Progress®

High Availability And Disaster Recovery in OpenEdge

Kevin Jones
Senior Database Administrator – Progress MDBA
March 2020

Who’s this Kevin Jones?

- Worked with Progress since early 1990’s (Has it been that long ?)
- Worked for Progress UK Professional Services (2000-2005)
- Spent a while consulting
- Now back with Progress
  - MDBA - Managed Database Services (2017)
- Looking after customers databases and environments worldwide
- Perform many database migration/moves

Agenda

- HA and DR – What it means
- HA options with OpenEdge
- DR options with OpenEdge
- OpenEdge 12 – HA (Getting to the five 9’s)
- OpenEdge 12 – Improved DR
- Upgrading to OpenEdge 12 - Plan to achieve minimum downtime

High Availability

Ability of a system to be continuously operational either forever or for a long length of time

- Measured in a % of time, Usually a % of a month / year
- Commonly Referred to a “Nines”

No of “Nines”	Availability	Monthly Downtime	Annual Downtime
1 “Nine”	90%	73 hours ~ 3 days	876 hours ~36.5 days
2 “Nines”	99%	7.3 hours	87.6 hours ~ 3.5 days
3 “Nines”	99.9%	~44 minutes	8.76 hours
4 “Nines”	99.99%	~4 minutes	~ 52 minutes
5 “Nines”	99.999%	~ 26 seconds	~ 5 minutes

Achieving High Availability

- Eliminate Single Points of Failure
  - Redundancy
    - Hot swap disks
    - SAN Mirroring
    - Virtual Machines running on a VM Server Farms
    - Host Machine Clustering
  - Load Balancing
    - Putting the redundancy to use
    - Multiple Network / Disk Interfaces
    - Splitting Client Processes over multiple servers
  - Check your processes. Things change so keep checking

Disaster Recovery

Process of returning your system, ie app & database to an operational state after a disaster

A disaster can happen at anytime , things like :

- Natural Events
  - Floods, Fire, Earthquakes
- Failures
  - Power, Storage, Network , Hardware, OS
- Mistakes
  - Operator Errors, Application bugs, 3rd Party Suppliers

Recovering from a disaster

- You need to document what you have and what you want to recover from
- You may need additional hardware
- You may need your software media and codes
- Build up you plan on how you will recover
- Configure any DR components
- Check it works. Things change to keep checking

Summary

- High Availability (HA)
  - How to keep your application and databases available , and running at all times
- Disaster Recovery (DR)
  - Bringing your application and databases back online after and unexpected or catastrophic event
  - Preferably in the shortest amount of time
  - Ideally with a little if not no data loss.

Progress®

HA options in OpenEdge

HA options in OpenEdge

- Applications consist of 2 components
  - Front End
    - GUI clients usually delivered as a published application
    - Load balance across multiple servers
  - Back End
    - Server(s)
    - AdminServer
    - Unified Brokers (Classic Appserver and Webspeed)
    - Progress Application Server for OpenEdge (PASOE)
    - Database(s)

High Availability - Server

Failover Clusters (active/passive nodes)

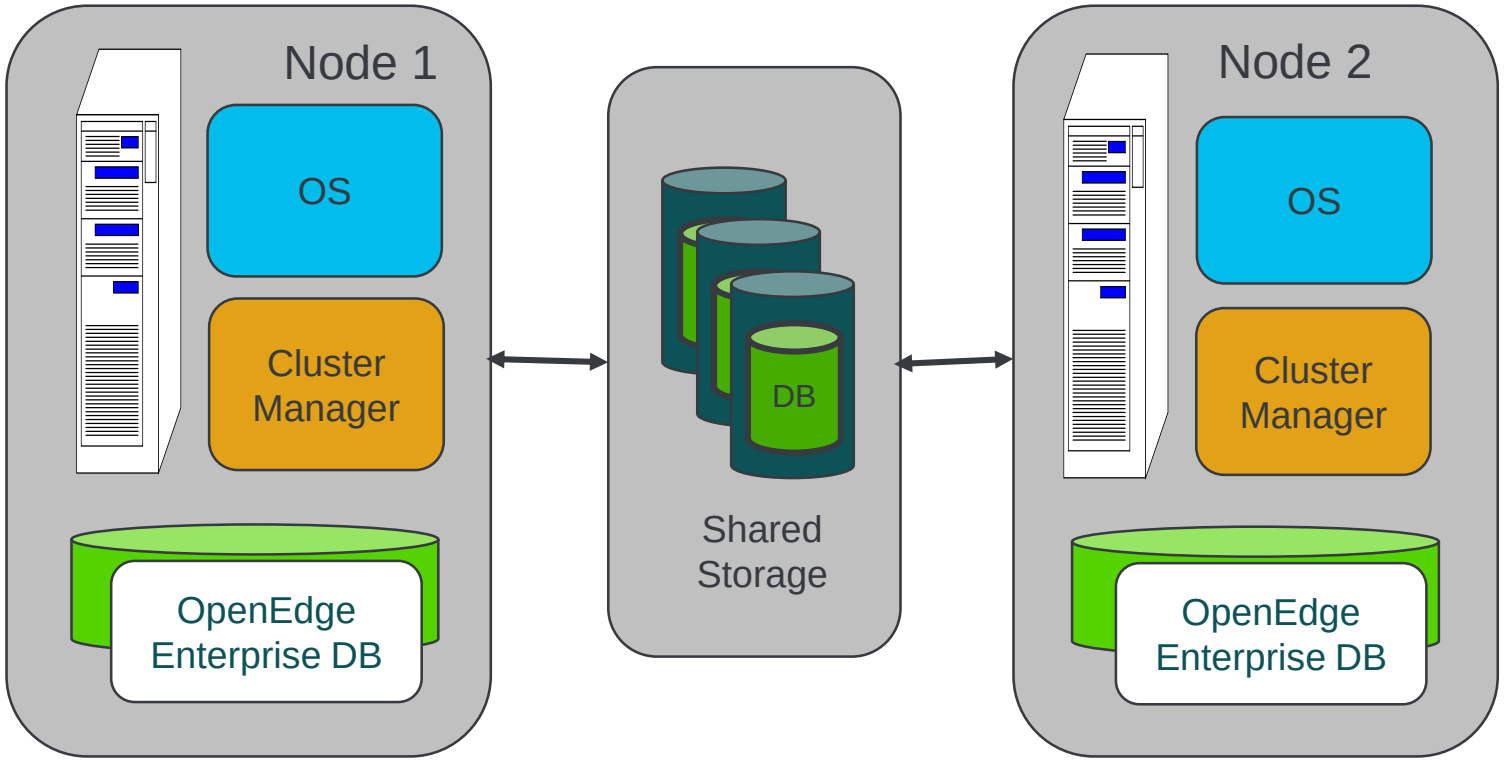

High Availability (Servers)

- Enables DB as cluster protected resource
  - Either manually or via procluster
- No active/active support
  - Still downtime
    - When any involved component fails
  - Users have to reconnect
    - No “client.pf” file change needed

High Availability (Server)

- Selected Cluster Managers are supported
  - Windows Cluster Failover
  - IBM HA Cluster (HACMP)
  - IBM PowerHA
  - HPUX Service Guard
- No support with Cluster Managers:
  - For Linux
  - For any from OpenEdge 12.2
  - But you can configure and script this yourselves

High Availability (Adminserver)

- Adminserver is a container for unified brokers and OEE/OEM
  - Nameserver
  - Appserver (Classic)
  - Webspeed brokers
  - OpenEdge Explorer/Management
- Is it a Java process
- If it stops then all the components running will terminate
- In OE11 the adminserver can stop and the other components remain running. New parameter -keepservers

High Availability (Appserver Classic and Webspeed)

- Appserver Webspeed brokers are a single point of failure
  - Multiple brokers can run the provide the same appservices
  - Need nameserver load balancing
  - Weights can be allocated to each broker
  - Nameserver allocated brokers in round robin based on weight

High Availability (Appserver Classic and Webspeed)

