



**ProTop**<sup>®</sup>

# **Back to Basics: After Imaging and the AI Management Daemon**

Rob Fitzpatrick  
Software Architect / DBA Consultant  
White Star Software

# About White Star Software

For over 30 years, we have been helping companies around the world simplify the job of managing and monitoring the world's best OpenEdge applications.

Our experts, combined with ProTop, the #1 OpenEdge monitoring and alerting tool, deliver unparalleled peace of mind for your OpenEdge environments.



# ProTop

**Monitor OpenEdge.  
Anticipate Problems.  
Avert Disasters.**

Prevent downtime, increase performance, and lower costs for cloud, on-premise, and hybrid environments with the best monitoring tool designed explicitly for OpenEdge.



[protop.com](https://protop.com)

The Best OpenEdge Performance,  
Monitoring, and Alerting Tool in the Galaxy!

# Agenda

- What is After-Imaging (AI)?
- AI Basics
- Why Use AI
- How to Configure AI
- How to Monitor AI
- Recovering with AI
- AI Archive File Retention
- AI: Useful Beyond DR

# Introduction

What is After-Imaging?

Quick Definition:

*A logging system that stores all information required to reapply all changes made to a database between two points in time.*

# Introduction

Many businesses only back up once per day

Without AI, any changes made since the last backup are lost in the event of a disaster

Best of all, AI is 100% free

Available with Enterprise, Workgroup, and Personal RDBMS licenses!

# Business Considerations

It is important for the business to define acceptable data loss.

Key concepts:

- Recovery point objective (RPO)
  - How much data can you afford to lose in a disaster?
- Recovery time objective (RTO)
  - How long can you take to recover from a disaster?

Both are a measure of time (seconds/hours/days)

# Range of Recovery Options

- RPO: how many minutes of data can you lose?
  - 1+ days
    - Daily backups
    - Is anyone realistically in this situation anymore?
  - <24 hours
    - After Imaging
  - <5 minutes
    - OpenEdge Replication Plus
      - Requires after-imaging
      - Extra-cost license





# AI Basics

# After-Imaging (AI) Basics

- RDBMS write-ahead logging rule:
  - **Transaction begin**
    - Change notes written to Before Image area (BI)
    - Change notes written to current After Image(AI) area
    - Changes are made to blocks in buffer pool
  - **Transaction end**
- Changes are written out asynchronously to the data extents

# After-Imaging (AI) Basics

What happens to all these BI/AI notes?

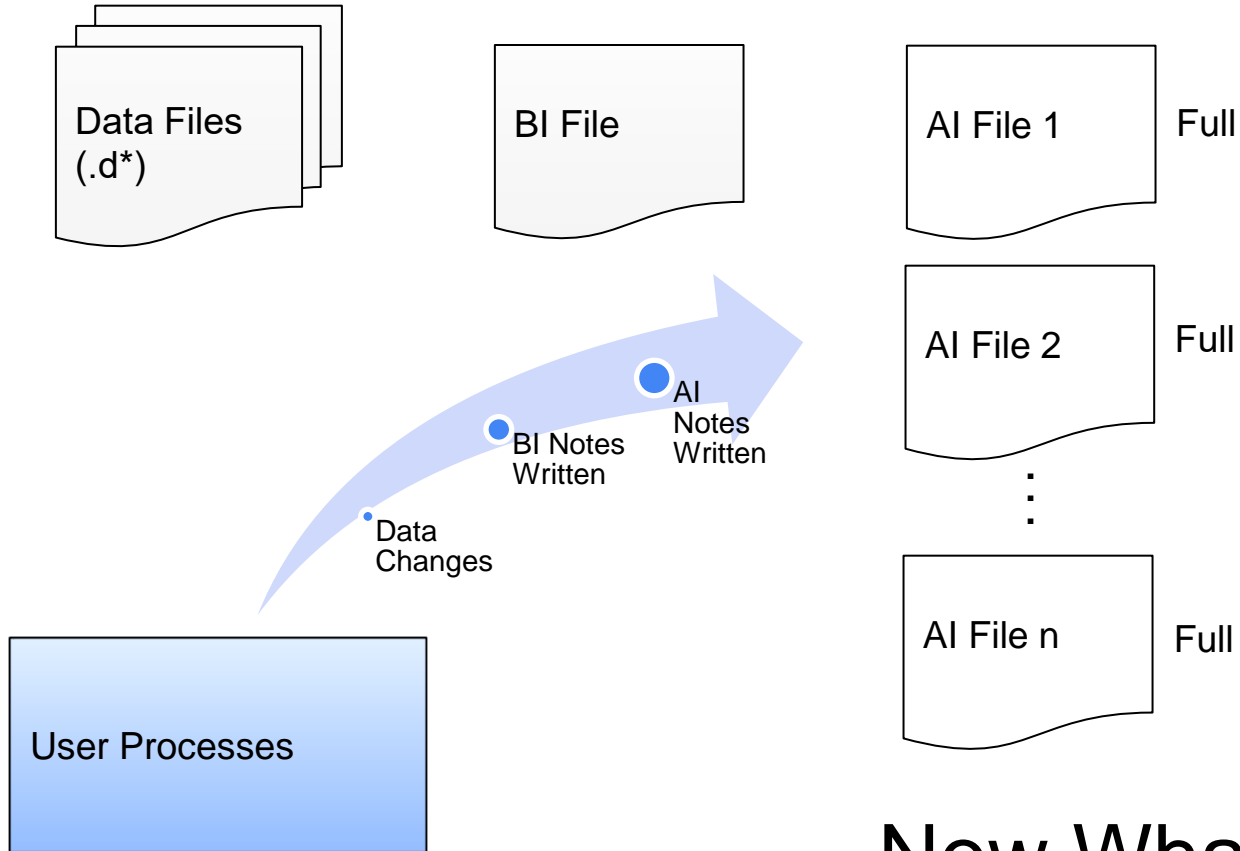
- OpenEdge automatically manages the re-use of BI file space
- AI files must be archived before allowing OpenEdge to reuse them

