


# WebSphere® Commerce V7 Feature Pack 5

## Search architecture updates



This presentation provides an overview of the architecture updates for the WebSphere Commerce search solution in Version 7 Feature Pack 5. You should be familiar with the existing WebSphere Commerce search solution before viewing this presentation.

## Table of contents

- Solution overview
- Facet management architecture
- Workspace preview support

This presentation begins with an overview of the new search capabilities. Following the overview is a discussion of the solution architecture updates for facet management and workspace preview.

## ***Solution overview***

This section introduces the WebSphere Commerce search solution enhancements.

## Facet management in attribute dictionary

Manage Attribute

Displayable

Comparable

Facetable

Note (United States English)

Unit of Measure (United States English)

Associated Keyword (United States English)

**Facetable properties**

Show facet in search results

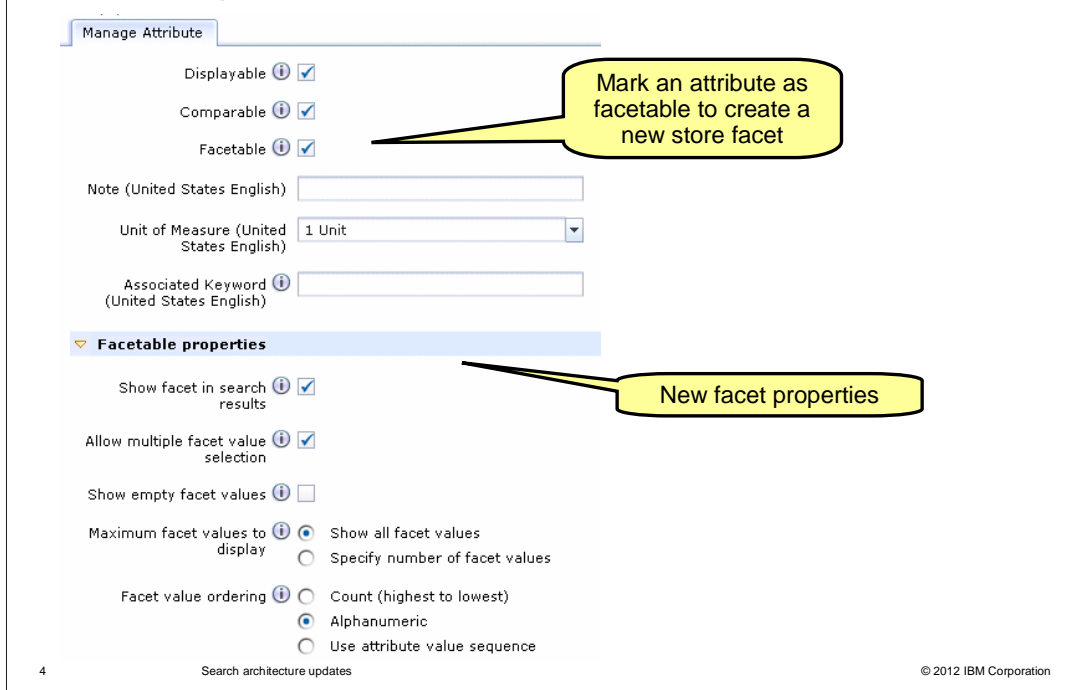
Allow multiple facet value selection

Show empty facet values

Maximum facet values to display  Show all facet values  
 Specify number of facet values

Facet value ordering  Count (highest to lowest)  
 Alphanumeric  
 Use attribute value sequence

4 Search architecture updates © 2012 IBM Corporation



The attribute dictionary tool in the Management Center Catalogs tool has been updated to allow business users to configure properties for storefront facets. Attributes can easily be enabled as facets by selecting the Facetable check box. Once an attribute is selected as facetable, a new properties section is added to allow business users to control the details of how the facet is displayed. For example, facets can be configured to allow multiple values to be selected in the storefront. The screen capture on this slide shows each of the available facet properties.

