

Supplement #55

USA: +1 603 924 8818 • techsupport@softlanding.com • <https://support.softlanding.com>

EMEA: +44 1293 872101 • softlandingsupport@macro4.com • <https://support.softlanding.com>

TURNOVER® FOR ISERIES DISTRIBUTION QUEUES

This document applies to FTP distribution in TURNOVER® Release 5.3 and later.

TURNOVER® Distribution Queues allow TURNOVER® to send multiple FTP distributions simultaneously, significantly reducing the time it takes to distribute a form to multiple systems using FTP.

REQUIREMENTS

This feature requires TURNOVER® Release 5.3 with a tape date of March 2004 or later.

OVERVIEW

Because of SNADS' inherent multi-threaded processing ability, TURNOVER®'s original distribution process was designed to be linear. SNADS processes TURNOVER® distribution requests in the background and quickly returns control to TURNOVER®, so it works fine for distributing to one system or to many systems.

Unlike SNADS, however, the FTP distribution process runs completely inline, so the next system in the distribution list is not processed until the previous FTP transfer is complete. This becomes a problem for customers with a large number of systems. Consider a scenario where there are 100 systems in the list and it takes roughly 2.5 minutes to send the save file. Under the current architecture, it would take over four hours for the whole process to complete!

TURNOVER® Distribution Queues address this problem. This feature brings the advantages of the background processing used by SNADS transmissions to FTP. However, for those who are familiar with SNADS, we should point out that this is nothing like SNADS distribution queues. TURNOVER® Distribution Queues are not tied to a particular system; they are instead a set of background worker jobs that you can have running that perform the actual FTP transmission of the TURNOVER® save file. You can start as many (or as few) of these jobs as you'd need to meet your requirements within the limits of the resources available on your server and network. A company with 100 systems could have 100 queues, or they could have 10.

CONFIGURATION

Configuring this feature requires two high-level steps, each of which has sub-steps:

1. Create TURNOVER® for iSeries distribution queues according to the instructions provided in this document, so that the distributions are actually processed. (Instructions begin below.)
2. Turn on the Distribution Queues feature in the FTP distribution override settings for each system with which you want to use TURNOVER® for iSeries distribution queues. (See page 5.)

Step 1) Setting up your TURNOVER® for iSeries distribution queues

Setting up the distribution queues consists of these activities:

- Setting up an IBM i subsystem for the jobs to run in.
- Adding as many autostart job entries to that subsystem as corresponds to the number of worker queues you want to have. For example, if you want to have up to 10 sends happening simultaneously, you would add 10 autostart jobs to the subsystem.
- Creating a data queue.
- Creating a data area.

UNICOM Systems, Inc. RECOMMENDS

We strongly recommend that you create a dedicated subsystem for this process, so that you have maximum control over tuning and managing the jobs.

A. Setting up a dedicated subsystem

What follows is a rough approximation of the steps you would go through to create a subsystem and add 10 worker jobs to it. You should adapt these instructions as needed for the number of jobs and the resource usage you deem best for your server and network.

1) Create a memory pool

This step is not strictly necessary; you *can* run these jobs in an existing memory pool such as *BASE or *INTERACT. However, if you want to be able to fine tune the amount of resources that these queues can utilize, then running the subsystem in a dedicated memory pool is the best way to do that. Use the following command:

```
CHGSHRPOOL POOL(*SHRPOOL1) SIZE(4096) ACTLVL(10) PAGING(*CALC)  
TEXT('Memory Pool for TURNOVER Distribution Queues')
```

TURNOVER® for iSeries Distribution Queues

This command establishes a memory pool named *SHRPOOL1, allocating to it 4 megabytes of memory and assigning an activity level of 10 jobs. This is just an example; you will need to work out what the right values for memory and activity level should be on your system.

2) Create a subsystem description

The main things to know in this step are the name you want to give the subsystem, and from which memory pool(s) you want it to draw resources.

```
CRTSBSD SBSD(QGPL/TDSTQUEUE) POOLS((1 *SHRPOOL1))  
TEXT('Subsystem for TURNOVER Distribution Queues')
```

This example names the subsystem TDSTQUEUE and stores its definition in QGPL.

