ONDEMAND NEWSLETTER

NEWS AND TIPS ABOUT IBM CONTENT MANAGER ONDEMAND

1ST QUARTER 2014

IN THIS ISSUE

ADDITIONAL INFORMATION

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 newsletter should be directed to odnews@us.ibm.com.

Server Fix Pack 9.0.0.3 Available

OnDemand server fix pack 9.0.0.3 is now available for IBM Content Manager OnDemand for Multiplatforms, IBM Content Manager OnDemand for z/OS, and IBM Content Manager OnDemand for i.

OnDemand for Multiplatforms

The fix pack files are available from the IBM Fix Central site.

OnDemand for z/OS

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

OnDemand for i

See Information APAR II14497 for a list of the required PTFs.

Note that for OnDemand for i, server fix pack 9.0.0.3 is available only for v7.1.

IBM Atlas Policy Suite v6.0.3 Available

IBM Atlas Policy Suite v6.0.3 is part of the IBM Information Lifecycle Governance (ILG) solution portfolio. The ILG solutions help organizations improve information economics and reduce risk by disposing of data debris and automating key governance processes to reflect information facts.

The Records and Retention Management solution enhancements in v6.0.3 include a new retention policy distribution connector that automates the distribution of retention schedules for IBM Content Manager OnDemand.

More information is available in IBM United States Announcement 213-544.

IBM Content Navigator v2.0.22 for Mobile Available on iTunes

IBM Content Navigator v2.0.22 for Mobile is now available for download on iTunes.

This release extends the value of IBM Datacap as a mobile capture platform for iPhone and iPod Touch – including automatic and manual image detection and correction; ability to adjust colors, brightness, contrast, rotation and compression; page ordering and removal; expanded metadata and batch creation capabilities; batch monitoring; and bug fixes.

CommonStore End of Service

IBM CommonStore will reach End of Service on April 30, 2014. If you are using CommonStore, you should be planning your migration to IBM Content Collector (ICC). You can also purchase extended service if you need more time to migrate to ICC.

DB2 v9 End of Service on z/OS

DB2 v9 on z/OS will reach End of Service on June 27, 2014.

See IBM United States Withdrawal Announcement 912-009 for more information.

Draft Redbook Available: Customizing and Extending IBM Content Navigator

IBM Content Navigator provides a unified user interface for your enterprise content management solutions. It also provides a robust development platform that allows you to build customized user interfaces and applications to deliver value and an intelligent, business-centric experience.

This IBM Redbooks® publication guides you through the Content Navigator platform, its architecture, and the available programming interfaces. It describes how you can configure and customize the user interface with the administration tools provided, and how you can customize and extend Content Navigator using available development options with sample code.

This book is intended for IT architects, application designers and developers working with IBM Content Navigator and IBM ECM products. It offers both a high-level description of how to extend and customize IBM Content Navigator and more technical implementation details with sample code.

The redbook, currently available in a draft version, can be found on the redbooks web site.

Tips — Cross Platforms

Oracle Java 7 Update 25 and the Line Data Viewer Applet

Overview

Starting in Oracle Java 7 Update 25, new security requirements have been introduced. When using the OnDemand Line Data Viewer v8.5.0.8 or older and v9.0.0.2 or older applet, the following security warning will be issued from Java:

```
This application is blocked in a Java security update because the JAR file
manifest does not contain the Permissions attribute. Please contact the
Publisher for more information.
```

Resolution

The Content Manager OnDemand Line Data Viewer Applet shipped with the Content Manager OnDemand Web Enablement Kit (ODWEK) was updated to comply with the new Java security requirements. In addition, the applet is now signed using an SHA-2 certificate; an SHA-1 certificate was used previously.

To obtain the updated applet, either:

1. Upgrade to ODWEK v9.0.0.3 and use the v9.0.0.3 applet.

2. Upgrade to ODWEK v8.5.0.8 and download the updated v8.5.0.8 applet from IBM Fix Central.

   Use the following Fix Central product selections:

   Product Group: Enterprise Content Management

   Select from Enterprise Content Management: Content Manager OnDemand for Multiplatforms

   Installed Version: 8.5.0.8

   Platform: Windows 64-bit, x86

   Download fix pack: 8.5.0.8-CMODMP-ODLineDataViewer2

Note: Your ODWEK and Content Manager OnDemand Line Data Viewer Applet version and fix pack must match.

This tip is adapted from support item 1660454.

