Progress®



Use OpenEdge® Database Replication AI Streaming

OpenEdge®

Copyright

© 2021 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.

These materials and all Progress® software products are copyrighted and all rights are reserved by Progress Software Corporation. The information in these materials is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear therein. The references in these materials to specific platforms supported are subject to change.

DataDirect Cloud, DataDirect Connect, DataDirect Connect64, DataDirect XML Converters, DataDirect XQuery, DataRPM, Defrag This, Deliver More Than Expected, DevReach (and design), Icenium, Inspec, Ipswitch, iMacros, Kendo UI, Kinvey, MessageWay, MOVEit, NativeChat, NativeScript, OpenEdge, Powered by Chef, Powered by Progress, Progress, Progress Software Developers Network, SequeLink, Sitefinity (and Design), Sitefinity, Sitefinity (and design), SpeedScript, Stylus Studio, Stylized Design (Arrow/3D Box logo), Styleized Design (C Chef logo), Stylized Design of Samurai, TeamPulse, Telerik, Telerik (and design), Test Studio, WebSpeed, WhatsConfigured, WhatsConnected, WhatsUp, and WS_FTP are registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and/or other countries.

Analytics360, AppServer, BusinessEdge, Chef Automate, Chef Compliance, Chef Desktop, Chef Habitat, Chef WorkStation, Corticon.js, Corticon Rules, Data Access, DataDirect Autonomous REST Connector, DataDirect Spy, DevCraft, Fiddler, Fiddler Everywhere, FiddlerCap, FiddlerCore, FiddlerScript, Hybrid Data Pipeline, iMail, JustAssembly, JustDecompile, JustMock, KendoReact, NativeScript Sidekick, OpenAccess, PASOE, Pro2, ProDataSet, Progress Results, Progress Software, ProVision, PSE Pro, Push Jobs, SafeSpaceVR, Sitefinity Cloud, Sitefinity CMS, Sitefinity Digital Experience Cloud, Sitefinity Feather, Sitefinity Insight, Sitefinity Thunder, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, Supermarket, SupportLink, Unite UX, and WebClient are trademarks or service marks of Progress Software Corporation and/or its subsidiaries or affiliates in the U.S. and other countries. Java is a registered trademark of Oracle and/or its affiliates. Any other marks contained herein may be trademarks of their respective owners.

July 2021

Product version: Progress OpenEdge 12.4

Updated: 2021/06/28

Table of Contents

Replication AI Streaming overview....7

Configure Replication AI Streaming....9
Observe Replication AI Streaming in the Replication Monitor....9
Observe Replication AI Streaming in VSTs....10
Manually apply after-image extents....11
Understand automatic transition with AI Streaming....13

Reference Materials....17

Agent properties....17
Server properties....19
Transition properties summary....22
DSRUTIL applyextent qualifier....28
DSRUTIL recovery qualifier....31
PROMON User Control option....32
Updated Replication Monitor screens....38
Replication agent status....38
Replication server status....46
Updated Virtual System Tables....51
_Repl-AgentAIStreaming VST field descriptions....51
_Repl-Agent VST field descriptions....52
_Repl-AgentActivity VST field descriptions....59
_Repl-Server VST field descriptions....63

1

Replication AI Streaming overview

DBAs want their disaster recovery system to have the most up to date database changes so that in the event of a failure, recovery can occur quicker, and retain as much data as possible. OpenEdge Replication AI Streaming is designed to meet this requirement.

In previous OpenEdge Replication implementations (without AI streaming), AI blocks are sent to the targets when they are generated, and are received by the agent and applied on the target databases. If AI data cannot be processed fast enough on the targets, it is queued on the source database machine. In case of a source crash, this data can be lost and lead to data consistency issues on the targets.

To minimize the data loss caused by AI data queued on the source, Replication AI Streaming streams the AI blocks to the targets as they are generated. The agent on the target is multi-threaded to improve performance. Replication AI Streaming allows the targets to stay as close as possible to the source database because source AI data is hardened on the targets once it is received. Replication AI Streaming is also designed to reduce the latency as AI data is less likely to be queued on the source side (PICA). The following figure demonstrates how AI Streaming works.

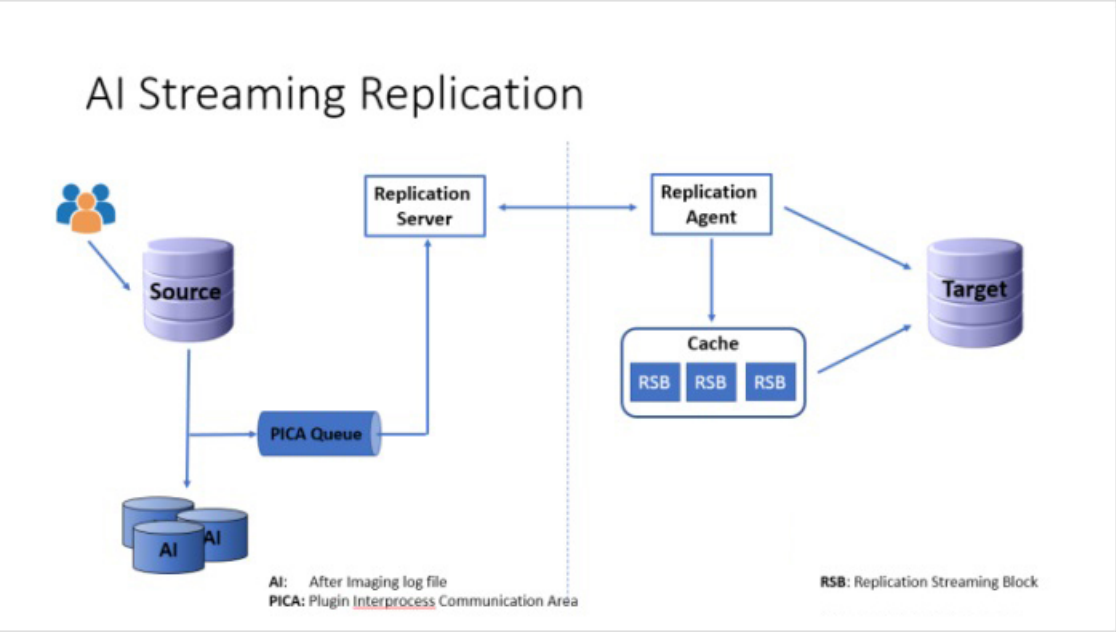

Streaming Replication Architecture

The Replication server sends AI data in the same way that prior releases did, but the Replication agent caches the incoming AI blocks in the Replication Stream Block (RSB) Cache. The blocks are then read from the RSB cache and applied to the target. Database admins have the ability to observe the Replication AI Streaming operations in the replication log with logging-level set to 3 on the replication agent.

Replication AI Streaming is enabled by default in OpenEdge. The topics that follow contain descriptions of how to configure your Replication Agent property file and how to observe the Replication AI Streaming operations in the Replication Monitor and through VSTs.

Note: For more information regarding OpenEdge Replication, see *Use Database Replication*.

For details, see the following topics:

- Configure Replication AI Streaming
- Observe Replication AI Streaming in the Replication Monitor
- Observe Replication AI Streaming in VSTs
- Manually apply after-image extents
- Understand automatic transition with AI Streaming

Configure Replication AI Streaming

Configuring Replication AI Streaming is controlled by properties. Add the following properties to the `[agent]` section of your Replication properties file(s):

- **`ai-streaming`**

  The `ai-streaming` property enables or disables this feature. The property is enabled by default. Set as follows:

  - **Enabled:** `ai-streaming=1`
  - **Not enabled:** `ai-streaming=0`

- **`rsb-cache-size`**

  The `rsb-cache-size` property sizes your RSB cache. Set as follows:

  `rsb-cache-size=`*n*

  Where *n* is the size of the RSB cache in units of Replication Stream Blocks. The default is 5000.

Observe Replication AI Streaming in the Replication Monitor

You can observe the functioning of Replication AI Streaming in the Replication Monitor. Start the monitor on your target machine, and select **Replication Agent Status**.

Observe the status of Replication AI streaming, and the cache, with the following lines of the output:

```
AI Streaming enabled:                       1
Replication streaming block cache:
    Cache size:                             5000
    Free cache entries:                     999
    Used cache entries:                     1
    Cache high water mark:                  175
    Number of times cache filled:           0
```

Observe the status of any latency between the target and the source, with the following lines of the output:

```
    AI Block Information:
        Source RDBMS Block (Seq / Block):       4 / 6905
        Last Received Block (Seq / Block):      4 / 6884
        Last Processed Block (Seq / Block):     4 / 6884
    Latency Information:
        Repl Server behind Source DB by:        0     second(s)
        Current Source Database Transaction:    2303
        Last Transaction Applied to Target:     2303
        Target Current as of (Target, Source):  Tue May 15 18:15:58 2018, Tue May 15
18:15:58 2018 with delta of 000:00:00
        Number of blocks to send on Source:     0 as of Tue May 15 18:16:40 2018
        Number of blocks to process on Target:  0 as of Tue May 15 18:16:40 2018
        Time delay to process the last block:   0 second(s)
```

For descriptions of each line of the **Replication Agent Status** output, see Replication agent status on page 38.

Observe Replication AI Streaming in VSTs

The _Repl-AgentAIStreaming table within the Replication Virtual System Tables (VSTs) provides insight into the configuration and functioning of Replication AI Streaming.

- _Repl-AgentAIStreaming

  The _Repl-AgentAIStreaming provides information about the RSB cache. It contains the following fields:

  - `_ReplAgtAIStr-AgentID`
  - `_ReplAgtAIStr-AgentName`
  - `_ReplAgtAIStr-TotalCacheEntries`
  - `_ReplAgtAIStr-FreeCacheEntries`
  - `_ReplAgtAIStr-UsedCacheEntries`
  - `_ReplAgtAIStr-CacheHWM`
  - `_ReplAgtAIStr-TimesCacheFilled`

  For detailed descriptions of the fields, including data types, see _Repl-AgentAIStreaming VST field descriptions on page 51.

- _Repl-Agent

  The _Repl-Agent VST features the following field:

  - `_ReplAgt-AIStreaming`

  For detailed descriptions of this field, and the rest of the _ReplAgent VST, see _Repl-Agent VST field descriptions on page 52.

- _Repl-AgentActivity

  The _Repl-AgentActivity VST features the following fields:

- _ReplAgtAct-LastRecvBlockSeq
- _ReplAgtAct-LastRecvBlockBlock
- _ReplAgtAct-ServerBlocksToSend
- _ReplAgtAct-LastBlockCountTime
- _ReplAgtAct-BlockProcDelay

For a detailed description of these fields, and the rest of the _ReplAgentActivity VST, see _Repl-AgentActivity VST field descriptions on page 59.