UNICOM Systems, Inc. WARNING

Do NOT create the subsystem in one of the TURNOVER® for iSeries libraries! If you do, you will lose it when you install future TURNOVER® for iSeries upgrades.

3) Add a routing entry

Subsystems must contain a routing entry in order to process jobs. This command adds a standard routing entry.

```
ADDRTGE SBSD(QGPL/TDSTQUEUE) SEQNBR(1) CMPVAL(*ANY) PGM(QCMD)
```

4) Create a job description

The job description you create is used during the next step to add your autostart jobs. You need only one job description to define the way these worker jobs will run.

```
CRTJOB JOB(QGPL/TDSTQUEUE) TEXT('JobD for TURNOVER  
Distribution Queue') USER(TSERVER) RTGDTA(QCMDB) RQSDTA('CALL  
TDSTQUE') INLLIBL(SOFTTURN SOFTTURN SOFTTURN) JOBMSGQFL(*WRAP)
```

TURNOVER® for iSeries Distribution Queues

There are three critical parameters in the job description:

- | | |
|---------|--|
| USER | Job descriptions for autostart jobs must specify a valid user. The user you specify is not too critical. TSERVER is a user profile we create on your system that is not used for anything but server jobs, so it is a good choice. You could create a unique user profile for this purpose if you really want to monitor what the user does. |
| RQSDTA | This is the program that will run when the autostart job starts. Specify a value of 'CALL TDSTQUE'. |
| INLLIBL | You must specify all of the TURNOVER® for iSeries libraries in the TURNOVER® for iSeries environment that these jobs are servicing ¹ . |

Note: If you run multiple TURNOVER® for iSeries environments, you would need one job description for each environment you need to service, and you would need to create one set of jobs for each environment. However, all of the jobs can share the same subsystem.

B. Add autostart jobs to the subsystem

In this step, you add to the subsystem the autostart jobs that will serve as the worker jobs to actually perform the FTP transfers. Run this command once for each job.

ADDAJE SBSD(QGPL/TDSTQUEUE) JOB(TDSTQUE001) JOBD(QGPL/TDSTQUEUE)

For our example, we would run the command, press F9 to retrieve it, change the number in the JOB parameter, and run it again – for a total of 10 times.

You can name the jobs whatever you want, but it does make sense to number them so that you have an idea which job is which and how many you have.

C. Create a Data Queue

The way the distribution queue process works is that distributions are placed on a data queue, and the worker jobs are idling, waiting for entries on the data queue. Data queues were designed for exactly this kind of situation, where there are multiple jobs waiting for instructions in the same place. Use this command as a guide:¹

**CRTDTAQ DTAQ(SOFTTURND/TDISTQ) MAXLEN(1500) AUTORCL(*YES)
TEXT('TURNOVER Distribution Queue') AUT(*CHANGE)**

The data queue must be created in your environment data library and with the values shown.¹ The presence of this queue in the data library helps TURNOVER® for iSeries decide at runtime whether to attempt to use distribution queues during an FTP distribution, in case the feature has been turned on without performing the accompanying setup.

¹ Substitute the name of your TURNOVER® for iSeries library or libraries as appropriate, if different from the default library name(s) we've used in our examples.

TURNOVER® for iSeries Distribution Queues

D. Create a data area

This data area is used as an indicator that you have configured the distribution queues, so that TURNOVER® for iSeries can decide at runtime whether to attempt to use them. The name and location are important – it should be called TDISTQ, and you should create it in the TURNOVER® for iSeries data library. Use this command as a guide:¹


```
CRTDTAARA DTAARA(SOFTTURND/TDISTQ) TYPE(*LGL) TEXT('TURNOVER  
Distribution Queues')
```

Step 2) Turning on the Distribution Queues feature

You can turn on the Distribution Queues feature using either the Administration area of the TURNOVER® for iSeries Client or a green screen session.

To reach the **FTP Method Details** panel (shown above) in the TURNOVER® for iSeries Client, click  **Administration**, then click  **Distribution Systems**, then right-click and choose  **New System**. On the **System Identity** page, supply a system name, system description, and serial number, and specify a distribution method of ***FTP**. When you click **Next**, you will see the **FTP Method Details** panel.

