ONDEMAND NEWSLETTER

NEWS AND TIPS ABOUT IBM CONTENT MANAGER ONDEMAND

4TH QUARTER 2014

IN THIS ISSUE

NEWS

About this Newsletter

This newsletter is designed to keep you better informed about IBM® Content Manager OnDemand on all platforms. The newsletter will be published quarterly.

Previous editions of this newsletter can be found on the OnDemand support web sites under the 'Featured links' heading. They are also available on the OnDemand Users Group web site under the heading 'Presentations, Newsletters, and such'.

Correspondence related to this bulletin should be directed to odnews@us.ibm.com.

OnDemand Version 9.5 available

Version 9.5 is now available for IBM Content Manager OnDemand for Multiplatforms, IBM Content Manager OnDemand for z/OS, and IBM Content Manager OnDemand for i.

Version 9.5 of OnDemand contains numerous new and enhanced functions. These include:

- Report Distribution (Multiplatforms) and ODF (z/OS) combined into one feature
- Ability to load data through ODWEK
- Support for multiple language names and descriptions for folders, cabinets, holds
- Non-root installation (UNIX)
- ODWEK support for multiple locales
- ACIF and PDF Indexer support for regular expressions to locate triggers and fields

- OS/390 Indexer enhancements
- XML Indexing
- ARSLOAD removal of duplicate records during the load process
- Report Wizard support of Full Text Indexing, Document Size, and Page Count

For more information, see IBM United States Software [Announcement 214-453](#) (MP) and IBM United States Software [Announcement 214-451](#) (z/OS).

OnDemand Windows clients

The V9.5 Windows clients (end-user, Administrator, and Monitor) are available from IBM Fix Central.

OnDemand for Multiplatforms

Version 9.5 is available from Passport Advantage.

OnDemand for z/OS

Version 9.5 is available from ShopZ.

OnDemand for i

See [Information APAR II14723](#) for a list of the required PTFs.

Note that for OnDemand for i, server level 9.5.0.0 is available only for V7.2.

Content Manager OnDemand Workshops

The IBM Content Manager OnDemand workshop is a series of complementary events held in cities around the world. The workshop features both informative and how-to sessions with demonstrations of the latest features IBM Content Manager OnDemand has to offer from IBM product experts.

More information, including the dates and locations of upcoming workshops can be found on the [workshop web site](#).

Content Navigator 2.0.3 available

IBM Content Navigator (ICN) 2.0.3 is now available. Also, IBM Content Navigator v2.0.3 for Mobile is now available on iTunes and Google Play.

IBM Content Navigator is available to customers of these products:

- IBM Content Manager
- IBM Content Manager OnDemand
- IBM FileNet Content Manager
- IBM Content Foundation

To learn about the new features of ICN, visit the '[What's new](#)' topic in the ICN Knowledge Center.

Get connected to Enterprise Content Management Support through social media

IBM ECM Support is now on LinkedIn, Facebook and Twitter.

We invite you to Connect with us on LinkedIn, Like us on Facebook and/or Follow us on Twitter. Becoming a member of the growing community of clients, developers and support engineers is a click away!

[LinkedIn](#)   [Facebook](#)   [Twitter](#)

In addition, a collection of useful links for OnDemand clients and servers can be found in the [IBM Client Success Essentials wiki](#).

ECM Redbook available

IBM Enterprise Content Management (ECM) solutions provide efficient and effective ways to capture content, manage the content and business processes, discover insights from the content, and derive actions to improve business processes, products, and services.

An IBM Redbooks publication, Building IBM Enterprise Content Management Solutions From End to End, introduces and highlights some of IBM ECM products that can be implemented and integrated together to create end-to-end ECM solutions, including:

- IBM Case Manager

- IBM Datacap

- IBM Content Manager OnDemand

- IBM Enterprise Records

- IBM Watson Content Analytics

- IBM Content Classification

For each product involved in the ECM solution, the redbook briefly describes what it is, its functions and capabilities, and provides step-by-step procedures on installing, configuring, and implementing it. In addition, the book provides procedures on integrating these products together to create an end-to-end ECM solution to achieve your overall solution objectives.

Not all the products are required to be integrated. Depending on your business requirement, you can choose a subset of these products to be built into your ECM solutions.

This book serves as a hands-on learning guide for IT specialists who plan to build ECM solutions from end-to-end, for a proof of concept, or proof of technology environment. For implementing a production-strength ECM solution, also refer to IBM Knowledge Center, IBM Redbooks publications, and IBM Software Services.

Did you know?

You are using OnDemand today! With millions of external licenses in production, if you are accessing your utility bills, phone bills, and credit card statements through a customer portal, odds are, it's coming from OnDemand!

IBM i 6.1 end of support

Effective September 30, 2015, IBM will withdraw support for IBM i V6.1. See IBM United States Withdrawal Announcement 914-181 for more details.

TIPS — CROSS PLATFORMS

PDF Indexer enhanced

This tip applies to the PDF Indexer on both Multiplatforms and IBM i. The PDF Indexer is deprecated on z/OS at V9.5.

