



ONDEMAND NEWSLETTER

NEWS AND TIPS ABOUT IBM CONTENT MANAGER ONDEMAND

1ST QUARTER 2013

IN THIS ISSUE

NEWS

- About This Newsletter 1
- Server Fix Pack 9.0.0.1 Available 1
- OnDemand Information Center Updates ... 2
- IBM Optical Libraries Withdrawn from Hardware Support 2
- Server Fix Pack 8.4.1.9 PTFs Available for OnDemand for i v6.1 2
- Server Fix Pack 9.0.0.1 PTFs Available for OnDemand for i v7.1 2

TIPS – CROSS PLATFORMS

- OnDemand v8.5 No Longer Supports the v2 Generic Index File Format 3
- Defined Macros in SQL 4
- New Wizard Available to Create Query Restrictions 5
- OnDemand v9.0 New Date Types 6
- OnDemand v9.0 LDAP Authentication Process 9
- ICN Support for Graphical Annotations ... 15

TIPS – MULTIPLATFORMS

- Upgrading to DB2 v10 18
- DB2 v10.1 Support at OnDemand v9.0.0.1. 19
- Upgrading to OnDemand v9.0 from OnDemand v8.4.x 19
- Update the OnDemand Report Distribution Definitions 19

TIPS – Z/OS

- OnDemand for z/OS Quick Hits 20
- How Do the MOS and the Object Size Interrelate? 21
- Introduction to OD for z/OS Sysplex Implementation 22

TIPS – IBM I

- Field Name Restriction 24
- Considerations for High Availability 24

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

Server Fix Pack 9.0.0.1 Available

OnDemand server fix pack 9.0.0.1 is now available for IBM Content Manager OnDemand for Multiplatforms, IBM Content Manager OnDemand for z/OS, and IBM Content Manager OnDemand for i.

More information on the enhancements in version 9.0 of OnDemand is available in the [Information Center Version 9.0](#), the 4th Quarter 2012 Newsletter, and in the Tips section of this newsletter.

Note that for OnDemand for i, server fix pack 9.0.0.1 is available only for v7.1.

OnDemand Information Center Updates

The information centers for Content Manager OnDemand Versions 8.5 and 9.0 were updated in January with corrections and minor updates. These updates apply only to the English version of the information center.

The updates will not be applied to the translated versions of the information centers until the next version of the information centers is released.

The information centers are available at:

[Information Center Version 9.0](#)

[Information Center Version 8.5](#)

IBM Optical Libraries Withdrawn from Hardware Support

Effective December 31, 2012, all IBM 3995 and 3996 models were withdrawn from hardware support.

For more information, see the following announcement letters:

United States: [911-106](#) dated June 14, 2011

Canada: [A11-0436](#) dated June 14, 2011

Asia Pacific: [WG11-0059](#) dated June 28, 2011

Server Fix Pack 8.4.1.9 PTFs Available for OnDemand for i v6.1

The following PTFs, or their superceding PTFs, are required to upgrade your v6.1 installation to server version 8.4.1.9.

Component	PTFs
Server	SI48435, SI48436
PDF Indexer	SI48652
ODWEK	SI48439

More information about OnDemand for i v6.1 is available in the [Read This First](#).

Server Fix Pack 9.0.0.1 PTFs Available for OnDemand for i v7.1

The following PTFs, or their superceding PTFs, are required to upgrade your v7.1 installation to server version 9.0.0.1.

Component	PTFs
Server	SI48916, SI49233, SI49234, SI49235
PDF Indexer	SI49343
ODWEK	SI49261
WebNav	SI48807

IBM Navigator for i (WebNav) is the web based alternative for System i Navigator. More information about the integration of OnDemand for i and IBM Navigator for i will be included in the next newsletter.

More information about OnDemand for i v7.1 is available in the [Read This First](#).

TIPS – CROSS PLATFORMS

OnDemand v8.5 No Longer Supports the v2 Generic Index File Format

Problem

After upgrading to OnDemand v8.5, my generic indexer files will not load.

For example, using the original generic index file format you might have these indexes:

```
FIELD NAMES BEGIN:
date
title
userID
FIELD NAMES END:
27.01.2010
End User Guide to Electronics
db50601
/db50601/uprirucka.doc
0
0
```

Answer

The original Generic Index file format introduced with OnDemand in v2 is no longer supported in v8.5 or higher. The original Generic Index file format was replaced in v7.1 with the current Generic Index file format documented in the Indexing Reference manual. Using the current Generic Index file format you might have these indexes:

```
CODEPAGE:00870
GROUP_FIELD_NAME:date
GROUP_FIELD_VALUE:27.01.2010
GROUP_FIELD_NAME:title
GROUP_FIELD_VALUE:End User Guide to Electronics
GROUP_FIELD_NAME:userID
GROUP_FIELD_VALUE:db50601
GROUP_OFFSET:00000000000000
GROUP_LENGTH:00000000000000
GROUP_FILENAME: /db50601/uprirucka.doc
```

Although later versions of OnDemand accepted input files that contained the original Generic Index file format in the past, this is no longer the case with v8.5. Users must ensure that applications are updated to produce the current supported Generic Index file format in order to load their data into OnDemand v8.5.

See the Generic indexer reference in the Indexing section of the OnDemand v8.5 Indexing Reference for more information.

This tip is adapted from [support item 1504191](#).

Defined Macros in SQL

Question

How can I write SQL statements to include current object values like application group name or user ID?

Cause

OnDemand v9.0 added macro support that can be used in SQL statements (including the query restriction) which allows the macro to be substituted by the appropriate value.

Answer

The new macros are:

Name	Description
\$ODUSERID	The userid used to log into OnDemand
\$ODALIAS	The alias defined to OnDemand for the user's session
\$ODAGNAME	The application group name
\$ODAGID	The application group internal identifier

The substitution does not include any necessary quotes for the macro, therefore you must ensure that you properly quote the macro, if required.

For instance, if you logon to OnDemand as USER1, this SQL

```
WHERE ag_field in (SELECT value FROM <customer_table> where userid =  
    '$ODUSERID')
```

becomes

```
WHERE ag_field in (SELECT value FROM <customer_table> where userid =  
    'USER1')
```

This tip is adapted from [support item 1611495](#).

New Wizard Available to Create Query Restrictions

This tip applies to the OnDemand v9.0 Administrator client when used with either server version 9.0 or server version 8.5.

To improve the usability of entering query restrictions, a new option has been added to the v9.0 Administrator client that allows the user to build the SQL string using a graphical user interface (GUI). The Build SQL option is available in the application group definition, on the Permissions tab.

The GUI provides the names of the database fields, the allowed operators, an example of a valid value, and validation of the string with error messages if an invalid string is generated.

