

In This Issue

News	1
About this newsletter	1
Server version 9.5.0.7 available	1
Unlocking your dark data archives	1
Tips – Cross Platform	2
Potential issue if multiple administrators are working on the same object type at the same time	2
Why doesn't the PDF indexer extract the images in my document?	3
Tips – z/OS	3
Quick Hits	3
Tips – IBM i	4
Enabling SSL for IBM Navigator for i	4
Updated component of IBM Navigator for i ...	4
How much space is Content Manager OnDemand using?	4
Additional Information	5

News

About this newsletter

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

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

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

This newsletter is formatted so that it is easier to read on wide screen devices. Use the full screen viewing option in Acrobat (Ctrl+L) for best results.

Server version 9.5.0.7 available

Multiplatforms

The fix pack 9.5.0.7 installation files are available from [IBM Fix Central](#). This includes the OnDemand Administrator and OnDemand Windows Client.

z/OS

To upgrade your system, choose the applicable PTF from the list in [support item 1260192](#).

IBM i

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

Release	PTF Group	Info APAR	Read This First
V7.3	SF99252	II14799	7046944
V7.2	SF99251	II14723	7041959
V7.1	SF99250	II14497	7018085

Unlocking your dark data archives

This item is adapted from an article by Brian Phelps, IBM Worldwide Sales Leader – Content Manager OnDemand Archive Solutions.

Sweeping clean your records floor

Have you ever moved homes from one location to another? If you have, you have probably gone through the daunting task of finding all your file cabinets and looking at all of the documents in the folders and deciding to keep them or throw them away. If you keep them, you will either put them into a modern file cabinet that isn't rusting or put them into a locked fireproof box if they are very important like birth certificates, passports, etc. For those documents getting thrown away, you decide to shred some because of personal information on those and for the rest you simply throw those documents away. Now imagine what an organization needs to go through to do the same thing at a corporate level, especially if mergers and acquisitions have occurred.

Why should an organization evaluate their archives that hold documents of record?

Do you even know what is in some of these old archives? If the answer is "no" then this means that you can't purge any records because if you don't know what is in a repository, it may need to be kept per your record retention schedule or it may be a part of an active legal situation (and you don't know it – yet). Don't worry though, the auditors and lawyers will let you know after they have completed their discovery process during a legal procedure.

Finding the keys to unlock your dark data archives:

Key#1 – Take an inventory: This sounds like a simple task but for most organizations it isn't. I have spoken with hundreds of organizations around the world in all industries and know that any decent-sized organization has around 20 archives containing documents of record. It is also common for most of these archives to be end-of-life and often also out-of-support.

Key#2 – Choose your target enterprise archive: Your target enterprise archive must have the following attributes:

- Proven electronic archive with large install base
- Proven for at least a decade with many customer references in your industry
- Active development roadmap with an active user group contributing input to future development
- Open architecture to deliver meaningful data sets to analytics to derive actionable business insights
- Encompasses modern Enterprise Content Management (ECM) capabilities enabling social, mobile, collaboration, enterprise application and data integration out of the box along with being Content Management Interoperability Service (CMIS) compliant and able to interact with other CMIS compliant repositories
- Architected for all content and formats past, present and future. Many of the older archives may have archived documents of record in the form of AFP, individual PDFs, TIFFs, etc. Today companies need to archive XML and EDI data sets as documents of record in addition to communications sent to customers and internal reports. Tomorrow it could be holograms, 3-D genomic sequences, and huge amounts of weather data sent in daily by millions of sensors around the world.
- An inherent capability for records retention and lock-down
- An ability to archive documents of record out of SharePoint/Collaborative and ECM systems once the collaboration and workflow is complete

Key#3 – Choose a partner that understands the old archives and the target archive: This is INCREDIBLY important. The partner must have a long history of migrations with a proven methodology and have developed toolkits for most/all of the old archives. Many of these old archives are listed in the 2006 Gartner Magic Quadrant for IDARS (Integrated Document Archive and Retrieval Systems). The only archive left from this Gartner report and being rapidly developed and deployed today is the IBM Content Manager OnDemand enterprise archive. If your old archives are running on a z/OS environment like BMC Control-D, ASG Mobius, CA View/Deliver/Dispatch, Systemware Xporter, Beta 92/93, IBM RDMS, etc., then you must verify that your migration partner is VERY skilled on the z/OS platform and ask for resumes that prove that fact.

