OnDemand Newsletter

News and Tips about IBM Content Manager OnDemand

1st Quarter 2012

In This Issue

News

About this Newsletter

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

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

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

OnDemand for Multiplatforms Fix Pack 8.5.0.4 Available

ODMP Fixpack 8.5.0.4 was released on December 15, 2011. No new functionality was added in this fixpack.

The fixpack files are available from the IBM Fix Central site.

This fix pack addresses the following APARs:

Server

- PM41540 PDF IDX can't find field if BASE=TRIGR has +1 word
- PM43548 ARSXML UPDATE abends if input file has +1 AG DB MAPPING
- PM46304 ARSLOAD crashes with PDF IDX and -j parameter file option

© Copyright IBM Corporation 2012

- PM46738 ARSXML group update resets GROUPOWNER information if parameter is missing
- PM47234 ARSDOC GET core dumps when using -X (Load ID)
- PM47805 LOAD of LINE data fails with can't convert ' ' to decimal
- PM50597 ARSXML UPDATE fails to update a no storageSet value
- PM52204 ARSSOCKD.EXE memory usage increases until OUT OF MEM crash
- PM53344 CFSOD=YES option does not stay selected

ODWEK

- PM46418 AFP2PDF stanza isn't being properly matched
- PM46904 AOIQS Cookie trace displays password in clear text
- PM48412 When ODWEK invokes exit program, stdout/stderr are not redirected

Windows Client

- PM46416 Scrolling fails if OVERSTRIKE=MERGE and NO overstrike lines
- PM46719 SQL QUERY truncated in ODF report bundle if 'space' in LIKE
- PM48345 Font used in printing annotation is too small to read
- PM46731 Silent install issue when non-English locale is used
- PM48014 HEB Client: NOTES are not placed in correct location
- PM48610 AFP VIEWER displays doc with custom FORMDEF incorrectly

OnDemand for z/OS Fix Pack 8.5.0.4 Available

OnDemand server level 8.5.0.4 for version 8.5 of OnDemand for z/OS® is now available.

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

IBM Web Interface for Content Management (WEBi) 1.0.4 Fix Pack 4

Fix pack 4 for WEBi 1.0.4 is now available. This fix pack contains support for:

- 64-bit operating systems
- Java 1.7
- Internet Explorer 9
- Firefox 6

WEBi 1.0.4 fix pack 4 is supported running on the following operating systems:

- AIX 7.1
- RHEL 6.1
- SuSE 11
- Windows 2008 R2

For more information see support item 4031531 and support item 7023303.

© Copyright IBM Corporation 2012

OnDemand for i Integration with Content Manager for iSeries

Integration with Content Manager for iSeries is now available for versions 5.4, 6.1, and 7.1 of OnDemand for i. To integrate OnDemand with Content Manager for iSeries you must install the following program temporary fixes (PTFs) or their superseding PTFs:

Version	Server
7.1	SI44870 – New sample programs SI44872 – Integration code SI44950 – Sample overlay
6.1	SI44680 – New sample programs SI44681 – Integration code
5.4	SI44867 – New sample programs SI44906 – Integration code

Using the integration software, you only need the Content Manager Client to access both images and archived spooled files.

More information about integration between OnDemand and CM for iSeries can be found in support item 7023237.

Update to the OnDemand for i Media Migration Facility

The enhancements in this update to the Media Migration Facility (MMF) include:

- MMF now generates more messages to aid problem determination.
- MMF can now bypass errors if they occur. A new data area QRLCSFAMMF can be created in QTEMP to enable this function.
- The MMF reports are now owned by the user who started the MMF job.

The PTFs required are:

Version	Server
7.1	SI45350
6.1	SI45366
5.4	SI45365

More information about the Media Migration Facility (MMF) can be found in support item 7008380.

© Copyright IBM Corporation 2012

Tips – Cross Platforms

The Security Certificate for the Content Manager OnDemand Applets has Expired

The Line Data Viewer and AFP2HTML applets included with the OnDemand Web Enablement Kit (ODWEK) have been refreshed for V8.4.1.8 to update the security certificate, which had expired.

The new certificate expires on Tuesday, October 01, 2013 at 17:59:59 Mountain Daylight Time.

To eliminate the expired security certificate warnings users may be experiencing, replace the following two files on the ODWEK server:

ODAfp2Html2.jar

ODLineDataViewer2.jar

Note: The V8.4.1.8 applets can be used with previous ODWEK versions from V8.4.1.0 to V8.4.1.7. Also, you might need to clear your browser cache to allow download of the new applets.

OnDemand for Multiplatforms

The V8.4.1.8 JAR files can be obtained by installing the ODWEK V8.4.1.8 fixpack or by downloading just the jar files from:
ftp://service.software.ibm.com/software/ondemand/fixes/v841/8.4.1.8/

OnDemand for z/OS

The V8.4.1.8 JAR files can be obtained by installing the ODWEK V8.4.1.8 fixpack or installing APAR PM47724.

OnDemand for i

The V8.4.1.8 JAR files can be obtained by installing the following PTFs:

Version	ODWEK PTF
7.1	SI44993
6.1	SI44992
5.4	SI44989

Your OnDemand server must be running V8.4.1.7 to install these PTFs. If your system does not already have the V8.4.1.7 PTFs applied, the V8.4.1.7 PTFs will be included when the applets PTF is requested. To determine the server version, log on to either the OnDemand

 © Copyright IBM Corporation 2012

Administrator or end-user client. Once logged on, the server version is shown on the message line in the lower right portion of the panel.

You should also clear your OnDemand cache directory on your IBM i server. The cache directory name will be:

/QIBM/UserData/OnDemand/www/cache

or

/QIBM/UserData/OnDemand/www/nnnn/cache

where nnnn is the CCSID of the ODWEK server if you have configured multilingual support.

This tip is adapted from support item 1567899 (ODMP); support item 1568050 (OD for z/OS); and support item 1568173 (OD for i).

 © Copyright IBM Corporation 2012

Create Customized Views for Line Data & SCS Reports

Introduction

In this tip you will learn how to use the OnDemand Administrator Client to create customized views for line data reports that contain columns of data. You can create customized views that reorder, duplicate, or remove columns. Also, one or more columns can be locked in place so that when horizontal scrolling is performed, the locked columns remain in view while other columns scroll into view.

Note that all references in this tip to line data also include SCS data.

Getting started

The first step to creating a customized view is to define header and field information. This information is located on the Logical View Fields tab of the application window that is shown in Figure 1. The header and field information can be specified on the tab, but it is much easier to view sample data and then define the fields by selecting text. This is done by clicking on the Sample Data button.

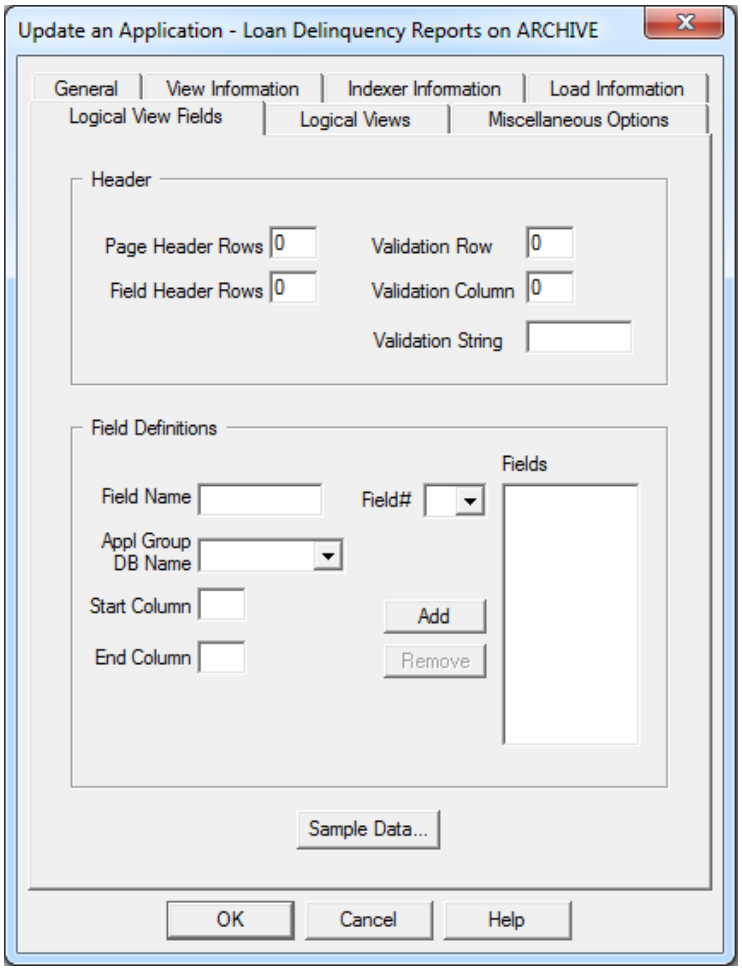

Figure 1. Logical View Fields page

 © Copyright IBM Corporation 2012

The sample data option allows you to select a file that contains the line data report, display it, and then select areas of the report to define the header and field information. To illustrate this, let's take a look at an example. Figure 2 is an example of a line data report that contains loan delinquency report information. The report contains 6 columns of data: loan number, customer name, loan amount, delinquent 30 days, delinquent 60 days, and delinquent 90 days. The report also includes several lines that contain page titles and field titles.

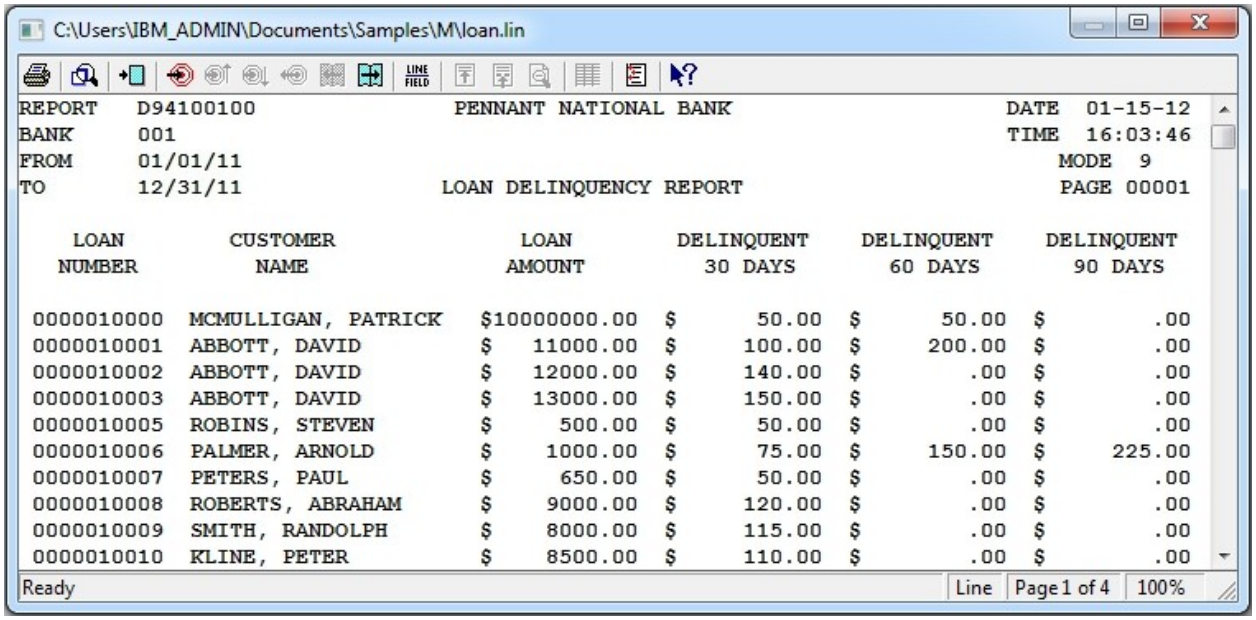

Figure 2. Loan delinquency report example

Define the page header

For line data reports that have title lines prior to the start of the column data, a page header or field header must be defined before the logical view fields are defined. These lines must be identified so that they can be differentiated from the lines that contain the columns of data. Page title lines can be locked in place. When the page title is locked in place and the page is scrolled in a vertical direction, only the columns of data are scrolled.