The four steps to enter a simple query restriction are:

1. Select a field, click on Insert Field Name.
2. Select a symbol, click on Insert Symbol.
3. Enter the values for the query restriction.
4. Click OK.

The screenshot shows the 'Define Query Restriction' dialog box. It has a title bar with a close button. The main area is divided into three columns: 'Fields', 'Symbols', and 'SQL'. The 'Fields' column has a list of database fields: ACCTNAME, ACCTNO (selected), ENDBAL, ENDDATE, TAXID, and VER. A red '1' is next to the list. Below the list is a text box with 'ACCTNO - String - Example: 'abc'' and an 'Insert Field Name' button. The 'Symbols' column has a list of operators: =, <>, <, <=, >, >=, BETWEEN (selected), NOT BETWEEN, IN, NOT IN, LIKE, NOT LIKE, AND, and OR. A red '2' is next to the list. Below the list is an 'Insert Symbol' button. The 'SQL' column has a text box containing 'ACCTNO BETWEEN '100000' and '199999'', with a red '3' next to it. Below the text box are 'Clear' and 'Restore' buttons. At the bottom left, there are two checkboxes: 'Include Segment Field in SQL' and 'Use BETWEEN operator'. Below these are two empty text boxes. At the bottom right, there is an 'Error Description' text box. At the very bottom, there are 'OK', 'Cancel', and 'Help' buttons, with a red '4' next to the 'OK' button.

OnDemand v9.0 New Date Types

In OnDemand v9.0, new date types have been introduced to allow for a broader range of values supported by the underlying database for Date, Date/Time and Date/Time (TZ).

In order to use the new types, *you must be using OnDemand v9.0 or later of the OnDemand clients*, for example: Windows® client, arsdoc, ODWEK, or CICS.

Field Names – Administrator client

In the OnDemand Administrator client, the previous types for Date, Date/Time, and Date/Time (TZ) were renamed as Date (old style), Date/Time (old style), and Date/Time (TZ) (old style). Although these types can still be used, they are no longer recommended. The Time field has been renamed Time (old style). The Time field type can not be used in new application groups.

Field Names - ARSXML

With ARSXML, the old names are not changed because that would break existing XML files or applications that generate XML files. So, the old names remain the same and the NEW names all have (native) appended.

Valid Values

The valid ranges for the Date/Time fields are as follows:

Field Type	Low Value	High Value
Date (old style)	1970-01-01	2058-12-31
Date	0001-01-01	9999-12-31
Date/Time (old style)	1970-01-01 00:00:00	2037-12-31 23:59:59
Date/Time	0001-01-01 00:00:00	9999-12-31 23:59:59
Date/Time (TZ) (old style)	1970-01-01 00:00:00	2037-12-31 23:59:59
Date/Time (TZ)	0001-01-01 00:00:00.000000	9999-12-31 23:59:59.999999

As can be seen, the new field Date/Time (TZ) now has the extended capability to store up to 6 digits of precision (microsecond) – which is now used by the OnDemand System Log and System Load application groups in order to represent more granular timing of activity on the system. In order to use these new types for the System Log/Load facility, use the -u option to the arssyscr command to update the existing application group(s). Although it is recommended that customers use these new fields, it is not required to update these application groups immediately upon installing OnDemand v9.0 – the update of the System Log and System Load can happen at any time. (Note that OnDemand for i changes the System Log and System Load automatically the first time each instance is started. You do not need to run the arssyscr command.)

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

You can create folder fields of type Date (old style) and map them to an application group field of type Date (old style) and/or Date. You can also create a folder field of type Date and map to an application group field of type Date (old style) and/or Date. It should also be noted that the only difference between the folder field types are that the minimum/maximum and default values that can be placed into those types must fall within the allowed range for that type. Same is true for Date/Time and Date/Time (old style) as well as Date/Time (TZ) and Date/Time (TZ) (old style).

What about existing application groups?

In OnDemand v9.0, the OnDemand Administrative client as well as the ARSXML command both support a way in which to update an existing application group from using old style formats to the new style formats.

In the OnDemand Administrative client, in an application group definition, on the General tab, click the Advanced button. There is an option called Create/Use new style date/time field types in place of existing old style date/time field types. Selecting this check box causes the application group to be updated to use the new application group date/time fields. This does not modify the previously defined old style field types, nor does it modify any data previously loaded to that application group.

In ARSXML, if you want to modify an existing application group with the new date/time fields, you should add an attribute for the <applicationGroup> object called createAndUseNewDates. Setting this attribute to Yes and updating the application group will result in the date fields being updated. For example, the XML might look like this:

```
<applicationGroup name="CHECKSTMTS"
  createAndUseNewDates="Yes" >
</applicationGroup>
```

However, it does cause the following actions:

1. The application group is updated with the new date/time fields. (This could include one or more depending on the application group definitions.) You will not see the additional field types defined in the application group. They are hidden to simplify the definitions. For example, if you had only a single Date value in the application group, you will continue to see only a single Date value defined using the OnDemand Administrator client, but looking at the application group data table directly, you would see two columns: one for the Date (old style) and one for the new Date field.
2. The existing application group data tables are altered such that a new column (with a value of NULL) is added for each new date/time field. The new column name takes the existing date/time database field name and adds _dt or _dt<num> to the end of the name.
3. All existing application group data tables are closed, so that no further data can be loaded into them. On z/OS and IBM i, it is possible to create an application group such that only a single application group data table is generated, however with these new fields types, the old data table is closed and a new table is created.
4. Once new data is attempted to be stored into the updated application group, a new table will be created and all new metadata values will be stored into the new date/time fields. The old style date/time field(s) will be stored with a value of -1.

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

5. Using any OnDemand v9.0 or later client guarantees that any query will transparently and appropriately be made across any application group data tables regardless of the updated field types. See the next section to understand the new OnDemand date/time query syntax.

New Query Syntax

OnDemand v9.0 introduces a new syntax, making it easier to create a query that includes date values. This new query syntax is also used to ensure that any updated application group is properly queried regardless of the existing or new date/time data in the application group. Previously, the SQL syntax for OnDemand date/time resulted in a query which used the database field name and the OnDemand internal date values.