Manually apply after-image extents

During normal conditions, the target database is automatically updated with data from the source database. The updates are transmitted as blocks of data from the source's AI transaction log. After a failure condition, it is possible that the most recent blocks were unable to reach the target. For example, a TCP/IP failure could lose packets that were awaiting transmission.

You can recover the missing data by manually applying the after-image extents that contain the data that has not already been applied to the target. However, you can do so only under the following conditions:

- The `transition` property (under the `[server]` section in the properties file) must be set to `manual`.
- The agent must be in a pre-transition state.
- The storage device that contains the source database AI extents must be accessible by the target machine. You cannot save AI extents on the same system where the source is running because the extents would not be accessible after a system crash. Network area storage (NAS) or storage area network (SAN) devices are ideal for storing AI extents.
- If Replication AI streaming, the `ai-streaming` property under the `[agent]` section in the properties file, is enabled, the RSB cache must be empty before applying extents. If it is not empty, attempting to apply an extent, results in the following message:

```
The extent cannot be applied. There are %d unapplied blocks remaining in the RSB
 cache. (19292)
```

Note: While the replication agent main thread is in the applying-extent state, you cannot run the `triggertranistion` command. With AI Streaming replication, you cannot run subsequent `applyextent` commands while the agent is in the applying-extent state. The writer thread must first empty the RSB cache and the main agent thread will move out of the applying-extent state.

To apply AI extents:

1. Display failure recovery information.

Use the following command:

```
dsrutil target-db-name -C recovery agent
```

2. Determine the AI extent number from the command output. Select the last-applied AI extent or, if the last-applied extent was completely processed, select the next available extent.

   The following is an example of the relevant section of the command output:

```
                            .
                            .
                            .
  Last AI Extent processed
   AIMAGE BEGIN date:              Tue Oct 18 13:33:31 2005
   AIMAGE NEW date:                Tue Oct 18 13:48:55 2005
   After Image File Number:        3
   File Last Opened:               Tue Oct 18 13:48:55 2005
   Completely Applied to Target:   No
```

   Determine the file number of the last-applied AI extent. In this example the last-applied extent is 3. Since the value after `Completely Applied to Target` is `No`, you will use 3 as the *`extent-name`* in the next step.

   If the value after `Completely Applied to Target` was `Yes`, you would use 4 as the *`extent-name`* in the next step, since extent 3 would have already been completely applied.

Note: The `After Image File Number` corresponds to the `Seqno` that appears in the extent list when you use the applyextent qualifier with DSRUTIL. For more information, see DSRUTIL applyextent qualifier on page 28.

---

3. Use the appropriate AI extent file number in the following command:

```
dsrutil target-db-name -C applyextent qualified-extent-name
```

   As the command executes, it performs the following validation:

   - The extent must exist and be valid.
   - The status of the extent must be `FULL`, `BUSY`, or `LOCKED`.
   - The AI extent file number either must be the same as the AI extent file number for the last AI block processed by the agent, or it must be the next extent file number (if the previous AI extent was completely processed).

   If the source database after-image extents cannot be applied to the target database before transition, the target database will not contain any of the unapplied source database transactions. In this case, unless the transactions can be recreated and re-entered after transition, you **will** lose an indeterminate number of transactions.

   After you apply the AI extent, you can begin a manual transition of the target. For more information, see Manually apply after-image extents on page 11.

Understand automatic transition with AI Streaming

The behavior of the data in your RSB cache, when an automatic transition occurs, is controlled by the transition property, `apply-streaming-data`. The `apply-streaming-data` property may be set to either 1 or 0. The setting indicates the following:

- **`apply-streaming-data=1`** — When the transition timeout expires, the replication agent waits for all entries in the RSB cache to be applied to the target database before continuing with auto-transition.
- **`apply-streaming-data=0`** — When the transition timeout expires, the replication agent discards all entries in the RSB cache and continues with auto-transition.

The `apply-streaming-data` property defaults to 0 unless explicity set to 1 in the properties file. Setting the `apply-streaming-data` property to values other than 1 is the equivalent of setting it to 0.

For all possible configurations of replication (one target, two target, replication set)[1], when `apply-streaming-data=1`, all the data in the RSB cache is written to the target database before transition.

Auto-transition for one-target or two-targets (no replication set)

In this configuration, the agents handle data in the RSB cache with the following steps:

1. The agent reaches time to transition.
2. Check the `apply-streaming-data` setting.
   - If 1, check if the RSB cache is empty.
     - **Yes** — initiate auto-transition, writing this message to the log file:

```
Agent's RSB cache is empty.  Ready for auto-transition.   (19317 )
```

     - **No** — reset the time to transition and return to **Step 1**, after writing these messages to the log file:

```
%l<num> blocks exist in RSB cache.  The Agent will transition after all blocks
 in the RSB cache have been applied to the database. (19300)
The transition time has been reset to allow for RSB cache to empty before auto
 transition. (19301)
```

   - If 0, check if the RSB cache is empty.
     - **Yes** — initiate auto-transition, writing this message to the log file:

```
Agent's RSB cache is empty.  Ready for auto-transition.   (19317 )
```

[1] Progress Software recommends configuring two-target replication environments as a replication set to avoid rebasing after transition.

- **No** — initiate auto-transition, writing this message to the log file:

```
Automatic transition is starting, there are up to %l<num> blocks in the RSB
cache that will not be processed. (19299)
```

Auto-transition for a two target replication set

In this configuration, the agents handle data in the RSB cache with the following steps. If at any time during the processing of the RSB cache prior to transition, the agents lose inter-agent communication, the replication set is broken and the agents transition as single agents, using the prior process.

1. The local agent reaches time to transition. (Each agent acts as a local agent, communicating with the other "remote" agent.)
2. Check if the RSB cache has been processed.
   - **Yes** — Proceed to **Step 3**.
   - **No** — Process the RSB cache.

     Check the `apply-streaming-data` setting:
     - **1** — Process RSB cache data before transition.

       Check if the RSB cache is empty.
       - **Yes** — Ping the remote agent that the RSB cache is empty, and the local agent is ready for transition. Write the following message to the log file:

```
Agent's RSB cache is empty.  Ready to transition with replication-set. (19311)
```

       - **No** — Ping the remote agent that the RSB cache is not empty, the local agent is processing the cache, and is not ready for transition. Write the following messages to the log file and repeat this step:

```
%l<num> blocks exist in RSB cache.  The Agent will transition after all blocks
 in the RSB cache have been applied to the database. (19300)
The transition time has been reset to allow for RSB cache to empty before
auto transition. (19301)
```

     - **0** — Do not process RSB cache data before transition.

       Check if the RSB cache is empty.
       - **Yes** — Ping the remote agent that RSB cache is empty and the local agent is ready for transition. Write the following message to the log file:

```
Agent's RSB cache is empty.  Ready to transition with replication-set.
(19311)
```

- **No** — Discard the remaining blocks in the RSB cache. Ping the remote agent that cache is empty and the local agent is ready for transition. Write the following message to the log file:

```
Agent is ready to transition with replication-set.  There are up to %l<num>
 blocks in the RSB cache that will not be processed.(19312)
```

Note: When the remote agent receives a ping from the local agent, it sets its time to transition to the current time to begin handling of its own RSB cache. The remote agent writes the following message to its log file:

```
Remote agent has signaled that it processed its RSB data and is ready for
auto-transition. (19316)
```

3. Check communications status with the remote agent.

   Has the remote agent sent a message indicating that it has completed processing its RSB cache and is ready to transition?

   - **No** — Ping the remote agent to confirm it is still connected. Reset the time to transition, and return to **Step 1** after writing the following message to the log file:

```
Waiting for remote agent to process its RSB data and signal that it is ready to
auto-transition.(19314)
```

   - **Yes** — Proceed to **Step 4**.

4. Confirm that both agents are ready to proceed with auto-transition.

   Both agents write the following message to their log file:

```
Both agents in the replication-set have processed their RSB data and are ready for
auto-transition.(19315)
```

   The priority agent initiates transition, and the other agent waits to be notified by the priority agent that transition processing is proceeding.

2

Reference Materials

The topics that follow contain new or previously released reference material, updated to support Replication AI Streaming. For complete reference details regarding OpenEdge Replication, see *Use Database Replication*.

For details, see the following topics:

- Agent properties
- Server properties
- Transition properties summary
- DSRUTIL applyextent qualifier
- DSRUTIL recovery qualifier
- PROMON User Control option
- Updated Replication Monitor screens
- Updated Virtual System Tables

Agent properties

Agent properties define the configuration of the local agent running for the target database.

Agent properties are found after the `[agent]` directive in the agent properties file. A properties file can only have one `[agent]` section. The properties file must be stored in the directory where the target database resides.

Note: Although you may create a source database without an `[agent]` section, it is recommended that every replication database include an `[agent]` section to allow the database to become a named target.

Property names and values in the section are separated by an equal sign (for example, `name=agent1`).

Note: When you add comments to a properties file, place them on a new line that begins with the pound sign (#). Do not specify a comment on the same line as a property value definition.

The following table provides a list of agent properties and their values.

Table 1: Agent properties

Property name	Value	Description
`ai-streaming`	`0` | `1`	Specifies if AI Streaming is enabled. • 0 — AI Streaming is disabled. • 1 — AI Streaming is enabled. This is the default. When AI Streaming is enabled, the replication agent utilizes multi-threading to pull the AI blocks from the source, and store and apply the data on the target. AI Streaming is enabled by default.
`connect-timeout`	*`seconds`*	Specifies how many seconds the OpenEdge Replication agent will wait for connection from the OpenEdge Replication server before the replication agent shuts itself down. *`seconds`* is an integer greater than or equal to 120 (which is the default) and less than or equal to 86,400. Using this property means you do not have to perform a forced shutdown on your target database. If the OpenEdge Replication agent does not receive a connection attempt from the OpenEdge Replication server before the number of seconds specified have elapsed, the OpenEdge Replication agent will terminate and allow some limited system-level target database connections.
`database`	*`db_name`*	Specifies the target database name.
`ipver`	`ipv4` | `ipv6`	Specifies whether to use TCP/IP IPv4 or TCP/IP IPv6 between the Replication server and the Replication agent. If the target database broker is using TCP/IP IPv6, OpenEdge Replication must also use IPv6.
`listener-maxport`	*`port-number`*	Specifies the maximum TCP port number for the agent. *`port-number`* must be greater than the value of *`listener-minport`*+1 and less than the maximum allowable port number on the system. For UNIX the maximum port number is 65534. For Windows the maximum port number is 32765.

