

# HA-Simple Failover Plugin User Guide

NexentaStor provides a variety of high availability (HA) and fault detection/correction features. This includes ZFS self-healing capability, advanced data protection, ZFS hotplug<sup>1</sup>, and other enterprise-grade features delivered along with the filesystem.

In addition, each appliance monitors other appliances on the network (including both ssh-bound hosts and dynamically discovered NexentaStor appliances).

The appliance's unique upgrade/checkpointing facility utilizes snapshots, provides system checkpoints, and is designed to support 24/7/365 availability.

In addition, there are currently two extension modules designed specifically to provide functionality that is typically associated with High Availability Clustering. One such module introduces a 'simple-failover' (or simply, 'failover') group of appliances. A failover group contains two or more appliances, which share certain common attributes (for instance, a so called "failover IP address") and replicated storage. The plugin integrates with Nexenta [AutoCDP](#) service, to support failover of NexentaStor volumes block-mirrored over IP network.

Please note that this plugin is installed on **all** appliances that are members of a 'simple-failover' group.

## Definitions

Failover group contains two or more appliances. None of the appliances in the group is specifically designated to be (what's often called) the "primary" or "active", as opposed to being "secondary" or "passive". In fact, **any appliance in the group may be replaced by (or more exactly, failed over to) any other appliance in this same group.**

The plugin supports up to two failover IP interfaces (defined and illustrated below). The 'simple-failover' group of appliances can be deployed with and without [AutoCDP](#) - the existence of replicated volumes is detected at failover and appliance startup time.



NexentaStor **simple-failover** playback demonstration is available on the website: see [Tutorials & Demos](#)

For information on NexentaStor plugins, please visit NexentaStor [F.A.Q.](#) pages, or see Section "[Frequently Asked Questions](#)" in this document.

Failover group contains two or more appliances, which share certain common attributes (for instance, a so called "failover IP address") and storage accessible via iSCSI.

Appliances - members of the failover group can also replicate their storage using NexentaStor data replication services including [auto-cdp](#) (a.k.a. AutoCDP); in the latter case the failover logic makes sure to fail over the replicated volumes as well.

## Network interface configuration

**Each appliance in the group must have a spare (logical or physical) network interface**, also termed "failover networking interface". Starting version 2.0, the plugin supports up to two failover interfaces. At failover time, one of those selected "failover"

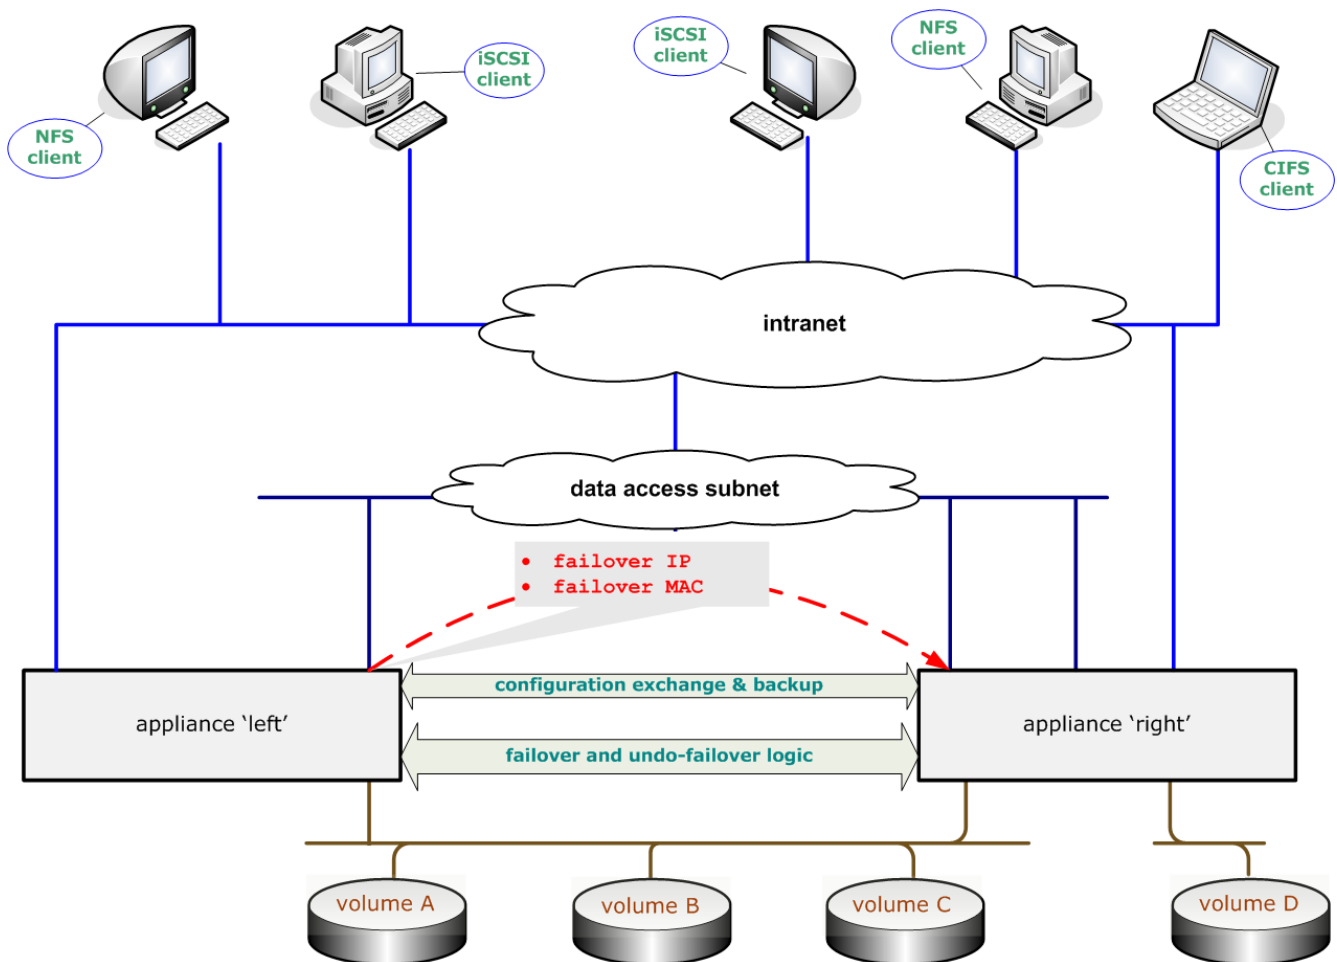
1 Assuming the appliance's data volumes are configured with hot spares, ZFS hotplug feature will make sure that faulted drives are automatically replaced by those spares. For more information, please see "[Note On Redundant Configurations](#)".

networking interfaces performs an important function - it gets configured with a single (per-group) triplet of static configuration:

- failover IP
- failover mask
- failover MAC address

The latter is a single **static** network interface configuration per given **simple-failover** group. The idea (and the requirement for the simple-failover configuration to work) is that the clients can access any of the appliances in the failover group via the configured { failover IP, failover mask, failover MAC }.

**Conceptually, the latter is the interconnect between the entire simple-failover group and the network "cloud" of the iSCSI/NFS/CIFS/etc. clients, as illustrated on the picture below:**



Failover is a compound operation that is either done as a whole, or not done. Once (and if) the failed appliance is brought back to service, the failover operation can be (atomically) un-done.



Note that, unlike [High Availability - HA Cluster](#) (next Section), **simple-failover** does not support shared storage, accessible from two appliances in the group. Detailed and the most current documentation on the [simple-failover](#) extension module is available on the corresponding [web page](#).