## Facet management for categories

- New facets list view
  - Control display sequence
  - Show or hide in the storefront
  - Display facet value counts

Sequence	Name	Display to customers	Attribute Code	Description	Values
1.0	Material	<input checked="" type="checkbox"/>	material	material	Cotton blend (17), 100% cotton (5), Synthetic (5), Silk (3), De
2.0	Color	<input checked="" type="checkbox"/>	swatchcolor	Color	Blue (25), Purple (16), Green (13), Gray (12), Brown (7), O (6), Black (5), Teal (4), Yellow Beige (2), Gold (2), Light Blue
0.0	Available Sizes	<input type="checkbox"/>	swatchSize	Available Sizes	L (39), M (39), S (39), XL (39)
3.0	Brand	<input checked="" type="checkbox"/>		The manufacturer	Albini (8), Gusso (8), Luigi Va (8), Hermitage Collection (7)
4.0	Price (USD)	<input checked="" type="checkbox"/>		Price-USD	{{* 100} 100} (31), {{100 200} 300} 300} (8), {{300 400} 400 500} (0), {{500 *} } (0)

5

Search architecture updates

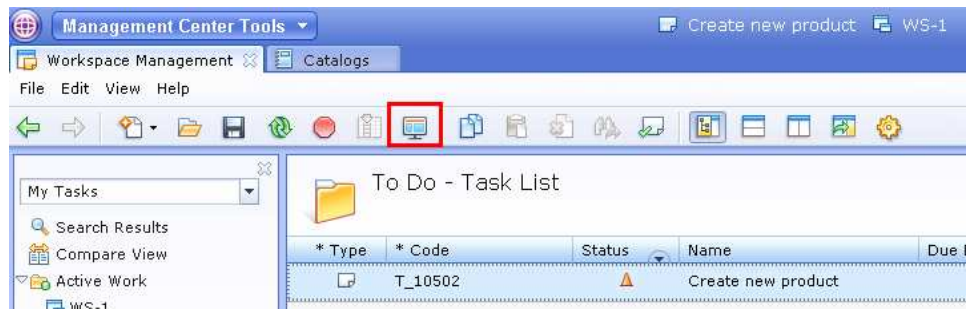
© 2012 IBM Corporation

Once you have attributes enabled as facets, another level of facet management is available in the Catalogs tool. A new category level list view allows you to see all the facets enabled for a category. From this list view, you can control the display order of facets within the category and choose to hide some facets from view in the storefront.

For a more detailed overview of the facet management solution, you can view the presentation and demonstration under the business user interface topic.

## Workspace preview support

- WebSphere Commerce search support added for workspace store preview
  - Preview changes within a task
  - Preview task group before approving



6

Search architecture updates

© 2012 IBM Corporation

Feature Pack 5 improves the usability of workspaces with WebSphere Commerce search by introducing full support for store preview within a workspace task. This was previously limited by not having task changes reflected in the search index. Contributors can now use store preview to verify changes made within a task and approvers can use it to verify changes before approving a task group.

## ***Facet management architecture***

This section covers the architecture updates for the facet management solution.