Changed Signing Requirements in the Java Runtime Environment

Abstract

Security warning messages or blocking for Java applets can occur in IBM Content Navigator, IBM FileNet Workplace XT, IBM FileNet Application Engine (Workplace), and IBM WEBi when the applications attempt to run a Java applet. This problem can occur with various versions of the Java Runtime Environment (JRE). Starting with JRE 1.7.0_51 (January 2014), applets might not run, in which case you must take certain actions to resolve the problem.

Symptom

For any Java applets in the IBM Enterprise Content Management (ECM) web clients, for example, when you open items in the viewer in IBM Content Navigator, you might see security warning messages. Starting in JRE 1.7.0_51, instead of warning messages, applets might be blocked unless you take the appropriate action.

Cause

Updates were made to the Java trust and signing requirements.

Environment

See the complete support item for details on the environments where this problem can occur.

Information for Content Manager OnDemand administrators

The Line Data Viewer Applet, which is used for viewing OnDemand documents from IBM Content Navigator, WEBi, and Workplace XT, might provide warnings with specific versions of the JRE. This viewer is provided by the OnDemand Web Enablement Kit (ODWEK); updates are provided as part of ODWEK deliverables. To resolve issues with the Line Data Viewer Applet, update your ODWEK installation to a version supported by your JRE. For example, if using JRE 1.7.0_51, upgrade to ODWEK 9.0.0.3 or download the updated files for ODWEK 8.5.0.8. For more information, see the previous tip in this newsletter.

Resolving the problem

Choose one of the following options to resolve the problem:

- Upgrade your IBM ECM web client version to a fix pack level that resolves the warnings.

- Implement an Oracle Java Rule Set, which allows an enterprise to declare certain applets as safe.

- Allow applets that have a trusted signature, for example, a valid International Business Machines Corporation (IBM) signature.

For more details on resolving this problem, and detail on unblocking a blocked Java applet, see the complete support item.

This tip is adapted from support item 1640073.

Preserving Case-switching Behavior When Migrating from OnDemand v8.4.x to v8.5 or v9.0

Question

After upgrading to OnDemand v8.5, some of my indexes store differently than they did at OnDemand v8.4.1. For example, the Eszett character is now stored as SS instead of ß.

Is there a method to instruct the OnDemand load process to use the old method of case switching of index values?

Cause

Before OnDemand v8.5, the load process used functions provided by the operating system to switch the case (upper-case to lower-case; lower-case to upper-case) of index values. These operating system functions were sensitive to the locale, and if the locale was not set correctly, the case was not switched. For example, in a US-English locale, the operating system functions would not switch a lower-case a-umlaut character to an upper-case A-umlaut character. The character remained a lower-case a-umlaut character and was saved into the DB2 database as a lower-case a-umlaut character. To ensure that the case was switched correctly, you needed to run the load process in a locale that specified both that an a-umlaut was a lower-case character, and that an A-umlaut was an upper-case character.

Beginning with OnDemand v8.5, the load process uses International Component for Unicode (ICU) functions, which are not sensitive to locale, to switch the case of characters. Therefore, a lower-case a-umlaut character was switched to an upper-case A-umlaut character, regardless of the locale. However, characters might have been produced that were not in the code page of the database. For example, the IBM-1047 code page has a lower-case y-umlaut character but no upper-case Y-umlaut character. Depending on the setting of ARS_ORIGINAL_CODEPAGE and the database code page, that nonexistent Y-umlaut was either converted to a sub character or passed as-is to DB2. If the character was passed as-is to DB2, DB2 would not save the load to the database, instead would report the SQLSTATE=01517 error code, and the load would fail.

For the Eszett character in our question, the ICU standard specifies that ß is a lower case character and that the corresponding upper case character is SS.

Answer

To retain the case switching behavior used by OnDemand prior to v8.5, specify the ARS_USE_OLD_STYLE_CASE=1 parameter in the ars.cfg file. This parameter instructs the load process to use the case-switching operating system functions that were used prior to OnDemand v8.5. If you omit the ARS_USE_OLD_STYLE_CASE parameter or set it to 0, the load process uses the ICU functions to switch the case of index values. After changing the ARS_USE_OLD_STYLE_CASE setting you must restart the server.

This tip is adapted from support item 1631713.

Local Server Failed While Accessing Application Group

Problem

I exported an application group to a local server. When I select the application group to view, update, or create a summary, I get this message: "The server failed while accessing Application Group data. Contact your system administrator."