Version 9.5 of OnDemand adds regular expression support to the PDF Indexer.

What are regular expressions?

A regular expression (abbreviated regex or regexp) is a sequence of characters that forms a search pattern, mainly for use in pattern matching with strings. A pattern consists of one or more character literals, operators, or constructs. Regular expressions are similar to masks but more flexible.

TRIGGER parameter

Regular expression support provides more flexibility in defining triggers.

For example:

```
TRIGGER1=UL(0.02,6.15),LR(0.36,7.63),*,REGEX='P[A-Z]{9}',(TYPE=GROUP)
```

Note the new sub-keyword REGEX.

This regular expression will match text that begins with the letter 'P' and is followed by nine uppercase letters, for example, "PRESCRIBER".

FIELD parameter

Regular expression support can be used to extract variable length text from a field. When a regular expression is used, a mask cannot be used.

For example:

```
FIELD1=UL(0.46,3.47),LR(0.82,7.46),0,(TRIGGER=1,BASE=0, REGEX='[A-Z][a-z]+
    [A-Z][a-z]+',DEFAULT='N/A')
```

If the FIELD1 parameter has extracted the text "Client: John Smith", then this regular expression will extract the name "John Smith" from the text.

Additional resources

More information on using regular expressions with the PDF Indexer can be found in the Content Manager OnDemand Version 9 Release 5 Indexing Reference.

There are many other excellent online resources which explain the syntax rules of regular expressions.

New XML Indexer

Introduction

Prior to OnDemand V9.5, XML documents had to be loaded using the Generic Indexer. Documents had to be broken up or offsets had to be calculated. There was no easy way to support resources.

With the introduction of the XML Indexer in V9.5, the indexing process for XML documents is now integrated with ARSLOAD and the Add Report to OnDemand (ADDRPTOND) command. Documents are extracted during indexing. Resources are extracted and archived with the data.

Configuration

To configure an application to use the XML Indexer, select a data type of XML and an indexer of XML.

Transform

Because XML data can contain any kind of information using an endless number of tags, you must indicate to the XML Indexer what documents within your data you would like stored along with the index values for each document. This is accomplished by transforming your XML input into an intermediate format. Tools like XSLT and XQuery can be used get your data ready for indexing.
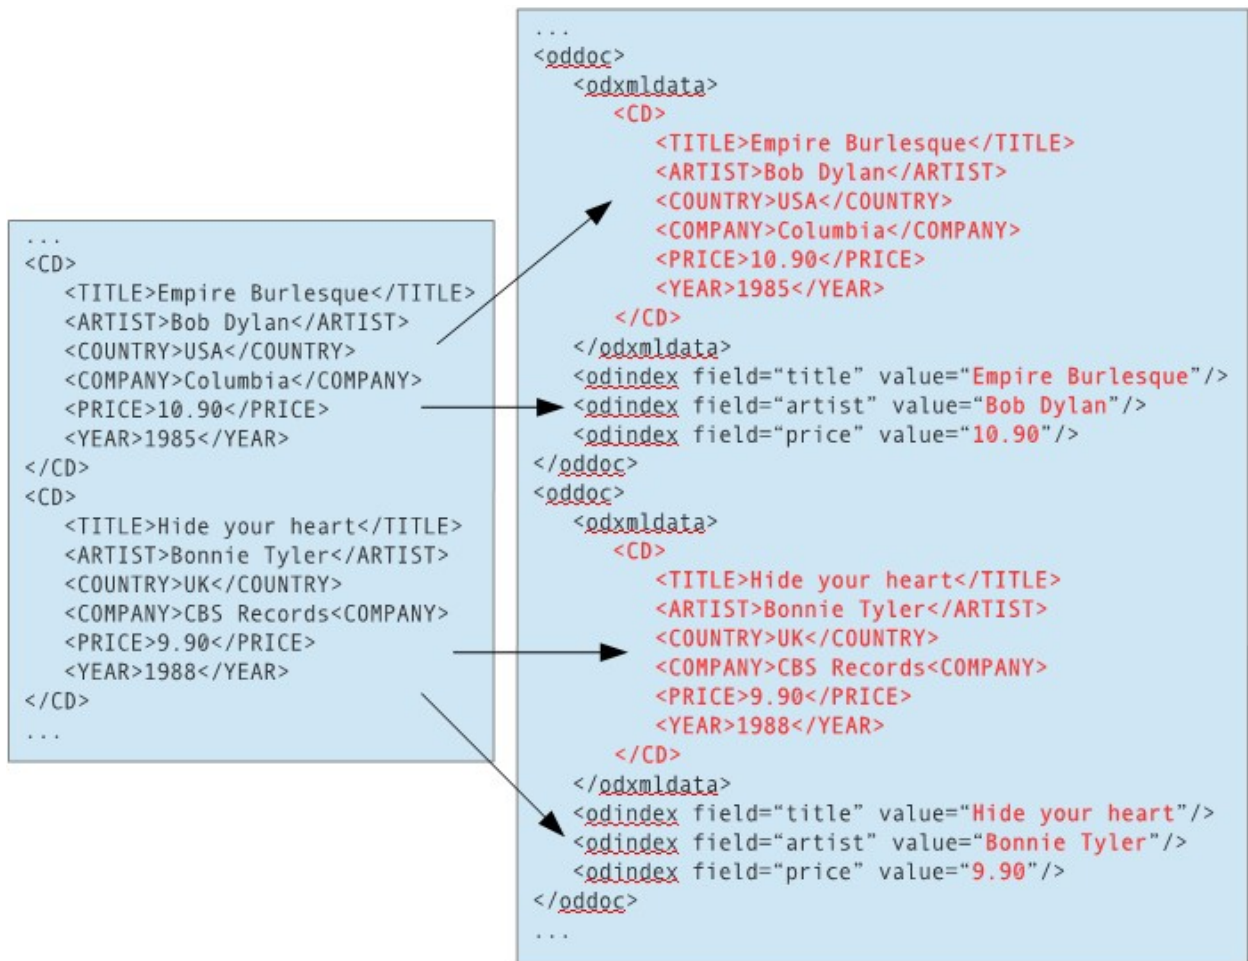
XSLT (Extensible Stylesheet Language Transformations) is a language for transforming XML documents into other XML documents, or even plain text. During the transformation, the original document is not changed; rather, a new document is created based on the content of an existing one. The basic processing paradigm is pattern matching. The XSLT style sheet defines what patterns to process and how to process them for output. While there are many processor implementations of XSLT, Saxon and Xalan are two of the more popular open source versions.

