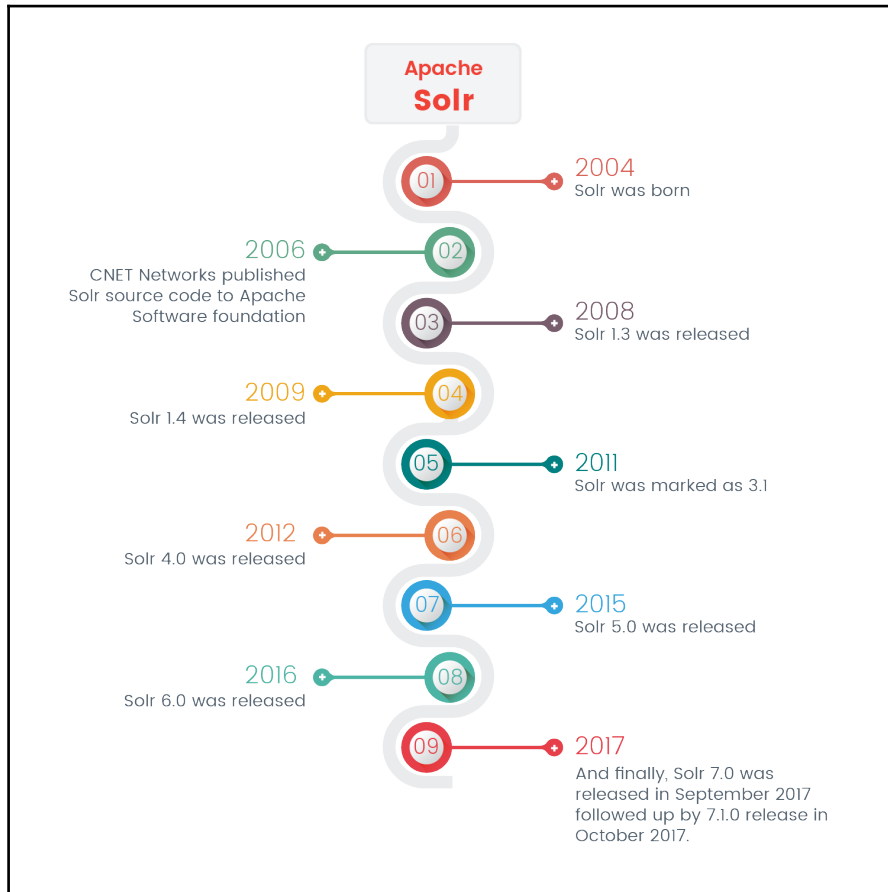
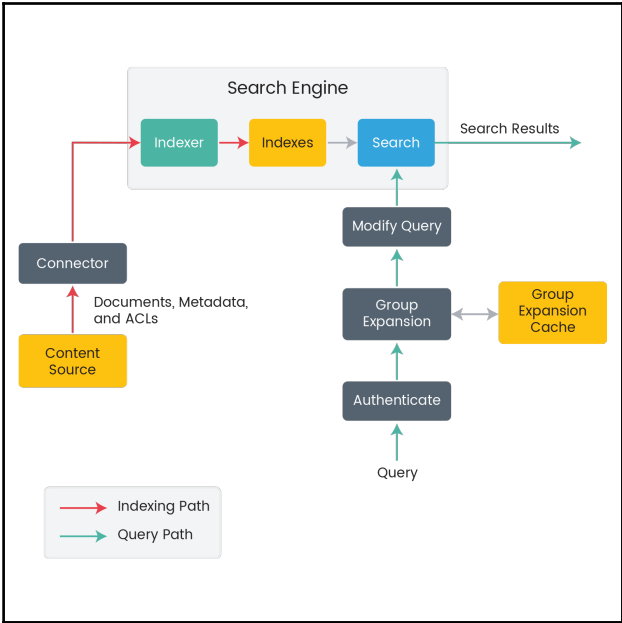


# Chapter 1: Introduction to Solr 7



## Solr v.s. relationalDB

Lucene	Solr	Relational DB
Text Search	Fast and sophisticated	Minimal and slow
Features	Few, targeted to text search	Many
Deployment Complexity	Medium	Medium
Administration Tools	Minimal open source projects	Many open source & commercial
Monitoring Tools	Weak	Very Strong
Scaling Tools	Automated, medium scale	Large scale
Support Availability	Weak	Strong
Schema Flexibility	Must in general rebuild	Changes immediately visible
Indexing Speed	Slow	Faster and adjustable
Query Speed	Text search is fast & predictable	Very dependent on design & use case
Row Addition/Extraction Speed	Slow	Fast
Partial Record Modification	No	Yes
Time to visibility after addition	Slow	Immediate
Access to internal data structures	High	None
Technical knowledge required	Java (minimal), web server deployment, IT	SQL, DB-specific factors, IT
Regular maintenance tasks		





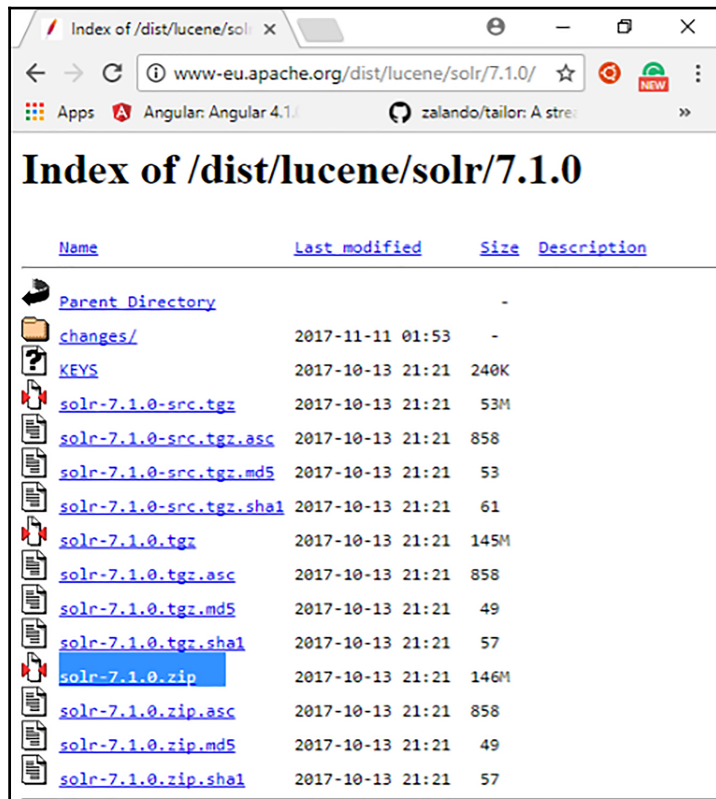
---

## Chapter 2: Getting Started

```
C:\WINDOWS\system32\cmd.exe
Microsoft Windows [Version 10.0.16299.19]
(c) 2017 Microsoft Corporation. All rights reserved.

C:\Users\chint>java -version
java version "1.8.0_151"
Java(TM) SE Runtime Environment (build 1.8.0_151-b12)
Java HotSpot(TM) 64-Bit Server VM (build 25.151-b12, mixed mode)

C:\Users\chint>
```



This PC > Local Disk (E:) > solr-7.1.0 > solr-7.1.0

Name	Date modified	Type	Size
bin	10/13/2017 4:16 PM	File folder	
contrib	10/13/2017 4:16 PM	File folder	
dist	10/13/2017 4:16 PM	File folder	
docs	10/13/2017 4:16 PM	File folder	
example	10/13/2017 4:16 PM	File folder	
licenses	10/13/2017 4:16 PM	File folder	
server	10/13/2017 4:16 PM	File folder	
CHANGES.txt	10/13/2017 4:16 PM	Text Document	738 KB
LICENSE.txt	10/13/2017 4:16 PM	Text Document	13 KB
LUCENE_CHANGES.txt	10/13/2017 4:16 PM	Text Document	661 KB
NOTICE.txt	10/13/2017 4:16 PM	Text Document	25 KB
README.txt	10/13/2017 4:16 PM	Text Document	8 KB

```
C:\Windows\System32\cmd.exe - solr -e cloud
E:\solr-7.1.0\solr-7.1.0\bin>solr -e cloud

Welcome to the SolrCloud example!

This interactive session will help you launch a SolrCloud cluster on your local workstation.
To begin, how many Solr nodes would you like to run in your local cluster? (specify 1-4 nodes) [2]:
```

```
C:\Windows\System32\cmd.exe - solr -e cloud
Welcome to the SolrCloud example!

This interactive session will help you launch a SolrCloud cluster on your local workstation.
To begin, how many Solr nodes would you like to run in your local cluster? (specify 1-4 nodes) [2]:
Ok, let's start up 2 Solr nodes for your example SolrCloud cluster.
Please enter the port for node1 [8983]:
Please enter the port for node2 [7574]:

Creating Solr home directory E:\solr-7.1.0\solr-7.1.0\example\cloud\node1\solr
Cloning E:\solr-7.1.0\solr-7.1.0\example\cloud\node1 into
E:\solr-7.1.0\solr-7.1.0\example\cloud\node2

Starting up Solr on port 8983 using command:
"E:\solr-7.1.0\solr-7.1.0\bin\solr.cmd" start -cloud -p 8983 -s "E:\solr-7.1.0\solr-7.1.0\example\cloud\node1\solr"

Waiting up to 30 to see Solr running on port 8983
Started Solr server on port 8983. Happy searching!

Starting up Solr on port 7574 using command:
"E:\solr-7.1.0\solr-7.1.0\bin\solr.cmd" start -cloud -p 7574 -s "E:\solr-7.1.0\solr-7.1.0\example\cloud\node2\solr" -z localhost:9983

Waiting up to 30 to see Solr running on port 7574
INFO - 2017-11-11 09:15:16.356; org.apache.solr.client.solrj.impl.ZkClientClusterStateProvider; Cluster at localhost:9983 ready

Now let's create a new collection for indexing documents in your 2-node cluster.
Please provide a name for your new collection: [gettingstarted]
Started Solr server on port 7574. Happy searching!
```

The screenshot displays the Solr Admin web interface in a browser window. The address bar shows the URL `localhost:8983/solr/#/`. The interface is divided into several sections:

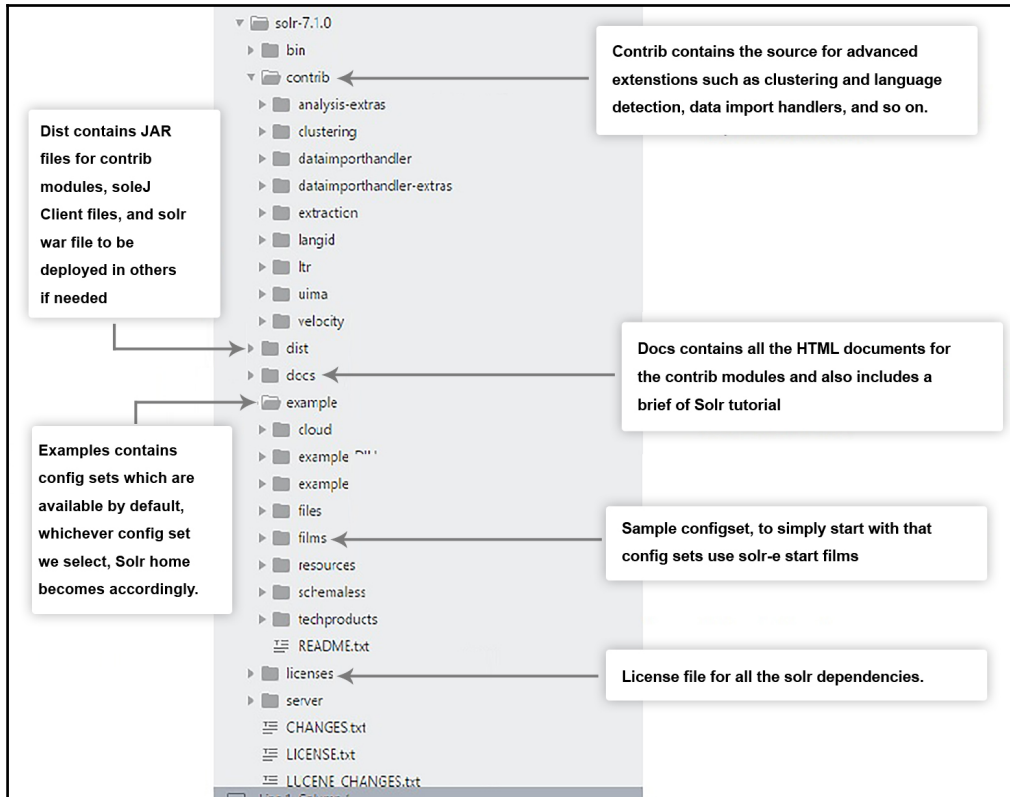
- Instance:** Shows the instance was started 3 minutes ago.
- Versions:** Lists the installed versions:
  - `solr-spec` 7.1.0
  - `solr-impl` 7.1.0 84c90ad2c0218156c840e19a64d72b8a38550659 - ubuntu - 2017-10-13 16:15:59
  - `lucene-spec` 7.1.0
  - `lucene-impl` 7.1.0 84c90ad2c0218156c840e19a64d72b8a38550659 - ubuntu - 2017-10-13 16:12:42
- JVM:**
  - Runtime:** Oracle Corporation Java HotSpot(TM) 64-Bit Server VM 1.8.0\_60 25.60-b23
  - Processors:** 4
  - Args:**

```

-DSTOP.KEY=solrlocks
-DSTOP.PORT=7983
-Djava.io.tmpdir=E:\solr-7.1.0\solr-7.1.0\server\tmp
-Djetty.home=E:\solr-7.1.0\solr-7.1.0\server
-Djetty.host=0.0.0.0
-Djetty.port=8983
-Dlog4j.configuration=file:E:\solr-7.1.0\solr-7.1.0\example\resources\log4j.properties
-Dsolr.default.confdir=E:\solr-7.1.0\solr-7.1.0\server\solr\configsets\_default\conf
-Dsolr.install.dir=E:\solr-7.1.0\solr-7.1.0
-Dsolr.log.dir=E:\solr-7.1.0\solr-7.1.0\server\logs
-Dsolr.log.muteconsole
-Dsolr.solr.home=E:\solr-7.1.0\solr-7.1.0\example\cloud\node1\solr
-Duser.timezone=UTC
-DzkClientTimeout=15000
-DzkRun
-XX:+CMSParallelRemarkEnabled

```
- System:**
  - Physical Memory:** 66.3% used (5.24 GB / 7.90 GB)
  - Swap Space:** 65.1% used (10.68 GB / 16.40 GB)
  - JVM-Memory:** 35.22 MB used (490.69 MB / 490.69 MB)

On the left sidebar, there are navigation options: Dashboard, Logging, Cloud, Collections, Java Properties, and Thread Dump. Below these are status indicators: "No collections available" and "No cores available".






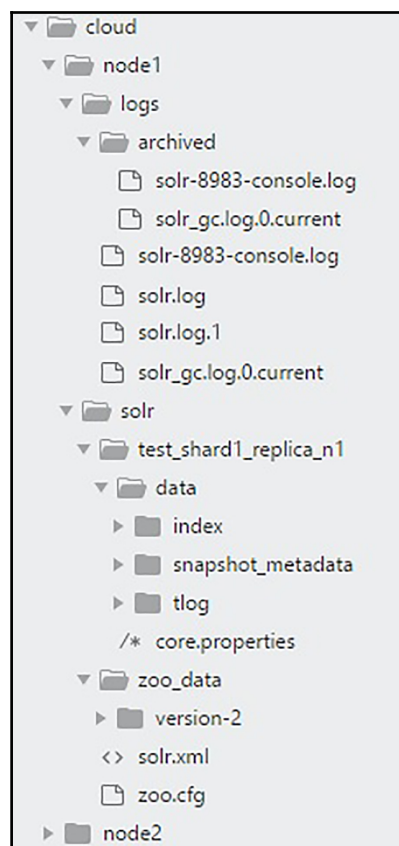
```
Select C:\Windows\System32\cmd.exe
Microsoft Windows [Version 10.0.15063]
(c) 2017 Microsoft Corporation. All rights reserved.

E:\solr-7.1.0\solr-7.1.0\example\exampledocs>java -jar post.jar -h
SimplePostTool version 5.0.0
Usage: java [SystemProperties] -jar post.jar [-h|-] [<file|folder|url|arg> [<file|folder|url|arg>...]]

Supported System Properties and their defaults:
-Dc=<core/collection>
-Durl=<base Solr update URL> (overrides -Dc option if specified)
-Ddata=files|web|args|stdin (default=files)
-Dtype=<content-type> (default=application/xml)
-Dhost=<host> (default: localhost)
-Dport=<port> (default: 8983)
-Dbasicauth=<user:pass> (sets Basic Authentication credentials)
-Dauto=yes|no (default=no)
-Drecursive=yes|no <depth> (default=0)
-Ddelay=<seconds> (default=0 for files, 10 for web)
-Dfiletypes=<type>[, <type>,...] (default=xml,json,jsonl, csv,pdf, doc,docx, ppt,pptx, xls,xlsx, odt,odp, ods,ott, otp,ots, rtf,htm,html,txt,log)
-Dparams="<key>=<value>[&<key>=<value>...]" (values must be URL-encoded)
-Dcommit=yes|no (default=yes)
-Doptimize=yes|no (default=no)
-Dout=yes|no (default=no)

This is a simple command line tool for POSTing raw data to a Solr port.
NOTE: Specifying the url/core/collection name is mandatory.
Data can be read from files specified as commandline args,
URLs specified as args, as raw commandline arg strings or via STDIN.
Examples:
java -Dc=gettingstarted -jar post.jar *.xml
java -Ddata=args -Dc=gettingstarted -jar post.jar '<delete><id>42</id></delete>'
java -Ddata=stdin -Dc=gettingstarted -jar post.jar < hd.xml
java -Ddata=web -Dc=gettingstarted -jar post.jar http://example.com/
java -Dtype=text/csv -Dc=gettingstarted -jar post.jar *.csv
java -Dtype=application/json -Dc=gettingstarted -jar post.jar *.json
java -Durl=http://localhost:8983/solr/techproducts/update/extract -Dparams=literal.id=pdf1 -jar post.jar solr-word.pdf
java -Dauto -Dc=gettingstarted -jar post.jar *
java -Dauto -Dc=gettingstarted -Drecursive -jar post.jar afolder
java -Dauto -Dc=gettingstarted -Dfiletypes=ppt,html -jar post.jar afolder

The options controlled by System Properties include the Solr
URL to POST to, the Content-Type of the data, whether a commit
or optimize should be executed, and whether the response should
be written to STDOUT. If auto=yes the tool will try to set type
automatically from file name. When posting rich documents the
```





Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

```
mysql> use parth;
Database changed
mysql> desc archives;
```

Field	Type	Null	Key	Default	Extra
category_id	int(11)	NO	PRI	NULL	auto_increment
category_name	varchar(150)	YES		NULL	
remarks	varchar(500)	YES		NULL	

```
3 rows in set (0.00 sec)

mysql> select * from archives;
```

category_id	category_name	remarks
1	parth	test test
2	ghiya	test1 test1
3	random	test1 test1
4	perplex	test1 test1
5	mystified	test1 test1

```
5 rows in set (0.00 sec)
```

**Solr**

Dashboard  
Logging  
Core Admin  
Java Properties  
Thread Dump

db

Overview  
Analysis  
Dataimport  
Documents  
Files  
Ping  
Plugins / Stats  
Query  
Replication  
Schema  
Segments info

**/dataimport**

Command: full-import

Verbose  
 Clean  
 Commit  
 Optimize  
 Debug

Entity: archives

Custom Parameters: key1=val1&key2=val2

**Execute** Refresh Status

Auto-Refresh Status

Last Update: 04:51:14  
✔ **Indexing completed. Added/Updated: 5 documents. Deleted 0 documents. (Duration: 01s)**  
Requests: 1 1/s, Fetched: 5 5/s, Skipped: 0, Processed: 5 5/s  
Started: about an hour ago

Raw Status-Output  
Configuration

4. You will be able to see Fetched Successfully Message.

2. Click on Entity and select Archives

3. After selecting full import and the entity name Click on Execute

1. Click on Dataimport

Documentation Issue Tracker IRC Channel Community forum

Solr Admin

Solr

Type of Search: **Simple** Spatial Group By

Find:

Boost by Price

**Field Facets** 32 results found in 277 ms Page 1 of 4

cat

- [electronics](#) (12)
- [currency](#) (4)
- [memory](#) (3)
- [connector](#) (2)
- [graphics\\_card](#) (2)
- [hard\\_drive](#) (2)
- [search](#) (2)
- [software](#) (2)
- [camera](#) (1)
- [copier](#) (1)
- [electronics\\_and\\_c...](#) (1)
- [electronics\\_and\\_s...](#) (1)
- [multifunction\\_pri...](#) (1)
- [music](#) (1)
- [printer](#) (1)
- [scanner](#) (1)
- [missing](#) (12)
- manu\_exact**

**Test with some GB18030 encoded characters** [More Like This](#)

Id: GB18030TEST

Price: 0.0,USD

Features: No accents here ... 这是一个功能 ... This is a feature (translated) ... 这份文件是很有光泽 ... This document is very shiny (trans

In Stock: true

**Samsung SpinPoint P120 SP2514N - hard drive - 250 GB - ATA-133** [More Like This](#) [Larger Map](#)

Id: SP2514N

Price: 92.0,USD

Features: 7200RPM, 8MB cache, IDE Ultra ATA-133 ... NoiseGuard, SilentSeek technology, Fluid Dynamic Bearing (FDB) motor

In Stock: true

**Maxtor DiamondMax 11 - hard drive - 500 GB - SATA-300** [More Like This](#) [Larger Map](#)

Id: 6H500F0

Price: 250.0,USD

The screenshot displays a Solr search interface with the following elements and annotations:

- Types of searches available:** A callout box pointing to the search type selection area.
- Do the nearest around me match:** A callout box pointing to the 'Distance (KM): 10' filter.
- Faceted Search:** A callout box pointing to the 'Field Facets' section on the left.
- Individual Search Result, the first is one whose score matches the most as per user query entered:** A callout box pointing to the first search result, 'Test with some GB18030 encoded characters'.

**Search Interface Details:**

- Search Type:** Simple (selected), Spatial, Group By.
- Find:** [Empty search bar]
- Filters:** Boost by Price (checkbox), Location Filter: No Filter, Distance (KM): 10.
- Field Facets:**
  - cat: electronics (12), currency (4), memory (3), connector (2), graphics\_card (2), hard\_drive (2), search (2), software (2), camera (1), copier (1), electronics\_and\_c... (1), electronics\_and\_s... (1), multifunction\_pri... (1), music (1), printer (1), scanner (1), missing (12).
- Search Results:** 32 results found in 14 ms Page 1 of 4.
  - Result 1:** Test with some GB18030 encoded characters [More Like This](#). Id: GB18030TEST, Price: 0.0,USD, Features: No accents here ... 这是一个功能 ... This is a feature (translated) ... 这份文件是很有光泽 ... This document is very shiny (translated), In Stock: true.
  - Result 2:** Samsung SpinPoint P120 SP2514N - hard drive - 250 GB - ATA-133 [More Like This](#). Id: SP2514N, Price: 92.0,USD, Features: 7200RPM, 8MB cache, IDE Ultra ATA-133 ... NoiseGuard, SilentSeek technology, Fluid Dynamic Bearing (FDB) motor, In Stock: true.
  - Result 3:** Maxtor DiamondMax 11 - hard drive - 500 GB - SATA-300 [More Like This](#).

