



REPSTANCE

Prepare to be Amazed

Repstance User Guide For Oracle

Version 3.0.1007

Collabcloud Limited

47 St Pauls Road, Staines-upon-Thames, Surrey, TW18 3HQ, England

All rights reserved. This product and document are protected by copyright and distributed under licenses restricting its use, copying, distribution and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Collabcloud Limited and its licensors, if any.

Third-party software, including font technology in this product, is protected by copyright and licensed from Collabcloud's Suppliers.

RESTRICTED RIGHTS LEGEND:

Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19. The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS:

The Repstance Name and Logo are trademarks or registered trademarks of Collabcloud Limited in the United Kingdom and may be protected as trademarks in other countries.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN, THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. COLLABCLOUD LIMITED MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAMS(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

TABLE OF CONTENTS

ABOUT REPSTANCE	4
What it is	4
How it works	4
GETTING STARTED WITH REPSTANCE IN AWS	6
Configure repcli to Remotely Connect	7
Configure Repstance Server to Accept Remote Connection	7
INTERACTING WITH A REPSTANCE SERVER	9
REST API	9
Command Line Interface (CLI)	9
Fully Interactive Mode	9
Command Mode	10
SUPPORTED DATABASE REQUIREMENTS	12
HOW TO USE REPSTANCE	13
HOW TO REMOVE REPSTANCE	16
Remove Apply Process	16
Remove Capture Process	16
Remove Target Database Objects	16
Remove Source Database Objects	17
COMMANDS TO BE USED	18
Prepare Source and Target Databases	18
Remove Repstance Database Objects	20
Validate Source and Target Database	21
Prepare Capture Process	23
Capture Objects Specification	26
Initial Load	27
Objects Mapping and Possible Transforms	29
Transformation Rules and Triggering Order	33
Alter Capture Process	34
Show Capture Process	36
Validate Capture Process	37
Remove Capture Process	38
Prepare Apply Process	40
Alter Apply Process	42
Show Apply Process	44
Validate Apply Process	46
Reset Apply Process	47
Remove Apply Process	48

Control Repstance Processes	49
Run any Capture or Apply Processes	50
Stop any Capture or Apply Processes	51
Status of any Capture or Apply Processes	52
REPSTANCE SERVER MAINTENANCE	56
How to Stop Repstance Services	56
How to Start Repstance Services	56
How to Check Repstance Service Status	56
Backup Repstance Files	56
Housekeeping	57
GLOSSARY	57

ABOUT REPSTANCE

What it is

Repstance is an entirely new data propagation product, which supports Oracle 10g-19c and any SQL Server that supports CDC. It includes the functionality to support DML and DDL operations, both of which can be automatically included in the replication stream and if desired use of sophisticated transformation abilities. Fast initial data loading and the ability to restart from any desired timestamp. Repstance gives a 360 view into the intricacies of the database replication process with the highest granular health check, along with reporting of - replication lags (with cause/consequences analysis), detailed DB replication statistics (including, but not limited to transactions count, objects modification count, DDLs count, average transaction rate calculation, etc.).

How it works

The Service is delivered via an “Amazon Machine Image” which runs within its own separate Amazon “Instance”, this is a Linux based instance and runs as a “Daemon” within this environment. Repstance has to have access to both Source and Target Databases, in addition it needs a set of minimal Database Objects to function. This functionality is added during the “Prepare Database” stage (see the [“Prepare Source and Target Databases”](#) chapter).

Repstance uses the running Daemon process to “Multi-task” – that is to say by using “threads” within the Daemon it is capable of running Multiple Capture and Apply Processes concurrently.

Both the Processes, while they depend on each other to supply or insert the necessary data they nonetheless run as independent threads.

The Capture Process is the means by which the data to be transferred is extracted from the Source Database. It puts this information into locally stored Trail Files which the Apply Processes can use. The data from a single Capture Process can therefore be used by many Apply Processes and as a result this data can therefore be propagated into multiple Target Databases.

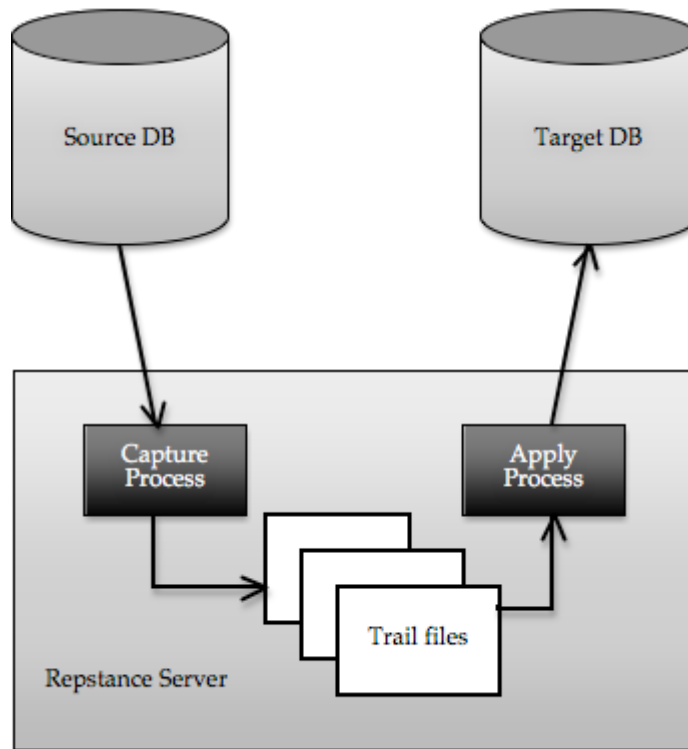
The data extracted by the Capture Process is written to these Trail Files in the same sequential order as the transactions occurred in the Source Database, which in turn allows the Apply Process to insert this data into the Target Database in the same order they were executed in the Source Database. This means that both Source and Target Databases will be synchronized that is to say in a “Consistent State”.

The Capture Process does not need to have a running or configured Apply Process as it will quite simply continue extracting data using the criteria supplied, conversely the Apply Process only needs valid Trail Files from a Capture Process to consume.

Both Capture and Apply Processes insert Checkpoints in the form of a LSN/SCN – this in turn means that there is the possibility for Repstance to be restarted from any specified point, in the circumstances where any unexpected event occurs that may

lead to data loss. This provides a robust form of data security and ensures that the data is, again, always “Consistent”.

The diagram below demonstrates the flow of data from Source to Target Databases.



GETTING STARTED WITH REPSTANCE IN AWS

Repstance Server is delivered as AMI instance and is accessible from the AWS Marketplace, using the “1 Click “ download process here:

<https://aws.amazon.com/marketplace/pp/B07PBF1CHK>

It requires you as a user to have your own AWS account and be subscribed to the product, and AWS instance of sufficient size and processing power to manage your database requirements (see vendor recommended instance size).

During AWS installation process you must ensure that TCP port “22” as a minimum, and in the case that you want remote access port “8797” is also enabled.

To connect to the operating system, use ssh protocol and the only user name that can be used is “**ec2-user**” along with the key that you provided during the “launch” of the installation process. See the following link for help:

https://docs.aws.amazon.com/en_us/AWSEC2/latest/UserGuide/AccessingInstancesLinux.html

Once it is installed in your environment it is ready to run. The product can be managed using the **repcli** interface from within this environment.

Now **repcli** can be invoked directly, and the user can start to configure Processes without the need for any further configuration.

Run **repcli** directly on the server:

```
$ repcli
$ repcli>help
```

Commands:

```
alter          This command is used to change both the
Capture and Apply process parameters
```

```
clear          Clear The Screen
```

```
exit           Exit The Program
```

```
help           Display Help
```

etc.

It can also be accessed remotely using the **repcli** client versions which are found in the following directories, once the product is installed:

- /opt/repstance/cli/macos – For MacOS Operating System
- /opt/repstance/cli/win – For Windows
- /opt/repstance/cli/linux – For Linux

Or they can be downloaded from here:

- <https://repstance.com/download/repcli/macos/repcli.zip>
- <https://repstance.com/download/repcli/win/repcli.zip>
- <https://repstance.com/download/repcli/linux/repcli.zip>

These remote clients are only able to connect using the HTTPS protocol and is supplied of with the following default configuration:

- HTTPS port is 8797
- For HTTPS initial use a Self signed certificate is provided
- HTTPS authorisation “Token” in the form of the Repstance AWS Instance ID

An alternative to either of the above communication methods is to call the REST API detailed in this document.

Configure repcli to Remotely Connect

In order to use **repcli** utility communicating with remote instance it has to be configured using the **repcli** command “configure”.

The following example demonstrates the command syntax:

```
$ repcli configure url=https://hostname:8797/ token=aws-  
instance-id verifyhttps=0
```

Note – The **verifyhttps** parameter should be set to **0** to avoid the Self-signed certificate error.

Configure Repstance Server to Accept Remote Connection

As the product is supplied with a Self-Signed Certificate it is highly recommended that the user replaces this with a security certificate of the required level themselves. The certificate and key needs to be supplied to the Repstance Server and the appropriate entries in the `/opt/repstance/conf/repstance.conf` file need to be edited to reflect these changes:

- **key** – Path to the key
- **crt** – Path to the certificate

Using this method the default HTTPS port can also be changed by modifying the **port** parameter.

The default token used to access the server can also be changed by modifying the **token** parameter in the same file.

If the user wishes to prevent any external access the **port** parameter should be set to value **0**.

INTERACTING WITH A REPSTANCE SERVER

There are currently only two possible ways to engage with a running Repstance Server Instance, they either via the Rest API or via the Command Line Interface (CLI).