XQuery is a query and functional programming language. XQuery can be used to query and transform customer XML data into the format required by OnDemand for XML indexing.

The XML Indexer will validate the structure and content of your input file using the XML Indexer schema file named odxmlidx.xsd.

The following figure shows an XML input file before and after transformation.

```
...
<oddoc>
   <odxmldata>
      <CD>
         <TITLE>Empire Burlesque</TITLE>
         <ARTIST>Bob Dylan</ARTIST>
         <COUNTRY>USA</COUNTRY>
         <COMPANY>Columbia</COMPANY>
         <PRICE>10.90</PRICE>
         <YEAR>1985</YEAR>
      </CD>
   </odxmldata>
   <odindex field="title" value="Empire Burlesque"/>
   <odindex field="artist" value="Bob Dylan"/>
   <odindex field="price" value="10.90"/>
</oddoc>
<oddoc>
   <odxmldata>
      <CD>
         <TITLE>Hide your heart</TITLE>
         <ARTIST>Bonnie Tyler</ARTIST>
         <COUNTRY>UK</COUNTRY>
         <COMPANY>CBS Records<COMPANY>
         <PRICE>9.90</PRICE>
         <YEAR>1988</YEAR>
      </CD>
   </odxmldata>
   <odindex field="title" value="Hide your heart"/>
   <odindex field="artist" value="Bonnie Tyler"/>
   <odindex field="price" value="9.90"/>
</oddoc>
...
```

```
...
<CD>
   <TITLE>Empire Burlesque</TITLE>
   <ARTIST>Bob Dylan</ARTIST>
   <COUNTRY>USA</COUNTRY>
   <COMPANY>Columbia</COMPANY>
   <PRICE>10.90</PRICE>
   <YEAR>1985</YEAR>
</CD>
<CD>
   <TITLE>Hide your heart</TITLE>
   <ARTIST>Bonnie Tyler</ARTIST>
   <COUNTRY>UK</COUNTRY>
   <COMPANY>CBS Records<COMPANY>
   <PRICE>9.90</PRICE>
   <YEAR>1988</YEAR>
</CD>
...
```

Indexing and loading

Indexing and loading are performed in a single pass. No intermediate files are created. You simply run ARSLOAD, for example:

```
arsload -nvf -I <instance> -u <user> -p <stash> -g <ag_name> <input_file>
```

On IBM i, you can also use the Add Report to OnDemand (ADDRPTOND) command to index and load the data.

```
ADDRPTOND APPGRP('<ag_name>') INPUT(*STMF) STMF('<input_file>')
   INSTANCE(<instance>)
```

Resources

Style sheets are specified with the <xsl-stylesheet> processing instruction. All style sheet references are copied to each XML document archived. References of type "file:" are gathered and archived as OnDemand resources. Resources must be in the indexer execution directory or a sub-directory. The resources are stored in a zip file format.

Viewing

Windows client

The OnDemand Windows client unzips the resources and launches the external viewer based on mime-type (application/xml).

ODWEK

The ODWEK CGI and Servlet will return just the XML document to the browser. No resources are returned.

The Java APIs return both the XML document and resources.

Content Navigator

Content Navigator unzips the resources and displays the document in the browser window.

Generic XML Index File Format (GXIFF)

The Generic XML Index File Format (GXIFF) provides a method for specifying the indexes in XML for any data type to be stored into OnDemand. GXIFF provides the same basic functions as the Generic Indexer, but using XML for the index file.