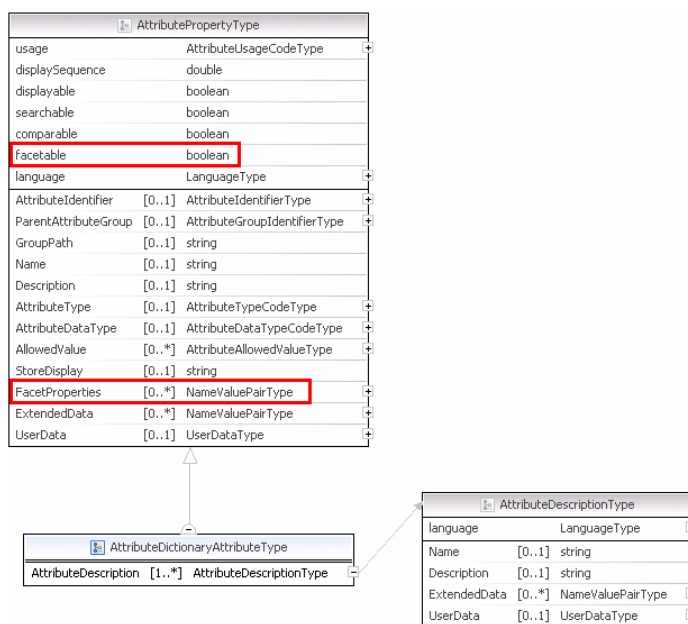
## Database tables

- Existing tables
  - ATTR
  - SRCHATTR
  - SRCHATTRPROP
- New tables
  - FACET
  - FACETDESC
  - FACETCATGRP

The database tables used to store facet information are listed on this slide. The attribute table has been updated with one new column to represent the facetable attribute flag. The search attribute and search attribute properties tables continue to be used to store information about attributes to be included in the search index. These include both searchable and facetable attributes. Three new tables have been added to support facet management. The facet table stores all of the facet display properties. The facet description table is only used for non-attribute based facets to store a language sensitive display name and description for the storefront. The facet to catalog group mapping table tracks category level facet settings such as whether to display the facet and the display sequence.



## Updates to AttributeDictionaryAttribute



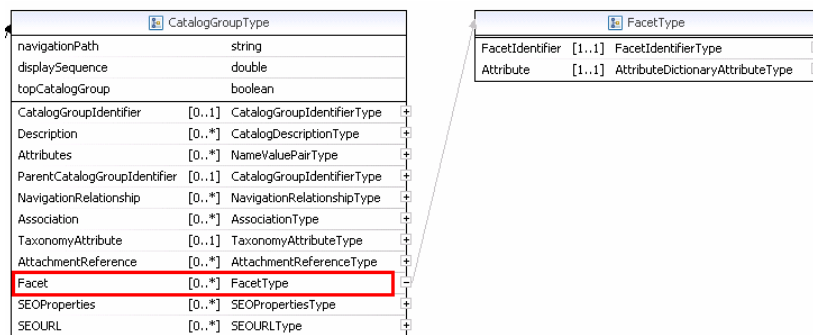
9

Search architecture updates

© 2012 IBM Corporation


The AttributeDictionaryAttribute noun has been updated to include the facetable flag and the FacetProperties noun part.

## Updates to CatalogGroup



The CatalogGroup noun has been updated to include a Facet noun part to store the category level facet settings.

## Updates to CatalogNavigationView



In the CatalogNavigationView noun, used by the storefront, the FacetView noun part has been updated. The FacetViewType and FacetEntryViewType include a new ExtendedData area to store facet properties needed for storefront display.

## Storefront expression builder updates

- Modified
  - getCatalogNavigationViewByCategory
    - Used by LeftNavigation widget
    - New parameters:

```
_wcf.search.facet='$facet$';_wcf.search.exclude.facet='$filterFacet$';_wcf.search.meta='$metaData$';_wcf.search.price.minimum='$minPrice$';_wcf.search.price.maximum='$maxPrice$'
```
- New
  - getCatalogNavigationBreadCrumbView
    - Used by BreadCrumb widget
  - getCatalogEntryViewAllWithoutAttachmentsByID
    - Used by ProductDescription widget

Feature Pack 5 includes new and updated expression builders for storefront use. The `getCatalogNavigationViewByCategory` expression builder is used to build the query for the facet navigation panel. This panel is implemented in the Aurora store by the `LeftNavigation` widget. New expression builders have been defined for use by the `BreadCrumb` and `ProductDescription` widgets.

