



IBM Software Group

# WebSphere Commerce - Case Study

## Debugging Common SOLR Support Questions

Brendon Baila ([brbaila@ca.ibm.com](mailto:brbaila@ca.ibm.com))  
Conway Chang ([changco@us.ibm.com](mailto:changco@us.ibm.com))  
Eric Rybczynski ([ericryb1@us.ibm.com](mailto:ericryb1@us.ibm.com))  
WebSphere Commerce L2 Support  
11 April 2011



WebSphere® Support Technical Exchange



# Agenda

- Architectural Overview
- Issue Scenarios
  - ▶ Architecture
  - ▶ Building Index
  - ▶ Queries
- Quick References
  - ▶ Location of logs and files
  - ▶ General technotes and APARs

# Architectural Overview

- An external Web server is recommended for use with a Solr application server, but is NOT required.
  - Connections may be made directly to the application server HTTP transport port.
  - Web server provides connection handling benefits during heavy load.
  - If an external Web server is not used, the 'SearchServerPort' value must be updated in SRCHCONF.CONFIG and SRCHCONFEXT.CONFIG for WCS indexing tools to work.
  - Recommended for production environments

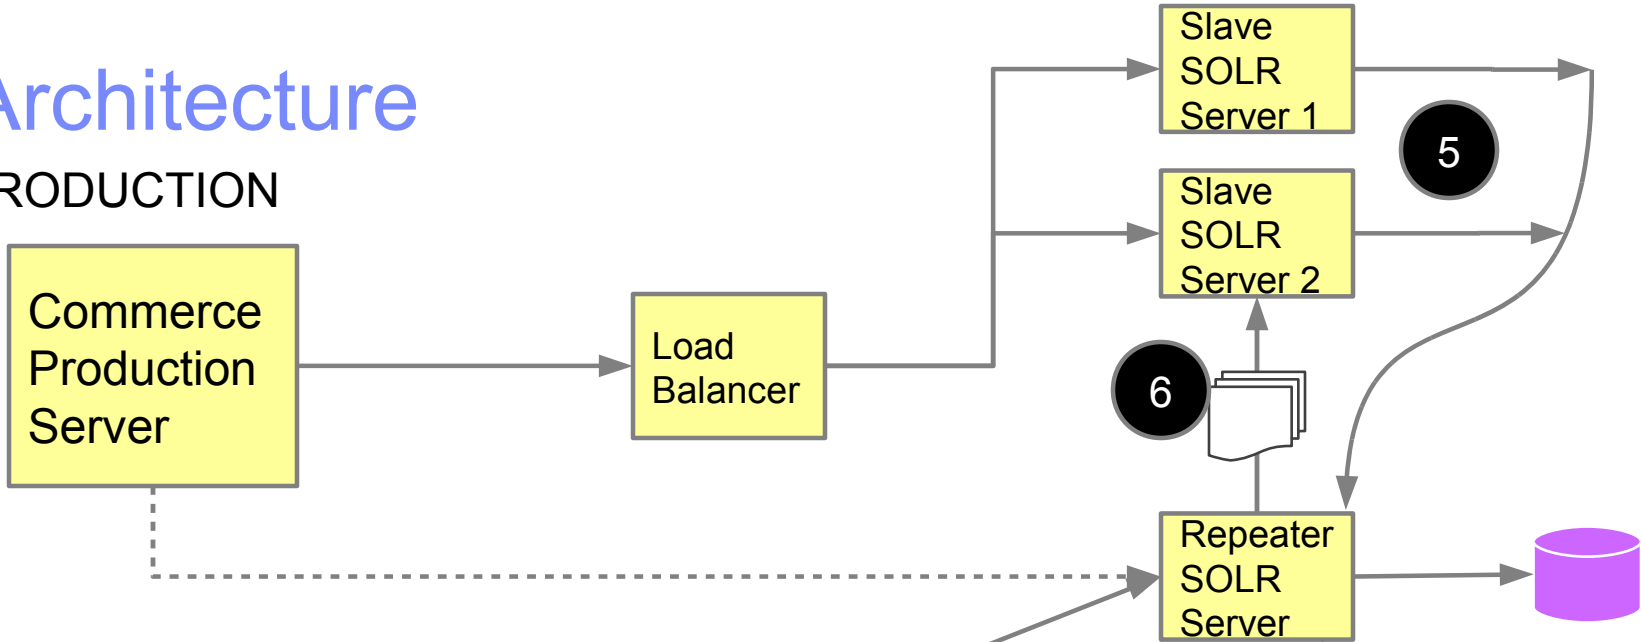
# Architectural Overview

- When clustering Solr search, the replication Master should be outside of the cluster to prevent slaves from attempting to replicate with themselves.
  - Caused by the Web server Plug-in routing application servers to all members of the cluster.
  - If the Master must be in the cluster, the slaves should be configured to replicate directly with the Master replica's HTTP transport port to bypass the Plug-in.