In the following example, two jpeg images are being loaded. In the application, the Data Type is specified as JPG and the Indexer is specified as XML.

```xml
<?xml version="1.0" encoding="UTF-8" ?>
<odidx xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
           xsi:noNamespaceSchemaLocation="odxmlidx.xsd">
<oddoc>
   <oddataref file="image1.jpg" offset="0" length="0">
   <odindex field="title" value="Empire Burlesque"/>
   <odindex field="artist" value="Bob Dylan"/>
   <odindex field="price" value="10.90"/>
</oddoc>
<oddoc>
   <oddataref file="image2.jpg" offset="0" length="0">
   <odindex field="title" value="Hide your heart"/>
   <odindex field="artist" value="Bonnie Tyler"/>
   <odindex field="price" value="9.90"/>
</oddoc>
...
</odidx>
```

GXIFF also supports loading multiple rows for a single document. The file name, offset, and length must be the same, as shown in the following example.

```xml
<?xml version="1.0" encoding="UTF-8" ?>
<odidx xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
            xsi:noNamespaceSchemaLocation="odxmlidx.xsd">
<oddoc>
    <oddataref file="image1.jpg" offset="0" length="0">
    <odindex field="title" value="Empire Burlesque"/>
    <odindex field="artist" value="Bob Dylan"/>
    <odindex field="price" value="10.90"/>
</oddoc>
<oddoc>
    <oddataref file="image1.jpg" offset="0" length="0">
    <odindex field="title" value="Hide your heart"/>
    <odindex field="artist" value="Bonnie Tyler"/>
    <odindex field="price" value="9.90"/>
</oddoc>
...
</odidx>
```

Negative mask values

All OnDemand customers should check their application groups for negative mask values. A negative mask value might cause a load failure with message ARS1130.

On z/OS, perform the following query:

```
SELECT AGID, NAME, MASK, HEX(MASK) FROM <dbowner>.ARSAGFLD WHERE MASK < 0
```

On Multiplatforms with DB2, perform the following query:

```
SELECT AGID, NAME, MASK, HEX(MASK) FROM <schema>.ARSAGFLD WHERE MASK < 0
```

On IBM i, perform the following query:

```
SELECT AGID, NAME, MASK, HEX(MASK) FROM <instance>.ARSAGFLD WHERE MASK < 0
```

If any data is returned, contact IBM software support.

Using stash files – the basics

A password stash file is an encrypted file used to store userids and passwords. Use of a stash file is recommended so that passwords do not have to be specified in a visible format in plain text (also known as "in the clear"). Stash file support was added to the OnDemand ARSxxx commands at V9. Stash file support was added to OnDemand for i commands at V7.2. Multiple stash files are supported.

A stash file can be specified in two ways. First, it can be specified as a parameter on the commands. Second, it can be specified in the instance stanza in the ARS.INI file.

A stash file can contain default userids and passwords for certain functions, such as ARSLOAD. It can also contain userids and passwords that can be specified as parameters on the commands.

About the ARSLOAD examples

Examples 1 and 2 below assume a user ID was added to the stash file that will be used by ARSLOAD when the user ID and password is not provided on the command line. To add the user ID "sysadmload" to the stash file as a default user ID for ARSLOAD, use the following command:

```
arsstash -a 3 -s stash_file -u sysadmload
```

-a 3   adds the user ID for the ARSLOAD command

-s   stash_file is the fully qualified path name of the stash file

-u sysadmload is the user ID that will be used to logon to the server when
ARSLOAD is executed without specifying a user ID and password (i.e. -u userID
and -p password aren't provided on command line)

In examples 1 and 2, if "sysadmload" was the user added to the stash file for ARSLOAD, ARSLOAD would use the user ID "sysadmload" to logon to the server with the password from the stash file.

Examples 3, 4, and 5 assume the user ID "bob" was added to the stash file prior to running the ARSLOAD command. To add "bob" to the stash file, use the following command:

```
arsstash -a 1 -s stash_file -u bob
```

-a 1   adds the user ID to the stash file

-s   stash_file fully qualified path name of the stash file

-u bob is the user ID that is added to the stash file (If -u bob is used on command line,
-p is not required since the password is obtained from the stash file.)

In example 5, the stash file name must be specified with the -p parameter because a TCP/IP address (9.99.111.222) was provided rather than an instance name, such as ARCHIVE or QUSROND. ARSLOAD will logon to the server using the user ID "bob". The password for "bob" is from the stash file shown as "stash_file" in the example.

ARSLOAD examples

1. The following shows how to run the ARSLOAD program to check the specified directory for input files to process. The input files must have a file type extension of .ARD or .PDF. The ARSLOAD program stores temporary work files in the location specified by the -c parameter. In this example, the ARSLOAD program uses the FORM part of the file name to determine the application group to load and the application group contains only one application; the ARSLOAD program logs on to the system and accesses the application group with the user ID and password stored in the stash file.