REST API

This is the detailed command structure for communicating with Repstance Services, it sends an HTTP/HTTPS request in the “JSON Format” to the services and adheres to the established protocols for the JSON format. List of the available Repstance commands and their specification is provided in chapter [“COMMANDS TO BE USED”](#).

Command Line Interface (CLI)

The Repstance Command Line Interface utility (**repcli**) is the interface with the Repstance environment and is used to configure and manage the environment as well as providing a reporting mechanism, and once connected to a Repstance server can be used with a predefined set of “Command” parameters, which are parsed by Repstance to ensure validity, and their specific use is detailed below.

It is delivered together with a Repstance server instance and available in the “**bin**” folder under Repstance home directory. By default it establishes network layer connection to the Repstance Server using IP address 127.0.0.1 (localhost).

The Repstance Server can be configured to accept an HTTPS connection by external means. The **repcli** utility can be configured to communicate with remote Repstance Servers also using the HTTPS protocol see [configure repcli for Remote Access](#).

The **repcli** utility can operate in two modes either “Command Mode” or “Fully Interactive Mode”. The two different ways of using the **repcli** interface have been developed to provide alternative methods of controlling the product.

If **repcli** is executed without any parameters it starts functioning in “Fully Interactive Mode”, by default, which means that it will wait for an input command, until user issues the **Exit** command at which point it returns to the console.

The “Fully Interactive Mode” is useful for a user to construct and/or configure Repstance, as it provides, at each stage of the process construction, a list of the possible commands and their applicable parameters, see the “Visuals” below. “Command mode” is more likely to be used when there is a need to “Embed” Repstance functions into a deployment Script for Automation processes within the users own environment.

Fully Interactive Mode

This is the utility's default behavior. As the user starts to provide a command for any Process and then "presses" the "Tab" key, **repcli** utility will supply the user with the expected possible parameters.

The following examples show the **repcli** interface displaying the list of all possible commands:

```
repcli>
status exit help prepare remove alter reset clear
validate show run stop
```

This example shows all commands starting with "s":

```
repcli>s
status show stop
```

And here it shows all the possible parameters for the "Alter" command:

```
repcli>alter process=capture id=1 dbtype=oracle connectiontype=tns

tnsname=          user=          password=         dbname=          name=
autostart=       debuglevel=   ddlinclude=     ddlexclude=     dmlinclude=
dmlexclude=     loadinclude=  loadexclude=    map=
```

The **repcli** has built in help facility, which can be invoked to provide detailed "Help" for each of the commands, and is shown below:

```
repcli>alter help
```

This command is used to change the Capture or Apply process parameters.

The Syntax is:

```
alter process=capture|apply id=processID [parameter=value, ...]
```

Note: The 'id' parameter can not be changed.

The 'name' parameter for a Capture process as well as the 'capturename' parameter for an Apply process can be changed only in the case that it has not been provided before - eg. it must be "Unique".

The process has to exist otherwise the command will fail.

To see available possible parameters, start to use the command.

Command Mode

This is the basic interface format, when used in this mode it will simply supply the result of the process defined and quit to the console:

```
$repcli status process=capture id=7
```

```
Capture Process 7, Name: oracap1, Connection: 192.168.56.18:1521/RAC.localdomain  
RUN (In Process), Waiting for Transactions  
  Last Source DB Change   : 2020/02/21 09:43:43.000 UTC (0x00000000000732FC8841)  
  Last Captured DB Change : 2020/02/21 09:43:43.000 UTC (0x00000000000732FC8841)  
Total since 2020/02/21 08:41:11.000 UTC (0x00000000000732FC7CDD):  
  Processing the Transactions : 00:00:11.810  
  Waiting for Transactions   : 00:01:56.046  
  Average Speed (ops)        : 4537.972  
  Transactions : 3 (DDL: 0, Delete: 30022, Insert: 23234, Update: 341)
```

SUPPORTED DATABASE REQUIREMENTS

Currently the product supports the following:

Oracle – versions 10g through 19c.

All of the above versions must be configured to run in ARCHIVELOG mode prior to use.

It is recommended practice to ensure that the archive logs are not “cleaned-up” until they are processed by the Repstance server.

Both “EC2” and “RDS” usage for Source and/or Target Databases are supported.

HOW TO USE REPSTANCE

Before running the replication both Source and Target Databases need to be prepared. For the Source Database it is necessary to enable “supplemental logging” and install Repstance’s objects. For the Target Database only the Repstance’s objects have to be installed. Repstance has built in functionality that allows it to prepare databases by simply running the CLI or REST API command (see the [“Prepare Source and Target Databases”](#) chapter).

Once it is done we recommend to validate both databases by running the Validate command (see the [“Validate Source and Target Database”](#) chapter) to make sure that the necessary functionality is installed.

The next step is to configure the replication processes. This is done in two parts, the first is to configure the Capture Process (see the [“Prepare Capture Process”](#) chapter), which is extracting the data from the specified Source Database, and the second is to configure the corresponding Apply Process (see the [“Prepare Apply Process”](#) chapter), which will insert the “captured” data into the specified Target Database.

The data to be replicated will be defined during the Capture Process configuration. Validation of a correctly constructed Capture or Apply Processes can be checked using the Validate command (see [“Validate Capture Process”](#) and [“Validate Apply Process”](#) chapters).

The first time the Capture Process is run, is the first point and the only point at which data will start to be captured by this process.

Assuming that the Capture Process has been properly configured and run, only at this point will it start to produce Trail Files, which can be used by an Apply Process.

In order for an Apply Process to succeed, of necessity the Trail Files of the corresponding Capture Process must exist.

In order to ensure that an Apply Process will start from the specified SCN, this SCN must exist in the Trail Files and the **Reset** command (see the [“Reset Apply Process”](#) chapter) can be used to point this Apply Process at the required SCN. If the Apply Process has not been reset, it will start processing data from the first available transaction found in the Trail File.

Once the Apply Process has been configured to use and the proper SCN (if needed) has been specified, the Apply Process can be run (see the [“Run any Capture or Apply Processes”](#) chapter).

At the point when both the Capture and Apply Processes are running the “Status” command (see the [“Status of any Capture or Apply Processes”](#) chapter) can be used to monitor them.

In the event it is necessary to “stop” any currently running processes use the “Stop” command (see the [“Stop any Capture or Apply Processes”](#) chapter).

The following example demonstrates how to setup simple DML and DDL replications for all objects located in the **scott** schema:

```
repcli>prepare database=source dbtype=oracle
connectiontype=ezconnect server=usoradb.ckppywxf9a8u.us-east-
1.rds.amazonaws.com port=1521 servicename=usoradb user=admin
password=repstance
[Completed]: Database has been prepared
```

```
repcli>prepare database=target dbtype=oracle
connectiontype=ezconnect server=ora19.ckppywxf9a8u.us-east-
1.rds.amazonaws.com port=1521 servicename=orcl user=admin
password=repstance
[Completed]: Database has been prepared
```

```
repcli>validate database=source dbtype=oracle
connectiontype=ezconnect server=usoradb.ckppywxf9a8u.us-east-
1.rds.amazonaws.com port=1521 servicename=usoradb user=admin
password=repstance
[Valid]: Database is valid
```

```
repcli>validate database=target dbtype=oracle
connectiontype=ezconnect server=ora19.ckppywxf9a8u.us-east-
1.rds.amazonaws.com port=1521 servicename=orcl user=admin
password=repstance
[Valid]: Database is valid
```

```
repcli>prepare process=capture id=1 dbtype=oracle
connectiontype=ezconnect server=usoradb.ckppywxf9a8u.us-east-
1.rds.amazonaws.com port=1521 servicename=usoradb user=admin
password=repstance name=rep1 dmlinclude=scott. % ddlinclude=scott. %
[Completed]: Capture 1 has been added
```

```
repcli>validate process=capture id=1
[Valid]: Capture process 1 is valid
```

```
repcli>run process=capture id=1
[In Process]: Command [RUN] has been sent. Use [status] command to
monitor the Capture 1 process
```

```
repcli>prepare process=apply id=1 dbtype=oracle
connectiontype=ezconnect server=ora19.ckppywxf9a8u.us-east-
1.rds.amazonaws.com port=1521 servicename=orcl user=admin
password=repstance capturename=rep1
[Completed]: Apply 1 has been added
```

```
repcli>validate process=apply id=1
[Valid]: Apply process 1 is valid
```

```
repcli>run process=apply id=1
[In Process]: Command [RUN] has been sent. Use [status] command to
monitor the Apply 1 process
```

```
repcli>status
```

Capture Process 1, Name: repl, Connection: usoradb.ckppywxf9a8u.us-east-1.rds.amazonaws.com:1521/usoradb
RUN (In Process), Waiting for Transactions
Last Source DB Change : 2020/02/21 17:40:52.000 UTC (0x0000000000000035963F)
Last Captured DB Change : 2020/02/21 17:40:52.000 UTC (0x0000000000000035963F)
Total since 2020/02/21 17:37:25.000 UTC (0x0000000000000035900C):
Processing the Transactions : 00:00:06.469
Waiting for Transactions : 00:03:39.452
Average Speed (ops) : 1545.916
Transactions : 2 (DDL: 1, Delete: 0, Insert: 10000, Update: 0)

Apply Process 1, Transactions Provider (Capture): repl, Connection: oral9.ckppywxf9a8u.us-east-1.rds.amazonaws.com:1521/orcl
RUN (In Process), Waiting for Transactions
Last Source DB Change : 2020/02/21 17:40:52.000 UTC (0x0000000000000035963F)
Last Applied DB Change : 2020/02/21 17:40:57.000 UTC (0x000000000000003595E9)
Lag : 00:00:00.000
Total since 2020/02/21 17:39:31.426 UTC (0x00000000000000000001):
Processing the Transactions : 00:00:00.075
Waiting for Transactions : 00:01:44.131
Average Speed (ops) : 132496.653
Transactions : 2 (DDL: 1, Delete: 0, Insert: 10000, Update: 0)

HOW TO REMOVE REPSTANCE

The steps involved in removing Repstance are detailed below.

