IBM Content Manager OnDemand Newsletter

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 is published quarterly.

Previous editions of this newsletter can be found in support item 7024130. They are also available on the OnDemand User Group web site under the heading 'Presentations, Newsletters, and such'.

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

We've changed the format of this newsletter so that it is easier to read on wide screen devices. Use the full screen viewing option in Acrobat (Ctrl+L) for best results.

Content Manager OnDemand for i V7.3 available

Content Manager OnDemand for i Version 7.3 was announced and became available in April. See announcement letters 216-216 (US), ZP16-0284 (EMEA), and AP16-204 (Asia Pacific).

Highlights of version 7.3 include:

- OnDemand Distribution Facility
- New, native, IBM i commands to query, print, and retrieve archived data
- OS/400 indexer enhancements
- Command enhancements
- IBM Navigator for i enhancements

For a list of all the enhancements, see the Knowledge Center topic What's New in Content Manager OnDemand for i Version 7.3. For more details on selected enhancements, see the IBM i tips section of this newsletter.

Feature pack 9.5.0.5 available

Multiplatforms

The 9.5.0.5 feature pack (formerly known as fix pack) installation files are available from IBM Fix Central. This includes the OnDemand Administrator and OnDemand Windows client. Support item 1981051 provides a shortcut to download all versions of the Content Manager OnDemand clients.

z/OS

To upgrade your system, choose the applicable PTF from the list in support item 1260192.

IBM i

See the Informational APAR for the PTF numbers for your release. You should also review the corresponding Read This First before installing the PTFs.

Release	Info APAR	Read This First
V7.3	II14799	7046944
V7.2	II14723	7041959
V7.1	II14497	7018085

Feature pack 9.0.0.7 available

Multiplatforms

The 9.0.0.7 feature pack (formerly known as fix pack) installation files are available from IBM Fix Central.

z/OS

To upgrade your system, choose the applicable PTF from the list in support item 1260192.

Feature pack 8.5.0.11 available

Multiplatforms

The 8.5.0.11 feature pack (formerly known as fix pack) installation files are available from IBM Fix Central.

z/OS

To upgrade your system, choose the applicable PTF from the list in support item 1260192.

IBM World of Watson Conference

With data and analytics at the heart of the cognitive revolution, we are thrilled to announce that IBM Insight is the cornerstone of the new IBM World of Watson Conference.

IBM World of Watson is coming to the Mandalay Bay, Las Vegas, NV from October 23 through October 27.

Discover how data and analytics can help you outperform the competition. Learn how cloud computing, mobile computing, and the Internet of Things are bound with the latest advances in analytics, content, data, governance, security and systems, and how the most successful organizations in the world are capitalizing on those connections.

You'll have the opportunity to get a sneak peek at the latest Watson technologies and even get your hands in the code. Stay tuned for more details.

Registration opens June 20 on the conference website.

Content Manager OnDemand Workshops

With 180 plus registrations, the multi-city IBM Content Manager OnDemand Workshops are gaining ground. Along with both informative and how-to sessions with demos of the latest features that IBM Content Manager OnDemand V9.5 has to offer, the workshops also feature Big Data readiness for customer communication data. More details and registration are available on the workshop website.

The dates and locations of the workshops are:

- Johannesburg, South Africa – July 13 & 14
- Palisades, NY – November 8 & 9
- Toronto, Canada – November 15 & 16
- Australia – January, 2017

News *(continued)*

OnDemand for i PTF Groups

To simplify ordering PTFs, we now have PTF Groups available for Content Manager OnDemand for i. When you order the PTF group for your release, the current feature pack (formerly known as fix pack) for that release, plus any pre-requisite and co-requisite PTFs, will be downloaded to your system. PTF groups can be ordered using IBM Fix Central, or using the Send PTF Order (SNDPTFORD) command.

Release	PTF Group	Group Level	Server
V7.3	SF99252	1	9.5.0.5
V7.2	SF99251	1	9.5.0.5
V7.1	SF99250	1	9.5.0.5

Note that installing a PTF Group for Content Manager OnDemand for i will typically result in a change to the server version on your system. You should read the Read This First document for your release for important information regarding the new server version.

You should also review the Informational APAR for your release to determine if any individual PTFs in addition to the PTF Group are required. Links to the Read This First documents and Informational APARs are included in the 9.5.0.5 news item on page 1 of this newsletter.

The availability of PTF Groups does not preclude Content Manager OnDemand for i PTFs from being included in cumulative PTF packages.

Business partner 10m workshop

10m, an IBM business partner, is holding a workshop in Koblenz, Germany on September 14 & 15.

"Back to the roots" is the motto of this year's workshop. By popular demand, this workshop will again be smaller and more technology-heavy. For the first time, 10m will use their new Koblenz office. The number of people will be limited this year to a maximum of 15 participants. The workshop will be taught in German and English.

The aim of the workshop is to offer a good mix between labs, tips, and tricks - straight from the Content Manager OnDemand lab. This is made possible thanks to the collaboration between Content Manager OnDemand clients and Content Manager OnDemand specialists from IBM, and 10m. For more information, see the workshop website.

Big data, analytics, and CMOD

The circle of life in big data, analytics, cloud, and CMOD

New on the IBM ECM blog is a post by Brian Phelps, Worldwide Sales Leader and Senior Certified Consultant, CMOD Solutions. Brian discusses how your Content Manager OnDemand archive, together with analytics partners such as Crawford Technologies and Datawatch, can feed analytics engines such as Watson, Cognos, and Spark, to derive deep business insight and achieve tremendous business value and customer experience.