## Limitations

- No data load support for facets
  - Upload new attributes and assign them to products
  - Set facetable attributes using Management Center tools
- Facet management is not workspace enabled


The facet management solution currently has two limitations you should take note of. First, the facetable flag and facet properties cannot be uploaded through the catalog upload tool or data load command line. You can upload attributes and assign them to products but marking attributes as facets must be done through the Management Center tools. The second limitation is facet management is not workspace enabled. If you are working on an authoring server, setting an attribute as facetable will apply directly to your production ready content.

## ***Workspace preview support***

This section covers the architecture updates for workspace preview support.

## Review of store preview in workspaces

- Workspace read schema combines task changes with production ready data



- Store preview accesses the read schema to display updated store before changes are committed



15

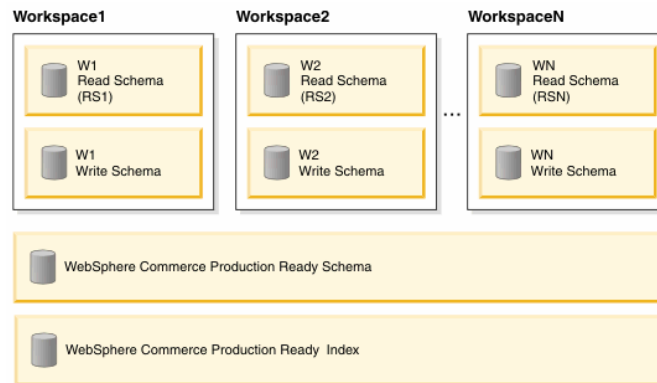
Search architecture updates

© 2012 IBM Corporation

In an authoring environment with workspaces enabled, business users are able to preview their changes before submitting them for approval. The store preview option in Management Center accesses the workspace read schema to provide a view of the store that combines the local workspace changes in the write schema with the existing production ready data.

## Pre-FEP 5 limitation with store preview

- Not all data is read from the database, some comes from the search index
- The search index contains only production ready data

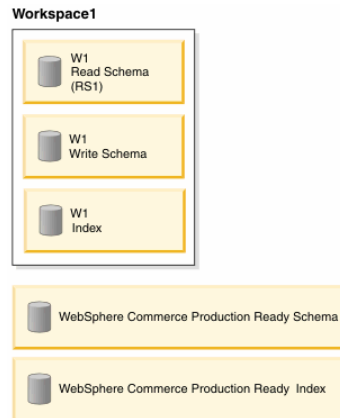


When WebSphere Commerce search is enabled, much of the catalog data displayed in the storefront comes from the search index, not the WebSphere Commerce database. Before Feature Pack 5, the authoring server had a single search index that contained only production ready data. This meant business users were not able to preview catalog changes made in a workspace since the changes were not reflected in the search index.



