

Ultimate Guide to OpenEdge Replication

OpenEdge Replication is a great product that does exactly what it's supposed to do, but there are plenty of important details that you cannot learn by reading the manual.

Come to this session to learn best practices and all the sneaky little tricks and tips you'll need to successfully deploy, monitor and maintain OE Replication.



The Ultimate Guide to OpenEdge Replication

Paul Koufalis, White Star Software
pk@wss.com



White Star
software

Why Are You Here?

- Who is already running OE Replication?
 - Never touched it?
- After-imaging?
- OER is a component of your BCP
 - You should have a comprehensive, tested Business Continuity Plan

About White Star Software

- The Oldest and Most Respected Independent Progress OpenEdge Consulting Firm
- 4 of the top OpenEdge DBAs in the world: Adam Backman, Tom Bascom, Dan Foreman and Paul Koufalis
- Our Performance, Monitoring and Alerting Tool, ProTop. An incredibly powerful single-pane-of-glass view of your entire OpenEdge ecosystems
- Real World Training From Real World DBAs

info@wss.com | wss.com



Paul Koufalis

- Progress DBA and UNIX admin since 1994
- Expert OpenEdge technical consulting
- Wide range of experience
 - Small 10 person offices to 4000+ concurrent users
 - AIX, HPUX, Linux, Windows...if Progress runs on it, I've worked on it
- Father to these two monkeys
- pk@wss.com



Menu

Appetizer

Terminology soup

Taster's platter

Main Dish

Replication stew

Dessert

Tips and tricks pie

Terminology

- Primary and Secondary Server
 - Production and Repl boxes
 - Do not change even in DR scenario
- Source and Target DBs
 - Source DB sends changes to Target DBs (max 2 targets)
- Replication Server and Agent
 - Server on source, sends changes to agent, which are applied on targets

Terminology

- DBService Queue (aka pica buffers)
 - Similar to AI/BI buffers
- Failover/Failback Transition
 - Move production from primary to secondary (failover) and back (failback)

How Does it Work?

- Replication notes are written to repl buffer
 - like AI notes to AI buffers and BI notes to BI buffers)
- rpServer transmits notes to rpAgent
- rpAgent writes changes to target DBs

- AI notes associated with repl notes cannot be archived until repl notes applied on target
 - In case RPServer goes down
 - Result: “locked” AI extents

Making the Replication Stew

Ingredients

- Repl / Repl+ license
- After-imaging
- AIMGT (optional in the way salt is optional)
- Dbname.repl.properties file

- Not required: down time

WAIT !!!

- Before you start, check bandwidth to targets
- Set AI archive interval to 15 minutes
- Graph file size over time
- Allow for 25% replication overhead
- Use PEAK bandwidth (largest archived AI file)
- Ex.: 1 GB / 15 min = 9 Mbits/sec bandwidth per target DB

WAIT...just a little more...

- Transition to normal? Or reverse? Or set?
- Normal: Same as disabling replication on target
- Reverse: secondary becomes source and accumulates changes
 - When primary comes back online, can replicate back
 - When sync'd, moving prod to primary quick and easy
- Set: target1 becomes source, and replicates to target2

What about Backups?

- Good news: you can backup one of the target DBs
- Bad news: you cannot apply source DB AI archives against tgt bkp
 - Need to enable AI on target too

And those Pesky AI Files ?

- Non-replicated AI notes (in AI files) accumulate on source
- AI rotation implies AI status = locked, not “full”
- AI archiver will archive file, but cannot empty it for reuse
- Replication rollout or rebase: plan for AI accumulation during target DB preparation
 - Backup src + transfer bkp file + restore target
 - Could be hours

Cooking Time



Prepare dbname.repl.properties

- Must be in same directory as dbname.db
- Use same file for all servers
 - With some minor mods
- 4 stanzas: server, agent, control-agent and transition

Server Stanza

- Used by source database

```
[server]
```

```
control-agents=MTLAgent, NYCName
```

```
database=dbname
```

```
transition>manual
```

```
transition-timeout=600
```

```
agent-shutdown-action=recovery
```

```
defer-agent-startup=10080 # minutes (1 wk)
```