For example, prior to OnDemand v9.0, to query for a date based on 2012/09/17, the query would look like:

```
WHERE crd_date = 15601
```

With OnDemand v9.0 or later, the same query will look like:

```
WHERE ODDAT_crd_date = '2012-09-17'
```

A query between two date ranges:

```
WHERE crd_date BETWEEN 15570 AND 15601
```

With OnDemand v9.0 or later:

```
WHERE ODDAT_crd_date BETWEEN '2012-08-17' AND '2012-09-17'
```

Using OnDemand tokenized query:

```
2;crd_date,crd_date; 15570,15601;WHERE crd_date BETWEEN ? AND ?
```

With OnDemand v9.0 or later:

```
2;ODDAT_crd_date,ODDAT_crd_date; '2012-08-17', '2012-09-17';WHERE  
ODDAT_crd_date BETWEEN ? AND ?
```

Regardless of whether the application group has old date types, new date types, or updated date types, the OnDemand server modifies the SQL such that it maps to the correct value(s). The significant change is that the SQL syntax requires that ODDAT_ is prepended to the database field name. There should always be a space between the field name, the operator and the value.

Expected value format

Date and/or Date (old style):

```
'YYYY-MM-DD'
```

Date/Time and /or Date/Time (old style):

```
'YYYY-MM-DD HH:MM:SS'
```

Date/Time (TZ) and /or Date/Time (TZ) (old style):

```
'YYYY-MM-DD HH:MM:SS.FFFFFFFF' where FFFFFFFF consists of 6 fractional digits
```


ONDEMAND NEWSLETTER - 1ST QUARTER 2013

If this syntax is not properly followed, it is possible for the SQL to get modified incorrectly by the OnDemand server which might result in either invalid SQL getting passed to the database, or a successful query that does not include all possible results. OnDemand modifies the new syntax such that any query made based on an old style date/time field is converted to the OnDemand internal date/time value. For new style date/time, OnDemand ensures the syntax matches the expectations of the underlying database.

For Date/Time (TZ) or Date/Time (TZ) (old style), the value(s) must always be in Coordinated Universal Time (UTC), sometimes known as GMT, as all values in the database are always in UTC.

This tip is adapted from [support item 7036188](#).

OnDemand v9.0 LDAP Authentication Process

Question

How does the OnDemand v9.0 LDAP authentication process work?

Answer

There are multiple logon scenarios possible when using LDAP authentication. These are described below.

Scenario #1 – OnDemand logon without LDAP authentication

In a very simplistic view, the logon process without the LDAP authentication can be described in the following steps.

1. The user types the user ID and password in the Logon to a Server panel of the OnDemand client.
2. The user ID and password are sent to the OnDemand server to be authenticated.
3. If the user ID exists in the OnDemand server database and the password is correct, the user logs on.

Note: If the user ID and password case sensitivity is disabled (the default), the user ID and password are converted to uppercase before they are validated. The user ID is converted to uppercase and then compared with the user IDs that are stored in the database. The password is converted to uppercase, encrypted, and compared with the stored version.

Scenario #2 – OnDemand logon with LDAP authentication and anonymous bind

When the OnDemand server is set to authenticate through an external LDAP server, the logon process can get a bit more complicated. Some LDAP servers are set up to allow "anonymous bind" and others are set up to allow only the "non-anonymous" bind. The term "non-anonymous" bind is explained in the next scenario.

The terminology "anonymous bind" can be a bit confusing. The bind here means the initial

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

connection to the LDAP server. An LDAP server that is set up to allow "anonymous bind" is "open" to the world. Anyone can connect to the LDAP server and search. It is like a building with its front doors unlocked. Anyone can walk in and look for a name from the directory in the lobby. In this example, even though the building is open, the individual offices might not be open. The individual offices are analogous to the entries in the LDAP server.

After a record is found, you might need to supply a password to access the information in that record (analogous to a key for a locked office). An example of an LDAP server that allows the "anonymous bind" is a company's intranet server. With this type of "open" LDAP server, the only logon information that an employee of that company would need is a user ID and password.

The anonymous bind LDAP authentication process can be described as follows:

- Connect to the LDAP server (called initial bind) without a user ID or password.
- If the initial bind is successful, search for an entry under the bind attribute name that is specified in `ARS_LDAP_BIND_ATTRIBUTE` at the location specified by `ARS_LDAP_BASE_DN`.
- If an entry is found, perform a second bind (logon) to it by using the user ID and password that was entered into the OnDemand logon panel.
- If the second bind is successful, locate the value under the mapped attribute name that is specified in `ARS_LDAP_MAPPED_ATTRIBUTE`.
- If the value is found, return it to OnDemand.

When the OnDemand LDAP authentication is configured with the anonymous bind set to true, the logon process can be described as follows:

1. The user types the user ID and password in the Logon to a Server panel of the OnDemand client.
2. The user ID and password are sent to the OnDemand server to be authenticated.
3. The OnDemand LDAP authentication component connects to the LDAP server. If the LDAP server is down or cannot be reached, the flow will be based on the setting of `ARS_LDAP_OD_AUTHORITY_FALLBACK`.

If `ARS_LDAP_OD_AUTHORITY_FALLBACK=FALSE` (default) then the logon will fail. If `ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE`, then the logon reverts back to the OnDemand logon and continues with Scenario #1, Step 2.

4. The OnDemand LDAP authentication component searches the LDAP server for the user ID that is under the attribute name specified in the `ARS_LDAP_BIND_ATTRIBUTE` configuration parameter. Note that this user ID can be an e-mail address. If the attribute name does not exist on the LDAP server, the logon to the OnDemand server fails.
5. If the attribute name is correct but the user ID does not exist on the LDAP server, then the flow will be based on the setting of `ARS_LDAP_OD_AUTHORITY_FALLBACK`.

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

6. If the user ID exists on the LDAP server and the password is wrong, the logon to the OnDemand server fails.
7. If the user ID exists on the LDAP server and the password is correct, the LDAP server returns a value to the OnDemand server. You tell the LDAP server what to return by setting the value on the configuration parameter `ARS_LDAP_MAPPED_ATTRIBUTE` to an attribute or field name that is known to the LDAP server. If this attribute name does not exist on the LDAP server, the logon to the OnDemand server fails.
8. If the attribute name is correct, a value is returned. For example, if the ID that is entered on the OnDemand logon screen is J12345, `ARS_LDAP_MAPPED_ATTRIBUTE` is set to `odid`, and the J12345 attribute `odid` is `jasson1`, then `jasson1` is returned. If the returned value matches a user ID that is stored in the OnDemand server database, the user logs on. Otherwise, the logon fails. The OnDemand password is not checked if the LDAP authentication is successful.

Note: If the user ID and password case sensitivity is disabled (the default), the user ID and password are converted to uppercase before they are sent to the LDAP server. The returned user ID string from the LDAP server is then converted to uppercase and compared with the user IDs that are stored in the database.

If the OnDemand user ID and password sensitivity is enabled, and the IDs that are stored in the OnDemand database are in uppercase (for example, as a result of case insensitivity in the past), the ID that is returned by LDAP must be in uppercase too. Otherwise, the logon will fail.