Turn archived information into big data and analytical assets

Many of the decision support, compliance, reporting and operational issues that businesses face are not efficiently solved with complex database query tools or data warehouses.

In this blog post, Dan Potter, Chief Marketing Officer of Datawatch Corporation, shares insights on how to turn an organization's transactional content into live data that can be used effectively for analysis, decision-making, presentation, and application integration.

CMOD – The bridge to big data analytics

In a pair of blog posts, Harvey Gross, V.P. World-Wide Sales, Crawford Technologies, Inc. explains how in today's world of analytics and big data projects, a case can be made that the original big data repository is, in fact, Content Manager OnDemand.

In the first blog post, he explains how CMOD is the archive for invoices, bills, statements and other vital business documents that are exact copies of business transactions, and that the information stored within those documents is true, highly accurate and beyond reproach.

In the second blog post, he covers how a major U. S. financial institution took advantage of the information stored within documents archived in CMOD to successfully answer a compliance audit.

CMOD education from enChoice

Whether you are a first timer with CMOD or a veteran with the product and its many features and functionalities, enChoice has courses designed just for your needs. Read the blog enChoice's IBM CMOD Education – Just Like Learning How to Ski to help you ascertain why and which course you should explore. For more information on CMOD education from enChoice, go to www.enchoice.com and click on the Services tab.

Tips – Cross Platform

Using Date (old style) field values in an SQL query string

Did you know that Date (old style) field values in an SQL query string are not required to be formatted using the internal date format?

Prior to server version 9.0, a date field value in an SQL query string had to be formatted using the internal date format. This required using the ARSDATE program to convert the date to a number. For example, the date 3/27/2013 is represented as the number 15792.

The ARSDATE syntax for this conversion is:

```
arsdate 3/27/13
3/27/2013 ->        15792
arsdate 3/27/16
3/27/2016 ->        16888
```

Below is an example of the ARSDOC command with a query string using internal dates 15792 and 16888 to retrieve documents.

```
arsdoc get -h odserver -u oduser -i "WHERE account='000-000-
    000' and crd_date BETWEEN 15792 AND 16888" -f "Credit
    Card Statements" -gcNa -o datafile -v
```

Beginning at server version 9.0, a date can be specified in the SQL query string in a specific format, enclosed in single quotes (i.e. 'YYYY-MM-DD'). The name of the date database field must be preceded by "ODDAT_" (i.e. ODDAT_crd_date). Specifying the date in this format is supported for both "Date (old style)" database fields and "Date" database fields. If you have "Date/Time (old style)" or "Date/Time" fields, the required format is 'YYYY-MM-DD HH:MM:SS'.

Below is an example of the ARSDOC command with a query string using dates 3/27/2013 and 3/27/2016 to retrieve documents.

```
arsdoc get -h odserver -u oduser -i "WHERE account='000-000-
    000' and ODDAT_crd_date BETWEEN '2013-03-27' AND '2016-
    03-27'" -f "Credit Card Statements" -gcNa -o datafile -v
```

Regardless of whether the application group has old date types, new date types, or has been updated from old to new date types, the Content Manager OnDemand server modifies the SQL such that it maps to the correct value(s).

ARSDOC GET output file options

When using ARSDOC GET, documents can either be retrieved and stored together in one file or stored individually in a separate file for each document.

The -o parameter is used to specify the name of the output file. If a separate file is created for each document, a numeric extension can be added to the name of the output file so that files with unique names are created.

As an alternative to simply specifying an output file name with the -o parameter, an output file name can instead be specified by providing one or more application group database field names enclosed in parentheses. In Example 1 below, a search is performed for documents that have an account number of 001200340056. Documents that match the search criteria will be retrieved and stored in separate files. In this example, since the database fields "account" and "name" will have the same value for each document, the documents will be written to the same file. To ensure each document is written to a separate file, a database field name needs to be provided that will result in unique file names. In Example 2, "sdate" is also specified. In this case, a separate file will be created for each document (assuming only one document per account number and name is available for each date in the database).

In some situations, it might not be possible to ensure unique file names because there isn't a database value that's unique for the specified query. To resolve this, an optional extension can be added that will ensure unique file names. In Example 3, the database field names are followed by ".SEQ". When ".SEQ" is specified, a numeric extension is added to the name of the file. For example, if the account number is '001200340056' and the name is 'JONES', one or more files will be created with file names: 001200340056.JONES.1, 001200340056.JONES.2, ... 001200340056.JONES.n.

Example 1:

```
arsdoc get -h ARCHIVE -u sysadmin -f "Student Information" -o
    "(account)(name)" -v -i "WHERE id='001200340056'"
```

Example 2:

```
arsdoc get -h ARCHIVE -u sysadmin -f "Student Information" -o
    "(sdate)(account)(name)" -v -i "WHERE id='001200340056'"
```
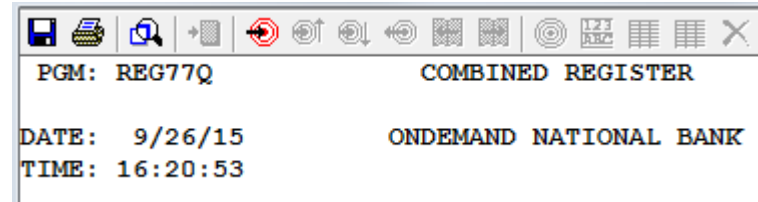
Example 3:

```
arsdoc get -h ARCHIVE -u sysadmin -f "Student Information" -o
    "(account)(name).SEQ" -v -i "WHERE id='001200340056'"
```

Indexing date and time fields

Question

I have been using separate date and time fields in my report definitions because the date and time are located on different lines on the report.

