

Platform Migration Mastery

A real-world guide to successful platform transitions



White Star
software

White Star Software

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

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

info@wss.com | wss.com



The speaker

Paul Koufalis (pk@wss.com)

- Speaker/teacher/OpenEdge DBA since 1994
- I have lost count of the number of migrations that I have personally worked on over the past 30 years



Agenda

- What is your current state?
- Where do you need (want) to go?
- How do you get there?
- Go live
- Q&A

Bonus: Top 10 hard lessons that we learned the hard way so that you don't have to.

Before we start...

Use a risk-based approach

- If you have a 10 user system with a 1 GB database, much of this content is overkill
- If it doesn't matter (too much) if you are not quite finished Monday morning, or if you can reboot if you need to...then again...overkill
- Think of this a bit like a menu
 - The more complex and critical the application, the more you should pay attention !!

Current state: what do we have to work with?

Take inventory:

- Infra: Hardware, operating system, network
- OpenEdge: version, components,
- Application: code, **interfaces**
- Databases: size, structure, current state
- 3rd party: **non-OE *stuff***

Infrastructure

- What everyone does: buys a BIG SERVER
 - Whatever the salesperson tries to sell to them
 - Often WRONG. Sometimes WORSE than what you have
 - Ask me about NUMA pitfalls during the Q&A
- Trend your current usage over a few weeks or months
 - CPU, memory*, disk*, network: peak vs *normal* usage
 - Longer if your business is growing fast**

Network and firewall

- Inventory all firewall and router ACLs that involve these servers
- Inventory all related IP addresses
 - **CAREFUL:** servers often have multiple IP addresses

OpenEdge

- What version of OE are you running now?
 - 32 or 64 bit? On the client? On the server?
- Draw out your architecture diagram
 - What products are installed and running, and where
 - **MOST OF YOU HAVE NEVER DONE THIS !!**
- Check your licence usage
 - We often see mistakes that are accidental but could be costly
 - Contact your local Progress representative for budget estimates

Application code

- Where is it installed?
- How is it deployed?
- Identify all the business critical workflows
 - Order, invoice, picking, shipping, customer service...
 - **The business users will have to test these end-to-end**

Database

- How big are the DBs?
- Are the structures optimal (type 2, etc) ?
- Any problems?
 - Fragmentation? Long RM chain? Corruption?
- Any special objects?
 - Very big tables or indexes
 - LOBs
 - RECID fields
- Don't forget users, sequence values, SQL, authentication domains, auditing, change data capture, security and related DB options, codepage...*

Third party tools

This is where the real trouble is!

- PDF generation
- Custom DLLs (whether Windows or UNIX)
- OCX and similar on Windows
- Email
- FTP/SFTP/SCP
- EDI
- ESB
- Telnet vs SSH vs xterm vs wyse60

STOP HERE

This is great.

- For the first time in 10 years, you have documented your entire OpenEdge environment!!
 - Ok - 99%. I'm certain that you forgot something...
- Now that you know where you are, you can decide where you want to go.

Target environment

- Which decisions are made by someone else?
 - On premises to Azure/AWS/Google/IBM/other Cloud
 - Hyper-V or ESXi
 - HPUX / AIX (both dropped by Progress) to Linux
- Which decisions are up to you?
 - Server specs, whether physical, virtual or cloud
 - Cloud provider
 - Operating system version

Target environment

- How much do you want to change?
 - Nothing: lift and shift
 - ISO8859-1 to UTF-8
 - root to non-root
 - Windows to Linux
 - Monolithic to distributed
 - Replication, CDC, Pro2...

On premises vs cloud

- A big step to move to “the cloud”
- Most of our customers are moving at least some of their OE components to the cloud
 - Example: disaster recovery servers
- Mostly to AWS

Virtualization

- Most of you are probably already virtualized
 - Who is using Hyper-V? ESXi?
 - Who is using bare metal?
- See my presentation from EMEA PUG Challenge 2022:
 - Benchmarking Bonanza: Cloud Edition
 - Compares bare metal to virtualization to Azure and AWS

Recommendations



Hardware - CPUs

- Get the fastest CPUs you can afford, NOT THE MOST CORES PER CPU !!!!
- Watch out for NUMA
 - You can only use one CPU per VM for database-related operations
- Watch out for vCPU counts vs core counts
 - ESXi typically configures one vCPU per hyperthread = 2 X #cores
 - Especially if you have core-based licenses
 - AIX could be SMT4 or SMT8

Hardware: RAM

- Most of you are not using most of the RAM in your current server
- Look at memAvailable in /proc/meminfo
 - numperm% on AIX
 - Sorry - I don't know the equivalent on Windows

BUT BUT BUT...

- Most of you under-allocate RAM to client processes
 - AppServer/Webspeed/PASOE/batch

Hardware: Disk

- I won't rehash RAID 5 arguments here
- An “all-flash SAN” is still a SAN
 - Flash disks are in a galaxy far, far away...
 - Will surely be good, but never great
- Remember

There ain't no such thing as a high performance SAN

Operating system

- **Make your life easy:** For your version of OpenEdge, use the latest supported Windows Server or Red Hat Enterprise Linux
- AIX, Solaris and HPUX are dead
- Other Linux are supported, but 99% of our Linux clients use RHEL

What does this mean in the cloud?

- Look for the latest generation of instance types
 - AWS C/M/R4: Intel(R) Xeon(R) CPU E5-2686 v4 @ 2.30GHz
 - AWS C/M/R5: Intel(R) Xeon(R) Platinum 8259CL CPU @ 2.50GHz
 - AWS C/M/R6i: Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90GHz
- The price is typically **less** on newer, faster types
- **Careful:** reserved instances and savings plans

What does this mean in the cloud?