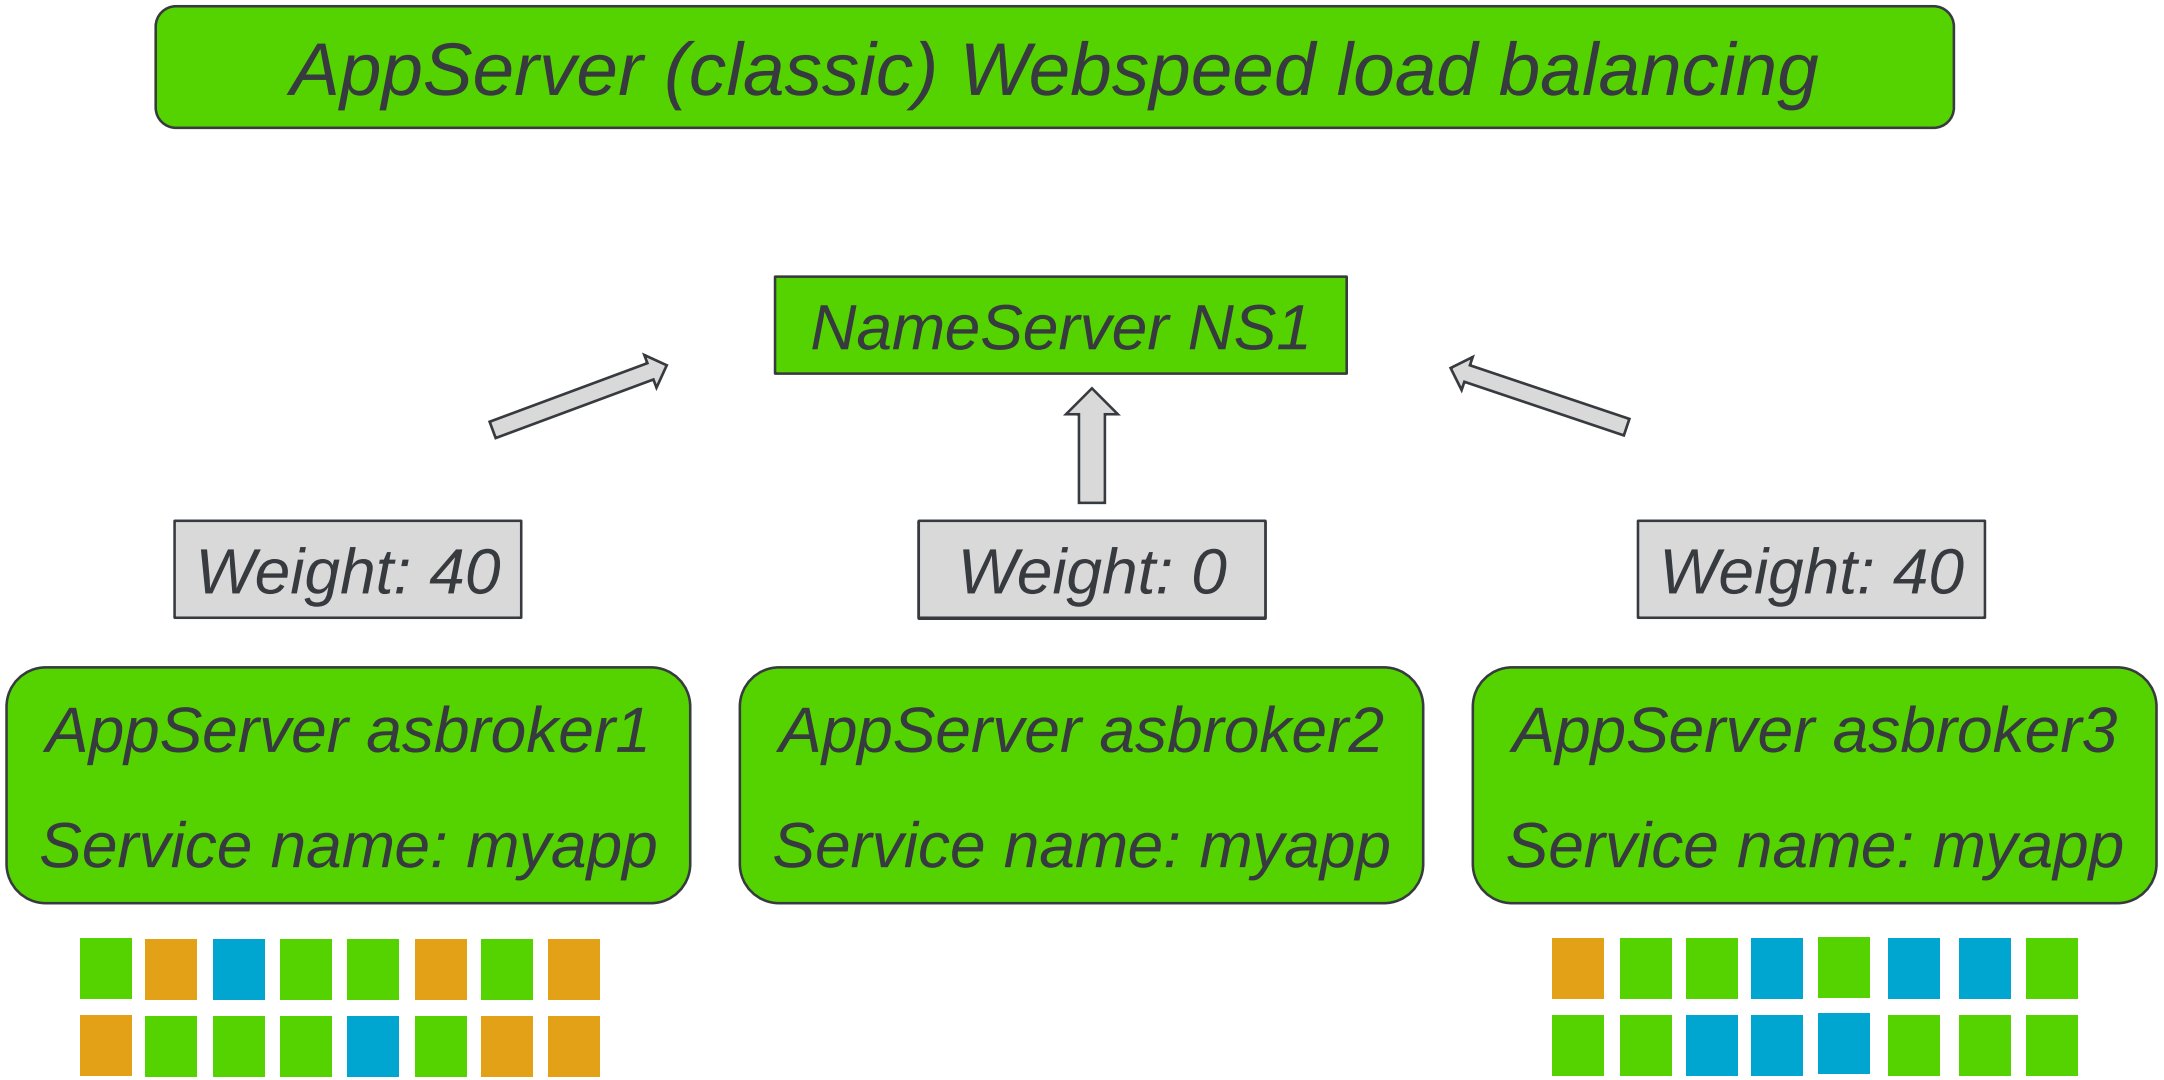

High Availability (Appserver Classic and Webspeed)

