

# Best Practices for Xpand Platform Configuration

This section describes best practices for the following subjects:

- [Network Configuration](#)
- [Storage Configuration](#)
- [Operating System Concerns](#)

This page is a list of Best Practices for when setting up and installing Xpand. See [Xpand Installation Guide Bare OS Instructions](#) to complete the initial installation and general configuration instructions.

## Network Configuration

### IP addressing and Hostnames

When setting up a new cluster it is best to maintain consistent ordering and numbering with IP address, hostnames, and node numbers.

For example:

- IP Address x.x.x.11 = hostname1 = node 1.
- IP Address x.x.x.12 = hostname2 = node 2.
- IP Address x.x.x.13 = hostname3 = node 3.

Currently, to achieve this, you have to add nodes one at a time to the cluster using either [ALTER CLUSTER Syntax - Flex Up](#) or the [Flex Up page on the XpandGUI](#). Otherwise, the node number is random as the numbers are assigned in the order that the nodes respond to the addition request.

This is only relevant for on-premises installations, as you don't typically have that much control over IP address allocation in cloud/hosted environments.

### Multi-homed Network Configuration

While Xpand runs fine on nodes with a single network connection, the recommended network topology is a setup with at least two ethernet connections; the front-end and back-end. The front-end ethernet is for external traffic to communicate with the cluster while the back-end ethernet should be used for internode communication only.

- By default, the database will listen on all addresses and so no additional settings will be needed for the front-end ethernet.
- The IP of the Back-end ethernet will need to be set using the initial software setup wizard and is listed as "8 - Private (Back-End) IP: x.x.x.x".

For more information on setting up the network, please see [Network Security with Xpand](#).

When adding nodes to a cluster in a cloud environment such as AWS or Rackspace it is important to use the **internal** IP address as using the external address can cause cluster creation to fail.

## Network Security

The 3 main approaches to security are as follows:

- You are in a secured/trusted environment, so iptables is not needed. ( **Recommended** )
- You are in a hosted cloud environment such as AWS, where **Security Group** or similar mechanisms provide IP security by limiting access to ports (thus iptables on each node would be redundant).
- Neither of the above, use **iptables** to properly secure the cluster.

Unable to render {include} The included page could not be found.

In a typical secure configuration, you will limit access to TCP Ports 80, 8080, and 22 to the network CIDR range that maps to the public IPs for your administrative clients (typically exposed through your firewall), and you will limit access to TCP Port 3306 to your administrative client CIDR range and also to the range of IPs used by your application servers (if they are outside your firewall).

For complete details on how to set up security, please see [Network Security with Xpand](#).

## Front-end Load Balancing

Xpand has been designed to take full advantage of a front-end load balancer. We recommend HAProxy. Please see [Load Balancing Xpand with HAProxy](#) for details.

## NTP


NTP should be running so that the nodes' clocks do not get out of sync; otherwise, you may get inconsistent timestamps when using `now()`, and this also makes log analysis much more difficult. See [Verify NTP is running on Xpand](#).

## Managing Internode SSH Access

There are three methods of internode SSH access, listed here in order of administrative simplicity:

1. Host Based Authentication. The Xpand installer will set up host based authentication if allowed. (This is Recommended.)
2. Key Pair Authentication. For information on configuring key pair authentication, see [Internode Administrative Connectivity via Ethernet](#).
3. Password.

If you use password-based authentication then the [CLX Command Line Tool](#) will require a password for most commands.

 For ease of use, we suggest having the same password for each node.

## Latency

Databases are not typically network-bound (as compared to a file server), however, a clustered database system does rely upon low latency links between nodes.

- Cluster nodes should always be on the same subnet, with no intermediate routers between. As mentioned above, a dedicated backend connection is ideal.

To learn more about how network latency can affect nodes in a cluster please see [Recognizing Platform Limits](#).

## Storage Configuration

### Storage

- As with MySQL, disk seeks can be a huge performance bottleneck, and as such we require Xpand data to be on local SSDs rather than on spinning disk.
- To efficiently use your SSDs, and also avoid log build-up filling your data directory, it is suggested that a single large spinning disk be allocated for the log files. The path for the logs folder is `/data/clustrix/log/`.
- To learn more about disk I/O and SSD vs spinning disk please visit the section: [Memory and Disk I/O](#).

### RAID

Below are some general best practices on RAID configuration for Xpand:

- RAID 0: For best performance, it is recommended to use RAID 0 and in the case of a failed disk rely on node level redundancy.
- RAID 5: Due to the performance hit for running RAID, it is not a supported configuration.
- RAID 10: Often chosen for redundancy (without the performance loss of RAID5), but as Xpand already offers data redundancy, Raid 0 is sufficient for most deployments.

For more information about the various RAID level please see [Wikipedia: RAID](#).

### File System

For optimal performance use the `ext4` format for your filesystem.

Xpand will work when using an `ext3` formatted volume, but it will generate warnings during startup and space allocation. Growing an `ext3` device file will take longer than with other volume formats.

You can see which format your filesystem is with the following:

```
shell> # df -T
Filesystem      Type      1K-blocks      Used Available Use% Mounted on
/dev/md1        ext3      14446128       2253276  11459028  17% /
tmpfs           tmpfs     164440012      131076   16308936   1% /dev/shm
/dev/md0        ext3       253807        27521   213182    12% /boot
/dev/md3        ext4     1060589952    993662712  13052376   99% /data/clustrix
/dev/sdh1       ext3     484535504     9458888  450657416   3% /data/clustrix/log
```

In the above example, we have most of the partitions running `ext3` and the main `/data` partition running `ext4`.

## Operating System Concerns

## General

Below are the general operating system best practices and concerns for running the Xpand software.

- Xpand supports deployment on Centos/RHEL 7.4+.
- You should always avoid running 3rd party software on a node that is running Xpand as the database expects to be able to use the majority of the system resources. Running 3rd party software can cause a node to behave in unexpected ways and is not officially supported. Running 3rd party software that makes heavy use of CPU can result in "slow kernel scheduling" warning messages, and if severe enough, can lead to group changes.
- It is recommended to add the following path to your user PATH: /opt/clustrix/bin. This will allow you run [The CLX Command-Line Administration Tool](#) with ease.
- There are several useful commands in The CLX Command-Line Administration Tool that can make administering a cluster much easier.

## Swap

Do not configure swap. Having the database process go to swap, even on SSD, will degrade performance.

## VMs/virtualization

- Shared CPU/Core architecture should be avoided if possible as the database is quite sensitive to wait time when attempting to access a core. Dedicated CPU/Cores are highly recommended.
- VMs on the same physical box may result in I/O contention especially if they're all trying to do bulk inserts.