- RAM vs CPU:
 - You can find many instance types that vary RAM vs CPU
 - R is for memory optimized
 - C is for compute optimized
 - and many more types...
- DISK: You almost certainly want GP3 disks
 - See *Benchmarking Bonanza* for Azure and other details
 - Ask me about ephemeral disks: Exciting !!

Planning your migration: what to change?

- Ideally, the least amount possible
- Remember the first slide: risk-based approach
- Things I am comfortable doing as part of a migration:
 - root to non-root
 - Adding OE Replication, CDC, Pro2...though you can enable these a couple of days later while production is running
 - ISO8859-1 to UTF - but needs a lot of testing

Planning your migration: timeline

- **MOST IMPORTANT INFO: How long do we have?**
 - Don't forget to take into account rollback
 - **Downtime window length is inversely proportional to the complexity of the migration**
 - The smaller the window, the more complex the migration
- **Define milestone goals with time markers**
 - This will make management happy
 - Will aid in GO/NO-GO decisions
- **Define success**
 - Count before-and-after table rows
 - Validated X, Y and Z business processes (ex.: order...ship...invoice)
 - Before-and-after financial reports
 - What else?

Planning your migration: components

- Inventory what data is dynamic vs static
 - Code is static: you can code freeze a week before go-live
 - Database data is dynamic
- rsync anything and everything that you can
- Create users, printers, etc. beforehand and freeze one week or more in advance
- Decide how network changes will be implemented
 - DNS entries
 - IP address re-assignments
 - Host aliases

Planning your migration: database

- Typically the longest part of the migration
- Identify your long pole tables
 - One HUGE table...especially in a type 1 storage area
 - Uploading data to the new server
 - You may need to come up with unconventional solutions!!
- I have done MANY dump&load presentations over the past 30 years
 - Email me for copies and check out <https://blog.wss.com>

Inventory what you need for go-live

- Temporarily need multiple servers
 - Example: backup and restore copies of DB to TEST and QA
 - Roll-forward AI files at GO-LIVE
 - Dump from three servers: TEST, QA and PROD
- Temporarily allocate more RAM
 - Accelerate idxbuild on new server
- You almost certainly will need extra disk space
 - Pre-allocate one week in advance to test performance
- Make sure that no 3rd party is hogging resources
 - Ex.: 100 TB Oracle data warehouse backup

Test, test and test some more...

- DBA: repeat the D&L process until it fits **very** comfortably into the allocated time
 - With room for surprises
- DEV: risk-based testing
 - Lift-and-shift to cloud might require minimal testing
 - OE version upgrade, AIX/HPUX to Linux, AppServer to PASOE...all require extensive testing
- USERS: test business processes
 - Yes I know they are busy...
- LOAD TESTING
 - Make sure to test a full typical workload, otherwise you will only find errors after go live

Document EVERYTHING

- GO LIVE IS NOT THE TIME TO THINK. IT IS THE TIME TO DO.
- Every step should be fully documented, including the server, the effective user, and the command line

GO-LIVE

- This should be the easiest part: follow your written plan
 - All the deep thinking happened in the weeks and months before
 - There will be surprises. Save your brain cells for those
- [TRY] and be objective about a NO-GO
 - It happens. Accept it. Learn and schedule another attempt
 - I know: our instinct is to fight the failure
- All hands on deck Monday morning
 - Make sure to have all your monitoring in place (i.e. ProTop) when the users start logging in

Questions?



White Star
software

BONUS: Hard Lessons Learned

No one notices your tears.
No one notices your sadness.
No one notices your pain...
But they all notice your mistakes.



- A few of the mistakes we have seen over the years
- Hopefully these can save you from our pain

BONUS: Migration Lessons #1

Running out of disk space

- Test had plenty, but prod did not...

BONUS: Migration Lessons #2

Prod hardware is not equal to test hardware

- Usually prod is more powerful than test, but not always
 - Especially for disk IO

BONUS: Migration Lessons #3

AWS peak versus sustained throughput

- Each instance type has IO limits, regardless of the underlying disk types
 - You are given a quota to exceed those limits: for example, 30 minutes per day

BONUS: Migration Lessons #4

SSH keys and known_hosts

- If you have automated transfers setup, and you change servers, the known_hosts entries will change

BONUS: Migration Lessons #5

DNS entries and IP addresses

- Watch out for local hosts files versus DNS aliases and hard-coded IP addresses in some devices

BONUS: Migration Lessons #6

Firewall rules in prod versus test

- Prod is often more restricted than test

BONUS: Migration Lessons #7

Azure proximity placement groups

- For client/server applications, make sure the instances are close to each other with low-latency

BONUS: Migration Lessons #8

Network (or SAN) bandwidth usage

- Make sure the 100 TB Oracle data warehouse is not being backed up and uploaded to the cloud when you expect to do your migration

BONUS: Migration Lessons #9

Over-provisioning in hypervisor

- Lucky you: you have 8 vCPUs and 128 GB of RAM
- Unlucky you: so do the other 20 VMs on your server!!

BONUS: Migration Lessons #10

Incorrectly provisioned cloud disks

- AWS: gp2 vs gp3 vs io2
- Azure: standard, premium, ultra...
- IOPS, throughput...

**Thank you for your
attention**



Monitor OpenEdge. Anticipate Problems. Avert Disasters.

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



535

Customer Sites

1.5

Petabytes of data

4,500

Databases

325,000

Connected users

