

Overview of the Juniper Networks Database

Each C-series Controller contains a Juniper Networks database. The database can store SRC data, SRC sample data, SRC configuration information, and a number of user profiles. You store subscriber data in another database.

The Juniper Networks database is designed to store a limited number of subscriber entries that may be shared among your subscribers. If you need to have dedicated entries for each subscriber, you can configure the SRC to use an external directory. We recommend that an external directory store the subscriber data in environments that have more than 1000 subscribers with an average of 3 subscriptions per subscriber.

When the C-series Controller starts for the first time, you must enable the Juniper Networks database. After the database is operational, you can load sample data and perform other configuration activities that use this database.

You can operate this database as a standalone database or as a member of a community of Juniper Networks databases. Typically, you run the database in standalone mode only in testing environments. In standalone mode, the database does not communicate with other Juniper Networks databases; there is no data distribution and no redundancy. In community mode, databases distribute data changes among specified databases. When you have two or more C-series Controllers, enable the Juniper Networks database to run in community mode, and assign a role to each database:

- **Primary role**—A database that provides read and write access to client applications. It replicates its data and distributes changes to any Juniper Networks databases configured as neighbors.

We recommend that you configure at least two databases to have a primary role.

- **Secondary role**—A database that provides read access to client applications. If client applications try to write data to this database, the database refers the client to a primary database.

Neighbors are Juniper Networks databases that receive data from another Juniper Networks database. When you configure a database to be a neighbor, you configure it as one of the following types:

- **Primary neighbor**—A database that propagates changes that it receives to other Juniper Networks databases configured as neighbors. A primary neighbor must be assigned a primary role.

We recommend that you configure at least two databases as primary neighbors.

- **Secondary neighbor**—A database that only receives database changes. A secondary neighbor must be assigned a secondary role.

When you configure neighbors for the databases, keep in mind the following guidelines:

- A database assigned a primary role can have primary and secondary neighbors.

- A database assigned a secondary role must have at least one primary neighbor, but no secondary neighbors. Because a secondary database cannot distribute changes to its neighbors, if you do configure a secondary neighbor for a secondary database, the software does not use the configuration for the secondary neighbor.

To share processing load, you can configure SRC components, such as SRC-ACP, NIC, or SAE, to use a specified database. In the local configuration for SRC components, you configure the URL of the directory.

Redundancy for a Juniper Networks Database

Protect SRC data by setting up a redundancy scheme for your Juniper Networks databases. Client applications control which database they connect to as their primary database and as their backup database.

Use the following guidelines to plan which databases are assigned primary or secondary roles, and which databases are primary or secondary neighbors:

- Each Juniper Networks database that is assigned a primary role should have at least one primary neighbor. If a database assigned a primary role become inoperable, a client application fails over to a primary neighbor.
- Each database that is assigned a secondary role should have at least two primary neighbors.
- Applications that frequently perform write operations to the database should connect to databases that have a primary role. Applications that perform frequent write operations are the C-Web interface, the SRC CLI, back-office applications that provision data, and in some cases the SRC-ACP.
- Applications that rarely perform updates, such as the NIC and SAE, can communicate with databases assigned a secondary role. For example, you could configure the NIC and SAE to communicate with the local directory on a C-series Controller, and configure the database on this system to have a secondary role.

Security for a Juniper Networks Database

You can secure connections to a Juniper Networks database by:

- Allowing only Secure Lightweight Directory Access Protocol (LDAPS) connections from remote systems. In this case, both database replication and remote SRC components connect through LDAPS. Restricting all remote connections to LDAPS is supported only on C-Series Controllers.
- Allowing only LDAPS connections for database replication, but LDAP or LDAPS connections for other applications. In this case, remote SRC components can connect through LDAP or LDAPS.

The type of secure connection you configure determines which ports are open to a Juniper Networks database:

- Remote component access through LDAP—Port 389
- Remote component access through LDAPS—Port 636

- Secure database access for replication—Port 636
- Database access without security for replication—Port 389
- Local component access through LDAP—Port 389

You can also increase the security of your Juniper Networks database by changing the passwords that SRC components use to communicate with the database.

Related Topics

- For information about configuring the SAE to access subscriber data, see [Configuring LDAP Access to Directory Data](#) .
- [Adding a Juniper Networks Database to an Established Community\(SRC CLI\)](#)
- [Configuration Statements for the Juniper Networks Database \(SRC CLI\)](#)
- [Viewing Statistics for the Juniper Networks Database](#)
- [Example: Configuration for a Database Community](#)