   ```
   arsload -h ARCHIVE -c /arsacif/acif1 -d /arsacif/acif2
   ```

2. The following shows how to run the ARSLOAD program to check the specified directory for input files to process. The input files must have a file type extension of .ARD or .PDF. The ARSLOAD program stores temporary work files in the location specified by the -c parameter. In this example, the ARSLOAD program uses the JOBNAME part of the file name to determine the application group to load and the DATASET part of the file name to determine the application to load; the ARSLOAD program logs on to the system and accesses the application group with the user ID and password stored in the stash file.

   ```
   arsload -h ARCHIVE -c /arsacif/acif1 -d /arsacif/acif2 -A DATASET -G
       JOBNAME
   ```

3. The following shows how to run the ARSLOAD program to load the specified file into the specified application group. The ARSLOAD program logs on to the system and accesses the application group by using the user ID from the -u parameter and the password for user ID "bob" from the stash file.

   ```
   arsload -h ARCHIVE -g BILLS -u bob PO3510
   ```

4. The following shows how to run the ARSLOAD program to load several input files into the specified application group. The application group name contains an embedded blank character, and must be quoted.

   ```
   arsload -h ARCHIVE -g "ABC Credit" -u bob RW7505 RW8505
   ```

5. The following shows how to run the ARSLOAD program to load the specified file into the specified application group. The TCP/IP address of the IBM Content Manager OnDemand library server is specified. This format of the command can be used to load input files from an object server and store the index data on the library server. The ARSLOAD program logs on to the system and accesses the application group by using the user ID from the -u parameter and the password for user ID "bob" from the stash file.

   ```
   arsload -h "9.99.111.222" -g BILLS -u bob -p stash_file ARSP0MST
   ```

Other ARSxxx examples

Commands such as ARSDOC and ARSXML can also use a stash file but the userid must always be specified as a parameter. For example:

1. The password is obtained from the stash file defined in the ARS.INI file (oduser must have

been added to the stash file.)

```
arsdoc get -h instance -u oduser
```

2. The password is obtained from the stash file specified with the -p parameter (oduser must have been added to the stash file.)

```
arsdoc get -h instance -u oduser -p stash_file
```

3. The password is obtained from the stash file specified with the -p parameter (oduser must have been added to the stash file.)

```
arsdoc get -h hostname  -u oduser -p stash_file
```

4. The password is obtained from the stash file defined in the ARS.INI file (oduser must have been added to the stash file.)

```
arsxml export -h instance -u oduser
```

5. The password is obtained from the stash file specified with the -p parameter (oduser must have been added to the stash file.)

```
arsxml export -h instance -u oduser -p stash_file
```

6. The password is obtained from the stash file specified with the -p parameter (oduser must have been added to the stash file.)

```
arsxml export -h hostname  -u oduser -p stash_file
```

IBM i commands

At V7.2, password stash file support was added to the following IBM i commands: Add Report to OnDemand (ADDRPTOND), Start Monitor for OnDemand (STRMONOND), and Merge Spooled Files (MRGSPLFOND).

When the STASHFILE parameter is specified on these commands, the ARSLOAD userid and password from the stash file are used.

32-bit versus 64-bit Java libraries

Note that this tip does not apply to Linux on System z, which includes only the 64-bit version.

Running a Java program that uses the ODWEK Java APIs

Beginning at ODWEK V8.5, 64-bit versions of the ODWEK shared library were made available for the Java APIs on all supported platforms.

To set up an environment where the shared library can be loaded, you must match up the 32-bit JVM with the 32-bit ODWEK Java API library and the 64-bit JVM with the 64-bit ODWEK Java API library.

The version that gets successfully loaded is determined by the JVM. In all attempts to run the

application, ODWEK will first try to load the 64-bit version and if that does not succeed, it will try to load the 32-bit version. If neither attempt is successful, a java.lang.UnsatisfiedLinkError error is issued and program execution is terminated.

Some error scenarios

The library path does not contain a directory whose contents contain the required library.

- The library path does not contain either the 32-bit or the 64-bit shared library.
- You are running a 64-bit JVM.
- You have a library path set to /path/to/libs.
- Neither the libars3w64.a nor the libars3w32.a libraries exist in /path/to/libs
- The JVM tries first to load libars3w64.a which fails because the file is not in /path/to/libs
- The JVM then tries to load libars3w32.a which fails because the file is not in/path/to/libs
- Result: Unsatisfied link error

The library path contains a mismatched version of the shared library.

- You are running a 64-bit JVM.
- You have your library path set to /path/to/libs
- /path/to/libs contains the libars3w32.a library.
- The JVM tries first to load libars3w64.a which fails because the file is not found in /path/to/libs
- The JVM then tries to load libars3w32.a which fails because the 64-bit JVM cannot load a 32-bit library.
- Result: Unsatisfied link error

Question

Will the JVM fail to launch the application successfully if the java.library.path contains both the 32-bit and the 64-bit version of the ODWEK library?

