

# Using a Load Balancer with the High Speed Add-On (HSAO) Option for Connect:Direct for UNIX

*Version 1.0*

© Copyright IBM Corporation 2016

## Overview

The Connect:Direct High Speed Add-On (HSAO) feature gives Connect:Direct for UNIX the ability to use IBM Aspera's FASP protocol to transfer files. Under certain network conditions common in MFT, the FASP protocol transfers data much faster than other reliable network protocols, such as TCP/IP, and it can significantly accelerate file transfers.

When Connect:Direct with the HSAO feature is used in a load balancer deployment, the load balancer configuration must ensure that the TCP and UDP connections for Connect:Direct HSAO sessions are routed to the same SNODE behind the load balancer.

This document provides information on how Connect:Direct HSAO works and how to configure a load balancer to work for Connect:Direct HSAO. These instructions only apply to Connect:Direct for UNIX, as Connect:Direct for Windows does not support a load balancer.

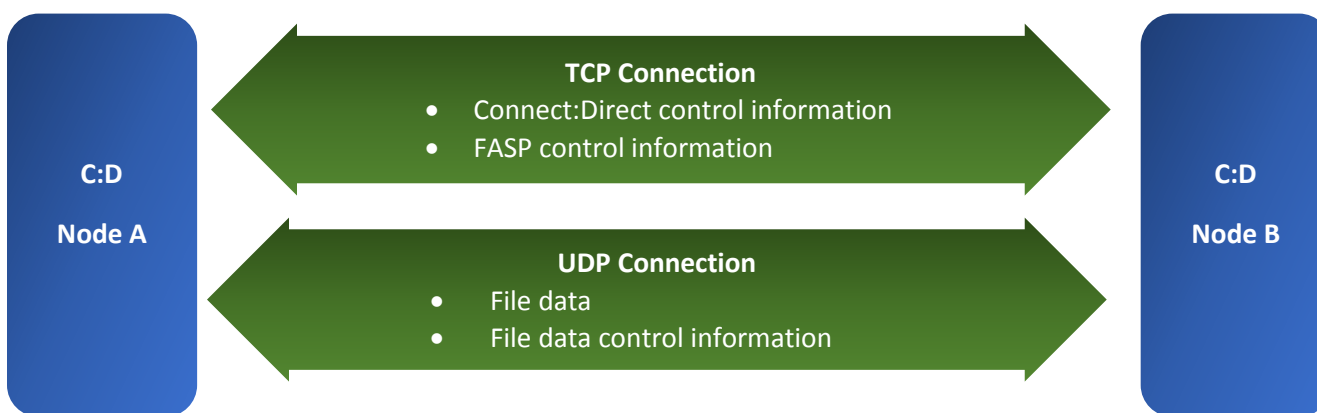
Two sample load balancer configurations are provided. These configurations were tested in the IBM lab with an F5 LTM device. However, the concepts in this document should be valid for other load balancers brands and models.

## Connect:Direct HSAO TCP and UDP connections

When Connect:Direct sends a file using HSAO, it uses the IBM Aspera FASP protocol, which utilizes both TCP and UDP connections.

The TCP connection is used for control information and session control, while the UDP connection is used for sending or receiving file data. Each copy step within the Connect:Direct process uses a UDP port. However, the UDP port number may not be the same for all copy steps using HSAO within the process.

The TCP control session and UDP connections for the Connect:Direct copy step must be routed to the same backend Connect:Direct server.



## Option 1 - Load Balancer with Source IP Persistence Configuration

One option to ensure the TCP and UDP connections are routed to the same backend Connect:Direct server is to configure the load balancer with persistence sessions based on the source IP address. This configuration ensures that:

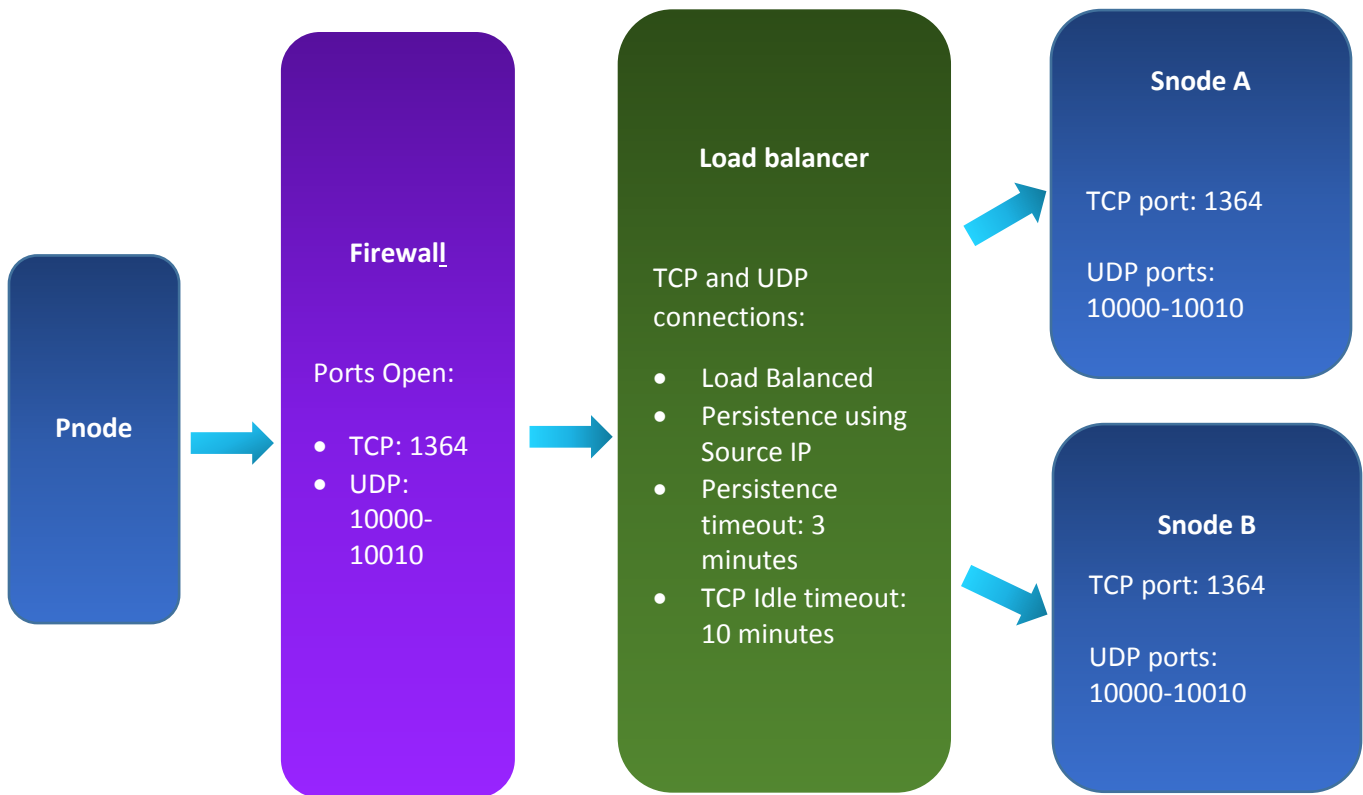
- Both UDP and TCP are persisted using the source IP address.
- Each backend Connect:Direct server use the same TCP and UDP port numbers.

### Advantages and disadvantages of using this configuration option

- Advantages
  - Easy to implement
  - Commonly used
  - Fewer UDP ports need to be open on the firewall
- Disadvantages
  - IP Source Persistence reduces the impact of load balancing when one Connect:Direct Pnode (or a very small number of Connect:Direct Pnodes) generates the majority of the load.
  - Proxies can cause multiple real IP sources to appear as the same source IP address

### TCP/UDP flow with a load balancer using Source IP persistence

1. The Connect:Direct Pnode initiates the TCP connection to the load balancer VIP.
2. The load balancer checks its source IP persistence table, and, if the source IP address (Pnode IP address) is in the table, the load balancer routes the connection to the persisted Snode.
3. If persistence has not been established for the source IP address, then the load balancer selects the appropriate Snode based on its load balancing configuration (such as round robin or least connected).
4. After the TCP connection is established, control information is exchanged between the Pnode and Snode.
5. When the Pnode encounters a COPY process statement, the Pnode and Snode negotiate a FASP or non-FASP file transfer.
6. If FASP is negotiated, the Snode sends the Pnode its UDP port. The Pnode uses the UDP port to send the file. This UDP port is selected from the Snode fasp.listen.ports initparm configuration setting.
7. The Connect:Direct Pnode initiates a UDP connection to the load balancer VIP.
8. The load balancer checks its source IP persistence table and finds the source IP (Pnode IP) is in the table (as a result of the TCP connection described in steps 2-4). The load balancer routes the connection to the persisted Snode.
9. When the file is sent, the UDP port is released.
10. If the Pnode encounters another COPY process statement, steps 5-9 are repeated.



## Load Balancer Configuration

VIP settings for TCP and UDP:

- Load balance method: Least connected
- Persistence profile: Source IP – virtual server (ties the TCP and UDP VIP together)

## Option 2 - Load Balancer with UDP Port Routing

Another option to ensure the TCP and UDP connections are routed to the same backend Connect:Direct server is to configure the load balancer to route UDP traffic based on the UDP port number. This configuration ensures that:

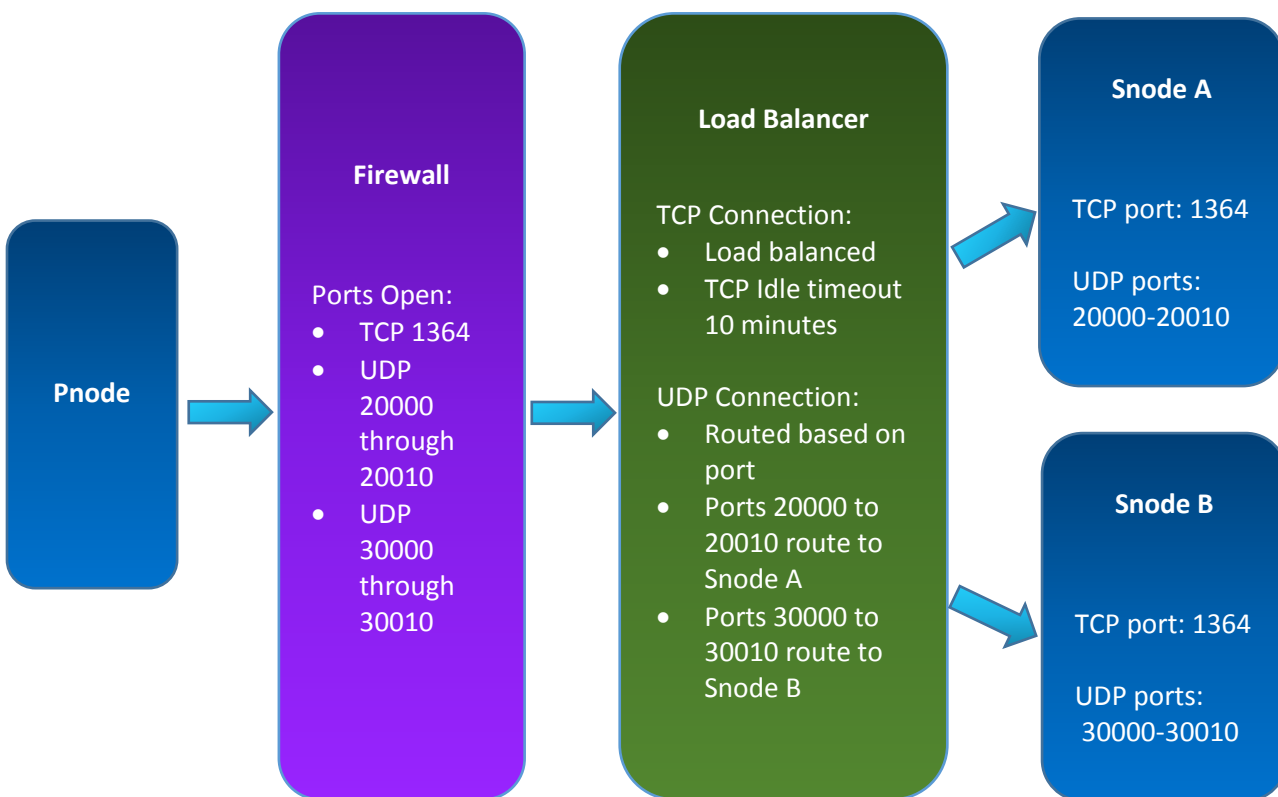
- TCP connections are load balanced using methods such as round robin or least connected. Each backend Connect:Direct server uses the same TCP port number.
- UDP messages are routed based on UDP port number. Each backend Connect:Direct server is configured with its own unique range of UDP ports to use.

## Advantages and disadvantages of using this configuration option

- Advantage
  - The work load is balanced even if the majority of the load is from a single Connect:Direct Pnode or a small number of Connect:Direct Pnodes
- Disadvantages
  - More UDP ports are open in the firewall.
  - This option is more difficult to configure in the load balancer and in Connect:Direct.

## TCP/UDP Flow with load balancer using UDP Port Routing Flow

1. The Connect:Direct Pnode initiates the TCP connection to the load balancer VIP.
2. The load balancer selects the appropriate Snode based on its load balancing configuration (such as round robin or least connected).
3. After the TCP connection is established, control information is exchanged between the Pnode and Snode.
4. When the Pnode encounters a COPY process statement, the Pnode and Snode negotiate a FASP or non-FASP file transfer.
5. If FASP is negotiated, the Snode sends the Pnode its UDP port. The Pnode uses the UDP port to send the file. This UDP port is selected from the Snode fasp.listen.ports initparm configuration setting.
6. Since each Connect:Direct Snode is configured to use a unique UDP port range, the load balancer routes the UDP connection to the appropriate Snode based on the configured UDP port range routing rules.
7. When the file is sent, the UDP port is released.
8. If the Pnode encounters another COPY process statement, steps 4-7 are repeated.



## Load Balancer/Firewall – TCP Idle timeout

Load balancers typically have a TCP Idle timeout feature. When no TCP traffic is detected during the idle timeout period, the TCP connection is dropped. This helps eliminate resources being allocated by inactive sessions.

During a Connect:Direct FASP file transfer, the data is transferred over the UDP connection. There is no TCP activity for the duration of the file transfer. If your file transfer takes longer than the TCP Idle timeout value, the load balancer disconnects the TCP connection, which interrupts the file transfer and the Connect:Direct process.

The TCP Idle timeout must be configured to a time that is greater than the time required to send your file. In the examples above, the TCP Idle timeout was set to 10 minutes. Your timeout value may need to be greater than 10 minutes.