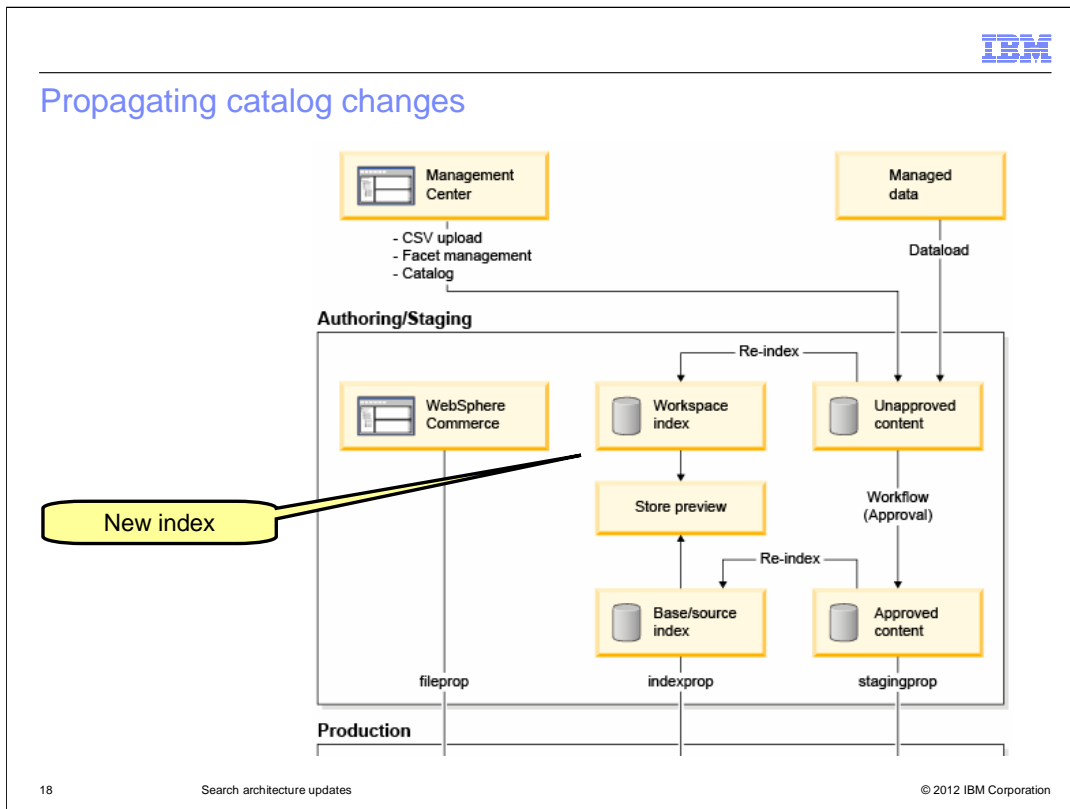
## Feature Pack 5 enhancements

- A search index is created for the workspace
  - A single workspace index is supported at a time
- Search results from the workspace index are merged with the production ready index



Feature Pack 5 solves the preview limitation by adding support for a workspace index. Similar to the write schema, the workspace index contains records for categories and catalog entries that have been added or changed within the workspace. Within the store preview, results from the workspace index are combined with the production ready index to allow the business user to see their changes in the storefront catalog and search results. While an authoring server might define multiple workspaces, only one workspace at a time can have an associated search index.

## Propagating catalog changes



In this diagram, you see the location of the new workspace index within the overall authoring server. Changes made through Management Center and data load are stored in the unapproved content database, also known as the write schema. When changes affecting the search index are made and a re-index is triggered, the workspace index is updated. When the re-index is complete, business users can verify changes using store preview. Store preview displays the combined results from both the workspace index and the base production ready index. Once changes are complete and the task group approved, the updates are pushed to the approved content database and the base index is updated. From there, the changes are ready to push to production during the next staging propagation.

## Index update triggers

- Store preview
- Show facet – CatalogEntry index only
- Task group commit - authoring servers only
- Scheduler job (UpdateSearchIndex) – production ready index only

As business users work in an authoring or staging environment, it is necessary for the search index to occasionally update to reflect changes made. The primary re-index trigger is launching the store preview from Management Center. Both the workspace index and production ready index are updated if changes are pending. Opening the new facet list view for a category also triggers an index update but only of the CatalogEntry index. In an authoring environment, when changes for a task group are committed to the production ready schema the index is updated. In addition to business user initiated updates, there is also a scheduler job, UpdateSearchIndex, which regularly updates the production ready index.


## Database updates

- Tables to store changes waiting to be indexed in each workspace
  - TI\_DELTA\_CATENTRY
  - TI\_DELTA\_CATGROUP
- Tables to store all changed records in each workspace
  - TI\_CATENTRY\_WS
  - TI\_CATGROUP\_WS

Four new tables help track changes that require updates to the workspace index. The TI\_DELTA\_CATENTRY and TI\_DELTA\_CATGROUP tables in the workspace write schema store objects that have changed since the last re-index. When an index update is triggered, the rows in these tables identify which records in the index to update. They also track whether a full or delta re-index is required. Once the index update is complete, the rows are removed from the tables. The TI\_CATENTRY\_WS and TI\_CATGROUP\_WS tables in the write schema track a cumulative list of all of the changes that require an index update. Unlike the delta tables, rows are not removed from these tables when the index is updated.

