



Disaster Recovery Best Practices for UC4 Operations Manager and UC4 Automation Engine

The intent of this document is to discuss the different types of Disaster Recovery that UC4 customers have used in the past, the best practices to set up basic Disaster Recovery, and lessons learned from past Disaster Recovery methods.

Definition of Disaster Recovery and Terminology Used

Wikipedia defines Disaster Recovery (DR) as “the process, policies and procedures that are related to preparing for recovery or continuation of technology infrastructure which are vital to an organization after a natural or human-induced disaster.” This is not to be confused with failover which is simply to cover a single server or set of servers failing and another server or set of servers taking over the processes from before. UC4 uses the same definition of Disaster Recovery where a large geographic region is impacted and servers in another region need to take over.

Original Processes

The terms “original Automation Engine processes” or “original Database” are referring to the processes or database that is running in the original region where a disaster occurs.

DR Processes

The terms “DR Automation Engine processes” or “DR Database” are referring to the processes or database that are running in the region that should take over when a Disaster occurs in the original region.

Assumptions used in this paper

Off-site database

This paper assumes that an off-site database is being used for Disaster Recovery. This is based off the idea that the servers running the Automation Engine nodes need to be no more than 20 kilometers from each other or the UC4 database.

Off-site Automation Engine Processes

This paper assumes that an off-site server(s) are being used for the Automation Engine Processes. This is based off the idea that the servers running the Automation Engine nodes be no more than 20 kilometers from each other or the UC4 database.

DNS Server

This paper assumes that a DNS server is available and network admins have set up replication and disaster recovery for the DNS server that will allow DNS changes automatically when a disaster occurs at the original site.

Setting up for Disaster Recovery

Licensing

Using a copy of the original database

If a replicated database is being used, this will already contain the licensing information needed.

Using a separate database

If a replicated database is not being used, a DR license will need to be loaded into the DR database. If you do not already have a DR license, please contact your Account Manager, your CRM, or UC4 Support for a DR license.

Database and AutomationEngine component builds

A best practice for Disaster Recovery or any sort of failover for any UC4 component is that the build (service pack and hotfix) is exactly the same for the components being used on each system. This will allow for a decrease in unexpected behaviors occurring between the two systems. Below are how to find what build each component has:

AutomationEngine

The build for the AutomationEngine components can be found either in the WP and CP logs by searching for the line that starts with "U0003400 UC4 Server ... Version"

This information can also be found by going to the System Overview, expanding the node showing <system name> - <client>, and clicking on the "Automation Engine" node. The "Version" column will have this information.

Database/InitialData

The build for the Database or Initialdata component can be found in the WP logs by searching for the line that starts with "U0003621 Version of UC4 database objects...".

This information can also be found by going to the System Overview, expanding the node showing <system name> - <client>, and clicking on the "Database" node. The information under the "InitialData Version" will show this information.

Agents

The build for each agent can be found in their respective log files by searching for the line that starts with "U2000000 Started program UC4 Agent..."

This information can also be found by going to the System Overview, expanding the node showing <system name> - <client>, and clicking on the "Agents" node. Each agent will show its respective build under the "Version" column.

NOTE: The build for each agent on the system may be different.

IP addresses used by default

By default, UC4 uses IP addresses to show where the Automation Engine processes reside. For example the database will show the following for WPs and CPs distributed across two servers:

Process Name	TCP/IP Address
UC4WA9#WP001	172.16.144.145
UC4WA9#WP002	172.16.144.145
UC4WA9#WP003	172.16.144.145
UC4WA9#CP001	172.16.144.145
UC4WA9#CP002	172.16.144.145
UC4WA9#WP004	172.16.144.146
UC4WA9#WP005	172.16.144.146
UC4WA9#WP006	172.16.144.146
UC4WA9#CP003	172.16.144.146
UC4WA9#CP004	172.16.144.146

The Agent ini files on startup will connect to the first CP and then will gather a list of all other CPs and put this into the ini file. It will show up as follows:

```
[CP_LIST]
2218=172.16.144.145
2219=172.16.144.145
2220=172.16.144.146
2221=172.16.144.146
```

The problem that occurs with this is that the DR servers will not always have the same IP addresses when a true Disaster occurs. At this point, agents will attempt to reconnect to the IP addresses which may not exist and job processing will stop beyond what the original disaster caused.

If the DR servers can quickly gain the same IP addresses as the servers the original processes were running on, then that can work instead of DNS.

Using DNS

For Disaster Recovery, it is recommended to use DNS instead of IP addresses. This is so that agents are able to automatically reconnect via DNS rather than IP. Please see the section "How to tell if your system already uses DNS rather than IP" to see if your system is already set up for using DNS for DR.