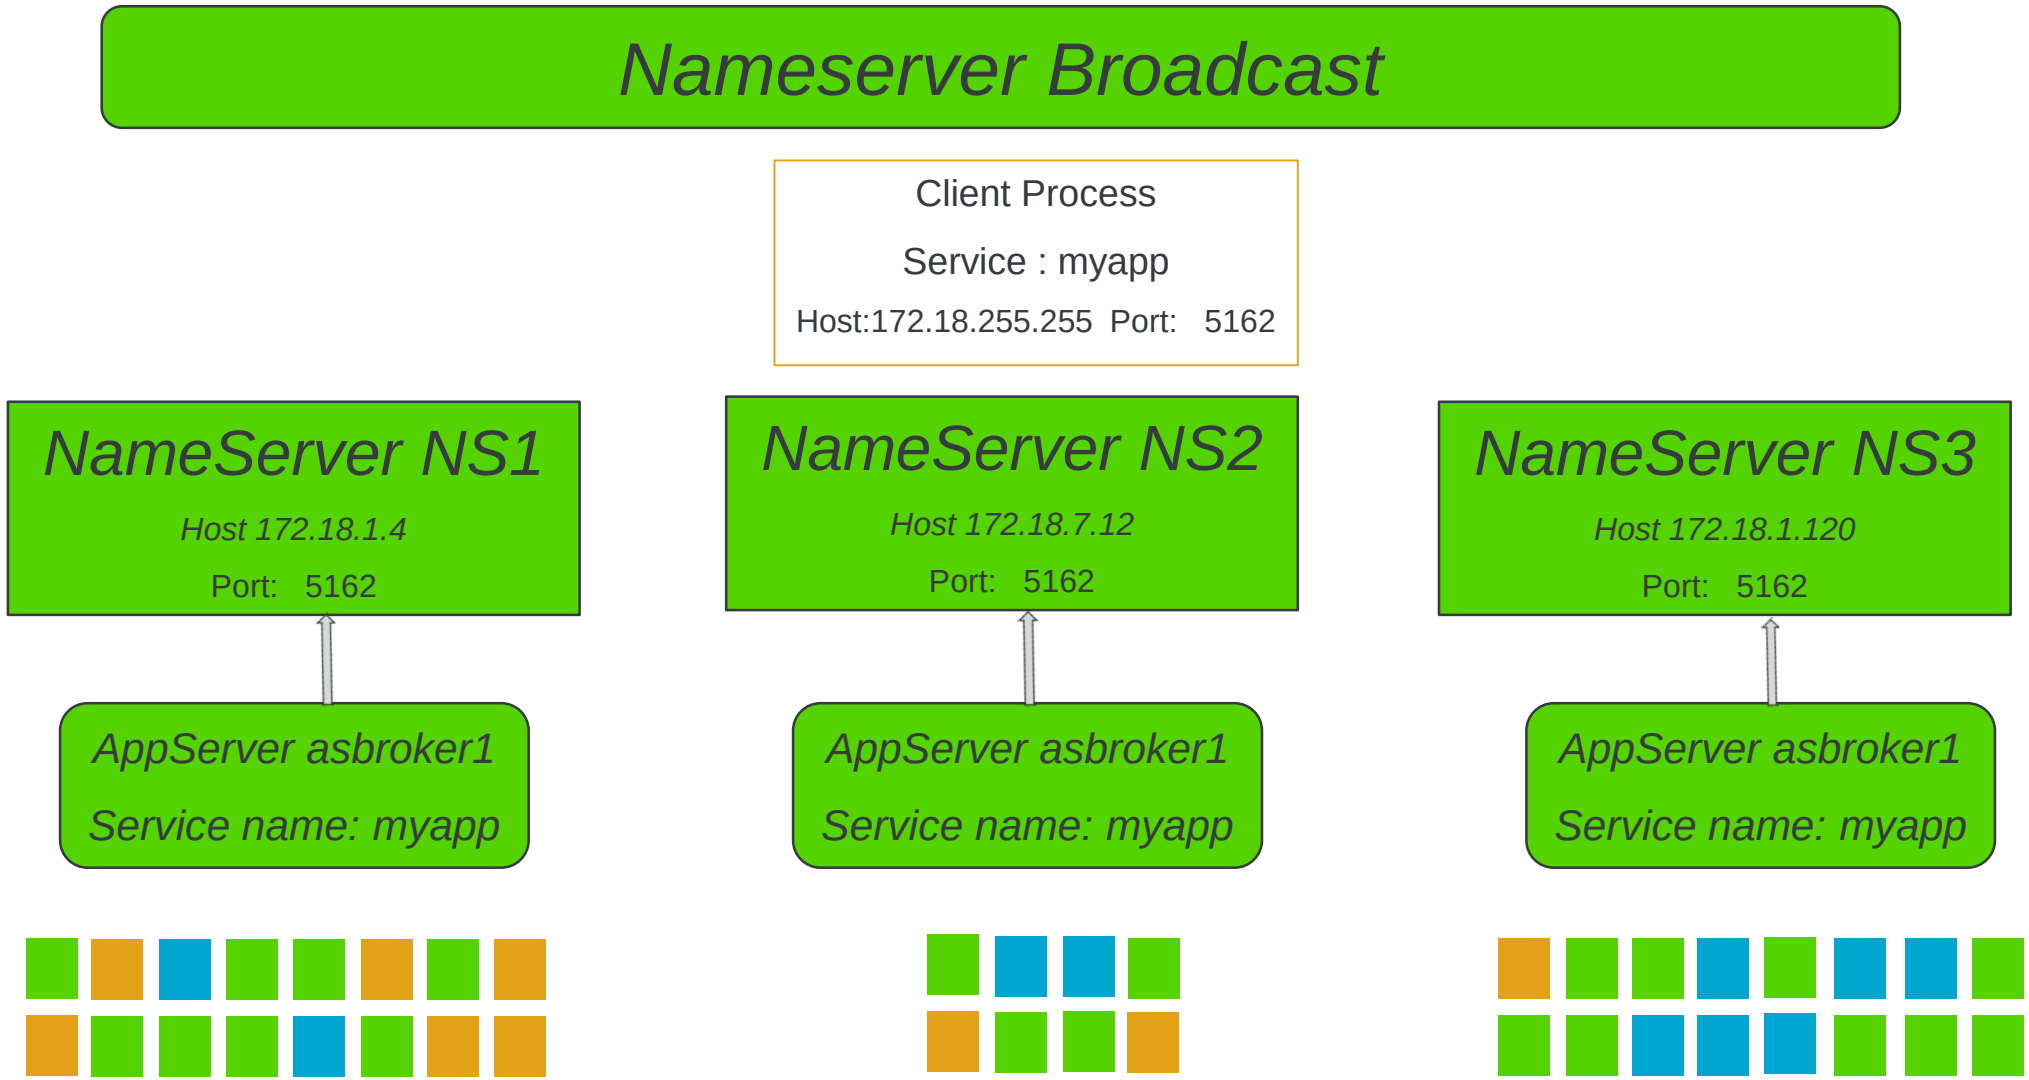

High Availability (PASOE)

- The only available Application Server from OpenEdge 12

- Tomcat based
  - Removed dependency on Progress Adminserver
  - Native Tomcat Load Balancing and Clustering
  - Use standard monitoring and security modules
  - Fronting Tomcat with Apache web server
    - Load balancing with F5

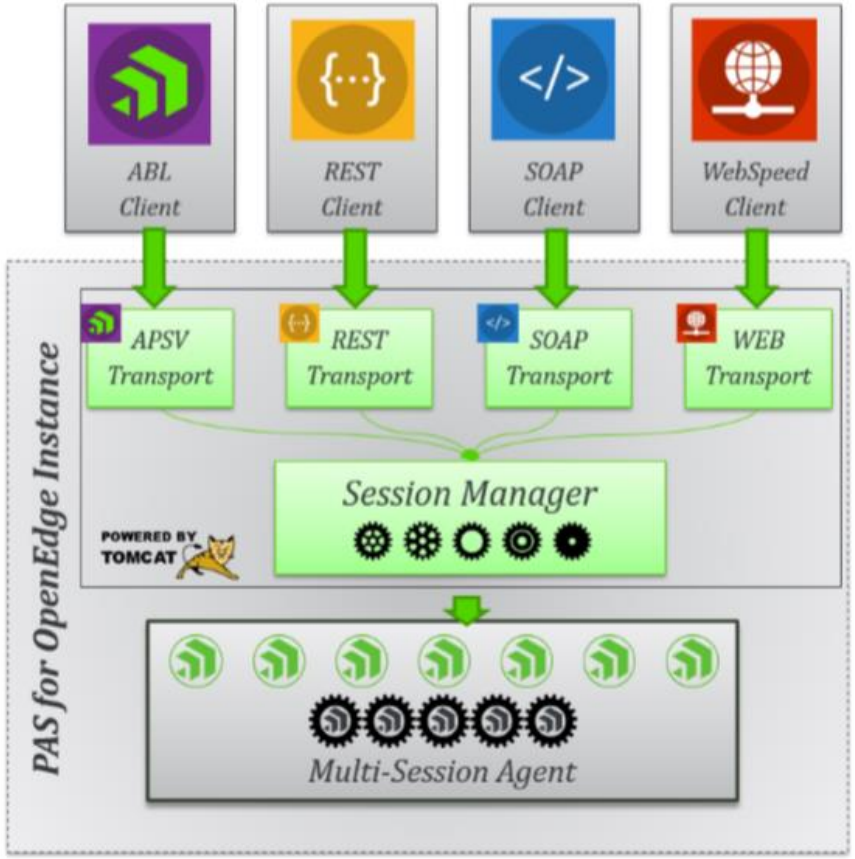

High Availability (PASOE)