Log4j (org.slf4j.impl.Log4jLoggerFactory)					
Time (Local)	Level	Level	Core	Logger	Message
11/11/2017, 1:37:00 PM	WARN	false		ClientCnxn	Client session timed out,&#8203; have not heard from server in 1184144ms for sessionId 0x15fa9dc7cb50000
11/11/2017, 1:37:00 PM	WARN	true		NIOServerCnxn	caught end of stream exception EndOfStreamException: Unable to read additional data from client sessionId 0x15fa9dc7cb50000, likely c1 at org.apache.zookeeper.server.NIOServerCnxn.doIO(NIOServerCnxn.java:239) at org.apache.zookeeper.server.NIOServerCnxnFactory.run(NIOServerCnxnFactory.java:283) at java.lang.Thread.run(Thread.java:745)
11/11/2017, 1:37:00 PM	WARN	false		ConnectionManager	Watcher org.apache.solr.common.cloud.ConnectionManager@4e85ac7 name: ZooKeeperConnection Watcher:localhost:9983 got event WatchedEvent state:Disconnected type:None path:null path: null type: None
11/11/2017, 1:37:00 PM	WARN	false		ConnectionManager	zkClient has disconnected
11/11/2017, 1:37:00 PM	WARN	false		ClientCnxn	Unable to reconnect to ZooKeeper service,&#8203; session 0x15fa9dc7cb50000 has expired
11/11/2017, 1:37:00 PM	WARN	false		ConnectionManager	Watcher org.apache.solr.common.cloud.ConnectionManager@4e85ac7 name: ZooKeeperConnection Watcher:localhost:9983 got event WatchedEvent state:Expired type:None path:null path: null type: None
11/11/2017, 1:37:00 PM	WARN	false		ConnectionManager	Our previous ZooKeeper session was expired. Attempting to reconnect to recover relationship with ZooKeeper...
11/11/2017, 1:37:01 PM	WARN	false		Overseer	Solr cannot talk to ZK,&#8203; exiting Overseer main queue loop
11/11/2017, 1:37:01 PM	WARN	false		OverseerTriggerThread	OverseerTriggerThread woken up but we are closed,&#8203; exiting.
11/11/2017, 1:37:01 PM	ERROR	false		Overseer	could not read the data
11/11/2017, 1:37:01 PM	WARN	false		DefaultConnectionStrategy	Connection expired - starting a new one...



The screenshot displays the Solr logging configuration page for Log4j. On the left, a sidebar contains navigation options: Dashboard, Logging, Level (selected), Cloud, Collections, Java Properties, and Thread Dump. Below these are dropdown menus for 'Collection Sele...' and 'Core Selector'.

The main area shows a tree view of the logging configuration hierarchy. The root is 'root' with level 'INFO'. Underneath are 'org' and 'apache'. The 'org' directory is expanded to show 'hadoop' (highlighted in bold) and 'http'. The 'hadoop' directory is further expanded to show 'client', 'conn', 'ssl', 'headers', and 'impl'. The 'impl' directory is expanded to show 'auth' and 'HttpAuthenticator'. The 'http' directory is expanded to show 'protocol', 'client', and 'conn'. The 'protocol' directory is expanded to show 'RequestAddCookies', 'RequestAuthCache', 'RequestClientConnControl', and 'ResponseProcessCookies'. The 'client' directory is expanded to show 'protocol', 'conn', and 'ssl'. The 'conn' directory is expanded to show 'ssl'. The 'ssl' directory is expanded to show 'AllowAllHostnameVerifier', 'BrowserCompatHostnameVerifier', 'DefaultHostnameVerifier', 'SSLConnectionSocketFactory', and 'StrictHostnameVerifier'. The 'headers' directory is expanded to show 'impl'. The 'impl' directory is expanded to show 'auth' and 'HttpAuthenticator'.

On the right side, there is a list of log levels: ALL, TRACE, DEBUG, INFO, WARN, ERROR, FATAL, OFF, and UNSET. The 'WARN' level is selected. A callout box points to this list with the text: 'You can select any level of logs for a given class'. Another callout box points to the 'hadoop' directory with the text: 'Items highlighted in bold indicate that any child level changes will not affect the parent level logs.' A third callout box points to the 'Level' button in the sidebar with the text: 'Click here to change log level settings for any class in Solr'.

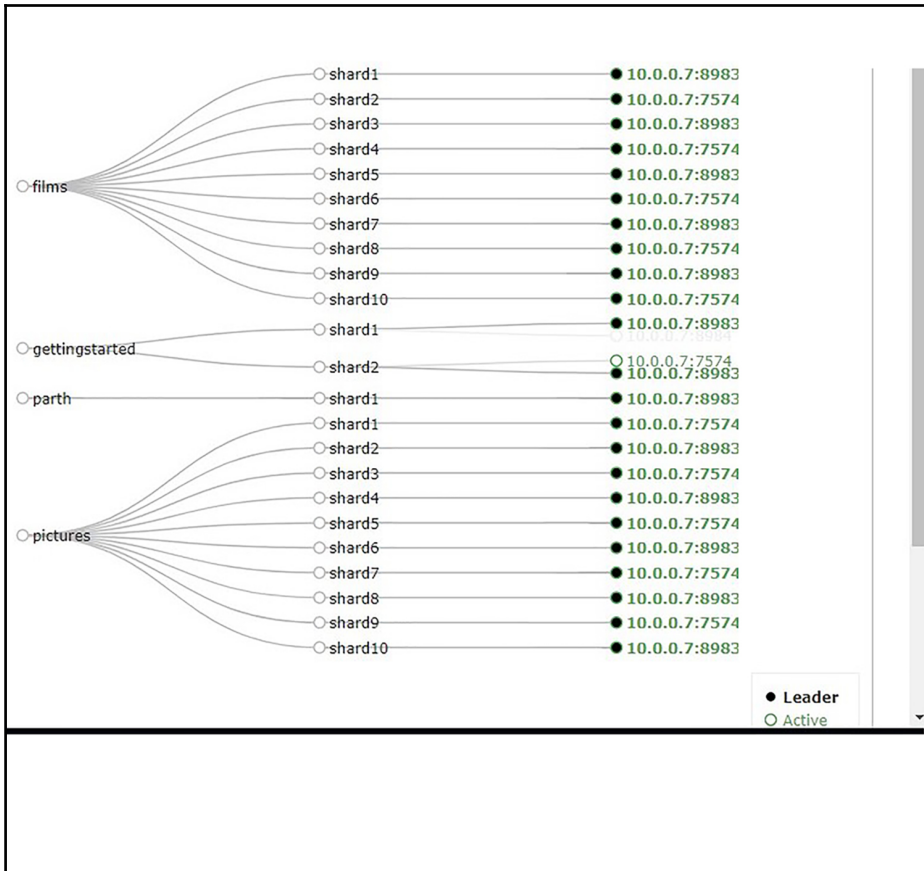
The image shows a ZooKeeper node with a directory tree on the left and a JSON configuration on the right. The directory tree includes folders like /autoscaling.json, /clusterstate.json, /collections, films, leader\_elect, leaders, shard1 through shard9, state.json, gettingstarted, parth, pictures, test11, /configs, and /live\_nodes. The JSON configuration contains metadata and a 'films' object with 'shards' and 'replicas' information. Annotations include 'Directory structure of data in ZooKeeper's' pointing to the left pane, and 'Shard and its replica Info.' pointing to the 'shard1' and 'replicas' fields in the JSON.

ctime	Sun Nov 12 20:56:47 UTC 2017 (1510520207720)
cversion	0
czxid	3312
dataLength	3700
ephemeralOwner	0
mtime	Mon Nov 13 01:52:04 UTC 2017 (1510537924913)
mzxid	4157
pzxid	3312
version	23

```

{"films":{
  "pullReplicas":"0",
  "replicationFactor":"1",
  "shards":{
    "shard1":{
      "range":"80000000-9998ffff",
      "state":"active",
      "replicas":{"core_node3":{
        "core":"films_shard1_replica_n1",
        "base_url":"http://10.0.0.7:8983/solr",
        "node_name":"10.0.0.7:8983_solr",
        "state":"active",
        "type":"NRT",
        "leader":"true"}}},
    "shard2":{

```



The screenshot shows the Solr Admin interface for the 'films' collection. The left sidebar contains navigation options: Dashboard, Logging, Cloud, Collections (selected), Java Properties, and Thread Dump. The main content area is divided into several sections:

- Collection: films**: Shows metadata including shard count (10), configName (schemaless), replicationFactor (1), maxShardsPerNode (-1), router (compositeId), and autoAddReplicas (false).
- Shard: shard1**: Shows details for the first shard, including its state (active), range, and core (core\_node3). A specific replica is highlighted with a callout: "Replica information for shard 1, which resides in core\_node3".
- Replica: core\_node3**: Shows details for a replica, including its state (active) and base URL (http://10.0.0.7:8983/solr).
- Available shards for Collection**: Lists shards from shard2 to shard7.
- Buttons**: Includes "Add Collection", "Delete", "Reload", "Create Alias", "Delete Alias", and "add replica".

Annotations on the screenshot include:

- "Available Collections" pointing to the left sidebar.
- "Based on the number of instances you are running, this will be Collections/Core Admin" pointing to the sidebar.
- "Metadata about films Collection" pointing to the collection metadata.
- "Available shards for Collection" pointing to the shard list.
- "Replica information for shard 1, which resides in core\_node3" pointing to the replica details.
- "As it's inactive, we can delete it" pointing to a red 'X' icon on the replica details.
- "Option to add more replica" pointing to the "add replica" button.

Angular application bundles  
[www.syntaxsuccess.com/viewarticle/angular-application-bundles](http://www.syntaxsuccess.com/viewarticle/angular-application-bundles)

STOP.KEY	solrrocks
STOP.PORT	6574
awt.toolkit	sun.awt.windows.WToolkit
file.encoding	Cp1252
file.encoding.pkg	sun.io
file.separator	\
java.awt.graphicsenv	sun.awt.Win32GraphicsEnvironment
java.awt.printerjob	sun.awt.windows.WPrinterJob
java.class.path	E:\solr-7.1.0\solr-7.1.0\server\lib\gmetric4j-1.0.7.jar E:\solr-7.1.0\solr-7.1.0\server\lib\javax.servlet-api-3.1.0.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-continuation-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-deploy-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-http-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-io-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-jmx-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-rewrite-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-security-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-server-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-servlet-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-servlets-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-util-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-webapp-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\jetty-xml-9.3.20.v20170531.jar E:\solr-7.1.0\solr-7.1.0\server\lib\metrics-core-3.2.2.jar E:\solr-7.1.0\solr-7.1.0\server\lib\metrics-ganglia-3.2.2.jar E:\solr-7.1.0\solr-7.1.0\server\lib\metrics-graphite-3.2.2.jar E:\solr-7.1.0\solr-7.1.0\server\lib\metrics-jetty9-3.2.2.jar E:\solr-7.1.0\solr-7.1.0\server\lib\metrics-jvm-3.2.2.jar E:\solr-7.1.0\solr-7.1.0\server\lib\ext\icl-over-slf4j-1.7.7.jar

Solr Admin

10.0.0.7:7574/solr/#/~-threads

Dashboard  
Logging  
Cloud  
Collections  
Java Properties  
Thread Dump

Collection Sele...  
Core Selector

Show all Stacktraces

name	cpuTime / userTime
searcherExecutor-77-thread-1 (80)	0.0000ms / 0.0000ms
searcherExecutor-22-thread-1 (52)	15.6250ms / 0.0000ms
searcherExecutor-9-thread-1 (46)	0.0000ms / 0.0000ms
searcherExecutor-8-thread-1 (45)	0.0000ms / 0.0000ms
searcherExecutor-7-thread-1 (44)	0.0000ms / 0.0000ms
Scheduler-610984013 (40)	0.0000ms / 0.0000ms
DestroyJavaVM (38)	2875.0000ms / 2125.0000ms
Thread-12 (35)	0.0000ms / 0.0000ms

There's some issue, and here is the stacktrace.

