

Connection Pools

Connection pools enable you to immediately get a host connection that is ready in a specific screen, the "initial screen". This saves the time of establishing the connection with the host and navigating to the relevant screen. See Navigation. This chapter applies to Web enablement and SOA scenarios running a path procedure; it covers the following topics:

- Creating a New Connection Pool
- Changing the Initialization Mode of a Connection Pool
- Changing the Log Level
- Defining an Initialization Path
- Defining a Termination Path
- Connection Pool Life Cycle
- Connection Pool States
- Connection Pool Troubleshooting

See also: Connection Pool and Connection Information Sets in the *Reference Guide* | Connection Information Sets in this manual.

Creating a New Connection Pool

➤ To create a new Connection Pool

1. Select the relevant application.
2. In the **File** menu select **New>Entity>Connection Pool**. The *New Connection Pool wizard* is displayed.
3. Enter a name for the screen, a suitable description and determine the folder where the screen is to be located. Click **Finish**. You have now created a new Connection Pool in the repository.
4. Edit the settings in the Editor area as detailed in the Connection Pool reference.

Changing the Initialization Mode of a Connection Pool

➤ To change the Initialization Mode of a Connection Pool

1. Double-click on the relevant Connection Pool.
2. Select the **General** tab.
3. In the **Initialization mode** field, select the relevant mode:

- **Manual:** An administrator manually initializes the connection pool (the connection is initiated by selecting the connection pool in the Management node, and then choosing Start Pool from the right-click shortcut menu).
- **When first accessed:** When Web applications or sessions request to connect to this connection pool.
- **Automatic:** Automatically initializes the connection pool when the application is loaded.

Changing the Log Level

Fine-tuning connection pool parameters or identifying problems are reasons you may want to change the level of detail that the log displays. It is possible to set a different detail of logging for each connection pool: None, Errors, Warnings, Information and Details. The server logger should be set to no less than the **Normal** log level, in order to see connection pool logs. Once the project is in the production phase, **Error** level is the recommended level to use.

Defining an Initialization Path

An initialization path is a path that can be executed on any new connection to the host that navigates the connection to one of the Initial Screens. In order to select a path from different folders, click the folder selection button.

Open the Navigation tab of the relevant Connection Pool and select the path/path procedure from the Initialization drop-down list.

Defining a Termination Path

A Termination path is a path that can be executed on any connection to the host (used or new) that should be performed before destroying a connection (for example, a host side logout procedure).

To define a Termination Path, open the Navigation tab of the relevant Connection Pool and select the path/path procedure from the Termination drop-down list.