Property name	Value	Description
`listener-minport`	*`port-number`*	Specifies the minimum TCP port number. The agent selects a port in a range between the values specified by *`listener-minport`* and *`listener-maxport`*. *`port-number`* must be greater than 1024 and less than the value of *`listener-maxport`*.
`maximum-polling-delay`	*`milliseconds`*	Specifies the maximum value, in milliseconds, for a polling delay. By default, the maximum polling delay is 500 ms. Set *`milliseconds`* to a value between 500 and 1000, inclusive.
`minimum-polling-delay`	*`milliseconds`*	Specifies the minimum value, in milliseconds, for a polling delay. By default, the minimum polling delay is 5 ms. Set *`milliseconds`* to a value between 1 and 30, inclusive. The polling delay starts at the minimum, specified by `minimum-polling-delay`, and increases by 5 ms increments, during periods of inactivity to the maximum delay, specified by `maximum-polling-delay`. Once activity occurs during polling, the delay is reset to the `minimum-polling-delay` value.
`name`	*`agent_name`*	Specifies the OpenEdge Replication agent name. By convention, the *`agent-name`* should match the name specified in the source database properties file. *`ALL`* is not allowed as an agent name. Each agent must have a unique name.
`rsb-cache-size`	*`cache-size`*	Specifies the size of the Replication Streaming cache in units of an AI block, when AI Streaming is enabled (`ai-streaming=1`). If AI Streaming is not enabled, this property is ignored. *`cache-size`* is a non-negative integer that determines the size of the cache. The default is 5000.

Server properties

Server properties are found after the `[server]` directive in the OpenEdge Replication source properties file. A properties file can only have one `[server]` section.

Property names and values are separated by an equal sign (for example, `transition=auto`).

Note: When you add comments to a properties file, place them on a new line that begins with the pound sign (#). Do not specify a comment on the same line as a property value definition.

The following table provides a list of server properties and their values.

Table 2: Server properties

Property name	Value	Description
`control-agents`	*`agent_name[`*`,` `agent_name``]`	Specifies a comma-separated list of OpenEdge Replication agent names. The comma must be followed by a space. The agent name must match the control agent name specified in the `[control-agent]` properties section head. `ALL` is not allowed as an agent name. Each agent must have a unique name.
`database`	*`db_name`*	Specifies the source database name.
`defer-agent-startup`	*`minutes`*	Specifies how long, in minutes, the server attempts to connect to an agent if the first connection attempt is unsuccessful. The default value is 0 (not active). Set *`minutes`* to 0, or a value between 5 and 10080, inclusive.
`increment-polling-delay`	*`milliseconds`*	Specifies the value, in milliseconds, to increase the polling delay. By default, the polling delay increment is 5 ms. For configurations with AI Streaming enabled, the default is 1 ms.[2] Set *`milliseconds`* to a value between 1 and 1000, inclusive. The value must be less than `maximum-polling-delay`.
`maximum-polling-delay`	*`milliseconds`*	Specifies the maximum value, in milliseconds, for the polling delay. By default, the maximum polling delay is 500 ms. Set *`milliseconds`* to a value between 500 and 1000, inclusive.

[2] The connect message between the agent and server informs the server when AI Streaming is enabled. In response, the server sets the `minimum-polling-delay` to 20 ms, and the `increment-polling-delay` to 1 ms. These values override any previous setting, including values set in the properties file. The change to the property settings is written to the server's database log file when the first AI Streaming agent is connected.

Property name	Value	Description
`minimum-polling-delay`	*`milliseconds`*	Specifies the minimum value, in milliseconds, for a polling delay. By default, the minimum polling delay is 5 ms. For configurations with AI Streaming enabled, the default is 20 ms.[2] Set *`milliseconds`* to a value between 1 and 30, inclusive. The polling delay starts at the minimum, specified by `minimum-polling-delay`, and automatically increases, by the amount specified by `increment-polling-delay`, during periods of inactivity to the maximum delay, specified by `maximum-polling-delay`. Once activity occurs during polling, the delay is reset to the `minimum-polling-delay` value.
`repl-Keep-Alive`	*`seconds`*	Specifies a time-out period for communications between a server and its agents. If a connection between the server and agent is not verified before the time-out expires, failure recovery begins. The minimum value is 90 seconds; there is no maximum value. The default is 300 seconds.
`schema-Lock-Action`	`wait | force`	Specifies the action an agent takes if an exclusive schema lock is not granted. Possible actions are: • `wait` — The agent waits until the exclusive schema lock is granted. The server blocks until the exclusive schema lock is granted. • `force` — The agent attempts to acquire the exclusive schema lock five times. If the fifth attempt fails, the agent disconnects all users from the target and makes another attempt. If the last attempt fails, the server and all agents terminate. When schema update activity completes, the server and target can be restarted. **Note:** Values are case-sensitive. Use `wait` or `force`, not `Wait` or `Force`.

Property name	Value	Description
agent-shutdown-action	recovery | normal	Specifies the action an agent takes during a shutdown when the replication server ends. Possible actions are: • recovery — The agent will remain active but in a standby state waiting for the replication server to reconnect. • normal — The agent will terminate; the target database will stay up.
transition	auto | manual	Specifies how to transition the target database to a normal database. If a critical agent cannot reconnect with the source database within the time specified in the transition-timeout property, the following occurs: • auto —The agent automatically transitions the target database. • manual —The agent listens for the server until it reconnects with the source database or until the DSRUTIL *db-name* -C transition command transitions the target database. **Note:** Values are case-sensitive. Use auto or manual, not Auto or Manual.
transition-timeout	*seconds*	Specifies the number of seconds the agent waits before it performs auto-transition. This property is ignored when transition=manual. The value is incremented by the sum of the connect-timeout for all configured agents.

Transition properties summary

The following table provides a summary of the transition properties. The table lists each property in the repl.properties file, identifies the property type, and provides a property description.

Table 3: Transition properties

Property name	Type and length	Description
`replication-set`	integer	Defines if source and target replicas transition together. Possible values are as follows: • **`0`** — The source and targets do not transition together (the default). • **`1`** — The source and targets transition together. For more information, see 'Transition of a Replication Set' in *Use Database Replication*.
`database-role`	character[15]	The new role of the database once it is transitioned. The possible values for this property are as follows: • **`reverse`** — The role of the database is reversed: A source database becomes a target database, and a target database becomes a source database. • **`normal`** — The role of the database becomes that of a normal database; the database is no longer enabled for replication once the transition is performed. This is the default value.
`responsibility`	character[15]	This property must contain one of the following values: • **`primary`** — This database is the primary database. • **`secondary`** — This database is the secondary database. This property is currently informational.

Property name	Type and length	Description
`restart-after-transition`	integer	The database can be automatically restarted after transition is performed. Valid values for the property are `0` and `1`. When the property is set to `1`, the following `*-startup-arguments` properties must be supplied, or the database startup will fail: • If the database role is normal, you must specify the `normal-startup-arguments`. • If the database role is reversed, you must specify the `source-startup-arguments` andthe `target-startup-arguments`. **Note:** The `restart-after-transition property` behaves differently during online transition. Setting the property value to `1` does not automatically restart the database. The value is ignored. For more information, see 'Plan for transition' in *Use Database Replication*.
`start-secondary-broker`	integer	Set to **1** to automatically start a secondary broker after transition.
`source-startup-arguments`	character[256]	If the database is transitioned to a source database, these arguments are used when the database is started. The arguments will be appended to the PROSERVE command (used to start the database). In most cases, the only argument specified here should be `-pf` followed by a parameter file name. For example:`-pf` *`db-name`*`.source.pf`
`target-startup-arguments`	character[256]	If the database is transitioned to a target database, these arguments are used when the database is started. The arguments will be appended to the PROSERVE command. In most cases, the only argument specified here should be `-pf` followed by a parameter file name. For example: `-pf` *`db-name`*`.target.pf` Because the database is a target, you must also specify the following arguments as an indication to the broker to listen on the TCP/IP port specified with -S *`port-number`* or *`service-name`*: `-S` { *`port-number`* | *`service-name`* }

Property name	Type and length	Description
`normal-startup-arguments`	character[256]	If the database is transitioned to a normal database, these arguments are used when the database is started. The arguments will be appended to the PROSERVE command. In most cases, the only argument specified here should be `-pf` followed by a parameter file name. For example: `-pf db-name.normal.pf`
`source-secondary-broker-startup-arguments`	character[256]	For secondary broker. See the description of `source-startup-arguments` for details.
`target-secondary-broker-startup-arguments`	character[256]	For secondary broker. See the description of `target-startup-arguments` for details.
`normal-secondary-broker-startup-arguments`	character[256]	For secondary broker. See the description of `normal-startup-arguments` for details.
`auto-begin-ai`	integer	Specify **1** to automatically begin AI after a target-to-source database transition. After-imaging can be started for a database that has AI areas.
`transition-to-agents`	character[256]	Transition to the first agent in the list when a failure occurs. If the first agent is not available, transition to the second agent in the list. If `replication-set=1`, the first agent in the `transition-to-agents` list is selected to become the new source. A valid value for this property is any configured agent name; separate the names by a comma if you are listing more than one. For example: `agent1,agent2` **Note:** If both `dr-transition-to-agents` and `transition-to-agents` are defined, `dr-transition-to-agents` is used in *recovery* transition and `transition-to-agents` is used for *failover* transition. If `dr-transition-to-agents` is not defined,`transition-to-agents` is used for both types of recovery.

Property name	Type and length	Description
dr-transition-to-agents	character[256]	**For Recovery Transition only.**Transition to the first agent in the list when a failure occurs. If the first agent is not available, transition to the second agent in the list. If replication-set=1, the first agent in the transition-to-agents list is selected to become the new source. A valid value for this property is any configured agent name; separate the names by a comma if you are listing more than one. For example: agent1,agent2 **Note:** If both dr-transition-to-agents and transition-to-agents are defined, dr-transition-to-agents is used in *recovery* transition and transition-to-agents is used for *failover* transition. If dr-transition-to-agents is not defined,transition-to-agents is used for both types of recovery.
auto-add-ai-areas	integer	This property instructs the transition process to automatically add AI areas to the database if: • It is transitioned to a source database. • There are currently no AI areas for the database. If the value specified for the property is **0**, AI areas are not added to the database. If the value specified is **1**, AI areas are automatically added to the database by using the structure file specified in the ai-structure-file property.
ai-structure-file	character[256]	Name of the structure file that contains the list of AI areas to add.

Property name	Type and length	Description
`backup-method`	character[15]	Specifies the backup method performed before AI is enabled. The following values are valid: • **`mark`** — Mark the database backed up by using the following command: `rfutil` *`db-name`* `-C mark backedup` Marking the database as backed up does not allow future AI extents to be used when recovering from a disaster. • **`full-offline`** — The database will be backed up offline by using the PROBKUP utility. The backup is performed in two steps: The first backup is a full backup, which is performed before AI is enabled for the database. The second backup is an incremental backup, which is performed after AI is enabled and after the role of the database is changed. • **`full-online`** — An online backup is performed after the database is restarted, which occurs after the database has been transitioned. **Note:** The `backup-method` property behaves differently during online transition. For more information, see 'Offline backup and online failover transition' in *Use Database Replication*.
`backup-arguments`	character[256]	The arguments required for the full online and offline backups that are performed for the database. At a minimum, the target file or device must be specified in these arguments for both online and offline backups. To avoid overwriting a backup, do not use the same target file or device for both the backup and the incremental backup. Do not use backup validation parameters (such as `-vp` and `-vf`). The arguments should begin with *`device-name`*.
`incremental-backup-arguments`	character[256]	The arguments required for the offline incremental backup performed after AI is enabled and the database's role is reversed. The arguments should begin with *`device-name`*.