- sun.misc.Unsafe.park(Native Method)
- java.util.concurrent.locks.LockSupport.park(LockSupport.java:175)
- java.util.concurrent.locks.AbstractQueuedSynchronizer\$ConditionObject.await(AbstractQueuedSynchronizer.java:2039)
- java.util.concurrent.LinkedBlockingQueue.take(LinkedBlockingQueue.java:442)
- java.util.concurrent.ThreadPoolExecutor.getTask(ThreadPoolExecutor.java:1067)
- java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1127)
- java.util.concurrent.ThreadPoolExecutor\$Worker.run(ThreadPoolExecutor.java:617)
- java.lang.Thread.run(Thread.java:745)

Executed successfully.

Solr Admin

10.0.0.7:7574/solr/#/films/analysis?analysis.fieldvalue=Harry%20Potter&analysis.query=Har%20pot&analysis.fieldname=\_text\_&verbose\_output=1

Apps Angular: Angular 4.1.1 Angular application Migrating my Angular What Angular is doing incubator-openwhisk The Apache Software Project Mosaic

Solr

Field Value (Index): Harry Potter

Field Value (Query): Har pot

Analyse Fieldname / FieldType:  Schema Browser  Verbose Output [Analyse Values](#)

ST	text	Harry	Potter	ST	text	Har	pot
	raw_bytes	[48 61 72 72 79]	[50 6f 74 74 65 72]		raw_bytes	[48 61 72]	[70 6f 74]
	start	0	6		start	0	4
	end	5	12		end	3	7
	positionLength	1	1		positionLength	1	1
	type	<ALPHANUM>	<ALPHANUM>		type	<ALPHANUM>	<ALPHANUM>
	termFrequency	1	1		termFrequency	1	1
	position	1	2		position	1	2

SE	text	Harry	Potter	SE	text	Har	pot
	raw_bytes	[48 61 72 72 79]	[50 6f 74 74 65 72]		raw_bytes	[48 61 72]	[70 6f 74]
	start	0	6		start	0	4
	end	5	12		end	3	7
	positionLength	1	1		positionLength	1	1
	type	<ALPHANUM>	<ALPHANUM>		type	<ALPHANUM>	<ALPHANUM>
	termFrequency	1	1		termFrequency	1	1
	position	1	2		position	1	2

LCF	text	harry	potter	SGF	text	Har	pot
	raw_bytes	[68 61 72 72 79]	[70 6f 74 74 65 72]		raw_bytes	[48 61 72]	[70 6f 74]
	start	0	6		start	0	4

Dashboard

Logging

Cloud

Collections

Java Properties

Thread Dump

films

Overview

Analysis

Dataimport

Documents

Files

Query

Stream

Schema

Core Selector

**Solr**

- Dashboard
- Logging
- Cloud
- Collections
- Java Properties
- Thread Dump

films

- Overview
- Analysis
- Dataimport
- Documents
- Files
- Query
- Stream
- Schema

Request-Handler (qt)  
/update

Document Type  
Document Builder

Field: id

Field Data: Enter your field text here and then click 'Add Field' to add the field to the document.

+ Add Field

Document(s)

```
{  
  "genre": "jazz",  
  "id": "1"  
}
```

Commit Within  
1000

Overwrite  
true

Submit Document

Documentation Issue Tracker IRC Channel Community forum Solr Query Syntax

Core Selector



The screenshot shows the Solr Admin web interface. On the left is a sidebar with navigation links: Dashboard, Logging, Cloud, Collections, Java Properties, Thread Dump, and a 'files' dropdown menu. Below the dropdown is a list of files in the 'lang' directory, including 'contractions\_ca.txt', 'contractions\_fr.txt', 'contractions\_ga.txt', 'contractions\_it.txt', 'hyphenations\_ga.txt', 'stemdict\_nl.txt', 'stoptags\_ja.txt', 'stopwords\_ar.txt', 'stopwords\_bg.txt', 'stopwords\_ca.txt', 'stopwords\_cz.txt', 'stopwords\_da.txt', 'stopwords\_de.txt', 'stopwords\_el.txt', 'stopwords\_en.txt', 'stopwords\_es.txt', 'stopwords\_eu.txt', 'stopwords\_fa.txt', 'stopwords\_fi.txt', 'stopwords\_fr.txt', 'stopwords\_ga.txt', 'stopwords\_gl.txt', 'stopwords\_hi.txt', 'stopwords\_hu.txt', 'stopwords\_hy.txt', and 'stopwords\_id.txt'. The main content area displays the XML schema for the selected file, starting with the root element <schema name="default-config" version="1.6">. The schema defines several field types and analyzers, including 'ancestors\_path', 'binary', 'boolean', 'booleans', 'delimited\_payloads\_float', 'delimited\_payloads\_int', and 'delimited\_payloads\_string'. Each field type is associated with a specific analyzer, such as 'solr.KeywordTokenizerFactory' or 'solr.WhitespaceTokenizerFactory'. The schema also includes filters like 'solr.DelimitedPayloadTokenFilterFactory' for float and integer payloads.

Solr Admin

10.0.0.7:7574/solr/#/films/query

Angular application bundles  
www.syntaxsuccess.com/viewarticle/angular-application-bundles

Request-Handler (qt) /select

— common —

q

fq

sort

start, rows  
0 0

fl

df

Raw Query Parameters  
key1=val1&key2=val2

wt

indent off

debugQuery

dismax

```
{
  "responseHeader": {
    "zkConnected": true,
    "status": 0,
    "QTime": 41,
    "params": {
      "q": "",
      "facet.field": "genre",
      "rows": "0",
      "facet": "on",
      "_": "1510544922787"
    }
  },
  "response": {
    "numFound": 1100,
    "start": 0,
    "maxScore": 1.0,
    "docs": []
  },
  "facet_counts": {
    "facet_queries": {},
    "facet_fields": {
      "genre": [
        "film", 793,
        "drama", 569,
        "comedy", 417,
        "romance", 270,
        "thriller", 266,
        "fiction", 263,
        "action", 208,
        "crime", 191,
        "cinema", 184,
        "adventure", 167,
        "world", 167,
        "indie", 144
      ]
    }
  }
}
```

The screenshot shows the Solr Admin web interface in a browser. The address bar displays `10.0.0.7:7574/solr/#/films/stream`. The left sidebar contains navigation options: Dashboard, Logging, Cloud, Collections, Java Properties, Thread Dump, and a dropdown menu for the 'films' collection. The 'Stream' option is selected in the sidebar. The main content area is titled 'Streaming Expression (expr)' and contains the text `add(5,5)`. Below this is an 'Execute' button and a checkbox for 'with explanation', which is checked. A URL bar shows `http://10.0.0.7:7574films/stream?explain=true&expr=add(5,5)`. A legend below the URL bar identifies components: Stream Decorator (purple), Stream Source (blue), Graph Source (green), and Datastore (dark blue). The 'tuple' component is selected. The JSON response is displayed in a code block:

```
{
  "result-set": {
    "docs": [
      {
        "return-value": 10
      },
      {
        "EOF": true,
        "RESPONSE_TIME": 0
      }
    ]
  }
}
```

At the bottom of the interface, there are links for Documentation, Issue Tracker, IRC Channel, Community forum, and Solr Query Syntax.

**Solr**

10.0.0.7:7574/solr/#/films/schema?field=genre

Buttons: Add Field, Add Dynamic Field, Add Copy Field

Field: genre

Field-Type: org.apache.solr.schema.TextField

P Gap: 100

Docs: 111

Copied to: genre\_str

Type: text\_general

Flags: Indexed, Tokenized, Stored, Multivalued

Properties	Indexed	Tokenized	Stored	Multivalued
Properties	✓	✓	✓	✓
Schema	✓	✓	✓	✓
Index	✓	✓	✓	

Index Analyzer: org.apache.solr.analysis.TokenizerChain

Query Analyzer: org.apache.solr.analysis.TokenizerChain

Unique Key Field: id

Global Similarity: SchemaSimilarity. Default: BM25(k1=1.2,b=0.75)

Load Term Info: 10 / 100 Top-Terms

N.B. Loaded from a single core - not from the whole collection.

Count	Term
81	film
65	drama
41	comedy
28	fiction
27	thriller
25	romance
22	action
20	adventure

Autoload

Histogram:

Count	Term
1	38
2	19
4	15
8	10
16	8
32	7
64	1
128	2

The screenshot shows the Solr Admin web interface. On the left is a navigation sidebar with the Solr logo and menu items: Dashboard, Logging, Cloud, Collections, Java Properties, Thread Dump, Collection Sele..., films\_shard10\_..., Overview, Ping (0ms), Plugins / Stats, and Segments info. The main content area is divided into two columns. The left column lists administrative categories: ADMIN, CACHE (selected), CORE, HIGHLIGHTER, OTHER, QUERY, QUERYPARSER, REPLICATION, SPELLCHECKER, and UPDATE. Below these are 'Watch Changes' and 'Refresh Values' buttons, followed by a note: 'NOTE: Only selected metrics are shown here. Full metrics can be accessed via /admin/metrics handler.' The right column displays the configuration for the 'documentCache' plugin, including its class, description, and a list of stats. At the bottom of the interface are links for Documentation, Issue Tracker, IRC Channel, Community forum, and Solr Query Syntax.

**Solr**

- Dashboard
- Logging
- Cloud
- Collections
- Java Properties
- Thread Dump
- Collection Sele...
- films\_shard10\_...
- Overview
- Ping (0ms)
- Plugins / Stats
- Segments info

**ADMIN**

- CACHE**
- CORE
- HIGHLIGHTER
- OTHER
- QUERY
- QUERYPARSER
- REPLICATION
- SPELLCHECKER
- UPDATE

Watch Changes

Refresh Values

*NOTE: Only selected metrics are shown here. Full metrics can be accessed via /admin/metrics handler.*

**documentCache**

class: org.apache.solr.search.LRUCache

description: LRU Cache(maxSize=512, initialSize=...

stats:

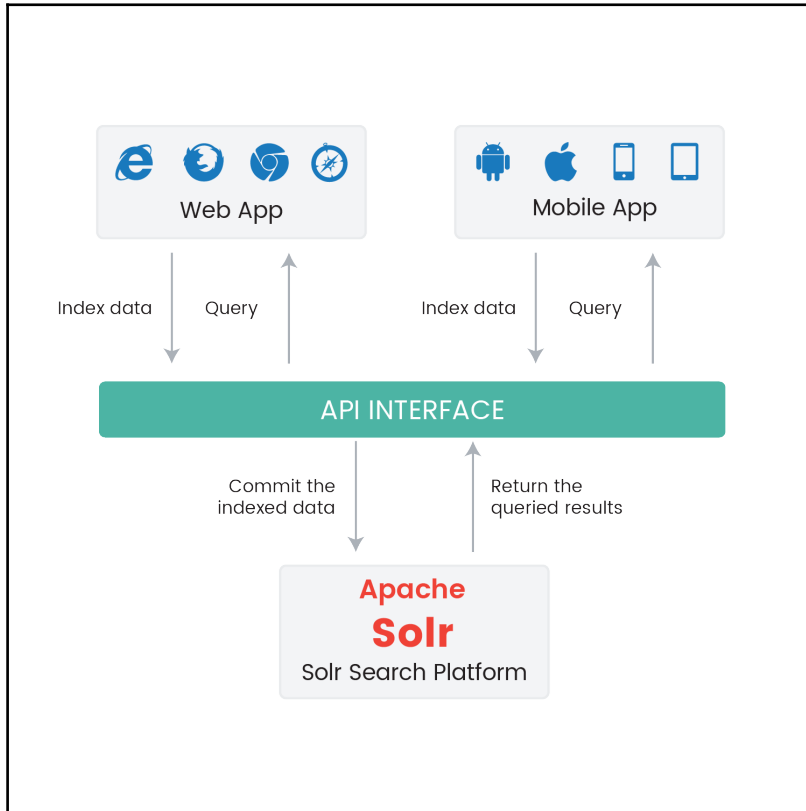
- CACHE.searcher.documentCache.cumulative\_ev
- CACHE.searcher.documentCache.cumulative\_hitra
- CACHE.searcher.documentCache.cumulative\_hits
- CACHE.searcher.documentCache.cumulative\_insert
- CACHE.searcher.documentCache.cumulative\_lookup
- CACHE.searcher.documentCache.evictions
- CACHE.searcher.documentCache.hitratio
- CACHE.searcher.documentCache.hits
- CACHE.searcher.documentCache.inserts
- CACHE.searcher.documentCache.lookups
- CACHE.searcher.documentCache.size
- CACHE.searcher.documentCache.warmupTime

