

# High-Availability (HA): how to setup

This article assumes you've already **read, understood, and are familiar with** the [general concepts explained here](#).

Once you have decided how to deploy your Synclify Server! high-availability set, going from theory to practice is actually quite simple: it's just a two-step process.

## Step #1: deploy the first node as if it were a single stand-alone server

This procedure is very well explained [here](#) (for Windows) and [here](#) (for Linux).

Before proceeding with Step #2, ensure that the **system clock on every node in your HA cluster** (including the one you are about to add) **is synchronized** with a reliable time source. The recommended way to achieve this is by enabling NTP on all nodes. If even a single node has an out-of-sync clock, **the entire HA cluster will not function**.

## Step #2: adding the second (or n-th) node to make your HA set

If you are familiar with the single-node setup of Synclify Server! (see step #1) then you surely have noticed that, at the beginning of the web-based part of the installation process, you're presented with the following choice:

## Syncplify Server!

- Setup this node as a standalone Syncplify Server!
- Add this node to another Syncplify Server! node to achieve high-availability (HA)

NEXT

# Setup

Now, please choose how to setup this node.

To deploy your HA set you simply have to install Syncplify Server! on another node (machine or VM) and, when presented with the above choice, you'll select the **"Add this node to another..."** option.

You will then choose a Node ID (and an optional description) as usual...

## Syncplify Server!

Machine/Node ID

romantic-lovelace

Description (optional)

My additional node...

NEXT

BACK

# Node ID

Now, please, choose a nice ID for this node, or have the software generate one for you.

And then you will type the **IP address and port** (separated by a single colon) of the first node, the one you previously installed as a single stand-alone server. For best performance and optimal routing you should always use the nodes' internal (LAN) IP addresses here.

## Syncplify Server!

IP Address or Hostname, and Port

10.0.0.169:6443

NEXT

BACK

# Join existing node

Enter the host name or IP address and the port of a remote existing Syncplify Server! node that you want this local node to join in a high-availability (HA) set.

**Do not try to be creative here**, do not type anything exotic like the nodes' NetBEUI names "just because Windows knows it" or similar oddities. It will not work. Please, just follow the instructions and type the LAN IP address and port of the existing node you want to join to create your high-availability set. Thank you!

## Step #3: safe-list your load balancer(s)

If your HA deployment includes a **load balancer** (or multiple load balancers), this step is **mandatory**.

Load balancers continuously probe SFTP nodes by opening connections and then closing them, often without completing authentication. This is normal and expected behavior. Syncplify Server!'s built-in *Protector!* mechanism, however, cannot distinguish between a load balancer health-check and a connection-flood attack. Without intervention, *Protector!* will eventually auto-block the load balancer's IP address, which will bring down the entire cluster from the clients' perspective.

The fix is straightforward: add the IP address of every load balancer to the **safe-list**. Safe-listed addresses are never auto-blocked by *Protector!*, regardless of how they behave. This is precisely the scenario the safe-list was designed for.

For a full explanation of the difference between the block-list, allow-list, and safe-list, see [this article](#).

---

Revision #5

Created 21 July 2023 23:16:52 by DevTeam

Updated 29 April 2026 09:54:26 by DevTeam