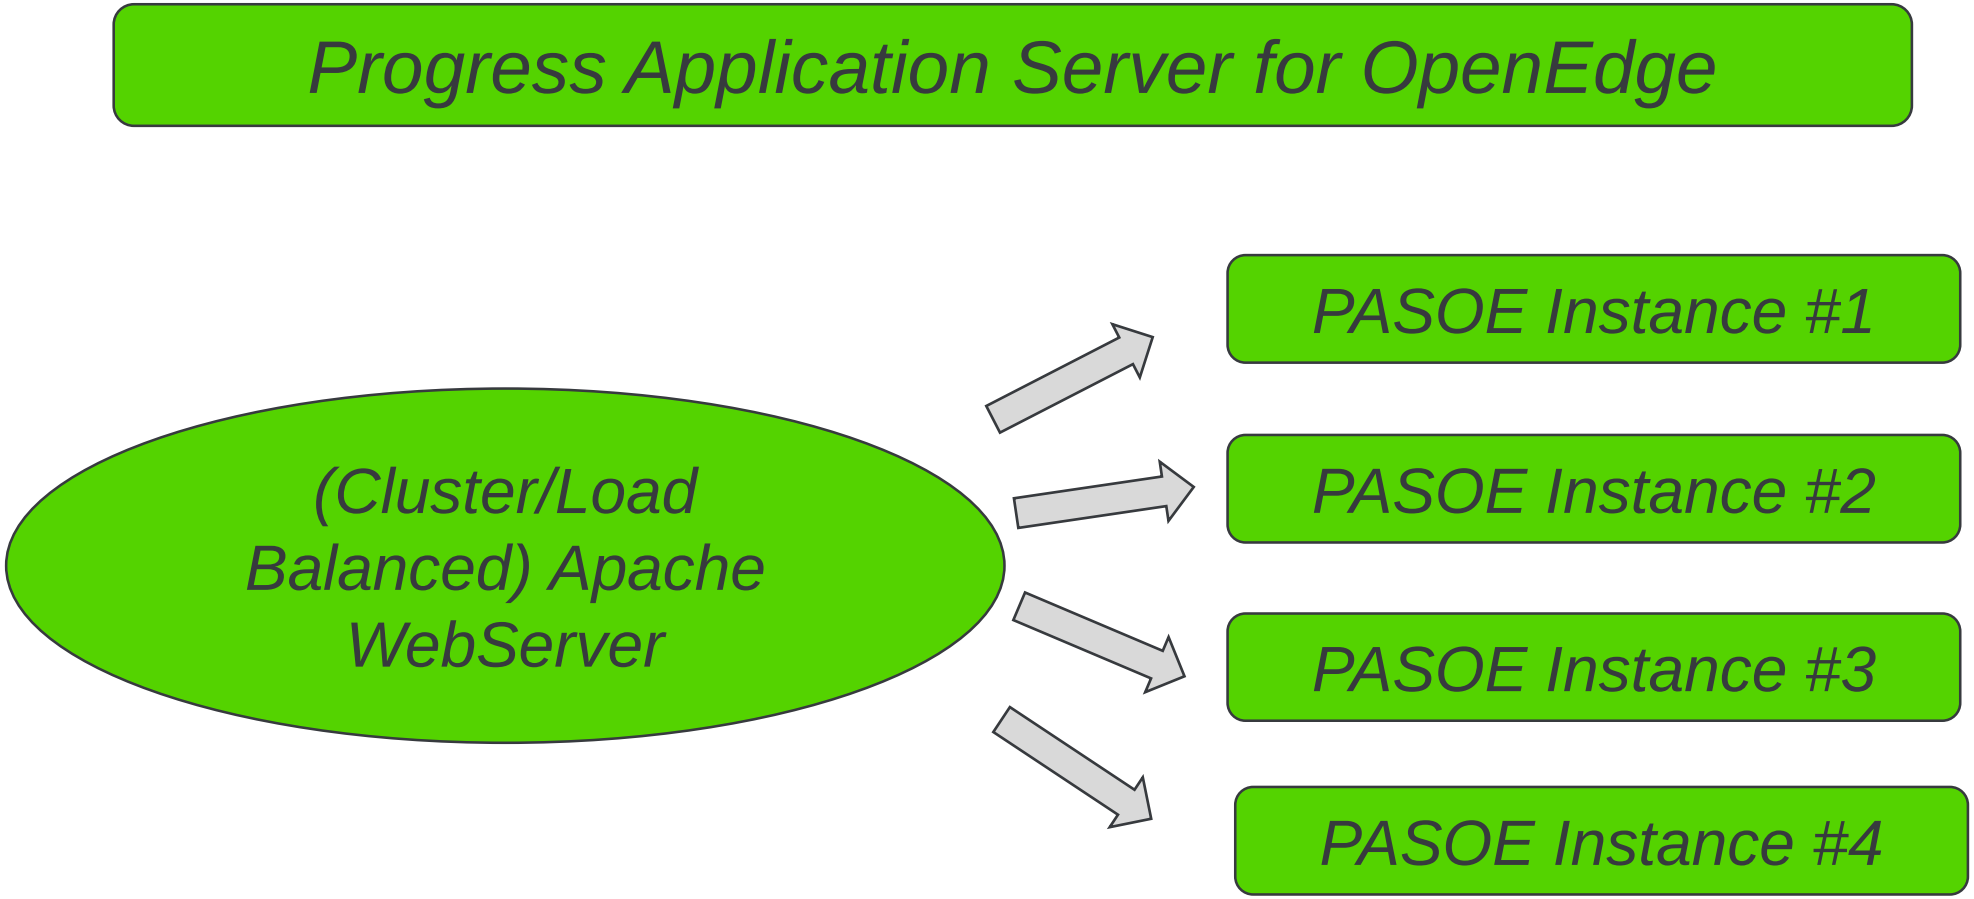

High Availability (Databases)

- OpenEdge DB does not have real HA
  - True – There is no active/active database, so in the event of a database failing we will have down time.
  - Even if there was the clients would have to reconnect to a different host. This is usually a startup parameter.

- We need to minimize
  - Possibility of the downtime
  - Length of the downtime
    - Can this be done fast enough without going to DR ?

High Availability (Databases)

Reasons for database downs or unavailable

- Truncate Before Image Files
- Service Pack Installs
- Dump and Load
- Database Broker Hangs
- All After Image Extents Are Full (and current AI extent is fixed)
- Full Filesystems
- Locktable overflow
- Schema Updates
- Increase parameters
- Add Database Extents
- Shared Memory Process crashes holding latch
- Accidental shutdown

High Availability (Databases - Monitoring)

Using VST's

- Lock table overflow
- Connection counts (-n) exceeded
- Fixed Extent Usage

Using OS commands

- After Image Extent Usage
- Disk space

High Availability (Databases – Monitoring VSTs)

Lock table overflow

Connection counts (-n) exceeded

```
find _startup no-lock .
find _dbstatus no-lock .
display _startup-LockTable
        _dbstatus-mostLocks
        _dbstatus-numLocks with title "Locks".
```

```
def var lv-server as int label "Active Servers" no-undo.
def var lv-users  as int label "Active Users" no-undo.

find _startup no-lock .
display _startup-maxusers _startup-maxservers.

for each _connect no-lock where  _connect-usr <> ? :
  if _connect-id = 1 then next.
  if _connect-id <= _startup-maxservers + 1 then
  do:
     lv-server = lv-server + 1.
     next.
  end.
  lv-users = lv-users + 1.
end.
display lv-users lv-server.
```

Locks

Current size of locking table (-L)	Lock table high water mark	Lock table entries in use
8192	3	0

Maximum number of users (-n)	Maximum number of servers (-Mn)	Active Users	Active Servers
21	4	4	0

_startup is replaced with _dbparams in OE12

High Availability (Databases – Monitoring VSTs)

Fixed Extent Usage

```
def var lv-highestextentinUse as int no-undo.
def var lv-%areaused          as dec format ">>9.99%" no-undo.
for each _area where _area-type = 6  and _area-num > 1 :
  find first _areastatus where _areastatus-Areanum = _area._area-number  no-lock .
  assign lv-highestextentinUse = int(substring(_areastatus-lastextent, r-index(_areastatus-lastextent,".") + 2 ))
       lv-%areaused  = _areastatus-hiwater / _areastatus-totblocks * 100 .
  display _area-num _area-name
       _areastatus-totblocks label "Total“   _areastatus-hiwater   label "HWM"
       lv-%areaused _areastatus-extents label "Extents” lv-highestextentinUse  label "InUse".
```

Area-number	Area-name	Total	HWM	lv-%areaused	Extents	InUse
6	Schema Area	622	401	64.47%	2	1
7	Employee	188	17	9.04%	4	1
8	Inventory	492	306	62.20%	4	2
9	Cust_Data	31292	42	0.13%	4	1
10	Cust_Index	6716	21	0.31%	4	1
11	Order	988	642	64.98%	4	2
12	Misc	188	26	13.83%	4	1
20	Type2	207	175	84.54%	1	1
250	Pro2RowID Area	13311	511	3.84%	1	1