Connection Pool Life Cycle

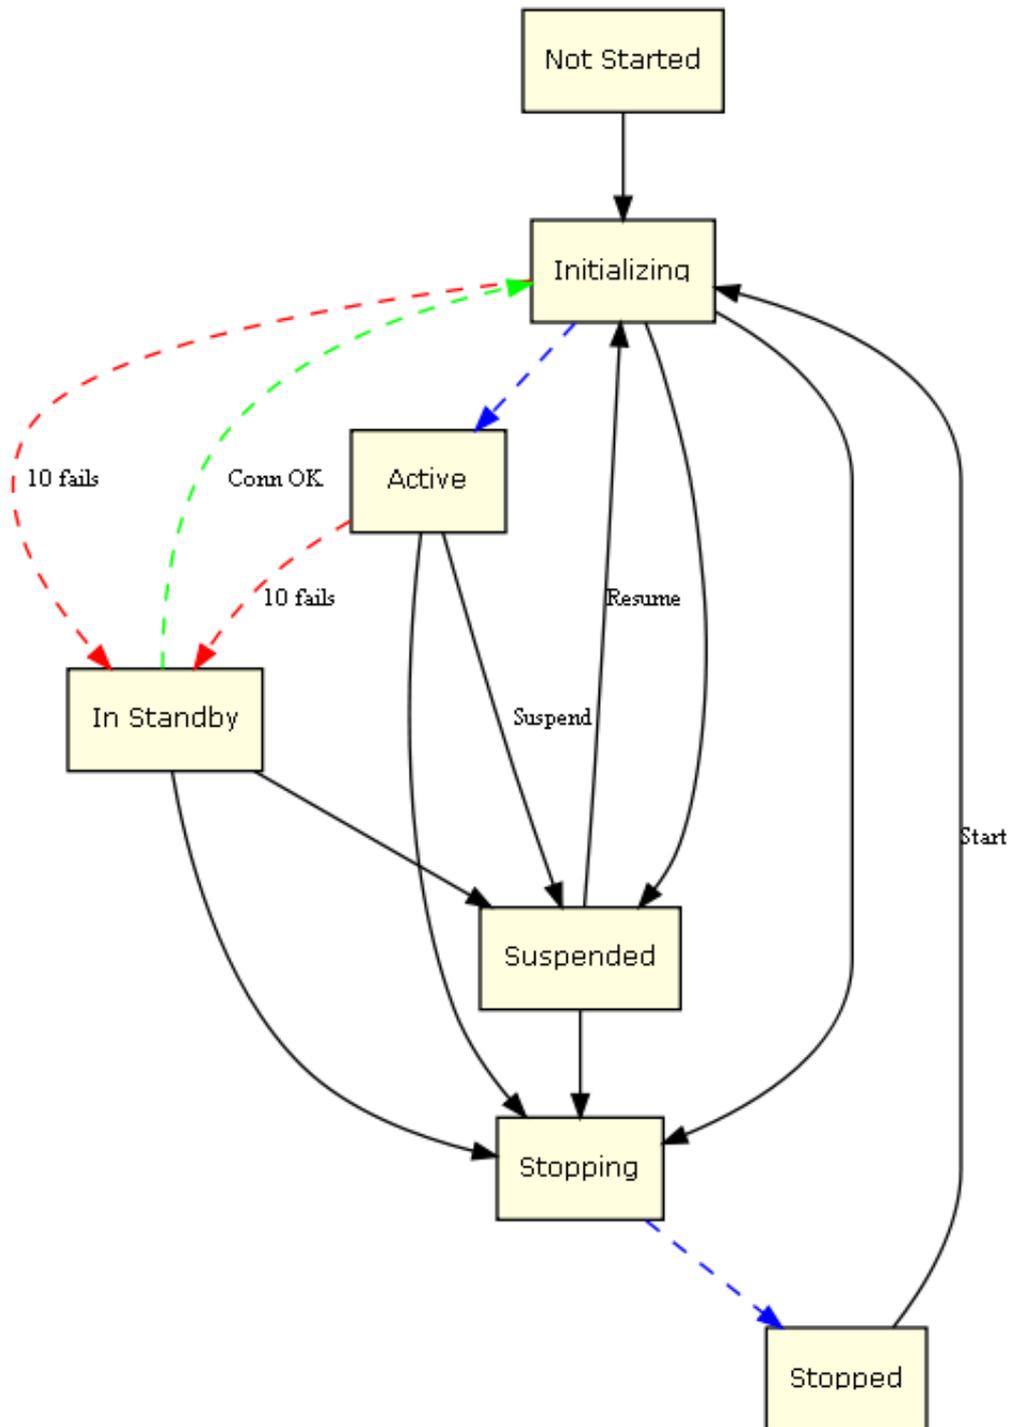
Connections in the pool automatically go through different states, according to the pool policy as defined in the Connection Pool Entity. The connection policy is built from:

- Pool size type: fixed, limited, flexible
- Keep-alive path and interval
- Run recycle anyway flag
- Disconnect after usage flag
- delays and timeouts

Other factors may also impact pool behavior.
See Pool Size Control Policy Considerations | License Key Concept | Connection Information Sets.

Connection Pool States

- The Connection Pool itself can have seven states:

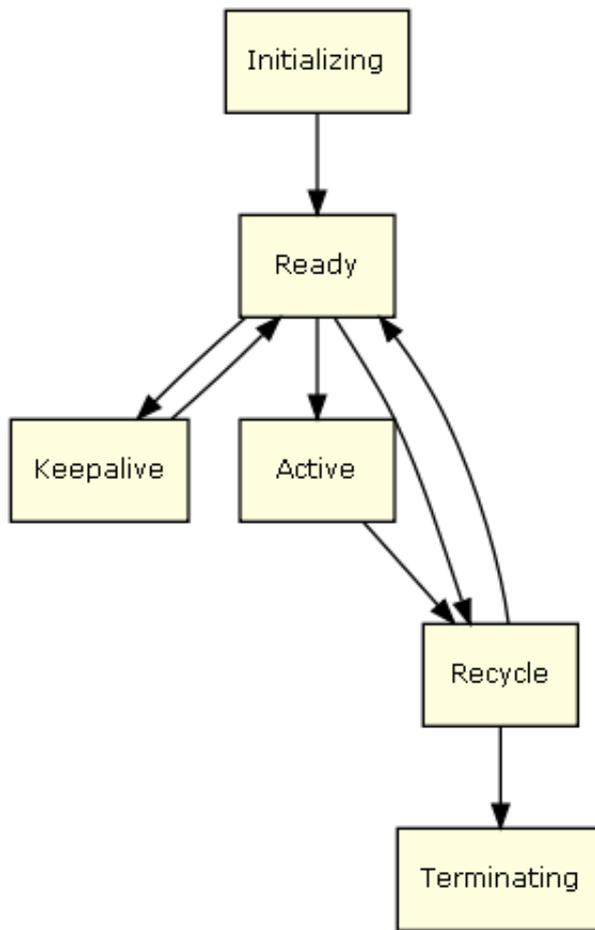


State	Description
Not Started	The connection pool was not initialized yet. If you request a connection, the connection pool will return an immediate error.
Initializing	The connection pool is trying to reach Active status, but does not have any "ready" connections yet. If you request a connection, the connection pool will return an immediate error.
Active	The connection pool is working and managed to create at least one connection to the host.
In Standby	The connection pool had several consecutive errors trying to create new connections to the host. The connection pool will continue to try to connect if user requests arrive, but will not initiate new connections otherwise. If a new connection is successfully created, the status will automatically change to Active.
Suspended	The connection pool is blocked for new users and does not maintain its connections. If you request a connection, the connection pool will return an immediate error.
Stopping	The connection pool is trying to reach Stopped status, but still has connections in different stages of termination. When all connections are down, the status will automatically change to Stopped. If you request a connection, the connection pool will return an immediate error.
Stopped	The connection pool does not have connections or maintenance. If you request a connection, the connection pool will return an immediate error.

Note:

the Stopped state will only happen after Stopping; the Initializing state will happen before Active. These are "in between" states, typically not lasting more than a couple of seconds.

- Host connections inside the connection pool can have six states:



State	Description
Initializing	Host connection will be in initializing state when created and until it reaches its initial screen, usually while running the initialization path. If you are using a connection info set: the parameters of a single record are retrieved.
Ready	Connection connected to the host but not yet serving any requests (session/user). Waiting on the initial screen for a request.
Active	Attached to a session/user, currently serving a request.
Keep-alive	If a keep-alive path is defined in the connection pool, the connection will be in this state while the keep-alive path is running (typically no longer that a few seconds).
Recycle	If a recycle path is defined in the connection pool, the connection will be in this state while the recycle path is running.
Terminating	Disconnecting the connection. If a termination path is defined in the connection pool, it will be executed. If using a connection info set: this will release the parameters of the record used.

Note:

Connections can also appear to be in Broken state. This is not a "real" state, it just indicates that the connection cannot get to a Ready state. Broken connections do not count with regards to license capacity and do not use any system resources. Broken connections will be automatically cleared after a certain time. See Connection Pool Troubleshooting for more information.

Connection Pool Troubleshooting

Symptom	Explanation	Possible Reasons	How to Check
All my connections are broken	Host sessions become <i>broken</i> if they cannot get to Ready state. Ready means that: <ol style="list-style-type: none"> 1. The connection is alive. 2. If an "initial screen" is defined by the pool, the session must get there. 	Host is unavailable.	Try to connect to the host without the pool.
		Initialization path needs to be modified.	Create a connection outside the pool, run the initialization path and verify it gets to the "initial screen".
Connections disconnect after use but I want to reuse them	The pool may be configured to allow reusing a host connection that was used by a previous user. This would require returning the connection to the initial screen, either by implementing a robust navigation logic inside the invoked procedure or in the recycle path.	The Disconnect after usage option is checked.	Observe the state of the Disconnect after usage checkbox in pool tab of the connection pool.
		The logic of the user activity or the recycle path causes a disconnection from the host.	Capture a trace file for the user activity and recycle path. Observe that the host session is not terminated during the invocation of user activity or the recycle path.
		When a connection is returned to the pool and it is not in the "initial screen", and there is no recycle path that successfully navigates to the initial screen.	Capture a trace file for the user activity and recycle path. Observe that at the end of the user activity and the recycle path invocation, the connection is in the initial screen.
Connections remain Active after use	Upon finishing the user activity on a pool connection, the state of the connection would change according to the pool configuration and the state of the user session. The state of the connection would remain active if the session used by the pool has not ended.	A flow procedure creates an emulation session on a pool connection, but the "end session" node is not called on that session.	Verify the existence of an "end session" node in the flow procedure that creates the emulation session. Verify that the "end session" node is reached by capturing a trace file or by logging the invocation of the flow procedure.

Symptom	Explanation	Possible Reasons	How to Check
Connections remain Terminating after use	A pool may be terminating a connection based on the pool's configuration, connection count and the status of the session. When doing so, the connection status would become "Terminating" and the termination path would run on the connection. After the terminate path is completed, the connection is removed from the pool.	An exception in the terminate path is causing the host connection to either get stuck or to terminate before completing the path.	Capture a trace file that includes the invocation of the terminate path and observe that the path has completed successfully. Search for exceptions in the server logs that occurred during the invocation of the terminate path.
A used connection is disconnecting while being used		The host is disconnecting the pool connection based on the user activity or due to the host state.	Capture a trace file for the user activity, observe the last packet transmissions between the server and the host. Look in the server log for an error indicating a socket close around this timestamp. Also ask the host administrator to inspect the host's log for disconnections.
		Another host connection is using the same host credentials/device name and hijacking the session.	When capturing trace files, you would notice a trace file created at the time when a previous trace file is closed. Both traces will include send sections containing the same credentials or device name.

Note:

When recording trace files to capture the symptoms mentioned above, we recommend using the following variables in the trace file name: connection ID (%c), session ID (%u), creation time of file (%t). See the Recording Trace Files for more information.