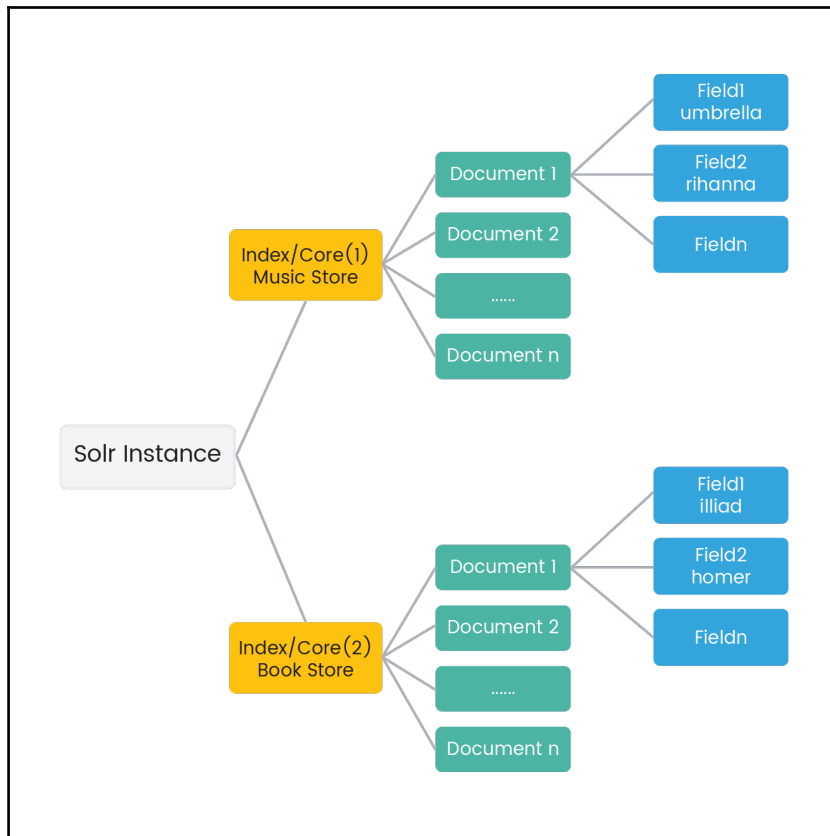
- fieldCache
- fieldValueCache
- filterCache
- perSegFilter
- queryResultCache

[Documentation](#)
[Issue Tracker](#)
[IRC Channel](#)
[Community forum](#)
[Solr Query Syntax](#)

---

## Chapter 3: Designing Schemas





# Chapter 4: Mastering Text Analysis Methodologies

The screenshot shows the Solr Admin interface. On the left is a navigation sidebar with options like Dashboard, Logging, Core Admin, Java Properties, Thread Dump, and a dropdown for 'techproducts'. The main area has two input fields: 'Field Value (Index)' containing 'Field value at index time' and 'Field Value (Query)' containing 'Field value at query time'. Below these is a dropdown for 'text\_en' and a 'Schema Browser' link. A 'Verbose Output' checkbox is checked. An 'Analyse Values' button is on the right. A diagram with arrows points from the 'text\_en' dropdown to a box labeled 'Configured field type' and from the 'techproducts' dropdown to a box labeled 'Selected Example'.

This screenshot shows the results of the analysis. The 'Field Value (Index)' field now contains 'The Nation of soccer' and the 'Field Value (Query)' field contains 'Famous country for football'. The 'text\_en' dropdown and 'Verbose Output' checkbox are still present. The 'Analyse Values' button is on the right. Two tables are displayed. The first table, labeled 'Final token stream for index', shows tokens: ST | The | Nation | of | soccer, SF | Nation | soccer, and LCE | nation | soccer. The 'LCE' row is highlighted with a red box. The second table, labeled 'Final token stream for query', shows tokens: ST | Famous | country | for | football, SF | Famous | country | football, SGF | Famous | nation | country | soccer | football, and LCE | famous | nation | country | soccer | football. The 'LCE' row is highlighted with a red box.



Field Value (Index)  
Running simple Solr analyzer through admin console

Field Value (Query)

Analyze Fieldname / FieldType: `text_en` Schema Browser

Verbose Output Analyse Values

WT Running simple Solr analyzer through admin console

**Splitted at white spaces and token stream generated by WhitespaceAnalyzer**

Documentation Issue Tracker IRC Channel Community forum Solr Query Syntax

Field Value (Index)  
The Host Country of Soccer World Cup 2018

Field Value (Query)  
The Host Nation of Football world cup 2018

Analyze Fieldname / FieldType: `text_en` Schema Browser

Verbose Output Analyse Values

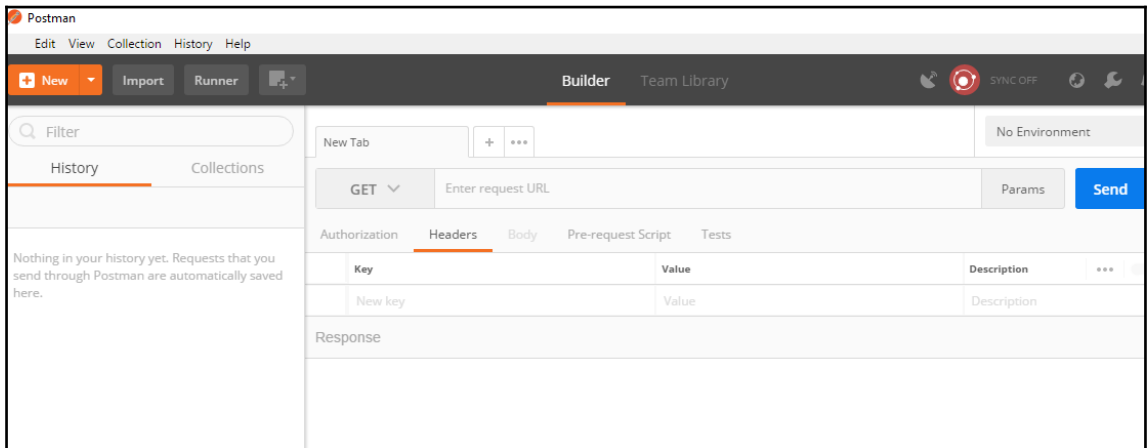

ST	The	Host	Country	of	Soccer	World	Cup	2018
SF		Host	Country		Soccer	World	Cup	2018
LCF		host	country		soccer	world	cup	2018

ST	The	Host	Nation	of	Football	world	cup	2018	
SF		Host	Nation		Football	world	cup	2018	
SGF		Host	country	Nation	soccer	Football	world	cup	2018
LCF		host	country	nation	soccer	football	world	cup	2018

Documentation Issue Tracker IRC Channel Community forum Solr Query Syntax

---

# Chapter 5: Data Indexing and Operations



```

C:\Windows\System32\cmd.exe

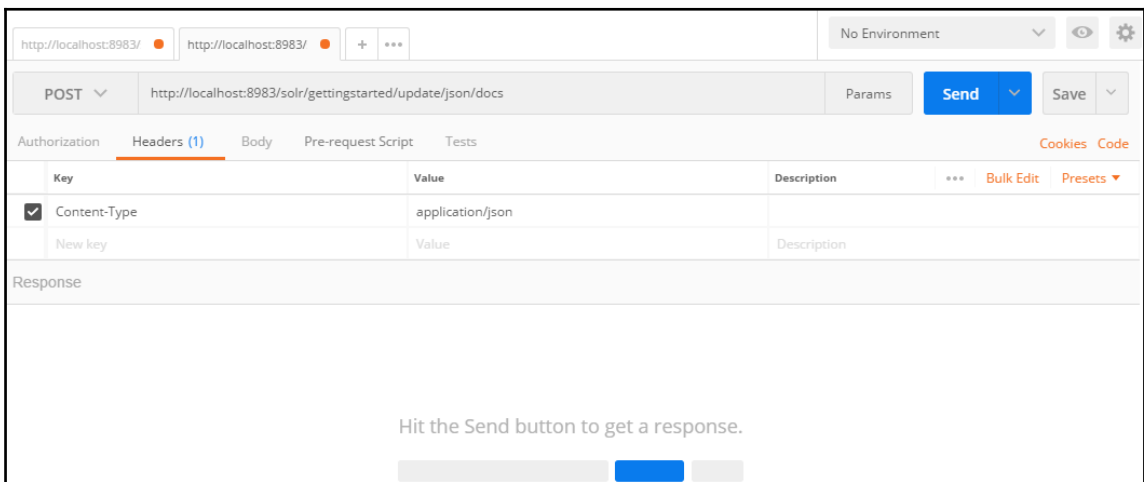
E:\book\solr\solr-7.2.0\example\exampledocs>java -jar post.jar -h
SimplePostTool version 5.0.0
Usage: java [SystemProperties] -jar post.jar [-h|-] [<file|folder|url|arg> [<file|folder|url|arg>...]]

Supported System Properties and their defaults:
-Dc=<core/collection>
-Durl=<base Solr update URL> (overrides -Dc option if specified)
-Ddata=files|web|args|stdin (default=files)
-Dtype=<content-type> (default=application/xml)
-Dhost=<host> (default: localhost)
-Dport=<port> (default: 8983)
-Dbasicauth=<user:pass> (sets Basic Authentication credentials)
-Dauto=yes|no (default=no)
-Drecursive=yes|no|<depth> (default=0)
-Ddelay=<seconds> (default=0 for files, 10 for web)
-Dfiletypes=<type>[,<type>,...] (default=xml,json,jsonl, csv,pdf,doc,docx,ppt,pptx,xls,xlsx,odt,odp,ods,ott,otp,ots,rtf,htm,html,txt,log)
-Dparams="<key>=<value>[&<key>=<value>...]" (values must be URL-encoded)
-Dcommit=yes|no (default=yes)
-Doptimize=yes|no (default=no)
-Dout=yes|no (default=no)

This is a simple command line tool for POSTing raw data to a Solr port.
NOTE: Specifying the url/core/collection name is mandatory.
Data can be read from files specified as commandline args,
URLs specified as args, as raw commandline arg strings or via STDIN.
Examples:
java -Dc=gettingstarted -jar post.jar *.xml
java -Ddata=args -Dc=gettingstarted -jar post.jar '<delete><id>42</id></delete>'
java -Ddata=stdin -Dc=gettingstarted -jar post.jar < hd.xml
java -Ddata=web -Dc=gettingstarted -jar post.jar http://example.com/
java -Dtype=text/csv -Dc=gettingstarted -jar post.jar *.csv
java -Dtype=application/json -Dc=gettingstarted -jar post.jar *.json
java -Durl=http://localhost:8983/solr/techproducts/update/extract -Dparams=literal.id=pdf1 -jar post.jar solr-word.pdf
java -Dauto -Dc=gettingstarted -jar post.jar *
java -Dauto -Dc=gettingstarted -Drecursive -jar post.jar afolder
java -Dauto -Dc=gettingstarted -Dfiletypes=ppt,html -jar post.jar afolder

The options controlled by System Properties include the Solr
URL to POST to, the Content-Type of the data, whether a commit
or optimize should be executed, and whether the response should
be written to STDOUT. If auto=yes the tool will try to set type
automatically from file name. When posting rich documents the
file name will be propagated as "resource.name" and also used
as "literal.id". You may override these or any other request parameter
through the -Dparams property. To do a commit only, use "-" as argument.
The web mode is a simple crawler following links within domain. Default delay=10s