Cause

The application group contains a database field with one or more mapping values that contain a comma in the Displayed Value of the Database Value/Displayed Value pair. These values are located on the Field Information tab of the application group. The comma is used as a delimiter to separate the Displayed Values on the local server. When the Displayed Values are processed on the local server, the values are not parsed correctly due to the extra commas.

Solution

To export and use an application group with mapping pairs on a local server, remove the commas from the Displayed Values before exporting the application group to the local server. If it is not feasible or desirable to remove the commas, the XML Batch System Administration program can be used to export the application group. The XML Batch System Administration program, called ARSXML, is a command line program (run in the QSHELL environment on IBM i) that runs on the OnDemand server and performs the same administrative tasks as the OnDemand Administrator client.

Note: In some cases, accessing an application group with a comma in the Displayed Value may result in the abnormal termination of the administrator client.

This tip is adapted from support item 1220081.

Tips – z/OS

OnDemand for z/OS Quick Hits

Enlarging a file system

Question: A small test of OnDemand was performed using only cache storage. The cache directory filled up, and even though there was room for SMS to extend the file system, the underlying zFS did not extend and an 'out of space error' was received. Why did this happen?

Answer: Out of all the operating systems that OnDemand runs on, the z/OS HFS is the only variant of a file system that allows multi-extend dynamic extension. Since the OnDemand base code is ported from AIX, the ability to force an extension doesn't exist. OnDemand currently checks to see if there is enough room in the zFS/HFS to fit the desired file in and, if not, fails the load. If some other activity (internal or external) causes the zFS/HFS to be extended prior to any potentially failing load, the space requirement will be satisfied.

Question: Working with zFS's can be a bit tricky. Is there a preferred method of extending the file system with the least possible disruption?

Answer: Yes. Under the ISPF Shell option for USS (PANELID BPXWP99) the following options are displayed across the top:

```
File  Directory  Special_file  Tools  File_systems  Options  Setup  Help
```

Tab the cursor to "File_systems" and press enter, and from the menu select "zFS aggregates..."

From that menu you can use the E option to easily extend any zFS file system.

Question: That's for zFS filesystems, what about HFS's?

a) It is time to convert to zFS's

b) Use the confighfs command and then see answer a) once again

Switching to the new PARM format

If you're on OnDemand 8.5 and above, it's time to switch to the new parm format if you haven't already.

The old format was similar to this:

```
//ARSSOCK2 EXEC PGM=ARSSOCKD,REGION=0M,TIME=NOLIMIT,
//  PARM='/VERBOSE ARCH850 ARSSOCKD'
```

The new format is similar to this:

```
//ARSSOCK2 EXEC PGM=ARSSOCKD,REGION=0M,TIME=NOLIMIT,
//  PARM='/-S -I ARCH850 -v'
```

Capital S means to start, capital I is the instance name. Lower case v means verbose. IBM recommends you always use the verbose option. It simply prints out a list of configuration information to the SYSOUT DD statement. There is absolutely no performance penalty.

Why do I receive messages about two different code pages?

In the ars.cfg I have this: ARS_ORIGINAL_CODEPAGE=37

However, in the startup messages for ARSSOCKD I see this message: ARS0220I Server code page is 1200. Why is the server code page different from the ARS_ORIGINAL_CODEPAGE?

There are actually two different code pages in use with OnDemand.

1. The DATABASE CODE PAGE, which is used by DB2.

2. The SERVER CODE PAGE, which is used by ARSSOCKD.

The server code page will always be 1200. The database code page can be any supported code page.

Preventive Service Planning

You should check the PSP "buckets" (as they are called) on a regular basis. This is the method by which development communicates HIPER (High Impact - PERvasive) PTFs to customers. In addition, the PSP buckets also include important documentation changes that have not made it into a revision.

Note the section at the bottom for CROSS PRODUCT DEPENDENCIES. Whenever you plan an OnDemand upgrade, check the OnDemand PSP buckets to see if there are any items which might affect you.

As of this writing, the following item has not yet been published in the CROSS PRODUCT DEPENDENCIES section for either OnDemand version 8.5 or OnDemand version 9.0:

```
    1. YY/MM/DD INTERDEPENDENT PRODUCT: Unix Systems Services

               PROBLEM: Hang

               USERS AFFECTED: All users of USS, not just CMOD

               RECOMMENDATION: INSTALL OA43472 ON VOLID 1312
```

Note that CMOD is an acronym for Content Manager OnDemand.

