


Db2 Web Query V2.2

High-Availability and Disaster Recovery Guide



Preface.....	3
Background.....	3
Components of Web Query HA/DR.....	4
Reporting on Data that Resides in an iASP	6
Moving Web Query Data onto an iASP	6
Setting-up the Primary	7
Step 1. Install Web Query and apply PTFs.....	7
Step 2. Start Web Query, verify it's working, and end it.....	7
Step 3. Move Web Query data onto the iASP	7
Step 4. Set up the Administrative Domain.....	8
Step 5. Edit job descriptions	9
Step 6. Start Web Query	10
Setting-up the Backup.....	10
Switching from Primary to Backup	11
Best Practices	12
References.....	12

Preface

This document explains the components and best practices for installing Web Query in an iASP-based cluster environment to ensure high-availability or disaster recovery for your business-critical reports and business intelligence (BI) analytics. A high-availability (HA) environment is defined as one in which the primary and secondary nodes of the cluster switch roles on a regular and sustained basis. A disaster recovery (DR) environment is differentiated as one that provides the recovery or continuation of vital data and operation following a disaster, such as an earthquake or fire. In the context of this document, high availability and disaster recovery are the same and the terms will be used interchangeably.

Background

IBM and independent service providers (ISVs) offer customers a choice of high-availability and disaster recovery products for IBM i. These products provide multi-system or multi-partition solutions that ensure data resiliency and business continuity. Clustering is the underlying infrastructure for most environments because it allows data to be replicated or switched between systems. In the event of an outage, clustering enables critical resources to be automatically available on backup systems.

Independent auxiliary storage pools (iASPs) are one of the key elements of an IBM i clustered environment. iASPs are user collections of disk units that are grouped together to form a single pool of storage. What makes them unique is that they can be used on a single system or switched between multiple systems or logical partitions, independently of other basic system auxiliary storage. They do require application changes in order for the application to see the name-space provided by the iASP and thereby to access the objects residing on the iASP. The act of making the contents of an iASP accessible is known as varying on the iASP.

You can configure an iASP to be used across multiple systems. This is referred to as a switchable iASP. When you switch an iASP to another system, the entire contents of the iASP can be accessed from that system. Everything that is contained in the iASP is switched from being owned by the first system to being owned by second system.

One consideration regarding high-availability solutions built upon independent disk pools is that not all objects can reside within an iASP. There are certain classes of objects that may not reside within an iASP, such as user profiles, job descriptions and authorization lists. These objects must reside in the system ASP.

IBM PowerHA for i is IBM's high-availability product offering. It supports multi-system iASP-based clusters in both mirrored and switched disk environments. Web Query can be installed in either of these environments, as well as more traditional logical replication environments. The differences between these environments are briefly described below.

Switched disk:

iASPs are centric to IBM i switched disk clustered environments. You can switch access to an iASP between the participating nodes within a cluster when required. You can switch the disks between two different systems or two logical partitions on the same system.

Mirroring:

There are several IBM i supported mirroring technologies that you can use to achieve disaster recovery and high availability. Geographic mirroring provides a high availability solution where your production iASP data is mirrored to a backup iASP that is attached to another, remote system. Geographic mirroring maintains a consistent backup copy of an iASP using internal or external storage. The benefits of geographic mirroring are essentially the same as that of basic switchable device solutions. The added advantage is that it provides disaster recovery to a second iASP at a dispersed location.

Logical replication:

This is a traditional approach for data replication in the IBM i environment. It is based on journaling, including the option of remote journaling. A key characteristic is that only those objects that are journaled by IBM i (Database, Integrated File System, data areas, and data queues) can be replicated in real time. All other objects are captured via the audit journal and then replicated to the target system. Software solutions built upon logical replication typically make use of journaling to capture information regarding transactional data changes made on the source system and then transport and replay those same changes logically to one or more target systems.

The following sections describe the components, setup, and best practices for Web Query for consideration in switched disk, mirroring, or logical replication environments.

Components of Web Query HA/DR

For high-availability or disaster recovery, Db2 Web Query has several components that need to be considered. They are grouped into three categories:

Licensed program product (LPP):

The Web Query product, 5733-WQX, is installed in the system ASP (SYSBASE). LPPs cannot be installed into an iASP. One reason is because LPPs do not allow dynamic

renaming, there's no way to override the product-build library and IFS directory names to change them based on where the objects get restored. Another reason is because the default location for restoring LPPs with the Restore License Program (RSTLICPGM) command is SYSBASE, and there is no way to change it.

