

ScyllaDB Production CHECKLIST

The goal of this document is to have a checklist that production customers can use to make sure their deployment adheres to ScyllaDB's recommendations.

Users should follow up on each of the main bullets described below to verify they comply with the recommendations provided.

[Prefix A pre-requirements](#)

[Prefix B compaction strategy](#)

1. Resiliency
 - a. [Replication factors](#)
 - b. [Consistency levels](#)
 - c. [Gossip configuration](#)
 - d. [Seed nodes redundancy](#)
2. Performance
 - a. [Io tune information](#)
 - b. [Perftune](#)
 - c. [MC format](#)
 - d. [Enable compression traffic between nodes:](#)
3. Connectivity
 - a. [Drivers settings \(connection pool\)](#)
4. Management
 - a. [Scylla Manager](#)
 - b. [Repair](#)
 - c. [Backup](#)
 - d. [Monitoring](#)
 - e. [Security](#)
 - f. [Configuration management](#)
5. [Security](#)
 - a. We recommend using hardened server machine and operating system
 - b. RBAC (Role Based Access Control)
 - c. Encryption in transit
6. Advanced Topics
 - a. [HA testing in single DC](#)
 - b. [HA testing in multi DC](#)

[Appendix A: Adding node procedure](#)

[Appendix B: Interesting read](#)

- **Prefix A pre-requirements**
 - [Scylla System Requirements](#)
 - [Scylla Getting Started](#)
 - NTP !!! (or chrony)
- **Prefix B compaction strategy**
 - For most use cases - [Size-tiered Compaction Strategy \(STCS\)](#) is recommended

1. Resiliency

- a. Replication factors
 - i. We recommend using replication factor greater than 2 . If you have a multi-datacenter architecture we would recommend to have RF=3 on each DC.
[Replication Factor](#)
[Scylla Architecture - Fault Tolerance](#)
[Scylla University replication factor](#)
- b. Consistency levels
 - i. We recommend to use LOCAL_QUORUM across the cluster and DC's
[Scylla Architecture - Fault Tolerance](#)
[Consistency Level Console Demo](#)
[Scylla University consistency level](#)
- c. Gossip configuration
 - i. Always use GossipingPropertyFileSnitch or Ec2MultiRegionSnitch:
[Gossip in Scylla](#)
[Scylla University gossip](#)
 - ii. Data Replication strategy: NetworkTopologyStrategy replication-strategy that support multi-DC for your Keyspaces
- d. Seed nodes redundancy - we are recommending to have 2~3 seed nodes per DC according to: [seed-nodes](#)

2. Performance

Please validate you have run `scylla_setup` in order to tuning ScyllaDB to your hardware.

If you are running on a physical hardware please take a look into the following configurations files:

- a. [io.conf](#) - Disk I/O tune information
- b. `perftune.yaml`
 - i. If you have more than 8cores or 16vcpu always use "mode: sq_split"
- c. Enabling MC format
 - i. Enable MC format in `scylla.yaml`:
[enable sstables mc format parameter as 'true' in scylla.yaml file](#)
- d. Enable compression traffic between nodes:
 - i. Enabling it in `scylla.yaml`:
Internode_compression = all
 - ii. *Verify your client driver is using compressed traffic when connected to Scylla. It is driver settings dependent, please check your client driver manual.*

3. Connectivity

- a. Drivers settings (connection pool)
 - i. Shard aware driver (if possible)
[Scylla Drivers](#)
 - ii. Configure connection pool - open more connections (>3 per shard) and/Or more clients.
[maximizing performance via concurrency](#)

4. Management

Remember **ALWAYS** to install Scylla Manager and Scylla Monitoring

- a. Scylla Manager
 - i. Scylla Manager enabling centralized cluster administration and database automation such as repair backup and health-check
[scylla manager](#)
- b. Repair
 - i. Run repair periodically - preferably once a week and from [Scylla Manager](#)
- c. Backup/Restore
 - i. Run a full weekly backup
 - ii. Run a daily backup from Scylla manager
 - iii. Check restore periodically
 - iv. Save backup to an S3 compatible API
<https://docs.scylladb.com/operating-scylla/manager/2.0/backup/>
- d. Monitoring
 - i. Scylla Monitoring helps you monitor everything about your Scylla cluster. We'll usually ask for monitoring metrics when you open a ticket
<https://docs.scylladb.com/operating-scylla/monitoring/>
- e. Configuration management -

Configuration settings coherency (how do you make sure the info in scylla.yaml is coherent among all nodes). Using tool such as Ansible / Chef / Puppet / Salt / Juju is recommended:
<https://www.softwaretestinghelp.com/top-5-software-configuration-management-tools/>

5. Security

- i. Enable authentication
 - ii. Work with RBAC
 - iii. Setup Encryption on Transit
- All and more can be found here:
[Scylla Security Checklist](#)

6. Advanced Topics

- a. HA testing in single DC - for example:
 - i. Shutdown one node from the cluster (Or scylla service if on the cloud) for 30min.
 - ii. Turn it back on
- b. HA testing in multi DC - for example:
 - i. Disconnect one DC from the other by stopping scylla service on all of these DC nodes.
 - ii. Reconnect the DC

Appendix A: Adding node procedure

- [Add node](#)
- [Repair](#)
- [Cleanup](#)

Appendix B: Interesting read

- [Scylla's Best Tips for Operational and Application Excellence](#)
- [How to be successful with Scylla](#)

- [Best Practices for Scylla Applications](#)
- [Max performance via concurrency \(with less timeouts\)](#)
- [Writing Applications for Scylla](#)
- [Ingestion flow control](#)
- [Advanced Monitoring and how to maximize performance](#)
- [Workload prioritization can reduce your DC footprint](#)
- Mgmt tools: [Scylla Monitoring 3.0](#) , [Scylla Manager 2.0](#)
- [Repair, Tombstones and Scylla Manager](#)