Following are links to the PSPs for the last four versions of OnDemand:

V9.0 Preventive Service Planning   V8.5 Preventive Service Planning

V8.4.1 Preventive Service Planning   V8.4.0 Preventive Service Planning

When upgrading OnDemand, you should check the PSP buckets not only of the release you are upgrading to, but also the present release you are running, and all intermediate releases.

TIPS — IBM I

Associating ASM Reports with the Job

When I start the Archived Storage Management (ASM) job by using the STRASMOND command, the ASM reports are associated with the instance user profile. Can I change this so that the ASM reports are associated with the ASM job?

By default, the Archived Storage Management reports are associated with the instance user profile. The reports can be found by issuing either of the following commands:

```
WRKSPLF SELECT(instancename)
```

or simply

```
WRKSPLF instancename
```

where instancename is the name of your OnDemand instance.

If you would prefer to have the Archived Storage Management reports associated with your ASM job rather than the instance user, you can change the printer files for the ASM reports to be associated with the job by issuing the following Change Printer File commands:

```
CHGPRTF FILE(QRDARS/QPRLCASM) SPLFOWN(*JOB)
CHGPRTF FILE(QRDARS/QPRLCASMA1) SPLFOWN(*JOB)
CHGPRTF FILE(QRDARS/QPRLCASMD1) SPLFOWN(*JOB)
CHGPRTF FILE(QRDARS/QPRLCASMU1) SPLFOWN(*JOB)
```

After making the change to the ASM printer files, the reports will now be associated with the ASM job and can be located by issuing the WRKJOB command for your ASM job. For example:

```
WRKJOB ASM
```

and then select the particular ASM job to work with, or

```
WRKJOB JOB(301172/MYUSER/ASM)
```

where 301172/MYUSER is the job number of the ASM job and the user profile name that submitted the STRASMOND command.

Note: The printer files will be reset to the default value of SPLFOWN(*CURUSRPRF) when Content Manager OnDemand is upgraded to a new version. Therefore, after upgrading Content Manager OnDemand, you will need to issue the CHGPRTF commands for the ASM printer files again if you wish to preserve this change.

This tip is adapted from support item 1654001.

IBM Navigator for i PTFs

To fully apply any PTF for the IBM Navigator for i OnDemand Administration component, the HTTP administrative server must be ended and restarted. Applying the PTF copies files into an install directory that the HTTP server will look at the next time is starts. If the HTTP server is not ended and started, the PTF is not activated for IBM Navigator for i and you will not be running the current level of the code.

1. ENDHTTPSVR   SERVER(*HTTP)   HTTPSVR(*ADMIN) (This will take several minutes to end completely.)

2. Clear the log files (Suggested - This is an optional step.)

  DEL   '/QIBM/UserData/OS/OSGi/LWISysInst/admin2/lwi/logs/*'

3. STRHTTPSVR   SERVER(*HTTP)   HTTPSVR(*ADMIN) (This will take several minutes to start completely.)

How to Save and Restore a UDFS

A question that is asked on a regular basis is 'How do I save the User-Defined File Systems used by OnDemand?' This tip explains the steps necessary to save and restore User-Defined File Systems (UDFS) data.

A UDFS is a file system that can be created and managed by user. Multiple UDFSs with unique names can be created on the system. The Auxiliary Storage Pool (ASP) number is determined by the directory that contains the UDFS.

For example, '/dev/QASP01/MyUdfs1.udfs' indicates that MyUdfs1 is in the system ASP.

Relating this to the UDFS used by OnDemand, the primary disk pool in the system ASP for instance QUSROND is:

  /dev/QASP01/ONDEMAND_QUSROND_PRIMARY_01.UDFS

A user-defined file system (UDFS) can be saved when it is mounted or unmounted. Ordinarily, you should unmount user-defined file systems before save and restore operations. Menu Options 21, 22, and 23 on the SAVE menu of the operating system provide an option to unmount UDFSs before the save operation. Note: To save UDFS when it is on an independent disk pool, you should ensure that the independent disk pool is varied on and that the UDFS is unmounted before you start the save operation.

Saving an unmounted UDFS

To save an unmounted UDFS, specify the following command:

  SAV OBJ(('/dev/QASP02/udfs_name.udfs'))

The system saves information about the UDFS (for example, the ASP number, authority, and

case sensitivity).

Saving a mounted UDFS

To save a mounted UDFS, specify the following command:

```
SAV OBJ(('/directory_mounted_over')
```

Note: In this example, the UDFS was mounted over directory '/directory_mounted_over' when directory '/directory_mounted_over' was saved.

When restoring a UDFS, different commands should be used if the UDFS was saved when mounted or unmounted.