In the example in Figure 2, there are 8 lines that precede the first line that contains the columns of data. The first 5 lines are the page title lines, and the next 3 lines are the field title lines. You can choose to define only a page header by selecting all 8 lines for the page header, or you can define a page header and a field header. In this example, we'll define both. The advantage of defining both headers is that the field titles will remain with the columns of data when header and field locking is used and the data is scrolled in a vertical or horizontal direction.

In Figure 3, the first 5 lines have been selected. These lines will be used to define the page header. To select one or more lines, the selection mode must be set to Line. The mode is selected by pressing the toolbar button that is circled in red in Figure 3. Pressing the toolbar button toggles the selection mode between Line and Field. The status bar in Figure 3 indicates the current selection mode is Line. Once the selection mode is set to Line, position the cursor on the first line at the top of the page, press and hold the left mouse button, and drag the mouse downward until the first 5 lines have been selected.

 © Copyright IBM Corporation 2012

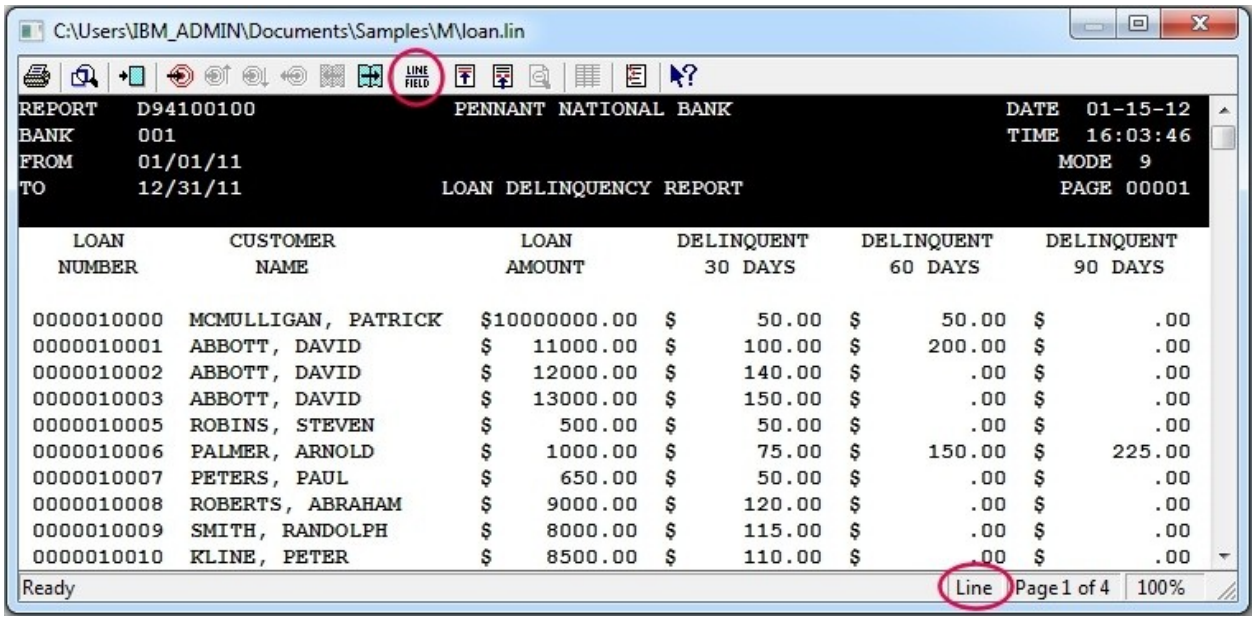

Figure 3. Page header lines

The fifth line that was selected doesn't contain any text. This line doesn't have to be a part of the page header but if it isn't, it must be selected when the field header is defined. The reason for this is the header definitions include only the number of lines for each header, and not the starting and ending line positions. As you can see in Figure 1, the two values for the header information are Page Header Rows and Field Header Rows.

Once the lines have been selected, click on the toolbar button that is circled in red in Figure 4. The Add Page Header window is displayed. The window shows the starting and ending line numbers as well as the number of lines selected. Click on the OK button to add the page header.

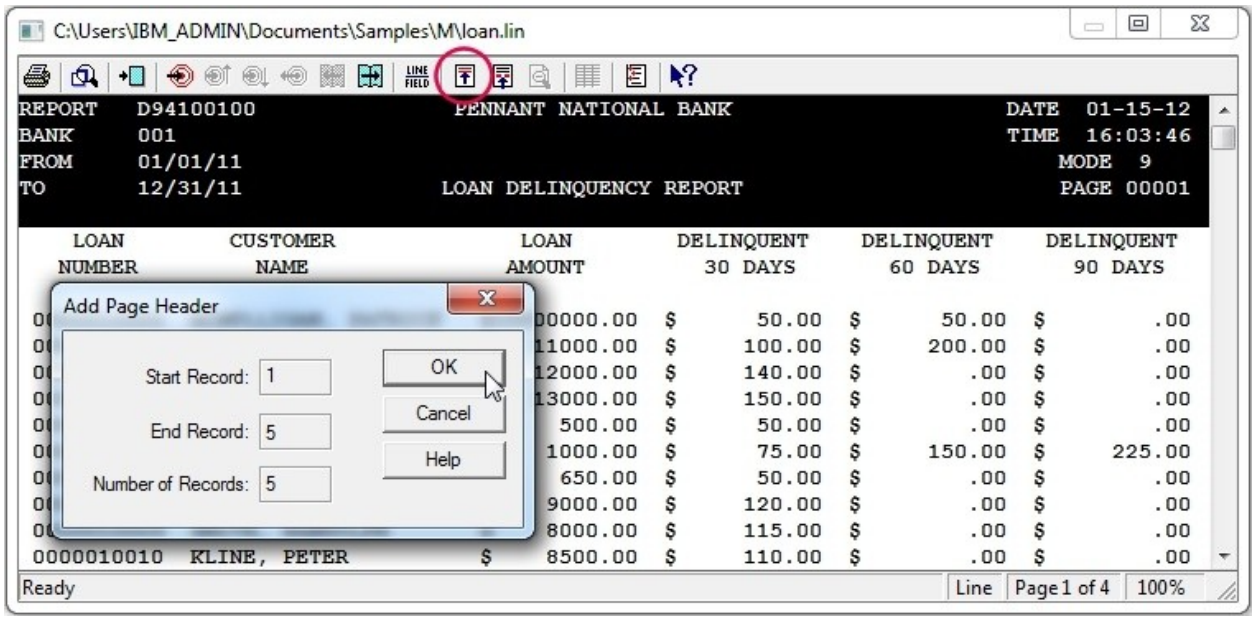

Figure 4. Add Page Header

 © Copyright IBM Corporation 2012

Define the field header

After the page header has been defined, the next step is to define the field header. Position the cursor on the sixth line (the line that immediately follows the last page header line), press and hold the left mouse button, and drag the mouse downward until 3 lines have been selected. The last line that was selected doesn't contain any text. Selecting this line is important because it provides spacing between the column titles and the columns of data.

Click on the toolbar button that is circled in red in Figure 5. The Add Field Header window is displayed. The window shows the starting and ending line numbers as well as the number of lines selected. Click on the OK button to add the field header.

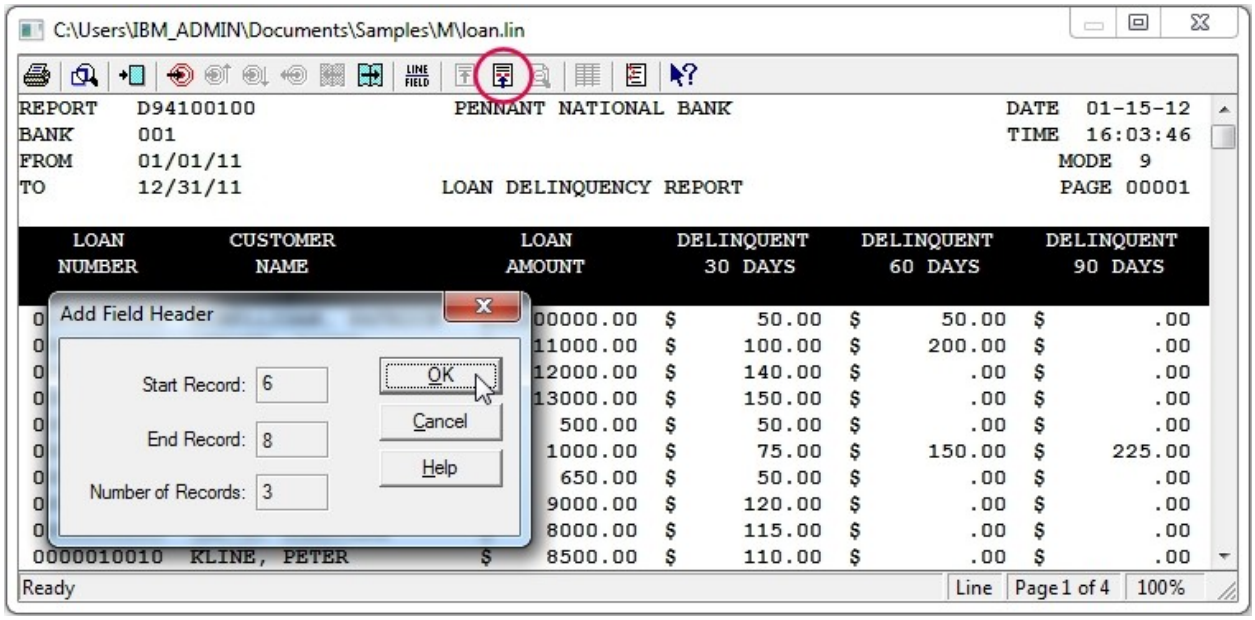

Figure 5. Add Field Header

Define the validation string

Once the page header and field header have been defined, the next step is to define the validation string. The validation string is used to determine whether or not the customized view can be used to format the page of data. If the validation string isn't found on the page, the data will be displayed in its original format. The validation string is defined by selecting a word or words that appear in the same place on every page. Up to eight characters can be selected. To select the validation string, the selection mode must be set to Field. The mode is selected by pressing the toolbar button that is circled in red in Figure 3.

Once the selection mode is set to Field, move the cursor to the left of the text to be selected. Press and hold the left mouse button and drag the mouse to the right until all of the characters have been selected. In this example, the word "REPORT" is used as the validation string since it appears in the same place on every page of the document. To define the validation string, click on the toolbar button circled in red in Figure 6. The Add Validation String window is displayed. The window shows the row number, column number, and selected string. Click on the OK button to add the validation string.