## Index updates

- Data Import Handler (DIH) accesses workspace read schema to populate workspace core
  - Configured in wc-data-config.xml
- Workspace Index Reader combines results from workspace core and production ready core
  - Based on Lucene MultiReader
  - Limited to two sub index readers
  - Assigns priority to documents from workspace core
  - Specified in solrconfig.xml



The new workspace index cores are created in the same location as the approved content cores. The folders for the workspace cores have the workspace identifier appended to the folder name as shown in the screen capture on the slide. The fields in the workspace index are the same as the approved content index however there are some configuration differences. For example, the data import handler accesses the workspace read schema to populate the workspace core. This is configured in the wc-data-config.xml file. Another important difference is the workspace index reader. This reader combines results from both the workspace core and the production ready core. It is based on the Lucene MultiReader but is limited to just the two sub index readers. When a document with the same ID is contained in both result sets, the one from the workspace core takes precedence. The workspace index reader is specified in the solrconfig.xml file.

## Utility changes

- Search index setup command
  - New parameters added to pass DBA user credentials
    - -dbauser: the DBA user account
    - -dbauserpwd: the password of DBA user account
  - Workspace must be allocated first
- Pre-process and index building
  - New parameter to specify a workspace schema
    - -workspace: the identifier of the workspace to operate on

The command line utilities for creating and building search indexes have been updated to support the new workspace index. The setup search index utility requires two new parameters, the dba user name and password, to setup the workspace index. The workspace should be allocated before running the creation script.

The pre-process and index building utilities also have a new parameter to specify a workspace ID when operating on the workspace index. This new parameter is primarily used by the automated index update processes.

## Summary

- Solution overview
- Facet management architecture
- Workspace preview support

This presentation began with an overview of the new search capabilities. Following the overview was a discussion of the solution architecture updates for facet management and workspace preview.

## References

- Workspaces in WebSphere Commerce search
  - <http://publib.boulder.ibm.com/infocenter/wchelp/v7r0m0/topic/com.ibm.commerce.developer.doc/concepts/csdmanagesearchworkspaces.htm>
- Setting up WebSphere Commerce search (updated for workspaces)
  - <http://publib.boulder.ibm.com/infocenter/wchelp/v7r0m0/topic/com.ibm.commerce.developer.doc/tasks/tsdsearchsetuplocal.htm>
- Indexing with staging propagation (updated for workspaces)
  - <http://publib.boulder.ibm.com/infocenter/wchelp/v7r0m0/topic/com.ibm.commerce.developer.doc/concepts/csdmanagesearchstage.htm>
- Deploying WebSphere Commerce search (new web server option)
  - <http://publib.boulder.ibm.com/infocenter/wchelp/v7r0m0/topic/com.ibm.commerce.developer.doc/tasks/tsdsearchbuilddeployfirst.htm>
- Common business tasks and their impact to the WebSphere Commerce search index
  - <http://publib.boulder.ibm.com/infocenter/wchelp/v7r0m0/topic/com.ibm.commerce.developer.doc/refs/rsdsearchindexhints.htm>
- Catalog upload operations in WebSphere Commerce search
  - <http://publib.boulder.ibm.com/infocenter/wchelp/v7r0m0/topic/com.ibm.commerce.management-center.doc/concepts/cpsearchcatuploadadmin.htm>

This slide contains some useful references on WebSphere Commerce search.



## Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

[mailto:iea@us.ibm.com?subject=Feedback\\_about\\_SearchArchitectureUpdates\\_FEP5.ppt](mailto:iea@us.ibm.com?subject=Feedback_about_SearchArchitectureUpdates_FEP5.ppt)

This module is also available in PDF format at: [../SearchArchitectureUpdates\\_FEP5.pdf]( ../SearchArchitectureUpdates_FEP5.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.



## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, and WebSphere are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.