Below are the settings that need to be changed. These are numbered so they can be referenced in other parts of this paper.

1) *Settings in the ucsrv.ini file(s)*

Each ucsrv.ini file being used will need to have the hostname= setting updated under the [TCP/IP] section. The hostname= setting should be the fully qualified dns name. An example is below:

```
[TCP/IP]
```

```
hostname=fully.qualified.dns.name.com
```

If this setting is updated or changed, a restart of the CP processes using this ucsrv.ini file will need to be done

2) *Settings in the agents' ini file(s)*

If the agent ini files shows an IP address rather than a dns name for the cp= setting (see step 2 of "How to tell if your system already uses DNS rather than IP" below), the agent ini file will need to have its cp= setting under the [TCP/IP] section updated to use the dns name. An example would be:

```
cp=fully.qualified.dns.name:2217
```

If the agent ini files show IP addresses rather than dns names (see step 3 of "How to tell if your system already uses DNS rather than IP" below), the [CP_LIST] section will need to be cleared out

Both changes to the agents' ini file(s) require the agent to be restarted after a change to the hostname= setting in the ucsrv.ini file has been made and the CPs have been restarted (see step 1 above). The [CP_LIST] will automatically fill in after the agent restart.

3) *Settings in the users' uc4config.xml file(s)*

If the uc4config.xml file(s) show IP addresses rather than dns names (see step 3 of "How to tell if your system already uses DNS rather than IP" below), the <connection> tag for that specific system in the uc4config.xml file will need to be changed to something like:

```
<connection name="UC4_PROD" system="UC4_PROD">  
  <cp ip="fully.qualified.dns.name.com" port="2217"/>  
</connection>
```

The rest of the cp information will be filled in on restart of the UI.

How to tell if your system already uses DNS rather than IP

Use the following to see if your system is already set up to use DNS

1) Run the following SQL statement against the original UC4 database:

```
select MQSRV_NAME, MQSRV_TCPIPADDR from mqsrv where  
mqsrv_type = 1;
```

If the results show the MQSRV_TCPIPADDR columns as DNS names rather than IP addresses, the Automation Engine processes are set up correctly. If they show IP addresses, refer to step 1 under the “Using DNS” section above.

- 2) Check all agents' cp= setting in their ini file. A correctly configured cp= setting will show something like:

```
cp=fully.qualified.dns.name.com:2217
```

An incorrectly configured cp= setting will show something like:

```
cp=10.10.10.10:2217
```

If their cp= settings show IP addresses, please refer to step 2 under the “Using DNS” section above.

- 3) Check all agents' [CP_LIST] sections in their ini files.

A correctly configured CP_LIST will show something like:

```
[CP_LIST]
2217=fully.qualified.dns.name.com
2218=fully.qualified.dns.name.com
```

An incorrectly configured [CP_LIST] will show something like:

```
2217=10.10.10.10
2218=10.10.10.10
```

If their [CP_LIST] sections show IP addresses, please refer to step 2 under the “Using DNS” section above.

- 4) For the UI, check users' uc4config.xml files under the <connection> tag.

A correctly configured <connection> tag will show something like:

```
<connection name="UC4_PROD" system="UC4_PROD">
  <cp ip="fully.qualified.dns.name.com" port="2217"/>
  <cp ip="fully.qualified.dns.name.com" port="2220"/>
</connection>
```

An incorrectly configured <connection> tag will show something like:

```
<connection name="UC4_PROD" system="UC4_PROD">
  <cp ip="10.10.10.10" port="2217"/>
  <cp ip="10.10.10.10" port="2220"/>
</connection>
```

If the <connection> tag is still showing IP addresses, refer to step 3 under the “Using DNS” section above.

Database replication

In regards to getting data to the DR database, there are two methods that can be used.

Full copy of the original database

The full copy of the database here means that the copy of the database was taken as a cold backup of the original database, a warm copy of the original database, or a real-time replication of the database.

In order to have a perfect replication of the database, the replication would need to be down to less than a millisecond to have everything exactly the same. Because of this, a COLD start will be required before starting the database (please see section “What should be done at Disaster Recovery” below).

Using a different database (not a direct copy)

It’s possible in this situation to use a completely different database and system that just has the objects transported over. This method requires a few extra steps.

DR License

If a different database is being used, a license needs to be acquired for the DR system that exactly mirrors the original system’s license. This is so that agents are able to reconnect.