# Architecture

## PRODUCTION



## STAGING



# Architecture con't

- 1) di-buildindex.sh
- 2) indexprop.sh
- 3) fetchindex
- 4) replicate
- 5) polling interval: 60 sec
- 6) replicate

	Location of Files/Commands
1	<b>Location:</b> CommerceServer70/bin/buildindex.sh
	<b>Run:</b> buildindex.sh WC_installdir/instances/<instance_name>/search/pre-processConfig/<MC_masterCatalogId>/<target_db>
2	<b>Run:</b> indexprop.sh -searchServerName <i>searchServerName</i> -searchServerPort <i>searchServerPort</i> -solrHome <i>solrHome</i> -log <i>log</i>
	<b>Location:</b> CommerceServer70/bin/indexprop.sh
3 4 5 6	<b>SolrConfig.xml:</b> WC_installdir/instances/ <i>instance_name</i> /search/solr/home/solrConfig.xml



# Issue Scenarios – Architecture

## SOLR Replication – Misconfiguration of the environment

### Symptoms

- A slave replica is able to detect that the master has an updated index, but the index is not updated.
- You confirm that replication is configured correctly in the solrconfig.xml files.
- You see messages like these in the slave Solr appserver SystemOut.log:

```
[12/8/11 16:25:00:140 EST] 00000038 SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex Number of
files in latest index in master: 9
[12/8/11 16:25:00:153 EST] 0000002a SolrCore I org.apache.solr.core.SolrCore execute
[MC_10001_CatalogGroup_en_US] webapp=/solr path=/replication
params={wt=filestream&file=_2.tis&checksum=true&indexversion=1321997905518&command=filecontent} status=0
QTime=2
[12/8/11 16:25:00:156 EST] 00000038 SnapPuller W org.apache.solr.handler.SnapPuller$FileFetcher fetchPackets No
content recieved for file: {size=27200, name=_2.tis, lastmodified=1323278104000}
[12/8/11 16:25:00:159 EST] 00000038 ReplicationHa E org.apache.solr.handler.ReplicationHandler doFetch SnapPull
failed org.apache.solr.common.SolrException: Unable to download _2.ti s completely. Downloaded 0!=27200
at org.apache.solr.handler.SnapPuller$FileFetcher.cleanup(SnapPuller.java:1026)
at org.apache.solr.handler.SnapPuller$FileFetcher.fetchFile(SnapPuller.java :898)
at org.apache.solr.handler.SnapPuller.downloadIndexFiles(SnapPuller.java:541)
at org.apache.solr.handler.SnapPuller.fetchLatestIndex(SnapPuller.java:294)
at org.apache.solr.handler.ReplicationHandler.doFetch(ReplicationHandler.java:264)
at org.apache.solr.handler.SnapPuller$1.run(SnapPuller.java:159)
at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:452)
at java.util.concurrent.FutureTask$Sync.innerRunAndReset(FutureTask.java:328)
...
```

## Issue Scenarios – Architecture

### SOLR Replication – Misconfiguration of the environment

#### Symptoms

- Seen in slave Solr's SystemOut.log:

```
[12/8/11 16:26:00:034 EST] 00000038 SnapPuller E org.apache.solr.handler.SnapPuller  
fetchFileList No files to download for indexversion: 1321997905518
```

```
[12/8/11 16:27:00:017 EST] 00000038 SnapPuller I org.apache.solr.handler.SnapPuller  
fetchLatestIndex Master's version: 1321997905518, generation: 4
```

```
[12/8/11 16:27:00:017 EST] 00000038 SnapPuller I org.apache.solr.handler.SnapPuller  
fetchLatestIndex Slave's version: 1323379476746, generation: 1
```



## Issue Scenarios – Architecture

### SOLR Replication – Misconfiguration of the environment

#### **Cause #1**

- Master and slaves are contained in same cluster
- Web server Plug-in routes replication requests to both master and slave Solr appservers.
- If both index check and download requests are routed to the master appserver, replication completes successfully.
- If the download request is routed to a slave appserver, the slave will not be able to download the updated master index.

#### **Solution #1**

- Configure the slave replicas to use the master's appserver HTTP transport port instead of the Web server's port 3737.
  - Example: `<str name="masterUrl">http://<master>:9081/solr/replication</str>`
- Change the 'SearchServerPort' in SRCHCONF.CONFIG and SRCHCONFEXT.CONFIG so the index tooling always uses the master.

# Issue Scenarios – Architecture

## SOLR Replication – Misconfiguration of the environment

### Cause #2

- The slave replica has both “master” and “slave” configurations in solrconfig.xml.
  - Example:

```
<requestHandler name="/replication" class="solr.ReplicationHandler">  
  <lst name="master">  
    <str name="replicateAfter">commit</str>  
    <str name="replicateAfter">startup</str>  
    <str name="confFiles">schema.xml,stopwords.txt</str>  
  </lst>  
  <lst name="slave">  
    <str name="masterUrl">http://<masterHostname>:9080/solr/replication</str>  
    <str name="pollInterval">00:00:60</str>  
  </lst>  
</requestHandler>
```

### Solution #2

- Comment out or remove the “master” section from the requestHandler configuration.
  - Example:

```
<requestHandler name="/replication" class="solr.ReplicationHandler">  
  <lst name="slave">  
    <str name="masterUrl">http://<masterHostname>:9080/solr/replication</str>  
    <str name="pollInterval">00:00:60</str>  
  </lst>  
</requestHandler>
```

# Issue Scenarios – Architecture

## SOLR Replication – Multiple Master Definitions

### Symptoms

- The slave Solr replica is able to detect that its index version and generation are different than the master's, but is unable to download the new index.
- The master and slave index version and generation can also change unexpectedly.
- You may see these messages logged in the slave's SystemOut.log:

```
[1/24/12 13:53:00:029 UTC] 00000045 SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex
  Master's version: 1325893509760, generation: 22
[1/24/12 13:53:00:031 UTC] 00000045 SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex
  Slave's version: 1325893509757, generation: 19
...
[1/24/12 14:01:00:371 UTC] 0000003c SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex
  Master's version: 1325893509760, generation: 22
[1/24/12 14:01:00:373 UTC] 0000003c SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex
  Slave's version: 1327413605266, generation: 1
...
[1/24/12 14:01:00:409 UTC] 0000003b SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex
  Number of files in latest index in master: 9
...
[1/24/12 14:01:02:507 UTC] 0000003b SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex Total
  time taken for download : 2 secs
...
[1/24/12 14:59:00:253 UTC] 00000045 SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex
  Master's version: 1325893509760, generation: 22
[1/24/12 14:59:00:256 UTC] 00000045 SnapPuller I org.apache.solr.handler.SnapPuller fetchLatestIndex
  Slave's version: 1325893509761, generation: 23
```