News (continued)

Key#4 – Take the Opportunity to Discover, Verify, Clean-Up and Modernize during the Migration: Most skilled migration partners bring all of the old archives into a common format, such as XML, before moving that information to the target archive in the target format. During this time your organization has a tremendous opportunity to do the following:

- Verify that all records being moved are accurate and unchanged to the bit level.
- Identify the records that can be purged and do not need to be moved.
- Identify records that contain inappropriate PII (Personally Identifiable Information) or full credit card information (PCI-DSS compliance). For those records, the migration partner should have the ability to redact or tokenize as those records are ingested into the target archive.
- Identify information that might also need to feed your data lake. Right now only 12% of corporate information is available for analysis. The remaining 88% live in these dark data archives. In addition, with a modern archive, like IBM's Content Manager OnDemand, all information archived is always available for secured analysis because of its modern and open architecture.
- Choose the format that best suits your organization to make sure that you "future proof" your documents of record while efficiently utilizing resources. XML and PDF using page piece dictionary are two very good choices. If you select XML, the original XML record can be kept unaltered and be associated with different style sheets for viewing such as for mobile devices or for users who might be sight impaired where the style sheet could invoke a reader.

Key#5 – Is This the Right Time to Go to the Cloud?: To the Cloud or not to the Cloud? That is the question. Almost all organizations are trying to answer this question and many have selected a "Cloud first" approach for all new applications. Whether or not you are going to Cloud today, your strategic enterprise archive should have a Cloud deployment option today in case your organization chooses a Cloud option in the future. Not all content is created equal. Some content can go to the Cloud now and some can stay on premise. You should also think about ways to bridge to the Cloud if you feel that a Cloud deployment is in your future. Here are some reasons organizations select a Cloud deployment:

- They want to focus on the business application and new use cases and avoid the hassles of database, app server, hardware and application level upgrades.
- Related to the first reason is the fact that they will always be on the latest releases of code to ensure that they can take advantage of enhancements and new features as they are delivered to the market.

- Rapid time to value. If an organization's IT department is backlogged with projects, a Cloud approach can be provisioned and delivered for a new application in as little as two weeks without the need to take the time of the IT department.
- Operating expense versus capital expense. This is can be an important issue especially in government organizations.
- Offload much of the work associated with failover and disaster recovery.

Click [here](#) for a short video on feeding analytics with IBM Content Manager OnDemand archive as the trusted data source (5min 30sec).

Tips – Cross Platform

Potential issue if multiple administrators are working on the same object type at the same time

Why am I getting "User ID or UID already exists" message when deleting and copying users with the OnDemand Administrator client or ARSXML batch administration program?

The scenario is this:

Two Content Manager OnDemand administrators log into the OnDemand Administrator client, each on their own workstation. Both administrators double-click on 'Users' to work with the list of defined users.

Administrator #1 deletes an existing user.

Administrator #2 copies any user and attempts to add a new user.

The result is that Administrator #2 receives a "User ID or UID already exists" message. The message is sent because the UID number that was used for the new user is the same as the UID number for the deleted user. After the deletion, the UID number becomes available on the server, but the UID number still exists in the user list on the workstation of Administrator #2. This list doesn't contain the latest information from the server so the user that was deleted by Administrator #1 is still in the list. The user list can be refreshed by pressing F5 or by clicking the View menu and selecting Refresh List from the menu when the user list is displayed. However, this does not guarantee that the list contains the latest information if multiple administrators are working with users at the same time.

There is nothing that stops multiple administrators from logging on to the same server with the OnDemand Administrator client at the same time. Multiple administrators can add, update, and delete different objects. The problem occurs when the administrators are trying to work with the same object type at the same time. Also, if two administrators are adding a user

at the same time, each user could get assigned the same UID number since the UID number is determined when the Add a User dialog is displayed. Depending on which administrator clicks OK first (to add the user), the other administrator will get an error message that says the user already exists.

As another example, both administrators generate a list of users on the OnDemand Administrator client. The list of users for each administrator does not contain the user name of NEWUSER. Then, both administrators try to add a new user called NEWUSER. Depending on which administrator clicks OK first (to add the user), the other administrator will get an error message that says the user already exists.