Scenario #3 – OnDemand logon with LDAP authentication and non-anonymous bind

In the previous scenario, the term "anonymous bind" was used. An LDAP server can also be set up to disallow the "anonymous bind", or to allow only the "non-anonymous bind." The Microsoft® Windows® Active Directory® (AD) server can only be set up using "non-anonymous bind.". When an LDAP server is set up this way, it is not "open" to the world. It is like a locked building that requires a key or badge to enter. In this example, the individual offices might or might not be locked. The individual offices are analogous to the entries in the LDAP server.

There are essentially two logons to an LDAP server that disallows the "anonymous bind" during the LDAP authentication process. The first logon (the initial bind) is required to gain connection to the LDAP server, and the second logon (the bind) is required to gain access to the entry or the record.

To connect to this type of LDAP server, you must supply a valid user ID and password at the initial bind time. It is like entering a locked building that requires a key or badge to enter. The initial bind ID and password are stored in the LDAP configuration under the parameters `ARS_LDAP_BIND_DN` and `ARS_LDAP_BIND_DN_PWD`, and are not entered by the user at logon time. When you are connected, you can perform the search. When you find a record, you might need to supply a password to access the record's information.

Starting in OnDemand v9.0, both parameters `ARS_LDAP_BIND_DN` and `ARS_LDAP_BIND_DN_PWD` have been moved to the instance stash file and are no longer

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

specified in the ars.cfg file, or in the Windows registry. If `ARS_LDAP_ALLOW_ANONYMOUS` is set to `FALSE`, both parameters must exist in the instance stash file or the LDAP authentication will fail. On Windows, use the OnDemand Configurator to enter or change those parameter values. On Unix, you must use the `arsstash` command to enter the values for `ARS_LDAP_BIND_DN` and `ARS_LDAP_BIND_DN_PWD`. The `arsstash` command can also be used on Windows to enter values for those parameters.

The non-anonymous bind LDAP authentication process can be described as follows:

- Connect to the LDAP server (called initial bind) with a user ID and password that are specified under the parameters `ARS_LDAP_BIND_DN` and `ARS_LDAP_BIND_DN_PW`.
- If the initial bind is successful, search for an entry under the bind attribute name that is specified in `ARS_LDAP_BIND_ATTRIBUTE` at the location specified by `ARS_LDAP_BASE_DN`.
- If an entry is found, perform a second bind (logon) to it using the user ID and password entered into the OnDemand logon panel.
- If the second bind is successful, locate the value under the mapped attribute name that is specified in `ARS_LDAP_MAPPED_ATTRIBUTE`.
- If a value is found, return it to OnDemand.

When the LDAP authentication is configured with the non-anonymous bind set to false, the logon process has one extra logon (or bind) involved (see Step 4).

1. The user types the user ID and password in the logon panel of the OnDemand client.
2. The user ID and password are sent to the OnDemand server to be authenticated.
3. The OnDemand LDAP authentication component connects to the LDAP server. If the LDAP server is down or cannot be reached, the flow will be based on the setting of `ARS_LDAP_OD_AUTHORITY_FALLBACK`.

If `ARS_LDAP_OD_AUTHORITY_FALLBACK=FALSE` (default) then the logon will fail. If `ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE`, then the logon reverts back to the OnDemand logon and continues with Scenario #1, Step 2.

4. If the LDAP server can be reached, the OnDemand LDAP authentication component connects to the LDAP server with the LDAP user ID and password that are specified in the `ARS_LDAP_BIND_DN` and `ARS_LDAP_BIND_DN_PWD` parameters. If the user ID and password are valid, the connection is established. Depending on the LDAP server setup, this user ID might need to have some administrative authority. For the Windows Active Directory server, the default is that any valid user ID that is a member of the domain can be used in this initial logon to the LDAP server. This is the first logon. If this logon fails, the logon to the OnDemand server fails.
5. After the connection is made, the OnDemand LDAP authentication component searches the

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

LDAP server for the user ID that is under the attribute name specified in the ARS_LDAP_BIND_ATTRIBUTE configuration parameter. Note that the user ID can be an e-mail address. If the attribute name does not exist on the LDAP server, the logon to the OnDemand server fails.

6. If the attribute name that is specified in the ARS_LDAP_BIND_ATTRIBUTE parameter is correct but the user ID does not exist on the LDAP server, then the flow will be based on the setting of ARS_LDAP_OD_AUTHORITY_FALLBACK. For example, if ARS_LDAP_BIND_ATTRIBUTE=mail and the user entered jscott@us.ibm.com in the OnDemand logon panel, but this e-mail address does not exist on the LDAP server.
7. If the user ID exists on the LDAP server and the password is wrong, the logon to the OnDemand server fails.
8. If the user ID exists on the LDAP server and the password is correct, the LDAP server returns a value to the OnDemand server. You tell the LDAP server what to return by setting the value on the configuration parameter ARS_LDAP_MAPPED_ATTRIBUTE to an attribute or field name that is known to the LDAP server. If this attribute name does not exist on the LDAP server, the logon to the OnDemand server fails.
9. If the attribute name is correct, a value is returned. If the returned value matches a user ID that is stored in the OnDemand server database, the user logs on. Otherwise, the logon fails. The OnDemand password is not checked if the LDAP authentication is successful.

Note: If the user ID and password case sensitivity is disabled (the default), the user ID and password are converted to uppercase before they are sent to the LDAP server. The returned user ID string from the LDAP server is then converted to uppercase and compared with the user IDs that are stored in the database.

If the OnDemand user ID that is stored is in uppercase, if the user ID case-sensitivity is enabled, and if the returned user ID is in lowercase, then the logon fails.

Other considerations

The LDAP server authentication is bypassed if one of the following conditions is true. The logon process is reverted back to the normal OnDemand logon:

The LDAP server is down or cannot be reached and
ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE

The user ID does not exist on the LDAP server for the bind attribute and
ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE

The LDAP authentication might fail if one of the following conditions is true:

- The initial bind fails because of the incorrect user ID, password, or both.
- The second bind fails because of the incorrect user ID, password, or both.
- The attribute name that is specified in the ARS_LDAP_BIND_ATTRIBUTE configuration parameter does not exist on the LDAP server.

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

- The returned value from the LDAP server does not match any existing OnDemand user ID.
- The returned value from the LDAP server matches an existing ID but it is in the wrong case and the OnDemand user ID case sensitivity has been turned on.
- When the OnDemand user ID case sensitivity option is off, the returned value is converted to the uppercase characters. If the stored OnDemand ID is in lower or mixed case, the logon can fail.

Depending on the LDAP server type and setup, the response you get from the LDAP authentication might not always be the same. The best way to diagnose the LDAP authentication problems is to get a detail trace and examine it. The trace can tell you which step went wrong in the LDAP authentication process.