 © Copyright IBM Corporation 2012

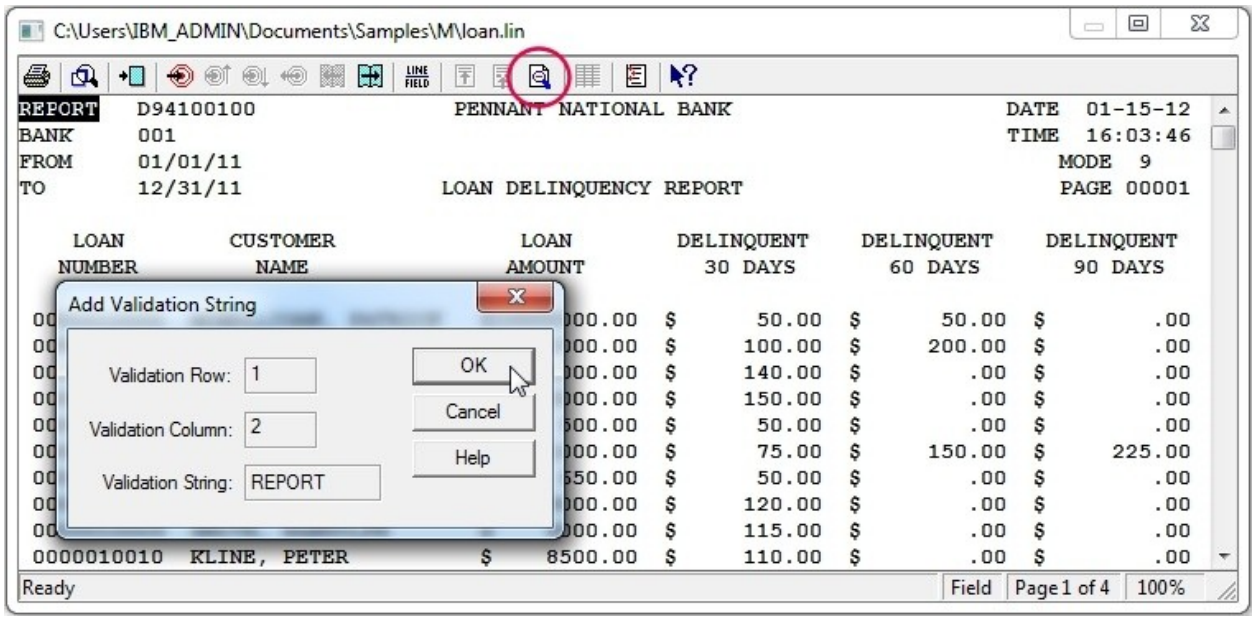

Figure 6. Add Validation String

Define logical view fields

Once the header information has been defined, the next step is to define the logical view fields. A logical view field defines the location and length of a column of data. Typically, a logical view field is defined for every column of data. This gives you the most flexibility for creating logical views with some or all of the columns of data. However, it isn't necessary to define a logical view field for every column of data. If there's a column of data that you want to exclude (for example, social security numbers, salary information, birth dates, and so on), don't define a logical view field for the column.

In this example there are six columns of data, so there will be six logical view fields to define. A logical view field is defined by selecting the characters that define a column of data. To select the logical view field, the selection mode must be set to Field. The mode is selected by pressing the toolbar button that is circled in red in Figure 3.

There are a few things to keep in mind when you are defining a logical view field:

- The length of the field should include one or more spaces that appear before or after the column of data. If the spaces aren't included, the columns of data will be displayed next to each other without any spacing in between the columns.
- If the columns are going to be rearranged, one or more spaces should be included in the logical view field for the last column so that if the last column is moved, there is still adequate spacing between this column and a column that follows it when the data is displayed.
- If a field header is used and the field title is longer than the data in the column, the logical view field should be defined so that the starting position of the logical view field is before the start of the field title and the ending position is after the end of the field title. In other words

 © Copyright IBM Corporation 2012

make sure the length of the logical view field is large enough to include the field title so it isn't truncated when it is displayed.

Once the selection mode is set to Field, move the cursor to the left of the text to be selected. In this example, the first logical view field to define is the column of data that contains the loan number. Since you must select text on the line that immediately follows the last line of the header line, position the cursor to the left of "000010000". Press and hold the left mouse button and drag the mouse to the right until the loan number and the two spaces to the right of the loan number have been selected.

To define the logical view field, click on the toolbar button circled in red in Figure 7. The Add a Logical View Field window is displayed. The window shows the starting column, the ending column, the field number, and the selected string. The logical view field is identified by providing a name. Type a name up to 20 characters in the Field Name field. In this example, "Loan_Number" is used. The name shouldn't include a space if you want to use the Expression Find function in the Windows client.

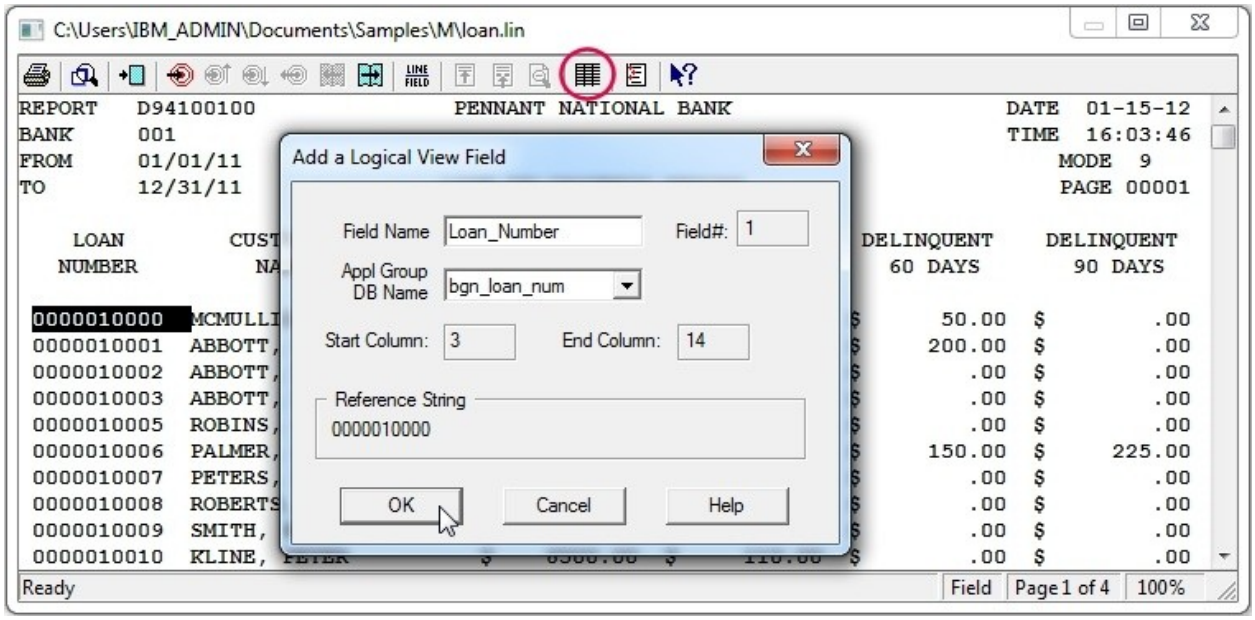

Figure 7. Add a Logical View Field

Optionally, the name of an application group database field can be specified by selecting one of the names from the Appl Group DB Name list. The only time an application group database field name should be associated with the logical view field is when you want to automatically display a document and highlight the line containing the database value. Some or all of the data in the column must have been stored in the database for the search to work. In most cases, it isn't practical to store every value in the database but it can be done. Typically, if the data is sorted in the column, the first and last value can be stored rather than every value in the column. When a search is performed, the database is searched by determining if the specified search value occurs between the starting and ending value. In this example, the column containing loan numbers is sorted in ascending order. Rather than storing every loan number in the database, the first and last loan number is stored in the database in two different database fields.

 © Copyright IBM Corporation 2012

The two database fields in this example are bgn_loan_num and end_loan_num. Since there are two database fields for this column of data, only one can be chosen. The database field that contains the starting value must be selected. For this example, bgn_loan_num contains the starting loan number. Click on the OK button to add the logical view field. Logical view fields for the remaining 5 columns are defined in a similar manner.

After all of the logical view fields have been defined, the header and field information is saved by clicking on the X in the upper right hand corner of the sample data window. Figure 8 shows the header and field information that have been defined for the loan delinquency report.

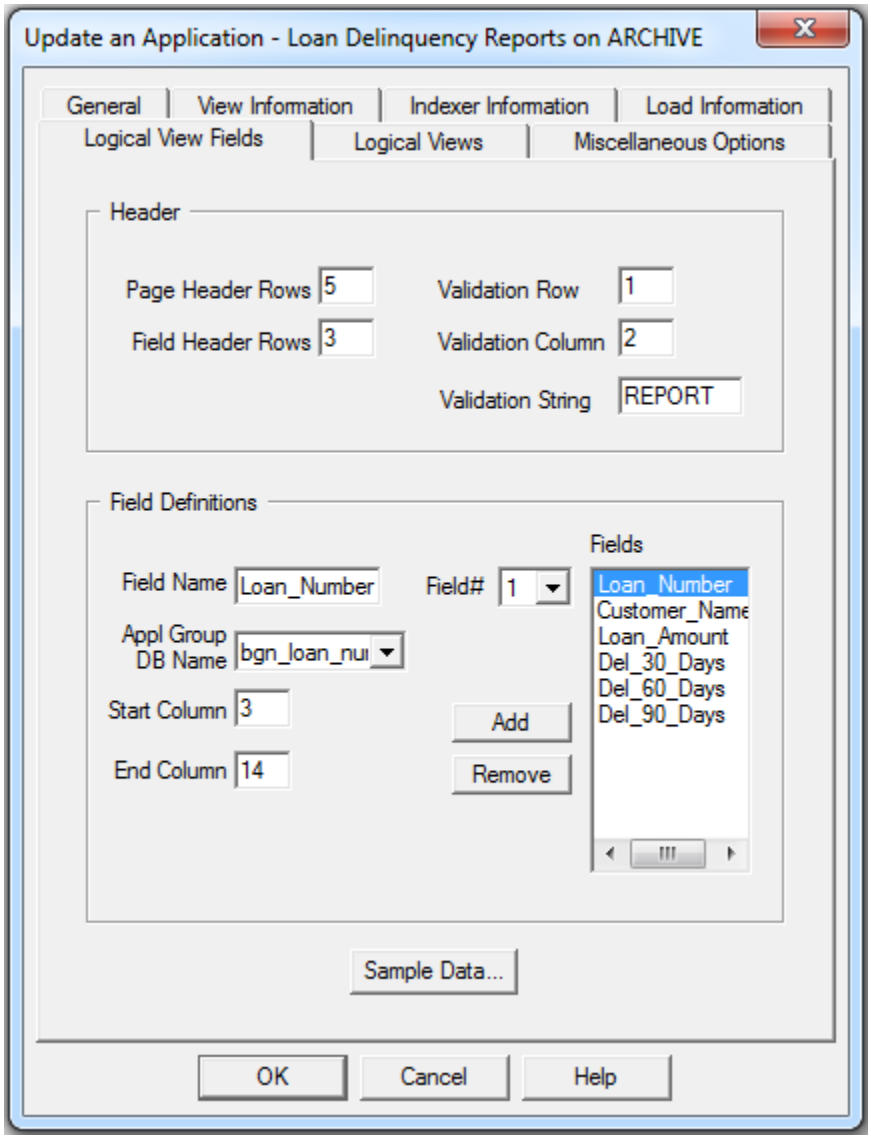

Figure 8. Logical View Fields page

You need to associate a logical view field with a database field so that when a search is performed using the Windows client, the value that is provided in the search criteria for the database field can be searched for in the column of data that is represented by the logical view field. If the

 © Copyright IBM Corporation 2012

value is found in the column of data, the document will be displayed and the line containing the value will be highlighted. For example, in Figure 9, a loan number is provided as part of the search criteria.