Property name	Type and length	Description
`recovery-backup- arguments`	character[256]	The arguments required for the full online and offline backups that are performed for the database. At a minimum, the target file or device must be specified in these arguments for both online and offline backups. To avoid overwriting a backup, do not use the same target file or device for both the backup and the incremental backup. Do not use backup validation parameters (such as `-vp` and `-vf`). The arguments should begin with *`device-name`*.
`apply-streaming-data`	integer	**For auto-transition only.** If Replication AI Streaming is enabled, and there is data in the RSB cache at the time of transition, indicate if the data should be applied or ignored before transition. The following values are valid: • **1** — Apply all data in the RSB cache before transition. • **0** — Do not apply any data in the RSB cache before transition. **Note:** Replication AI Streaming is enabled by default.

DSRUTIL applyextent qualifier

Allows you to apply AI extents, generated by your source database, directly to a target database. This is useful when there is a source failure and there are AI blocks in transit between the server and agents.

Syntax

```
dsrutil db-name -C applyextent extent-name
```

Parameters

`db-name`

The name of the target database.

`extent-name`

The fully qualified path name of the extent to apply. See the topic below for details on how to identify an extent to apply.

Your Replication configuration must meet the following requirements in order to apply extents:

- The agent must be in pre-transition state.
- The `transition` property must be set to `manual`.
- The source must save AI extents to remote storage that is accessible to the target's agent.
- If you have AI streaming enabled, your RSB cache must be empty. The `applyextent` command will cause the AI blocks contained in the extent to be written to the RSB cache. If it is not empty, attempting to apply an extent, results in the following message:

```
The extent cannot be applied. There are %d unapplied blocks remaining in the RSB
 cache. (19292)
```

Note: While the replication agent main thread is in the applying-extent state, you cannot run the `triggertranistion` command. With AI Streaming replication, you cannot run subsequent `applyextent` commands while the agent is in the applying-extent state. The writer thread must first empty the RSB cache and the main agent thread will move out of the applying-extent state.

Identifying the AI extent to apply

Identifying the AI extent to apply is a multi-step process.

- First, determine which extent to apply by using the following command:

```
dsrutil db-name -C recovery Agent
```

The recovery output provides the output **`After Image File Number`** and **`Completely Applied to Target`**, as shown:

```
                            .
                            .
                            .
  Last AI Extent processed
   AIMAGE BEGIN date:              Tue Oct 18 13:33:31 2005
   AIMAGE NEW date:                Tue Oct 18 13:48:55 2005
   After Image File Number:        3
   File Last Opened:               Tue Oct 18 13:48:55 2005
   Completely Applied to Target:   No
```

If **`Completely Applied to Target`** is **`No`**, use the **`After Image File Number`** provided. If **`Completely Applied to Target`** is **`Yes`**, increment the **`After Image File Number`** value by 1.

- After determining the AI Image File number, locate the file, as follows:
  - If the source database is available, use the following command to generate a list of after-image extents for the source database:

```
rfutil source-db-name -C aimage list
```

Executing the command produces the following output:

```
Extent:  1
Status:  Full
  Type:  Variable Length
  Path:  /vobs_repl/solaris/bin/ks1.a1
  Size:  120
  Used:  1
 Start:  Wed Oct 26 08:32:12 2005
 Seqno:  1

Extent:  2
Status:  Full
  Type:  Variable Length
  Path:  /vobs_repl/solaris/bin/ks1.a2
  Size:  4728
  Used:  4534
 Start:  Wed Oct 26 08:32:14 2005
 Seqno:  2

Extent:  3
Status:  Busy
  Type:  Variable Length
  Path:  /vobs_repl/solaris/bin/ks1.a3
  Size:  4728
  Used:  4456
 Start:  Wed Oct 26 08:33:29 2005
 Seqno:  3

Extent:  4
Status:  Empty
  Type:  Variable Length
  Path:  /vobs_repl/solaris/bin/ks1.a4
  Size:  120
  Used:  0
 Start:  N/A
 Seqno:  0
```

Match the Seqno from this output to the **`After Image File Number`** provided by the DSRUTIL recovery output. Apply all BUSY and FULL extents beginning with this extent.

- If the source database is unavailable, but its after-image extents are available on a SAN or NAS device that the target machine has access to, you can determine the first after-image extent to apply, by examining the AI files.

  1. Change to the directory where the source after-image extents are stored, and then execute the following command:

     ```
     rfutil db-name -C aimage scan -a after-image-extent-name
     ```

     The database you specify in the *db-name* location in the command can be any database provided it has AI enabled and an AI block size that is the same as the original source database. However, this command must be run offline, so the database specified must be offline.

The command produces the following output:

```
After-image dates for this after-image file: (1633)
    Last AIMAGE BEGIN Wed Oct 26 08:32:12 2005 (1640)
    Last AIMAGE NEW Wed Oct 26 08:33:29 2005 (1641)
    This is aimage file number 3 since the last AIMAGE BEGIN. (1642)
    This file was last opened for output on Wed Oct 26 08:33:29 2005. (1643)


41706 notes were processed. (1634)
0 in-flight transactions. (3785)
614 transactions were started. (1635)
614 transactions were completed. (11138) At the end of the .ai file, 0
transactions were still active. (1636)
```

2. Match the **This is aimage file number *n*** from this output to the **After-Image File Number** provided by the DSRUTIL recovery output. Apply all BUSY and FULL extents beginning with this extent.

For more information on applying extents, see Manually apply after-image extents on page 11.

DSRUTIL recovery qualifier

Displays the replication recovery information.

Syntax

```
dsrutil db-name -C recovery { server | agent }
```

Parameters

db-name

The name of the database to perform the action on. The name of the database must be the first argument and must be a valid name.

The replication recovery information for the server looks similar to the following:

```
Replication version:              4.0
  Date created:                     Wed Oct 26 08:32:58 2005
  Date last written:                Wed Oct 26 08:33:46 2005

  Replication server information:
    Number of agents:               1
    Number of unused agents:        1
    Last modified:                  Wed Oct 26 08:32:08 2005
    Master block update count:      21

  Remote Agent information:
    Remote Agent 1
      Identification:               1
      Agent name:                   agent1
      Last AI block acknowledged:   area: 14, seq: 3, loc: 4259840, offset: 3189
      Last modified:                Wed Oct 26 08:33:29 2005
      Last AI block ACK time:       Wed Oct 26 08:33:46 2005
      Remote agent host:            localhost
      Remote agent database:        ks2
```

The replication recovery information for the agent looks similar to the following:

```
Replication version:              4.0
  Date created:                     Wed Oct 26 08:32:57 2005
  Date last written:                Wed Oct 26 08:34:36 2005

  Replication local agent information:
    Last Block:                     Complete
    Last block received location:   area: 14, seq: 3, loc: 4554752, offset: 0
    Last block processed location:  area: 0, seq: 0, loc: 0, offset: 0
    Last block ACKed location:      area: 14, seq: 3, loc: 4259840, offset: 3189
    Last block received:            no date
    Last block ACKed:               no date
    ID of the last TX begin:        4345
    ID of the last TX end:          4345
    Time of last TX end:            Wed Oct 26 08:33:46 2005
    Last AI Extent processed
      AIMAGE BEGIN date:            Wed Oct 26 08:32:12 2005
      AIMAGE NEW date:              Wed Oct 26 08:33:29 2005
      After-Image File Number:      3
      File Last Opened:             Wed Oct 26 08:33:29 2005
      Completly Applied to Target:  No
```

Note that the **After-Image File Number** is the same number that is returned as **Seqno** by the following command:

```
rfutil source-db-name -C aimage list
```

PROMON User Control option

Displays information for each database process.

Sample output for PROMON User Control option

The following shows an example of this option's output. Not all fields in the User Control display fit in a standard window width, but all are described below.

```
User Control: by user number

  Usr:Ten  Name      Domain     Type     Wait  Table:Part  Dbkey  Trans    PID  Sem Srv   SrvUSr
 Login  Time IVP#  Remote Address
    0:0    dba                  BROK     --       0:0          0      0   6492   0    0   0
06/24/14
    5:3    lobtenant lobtenant  SELF/ABL --       0:0          0      0   7364   0    0   0
06/24/14
    6:4    lobtenant lobtenant  SELF/ABL --       0:0          0      0   4780   0    0   0
06/24/14
    7:2    tenant2   tenant2    SELF/ABL --       0:0          0      0   6932   0    0   0
06/24/14
    8:1    tenant1   tenant1    SELF/ABL --       0:0          0      0   3616   0    0   0
06/24/14
    9:0    dba                  MON      --       0:0          0      0   7892   0    0   0
06/24/14
   23:0                   0     REMC/ABL --      31:47         0  98676      0  11   2    0
06/24/14 01:05 IPV4 172.16.18.92

RETURN - repeat, U - continue uninterrupted, Q - quit
```

`Usr:Ten`

Displays the user number and tenant ID.

User numbers are assigned sequentially to database processes as they are started. The broker process is always user (0).

`Name`

The user name for end-user (client) processes.

`Domain`

The user's domain name.

`Type`

The table below lists the distinct types of database processes that can appear in the Type field.

Table 4: Process type values

Value	Process type
AIMD	After-image file management utility daemon
AIW	After-image writer (AIW) process
APW	Asynchronous page writer
AUDA	Audit archive utility
AUDL	Audit load utility

Value	Process type
BAT	Batch user
BIND	Binary dump utility
BINL	Binary load utility
BIW	Before-image writer
BKUP	Online backup
BROK	Broker process
FMA	OpenEdge Management (formerly Fathom™ Management) Monitor
IACT	Index activate utility
IDXC	Index compact utility
IDXF	Index fix utility
IDXM	Index move utility
MON	Monitor process
PASA	Progress Application Server for OpenEdge Session Manager
PASN	Progress Application Server for OpenEdge Session
PSA	PROSTRCT ADD utility
QUIE	PROQUIET utility
REMC	Remote client
RFUT	RFUTIL utility
RPLA	OpenEdge Replication Site Replication Agent
RPLE	OpenEdge Replication Site Replication online enable
RPLS	OpenEdge Replication Site Replication Server
RPLT	OpenEdge Replication Site Replication Transition
RPLU	OpenEdge Replication Site Replication Utilities
SELF	Self-service client
SERV	Remote-user server
SHUT	Database shutdown process

Value	Process type
SQSC	SQLSCHEMA utility
SQDP	SQLDUMP utility
SQLD	SQLLOAD utility
SQSV	SQL server
TBLM	Table move utility
WDOG	Watchdog

`Wait`

Indicates that a process is waiting for a lock or latch. The table below lists common values and describes the resource or event associated with each value.