# After-Imaging (AI) Basics

AI extent statuses:


- Empty
  - Extent is ready to be written to
- Busy
  - The current extent where notes are written
- Locked
  - Specific to OE Replication
  - Full; notes not yet replicated to all target DBs
- Full
  - Previously-busy extent that has been closed

# After-Imaging (AI) Basics



Now What?

# After-Imaging (AI) Basics



# After-Imaging (AI) Basics

- Sounds simple enough:
  - When full, copy/archive AI files to a safe location
  - Mark them as empty so OpenEdge can reuse the space in the files



# Why Use AI



# Why Use AI?

- To ensure compliance with business-defined RPO and RTO
- In 2024, re-keying lost data from paper is no longer an option
- It's not just about the "big D" disaster
- Cost/benefit value of AI is clear

# Cold vs. Warm vs. Hot Spare DB PROTOP.COM

18

- Cold spare:
  - Restore DB from backup, on demand
  - Often from off-site data warehouse
- Warm spare:
  - AI log-shipping, continually updating
  - DB is hours to minutes behind prod (typically)
- Hot Spare:
  - Near real-time updating
  - OpenEdge Replication Plus

# AI vs. Infrastructure Redundancy PROTOP.COM

- If I use:
  - RAID 1/5/6/10
  - SAN mirroring
  - VM snapshots
  - File system replication
  - Veeam etc.
- Do I still need AI?
  - **Yes!** These technologies will happily and immediately replicate logical corruption or data loss to a replica.



# How To Configure AI

# Add AI Extents

- Can be done offline or online as of 10.1A
- Create add.st file

```
# Sixteen variable length AI areas in /ai/prod  
a /ai/prod  
<...>  
a /ai/prod
```

- Ideally, use a different volume/device from the database

- Apply to database

```
$ prostrct add sports add.st  
OR  
$ prostrct addonline sports add.st
```

# Enable AI, AIMD Offline

- Take a full backup of your database

```
$ probkup sports /backup/sports.probkp -com
```

- Enable after-imaging

```
$ rfutil sports -C aimage begin
```

- Enable AI File Management Daemon (AIMD)

```
$ rfutil sports -C aiarchiver enable
```

# Enable AI, AIMD Offline

- Start your database
  - Be sure to include the `-aiarcdir` and, if required, `-aiarcinterval` startup parameters

```
$ proserve sports -pf sports.pf  
                  -aiarcdir /aiarch,/aiarch2  
                  -aiarcinterval 900
```

```
$ proaiw sports
```

- `-aiarcdir` can be a comma-separated list of directories; relative or absolute path

