NOCTION

Noction Flow Analyzer

ClickHouse Multi-Shard and Multi-Replica Setup

Copyright ©2026 Noction Inc.

NFA ClickHouse Multi-Shard and Multi-Replica Setup

Customer configuration guide | ClickHouse 26.3.x | Cluster: nfa

This guide covers adding shards or replicas to an NFA ClickHouse cluster. Configure all XML files first, then restart ClickHouse Server once per existing node. Run database initialization only from the NFA node.

Example topology

```
Shard 01 / Replica 01  shard1-node.example.com  (NFA + ClickHouse + Keeper)
Shard 02 / Replica 01  shard2-node.example.com  (ClickHouse)
```

Connection	Port	Requirement
ClickHouse native	TCP 9000	Reachable between all ClickHouse nodes
ClickHouse Keeper	TCP 9181	Reachable from every ClickHouse node

Important: Use real hostnames and addresses throughout. Restrict these ports to trusted hosts and networks.

1. Install ClickHouse on additional nodes

Install NFA normally on the primary node. On additional nodes, install the same ClickHouse 26.3.x version as the NFA node. Example for Ubuntu/Debian:

```
cat >/etc/apt/preferences.d/nfa-clickhouse <<'EOF'
Package: clickhouse-*
Pin: version 26.3.*
Pin-Priority: 1001
EOF
apt-get update
apt-get install clickhouse-server clickhouse-client clickhouse-common
```

Do not install NFA or nfaadm on additional ClickHouse nodes. Keep a newly installed ClickHouse Server stopped until all configuration files are ready, if possible.

2. Configure name resolution

All nodes must resolve all cluster hostnames. Use DNS or add the following entries to /etc/hosts on every node:

```
192.0.2.10  shard1-node.example.com
192.0.2.11  shard2-node.example.com
```

Replace the example addresses with the actual reachable IP addresses. Verify connectivity to ClickHouse port 9000 on existing running nodes and Keeper port 9181 on the NFA node before applying changes. Newly installed nodes can be checked after startup.

3. Configure ClickHouse on every node

Create or update the following files on every ClickHouse node. Existing NFA-provided configuration files may already exist; update them rather than creating conflicting copies.

3.1. Listen interface

/etc/clickhouse-server/config.d/listen.xml

```
<?xml version="1.0"?>
<clickhouse>
    <listen_host>0.0.0.0</listen_host>
</clickhouse>
```

3.2. Keeper connection

/etc/clickhouse-server/config.d/keeper.xml - identical on all nodes:

```
<?xml version="1.0"?>
<clickhouse>
    <zookeeper>
        <node>
            <host>shard1-node.example.com</host>
            <port>9181</port>
        </node>
    </zookeeper>
</clickhouse>
```

3.3. Shard and replica identifiers

/etc/clickhouse-server/config.d/macros.xml - unique to each node. Example for Shard 01 / Replica 01:

```
<?xml version="1.0"?>
<clickhouse>
    <macros>
        <shard>01</shard>
        <replica>01</replica>
    </macros>
</clickhouse>
```

On shard2-node.example.com, set shard=02 and replica=01. Replicas within one shard share a shard ID but must have distinct replica IDs.

3.4. Cluster topology

/etc/clickhouse-server/config.d/cluster.xml - identical on all nodes:

```
<?xml version="1.0"?>
<clickhouse>
    <remote_servers>
        <nfa>
            <shard>
                <internal_replication>true</internal_replication>
                <replica>
                    <host>shard1-node.example.com</host>
                    <port>9000</port>
                </replica>
            </shard>
            <shard>
                <internal_replication>true</internal_replication>
                <replica>
                    <host>shard2-node.example.com</host>
                    <port>9000</port>
                </replica>
            </shard>
        </nfa>
    </remote_servers>
</clickhouse>
```

3.5. Distributed query setting (multi-shard installations)

/etc/clickhouse-server/users.d/n_distributed_product_mode.xml

```
<?xml version="1.0"?>
<clickhouse>
    <profiles>
        <default>
            <distributed_product_mode>global</distributed_product_mode>
        </default>
    </profiles>
</clickhouse>
```

3.6. Inter-node access

/etc/clickhouse-server/users.d/n_cluster_networks.xml - specify the actual trusted IP addresses of every ClickHouse node:

```
<?xml version="1.0"?>
<clickhouse>
    <users>
        <default>
            <networks>
                <ip>127.0.0.1</ip>
                <ip>::1</ip>
                <ip>192.0.2.10</ip>
                <ip>192.0.2.11</ip>
            </networks>
        </default>
    </users>
</clickhouse>
```

Important: NFA uses /etc/noction/clickhouse/users.xml as the primary user configuration. Do not edit it for this procedure. Preserve any additional existing network permissions needed by the installation. The default account shown here has no password: allow only trusted addresses.