## Issue Scenarios – Architecture

### SOLR Replication – Multiple Master Definitions

#### Cause

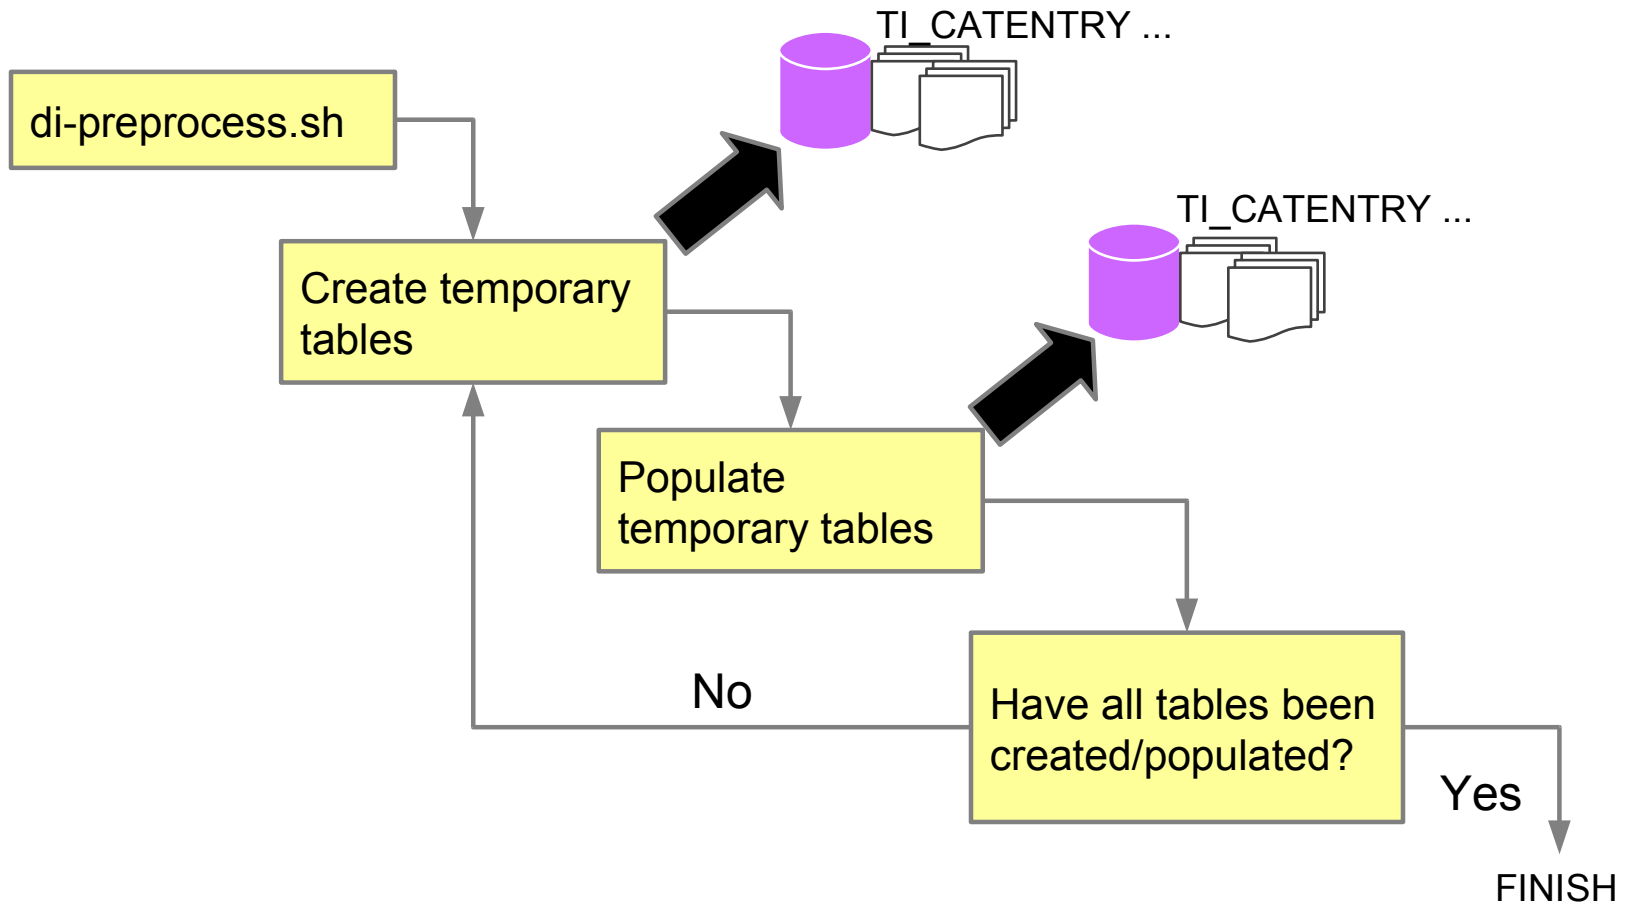
- The master ReplicationHandler configuration is specified in solrconfig.xml more than once.