The product library is QWEBQRY and the product directory is /qibm/ProdData/qwebqry. These must reside in SYSBASE. The two other libraries shipped with Web Query are QWQCENT and QWQRETAL. They contain sample data and should also reside in SYSBASE:

- QWQCENT is the Century database and is used with the Redbook tutorials.
- QWQRETAIL is the Retail Samples database used with the Retail Samples reports and demo videos.

When installing Web Query into an HA/DR environment, the product must be installed separately onto each system. It is strongly recommended that the Web Query release and group PTF levels be managed so they are at the same levels on all systems.

Web Query data, objects, and user licenses:

Web Query data: Web Query generates operational data and stores it in two locations:

- **IFS:** The user-created synonyms, also referred to as metadata, and configuration settings are stored in the Integrated File System (IFS). They are stored in directory /QIBM/UserData/qwebqry.
- **Db2:** All reports, documents, dashboards, schedules, distribution lists, and runtime environments (RTEs) reside in Db2, in the QWQREPOS library.

In iASP environments, both the IFS and Db2 data can be moved onto the iASP to ensure recovery of critical reports. In replication environments, both the userdata directory and QWQREPOS library should be replicated.

System objects: There are classes of IBM i objects that Web Query creates or modifies that cannot reside on an iASP. These objects include user profiles, job descriptions, and authorization lists. In HA/DR environments, these objects must be replicated and synchronized across systems.

User profiles. There are two Web Query profiles:

- QWQADMIN is the Web Query administrative profile.
- QWQADMGRP is the group profile for the Web Query administrators group in Security Center.

These profiles must be replicated to ensure the CCISD, locale, and group ID values are synchronized across nodes.

Job descriptions for users who are accessing data on the iASP, must be changed to specify the iASP name as the initial ASP group. If the job descriptions are not automatically synchronized by the high-availability product, they must be manually changed on each node in the cluster to specify the iASP name.

Authorization lists are used to secure the metadata directories and their files in the IFS. Some high-availability products have ability to automatically synchronize authorization lists across nodes. PowerHA has a cluster administrative domain that can be configured to help synchronize authorization lists in IBM i 7.2 or later.

Licenses: Web Query's operational environment is regulated by user-based licensing for Developers, Developer Workbench users, and Runtime Groups. Licenses are not IBM i objects and are not handled by the PowerHA cluster administrative domain. For PowerHA and other high-availability products that do not replicate licenses, the licenses must be manually maintained on each node via Web Query's Security Center to ensure a consistent

operational environment.

Reported-on data:

This is the target data that Web Query accesses when running reports. Although it's possible for the reported-on data to reside in SYSBASE, in iASP-based cluster environments, it commonly resides on the iASP.

In some iASP environments, Web Query and its data reside in SYSBASE, but the reported-on data resides on the iASP. It's more common to move the Web Query data to the iASP so that it is replicated or switchable along with the target data. The next sections in this document describe requirements and setup for both these environments.

Reporting on Data that Resides in an iASP

For the iASP environment where Web Query and its data reside in SYSBASE but the metadata and reports are rann against data on an iASP, there is minimal setup needed for Web Query. The only requirement is to change the job description of any user who will be creating metadata or running against the data on the iASP. Change the user's job description to specify the iASP name as the initial ASP group. You can do that using the CHGJOB command, as shown in this example using the QDFTJOB job description:

```
CHGJOB JOB(QGPL/QDFTJOB) INLASPGRP(your_iasp_name)
```

Note: To move the Web Query data (in the IFS and in Db2) to the iASP, see the following section. This section applies to those environments where Web Query resides fully in SYSBASE and only the reported-on data resides on the iASP.

Moving Web Query Data onto an iASP

This section explains how to move the Web Query data to the iASP and to set up Web Query in a geographic mirroring or switched disk environment, specifically:

- Storing the Web Query IFS objects in iASP directories and using the environment on the second node when the iASP is allocated to that machine.
- Storing the Web Query Db2 repository in the iASP and synchronizing it with the Security Center and license manager on the second system when the iASP is allocated to that machine.
- Creating reports on data that is stored in the iASP database and executing them alternately from the system to which the iASP is allocated.