Now that the time field type is no longer available for new report definitions, how can I capture the time from the report?

Answer

You can capture the date in one field, the time in another field, and combine both fields into one index defined as a Date/Time or Date/Time (TZ) field type. This technique will work with the OS/390 and OS/400 indexers, and, with some limitations, ACIF, and the PDF indexer.

When using ACIF or the PDF indexer, both the date and time fields must be located using group triggers. Fields located using float triggers cannot be combined into a single index when using these indexers.
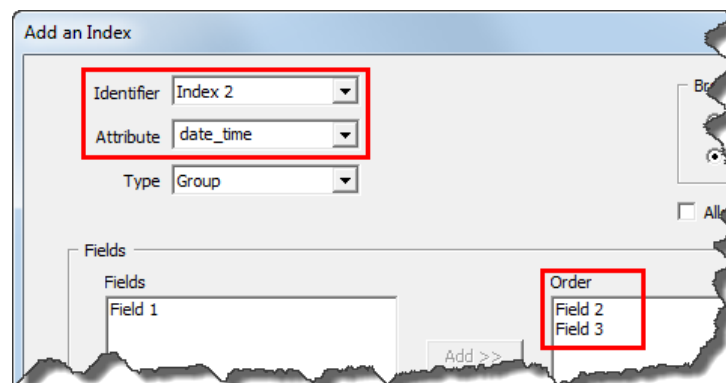
Graphical markup

Using our sample data, in the graphical indexer we mark up one trigger and three fields, as shown below:
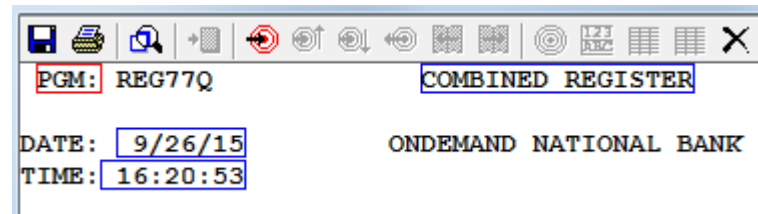
The extra space at the beginning of the time field provides a separator character between the combined date and time.

If you are not using the Report Wizard, you can specify two fields for one index while using the graphical indexer, as shown below.

If you are using the Report Wizard, you will have to make manual changes after you create the report definitions.

Indexer parameters

In our final indexer parameters, the field capturing the date (FIELD2) and the field capturing the time (FIELD3) are combined into one index (INDEX2).

```
TRIGGER1=*,3,X'D7C7D47A',(TYPE=GROUP)
    /* PGM:         */
FIELD1=0,27,17,(TRIGGER=1,BASE=0)
FIELD2=1,8,8,(TRIGGER=1,BASE=0)
FIELD3=2,7,9,(TRIGGER=1,BASE=0)
INDEX1=X'9985979699A36D95819485',FIELD1,(TYPE=GROUP,BREAK=NO)
    /* report_name */
INDEX2=X'8481A3856DA3899485',FIELD2,FIELD3,(TYPE=GROUP,BREAK=NO)
    /* date_time   */
```

Application changes

On the Application, Load Information tab, you should verify that the Format specified matches the data captured from the report. You should also specify to remove leading blanks.
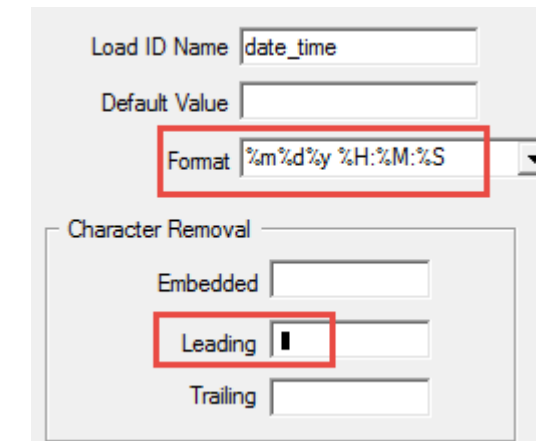
End result

In the OnDemand client document list, both the date and time are displayed.

Tips – z/OS

Was a document stored as a "large object"?

To determine if a document was stored using large object support, check for a $ at the end of the document name field (DOC_NAME) in the application group data table. The $ means the document was stored as a large object document. For example, if the instance name is ARCHIVE and the application group data table name is PIA1, you could use the following SQL statement:

```
SELECT DOC_NAME FROM ARCHIVE.PIA1
```

The results might be similar to the following:

```
DOC_NAME
34FAAA
35FAAA$ <--- large object
36FAAA
```

Add documents using Content Navigator

Together with IBM Content Navigator version 2.0.3.6 or higher and IBM Content Manager OnDemand version 9.5.0.4 or higher, users can now add documents using IBM Content Navigator. This is in addition to adding documents using the traditional batch processing methods using the ARSLOAD program (all platforms) or Start Monitor command (IBM i only).

Support for CMIS write interface

Previous versions of IBM Content Manager OnDemand offered support for only "read" interfaces of the Content Management Interoperability Services (CMIS) standard. Now, with version 9.5.0.4 or higher, IBM Content Manager OnDemand supports "write" interfaces as well, enabling CMIS-compliant programs to add content to IBM Content Manager OnDemand as well as read content from IBM Content Manager OnDemand.

