



SolrCloud Configuration Guide

CONFIDENTIAL INFORMATION

The information herein is the property of Ex Libris Ltd. or its affiliates and any misuse or abuse will result in economic loss. DO NOT COPY UNLESS YOU HAVE BEEN GIVEN SPECIFIC WRITTEN AUTHORIZATION FROM EX LIBRIS LTD.

This document is provided for limited and restricted purposes in accordance with a binding contract with Ex Libris Ltd. or an affiliate. The information herein includes trade secrets and is confidential.

DISCLAIMER

The information in this document will be subject to periodic change and updating. Please confirm that you have the most current documentation. There are no warranties of any kind, express or implied, provided in this documentation, other than those expressly agreed upon in the applicable Ex Libris contract. This information is provided AS IS. Unless otherwise agreed, Ex Libris shall not be liable for any damages for use of this document, including, without limitation, consequential, punitive, indirect or direct damages.

Any references in this document to third-party material (including third-party Web sites) are provided for convenience only and do not in any manner serve as an endorsement of that third-party material or those Web sites. The third-party materials are not part of the materials for this Ex Libris product and Ex Libris has no liability for such materials.

TRADEMARKS

"Ex Libris," the Ex Libris bridge, Primo, Aleph, Alephino, Voyager, SFX, MetaLib, Verde, DigiTool, Rosetta, URM, Voyager, ENCompass, Endeavor eZConnect, WebVoyage, Citation Server, LinkFinder and LinkFinder Plus, and other marks are trademarks or registered trademarks of Ex Libris Ltd. or its affiliates.

The absence of a name or logo in this list does not constitute a waiver of any and all intellectual property rights that Ex Libris Ltd. or its affiliates have established in any of its products, features, or service names or logos.

Trademarks of various third-party products, which may include the following, are referenced in this documentation. Ex Libris does not claim any rights in these trademarks. Use of these marks does not imply endorsement by Ex Libris of these third-party products, or endorsement by these third parties of Ex Libris products.

Oracle is a registered trademark of Oracle Corporation.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Ltd.

Microsoft, the Microsoft logo, MS, MS-DOS, Microsoft PowerPoint, Visual Basic, Visual C++, Win32,

Microsoft Windows, the Windows logo, Microsoft Notepad, Microsoft Windows Explorer, Microsoft Internet Explorer, and Windows NT are registered trademarks and ActiveX is a trademark of the Microsoft Corporation in the United States and/or other countries.

Unicode and the Unicode logo are registered trademarks of Unicode, Inc.

Google is a registered trademark of Google, Inc.

Copyright Ex Libris Limited, 2023. All rights reserved.

Document released: June 2023

Web address: <http://www.exlibrisgroup.com>

Table of Contents

1	SolrCloud Configuration in Rosetta	4
2	Configuring IDX Servers to SolrCloud Mode	4
3	Adding Replicas/Shards to an Existing SolrCloud Cluster	5
4	Rollback to Non-SolrCloud Mode	6

SolrCloud Configuration in Rosetta

SolrCloud is a cluster of Solr servers that combines fault tolerance and high availability for flexible distributed search and indexing. These servers are managed by the ZooKeeper service. This document explains how to configure SolrCloud for Rosetta.

The following terms are referred to in this document:

- **Legacy (non-SolrCloud) mode** – Includes distributed indexing and searching. When an IDX server or a Solr instance is down, part of the data is unavailable.
- **Shard** – A logical piece of a collection. Each shard is made up of one or more replicas.
- **Replica** – One copy of a shard.
- **SolrCloud mode** – Includes index replication, failover, load balancing, and distributed queries with the help of ZooKeeper. Every index update is distributed across the cluster to every shard and replica of the cluster.
- **Replication Factor** – The number of physical replicas to be created for each logical shard of the collection.

Note: A list of your Rosetta IDX servers is available in the Administration module under “Index Status”.

Configuring IDX Servers to SolrCloud Mode

To configure Rosetta to use SolrCloud mode, confirm Rosetta is down on all IDX servers and perform the following steps on each one:

- 1 Open `global.properties (/exlibris/dps/d4_1/system.dir/conf)` and locate the following properties:

```
zk.leader.port=5811
zk.election.port=5821
zk.client.port=5831
```

Three ports are assigned: One is the client port and the other two are used by the Zookeeper servers to communicate with each other. Assign different ports if necessary or skip to the next step.

- 2 Enter the hostnames (**the same name as in the “server.name” property**) of three Rosetta servers to run the Zookeeper ensemble (the servers do not have to be IDX servers). For example:

```
zk.server.1=rosetta1.myinstitution.com
zk.server.2=rosetta2.myinstitution.com
zk.server.3=rosetta3.myinstitution.com
```

- 3 Run **rosettings.csh**. The following menu appears:

```
>dps_bin
>rosettings.csh

[1] Update gateway details
[2] Update server name
[3] Update pds server details
[4] Update database passwords
[5] Setup Solr
[6] Exit.
```

- 4 Select option **5 Setup Solr**. The following appears:

```
Solr Cloud options:

[1] Setup SolrCloud
[2] Setup as zookeeper server
[3] Update SolrCloud Cluster Architecture
[4] Rollback to Solr Legacy-mode
```

- 5 Select option **1 Setup SolrCloud**. Enter a replication factor (the number of replicas for each shard). The number should be between 1 and the total number of IDX servers.

Repeat the previous command for each IDX server.

- 6 On each ZooKeeper server (defined in #2) run **rosettings.csh**, select **5 Setup Solr**, and then **2 Setup as ZooKeeper server**.

Note: This step must be performed on all the ZK environments in a time range of up to 10 minutes, in order for the configurations to be uploaded.

- 7 Restart all IDX servers.

Adding Replicas/Shards to an Existing SolrCloud Cluster

Adding a shard is required in case an IDX server is added to the Rosetta cluster (either a new server or reassignment of a non-IDX server). Adding a replica means using an existing IDX server to replicate another IDX server.

Confirm Rosetta is down.

- 1 If you are adding a replica, skip to the next step. If you are adding a shard, open `global.properties (/exlibris/dps/d4_1/system.dir/conf)`, locate the `app.role` property, and add 'IDX':

```
app.role=DEP,REP,DEL,PER,IDX
```

Save changes and run `$dps_dev/system.dir/bin/set_globals.sh`.

- 2 Run `rosettings.csh` and select option **5 Setup Solr**. The following appears:

```
Solr Cloud options:
[1] Setup SolrCloud
[2] Setup as zookeeper server
[3] Update SolrCloud Cluster Architecture
[4] Rollback to Solr Legacy-mode
```

- 3 Select **3 Update SolrCloud Cluster Architecture**. Enter a replication factor (the number of replicas for each shard). The number should be between 1 and the total number of IDX servers).
- 4 Rerun `rosettings.csh`, select **2 Setup as ZooKeeper server** for the Zookeeper servers to update the ZooKeeper ensemble with the change.

Note: This step must be performed on all the ZK environments in a time range of up to 10 minutes, in order for the configurations to be uploaded.

Rollback to Non-SolrCloud Mode

To roll back to non-SolrCloud mode:

- 1 run `rosettings.csh` and select **5 Setup Solr**:

```
Solr Cloud options:
[1] Setup SolrCloud
[2] Setup as zookeeper server
[3] Update SolrCloud Cluster Architecture
[4] Rollback to Solr Legacy-mode
```

- 2 Select **4 Rollback to Solr Legacy-mode**.

Repeat for each IDX server.