We assume that the iASP-based geo-mirroring or switched disk cluster is already configured and fully operational and that it includes a common IP interface that allows the user to use the same URL to connect to Web Query, regardless of the system that has allocated the iASP at a given moment.

Geo-mirroring is a multi-system solution. Switched disk solutions can be single system where the iASP is switched between partitions, or it can be a multi-system cluster. In the following sections of this document,

we refer to the production node for which the iASP is currently varied on as the *primary* node on the cluster. We refer to a system or partitions that has a mirrored copy of the iASP, or a switchable disk configuration for the iASP, as the *backup* node.

Setting-up the Primary

With the iASP device varied on, perform the following tasks on the primary node:

1. Install Db2 Web Query and apply the latest PTFs.
2. Start Web Query, verify it's working, and end it.
3. Move Web Query data to the iASP.
4. Set up the administrative domain to synchronize SYSBASE objects
5. Set the initial ASP group in the job descriptions.

These tasks are described in more detail in the following sections.

Step 1. Install Web Query and apply PTFs

Install Db2 Web Query 5733-WQX. Instructions for installing the product can be found on the Installation page of the product wiki at <http://ibm.co/db2wqinstallation>. It is recommended that you also install the latest Web Query PTFs. The Web Query group PTF numbers and their latest levels can be found at <http://ibm.biz/BdsuBX>. To request an EZ-install package that already has the latest PTFs bundled in, send a request to qu2@us.ibm.com.

Step 2. Start Web Query, verify it's working, and end it

Start Web Query using the Work Web Query (WRKWEBQRY) or Start Web Query (STRWEBQRY) command. Verify on the WRKWBQRY panel that all ports are active, then verify that you can successfully log into the Web Query portal using the QWQADMIN profile. The portal URL is <http://<your system>:12331/webquery>. Sign out. End Web Query using the WRKWEBQRY or End Web Query (ENDWEBQRY) command.

It is important to start and end Web Query before moving on to the next steps. One reason is because the first startup completes the installation process. Another reason is because startup validates the underlying environment for Web Query.

Step 3. Move Web Query data onto the iASP

Move the Web Query metadata and repository to the iASP. This will ensure it is replicated or switchable along with your other critical business data. The metadata and repository are paired; if you move one to the iASP, then you must move both.

Move the QWQREPOS library.

Web Query's repository resides in the QWQREPOS library. It contains reports, dashboards,

schedules, distribution lists, runtime environments, user permissions, and more. To move it onto the iASP, use these commands:

1. CRTSAVF FILE(QGPL/QWQREPOS)
2. SAVLIB LIB(QWQREPOS) DEV(*SAVF) SAVF(QGPL/QWQREPOS) PVTAUT(*YES)
3. STRSQL
4. DROP SCHEMA QWQREPOS
5. RSTLIB SAVLIB(QWQREPOS) DEV(*SAVF) SAVF(QGPL/QWQREPOS)
RSTASPDEV(*your_iasp_name*) PVTAUT(*YES)

Move the Userdata.

Web Query's userdata directory contains the metadata that is critical for creating and running reports over Db2 or other data. The directory also contains configuration settings and other files for the Web Query installation. To move it onto the iASP, you will need to create and mount a user-defined file system (UDFS). Use these commands:

1. CRTSAVF FILE(QGPL/WEBQRYIFS) TEXT('WebQuery IFS USERDATA Objects')
2. SAV DEV('/qsys.lib/qgpl.lib/WEBQRYIFS.file') OBJ('/qibm/UserData/qwebqry')
DTACPR(*YES) PVTAUT(*YES)
3. Clear the /qibm/UserData/qwebqry directory by entering these commands:
QSH
cd /qibm/UserData/qwebqry
rm -rf *
Press F12 to exist QSH.
4. CRTUDFS UDFS('/dev/*your_iasp_name*/qwebqry.udfs')
5. ADDMFS TYPE(*UDFS) MFS('/dev/*your_iasp_name*/qwebqry.udfs')
MNTOVRDIR('/qibm/UserData/qwebqry')
6. RST DEV('/qsys.lib/qgpl.lib/WEBQRYIFS.file') OBJ('/qibm/UserData/qwebqry')
PVTAUT(*YES)

Note: An alternative approach for the userdata would be to restore it into the default UDFS for the iASP, and then to create symbolic links (symlink) from the SYSBASE userdata to the iASP userdata. This isn't recommended for two reasons. First, there is improved performance with the Web Query UDFS, as opposed to repeatedly resolving the symlinks. Second, you cannot use the restore (RST) command on a symlink. Web Query performs maintenance on certain userdata files when PTFs are applied, and any related RST over a symlink will fail.