Running multiple ARSSOCKDs on the same LPAR

Recently, several customers have requested instructions for setting up two started tasks accessing the same database so that they could split the load between the two. The following examples show how to accomplish this.

First, add a new server stanza to the ARS.INI file. Here is what that file might look like:

Existing server stanza

```
[@SRV@_ABC900]
HOST=mvs222
PROTOCOL=2
PORT=1451
SRVR_INSTANCE=ABC900DB
SRVR_INSTANCE_OWNER=ODPROD
SRVR_OD_CFG=/etc/ars/V900/ars.abc900.cfg
SRVR_SM_CFG=/etc/ars/V900/ars.abc900.cache
SRVR_FLAGS_FOLDER_APPLGRP_EXIT=0
SRVR_FLAGS_SECURITY_EXIT=0
SRVR_FLAGS_FORCE_SECURITY=0
```

New server stanza

```
[@SRV@_ABC902]    <------ this is different
HOST=mvs222
PROTOCOL=2
PORT=2451         <------ this is different
SRVR_INSTANCE=ABC900DB
SRVR_INSTANCE_OWNER=ODPROD
SRVR_OD_CFG=/etc/ars/V900/ars.abc900.cfg
SRVR_SM_CFG=/etc/ars/V900/ars.abc900.cache
SRVR_FLAGS_FOLDER_APPLGRP_EXIT=0
SRVR_FLAGS_SECURITY_EXIT=0
SRVR_FLAGS_FORCE_SECURITY=0
```

Note, only the server name and port are different.

Next, create two started tasks, which you can name anything you choose.

Original started task

```
//ARSSOCK0 PROC
//ARSSOCK0 EXEC PGM=ARSSOCKD,REGION=0M,TIME=NOLIMIT,
//* PARM='/-S -I ABC900 -v'
...
```

New started task

```
//ARSSOCK2 PROC
//ARSSOCK2 EXEC PGM=ARSSOCKD,REGION=0M,TIME=NOLIMIT,
//* PARM='/-S -I ABC902 -v'    <------ the only difference
...
```

Start both of the tasks. Using the port number, sign on to each of them with a client to see that you can independently use each one. Submit a simple batch job, changing only the instance name from ABC900 to ABC902.

Here's a sample job you can submit. It displays a list of the threads associated with the server task.

```
//   JOB
//TMP1  EXEC  PGM=IKJEFT01,REGION=0M,
//          DYNAMNBR=200
//SYSPROC  DD  DSN=SYS1.SBPXEXEC,DISP=SHR
//*
//SYSTSPRT DD  SYSOUT=*
//*
//SYSTSIN  DD  *
 oshell /usr/lpp/ars/V9R0M0/bin/arssockd -I ABC900 -px
//*
```

By using the -I (or -h) parameter for your load jobs, you can direct loading to one started task or the other as desired.

MidServer does not work with Java 8

This tip is for customers using the Content Manager OnDemand for z/OS MidServer.

The introduction of Java 8 deprecated some previously available functions. If you try to start the MidServer with Java 8, you will see the following in STDERR:

```
Exception in thread "main" java.lang.NoClassDefFoundError:
    sun.io.Converters
...
Caused by: java.lang.ClassNotFoundException:
    sun.io.Converters
```

Customers using the MidServer should continue to use any of the previous supported levels of Java. Information on z/OS support for Java can be found online.

Content Manager OnDemand development recommends using the 64-bit Java versions.

Changes in ARSLOAD -E processing

Starting at version 9.5, for Content Manager OnDemand for z/OS, the behavior of ARSLOAD –E behavior has been modified such that it is now consistent across all the CMOD platforms that support the –E parameter.

When specified, the -E parameter causes ARSLOAD to execute the ARSUUPDT user exit program. ARSUUPDT is the report specifications archive definition exit, which provides the ability to alter some of the parameters used by Content Manager OnDemand when document data is being captured (loaded) by the ARSLOAD program.

At server version 9.5.0.x and higher, when using the -E parameter, the filename resolution process for ARSLOAD is illustrated in the following examples:

Example 1:
```
arsload -h ARCHIVE -E … (more parameters) …
      /tmp/IBMCM.ODKREL.D2I
```
1) Looks for filename IBMCM.ODKREL.D2I.
2) Appends .IND to the filename, so it looks for IBMCM.ODKREL.D2I.ind.
3) Reads the .IND file and looks for the filename specified in the GROUP_FILENAME entry, for example IBMCM.ODKREL.D2I.out. So it looks for filename IBMCM.ODKREL.D2I.out.
4) If any of the files do not exist, then the ARSLOAD process fails.

Example 2:
```
arsload -h ARCHIVE -E … (more parameters) … /tmp/REGISTER
```
1) Looks for filename REGISTER.
2) Appends .IND to the filename, so it looks for REGISTER.ind.
3) Reads the .IND file and looks for the filename specified in the GROUP_FILENAME entry, for example KEY3REG.TXT. So it looks for filename. KEY3REG.TXT.
4) If any of the files do not exist, then the ARSLOAD process fails.

Example 3:
```
arsload -h ARCHIVE -E … (more parameters) … /tmp/REGISTER
```
1) Looks for filename REGISTER.
2) Appends .IND to the filename, so it looks for REGISTER.ind.
3) Reads the .IND file and looks for the filename specified in the GROUP_FILENAME entry, for example REGISTER. So it looks for filename REGISTER.
4) If any of the files do not exist, then the ARSLOAD process fails.