Table 5: Possible wait column values

Value	Description
Lock waits	
REC	Row (record) lock
IX	Index lock
SCHE	Schema lock
TRAN	Transaction commit
TXB	Transaction end update mode lock
TXE	Transaction end commit mode lock
TXS	Transaction end share lock
TXX	Transaction end exclusive lock
Resource waits	
AIRD	After-image buffer read from disk
AIWR	After-image buffer write to disk
BIRD	Before-image buffer read from disk
BIWR	Before-image buffer write to disk
BKEX	Exclusive lock on database buffer

Value	Description
BKSH	Share lock on database buffer
BKSI	Share lock with intent to update on database buffer
BUFF	Shared buffer
DB	Database server
DBBK	Database buffer being backed up
DBRD	Database buffer read from disk
DBWR	Database buffer write to disk
DEAD	Process is waiting to be removed from the process list[3]
RGET	Row get lock
SRPL	Replication agent acknowledgment
Latch waits	
AIB	After-image buffer latch
BFP	Buffer pool latch
BHT	Buffer pool hash table latch
BUF	Buffer pool buffer latch
CPQ	Checkpoint queue latch
DLC	Index delete chain
GST	Global storage pool latch
IXD	Index delete chain latch
LKF	Lock table free chain latch
LKP	Lock table purge chain latch
LKT	Lock table latch
LRS	Secondary LRU chain (not currently implemented)
LRU	Buffer pool LRU chain latch
MTL_RPL	Replication latch

[3] Seen when a process is terminated with the UNIX `kill-9` command

Value	Description
MTX	Log space allocation latch
PWQ	Page writer queue latch
SCH	Shared schema cache latch
SEQ	Sequence cache latch
TXQ	Transaction queue latch
TXT	Transaction table latch
USR	User table latch

You may also see the following values listed in the **Wait** column:

- BF2
- BF3
- BF4
- BIB
- LRI
- LRX
- OM

The previous values are for short-duration (micro- to milliseconds) latches used internally by the OpenEdge RDBMS.

`Table:Part Dbkey`

If the user has a REC wait, the combination of the table number, partition number, and the db key, identifies the record. If the table is not table partitioned, the partition number is 0.

`Trans`

Transaction (task) number, if one is active. After the broker starts, numbers are assigned sequentially to each transaction that accesses the database.

`PID`

The process ID as assigned by the operating system. The PID column typically displays (0) for remote clients.

`Sem`

The number of the semaphore the process is using. Each process uses exactly one semaphore. The database engine uses two semaphores (numbers 1 and 2) internally and assigns the remaining semaphores one at a time (starting with the broker) as database processes log in.

`Srv`

For remote clients (**REMC**), the server number of the server the client runs against.

`SrvUsr`

For remote clients (**REMC**), the user number of the server the client runs against.

`Login`

The date the process started.

`Time`

The time the process started.

`IPV#`

The IP version.

`Remote Address`

Updated Replication Monitor screens

The topics that follow describe the Replication Monitor screens updated for Replication AI Streaming.

Replication agent status

If you select the **Replication agent status** from the **Startup** menu of the DSRUTIL Monitor Utility, the Figure 1: OpenEdge Replication agent status screen on page 38 is displayed.

Figure 1: OpenEdge Replication agent status screen

```
OpenEdge Replication Monitor                  Page 1
Database:  /scratch/doc/120/demo/repl2

    Database is enabled as OpenEdge Replication:  Target

    Agent:
        Name:                                   agent1
        ID:                                     1
        Host name:                              127.0.0.1
        State:                                  Normal Processing
        Ready:                                  Yes
        Critical:                               No
        Method:                                 Asynchronous
    Agent is waiting for:                       Nothing
    Maximum bytes in TCP/IP message:            8496
    Server/Agent connection time:               Tue May 15 18:08:41 2018
    Delay Interval (current / min / max):       10 / 5 / 500
    Transition information:
        Type:                                   Manual
    The last block received at:                 Tue May 15 18:16:40 2018
    Activity information:
        Blocks received:                        1857
        Blocks processed:                       1858
        Blocks acknowledged:                    0
        Notes processed:                        184042
        Transactions started:                   0
        Transactions ended:                     0
        Synchronization points:                 29
```

```
    AI Block Information:
        Source RDBMS Block (Seq / Block):       4 / 6905
        Last Received Block (Seq / Block):      4 / 6884
        Last Processed Block (Seq / Block):     4 / 6884
    Latency Information:
        Repl Server behind Source DB by:        0     second(s)
        Current Source Database Transaction:    2303
        Last Transaction Applied to Target:     2303
        Target Current as of (Target, Source):  Tue May 15 18:15:58 2018, Tue May 15
18:15:58 2018 with delta of 000:00:00
        Number of blocks to send on Source:     0 as of Tue May 15 18:16:40 2018
        Number of blocks to process on Target:  0 as of Tue May 15 18:16:40 2018
        Time delay to process the last block:   0 second(s)
    Connect timeout:                            1800
    Listener port range:                        1025-65500
    Current listener port:                      1037
    Additional transition information:
        Replication set:                        1
        Database role:                          Reverse
        Transition to agents:                   agent1,agent2
        Recovery transition to agents:
        Restart after transition:               1
        Automatically begin AI:                 ?
        Automatically add AI areas:             1
        AI structure file:                      addai.st
        Backup method:                          Full-offline
        Transition end trigger:                 0
        AI Streaming enabled:                   1
        Apply Streaming Data:                   1
    Replication streaming block cache:
        Cache size:                             1000
        Free cache entries:                     999
        Used cache entries:                     1
        Cache high water mark:                  175
        Number of times cache filled:           0
    SSL Enabled:                                Yes
    SSL No Host Verify:                         Yes
    SSL Key Alias:                              repltest
    SSL Key Alias Password:                     12345f6d23456789
    SSL Key Store Path:                         /bigdisk/jlopez/REPLSETUP_1TARG/keys
    SSL Cert Store Path:                        /bigdisk/jlopez/REPLSETUP-1TARG/certs
    SSL Cipher Suites:
AES128-SHA256,AES256-SHA256,DHE-RSA-AES256-GCM-SHA384
    SSL Protocols:                              TLSv1.1,TLSv1.2,TLSv1.3
```

As shown in the previous figure, the first two lines of the **Replication agent status** screen are the utility title line and the fully qualified name of the database specified by the user. The remaining lines of the Replication Agent Status screen are defined in the table that follows.

Table 6: Replication Agent Status details

Field	Description
`Database is enabled as OpenEdge Replication`	Indicates whether the database is enabled as a replication source or a target.
`Agent`	Lists basic remote agent information.
`Name`	Lists the name of the remote agent.

Field	Description
`ID`	Lists the remote agent identification number.
`Host name`	Lists the host on which the target database associated with the remote agent resides.
`State`	Shows agent processing information: • **Normal processing**—The server and the agent are performing normal processing. • **Initial connection**—The agent is waiting for the initial connection from server. • **Startup synchronization**—The server and the agent are synchronizing. • **Initialization**—The agent is being initialized by the server. • **Initial connection failed**—The server could not connect to an agent. • **Invalid target database configuration**—Something in the target database does not match the source. • **Agent terminated**—The target database shut down or the agent terminated using `Terminate agent`. • **Online backup of target database** —An online target database backup is being performed. • **Recovery synchronization**— Recovery synchronization is being performed. • **Recovery failed**—Recovery was unsuccessful. • **Unknown**—The agent is in an unknown state.
`Critical`	Shows whether the remote agent is a critical agent (**Yes**) or is not (**No**).
`Method`	Lists the replication method: **asynchronous**.
`Agent is waiting for`	Shows why the agent might be waiting. The field value might be one of the following: • **Nothing**—Indicates that no waiting is taking place • **Schema lock request**—Indicates that the server is waiting for the agent to acquire the schema lock so the database schema can be changed • **Encryption objects**—For an encrypted database, the agent's wait status, or what the agent is waiting for
`Maximum bytes in TCP/IP message`	Shows the maximum number of bytes used for the TCP/IP communication messages.

Field	Description
Server/Agent connection time	Shows the date and time at which the server and agent connected.
Delay Interval (current / min / max)	Shows the amount of time, in milliseconds, the agent will wait between polling the TCP/IP connection for information to be replicated. **Current** is the current value, **min** is the minimum value, and **max** is the maximum value. Polling is used to increase the performance of the server, and the wait is used to limit the amount of overhead when no data is available to be replicated.
Transition information	Shows a group of remote agent-related transition information.
Type	Shows the type of transition to be performed: • **Manual**—Indicates that intervention is required to complete the transition of a target database to a normal OpenEdge database • **Automatic**—Indicates that a transition from a target database to a normal OpenEdge database will take place without intervention
The last block received at	Shows the date and time the last block was sent to the agent.
Activity information	Shows a selection of activity information.
Blocks received	Shows the number of blocks received from the server.
Blocks processed	Shows the number of blocks processed.
Blocks acknowledged	Shows the number of blocks acknowledged.
Notes processed	Shows the number of AI transaction notes processed by the agent.
Transactions started	Shows the number of transactions started on the agent.
Transactions ended	Shows the number of transactions completed on the agent.
Synchronization points	Shows the number of synchronization points that occurred.
AI Block Information	Identifies the current AI block that is being written by the source database and provides the last AI block that was sent by Replication.
Source RDBMS Block (Seq / Block)	Provides the current source database block that the OpenEdge Replication agent is processing.
Last Received Block (Seq / Block)	Provides the AI location, defined by AI sequence number and AI block number, of the last block received by the replication agent.
Last Processed Block (Seq / Block)	Provides the AI location, defined by AI sequence number and AI block number, of the last processed block by the replication agent.

Field	Description
`Latency Information`	Provides latency information that shows how far behind OpenEdge Replication is in updating the target database. This is important if the target database is transitioned due to a source database failure.
`Repl Server behind Source DB by`	Provides, in seconds, the difference between the `Target Current As Of` field and the timestamp of the last AI block flush on the source database, at the time when the last AI block was sent from the replication server to the replication agent.
`Current Source Database Transaction`	Provides the current source database transaction that the OpenEdge Replication agent is processing.
`Last Transaction Applied to Target`	Provides the last source database transaction that was applied to the target database.
`Target Current As Of`	Provides the date of the last transaction applied to the target database by the OpenEdge Replication agent based on source database time.
`Number of blocks to send on Source`	Provides the number of AI blocks to send on the source database.
`Number of blocks to process on Target`	Provides the number of AI blocks to process on the target database.
`Time delay to process the last block`	Provides the time it takes for an AI block to be processed by the agent after it was received.
`Connect timeout`	Number of seconds during which the agent attempts to reconnect to the server.
`Listener port range`	Range of ports for the agent.
`Current listener port`	Current port in use.
`Replication set`	Set to 1 if Replication set is enabled: 0 otherwise.
`Database role`	Role of agent database in transition.
`Transition to agents`	Displays the names of the agents for transition. If the next line (`Recovery transition to agents`) is set, then this field represents the agent list for *failover* transition. If not set, then this line represents the agent list for both *failover* and *recovery* transition. If `replication-set=1`, then the first agent in this list is selected to become the new source.
`Recovery transition to agents`	Displays the names of the agents for *recovery* transition.
`Restart after transition`	Indicates whether a restart should occur after transition. Set to 1 for yes or 0 for no.