Additionally, the following tools can be useful to examine the LDAP setup and to help with problem determination:

[AD Explorer](#) (freeware)

[Softerra LDAP Browser](#) (note that only the browser is freeware)

Important: The information in this document concerning non-IBM products was obtained from the suppliers of those products. IBM has not tested such products and cannot confirm the accuracy of the performance, compatibility or any other claims related to non-IBM products. Questions about the capabilities of non-IBM products should be addressed to the suppliers of those products.

OnDemand IDs exempt from the LDAP authentication

The OnDemand user ID of admin will never be subjected to the LDAP authentication. OnDemand administrators depend on this user ID to log on to OnDemand. It is, however, affected by the case-sensitivity settings just as any other user ID is.

In OnDemand v9.0, an optional ARS_LDAP_IGN_USERIDS parameter has been added and can be used to specify additional IDs that will be exempt from the LDAP authentication. On Windows, the OnDemand Configurator can be used to specify these values.

If the parameter does not exist or no value is specified, only the ID admin is exempt. Otherwise, up to 10 user IDs can be specified with commas used as delimiter. Here is an example of three user IDs that are exempt from the LDAP authentication:

```
ARS_LDAP_IGN_USERIDS=user1,user2,user3
```

The exempted user IDs should not contain any comma characters as comma is used as the delimiter on the exempted list.

More detail, including configuration parameters, sample messages, and sample configurations can be found in [support item 1597246](#).

ICN Support for Graphical Annotations

IBM Content Navigator (ICN) v2.01 now supports graphical annotations for line data when used with the OnDemand v9.0 Line Data Viewer Applet.

Requirements:

- OnDemand Server v9.0
- OnDemand Web Enablement Kit (ODWEK) v9.0
- IBM Content Navigator v2.0.1

What's New

Graphical annotations include:

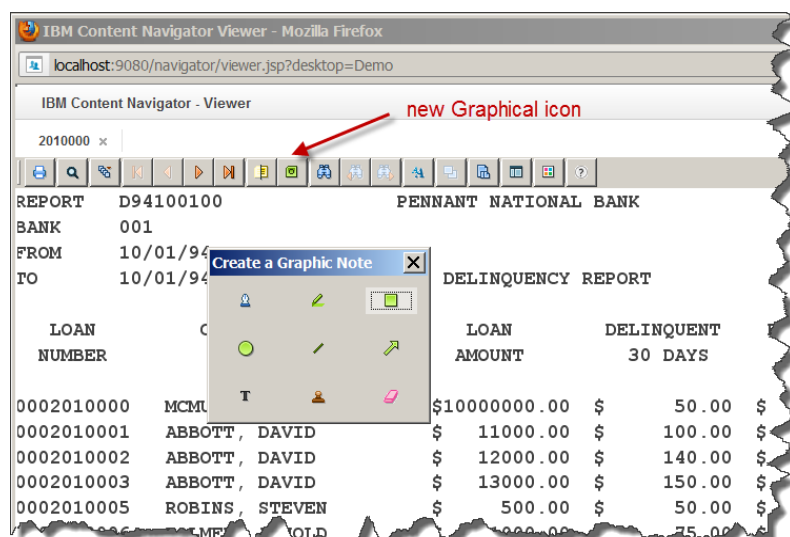
- Pen (free hand drawing)
- Highlighter
- Box
- Circle
- Line
- Arrow
- Stamp
- Text
- Eraser (used to delete other graphical annotations)

Text annotations (notes) are still available and positional.

Note that graphical annotations are not available for AFP documents.

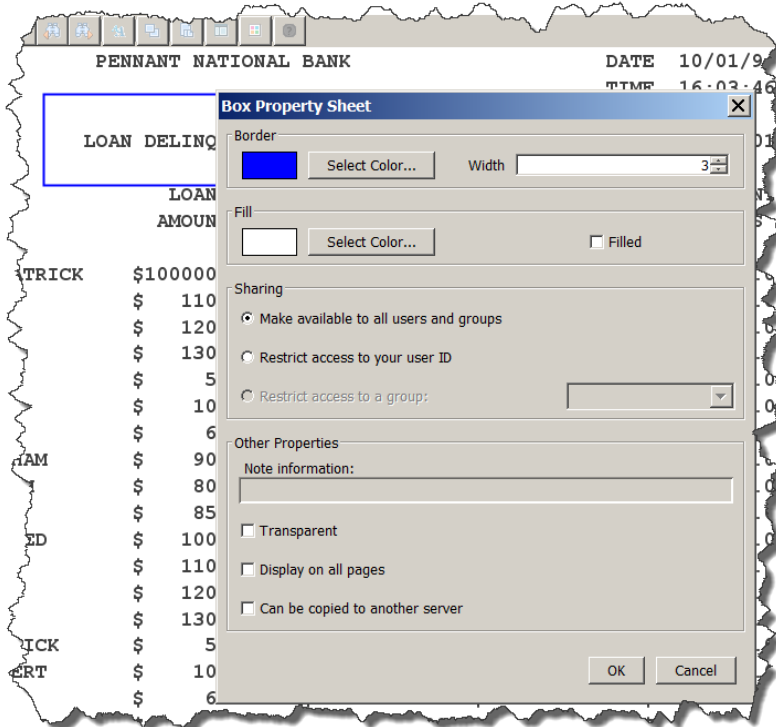
Creating a graphical annotation

To add a graphical annotation to your line data document (line data only), click on the new Graphical Annotation icon on the Line Data Applet tool bar.



ONDEMAND NEWSLETTER - 1ST QUARTER 2013

The Create a Graphic Note panel appears. Select the type of graphical annotation you wish to use (pen, line, box, etc.) by clicking the shape. Next right click the document where you would like the object to appear, and release the mouse when you have created the desired size. A new panel will appear with all the options for that shape, as shown in the example below:



These options allow you to choose the color, border, sharing, transparency, display on all pages and other options. Below are some samples of graphical annotations:

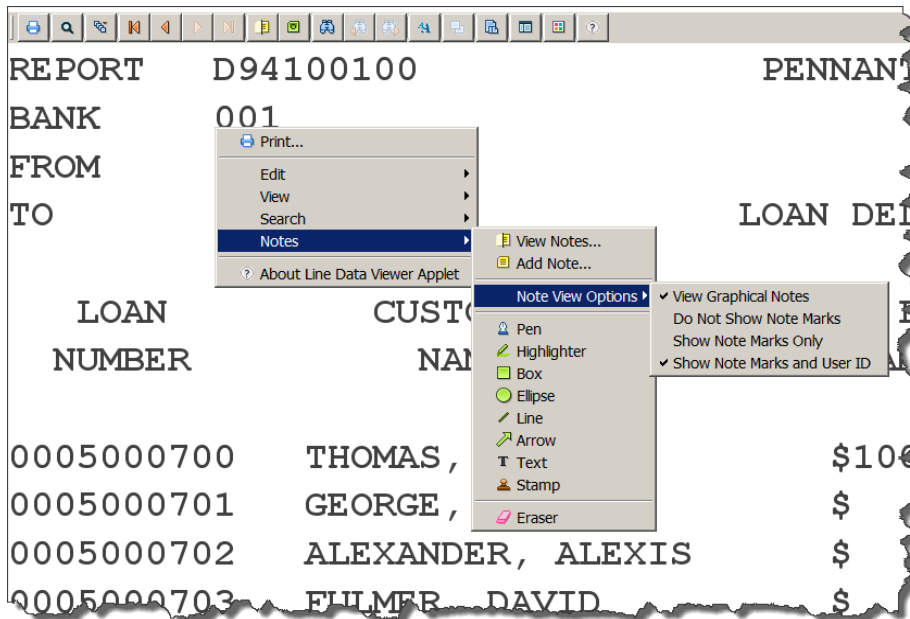
The screenshot shows a report titled 'LOAN DELINQUENCY REPORT' from Pennant National Bank. The report data is as follows:

LOAN NUMBER	CUSTOMER NAME	LOAN AMOUNT	DELINQUENT 30 DAYS	DELINQUENT 60 DAYS
0002010000	MCMULLIGAN, PATRICK	\$10000000.00	\$ 50.00	\$ 50.00
0002010001	ABBOTT, DAVID	\$ 11000.00	\$ 100.00	\$ 200.00
0002010002	ABBOTT, DAVID	\$ 10000.00	\$ 140.00	\$
0002010003	ABBOTT, DAVID	\$ 13000.00	\$ 150.00	\$
0002010005	ROBINS, STEVEN	\$ 500.00	\$ 50.00	\$
0002010006	PALMER, ARNOLD	\$ 1000.00	\$ 75.00	\$ 150.00
0002010007	PETERS, PAUL	\$ 650.00	\$ 50.00	\$
0002010008	ROBERTS, ALAN	\$ 9000.00	\$ 120.00	\$
0002010009	SMITH, ANDREW	\$ 8000.00	\$ 115.00	\$
0002010010	SMITH, PETER	\$ 8500.00	\$ 110.00	\$
0002010017	WILLIAMS, ALFRED	\$ 10000.00	\$ 50.00	\$ 50.00
0002010019	JAMES, TIMOTHY	\$ 11000.00	\$ 100.00	\$ 200.00
0002010022	THOMAS, JAMES	\$ 12000.00	\$ 140.00	\$

Annotations in the image include a blue box around the report title, a yellow box around the 'ABBOTT, DAVID' row, a red box around the 'SMITH, PETER' row, and a red box around the 'WILLIAMS, ALFRED' row. A red 'Graphical' watermark is overlaid on the table. A yellow arrow points to the 'DELINQUENT 30 DAYS' column, and a pink arrow points to the 'LOAN AMOUNT' column.

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

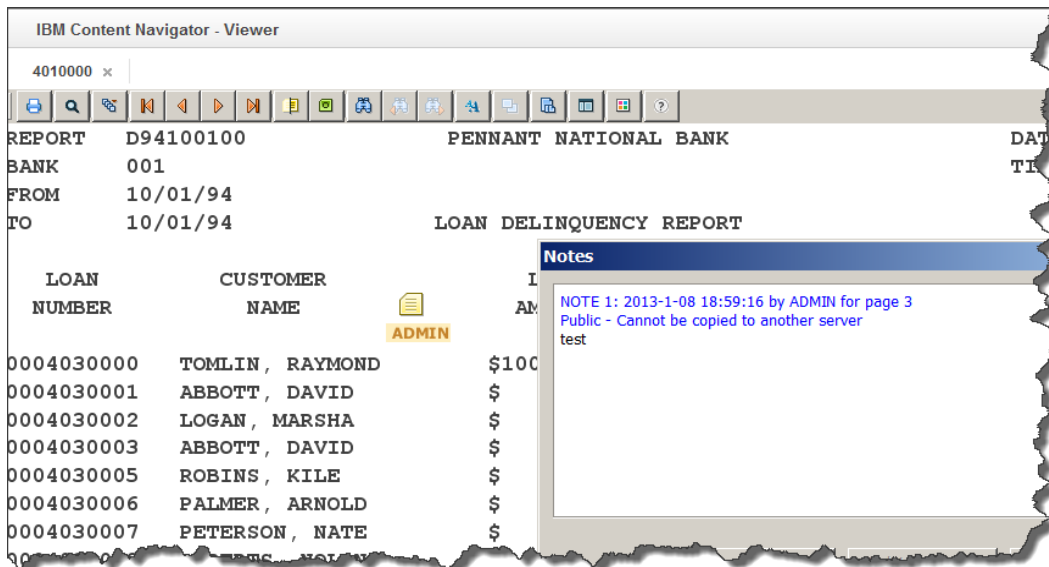
Another method to invoke the tool bar icons is to right click anywhere on the document as in example below:



Creating positional notes

To create a positional note, first use the mouse to select a small area on the document where you want the note to be pinned. When you select an area on the document it will create a highlighted area. Next, click on the Notes icon on the tool bar and create your note. If you do not select an area first, then the note will be pinned to the first page of the document in the upper left corner.

After you create a positional note, you can drag it to another location if needed. Here is a view of a note created by a user named Admin:



TIPS – MULTIPLATFORMS

Upgrading to DB2 v10

Problem

OnDemand uses the db2Runstats API to perform runstats on newly created tables as well as existing tables that have been updated since their last runstats. This API uses a structure that has changed in DB2 v10. If the db2Runstats API fails, the OnDemand server will rollback the table creation and issue an error.

Symptom

You might see messages on the console when creating new tables. For example:

```
RUNSTATS ON TABLE archive.SL21 -- SQLSTATE=Not Defined, SQLCODE=1,  
File=arssys.c, Line=1697
```

or

```
> arsdm -sv  
ARS4042I Updating runstat statistics for table root.arsag  
arsdb: ARS4042I Updating runstat statistics for table  
ARS4090E Unable to do a runstats on table arsag. err=1  
arsdb: ARS4043E Unable to update runstat statistics on table root.arsag
```

Cause

DB2 v10 has added a new parameter to the db2 runstats structure for the Index Sampling Option. This option did not exist in prior versions of DB2. The v10 db2Runstats API is not recognizing the older version of the structure and issuing an SQL1197N Invalid value specified for keyword "INDEX SAMPLE" in command "RUNSTATS". error.

