

Xindice 1.0 Administration Guide

\$Revision: 1.2 \$

by Kimbro Staken

1. Database Administration

Database administration of Xindice is accomplished from the command line using the `xindice` command. This command allows you to view and alter the database configuration on the fly on a running system. A complete list of available commands and more detail about each command can be found in the [Command Line Tools Reference Guide](#).

1.1. Managing Collections

1.1.1. Adding a Collection

Adds a collection named `products` under the collection `/db/data`.

```
xindice add_collection -c /db/data -n products
```

1.1.2. Deleting a Collection

Deletes the collection named `products` from the collection `/db/data`.

```
xindice delete_collection -c /db/data/products
```

1.1.3. Listing the Collections

This will display a list of all child collections under the collection `/db/data`

```
xindice list_collections -c /db/data
```

1.2. Managing Indexes

The Xindice indexing system allows you to define indexes to speed performance on commonly used XPath queries. If no indexes are defined you can still execute queries but performance will suffer because the query engine will need to scan the entire collection to create the result node-set.

Indexes can be added using the `xindice` command.

1.2.1. Adding an Index

Using this simple XML file you might want to index the `product_id` element because searches for products by `product_id` are common.

```
<?xml version="1.0"?>
<product>
  <product_id>120320</product_id>
  <description>Glazed Ham</description>
</product>
```

This can be accomplished by running the following command. This will create an index named `idindex` on all `product_id` elements in the collection `/db/data/catalog`.

```
xindice add_indexer -c /db/data/catalog -n idindex -p product_id
```

Once this is done the query engine will now use this index to help resolve XPath queries that involve restriction on the value of the `product_id` element.

The `-p` parameter to the command specifies the pattern to use in the index. These patterns are used by the Indexing system to determine best-fit and match-based Indexers for queries and index updating. The pattern used **MUST** resemble the following scheme.

Pattern	Description
=====	=====
elem	The value of the named element
elem@attr	The value of the attribute for the named element
*	The value for all elements
*@attr	The value of the named attribute for all elements
elem@*	The value of all attributes for the named element
@	The value of all attributes for all elements

Note: In order to index a namespace other than the default namespace, you must prepend your pattern components with a URI placed in square brackets. Example:

```
[http://www.world.org/People]person
*@[http://www.world.org/People]id
[http://www.world.org/People]person@[http://www.world.org/People]id
```

Do not include a prefix in these patterns, as the indexing system, like most Namespace processing applications, processes namespaced elements and attributes independently of the prefix that is used.

1.2.2. Indexing both Elements and Attributes

Because the patterns recognize either an element or an attribute, and not both, in order to index all element and attribute values in a collection, you'd have to create two index entries. The `*` pattern will index all elements and the `*@*` pattern will index all attributes of all elements.

```
xindice add_collection_indexer -c /db/data/catalog -n idindex -p '*'
xindice add_collection_indexer -c /db/data/catalog -n idindex -p '*@*'
```

Excessive use of wildcard indexes can adversely affect the performance of the indexing system. Best practice would be to use specific element or attribute indexes whenever possible, and only define wildcard indexes when it is absolutely necessary.

2. Server Administration

2.1. Starting the Server

The server must be started from within the Xindice directory. A future revision of the server will fix this limitation.

2.1.1. Starting the Server on UNIX

```
cd Xindice
./start
```

2.1.2. Starting the Server on Windows

```
cd Xindice
startup
```

2.2. Stopping the Server

The Xindice server can be easily shutdown from the command line. You must provide the name of the server instance to shutdown.

2.2.1. Stopping the Server

This example assumes that the Xindice/bin directory is in your path.

```
xindice shutdown -c /db
```

2.3. Backing up Your Data

Currently backing up Xindice consists of simply shutting down the server and copying the entire contents of the Xindice/db directory to the backup media.

2.3.1. Backing up the server

This example assumes that the Xindice/bin directory is in your path.

```
cd Xindice
xindice shutdown
cp -pr db /backup/db
./start
```

2.3.2. Restoring the Data

Restoring the data is simply removing the current database and reversing the backup process. This example assumes that the Xindice/bin directory is in your path.

```
cd Xindice
xindice shutdown
rm -rf db
cp -pr /backup/db db
./start
```

2.4. Exporting the Contents of the Database

Xindice includes tools to export data to a directory hierarchy and to also import data from a directory hierarchy. Each directory in the hierarchy corresponds to a collection in Xindice. Each XML document is stored in a separate file named with the key from the database.

2.4.1. Exporting the database

This example assumes that the Xindice/bin directory is in your path.

```
xindice export -c /db/root -f /path/to/data
```

The entire contents of the collection /db/root will be exported to the directory /path/to/data.

2.4.2. Importing the database

This example assumes that the Xindice/bin directory is in your path.

Xindice 1.0 Administration Guide

```
xindice import -c /db -f /path/to/data/root
```

Each directory under /path/to/data will be used to create a collection and all XML documents in the hierarchy will be imported in to the database. You can also restrict the documents that are imported by adding -i and the extension of the files you want to import.