Field	Description
Automatically begin AI	For the replication set only, indicates whether after-imaging should automatically be started on transition. Set to 1 to begin after-imaging, or 0 to not. If the replication set is not enabled, set to the unknown value, "?"
Automatically add AI areas	Indicates whether after-image areas should be automatically added to the target database that is transitioning to a source. Set to 1 to add after-image areas or 0 not to add AI.
AI structure file	The name of the structure file to use if automatically adding AI areas. If **Automatically add AI** is not enabled, then set to the unknown value, "?"
Backup method	The backup method for the database during transition. Valid entries include: • Full offline • Full online • Mark
Transition end trigger	Indicates whether a transition end trigger should run at the end of transition. Set to 1 to indicate that the trigger property is enabled, and set to 0 to indicate that there is no transition end trigger.
AI Streaming enabled	Indicates whether **Replication AI Streaming** is enabled. 1 indicates enabled; 0 indicates not enabled.
Apply Streaming Data	If **Replication AI Streaming** is enabled, then this field indicates how to treat data in the RSB cache during automatic transition. 1 indicates that data in the RSB cache is applied before transition; 0 indicates that data in the RSB cache is not applied before automatic transition.
Replication streaming block cache	Shows a group of related information about the cache for Replication AI Streaming.
Cache size	Displays the cache size in units of Replication Streamed Blocks (RSB).
Free cache entries	Displays the number of free cache entries in the RSB cache.
Used cache entries	Displays the number of used cache entries in the RSB cache.
Cache high water mark	Displays the maximum number of cache entries that were in use at one point in the life cycle of the cache.
Number of times cache filled	Displays the number of times the RSB cache filled.

Field	Description
SSL Enabled	Specifies if SSL use is enabled for the replication database. • 0—SSL use is disabled. This is the default. • 1—SSL use is enabled. The property affects all replication processes that can run against the replication database.
SSL No Host Verify	A comma-separated list of the SSL protocol versions that can be used for any OpenEdge Replication SSL network connections for the replication database. Supported SSL protocols are: • TLSv1.1 • TLSv1.2 • TLSv1.3
SSL Key Alias	For an SSL server, the key-alias of the SSL key that will be used in OpenEdge Replication SSL.
SSL Key Alias Password	For an SSL server, the key-alias-password of the SSL key that is used in OpenEdge Replication SSL. The value of this property is encrypted to avoid exposing the actual password to unauthorized viewers of the properties file. You can generate the encrypted form of the password using the `genpassword` utility.
SSL Key Store Path	For an SSL server, the absolute path of the directory containing the SSL key and certificate for OpenEdge Replication SSL.
SSL Cert Store Path	For an SSL client, the absolute path of the directory containing the SSL certificate for OpenEdge Replication SSL.

Field	Description
SSL Cipher Suites	A comma separated list of the SSL cipher suites that may be used for any OpenEdge Replication SSL network connections for the replication database. Cipher suites are: • AES128-SHA256 • AES256-SHA256 • DHE-RSA-AES256-GCM-SHA384
SSL Protocols	A comma separated list of the SSL protocol versions that may be used for any OpenEdge Replication SSL network connections for the replication database. Supported SSL protocols are: • TLSv1.1 • TLSv1.2 • TLSv1.3

Replication server status

If you select **Replication server status** from the **Startup** menu of the DSRUTIL Monitor Utility, then the Figure 2: OpenEdge Replication server status screen on page 46 is displayed.

Figure 2: OpenEdge Replication server status screen

```
OpenEdge Replication Monitor                    Page 1

  Database:  C:\OE120WRK\repl1

  Database is enabled as OpenEdge Replication:  Source

  Server is:                                    In Normal Processing
  Number of configured agents:                  2
  Delay Interval (current / min / max):         500 / 5 / 500
  Recovery information:
      State:                                    No recovery being performed
      Agents needing recovery:                  0
      Agents connected:                         0
      Agents in synchronization:                0
  Transition information:
      Type:                                     Manual
      Replication set:                          1
      Database role:                            Reverse
      Transition to agents:                     agent1, agent2
      Recovery transition to agents:            agent1, agent2

      Restart after transition:                 1
      Automatically begin AI:                   ?
      Automatically add AI areas:               1
      AI structure file:                        addai.st
      Backup method:                            Mark
      Transition end trigger:                   0
  Configured agents:                            agent1, agent2
  Repl keep alive:                              300
  Schema lock action:                           Wait
  Agent shutdown action:                        Recovery
  SSL Enabled:                                  Yes
  SSL No Host Verify:                           Yes
  SSL Cert Store Path:                          /.../SSL/124/DiffHostCerts/certs
  SSL Cipher Suites:                            AES256-SHA256
  SSL Protocols:                                TLSv1.3
```

If you enabled deferred agent startup, then you will see the following additional lines after the `Number of configured agents` line.

```
  Defer Agent Startup :
      Continue connection attempts until:     Mon Feb 26 16:37:46 2018
      Deferred Agent startup will expire in : 9 Hr 58 Min 37 Sec
      Next connection attempt in :            4 Min 41 Sec
      Connection attempts performed :         3
      Agent(s) currently connected :          0
```

As shown in the Figure 2: OpenEdge Replication server status screen on page 46, the first two lines of the **Replication server status** screen are the utility title line and the fully qualified name of the database you specify. The remaining lines are defined in Table 7: Replication server status on page 47 .

Table 7: Replication server status

Field	Description
`Database is enabled as OpenEdge Replication`	Indicates whether the database is enabled as a replication source or target.
`Server is`	Describes how the server is processing information: • **Normal processing**—The server is processing information in the normal fashion. • **Performing initialization**—The server is performing initialization. • **Performing startup** • **synchronization**—The server is in the process of synchronizing the target databases with the source database. • **Connection**—The server connection to agent(s). • **Configuring connected agents**—The server is performing a handshake with the agent. • **Performing failure recovery**—The server is attempting to recover from a connection failure. • **Connecting to Agents**—The server is in deferred agent startup mode, and waiting to connect to agents. • **Unknown**—The server is in an unknown state.
`Number of configured agents`	Shows the number of agents currently configured to operate with the server.
`Defer Agent Startup`	Shows information related to deferred agent startup.
`Continue connection attempts until`	Shows the time, day, and date when connection attempts will stop.
`Deferred Agent Startup will expire in`	Shows the remaining duration of deferred agent startup.
`Next connection attempt in`	Shows when the next connection attempt starts.
`Connection attempts performed`	Shows how many connection attempts have occurred.
`Agent(s) currently connected`	Shows the number of connected agents.

Field	Description
`Delay Interval (current / min / max)`	Shows the amount of time, in milliseconds, the server will wait between polling the database for information to be replicated. **Current** is the current value, **min** is the minimum value, and **max** is the maximum value. Polling is used to increase the performance of the server, and the wait is used to limit the amount of overhead when no data is available to be replicated.
`Recovery information`	Shows failure-recovery-related details.
`Agents needing recovery`	Shows the number of remote agents requiring failure recovery.
`Agents connected`	Shows the number of remote agents currently connected to the server. (Note that since the database is not going through recovery, the number of **Agents connected** indicated under **Recovery information** is correctly reflected as zero.)
`Agents in synchronization`	Shows the number of remote agents in the process of being brought up to date.
`Transition information`	Shows Replication transition details.
`Type`	Shows the type of transition to be performed: • **Manual**—Indicates intervention is required to complete the transition of a target database to a normal OpenEdge database • **Automatic**—Indicates a transition from a target database to a normal OpenEdge database will take place without intervention
`Transition timeout limit`	Shows the maximum amount of time that will elapse before the transition of a target database to a normal OpenEdge database. **Note:** This line does not appear in manual transition, as the timeout limit is shown for automatic transition only.
`Replication set`	Indicates if Replication is configured for Replication sets. Set to 1 if configured, 0 if not.
`Database role`	Indicates the role of the source database in transition. Valid values include: • Normal • Reverse

Field	Description
`Transition to agents`	Displays the names of the agents for transition. If the next line (`Recovery transition to agents`) is set, this field represents the agent list for *failover* transition; if not set, this line represents the agent list for both *failover* and *recovery* transition. If `replication-set=1`, the first agent in this is selected to become the new source.
`Recovery transition to agents`	Displays the names of the agents for recovery transition.
`Restart after transition`	Displays if the source database is configured to be restarted after transition. Set 1 for restart, 0 for no restart.
`Automatically begin AI`	When Replication Set =1, indicates whether to start after-imaging on a target database that is transition to a source. Set 1 to start after-imaging, 0 to not start. If Replication set is not enabled, set to the unknown value, "?"
`Automatically add AI areas`	Indicates if after-image areas should be automatically added to the target database that is transitioning to a source. Set to 1 to add after-image areas, 0 not to add AI.
`AI structure file`	The name of the structure file to use if automatically adding AI areas. If Automatically add AI is not enabled, set to the unknown value, "?"
`Backup method`	The backup method used during transition. Valid entries include: • Full-offline • Full-online • Mark
`Configured agents`	Displays the names of the configured replication target agents
`Repl keep alive`	Displays the time, in seconds, that replication will wait and attempt to re-establish communication before starting transition.
`Schema lock action`	Action to take if schema lock exists. Valid entries include: • Force • Wait

Field	Description
`Agent shutdown action`	Action to take when agent shuts down. Valid entries include: • Normal • Recovery • Wait
`SSL Enabled`	Specifies if SSL use is enabled for the replication database. • 0 — SSL use is disabled. This is the default. • 1 — SSL use is enabled. The property affects all replication processes that may run against the replication database.
`SSL No Host Verify`	A comma separated list of the SSL protocol versions that may be used for any OE Replication SSL network connections for the replication database. Supported SSL protocols are: • TLSv1.1 • TLSv1.2 • TLSv1.3
`SSL Cert Store Path`	For an SSL client, the absolute path of the directory containing the SSL certificate of OE Replication SSL.

Field	Description
SSL Cipher Suites	A comma separated list of the SSL cipher-suites that may be used for any OpenEdge Replication TLS network connections for the replication database. Cipher suites are: • AES128-SHA256 • AES256-SHA256 • DHE-RSA-AES256-GCM-SHA384
SSL Protocols	A comma separated list of the SSL protocol versions that may be used for any OE Replication SSL network connections for the replication database. Supported SSL protocols are: • TLSv1.1 • TLSv1.2 • TLSv1.3

Updated Virtual System Tables

The topics that follow describe the Virtual System Tables updated for Replication AI Streaming.