As a final example, Administrator #1 selects USER1 to update. The OnDemand Administrator client queries the server for the attributes of USER1 and displays them in the Update a User dialog. Administrator #2

also selects USER1 to update. The OnDemand Administrator client queries the server for the attributes of USER1 and displays them in the Update a User dialog for Administrator #2. If Administrator #1 makes changes to USER1 and saves the changes by pressing OK before Administrator #2 has selected USER1 for update, Administrator #2 will see the changes to USER1 that Administrator #1 made. However, if Administrator #2 selects USER1 for update before Administrator #1 has pressed OK, the information shown for USER1 will not be the latest information. If Administrator #2 makes changes to USER1, the changes will override what changes were made by Administrator #1.

A similar issue can occur when using the ARSXML batch administration program. If two administrators run ARSXML to add a user at the same time and both input XML files reference the same user ID or UID number, one administrator will add the user and the other administrator will get the message that the user already exists. Similarly, if two administrators run ARSXML to update a user at the same time and both input XML files reference the same user ID, both administrators will update the user. However, the administrator that runs ARSXML last will overwrite the updated information from the administrator that ran ARSXML first.

Unexpected results can occur if multiple administrators are working with any Content Manager OnDemand object type (such as groups, applications, application groups, folders, etc.) at the same time. Therefore, it is recommended that different object types are maintained by different administrators. For example, one administrator can be designated as the user administrator and another administrator can be designated as the report administrator to maintain applications, application groups, folders, and cabinets.

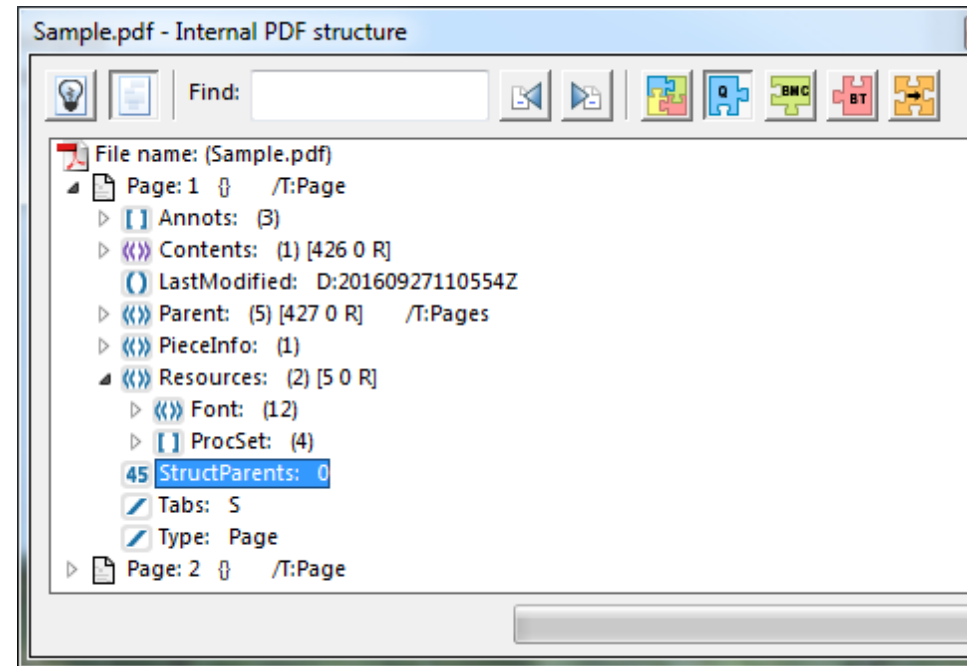
This tip is adapted from [support item 1995148](#).

Why doesn't the PDF indexer extract the images in my document?