Remove Apply Process

To remove an Apply Process use the “Remove” command (see the [“Remove Apply Process”](#) chapter).

After using this command it will remove **all** information about the Apply Process, namely configuration files, and the checkpoint details, together with historical information. If the Target Database **is** reachable the “Remove” command will clean-up all the checkpoint information on the Target Database.

As the **Remove** command can only accept a single “Apply parameter” this action can only be used for one Apply Process at a time.

Remove Capture Process

To remove a Capture Process use the “Remove” command (see the [“Remove Capture Process”](#) chapter).

After using this command it will remove **all** information about the Capture Process, namely configuration files, checkpoint details, and the Trail Files.

As the “Remove” command can only accept a single “Capture parameter” this action can only be used for one Capture Process at a time.

Note – If there are still any Apply Processes using Trail Files from this Capture Process, they will fail.

The steps involved in removing Repstance Server Database Objects are detailed below.

Remove Target Database Objects

To remove Target Database Objects use the “Remove” command (see the [“Remove Repstance Database Objects”](#) chapter).

After using this command it will remove all the Data Objects that were required by the Apply Process. If there are still Apply Processes inserting data, using these objects they will immediately fail.

Note – **It is strongly recommended that before running this command** you run the “Status” command and make a record the SCN in the “Last Applied DB Change”

section for each Apply Process running on this database, in this way it will be possible to reverse this action for this database at a later point in time, assuming that the correct Trail Files still exist and can be accessed.

Remove Source Database Objects

To remove Source Database Objects use the “Remove” command (see the “[Remove Repstance Database Objects](#)” chapter).

Note – When using this command it MUST be understood that there will be no way to “Recover” from its effects.

After using this command it will remove all the Data Objects that were required by the Capture Processes. If there are still Capture Processes extracting data, using these objects they will immediately fail. At this point Repstance will no longer be able to record or extract data from this database.

COMMANDS TO BE USED

Prepare Source and Target Databases

The “Prepare” database command is used to configure a database for replication. The command must be executed on both Source and Target Databases before running any replication processes. In the Source Database it enables the necessary level of supplemental logging and creates the database objects necessary for any Capture Process. In the Target Database it creates the necessary database objects for any Apply Process.

Note – In order to use the “Prepare” database command the database user must have sufficient privileges.

There are two possible connection types - EZCONNECT and TNS. Depending on which one is chosen, there will be a different set of possible parameters.

REST API:

- **Endpoint:** https://repstance_url/configure/database
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** token
- **Body:**

```
{
  "command": "prepare",
  "parameters": [
    ["database", "source|target"],
    ["dbtype", "oracle"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "database_name"],
    ["user", "db_user_name"],
    ["password", "db_user_password"],
    ["tablespace", "user_tablespace"]
  ]
}
```

Server response:

- **HTTP Status** - the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs

- **Body:**

```
{
  "Status": "{Completed/Failed}",
  "Message": "The command execution details"
}
```

CLI Syntax:

```
repcli prepare database=source|target \
  dbtype=oracle \
  connectiontype=tns|ezconnect \
  tnsname=tns_alias \
  server=databaseHost port=databasePort \
  servicename=service_or_SID
  user=username password=password \
  dbname=databaseName \
  tablespace=user_tablespace \
```

The input parameters are:

- **database** – database role. The possible values are:
 - *source* – to insert Capture Process objects
 - *target* – to insert Apply Process objects
- **dbtype** – type of RDBMS. The possible values are:
 - *oracle*
 - *mssql* – see SQL Server documentation for the details
- **connectiontype** – Specifies the Oracle connection method. The possible values are:
 - **tns** – Local Naming Method to be used
 - **ezconnect** – EZCONNECT to be used
- **tnsname** – Name of the TNS alias. Valid only if **connectiontype=tns**
- **server** – host name or IP address of the database server. Valid only if **connectiontype=ezconnect**
- **port** – database port number. Valid only if **connectiontype=ezconnect**
- **servicename** – database service name or SID. Valid only if **connectiontype=ezconnect**
- **dbname** – name of either container or pluggable database. Valid only if **database =source** and for any Oracle versions 12c-19c but excluding RDS instances
- **user** – database user name
- **password** – database user password
- **tablespace** – name of the tablespace that the Repstance’s objects are to be installed in. The default is the USERS tablespace if no alternative has been specified

This action differs slightly for the Source Database as it not only inserts the database objects, but in addition sets supplemental logging to the necessary level, whereas in the Target Database it only inserts the necessary database objects.

Remove Repstance Database Objects

The “Remove” database objects command is used to remove Repstance database replication objects.

The successful result of using this command will be:

- For the Source Database, it will remove the previously inserted Database Objects;
- For the Target Database, it will only remove the Repstance Database Objects that were created by the “Prepare Command”.

After successful completion of the command any remaining running processes will fail.

REST API:

- **Endpoint:** https://repstance_url/configure/database
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** token
- **Body:**

```
{
  "command": "remove",
  "parameters": [
    ["database", "source|target"],
    ["dbtype", "oracle"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "database_name"],
    ["user", "db_user_name"],
    ["password", "db_user_password"]
  ]
}
```

Server response:

- **HTTP Status** - The status of command. The possible codes are:
 - **200** - if no error occurs
 - **422** - if error occurs
- **Body:**

```
{
  "Status": "{Completed/Failed}",
  "Message": "The command execution details"
}
```

CLI Syntax:

```
repcli remove database=source|target \  
  connectiontype=tns|ezconnect \  
  tnsname=tns_alias \  
  server=databaseHost port=databasePort \  
  servicename=service_or_SID \  
  user=username password=password \  
  dbname=databaseName
```

The input parameters are:

- **database** – database role. The possible values are:
 - *source* – To remove Capture Process objects
 - *target* – To remove Apply Process objects
- **dbtype** – type of RDBMS. The possible values are:
 - *oracle*
 - *mssql* – see SQL Server documentation for the details
- **connectiontype** – Specifies the Oracle connection method. The possible values are:
 - **tns** – Local Naming Method to be used
 - **ezconnect** – EZCONNECT to be used
- **tnsname** – Name of the TNS alias. Valid only if **connectiontype=tns**
- **server** – host name or IP address of the database server. Valid only if **connectiontype=ezconnect**
- **port** – database port number. Valid only if **connectiontype=ezconnect**
- **servicename** – database service name or SID. Valid only if **connectiontype=ezconnect**
- **dbname** – name of either container or pluggable database. Valid only if **database=source** and for any Oracle versions 12c-19c but excluding RDS instances
- **user** – database user name
- **password** – database user password

Validate Source and Target Database

The “Validate” database objects command is used to validate that the correctly initialised database objects exist.

It is considered to be “Good Practice” to run this command after either, initial configuration has been completed or if changes have been made to either a Capture or Apply Process.

REST API:

- **Endpoint:** https://repstance_url/configure/database
- **Method:** POST
- **Header:**

- Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "validate",
  "parameters": [
    ["database", "source|target"],
    ["dbtype", "oracle"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "database_name"],
    ["user", "db_user_name"],
    ["password", "db_user_password"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Valid/Invalid}",
  "Message": "The command execution details"
}
```

CLI Syntax:

```
repcli validate database=source|target \
dbtype=oracle \
connectiontype=tns|ezconnect \
tnsname=tns_alias \
server=databaseHost port=databasePort \
servicename=service_or_SID
user=username password=password \
dbname=databaseName
```

The input parameters are:

- **database** – database role. The possible values are:
 - *source* – To remove Capture Process objects
 - *target* – To remove Apply Process objects
- **dbtype** – type of RDBMS. The possible values are:
 - *oracle*
 - *mssql* – see SQL Server documentation for the details
- **connectiontype** – Specifies the Oracle connection method. The possible values are:
 - **tns** – Local Naming Method to be used

- **ezconnect** – EZCONNECT to be used
- **tnsname** – Name of the TNS alias. Valid only if **connectiontype=tns**
- **server** – host name or IP address of the database server. Valid only if **connectiontype=ezconnect**
- **port** – database port number. Valid only if **connectiontype=ezconnect**
- **servicename** – database service name or SID. Valid only if **connectiontype=ezconnect**
- **dbname** – name of either container or pluggable database. Valid only if **database=source** and for any Oracle versions 12c-19c but excluding RDS instances
- **user** – database user name
- **password** – database user password

Prepare Capture Process

The “Prepare” Capture Process command is used to add new Capture Process.

The “Prepare” command cannot be used to validate the database connection, or the presence of the necessary data of captured objects. In order to validate that the process is configured properly and it is able to run on the Source Database, the “[Validate](#)” command must be used.

The “Prepare” command does not enable the database to start capturing data on the configured objects, until the specified Capture Process runs (see the “[Run any Capture or Apply Processes](#)” chapter for the details).

Parameter **id** is mandatory. All others parameters are optional. All parameters except **id** can be changed using “[Alter](#)” Capture Process command. The **name** parameter can be altered only in the case if it has not been provided earlier. If there is another Capture Process which has the same **id** or the same **name** then the “Prepare” command will fail.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** token
- **Body:**

```
{
  "command": "prepare",
  "parameters": [
    ["process", "capture"],
    ["id", "captureID"],
    ["dbtype", "oracle"],
    ["name", "captureName"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
```