Step 4. Set up the Administrative Domain

For PowerHA, the cluster administrative domain is a feature that allows you to identify objects in SYSBASE (i.e. that do not reside on the iASP) that need to be maintained consistently across systems. Because the administrative domain synchronizes these

objects, the application readily has the resources and operational environment it needs to function on each node in the cluster for failover situations.

A monitored resource is a system resource that is managed by a cluster administrative domain. Any changes you make to a monitored resource are synchronized across the nodes in the cluster administrative domain and are applied to the resource on each of the active nodes. Monitored resources are represented in the cluster administrative domain as monitored resource entries (MREs).

For Web Query, an MRE should be created in the administrative domain for each of the following:

- **QWQADMIN user profile:**
This is the administrative profile for Web Query. The profile's password, CCSID, and locale must be synchronized.
- **QWQADMGRP user profile**
This is the group profile for Web Query administrators. The group profile ID must be synchronized.
- **User profiles for each user who will be a Web Query Administrator**
The Web Query Administrators group is managed in Web Query's Security Center. When a user is added or removed from this group, their user profile is updated to add/remove them from the QWQADMGRP group or supplemental group. This membership in the QWQADMGRP profile must be synchronized.
- **User job descriptions:**
Users who will be accessing files on the iASP via Web Query must change their job descriptions to specify the name of the initial iASP group, if not already done. Any changed job descriptions should be monitored so that the change is synchronized across nodes.

The administrative domain and MRE's can be administered through graphical user interfaces or CL commands. For details, refer to the [Implementing PowerHA for IBM i](#) redbook (SG24-7405-00).

The administrative domain is an optional function of PowerHA. If you choose to not monitor the profiles or job descriptions, or if you are using another high-availability product that does not replicate these object types, then you must manually maintain them before switching to a backup node.

Step 5. Edit job descriptions

There are two considerations for job descriptions.

First, you must update QWQADMIN's job description to set your iASP name as the value for the

initial ASP group parameter. This is required. If you do not do this, Web Query will fail to start and there will be other functional issues. To change it, enter this command:

```
CHGJOB JOB(QWEBQRY/QWEBQRYJOB) INLASPGRP(your_iasp_name)
```

Second, to access data that resides on the iASP, individual users will need to update their job description to specify *the initial ASP group*. This step is applicable to users who run reports against data that is located on the iASP or who create metadata over it.

Note: If *all* data that is reported on by Web Query resides in SYSDATABASE, you may skip this step for the users.

To change a user's job description to set the iASP name as the *initial ASP group* value, follow these steps:

1. Display the user's profile to determine what job description it is using:
DSPUSRPRF *user_name*
2. If the profile is using a job description that does not already have the initial ASP group set, you may edit that job description or create a new one. To create a new one based on the existing one, use the WRKJOB command and select option 3 to copy it.
3. Set the *initial ASP group* parameter value to your iASP name. Here is an example command for the job description QGPL/QWEBQRYJD:
CHGJOB JOB(QGPL/QWEBQRYJD) INLASPGRP(*your_iasp_name*)
4. If a new job description was created, edit the user's profile to specify the name of the job description. Here is an example command for job description QGPL/QWEBQRYJD:
CHGUSRPRF USRPRF(*user_name*) JOB(QGPL/QWEBQRYJD)

Step 6. Start Web Query

Setup on the primary system is now complete. You may start Web Query by entering the WRKWEBQRY or STRWEBQRY command.

Setting-up the Backup

This section lists the steps required to set up Web Query on a backup node when the Web Query data is on a switchable or replicated iASP. The iASP should be varied off for the node when performing these steps.