#### Solution

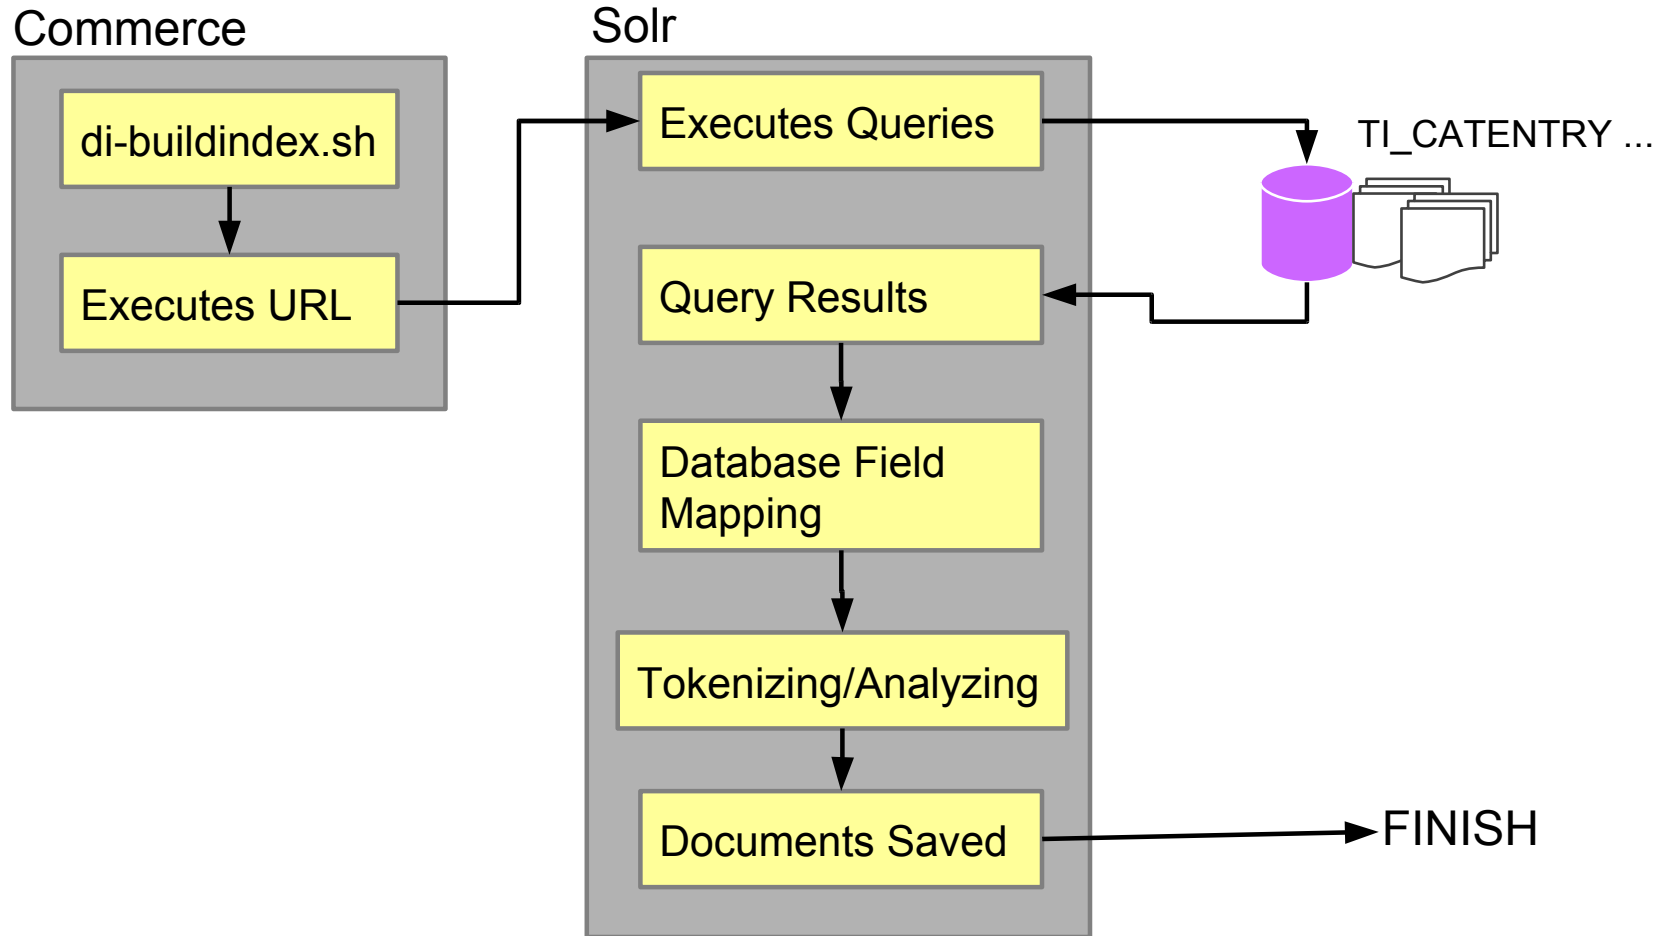
- Remove or comment out the extra master ReplicationHandler configuration(s).
  - Several sample replication configurations in solrconfig.xml are commented out
  - Make sure only one master configuration is active.
  - In FEP3 or higher, there are two solrconfig.xml's, so both must be configured correctly:
    - *solrhome/MC\_masterCatalogId/en\_US/CatalogEntry/conf/solrconfig.xml*
    - *solrhome/MC\_masterCatalogId/en\_US/CatalogEntry/unstructured/conf/solrconfig.xml*
    - *solrhome/MC\_masterCatalogId/en\_US/CatalogGroup/conf/solrconfig.xml*

```
<requestHandler name="/replication" class="solr.ReplicationHandler">  
  <lst name="master">  
    <str name="replicateAfter">commit</str>  
    <str name="replicateAfter">startup</str>  
    <str name="confFiles">schema.xml,stopwords.txt</str>  
  </requestHandler>
```