_Repl-AgentAIStreaming VST field descriptions

The table below contains the field descriptions of the OpenEdge Replication VST, _Repl-AgentAIStreaming.

Table 8: Virtual system table _Repl-AgentAIStreaming field descriptions

Field name	Field number	Data type	Description
_ReplAgtAIStr-AgentID	2	INT64	Agent number The ID of the OpenEdge Replication agent Format ->>>>>>>>>>9

Field name	Field number	Data type	Description
_ReplAgtAIStr-AgentName	3	CHARACTER	Agent Name Name of the OpenEdge Replication agent as configured by the name property in the [agent] section of the repl.properties file Format X(32)
_ReplAgtAIStr-TotalCacheEntries	4	INT64	Total number of entries in the Replication Streaming Block cache in units of Replication Stream Blocks (RSB) The maximum size of the cache is specified by the rsb-cache-size property in the [agent] section of the repl.properties file Format ->>>>>>>>>>9
_ReplAgtAIStr-FreeCacheEntries	5	INT64	Free cache entries in the Replication Streaming Block cache in units of RSBs The number of RSB cache entries currently not in use Format ->>>>>>>>>>9
_ReplAgtAIStr-UsedCacheEntries	6	INT64	Used cache entries in the Replication Streaming Block cache in units of RSBs The number of RSB cache entries currently in use Format ->>>>>>>>>>9
_ReplAgtAIStr-CacheHWM	7	INT64	AI Stream Cache high water mark The maximum number of RSB cache entries that have been in use at one point in the life cycle of the cache Format ->>>>>>>>>>9
_ReplAgtAIStr-TimesCacheFilled	8	INT64	Number of times the Replication Streaming Block cache has filled Format ->>>>>>>>>>9

_Repl-Agent VST field descriptions

Table 1 contains the field descriptions of the OpenEdge Replication virtual system table (VST) _Repl-Agent.

Table 9: Virtual system table _Repl-Agent field descriptions

Field name	Field number	Data type	Description
_ReplAgt-AgentID	2	INT64	Agent number The ID of the OpenEdge Replication agent Format ->>>>>>>>>>9
_ReplAgt-AgentName	3	CHARACTER	Agent name Name of the OpenEdge Replication agent as configured in the repl.properties file Format X(32)
_ReplAgt-ConnectTime	4	CHARACTER	Agent connect time Date and time when the OpenEdge Replication server successfully connected to the OpenEdge Replication agent Format X(24)
_ReplAgt-DBName	5	CHARACTER	Database name The fully qualified database name that the OpenEdge Replication agent is connected to Format X(128)
_ReplAgt-ServerHost	6	CHARACTER	Server host name The name of the host where the OpenEdge Replication server is running Format X(128)
_ReplAgt-Port	7	INTEGER	Port number The port number that the OpenEdge Replication server uses to connect to this OpenEdge Replication agent Format ->>>>>>>>>>9
_ReplAgt-BlocksReceived	8	INT64	Number of blocks received The number of after-image (AI) blocks that this OpenEdge Replication server received Format ->>>>>>>>>>9

Field name	Field number	Data type	Description
_ReplAgt-BlocksProcessed	9	INT64	Number of blocks processed. The number of AI blocks that this OpenEdge Replication agent processed Format ->>>>>>>>>>9
_ReplAgt-BlocksACK	10	INTEGER	Number of block acknowledgements sent by this agent The number of block acknowledgements that this OpenEdge Replication agent sent to the OpenEdge Replication server Format ->>>>>>>>>>9
_ReplAgt-NotesProcessed	11	INTEGER	Number of notes processed The number of AI notes that this OpenEdge Replication agent processed Format ->>>>>>>>>>9
_ReplAgt-LastTRID	12	INTEGER	Transaction ID of the last transaction that began The transaction ID of the last Transaction Begin encountered Format ->>>>>>>>>>9
_ReplAgt-Method	13	CHARACTER	Replication method The type of replication that is being performed with this OpenEdge Replication agent, **A** — Asynchronous Format X(1)
_ReplAgt-Status	14	INTEGER	A value from the **Value** column, as reported by the following command: `DSRUTIL -C status -detail` Format >>>>>>>>>9 For more information, see 'When the detail or verbose argument is used' in *Use Database Replication*.

Field name	Field number	Data type	Description
_ReplAgt-CommStatus	15	INTEGER	Agent's communication status The OpenEdge Replication agent's TCP communication status: • **1** — Connected • **2** — Disconnected Format >>>>>>>>>9
ReplAgt-DatabaseRole	16	CHARACTER	Database role. Valid values are: • Replication source • Replication target Format X(8)
ReplAgt-Critical	17	INTEGER	Critical agent flag. Valid values are: • **1** — Critical • **0** — Not critical Format >>>>>>>>>9
ReplAgt-MaxMsgSize	18	INT64	Maximum number of bytes in a TCP message Format ->>>>>>>>>>>>>>>>>>9
ReplAgt-MinDelayInterval	19	INTEGER	Minimum message delay interval (in seconds) Format >>>>>>>>>9
ReplAgt-MaxDelayInterval	20	INTEGER	Maximum message delay interval (in seconds) Format >>>>>>>>>9
ReplAgt-TransitionType	21	CHARACTER	Transition type. Valid values are: • Manual • Automatic Format X(24)
ReplAgt-TimeoutLimit	22	INTEGER	Time (in seconds) that the agent tries to reconnect before entering transition Format >>>>>>>>>9

Field name	Field number	Data type	Description
_ReplAgt-ConnectTimeout	23	INT64	Agent's connect timeout. Agent tries to connect to the server for this many seconds Format ->>>>>>>>>>>>>>>>>>>9
_ReplAgt-MinPort	24	INT64	Minimum listener port number Format ->>>>>>>>>>>>>>>>>>>9
_ReplAgt-MaxPort	25	INT64	Maximum listener port number Format ->>>>>>>>>>>>>>>>>>>9
_ReplAgt-ReplicationSet	26	INTEGER	Replication set enabled Format >>>9
_ReplAgt-TransDatabaseRole	27	CHARACTER	Role database takes after transition. Valid values include: • Normal • Reverse Format X(32)
_ReplAgt-TransitionToAgents	28	CHARACTER	Transition to agents Format X(100)
_ReplAgt-RestartAfterTrans	29	BOOLEAN	Restart database after transition Format yes/no
_ReplAgt-AutoBeginAI	30	BOOLEAN	Automatically begin after-imaging during transition Format yes/no
_ReplAgt-AutoAddAI	31	BOOLEAN	Automatically add after-image files during transition Format yes/no
_ReplAgt-AIFile	32	CHARACTER	AI structure file name for adding after-imaging during transition Format X(256)

Field name	Field number	Data type	Description
_ReplAgt-BackupMethod	33	CHARACTER	Backup method during transition. Valid values include: • Full-offline • Full-online • Mark Format X(32)
_ReplAgt-ReservedChar	34	CHARACTER	Reserved
_ReplAgt-ReservedInt	35	INTEGER	Reserved
_ReplAgt-TransEndTrigger	36	INTEGER	Indicates whether a transition end trigger is to execute on transition. Valid values include: • **0** — Execute a transition end trigger • **1** — Do not execute a transition end trigger Format >>>9
_ReplAgt-AIStreaming	37	BOOLEAN	Indicates whether AI streaming is enabled. Valid values include: • **yes** — Enabled • **no** — Disabled Format yes/no
_ReplAgt-DRTransitionToAgents	38	CHARACTER	Recovery transition to agents list. Format X(100)
_ReplAgt-ApplyStreamingData	39	BOOLEAN	Indicates whether AI streaming data residing in the RSB cache should be applied to the target before auto-transition. • **yes** — Apply data • **no** — Do not apply data Format yes/no

Field name	Field number	Data type	Description
_ReplAgt-SSLEnabled	40	BOOLEAN	Specifies whether or not SSL use is enabled for the replication database. • 0—SSL use is disabled. This is the default. • 1—SSL use is enabled. The property affects all replication processes that may run against the replication database.
_ReplAgt-SSLKeyAlias	41	CHARACTER	For an SSL server, the key alias of the SSL key that will be used in OpenEdge Replication SSL.
_ReplAgt-SSLKeyAliasPassword	42	CHARACTER	For an SSL server, the key alias password of the SSL key that will be used in OpenEdge Replication SSL. The value of this property is encrypted to avoid exposing the actual password to unauthorized viewers of the properties file. You can generate the encrypted form of the password using the genpassword utility.
_ReplAgt-SSLKeyStorePath	43	CHARACTER	For an SSL server, the absolute path of the directory containing the SSL key and the certificate for OpenEdge Replication SSL.

Field name	Field number	Data type	Description
_ReplAgt-SSLCipherSuites	44	CHARACTER	A comma-separated list of the SSL cipher suites that can be used for any OpenEdge Replication SSL network connections for the replication database. Cipher suites are: • AES128-SHA256 • AES256-SHA256 • DHE-RSA-AES256-GCM-SHA384
_ReplAgt-SSLProtocols	45	CHARACTER	Specifies whether or not SSL use is enabled for the replication database. • 0—SSL use is disabled. This is the default. • 1—SSL use is enabled. The property affects all replication processes that may run against the replication database.

Note

- If the replication agent is not running, the _replagent VST returns no records.

_Repl-AgentActivity VST field descriptions

The table below contains the field descriptions of the OpenEdge Replication VST, _ReplAgentActivity.