Control-Agent Stanza

- Used by source to locate targets
 - One stanza per agent

```
[control-agent.MtlAgent]
  name=MtlAgent
  database=dbname
  host=MtlServer
  port=db_4GL_port
  connect-timeout=86400
  replication-method=async
  critical=0
```

Watch
out for
OEE bug

seconds (24h)

Agent Stanza

- Used by agent on target DB

```
[agent]
```

```
name=Mt1Agent
```

```
database=dbname
```

```
listener-minport=20000
```

```
listener-maxport=22000
```

Transition Stanza - normal

- Used by source and/or target when transition occurs
 - Same as disable replication

```
[transition]
```

```
database-role=normal
```

```
restart-after-transition=0
```

Transition Stanza - reverse

- Used by source and/or target when transition occurs

```
[transition]
```

```
database-role=reverse
```

```
restart-after-transition=1
```

```
auto-begin-ai=1
```

```
transition-to-agents=MtlAgent
```

```
source-startup-arguments=-DBService replserv -S port
```

```
target-startup-arguments=-DBService replagent -S port
```

Transition Stanza - set

- Used by source and/or target when transition occurs

```
[transition]
```

```
replication-set=1                # True
```

```
database-role=reverse
```

```
restart-after-transition=1
```

```
auto-begin-ai=1
```

```
transition-to-agents=MTLAgent,NYCAgent
```

```
source-startup-arguments=-DBService replserv -S port
```

```
target-startup-arguments=-DBService replagent -S port
```

Put it in the Oven



Steps to Enable Replication

```
$ proutil db -C enablesitereplication source
```

```
$ probkup online db db.repl.bkp -  
REPLTargetCreation
```

```
$ scp db.repl.bkp MtlServer:/backup
```

```
$ dsrutil db -C restart server
```

```
$ prorest tgtodb /backup/db.repl.bkp
```

```
$ proutil tgtodb -C enablesitereplication  
target
```

```
$ _mprosrv tgtodb ... -DBService replagent
```



It Didn't Work

- Turn off the automatic probkup job
 - Any probkup between the ReplTargetCreation and initial server-agent handshake will break replication
- Fix the inevitable errors in dbname.repl.properties
 - Wrong server name, wrong port number, wrong something...
- Check if rpserver stopped
 - Wrong parameters in repl.properties or something else
 - Restart server (dsrutil db -C restart server)

It Didn't Work

- Increase the replication logging level in \$DLC/pmmgr.properties

```
[replserv]
```

```
install-dir=C:\Apps\OpenEdge\OE116
```

```
install-date=Wed Jun 15 21:18:32 2016
```

```
version=11.6
```

```
arguments=-logging 5
```

```
[replagent]
```

```
install-dir=C:\Apps\OpenEdge\OE116
```

```
install-date=Wed Jun 15 21:18:32 2016
```

```
version=11.6
```

```
arguments=-logging 5
```

Production Monitoring

- At a minimum:
 1. AI files status (How many empty/locked/full)
 2. AIMGT status
 3. AI Archive Directories
 4. RPServer status (6021)
 5. RPAgent status (3049)
 6. PICA Buffers
 1. Screen scrape pre-11.6
 2. VST 11.6+
- Standard in ProTop (\$\$)



Monitoring

Alerts Feed



10:00:57

(9407) Connection failure for host [REDACTED] port 4389 transport TCP.

[10/09/17 09:54:48]

[REDACTED]-PROGRESS01]

[2017/10/09@09:54:43.921-0400] P-5048 T-4896 I RPLS 62: (9407) Connection failure for host [REDACTED]-progress02 port 4389 transport TCP.
1x

OSRd OS reads/sec 23647 > 20000

(1362) Full backup started.

agentCommStat Replication Agent status 2 <> 1

agentCommStat Replication Agent status 2 <> 1

agentStatus Repl agent status -1 <> 3049

agentStatus Repl agent status -1 <> 3049

agentCommStat Replication Agent status 2 <> 1

agentCommStat Replication Agent status 2 <> 1

agentStatus Repl agent status -1 <> 3049

agentStatus Repl agent status -1 <> 3049

(10662) Recovery completed normally for agent progress02a...