Environment

This problem affects all versions of OnDemand for Multiplatforms prior to v9 fix pack 1 (v9.0.0.1).

Diagnosing the problem

To determine if you are affected by this problem, check the system console or run one of the arsdm database maintenance commands that perform a runstats.

Resolving the problem

Affected customers need to download and apply a DB2 APAR Patch for IC89492. This patch is part of DB2 v10 fix pack 2.

This tip is adapted from [support item 1623537](#).

DB2 v10.1 Support at OnDemand v9.0.0.1

Starting at OnDemand fix pack 9.0.0.1, you can create new databases using DB2 v10.1.

Prior to fix pack 9.0.0.1, you need the DB2 fix described in the previous tip.

Upgrading to OnDemand v9.0 from OnDemand v8.4.x

Upgrading from OnDemand v8.4.0 directly to OnDemand v9.0 is supported starting with server fix pack 9.0.0.1. You no longer need to install v8.4.1 and v8.5.0 before upgrading to v9.0. Likewise, you can upgrade from v8.4.1 directly to v9.0 without first installing v8.5.0.

To upgrade directly from 8.4.0 or 8.4.1, first install the base 9.0.0.0 code, then install the 9.0.0.1 fix pack. Then follow the detailed upgrade steps in the [Upgrade Guide for Content Manager OnDemand Server v9.0.0](#) for your initial server level.

Update the OnDemand Report Distribution Definitions

Customers who are using the OnDemand Report Distribution Feature should run the arssyscr command to redo the OnDemand Report Distribution definitions. This will correct the folder field to application group field mapping. This correction is required by code changes in fix pack 9.0.0.1.

To redo the definitions, run:

```
arssyscr -I <INSTANCE_NAME> -r
```

TIPS – z/OS

OnDemand for z/OS Quick Hits

Still running OnDemand for z/OS Version 8.4.0 or 8.4.1?

Just a reminder that End of Support for OnDemand for z/OS Version 8.4.x is April 30th, 2013.

See [US announcement letter 911-147](#) for details.

Essentials for a new OnDemand technical support member

Having a library of Installation Verification Procedures (IVP) is beneficial in three ways:

- 1) It provides a simple, repeatable set of jobs that can be used for training.
- 2) They can be used to verify upgrades.
- 3) They can be used to identify and isolate problems.

Here is a list of essentials:

- 1) The OnDemand IVPs themselves
- 2) From DB2 DSNTEJ8, the ODBC 31-bit sample C/C++ application
- 3) From DB2 DSNTEJ8E, the ODBC 64-bit sample C/C++ application
- 4) From DB2 DSNTEJ80, verifies Utility Stored Procedure (DSNUTILS) caller using C and ODBC
- 5) From OAM SYS1.SAMPLIB(CBRSAMIV) - this is extremely useful for verifying any OAM changes on your system
- 6) If you run LDAP, the Tivoli Directory Server [LDAP IVP](#)
- 7) If you use the ODF component of OnDemand, then of course the ODF IVP is extremely useful as it breaks everything down into very simple steps.

Sample JCL for running some of these in batch are available from the [Tips And Tricks](#) section of the OnDemand User Group web site.

If any IVP is not working correctly open up a support ticket to the owning IBM group.

How Do the MOS and the Object Size Interrelate?

Question

How do the MOS (Maximum Object Size) parameter in OAM and the Object Size parameter in the OnDemand Administrator client interrelate?

Answer

First, go to the OnDemand Administrator client setting on the Application Groups -> Storage Manager -> Advanced panel.

What does the field Object Size(K) define? The help panel says it means "The size of a storage object for the application group, in kilobytes." That's actually a little misleading and has been flagged for update in a future release.

What it really means is "When OnDemand reaches this minimum threshold of object size, the object will then be written to the storage manager." Note that it is not a limit, it is a minimum size.

For example, consider what will happen if:

- a) Object Size(K) is set to 40,000K (40M)
- b) There are 4 documents with sizes of 10M, 10M, 10M, and 50M

The first 10M document is written to the object. Is the object bigger than 40M? No, go to the next document.

Add the second 10M document to the object. Is the object bigger than 40M? No, go to the next document.

Add the third 10M document to the object. Is the object bigger than 40M? No, go to the next document.

Add the 50M document to the object. Is the object bigger than 40M? Yes, store the 80M object in the storage manager repository.

The storage manager must handle the object size or the load will fail. If the storage manager is OAM then for this store to succeed MOS in OAM must be set to a minimum of 80M.

Introduction to OD for z/OS Sysplex Implementation

OnDemand for z/OS supports the Parallel Sysplex, multiple Logical partition (LPAR) configuration. In a parallel sysplex environment, there are multiple ARSSOCKD combined library and object server modules. This type of configuration is unique to z/OS and is based on the Parallel Sysplex technology. Although the library and object server are installed in separate LPARs, the sysplex technology allows the servers in each separate LPAR to share the same configuration files, databases, JES, HFS and archive. For performance reasons, all HFS directories that are used for temporary storage are configured as unique to each ARSSOCKD server.

There are three major components that must be configured to support an OnDemand for z/OS Sysplex implementation:

1. Parallel Sysplex with XCF – Cross Coupling Facility
2. DB2 Data Sharing
3. Object Access Method -OAMplex

Parallel Sysplex with XCF

A Sysplex is a group of z/OS systems that communicate and cooperate with one another using specialized hardware and software. They are connected and synchronized through a Sysplex Timer® or System z™ Server Time Protocol (STP), and enterprise systems connection (ESCON®) or fiber connection (FICON®) channels. A Parallel Sysplex is a Sysplex that uses one or more coupling facilities (CFs), which provide high-speed caching, list processing, and lock processing for any applications on the Sysplex. Parallel Sysplex combines the capability of parallel processing and enabling read/write data sharing across multiple systems with full data integrity.

DB2 Data Sharing

A collection of one or more DB2 subsystems that share DB2 data is called a data sharing group. DB2 subsystems that access shared DB2 data must belong to a data sharing group. A DB2 subsystem that belongs to a data sharing group is a member of that group.

Each member can belong to one, and only one, data sharing group. All members of a data sharing group share the same DB2 catalog and directory, and all members must reside in the same Parallel Sysplex. Currently, the maximum number of members in a data sharing group is 32.