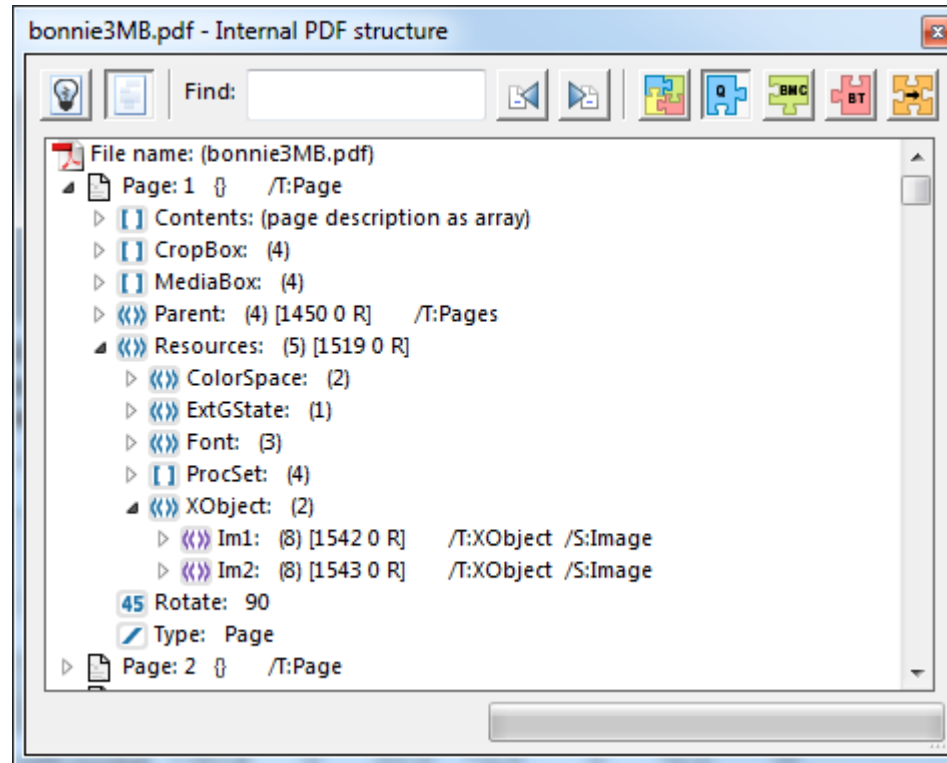
The PDF indexer only extracts images that have a subtype set to Image. Some PDF files contain images that do not have a subtype of image. For example, images can be created with Photoshop and be part of the content stream for the page, as opposed to a separate image resource.

To verify whether your document contains images that the PDF indexer can extract, use Acrobat Professional to display the internal structure of the PDF document. If you have Acrobat X or XI, click Tools (on right side) ->Print Production ->Preflight. Click Options in the upper right corner of the dialog box, then click Browse internal PDF structure. If you have Acrobat DC, type "Preflight" into the Search Tools box on the right side and then click Preflight to bring up the dialog box.

Here is a screen shot of the internal structure of a document that contains images that do not have a subtype of Image:



And here is a screen shot of another file which contains the images that the PDF indexer can extract. Note that the images have a Type of XObject and a Subtype of Image.



If you want the PDF indexer to be able to extract the images in your document, talk to your PDF producer about creating the correct type of images.

Tips – z/OS

Quick Hits

Hierarchy of userid and password checking by ARSLOAD

The ARSLOAD program first looks to see if there is a userid and password specified with the ARSLOAD command. If there is not, then it tries to use the unified login exit. If it does not find the exit, then it looks for a stash file.

It's time to upgrade to a new server release - How do you plan your Content Manager OnDemand client software upgrades?

The first thing you should do is check the compatibility matrix for the Content Manager OnDemand clients and servers available in [support item 1392275](#).

Note that the newer clients are backwards compatible with older releases of the server so upgrading clients before you upgrade your server is a good idea.

Is there a way to inventory what clients are being used to access a given Content Manager OnDemand server?

Yes! The ARSRPT program to the rescue!

Here's sample JCL to run the ARSRPT program:

```
//TMP1 EXEC PGM=IKJEFT01,
//          DYNAMNBR=200
//SYSPROC DD DSN=SYS1.SBPXEXEC,DISP=SHR
//*
//SYSTSPRT DD SYSOUT=*
//*
//SYSTSIN DD *
  oshell logger -d1 starting arsrpt run
  oshell /usr/lpp/ars/V9R5M0/bin/arsrpt $A1 $A2
  oshell logger -d1 ending arsrpt run
/*
//STDENV DD *
  _BPX_SHAREAS=YES
  _BPX_BATCH_SPAWN=YES
  A1= -h ARCH950 -u myuserid -p mypw
  A2= -t 2016-11-16 -d /tmp/somewhere -o anynameyouwant
/*
//OSHOOT1 DD SYSOUT=*,DCB=(RECFM=F,LRECL=255)
//STDOUT DD SYSOUT=*
//STDERR DD SYSOUT=*
```

Change the -h to your instance name, -u and -p are self-explanatory, -t is a particular date, -d is a writeable directory and -o is the output file name.