```

    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "databaseName"],
    ["autostart", "{0/1}"],
    ["debuglevel", "{0-15}"],
    ["maxtrailsize", "{10000000-9999999999}"],
    ["ddlinclude", "objectsMask"],
    ["ddlexclude", "objectsMask"],
    ["dmlinclude", "objectsMask"],
    ["dmlexclude", "objectsMask"],
    ["loadinclude", "objectsMask"],
    ["loadexclude", "objectsMask"],
    ["map=mapID", "mappingClause"]
  ]
}

```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```

{
  "Status": "{Failed/Completed}",
  "Message": "The command execution details"
}

```

CLI Syntax:

```

repcli prepare process=capture id=captureID \
  dbtype=oracle name=captureName \
  connectiontype=tns|ezconnect \
  tnsname=tns_alias \
  server=databaseHost port=databasePort \
  servicename=service_or_SID
  user=username password=password \
  dbname=databaseName autostart={0/1} \
  debuglevel={0-15} maxtrailsize={10000000-9999999999} \
  ddlinclude=objectsMask ddlexclude=objectsMask \
  dmlinclude=objectsMask          dmlexclude=objectsMask          \
  loadinclude=objectsMask loadexclude=objectsMask \
  map=mapID,mappingClause

```

The input parameters are:

- **process** – *capture*
- **id** – Capture Process id
- **dbtype** – type of RDBMS. The possible values are:
 - *oracle*
 - *mssql* – see SQL Server documentation for the details
- **connectiontype** – Specifies the Oracle connection method. The possible values are:

- **tns** – Local Naming Method to be used
- **ezconnect** – EZCONNECT to be used
- **tnsname** – Name of the TNS alias. Valid **only** if **connectiontype=tns**
- **server** – host name or IP address of the database server. Valid **only** if **connectiontype=ezconnect**
- **port** – database port number. Valid **only** if **connectiontype=ezconnect**
- **servicename** – database service name or SID. Valid **only** if **connectiontype=ezconnect**
- **dbname** – name of either container or pluggable database. Valid **only** if **database=source** and for any Oracle versions 12c-19c but excluding RDS instances
- **name** – name of the Capture Process
- **user** – database user name
- **password** – database user password
- **autostart** – determine if the process must be run automatically. The possible values are:
 - **0** – do not run the process automatically
 - **1** – run the process automatically
- **debuglevel** – the level of debugging. Possible values are 0-15
- **ddlinclude** – mask of the DDL objects to be captured. See “[Capture Objects Specification](#)” chapter for the details
- **ddlexclude** – mask of the DDL objects to be skipped by the Capture Process. See “[Capture Objects Specification](#)” chapter for the details
- **dmlinclude** – mask of the DML objects to be captured. See “[Capture Objects Specification](#)” chapter for the details
- **dmlexclude** – mask of the DML objects to be skipped by the Capture Process. See “[Capture Objects Specification](#)” chapter for the details
- **loadinclude** – Mask of the objects to be included into the Initial Load (see chapter [Initial Load](#) for the details)
- **loadexclude** – Mask of the objects to be skipped during Initial Load (see chapter [Initial Load](#) for the details)
- **map** – the set of the parameters to determine objects’ transformation. See “[Prepare Capture Process Objects Mapping](#)” chapter for the details.

The following example demonstrates how to configure a new Capture Process having id=2. The process is configured to capture objects in the **ora19** database, which is running on the **ora19.ckppywxf9a8u.us-east-1.rds.amazonaws.com:1521** server. The process has now been configured to run, once the Repstance service has been started. It captures DML and DDL operations for the tables located in the **scott** schema.

```
{
  "command": "prepare",
  "parameters": [
    ["process", "capture"],
    ["id", "2"],
    ["dbtype", "oracle"],
    ["connectiontype", "ezconnect"],
    ["server", "ora19.ckppywxf9a8u.us-east-1.rds.amazonaws.com"],
    ["port", "1433"],
```

```

    ["servicename", "ORCL"],
    ["name", "cap2"],
    ["user", "dbusr"],
    ["password", "dbpwd1"],
    ["autostart", "1"],
    ["ddlinclude", "scott.%"],
    ["dmlinclude", "scott.%"]
  ]
}

```

Capture Objects Specification

Repstance allows flexibility to determine a set of tables, which are to be captured. The tables for DML and DDL operations are specified by the different parameters. The following parameters are used for DML replication:

- `dmlinclude` – list of the tables to be included into the DML replication
- `dmlexclude` – list of the tables to be excluded from the DML replication

The following parameters are used for DDL replication:

- `ddlinclude` – list of the tables to be included into the DDL replication
- `ddlexclude` – list of the tables to be excluded from the DDL replication

Note – These criteria are mutually exclusive. The table will only be replicated in the case that the name matches the “include” criteria and does not match the “exclude” criteria.

The table is specified in *schema_name.table_name* format. In order to specify the list of the tables, each table must be separated by comma. The % symbol may be used in order to match any number of characters. It may also be used in both *schema_name* and *table_name* parts.

The following examples show various ways of defining tables:

- **Customers and employees** tables, which are in the **scott** schema:

```
scott.customers,scott.employees
```

- Any tables in the **scott** schema:

```
scott.%
```

- Any tables having name started from **rep** or **tmp**:

```
%.rep%,%.tmp%
```

- Any table in report schema ends with #:

```
report.##
```

The following examples show how to use the **DML/DDL** parameters:

- To capture DML changes for all tables in **scott** schema except where tables start with **rep**:

```
dmlinclude=scott.%
dmlexclude=scott.rep%
```

- To capture DML changes for all tables in the **cap1**, **cap2** and **cap3** schemas except the **report1** table from the **cap1** schema:

```
dmlinclude=cap1.%,cap2.%,cap3.%
dmlexclude=cap1.report1
```

- To capture DML and DDL operations for **any** tables in the **scott** schema:

```
dmlinclude=scott.%
ddlinclude=scott.%
```

- To capture only DDL operations for objects in the **report** schema:

```
ddlinclude=report.%
```

These parameters influence the newly created tables as well. The description below shows the capture behaviour of the tables created:

Tables that match DML criteria	Tables that match DDL criteria	Capture Behavior
true	true	This create statement is captured and any further DML statements will be captured as well.
true	false	This create statement is not captured but any DML statements will be captured.
false	true	This create statement is captured but DML statements are not captured. However any further DDL statements are captured.
false	false	This create statement is not captured and neither DML or DDL statements will be captured.

Initial Load

Repstance has built-in functionality to load objects' data, and start the replication process from the timestamp at which the data has been created. Initial Loading is performed by the Capture Process at the "Run" stage (see chapter [Run any Capture or Apply Processes](#) for the details).

This functionality is used when it is necessary to synchronise data between the Source and Target Databases before starting the replication process. The Initial Load parameters can be configured to either clean up, or preserve data in the Target Database before loading from the Source Database. Two further options are to create table in the Target Database if it doesn't exist and recreate this table in the Target Database if it does exist.

The following parameters are used to configure the objects for Initial Loading:

- `loadinclude` – List of the tables to be included into the Initial Load
- `loadexclude` – List of the tables to be excluded from the Initial Load

Note – These parameters MUST be used in conjunction with the `dmlinclude/dmlexclude` object definitions as well, otherwise the Capture Process will skip them.

The format to be specified takes the following form:

```
schema_name1.table_name1:[loadOption1],  
schema_name2.table_name2:[loadOption2], ...
```

In order to define the list of the tables to be part of this process, and where there are more than one tables mask to be defined each of these must be separated by comma. The `%` symbol may be used in order to match any number of characters. It may also be used in both `schema_name` and `table_name` parts.

The *loadOption* is an optional parameter, and it used to perform the action detailed below on the data before insertion into the Target Database, and can have the following values:

- A – Preserve the data which is the **DEFAULT** value
- T – Use **truncate** statement to clean up the data
- D – Use **delete** statement to clean up the data
- C – Create table if it doesn't exist
- R – Recreate table if it already exist

Note – Once the Initial Load has been completed by Capture Process the values of `loadinclude` and `loadexclude` parameters will be removed automatically. The details of Initial Load criteria can be found in the Capture log files.

The following examples show various ways of defining the Initial Loading parameters:

- To reload data for the **emp** and **emp_audit_trail** tables, which are in the **scott** schema and **truncate** these tables before insertion:

```
loadinclude=scott.emp:T,scott.emp_audit_trail:T
```

- Any data from tables starting with the **rep_** prefix in the **scott** schema which need to be **deleted** before insertion:

```
loadinclude=scott.rep_%:D
```

- To reload all tables started with **emp** prefix except the **employee_archive** tables:

```
loadinclude=scott.emp%:T  
loadexclude=scott.employee_archive
```

- To reload **customer** table and create it if it doesn't exist in the Target Database:

```
loadinclude=scott.customer:C
```

Objects Mapping and Possible Transforms

Transforms are primarily used where there is a need to reformat any statements. The transformation can only be configured for a Capture Process as the Apply function is inherent in the use of transformation. Therefore it is not possible nor is it necessary to construct an Apply Process for a Transform.