Notes:
- In examples 1 and 2, the files IBMCM.ODKREL.D2I.out. and KEY3REG.TXT contain the data to be loaded.
- In Example 3, the file REGISTER contains the data to be loaded.

- In Examples 1 and 2, unless the ARSUUPDT user exit program reads the files IBMCM.ODKREL.D2I and REGISTER, then from an ARSLOAD perspective, these files do not need to contain any data. In other words, a touch of "filename" to create the file is sufficient.

On z/OS, at server versions 9.0.0.x and lower, the base filename (i.e. IBMCM.ODKREL.D2I in example 1, or REGISTER in example 2) did not need to exist.

For more information on ARSLOAD, and on loading generic indexer files, see the Content Manager OnDemand for z/OS Knowledge Center.

Tips – Multiplatforms

Windows TCP/IP tuning

A known Windows configuration setting might affect performance when you connect to a Content Manager OnDemand server. During repeated searches and retrievals on a Content Manager OnDemand server, many Windows sockets are opened and closed. Other Content Manager OnDemand functions that typically open and close many Windows sockets are ARSLOAD and ARSXML. There are two default Windows settings that might affect heavy traffic between the client and the Content Manager OnDemand server.

- When an application closes a Windows socket, Windows places the socket's port into TIME_WAIT status for 240 seconds. During this time, the port cannot be reused.

- Windows limits the number of ports that an application can use to 5000.

To avoid the problems that might result from running out of sockets, change the values for the timeout wait time and number of ports by editing the Windows registry.

Important: This solution includes changing the Windows registry with Windows registry editor (regedit). It is strongly recommended that you back up the Windows registry before you make any changes.

- Reduce the timeout wait time for closed sockets. Add registry key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TcpTimedWaitDelay` and set its value to a smaller number than the default of 240, for example, 40. The valid range is 30-300.

- Increase the maximum port available to applications. Add registry key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\MaxUserPort` and set its value to a larger number than the default of 5000, example, 65000. The valid range is 5000-65534.

After making these changes, restart the Windows system for the changes to take effect. For more information about these two parameters, you can search for TcpTimedWaitDelay and MaxUserPort on Microsoft Technet.

Tips – Multiplatforms & IBM i

Managing PDF resources

A problem was recently discovered when retrieving very large storage objects. Content Manager OnDemand is optimized for retrieving objects of 10MB or less. The customer that originally reported the problem was retrieving PDF documents of approximately 10KB or less and then a PDF resource file that was 300MB or larger.

A fix to the retrieval problem was made at server version 9.5.0.3.

In addition to installing server version 9.5.0.3 or higher, customers should review the size of the resource files created by the PDF indexer. A very large resource file size can significantly impact overall performance because that resource file is retrieved repeatedly.

Removing resources

The overall size of the PDF resource file can be greatly reduced by using the indexer parameter REMOVERES=YES. The REMOVERES parameter indicates whether or not to remove unused resources before the indexer collects resources and creates the indexes. The input file is examined and a new copy is saved in the Content Manager OnDemand temporary directory. This new copy is then used for processing, and the original input file is not changed. If a file contains many unused resources, you can greatly reduce the size of the resource file and speed up the indexing process by using this parameter. If a file does not contain any unused resources, then do not specify this parameter. You can use this parameter without resource collection. The REMOVERES parameter is available at server version 8.4.1 and higher.

In the original problem scenario, the PDF resource file was 30,000 times the size of the document data. After specifying the REMOVERES=YES parameter, the resource file size was reduced by over 500%.

Optimizing resources

The size of the PDF resource file can also be reduced by managing the resources (fonts and images) embedded when the PDF document is created.

The base 14 fonts are not embedded in PDF documents because their font metrics and suitable substitution fonts are guaranteed to be available to Acrobat viewer programs. Using the base 14 fonts will result in smaller PDF resource files.

Minimizing the size and number of images embedded in the PDF document will also reduce the input file size and the resource file size.

Base 14 Fonts			
Courier	Helvetica	Times-Roman	Symbol
Courier-Bold	Helvetica-Bold	Times-Bold	ZapfDingbats
Courier-BoldOblique	Helvetica-BoldOblique	Times-Italic	
Courier-Oblique	Helvetica-Oblique	Times-BoldItalic	

Tips – IBM i

Considerations when upgrading to V7.3

IBM Content Manager OnDemand for i upgrade requirements

If you are upgrading from a previous version of IBM Content Manager OnDemand, you must be running Content Manager OnDemand server version 8.4.1.3 (or higher) before you upgrade to Content Manager OnDemand for i V7.3. For instructions on how to determine your current server version, see the Content Manager OnDemand for i: Common Server Planning and Installation Guide.

Content Manager OnDemand client software must be upgraded to version 8.5.0.5 (or higher) before you upgrade to Content Manager OnDemand for i V7.3. This includes, but is not limited to, the OnDemand Windows (end-user) client and the ODWEK CGI, Servlet, and Java APIs. If you use IBM Content Navigator (ICN), see the appropriate documentation to determine minimum software requirements. WEBi and II4C are not supported at V7.3. You should upgrade to IBM Content Navigator (ICN) either before or at the same time you upgrade to V7.3.

The OnDemand Administrator client must be at the same version or higher as the Content Manager OnDemand server. For V7.3, the OnDemand Administrator client must be at version 9.5.0.4 or higher.

See support item 1515197 for important server upgrade considerations.

