Only a path to an existing location on the client can be specified.</p> <p>Select Delete empty source subdirectories to remove empty subdirectories from the source once the files that they contain transfer successfully. This option is usually used to clean up the Hot Folder when source files are moved or deleted after transfer.</p>

12. Click **OK** to save your changes.

Changes are not saved until you click **OK**. Selecting **Cancel** discards any unsaved changes that are made in the Connection Manager, including the addition and removal of connections.

13. Connect to the remote host.

Double-click the connection name, or select it and click **Connect**.

Editing and Deleting Connections

Click **Connections** and select the connection that you want to edit or delete. Edit the settings or click **-** to delete the connection. Deleting connections cannot be undone. When in doubt, export the connections as a backup before deleting a connection.

Exporting and importing connections

Connections, and optionally their passwords, can be exported and imported as a text file, and the text file can be password protected.

Usage notes:

- If you are exporting a connection that uses SSH key authentication, back up the keys manually and import separately. For instructions, see [“Creating SSH keys in the GUI” on page 14](#).
- A shared connection that is exported or imported by a nonadministrator is imported as a regular connection (not as a shared connection).
- Email templates are not exported with the connection.

Export connections

1. Right-click the remote server panel and click **Export**.
2. Enter the following information:
 - **Destination:** Enter or browse to the location on your computer where to save the file.
 - **Options:** The passwords that are associated with your connections can be exported. Select if you do not want to export passwords, export passwords without obscuring the passwords (**Export passwords in clear**), or export encrypted passwords (**Encrypt passwords**).
 - **Password:** Required if **Encrypt passwords** is selected. When the connections are imported, use the password to decrypt the connection passwords.
3. Click **OK** to export your connection information to the file.

Import connections

1. Right-click the remote server panel and select **Import**.
2. Enter the following information:
 - **Source file:** The file with the exported connections.
 - **Options:** Select the appropriate option, depending on how the connections were exported.
 - **Password:** If you select the **Passwords are encrypted** option, enter the password that was set when the connections were exported.
3. Click **OK** to import the connection information.
4. Before deleting the source file, confirm that the import process was successful by testing your connections.

Creating SSH keys in the GUI

Public key authentication (SSH Key) is a more secure alternative to password authentication that allows users to avoid entering or storing a password, or sending it over the network. The client creates an SSH key pair (a public key and a private key) and then sends the public key to the server's administrator. Once the admin configures the server with the client's public key, the client can authenticate connections to the server with their private key.

You can use the application GUI to generate key pairs and to import existing key pairs. You can also generate key pairs by using the command line; for instructions, see [“Creating SSH keys ” on page 53](#).

1. Start the application.
2. In the menu bar, click **Tools > Manage Keys**.
3. In the SSH Keys dialog, click **+** to create a new key pair.

The SSH Keys dialog is also available from the **Connection** tab in the Connection Manager. When you select **Public Key** for authentication, the **Manage Keys** button appears; clicking it opens the SSH Keys dialog.

4. In the **New SSH Key Pair** window, enter the requested information.

Field	Description
Identity	Name your key pair, such as with your username.
Passphrase	Set a passphrase on your SSH key, which is prompted for whenever it needs to use the key. If you don't want the user to be prompted for passphrase when logging in, leave this field blank.
Type	Select RSA (default) or ECDSA key.
Access	When sharing a connection with public key authentication, or a connection that is has a Hot Folder (on Windows machines), this option must be checked.

5. Click **OK** to create the key.

The public key is displayed in the window and you can copy it to a clipboard or export it to a file that is easy to locate. The key is automatically saved as a file named `id_key_type.pub` in the following location. In the example, the public key file name is `id_rsa.pub`:

```
/home/username/.ssh/id_rsa.pub
```

6. Distribute the public key.

Provide the public key file to your server administrator so that it can be set up for your server connection.


To copy or export the public key, select the key in the **SSH Keys** window, click **Copy Public Key to Clipboard**, and paste the string into an email to send to the server administrator, or click **Export to File** and save the public key as a file.


7. Set up connections by using public key authentication.

Note: Your public key must be configured on the server before you can connect with it.

- a) Click **Connections** to open the Connection Manager.
- b) Select the connection for which you want to set up the key.
- c) In the **Connection** tab, select the **Public Key** Authentication option and select the key from the drop-down menu.

Importing keys:

To import keys created outside the GUI, go to **Tools > Manage Keys** to open the **SSH Keys** dialog. Click the  button in the upper-left corner of the dialog to open a file browser. You can import the key pair by selecting either the private key or the public key; this copies both keys into the user's `.ssh` directory. You cannot import a key pair if a key pair with the same identity exists in the `.ssh` directory.

Imported key pairs can be shared with other users. In the SSH Keys dialog, select a key and click the  button to open the **Edit SSH Key Pair** dialog. Select **Access** to allow shared connections to use this key. Shared keys are moved to the Aspera etc directory.

Transferring content

The GUI provides an easy, intuitive way to transfer content between the local computer and a remote host.

Note: Do not use the following characters in file or folder names:

```
/ \ " : ' ? > < & * |
```

They can produce unpredictable transfer behavior.

1. Create a connection.

For instructions, see [“Adding and editing connections” on page 7](#).

2. Select the remote server under **Connection Name**.

3. For uploads, if the target directory is correct, then you can select the content to upload from the local file tree and either drag-and-drop the content into the connection pane, or click the upload arrow. If you want to browse the remote file system or download content from it, go on to the next step.

4. Connect to the remote server by either double-clicking the connection name, or select it and click **Connect**.

5. Select the content to transfer (from the local or remote file system) and do any of the following:

- Click the upload or download arrow.
- Drag and drop the files between the windows.
- Copy and paste the files between the windows.


6. Once a transfer is started, you can manage the transfer rate, transfer policy, and priority. For information, see [“Managing transfers” on page 16](#).


Managing transfers

The Desktop Client GUI enables the option to start, stop, and reorder transfers, and adjust the transfer rates and policies, and configure transfer preferences.

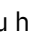

The transfers panel: Start, stop, and reorder transfers

Once the transfer starts, a progress bar appears in the **Transfers** panel. You can manage transfer behavior with the following actions:

Click  to start the selected transfer.

Click  to stop the selected transfer.

Click  to delete the selected transfer.

If you have multiple ongoing transfers, use the  and  to change the selected transfer's priority. The # field indicates the transfer's order in the queue.

The details view: Adjust transfer rates and policies of active transfers

The **Details** button provides additional oversight and control (if you have permission) over transfers. Select a transfer session from the **Transfers** panel and click **Details** to view details and adjust settings.

The **Details** display shows the following information:

Item	Name	Description
A	Details (tab)	Transfer details, including status (rate and ETA) and statistics (session size, files transferred versus total files to be transferred, average speed, time elapsed, RTT delay and average loss in percent).
B	Files (tab)	All files being transferred in this session, along with each files' size and transfer progress.
C	Transfer controls	Set the FASP transfer policy and transfer rate, if allowed. <ul style="list-style-type: none">• high - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, the transfer rate is twice as fast as a fair-policy transfer. The high policy requires maximum (target) and minimum transfer rates.

Item	Name	Description
		<ul style="list-style-type: none"> • fair - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, bandwidth is shared fairly by transferring at an even rate. The fair policy requires maximum (target) and minimum transfer rates. • low - Adjust the transfer rate to use the available bandwidth up to the maximum rate. Similar to fair mode, but less aggressive when the bandwidth is shared with other network traffic. When congestion occurs, the transfer rate is reduced to the minimum rate until other traffic decreases. • fixed - Attempt to transfer at the specified target rate, regardless of network or storage capacity. This can decrease transfer performance and cause problems on the target storage. Use the fixed policy only for specific contexts, such as bandwidth testing, otherwise, avoid the use of this policy. The fixed policy requires a maximum target rate. • aggressiveness - The aggressiveness of transfers that are authorized by this access key in claiming available bandwidth. Value can be 0.00-1.00. For example, these values correspond to the policy option where a policy of high approximates to aggressiveness of 0.75, fair to 0.50 and low to 0.25. Aggressiveness can be used if you need to fine-tune the transfer policy. <p>Important: If <code>--policy</code> is not set, ascp uses the server-side policy setting (fair by default).</p>
D	Transfer Monitor	The transfer graph. Use the sliders on the vertical axis to adjust the transfer rate up or down (if allowed).

Configuring transfer preferences

If you have administrator privileges, you can set the target transfer rate for all users from the **Global Preferences** dialog. As an individual user, you can override the global settings from **My Preferences**.

To update these settings, go to **Tools > Global Preferences** or **Tools > Preferences**. You can also open **My Preferences** from the **Preferences** button in the upper-right corner of the application's main window; from there you can also reach the **Global Preferences** dialog by clicking **Global Preferences**.

The following options are available under the **Transfers** tab:

Item	Description
Global Bandwidth Limits	The aggregated bandwidth cap for all FASP transfers on this computer.
Default Target Rate	The initial download and upload rates for all transfers.
Maximum Active Transfers	The maximum number of concurrent upload transfers and download transfers.

For information about **Email** settings, see [“Configuring transfer notifications”](#) on page 18.

Scheduling and customizing transfers in advanced mode

You can start a transfer in advanced mode to set per-session transfer options such as filters, security, which overrides the default transfer settings. You can also schedule the transfer as a one-time transfer or recurring.

1. In the GUI, right-click a file or folder to open the context menu and select **Upload** (in the client panel) or **Download** (in the server panel).

2. Configure the transfer settings, as needed.

The advanced transfer configuration options except **Scheduling** are identical to those in the **Connection Manager**. For more information about these tabs, see [“Adding and editing connections” on page 7](#). The **Scheduling** tab is described in the next step.

Tab	Description
Transfer	The transfer session-related options, such as the transfer speed and retry rules.
Tracking	Options for tracking the transfer session, including the confirmation receipt and the email notifications.
Filters	Create filters to skip or include files that match certain patterns.
Security	Enable the transfer encryption and the content protection.
File Handling	Set up resume rule, preserve transferred file attributes, and remove or move source files.
Scheduling	Schedule the transfer.

3. Schedule the transfer.

To enable transfer scheduling, select **Schedule this transfer**.

The following scheduling options are available in the **Transfer repeats** drop-down menu:

Does not repeat

Set the time and date for a single transfer.

Daily

Set the time for a daily transfer. For **End repeat**, select **Never** to continue daily transfers indefinitely, or **On** and set an end date and time.

Monday-Friday

Set the time for a daily transfer only on weekdays. For **End repeat**, select **Never** to continue daily transfers indefinitely, or **On** and set an end date and time.

Weekly

Select the time and days of the week for a repeating transfer. For **End repeat**, select **Never** to continue weekly transfers indefinitely, or **On** and set an end date and time.

Periodically

Set the frequency to repeat the transfer, in minutes.

4. Click **Transfer** to submit the scheduled transfer.

The transfer is then listed under the Transfers tab, along with an icon (📁) under the # column.

5. To modify the transfer, right-click the row and click **Edit**

Configuring transfer notifications

Transfer notification emails are triggered by three transfer session events: start, completion, and error. Transfer notification emails can be enabled and configured globally and by each user. The emails are generated from mail templates that can be customized.


Note: The GUI must remain open for transfer notification emails to send. Closing the GUI stops email notifications.

Enable email notifications

1. Start Desktop Client.
2. To configure global email notification settings:
 - a) Click **Tools > Global Preferences**.

- b) Click **Mail**.
 - c) To turn on email notifications for all users, select **Enable email notifications**.
Enter the email address from which the notifications are sent in the **From Address** field and enter the outgoing email server hostname in the **Host** field. The other values are optional.
 - d) To test your settings, click **Send test email**, which sends a test message to the **From Address**.
3. Set your personal mail preferences.
Personal mail preferences override global settings.
- a) Click **Preferences**.
 - b) Click **Mail** and edit the inherited global default values.
To restore your settings to global values, click **Restore Defaults**.

Configure email templates

1. Open the **Mail Templates** window by clicking **Tools > Mail Templates**.
2. To create a new template, click **+**, or to edit an existing template, select the template and click .
3. For new templates, name the template and select its base template.
Select an existing template from the **Based On** menu. Click **OK**.
4. Edit the template text.
The **Edit Template** window has four fields:

Field	Description
Name	The template name.
HTML	The HTML mail body. Click Insert Image to insert an image into the template. The image is copied to the template directory. Preview the template by clicking Preview .
Text	The plain text mail body. Preview the template by clicking Preview .
Access	Select Share this template with all users on this computer to allow other system users to access this template.

The mail template supports MIME (Multipurpose Internet Mail Extensions) multipart messages. You can edit both the HTML and plain text versions of the mail body. The templates are rendered by Apache Velocity (for more information, see the Apache Velocity User Guide at <http://velocity.apache.org/>). Templates use two predefined variables:

- `$formatter` - Contains some utility methods
- `$notifications` - Holds the transfer notifications

To iterate over notifications, use a `foreach` loop. A `foreach` loop generates content for each iteration of the loop. In the following example, a local `$event` variable is declared for use within the `foreach` loop:

```
#foreach ($event in $notifications.getEvents())
...
#end
```

To generate content only under specific conditions, use a conditional statement. To construct a conditional statement, use `#if`, `#else`, and `#end`, with the following syntax:

```
#if
...
#else
...
#end
```

All conditional statements are categorized in four parts: the conditional (what must occur to trigger the action), session information (what action is triggered), time, and statistics.

Conditional

Use conditional tests in an `if` statement. For example,

```
#if ($event.isFailed())
...
#end
```

Statement	Description
<code>\$event.isStarted()</code>	If the transfer session is started.
<code>\$event.isCompleted()</code>	If the transfer session is completed.
<code>\$event.isEnded()</code>	If the transfer session is ended.
<code>\$event.isFailed()</code>	If the transfer session is failed.

Session Information

Statement	Description
<code>\$event.getSourceHost()</code>	The source hostname (or host address if the hostname is not discoverable).
<code>\$event.getSourceHostAddress()</code>	The source host address.
<code>\$event.getSourcePaths()</code>	The source file path.
<code>\$event.getDestinationHost()</code>	The destination hostname (or host address if the hostname is not discoverable).
<code>\$event.getDestinationHostAddress()</code>	The destination host address.
<code>\$event.getDestinationPath()</code>	The destination file path.
<code>\$event.getInitiatingHost()</code>	The mail template supports MIME The session-initiating hostname (or host address if the hostname is not discoverable).
<code>\$event.getInitiatingHostAddress()</code>	The session-initiating host address.
<code>\$event.getId()</code>	The session ID.
<code>\$event.getName()</code>	The session name.
<code>\$event.getType().getDescription()</code>	The session state. Three outputs: STARTED, FAILED, and COMPLETED.
<code>\$event.getUser()</code>	The transfer login.
<code>\$event.GetFiles()</code>	The files that are being transferred. Use this statement in a <code>foreach</code>

Statement	Description
	<p>loop: (Any text after ## is a comment)</p> <pre>#foreach (\$file in \$event.GetFiles()) ## \$file is a new variable visible in this foreach loop. ## \$file holds the complete file path and file name. ## \$formatter.decodePath() is used to ensure a correct string decoding. \$formatter.decodePath(\$file) #end</pre> <p>Use the counter \$velocityCount in an if statement to limit the output file count. For example, to list only the first 10 files:</p> <pre>#foreach (\$file in \$event.GetFiles()) #if (\$velocityCount > 10) #break #end \$file #end</pre>
\$event.getMessage()	The message that is entered in the email Message field.
\$event.getError()	The error message.

Time

Statement	Description
\$formatter.date(<i>var</i>, "<i>lang</i>", "<i>format</i>")	<p>Formatting the date and time output. Enter three values in the parenthesis:</p> <ul style="list-style-type: none"> <i>var</i> is either \$event.getStartTime() or \$event.getEndTime() <i>lang</i> is an abbreviated language name; for example, en for English. <i>format</i> is the display format. Use these symbols: <ul style="list-style-type: none"> – yyyy - The year; for example, 2022. – MM - Month of the year; for example, 03. – dd - Day of the month; for example, 26. – HH - Hour of the day; for example, 16. – mm - Minute. – ss - Second. – z - Time zone. – EEE - The abbreviated weekday name; for example, Fri. <p>For example,</p> <pre>"EEE, yyyy-MM-dd HH:mm:ss z"</pre> <p>shows Fri, 2010-03-26 16:19:01 PST.</p>
\$event.getStartTime()	The session start time.

Statement	Description
<code>\$event.getEndTime()</code>	The session end time.

Statistics


Statement	Description
<code>\$event.getSourceFileCount()</code>	The number of source files.
<code>\$event.getCompletedFileCount()</code>	The number of files that successfully transferred.
<code>\$event.getFailedFileCount()</code>	The number of files that failed to transfer.
<code>\$event.getAverageRatePercentage()</code>	The average transfer rate in bps. Enclose this statement with <code>\$formatter.formatRate()</code> to simplify the output.
<code>\$event.getAverageLossPercentage()</code>	The average packet loss percentage.
<code>\$event.getSourceSizeB()</code>	The source file size. Enclose this statement with <code>\$formatter.toBestUnit()</code> to simplify the output.
<code>\$event.getTransferredB()</code>	The transferred file size. Enclose this statement with <code>\$formatter.toBestUnit()</code> to simplify the output.
<code>\$event.getWrittenB()</code>	The destination file size. Enclose this statement with <code>\$formatter.toBestUnit()</code> to simplify the output.

5. Click **OK** to save your changes.

Apply the notifications to a specific connection host or a transfer session. You can also customize the subject line of the notification emails. For details, see [“Using transfer notifications” on page 22](#).

Using transfer notifications

Transfer notifications can be emailed to a set list of recipients upon transfer start, complete, and error. The email templates can be fully customized. These instructions describe how to configure email notifications for all transfers to and from a specific connection.

1. Preview existing mail templates and create new ones, if needed.
 - a) Click **Tools > Mail Templates** to open the **Mail Template** window.
 - b) Select an existing template and click .
 - c) In the **Edit Template** window, click **Preview** to view the template's output example.
For instructions on how to create a new template or edit an existing one, see [“Configuring transfer notifications” on page 18](#).
2. Enable email notifications for connections.
 - a) Click **Connections** on the main page of the application, select the connection that you want to configure with email notifications, and go to the **Tracking** tab.
 - b) Select **Send email notifications**, and configure the following settings:

Item	Description
When	Check the events for which to send notifications.
Subject	Customize the subject line, which can use the same template fields as described in “Configuring transfer notifications” on page 18 .
To	Enter the recipients, comma-separated.
Template	Select a mail template.

Item	Description
Message	Enter a message to include in the notifications (optional).

c) Click **OK** to save your changes.

ascp: Transferring from the command line

Ascp is a scriptable FASP transfer binary that enables the transfer to and from Aspera transfer servers to which you have authentication credentials. Transfer settings are customizable and can include file manipulation on the source or destination, filtering of the source content, and client-side encryption-at-rest.

Ascp command reference

The **ascp** executable is a command line FASP transfer program. This reference describes ascp syntax, command options, and supported environment variables.

For examples of ascp commands, see the following topics:

- [“Ascp general examples” on page 38](#)
- [“Ascp file manipulation examples” on page 41](#)

Another command line FASP transfer program is Ascp4, which is optimized for transfers of many small files. It has many of the same capabilities as ascp, and also its own features. For more information, see [“Introduction to ascp4” on page 63](#) and [“Comparison of ascp and ascp4 options” on page 57](#).