Sample output from the ARSRPT program:

Client	Version	Operating System	Logins	SSL Logins	Failed Logins	SSL Failed
ADMIN GUI	9.5.0.4	Windows	0	1	0	0
ADMIN GUI	9.5.0.5	Windows	1	0	0	0
ADMIN GUI	9.5.0.7	Windows	4	0	0	0
ARSDOC	9.5.0.7	z/OS	9	0	0	0
ARSLOAD	9.5.0.7	z/OS	2	0	0	0
ARSMANT	9.5.0.7	z/OS	1	0	0	0
CLIENT GUI	9.5.0.4	Windows	0	4	0	0
CLIENT GUI	9.5.0.7	Windows	2	0	0	0
ODF	9.5.0.7	z/OS	7	0	0	0
ODWEK JAVA API	9.5.0.7	z/OS	4	0	0	0

Recommended DB2 service when running Content Manager OnDemand

Level 2 software support keeps a running thread out on the OnDemand User Group website called [Recommended DB2 on z/OS Service, V10 and V11](#).

Be sure to check it whenever performing a DB2 upgrade!

Tips – IBM i

Enabling SSL for IBM Navigator for i

Problem

After installing the HTTP Group PTF from December 2016 (see table below), you will need to manually configure IBM Navigator for i to use SSL.

Release	Group PTF	Level
V7.3	SF99722	5
V7.2	SF99713	18
V7.1	SF99368	44

URL change

If you use the URL `http://hostname:2001` to access IBM Navigator for i, the URL will now resolve to `http://hostname:2004/ibm/console/logon.jsp` (non-SSL logon). If you use the URL

`http://hostname:2005/ibm/console/logon.jsp` to access IBM Navigator for i, you will be unable to connect to Navigator until you enable SSL.

Resolving the problem

You can enable SSL (https) by either using the Keystore created by the Level 1 Console (L1C) or by using the Digital Certificate Manager.

See [support item N1021726](#) for details on enabling SSL. Note that if you choose not to enable SSL, you can still use IBM Navigator for i without SSL.

Updated component of IBM Navigator for i

The Content Manager OnDemand component of IBM Navigator for i has been updated on V7.1, V7.2, and V7.3 to resolve a possible authority problem.

The following message might be received in some cases when using the Content Manager OnDemand component of IBM Navigator for i to create a migration policy, even though the user has the proper authority:

```
You do not have authority to work with Content Manager
OnDemand storage sets and storage nodes that are
associated with migration policies in this instance.
```

The authority check for creating a migration policy is corrected by the PTFs listed in the table below.

In addition, the Content Manager OnDemand user ID that is associated with your IBM i user profile (which have identical names if you have not disabled the default security exit) must now have System Administrator authority to perform any of the operations provided by the Content Manager OnDemand component of IBM Navigator for i.

If you have disabled the default security exit such that your Content Manager OnDemand user ID and your IBM i user profile are not linked to each other, you need to ensure that a Content Manager OnDemand user ID with the same name as your IBM i user profile exists, and has System Administrator authority, in order to perform any of the operations provided by the Content Manager OnDemand component of IBM Navigator for i.

Prior to applying this PTF, System Administrator authority was only required to create migration policies.

Release	PTF
V7.3	SI63535
V7.2	SI63534
V7.1	SI63533

How much space is Content Manager OnDemand using?

Question

I am trying to determine how much disk space is currently consumed by my Content Manager OnDemand for i installation. I understand libraries and IFS directories are both involved in storing report data and indexes. Is there an accurate way to determine the current space taken up so I can do some forecasting for the future?

