

MultinodeCluster06

Creating a multinode cluster

The default storage-conf.xml provided with cassandra is great for getting up and running on a single node. However, it is inappropriate for use in a multi-node cluster. The configuration and process here are the *simplest* way to create a multi-node cluster, but may not be the *best* way in production deployments.

Preparing the first node

The default storage-conf.xml uses the local, loopback address as its listen (inter-node) and Thrift (client access) addresses:

```
<ListenAddress>localhost</ListenAddress>

<ThriftAddress>localhost</ThriftAddress>
```

As the listen address is used for intra-cluster communication, it must be changed to a routable address so the other nodes can reach it. For example, assuming you have an Ethernet interface with address 192.168.1.1, you would change the listen address like so:

```
<ListenAddress>192.168.1.1</ListenAddress>
```

The Thrift interface can be configured using either a specified address, like the listen address, or using the wildcard 0.0.0.0, which causes cassandra to listen for clients on all available interfaces. Update it as either:

```
<ThriftAddress>192.168.1.1</ThriftAddress>
```

Or:

```
<ThriftAddress>0.0.0.0</ThriftAddress>
```

If the DNS entry for your host is correct, it is safe to use a hostname instead of an IP address. Similarly, the seed information should be changed from the loopback address:

```
<Seeds>
  <Seed>127.0.0.1</Seed>
</Seeds>
```

Becomes:

```
<Seeds>
  <Seed>192.168.1.1</Seed>
</Seeds>
```

Once these changes are made, simply restart cassandra on this node. Use netstat to verify cassandra is listening on the right address. Look for a line like this:

```
tcp4  0  0  192.168.1.1.7000  *.*  LISTEN
```

If netstat still shows cassandra listening on 127.0.0.1.7000, then either the previous cassandra process was not properly killed or you are not editing the storage-conf.xml file cassandra is actually using.

Preparing the rest of the nodes

The other nodes in the ring will use a storage-conf.xml almost identical to the one on your first node, so use that configuration as the base for these changes rather than the default storage-conf.xml. The first change is to turn on automatic bootstrapping. This will cause the node to join the ring and attempt to take control of a range of the token space:

```
<AutoBootstrap>true</AutoBootstrap>
```

The second change is to the listen address, as it must also not be the loopback and cannot be the same as any other node. Assuming your second node has an Ethernet interface with the address 192.168.2.34, set its listen address with:

```
<ListenAddress>192.168.2.34</ListenAddress>
```

Finally, update the Thrift address to accept client connections, as with the first node, either with a specific address or the wildcard:

```
<ThriftAddress>192.168.2.34</ThriftAddress>
```

Or:

```
<ThriftAddress>0.0.0.0</ThriftAddress>
```

Note that you should leave the Seeds section of the configuration as is so the new nodes know to use the first node for bootstrapping. Once these changes are made, start cassandra on the new node and it will automatically join the ring, assign itself an initial token, and prepare itself to handle requests.

<https://c.statcounter.com/9397521/0/fe557aad/1/> | stats