Ascp syntax

```
ascp options [[username@]src_host:]source1[ source2 ...] [[username@]dest_host:]dest_path
```

username

The username of the Aspera transfer user can be specified as part of the source or destination, whichever is the remote server. It can also be specified with the **--user** option. If you do not specify a username for the transfer, the local username is authenticated by default.

Note: If you are authenticating on a Windows computer as a domain user, the transfer server strips the domain from the username. For example, Administrator is authenticated rather than DOMAIN\Administrator. For this reason, you must specify the domain explicitly.

src_host

The name or IP address of the computer where the files or directories to be transferred reside.

source

The file or directory to be transferred. Separate multiple arguments with spaces.

The *growing files* feature can be used with the *source* option to start transferring files to the target directory while they are still being written to the source directory.

Note: To use the growing files feature, the source file must be on a native file system. Growing files cannot be larger than 8 exabytes - 1 (9,223,372,036,854,775,807 bytes). However, the maximum file size of the file system overrides a setting that is larger.

Growing files use can also be configured statically with `aspera.conf`, see [aspera.conf - File system configuration](#). See also [“Ascp general examples” on page 38](#).

To use the growing files feature with ascp, the *source* parameter can be used with the following syntax:

```
source?grow=wait_time[&wait_start=[mtime | null_read]][&confirm_stop=[ true | false ]]
```

A file transfer is deemed to be complete when the time since last update to the source file reaches the *wait_time* value (in seconds). However, the time is only sampled when all currently available source

data was transferred. In other words, if more data arrives after the wait time has elapsed, but the transfer is still in progress, the additional data is still transferred.

grow

Enables the growing file feature and is used to set the wait time in seconds. The wait time is the amount of time that is allowed to pass before a file that is not changing is treated as complete. The `grow` element must be set to a nonnegative integer to define wait time. If it is set to a nonnumeric string, the default wait time of 10 seconds is used.

wait_start

Can be used to specify how time is calculated to determine whether a file is complete. If `mtime` is used, then the file modification time is used to calculate the wait time. If `null_read` is used, then the time of a file read that returns zero bytes is used. The default is `mtime`.

confirm_stop

Can be used to indicate when all the data was added to the source file and to prevent any additional wait time followed by a read of EOF.

Note that `confirm_stop` is ignored if `wait_start` is set to `null_read`.

To use `confirm_stop`, set it to `true` (the default value is `false`). Then, use an external program to adjust the source file `mtime` upon completion of writing data to the source file, by using the following criteria:

```
mtime < current_system_time - wait_time, mtime != 0
```

Any value for `mtime` that meets this criteria is acceptable to flag this condition except `mtime = 0`, which is used to flag a file error. You can, for example, use **touch 1**. If `mtime` is not adjusted before the timeout is reached, an error is generated.

An alternative method for defining the `wait_time` value is to use modifiers for powers of 1,024. However, the value must be less than $8 * 2^{60}$. The modifier must consist of an integer, and a unit specifier. The unit specifiers are:

- k or K for $1 * 1024$
- m or M for $1 * 1024 * 1024$
- g or G for $1 * 1024 * 1024 * 1024$

dest_host

The name or IP address of the computer where the source files or directories are to be transferred.

dest_path

The destination directory where the source files or directories are to be transferred.

- If the source is a single file, the destination can be a file name. However, if there are multiple source arguments, the destination must be a directory.
- To transfer to the transfer user's docroot, specify `.` as the destination.
- If the destination is a symbolic link, then the file or directory is written to the target of the symbolic link.

Specifying files, directories, and paths

- Specify paths on the remote computer relative to the transfer user's docroot. If the user has a restriction instead of a docroot, specify the full path, which must be allowed by the restriction.
- Avoid the following characters in file and directory names: `/ \ " : ' ? > < & * |`
- Specify paths with forward-slashes, regardless of the operating system.
- If directory or file arguments contain special characters, specify arguments with single quotation marks (`'`) to avoid interpretation by the shell.

URI paths

URI paths are supported, but with the following restrictions:

- If the source paths are URIs, they must all be in the same cloud storage account. No docroot (download), local docroot (upload), or source prefix can be specified.
- If a destination path is a URI, no docroot (upload) or local docroot (download) can be specified.
- The special schemes `stdio://` and `stdio-tar://` are supported on the client side only. They cannot be used for specifying an upload destination or download source.
- If required, specify the URI passphrase as part of the URI or set it as an environment variable (`ASPERA_SRC_PASS` or `ASPERA_DST_PASS`, depending on the transfer direction).

UNC paths

If the server is Windows and the path on the server is a UNC path (a path that points to a shared directory or file on Windows), it can be specified in an `ascp` command by using one of the following conventions:

- As an UNC path that uses backslashes (`\`): If the client side is a Windows computer, the UNC path can be used with no alteration. For example, `\\192.168.0.10\temp`. If the client is not a Windows computer, every backslash in the UNC path must be replaced with two backslashes. For example, `\\192.168.0.10\\temp`.
- As an UNC path that uses forward slashes (`/`): Replace each backslash in the UNC path with a forward slash. For example, if the UNC path is `\\192.168.0.10\temp`, change it to `//192.168.0.10/temp`. This format can be used with any client-side operating system.

Testing paths

To test `ascp` transfers, use a `faux://` argument in place of the source or target path to send random data without writing it to disk at the destination. For more information, see [. For examples, see “Ascp general examples” on page 38.](#)

WebSocket protocol

The WebSocket protocol can be used instead of SSH or HTTPS for client connections with the transfer server. To use it, you must configure `aspera.conf` specifically for WebSocket use. Then, for transfers, you must use the `ascp -ws-connect` option to specify by using WebSocket, and the `-P` option to specify the WebSocket port (9093).

Required file access and permissions

- Sources (for downloads) or destinations (for uploads) on the server must be in the transfer user's docroot or match one of the transfer user's file restrictions, otherwise the transfer stops and returns an error.
- The transfer user must have sufficient permissions to the sources or destinations, for example write access for the destination directory, otherwise the transfer stops and returns a permissions error.
- The transfer user must have authorization to do the transfer (upload or download), otherwise the transfer stops and returns a "management authorization refused" error.
- Files that are open for write by another process on a Windows source or destination cannot be transferred and return a "sharing violation" error. On Unix-like operating systems, files that are open for write by another process are transferred without reporting an error, but might produce unexpected results that depend on what data in the file is changed and when relative to the transfer.

Environment variables

The following environment variables can be used with the `ascp` command. The total size for environment variables depends on your operating system and transfer session. Each environment variable value must not exceed 4096 characters.

ASPERA_DST_PASS=password

The password to authenticate a URI destination.

ASPERA_LOCAL_TOKEN=token

A token that authenticates the user to the client (in place of SSH authentication).

Note: If the local token is a basic or bearer token, the access key settings for cipher and `preserve_time` are not respected and the server settings are used. To set the cipher and timestamp preservation options as a client, set them in the command line.

ASPERA_PROXY_PASS=*proxy_server_password*

The password for an Aspera Proxy server.

ASPERA_SCP_COOKIE=*cookie*

A cookie string that you want associated with transfers.

ASPERA_SCP_DOCROOT=*docroot*

The transfer user docroot. Equivalent to use **--apply-local-docroot** when a docroot is set in `aspera.conf`.

ASPERA_SCP_FILEPASS=*password*

The passphrase to be used to encrypt or decrypt files. For use with **--file-encrypt**.

ASPERA_SCP_KEY="-----BEGIN RSA PRIVATE KEY..."

The transfer user private key. Use instead of the **-i** option.

ASPERA_SCP_PASS=*password*

The password for the transfer user.

ASPERA_SCP_TOKEN=*token*

The transfer user authorization token. Overridden by **-W**.

ASPERA_SRC_PASS=*password*

The password to authenticate to a URI source.

Ascp options

-6

Enable IPv6 address support. When you specify an IPv6 numeric host for `src_host` or `dest_host`, write it in brackets. For example, `username@[2001:0:4137:9e50:201b:63d3:ba92:da]:/path` or `--host=[fe80::21b:21ff:fe1c:5072%eth1]`.

-@ *range_start:range_end*

Transfer only part of a file: `range_start` is the first byte to send, and `range_end` is the last. If either position is unspecified, the file's first and last bytes (respectively) are assumed. This option only works for downloads of a single file and does not support transfer resume.

-A, --version

Display version and license information.

--apply-local-docroot

Apply the local docroot that is set in `aspera.conf` for the transfer user. Use to avoid entering object storage access credentials in the command line. This option is equivalent to setting the environment variable `ASPERA_SCP_DOCROOT`.

-C *nodeid:nodecount*

Enable multi-session transfers (also known as parallel transfers) on a multi-node and multi-core system. A node ID (`nodeid`) and count (`nodecount`) are required for each session. `nodeid` and `nodecount` can be 1-128, but `nodeid` must be less than or equal to `nodecount`, such as 1:2, 2:2. Each session must use a different UDP port that is specified with the **-O** option. Large files can be split across sessions, see **--multi-session-threshold**. For more information, see [IBM Aspera High-Speed Transfer Server](#).

-c *cipher*

Encrypt in-transit file data by using the specified cipher. Aspera supports three sizes of AES cipher keys (128, 192, and 256 bits) and supports two encryption modes, Cipher Feedback mode (CFB) and Galois Counter Mode (GCM). The GCM mode encrypts data faster and increases transfer speeds compared to the CFB mode, but the server must support and permit it.

Cipher rules

The encryption cipher that you are allowed to use depends on the server configuration and the version of the client and server:

- When you request a cipher key that is shorter than the cipher key that is configured on the server, the transfer is automatically upgraded to the server configuration. For example, when the server setting is AES-192 and you request AES-128, the server enforces AES-192.
- When the server requires GCM, you must use GCM (requires version 3.9.0 or newer) or the transfer fails.
- When you request GCM and the server is older than 3.8.1 or explicitly requires CFB, the transfer fails.
- When the server setting is any, you can use any encryption cipher. The only exception is when the server is 3.8.1 or older and does not support GCM mode; in this case, you cannot request GCM mode encryption.
- When the server setting is none, you must use none. Transfer requests that specify an encryption cipher are refused by the server.

Cipher values

Value	Description	Support
aes128 aes192 aes256	Use GCM or CFB encryption mode, depending on the server configuration and version (see cipher negotiation matrix).	All client and server versions.
aes128cfb aes192cfb aes256cfb	Use CFB encryption mode.	Clients version 3.9.0 and newer, all server versions.
aes128gcm aes192gcm aes256gcm	Use GCM encryption mode.	Clients and servers.
none	Do not encrypt data in transit. For security and to keep the file integrity, avoid the use of this setting.	All client and server versions.

Client/Server Cipher negotiation

The following table shows which encryption mode is used depending on the server and client versions and settings:

	Server AES-XXX-GCM	Server AES-XXX-CFB	Server AES-XXX	Server AES-XXX
Client AES-XXX-GCM	GCM	Server refuses transfer	GCM	Server refuses transfer
Client AES-XXX-CFB	Server refuses transfer	CFB	CFB	CFB
Client AES-XXX	GCM	CFB	CFB	CFB
Client AES-XXX	Server refuses transfer	CFB	CFB	CFB

--check-sshfp=fingerprint

Compare *fingerprint* to the server SSH host key fingerprint that is set with `<ssh_host_key_fingerprint>` in `aspera.conf`. The Aspera fingerprint convention is to use a hex string without the colons; for example, `f74e5de9ed0d62feaf0616ed1e851133c42a0082`. For more information on SSH host key fingerprints, see [IBM Aspera High-Speed Transfer Server](#).

Note: If HTTP fallback is enabled and the transfer falls back to HTTP, this option enforces server SSL certificate validation (HTTPS). Validation fails if the server has a self-signed certificate; a properly signed certificate is required.

-D | -DD | -DDD

Log at the specified debug level. With each **D**, an additional level of debugging information is written to the log.

-d

Create the destination directory, if it does not exist. This option is automatically applied to uploads to object storage.

--delete-before-transfer

Before transfer, delete any files that exist at the destination but not also at the source. The source and destination arguments must be directories that have matching names. Do not use with multiple sources, keepalive, URI storage, or HTTP fallback. The **asdelete** tool provides the same capability.

--dest64

Indicate that the destination path or URI is base64 encoded.

-E 'pattern'

Exclude files or directories from the transfer based on matching the specified pattern to file names and paths (to exclude files by modification time, use `--exclude-newer-than` or `--exclude-older-than`). Use the `-N` option (include) to specify exceptions to `-E` rules. Rules are applied in the order in which they are encountered, from left to right. The following symbols can be used in the pattern:

- ***** (asterisk) represents zero or more characters in a string, for example `*.tmp` matches `.tmp` and `abcde.tmp`.
- **?** (question mark) represents a single character, for example, `t?p` matches `tmp` but not `temp`.

For details and examples, see [“Using filters to include and exclude files”](#) on page 45.

Note: When filtering rules are found in `aspera.conf`, they are applied *before* rules given on the command line (`-E` and `-N`).

-e prepost_script

Run the specified pre-post script as an alternative to the default `aspera-prepost` script. Specify the full path to the pre-post script. Use pre-post scripts to run custom commands such as shell scripts, Perl scripts, Windows batch files, and executable binaries that can invoke various environment variables. For instructions, see [IBM Aspera High-Speed Transfer Server](#).

--exclude-newer-than=mtime, --exclude-older-than=mtime

Exclude files but not directories from the transfer, based on when the file was last modified. Positive `mtime` values are used to express time, in seconds, since the original system time (usually `1970-01-01 00:00:00`). Negative `mtime` values (prefixed with `-`) are used to express the number of seconds before the current time.

-f config_file

Read Aspera configuration settings from `config_file` rather than `aspera.conf` (the default).

--file-checksum=hash

Enable checksum reporting for transferred files, where `hash` is the type of checksum to calculate: `sha1`, `md5`, `sha-512`, `sha-384`, `sha-256`, or `none` (the default). When the value is `none`, the checksum that is configured on the server, if any, is used. For more information about checksum reporting, see [“Reporting checksums”](#) on page 54.

Important: When checksum reporting is enabled, transfers of very large files (>TB) take a long time to resume because the entire file must be reread.

--file-crypt={encrypt|decrypt}

Encrypt files (when sending) or decrypt files (when receiving) for client-side encryption-at-rest (EAR). Encrypted files have the file extension `.aspera-env`. This option requires the encryption and decryption passphrase to be set with the environment variable `ASPERA_SCP_FILEPASS`. If a client-side encrypted file is downloaded with an incorrect password, the download is successful, but the file remains encrypted and still has the file extension `.aspera-env`. For more information, see [“Client-Side Encryption-at-Rest \(EAR\)”](#) on page 56.

--file-list=file

Transfer all source files and directories that are listed in *file*. Each source item is specified on a separate line. UTF-8 file format is supported. Only the files and directories are transferred. Path information is not preserved at the destination. To read a file list from standard input, use `-` in place of *file*.

For example, if `list.txt` contains the following list of sources:

```
/tmp/code/compute.php
doc_dir
images/iris.png
images/rose.png
```

And the following command is run:

```
# ascp --file-list=list.txt --mode=send --user=username --host=ip_addr .
```

Then, the destination in this case, the transfer user's docroot, contains:

```
compute.php
doc_dir (and its contents)
iris.png
rose.png
```

Restrictions:

- The command line cannot use the `user@host:source` syntax. Instead, specify this information with the options `--mode`, `--host`, and `--user`.
- Paths that are specified in the file list cannot use the `user@host:source` syntax.
- Because multiple sources are being transferred, the destination must be a directory.
- Only one **--file-list** or **--file-pair-list** option is allowed per ascp session. If multiple lists are specified, only the last one is used.
- Only files and directories that are specified in the file list are transferred. Any sources that are specified on the command line are ignored.
- If the source paths are URIs, the size of the file list cannot exceed 24 KB.

To create a file list that also specifies destination paths, use **--file-pair-list**.

--file-manifest={none|text}

Generate a list of all transferred files when set to `text`. Requires **--file-manifest-path** to specify the location of the list. (Default: `none`)

--file-manifest-path=directory

Save the manifest file to the specified location when you use **--file-manifest=text**. Manifests file must be stored locally. For cloud or other nonlocal storage, specify a *local* manifest file path.

--file-manifest-inprogress-suffix=suffix

Apply the specified suffix to the file manifest's temporary file. For use with **--file-manifest=text**. (Default suffix: `.aspera-inprogress`)

--file-pair-list=file

Transfer files and directories that are listed in *file* to their corresponding destinations. Each source is specified on a separate line, with its destination on the line that follows it.

Specify destinations relative to the transfer user's docroot. Even if a destination is specified as an absolute path, the path at the destination is still relative to the docroot. Destination paths that are specified in the list are created automatically if they do not exist.

For example, if the file `pairlist.txt` contains the following list of sources and destinations:

```
Dir1
Dir2
my_images/iris.png
project_images/iris.png
/tmp/code/compute.php
/tmp/code/compute.php
/tmp/tests/testfile
testfile2
```

And the following command is run:

```
# ascp --file-pair-list=pairlist.txt --mode=send --user=username --host=ip_addr .
```

Then, the destination in this case, the transfer user's docroot, now contains the following:

```
Dir2 (and its contents)
project_images/iris.png
tmp/code/compute.php
testfile2
```

Restrictions:

- The command line cannot use the `user@host:source` syntax. Instead, specify this information with the options `--mode`, `--host`, and `--user`.
- The `user@host:source` syntax cannot be used with paths that are specified in the file list.
- Because multiple sources are being transferred, the destination that is specified on the command line must be a directory.
- Only one `--file-pair-list` or `--file-list` option is allowed per `ascp` session. If multiple lists are specified, only the last one is used.
- Only files from the file pair list are transferred; any additional source files that are specified on the command line are ignored.
- If the source paths are URIs, the file list cannot exceed 24 KB.

For additional examples, see [“Ascp general examples”](#) on page 38.

-G *write_size*

If the transfer destination is a server, use the specified write-block size, which is the maximum number of bytes that the receiver can write to disk at a time. Default: 256 KB, Range: up to 500 MB. This option accepts suffixes M or m for *mega* and K or k for *kilo*, such that a *write_size* of 1M is 1 MB.

This is a performance-tuning option that overrides the `write_block_size` set in the client's `aspera.conf`. However, the `-G` setting is overridden by the `write_block_size` set in the server's `aspera.conf`. The receiving server never uses the `write_block_size` set in the client's `aspera.conf`.

-g *read_size*

If the transfer source is a server, use the specified read-block size, which is the maximum number of bytes that the sender reads from the source disk at a time. Default: 256 KB, Range: up to 500 MB. This option accepts suffixes M or m for *mega* and K or k for *kilo*, such that a *read_size* of 1M is 1 MB.

This is a performance-tuning option that overrides the `read_block_size` set in the client's `aspera.conf`. However, the `-g` setting is overridden by the `read_block_size` set in the server's `aspera.conf`. When set to the default value, the read size is the default internal buffer size of the server, which might vary by operating system. The sending server never uses the `read_block_size` set in the client's `aspera.conf`.

-h, --help

Display the help text.

--host=hostname

Transfer to the specified hostname or address. Requires **--mode**. This option can be used instead of specifying the host with the *hostname:file* syntax.