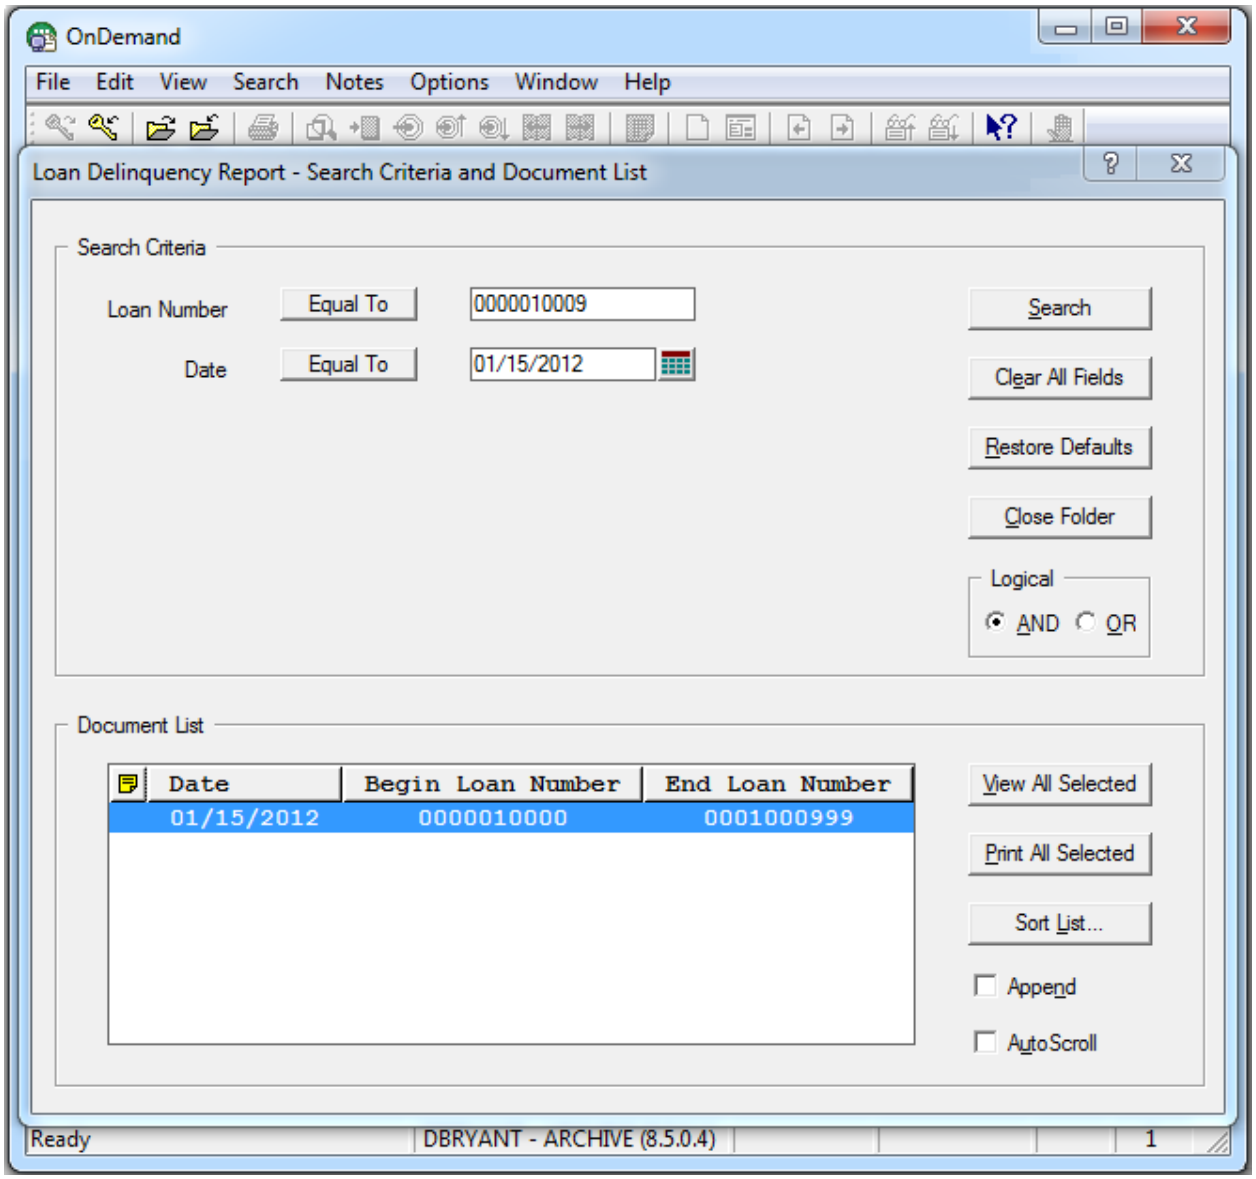

Figure 9. Windows client search example

When the search is performed, one document matches the search criteria. The document is automatically retrieved and displayed and the line that contains the loan number is highlighted. In order for this to work, the AutoView setting on the Options menu must be set to Single Document. An AutoView setting of First Document will also work if only one document matches the search criteria or if the document containing the loan number is the first document in the list. Figure 10 shows a loan delinquency report with the line containing loan number "000010009" highlighted.

 © Copyright IBM Corporation 2012

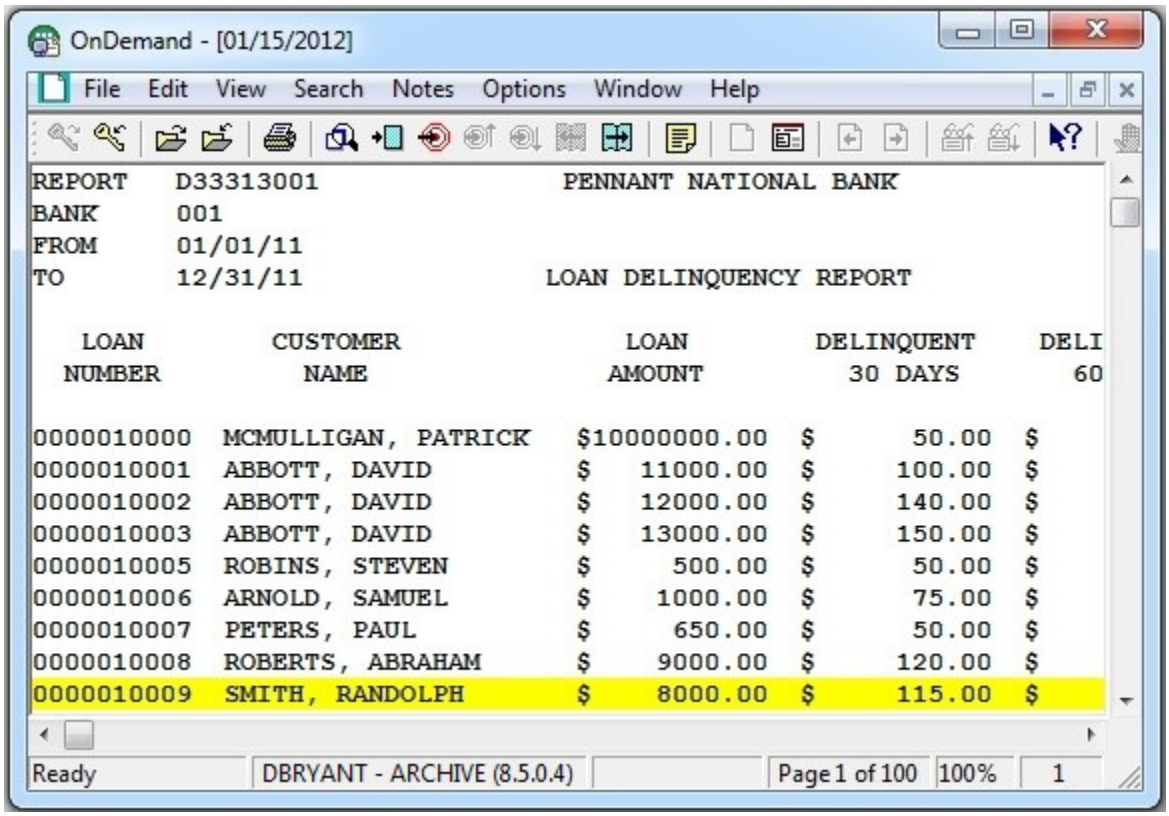

Figure 10. Document displayed in the Windows client

Defining columns for multi-up reports

In some cases, a document may contain a column of data that occurs multiple times on the page. To search all of the related columns for the specified search value, you'll need to associate the same database field with each column that contains the data.

For example, suppose the loan delinquency report contains 6 columns where the first 3 columns are loan number, customer name, and loan amount. The last 3 columns are also loan number, customer name, and loan amount. Six logical view fields are defined (for example, loan_num1, name1, amount1, loan_num2, name2, amount2). In this case, both logical view fields loan_num1 and loan_num2 will be defined by specifying the same database field bgn_loan_num so that the search will be performed in both columns. Suppose a search is performed with a loan number that is located in the second column of loan numbers. In Figure 11 a search is being performed where the loan number "0000010050" is specified.

 © Copyright IBM Corporation 2012

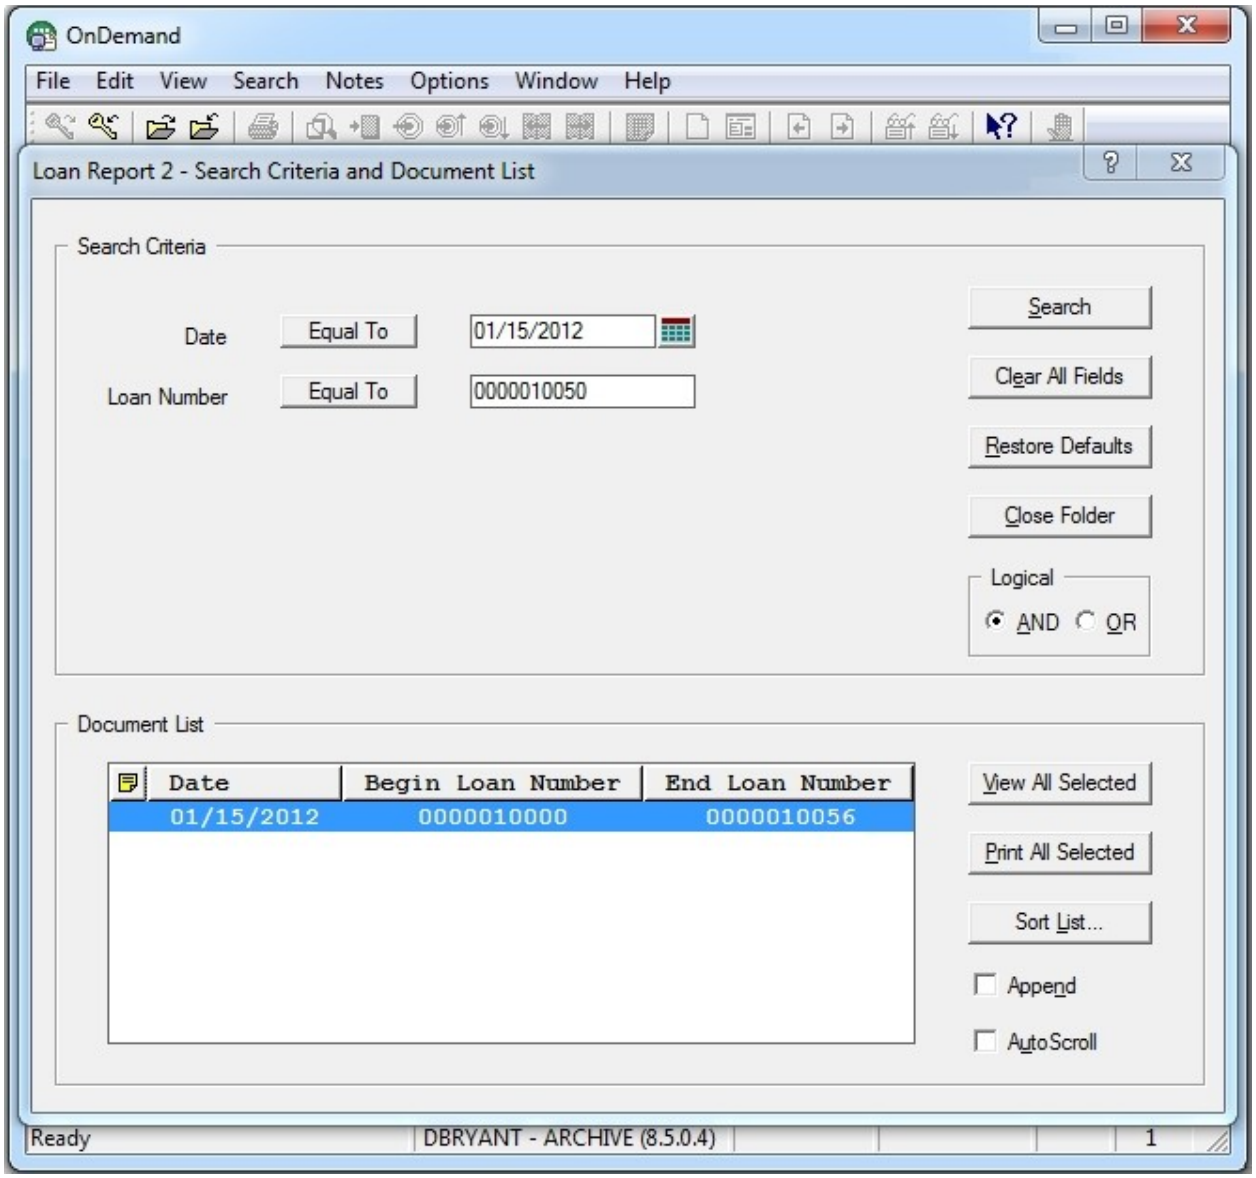

Figure 11. Windows client search example

When OnDemand performs the search, one document matches the search criteria. The document is automatically retrieved and displayed, and the line that contains the loan number is highlighted. The loan number is found in the second column that contains loan numbers. Figure 12 shows a loan delinquency report with the line containing loan number "000010050" highlighted.