Repstance supports transformation rules, which can be applied to the following objects:

- schema name
- table name
- column name
- "primary key" attribute
- data type
- data type length
- data

Transformation rules can be triggered based on:

- schema name mask
- table name mask
- column name mask
- "primary key" attribute
- data type mask
- data type length

Transformation rules can be used to exclude the following objects:

- table
- column
- "primary key" attribute

Transformation rules can be used to add the following objects:

- column
- "primary key" attribute

Transformation rules can be used to change the following objects:

- ❑ schema name
- ❑ table name
- ❑ column name
- ❑ “primary key” attribute
- ❑ data type
- ❑ data type length
- ❑ data

Note – Data can be changed to the :

- ❑ constant value
 - ❑ original value of another column from **this** table
 - ❑ result of SQL Server function execution. The original values of any columns from this table can be passed to the function as parameters. **The Function must exist in the Target Database, where the Apply Process is to be run.**
-

There are normally three separate parts that comprise a Transformation. They are:

- ❑ **id** – the id, which is unique to each Transform has the purpose of determining the order in which the Transform is executed by the process
- ❑ **rule** – specifies the object to be transformed and specifies how it is to be transformed
- ❑ **description** – a meaningful description of this Transformation

The syntax used is:

```
map=id,rule=(CaptureCriteria:TransformCriteria),
description="description"
```

where:

```
CaptureCriteria=schema_mask.table_mask.column_mask.[attribute_specification]
```

```
attribute_specification=PK=PK,TYPE="type",ID=id,LEN=len
```

The CaptureCriteria consists of:

- **schema_mask** – the mask is used to specify schema name
- **table_mask** – the mask is used to specify table name
- **column_mask** – the mask is used to specify column name

The attribute_specification is:

- **PK** – the specification of PK attribute. It accepts either **N** – to determine column without PK and **U** – to include PK column
- **Type** – the name of the data type
- **LEN** – length of data type if applicable

```
TransformCriteria=schema_mask.table_mask.column_mask.[attribute_specification]
```

```
attribute_specification=PK=PK,TYPE="type",ID=id,LEN=len,DATA="data_specification"
```

`data_specification=DATA="SQL_function_or_operation"`

The `TransformCriteria` consists of:

- **schema_mask** – the mask is used to specify schema name
- **table_mask** – the mask is used to specify table name
- **column_mask** – the mask is used to specify column name

The `attribute_specification` is:

- **PK** – the specification of PK attribute. It accepts either **N** – to determine column without PK and **U** – to include PK column
- **Type** – the name of the data type
- **LEN** – length of data type if applicable
- **DATA** – this is the data that will result from the transformation specified on the captured data. This can be:
 - Predefined value. The format is "*predefined_value*"
 - Any operations based on the values of this column. In order to use the values found in the column the `%` needs to be used
 - Any operations based on the values of any columns in this table. In order to use the values of any column the following format is used "*C:number_of_column*". Table name can be passed as parameter as well. The format is "*T:0*".

To **exclude columns** the following syntax is used:

`schema_mask.table_mask.NULL`

where:

- **schema_mask** – the mask is used to specify schema name
- **table_mask** – the mask is used to specify table name.

To **include any columns that were not represented in the source table** the following syntax is used:

`schema_mask.table_mask.column_name.[TYPE="type",LEN=len,
PK=pk,DATA="data_specification"]`

where:

- **schema_mask** – the mask is used to specify schema name
- **table_mask** – the mask is used to specify table name
- **column_name** – column name (this must be included and can not use the `%` symbol)
- **type** – data type definition of the column to be created (this must be included and can not use the `%` symbol)
- **len** – length of data type, where necessary must be specified here
- **PK** – the specification of PK attribute. It accepts either **N** – to determine column without PK and **U** – to include PK column
- **DATA** – data specification

The following examples show how to use the Transformation Rules:

- For all tables located in the **scott** schema change the schema name to **report**:

```
map=1,rule=(scott.:%:report.%),description="sample 1"
```
- Rename **scott.emp** table to **report.employees** table:

```
map=2,rule=(scott.emp:report.employees),description="sample 2"
```
- Exclude **info** column in the **scott.customer** table from the replication:

```
map=3,rule=(scott.customer.info:scott.customer.NULL),description="sample 3"
```
- For any tables in the **scott** schema having name starting with **rep_** replace the prefix to the **report_**:

```
map=4,rule=(scott.rep_%:scott.report_%),description="sample 4"
```
- Exclude any columns having name **id** and defined as primary key from the replication:

```
map=5,rule=(scott.%.id.[PK=U]:scott.%.id.NULL),description="sample 5"
```
- For any tables having column **id** of **int** type handle this column as primary key. Any update and delete statements will be built to use **id** column as primary key or the part of composite primary key:

```
map=6,rule=(scott.%.id.[TYPE=int]:scott.%.id.[PK=P]),description="sample 6"
```
- Put the value of the **desc** column in the **branch** table into the uppercase:

```
map=7,rule=(scott.branch.desc:scott.branch.desc.[DATA="\upper(%)\"])
```
- Add new column named **tag** of the **varchar2(100)** type into the **customer** table contains predefined value **rep**:

```
map=8,rule=(scott.customer.NULL:scott.customer.tag.[TYPE=varchar2;LEN=100;DATA="'rep'\"])
```
- Add new column named **name** of the **varchar2(255)** type to the **person** table contains concatenation of the **firstname** and **lastname** columns (in the table **firstname** has the position 2 and **lastname** has the position 3):

```
map=9,rule=(scott.person.NULL:scott.person.name.[TYPE=varchar2;LEN=255;DATA="\C:2+C:3\"])
```

Transformation Rules and Triggering Order

It is imperative that you understand that the order in which the rules are applied can significantly impact on the expected results. The following logic is used to handle rule order processing. The specified Capture Process will go and look for any possible data that meets the first of the configured rules, to apply. In the case where there is more than one rule to be processed, the rule **which will** “take precedence” will be the rule which best fulfills the data found. If the objects found matches the criteria given in the **first rule** to be processed any subsequent rules will be ignored because the previous criteria have been met.

The logic used in transformation does not necessarily follow the rules for human logic, in that, the least significant amount of data to be processed should be the **first rule**, and the rule which affects the largest amount of data should be the **last rule**. In this way the user will be able to apply the requirements of the transform in the desired way.

The following example demonstrates this program logic. Assume that you need to rename all tables in the **scott** schema having name starting from **emp** into the schema **hr** and keep the original table name. Another requirement is to rename **ident** columns be the **id** for any tables in the **scott** schema. In the case we have **employee** table contains **ident** column it matches to both rules. The following examples demonstrate replication behaviour depending on the rules order:

- In this case you can't expect the **ident** column to be renamed into **id** column together with renaming **scott** schema in **hr** schema. The **scott** schema for **employee** table will be changed to **hr** schema but the **ident** column will not be changed:

```
map=1,rule=(scott.emp%:hr.emp%)
map=2,rule=(scott.%.ident:scott.%.id)
```

- In this case the **scott** schema will not be changed into the **hr** schema with renaming the **ident** column to the **id** column. The **scott** schema for **employee** table will not be changed but the **ident** column will be changed to **id**:

```
map=1,rule=(scott.%.ident:scott.%.id)
map=2,rule=(scott.emp%:hr.emp%)
```

- In this case the **scott** schema for **employee** table will be changed and the **ident** column will be changed to **id**:

```
map=1,rule=(scott.emp%.ident:hr.emp%.id)
map=2,rule=(scott.emp%:hr.emp%)
map=3,rule=(scott.%.ident:scott.%.id)
```

Alter Capture Process

The "Alter" Capture command is used to change the Capture Process parameters. The **id** parameter can not be changed. The **name** parameter can be changed only in the case that it has not been provided before - eg. it must be "Unique".

The Capture Process has to exist otherwise the command will fail.

Note – You can alter the Capture Process even if it is running, however these changes will only be implemented after the Capture Process has been stopped and rerun.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "alter",
  "parameters": [
    ["process", "capture"],
    ["id", "captureID"],
    ["dbtype", "oracle"],
    ["name", "captureName"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "databaseName"],
    ["autostart", "{0/1}"],
    ["debuglevel", "{0-15}"],
    ["maxtrailsize", "{10000000-9999999999}"],
    ["ddlinclude", "objectsMask"],
    ["ddlexclude", "objectsMask"],
    ["dmlinclude", "objectsMask"],
    ["dmlexclude", "objectsMask"],
    ["loadinclude", "objectsMask"],
    ["loadexclude", "objectsMask"],
    ["map=mapID", "mappingClause"]
  ]
}
```

Server response:

- **HTTP Status** – The status of the command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Failed|Completed}",
  "Message": "The command execution details"
}
```

CLI Syntax:

```
repcli alter process=capture id=captureID \
  dbtype=oracle name=captureName \
  connectiontype=tns|ezconnect \
  tnsname=tns_alias \
  server=databaseHost port=databasePort \
  servicename=service_or_SID
  user=username password=password \
  dbname=databaseName autostart={0/1} \
  debuglevel={0-15} maxtrailsize={10000000-999999999} \
  ddlinclude=objectsMask ddlexclude=objectsMask \
  dmlinclude=objectsMask          dmlexclude=objectsMask \
  loadinclude=objectsMask loadexclude=objectsMask \
  map=mapID, mappingClause
```

The input parameters are:

- **process** – *capture*
- **id** – Capture Process id
- **dbtype** – type of RDBMS. The possible values are:
 - *oracle*
 - *mssql* – see SQL Server documentation for the details
- **connectiontype** – Specifies the Oracle connection method. The possible values are:
 - **tns** – Local Naming Method to be used
 - **ezconnect** – EZCONNECT to be used
- **tnsname** – Name of the TNS alias. Valid **only** if **connectiontype=tns**
- **server** – host name or IP address of the database server. Valid **only** if **connectiontype=ezconnect**
- **port** – database port number. Valid **only** if **connectiontype=ezconnect**
- **servicename** – database service name or SID. Valid **only** if **connectiontype=ezconnect**
- **dbname** – name of either container or pluggable database. Valid **only** if **database=source** and for any Oracle versions 12c-19c but excluding RDS instances
- **name** – name of the Capture Process
- **user** – database user name
- **password** – database user password
- **autostart** – determine if the process must be run automatically. The possible values are:
 - **0** – do not run the process automatically
 - **1** – run the process automatically
- **debuglevel** – the level of debugging. Possible values are 0-15
- **ddlinclude** – mask of the DDL objects to be captured. See “[Capture Objects Specification](#)” chapter for the details

- **ddlexclude** – mask of the DDL objects to be skipped by the Capture Process. See “[Capture Objects Specification](#)” chapter for the details
- **dmlinclude** – mask of the DML objects to be captured. See “[Capture Objects Specification](#)” chapter for the details
- **dmlexclude** – mask of the DML objects to be skipped by the Capture Process. See “[Capture Objects Specification](#)” chapter for the details
- **loadinclude** – Mask of the objects to be included into the Initial Load (see chapter [Initial Load](#) for the details)
- **loadexclude** – Mask of the objects to be skipped during Initial Load (see chapter [Initial Load](#) for the details)
- **map** – the set of the parameters to determine objects’ transformation. See “[Prepare Capture Process Objects Mapping](#)” chapter for the details.