Agent connection

If a different database is being used, the agents will need to have been connected to the DR system before they can reconnect from the original system so that they can be allowed in the correct clients. An alternative and more preferred method is to use an HSTA object in client 0 that will do this automatically. For more information on HSTA objects, please see the UC4 documentation under User Guide, Objects, Alphabetical Listing, Agent/Client Assignment, Agent/Client Assignment.

Copying objects

This method requires a client copy or transport case be used to transfer the objects.

More information on using the Client Copy utility can be found in the UC4 documentation under Administration Guide, Utilities, UC4.DB Client Copy, Copying and Deleting Clients – UC4.DB Client Copy.

More information on Transport Cases can be found in the UC4 documentation under Administration Guide, User Guide, UserInterface, UC4 Explorer, UC4 Transport Case.

Keystore issues

A security feature introduced with UC4 Operations Manager 8.00A is the use of “Transfer keys”. These are entries in the UC4 database coupled with a file on the agent system that need to match each other. If they do not match, an error will be thrown on agent start in the agent’s log file. An example of this is “U2000099 Transfer key could not be loaded...” or “U0003358 Transfer key for Agent '<agent name>' has not yet been loaded.” These errors can occur under the following circumstances:

- An Agent is started on a different system before starting on the current system
- A newly installed Agent is started with the same name as another Agent in a different directory
- A newly installed Agent\bin directory is used with the same agent name as one already started on that system
- An already-existing Agent is started with a new name

In a DR situation where a new database is being used, the agents will match the first transfer key issue described above where an agent is started on a different system (here, the original system) before starting on the current system (here, the DR system).

To resolve this issue, an extra step is needed before the agents reconnect to the database – once the DR system is up and running, go into Client 0 on the DR system, go to System Overview, expand the <system name> - <client> node, click on the “Agents” node, click in the right panel, and select all agents using “Ctrl + A”. Right click one of these agents, and choose the option “Renew Transfer Key”. This will allow the agents to start on the new system.

What should be done at Disaster Recovery

There are a number of steps that must be taken when the DR system takes over. Some of these are described above for specific circumstances. The below steps should be taken on any DR system for it to take over processing

NOTE: In any disaster, there should be an expectation that some objects will need to be restarted. During any DR instance, all activities should be monitored closely. The steps below are best practices that may not cover all scenarios.

COLD start

The first step that must be taken is a COLD start of the UC4 system. This will clear out temporary tables that UC4 uses to see what routines should be run next. To perform a COLD start of a UC4 system the following steps should be used:

- 1) one of the ucsrv.ini files must have the setting:

StartMode=COLD

- 2) Start a single WP that uses that ini file – it is very important that no other WP or CP processes are running at the time or the COLD start will not be effective
- 3) View the WP’s log file and wait until the message below shows:

```
"U0003471 The Server '<system name>#WP001' has successfully been initialized, ***  
READY FOR RUN ***"
```

- 4) Start all other WP and CP processes that should be started.

Since the DR system should usually be down except in the cases where it being tested or a disaster has occurred that required the DR system take over processing, a best practice is to always keep the "StartMode" setting in one of the ucsrv.ini files on the DR system set to "COLD" so this step is ready as soon as DR is necessary.

DNS/IP switch over

The network admins should make it so that the either the DNS switch has taken place to point to the IP addresses for the DR servers, or that the DR servers have the same IP addresses as the original servers had (please see section "Setting up for Disaster Recovery" above).

Agent reconnection

Assuming the DNS/IP switch over has taken place and any transfer key issues have already been resolved, the agents should reconnect to the system automatically within their reconnect period (by default 10 minutes). There is a possibility that some agents will be unable to reconnect for whatever reason (transfer key, network issues, a disaster in the area where the agent resides, etc...). Agent reconnection should be monitored closely until all needed agents have been confirmed as reconnected.

Objects that should be in the activities window

A step that should be taken is confirming that the correct objects are in the activities window. A recommendation here would be to make sure that only "scheduling" objects are in the activities window. This would include "Schedule", "Event", and "C_PERIOD" objects. The reason for this is that other objects may lose information during the switch over to the DR system.

Testing Disaster Recovery

It is highly recommended that any Disaster Recovery system implementation is fully tested to be sure everything runs smoothly if a disaster does occur. Please contact UC4 Support regarding Support that will be provided during a DR test before going forward with it.

In Case a Disaster Does Occur

In the case of an actual disaster that requires a Disaster Recovery system to take over processing, if any issues do arise that leave your Production processes down or stop them from running and there is no viable workaround, contact UC4 Support immediately. Be sure to explain that the current system is your Disaster Recovery system which was just brought up after a disaster and that your production processes are currently unable to run.