Warning message requires reply when you start the first Content Manager OnDemand instance server

When you start your first Content Manager OnDemand instance (by using the Start TCP/IP Server (STRTCPSVR) command, or by calling program QRLMCTL) after upgrading to V7.3, you might receive message OND0156 in the QSYSOPR message queue that requires a response. The message is sent if you are upgrading to V7.3 from a Content Manager OnDemand server version prior to 9.5.0.4. For example, if you are running V7.2 with server version 9.0.0.3 and you upgrade to V7.3 which runs server version 9.5.0.4, you will receive the message: *The Content Manager OnDemand server will be upgraded to server version 9504. (C G).*

If you have reviewed the Content Manager OnDemand for i Read This First document for V7.3 and are ready to upgrade to version 9.5.0.4, reply G to Go. Your instance servers will not start until you respond with a G. If you respond with a C to Cancel, the instance server will not start and the message will continue to be sent each time you attempt to start an instance server until you respond with a G to Go.

See support item 1698611 for more information on the server upgrade message.

Command enhancements

Start Archived Storage Management – Force end

At V7.3, the Start Archived Storage Management (STRASMOND) command has a new parameter named *Force ASM to end* (ENDASM) that specifies whether you want to force the ASM process to end, and if so, when it should end. You can force ASM to end after a specified number of hours or at a specified time. If ASM completes the movement of all eligible data before the specified ending time, it will end before the specified number of hours or end time occurs. This parameter is optional.

As an example, you might use the *Force ASM to end* parameter to ensure that ASM ends before your nightly backup begins.

Start Archived Storage Management – Starting the server

A change is made at V7.3 to how the Start Archived Storage Management (STRASMOND) command handles starting the server.

ASM will start the server if it is not running in order to place a copy of the Archived Storage Management reports into the Content Manager OnDemand System Log folder after the ASM process completes. If the server is not running, and if the STRSVR parameter is *NO, the server is started for ASM, and then ended after ASM completes. If the server is running when ASM starts, it remains running when ASM completes.

This enhancement is also available at V7.1 and V7.2 at server version 9.5.0.5 and higher.

Start Monitor – Check interval

At V7.3, the Start Monitor (STRMONOND) command has a new *Check interval* (CHKITV) parameter that specifies, in seconds, how long the monitor will wait before it checks the monitored directory for a file to process. The wait time specified only applies when there are no more files to process in the monitored directory. This parameter only applies when the TYPE parameter is set to *DIR or *DIR2. If this parameter is not specified, the default wait time is 90 seconds. This parameter is optional.

This enhancement is also available at V7.1 and V7.2 at server version 9.5.0.3 and higher.

End TCP/IP Server – Unmounting file systems

A change is made at V7.3 to how mounted file systems are handled when the instance server is ended.

When the instance server is ended using the End TCP/IP Server (ENDTCPSVR) command, or by calling program QRLMCTL, all mounted file systems are unmounted. If you want the file systems to remain mounted, you must create data area QRLCUFS in the instance library.

The data area must be type *CHAR with a length of 10, and contain the value *YES. If the data area does not exist, or if the value is not *YES, the file systems will be unmounted.

This enhancement is also available at V7.1 and V7.2 at server version 9.5.0.5 and higher.

OS/400 indexer 'regular expression' support

At V7.3, the OS/400 indexer has added support for regular expressions. A regular expression (regex) is a pattern that is used to match characters in a string.

Regular expression support for triggers

The OS/400 indexer can use a regular expression in the TRIGGER parameter. The regular expression specifies the pattern for which to search.

The following trigger parameter causes the OS/400 indexer to search for a string containing four uppercase letters followed by three digits. The regular expression must match the text starting in column 10.

```
TRIGGER2=*,10,REGEX='[A-Z]{4}[0-9]{3}', (TYPE=FLOAT)
```

Another case where a regular expression might be useful is when the spooled file contains multiple languages. For example, if you want to use the word Name as a trigger, but the spooled file contains both English (Name) and Spanish (Nombre), you could use this regular expression to match either word as a trigger:

```
TRIGGER2=*,12,REGEX='[N][a,o][m][a-z]+', (TYPE=FLOAT)
```

This regular expression will match text starting in column 12 that begins with N followed by either a or o, then followed by m, and ending with one or more of the letters a through z.

Regular expression support for fields

The OS/400 indexer can use a regular expression in the FIELD parameter. The regular expression is applied to the characters which have been extracted from the field in a way similar to using a mask.

```
FIELD2=0,12,20,(TRIGGER=2,BASE=0,REGEX='[A-Z]+\s[A-Z]+')
```

In the following example, the regular expression will match a string containing two numbers, a dash, two numbers, a dash, three numbers, a dash, and a single number.

```
FIELD3=0,34,11,(TRIGGER=2,BASE=0,REGEX='[0-9]{2}-[0-9]{2}-[0-9]{3}-[0-9]')
```

(continued)

Regular expression examples

The following examples show some common regular expressions:

Pattern	Meaning
[A-Z]	Finds one uppercase letter
[A-Z]{3}	Finds three consecutive uppercase letters
[0-9]{5}	Finds five consecutive digits
[0-9]+	Finds one or more digits
[^a-z]	Finds everything except lower case a to z
\s	Finds one white space character; space, tab, and so on
\S	Finds any character except for white space

There are many online resources that explain the syntax rules of regular expressions.

More OS/400 indexer enhancements

The following additional enhancements have been made to the OS/400 indexer at V7.3.