Show Capture Process

The “Show” Capture command is used to show a Capture Process configuration. This command will show the latest “Set” of configured processes, regardless of whether or not the Capture Process has been “stopped and re-applied”, this will not necessarily be the configuration of the currently running Capture Process.

Note – If the configuration has been changed when the Capture Process is running, the “Show” command will still display the currently configured process, not the “running” one.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** *token*
- **Body:**

```
{
  "command": "show",
  "parameters": [
    ["process", "capture"],
    ["id", "captureID"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Failed/Completed}",
  "Message": "The command execution details.",
  "parameters": [
```

```

    ["id", "captureID"],
    ["dbtype", "oracle"],
    ["name", "captureName"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "databaseName"],
    ["autostart", "{0/1}"],
    ["debuglevel", "{0-15}"],
    ["maxtrailsize", "{10000000-9999999999}"],
    ["ddlinclude", "objectsMask"],
    ["ddlexclude", "objectsMask"],
    ["dmlinclude", "objectsMask"],
    ["dmlexclude", "objectsMask"],
    ["loadinclude", "objectsMask"],
    ["loadexclude", "objectsMask"],
    ["map=mapID", "mappingClause"]
  ]
}

```

CLI Syntax:

```
repcli show process=capture id=captureID
```

The input parameters are:

- **process** – the value is **capture**. This is a “constant” value and **can only be capture**
- **id** – the process identifier

Note – Only if the command completes successfully will the response contain all the “parameter” values and at the same time this information will be displayed in Plain Text in the “Message” display.

Validate Capture Process

The “Validate” Capture Process command is used to perform a set background checks to ensure that all the necessary parameters have been supplied to enable the functionality of a Capture configuration.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** token
- **Body:**

```

{
  "command": "validate",
  "parameters": [
    ["process", "capture"],
    ["id", "captureID"]
  ]
}

```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```

{
  "Status": "{Valid/Invalid}",
  "Message": "The command execution details",
}

```

CLI Syntax:

```

repcli validate process=capture id=captureID

```

The input parameters are:

- **process** – the value is **capture**. This is a “constant” value and **can only be capture**
- **id** – the process identifier

In the case the **Status** is **Invalid** the **Message** will contain detailed information about the validation error.

Remove Capture Process

The “Remove” Capture Process command is used to remove all of the Capture Process- configuration files, Trail Files and the checkpoint information.

The command can only be executed if the specified Capture Process is not running.

It is strongly recommended that you backup all trail and configuration files before running this command (see the “[Backup Repstance Files](#)” chapter).

Note – After removing the Capture Process it is not possible to reverse this action.

Any Apply Processes using data from this Capture will cease to work.

This command cannot be applied to a group of “Configures” at the same time, it must be used individually for each of the configured Capture Processes.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{  
  "command": "remove",  
  "parameters": [  
    ["process", "capture"],  
    ["id", "captureID"]  
  ]  
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{  
  "Status": "{Completed/Error/Warning}",  
  "Message": "The command execution details",  
}
```

CLI Syntax:

```
repcli remove process=capture id=captureID
```

The input parameters are:

- **command** – remove
- **process** – capture
- **id** – the identifier of the Capture Process

The output parameters are:

- **Status** – status of the command. The possible values are:
 - **Error** – the process has failed to run
 - **Warning** – the process has failed to remove all the Capture objects
 - **Completed** – the process has completed successfully
- **Message** – the details of command execution

The following example demonstrates how to remove Capture 2:

```
repcli remove process=capture id=2
```


Prepare Apply Process

The “Prepare” Apply command is used to add a new Apply Process. The “Prepare” command can not be used to validate the database connection, and availability of the Trail Files or the presence of the necessary database objects to be used by the Apply Process.

In order to validate that the process is configured properly and it is able to run on the Target Database, the “[Validate](#)” command must be used.

The “Prepare” command does not enable the Apply Process to write data into the Target Database until the specified Apply Process is run (see the “[Run any Capture or Apply Processes](#)” chapter for the details).

Parameter **id** is mandatory. All others parameters are optional. All parameters except **id** can be changed using “[Alter](#)” Apply Process command.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** token
- **Body:**

```
{
  "command": "prepare",
  "parameters": [
    ["process", "apply"],
    ["id", "applyID"],
    ["dbtype", "oracle"],
    ["capturename ", "captureName"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "databaseName"],
    ["user", "username"],
    ["password", "password"],
    ["autostart", "{0/1}"],
    ["ddlcreate", "objectMask:options"],
    ["ddldrop", "objectMask:options"],
    ["debuglevel", "{0-15}"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs

- **Body:**

```
{
  "Status": "{Failed/Completed}",
  "Message": "The command execution details"
}
```

CLI Syntax:

```
repcli prepare process=apply id=applyID capturename=captureName \
  server=databaseHost port=databasePort \
  user=username password=password \
  autostart={0/1} \
  ddlcreate=objectMask:options ddldrop=objectMask:options \
  debuglevel={0-15}
```

The input parameters are:

- **process** – apply
- **id** – the process id
- **dbtype** – type of RDBMS. The possible values are:
 - *oracle*
 - *mssql* – see SQL Server documentation for the details
- **connectiontype** – Specifies the Oracle connection method. The possible values are:
 - **tns** – Local Naming Method to be used
 - **ezconnect** – EZCONNECT to be used
- **tnsname** – Name of the TNS alias. Valid only if **connectiontype=tns**
- **server** – host name or IP address of the database server. Valid only if **connectiontype=ezconnect**
- **port** – database port number. Valid only if **connectiontype=ezconnect**
- **servicename** – database service name or SID. Valid only if **connectiontype=ezconnect**
- **capturename** – the name of the Capture Process which is providing the data to be used
- **user** – database user name
- **password** – database user password
- **autostart** – determine if the process must be run automatically. Possible values are:
 - **0** – do not run the process automatically
 - **1** – run the process automatically
- **ddlcreate** – used to define the Apply Process behavior when creating a table, if a table with the same name is found. The table is specified in *schema_name.table_name* format. In order to specify the list of the tables, each table must be separated by comma. The % symbol may be used in order to match any number of characters. It has the following options:
 - **skip** – do not run DDL create command if the table already exists
 - **recreate** – drop the existing table and recreate it
- **ddldrop** – used to define the Apply Process behavior when dropping a table, if a table with the same name is not found. The table is specified in *schema_name.table_name* format. In order to specify the list of the tables,

each table must be separated by comma. The % symbol may be used in order to match any number of characters. It has the following options:

- **skip** – do not run DDL drop command if the table doesn't exist
- **debuglevel** – the level of debugging. Possible values are 0-15

The following example demonstrates how to configure a new Apply Process having **id=2**. The Apply Process is configured to read data provided by the **cap2** process and apply it into the **usoradb** database, which is running on the **usoradb.ckppywxf9a8u.us-east-1.rds.amazonaws.com:1521** server.

The Process has now been configured to run automatically, once the Repstance service has been started:

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** token
- **Body:**

```
{
  "command": "prepare",
  "parameters": [
    ["process", "apply"],
    ["id", "2"],
    ["dbtype", "oracle"],
    ["connectiontype", "ezconnect"],
    ["capturename", "cap2"],
    ["server", "usoradb.ckppywxf9a8u.us-east-1.rds.amazonaws.com"],
    ["port", "1521"],
    ["servicename", "usaoradb"],
    ["user", "dbusr"],
    ["password", "dbpwd1"],
    ["autostart", "1"]
  ]
}
```

Alter Apply Process

The “Alter” Apply command is used to change the Apply Process parameters. The **id** parameter can not be changed. The Process has to exist otherwise the command will fail. The **capturename** parameter can be altered only in the case that it has not been used earlier.

Note – You can alter the Apply Process even if it is running, however **these changes will only be implemented after this Apply Process has been stopped and rerun.**

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "alter",
  "parameters": [
    ["process", "apply"],
    ["id", "applyID"],
    ["dbtype", "oracle"],
    ["capturename ", "captureName"],
    ["connectiontype", "tns|ezconnect"],
    ["tnsname", "tns_alias"],
    ["server", "host_name"],
    ["port", "port_number"],
    ["servicename", "service_or_SID"],
    ["dbname", "databaseName"],
    ["user", "username"],
    ["password", "password"],
    ["autostart", "{0/1}"],
    ["ddlcreate", "objectMask:options"],
    ["ddldrop", "objectMask:options"],
    ["debuglevel", "{0-15}"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Failed/Completed}",
  "Message": "The command execution details"
}
```

CLI Syntax:

```
repcli alter process=apply id=applyID capturename=captureName \
server=databaseHost port=databasePort \
user=username password=password \
autostart={0/1} \
ddlcreate=objectMask:options ddldrop=objectMask:options \
debuglevel={0-15}
```