# Preprocess



# Build Index



# Issue Scenarios – Building the Index

## Building the Index – Gathering Trace Data

- Modify
  - Server:  
`<WCS>/instances/<instanceName>/xml/config/dataimport/logging.properties`
  - Developer toolkit:  
`<WCDE>\workspace\WC\xml\config\dataimport\logging.properties`
- Change occurrences of INFO to FINEST
- Modify the line: `java.util.logging.FileHandler.count=1`
  - Increase the value from 1 to a larger number, such as 10
  - FINEST trace will log more data, more log files are needed to retain details
- Log file size can be changed by modifying the following:
  - `java.util.logging.FileHandler.limit=25000000`



# Issue Scenarios – Building the Index

## Pre-processing – Mapping to multiple categories

### Symptom

- INFO: Error for batch element #1: DB2 SQL Error: **SQLCODE=-803, SQLSTATE=23505, SQLERRMC=1;DB2INST1.TI\_APGROUP\_0, DRIVER=4.12.55**

### Cause

- Mappings for a single catalog entry to multiple categories
- Mappings for a single category mapped to multiple parent categories



# Issue Scenarios – Building the Index

## Pre-processing – Mapping to multiple categories

### **Solution #1**

- Mappings for a single catalog entry to multiple categories
  - Run SQL:  
select catentry\_id, catgroup\_id from catgpenrel where catentry\_id in (select catentry\_id from catgpenrel where catalog\_id = <catalogId> group by catentry\_id having count(catentry\_id) > 1)
  - If results are returned
    - A catalog entry can only be mapped to one category
    - Remove the extra mappings (product → category mapping)



# Issue Scenarios – Building the Index

## Pre-processing – Mapping to multiple categories

### Solution #2

- Mappings for a single category mapped to multiple parent categories
  - Enable FINEST trace, and observe:

```
0 Aug 11, 2011 6:53:25 PM com.ibm.commerce.foundation.dataimport.preprocess.CatalogHierarchyDataPreProcessor logCacheContents FINEST:
```

```
-----  
Cache contents:
```

```
-----  
25076=10001_12021;10001_10001,10001_10503  
-----
```

```
...
```

```
0 Aug 11, 2011 7:05:33 PM com.ibm.commerce.foundation.dataimport.preprocess.CatalogHierarchyDataPreProcessor logCacheContents FINEST:
```

```
-----  
Cache contents:
```

```
-----  
25076=10001_12021;10001_10001,10001_10504  
-----
```

- A category can only have one parent category
- Remove the extra mappings (child category → parent category)

# Issue Scenarios – Building the Index

## Pre-processing – Out Of Memory

### Symptom

- di-preprocess runs indefinitely
- di-preprocess step fails repeatedly with java/lang/OutOfMemoryError
- Javacore data shows repeating stack:  
4XESTACKTRACE           at  
com/ibm/commerce/foundation/dataimport/preprocess/CatalogHierarchyDataPreProcessor.cacheChildren(Catalog  
HierarchyDataPreProcessor.java:434(Compiled Code))  
4XESTACKTRACE           at  
com/ibm/commerce/foundation/dataimport/preprocess/CatalogHierarchyDataPreProcessor.cacheChildren(Catalog  
HierarchyDataPreProcessor.java:434(Compiled Code))  
4XESTACKTRACE           at  
com/ibm/commerce/foundation/dataimport/preprocess/CatalogHierarchyDataPreProcessor.cacheChildren(Catalog  
HierarchyDataPreProcessor.java:434(Compiled Code))  
4XESTACKTRACE           at  
com/ibm/commerce/foundation/dataimport/preprocess/CatalogHierarchyDataPreProcessor.cacheChildren(Catalog  
HierarchyDataPreProcessor.java:434(Compiled Code))

# Issue Scenarios – Building the Index

## Pre-processing – Out Of Memory

### **Cause**

- Cyclical reference(s) in category mapping  
Pots → Pots & Pans → Cookware → Pots

### **Solution**

- Remove cyclical reference (Cookware → Pots)

# Issue Scenarios – Building the Index

## Pre-processing – Column size too small

### Symptom

- INFO: [jcc][t4][20111][11366][4.9.78] The value of a host variable is too large for its corresponding use. Host variable=2. **ERRORCODE=-4461, SQLSTATE=42815**
- Change trace to FINEST to identify which table encounters the import issue
- Note the line before the error to determine the data attempted to be imported:

```
01/01/2012 2:23:45 PM com.ibm.commerce.foundation.dataimport.preprocess.AbstractDataPreProcessor
insertColumns(PreparedStatement, Object, Map, List) FINER: ENTRY 1,234,567
{PRODUCTSET=100551;100552;100008;100009;100029;100030;100031;100032;100033;100034;100035;100036;
100037;100038;100039;100040;100041;100042;100043;100044;100045;100046;100047;100048;100049;100050;
100051;100052;100053;100054;100055;100056;100057;100058;100059;100060;100061;100062;100063;100064;
100065;100066;100067;100068;100069;100070;100071;100072;100073;100074;100075;100076;100077;100078;
100079;100080;100081;100082;100083;100084;100085;100086;100087;100088;100089;100090;100091;100092;
100093;100094;100095;100096;100097;100098;510108;510109;511013;511014;512517;512518}
```

### Cause

- A value being populated into a column is too large to fit into the limited column size