32-bit JVM

- You are running a 32-bit JVM.
- You have your library path set to /path/to/libs.
- Both the libars3w64.a and the libars3w32.a libraries exist in /path/to/libs.
- The JVM tries first to load libars3w64.a. It is successfully found in /path/to/libs but because the 32-bit JVM cannot load a 64-bit library, it fails, but does not cease execution.

- The JVM then tries to load libars3w32.a. It is successfully found in /path/to/libs and since the 32-bit JVM can load a 32-bit library.

- Result: Success!

64-bit JVM

- You are running a 64-bit JVM.

- You have your library path set to /path/to/libs.

- Both the libars3w64.a and the libars3w32.a libraries exist in /path/to/libs.

- The JVM then tries to load libars3w64.a. It is successfully found in /path/to/libs and since the 64-bit JVM can load a 64-bit library.

- Result: Success!

Windows Client: AND and OR logical operator options

Question

Can the AND and OR logical operator options be:

1. limited to a particular OnDemand user?

2. saved in a public query?

Note: The AND and OR logical operator options are located in the OnDemand Windows client Search Criteria and Document List window, just below Close Folder.

Answer

The AND and OR logical operator options cannot be limited to particular OnDemand users. You can set them only in the Search Criteria and Document List window of the client, and for this reason, they cannot be saved like query restrictions.

Note: The AND and OR operators apply to all fields *except* an application group's Segment Date field, which always uses AND with the other fields.

This tip is adapted from support item 1198445.

Tips — Multiplatforms

ACIF enhanced

This tip applies only to ACIF on Multiplatforms.

Version 9.5 of OnDemand adds several enhancements to ACIF.

Regular expression support

Regular expression support in ACIF is the same as that in the PDF Indexer.

What are regular expressions?

A regular expression (abbreviated regex or regexp) is a sequence of characters that forms a search pattern, mainly for use in pattern matching with strings. A pattern consists of one or more character literals, operators, or constructs. Regular expressions are similar to masks but more flexible.

TRIGGER parameter

Regular expression support provides more flexibility in defining triggers.

For example:

```
TRIGGER1=*,1,REGEX='P[A-Z]{3}',(TYPE=GROUP)
```

Note the new sub-keyword REGEX.

This regular expression will match text that begins with the letter 'P' and is followed by three uppercase letters, for example, "PAGE".

FIELD parameter

Regular expression support can be used to extract variable length text from a field.

For example:

```
FIELD1=0,13,31,(REGEX='[A-Z][a-z]+ [0-9]+, [0-9]{4}',DEFAULT='January 1,
    1970')
```

If the FIELD1 parameter has extracted the text "during the July 4, 1956 holiday", then this regular expression will extract the date "July 4, 1956" from the text.

Additional resources

More information on using regular expressions with ACIF can be found in the Content Manager OnDemand Version 9 Release 5 Indexing Reference.

There are many other excellent online resources which explain the syntax rules of regular expressions.

New parameter BREAKYES

When multiple index parameters are specified with the BREAK=YES sub-parameter, the BREAKYES parameter can be used to determine how these multiple index parameters are interpreted. They can either be logically AND'ed or OR'ed together.

When BREAKYES = AND is specified, the multiple index parameters with BREAK = YES are all AND'ed together. In other words, all index values must change for the current page to be considered the start of a new document.

When BREAKYES = OR is specified, the multiple index parameters with BREAK = YES are all OR'ed together. In other words, a change to any index value causes the current page to be considered the start of a new document.

The default is BREAKYES = OR.

New parameter LINEOFFSET

The LINEOFFSET parameter determines whether ANSI carriage-control characters are used to calculate the record offsets when determining the location of the fields.

Only the 0 (double space) and the dash (triple space) are supported. The + (overstrike) character is not supported. Machine carriage controls are not supported.

When LINEOFFSET=ASREAD is specified, ANSI carriage controls are not used to calculate the record offsets for the fields. The offsets are relative to the lines as they are read from the load file.

When LINEOFFSET=ASPRINTED is specified, ANSI carriage controls are used to calculate the record offsets for the fields. The offsets are relative to the line spacing that occurs when the lines are printed.

The default is LINEOFFSET = ASREAD.

For example, the first three records of an input file contain the following. The first character is an ANSI carriage control.

```
1REPORT
-ACCOUNT 777777
0JOHN SMITH
```

Using ASREAD (the default), the indexing parameters to collect the account number and name would be as follows:

```
TRIGGER1=*,1,'REPORT'
FIELD1=1,10,6
FIELD2=2,2,10
```

Using ASPRINTED, the indexing parameters to collect the account number and name would be as follows:

```
TRIGGER1=*,1,'REPORT'
FIELD1=3,10,6
FIELD2=5,2,10
```

This feature is meant to be used as a migration aid if you have existing indexing parameters that were created to take the carriage controls into account. You should use the default method for creating new ACIF parameters.

Tips – z/OS

OnDemand for z/OS Quick Hits

Exit precedence and the OS/390 Indexer

An input exit (INPEXIT) cannot be used if you are using an AnyStore exit (ANYEXIT). You also cannot use an input exit if the index style is AFP (INDEXSTYLE=AFP).