```



http://localhost:8983/ ● http://localhost:8983/ ● + ... No Environment v 👁 ⚙

POST v http://localhost:8983/solr/gettingstarted/update/json/docs Params v Send v Save v

Authorization v **Headers (1)** v Body v Pre-request Script v Tests Cookies v Code v

	Key	Value	Description
<input checked="" type="checkbox"/>	Content-Type	application/json	
<input type="checkbox"/>	New key	Value	Description

Response

Hit the Send button to get a response.

http://localhost:8983/ No Environment

POST http://localhost:8983/solr/gettingstarted/update Params Send Save

Authorization Headers (1) Body Pre-request Script Tests Cookies

form-data x-www-form-urlencoded raw binary XML (text/xml)

```
1 <add>
2 <doc>
3 <field name="title">Harry Potter and Prisoner of Azkaban</field>
4 <field name="author">J.K. Rowling</field>
5 </doc>
6 </add>
```

Body Cookies Headers (2) Test Results Status: 200 OK Time: 21135 ms Size: 245 B

Pretty Raw Preview XML

```
1 <?xml version="1.0" encoding="UTF-8"?>
2 <response>
3 <lst name="responseHeader">
4 <int name="status">0</int>
5 <int name="QTime">21093</int>
6 </lst>
7 </response>
```

The image shows the Solr Admin UI interface. On the left is a navigation sidebar with the Solr logo at the top. The sidebar contains the following items: Dashboard, Logging, Cloud, Collections, Java Properties, Thread Dump, a dropdown menu currently showing 'gettingstarted', Overview, Analysis, Dataimport, Documents, Files, and Query (which is highlighted). The main content area on the right is titled 'Request-Handler (qt)' and contains several input fields: a text field with '/select', a section labeled 'common' with a 'q' field containing '\*:\*', an 'fq' field with a red minus and green plus icon, a 'sort' field, a 'start, rows' section with '0' and '10' fields, an 'fl' field, and a 'df' field. At the bottom of this section is a 'Raw Query Parameters' field containing 'key1=val1&key2=val2'.

The screenshot shows a web browser's developer tools interface. The top section displays the request details for a POST method to the URL `http://localhost:8983/solr/gettingstarted/update`. The request body is shown in XML format, containing a `<delete>` element with an `<id>` attribute and a `<query>` element.

```
1 <delete>
2   <id>0bb62d53-20d0-4df2-a385-c08b8658dc28</id>
3   <query>title:azkaban</query>
4 </delete>
```

The bottom section displays the response details, showing the response body in XML format. The response includes an XML declaration, a `<response>` element, and a `<lst name="responseHeader">` element containing `<int name="status">` and `<int name="QTime">` elements.

```
1 <?xml version="1.0" encoding="UTF-8"?>
2 <response>
3   <lst name="responseHeader">
4     <int name="status">0</int>
5     <int name="QTime">165</int>
6   </lst>
7 </response>
```

localhost:8983/solr/#/gettingstarted/query

Solr

Request-Handler (qt)  
/select

— common —

q  
\*:\*

fq

sort

start, rows  
0 10

fl

df

```

http://localhost:8983/solr/gettingstarted/select?q=*:*
{
  "responseHeader":{
    "zkConnected":true,
    "status":0,
    "QTime":16,
    "params":{
      "q":"*:*",
      "_:":1514434725886}},
  "response":{"numFound":1,"start":0,"maxScore":1.0,"docs":[
    {
      "title":["Harry Potter and Philosophers Stone"],
      "author":["J.K. Rowling"],
      "id":"dcbfd369-2a4b-4e9d-9713-1af373849438",
      "title_str":["Harry Potter and Philosophers Stone"],
      "author_str":["J.K. Rowling"],
      "_version_":1587892903981613056}]}
}

```

Dashboard  
Logging  
Cloud  
Collections  
Java Properties  
Thread Dump  
gettingstarted  
Overview  
Analysis  
Dataimport  
Documents

http://localhost:8983/ http://localhost:8983/ + ... No Environment

POST http://localhost:8983/solr/gettingstarted/update/json/docs Params Send Save

Authorization Headers (1) Body Pre-request Script Tests Cookies Code

Key	Value	Description
<input checked="" type="checkbox"/> Content-Type	application/json	
New key	Value	Description

Response

Hit the Send button to get a response.

http://localhost:8983/ http://localhost:8983/ No Environment

POST http://localhost:8983/solr/gettingstarted/update/json/docs Params Send Save

Authorization Headers (1) Body Pre-request Script Tests Cookies Code

form-data x-www-form-urlencoded raw binary JSON (application/json)

```

1 {
2   "title": "Jonathan Livingston Seagull",
3   "author": "Richard Bach"
4 }

```

Body Cookies Headers (2) Test Results Status: 200 OK Time: 40 ms Size: 134 B

Pretty Raw Preview Text

```

1 {
2   "responseHeader": {
3     "status": 0,
4     "QTime": 8
5 }

```

**Solr**

- Dashboard
- Logging
- Cloud
- Collections
- Java Properties
- Thread Dump
- gettingstarted
  - Overview
  - Analysis
  - Dataimport
  - Documents
  - Files
  - Query
  - Stream

Request-Handler (qt) /select

— common —

q \*:\*

fq

sort

start, rows 0 10

fl

df

Raw Query Parameters key1=val1&key2=val2

wt

http://localhost:8983/solr/gettingstarted/select?q=\*:\*

```

{
  "responseHeader": {
    "zkConnected": true,
    "status": 0,
    "QTime": 15,
    "params": {
      "q": "*:*",
      "_": "1514453950466"
    },
    "response": { "numFound": 2, "start": 0, "maxScore": 1.0, "docs": [
      {
        "title": ["Harry Potter and Philosophers Stone"],
        "author": ["J.K. Rowling"],
        "id": "dcbfd369-2a4b-4e9d-9713-1af373849438",
        "title_str": ["Harry Potter and Philosophers Stone"],
        "author_str": ["J.K. Rowling"],
        "_version_": 1587892903981613056
      },
      {
        "title": ["Jonathan Livingston Seagull"],
        "author": ["Richard Bach"],
        "id": "936df16b-f6fe-4147-874d-ad4525bf7e9f",
        "title_str": ["Jonathan Livingston Seagull"],
        "author_str": ["Richard Bach"],
        "_version_": 1588019816078245888
      }
    ]
  }
}

```



http://localhost:8983/ | http://localhost:8983/ | No Environment

POST | http://localhost:8983/solr/gettingstarted/update/ | Params | Send | Save

Authorization | Headers (1) | **Body** | Pre-request Script | Tests | Cookies | Code

form-data | x-www-form-urlencoded | **raw** | binary | JSON (application/json)

```
1 [
2   {
3     "title": "Red sails to Capri",
4     "author": "Ann Weil"
5   },
6   {
7     "title": "Five Point Someone",
8     "author": "Chetan Bhagat"
9   }
10 ]
```

Body | Cookies | Headers (2) | Test Results | Status: 200 OK | Time: 74 ms | Size: 135 B

Pretty | Raw | Preview | Text |

```
1 {
2   "responseHeader": {
3     "status": 0,
4     "QTime": 49}
5 }
```

localhost:8983/solr/#/gettingstarted/query

Java Properties

Thread Dump

gettingstarted

Overview

Analysis

Dataimport

Documents

Files

Query

Stream

Schema

Core Selector

sort

start, rows

0 10

fl

df

Raw Query Parameters

key1=val1&key2=val2

wt

-----

indent off

debugQuery

dismax

edismax

hl

facet

spatial

spellcheck

Execute Query

```
"response":{"numFound":4,"start":0,"maxScore":1.0,"docs":[{"title":["Red sails to Capri"],"author":["Ann Weil"],"id":["3813c196-0341-4708-8b05-f47225eacb12"],"title_str":["Red sails to Capri"],"author_str":["Ann Weil"],"_version_":1588021679301328896}, {"title":["Harry Potter and Philosophers Stone"],"author":["J.K. Rowling"],"id":["dcbfd369-2a4b-4e9d-9713-1af373849438"],"title_str":["Harry Potter and Philosophers Stone"],"author_str":["J.K. Rowling"],"_version_":1587892903981613056}, {"title":["Jonathan Livingston Seagull"],"author":["Richard Bach"],"id":["936df16b-f6fe-4147-874d-ad4525bf7e9f"],"title_str":["Jonathan Livingston Seagull"],"author_str":["Richard Bach"],"_version_":1588019816078245888}, {"title":["Five Point Someone"],"author":["Chetan Bhagat"],"id":["68fb69cd-ad31-441b-be02-eb3339afdd42"],"title_str":["Five Point Someone"],"author_str":["Chetan Bhagat"],"_version_":1588021679292940288}]}}
```

http://localhost:8983/ http://localhost:8983/ + ... No Environment

POST http://localhost:8983/solr/gettingstarted/update/ Params Send Save

```
1 {
2   "add": {
3     "doc": {
4       "title": "Liferay Beginner's Guide",
5       "author": "Samir Bhatt"
6     }
7   },
8   "delete": { "id": "3813c196-0341-4708-8b05-f47225eacb12" },
9   "delete": { "query": "author:rowling" }
10 }
```

Body Cookies Headers (2) Test Results Status: 200 OK Time: 79 ms Size: 135 B

Pretty Raw Preview Text

```
1 {
2   "responseHeader": {
3     "status": 0,
4     "QTime": 56 }
5 }
```

http://localhost:8983/ http://localhost:8983/ http://localhost:8983/ + ... No Environment

POST http://localhost:8983/solr/gettingstarted/update/extract?literal.id=doc1&commit=true Params Send Save

Authorization Headers **Body** Pre-request Script Tests Cookies Code

form-data  x-www-form-urlencoded  raw  binary

Choose Files solr-word.pdf

Body Cookies Headers (2) Test Results Status: 200 OK Time: 7465 ms Size: 143 B

Pretty Raw Preview JSON

```
1 {
2   "responseHeader": {
3     "status": 0,
4     "QTime": 7360
5   }
6 }
```

# Chapter 6: Advanced Queries – Part I

Type of Search:

Find:

Boost by Price

**Field Facets** 46 results found in 163 ms Page 1 of 5

**cat**

- [electronics](#) (12)
- [currency](#) (4)
- [memory](#) (3)
- [cat1\\_cat2](#) (2)
- [connector](#) (2)
- [graphics card](#) (2)
- [hard drive](#) (2)
- [search](#) (2)
- [software](#) (2)
- [camera](#) (1)
- [copier](#) (1)
- [electronics\\_and\\_c...](#) (1)
- [electronics\\_and\\_s...](#) (1)
- [multifunction pri...](#) (1)
- [music](#) (1)
- [printer](#) (1)
- [scanner](#) (1)
- [missing](#) (24)

**manu\_exact**

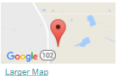
- [Apache Software F...](#) (2)
- [Belkin](#) (2)
- [Canon Inc.](#) (2)
- [Corsair Microsvst...](#) (2)

**Test with some GB18030 encoded characters** [More Like This](#)

Id: GB18030TEST  
Price: 0.0,USD  
Features: No accents here ... 这是一个功能 ... This is a feature (translated) ... 这份文件是很有光泽 ... This document is very shiny (translated)  
In Stock: true


**Samsung SpinPoint P120 SP2514N - hard drive - 250 GB - ATA-133** [More Like This](#)

Id: SP2514N  
Price: 92.0,USD  
Features: 7200RPM, 8MB cache, IDE Ultra ATA-133 ... NoiseGuard, SilentSeek technology, Fluid Dynamic Bearing (FDB) motor  
In Stock: true