# Issue Scenarios – Building the Index

## Pre-processing – Column size too small

### **Solution**

- Modify the schema definition:
- `<WC_Install_Dir>/instances/<instance>/search/pre-processConfig/<store_id>/<db_type>/wc-dataimport-preprocess-*.xml`
- `<WC_Install_Dir>/components/foundation/samples/dataimport/catalog/<db_type>/wc-dataimport-preprocess-*.xml`
- Increase VARCHAR size, or use a CLOB

# Issue Scenarios – Building the Index

## Pre-processing – Column size too small

### Solution

- If changing to a CLOB, modify:
  - wc-data-config.xml:
    - `<field column="DATABASECOLUMN" clob="true"/>`  
`<field column="INDEXFIELD" splitBy=";"`  
`sourceColName="DATABASECOLUMN"/>`
    - For example:  
`<field column="CAS_F1ATTR" clob="true"/>`  
`<field column="cas_f1" splitBy=";" sourceColName="CAS_F1ATTR"/>`

# Issue Scenarios – Building the Index

## Building the Index – Unstructured Index

### **Symptom**

- di-buildindex fails – examination of SOLR trace shows CatalogEntry and CatalogGroup index building completes
- Only unstructured index building fails to complete
- No error message in the logs
- Last message for unstructured build in the trace:

```
JdbcDataSourc 1 org.apache.solr.handler.dataimport.JdbcDataSource$ResultSetIterator  
<init> Executing SQL: SELECT I_CATENTATCHT_1.CATENTRY_ID,  
I_CATENTATCHT_1.ATCHAST_ID, ...JOIN CATENTRY ON  
(CATENTRY.CATENTRY_ID=I_CATENTATCHT_1.CATENTRY_ID)
```

```
JdbcDataSourc 3 org.apache.solr.handler.dataimport.JdbcDataSource$ResultSetIterator  
<init> Time taken for sql :8
```



# Issue Scenarios – Building the Index

## Building the Index – Unstructured Index

### Cause

- Unstructured Index indexes catalog entry attachments such as PDF or .txt files
  - Uses TikaEntityProcessor
- Javacore of index building thread:

```
at java/net/PlainSocketImpl.socketConnect(Native Method)
at java/net/PlainSocketImpl.doConnect(PlainSocketImpl.java:383)
...
at sun/net/www/protocol/http/URLConnection.connect(URLConnection.java:846)
at sun/net/www/protocol/http/URLConnection.getInputStream(URLConnection.java:1182)
at sun/net/www/protocol/http/URLConnection.getHeaderField(URLConnection.java:2314)
at java/net/URLConnection.getContentType(URLConnection.java:509)
at org/apache/solr/common/util/ContentStreamBase$URLStream.<init>(ContentStreamBase.java:80)
at com/ibm/commerce/solr/handler/TikaEntityProcessor.initConnection(TikaEntityProcessor.java:176)
at com/ibm/commerce/solr/handler/TikaEntityProcessor.nextRow(TikaEntityProcessor.java:238)
at org/apache/solr/handler/dataimport/EntityProcessorWrapper.nextRow(EntityProcessorWrapper.java:237)
at org/apache/solr/handler/dataimport/DocBuilder.buildDocument(DocBuilder.java:357)
at org/apache/solr/handler/dataimport/DocBuilder.buildDocument(DocBuilder.java:383)
```

### Solution

- Caused by a firewall between SOLR and content servers
- Opened ports between SOLR and Content server, index build finished

# Issue Scenarios – Building the Index

## Building the Index – Droid Crawler

### Symptom

- crawler exception in crawlerlog.txt :  
Nov 3, 2011 3:45:37 PM  
com.ibm.commerce.search.crawler.task.SingleThreadTaskMaster processAllTasks\  
WARNING: org.apache.droids.protocol.http.ContentTooLongException\  
Nov 3, 2011 3:45:37 PM org.apache.droids.robot.crawler.CrawlingDroid finished\  
INFO: FINISHED!!!

### Cause

- Hardcoded limit to how much it can crawl – 512KB

### Solution

- iFix available to make configurable limit (APAR IV10446)
- When patching the Solr application, the iFix will be delivered as a zip file to be deployed as a partial application update.



# Issue Scenarios – Queries

## Symptom

- Issues seen once the index is built/replicated:
  - Unexpected Search Results
  - Blank categories
  - Missing data (ex: prices)
  - Generic Error page
- Error seen on the Commerce Server is often:
  - *CMN0409E: The following error occurred during processing:  
"com.ibm.commerce.catalog.facade.client.CatalogNavigationViewException".*
  - Generic error, simply reporting a problem with the service call
- A good place to start with search errors is the SystemOut.log/trace.log on the SOLR server
- Need further tracing enabled on the Commerce side to troubleshoot issue

# Issue Scenarios – Queries: Maximum Boolean Clause

## Symptoms

- Error on the Commerce Server:  
com.ibm.commerce.catalog.facade.client.CatalogNavigationViewException

- Error on the SOLR server:

*SolrCore E org.apache.solr.common.SolrException log org.apache.solr.common.SolrException:  
org.apache.lucene.queryParser.ParseException:*

*Cannot parse 'catentry\_id:(1278505 "102864" .... "1563276)": **too many boolean clauses***

## Solution

- Increase the Max boolean clause in the solrconfig.xml:
  - Default FEP3:<maxBooleanClauses>3072</maxBooleanClauses>
  - Each index has it's own solrconfig.xml, maxBooleanClauses tag
- Usually caused by a large number of SKUs per product, during product entitlement
- APARs JR41115 and JR41116 improve product entitlement performance/queries

# Issue Scenarios – Queries

## Encoding metadata of facets

### Symptom

- Error page on Commerce
- CatalogNavigationViewException in Commerce logs
- A parse error is seen in the SOLR logs, with strange looking characters:

Error 400: org.apache.lucene.queryParser.**ParseException**: Cannot parse '<MT@SP> ads\_f5\_ntk\_cs: ("Reflectivity")<MT@SP~YÃ—Ã™ÂWÃ>ÂÃ—Ã~ÃœÃŽÂ â€¢Ãš[â„¢  
\\â€°Ã>Ã™!': **Encountered "<EOF>"** at line 1, column 91



# Issue Scenarios – Queries

## Encoding metadata of facets

### Cause

- When base64 encoding metaData, characters such as '+' can be in the encoded string
- Before encoding:  
`<MT@SP>+ads_f5_ntk_cs:("Reflectivity")<MT@SP>+ads_f5_ntk_cs:("Wind\proof")`
- After encoding:  
`PE1UQFNQPithZHNfZjVfbnRrX2NzOigiUmVmbGVjdGI2aXR5lik8TVRAU1A+K2Fkc19mNV9udGtfY3M6KCJXaW5kXC1wcm9vZilp`
- The link on the page would look like:

<http://localhost/webapp/wcs/stores/servlet/SearchDisplay?searchTermScope=&searchType=1000&filterTerm=&maxPrice=&showResultsPage=true&langId=-1&beginIndex=0&sType=SimpleSearch&metaData=PE1UQFNQPithZHNfZjVfbnRrX2NzOigiUmVmbGVjdGI2aXR5lik8TVRAU1A+K2Fkc19mNV9udGtfY3M6KCJXaW5kXC1wcm9vZilp&pageSize=12&manufacturer=&resultCatEntryType=2&catalogId=10001&pageView=image&searchTerm=Shoe&minPrice=&urlLangId=-1&categoryId=10056&storeId=10001>

- Commerce runtime interprets the '+' as a space
- Metadata loses the '+':  
`PE1UQFNQPithZHNfZjVfbnRrX2NzOigiUmVmbGVjdGI2aXR5lik8TVRAU1A K2Fkc19mNV9udGtfY3M6KCJXaW5kXC1wcm9vZilp`
- When decrypted, becomes corrupt with strange characters, which is then sent to SOLR server

# Issue Scenarios – Queries

## Encoding metadata of facets

### Solution

- Feature Pack 4 URL encodes the base64 encoded data
- Available as an APAR for FEP2 and FEP3: JR41857

### Workaround

- Manually URL encode in the JSP:

```
<%  
    String metaData = (String) request.getAttribute("metaData");  
    if (metaData != null) {  
        metaData = java.net.URLEncoder.encode(metaData,"UTF-8");  
        request.setAttribute("metaData", metaData);  
    }  
%>
```

- Best place for this is in the SearchSetup.jspf

# Issue Scenarios – Queries

## Unexpected Search Results

### Symptom

- Search for ‘torch welder’, results return documents that do not seem to contain the word ‘welder’

Query:

```
com.ibm.commerce.foundation.internal.server.services.search.processor.solr.SolrSearchExpressionProcessor
getEntityObjects Final Solr query expression: q (+(name:(torch welder) defaultSearch:(torch welder)
shortDescription:(torch welder)))&fq=catalog_id:10051&fq=storeent_id:(10152 10051)&fq=published:1
```

Result Set:

```
SolrDocument[{shortDescription=Gas Cyl 40, catalog_id=[10001, 10051], storeent_id=10051, parentCatgroup_id_facet=[10001_10001,
10051_10520],parentCatgroup_id_search=[ 10051_10507], catenttype_id_ntk_cs=ItemBean, name=Gas Cylinder
cas_f1_ntk_cs=Gas Cylinder, mfName=MyCompany, mfName_ntk_cs=MyCompany, mfName_ntk=MyCompany, buyable=1, published=1, catentry_id=71234
partNumber_ntk=Cylinder1}],
```

```
SolrDocument[{shortDescription=Gas S, catalog_id=[10001, 10051], storeent_id=10051, parentCatgroup_id_facet=[10001_10001, 10051_10520],
parentCatgroup_id_search=[10051_10507], catenttype_id_ntk_cs=ItemBean, name=GAS S,
cas_f1_ntk_cs=Gas Cylinder, mfName=Company, mfName_ntk_cs=Company, mfName_ntk=Company,buyable=1, published=1, catentry_id=98765,
partNumber_ntk=Cylinder2}]
```

- No obvious match



## Issue Scenarios – Queries

### Unexpected Search Results

- Debug trace enabled SOLR side when tracing enabled
- Shows in which field the match occurred, and on which search term:

```
getEntityObjects Query response: {responseHeader={status=0,QTime=1093,params={q=  
....  
response={numFound=39,start=0,docs=...  
...  
explain={72820=  
0.18911405 = (MATCH) sum of:  
    0.18911405 = (MATCH) weight(defaultSearch:welder in 5115),
```

- Can see match is occurring in the field 'defaultSearch', for the term 'welder'
- 'defaultSearch' field was not returned in the search results

# Issue Scenarios – Queries

## Unexpected Search Results

### Cause

- Fields are not returned in document for two possible reasons:
  - Commerce side: limited by 'fl'
    - fl=name+member\_id+mfName
  - SOLR side: limited by stored property in schema.xml
    - `<field name="defaultSearch" type="wc_text" indexed="true" stored="false" multiValued="true"/>`
- DefaultSearch defined as a combination of multiple fields in schema.xml:

```
<!-- Copy fields for default search field -->
<copyField source="name" dest="defaultSearch"/>
<copyField source="shortDescription" dest="defaultSearch"/>
<copyField source="partNumber_ntk" dest="defaultSearch"/>
<copyField source="keyword" dest="defaultSearch"/>
<copyField source="cas_f*" dest="defaultSearch"/>
<copyField source="cai_f*" dest="defaultSearch"/>
<copyField source="caf_f*" dest="defaultSearch"/>
<copyField source="ads_f*" dest="defaultSearch"/>
<copyField source="adj_f*" dest="defaultSearch"/>
<copyField source="adf_f*" dest="defaultSearch"/>
```