There are no restrictions on the use of an index exit (INDXEXIT); it can be used under any circumstances. It can be used when either an input exit or an AnyStore is used. The index exit is called "last", just before the index values are stored, and after the document is extracted from the load file (either by the base program, the input exit, or the AnyStore exit.) The purpose of the index exit is to modify the index values before they are loaded into the database.

For more information see support item 1681258.

Content Manager OnDemand and SYSPLEX

Does OnDemand exploit a parallel sysplex? The short answer is, while nothing from OnDemand gets placed into the Coupling Facility, there are benefits to running OnDemand in a SYSPLEX environment. The benefits are discussed in support item 7037914, titled "Best Practices for Implementing IBM Content Manager OnDemand for z/OS in a Parallel Sysplex."

One caveat: When upgrading, whether to a new version or new fixpack (PTF) level, all ARSSOCKD started tasks sharing a back end must be upgraded at the same time. Mixing levels of OnDemand in a shared environment is not supported.

OS/390 Indexer enhanced

This tip applies to the OS/390 Indexer on both z/OS and AIX.

Version 9.5 of OnDemand adds several enhancements to the OS/390 Indexer.

FIELD indexing parameter

The FIELD indexing parameter now supports BASE=TRIGGER. With BASE=TRIGGER support, the column of the index value is located by adding the column value in the FIELD parameter to the column in which the trigger value was found. This allows the index value to float across the record, based on the trigger value which can also float across the record.

```
FIELDn=record,column,length,(TRIGGER=n,BASE={0|TRIGGER})
```

If **BASE=0**, the indexer adds zero to the field *column* offset.

If **BASE=TRIGGER**, the indexer adds the starting *column* number of the trigger string value to the field column offset.

Examples

BASE=0

```
TRIGGER1=*,1,X'F1',(TYPE=GROUP)  /* 1 */

FIELD1=0,83,8,(TRIGGER=1,BASE=0)
```

locates field values that begin in column 83

BASE=TRIGGER

```
TRIGGER2=*,*,X'E2A482A3966A38193',(TYPE=FLOAT) /* Subtotal */

FIELD2=0,10,12,(TRIGGER=2,BASE=TRIGGER)
```

locates field values that begin 10 columns after the "S" in the trigger string value of "Subtotal".

New message ARS5481I

System log messages 87 and 88 now contain message ARS5481I, which lists the indexing parameters used.

```
ARS5481I FILEFORMAT=RECORD,133
ARS5481I GROUPMAXPAGES=100
ARS5481I TRIGGER1=*,1,X'F1',(TYPE=GROUP)
ARS5481I FIELD1=0,2,4,(TRIGGER=1,BASE=0)
ARS5481I INDEX1=x'C69384F1',FIELD1,(TYPE=GROUP,BREAK=YES)
ARS5481I INDEXSTYLE=DOC
ARS5481I BREAKYES=OR
ARS5481I LINEOFFSET=ASREAD
ARS5481I INDEXSTYLE=DOC
ARS5481I CPGID=500
```

The original indexing parameters in the OnDemand application definition might be slightly different due to one or more of the following:

- All comments are removed from the parameters.

- All extraneous imbedded blanks are removed.

- The sequence numbers of the trigger, field and index parameters are changed to remove any gaps in the numbering sequence.

Large Object support

Prior to version 9.5, if you wanted the OS/390 Indexer to use the Large Object support, you

needed to include the INDEXOBJ=ALL parameter. In addition, you needed to have the Large Object check box checked on the Load Information tab of the application definition as shown below.

Starting with V9.5, you only need to have the check box checked. The INDEXOBJ parameter is now ignored by the indexer. There is no need to change or remove this parameter from existing applications.
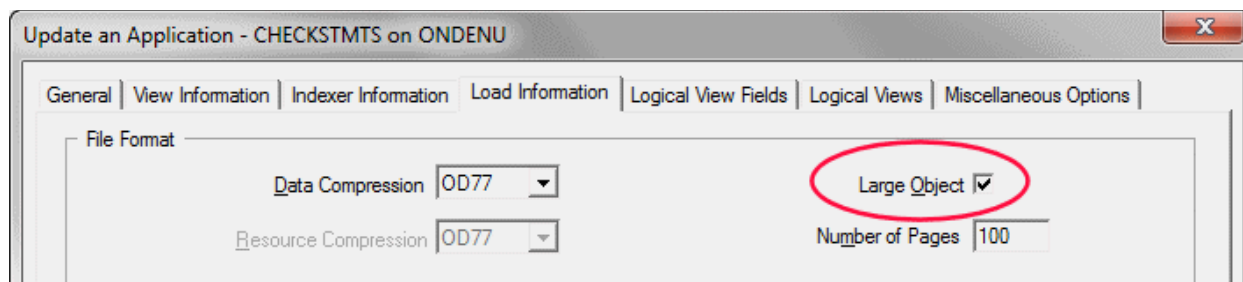
New Indexing Parameter LINEOFFSET

Version 9.5 adds a new indexing parameter, LINEOFFSET. The new LINEOFFSET parameter is used in conjunction with the FIELD parameter. It specifies how the indexer counts the records in the load file when determining the offset number of lines from the trigger record to the record that contains the index value.

```
LINEOFFSET={ASREAD | ASPRINTED}
```

ASREAD

Carriage controls are not used to calculate the record offsets for the fields. The offsets are relative to the lines as they are read from the load file. This corresponds to how the indexer worked prior to V9.5.