(10662) Recovery completed normally for agent progress02a...

agentCommStat Replication Agent status 2 <> 1

agentCommStat Replication Agent status 2 <> 1

agentStatus Repl agent status -1 <> 3049

agentStatus Repl agent status -1 <> 3049

(10661) The Fathom Replication Server is beginning recovery...

(10661) The Fathom Replication Server is beginning recovery...

(10492) A communications error -157 occurred in function ...

10:00:26

09:58:31

09:58:31

09:58:25

09:58:25

09:55:04

09:55:04

09:54:59

Transition

- Three types (my terms)
 - To normal (effectively disable replication)
 - Forced (prod crash)
 - Controlled (reverse, set)
- Can be tricky – test test test
- High pmmgr logging level very useful

Transition to Normal

```
[transition]  
database-role=normal
```

- Essentially same as “proutil db -C disablesitereplication target”

Forced Transition - Reverse

- Production crash
- Target becomes source
- AI/repl notes accumulate on secondary source
- Intend to resurrect primary quickly
- At some point, you may have to disable repl on secondary
 - AI notes accumulating
- No DR/BCP until primary comes back

Forced Transition – Set

- New in 11.7
- Target 1 becomes source
- Target 2: no status change
 - Before: target of primary source
 - After: target of new source on secondary server
- Never lose DR/BCP
- No issue regarding AI accumulation
- Need to rebase to return to 3-way set

Controlled Transition

- Reverse or Set transition
- All pieces up and in “Normal Processing”
- During transition source and targets communicate
- Reverse transition:
 - Source -> target and target -> source
- Set transition
 - Source 1 -> target 2, target 1 -> source, target 2 -> target 1
 - Triangle of replication

Example

```
D:\TEMP\repl>dsrutil sports -C transition failover -logging 2
```

```
Transitioning database D:\TEMP\repl\sports
```

```
-----  
23:30:22 Opening database : Succeeded  
23:30:22 Setting up transition : Succeeded  
23:30:24 Shutting down database : Shutdown is executing. (1613)  
Shutdown complete. (1614)  
Succeeded  
23:30:59 Truncating BI : Succeeded  
23:31:03 Starting database in Cur Role : 23:31:03 BROKER This broker will  
terminate when session ends. (5405)  
23:31:03 BROKER The startup of this database requires 17Mb of shared memory. Maximum  
segment size is 1024Mb.  
23:31:03 BROKER 0: Multi-user session begin. (333)
```



Example

<snip>

```
23:31:18 Synchronization in process           : Succeeded
23:31:23 Replication Server processing       : Succeeded
23:31:26 Preparing to transition Target DB    : Succeeded
23:31:58 Shutting down database              : Shutdown is executing. (1613)
```

Shutdown complete. (1614)

Succeeded

```
23:32:29 Target transition being performed   : Succeeded
23:33:13 Switching AI Extents                : Succeeded
23:33:17 Switching database role            : Succeeded
23:33:17 Updating database master block     : Succeeded
23:33:17 Comparing databases                 : Succeeded
23:33:17 Completing Target transition       : Succeeded
```

<snip>



Tips and Tricks

- Most important: pmmgr logging level
 - Otherwise information is sparse and cryptic
- Failover transition
 - Stop and restart DBs before running transition failover
 - Verify 6021/3049 status
 - Auto-restart required, but use minimal parameters
 - After successfully completed, restart DBs with real parameters

Tips and Tricks

- Monitor both source and target
 - Bug (fixed? Don't know) where source incorrectly returns 6021
- Monitor AI files Full/Locked
 - Should never be more than 1-2 full/locked
- Windows: restart server/agent through Scheduled Tasks
 - Not linked to login session
 - Export/import task.xml to fix priority issue (BELOW_NORMAL_PRIORITY)



protop

THE BEST OPENEDGE PERFORMANCE, MONITORING, AND ALERTING TOOL IN THE GALAXY! | [WSS.COM/PROTOP](https://wss.com/protop)



Detect and correct
issues before they affect
your critical business processes



Questions?



White Star
software