Answer

Yes, there are a couple of methods that can be used to calculate the space taken for the Content Manager OnDemand for i implementation on your system. The product uses several libraries and IFS paths to store programs and data, and the number of instances you have created will impact the number of libraries and IFS paths involved in the calculation. These product components are listed as follows:

- QRDARS – Program library
- QUSRRDARS – Shared user data library
- QUSROND – Instance library with instance specific data
Note: While QUSROND is referred to as the default instance name, you may have used a different instance name and may have more than one instance on the system. To determine all the instances on a system, run the command: `WRKOBJ OBJ(*ALL/ARSAG) OBJTYPE(*FILE)`. All libraries listed are instance libraries. When checking instance libraries or directories, use each library listed in place of or in addition to QUSROND.
- /QIBM/ProdData/OnDemand/ – Program objects
- /QIBM/UserData/OnDemand/ – Data for all instances, including shared data such as configuration files

- /QIBM/UserData/RDARS/ – If you previously used Spool File Archive at V5.4 or earlier, and depending on how you migrated your data, there might still be data in this path.
1. In order to calculate the space consumed by the libraries involved, run the following Display Library (DSPLIB) command for each library. You should have all object (*ALLOBJ) special authority to ensure that all objects are included in the totals.

```
DSPLIB LIB(QUSROND) OUTPUT(*PRINT)
```

Go to the bottom of the resulting QPDSPLIB spooled file and check the total size, which will be similar to "Total size: 140922112". In this case, the instance library takes up 140Mb.

2. In order to calculate the space used by your Content Manager OnDemand for i reports in the IFS, you should use the Retrieve Directory Information (RTVDIRINF) command.

```
RTVDIRINF DIR('/QIBM/UserData/OnDemand') SUBTREE(*ALL)
INFFILEPFX(*GEN) INFLIB(QUSRSYS)
```

You will receive messages similar to:

```
Database files QAEZD00010 and QAEZD0001D created in
library QUSRSYS.
```

```
Directory information was successfully retrieved.
```

You can then query the first file, QUSRSYS/ QAEZD00010 in our example. Field QEZALCSIZE contains the allocated size of the object, and QEZDTASIZE contains the data size of the object.

```
SELECT SUM( QEZALCSIZE ) AS TOTAL_ALLOCATED, SUM( QEZDTASIZE
) AS TOTAL_DATA_SIZE FROM QUSRSYS/QAEZD00010
```

```
TOTAL_ALLOCATED          TOTAL_DATA_SIZE
21,687,711,744           21,632,151,940
```

Both numbers are about the same, but to be conservative, we suggest that you use the "Total Allocated" figure since it is the highest.

You can also run Print Directory Information (PRTDIRINF) over the files created by RTVDIRINF.

```
PRTDIRINF RPTTYPE(*DIR) INFFILEPFX(*LAST) INFLIB(*LAST)
SUBTREE(*ALL)
```

```
Database files QAEZD00010 and QAEZD0001D used from library
QUSRSYS
```

```
Printer output created.
```

The output from PRTDIRINF is in spooled file QPEZDIR. Within the spooled file, locate directory /QIBM/UserData/OnDemand, then check the 'Total space used by objects in this directory and subdirectories'.

There are some restrictions for the RTVDIRINF and PRTDIRINF commands:

- Directory information can be retrieved only for mounted file systems.
- Directory information can be retrieved only for local file systems.
- You must have all object (*ALLOBJ) special authority to run these commands.

You should perform RTVDIRINF for each of the IFS paths listed above. Paths without active instance data, such as the /QIBM/UserData/RDARS directory, should not be growing. The size of the RDARS directory should be stable or declining (as expiration occurs).

If you want to check the directory size for a specific instance, run the RTVDIRINF command with path /QIBM/UserData/OnDemand/QUSROND where QUSROND is the name of your instance.

The directions above do not include any space used by objects on virtual optical or virtual tape, or objects on network file systems (NFS).

This tip adapted from [support item 1698970](#).

Additional Information

Knowledge Centers

Content Manager OnDemand for Multiplatforms [Knowledge Centers](#)

Content Manager OnDemand for z/OS [Knowledge Centers](#)

Content Manager OnDemand for i [Knowledge Centers](#)

Publication Libraries - PDF versions of the documentation

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

Product System Requirements

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

More Enterprise Content Management web sites

Content Manager OnDemand [Product Overview](#)

[Compatibility Matrix](#) for the Content Manager OnDemand clients and servers

Content Navigator [Product Overview](#)

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

Social Media

Follow [@IBM_ECM](#) on twitter for almost real-time news & updates on Content Manager OnDemand. For more information and links to all social media sites, see [support item 1673183](#).

OnDemand User Group

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

Copyright and trademark information

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.

A current list of IBM trademarks is available on the Web at "[Copyright and trademark information](#)".