4. Configure Keeper on the NFA node

In /etc/clickhouse-keeper/keeper_config.xml, ensure Keeper listens on an address reachable by the other nodes:

```
<listen_host>0.0.0.0</listen_host>
```

If a change was needed, apply it once:

```
systemctl restart clickhouse-keeper
```

From each additional node, confirm Keeper connectivity:

```
nc -vz shard1-node.example.com 9181
```

If Keeper already accepts remote connections, do not restart it.

5. Apply configuration once

After all XML changes and network rules are in place, restart ClickHouse Server once per existing node, preferably one node at a time:

```
systemctl restart clickhouse-server
```

On a new node where ClickHouse has not yet been started:

```
systemctl start clickhouse-server
```

Confirm that port 9000 is reachable between nodes. No separate restarts are needed for individual XML files.

6. Verify the cluster

Run from the NFA node:

```
clickhouse-client -q "
SELECT shard_num, replica_num, host_name
FROM system.clusters
WHERE cluster = 'nfa'
ORDER BY shard_num, replica_num"
```

```
clickhouse-client -q "
SELECT hostName(), version()
FROM clusterAllReplicas('nfa', system.one)
ORDER BY hostName()"
```

All configured nodes must appear without errors. On each new node, verify that ClickHouse can communicate with Keeper:

```
clickhouse-client -q "
SELECT name FROM system.zookeeper WHERE path = '/'"
```

Important: Do not proceed if any node is unreachable, inter-node authentication fails, or Keeper queries fail.

7. Initialize the NFA database

Run the following only on the NFA node, after cluster connectivity is confirmed:

```
nfaadm --init ts
```

This step is mandatory after adding a shard or replica. Wait for successful completion.

8. Basic post-initialization checks

Confirm the NFA database exists on all nodes:

```
clickhouse-client -q "
SELECT hostName(), name
FROM clusterAllReplicas('nfa', system.databases)
WHERE name = 'nfa'
ORDER BY hostName()"
```

Check replica availability:

```
clickhouse-client -q "
SELECT hostName(), table, is_readonly, active_replicas, total_replicas
FROM clusterAllReplicas('nfa', system.replicas)
WHERE database = 'nfa'
ORDER BY table, hostName()"
```

Expected: is_readonly = 0 and, once synchronization finishes, active_replicas = total_replicas. Initial synchronization of a new replica may take time.

9. Restart NFA services

```
systemctl restart nfa.target
```

The expanded cluster is ready for normal NFA operation once initialization and the checks above complete successfully.

Adding a replica to an existing shard

Follow the same steps. Assign the new node the existing shard ID and a new replica ID, then update cluster.xml on every node. For example, adding Replica 02 to Shard 01:

```
Shard 01: shard1-node.example.com     (replica 01)
          shard1-replica2.example.com (replica 02)
Shard 02: shard2-node.example.com     (replica 01)
```

The new replica uses shard=01 and replica=02 in macros.xml. Add a second <replica> entry inside the first <shard> in the shared cluster.xml:

```
<replica>
    <host>shard1-replica2.example.com</host>
    <port>9000</port>
</replica>
```

Ensure hostname resolution, Keeper access, trusted-network permissions, and matching ClickHouse versions. Apply the configuration once, run nfaadm --init ts from the NFA node, and allow time for synchronization.