-i private_key_file | cert_file

The **-i** option can be used to specify either:

- An SSH private key file, for authenticating a transfer by using public key authentication with the specified SSH private key file. The argument can be just the file name if the private key is located in *user_home_dir/.ssh/* because ascp automatically searches for key files there. Multiple private key files can be specified by repeating the **-i** option. The keys are tried in order and the process ends when a key passes authentication or when all keys were tried without success, at which point authentication fails.
- A **Certificate Authority** certificate, for use with fallback transfers or with WebSocket use, when you do not want to use the system default certificate.

-K probe_rate

Measure bottleneck bandwidth at the specified probing rate (Kbps). (Default: 100 Kbps).

-k {0|1|2|3}

Enable the resuming of partially transferred files at the specified resume level. (Default: 0).

Specify this option for the first transfer or it doesn't work for subsequent transfers. Resume levels:

- k 0 – Always retransfer the entire file.
- k 1 – Compare file attributes and resume if they match, and retransfer if they do not.
- k 2 – Compare file attributes and the sparse file checksum; resume if they match, and retransfer if they do not.
- k 3 – Compare file attributes and the full file checksum; resume if they match, and retransfer if they do not.

If a complete file exists at the destination (no .aspx), the source and destination file sizes are compared. If a partial file and a valid .aspx file exist at the destination, the source file size and the file size that is recorded in the .aspx file are compared.

Note: If the destination is a URI path, then the only valid options are **-k 0** and **-k 1** and no .aspx file is created.

-L local_log_dir[:size]

Log to the specified directory on the client computer rather than the default directory. Optionally, set the size of the log file (Default: 10 MB). See also **-R** for setting the log directory on the server.

-l max_rate

Transfer at rates up to the specified target rate. (Default: 10000 Kbps) This option accepts suffixes G or g for *giga*, M or m for *mega*, K or k for *kilo*, and P, p, or % for percentage. Decimals are allowed. If this option is not set by the client, the setting in the server's *aspera.conf* is used. If a rate cap is set in the local or server *aspera.conf*, the rate does not exceed the cap.

-m min_rate

Attempt to transfer no slower than the specified minimum transfer rate. (Default: 0) If this option is not set by the client, then the server's *aspera.conf* setting is used. If a rate cap is set in the local or server *aspera.conf*, then the rate does not exceed the cap.

--mode={send|recv}

Transfer in the specified direction: send or recv (receive). Requires **--host**.

--move-after-transfer=archivedir

Move source files and copy source directories to *archivedir* after they are successfully transferred. Because directories are copied, the original source tree remains in place. The transfer user must have write permissions to the *archivedir*. The *archivedir* is created, if it does not exist. If the archive directory cannot be created, the transfer proceeds and the source files remain in their original location.

To preserve portions of the file path above the transferred file or directory, use this option with **--src-base**. For an example, see [“Ascp file manipulation examples” on page 41](#).

To remove empty source directories (except those specified as the source to transfer), use this option with **--remove-empty-directories**.

Restrictions:

- *archivedir* must be on the same file system as the source. If the specified archive is on a separate file system, it is created (if it does not exist), but an error is generated and files are not moved to it.
- For cloud storage, *archivedir* must be in the same cloud storage account as the source and must not already contain files with the same name (the existing files cannot be overwritten and the archiving fails).
- If the source is on a remote system (ascp is run in receive mode), *archivedir* is subject to the same docroot restrictions as the remote user.
- **--remove-after-transfer** and **--move-after-transfer** are mutually exclusive. Using both in the same session generates an error.
- Empty directories are not saved to *archivedir*.
- When used with **--remove-empty-directories** and **--src-base**, scanning for empty directories starts at the specified source base and proceeds down any subdirectories. If no source base is specified and a file path (as opposed to a directory path) is specified, then only the immediate parent directory is removed (if empty) after the source files were moved.

--multi-session-threshold=threshold

Split files across multiple ascp sessions if their size is greater than or equal to *threshold*. Use with **-C**, which enables multi-session transfers.

Files whose sizes are less than *threshold* are not split. If *threshold* is set to 0 (the default), no files are split.

If **--multi-session-threshold** is not used, the threshold value is taken from the setting for `<multi_session_threshold_default>` in the `aspera.conf` file on the client. If not found in `aspera.conf` on the client, the setting is taken from `aspera.conf` on the server. The command line setting overrides any `aspera.conf` settings, including when the command line setting is 0 (zero).

Multi-session uploads to cloud storage are supported for S3 only and require additional configuration. For more information, see the [IBM Aspera High-Speed Transfer Server](#).

-N 'pattern'

Include files or directories in the transfer based on matching the specified pattern to file names and paths. Rules are applied in the order in which they are encountered, from left to right, such that **-N** rules protect files from **-E** rules that follow them.

Note: An include rule **must** be followed by at least one exclude rule, otherwise all files are transferred because none are excluded. To exclude all files that do not match the include rule, use **-N '/**/'** **-E '/**'** at the end of your filter arguments.

The following symbols can be used in the pattern:

- ***** (asterisk) represents zero or more characters in a string, for example `*.tmp` matches `.tmp` and `abcde.tmp`.
- **?** (question mark) represents any single character, for example `t?p` matches `tmp` but not `temp`.

For more information see [“Using filters to include and exclude files”](#) on page 45.

Note: Filtering rules can also be specified in `aspera.conf`. Rules that are found in `aspera.conf` are applied *before* any **-E** and **-N** rules that are specified on the command line.

-O fasp_port

Use the specified UDP port for FASP transfers. (Default: 33001).

--output-file-progress

Can be used to write the individual file transfer progress to the stdout file descriptor.

--overwrite={never|always|diff|diff+older|older}

Overwrite destination files with source files of the same name. Default: diff. This option takes the following values:

- **never** - Never overwrite the file. However, if the parent folder is not empty, its access, modify, and change times might still be updated.
- **always** - Always overwrite the file.
- **diff** - Overwrite the file if different from the source. If a complete file at the destination is the same as a file on the source, it is not overwritten. Partial files are overwritten or resumed depending on the resume policy.
- **diff+older** - Overwrite the file if older and also different than the source. For example, if the destination file is the same as the source, but with a different timestamp, it is not overwritten. Plus, if the destination file is different than the source, but newer, it is overwritten.
- **older** - Overwrite the file if its timestamp is older than the source timestamp.

Interaction with resume policy (-k): If the overwrite method is **diff** or **diff+older**, difference is determined by the resume policy (**-k {0|1|2|3}**). If **-k 0** or no **-k** is specified, the source and destination files are always considered different and the destination file is always overwritten. If **-k 1**, the source and destination files are compared based on file attributes (currently file size). If **-k 2**, the source and destination files are compared based on sparse checksum. If **-k 3**, the source and destination files are compared based on full checksum.

-P ssh-port | websockets-port

Use the specified TCP port to initiate the FASP transfer session, by using the port number that is appropriate for your use of either SSH or WebSocket.

-p

Preserve file timestamps for access and modification time. Equivalent to setting **--preserve-modification-time** and **--preserve-access-time** (and **--preserve-creation-time** on Windows). Timestamp support in object storage varies by provider; consult your object storage documentation to determine which settings are supported.

On Windows, modification time might be affected when the system automatically adjusts for Daylight Savings Time (DST). For more information, see [Daylight saving time help and support](#).

--partial-file-suffix=suffix

Enable the use of partial files for files that are in transit, and set the suffix to add to names of partial files. (The suffix does not include a ". ", as for a file extension, unless explicitly specified as part of the suffix). This option only takes effect when set on the receiver side. When the transfer is complete, the suffix is removed. (Default: suffix is null; use of partial files is disabled).

--policy={high|fair|low|fixed}

Set the FASP transfer policy.

- **high** - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, the transfer rate is twice as fast as a fair-policy transfer. The **high** policy requires maximum (target) and minimum transfer rates.
- **fair** - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, bandwidth is shared fairly by transferring at an even rate. The **fair** policy requires maximum (target) and minimum transfer rates.
- **low** - Adjust the transfer rate to use the available bandwidth up to the maximum rate. Similar to fair mode, but less aggressive when the bandwidth is shared with other network traffic. When congestion occurs, the transfer rate is reduced to the minimum rate until other traffic decreases.
- **fixed** - Attempt to transfer at the specified target rate, regardless of network or storage capacity. This can decrease transfer performance and cause problems on the target storage. Use the **fixed** policy only for specific contexts, such as bandwidth testing, otherwise, avoid the use of this policy. The **fixed** policy requires a maximum target rate.
- **aggressiveness** - The aggressiveness of transfers that are authorized by this access key in claiming available bandwidth. Value can be 0.00-1.00. For example, these values correspond to the

policy option where a policy of high approximates to aggressiveness of 0.75, fair to 0.50 and low to 0.25. Aggressiveness can be used if you need to fine-tune the transfer policy.

If **--policy** is not set, ascp uses the server-side policy setting (fair by default). If the server does not allow the selected policy, the transfer fails.

--precalculate-job-size

Calculate the total size before a transfer starts. The server-side `pre_calculate_job_size` setting in `aspera.conf` overrides this option.

--preserve-access-time

Preserve the source-file access timestamps at the destination. Because source access times are updated by the transfer operation, the timestamp that is preserved is the one just *before* to the transfer. (To prevent access times at the source from being updated by the transfer operation, use the **--preserve-source-access-time** option).

--preserve-acls={native|metafile|none}

Preserve Access Control Lists (ACL) data for macOS, Windows, and AIX files. To preserve ACL data for other operating systems, use **--preserve-xattrs**. See also **--remote-preserve-acls**. Default: none.

- **native** - Preserve attributes by using the native capabilities of the file system. This mode is only supported for Windows, macOS, and AIX. If the destination and source do not support the same native ACL format, **async** reports an error and exits.
- **metafile** - Preserve file attributes in a separate file, named `filename.aspera-meta`. For example, attributes for `readme.txt` are preserved in a second file named `readme.txt.aspera-meta`. These metafiles are platform independent and can be copied between hosts without loss of information. This mode is supported on all file systems.
- **none** - Do not preserve attributes. This mode is supported on all file systems.

Important usage information:

- ACLs are not preserved for directories.
- Both **--preserve-acls** and **--remote-preserve-acls** must be specified in order for the target side of a pull (Ascp with `--mode=recv`) to apply the ACLs.
- Old versions of ascp do not support values other than none, and transfers that use native or metafile fail with an error that reports incompatible FASP protocol versions.

--preserve-creation-time

(Windows only) Preserve source-file creation timestamps at the destination. Only Windows systems retain information about creation time. If the destination is not a Windows computer, this option is ignored.

--preserve-file-owner-gid, --preserve-file-owner-uid

(Linux, UNIX, and macOS only) Preserve the group information (gid) or owner information (uid) of the transferred files. These options require the transfer user to be authenticated as a superuser.

--preserve-modification-time

Set the modification time, the last time a file or directory was modified (written), of a transferred file to the modification of the source file or directory. Preserve source-file modification timestamps at the destination.

On Windows, modification time might be affected when the system automatically adjusts for Daylight Savings Time (DST). For details, see [Daylight saving time help and support](#).

--preserve-source-access-time

Preserve the access times of the original sources to the last access times before to transfer. This prevents access times at the source from being updated by the transfer operation. Typically used with the **--preserve-access-time** option.

--preserve-xattrs={native|metafile|none}

Preserve extended file attributes data (xattr). Default: none. See also **--remote-preserve-xattrs**.

- **native** - Preserve attributes by using the native capabilities of the file system. This mode is supported only on macOS and Linux. If the destination and source do not support the same native xattr format, **async** reports an error and exits. If the Linux user is not root, some attributes such as system group might not be preserved.
- **metafile** - Preserve file attributes in a separate file, named *filename.aspera-meta*. For example, attributes for *readme.txt* are preserved in a second file named *readme.txt.aspera-meta*. These metafiles are platform independent and can be copied between hosts without loss of information. This mode is supported on all file systems.
- **none-preserve-acls={native|metafile|none}** - Do not preserve attributes. This mode is supported on all file systems.

Important usage information:

- Extended attributes are not preserved for directories.
- If Ascp is run by a regular user, only user-level attributes are preserved. If run as superuser, all attributes are preserved.
- The amount of attribute data per file that can be transferred successfully is subject to ascp's internal PDU size limitation.
- Old versions of Ascp do not support values other than none, and transfers that use native or metafile fail with an error that reports incompatible FASP protocol versions.

--proxy=proxy_url

Use the proxy server at the specified address. *proxy_url* must be specified with the following syntax:

```
dnat[s]://proxy_username:proxy_password@server_ip_address:port
```

The default ports for DNAT and DNATS protocols are 9091 and 9092. For a usage example, see [“Ascp general examples”](#) on page 38.

-q

Run ascp in quiet mode (disables the progress display).

-R remote_log_dir

Log to the specified directory on the server rather than the default directory.

Note: Client users restricted to `aspshe11` are not allowed to use this option. To specify the location of the local log, use **-L**.

--remote-preserve-acls={native|metafile|none}

Like, `--preserve-acls` but used when ACLs are stored in a different format on the remote computer. Defaults to the value of `--preserve-acls`.

Note: Both `--preserve-acls` and `--remote-preserve-acls` must be specified in order for the target side of a pull (Ascp with `-mode=recv`) to apply the ACLs.

--remote-preserve-xattrs={native|metafile|none}

Like, `--preserve-xattrs` but used when attributes are stored in a different format on the remote computer. Defaults to the value of `--preserve-xattrs`.

--remove-after-transfer

Remove all source files, but not the source directories, once the transfer is completed successfully. Requires write permissions on the source.

--remove-empty-directories

Remove empty source directories once the transfer is completed successfully, but do not remove a directory that is specified as the source argument. To also remove the specified source directory, use `--remove-empty-source-directory`. Directories can be emptied by using `--move-after-transfer` or `--remove-after-transfer`. Scanning for empty directories starts at the `src-base` and proceeds down any subdirectories. If no source base is specified and a file path (as opposed to a directory path) is specified, then only the immediate parent directory is scanned and removed if it's empty following the move of the source file.

Note: Do not use this option if multiple processes (ascp or other) might access the source directory at the same time.

--remove-empty-source-directory

Remove directories that are specified as the source arguments. For use with **--remove-empty-directories**.

-S remote_ascp

Use the specified remote ascp binary, if different than ascp.

--save-before-overwrite

Save a copy of a file before it is overwritten by the transfer. A copy of `filename.ext` is saved as `filename.yyyy.mm.dd.hh.mm.ss.index.ext` in the same directory. `index` is set to 1 at the start of each second and incremented for each additional file that is saved during that second. The saved copies retain the attributes of the original. Not supported for URI path destinations.

-SSH

Use an external SSH program instead of the built-in libssh2 implementation to establish the connection with the remote host. The wanted SSH program must be defined in the environment's PATH variable. To enable debugging of the SSH process, use the `-DD` and `--ssh-arg=-vv` options with ascp.

--ssh-arg=ARG

Add `ARG` to the command line arguments passed to the external SSH program (implies that you use `-SSH`). This option might be repeated as needed to supply multiple separate SSH arguments. The order is preserved. The `ARG` elements are inserted before any key files supplied to ascp, and before the user and host argument.

--skip-special-files

Skip special files, such as devices and pipes, without reporting errors for them.

--source-prefix=prefix

Add a `prefix` to each source path. The prefix can be a conventional path or a URI; however, URI paths can be used only if no `docroot` is defined.

--source-prefix64=prefix

Add a base64-encoded `prefix` to each source path. If **--source-prefix=prefix** is also used, the last option takes precedence.

--src-base=prefix

Strip the specified path prefix from the source path of each transferred file or directory. The remaining portion of the path remains intact at the destination.

Without `--src-base`, source files and directories are transferred without their source path. (However, directories do include their contents).

Note: Sources that are located outside the source base are not transferred. No errors or warnings are issued, but the skipped files are logged.

Use with URIs: The `--src-base` option performs a character-to-character match with the source path. For object storage source paths, the prefix must specify the URI in the same manner as the source paths. For example, if a source path includes an embedded passphrase, the prefix must also include the embedded passphrase otherwise it does not match.

For examples, see [“Ascp file manipulation examples”](#) on page 41.

--src-base64=<base64-encoded src-base>

An alternative to `--src-base`, with the same value except base64-encoded to help avoid character conversion issues for nonascii character sets. If both `--src-base` and `--src-base64` are specified, then the last argument on the command line is used.

--symbolic-links={follow|copy|copy+force|skip}

Handle-symbolic links that use the specified method, as allowed by the server. For more information, see [“Symbolic link handling”](#) on page 51. On Windows, the only method is `skip`. On other operating systems, any of the following methods can be used:

- `follow` - Follow symbolic links and transfer the linked files. (Default).

- `copy` - Copy only the alias file. If a file with the same name is found at the destination, the symbolic link is not copied.
 - `copy+force` - Copy only the alias file. If a file (not a directory) with the same name is found at the destination, the alias replaces the file. If the destination is a symbolic link to a directory, it is not replaced.
 - `skip` - Skip symbolic links. Do not copy the link or the file it points to.
- T**
Disable in-transit encryption for maximum throughput.
- tags *string***
Metatags in JSON format. The value is limited to 4 Kb.
- tags64 *string***
Metatags in JSON format and base64 encoded. The value is limited to 4 Kb.
- u *user_string***
Define a user string for Lua scripts that can be run with transfer events. For more information, see [Transfer session data accessible to scripts](#).
- user=*username***
Authenticate the transfer by using the specified username. Use this option instead of specifying the username as part of the destination path (as `user@host:file`).
- Note:** If you are authenticating on a Windows computer as a domain user, the transfer server strips the domain from the username. For example, `Administrator` is authenticated rather than `DOMAIN\Administrator`. For this reason, you must specify the domain explicitly.
- v**
Run `ascp` in verbose mode. This option prints connection and authentication debug messages in the log file. For more information, see [“Log files” on page 75](#).
- W {*token_string*|@*token_file*}**
Authenticate by using the authorization token string for the transfer, either as the string itself or when preceded with an `@`, the full path to the token file. This option takes precedence over the setting for the `ASPERA_SCP_TOKEN` environment variable.
- w*r*, -w*f***
Measure and report bandwidth from server to client (**-w*r***) or client to server (**-w*f***) before the transfer.
- ws-connect**
Use WebSocket instead of SSH for client connections with the transfer server.
- X *rexmsg_size***
Limit the size of retransmission requests to no larger than the specified size, in bytes. (Max: 1440).
- Z *dgram_size***
Use the specified datagram size (MTU) for FASP transfers. Range: 296 - 65535 bytes. (Default: the detected path MTU).
- As of version 3.3, datagram size can be specified on the server by setting `<datagram_size>` in `aspera.conf`. The server setting overrides the client setting, unless the client is using a version of `ascp` that is older than 3.3, in which case the client setting is used. If the pre-3.3 client does not set **-Z**, the datagram size is the discovered MTU and the server logs the message `LOG Peer client does not support alternative datagram size`.

Ascp Options for HTTP Fallback

HTTP fallback serves as a secondary transfer method when the internet connectivity required for Aspera FASP transfers (UDP port 33001, by default) is unavailable. When HTTP fallback is enabled and UDP connectivity is lost or cannot be established, the transfer continues over the HTTP/S protocol.

Limitations:

- HTTP fallback must be enabled on the server.