ASPRINTED

Carriage controls are used to calculate the record offsets for the fields. The offsets are relative to the line spacing that occurs when the lines are printed.

If this parameter is not specified, the ASREAD method is used.

Examples

If this is the data in the report file,

```
1REPORT
-ACCOUNT 777777
0JOHN SMITH
```

using ASREAD (the default), the indexing parameters to collect the account number and name are as follows:

```
TRIGGER1=*,2,'REPORT',(TYPE=GROUP)
FIELD1=1,10,6,(TRIGGER=1,BASE=0)
```

```
FIELD2=2,2,10,(TRIGGER=1,BASE=0)
```

using ASPRINTED, the indexing parameters to collect the account number and name are as follows:

```
TRIGGER1=*,2,'REPORT',(TYPE=GROUP)

FIELD1=3,10,6,(TRIGGER=1,BASE=0)

FIELD2=5,2,10,(TRIGGER=1,BASE=0)
```

Supported carriage controls used by ASPRINTED

The following table lists the carriage control characters recognized by the LINEOFFSET= ASPRINTED parameter. Both ANSI and Machine carriage control are supported. Any carriage control values found in the load file that are not shown in this table are treated, for the purposes of the ASPRINTED support, as if they are the "space 1 line" character.

Action	ANSI Act, then print	Machine Print, then act	Machine Act without printing
Overprint	+	x'01'	
Space 1 Line	blank	x'09'	
Space 2 Lines	0	x'11'	
Space 3 Lines	-	x'19'	
Skip to Channel 1	1		x'8B'

Notes:

- When the OS/390 Indexer encounters a record in the load file containing a x'89' carriage control character (skip to channel 1 for "print, then act"), the indexer changes the carriage control character to a x'09'. It then inserts a new record containing only a x'8B' carriage control character. (The OS/390 Indexer has always worked this way.)

- For purposes of determining the record offset values for both LINEOFFSET=ASREAD and ASPRINTED, the x'8B' carriage control character is treated as a counted line.

TIPS — IBM I

Save and restore of disk pools used by OnDemand

The simplest approach to make sure that the disk pool for ASP1 is saved and restored properly for instance QUSROND is to do the following:

To perform the save:

1. Unmount the file system

   ```
   CALL QRDARS/QRLCASMUFS PARM('QUSROND')
   ```

2. Save the directory structure

   ```
   /dev/QASP01/ONDEMAND_QUSROND*
   ```

3. Save the directory structure

   ```
   /QIBM/UserData/OnDemand/QUSROND
   ```

To perform the restore:

1. Unmount the file system

   ```
   CALL QRDARS/QRLCASMUFS PARM('QUSROND')
   ```

2. Restore the directory structure

   ```
   /dev/QASP01/ONDEMAND_QUSROND*
   ```

3. Restore the directory structure

   ```
   /QIBM/UserData/OnDemand/QUSROND
   ```

The file system will be remounted automatically the first time the OnDemand server attempts to access it.

Note: The instance name (QUSROND) and Auxiliary Storage Pool (ASP) name (QASP01) in the examples must be changed to match your actual environment.

Additional Information

Knowledge Centers

OnDemand for Multiplatforms Knowledge Centers

OnDemand for z/OS Knowledge Centers

OnDemand for i Knowledge Centers

Publication Libraries - Containing all PDF versions of the documentation

OnDemand for Multiplatforms	Version 8.5	Version 9.0	Version 9.5
OnDemand for z/OS	Version 8.5	Version 9.0	Version 9.5
OnDemand for i	Version 6.1	Version 7.1	Version 7.2

Product System Requirements

OnDemand for Multiplatforms	Version 8.5	Version 9.0	Version 9.5
OnDemand for z/OS	Version 8.5	Version 9.0	Version 9.5
OnDemand for i	Version 6.1	Version 7.1	Version 7.2

More Enterprise Content Management Web Sites

Content Manager OnDemand Product Overview

Compatibility Matrix for the OnDemand Clients and Servers

Content Navigator Product Overview

Social Media

IBM OnDemand is on social media, including developerWorks, Facebook, and the IBM Client Success Essentials wiki. For more information and links to all social media sites, see support item 1673183.

OnDemand User Group

The primary objective of the OnDemand User Group (ODUG) is to create an environment and network encouraging the exchange and development of information regarding Content Manager OnDemand and its associated products.

Copyright and Trademark Information

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

IBM, the IBM logo and ibm.com are 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 IBM trademarks is available on the Web at "Copyright and trademark information".