OAMplex

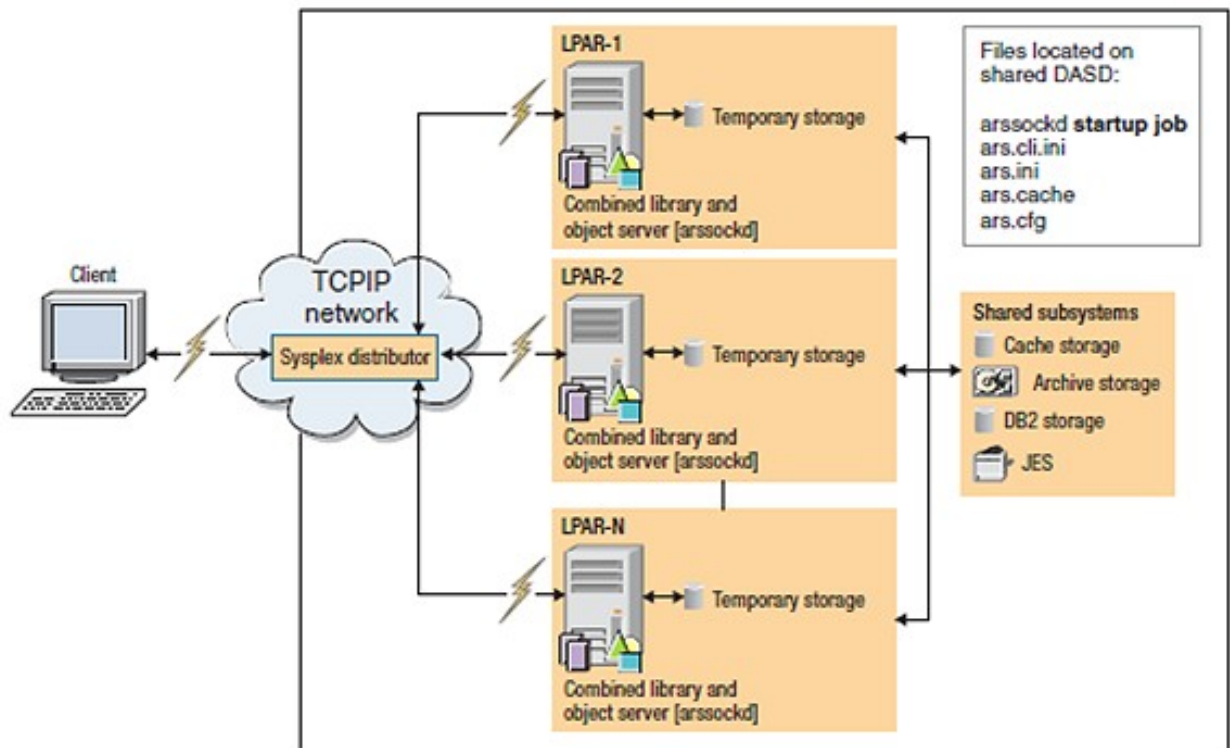
An OAMplex consists of one or more instances of OAM running on systems that are part of a Parallel Sysplex. An OAMplex has a one-to-one correlation to an XCF group in a Parallel Sysplex. The XCF group associated with an OAMplex consists of instances of OAM address spaces, running on separate LPARs in a Parallel Sysplex, sharing a common OAM database in a DB2 sharing group. Each instance of OAM is a member of the same XCF group. Also, the DB2

ONDEMAND NEWSLETTER - 1ST QUARTER 2013

subsystems connected to these instances of OAM belong to the same DB2 data sharing group. These instances of OAM belonging to the same XCF group have the ability to communicate with each other through the XCF services. The DB2 data sharing group shares the DB2 database information (OCDB, OAMADMIN, and object databases) among OAMs belonging to the OAMplex. When different OAMs sharing a common database on DB2 join an XCF group to become an OAMplex, all object data and configuration information is known to all instances of OAM in the OAMplex. Any object, regardless of which OAM stored the object, can be retrieved by any instance of OAM in the OAMplex.

Business Justification

This setup has the advantages of vertical scalability within the parallel sysplex as well as high availability. If there is failure of any single LPAR, it will only affect the current transaction(s). A retry request sent by the client will automatically be routed to the active LPAR. The clients connect to the system using a single IP address and port that are associated with a device or program that routes the request to an appropriate ARSSOCKD server. Some systems can route the request according to the z/OS Work Load Manager suggestions, while other devices simply route the requests following a round robin routine.



More details of a Parallel Sysplex implementation of OnDemand can be found in the [CMOD Sysplex Cookbook](#).

TIPS – IBM i

Field Name Restriction

At 8.5.0.6 and 9.0.0.1 or higher of the OnDemand Administrator client, the name "version" (in any case) can no longer be used as an application group field name when creating new application groups in OnDemand for i. This is a reserved keyword on IBM i.

Existing application groups that use the name "version" as a field name are not affected.

Considerations for High Availability

In a high availability environment on IBM i, we recommend the following:

Do not replicate any work directories

High Availability software will lock the stream files during replication. This can result in failed loads, or orphan files left in the Integrated File System (IFS) on your IBM i system.

Archiving

The home directory of the user archiving data is used by the Add Report to OnDemand (ADDRPTOND) command for temporary files. The root directory (/) is used if the user's home directory does not exist. For example, for user DBRYANT:

```
/home/dbryant
```

Server Printing

Server printing uses directory PRTTMP in the instance directory. For example:

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

Error and log files

Some log files go to directory TMP in the instance directory. For example:

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

ODWEK cache

Files will be cached again on backup system, if necessary. For example:

```
/QIBM/UserData/OnDemand/www/1141/CACHE
```

The system TMP directory

Some work files go to the system TMP directory. For example:

```
/TMP
```


ONDEMAND NEWSLETTER - 1ST QUARTER 2013

Kofax Ascent Capture release directories

Kofax Ascent Capture uses a temporary directory for its release process. For example:

```
/Ascent/Scan01
```

Mounting file systems

In a high availability setup, you will have a source system and a target system. On the source system, each time the instance is started, Archive Storage Manager (ASM) will mount the file systems the first time data is migrated to or retrieved from a disk pool. On the target system, the instance will not be active, so you must mount the file system. This can be done manually or by your high availability software.

An example of the MOUNT command for the primary disk pool in ASP01 is given below:

```
MOUNT TYPE(*UDFS) MFS('/dev/qasp01/ondemand_qusrond_primary_01.udfs')  
MNTOVRDIR('/qibm/userdata/ondemand/QUSROND/ASMASP01/PRIMARY')
```

You must ensure that the disk pools are always mounted on both the source and target systems when replication is running.

Replicating Libraries

In addition to your instance libraries, you should also replicate library QUSRRDARS if you are at v5.4 or v6.1. At those releases, library QUSRRDARS contains information used by ASM. At v7.1, the only file in QUSRRDARS that must be replicated is QARLCASP, and then only if you are using Independent Auxiliary Storage Pools (IASPs) with OnDemand.

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](#)".