# Enable AI, AIMD Online

- Enabling AI & the AIMD *offline* required three distinct steps
- As of 10.1B, do it *online* all in one command with probkup



# Enable AI, AIMD Online

```
$ probkup online sports /backup/sports.probkp  
  enableai  
  enableaiarchiver  
  -aiarcdir /aiarch,/aiarch2  
  -aiarcinterval 900
```

# Enable AIMD Online

- In 12.7+, it is possible to enable the AIMD online without taking a backup:

```
$ rfutil db -C aiarchiver enable -aiarcdir dir-list  
  [ -aiarcinterval n ] [ -aiarcdircreate ]
```

- Great benefit for very large databases
  - No long wait for enablement

# AI Extent Switching

- Extent switching will happen in one of five ways:
  - Time interval: AI file rotation occurs every x seconds
    - Min: 2 minutes
    - Max: 24 hours
  - Full AI File: Rotation occurs when the currently active AI file is filled
    - Used only with fixed-length AI extents
  - On demand
    - `rfutil db -C aimage new`
  - At the start of a backup (online or offline)
  - Proquiet enable

# AIMD Configuration

- Primary broker configuration parameters:
  - **-aiarcdir:**
    - Directory where full AI extents will be copied
    - Can specify multiple directories in case writes to the first fail
    - Do NOT specify a remote location, e.g. NFS, Samba
  - **-aiarcinterval:**
    - Time between AI file rotations (seconds)
    - If not specified, rotation occurs when extent fills
  - **-aiarcdircreate**
    - AIMD will create the -aiarcdir directories if they don't exist
    - We don't use this parameter

# Changing Parameters Online

- Both parameters (-aiarcdir and -aiarcinterval) can be changed online

```
$ rfutil db -C aiarchiver setdir /aiarch
```

```
$ rfutil db -C aiarchiver setinterval 1800
```

- The second option is particularly useful
  - Decrease interval during high activity
  - Increase interval during quiet times

# Stopping the AI Archiver Daemon

- To permanently disable the AI Archiver

```
$ rfutil sports -C aiarchiver disable
```

- To stop the daemon temporarily without disabling the functionality
  - i.e. for maintenance

```
$ rfutil sports -C aiarchiver end
```

# Starting the AI Archiver Daemon

- To restart the AI Archiver

```
$ _mprshut sports -C aimgt -aiarcdir <dir>  
-aiarcinterval <n sec>
```

- There is no “rfutil db -C aiarchiver begin”

# Starting the AI Archiver Daemon

- Note the similarity to other helper processes
  - APW, AIW, BIW and WDOG

```
root 2906 1 0 00:40 pts/1 00:00:00 _mprshut atm -C apw
root 2907 1 0 00:40 pts/1 00:00:00 _mprshut atm -C biw
root 2910 1 0 00:40 pts/1 00:00:00 _mprshut -db atm -C aimgt
root 2929 1 0 00:40 pts/0 00:00:00 _mprshut atm -C aiw
```



# AI Archiver Status

- **Activated?**

- `$ proutil <dbname> -C describe`

Database Features

ID	Feature	Active	Details
----	-----	-----	-----
8	After Image Mangement/Archiver	Yes	Timed 3600
9	64 Bit DBKEYS	Yes	
10	Large Keys	Yes	
11	64 Bit Sequences	Yes	



# How to Monitor AI

# Monitoring AI

- What to monitor
  - AI extent statuses
  - AIMD running?
  - AIW running?
  - AI archive directory statuses

# Monitoring AI Extent Statuses

- Number of extents in each status:
  - Empty, busy, locked, full
- Only one extent will be busy
- More than one locked or full extent warrants investigation
  - This may occur temporarily when the database is very busy
- Almost all extents should be empty
  - Give yourself lots of time to react to alerts

# Monitoring AI Extent Statuses

- `rfutil db -C aimage list`

...

Extent: 2

Status: Busy

Type: Variable Length

Path: /data/rob/db/12/sp/sp.a2

Size: 120

Used: 1

Start: Tue Sep 17 19:15:24 2024

LastOp: N/A

Seqno: 10

...