Column range for triggers

At V7.3, you can now specify a beginning and ending column range and separate them by a colon. If you specify a column range, the beginning column cannot be zero, and the ending column must be greater than the beginning column. In the example below, trigger 1 will be searched for in columns 33 through 39.

```
TRIGGER1=*,33:39,X'D7C1C7C54040F1',(TYPE=GROUP) /* PAGE  1 */
```

Break on field value change can be AND'ed

At V7.3, the BREAKYES parameter determines how the OS/400 indexer starts a new document.

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

If the parameter value is OR, 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.

If this parameter is not specified, the default value is BREAKYES=OR, which provides the same functionality as previous releases.

New document commands

To make printing, querying, and retrieving archived documents easier, new commands have been added to Content Manager OnDemand for i at V7.3. These commands make most of the functions of the ARSDOC PRINT, QUERY, and GET qshell commands available in easy-to-use IBM i commands. The three commands are Print Document (PRTDOCOND), Query Document (QRYDOCOND), and Retrieve Document (RTVDOCOND).

The documents to be printed, queried, or retrieved are selected using one of three methods:

SQL Query String (**SQLQRY**): Specifies a valid SQL query that includes the name of one or more application group database fields, index values, and operators.

Public Named Query (**PUBNAMQRY**): Specifies the name of a public named query for the folder named with the FOLDER parameter. A named query is a set of search criteria previously saved on the server that can be recalled by name to search a folder.

Report ID (**RPTID**): Specifies the report ID of the documents to be printed. The report ID is also known as the Load ID.

All of the commands support the segment date range parameter.

Segment Date Range (**SEGDATE**): Specifies a date range that the command uses to limit a search to specific tables. When you specify this parameter, the command searches only the tables that contain a segment date within the specified date range.

Print Document (PRTDOCOND)

The Print Document (PRTDOCOND) command prints the specified Content Manager OnDemand documents to a server printer. Other parameters of note include:

Order By (**ORDBY**): Specifies which database field is used to sort the document list in ascending order. Note that this parameter cannot be specified if a report ID (RPTID) is specified.

Query Document (QRYDOCOND)

The Query Document (QRYDOCOND) command queries the specified Content Manager OnDemand documents and returns the results in a delimited stream file. Other parameters of note include:

Field Delimiter (**FLDDLM**): Specifies a one-character delimiter to use as a separator between values in the output file. By default, the values in the output are separated with a comma.

Add Document Handle (**ADDDOCHNDL**): Specifies whether or not to append the document handle to the end of each line in the output file.

Add Header Record (**ADDHDR**): Specifies whether or not to generate a header record in the output file. The header record contains the application group field names.

Add Report ID (**ADDRPTID**): Specifies whether or not to append the report ID to each line in the output file.

Order By (**ORDBY**): Specifies which database field is used to sort the document list in ascending order. Note that this parameter cannot be specified if a report ID (RPTID) is specified.

Retrieve Document (RTVDOCOND)

The Retrieve Document (RTVDOCOND) command retrieves the specified Content Manager OnDemand documents and optionally the related resources, indexes, and annotations. Other parameters of note are:

Concatenate Documents (**CONCAT**): Specifies whether or not to concatenate all of the documents that match the query into one output file.

Include Resource Group (**RSCGRP**): When retrieving AFP or XML documents, specifies whether or not to include resources with the documents that are retrieved.

Create Index File (**CRTIDXF**): Specifies whether or not to generate generic indexer data for the documents that match the selection criteria.

Retrieve Annotations (**RTVANNOT**): Use this parameter to retrieve annotations associated with the documents that are retrieved.

Add Name (**ADDNAM**): Specifies whether or not to add the resource identifier, application group name, and application name to the output file name.

IBM Navigator for i enhancements

The following enhancements have been added to the Content Manager OnDemand component of IBM Navigator for i at V7.3.

There is now a *Start current instance* button on the initial menu dialog. This means that you can easily start the instance before creating migration policies, which require the instance to be started before creating them.
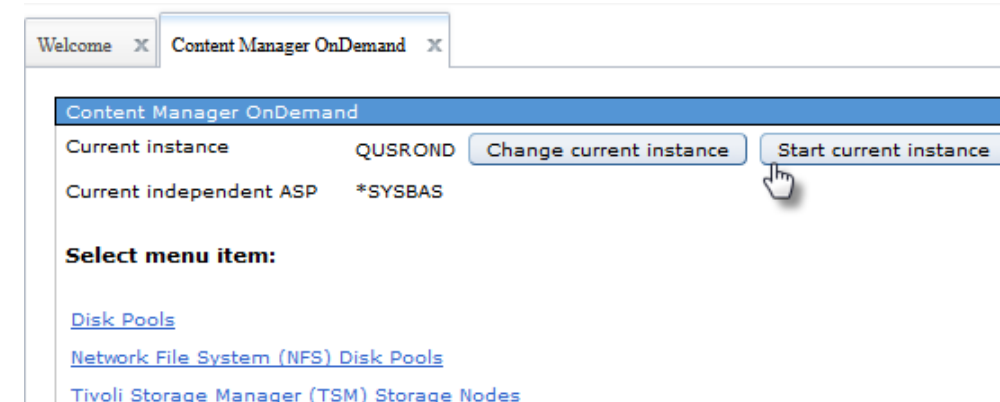
(continued)