- Folders that are symbolic links cannot be downloaded directly by using HTTP fallback. Folders that are symbolic links are processed correctly when their parent folder is the source.
- HTTP fallback can only follow symbolic links. Settings in `aspera.conf` or in the command line are ignored.
- HTTP fallback attempts to transfer at the target rate but is limited by TCP.
- HTTP fallback does not support automated execution of Lua scripts ([Automated execution of Lua scripts with transfer Events](#)).

Options:

-i *cert_file*

By default ascp uses the system certificates. However, the **-i** option can be used to use the specified **Certificate Authority** certificate for fallback transfers, and for WebSocket.

-t *port*

Transfer through the specified server port for HTTP fallback.

-x *proxy_server*

Transfer to the specified proxy server address for HTTP fallback.

-Y *key_file*

Certify HTTPS fallback transfers by using the specified HTTPS transfer key.

-y *{0|1}*

If set to 1, use the HTTP fallback transfer server when a UDP connection fails. (Default: 0).

Ascp general examples

Use the following Ascp examples to craft your own transfers.

To describe file paths, use single quotation marks (' ') around the file path string, and forward-slashes (/) on all platforms. Avoid the following characters in file names: / \ " : ' ? > < & * |

• Growing Files

The growing files feature that allows to start transferring files to the target directory while they are still being written to the source directory.

Download the growing file `myfile` with a wait period of 120 seconds that uses a 0 bytes read that calculates the wait time.

```
ascp --mode=recv --user=root --host=10.0.0.2 "file:///tmp/myfile?grow=120&wait_start=null_read" file:///tmp2/mylocalfile
```

To support this command, the `ascp.conf` file must include the following configuration:

```
<default>
  <file_system>
    <access>
      <paths><path><absolute>
        file:///tmp?grow=120;wait_start=null_read
      </absolute></path></paths>
    </access>
  </file_system>
</default>
```

For more information, see [aspera.conf - File system configuration](#).

• Using the WebSocket Protocol

This example shows how to use the WebSocket protocol for a transfer. The Aspera Node Service provides a WebSocket server, which must be enabled. Because the ascp client supports only a secure WebSocket transfer (HTTPS), the Aspera Node Service must be configured for HTTPS, or must use a reverse proxy to terminate the secure connection.

A basic token, bearer token, or transfer token must be used with a WebSocket connection.

The following ascp options are required for the use of a WebSocket:

--ws-connect

Specifies the use of a WebSocket.

-P

Specifies the WebSocket port (9093).

```
# ascp -L- --ws-connect -P 9093 --host=www.example.com --mode=send --user=xeno c:/Users/xeno/Desktop/myfile /Desktop/ dest
```

- **Fair-policy transfer**

Fair-policy transfer with maximum rate 100 Mbps and minimum at 1 Mbps, without encryption, transfer all files in \local-dir\files to 10.0.0.2:

```
# ascp --policy=fair -l 100m -m 1m /local-dir/files root@10.0.0.2:/remote-dir
```

- **Fixed-policy transfer**

Fixed-policy transfer with target rate 100 Mbps, without encryption, transfer all files in \local-dir\files to 10.0.0.2:

```
# ascp -l 100m /local-dir/files root@10.0.0.2:/remote-dir
```

- **Specify UDP port for transfer**

Transfer by using UDP port 42000:

```
# ascp -l 100m -o 42000 /local-dir/files user@10.0.0.2:/remote-dir
```

- **Public key authentication**

Transfer with public key authentication by using the key file <home dir>/.ssh/aspera_user_1-key local-dir/files:

```
$ ascp -l 10m -i ~/.ssh/aspera_user_1-key local-dir/files root@10.0.0.2:/remote-dir
```

- **Username or file path contains a space**

Enclose the target in double quotation marks (" ") when spaces are present in the username and remote path:

```
# ascp -l 100m local-dir/files "User Name@10.0.0.2:/remote directory"
```

- **Content is specified in a file pair list**

Specify source content to transfer to various destinations in a file pair list. Source content is specified by using the full file or directory path. Destination directories are specified relative to the transfer user's docroot, which is specified as a "." at the end of the **ascp** command. For example, the following is a simple file pair list, `filepairlist.txt` that lists two source folders, `folder1` and `folder2`, with two destinations, `tmp1` and `tmp2`:

```
/tmp/folder1
tmp1
/tmp/folder2
tmp2
```

```
# ascp --user=user_1 --host=10.0.0.2 --mode=send --file-pair-list=/tmp/filepairlist.txt .
```

This command and file pair list create the following directories within the transfer user's docroot on the destination:

```
/tmp1/folder1
/tmp2/folder2
```

- **Network shared location transfer**

Send files to a network shares location \\1.2.3.4\nw-share-dir, through the computer 10.0.0.2:

```
# ascp local-dir/files root@10.0.0.2:"//1.2.3.4/nw-share-dir/"
```

- **Parallel transfer on a multi-core system**

Use parallel transfer on a dual-core system, together transferring at the rate 200 Mbps, by using UDP ports 33001 and 33002. Two commands are run in different Terminal windows:

```
# ascp -C 1:2 -O 33001 -l 100m /file root@10.0.0.2:/remote-dir &  
# ascp -C 2:2 -O 33002 -l 100m /file root@10.0.0.2:/remote-dir
```

- **Upload with content protection**

Upload the file **local-dir/file** to the server 10.0.0.2 with password protection (*password: secRet*):

```
# export ASPERA_SCP_FILEPASS=secRet ascp -l 10m --file-crypt=encrypt local-dir/file  
root@10.0.0.2:/remote-dir/
```

The file is saved on the server as `file.aspera-env`, with the extension that indicates that the file is encrypted. See the next example for how to download and decrypt an encrypted file from the server.

- **Download with content protection and decryption**

Download an encrypted file, `file.aspera-env`, from the server 10.0.0.2 and decrypt while the transfer:

```
# export ASPERA_SCP_FILEPASS=secRet; ascp -l 10m --file-crypt=decrypt root@10.0.0.2:/remote-  
dir/file.aspera-env /local-dir
```

- **Decrypt a downloaded, encrypted file**

If the password-protected file **file1** is downloaded on the local computer without decrypting, decrypt **file1.aspera-env** (the name of the downloaded and encrypted version of **file1**) to **file1**:

```
$ export ASPERA_SCP_FILEPASS=secRet; /opt/aspera/bin/asunprotect -o file1 file1.aspera-env
```

- **Download through Aspera forward proxy with proxy authentication**

User Pat transfers the file `/data/file1` to `/Pat_data/` on 10.0.0.2, through the proxy server at 10.0.0.7 with the proxy username `aspera_proxy` and password `pa33w0rd`. After the command is run, Pat is prompted for the transfer user's (Pat's) password.

```
# ascp --proxy dnats://aspera_proxy:pa33w0rd@10.0.0.7 /data/file1 Pat@10.0.0.2:/Pat_data/
```

Test transfers that use faux://

For more information, see .

- **Transfer-random data (no source storage required)**

Transfer 20 GB of random data as user `root` to file `newfile` in the directory `/remote-dir` on 10.0.0.2:

```
#ascp --mode=send --user=root --host=10.0.0.2 faux:///newfile?20g /remote-dir
```

- **Transfer a file but do not save results to disk (no destination storage required)**

Transfer the file `/tmp/sample` as user `root` to 10.0.0.2, but do not save results to disk:

```
#ascp --mode=send --user=root --host=10.0.0.2 /tmp/sample faux://
```

- **Transfer-random data and do not save result to disk (no source or destination storage required)**

Transfer 10 MB of random data from 10.0.0.2 as user `root` and do not save result to disk:

```
#ascp --mode=send --user=root --host=10.0.0.2 faux:///dummy?10m faux://
```

Ascp file manipulation examples

Ascp can manipulate files and directories as part of the transfer, such as upload only the files in the specified source directory but not the directory itself, create a destination directory, and move or delete source files after they are transferred.

- **Upload a directory**

Upload the directory `/data/` to the server at `10.0.0.1`, and place it in the `/storage/` directory on the server:

```
# ascp /src/data/ root@10.0.0.1:/storage/
```

- **Upload only the contents of a directory (not the directory itself) by using the `--src-base` option**

Upload only the contents of `/data/` to the `/storage/` directory at the destination. Strip the `/src/data/` portion of the source path and preserve the remainder of the file structure at the destination:

```
# ascp --src-base=/src/data/ /src/data/ root@10.0.0.1:/storage/
```

- **Upload a directory and its contents to a new directory by using the `-d` option**

Upload the `/data/` directory to the server and if it doesn't already exist, create the new folder `/storage2/` to contain it, resulting in `/storage2/data/` at the destination.

```
# ascp -d /src/data/ root@10.0.0.1:/storage2/
```

- **Upload the contents of a directory, but not the directory itself, by using the `--src-base` option**

Upload all folders and files in the `/clips/out/` folder, but not the `out/` folder itself, to the `/in/` folder at the destination.

```
# ascp -d --src-base=/clips/out/ /clips/out/ root@10.0.0.1:/in/
```

Result, the source folders and their content appear in the `in` directory at the destination:

Source	Destination	Destination without <code>--src-base</code>
<code>/clips/out/file1</code>	<code>/in/file1</code>	<code>/in/out/file1</code>
<code>/clips/out/folderA/file2</code>	<code>/in/folderA/file2</code>	<code>/in/out/folderA/file2</code>
<code>/clips/out/folderB/file3</code>	<code>/in/folderB/file3</code>	<code>/in/out/folderB/file3</code>

Without `--src-base`, the example command transfers not only the contents of the `out/` folder, but the folder itself.

Note: Sources that are located outside the source base are not transferred. No errors or warnings are issued, but the skipped files are logged. For example, if `/clips/file4` were included in the example sources, it would not be transferred because it is located outside the specified source base `/clips/out/`.

- **Upload only the contents of a file and a directory to a new directory by using `--src-base`**

Upload a file, `/monday/file1`, and a directory, `/tuesday/*`, to the `/storage/` directory on the server, while stripping the `src-base` path and preserving the rest of the file structure. The content is saved as `/storage/monday/file1` and `/storage/tuesday/*` on the server.

```
# ascp --src-base=/data/content /data/content/monday/file1 /data/content/tuesday/ root@10.0.0.1:/storage
```

- **Download only the contents of a file and a directory to a new directory by using `--src-base`**

Download a file, /monday/file1, and a directory, /tuesday/*, from the server, while stripping the src-base path and preserving the rest of the file structure. The content is saved as /data/monday/file1 and /data/tuesday/* on the client.

```
# ascp --src-base=/storage/content root@10.0.0.1:/storage/content/monday/file1 root@10.0.0.1:/storage/content/tuesday/ /data
```

- **Move the source file on the client after it is uploaded to the server by using --move-after-transfer**

Upload file0012 to Pat's docroot on the server at 10.0.0.1, and move (not copy) the file from C:/Users/Pat/srcdir/ to C:/Users/Pat/Archive on the client.

```
# ascp --move-after-transfer=C:/Users/Pat/Archive C:/Users/Pat/srcdir/file0012 Pat@10.0.0.1:/
```

- **Move the source file on the server after it is downloaded to the client by using --move-after-transfer**

Download srcdir from the server to C:/Users/Pat on the client, and move (not copy) srcdir to the archive directory /Archive on the server.

```
# ascp --move-after-transfer=Archive Pat@10.0.0.1:/srcdir C:/Users/Pat
```

- **Move the source file on the client after it is uploaded to the server and preserve the file structure one level above it by using --src-base and --move-after-transfer**

Upload file0012 to Pat's docroot on the server at 10.0.0.1, and save it as /srcdir/file0012 (stripped of C:/Users/Pat). Also, move file0012 from C:/Users/Pat/srcdir/ to C:/Users/Pat/Archive on the client, where it is saved as C:/Users/Pat/Archive/srcdir/file0012.

```
# ascp --src-base=C:/Users/Pat --move-after-transfer=C:/Users/Pat/Archive C:/Users/Pat/srcdir/file0012 Pat@10.0.0.1:/
```

- **Delete a local directory once it is uploaded to the remote server by using --remove-after-transfer and --remove-empty-directories**

Upload /content/ to the server, then delete its contents (excluding partial files) and any empty directories on the client.

```
# ascp -k2 -E "*.partial" --remove-after-transfer --remove-empty-directories /data/content root@10.0.0.1:/storage
```

- **Delete a local directory once its contents were transferred to the remote server by using --src-base, --remove-after-transfer, and --remove-empty-directories**

Upload /content/ to the server, while stripping the src-base path and preserving the rest of the file structure. The content is saved as /storage/* on the server. On the client, the contents of /content/, including empty directories but excluding partial files, are deleted.

```
# ascp -k2 -E "*.partial" --src-base=/data/content --remove-after-transfer --remove-empty-directories /data/content root@10.0.0.1:/storage
```

Using standard I/O as the source or destination

Ascp can use standard input (stdin) as the source or standard output (stdout) as the destination for a transfer. The syntax depends on the number of files in your transfer; for single files use `stdio://` and for multiple files use `stdio-tar://`. The transfer is authenticated by using SSH or a transfer token.

Named pipes

A named pipe can be specified as a `stdio` destination, with the syntax `stdio:///path` for single files, or `stdio-tar:///path` for multiple files, where *path* is the path of the named pipe. If a docroot is configured on the destination, then the transfer goes to the named pipe `docroot/path`.

Note: Do not use `stdio:///path` to transfer multiple files. The file data is asynchronously concatenated in the output stream and might be unusable. Use `stdio-tar:///path` instead, which demarcates multiple files with headers.

Note: Do not use 0-byte files with standard I/O transfers.

Single file transfers

To upload data that is piped into `stdin`, set the source as `stdio:///?fsize`, where *fsize* is the number of bytes (as a decimal) that are received from `stdin`. The destination is set as the path and file name. The file modification time is set to the time at which the upload starts. Standard input must transfer the exact amount of data that is set by *fsize*. If more or less data is received by the server, an error is generated.

To download data and pipe it into `stdout`, set the destination as `stdio://`.

Restrictions:

- `stdio://` cannot be used for persistent sessions. Use `stdio-tar://` instead.
- Only `--overwrite=always` or `--overwrite=never` are supported by `stdio://`. The behavior of `--overwrite=diff` and `--overwrite=diff+older` is undefined.
- When the size (*?fsize*) is not specified, the file is managed by the growing files feature where an End-of-File (EOF) on reading `stdin` will signify that the growing file is complete.

Single-file Transfer Examples:

- Upload 1025 bytes of data from the client `stdin` to `/remote-dir` on the server at 10.0.0.2. Save the data as the file `newfile`. Transfer at 100 Mbps.

```
cat myfile | ascp -l 100m --mode=send --user=username --host=10.0.0.2 stdio:///newfile?1025 /remote-dir
```

- Download the file `remote_file` from the server at 10.0.0.2 to `stdout` on the client. Transfer at 100 Mbps.

```
ascp -l 100m --mode=recv --user=username --host=10.0.0.2 remote_file stdio://
```

- Upload the file `local_file` to the server at 10.0.0.2 to the named pipe `/tmp/outpipe`. Transfer at 100 Mbps.

```
ascp -l 100m --mode=send --user=username --host=10.0.0.2 local_file stdio:///tmp/outpipe
```

Multi-File transfers

Ascp can transfer one or more files in an encoded, streamed interface, similar to single file transfers. The primary difference is that the stream includes headers that demarcate data from individual files.

To upload files that are piped into `stdin`, set the source as `stdio-tar://`. The file modification time is set to the time at which the upload starts.

The files in the input stream must be encoded in the following format. *File* can be the file name or file path, *Size* is the size of the file in bytes, and *Offset* is an optional parameter that sets where in the destination file to begin overwriting with the raw inline data:

```
[0 - n blank lines]
File: /path/to/file_1
Size: file_size
Offset: bytes

file_1 data
[0 - n blank lines]
File: /path/to/file_2
Size: file_size

file 2 data
...
```

To download one or more files to stdout, set the destination as `stdio-tar://`. Normal status output to stdout is suppressed during downloads because the transfer output is streamed to stdout. The data sent to stdout has the same encoding as described for uploads.

To download to a named pipe, set the destination to `stdio-tar:///path`, where *path* is the path of the named pipe.

When an offset is specified, the bytes that are sent replace the existing bytes in the destination file (if it exists). The bytes added to the destination file can extend beyond the current file size. If no offset is set, the bytes overwrite the file if overwrite conditions are met.

Restrictions:

- When downloading to `stdio-tar://`, the source list must consist of individual files only. Directories are not allowed.
- Only `--overwrite=always` or `--overwrite=never` are supported by `stdio-tar://`. The behavior of `--overwrite=diff` and `--overwrite=diff+older` is undefined.
- Offsets are only supported if the destination files are located in the native file system. Offsets are not supported for cloud destinations.

Multi-file transfer examples:

- Upload two files, `myfile1` (1025 bytes) and `myfile2` (20 bytes), to `/remote-dir` on the server at 10.0.0.2. Transfer at 100 Mbps.

```
cat sourcefile | ascp -l 100m --mode=send --user=username --host=10.0.0.2 stdio-tar:// /remote-dir
```

Where `sourcefile` contains the following:

```
File: myfile1
Size: 1025

<< 1025 bytes of data>>
File: myfile2
Size: 20

<<20 bytes of data>>
```

- Uploading multiple files from stdin by using a persistent session is the same as a nonpersistent session.
- Update bytes 10-19 in file `/remote-dir/myfile1` on the server at 10.0.0.2 at 100 Mbps.

```
cat sourcefile | ascp -l 100m --mode=send --user=username --host=10.0.0.2 stdio-tar:// /remote-dir
```

Where `sourcefile` contains the following:

```
File: myfile1
Size: 10
Offset: 10

<< 10 bytes of data>>
```

- Upload two files, `myfile1` and `myfile2`, to the named pipe `/tmp/mypipe` (streaming output) on the server at 10.0.0.2. Transfer at 100 Mbps.

```
ascp -l 100m --mode=send --user=username --host=10.0.0.2 myfile1 myfile2 stdio-tar:///tmp/mypipe
```

This sends an encoded stream of `myfile1` and `myfile2` (with the format of `sourcefile` in the upload example) to the pipe `/tmp/mypipe`. If `/tmp/mypipe` does not exist, it is created.

- Download the files from the previous example from 10.0.0.2 to stdout. Transfer at 100 Mbps.

```
ascp -l 100m --mode=recv --user=username --host=10.0.0.2 myfile1 myfile2 stdio-tar://
```

Standard output receives data identical to sourcefile in the upload example.

- Download /tmp/myfile1 and /tmp/myfile2 to stdout by using a persistent session. Start the persistent session, which listens on management port 12345:

```
ascp -l 100m --mode=recv --keepalive -M 12345 --user=username --host=10.0.0.2 stdio-tar://
```

Send the following in through management port 12345:

```
FASPMGR 2
Type: START
Source: /tmp/myfile1
Destination: mynewfile1

FASPMGR 2
Type: START
Source: /tmp/myfile2
Destination: mynewfile2

FASPMGR 2
Type: DONE
```

The destination must be a file name; file paths are not supported.

Standard out receives the transferred data with the following syntax:

```
File: mynewfile1
Size: file_size

mynewfile1_data
File: mynewfile2
Size: file_size

mynewfile2_data
```

- Upload two files, myfile1 and myfile2, to named pipe /tmp/mypipe on the server at 10.0.0.2. Transfer at 100 Mbps.

```
ascp -l 100m --mode=send --user=username --host=10.0.0.2 myfile1 myfile2 stdio-tar:///tmp/
mypipe
```

If file/tmp/mypipe does not exist, it is created.