TURNOVER® for iSeries Distribution Queues



To reach the *Change Distribution Method Override* panel (shown above), use Main Menu option 11 (*Distribution Menu*) and then option 6 (*Override network distribution method*), and then select a system.

Once this parameter is turned on, TURNOVER® for iSeries checks at runtime to make sure that distribution queues actually exist. If they don't, then TURNOVER® for iSeries behaves as if this parameter is turned off. If you are already using FTP for distribution, then after you set up your distribution queues, you must explicitly turn this parameter on as instructed on page 5.


MANAGING THE DISTRIBUTION QUEUES

The distribution queues are controlled by the subsystem you created. To start them, just use the **STRSBS** command:

STRSBS SBSD(QGPL/TDSTQUEUE)

You might want to add this command to your IPL startup program so that the subsystem is always running.

If you run the command **WRKSBSJOB**, you should see something like this:



To end the distribution queues, run this command:

ENDSBS SBS(TDSTQUEUE) OPTION(*CNTRLD) DELAY(*NOLIMIT

The worker jobs check to see if they should end every 30 seconds, so if you end the subsystem this way, all of the jobs that are idle will end within 30 seconds. Any jobs that are actively sending will end when they finish with what they are currently sending.

Likewise, if you end an individual job instead of the entire subsystem, you should use the ***CNTRLD** option with a time of 999999 to give an active job time to finish sending what it is sending.

Using multiple subsystems

If you perform most of your distributions in the evening when the system is less busy, you could create two subsystems. The first subsystem would always be running, and would have a limited number of jobs so as not to overload your server or network during business hours. You could then use a Job Schedule Entry to start and stop a second subsystem that runs only after hours. This subsystem could add more worker jobs to the mix, allowing your distributions to happen more quickly.

CHANGES TO TURNOVER® FOR ISERIES PROCESSING

People accustomed to using TURNOVER® for iSeries might notice some differences once they've installed this change. In many cases, the differences only appear if you're using this feature.

Deleting the distribution save file

This is an invisible change. In the past when the distribution job runs, it has always deleted the save file it creates at the end since it will have already sent it to all of the systems. Now, the distribution job and the worker jobs have to coordinate with each other to determine when the save file should be deleted. The last process to use the save file deletes it; typically, this will be one of the worker jobs.

Updating the form status

When the distribution job ends, it typically updates the form status from DIST-PND to either DISTRIB or PARTIAL. However, if distribution queues are in use and pending distributions are still in progress when the job ends, then the form status is updated when the final worker job has sent the save file. This ensures that the form status is changed properly to reflect any errors that may have been encountered by the worker jobs.

Name and location of FTP log spooled file

Previously, an FTP log was always produced within the distribution job for each system to which it was sent. The name of the spooled file was "FTPLOG," with the system name in the USRDTA of the spooled file. Because the FTP processing occurs within the worker jobs, those jobs produce the spooled files. If you were sending to hundreds of systems and one of them

TURNOVER® for iSeries Distribution Queues


failed, you'd find it virtually impossible to locate the right log with this new system, so we have made the following changes.

1. The name of the spooled file is now the target system name.
2. The USRDTA contains the name of the distribution job.

With these changes in place, you can run the following command to see all the logs for the job in one place:

WRKSPLF SELECT(*ALL *ALL *ALL DISJOBNAME)

This processing has also been integrated into the *Work with Distribution Status* panel² in the form of a new option, **4=Spooled files**, to make it easy for you to retrieve information about the jobs.



All TURNOVER® for iSeries users will see this change, whether they use the new feature or not.

RECOMMENDATIONS

We do recommend that anyone using FTP to distribute changes to multiple systems take the time to configure and use this feature, as it improves the speed and efficiency of the process.

² From *Work with Forms*, run option 11 against a form to see the *Work with Distribution Status* panel.

TURNOVER® for iSeries Distribution Queues

We expect that people will want some tuning recommendations as to how many jobs they should run, and how much memory to allocate. Right now we don't have this information, but as we work with customers that implement this feature, we will gather it and update this document accordingly.