High Availability

After Image Extent Usage

```
rfutil <dbname> -C aimage list
```

- Count number of Full
- Count number of Locked
- Count number of Empty
- Is current extent fixed length and next extent full?

```
Extent:  1
Status:  Busy
  Type:  Fixed Length
  Path:  C:\OpenEdge\WRK\11.6\sports2000.a1
  Size:  102400
  Used:  1
 Start:  Mon Feb 17 15:52:07 2020
 Seqno:  8

Extent:  2
Status:  Full
  Type:  Fixed Length
  Path:  C:\OpenEdge\WRK\11.6\sports2000.a2
  Size:  102400
  Used:  102400
 Start:  N/A
 Seqno:  6

Extent:  3
Status:  Full
  Type:  Fixed Length
  Path:  C:\OpenEdge\WRK\11.6\sports2000.a3
  Size:  102400
  Used:  102400
 Start:  N/A
 Seqno:  7
```

High Availability (Databases – OS Commands)

Disk Space

wmic logicaldisk get size,freespace,caption

```
proenv>wmic logicaldisk get size,freespace,caption
Caption  FreeSpace     Size
C:       103345971200  250874839040
D:       880901611520  1000068870144
E:       538471653376  1000068870144
```

df -h

```
-sh-4.1$ df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/mapper/vg_root-lv_root
                       40G   20G   19G  52% /
tmpfs                 127G  228K  127G   1% /dev/shm
/dev/sda2             488M   75M  389M  17% /boot
/dev/sda1             512M  304K  512M   1% /boot/efi
/dev/mapper/vg_root-lv_tmp
                       20G  541M   19G   3% /tmp
/dev/mapper/vg_root-lv_usr
                       20G   14G  4.9G  75% /usr
/dev/mapper/vg_root-lv_var
                       20G   11G  7.8G  59% /var
/dev/mapper/vg_data-lv_ai
                       99G   46G   48G  49% /ai
/dev/mapper/vg_data-lv_bi
                       99G  1.3G   93G   2% /bi
/dev/mapper/vg_data-lv_pmdb
                      7.9T  3.6T  4.0T  48% /db
/dev/mapper/vg_data-lv_oedev
                       50G   28G   20G  60% /oedev
/dev/mapper/vg_data-lv_logs
                      4.8G  199M  4.4G   5% /logs
/dev/mapper/vg_data-lv_promon
                       50G   41G  5.9G  88% /apps
/dev/mapper/vg_data-lv_bkp
                      3.0T  1.6T  1.3T  57% /backup
```

You may require downtime

- Schema Updates
- Increase parameters
- Add Database Extents
- Shared Memory Process crashes holding latch

- We can perform selected schema updates online
- proutil <dbname> -C increaseto for selected parameters
- We can add database extents online

High Availability

- Is achieved by being proactive
- Monitoring the conditions that cause outages in services
- Taking action before it gets to the outage
- Can depend on the version you are using OE 11 supports more than OE 10 and OE12 is even better
- Restrict access to functionality and scripts and log when they are run

Minimising downtime

- Our server has crashed
- But it is ok, we have rebooted and it has powered up
- Hardware checks are ok
- So why is it taking so long for my database to start up ?
- More importantly, how much longer ?

- Long running transactions
- Monitor using promon or _Trans VST

Progress®

DR options in OpenEdge

Option 1 – Absolutely Nothing

Option 2 – Backup

- Point in time backup. Could you business really recover from a backup potentially days old ?
- Using an OpenEdge Backup (probkup online)
- Using a Third Party Tool (Database must be offline or have its quiet point enabled)
- *https://knowledgebase.progress.com/articles/Article/P11836*

We do not support issues related to restoration of databases using backup files that were not generated by PROBKUP or taken online without a PROQUIET raised and entire database files included. Neither do we support integrity issues that may arise from 3rd Party Replication technologies. The only supported solutions are OpenEdge Replication and PRO2.

Option 3 Backup and AI Roll forward

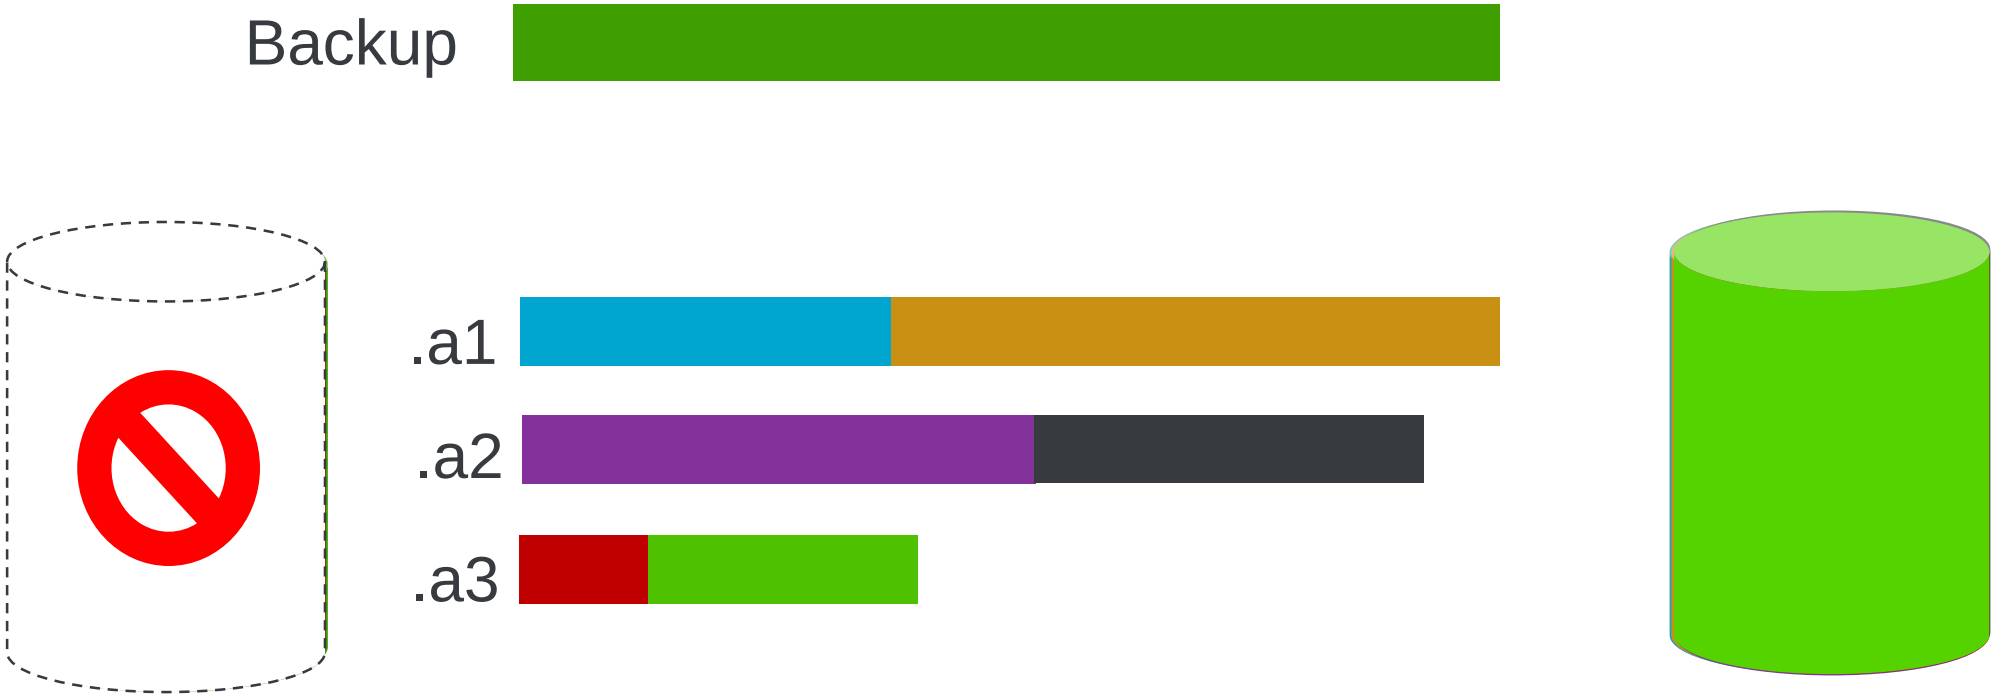

- Use the AI Archiver to manage your AI files

Option 4 OpenEdge Replication (OER)