- Upload two files, myfile1 (1025 bytes) and myfile2 (20 bytes) from stdio and regenerate the stream on the destination to send out through the named pipe /tmp/mypipe on the server at 10.0.0.2. Transfer at 100 Mbps.

```
cat sourcefile | ascp -l 100m --mode=send --user=username --host=10.0.0.2 stdio-tar:// stdio-
tar:///tmp/pipe
```

Where sourcefile contains the following:

```
File: myfile1
Size: 1025

<< 1025 bytes of data>>
File: myfile2
Size: 20

<<20 bytes of data>>
```

Using filters to include and exclude files

For example, on Windows FAT or NTFS (or directories) to transfer by indicating which to skip or include based on name matching. When no filtering rules are specified by the client, **ascp** transfers all source files in the transfer list; servers cannot filter client uploads or downloads.

Command line syntax

-E '*pattern*' Exclude files or directories with names or paths that match *pattern*.

-N '*pattern*' Include files or directories with names or paths that match *pattern*.

Where:

- *pattern* is a file or directory name, or a set of names expressed with UNIX *glob* patterns.
- Surround patterns that contain wildcards with single quotation marks (' ') to prevent filter patterns from being interpreted by the command shell. Patterns that do not contain wildcards can also be in single quotation marks (' ').

Basic usage

- Filtering rules are applied to the transfer list in the order that they appear on the command line. If filtering rules are configured in `aspera.conf`, they are applied before the rules on the command line.
- Filtering is a process of exclusion, and -N rules override -E rules that follow them. -N cannot add back files that are excluded by a preceding exclude rule.
- An include rule **must** be followed by at least one exclude rule, otherwise all files are transferred because none are excluded. To exclude all files that do not match the include rule, use -N '/**/' -E '/**/' at the end of your filter arguments.
- Filtering operates only on the set of files and directories in the transfer list. An include rule (-N) cannot add files or directories that are not already part of the transfer list.

Example	Transfer Result
-E ' <i>rule</i> '	Transfer all files and directories except those with names that match <i>rule</i> .
-N ' <i>rule</i> '	Transfer all files and directories because none are excluded. To transfer only the files and directories with names that match <i>rule</i> use: <pre>ascp -N '<i>rule</i>' -N '/**/' -E '/**'</pre>
-N ' <i>rule1</i> ' -E ' <i>rule2</i> '	Transfer all files and directories with names that match <i>rule1</i> , and all other files and directories except those with names that match <i>rule2</i> .
-E ' <i>rule1</i> ' -N ' <i>rule2</i> '	Transfer all files and directories except those with names that match <i>rule1</i> . All files and directories that are not already excluded by <i>rule1</i> are included because no additional exclude rule follows -N ' <i>rule2</i> '. To transfer only the files and directories with names that do not match <i>rule1</i> but do match <i>rule2</i> use: <pre>ascp -E '<i>rule1</i>' -N '<i>rule2</i>' -N '/**/' -E '/**'</pre>

Filtering rule application

Filters can be specified on the `ascp` command line and in `aspera.conf`. `Ascp` applies filtering rules that are set in `aspera.conf` *before* it applies rules on the command line.

Filtering order

Filtering rules are applied to the transfer list in the order that they appear on the command line.

1. **Ascp** compares the first file (or directory) in the transfer list to the pattern of the first rule.
2. If the file matches the pattern, `Ascp` includes it (-N) or excludes it (-E) and the file is immune to any following rules.

Note: When a directory is excluded, directories and files in it are also excluded and are not compared to any following rules. For example, with the command line options -E '/images/' -N '/images/icons/', the directory `/images/icons/` is not included or considered because `/images/` was already excluded.

3. If the file does not match, Ascp compares it with the next rule and repeats the process for each rule until a match is found or until all rules are tried.
4. If the file never matches any exclude rules, it is included in the transfer.
5. The next file or directory in the transfer list is then compared to the filtering rules until all eligible files are evaluated.

Example:

Consider the following command:

```
# ascp -N 'file2' -E 'file[0-9]' /images/icons/ user1@examplehost:/tmp
```

Where `/images/icons/` is the source.

If `/images/icons/` contains `file1`, `file2`, and `fileA`, the filtering rules are applied as follows:

1. `file1` is compared with the first rule (`-N 'file2'`) and does not match so filtering continues.
2. `file1` is compared with the second rule (`-E 'file[0-9]'`) and matches, so it is excluded from the transfer.
3. `file2` is compared with the first rule and matches, so it is included in the transfer and filtering stops for `file2`.
4. `fileA` is compared with the first rule and does not match so filtering continues.
5. `fileA` is compared with the second rule and does not match. Because no rules exclude it, `fileA` is included in the transfer.

Note: If the filtering rules ended with `-N '/**/'` `-E '/**'`, then `fileA` would be excluded because it was not protected by an include rule.

Rule patterns

Rule patterns (globs) use standard globbing syntax that includes wildcards and special characters, and several Aspera extensions to the standard.

- **Character case:** Case always matters, even if the file system does not enforce such a distinction. For example, on Windows FAT or NTFS file systems and macOS HPFS+ a file system search for `DEBUG` returns files `Debug` and `debug`. In contrast, **ascp** filter rules use exact comparison, such that `debug` does not match `Debug`. To match both, use `[Dd]ebug`.
- **Partial matches:** With globs, unlike standard regular expressions, the entire file name or directory name must match the pattern. For example, the pattern `abc*f` matches `abcdef` but not `abcdefg`.

Standard globbing: Wildcards and special characters

<code>/</code>	The only recognized path separator.
<code>\</code>	Quotes any character literally, including itself. <code>\</code> is exclusively a quoting operator, not a path separator.
<code>*</code>	Matches zero or more characters, except <code>/</code> or the <code>.</code> in <code>..</code>
<code>?</code>	Matches any single character, except <code>/</code> or the <code>.</code> in <code>..</code>
<code>[...]</code>	Matches exactly one of a set of characters, except <code>/</code> or the <code>.</code> in <code>..</code>
<code>[^...]</code>	When <code>^</code> is the first character, matches exactly one character <i>not</i> in the set.
<code>[!...]</code>	When <code>!</code> is the first character, matches exactly one character <i>not</i> in the set.
<code>[x-x]</code>	Matches exactly one of a range of characters.
<code>[:xxxxx:]</code>	For details about this type of wildcard, see any POSIX-standard guide to globbing.

Globber Extensions: Wildcards and special characters

no / or * at end of pattern	Matches files only.
/ at end of pattern	Matches directories only. With -N, no files under matched directories or their subdirectories are included in the transfer. All subdirectories are still included, although their files are not included. However, with -E, excluding a directory also excludes all files and subdirectories under it.
* or /** at end of pattern	Matches both directories and files.
/**	Like * but also matches / and the . in /..
/ at start of pattern	Must match the entire string from the root of the transfer set. (Note: The leading / does not refer to the system root or the docroot.)

Standard globbing examples

Wildcard	Example	Matches	Does Not Match
/	abc/def/xyz	abc/def/xyz	abc/def
\	abc\?	abc?	abc\? abc/D abcD
*	abc*f	abcdef abc.f	abc/f abcefg
?	abc??	abcde abc.z	abcdef abc/d abc/.
[...]	[abc]def	adef cdef	abcdef ade
[^...]	[^abc]def	zdef .def 2def	bdef /def /.def
[!...]	[!abc]def	zdef .def 2def	cdef /def /.def
[:xxxxx:]	[:lower:]def	cdef ydef	Adef 2def .def

Globber extension examples

Wildcard	Example	Matches	Does Not Match
/**	a/**/f	a/f a/.z/f a/d/e/f	a/d/f/ za/d/f
* at end of rule	abc*	abc/ abcfile	
/** at end of rule	abc/**	abc/.file abc/d/e/	abc/
/ at end of rule	abc*/	abc/dir	abc/file
no / at end of rule	abc	abc (file)	abc/
/ at start of rule	/abc/def	/abc/def	xyz/abc/def

Testing your filter rules

You can use this procedure to test your filtering rules.

1. On your computer, create a set of directories and files (size can be small) that approximate a typical transfer file set. In the following example, the file set is in /tmp/sic.
2. Upload the file set to a server. For example,

```
# ascp /tmp/sic my_user_name@my_demo.example.com:Upload/
```

Where the user is my_user_name, and the target is the Upload directory.

At the prompt, enter my_user_name user password.

3. Create a destination directory on your computer, for example, /tmp/dest.
4. Download your files from the demo server to /tmp/dest to test your filtering rules. For example,

```
# ascp -N 'wxy/**' -E 'def' my_user_name@my_demo.example.com:Upload/src/ /tmp/dest
```

5. Compare the destination directory with the source to determine whether the filter behaved as expected.

```
$ diff -r dest/ src/
```

The **diff** output shows the missing files and directories (those that were not transferred).

Example filter rules

The example rules are based on running a command such as the following to download a directory AAA from my_demo.example.com to /tmp/dest:

```
# ascp rules aspera@my_demo.example.com:Upload/AAA /tmp/dest
```

The examples use the following file set:

```
AAA/abc/def
AAA/abc/.def
AAA/abc/.wxy/def
AAA/abc/wxy/def
AAA/abc/wxy/.def
AAA/abc/wxy/tuv/def
AAA/abc/xyz/def/wxy
AAA/wxyfile
AAA/wxy/xyx/
AAA/wxy/xyxfile
```

Key for interpreting example results:

```
< xxx/yyy = Excluded
xxx/yyy = Included
zzz/ = directory name
zzz = file name
```

1. Transfer everything except files and directories that starts with ".":

```
-N '*' -E 'AAA/**'
```

Results:

```
AAA/abc/def
AAA/abc/wxy/def
AAA/abc/wxy/tuv/def
AAA/abc/xyz/def/wxy
AAA/wxyfile
AAA/wxy/xyx/
AAA/wxy/xyxfile
< AAA/abc/.def
< AAA/abc/.wxy/def
< AAA/abc/wxy/.def
```

2. Exclude directories and files whose names start with wxy:

```
-E 'wxy*'
```

Results:

```
AAA/abc/def
AAA/abc/.def
AAA/abc/.wxy/def
AAA/abc/xyz/def/
< AAA/abc/wxy/def
< AAA/abc/wxy/.def
< AAA/abc/wxy/tuv/def
< AAA/abc/xyz/def/wxy
< AAA/wxyfile
< AAA/wxy/xyx/
< AAA/wxy/xyxfile
```

3. Include directories and files that start with wxy if they fall directly under AAA:

```
-N 'wxy*' -E 'AAA/**'
```

Results:

```
AAA/wxy/
AAA/wxyfile
< AAA/abc/def
< AAA/abc/.def
< AAA/abc/.wxy/def
< AAA/abc/wxy/def
< AAA/abc/wxy/.def
< AAA/abc/wxy/tuv/def
< AAA/abc/xyz/def/wxy
< AAA/wxy/xyx/
< AAA/wxy/xyxfile
```

4. Include directories and files at any level that start with wxy, but do not include dot-files, dot-directories, or any files under the wxy directories (unless they start with wxy). However, subdirectories under wxy are included:

```
-N '*/wxy*' -E 'AAA/**'
```

Results:

```
AAA/abc/wxy/tuv/
AAA/abc/xyz/def/wxy
AAA/wxyfile
AAA/wxy/xyx/
< AAA/abc/def
< AAA/abc/.def
< AAA/abc/.wxy/def
< AAA/abc/wxy/def *
< AAA/abc/wxy/.def
< AAA/abc/wxy/tuv/def
< AAA/wxy/xyxfile
```

* Even though wxy is included, def is excluded because it is a file.

5. Include wxy directories and files at any level, even those that start with ".":

```
-N '*/wxy*' -N '*/wxy/**' -E 'AAA/**'
```

Results:

```
AAA/abc/wxy/def
AAA/abc/wxy/.def
AAA/abc/wxy/tuv/def
AAA/abc/xyz/def/wxy
AAA/wxyfile
AAA/wxy/xyx/
AAA/wxy/xyxfile
< AAA/abc/def
< AAA/abc/.def
< AAA/abc/.wxy/def
```

6. Exclude directories and files that start with wxy, but only those found at a specific location in the tree:

```
-E '/AAA/abc/wxy*' 
```

Results:

```
AAA/abc/def
AAA/abc/.def
AAA/abc/.wxy/def
AAA/abc/xyz/def/wxy
AAA/wxyfile
AAA/wxy/xyx/
AAA/wxy/xyxfile
< AAA/abc/wxy/def
< AAA/abc/wxy/.def
< AAA/abc/wxy/tuv/def
```

7. Include the wxy directory at a specific location, and include all its subdirectories and files, including those starting with ".":

```
-N 'AAA/abc/wxy/**' -E 'AAA/**' 
```

Results:

```
AAA/abc/wxy/def
AAA/abc/wxy/.def
AAA/abc/wxy/tuv/def
< AAA/abc/def
< AAA/abc/.def
< AAA/abc/.wxy/def
< AAA/abc/xyz/def/wxy
< AAA/wxyfile
< AAA/wxy/xyx/
< AAA/wxy/xyxfile
```

Surround patterns that contain wildcards with single quotation marks (' ') to prevent filter.

Symbolic link handling

When transferring files by using FASP (**ascp**, **ascp4**, or **async**), you can configure how the server and client handle symbolic links.

Note: Symbolic links are not supported on Windows. Server settings are ignored on Windows servers. If the transfer destination is a Windows computer, the only supported option that the client can use is **skip**.

Symbolic link-handling options and their behavior

- **Follow:** Follow a symbolic link and transfer the contents of the linked file or directory when the link target is in the user's docroot.
- **Follow_wide** (Server only): For downloads, follow a symbolic link and transfer the contents of the linked file or directory **even if the link target is outside of the user's docroot**. Use caution with this setting because it might allow transfer users to access sensitive files on the server.
- **Create** (Server only): If the client requests to copy symbolic links in an upload, create the symbolic links on the server.
- **None** (Server only): Prohibit clients from creating symbolic links on the server. With this setting, clients can request only to follow or skip symbolic links.
- **Copy** (Client only): Copy only the symbolic link. If a file with the same name exists at the destination, **the symbolic link does not replace the file**.
- **Copy+force** (Client only): Copy only the symbolic link. If a file with the same name exists at the destination, **the symbolic link replaces the file**. If the file of the same name at the destination is a symbolic link to a directory, it is not replaced.

Note: Ascp4 and Sync do not support the `copy+force` option.

- **Skip** (Client only): Skip-symbolic links. The link or the file to which it points are not transferred.

Symbolic link handling depends on the server configuration, the client-handling request, and the direction of transfer, as described in the following tables. Multiple values can be set on the server as a comma-delimited list, such as the default `follow, create`. In this case, the options are logically ordered based on the client's handling request.

Send from client to Server (Upload)

	Server setting = create, follow (default)	Server setting = create	Server setting = follow	Server setting = follow_wide	Server setting = none
Client setting = follow (default for ascp and ascp4)	Follow	Follow	Follow	Follow	Follow
Client setting = copy (default for async)	Copy	Copy	Skip	Skip	Skip
Client setting = copy+force	Copy and replace any existing files.	Copy and replace any existing files.	Skip	Skip	Skip
Client setting = skip	Skip	Skip	Skip	Skip	Skip

Receive to client from Server (Download)

	Server setting = create, follow (default)	Server setting = create	Server setting = follow	Server setting = follow_wide	Server setting = none
Client setting = follow (default for ascp and ascp4)	Follow	Skip	Follow	Follow even if the target is outside the user's docroot.	Skip
Client setting = copy (default for async)	Copy	Copy	Copy	Copy	Copy
Client setting = copy+force	Copy and replace any existing files.	Copy and replace any existing files.	Copy and replace any existing files.	Copy and replace any existing files.	Copy and replace any existing files.
Client setting = skip	Skip	Skip	Skip	Skip	Skip

Server and client configuration

Server configuration

To set symbolic link handling globally or per user, run the appropriate command:

```
# asconfigurator -x "set_node_data;symbolic_links,value"
# asconfigurator -x "set_user_data;user_name,username;symbolic_links,value"
```

For more information, see .

Client configuration

To specify symbolic link handling on the command line (with `ascp`, `ascp4`, or `async`), use `--symbolic-links=option`.

Creating SSH keys

Public key authentication (SSH Key) is a more secure alternative to password authentication that allows users to avoid entering or storing a password, or sending it over the network. Public key authentication uses the client computer to generate the key pair (a public key and a private key). The public key is then provided to the remote computer's administrator to be installed on that machine.

1. Create a `.ssh` directory in your home directory if it does not exist:

```
$ mkdir /home/username/.ssh
```

Go to the `.ssh` folder:

```
$ cd /home/username/.ssh
```

2. Generate an SSH key pair.

In the `.ssh` folder, use the **ssh-keygen** command to create a key pair.

```
# ssh-keygen -m key_format -t key_type
```

- For *key_format*, specify a format that is supported by the SSH server.
- For *key_type*, specify either RSA (`rsa`) or ECDSA (`ecdsa`).

At the prompt that appears for the key pair's file name, press ENTER to use the default name `id_rsa` or `id_ecdsa`, or enter a different name, such as your username. For a passphrase, either enter a password, or press return twice to leave it blank.

Note: When you run `ascp` in FIPS mode (`<fips_enabled>` is set to `true` in `aspera.conf`), and you use passphrase-protected SSH keys, you must either:

- a. Use keys that are generated by running **ssh-keygen** in a FIPS enabled system. Or,
- b. Convert existing keys to a FIPS compatible format by using a command such as:

```
# openssl pkcs8 -topk8 -v2 aes128 -in id_rsa -out new-id_rsa
```

3. As the root user, make sure that the SSH key is owned by the transfer user and that proper restrictive permissions are set. SSH keys must be readable only by the key owner.

Use the following command syntax, where *username* is the transfer username and *id_rsa* is the key pair's file name.

```
chown username /home/username/.ssh/id_rsa  
chmod 600 /home/username/.ssh/id_rsa
```

4. Retrieve the public key file.

The key pair is generated to your home directory's `.ssh` folder.

For example, assuming you generated the key with the default name `id_rsa`:

```
/home/username/.ssh/id_rsa.pub
```

Provide the public key file (for example, `id_rsa.pub`) to your server administrator so that it can be set up for your server connection.

5. Start a transfer by using public key authentication with the `ascp` command.

To transfer files by using public key authentication on the command line, use the option **-i** *private_key_file*.

For example,

```
$ ascp -T -l 10M -m 1M -i ~/.ssh/id_rsa myfile.txt jane@10.0.0.2:/space
```

In this example, you are connecting to the server (10.0.0.2, directory /space) with the user account jane and the private key ~/.ssh/id_rsa.

Reporting checksums

File checksums are useful for troubleshooting file corruption, allowing to determine at what point in the transfer file corruption occurred. Aspera servers can report source file checksums that are calculated in real time during transfer and then sent from the source to the destination.

To support checksum reporting, the transfer must meet both of the following requirements:

- Both the server and client computers must be running HSTS or HSTE.
- The transfer must be encrypted. Encryption is enabled by default.