Restoring a UDFS if the backup was done with the file system unmounted

If the backup was done when file systems were unmounted (which is recommended), the UDFS should be unmounted during restore operation:

1) Check if UDFS is mounted:

```
DSPUDFS UDFS('/dev/QASPxx/MyUdfs.UDFS')
```

2) If UDFS is mounted, use command UNMOUNT for each UDFS that needs to be restored:

```
UNMOUNT TYPE(*UDFS) MNTOVRDIR('/directory_mounted_over')
```

3) Restore UDFS using the following command:

```
RST OBJ(('/dev/QASPxx/MyUdfs.UDFS'))
```

If the UDFS does not exist on the system, the system creates the block special file (*BLKSF) object that represents the UDFS. If the UDFS does exist, objects from the save media overlay objects on the system. During disaster recovery, auxiliary storage pools (ASPs) that contain the UDFS must be created before the restore operation. If ASPs is not created, the system does not restore the UDFSs.

Restoring an individual object if UDFS was unmounted during backup

Even if the UDFS was unmounted during backup, it is possible to restore individual objects.

To perform the task, the name of the object that needs to be restored should be specified. The parent directory of the new name must exist in an accessible file system.

To restore an object from the UDFS that was saved unmounted into a "normal" (for example, non-UDFS) directory, use the following command(s):

1) Create target directory:

```
CRTDIR DIR('/MyRestoreDir')
```

2) Restore object named "obj" using the following command:

```
RST OBJ(('/dev/QASPxx/MyUdfs.UDFS/obj' *INCLUDE '/MyRestoreDir/obj'))
```

To restore an object back into the UDFS, the UDFS must be mounted before the restore:

1) Check if UDFS is mounted:

```
DSPUDFS UDFS('/dev/QASPxx/MyUdfs.UDFS')
```

2) If UDFS is unmounted, use following command:

```
MOUNT TYPE(*UDFS) MFS('/dev/QASPxx/MyUdfs.UDFS')
    MNTOVRDIR('/directory_mounted_over')
```

3) If target object is already present, use command:

```
RMVLNK OBJLNK('/directory_mounted_over/obj')
```

4) Restore object named "obj" using following command:

```
RST OBJ(('dev/QASPxx/MyUdfs.UDFS/obj' *INCLUDE
    '/directory_mounted_over/obj'))
```

Restoring a UDFS if the backup was done with the file system mounted

If the backup was done when file systems were mounted, the UDFS must be mounted during restore.

If the file system is not mounted, the directories are restored to the SYSBAS instead of IASP.

1) Check if the file system is mounted using command DSPUDFS:

```
DSPUDFS UDFS('/dev/QASPxx/MyUdfs.UDFS')
```

2) If UDFS is unmounted, use following command:

```
MOUNT TYPE(*UDFS) MFS('/dev/QASPxx/MyUdfs.UDFS')
    MNTOVRDIR('/directory_mounted_over')
```

3) To restore the objects within the mounted UDFS, specify the following command:

```
RST OBJ(('/directory_mounted_over'))
```

The system restores the objects into the file system of the parent directory to which the objects are restored. The system does not restore UDFS 'dev/QASPxx/MyUdfs.udfs'.

If the UDFS is not present during the restore operation, you should use the RBDMFS(*UDFS) parameter on the RST command to rebuild the mounted UDFS during the restore operation; otherwise, only the objects contained in the directory are restored, and none of the UDFS information is restored:

```
RST OBJ(('/directory_mounted_over')) RBDMFS(*UDFS)
```

This tip is adapted from support item N1010543.

Additional Information

Additional information about IBM Content Manager OnDemand can be found at the following web sites.

Information Centers

OnDemand Version 9.0 Information Center

OnDemand Version 8.5 Information Center

OnDemand Version 8.4.1 Information Center

Publication Libraries - Containing all PDF versions of the documentation

OnDemand for Multiplatforms Version 9.0 Publication Library

OnDemand for Multiplatforms Version 8.5 Publication Library

OnDemand for z/OS Version 9.0 Publication Library

OnDemand for z/OS Version 8.5 Publication Library

OnDemand for i Version 7.1 Publication Library

Product System Requirements

OnDemand for Multiplatforms Version 9.0 System Requirements

OnDemand for z/OS Version 9.0 System Requirements

OnDemand for i Version 7.1 System Requirements

More OnDemand Web Sites

OnDemand Product Overview

OnDemand Information Roadmap

Compatibility Matrix for the OnDemand Client and Servers

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".