 © Copyright IBM Corporation 2012

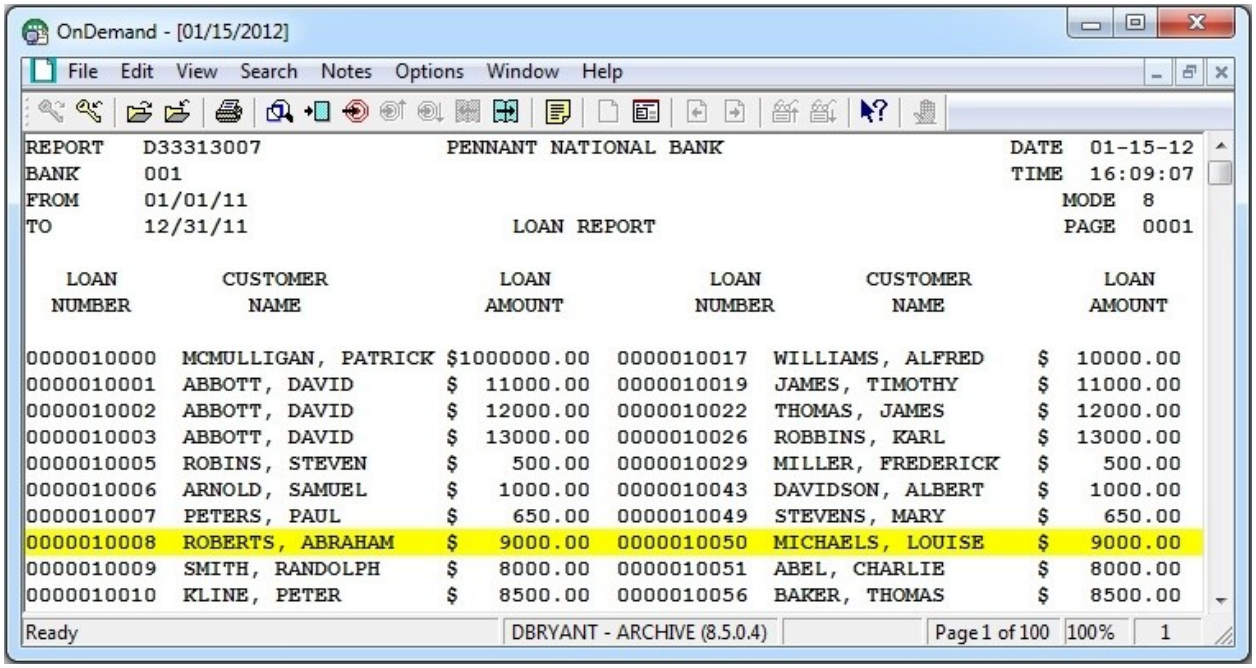

Figure 12. Document displayed in the Windows client

Define logical views with fields

Now that the logical view fields have been defined, you can create customized views of the line data. Customized views or logical views are created and stored on the Logical Views tab of the application window shown in Figure 13. In addition to specifying which columns of data are included when the data is displayed, you can also specify how the data is displayed by selecting values for attributes such as background color, heading color, zoom, text fidelity, overstrike, and selected area color.

 © Copyright IBM Corporation 2012

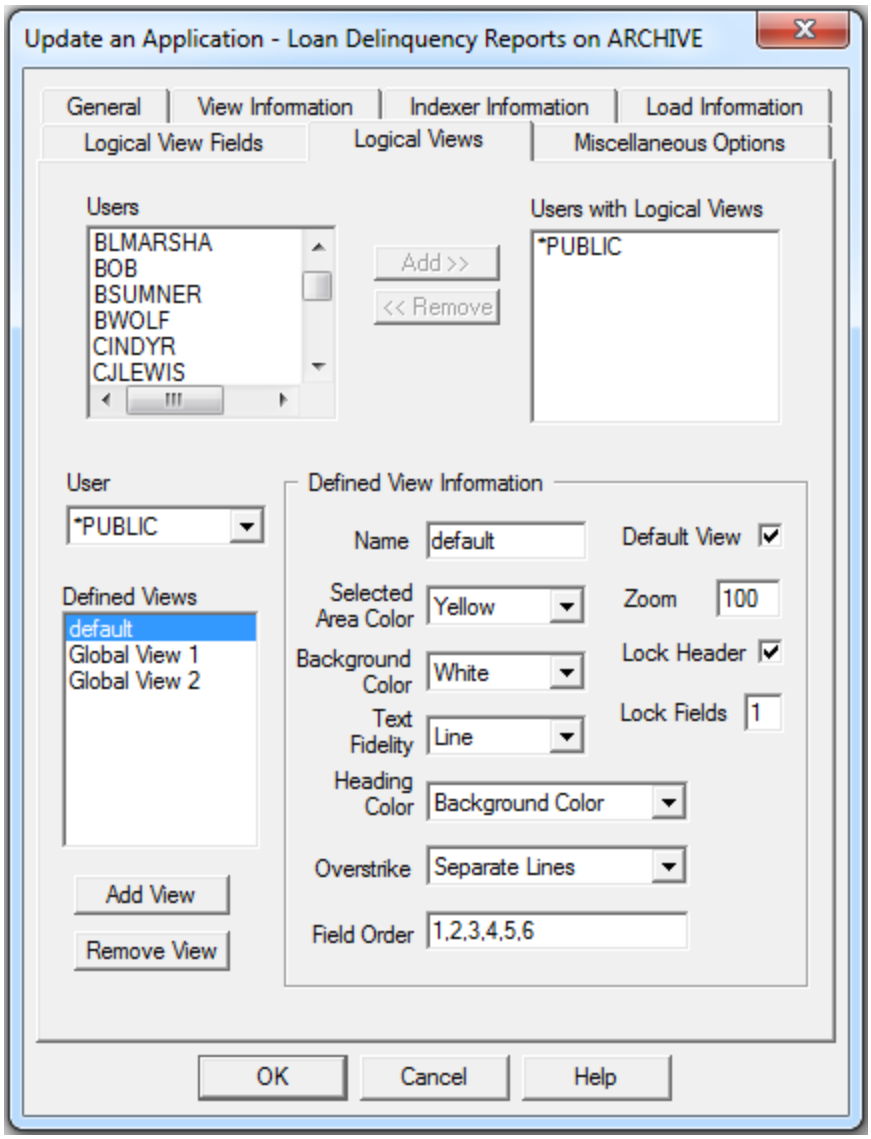

Figure 13. Logical Views page

To specify which columns of data to include and what order to display the columns, you specify a comma delimited list of logical view field numbers in the Field Order field shown in Figure 13. For each logical view field defined, there is a field number associated with it. To determine the field number, select the logical view field on the Logical View Fields tab. The field number will be shown in the list called Field#. For example, in Figure 8, the logical view field Loan_Amount is selected. The field number is 3 for this field.

In Figure 13, the value for the Field Order is "1,2,3,4,5,6". This means that all 6 of the logical view fields are included in the view called "default". The columns are displayed in the same order they appear in the original data since the logical view fields were defined in the same order the columns appear in the data. In addition to setting the column order, the view also specifies that the page and field titles are locked in place as well as the first column (i.e. Loan_Number). This is

 © Copyright IBM Corporation 2012

accomplished by placing a check mark in the Lock Header box and setting the Lock Fields value to 1. When the data is displayed, the page and field titles will remain in view when the page is scrolled in a vertical direction. When the user scrolls the page in a horizontal direction, the first column will remain in place while the remaining columns move to the left or right.

In Figure 14 the loan delinquency report is displayed using the logical view described in Figure 13. The document has been scrolled to the bottom of page 1 and scrolled to the left. Notice the page and field titles are displayed on the page even though the page has been scrolled in a vertical direction. Also, the loan number column is still in view even though the page has been scrolled to the left.

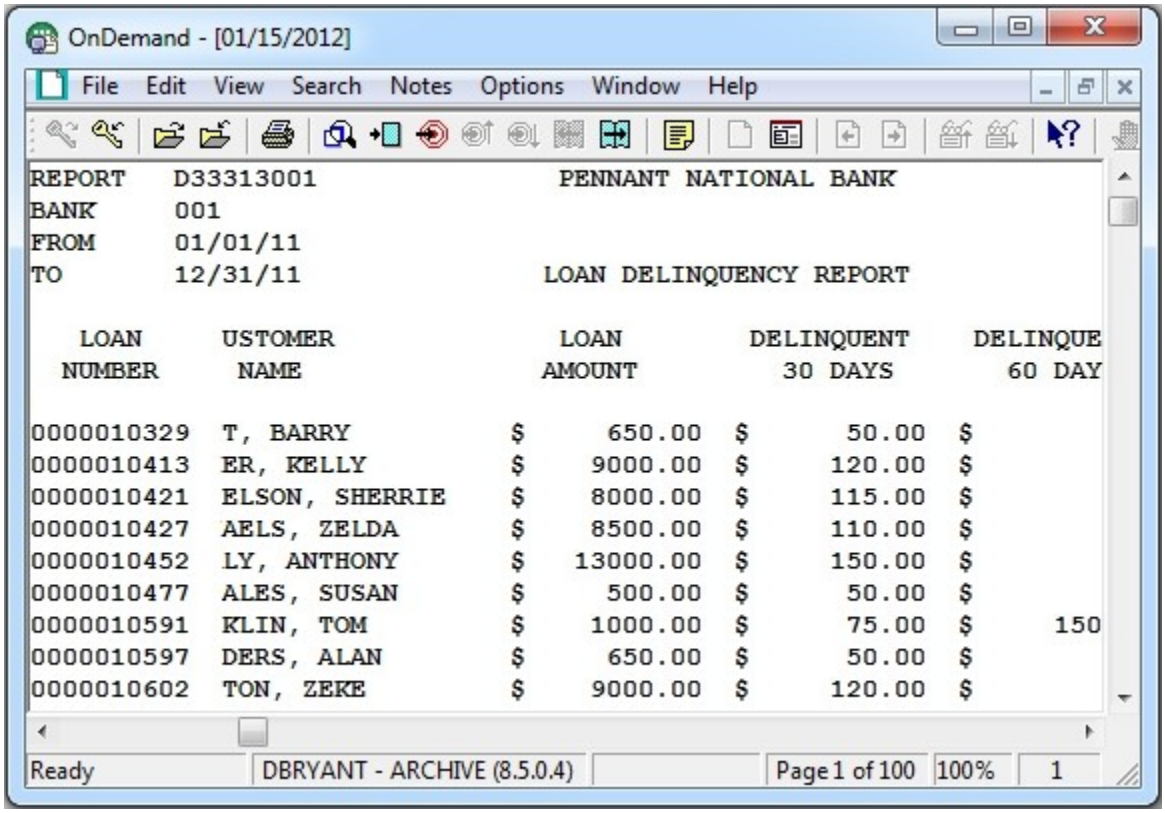

Figure 14. Loan delinquency report with locked headers and locked fields

You can also create a customized view by arranging the columns in a different order or by including a subset of the columns. In Figure 15, the value for the Field Order is "2,1,3,6". Only 4 of the 6 logical view fields are included in the view. The columns are displayed in the following order: customer name, loan number, loan amount, and delinquent 90 days.