The input parameters are:

- **process** – apply
- **id** – the process id

- **dbtype** – type of RDBMS. The possible values are:
 - *oracle*
 - *mssql* – see SQL Server documentation for the details
- **connectiontype** – Specifies the Oracle connection method. The possible values are:
 - **tns** – Local Naming Method to be used
 - **ezconnect** – EZCONNECT to be used
- **tnsname** – Name of the TNS alias. Valid only if **connectiontype=tns**
- **server** – host name or IP address of the database server. Valid only if **connectiontype=ezconnect**
- **port** – database port number. Valid only if **connectiontype=ezconnect**
- **servicename** – database service name or SID. Valid only if **connectiontype=ezconnect**
- **capturename** – the name of the Capture Process which is providing the data to be used
- **user** – database user name
- **password** – database user password
- **autostart** – determine if the process must be run automatically. Possible values are:
 - **0** – do not run the process automatically
 - **1** – run the process automatically
- **ddlcreate** – used to define the Apply Process behavior when creating a table, if a table with the same name is found. The table is specified in *schema_name.table_name* format. In order to specify the list of the tables, each table must be separated by comma. The % symbol may be used in order to match any number of characters. It has the following options:
 - **skip** – do not run DDL create command if the table already exists
 - **recreate** – drop the existing table and recreate it
- **ddldrop** – used to define the Apply Process behavior when dropping a table, if a table with the same name is not found. The table is specified in *schema_name.table_name* format. In order to specify the list of the tables, each table must be separated by comma. The % symbol may be used in order to match any number of characters. It has the following options:
 - **skip** – do not run DDL drop command if the table doesn't exist
- **debuglevel** – the level of debugging. Possible values are 0-15

Show Apply Process

The “Show” Apply Process command is used to show an Apply configuration. This command will show the latest “Set” of configured Process’s, regardless of whether or not the Apply Process has been “stopped and re-applied”, this will not necessarily be the configuration of the currently running Apply Process.

Note – If the configuration has been changed when the Apply Process is running, the Show command will still display the currently configured process, and not the “running” one.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{  
  "command": "show",  
  "parameters": [  
    ["process", "apply"],  
    ["id", "applyID"]  
  ]  
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{  
  "Status": "{Failed/Completed}",  
  "Message": "The command execution details.",  
  "parameters": [  
    ["process", "apply"],  
    ["id", "applyID"],  
    ["dbtype", "oracle"],  
    ["capturename ", "captureName"],  
    ["connectiontype", "tns|ezconnect"],  
    ["tnsname", "tns_alias"],  
    ["server", "host_name"],  
    ["port", "port_number"],  
    ["servicename", "service_or_SID"],  
    ["dbname", "databaseName"],  
    ["user", "username"],  
    ["password", "password"],  
    ["autostart", "{0/1}"],  
    ["ddlcreate", "objectMask:options"],  
    ["ddldrop", "objectMask:options"],  
    ["debuglevel", "{0-15}"]  
  ]  
}
```

CLI Syntax:

```
repcli show process=apply id=applyID
```

The input parameters are:

- **process** – apply

- **id** – the process id

Note – Only in the case that the command completes successfully will the response contain all the “parameter” values and at the same time this information will be displayed in Plain Text in the “Message” display.

Validate Apply Process

The “Validate” Apply Process command is used to perform a set background checks to ensure that all the necessary Parameters have been supplied to enable the functionality of an Apply configuration on the Target Database. This command will also provide the “last” successfully applied transaction’s SCN (in hexadecimal format) but **only** if the Apply Process has previously been run or reset.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "validate",
  "parameters": [
    ["process", "apply"],
    ["id", "applyID"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Valid/Invalid}",
  "Message": "The command execution details",
}
```

CLI Syntax:

```
repcli validate process=apply id=applyID
```

The input parameters are:

- **process** – apply
- **id** – the process id

In the case the **Status** is **Invalid** the **Message** will contain detailed information about the validation error.

Reset Apply Process

This command can **only** be used for an Apply Process – it is used where there is a need to force the Apply Process to start from specific SCN or after this SCN.

It is primarily used to “Set” or “Change” the “Startpoint” for an Apply Process.

REST API:

- **Endpoint:** https://repstance_url/reset/process
- **Method:** POST
- **Header:**
 - **Content-Type:** application/json
 - **X-Token:** *token*
- **Body:**

```
{
  "command": "reset",
  "parameters": [
    ["process", "apply"],
    ["id", "applyID"],
    ["LSN", "SCN (in hexadecimal format)"],
    ["skip", "{0/1}"],
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Completed/Failed}",
  "Message": "The command execution details",
}
```

CLI Syntax:

```
repcli reset process=apply id=applyID LSN=SCN skip={0/1}
```

The input parameters are:

- **process** – apply
- **id** – the process id
- **SCN** – SCN in hexadecimal format at which we need to start reprocessing data.
- **skip** – determine if we need to jump over or not, a specific SCN

The following example demonstrates how to reset an Apply Process to start from the first transaction found, after SCN = 0x000000450000007B0004:

```
repcli reset process=apply id=1 lsn=0x000000450000007B0004
      skip=1
```

The following example demonstrates how to reset an Apply Process to start using the transactions from the SCN = 0x000000450000007F0004:

```
repcli reset process=apply id=1 lsn=0x000000450000007F0004
      skip=0
```

Remove Apply Process

The “Remove” Apply Process command is used to remove **all** of the Apply Process – configuration files and the checkpoint details, together with historical information. If the Target Database is **reachable** the “Remove” command will clean-up all the checkpoint information on the Target Database.

The command can only be executed if the specified Apply Process is not running.

It is strongly recommended that you backup configuration files and last successfully processed SCN before running this command (see the “[Backup Repstance Files](#)” chapter).

This command cannot be applied to a group of Apply Processes at the same time, it must be used individually for each of the configured Apply Processes.

REST API:

- **Endpoint:** https://repstance_url/configure/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "remove",
  "parameters": [
    ["process", "apply"],
    ["id", "applyID"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs

- **Body:**

```
{
  "Status": "{Completed/Error/Warning}",
  "Message": "The command execution details",
}
```

CLI Syntax:

```
repcli remove process=apply id=applyID
```

The input parameters are:

- **command** – remove
- **process** – apply
- **id** – the identifier of the Apply Process

The output parameters are:

- **Status** – status of the command. The possible values are:
 - **Error** – the command has failed to run
 - **Warning** – the process has failed to remove all the Apply Objects
 - **Completed** – the command has completed successfully
- **Message** – the details of command execution

The following example demonstrates how to remove Apply Process having id=2:

```
repcli remove process=apply id=2
```

Control Repstance Processes

There are three commands are used to control the Capture and Apply Processes. They are :

- run
- stop
- status

Each of these commands can be used to control a single Process, a group of Processes or all Processes. **The default is to run “All” the Processes that are available to be run, stopped or get the status of, if no additional “Parameters” are specified.**

The syntax used for these commands is:

```
repcli run|stop|status [process=(capture|apply)
[id=processID]]
```

Here are some **examples** of using “Run” command.

This command will “run” everything (all existing Capture and Apply Processes):

```
repcli run
```

This command will “run” ALL Capture Processes:

```
repcli run process=capture
```

This command will “run” a single specified Capture Process:

```
repcli run process=capture id=1
```

The command format is identical in use for both the “Stop” and “Status” commands.

Run any Capture or Apply Processes

This command is used to “Run” either a single Capture or Apply Process or a group of either Capture, or Apply Processes or all configured Processes. **The default is to run all Processes that are available to be run, if no additional “Parameters” are specified.**

As part of the “Run” command, Repstance carries out background checks to ensure that Source and/or Target Databases are properly configured with the necessary functionality to enable the Capture and/or Apply Processes to run. In the event that either of the Databases are not properly configured, it will display an error message listing the inconsistencies.

If this is the “First time” a Capture Process has been run, at this point the first SCN will be written to the Trail File generated by this Capture Process.

If the previously defined Objects (see the “[Capture Objects Specification](#)” chapter) are changed or the Transformation Rules (see the “[Transformation Rules and Triggering Order](#)” chapter) are altered, at this point all the DDL and DML changes will be re-implemented using the new criteria.

If any Objects are no longer configured for Capture Process the “Run” command will ensure that these object are no longer be written to the Trail Files – the reverse is true i.e. if new Objects are specified then the “Run” command will start writing these newly added objects to the Trial Files.

If this is the “First time” an Apply Process has been run, it is **ONLY** at this point that data from the specified Capture Process will be written to the Target Database using the first available SCN in the Trail File generated by this Capture Process.

The Apply Process can be configured to use a specific SCN from a Capture Process as a start point, i.e. it can be configured to use any available SCN from a Trail File, but this is done by using the “Reset” command (see the “[Reset Apply Process](#)” chapter).

If the specified Capture Process either is not configured or has never been run, the Apply Process will fail.