Table 10: Virtual system table _ReplAgentActivity field description

Field name	Field number	Data type	Description
_ReplAgtAct-AgentID	2	INT64	Agent number Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-AgentName	3	CHARACTER	Agent name Format X(32)
_ReplAgtAct-ConnectTime	4	CHARACTER	Date and time of replication agent/server connect Format X(24)

Field name	Field number	Data type	Description
_ReplAgtActBlocksReceived	5	INT64	The number of after-image (AI) blocks this agent has received Format >>>>>>>>9
_ReplAgtAct-BlocksProcessed	6	INT64	The number of AI blocks this agent has processed Format >>>>>>>>9
_ReplAgtAct-BlocksACK	7	INTEGER	The number of block acknowledgments this agent has sent to the server Format >>>>>>>>9
_ReplAgtAct-NotesProcessed	8	INTEGER	The number of AI notes this agent has processed Format >>>>>>>>9
_ReplAgtAct-LastTRID	9	INTEGER	The TRID of the last transaction on the target Format >>>>>>>>9
_ReplAgtAct-Status	10	INTEGER	Current status of this replication agent Format >>>>>>>>9
_ReplAgtAct-CommStatus	11	INTEGER	Current TCP communication status of this agent. Valid values are: • **1** — Connected • **2** — Disconnected Format >>>>>>>>9
_ReplAgtAct-State	12	CHARACTER	Agent state. Valid values are: • Terminating • In Normal Processing • In Synchronization Processing • In Forced Transition • In Pre-Transition • Ready For Manual Transition Format X(48)

Field name	Field number	Data type	Description
_ReplAgtAct-Ready	13	INTEGER	Agent ready status. Valid values are: • **0** — no • **1** — yes Format >>>>>>>>9
_ReplAgtAct-WaitingFor	14	CHARACTER	What this agent is waiting for. Valid values include: • Nothing • Schema Lock Request • Encryption Objects • Schema Object state change Format X(48)
_ReplAgtAct-CurDelayInterval	15	INTEGER	Current message delay interval (in seconds) Format >>>>>>>>9
_ReplAgtAct-TimeToTransition	16	CHARACTER	Time remaining before transition Format X(24)
_ReplAgtAct-LastBlockReceived	17	CHARACTER	Last block received on/at Format X(24)
_ReplAgtAct-TransStarted	18	INT64	Number of transactions started Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-TransEnded	19	INT64	Number of transactions ended Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-SyncPoints	20	INT64	Number of synchronization points Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-SrcRDBMSSeq	21	INT64	The AI sequence number of the most recent AI block on the source database. -1 if not available Format ->>>>>>>>>>>>>>>>>>9

Field name	Field number	Data type	Description
_ReplAgtAct-SrcRDBMSBlock	22	INT64	The AI block number of the most recent AI block on the source database. -1 if not available Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-LastRecvBlockSeq	23	INT64	The AI sequence number of the last received AI block on the target database. -1 if not available Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-LastRecvBlockBlock	24	INT64	The AI block number of the last received AI block on the target database. -1 if not available Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-LastBlockSeq	25	INT64	Last processed AI block sequence number -1 if not available Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-LastBlockBlock	26	INT64	Last processed AI block -1 if not available Format ->>>>>>>>>>>>>>>>>>9
_ReplAgtAct-ServerLag	27	INTEGER	The difference, in seconds, between the `Target Current As Of` field and the timestamp of the last AI block flush on the source database, at the time when the last AI block was sent from the replication server to the replication agent. -1 if not available Format >>>>>>>>9
_ReplAgtAct-SourceTRID	28	INTEGER	Last TRID from the source database -1 if not available Format >>>>>>>>9
_ReplAgtAct-LastSyncTimeTgt	29	CHARACTER	Time that the target was last current Format X(24)

Field name	Field number	Data type	Description
_ReplAgtAct-LastSyncTimeSrc	30	CHARACTER	Time that the source was last current Format X(24)
_ReplAgtAct-LastSyncDelta	31	CHARACTER	Delta in Source/Target Sync Format X(24)
_ReplAgtAct-ServerBlocksToSend	32	INTEGER	Number of AI Blocks in the Replication server's PICA queue at time of the last AIBLOCK message Format >>>>>>>>9
_ReplAgtAct-LastBlockCountTime	33	CHARACTER	Date/time the Replication server last sent the PICA queue information to the replication agent Format X(24)
_ReplAgtAct-BlockProcDelay	34	INTEGER	The elapsed time between when the last timestamped block was received by the Replication agent, and when it was applied to the target database Format >>>>>>>>9
_ReplAgtAct-ReservedChar	35	CHARACTER	Reserved for future use
_ReplAgtAct-ReservedInt	36	INTEGER	Reserved for future use

_Repl-Server VST field descriptions

The table below contains the field descriptions of the OpenEdge Replication virtual system table, _Repl-Server.

Table 11: Virtual system table _Repl-Server field description

Field name	Field number	Data type	Description
_ReplSrv-id	2	INT64	Server Number Format ->>>>>>>>>>>>>>>>>>9
_ReplSrv-AgentCount	3	INTEGER	Agents Count Number of agents that the OpenEdge Replication server is controlling Format >>>>>>>>>9

Field name	Field number	Data type	Description
_ReplSrv-BlocksSent	4	INTEGER	After-image(AI) Blocks Sent Number of AI Blocks the OpenEdge Replication server has sent to connected OpenEdge Replication agents Format >>>>>>>>>9
_ReplSrv-StartTime	5	CHARACTER	Server Start Time Date and time when the OpenEdge Replication server was started Format X(24)
_ReplSrv-LastBlockSentAt	6	CHARACTER	Last Block Sent At Date and time when the OpenEdge Replication server was started Format X(24)
ReplSrv-Database	7	CHARACTER	Database file name Format X(24)
ReplSrv-DatabaseRole	8	CHARACTER	Database role. Valid values are: • Replication source • Replication target Format X(8)
ReplSrv-ServerState	9	CHARACTER	Server state. Valid values are: • Performing Initialization • Connecting to Agents • Configuring Connected Agents • Performing Failure Recovery • Performing Startup Synchronization • In Normal Processing • Unknown Format X(36)
ReplSrv-DeferAgentConnUntil	10	CHARACTER	Date and time that deferred agent connection attempts will cease Format X(24)

Field name	Field number	Data type	Description
ReplSrv-DeferAgentExpire	11	CHARACTER	Time remaining until deferred agent startup expires Format X(24)
ReplSrv-NextConnectIn	12	CHARACTER	Time until next connection attempt Format X(24)
ReplSrv-ConnectAttempts	13	INT64	Number of connection attempts performed Format ->>>>>>>>>>>>>>>>>>9
ReplSrv-AgentsConnected	14	INTEGER	Number of agents currently connected Format >>>>>>>>>9
ReplSrv-CurDelayInterval	15	INTEGER	Current message delay interval (in seconds) Format >>>>>>>>>9
ReplSrv-MinDelayInterval	16	INTEGER	Minimum message delay interval (in seconds) Format >>>>>>>>>9
ReplSrv-MaxDelayInterval	17	INTEGER	Maximum message delay interval (in seconds) Format >>>>>>>>>9
ReplSrv-RecoveryState	18	CHARACTER	Server recovery state. Valid values are: • No Recovery • Just Entered Recovery • Server Attempting Connection • Initialize Synchronization with Agents • Synchronizing Agents • Recovery Complete Format X(42)
ReplSrv-RecAgentsToRecover	19	INTEGER	Number of agents needing recovery Format >>>>>>>>>9

Field name	Field number	Data type	Description
ReplSrv-RecAgentsConnected	20	INTEGER	Number of agents connected and not needing recovery Format >>>>>>>>>9
ReplSrv-AgentsInSync	21	INTEGER	Number of agents in synchronization Format >>>>>>>>>9
ReplSrv-TransitionType	22	CHARACTER	Transition type. Valid values are: • Manual • Automatic Format X(32)
ReplSrv-TransitionTimeout	23	INTEGER	For automatic transition, the maximum amount of time (in seconds) that the target database will wait until it transitions to a normal database. Format >>>>>>>>>9
_ReplSrv-ReplicationSet	24	INTEGER	Replication set enabled Format >>>9
_ReplSrv-TransDatabaseRole	25	CHARACTER	Role the database takes after transition. Valid values are: • Normal • Reverse Format X(32)
_ReplSrv-TransitionToAgents	26	CHARACTER	Transition to agents list Format X(100)
_ReplSrv-RestartAfterTrans	27	BOOLEAN	Restart after transition Format yes/no
_ReplSrv-AutoBeginAI	28	BOOLEAN	Automatically begin after-imaging during transition Format yes/no
_ReplSrv-AutoAddAI	29	BOOLEAN	Automatically add after-imaging files during transition Format yes/no

Field name	Field number	Data type	Description
_ReplSrv-AIFile	30	CHARACTER	AI structure filename for automatically adding after-imaging Format X(256)
_ReplSrv-BackupMethod	31	CHARACTER	Backup method during transition Format X(32)
_ReplSrv-ConfiguredAgents	32	CHARACTER	List of configured agents Format X(68)
_ReplSrv-ReplKeepAlive	33	INT64	Replication keep alive property (in seconds) Format ->>>>9
_ReplSrv-SchemaLockAction	34	CHARACTER	Method server uses to perform schema locking. Valid values include: • Force • Wait Format X(32)
_ReplSrv-AgentShutdownAction	35	CHARACTER	Action replication agent takes if connection to server is lost. Valid values include: • Normal • Recovery • Wait Format X(32)
_ReplSrv-ReservedChar	36	CHARACTER	Reserved
_ReplSrv-ReservedInt	37	INTEGER	Reserved
_ReplSrv-TransEndTrigger	38	INTEGER	Indicates if a transition end trigger is to execute on transition. Valid values include: • **0** — Execute a transition end trigger • **1** — Do not execute a transition end trigger Format >>>9

Field name	Field number	Data type	Description
_ReplSrv-DRTransitionToAgents	39	CHARACTER	Recovery transition to agents list. Format X(100)
_ReplSrv-TransToDBName	40	CHARACTER	Remote Database Name Fully qualified database name to which the OpenEdge Replication agent is connected. Target database for the primary transition to agent in the event of failover transition. Format X(128)
_ReplSrv-TransToDBHost	41	INTEGER	Remote Host Name The hostname of the primary transition to target database identified by the _replserv-TransToDBName field. Format X(128) **Note:** The hostname is the IP address instead of the hostname specified in repl properties or the startup (-H) on the target.
_ReplSrv-TransToDBPort	42	INTEGER	Connected to Agent on Port Number The port number the primary transition to target database identified by the _replserv-TransToDBName field as specified by the -S startup parameter. Format >>>>>>>>>9
_ReplSrv-SSLEnabled	43	BOOLEAN	Specifies whether or not SSL use is enabled for the replication database. • 0—SSL use is disabled. This is the default. • 1—SSL use is enabled. The property affects all replication processes that may run against the replication database.

Field name	Field number	Data type	Description
_ReplSrv-SSLNoHostVerify	44	BOOLEAN	Specifies whether or not the client verifies the SSL server's machine name. During the SSL handshake process, the SSL client may want to verify that the machine name where the SSL server runs matches what is in the server's certificate. • 0—Enables verification. • 1—Disables verification. In OpenEdge Replication, it is not guaranteed that the SSL server's host machine name will match what is stored in the server's certificate. For this reason, it is a best practice to disable this property.
_ReplSrv-SSLCertStorePath	45	CHARACTER	For an SSL client, the absolute path of the directory containing the SSL certificate for OpenEdge Replication SSL.
_ReplSrv-SSLCipherSuites	46	CHARACTER	A comma-separated list of the SSL cipher suites that can be used for any OpenEdge Replication SSL network connections for the replication database. Cipher suites are: • AES128-SHA256 • AES256-SHA256 • DHE-RSA-AES256-GCM-SHA384
_ReplSrv-SSLProtocols	47	CHARACTER	A comma-separated list of the SSL protocol versions that can be used for any OpenEdge Replication SSL network connections for the replication database. Supported SSL protocols are: • TLSv1.1 • TLSv1.2 • TLSv1.3

Note

- If the replication server is not running, the _replsrv VST returns no records.