 © Copyright IBM Corporation 2012

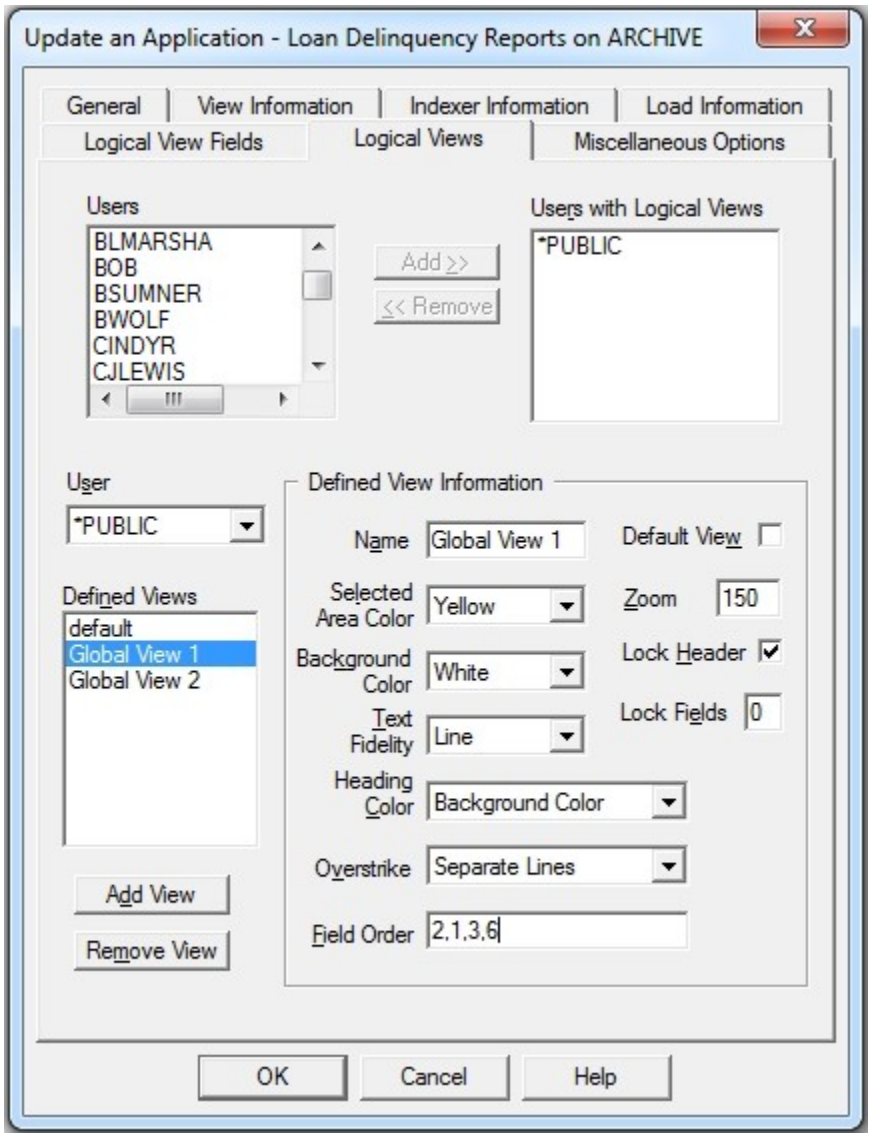

Figure 15. Logical Views page

When the loan delinquency report is displayed, the document is either displayed in its original format or if a default logical view is specified, it will be displayed using the default logical view. After the document is displayed, you can choose to display the document using a different logical view. In Figure 16 the document is displayed using the logical view shown in Figure 15. As you can see, only 4 of the original 6 columns are included.

 © Copyright IBM Corporation 2012

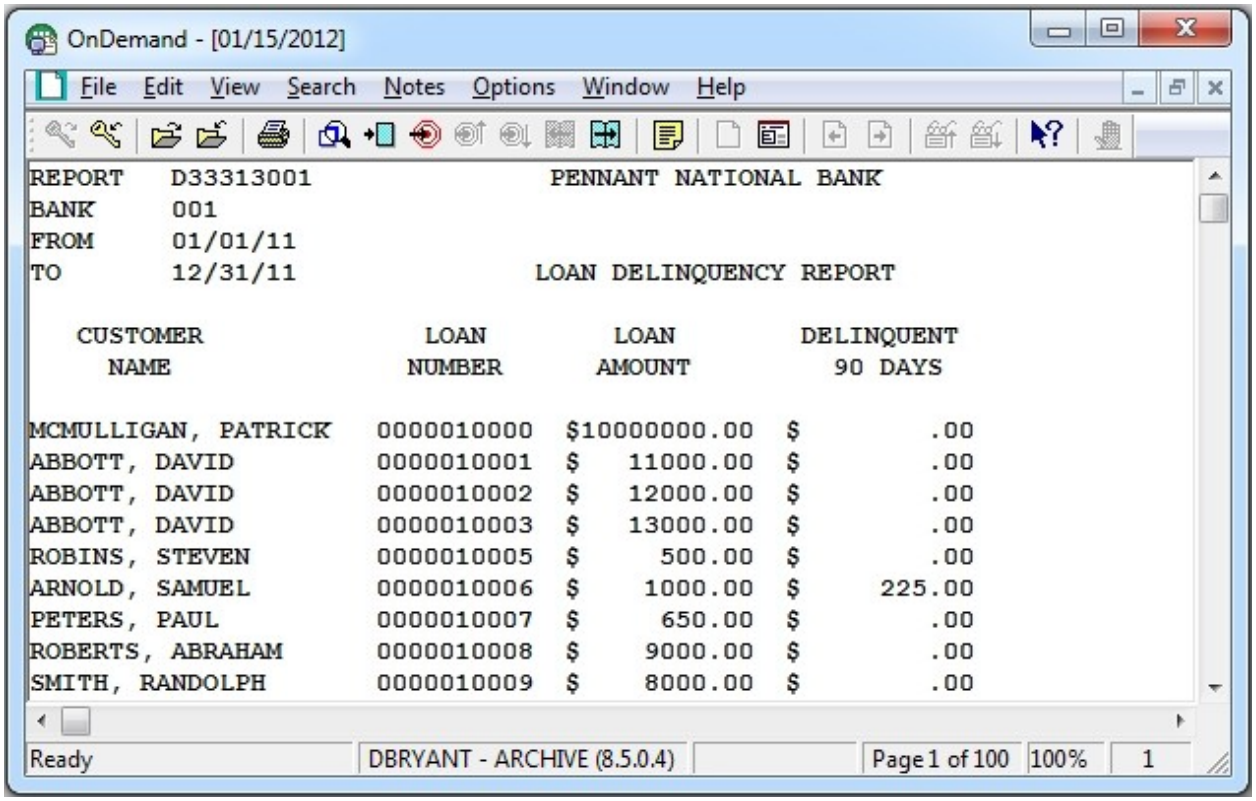

Figure 16. Loan delinquency report using a logical view with 4 logical view fields

Hints and tips

- If you use a page header, it must start with the first line on the page, and you must define it before defining the field header and logical view fields.
- If you use a field header, it must start with the line that immediately follows the last line of the page header, and you must define it before defining logical view fields. If you don't use a page header, the field header must start with the first line on the page.
- If you want to change which lines are used for the page header, you should delete both the page header and field header because the field header will automatically be changed to the first line on the page after the page header is deleted.
- When deciding which string to use for a validation string, select a string that appears in the same place on every page. Pages where the validation string can't be found will be displayed in their original format.
- Define a validation string by selecting a string with 8 characters or less. If more than 8 characters are selected, a validation string cannot be defined.
- When defining logical view fields, you must select text on the line that immediately follows the header lines. If there is a page header and a field header, select text on the line that

 © Copyright IBM Corporation 2012

immediately follows the last line of the field header. If only a page header is defined, select text on the line that immediately follows the last line of the page header.

- When the data is displayed using the Windows client, you can use an option called Expression Find to search for the occurrence of a specified string in one of the columns of data. To use this option, the names of the logical view fields cannot contain a space in the name. For example, if the logical view field that has been defined for the loan number column is "Loan Number", a search cannot be performed for text located in this column of data. If however, the name is "Loan_Number", "LoanNumber", and so forth, a search can be performed.

- When one or more columns of data are removed from the displayed document, the columns can be re-displayed if a different logical view is selected or Reset View is selected from the View menu in the Windows client. This may not be desirable if the columns of data contain sensitive information. To ensure the columns of data cannot be displayed, create a logical view without the columns of data, set the logical view as the default view, and remove the Logical Views permission from the application group permissions for the user or group of users.

Conclusion

As you have seen, creating customized views is as simple as viewing sample data and visually identifying page and field titles and columns of data. Logical views can then be set up to include some or all of the columns of data.

In summary, here are the steps required to create a customized view:

1. If the data contains a page title, define a page header by selecting lines, starting with the first line on the page.

2. If the data contains field titles, define a field header by selecting lines that start with the line that immediately follows the page header or the first line on the page if there isn't a page title.

3. Define the validation string.

4. Create one or more logical view fields.

5. Create logical views that include one or more logical view fields.

This tip is adapted from a developerWorks technical article.

 © Copyright IBM Corporation 2012

Tips – Multiplatforms

How Can I Get a Status Listing of OnDemand Server Threads?

Question

Since OnDemand for Multiplatforms has changed from a process based model to being threaded, how can I get a listing of the thread status for the OnDemand server process?

Cause

In previous versions of OnDemand that used a process based model, a user could issue the command "ps -ef | grep arssockd" to see all the processes associated to the OnDemand Server. That output contained information about what each process was doing. Now that OnDemand is threaded it is not possible to distinguish what the server threads are doing.

Answer

Starting with OnDemand for Multiplatforms Version 8.5.0.3, new command syntax has been added to the arssockd command to return information about the OnDemand Server.