REST API:

- **Endpoint:** https://repstance_url/control/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "run",
  "parameters": [
    ["process", "apply/capture"],
    ["id", "processID"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Completed/Error}",
  "Message": "The command execution details",
}
```

CLI Syntax:

```
repcli run process=apply|capture id=processID
```

The input parameters are:

- **command** – run
- **process** – constant value **apply** for the Apply Processes or constant value **capture** for the Capture Processes
- **id** – the identifier of the Process

The output parameters are:

- **Status** – status of the command. The possible values are:
 - **Error** – the Process has failed to run
 - **Completed** – the Process has completed successfully
- **Message** – the details of command execution

Stop any Capture or Apply Processes

This command is used to stop either a single Capture or an Apply Process or a group of either Capture, or Apply Processes or all configured Processes. **The default is to stop everything that is running, if no additional parameters are specified.**

REST API:

- **Endpoint:** https://repstance_url/control/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "stop",
  "parameters": [
    ["process", "apply|capture"],
    ["id", "processID"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
{
  "Status": "{Completed/Error}",
  "Message": "The command execution details",
}
```

CLI Syntax:

```
repcli stop process=apply|capture id=processID
```

The input parameters are:

- **command** – stop
- **process** – constant value **apply** for the Apply Process or constant value **capture** for the Capture Process
- **id** – the identifier of the Process

The output parameters are:

- **Status** – status of the command. The possible values are:
 - **Error** – the command has failed to run
 - **Completed** – the command has completed successfully
- **Message** – the details of command execution

Status of any Capture or Apply Processes

This command is used to show the status of either a single Capture or an Apply Process or a group of either Capture, or Apply Processes or all configured Processes. **The default is to show the status of every Process that is running, if no additional parameters are specified.**

REST API:

- **Endpoint:** https://repstance_url/control/process
- **Method:** POST
- **Header:**
 - Content-Type: application/json
 - X-Token: token
- **Body:**

```
{
  "command": "status",
  "parameters": [
    ["process", "apply/capture"],
    ["id", "processID"]
  ]
}
```

Server response:

- **HTTP Status** – the status of command. The possible codes are:
 - **200** – if no error occurs
 - **422** – if error occurs
- **Body:**

```
[
  {
    "Process": "capture|apply",
    "ID": "processID",
    "Parameters": {
      "Name|CaptureName": "captureName",
      "Server": "databaseHost",
      "Port": "databasePort",
      "DBname": "databaseName"
    },
    "Command": "INITIALISED|RUN|STOP|RESET",
    "CommandStatus": "Completed|In Process|Failed",
    "StatusDetails": "Waiting for the Next Command|
      Processing the Transactions|
      Waiting for Transactions",
    "CommandStartTime": "Time_in_UTC",
    "LastActivityTime": "Time_in_UTC",
    "StartLSN": "SCN",
    "StartLSNTime": "Time_in_UTC",
    "CheckpointLSN": "SCN",
    "CheckpointLSNTime": "Time_in_UTC",
    "LastDBLSN": "SCN",
    "LastDBLSNTime": "Time_in_UTC",
    "WAITING": "time_in_microseconds",
    "PROCESSING": "time_in_microseconds",
    "Message": "The command execution details",
    "CNTD": "Number_of_deletes",
    "CNTI": "Number_of_inserts",
    "CNTU": "Number_of_updates",
    "CNTDDL": "Number_of_ddls",
    "CNTTX": "Number_of_transactions"
  },
  {

```

```

    Next_Process_description
  }, ...
]

```

CLI Syntax:

```

repcli status process=apply|capture id=processID

```

The input parameters are:

- **command** – **status**, the constant value
- **process** – the **apply** value for an Apply Process or the **capture** value for a Capture Process
- **id** – the identifier of the Process

The output parameters are:

- **Process** – name of the configured Process. The possible values are:
 - **capture**
 - **apply**
- **ID** – the identifier of the Process
- **Name** – the name of a Capture Process (only for Capture Processes)
- **CaptureName** – the name of Capture Process being used by the Apply Process (only for Apply Processes)
- **Connection** – the current connection details (if the connectiontype was specified as ezconnect):
 - **Server** – host name or IP address of the database server
 - **Port** – database port
 - **ServiceName** – Service Name
 - **Dbname** – database name if available
- **TNS** – TNS names (if the connectiontype was specified as tns)
- **Command** – the last executed command on the listed Process. The possible values are:
 - **INITIALISED** – means that the Process exists but currently has no command executed
 - **RUN** – means that the Process is running
 - **STOP** – means that the Process is stopping
 - **RESET** – means that the Apply Process is resetting to the specified SCN
- **CommandStatus** – the status of the command. The possible values are:
 - **Waiting for the Next Command**
 - **Processing the Transactions**
 - **Waiting for Transactions**
- **CommandStartTime** – the time in UTC format when the command was initialised (the parameter is available only in JSON response)
- **LastActivityTime** – the time in UTC format when the Process activity changed (the parameter is available only in JSON response)
- **StartLSN** – the first SCN (in hexadecimal format) when a transaction was processed (in **repcli** the parameter is displayed in the “Total since” section)
- **StartLSNTime** – the time of the transaction in UTC format the Process has started with (in **repcli** the parameter is displayed in the “Total since” section)
- **CheckpointLSN** – the SCN (in hexadecimal format) of the last transaction that has been successfully processed (in **repcli** the parameter is displayed in the “Last Captured/ Applied DB Change” section)

- **CheckpointLSNTime** – the time in UTC format of the last transaction that has been successfully processed (in **repcli** the parameter is displayed in the “Last Captured/Applied DB Change” section)
- **LastDBLSN** – the last known transaction SCN (in hexadecimal format) in the Source Database. The parameter is provided by Capture Process and valid only in the case the Capture Process is running (in **repcli** the parameter is displayed in the “Last Sourced DB Change” section)
- **LastDBLSNTime** – the time in UTC format of the last known transaction SCN (in hexadecimal format) in the Source Database. The parameter is provided by Capture Process and valid only in the case the Capture Process is running (in **repcli** the parameter is displayed in the “Last Sourced DB Change” section)
- **WAITING** – number in nanoseconds the Process has been waiting for data to process (in **repcli** the parameter is displayed in HH:MM:SS.FFF format)
- **PROCESSING** – number in nanoseconds the Process has been processing the data (in **repcli** the parameter is displayed in HH:MM:SS.FFF format)
- **Message** – the details of command execution
- **CNTD** – number of Delete Operations currently performed by the Process
- **CNTI** – number of Insert Operations currently performed by the Process
- **CNTU** – number of Update Operations currently performed by the Process
- **CNTDDL** – number of DDL statements currently processed by the Process
- **CNTTX** – number of Transactions currently processed by the Process

The **Average Speed** parameter is average number of records processed per second. The **Lag** is the latency between last record processed and timestamp of the last transaction in the Source Database. The **Lag** is valid only in the case the corresponding Capture Process is running. In **repcli** it is displayed in HH:MM:SS.FFF format.

The following example shows statuses of all existing Processes:

repcli status

```
Capture Process 1, Name: db1, TNS: RAC.local
  RUN (In Process), Waiting for Transactions
    Last Source DB Change      :      2019/02/20  22:17:00.630  UTC
(0x0000067E000019F20004)
    Last Captured DB Change    :      2019/02/20  22:17:00.630  UTC
(0x0000067E000019F20004)
  Total since 2019/02/10 13:36:26.653 UTC (0x000006600000000B10003):
    Processing the Transactions : 00:00:15.659
    Waiting for Transactions    : 00:08:19.490
    Average Speed (ops)        : 1.149
    Transactions: 18 (DDL: 0, Delete: 9, Insert: 9, Update: 0)

Apply Process 1, Transactions Provider (Capture): db1, Connection:
192.168.0.55:1521/ORCL
  RUN (In Process), Waiting for Transactions
    Last Source DB Change      :      2019/02/20  22:17:00.630  UTC
(0x0000067E000019F20004)
    Last Applied DB Change     :      2019/02/20  22:17:00.630  UTC
(0x0000067E000019F20004)
    Lag                        : 00:00:00.000
  Total since 2019/02/10 13:36:26.653 UTC (0x000006600000000B10003):
    Processing the Transactions : 00:00:11.039
```



```
Waiting for Transactions      : 00:08:25.287
Average Speed (ops)         : 1.630
Transactions: 18 (DDL: 0, Delete: 9, Insert: 9, Update: 0)
```

REPSTANCE SERVER MAINTENANCE

As Repstance is a pre-configured Linux service it starts/stops together with the Server and MUST be maintained using the Linux Service Command.

How to Stop Repstance Services

In order to stop the Repstance Server the “service repstance stop” must be used:

```
# service repstance stop
Stopping repstance:          [ OK ]
```

How to Start Repstance Services

In order to start Repstance Server the “service repstance start” must be used:

```
# service repstance start
Starting repstance:         [ OK ]
```

How to Check Repstance Service Status

In order to check Repstance Server status the “service repstance status” must be used:

```
# service repstance status
repstance (pid 3033) is running...
```

Backup Repstance Files

Regularly backup the Trail Files which are located in /opt/repstance/trail directory along with the Configuration Files, which are located in /opt/repstance/conf and Capture Checkpoint Files in /opt/repstance/captureckpt.* and lastly make sure that you record the last applied checkpoint in a file of your choosing.

Housekeeping

Schedule regular jobs which are responsible for cleaning up “old” Trail and Log Files which are found under these directories:

- /opt/repstance/trail – Trail Files
- /opt/repstance/log – Log Files

GLOSSARY

Source Database – is the Database from which you wish to extract/copy Data from.

Target Database – is the Database into which you wish to insert/copy Data to.

Capture Process – is the “Process” that extracts Data from the specified Source Database.

Apply Process – is the “Process” that inserts the extracted Data into the specified Target Database.

Process – either Capture or Apply Process.

Trail Files – these are the “Source Files” – containing Data extracted by a Capture Process and used by an Apply Process.

SCN – every record in the Trail Files is identified by a System Change Number generated by Oracle.

Capture Checkpoint – the last SCN identifier for data change captured.

Apply Checkpoint – the last SCN identifier for data change applied.

DML – a type of the Database Statement, that is used to modify Data in the tables.

DDL – a type of the Database Statement, that is used to provide or modify a Dictionary Definition of a database object.

Transformation Rules – these are the rules provided by user to overwrite transaction processing behaviour.