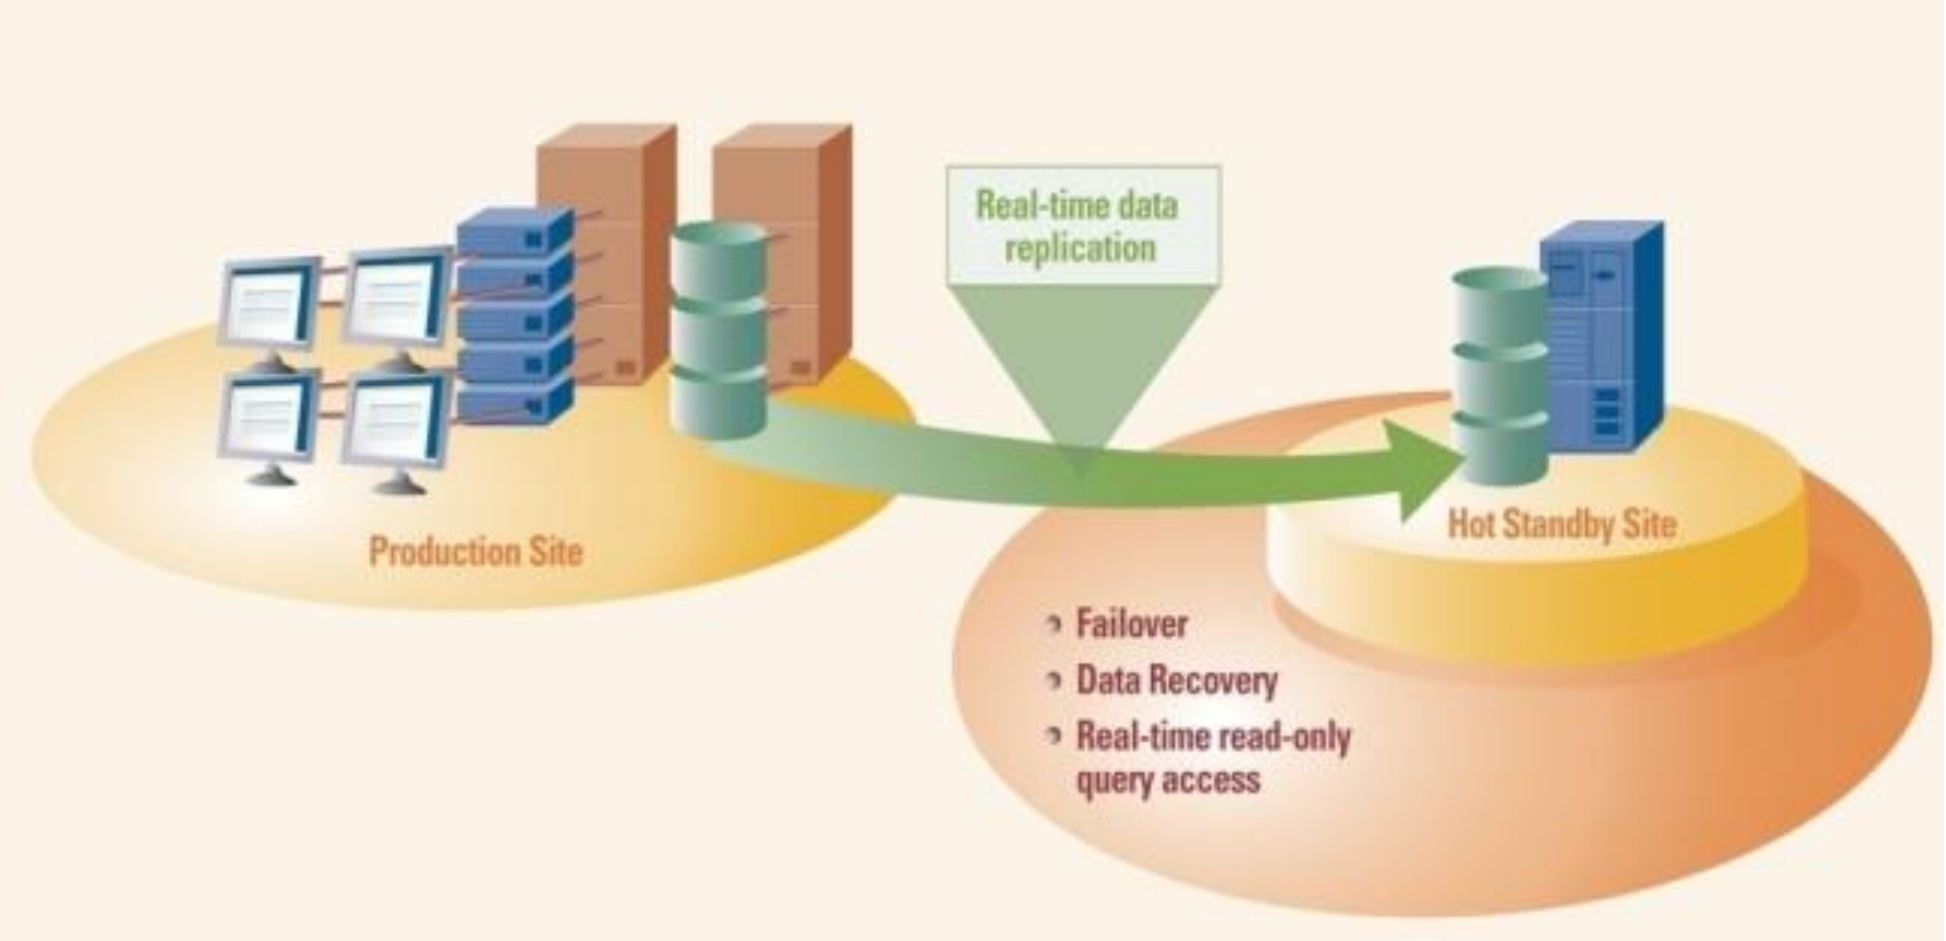

Option 4 OpenEdge Replication (OER)