The user on the destination can calculate a checksum for the received file and compare it (manually or programmatically) to the checksum reported by the sender. The checksum reported by the source can be retrieved in the destination logs, a manifest file, in IBM Aspera Console, or as an environment variable. Instructions for comparing checksums follow the instructions for enabling checksum reporting.

Checksum reporting is disabled by default. Enable and configure checksum reporting on the server by using the following methods:

- Edit aspera.conf with **asconfigurator**.
- Set ascp command line options (per-transfer configuration).

Command line options override the settings in aspera.conf.

Important: When checksum reporting is enabled, transfers of very large files (>TB) take a long time to resume because the entire file must be reread.

Overview of checksum configuration options

asconfigurator Option ascp Option	Description
file_checksum --file-checksum= <i>type</i>	<p>Enable checksum reporting and specify the type of checksum to calculate for transferred files.</p> <p>any - Allow the checksum format to be whichever format the client requests. (Default in aspera.conf) md5 - Calculate and report the MD5 checksum. sha1 - Calculate and report the SHA-1 checksum. sha256 - Calculate and report the SHA-256 checksum. sha384 - Calculate and report the SHA-384 checksum. sha512 - Calculate and report the SHA-512 checksum.</p> <p>Note: The default value for the ascp option is none, in which case the reported checksum is the one configured on the server, if any.</p>
file_manifest --file_manifest= <i>output</i>	<p>The file manifest is a file that contains a list of content that was transferred in a transfer session. The file name of the file manifest is automatically generated from the transfer session ID.</p> <p>When set to none, no file manifest is created. (Default)</p>

asconfigurator Option ascp Option	Description
	When set to text, a text file is generated that lists all files in each transfer session.
file_manifest_path --file_manifest_path= <i>path</i>	The location where manifest files are written. The location can be an absolute path or a path relative to the transfer user's home directory. If no path is specified (default), the file is generated under the destination path at the receiver, and under the first source path at the sender. Note: File manifests can be stored only locally. Thus, if you are using S3 or other nonlocal storage, you must specify a local manifest path.

Enabling checksum reporting by editing aspera.conf

To enable checksum reporting, run the following command:

```
# asconfigurator -x "set_node_data;file_checksum,checksum"
```

To enable and configure the file manifest where checksum report data is stored, run the following commands:

```
# asconfigurator -x "set_node_data;file_manifest,text"
# asconfigurator -x "set_node_data;file_manifest_path,filepath"
```

These commands create lines in `aspera.conf` as shown in the following example, where checksum type is **md5**, file manifest is enabled, and the path is `/tmp`.

```
<file_system>
  ...
  <file_checksum>md5</file_checksum>
  <file_manifest>text</file_manifest>
  <file_manifest_path>/tmp</file_manifest_path>
  ...
</file_system>
```

Enabling checksum reporting in an ascp session

To enable checksum reporting on a per-transfer-session basis, run `ascp` with the **--file-checksum=hash** option, where *hash* is `sha1`, `md5`, `sha-512`, `sha-384`, `sha-256`, or `none` (the default).

Enable the manifest with **--file-manifest=output** where *output* is either `text` or `none`. Set the path to the manifest file with **--file-manifest-path=path**.

For example,

```
# ascp --file-checksum=md5 --file-manifest=text --file-manifest-path=/tmp file
aspera_user_1@189.0.202.39: /destination_path
```

Setting up a processing script

An alternative to enabling and configuring the file manifest to collect checksum reporting is to set up a script to report the values. See [Automated execution of Lua scripts with transfer Events](#).

Comparing checksums

If you open a file that you downloaded with Aspera and find that it is corrupted, you can determine when the corruption occurred by comparing the checksum that is reported by Aspera to the checksums of the files on the destination and on the source.

1. Retrieve the checksum that was calculated by Aspera as the file was transferred.
 - If you specified a file manifest and file manifest path as part of an ascp transfer or Lua transfer event script, the checksums are in that file in the specified location.
 - If you specified a file manifest and file manifest path in the GUI or `aspera.conf`, the checksums are in a file that is named `aspera-transfer-transfer_id-manifest.txt` in the specified location.
2. Calculate the checksum of the corrupted file. This example uses the MD5 checksum method; replace MD5 with the appropriate checksum method if you use a different one.

```
# md5sum filepath
```

3. Compare the checksum reported by Aspera with the checksum that you calculated for the corrupted file.
 - If they do not match, then corruption occurred as the file was written to the destination. Download the file again and confirm that it is not corrupted. If it is corrupted, compare the checksums again. If they do not match, investigate the write process or attempt another download. If they match, continue to the next step.
 - If they match, then corruption might have occurred as the file was read from the source. Continue to the next step.
4. Calculate the checksums for the file on the source. These examples use the MD5 checksum method; replace MD5 with the appropriate checksum method if you use a different one.

Windows:

```
> CertUtil -hashfile filepath MD5
```

Mac OS X:

```
$ md5 filepath
```

Linux and Linux on z Systems:

```
# md5sum filepath
```

AIX:

```
# csum -h MD5 filepath
```

Solaris:

```
# digest -a md5 -v filepath
```

5. Compare the checksum of the file on the source with the one reported by Aspera.
 - If they do not match, then corruption occurred when the file was read from the source. Download the file again and confirm that it is not corrupted on the destination. If it is corrupted, continue to the next step.
 - If they match, confirm that the source file is not corrupted. If the source file is corrupted, replace it with an uncorrupted one, if possible, and then download the file again.

Client-Side Encryption-at-Rest (EAR)

Aspera clients can set their transfers to encrypt content that they upload to a server while it is in transit and stored on the server, a process known as client-side encryption-at-rest (EAR). The client specifies an encryption password and the files are uploaded to the server with a `.aspera-env` extension. Anyone

downloading these .aspera-env files must have the password to decrypt them, and decryption can occur as the files are downloaded or later once they are physically moved to a computer with no network connection.

Implementation Notes:

- Client-side and server-side EAR can be used simultaneously, in which case files are doubly encrypted on the server.
- Servers can require client-side encryption. In this case, transfers that do not use client-side EAR fail with the error message "Error: Server aborted session: Server requires content protection".
- Client-side encryption-at-rest is supported only for ascp transfers, and is not supported for ascp4 or async transfers.

Using Client-Side EAR

Client-side EAR can be set in the ascp command line.

First, set the encryption and decryption password as the environment variable ASPERA_SCP_FILEPASS:

```
# export ASPERA_SCP_FILEPASS=password
```

For uploads (--mode=send), use --file-crypt=encrypt. For downloads (--mode=recv), use --file-crypt=decrypt.

```
# ascp --mode=send --file-crypt=encrypt source_file user@host:/remote_destination
# ascp --mode=recv --file-crypt=decrypt user@host:/source_path/file.aspera-env local_destination
```

For more command line examples, see [“Ascp general examples” on page 38](#).

Note: When a transfer to HSTS falls back to HTTP or HTTPS, client-side EAR is no longer supported. If HTTP fallback occurs while uploading, then the files are NOT encrypted. If HTTP fallback occurs while downloading, then the files remain encrypted.

Encrypting and decrypting files outside of a transfer

For sensitive content, do not store unencrypted content on any computer with network access. Use an external drive to physically move encrypted files between computers. Desktop Client include the **asprotect** and **asunprotect** command line tools that can be used to encrypt and decrypt files.

- To encrypt a file before you move it to a computer with network access, run the following command:

```
# export ASPERA_SCP_FILEPASS=password;/opt/aspera/bin/asprotect -o file1.aspera-env file1
```

- To download client-side-encrypted files without decrypting them immediately, run the transfer without decryption enabled (do not specify --file-crypt=decrypt on the ascp command line).
- To decrypt encrypted files once they are on a computer with no network access, run the following command:

```
# export ASPERA_SCP_FILEPASS=password;/opt/aspera/bin/asunprotect -o file1 file1.aspera-env
```

Comparison of ascp and ascp4 options

Many command line options are the same for **ascp** and **ascp4**; however, some options are available for only one or the behavior of an option is different. The following table lists the options that are available only for **Ascp** or **Ascp4**, and the options that are available with both. If the option behavior is different, the **ascp** option has ** added to the end and the difference is described following the table.

Ascp	Ascp4
-6	

Ascp	Ascp4
-@[<i>range_low:range_high</i>]	
-A, --version	-A, --version
--apply-local-docroot	
-C <i>nodeid:nodecount</i>	
-c <i>cipher</i>	-c <i>cipher</i>
--check-sshfp= <i>fingerprint</i>	
	--chunk-size= <i>bytes</i>
	--compare= <i>method</i>
	--compression= <i>method</i>
	--compression-hint= <i>num</i>
-D -DD -DDD	
-d	
	--delete-before
--delete-before-transfer**	--delete-before-transfer**
--dest64	--dest64
-E <i>pattern</i>	-E <i>pattern</i>
-e <i>prepost_filepath</i>	
	--exclude-newer-than= <i>mtime</i>
	--exclude-older-than= <i>mtime</i>
-f <i>config_file</i>	-f <i>config_file</i>
	--faspmgr-io
--file-checksum= <i>hash</i>	
--file-list= <i>filepath**</i>	--file-list= <i>filepath**</i>
--file-pair-list= <i>filepath</i>	
-G <i>write_size</i>	
-g <i>read_size</i>	
-h, --help	-h, --help
-i <i>private_key_file_path**</i>	-i <i>private_key_file_path</i>
-K <i>probe_rate</i>	
-k {0 1 2 3}	-k {0 1 2 3}
--keepalive	--keepalive
-l <i>max_rate</i>	-l <i>max_rate</i>
-L <i>local_log_dir[:size]</i>	-L <i>local_log_dir[:size]</i>
-m <i>min_rate</i>	-m <i>min_rate</i>
	--memory= <i>bytes</i>

Ascp	Ascp4
	--meta-threads= <i>num</i>
--mode={send recv}	--mode={send recv}
--move-after-transfer= <i>archivedir</i>	
--multi-session-threshold= <i>threshold</i>	
-N <i>pattern</i>	-N <i>pattern</i>
	--no-open
	--no-read
	--no-write
-O <i>fasp_port</i>	-O <i>fasp_port</i>
--overwrite= <i>method**</i>	--overwrite= <i>method**</i>
-P <i>ssh-port</i>	-P <i>ssh-port</i>
-p	-p
--partial-file-suffix= <i>suffix</i>	--partial-file-suffix= <i>suffix</i>
--policy={fixed high fair low}	--policy={fixed high fair low}
--precalculate-job-size	
--preserve-access-time	
--preserve-acls= <i>mode</i>	
--preserve-creation-time	
--preserve-file-owner-gid	--preserve-file-owner-gid
--preserve-file-owner-uid	--preserve-file-owner-uid
--preserve-modification-time	
--preserve-source-access-time	
--preserve-xattrs= <i>mode</i>	
--proxy= <i>proxy_url</i>	
-q	-q
-R <i>remote_log_dir</i>	-R <i>remote_log_dir</i>
	--read-threads= <i>num</i>
	--remote-memory= <i>bytes</i>
--remote-preserve-acls= <i>mode</i>	
--remote-preserve-xattrs= <i>mode</i>	
--remove-after-transfer	
--remove-empty-source-directory	
	--resume (similar to -k)
--retry-timeout= <i>secs</i>	
-S <i>remote_ascp</i>	

Ascp	Ascp4
--save-before-overwrite	
	--scan-threads= <i>num</i>
--source-prefix= <i>prefix</i>	
--source-prefix64= <i>prefix</i>	
	--sparse-file
--src-base= <i>prefix</i>	--src-base= <i>prefix</i>
--symbolic-links= <i>method</i> **	--symbolic-links= <i>method</i> **
-T	-T
-u <i>user_string</i>	-u <i>user_string</i>
--user= <i>username</i>	--user= <i>username</i>
-v	
-W <i>token_string</i> @ <i>token_filepath</i>	
-w{ <i>r</i> <i>f</i> }	
-X <i>rexmsg_size</i>	-X <i>rexmsg_size</i>
-Z <i>dgram_size</i>	-Z <i>dgram_size</i>

Differences in option behavior

--delete-before-transfer

With **ascp4**, **--delete-before-transfer** can be used with URI storage. URI storage is not supported for this option in **ascp**.

--file-list

ascp automatically applies **-d** if the destination folder does not exist. With **ascp4**, you must specify **-d**. Otherwise, all the files in the file list are written to a single file.

-i (SSH key authentication)

With **ascp**, the argument for **-i** can be just the file name of the private key file and **ascp** automatically looks in the `.ssh` directory of the user's home directory. With **ascp4**, the full or relative path to the private key file must be specified.

--overwrite=*method*

The default overwrite method is `diff` for **ascp** and always for **ascp4**.

--symbolic-links

Both **ascp** and **ascp4** support `follow`, `copy`, and `skip`, but only **ascp** supports `copy+force`.

Ascp FAQs

Answers to some common questions about controlling transfer behavior, such as bandwidth usage, resuming files, and overwriting files.

1. How do I control the transfer speed?

You can specify a transfer policy that determines how a FASP transfer uses the network resource, and you can specify target and minimum transfer rates where applicable. In an **ascp** command, use the following flags to specify transfer policies that are fixed, fair, high, or low:

Policy	Command template
Fixed	<code>--policy=fixed -l target_rate</code>
Fair	<code>--policy=fair -l target_rate -m min_rate</code>
High	<code>--policy=high -l target_rate -m min_rate</code>
Low	<code>--policy=low -l target_rate -m min_rate</code>

The policies have the following characteristics:

- **high** - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, the transfer rate is twice as fast as a fair-policy transfer. The high policy requires maximum (target) and minimum transfer rates.
- **fair** - Adjust the transfer rate to fully use the available bandwidth up to the maximum rate. When congestion occurs, bandwidth is shared fairly by transferring at an even rate. The fair policy requires maximum (target) and minimum transfer rates.
- **low** - Adjust the transfer rate to use the available bandwidth up to the maximum rate. Similar to fair mode, but less aggressive when the bandwidth is shared with other network traffic. When congestion occurs, the transfer rate is reduced to the minimum rate until other traffic decreases.
- **fixed** - Attempt to transfer at the specified target rate, regardless of network or storage capacity. This can decrease transfer performance and cause problems on the target storage. Use the fixed policy only for specific contexts, such as bandwidth testing, otherwise, avoid the use of this policy. The fixed policy requires a maximum target rate.
- **aggressiveness** - The aggressiveness of transfers that are authorized by this access key in claiming available bandwidth. Value can be 0.00-1.00. For example, these values correspond to the policy option where a policy of high approximates to aggressiveness of 0.75, fair to 0.50 and low to 0.25. Aggressiveness can be used if you need to fine-tune the transfer policy.

2. What transfer speed must I expect? How do I know if something is wrong with the speed?

Aspera FASP transport has no theoretical throughput limit. Other than the network capacity, the transfer speed might be limited by rate settings and resources of the computers. To verify that your system's FASP transfer can fulfill the maximum bandwidth capacity, prepare a client computer to connect to a server, and test the maximum bandwidth.

Note: This test typically occupies most of a network's bandwidth. Perform this test on a dedicated file transfer line or during a time of low network activity.

On the client computer, start a transfer with fixed bandwidth policy. Start with a lower transfer rate and gradually increase the transfer rate toward the network bandwidth (for example, 1 MB, 5 MB, 10 MB, and so on). Monitor the transfer rate; at its maximum, it must be slightly below your available bandwidth:

```
$ ascp -l 1m source-file destination
```

To improve the transfer speed, also consider upgrading the following hardware components:

Component	Description
Hard disk	The I/O throughput, the disk bus architecture (such as RAID, IDE, SCSI, ATA, and Fibre Channel).
Network I/O	The interface card, the internal bus of the computer.
CPU	Overall CPU performance affects the transfer, especially when encryption is enabled.

3. How do I ensure that if the transfer is interrupted or fails to finish, it resumes without retransferring the files?

Use the **-k** flag to enable resume, and specify a resume rule:

- k 0** – Always retransfer the entire file.
- k 1** – Compare file attributes and resume if they match, and retransfer if they do not.
- k 2** – Compare file attributes and the sparse file checksum; resume if they match, and retransfer if they do not.
- k 3** – Compare file attributes and the full file checksum; resume if they match, and retransfer if they do not.

Corruption or deletion of the `.asp-meta` file that is associated with an incomplete transfer often result in a permanently unusable destination file even if the file transfer resumed and successfully transferred.

4. How does Aspera handle symbolic links?

The **ascp** command follows symbolic links by default. This can be changed by using `--symbolic-links=method` with the following options:

- `follow` - Follow symbolic links and transfer the linked files. (Default).
- `copy` - Copy only the alias file. If a file with the same name is found at the destination, the symbolic link is not copied.
- `copy+force` - Copy only the alias file. If a file (not a directory) with the same name is found at the destination, the alias replaces the file. If the destination is a symbolic link to a directory, it is not replaced.
- `skip` - Skip symbolic links. Do not copy the link or the file it points to.

Important: On Windows, the only option is `skip`.

Symbolic link handling also depends on the server configuration and the transfer direction. For more information, see [“Symbolic link handling”](#) on page 51.

5. What are my choices for overwriting files on the destination computer?

In **ascp**, you can specify the `--overwrite=method` rule with the following method options:

- `never` - Never overwrite the file. However, if the parent folder is not empty, its access, modify, and change times might still be updated.
- `always` - Always overwrite the file.
- `diff` - Overwrite the file if different from the source. If a complete file at the destination is the same as a file on the source, it is not overwritten. Partial files are overwritten or resumed depending on the resume policy.
- `diff+older` - Overwrite the file if older and also different than the source. For example, if the destination file is the same as the source, but with a different timestamp, it is not overwritten. Plus, if the destination file is different than the source, but newer, it is overwritten.
- `older` - Overwrite the file if its timestamp is older than the source timestamp.

Interaction with resume policy (-k): If the overwrite method is `diff` or `diff+older`, difference is determined by the resume policy (**-k {0|1|2|3}**). If **-k 0** or no **-k** is specified, the source and destination files are always considered different and the destination file is always overwritten. If **-k 1**, the source and destination files are compared based on file attributes (currently file size). If **-k 2**, the source and destination files are compared based on sparse checksum. If **-k 3**, the source and destination files are compared based on full checksum.

ascp4: Transferring from the command line

Ascp4 is a FASP transfer binary similar to **ascp** but it has different strengths as well as capabilities that are unavailable with **Ascp**.

Introduction to ascp4

Ascp4 is a FASP transfer binary that is optimized for sending large sets of individual files. The executable, `ascp4`, is similar to `ascp` and shares many of the same options and capabilities, in addition to data streaming capabilities.

Both `ascp4` and `Ascp` are automatically installed with IBM Aspera High-Speed Transfer Server, IBM Aspera High-Speed Transfer Endpoint, and IBM Aspera Desktop Client.

Ascp4 command reference

Supported environment variables, the general syntax, and command options for **ascp4** are described in the following sections. `ascp4` exits with a 0 on success or a 1 on error. The error code is logged in the `ascp4` log file.

ascp4 syntax

```
ascp4 options [[user@]srcHost:]source_file1[,source_file2,...] [[user@]destHost:]dest_path
```

User

The username of the Aspera transfer user can be specified as part of the as part of the source or destination, whichever is the remote server or with the `--user` option. If you do not specify a username for the transfer, the local username is authenticated by default.

Note: If you are authenticating on a Windows machine as a domain user, the transfer server strips the domain from the username. For example, `Administrator` is authenticated rather than `DOMAIN\Administrator`. Thus, you must specify the domain explicitly.