1. Install Web Query and apply PTFs. The primary and backup systems should be at the same Web Query release and fix levels.
2. Start Web Query. Login to the portal to verify the environment. End Web Query.
3. Clear the /qibm/UserData/qwebqry directory by entering these commands:
qsh
cd /qibm/UserData/qwebqry
rm -rf *

4. Delete the repository using these commands:
STRSQL
DROP SCHEMA QWQREPOS
5. Update the job description for QWQADMIN to set the iASP name using the following command:
CHGJOB JOB(QWEBQRY/QWEBQRYJOB) INLSPGRP(*your_iasp_name*)

Switching from Primary to Backup

When switching the primary (production) node to a backup, or when switching an iASP from the primary node to a backup node, you must mount the UDFS, then start Web Query. Additionally, you must add or remove user-based licenses and Web Query Administrators, as needed, to synchronize with the Web Query repository. Follow these steps.

1. Mount the Web Query userdata in your iASP:
ADDMFS TYPE(*UDFS) MFS('/dev/*your_iasp_name*/qwebqry.udfs')
MNTTOVRDIR('/qibm/UserData/qwebqry')
2. Log into the IBM i with the QWQADMIN profile and start Web Query:
STRWEBQRY
3. Add or remove user licenses:
After switching systems, the users and their roles will be known to Web Query on the backup (switched-to) system, but the underlying user licenses might not be synchronized. This situation can occur if licenses were added or removed for any Developers, Developer Workbench users, or Runtime Groups since the last time the iASP was varied onto this backup node.

To **add Developer user licenses** that are missing on the backup system, log into the Web Query Security Center using the QWQADMIN profile. Click the Add User icon, and import the user(s) from the import candidate list. To **remove a Developer user license**, first click the Add User icon to re-add the user, and then delete the user.

To **add Developer Workbench user license** on the backup system, remove the user from the Developer Workbench folder group and re-add them. To **remove a Developer Workbench user license**, add the user to the Developer Workbench folder group and then remove them again.

To **add a Runtime Enablement group license** on the backup system, click the run-only folder group that the group profile is authorized to, then click Edit in the pulldown. Remove the group profile and re-add it. Note that if the group profile is authorized to more than one run-only group, you must remove it from all the folder groups before re-adding it. To **remove a Runtime Enablement group license**, add the group profile to a run-only group and then remove it.

Tip: If the production (switched-from) node is accessible before the switch, you can use the Work License Info (WRKLCINF) command and select option 8 on features 5104 (Developer Users), 5105 (Developer Workbench Users), and 5106 (Runtime Enablement Groups) to get a list of currently licensed profiles. Compare those lists to the WRKLCINF results on the backup node to determine what licenses, if any, need to be added or removed on the backup.

4. Synchronize user profiles:
If your high-availability product does not automatically synchronize user profiles, use the Security Center to manually synchronize members of the QWQADMGRP group profile with the Web Query Administrators group in Security Center. (There should be a one-to-one correspondence between the group profile members and the folder group members.) You will need to synchronize them if any Web Query administrators were added or removed on the production system since the last time the iASP was varied onto the backup node you are switching to.

In Security Center, click the Web Query Administrators group to list its members. For each user in the list, remove the user from the Web Query Administrators group and re-add them. This will update the user's profile to add the QWQADMGRP group.

Now use the `DSPUSRPRF QWQADMGRP TYPE(*GRPMBR)` command to list the members of the QWQADMGRP group profile and verify that there are no extraneous members who are not listed in the Web Query Administrators group. If there are, use the Change User Profile (`CHGUSRPRF`) command to remove the group from the user's profile

Best Practices

It is recommended that nodes in a multi-system or multi-partition cluster always be at the same Web Query release and PTF fix level. Web Query frequently introduces new function in the group PTFs, and this can cause incompatibilities, particularly in the repository and in configuration settings, when switching to down level nodes.

After initial setup of the primary and backup nodes, always vary on the iASP for a node before starting Web Query on that node. If Web Query userdata directory resides on an iASP and if the userdata is not accessible at startup, Web Query will fail.

References

- For more information about Web Query, refer to the product wiki at <http://ibm.co/db2wqwiki>.
- For more information on the process of selecting, configuring, and managing a high-availability solution using IBM PowerHA, refer to the Implementing PowerHA for IBM i redbook (SG24-7405-00).
- For more information about iASPs, refer to the topic Disk Pools in the IBM i Information Center.