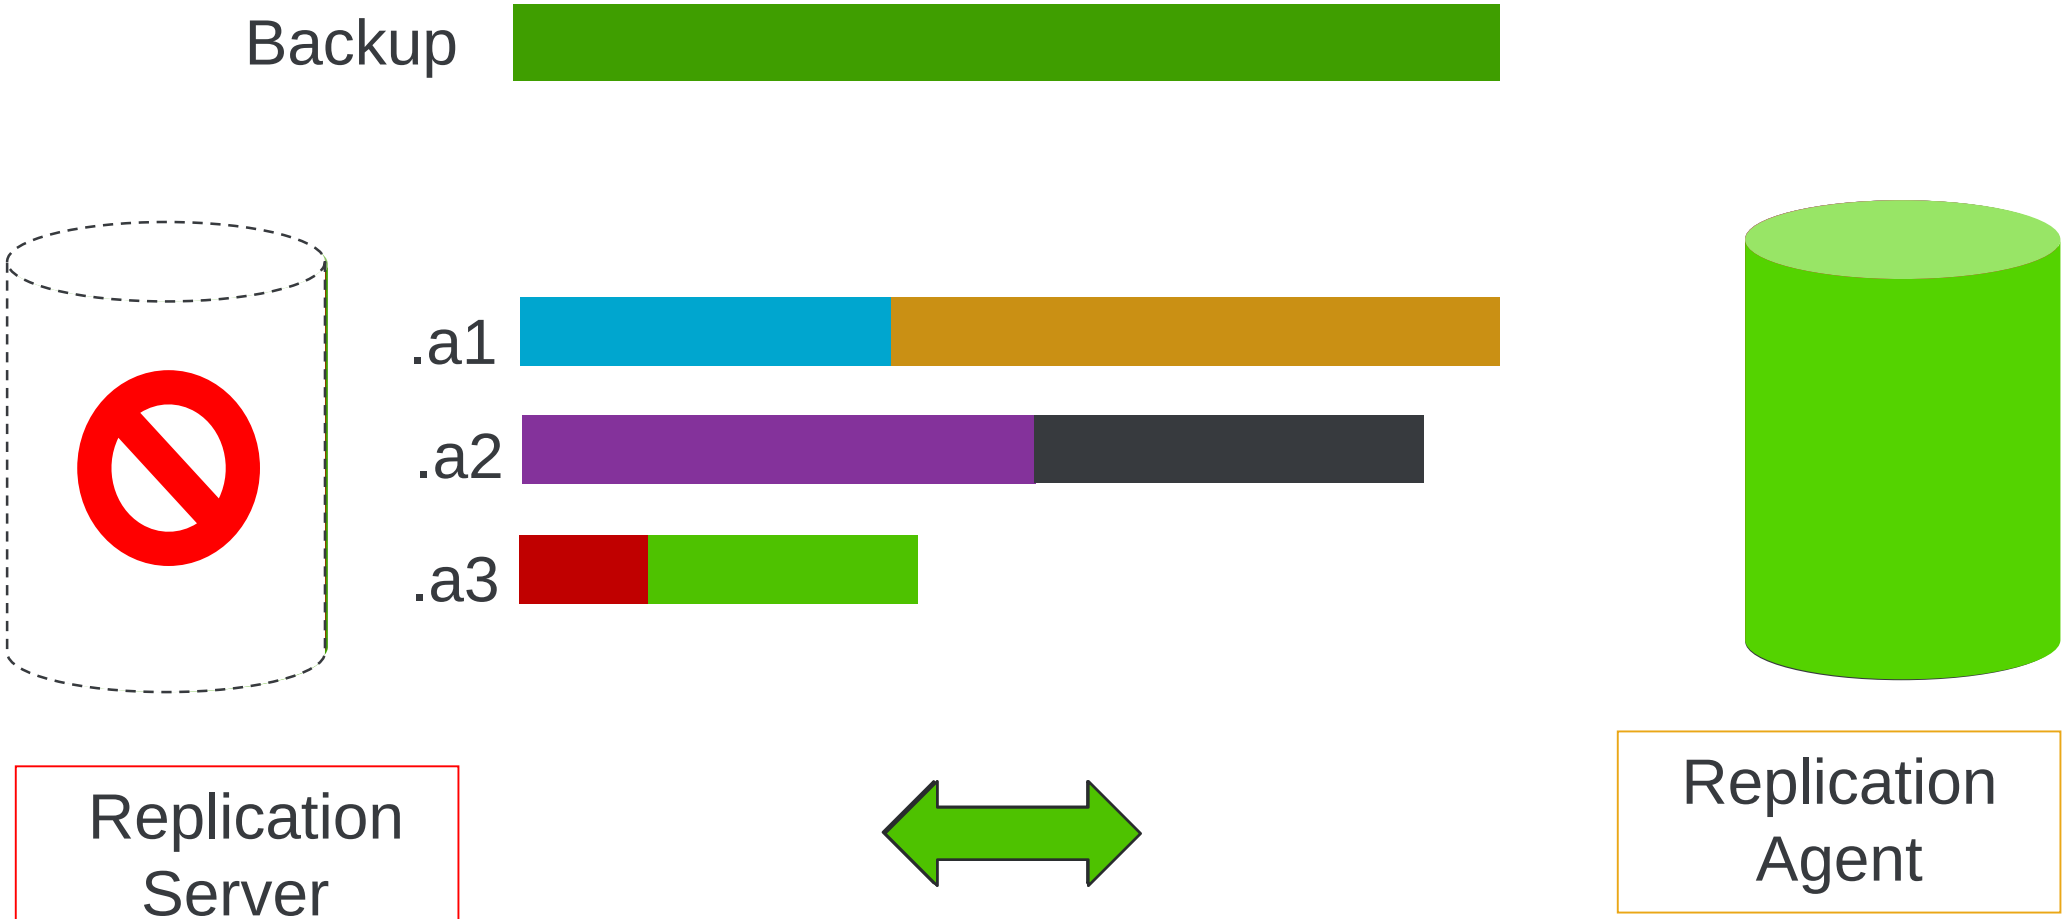

Option 4 OpenEdge Replication (OER)

OpenEdge Replication

- Failover
- Failback
- Replication Sets
  - Available from OpenEdge 11.7
  - Multi target environment
  - Allows to define “next” source after failover
    - Automatic role switch

Database covered anything else ?

- Application code and dependent files
- OpenEdge folder (especially properties)
- PASOE config
- Users
- Cron jobs / Schedule Task
- Printers …

Anything your system requires to run has to be in sync between Live and DR box or boxes

Progress®

OpenEdge 12
Getting to the 5 Nines

Availability

Progress®

The cause downtime

- Schema Updates
- Increase parameters
- Add Database Extents
- Shared Memory Process crashes holding latch

OpenEdge 12.0 – Availability

- Schema changes:
  - Non Structural field changes can be done online
  - Integers can be changed to int64 online
  - SQL Drop Index, Column or Trigger
  - SQL Rename Index or Column
  - SQL – JVM management and configuration

OpenEdge 12.0 – Availability

- Database File Maintenance:
  - Convert variable-length Extents to a fixed length
  - Database log file archive and truncate (OE 11.7.3)

OpenEdge 12.0 – Availability

Parameter updates

- A VST _dbparams was introduced in OE11. It replaces _startup which has been removed in OE12.
- Many parameters can not be modified (_dbparams-is-modifiable) online in real time by changing the _dbparams-value field in this table.
- The number of modifiable fields has increased from 42 in OE 11 to 85 in OE12.
- Other parameters can be increased using proutil –C increase to

OpenEdge 12.0 – Client Server Performance

Client Server Performance helps make applications distributed and scalable. It also means that connections cannot die holding share memory locks.

- BHT improvements
- Multi Threaded Remote Server Processes
- Server Side Query Resolution