Usage: arssockd [options]

```
Version: 8.5.0.3
-h <od_inst> OnDemand instance name or host name (same as -I)
-I <od_inst> OnDemand instance name or host name (same as -h)
-p Display process usage information for the given instance
-q Display configuration and version information for the given instance
-r <iterations> Number of iterations (defaults to 1)
-s <seconds> Number of seconds between iterations (defaults to 1)
-S Start the OnDemand server for the given instance
-T Stop the OnDemand server for the given instance
-v Verbose output
-x Extended information (when used with -p)
-1 <trace_file> Fully-qualified trace file name
-2 <level> Trace level number
The values are additive. (Default: 3)
1: Errors
2: Warnings
4: Info
8: Flow
```

 © Copyright IBM Corporation 2012

Examples:

To start the server

```
arssockd -I <instance> -S
arssockd
arssockd start <instance>
```

Where instance defaults to ARCHIVE if not specified

To stop the server

```
arssockd -I <instance> -T
arssockd stop <instance>
```

Where instance defaults to ARCHIVE if not specified

To display configuration information of the instance

```
arssockd -I <instance> -q
```

To display process usage information

```
arssockd -I <instance> -p
  PID TID START TIME CPU MEM STYPE USERID INFO
  1138878 - 07/04/11 04:09:12 1:25.049844 36724 Program - ARCHIVE
  1138878 1 07/04/11 04:09:13 0:1.790156 - Main - Accepting
  1138878 258 07/04/11 04:09:14 0:0.356935 - License - Cur(0)
  1138878 515 07/04/11 04:09:15 0:10.630209 - DB - Idle
  1138878 772 07/04/11 04:09:15 0:9.086973 - DB - Idle
  1138878 1029 07/04/11 04:09:15 0:15.093788 - DB - Idle
  1138878 1286 07/04/11 04:09:15 0:14.202843 - DB - Idle
```

where

PID is the process id

TID is the thread id

START TIME is when the process/thread was created

CPU is the amount of CPU in MINS:SECS.MICROSECS spent for the process/thread

MEM is the amount of memory used by the process/thread

USERID is the name of the user currently active on the thread

INFO is additional information used to display current activity of the process/thread. For

 © Copyright IBM Corporation 2012

process this will be the name of the OnDemand instance

Note that the process status does not display all threads, only those we consider to be relevant. The CPU and MEM columns are informational only and are not always accurate at a given point in time. They are updated only when the update will not effect performance.

To monitor process usage information

```
arssockd -I <instance> -p -r 5 -s 5
```

Output similar to that shown above will be collected 5 times at 5 second intervals.

This tip is adapted from support item 1577376.

Defaulted Values for Added Columns Can Cause Uninitialized Memory to be Read

Upgrading from a previous version of OnDemand to Version 8.5 can cause uninitialized memory to be read from database columns previously added during migration steps. There is no data loss, however invalid values can get populated into the OnDemand Server structures. The following symptoms are observed:

1. Error messages when trying to read an invalid length from a resource file from cache
2. The OnDemand server arssockd.exe file crashes
3. Users are forced to reset their password
4. An object stored in cache can have an incorrect expiration value

This problem only effects customers running OnDemand 8.5 (including all fix packs up to and including 8.5.0.3) on a Windows platform, using DB2 for their OnDemand database, and that are upgrading their OnDemand Server instance(s) that were created prior to 8.4.1. Any OnDemand instance created with 8.4.1 or 8.5 does not have this issue.

The problem is due to OnDemand Version 8.5 reading uninitialized memory from columns previously added to the OnDemand System Tables during upgrades from OnDemand instances created prior to version 8.4.1.

The issue was due to the OnDemand server not properly handling NULL database values in tables which resulted in OnDemand assuming the read value was accurate, when in actuality the value was read from uninitialized memory.

Customers running OnDemand Version 8.5 on Windows with DB2 should download and install the PM48739 hot fix from FixCentral that is now posted as part of the ODMP fix pack 8.5.0.3.

This tip is adapted from support item 1516916.

 © Copyright IBM Corporation 2012

Tips – z/OS

OnDemand for z/OS Quick Hits

Still running DB2 for z/OS Version 8?

Just a reminder that End of Support for DB2 for z/OS Version 8 is April 30th, 2012.

See US announcement letter ENUS910-169 for details.

Would you like a sample of how to run the 8.5.0.4 ACTION HOLD in batch? Or samples for executing other OnDemand for z/OS functions in batch?

The OnDemand User Group forum section on Tips and Tricks is fast becoming a repository for executing various OnDemand for z/OS functions in batch.

Be sure to check in there regularly for samples and updates.

The OAM parameter Maximum Object Size - what are the pros and cons of raising the limit?

The Maximum Object Size (MOS) that could be stored in OAM when it was first written was 50 Megabytes. That size remains the default to this day. In 2004 the maximum possible size was raised to 256 Megabytes. IBM Software Support is seeing more and more customers hitting the 50 Megabyte limit as the size of objects being stored creeps upward.

Q: When this parameter needs to be raised, an IPL is required, so should everyone just increase the value of this parameter now?

A: Probably not.

If you see the following in an error message:

```
SM Error: ARSMVSST: 0000000b(00000008-24020202)
```

Check support item S1002890. The 24020202 is an indication that the MOS value needs to be increased.

In some cases a loadfile/index object blob (Binary Large OBject) which contains all the index values for a report can get very large for large reports with an abundance of index values.

In either of the above cases opening a PMR to IBM Software Support is recommended.

However, if you have a single document that even after compression is greater than 50 Mb, then further investigation is required. This could also be a sign that indexing parameters have been set up incorrectly and a report is unintentionally being stored as one large document. Also, there is a virtual storage consideration of not being able to simultaneously load a number of these large objects into the ARSSOCKD address space.

So consider carefully whether or not you want to pro-actively raise the MOS limit.

Note that in z/OS 1.12 the maximum object size has been raised to 2000 Megabytes, but OnDemand for z/OS still only supports up to 256 Megabytes.

 © Copyright IBM Corporation 2012

Migrating From z/OS 1.11 and Below to 1.12 and Above, Change in Job Output

Do you use any ACIF exits? A change in z/OS 1.12 may require a change.

z/OS 1.12 changed the output messages at the end of every job. Messages IEF374I and IEF376I have been replaced by messages IEF032I and IEF033I, respectively, which provide more information.

z/OS 1.11

```
IEF142I BPXAS BPXAS - STEP WAS EXECUTED - COND CODE 0000
IEF373I STEP/        /START 2011208.0000
IEF374I STEP/        /STOP  2011208.0030 CPU    0MIN 00.00SEC SRB    0MIN 00.00S
IEF375I  JOB/BPXAS   /START 2011208.0000
IEF376I  JOB/BPXAS   /STOP  2011208.0030 CPU    0MIN 00.00SEC SRB    0MIN 00.00S
```

z/OS 1.12

```
IEF142I BPXAS BPXAS - STEP WAS EXECUTED - COND CODE 0000
IEF373I STEP/        /START 2011208.1128
IEF032I STEP/        /STOP  2011208.1218
        CPU:     0 HR  00 MIN  00.00 SEC    SRB:     0 HR  00 MIN  00.00 SEC
        VIRT:     0K  SYS:     0K  EXT:        0K  SYS:        0K
IEF375I  JOB/BPXAS   /START 2011208.1128
IEF033I  JOB/BPXAS   /STOP  2011208.1218
        CPU:     0 HR  00 MIN  00.00 SEC    SRB:     0 HR  00 MIN  00.00 SEC
```

The ARSSPVIN shipped with OnDemand keyed off the IEF376I message, so OnDemand for z/OS development shipped a new sample exit for this change:

```
PM44732 -
  ************************************************************
  * USERS AFFECTED: All V8 OnDemand for z/OS Users.          *
  ************************************************************
  * PROBLEM DESCRIPTION: ENHANCE ARSSPVIN SAMPLE TO SUPPORT  *
  *                      IEF033I MESSAGE.                    *
  ************************************************************
  * RECOMMENDATION:                                          *
  ************************************************************
  Message IEF033I replaced IEF376I in z/OS 1.12, new message
  added to sample ACIF exit.
```

 © Copyright IBM Corporation 2012

OS/390 Indexer Usage Notes: The INDEXSTYLE Parameter - Part 1 of 4

This is the first of a four part series on the use of the INDEXSTYLE indexing parameter in the OS/390 Indexer. The OS/390 Indexer is available on the z/OS platform. This part covers the INDEXSTYLE=NODX usage. Later parts will cover the PDOC, PAGE and AFP usages. Refer to the Indexing Reference manual, Part 5 OS/390 Indexer Reference, Chapter 29 OS/390 Indexer Parameters, for more information on the INDEXSTYLE indexing parameter.

INDEXSTYLE=NODX

Application Groups having the NODX (no index) INDEXSTYLE are ones which either do not have obvious index values, or which are very short and do not need to be broken up into documents. The first three logical indexes must be Segment Number, Report Date and Page Number. The Segment Number is calculated by the capture program, and is just a sequential number for each document stored. The Report Date is the run date in the format of MM/DD/YY, stored as a string. The Page Number is calculated by the capture program, and is the first page number in each document.

The Page Number is a transaction (or range) index. The Application Group must be defined with an additional index value for the last Page Number in each document. The Folder Mapping definitions need to be set just right to get the Page Number searches to work as desired.

Additional index values can be defined, but must come after these in the Application Group definition.

Note that in the following examples parameters and values not relevant to understanding the INDEXSTYLE parameter are replaced with an ellipsis (…).

Application Group definition

The following summarizes the definition of a NODX Application Group.

```
APPLICATION GROUP NAME:  NODXMG  01

  *** GENERAL ***
    DESCRIPTION:                                NODX TEST
    …

    *** APPLICATIONS ***
    APPLICATION NAME:  NODXMG  01

  *** FIELD INFORMATION ***

    NAME:  SEGMENT_NUMBER
      SEGMENT:                 No
      DATA TYPE:               Integer
      TYPE:                    Index
      …
```

 © Copyright IBM Corporation 2012

```
NAME:  REPORT_DATE
  SEGMENT:                No
  DATA TYPE:              String
  TYPE:                   Index
  …

NAME:  PAGE_NUMBER
  SEGMENT:                No
  DATA TYPE:              Integer
  TYPE:                   Index
  …

NAME:  PAGE_NO_END
  SEGMENT:                No
  DATA TYPE:              Integer
  TYPE:                   Index
  …

NAME:  POSTING_DATE
  SEGMENT:                Yes
  DATA TYPE:              Date
  TYPE:                   Filter
  …

NAME:  VERSION
  SEGMENT:                No
  DATA TYPE:              String
  TYPE:                   Filter
  APPLICATION ID FIELD:   Yes
  …
```

Notice that the Posting Date field is defined as the Segmentation Field and the Version field is defined as the Application ID Field.

The SEGMENT_NUMBER field is defined as an integer. This facilitates the searching by segment number.

The Page Number field is actually defined as two fields. We must keep track, for our purposes, of both the first page number and the last page number in each document. Both the PAGE_NUMBER and PAGE_NO_END are defined as integers.

 © Copyright IBM Corporation 2012

Application Definition

The following summarizes the definition of a NODX Application.

```
APPLICATION NAME:  NODXMG  01

  *** GENERAL ***
    APPLICATION GROUP NAME:  NODXMG  01
    IDENTIFIER:              01
    DESCRIPTION:             NODX TEST

    *** INDEXER PARAMETERS ***

    INDEXER:  OS/390

  CC=YES
  CCTYPE=A
  CONVERT=NO
  MCF2REF=CPCS
  CPGID=500
  FILEFORMAT=RECORD,90
  GROUPMAXPAGES=50
  TRIGGER1=*,1,X'F1',(TYPE=GROUP)  /* 1 */
  FIELD1=0,83,8,(TRIGGER=1,BASE=0)
  INDEX1=X'E2C5C7D4C5D5E36DD5E4D4C2C5D9',FIELD1,(TYPE=GROUP,BREAK=NO)
      /* SEGMENT_NUMBER */
  INDEX2=X'D9C5D7D6D9E36DC4C1E3C5',FIELD1,(TYPE=GROUP,BREAK=NO)
      /* REPORT_DATE */
  INDEX3=X'D7C1C7C56DD5E4D4C2C5D9',FIELD1,(TYPE=GROUPRANGE)
      /* PAGE_NUMBER */
  INDEX4=X'D7D6E2E3C9D5C76DC4C1E3C5',FIELD1,(TYPE=GROUP,BREAK=NO)
      /* POSTING_DATE */
  INDEXSTYLE=NODX
  …

  *** LOAD INFORMATION ***
  …

    APPLICATION GROUP DB NAME:  SEGMENT_NUMBER
    …

    APPLICATION GROUP DB NAME:  REPORT_DATE
    …

    APPLICATION GROUP DB NAME:  PAGE_NUMBER
    LOAD ID NAME:               PAGE_NUMBER
    …
```

 © Copyright IBM Corporation 2012

```
APPLICATION GROUP DB NAME:  PAGE_NO_END
LOAD ID NAME:               PAGE_NUMBER
...

APPLICATION GROUP DB NAME:  POSTING_DATE
...
```

In the INDEXER PARAMETERS section are the following INDEXn values:

```
INDEX1=X'E2C5C7D4C5D5E36DD5E4D4C2C5D9',FIELD1,(TYPE=GROUP,BREAK=NO)
    /* SEGMENT_NUMBER */
INDEX2=X'D9C5D7D6D9E36DC4C1E3C5',FIELD1,(TYPE=GROUP,BREAK=NO)
    /* REPORT_DATE */
INDEX3=X'D7C1C7C56DD5E4D4C2C5D9',FIELD1,(TYPE=GROUPRANGE)
    /* PAGE_NUMBER */
INDEX4=X'D7D6E2E3C9D5C76DC4C1E3C5',FIELD1,(TYPE=GROUP,BREAK=NO)
    /* POSTING_DATE */
```

Note the following:

The values in hexadecimal exactly match the Application Group field names. This is not a requirement, but it does help simplify things.