**Maxtor DiamondMax 11 - hard drive - 500 GB - SATA-300** [More Like This](#)

Id: 6H500F0  
Price: 350.0,USD  
Features: SATA 3.0Gb/s, NCO ... 8.5ms seek ... 16MB cache  
In Stock: true



**Solr**

Request-Handler (qt) /select

common

q: **SP2514N**

fq: [ ]

sort: [ ]

start, rows: 0 10

fl: [ ]

df: [ ]

Raw Query Parameters: key1=val1&key2=val2

wt: **xml**

indent off

debugQuery

dismax

edismax

hl

facet

spatial

spellcheck

**Execute Query**

http://localhost:8983/solr/techproducts/select?q=SP2514N&wt=xml

```
<?xml version="1.0" encoding="UTF-8"?>
<response>
  <lst name="responseHeader">
    <int name="status">0</int>
    <int name="QTime">0</int>
    <lst name="params">
      <str name="q">SP2514N</str>
      <str name="wt">xml</str>
      <str name="_">1514797189592</str>
    </lst>
  </lst>
  <result name="response" numFound="1" start="0">
    <doc>
      <str name="id">SP2514N</str>
      <str name="name">Samsung SpinPoint P120 SP2514N - hard drive - 250 GB - ATA-133</str>
      <str name="manu">Samsung Electronics Co. Ltd.</str>
      <str name="manu_id_s">samsung</str>
      <arr name="cat">
        <str>electronics</str>
        <str>hard drive</str>
      </arr>
      <arr name="features">
        <str>7200RPM, 8MB cache, IDE Ultra ATA-133</str>
        <str>NoiseGuard, SilentSeek technology, Fluid Dynamic Bearing (FDB) motor</str>
      </arr>
      <float name="price">92.0</float>
      <str name="price_c">92.0,USD</str>
      <int name="popularity">6</int>
      <bool name="inStock">>true</bool>
      <date name="manufacturedate_dt">2006-02-13T15:26:37Z</date>
      <str name="store">35.0752,-97.032</str>
      <long name="_version_">1583131913641525248</long>
      <long name="price_c___l_ns">9200</long>
    </doc>
  </result>
</response>
```

The screenshot displays the Solr Admin interface. On the left, the navigation menu includes Dashboard, Logging, Core Admin, Java Properties, Thread Dump, and a selected 'techproducts' section. The 'Query' option is highlighted in the Plugins / Stats menu. The main form contains the following fields:

- Request-Handler (qt): /select
- common: [empty]
- q: SP2514N
- fq: [empty]
- sort: [empty]
- start, rows: 0, 10
- fl: [empty]
- df: [empty]
- Raw Query Parameters: key1=val1&key2=val2
- wt: json
- indent off
- debugQuery

Below the form, a list of facets is shown, with 'dismax' selected. Other facets include q.alt, qf, mm, pf, ps, qs, tie, bq, bf, edismax, hl, facet, spatial, and spellcheck. The 'Execute Query' button is at the bottom.

The browser address bar shows the URL: `http://localhost:8983/solr/techproducts/select?defType=dismax&q=SP2514N&wt=json`

```
{
  "responseHeader": {
    "status": 0,
    "QTime": 0,
    "params": {
      "q": "SP2514N",
      "defType": "dismax",
      "wt": "json",
      "_": "1515607090788"
    },
    "response": {
      "numFound": 1, "start": 0, "docs": [
        {
          "id": "SP2514N",
          "name": "Samsung SpinPoint P120 SP2514N - hard drive - 250 GB - ATA-133",
          "manu": "Samsung Electronics Co. Ltd.",
          "manu_id_s": "samsung",
          "cat": ["electronics", "hard drive"],
          "features": ["7200RPM, 8MB cache, IDE Ultra ATA-133", "NoiseGuard, SilentSeek technology, Fluid Dynamic Bearing (FDB) motor"],
          "price": 92.0,
          "price_c": "92.0,USD",
          "popularity": 6,
          "inStock": true,
          "manufacturedate_dt": "2006-02-13T15:26:37Z",
          "store": "35.0752,-97.032",
          "_version_": "1503131913641525248",
          "price_c__l_ns": "9200"
        }
      ]
    }
  }
}
```

The screenshot displays the Solr Admin interface. On the left, the navigation menu includes 'techproducts' and 'Query'. The main configuration area shows the following settings:

- Request-Handler (qt): /select
- q: SP2514N
- wt: json
- Execute Query: [button]

The right pane shows the browser address bar with the URL: `http://localhost:8983/solr/techproducts/select?q=SP2514N&wt=json`. Below the address bar, the JSON response is displayed:

```
{
  "responseHeader": {
    "status": 0,
    "QTime": 0,
    "params": {
      "q": "SP2514N",
      "wt": "json",
      "_: "1515346597997"}},
  "response": {
    "numFound": 1, "start": 0, "docs": [
      {
        "id": "SP2514N",
        "name": "Samsung SpinPoint P120 SP2514N - hard drive - 250 GB - ATA-133",
        "manu": "Samsung Electronics Co. Ltd.",
        "manu_id_s": "samsung",
        "cat": ["electronics", "hard drive"],
        "features": ["7200RPM, 8MB cache, IDE Ultra ATA-133", "NoiseGuard, SilentSeek technology, Fluid Dynamic Bearing (FDB) motor"],
        "price": 92.0,
        "price_c": "92.0,USD",
        "popularity": 6,
        "inStock": true,
        "manufacturedate_dt": "2006-02-13T15:26:37Z",
        "store": "35.0752,-97.032",
        "_version_": 1583131913641525248,
        "price_c___l_ns": 9200}
    ]
  }
}
```

The screenshot displays the Solr Admin interface. On the left is a navigation sidebar with the Solr logo and menu items: Dashboard, Logging, Core Admin, Java Properties, Thread Dump, techproducts (selected), Overview, Analysis, Dataimport, Documents, Files, Ping, Plugins / Stats, Query (highlighted), Replication, Schema, and Segments info. The main panel is titled 'Request-Handler (qt)' and shows a search form with the following fields: '/select', 'common', 'q' (containing 'SP2514N'), 'fq', 'sort', 'start, rows' (0, 10), 'fi', 'df', 'Raw Query Parameters' (key1=val1&key2=val2), and 'wt' (set to 'xml'). Below these fields are checkboxes for 'indent off', 'debugQuery', 'dismax', 'edismax', 'hl', 'facet', 'spatial', and 'spellcheck'. A blue 'Execute Query' button is at the bottom. The right pane shows the URL 'http://localhost:8983/solr/techproducts/select?q=SP2514N&wt=xml' and the resulting XML response, which includes product details for a Samsung SpinPoint P120 SP2514N hard drive.

```
Request-Handler (qt)
/select
common
q
SP2514N
fq
sort
start, rows
0 10
fi
df
Raw Query Parameters
key1=val1&key2=val2
wt
xml
 indent off
 debugQuery
 dismax
 edismax
 hl
 facet
 spatial
 spellcheck
Execute Query

http://localhost:8983/solr/techproducts/select?q=SP2514N&wt=xml
<?xml version="1.0" encoding="UTF-8"?>
<response>
<lst name="responseHeader">
  <int name="status">0</int>
  <int name="QTime">0</int>
  <lst name="params">
    <str name="q">SP2514N</str>
    <str name="wt">xml</str>
    <str name="_">1514797189592</str>
  </lst>
</lst>
<result name="response" numFound="1" start="0">
  <doc>
    <str name="id">SP2514N</str>
    <str name="name">Samsung SpinPoint P120 SP2514N - hard drive - 250 GB - ATA-133</str>
    <str name="manu">Samsung Electronics Co. Ltd.</str>
    <str name="manu_id_s">samsung</str>
    <arr name="cat">
      <str>electronics</str>
      <str>hard drive</str>
    </arr>
    <arr name="features">
      <str>7200RPM, 8MB cache, IDE Ultra ATA-133</str>
      <str>NoiseGuard, SilentSeek technology, Fluid Dynamic Bearing (FDB) motor</str>
    </arr>
    <float name="price">92.0</float>
    <str name="price_c">92.0,USD</str>
    <int name="popularity">6</int>
    <bool name="inStock">>true</bool>
    <date name="manufacturedate_dt">2006-02-13T15:26:37Z</date>
    <str name="store">35.0752,-97.032</str>
    <long name="_version_">1583131913641525248</long>
    <long name="price_c____l_ns">9200</long>
  </doc>
</result>
</response>
```



**Solr**

- Dashboard
- Logging
- Core Admin
- Java Properties
- Thread Dump
- techproducts**
- Overview
- Analysis
- Dataimport
- Documents
- Files
- Ping
- Plugins / Stats
- Query**
- Replication
- Schema
- Segments info

Request-Handler (qt)  
/select  
— common

q  
SP2514N

fq

sort

start, rows  
0 10

fl

df

Raw Query Parameters  
key1=val1&key2=val2

wt  
csv

indent off  
 debugQuery

dismax  
 edismax  
 hl  
 facet  
 spatial  
 spellcheck

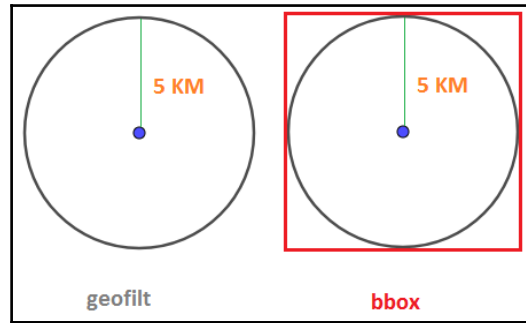
**Execute Query**

http://localhost:8983/solr/techproducts/select?q=SP2514N&wt=csv

```
payloads,manu_id_s,manu,address_s,weight,includes,incubationdate_dt,store,manufacturedate_dt,features,price_c,price,cat,popularity,
"",samsung,Samsung Electronics Co. Ltd.,,,,"35.0752",-97.032",2006-02-13T15:26:37Z,"7200RPM", 8MB cache\, IDE Ultra ATA-133,NoiseI
```

---

## Chapter 7: Advanced Queries – Part II



# Chapter 8: Managing and Fine-Tuning Solr

The screenshot shows the Solr Admin UI dashboard. The left sidebar contains navigation options: Dashboard, Logging (circled in red), Cloud, Collections, Java Properties, and Thread Dump. Below the sidebar are dropdown menus for 'Collection Sele...' and 'Core Selector'. The main content area is divided into three sections: 'Instance' (Start: about an hour ago), 'Versions' (listing solr-spec, solr-impl, and lucene-spec/impl), and 'JVM' (Runtime: Oracle Corporation Java HotSpot(TM) 64-Bit Server VM 1.8.0\_92 25.92-b14, Processors: 4, Args: -DSTOP.KEY=solrlocks). On the right, the 'System' section displays three bar charts: Physical Memory (30.4% used, 3.98 GB of 7.90 GB), Swap Space (65.4% used, 5.98 GB of 9.15 GB), and JVM-Memory (19.8% used, 97.07 MB).