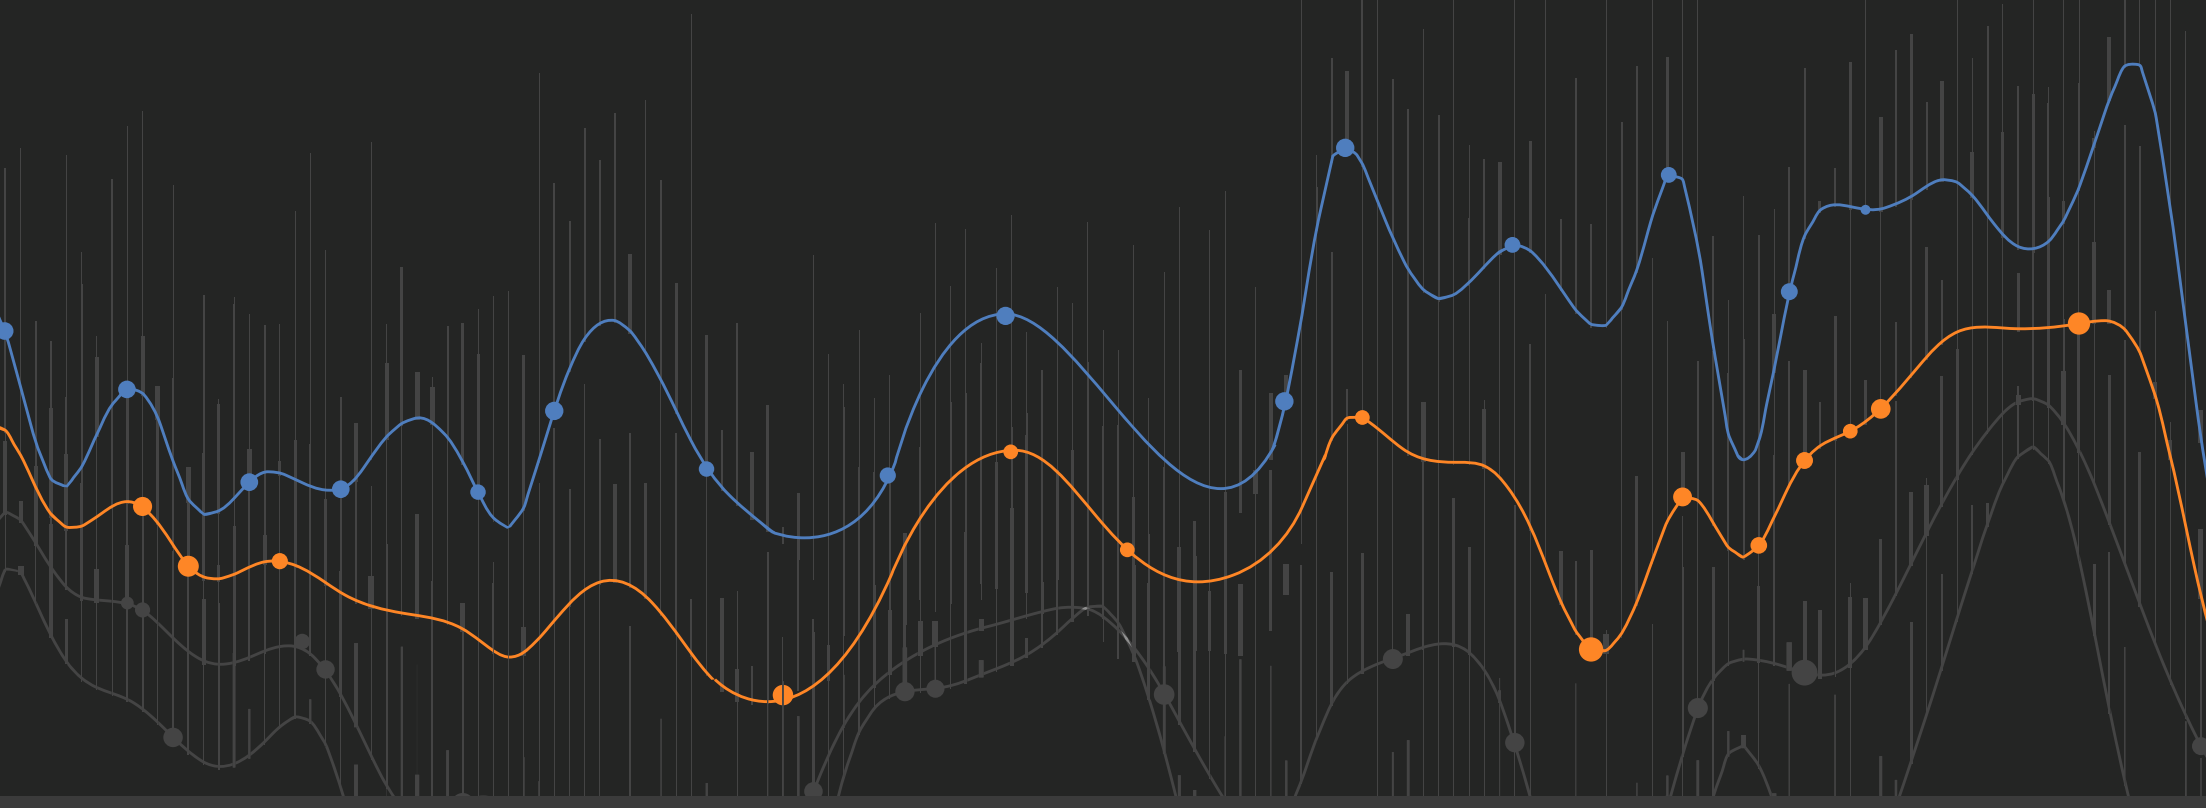
Completion checklist

- ☐ Matching ClickHouse 26.3.x versions on all nodes
- ☐ Hostname resolution and TCP 9000 / 9181 connectivity
- ☐ Correct macros and identical cluster.xml on all nodes
- ☐ Trusted inter-node access and distributed query setting
- ☐ Cluster query returns every configured node
- ☐ nfaadm --init ts completed on the NFA node
- ☐ NFA database present on all nodes; replicas active
- ☐ NFA services restarted

Note: Adding a shard does not automatically rebalance historical data. Adding a replica synchronizes data for its existing shard.

NOCTION





Noction Flow Analyzer

ClickHouse Multi-Shard and Multi-Replica Setup

Copyright ©2026 Noction Inc., All Rights Reserved. Noction logos, and trademarks or registered trademarks of Noction Inc. or its subsidiaries in the United States and other countries.

Other names and brands may be claimed as the property of others. Information regarding third party products is provided solely for educational purposes.

Noction Inc. is not responsible for the performance or support of third party products and does not make any representations or warranties whatsoever regarding quality, reliability, functionality, or compatibility of these devices or products.

Copyright ©2026 Noction Inc.