The *Check interval* and *Password stash file* parameters have been added to the directory monitor panel. The *Check interval* parameter specifies, in seconds, how long the monitor will wait before it checks the monitored directory for a file to process. The *Password stash file* parameter specifies the path and file name of the stream file that contains the stashed user profile name and password used to run the monitor job.
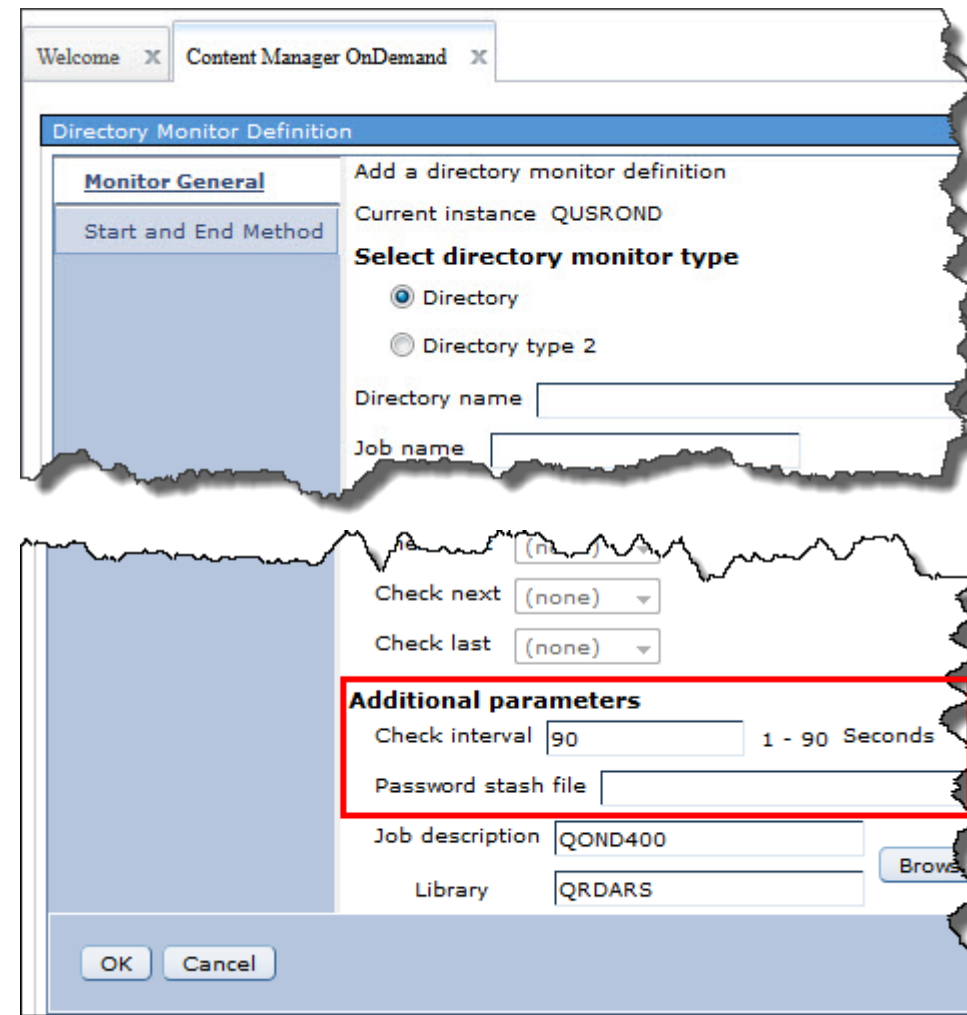
The *Password stash file* parameter has also been added to the output queue monitor panel.

ASM aggregation changes - all releases

Changes have been made to the supported aggregate sizes. The following aggregation size limits for Content Manager OnDemand migration policies have been deprecated: 100 kilobytes, 500 kilobytes, 50 megabytes and 100 megabytes.

In IBM Navigator for i, the possible values are restricted to 1, 5 or 10 megabytes on IBM i release 7.1, and 1, 5, 10 or 25 megabytes on IBM i releases 7.2 and higher. Other values supported in the past are 100 and 500 kilobytes, and 50 and 100 megabytes. The smaller values provide little or no benefit compared to not using aggregation at all, and the larger values are likely to result in individual archived files remaining in archive storage longer than specified in the Application Group definition.

After application of the PTFs that provide these changes, each time the server is started, the aggregation size limit will be forced to be at least 1MB and no larger than 10MB on V7.1, and at least 1MB and no larger than 25MB on V7.2 and V7.3.

Release	PTFs
V7.3	SI60714
	SI60705
V7.2	SI60713
	SI60704
V7.1	SI60712
	SI60703

Additional Information

Knowledge Centers

Content Manager OnDemand for Multiplatforms Knowledge Centers

Content Manager OnDemand for z/OS Knowledge Centers

Content Manager OnDemand for i Knowledge Centers

Publication Libraries - PDF versions of the documentation

Multiplatform	Version 9.0	Version 9.5	
z/OS	Version 9.0	Version 9.5	
IBM i	Version 7.1	Version 7.2	Version 7.3 - Use the Knowledge Center

Product System Requirements

Multiplatform	Version 9.0	Version 9.5	
z/OS	Version 9.0	Version 9.5	
IBM i	Version 7.1	Version 7.2	Version 7.3

More Enterprise Content Management web sites

Content Manager OnDemand Product Overview

Compatibility Matrix for the CMOD clients and servers

Content Navigator Product Overview

IBM Software Support Lifecycle Policies (search for Content Manager OnDemand)

Social Media

Follow @IBM_ECM on Twitter for almost real-time news & updates on Content Manager OnDemand. 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