The TYPE value for the PAGE_NUMBER index is set to GROUPRANGE. This indicates that this is a range index, which will be mapped to two Application Group Data Table columns.

Another Indexer Parameter is GROUPMAXPAGES. This specifies the maximum number of pages to put into each document.

In the LOAD INFORMATION section, the Application Group DB Names are the Name field values from the Application Group definition, and will become the DB2 Column Names for the Application Group Data Table. The Load ID Name is the name referenced in the INDEXn= value (entered in hexadecimal). If no Load ID Name is specified, it defaults to the Application Group DB Name.

In this example, the Application Group DB Name values of both PAGE_NUMBER and PAGE_NO_END are mapped to the same Load ID Name of PAGE_NUMBER. While there are four INDEXn parameters, the capture program actually generates five index values - the fifth being what gets mapped to the PAGE_NO_END Application Group DB Name.

The TRIGGER1 parameter is used by the OS/390 Indexer to establish the “start of report”. The FIELD1 parameter is used to locate the Posting Date. The references to FIELD1 on the other INDEXn parameters is for syntax checking purposes only, and is ignored by the OS/390 Indexer.

 © Copyright IBM Corporation 2012

Folder Definition

The following summarizes the definition of a NODX folder:

```
FOLDER NAME:  NODXMG FOLDER

  *** GENERAL  ***
    DESCRIPTION:                NODX TEST
    NOTE SEARCH:                Retrieve
    DISPLAY DOCUMENT LOCATION:  No
    DISPLAY DOCUMENT HOLD:      No
    APPLICATION GROUP NAME:  NODXMG  01
      APPLICATION NAME:  NODXMG  01

  *** FIELD INFORMATION ***

    ID:  *PUBLIC

      NAME:              SEGMENT NUMBER
      ...

      NAME:              REPORT DATE
      ...

      NAME:              PAGE NUMBER
      TYPE:              Integer
      MAPPING TYPE:      Single
      QUERY:             0
      HIT LIST:          4
      ...

      NAME:              POSTING_DATE
      ...

      NAME:              VERSION
      ...

      NAME:              PAGE NUMBER2
      TYPE:              Integer
      MAPPING TYPE:      Single
      QUERY:             0
      HIT LIST:          5
      ...

      NAME:              PAGE NUMBER.
      TYPE:              Integer
      MAPPING TYPE:      Range
```

 © Copyright IBM Corporation 2012

```
    QUERY:              3
    HIT LIST:           0
    ...

*** FIELD MAPPING  ***

  FOLDER FIELD NAME :: APPLICATION GROUP NAME: DATABASE NAME

  SEGMENT NUMBER :: NODXMG  01: SEGMENT_NUMBER
  REPORT DATE    :: NODXMG  01: REPORT_DATE
  PAGE NUMBER    :: NODXMG  01: PAGE_NUMBER
  POSTING_DATE   :: NODXMG  01: POSTING_DATE
  VERSION        :: NODXMG  01: VERSION
  PAGE NUMBER2   :: NODXMG  01: PAGE_NO_END
  PAGE NUMBER.   :: NODXMG  01: PAGE_NUMBER
  PAGE NUMBER.   :: NODXMG  01: PAGE_NO_END
```

Now, things get interesting. Back on the Application Group definition, we had two fields defined for Page Number: PAGE_NUMBER to hold the first page number value in each document, and PAGE_NO_END to hold the last page number value in each document.

Now we have three Page Number fields in the Folder. These field names are what the customers see on their displays. You can name them anything you want.

PAGE NUMBER - only appears on the hit list, not on the search screen. This field is mapped to the PAGE_NUMBER field of the Application Group.

PAGE NUMBER2 - only appears on the hit list, not on the search screen. This field is mapped to the PAGE_NO_END field of the Application Group.

PAGE NUMBER. - (note the period after NUMBER) only appears on the search screen, not on the hit list. This field must be mapped to both the PAGE_NUMBER and PAGE_NO_END fields. When you define this field, the Mapping Type must be set to RANGE.

Summary of NODX fields

Application Group:

```
Database field1 = SEGMENT_NUMBER
Database field2 = REPORT_DATE
Database field3 = PAGE_NUMBER
Database field4 = PAGE_NO_END
Database field5 = POSTING_DATE
```

Application:

```
Index1 = SEGMENT_NUMBER
Index2 = REPORT_DATE
Index3 = PAGE_NUMBER
```

 © Copyright IBM Corporation 2012

```
Index4 = POSTING_DATE

    (index1) SEGMENT_NUMBER →  SEGMENT_NUMBER (DB field1)
    (index2) REPORT_DATE    →  REPORT_DATE    (DB field2)
    (index3) PAGE_NUMBER    →  PAGE_NUMBER    (DB field3)
    (index3) PAGE_NUMBER    →  PAGE_NO_END    (DB field4)
    (index4) POSTING_DATE   →  POSTING_DATE   (DB field5)
```

Folder:

```
Folder Field1 = SEGMENT NUMBER
Folder Field2 = REPORT DATE
Folder Field3 = PAGE NUMBER
Folder Field4 = POSTING_DATE
Folder Field5 = PAGE NUMBER2
Folder Field6 = PAGE NUMBER.

    (DB Field1) SEGMENT_NUMBER →  SEGMENT NUMBER (folder field1)
    (DB Field2) REPORT_DATE    →  REPORT DATE    (folder field2)
    (DB Field3) PAGE_NUMBER    →  PAGE NUMBER    (folder field3 - Hitlist)
    (DB Field4) PAGE_NO_END    →  PAGE NUMBER2   (folder field5 – Hitlist)
    (DB Field3) PAGE_NUMBER    →  PAGE NUMBER.   (folder field6 - Query)
    (DB Field4) PAGE_NO_END    →  PAGE NUMBER.   (folder field6 - Query)
    (DB Field5) POSTING_DATE   →  POSTING_DATE   (folder field4)
```

 © Copyright IBM Corporation 2012

Tips – IBM i

Important Upgrade Considerations for OnDemand for i

Customers with more than one OnDemand instance defined

The first time you start your OnDemand servers after a server upgrade or new release installation, you should start each instance's server individually. After the first server starts, test the storing and viewing of data to ensure that no problems exist. Then start a second server, and perform the same tests. If you have a test instance and a production instance defined, you might consider working with the test instance first. This approach will greatly reduce the time it takes to resolve any issues that might arise that are unique to your environment.

If the OnDemand server fails to start

The OnDemand server might fail to start after applying server upgrade PTFs if you have non-IBM files in the instance library that are not owned by the instance user profile

When upgrading to a new OnDemand server version, such as server version 8.4.1.7, the server might not start after applying the server upgrade PTFs if you have unexpected files in the instance library. If you have added a file to the instance library with a file name that begins with 'ars' but that file is not owned by the instance user profile, the server will fail to start.

For example, if you use an OnDemand instance named QUSROND, the instance library and instance user profile are also named QUSROND. You might make a backup copy of the arsapp file and call it arsappbk and place it in the QUSROND library. The backup file might be owned by the user profile you were signed on with when you created it, for example MYUSER. In this example, the QUSROND instance server will fail to start after applying the server upgrade PTFs.

The following error messages will be found in the job log of the server that failed to start, where QUSROND is the name of the instance being started:

```
OND2798 arsdb: CHGAUT for library QUSROND failed with exception CPF223A
OND2798 arsdb: CHGAUT *PUBLIC for library QUSROND failed with exception
   CPF223A
OND2798 arsdb: CHGOWN for library QUSROND failed with exception CPF223A
OND2798 DATABASE CHANGES FAILED FOR INSTANCE QUSROND. RETURN CODE -001. SEE
   PREVIOUS MESSAGES IN THE JOB LOG FOR MORE INFORMATION
```

To correct the problem, look in the instance library for file names beginning with 'ars' that are not owned by the instance user profile. If you find a file beginning with 'ars' that is not owned by the instance user profile, do one of the following:

- move the file to a different library
- delete the file if it is no longer needed
- change the owner of the file to be the instance user profile

After one of these steps is completed, restart the server.

This tip is adapted from support item 1515197.

 © Copyright IBM Corporation 2012

Resolving Media Migration Facility Errors OND0524 and CPF4160

Question

How do I resolve error messages OND0524 and CPF4160 while running the Media Migration Facility (MMF)?

Cause

The error messages indicate that program QRLRTACT is not found. This occurs if OnDemand program product option 1 (Spool File Archive) has been removed from the system.

Answer

After migrating from Spool File Archive to Common Server, when using MMF to then migrate the actual data that was originally stored into Spool File Archive into the Common Server Archived Storage Manager (ASM), you might receive message OND0524 and CPF4160.

For example, running this command at version 6.1 or above:

```
MGRMEDRDAR SRC(CHECKSTMTS) SRCTYPE(*REPORT) TGT(*ASM) POLICY(LONGTERM)
   INSTANCE(QUSROND)
```

or calling this program at version 5.4:

```
CALL QRDARS/QRLCSFAMMF PARM(CHECKSTMTS *REPORT *ASM QUSROND LONGTERM)
```

might result in the following message:

```
OND0524 SQL error occurred in procedure: updateFiles at statement:
   0000000064.
```

If this message occurs, look for a job log for job QRLCSFADT1 that contains the following message:

```
CPF4160 Program QRLRTACT in QRDARS not found.
```

The reason for the errors is that, once the report data has been successfully moved into ASM in Common Sserver, the report data is deleted from its original location in Spool File Archive. The job deleting the report data receives the error message that program QRLRTACT is not found if OnDemand program product option 1 (Spool File Archive) has been removed from the system. Many customers remove Spool File Archive from their system after migration from Spool File Archive to Common Server has been completed. Option 1 might also be removed during an upgrade from 5.4 to 6.1 or 7.1.

To resolve this problem, simply run the following command to disable the trigger program on the file QARLRACT in library QUSRRDARS:

```
CHGPFTRG FILE(QUSRRDARS/QARLRACT) TRG(*ALL) TRGLIB(QUSRRDARS)
   STATE(*DISABLED)
```

This tip is adapted from support item 1574579.

 © Copyright IBM Corporation 2012

Additional Information

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

Information Centers

OnDemand Version 8.5 Information Center

OnDemand Version 8.4.1 Information Center

OnDemand Version 8.4 Information Center

Publication Libraries - Containing all PDF versions of the documentation

OnDemand for Multiplatforms Version 8.5 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 8.5 System Requirements

OnDemand for z/OS Version 8.5 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".