Progress®

OpenEdge 12
Improved DR

Current AI Replication: Pica Queue

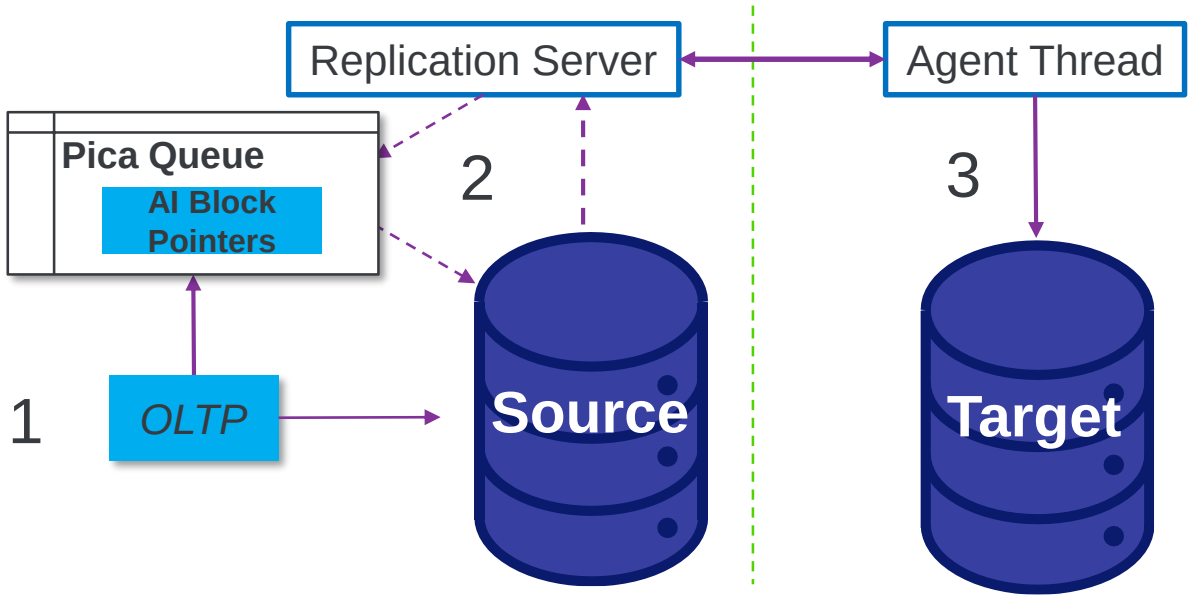

pica: Plugin Communication Area

RSB: Replication Streamed Block

AI Streaming Replication: RSB Cache / Queue

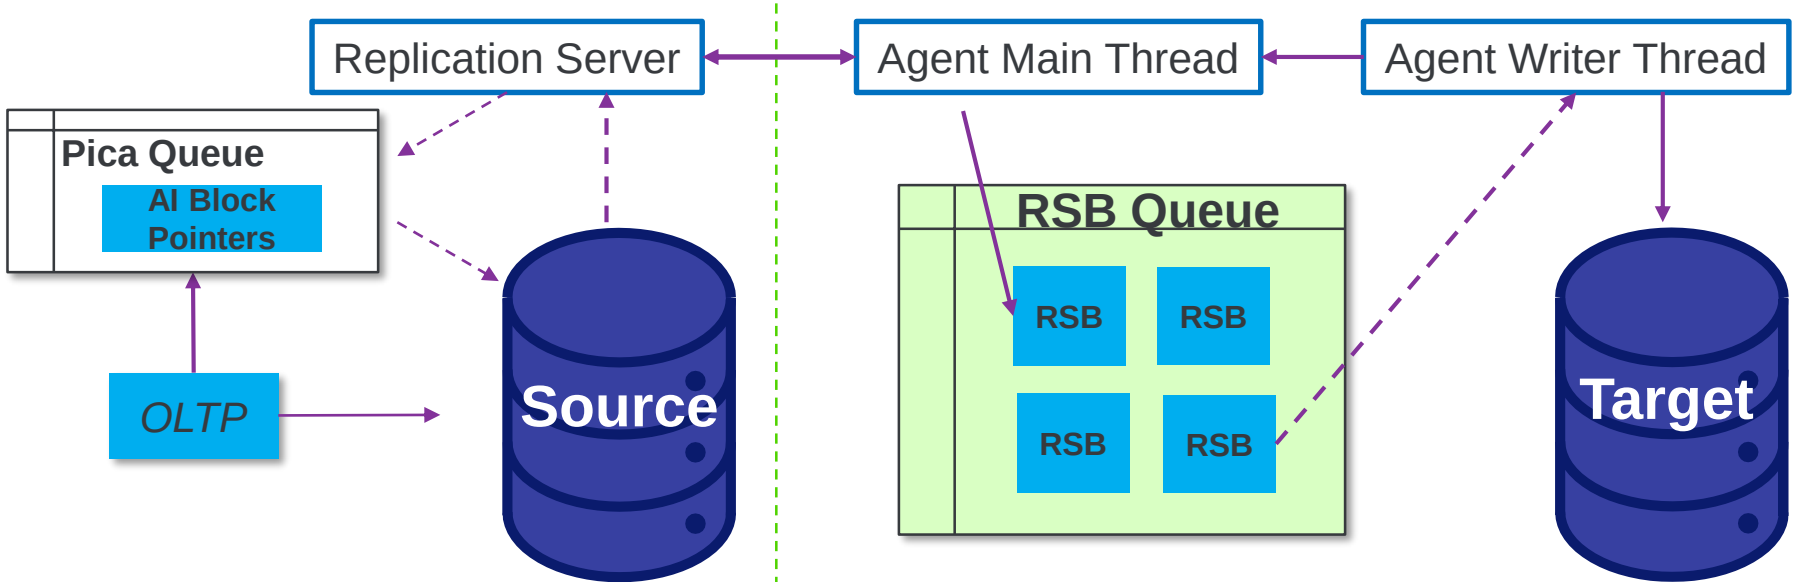

pica: Plugin Communication Area

RSB: Replication Streamed Block

Replication Latency by Time

Source Machine

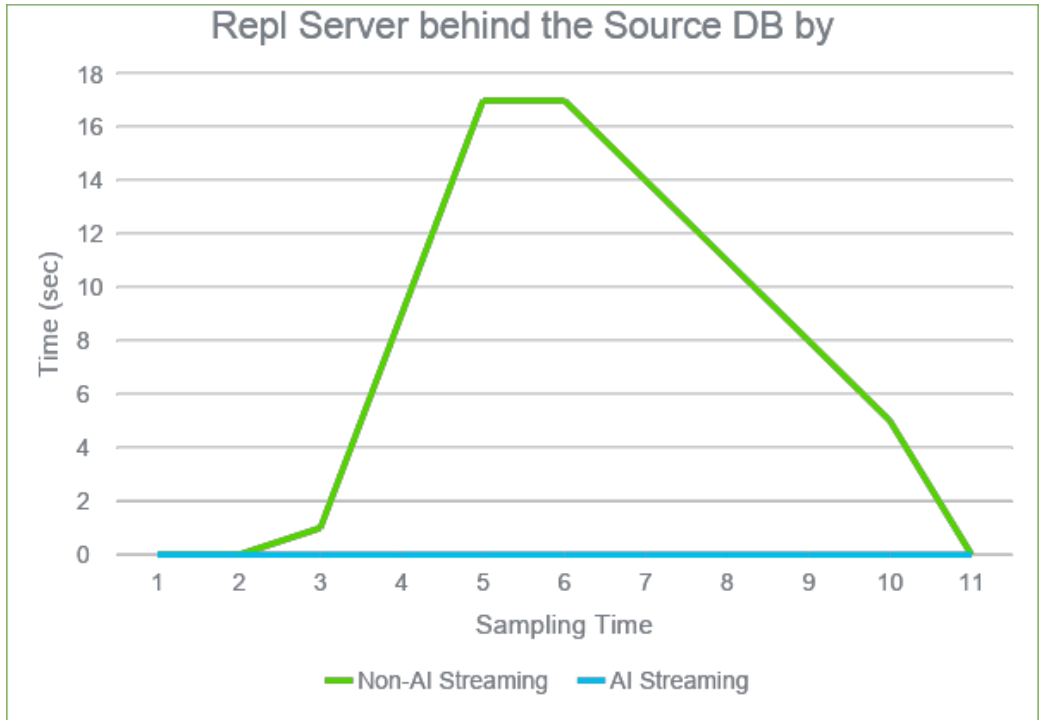

Target Machine

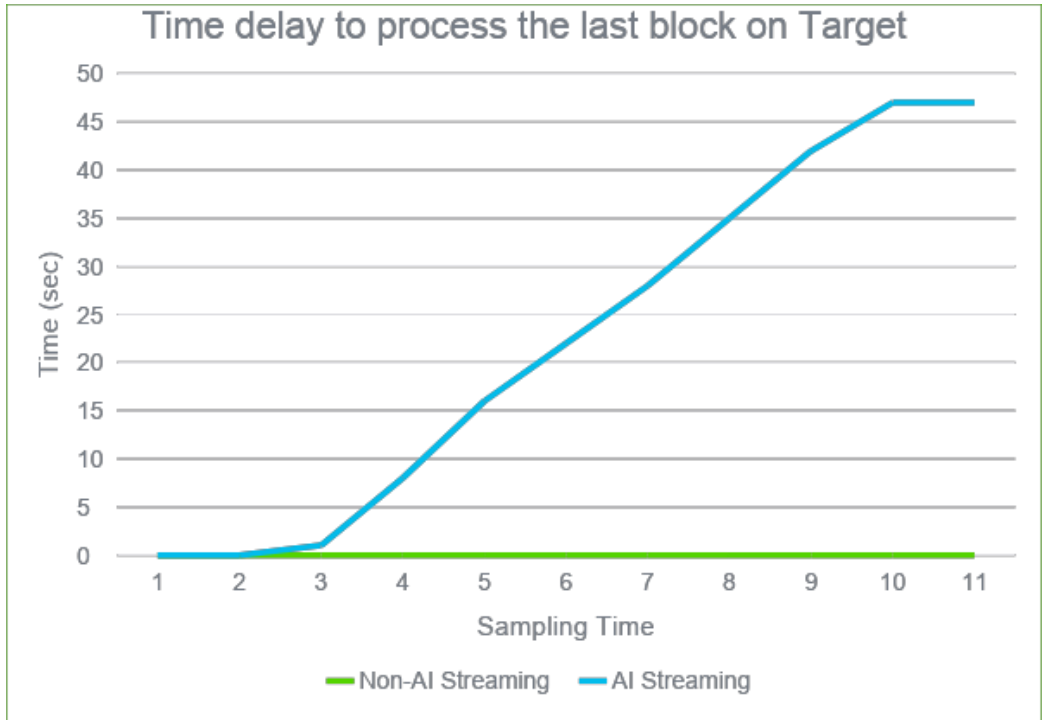

Progress®

Upgrading to OpenEdge 12
Plan to achieve minimum downtime

OpenEdge 12.0 – Upgrading

Knowledge is everything

- https://docs.progress.com/bundle/openedge-upgrade/page/Overview-of-Upgrading-to-OpenEdge-12.html
- There is a conv1112, it does not take long to run
- If you are running TDE then RC4-encryption is not supported in OE12. You must change the encryption and ensure the object(s) are fully encrypted with the new algorithm
- If you can prepare the database for conversion using the older version, disable replication, end after imaging, disable 2phase commit, truncate BI. However there are OE10 and OE11 DBTools with OE12.
- Large file support is automatically enabled/disabled in OE12. It will fail to disable if there are extents larger than 2GB
- PASOE is the only option , no classic Appserver/Webspeed. This is sorted out in development, but does require a new deployment of the application
- Only bcrypt encryption is supported in OE12 PASOE for encrypted passwords, so if you are using anything else you must change it before as bcrypt is supported in OE11.7 too.
- Fathom/OE Explorer runs as a separate service. It is a PASOE instance , and will not stop with the Adminserver


Progress®