Source and destination paths

- If there are multiple source arguments, then the target path must be a directory.
- To describe file paths, use single quotation marks (' ') and forward slashes (/) on all platforms.
- To transfer to the transfer user's docroot, specify "." as the destination.
- Avoid the following characters in file names: / \ " : ' ? > < & * |.
- If the destination is a symbolic link, then the file is written to the target of the symbolic link. However, if the symbolic link path is a relative path to a file (not a directory) and a partial file name suffix is configured on the receiver, then the destination path is relative to the user's home directory. Files within directories that are sent to symbolic links that use relative paths are not affected.

URI paths: URI paths are supported, but only with the following restrictions:

- If the source paths are URIs, they must all be in the same cloud storage account. No docroot (download), local docroot (upload), or source prefix can be specified.
- If a destination path is a URI, no docroot (upload) or local docroot (download) can be specified.
- The special schemes `stdio://` and `stdio-tar://` are supported only on the client side. They cannot be used as an upload destination or download source.
- If required, specify the URI passphrase as part of the URI or set it as an environment variable (`ASPERA_SRC_PASS` or `ASPERA_DST_PASS`, depending on the direction of transfer).

UNC paths: If the server is Windows and the path on the server is a UNC path (a path that points to a shared directory or file on Windows operating systems) then it can be specified in an `ascp4` command by using one of the following conventions:

1. UNC path that uses backslashes (\)

If the client side is a Windows machine, the UNC path can be used with no alteration. For example, \\192.168.0.10\temp. If the client is not a Windows machine, every backslash in the UNC path must be replaced with two backslashes. For example, \\192.168.0.10\temp.

2. UNC path that uses forward slashes (/)

Replace each backslash in the UNC path with a forward slash. For example, if the UNC path is \\192.168.0.10\temp, change it to //192.168.0.10/temp. This format can be used with any client-side operating system.

Required file access and permissions

- Sources (for downloads) or destinations (for uploads) on the server must be in the transfer user's docroot or match one of the transfer user's file restrictions, otherwise the transfer stops and returns an error.
- The transfer user must have sufficient permissions to the sources or destinations, for example write access for the destination directory, otherwise the transfer stops and returns a permissions error.
- The transfer user must have authorization to do the transfer (upload or download), otherwise the transfer stops and returns a "management authorization refused" error.
- Files that are open for write by another process on a Windows source or destination cannot be transferred and return a "sharing violation" error. On Unix-like operating systems, files that are open for write by another process are transferred without reporting an error, but might produce unexpected results that depend on what data in the file is changed and when relative to the transfer.

Environment variables

If needed, you can set the following environment variables for use with an ascp4 session. The total size for environment variables depends on your operating system and transfer session. Each environment variable value must not exceed 4096 characters.

ASPERA_SCP_PASS=*password*

The password that is used for SSH authentication of the transfer user.

ASPERA_SCP_TOKEN=*token*

Set the transfer user authorization token. Ascp4 currently supports transfer tokens, which must be created by using **astokengen** with the **--full-paths** option. For more information, see Transfer Token Generation (astokengen) in the [IBM Aspera High-Speed Admin Guide](#)

ASPERA_SCP_COOKIE=*cookie*

A cookie string that is passed to monitoring services.

ASPERA_SRC_PASS=*password*

The password that is used to authenticate to a URI source.

ASPERA_DST_PASS=*password*

Set the password that is used to authenticate to a URI destination.

ASPERA_LOCAL_TOKEN=*token*

A token that authenticates the user to the client (in place of SSH authentication).

Note: If the local token is a basic or bearer token, the access key settings for cipher and preserve_time are not respected and the server settings are used. To set the cipher and timestamp preservation options as a client, set them in the command line.

Ascp4 options

-A, --version

Display version and license information.

-c {aes128|aes192|aes256|none}

Encrypt in-transit file data by using the specified cipher. This option overrides the <encryption_cipher> setting in aspera.conf.

--check-sship=fingerprint

Compare *fingerprint* to the server SSH host key fingerprint that is set with `<ssh_host_key_fingerprint>` in `aspera.conf`. Aspera fingerprint convention is to use a hex string without the colons; for example, `f74e5de9ed0d62feaf0616ed1e851133c42a0082`. For more information on SSH host key fingerprints, see [.IBM Aspera High-Speed Admin Guide](#).

--chunk-size=bytes

Perform storage read and write operations with the specified buffer size. Also, use the buffer size as an internal transmission and compression block. Valid range: 4 KB - 128 MB. For transfers with object storage, use `--chunk-size=1048576` if chunk size is not configured on the server to ensure that the chunk size of `ascp4` and `Trapd` match.

--compare={size|size+mtime|md5|md5-sparse|sha1|sha1-sparse}method

When using `--overwrite` and `--resume`, compare files with the specified method. If the `--overwrite` method is `diff` or `diff+older`, the default `--compare` method is `size`.

--compression={none|zlib|lz4}

Compress file data inline. Default: `lz4`. If set to `zlib`, `--compression-hint` can be used to set the compression level.

--compression-hint=num

Compress file data to the specified level when `--compression` is set to an option that accepts compression level settings (currently only `zlib`). A lower value results in less, but faster, data compression (0 = no compression). A higher value results in greater, slower compression. Valid values are -1 to 9, where -1 is "balanced". Default: -1.

-D | -DD | -DDD

Log at the specified debug level. With each **D**, an additional level of debugging information is written to the log. This option is not supported if the transfer user is restricted to `aspsell`.

--delete-before, --delete-before-transfer

Before transfer, delete files that exist at the destination but not at the source. The source and destination arguments must be directories that have matching names. Objects on the destination that have the same name but different type or size as objects on the source are not deleted. Do not use with multiple sources or `--keepalive`.

--dest64

Indicate that the destination path or URI is base64 encoded.

-E pattern

Exclude files or directories from the transfer based on the specified pattern. Use the `-N` option (include) to specify exceptions to `-E` rules. Rules are applied in the order in which they are encountered, from left to right. The following symbols can be used in the pattern:

- `*` (asterisk) represents zero or more characters in a string, for example `*.tmp` matches `.tmp` and `abcde.tmp`.
- `?` (question mark) represents a single character. For example, `t?p` matches `tmp` but not `temp`.

Note: When filtering rules are found in `aspera.conf`, they are applied *before* rules that are given on the command line (`-E` and `-N`).

--exclude-newer-than=mtime

--exclude-older-than=mtime

Exclude files (but not directories) from the transfer based on when the file was last changed. Positive *mtime* values are used to express time, in seconds, since the original system time (usually `1970-01-01 00:00:00`). Negative *mtime* values (prefixed with `"-"`) are used to express the number of seconds before the current time.

-f config_file

Read Aspera configuration settings from *config_file* rather than `aspera.conf` (the default).

--faspmgr-io

Run `ascp4` in API mode by using FASP manager I/O. `Ascp4` reads FASPMGR4 commands from management and runs them. The FASPMGR4 commands are `PUT/WRITE/STOP` to open/write/close on a file on the server.

--file-crypt={encrypt|decrypt}

Encrypt files (when sending) or decrypt files (when receiving) for client-side encryption-at-rest (EAR). Encrypted files have the file extension `.aspera-env`. This option requires the encryption and decryption passphrase to be set with the environment variable `ASPERA_SCP_FILEPASS`. If a client-side encrypted file is downloaded with an incorrect password, the download is successful, but the file remains encrypted and still has the file extension `.aspera-env`. For more information, see [“Client-Side Encryption-at-Rest \(EAR\)”](#) on page 56.

--file-list=filepath

Transfer the files and directories that are listed in *filepath*. Only the files and directories are transferred. The path information is not preserved at the destination. Each source must be specified on a separate line, for example:

```
sic
sic2
...
sicN
```

To read a file list from standard input, use `"-"` in place of *filepath* (as `ascp4 --file-list=- ...`). UTF-8 file format is supported. Use with `-d` if the destination folder does not exist.

Restrictions:

- Paths in file lists cannot use `user@host:filepath` syntax. You must use `--user` with `--file-list`.
- Only one `--file-list` option is allowed per `ascp4` session. If multiple file lists are specified, all but the last are ignored.
- Only files and directories from the file list are transferred, and any additional source files or directories that are specified on the command line are ignored.
- If more than one read thread is specified (default is 2) for a transfer that uses `--file-list`, the files in the file list must be unique. Duplicates can produce unexpected results on the destination.
- Because multiple sources are being transferred, the destination must be a directory.
- If the source paths are URIs, the size of the file list cannot exceed 24 KB.

For large file lists (~100 MB+), use with `--memory` to increase available buffer space.

--file-manifest={none|text}

Generate a list of all transferred files when set to `text`. Requires `--file-manifest-path` to specify the location of the list. (Default: `none`)

--file-manifest-path=directory

Save the file manifest to the specified location when using `--file-manifest=text`. File manifests must be stored locally. For cloud or other nonlocal storage, specify a *local* manifest path.

--file-manifest-inprogress-suffix=suffix

Apply the specified suffix to the file manifest's temporary file. For use with `--file-manifest=text`. (Default suffix: `.aspera-inprogress`)

-h, --help

Display the usage summary.

--host=host

Transfer to the specified hostname or address. Requires `--mode`. This option can be used instead of specifying the host as part of the file name (as `hostname:filepath`).

-i private_key_file

Authenticate the transfer by using public key authentication with the specified SSH private key file (specified with a full or relative path). The private key file is typically in the directory `$HOME/.ssh/`. If multiple `-i` options are specified, only the last one is used.

-k {0|1|2|3}

Enable the resumption of partially transferred files at the specified resume level. Default: 0. This option must be specified for your first transfer or it does not work for subsequent transfers. Resume levels:

- -k 0: Always retransfer the entire file (same as `--overwrite=always`).
- -k 1: Compare file modification time and size and resume if they match (same as `--overwrite=diff --compare=size --resume`).
- -k 2: Compare-sparse checksum and resume if they match (same as `--overwrite=diff --compare=md5-sparse --resume`).
- -k 3: Compare-full checksum and resume if they match (same as `--overwrite=diff --compare=md5 --resume`).

--keepalive

Enable ascp4 to run in persistent mode. This option enables a persistent session that does not require that source content and its destination are specified at execution. Instead, the persistent session reads source and destination paths through `mgmt` commands. Requires `--mode` and `--host`.

-L local_log_dir[size]

Log to the specified directory on the client machine rather than the default directory. Optionally, set the size of the log file (default 10 MB).

-l max_rate

Transfer at rates up to the specified target rate. Default: 10 Mbps. This option accepts suffixes G/g for Giga, M/m for Mega, K/k for Kilo, and P/p/% for percentage. Decimals are allowed. If this option is not set by the client, the server target rate is used. If a rate cap is set in the local or server `aspera.conf`, then the rate does not exceed the cap.

For streaming, the value must be equal to or greater than the bit rate of the video.

-m min_rate

Attempt to transfer no slower than the specified minimum transfer rate. Default: 0. If this option is not set by the client, then the server's `aspera.conf` setting is used. If a rate cap is set in the local or server `aspera.conf`, then the rate does not exceed the cap.

--memory=bytes

Allow the local ascp4 process to use no more than the specified memory. Default: 256 MB. See also `--remote-memory`.

--meta-threads=num

Use the specified number of directory creation threads (receiver only). Default: 2.

--mode={send|recv}

Transfer in the specified direction: `send` or `recv` (receive). Requires `--host`.

--move-after-transfer=archivedir

Move source files and copy source directories to `archivedir` after they are successfully transferred. Because directories are copied, the original source tree remains in place. The transfer user must have write permissions to the `archivedir`. The `archivedir` is created if it does not already exist. If the archive directory cannot be created, the transfer proceeds and the source files remain in their original location.

To preserve portions of the file path above the transferred file or directory, use this option with `--src-base`.

To remove empty source directories (except those specified as the source to transfer), use this option with `--remove-empty-directories`.

Restrictions:

- `archivedir` must be on the same file system as the source. If the specified archive is on a separate file system, it is created (if it does not exist), but an error is generated and files are not moved to it.
- For cloud storage, `archivedir` must be in the same cloud storage account as the source and must not already contain files with the same name (the existing files cannot be overwritten and the archiving fails).
- If the source is on a remote system (ascp is run in receive mode), `archivedir` is subject to the same docroot restrictions as the remote user.

- **--remove-after-transfer** and **--move-after-transfer** are mutually exclusive. Using both in the same session generates an error.
- Empty directories are not saved to *archivedir*.
- When used with **--remove-empty-directories** and **--src-base**, scanning for empty directories starts at the specified source base and proceeds down any subdirectories. If no source base is specified and a file path (as opposed to a directory path) is specified, then only the immediate parent directory is removed (if empty) after the source files were moved.

-N pattern

Protect (include) files or directories from exclusion by any **-E** (exclude) options that follow it. Files and directories are specified by using *pattern*. Each option-plus pattern is a *rule*. Rules are applied in the order (left to right) in which they're encountered. Thus, **-N** rules protect files only from **-E** rules that follow them. Create patterns by using standard globbing wildcards and special characters such as:

- ***** (asterisk) represents zero or more characters in a string, for example, `*.tmp` matches `.tmp` and `abcde.tmp`.
- **?** (question mark) represents any single character, for example, `t?p` matches `tmp` but not `temp`.

Note: Filtering rules can also be specified in `aspera.conf`. Rules that are found in `aspera.conf` are applied *before* any **-E** and **-N** rules that are specified on the command line.

--no-open

In test mode, do not open or write the contents of destination files.

--no-read

In test mode, do not read the contents of source files.

--no-write

In test mode, do not write the contents of destination files.

-O fasp_port

Use the specified UDP port for FASP transfers. Default: 33001.

--overwrite={always|never|diff|diff+older|older}

Overwrite files at the destination with source files of the same name based on the *method*. Default: `always`. Use with `--compare` and `--resume`. *method* can be the following:

- `always` – Always overwrite the file.
- `never` – Never overwrite the file. If the destination contains partial files that are older or the same as the source files and `--resume` is enabled, the partial files resume transfer. Partial files with checksums or sizes that differ from the source files are not overwritten.
- `diff` – Overwrite the file if it is different from the source, depending on the compare method (default is `size`). If the destination is object storage, `diff` has the same effect as `always`.

If `resume` is not enabled, partial files are overwritten if they are different from the source, otherwise they are skipped. If `resume` is enabled, only partial files with different sizes or checksums from the source are overwritten; otherwise, files resume.

- `diff+older` – Overwrite the file if it is older and different from the source, depending on the compare method (default is `size`). If `resume` is not enabled, partial files are overwritten if they are older and different from the source, otherwise they are skipped. If `resume` is enabled, only partial files that are different and older than the source are overwritten, otherwise they are resumed.
- `older` – Overwrite the file if its timestamp is older than the source timestamp.

-P ssh-port

Use the specified TCP port to initiate the FASP session. (Default: 22).

-p

Preserve file timestamps for access and modification time. Equivalent to setting `--preserve-modification-time`, `--preserve-access-time`, and `--preserve-creation-time`. Timestamp support in object storage varies by provider; consult your object storage documentation to determine which settings are supported.

On Windows, modification time might be affected when the system automatically adjusts for Daylight Savings Time (DST). For more information, see [Daylight saving time help and support](#).

--partial-file-suffix=suffix

Enable the use of partial files for files that are in transit, and set the suffix to add to names of partial files. (The suffix does not include a ".", as for a file extension, unless explicitly specified as part of the suffix). This option takes effect only when set on the receiver side. When the transfer is complete, the suffix is removed. (Default: suffix is null; use of partial files is disabled).

--policy={fixed|high|fair|low}

Transfer according to the specified policy:

- **fixed** – Attempt to transfer at the specified target rate, regardless of network capacity. Content is transferred at a constant rate and the transfer finishes in a guaranteed time. The **fixed** policy can consume most of the network's bandwidth and is not recommended for most types of file transfers. This option requires a maximum (target) rate value (-l).
- **high** – Adjust the transfer rate to fully utilize the available bandwidth up to the maximum rate. When congestion occurs, the transfer rate is twice as fast as transfer with a fair policy. This option requires maximum (target) and minimum transfer rates (-l and -m).
- **fair** – Adjust the transfer rate to fully utilize the available bandwidth up to the maximum rate. When congestion occurs, bandwidth is shared fairly by transferring at an even rate. This option requires maximum (target) and minimum transfer rates (-l and -m).
- **low** – Adjust the transfer rate to use the available bandwidth up to the maximum rate. Similar to fair mode, but less aggressive when sharing bandwidth with other network traffic. When congestion occurs, the transfer rate is reduced to the minimum rate until other traffic decreases.

If --policy is not set, ascp4 uses the server-side policy setting (fair by default).

--precalculate-job-size

Calculate the total size before starting the transfer. The server-side `pre_calculate_job_size` setting in `aspera.conf` overrides this option.

--preserve-access-time

Preserve the file timestamps (currently the same as -p).

--preserve-creation-time

Preserve the file timestamps (currently the same as -p).

--preserve-file-owner-gid

--preserve-file-owner-uid

(Linux, UNIX, and macOS only) Preserve the group information (gid) or owner information (uid) of the transferred files. These options require that the transfer user is authenticated as a superuser.

--preserve-modification-time

Preserve the file timestamps (currently the same as -p).

--preserve-source-access-time

Preserve the file timestamps (currently the same as -p).

-q

Run ascp4 in quiet mode. This option disables the progress display.

-R remote_log_dir

Log to the specified directory on the remote host rather than the default directory.

Note: Client users that are restricted to `aspshe11` are not allowed to use this option.

--read-threads=num

Use the specified number of storage read threads (sender only). Default: 2. To set write threads on the receiver, use `--write-threads`.

Note: If more than one read thread is specified for a transfer that uses `--file-list`, the files in the file list must be unique. Duplicates can produce unexpected results on the destination.

--remote-memory=bytes

Allow the remote ascp4 process to use no more than the specified memory. Default: 256 MB.

--remove-empty-directories

Remove empty source directories once the transfer is completed successfully, but do not remove a directory that is specified as the source argument. To also remove the specified source directory, use **--remove-empty-source-directory**. Directories can be emptied by using **--move-after-transfer** or **--remove-after-transfer**. Scanning for empty directories starts at the src-base and proceeds down any subdirectories. If no source base is specified and a file path (as opposed to a directory path) is specified, then only the immediate parent directory is scanned and removed if it's empty following the move of the source file.

Note: Do not use this option if multiple processes (ascp4 or other) might access the source directory at the same time.

--resume

Resume a transfer rather than retransferring the content if partial files are present at the destination and they do not differ from the source file based on the **--compare** method. If the source and destination files do not match, then the source file is retransferred. See **-k** for another way to enable resume.

--scan-threads=num

Use the specified number of directory scan threads (sender only). Default: 2.

-SSH

Use an external SSH program instead of the built-in libssh2 implementation to establish the connection with the remote host. The wanted SSH program must be defined in the environment's PATH variable. To enable debugging of the SSH process, use the **-DD** and **--ssh-arg=-vv** options with ascp4.

--ssh-arg=ARG

Add ARG to the command line arguments passed to the external SSH program (implies that you use -SSH). This option might be repeated as needed to supply multiple separate SSH arguments. The order is preserved. The ARG elements are inserted before any key files supplied to ascp4, and before the user and host argument.