# Issue Scenarios – Queries

## Unexpected Search Results

### Solution

- Queried the individual fields with the problematic search term 'welder':

http://<hostname>:3737/solr/MC\_10001\_CatalogEntry\_en\_US/select?q=partNumber\_ntk:welder

```
<response>
- <lst name="responseHeader">
  <int name="status">0</int>
  <int name="QTime">9</int>
  - <lst name="params">
    <str name="q">partNumber_ntk:welder</str>
  </lst>
</lst>
<result name="response" numFound="0" start="0"/>
</response>
```

http://<hostname>:3737/solr/MC\_10001\_CatalogEntry\_en\_US/select?q=keyword:welder

```
- <result name="response" numFound="39" start="0">
```

- Keyword contained 'gas used with welder'
- Remove term from keyword, or remove keyword from defaultSearch

# Quick Reference

Activity	Logs	Files
Installing Feature pack	WC_installdir/logs/FEP2/*	
Enabling foundation feature	WC_installdir/instances/instance_name/logs/enablefoundation_timestamp.log	
setupsearchIndex.sh on Commerce Server	WC_installdir/components/foundation/subcomponents/search/log/wc-search-index-setup.log	
setupsearchIndex.sh on Remote SearchServer	working_dir/search/log/wc-search-index-setup.log	
di-preprocess.sh	WC_installdir/logs/wc-dataimport-preprocess.log	WC_installdir/instances/instance_name/search/pre-ProcessConfig/MC_catalog_ID/database_type/wc-dataimport-preprocess-*.xml
di-buildindex..sh	WC_installdir/logs/wc-dataimport-buildindex.log	../solr/home/<MC_catalog_ID>/<language>/CatalogEntry/conf/schema.xml ../solr/home/<MC_catalog_ID>/<language>/CatalogEntry/conf/wc-data-cōnfig.xml
Replication	WAS_installdir/profiles/search_profile_name/logs/search_server_name/Trace.log	Solr/home/<MC_masterCatalogId>/language/<index type>/conf/solrcōnfig.xml from all search servers/indexes
Querying	Commerce Side trace: com.ibm.commerce.foundation.*all:com.ibm.commerce.catalog.*=all:com.ibm.commerce.search.*=all WAS_installdir/profiles/commerce_profile_name/logs/commerce_server_name/trace.log  SOLR Side trace: org.apache.solr.*=all WAS_installdir/profiles/search_profile_name/logs/search_server_name/trace.log	WC_ear_dir/xml/config/com.ibm.commerce.catalog-fep/wc-search.xml



## Quick Reference

### Technotes

- DB2 error SQLCODE=-803, SQLSTATE=23505 occurs during di-preprocess
  - <http://www.ibm.com/support/docview.wss?uid=swg21508894>
- Crawler.sh fails with NoClassDefFound error
  - <http://www.ibm.com/support/docview.wss?uid=swg21578873>
- 'Top Search Hits' and 'Top Search Misses' pages in Management Center are blank
  - <http://www.ibm.com/support/docview.wss?uid=swg21578874>
- Attribute Facets not working - Corrupt metadata
  - <http://www.ibm.com/support/docview.wss?uid=swg21584798>
- NullPointerException on WebSphere Commerce search server startup
  - <http://www.ibm.com/support/docview.wss?uid=swg21508535>

### APARs

- JR41404 – FEP3 - reduce SQL impact in B2B/esites environments
- JR40305 & JR40304 – Cumulative fix for FEP2
- JR39904 – FEP2 & 3 – properly merges changes to wc-search.xml & wc-component.xml in *catalog-ext* directory

# Summary

- Reviewed Architectural Overview
- Learned basic troubleshooting techniques for common issues
  - ▶ Architecture
  - ▶ Building Index
  - ▶ Queries

# Additional WebSphere Product Resources

- Learn about upcoming WebSphere Support Technical Exchange webcasts, and access previously recorded presentations at:  
[http://www.ibm.com/software/websphere/support/supp\\_tech.html](http://www.ibm.com/software/websphere/support/supp_tech.html)
- Discover the latest trends in WebSphere Technology and implementation, participate in technically-focused briefings, webcasts and podcasts at:  
<http://www.ibm.com/developerworks/websphere/community/>
- Join the Global WebSphere Community:  
<http://www.websphereusergroup.org>
- Access key product show-me demos and tutorials by visiting IBM® Education Assistant:  
<http://www.ibm.com/software/info/education/assistant>
- View a webcast replay with step-by-step instructions for using the Service Request (SR) tool for submitting problems electronically:  
<http://www.ibm.com/software/websphere/support/d2w.html>
- Sign up to receive weekly technical My Notifications emails:  
<http://www.ibm.com/software/support/einfo.html>

# Connect with us!

## 1. Get notified on upcoming webcasts

Send an e-mail to [wsehelp@us.ibm.com](mailto:wsehelp@us.ibm.com) with subject line “wste subscribe” to get a list of mailing lists and to subscribe

## 2. Tell us what you want to learn

Send us suggestions for future topics or improvements about our webcasts to [wsehelp@us.ibm.com](mailto:wsehelp@us.ibm.com)

## 3. Be connected!

Connect with us on [Facebook](#)

Connect with us on [Twitter](#)





# Questions and Answers