# Monitoring AI Extent Statuses

## On 11.7+ query the `_AreaStatus` VST:

```
for each _AreaStatus no-lock where _AreaStatus-Type = "AI":
display
  _AreaStatus-AreaName
  _AreaStatus-AreaNum
  _AreaStatus-AI-Seq
  _AreaStatus-State
  _AreaStatus-AIActivated.
end.
```

<u>AreaStatus-Areaname</u>	<u>AreaStatus-Areanum</u>	<u>AI</u>	<u>Seq</u>	<u>State</u>	<u>AI Activated</u>
After Image Area 1	12	0	Empty	?	
After Image Area 2	13	10	Busy	Tue Sep 17 19:15:24 2024	
After Image Area 3	14	0	Empty	?	
After Image Area 4	15	0	Empty	?	

# Monitoring AIMD & AIW

## Query the \_Connect VST:

```
for each _connect no-lock where
  _connect-type = 'AIMD' or _connect-type = 'AIW':
  display
    _connect-type
    _connect-time.
end.
```

```
Type Login time
```

---

```
AIMD Wed Sep 18 11:21:31 2024
AIW   Wed Sep 18 11:22:14 2024
```

Or use promon 1, 1 (User Control)  
or promon R&D, 1, 4, 7 (Background proc.)

# AIW & AI Archiver Status

promon R&D, 1, 4,7

10/08/24  
08:28:15

Status: Background Processes by user number

Usr:Ten	Name	Domain	Type	Start time
0	rob	0	BROK	10/07/24 13:32
1	rob	0	BIW	10/08/24 07:33
2	rob	0	BIM	10/08/24 07:34
3	rob	0	APW	10/08/24 07:34
4	rob	0	WDOG	10/08/24 07:34
5	rob	0	AIMD	10/08/24 08:18
6	rob	0	AIW	10/08/24 08:28
7	rob	0	MON	10/08/24 08:28



# Monitoring AIMD & AIW

- `ps -ef | grep db | egrep "aimgt|aiw"`

```
root      7837      1 /usr/dlc128/dlc/bin/_mprshut -db db -C aimgt
root      8252      1 /usr/dlc128/dlc/bin/_mprshut db -C aiw
```

- `proshut db -C list | grep " AIW "`

```
2      8252      8252      0      0 Wed Sep 18 11:22:14 2024      rob
AIW      /dev/pts/0      no
```

**Note: proshut does not list the AIMD process**

# Monitoring AI Archive Directory

- Monitor the database log file for AIMGT errors or info messages, e.g.:

(13621) The directory /data/rob/db/10/aiarcdir/a appears to have been removed. The After-image Extent Management will switch to the next directory specified in the destination list.

(13231) From this point forward all after-image extents will be archived to /data/rob/db/10/aiarcdir/b.

(13198) The after-image management daemon cannot find a destination in its current list with enough space to continue archiving extents.

(13187) The directory <name> specified with -aiarcdir does not exist.

# Why Monitoring Is Vital

- If the AIMD stops switching, or it cannot write to any archive directory...
  - You will eventually run out of empty extents
  - Your busy extent will grow (or fill, if fixed)
- If you run out of empty extents
  - You will be unable to switch extents
  - Online backups will fail!
  - Warm spare falls behind
  - Your DR plan is up in flames



# Recovering with AI

# Recovering with AI

- This process is called *rolling forward*
- Restore the database from backup
- Identify the first AI extent archived after the backup began
- Apply each extent, in order, to the database

# Recovering with AI

- Roll forward one extent:

```
rfutil db -C roll forward -a ai-extent-name
```

- Roll forward a list of extents (more efficient):

```
rfutil db -C roll forward -ailist ai-extent-listfile
```

- Other options

## To point in time:

```
rfutil db -C roll forward endtime yyyy:mm:dd:hh:mm:ss
```

## To a transaction ID:

```
rfutil db -C roll forward endtrans transaction-number
```



# AI Archive File Retention

# AI Archive File Retention

- How many archived AI files should you retain?
- This is a business decision: a trade-off between cost (disk space used), time to recover (media type/location) and recovery capability
  - You can't recover with files you don't have!
- (war story)





# AI: Useful Beyond DR

# AI: Useful Beyond DR

- Common scenario: refresh dev/test/UAT with up-to-date data
- Feed a reporting DB, when you don't have a replication process
- Forensic investigation
  - See who changed what, in detail, months ago  
`rfutil db -C aimage scan verbose -a ai-extent-name`
- Correcting data issues
  - E.g. user-deleted data, bad code

Questions?





**ProTop**<sup>®</sup>

**Thank you!**