--sparse-file

Enable ascp4 to write sparse files to disk. This option prevents ascp4 from writing zero content to disk for sparse files; **ascp4** writes a block to disk if even 1 bit is set in that block. If no bits are set in the block, ascp4 does not write the block (ascp4 blocks are 64 KB by default).

--src-base=prefix

Strip the specified prefix from each source path. The remaining portion of the source path is kept intact at the destination. Available only in send mode.

Use with URIs: The **--src-base** option performs a character-to-character match with the source path. For object storage source paths, the prefix must specify the URI in the same manner as the source paths. For example, if a source path includes an embedded passphrase, the prefix must also include the embedded passphrase otherwise it does not match.

--symbolic-links={follow|copy|skip}

Handle symbolic links by using the specified method. On Windows, the only option is skip. On other operating systems, this option takes following values:

- **follow** – Follow symbolic links and transfer the linked files. (Default)
- **copy** – Copy only the alias file. If a file with the same name exists on the destination, the symbolic link is not copied.
- **skip** – Skip symbolic links. Do not copy the link or the file it points to.

-T

Disable in-transit encryption for maximum throughput.

-u *user_string*

Define a user string for Lua scripts that can be run with transfer events. See [Transfer session data accessible to scripts](#).

--user=*username*

Authenticate the transfer by using the specified username. Use this option instead of specifying the username as part of the destination path (as *user@host:file*).

Note: If you are authenticating on a Windows machine as a domain user, the transfer server strips the domain from the username. For example, Administrator is authenticated rather than DOMAIN\Administrator. Thus, you must specify the domain explicitly.

--worker-threads=*num*

Use the specified number of worker threads for deleting files. On the receiver, each thread deletes one file or directory at a time. On the sender, each thread checks for the presences of one file or directory at a time. Default: 1.

--write-threads=*num*

Use the specified number of storage write threads (receiver only). Default: 2. To set read threads on the sender, use `--read-threads`.

For transfers to object or HDFS storage, write threads cannot exceed the maximum number of jobs that are configured for Trapd. Default: 15. To use more threads, open `/opt/aspera/etc/trapd/trap.properties` on the server and set `aspera.session.upload.max-jobs` to a number larger than the number of write threads. For example,

```
# Number of jobs allowed to run in parallel for uploads.
# Default is 15
aspera.session.upload.max-jobs=50
```

-X *rexmsg_size*

Limit the size of retransmission requests to no larger than the specified size, in bytes. Max: 1440.

-Z *dgram_size*

Use the specified datagram size (MTU) for FASP transfers. Range: 296 - 65535 bytes. Default: the detected path MTU.

As of version 3.3, datagram size can be specified on the server by setting `<datagram_size>` in `aspera.conf`. The server setting overrides the client setting, unless the client uses a version of ascp that is older than 3.3, in which case the client setting is used. If the pre-3.3 client does not set **-Z**, the datagram size is the discovered MTU and the server logs the message "LOG Peer client doesn't support alternative datagram size".

Ascp4 transfers with object storage

Files that are transferred with object storage are sent in chunks through the Aspera Trapd service. By default, **ascp4** uses 64 KB chunks and Trapd uses 1 MB chunks; this mismatch in chunk size can cause **ascp4** transfers to fail.

To avoid this problem, take one of the following actions:

1. Set the chunk size (in bytes) in the server's `aspera.conf`. This value is used by both **ascp4** and **Trapd**, so the chunk sizes match.

To set a global chunk size, run the following command:

```
# asconfigurator -x "set_node_data;transfer_protocol_options_chunk_size,value"
```

Where *value* is between 256 KB (262144 bytes) and 1 MB (1048576 bytes).

To set a chunk size for the user, run the following command:

```
# asconfigurator -x "set_user_data;user_name,
username;transfer_protocol_options_chunk_size,value"
```

2. Set the chunk size in the client's `aspera.conf` to the **Trapd** chunk size.

If Trapd is using the default chunk size, run the following command to set the chunk size to 1 MB:

```
# asconfigurator -x "set_node_data;transfer_protocol_options_chunk_size,1048576"
```

3. Set the chunk size in the client command line.

Run the ascp4 session with the chunk size setting: `--chunk-size=1048576`.

Ascp4 examples

The command options for **ascp4** are generally similar to those for **ascp**. The following examples demonstrate options that are unique to Ascp4. These options enable reading management commands and enable read/write concurrency.

For Ascp examples, see “Ascp command reference” on page 23. See “[Comparison of ascp and ascp4 options](#)” on page 57 for differences in option availability and behavior.

- **Read FASP4 management commands**

Read management commands V4 from management port 5000 and run the management commands. The management commands versions 4 are PUT, WRITE, and CLOSE.

```
# ascp4 -L /tmp/client-logs -R /tmp/server-logs --faspmgr-io -M 5000 localhost:/tmp
```

- **Increase concurrency**

The following command runs ascp4 with two scan threads and eight read threads on the client, and eight meta threads and 16 write threads on the server.

```
# ascp4 -L /tmp/logs -R /tmp/logs -l1g --scan-threads=2 --read-threads=8 --write-threads=16 --meta-threads=8 /data/100K aspera@10.0.113.53:/data
```

Built-in I/O provider

Input/output provider are library modules that abstract I/O scheme in **Ascp4** architecture. **Ascp4** has the following built-in I/O provider:

- File (as a simple path or `file://path`)

File provider

The local disk can be specified for ascp4 I/O by using a simple path or URL that starts with `file`. The following paths identify the same file (`/test/ascp4.log`) on the disk:

```
file:///test/ascp4.log
/test/ascp4.log
file://localhost:/test/ascp4.log
```

Similarly, the following URLs identify the same file (`test/ascp4.log`) on the disk:

```
file:///test/ascp4.log
test/ascp4.log
```

Appendix

Restarting Aspera services

When you change product settings, you might need to restart certain Aspera services in order for the new values to take effect.

IBM Aspera Central

If `asperacentral` is stopped, or if you modified the `<central_server>` or `<database>` sections in `aspera.conf`, then you need to restart the service.

Run the following command in a Terminal window to restart `asperacentral`:

```
# systemctl restart asperacentral
```

Or for Linux systems that use `init.d`:

```
# service asperacentral restart
```

IBM Aspera NodeD

Restart `asperanoded` if you modified any setting in `aspera.conf`.

Run the following commands to restart `asperanoded`:

```
# systemctl restart asperanoded
```

Or for Linux systems that use `init.d`:

```
# service asperanoded restart
```

Testing and optimizing transfer performance

To verify that your system's FASP transfer is reaching the target rate and can use the maximum bandwidth capacity, prepare a client to connect to an Aspera server. For these tests, you can transfer an existing file or file set, or you can transfer uninitialized data in place of a source file, which you can delete at the destination, eliminating the need to read from or write to disk and saving disk space.

Using `faux:///` as a test source or destination

You can use `faux:///` as the argument for the source or destination of an `ascp` session to test data transfer without reading from disk on the source and writing to disk on the target. The argument takes different syntax that depends on if you are using it as a mock source file or mock source directory.

Note: If you set large file sizes (> PB) in a `faux:///` source, use `faux://` as a target on the destination because most computers do not have enough system memory available to handle files of this size, and your transfer might fail.

Faux Source File

To send random data in place of a source file (do not read from the source), you can specify the file as `faux:///fname?fsize`. `fname` is the name assigned to the file on the destination and `fsize` is the number of bytes to send. `fsize` can be set with modifiers (k/K, m/M, g/G, t/T, p/P, or e/E) to a maximum of 7×2^{60} bytes (7 EiB).

For example,

```
# ascp --mode=send --user=username --host=host_ip_address faux:///fname?fsize target_path
```

Faux source directory

In some cases, you might want to test the transfer of an entire directory, rather than a single file. Specify the faux source directory with the following syntax:

```
faux:///dirname?file=file&count=count&size=size&inc=increment&seq=sequence&buf_init=buf_option
```

Where:

- *dirname* is a name for the directory (required).
- *file* is the root for file names. The default is *file* (optional).
- *count* is the number of files in the directory (required).
- *size* is the size of the first file in the directory, default 0 (optional). *size* can be set with modifiers (k/K, m/M, g/G, t/T, p/P, or e/E) to a maximum of 7×2^{60} bytes (7 EiB).
- *increment* is the increment of bytes to use to determine the file size of the next file, default 0 (optional).
- *sequence* is how to determine the size of the next file: "sequential" or "random". Default is "sequential" (optional). When set to "sequential", file size is calculated as:

```
size + ((N - 1) * increment)
```

Where *N* is the file index; for the first file, *N* is one.

When set to "random", file size is calculated as:

```
size +/- (rand * increment)
```

Where *rand* is a random number between zero and one. If necessary, *increment* is automatically adjusted to prevent the file size from being negative.

For both options, *increment* is adjusted to prevent the file size from exceeding 7×2^{60} bytes.

- *buf_option* is how faux source data is initialized: "none", "zero", or "random". Default is "zero". "none" is not allowed for downloads (Ascp run with `--mode=recv`).

When the defaults are used, Ascp sends a directory that is named *dirname* and that contains *count* number of zero-byte files that are named *file_count*.

For example, to transfer a faux directory (*mydir*) that contains 1 million files to */tmp* on 10.0.0.2, and the files in *mydir* are named "**testfile**" and file size increases sequentially from 0 - 2 MB by an increment of 2 bytes:

```
# ascp --mode=send --user=username --host=10.0.0.2 faux:///mydir?
file=testfile&count=1m&size=0&inc=2&seq=sequential /tmp
```

Faux target

To send data but not save the results to disk at the destination (do not write to the target), specify the target as `faux://`.

For example, to send a real file to a faux target, run the following command:

```
# ascp --mode=send --user=username --host=host_ip_address source_file1 faux://
```

To send random data to a faux target, run the following command:

```
# ascp --mode=send --user=username --host=host_ip_address faux:///fname?fsize faux://
```

Testing transfer performance

1. Start a transfer with fair transfer policy and compare the transfer rate to the target rate.

On the client computer, open the user interface and start a transfer (either from the GUI or command line). Click **Details** to open the **Transfer Monitor**.

To leave more network resources for other high-priority traffic, use the **Fair** policy and adjust the target rate and minimum rate by sliding the arrows or entering values.

2. Test the maximum bandwidth.

Note: This test typically occupies most of the network's bandwidth. Perform it on a dedicated file transfer line or during a time of low network activity.

Use **Fixed** policy for the maximum transfer speed. Start with a lower transfer rate and increase gradually toward the network bandwidth.

Hardware upgrades for better performance

To improve the transfer speed, you can also upgrade the related hardware components:

Component	Description
Hard disk	The I/O throughput, the disk bus architecture (such as RAID, IDE, SCSI, ATA, and Fibre Channel).
Network I/O	The interface card, the internal bus of the computer.
CPU	Overall CPU performance affects the transfer, especially when encryption is enabled.

Log files

The application log file includes detailed transfer information and can be useful for review and support requests. You can configure log rotation and redirect Aspera logging so that it is not recorded in the system log file.

Viewing logs and setting log preferences

To view the log, from the GUI, click **Tools > View Log**.



Note: To view logs from the command line in Linux, you must have a functional web-browser or other default application for opening HTML files.

To set the logging level for transfers, open the **My Preferences** dialog by clicking **Tools > Preferences** or by clicking **Preferences** in the menu bar of the application window.



The five logging levels to select from are: **Off**, **Error**, **Warn**, **Info**, and **Debug**. The system default is **Info**.



Redirecting Aspera logging to a different location

On Linux systems, the application transfer logs are recorded in the system log file. Instead of mixing Aspera logging with system logging, you might want to redirect Aspera logging to a separate log file of your choice.

RedHat, CentOS, and Debian

On RedHat, CentOS, and Debian, the transfer logs are recorded in the following log file: `/var/log/messages`

To redirect Aspera logging, modify `/etc/syslog.conf` (`/etc/rsyslog.conf` in the case of Red Hat or CentOS 6.XA) and add `local2.none` to the `/var/log/messages` line. For example, if you have the following line:

```
*.info;mail.none;authpriv.none;cron.none /var/log/messages
```

Change it to:

```
*.info;mail.none;authpriv.none;cron.none;local2.none /var/log/messages
```

Next, forward `local2.info` log messages to your new file. For example, to write to `/var/log/aspera.log`, add the following line just below the line you modified in the prior step:

```
local2.info -/var/log/aspera.log
```

The log file name must be separated from the log facility (**local2.info**) by tab characters, not spaces and be preceded by a hyphen. The hyphen before the log file name allows for asynchronous logging.

Next, restart the syslog daemon to have it load the new configuration:

```
# service syslog restart
```

In the case of Red Hat or CentOS 6.X:

```
# service rsyslog restart
```

Your Aspera log messages now appear in `/var/log/aspera.log` instead of `/var/log/messages`.

SLES (Suse) systems

On SLES (Suse) systems, the transfer logs are recorded in the following system log file: `/var/log/localmessages`

To redirect Aspera logging, locate the following section in `/etc/syslog-ng/syslog-ng.conf`:

```
filter f_local { facility(local0, local1, local2, local3, local4, local5, local6, local7); };
destination localmessages { file("/var/log/localmessages"); };
log { source(src); filter(f_local); destination(localmessages); };
```

Modify the section as follows:

```
filter f_local { facility(local0, local1, local3, local4, local5, local6, local7); };
filter f_aspera { facility(local2); };

destination localmessages { file("/var/log/localmessages"); };
log { source(src); filter(f_local); destination(localmessages); };

destination asperalog { file("/var/log/aspera.log"); };
log { source(src); filter(f_aspera); destination(asperalog); };
```

Then, run the following command:

```
# rcsyslog restart
```

Your Aspera log messages now appear in `/var/log/aspera.log` instead of `/var/log/localmessages`.

To test this, run the following commands:

```
# logger -p local2.info aspera test
# cat /var/log/aspera.log
```

The **cat** command must display something similar to the following:

```
Jun 13 10:30:33 linux-kua5 root: aspera test
```

Rotating your Aspera log file

There are several ways to rotate Aspera logs in Linux:

1. Add `/var/log/aspera.log` to `/etc/logrotate.d/syslog`.
2. Create an entry for `aspera.log` in `/etc/logrotate.conf`.
3. Create a separate configuration file for `aspera.log` in `/etc/logrotate.d/`.

The first option rotates your logs with the system logs (usually once a week, is compressed, and saving the last 10 logs). On some servers, there is so much traffic that the logs need to be rotated more often than once a week, in which case option 2 or 3 must be used.

1. Add `/var/log/aspera.log` to the entries in `/etc/logrotate.d/syslog`, as follows:

```
/var/log/messages /var/log/secure /var/log/maillog /var/log/spooler /var/log/boot.log /var/log/
cron /var/log/aspera.log
{
    sharedscripts
    postrotate
        /bin/kill -HUP `cat /var/run/syslogd.pid 2> /dev/null` 2> /dev/null || true
        /bin/kill -HUP `cat /var/run/rsyslogd.pid 2> /dev/null` 2> /dev/null || true
    endscript
}
```

2. Edit `/etc/logrotate.conf` by adding the configuration after the line `"# system-specific logs may also be configured here."` The following example compresses and rotates 10 logs whenever `/var/log/aspera.log` reaches 100 MB. After log rotation is complete, it runs whatever scripts are specified by `postrotate ... endscript`.

```
/var/log/aspera.log {
    rotate 10
    size 100M
    create 664 root
    postrotate
        /usr/bin/killall -HUP syslogd
    endscript
    compress
}
```

The following example compresses and rotates 10 logs once daily. Instead of moving the original log file and creating a new one, the **copytruncate** option tells **logrotate** to first copy the original log file, then truncate it to zero bytes.

```
/var/log/aspera.log {
    daily
    rotate 10
    copytruncate
    compress
}
```

3. Create a separate `/etc/logrotate.d/aspera` configuration file that contains the same information as option 2.

Preserving IBM Spectrum Scale ACLs of transferred files

Ascp and Aspera Sync can preserve NFSv4 and POSIX ACLs and immutability attributes when transferring files from an IBM Spectrum Scale (formerly GPFS) cluster to another cluster.

Preserving Spectrum Scale ACLs

To preserve file attributes and permissions when transferring from one Spectrum Scale cluster to another, run **ascp** or **async** with the `--preserve-xattrs=native` option.

Preserving expiration attributes of immutable file sets

To preserve expiration attributes, use timestamp preservation options:

- **Ascp:** Use `--preserve-access-time` to preserve only the expiration attributes or `-p` to preserve all timestamps.
- **Aspera Sync:** Use `--preserve-access-time`.

Important behavior notes

Attribute preservation:

- Pools, replication, and cloning attributes are not preserved from source to destination.
- In Aspera Sync transfers, security attributes are preserved only when the user on both endpoints is root.

Immutable files and directories:

- Immutable directories on the source are not set as immutable on the destination. This ensures that the contents of the directory can be transferred.
- When the destination file exists and the source file is changed to immutable (`immutable: yes`), ascp and Aspera Sync update the destination file from mutable to immutable. However, if the source file is changed back to mutable (`immutable: no`), the change cannot be applied to the destination file because it is still immutable. Manually change the destination file after which ascp and Aspera Sync can write content and permissions changes.
- When the destination file exists and is immutable, ascp and Aspera Sync return an error for the immutable file (`Destination: Read-only file system`) and then transfers any other files in the transfer.

Append-only files:

- When the source file is an append-only file (`appendOnly: yes`), it can be transferred once to the destination. After the initial transfer, it cannot be updated by ascp or Aspera Sync because FASP must overwrite the file.
- When the destination file exists and is an append-only file, ascp and Aspera Sync either stop the rest of the transfer (for Spectrum Scale 5.0.1), or return an error for the file (`Destination: Read-only file system`) and then transfers any other files (for Spectrum Scale 5.0.0). To avoid transfer failures, exclude append-only files that exist on the destination from consideration on the source.

Logging client file system activity on HSTS

HSTS can be configured to log operations on the server's file system that are performed from client applications (such as the HSTS in client mode, or Console).

The logging of specific file system operations is controlled with an `<ascmd>` element in `aspera.conf`, within which logging can be set to yes or no for each operation.

```
<ascmd>
  <log_cmd>
    <as_info>no</as_info>
    <as_ls>no</as_ls>
    <as_rm>no</as_rm>
    <as_du>no</as_du>
    <as_df>no</as_df>
    <as_mkdir>no</as_mkdir>
    <as_cp>no</as_cp>
    <as_mv>no</as_mv>
    <as_md5sum>no</as_md5sum>
  </log_cmd>
</ascmd>
```

As an example of **asconfigurator** usage, the following command specifies that any deletions from the server file system by user **xeno** are logged:

```
asconfigurator -x "set_user_data;user_name,xeno;ascmd_log_cmd_as_rm,yes"
```

The command generates this `<ascmd>` element in `aspera.conf`:

```
<ascmd>
  <log_cmd>
```



```
<as_im>yes</as_im>  
</log_cmd>  
</ascmd>
```

Product limitations

Describes any limitations that currently exist for Aspera transfer server and client products.

- **Path Limit:** The maximum number of characters that can be included in *any* path name is 512 on Windows and 4096 on Unix-based platforms.
- **Illegal Characters:** Avoid the following characters in file names - / \ " : ' ? > < & * |.
- **Environment Variables:** The total size for environment variables depends on your operating system and transfer session. Each environment variable value must not exceed 4096 characters.