The screenshot shows the Solr Admin UI logging configuration page. The left sidebar has 'Level' selected under the 'Logging' section. The main content area displays the Log4j configuration for 'org.slf4j.impl.Log4jLoggerFactory'. The configuration is a tree structure with the following levels and their respective log levels:

Path	Level
root	INFO
/solr	null
com	null
codahale	null
metrics	null
JmxReporter	null
org	null
apache	null
hadoop	WARN
http	null
client	null
protocol	null
RequestAddCookies	null
RequestAuthCache	null
RequestClientConnControl	null
ResponseProcessCookies	null
conn	null

**Logging**

- Level
- Cloud
- Collections
- Java Properties
- Thread Dump

Collection Sele... ▾

Core Selector ▾

metrics	<i>null</i>
JmxReporter	<i>null</i>
org	<i>null</i>
apache	<i>null</i>
<b>hadoop</b>	<b>WARN</b>
http	<i>null</i>
client	<i>null</i>
protocol	<i>null</i>
RequestAddCookies	<i>null</i>
RequestAuthCache	<i>null</i>
RequestClientConnControl	<i>null</i>
ResponseProcessCookies	<i>null</i>
conn	<i>null</i>
ssl	<i>null</i>
AllowAllHostnameVerifier	<i>null</i>
BrowserCompatHostnameVerifier	<i>null</i>
DefaultHostnameVerifier	<i>null</i>
SSLConnectionSocketFactory	<i>null</i>
StrictHostnameVerifier	<i>null</i>
headers	<i>null</i>
impl	<i>null</i>
auth	<i>null</i>

ALL x  
 TRACE  
 **DEBUG**  
 INFO  
 WARN  
 ERROR  
 FATAL  
 OFF  
 UNSET

```

C:\Windows\System32\cmd.exe

Ok, let's start up 2 Solr nodes for your example SolrCloud cluster.
Please enter the port for node1 [8983]:

Please enter the port for node2 [7574]:

Creating Solr home directory E:\book\solr\solr-7.2.0\example\cloud\node1\solr
Cloning E:\book\solr\solr-7.2.0\example\cloud\node1 into
E:\book\solr\solr-7.2.0\example\cloud\node2

Starting up Solr on port 8983 using command:
'E:\book\solr\solr-7.2.0\bin\solr.cmd' start -cloud -p 8983 -s "E:\book\solr\solr-7.2.0\example\cloud\node1\solr"

Waiting up to 30 to see Solr running on port 8983
Started Solr server on port 8983. Happy searching!

Starting up Solr on port 7574 using command:
'E:\book\solr\solr-7.2.0\bin\solr.cmd' start -cloud -p 7574 -s "E:\book\solr\solr-7.2.0\example\cloud\node2\solr" -z localhost:9983

Waiting up to 30 to see Solr running on port 7574
Started Solr server on port 7574. Happy searching!
INFO - 2018-01-28 20:49:55.896; org.apache.solr.client.solrj.impl.ZkClientClusterStateProvider; Cluster at localhost:9983 ready

Now let's create a new collection for indexing documents in your 2-node cluster.
Please provide a name for your new collection: [gettingstarted]

How many shards would you like to split gettingstarted into? [2]

How many replicas per shard would you like to create? [2]

Please choose a configuration for the gettingstarted collection, available options are:
default or sample_techproducts_configs [_default]

Created collection 'gettingstarted' with 2 shard(s), 2 replica(s) with config-set 'gettingstarted'

Enabling auto soft-commits with maxTime 3 secs using the Config API

POSTing request to Config API: http://localhost:8983/solr/gettingstarted/config
{"set-property":{"updateHandler.autoSoftCommit.maxTime":"3000"}}
Successfully set-property updateHandler.autoSoftCommit.maxTime to 3000

SolrCloud example running, please visit: http://localhost:8983/solr

```

The screenshot shows the Solr Admin UI for a collection named 'gettingstarted'. The left sidebar contains navigation links for Dashboard, Logging, Cloud, Collections, Java Properties, and Thread Dump. The main content area is divided into two sections: 'Collection: gettingstarted' and 'Shards'.

**Collection: gettingstarted**

- Config name: gettingstarted
- Max shards per node: -1
- Replication factor: 2
- Auto-add replicas:
- Router name: compositeId

**Shards**

- shard1**
  - Range: 80000000-ffffff
  - Active:
  - Replicas: gettingstarted\_shard1\_replica\_n1, gettingstarted\_shard1\_replica\_n2
- shard2**
  - Range: 0-7ffffff
  - Active:
  - Replicas: gettingstarted\_shard2\_replica\_n4, gettingstarted\_shard2\_replica\_n6

```

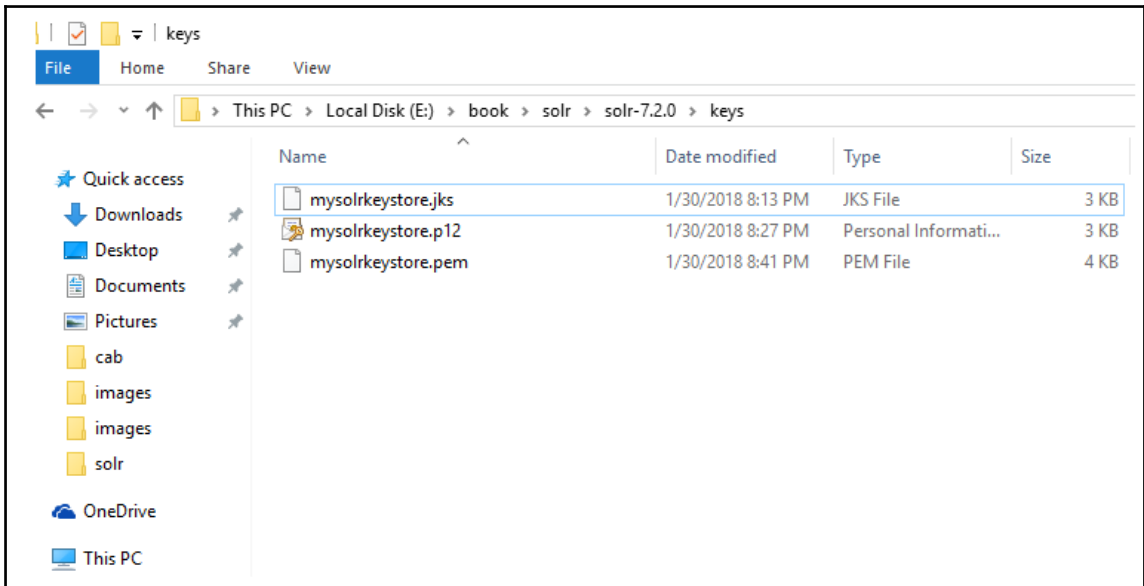
E:\book\solr\solr-7.2.0\keys>keytool -genkeypair -alias mysolr -keyalg RSA -keysize 2048 -keypass solrpass -storepass solrpass -validity 3650 -keystore mysolrkeystore.jks
What is your first and last name?
[Unknown]: mastering solr
What is the name of your organizational unit?
[Unknown]: knowarth
What is the name of your organization?
[Unknown]: knowarth
What is the name of your City or Locality?
[Unknown]: ahmedabad
What is the name of your State or Province?
[Unknown]: gujarat
What is the two-letter country code for this unit?
[Unknown]: in
Is Cmastering solr, OU=knowarth, O=knowarth, L=ahmedabad, ST=gujarat, C=in correct?
[no]: y

```

```

E:\book\solr\solr-7.2.0\keys>keytool -importkeystore -srckeystore mysolrkeystore.jks -destkeystore mysolrkeystore.p12 -srcstoretype jks -deststoretype pkcs12
Enter destination keystore password:
Re-enter new password:
Enter source keystore password:
Entry for alias mysolr successfully imported.
Import command completed: 1 entries successfully imported, 0 entries failed or cancelled

```



```

97 REM set SOLR_PORT=8983
98
99 REM Enables HTTPS. It is implicitly true if you set SOLR_SSL_KEY_STORE. Use this config
100 REM to enable https module with custom jetty configuration.
101 REM set SOLR_SSL_ENABLED=true
102 REM Uncomment to set SSL-related system properties
103 REM Be sure to update the paths to the correct keystore for your environment
104 REM set SOLR_SSL_KEY_STORE=etc/solr-ssl.keystore.jks
105 REM set SOLR_SSL_KEY_STORE_PASSWORD=secret
106 REM set SOLR_SSL_KEY_STORE_TYPE=JKS
107 REM set SOLR_SSL_TRUST_STORE=etc/solr-ssl.keystore.jks
108 REM set SOLR_SSL_TRUST_STORE_PASSWORD=secret
109 REM set SOLR_SSL_TRUST_STORE_TYPE=JKS
110 REM set SOLR_SSL_NEED_CLIENT_AUTH=false
111 REM set SOLR_SSL_WANT_CLIENT_AUTH=false
112
113 REM Uncomment if you want to override previously defined SSL values for HTTP client
114 REM otherwise keep them commented and the above values will automatically be set for HTTP clients
115 REM set SOLR_SSL_CLIENT_KEY_STORE=
116 REM set SOLR_SSL_CLIENT_KEY_STORE_PASSWORD=
117 REM set SOLR_SSL_CLIENT_KEY_STORE_TYPE=
118 REM set SOLR_SSL_CLIENT_TRUST_STORE=
119 REM set SOLR_SSL_CLIENT_TRUST_STORE_PASSWORD=
120 REM set SOLR_SSL_CLIENT_TRUST_STORE_TYPE=
121

```

The screenshot shows a web browser window with the address bar containing `https://localhost:8984/solr/#/`, which is circled in black. The page displays the Solr administration dashboard. On the left is a navigation menu with options: Dashboard, Logging, Core Admin, Java Properties, Thread Dump, and No cores available (Go and create one). The main content area is divided into three sections: Instance, Versions, and JVM. The Instance section shows the service started 6 minutes ago. The Versions section lists the following components and their details:

Component	Version	Details
solr-spec	7.2.0	
solr-impl	7.2.0	bca54cad5a9f6a80800944fd5bd585b68acde8c8 - jpountz - 2017-12-14 1
lucene-spec	7.2.0	
lucene-impl	7.2.0	bca54cad5a9f6a80800944fd5bd585b68acde8c8 - jpountz - 2017-12-14 1

The JVM section shows the runtime as Oracle Corporation Java HotSpot(TM) 64-Bit Server VM 1.8.0\_92 25.92-b14 and 4 processors.

The screenshot shows the Solr Admin interface in a browser window. The address bar indicates the URL is `https://localhost:8984/solr/#/`. The page title is "Solr". The left sidebar contains navigation links: Dashboard, Logging, Core Admin, Java Properties, and Thread Dump. The main content area shows the "Instance" status (Start 6 minutes ago) and a "Versions" table:

Component	Version	Checksum
solr-spec	7.2.0	
solr-impl	7.2.0	bca54cad5a9f6a80800944fd5bd585b68acde8c8 - jpo...
lucene-spec	7.2.0	
lucene-impl	7.2.0	bca54cad5a9f6a80800944fd5bd585b68acde8c8 - jpo...

Below the versions table, a "Security overview" section displays a warning: "This page is not secure (broken HTTPS)." and "Subject Alternative Name missing".

A "Certificate" dialog box is open, showing details for a certificate. The "Certification Path" tab is active, displaying the following fields:

Field	Value
Version	V3
Serial number	75 b2 65 e9
Signature algorithm	sha256RSA
Signature hash algorithm	sha256
Issuer	mastering solr, knowarth, kno...
Valid from	Tuesday, January 30, 2018 8:...
Valid to	Friday, January 28, 2028 8:13:...
Subject	mastering solr, knowarth, kno...

The dialog box also shows the certificate's DN: `CN = mastering solr, OU = knowarth, O = knowarth, L = ahmedabad, S = gujarat, C = in`. Buttons for "Edit Properties..." and "Copy to File..." are visible, along with an "OK" button at the bottom right.



---

## Chapter 9: Client APIs – An Overview

ipod

**Total Found: 3**

```
{ "responseHeader": { "status": 0, "QTime": 1, "params": { "q": "ipod", "fl": "id,name" } }, "response": { "numFound": 3, "start": 0, "docs": [ { "id": "IW-02", "name": "iPod & iPod Mini USB 2.0 Cable", { "id": "F8V7067-APL-KIT", "name": "Belkin Mobile Power Cord for iPod w/ Dock", { "id": "MA147LL/A", "name": "Apple 60 GB iPod with Video Playback Black" } } ], "spellcheck": { "suggestions": [], "correctlySpelled": false, "collations": [] } }
```