



Set it and Forget it – Allow Progress OpenEdge to Take Care of Your Environment

Adam Backman
White Star Software

About the Speaker

- **President – White Star Software, LLC**

One of the oldest and most respected consulting and training companies in the OpenEdge Community

- **Vice President – DBAppraise, LLC**

Managed Database services from the experts. Get an expert, not someone off the bench to manage your database

- **Author - Progress Expert Series Books**



About This Presentation

- I will tell you things that will help maintain the integrity of your database
- Self-healing and self-protection are more about good setup and maintenance than just adding a few switches on database startup
- Secret: The name of the presentation was chosen to “fool” marketing into selecting it and it appears to have worked



Job Responsibilities of the Database Administrator

- Protect the data
- Plan for the future
- Keep the environment predictable



Job Responsibilities of the Database Administrator

- Backup and recovery - Protect
 - Maintain backups
 - Test recovery procedures
- Database security - Protect
 - Both OS and application
- Storage and capacity planning - Plan
- Performance monitoring and tuning - Predictable
- Software maintenance and control - Predictable
 - Manage software release process



Agenda

- Build the right foundation
 - Hardware
 - Operating System
 - Database
 - Application
- Reliable high performance
- OpenEdge startup options
- Maintenance and upkeep



Build The Right Foundation

- No single points of failure
 - Review your processes from end to end
- Disks (RAID 10 is your friend)
- Use journal-enabled file systems
- Memory (have enough and avoid interleaved/non-mirrored)
- CPUs
 - More than one CHIP not only more than one core
 - Note on NUMA (Don't) but if you do



Build the Right Foundation

- Infrastructure
 - Network (both internal and Internet)
 - Phones
 - Fax, EDI, Web, ...
- Property, plant and equipment
 - Where do people go if we lose the building?
 - How do they find out?
 - Where can they check in?



No Single Point of Failure

- If one fails, there is another there to take over transparently from a user's perspective
- Build for high availability because everyone requires the system to be available 100% of the time during their business hours
- Think about the whole system (end to end)
- Determine what “touches” your customer
- This is a mindset that everyone needs when acquiring any new technology (every project, every person on every project)



Storage Considerations

- RAID 10 is still recommended for database storage
 - Database
 - Before and After Image journals
 - Temporary files
- Disks are where to spend your money
- Most people overbuy CPU and neglect storage
- Why no RAID 5?



RAID – Why RAID 10 vs. RAID 5

- RAID 10 – isn't 10 times as good as 5
 - Two paths to the data
 - No calculation of parity
 - Optimized reads
- RAID 5
 - Great most of the time
 - Penalty for massive writes – like backup and recovery
 - Fine for test and development



Network Reliability

- Most companies cannot survive without Internet
 - Remote users
 - Customers
- All critical portions of the network must have redundancy

All data center clients should have multiple routes and end to end redundancy (NICs to Switches to Cables)
- Dual Internet feeds are becoming required for most clients
 - Physical – No shared wire (poles, ...)
 - Logical – multiple routes to Internet and intranet



Foundational Reliability

- View processes from end to end
 - Know your applications
 - Know the relationship between applications
- Prioritize items that directly affect your customers first
 - Know your business
 - Know your customers
 - Internal
 - External
- Next are things that indirectly affect customers
- Then come operationally critical items
- Last comes new development and reporting



Predictable Performance

- Performance should be essentially the same regardless of the load on the system
- The goals of data integrity and application performance are sometimes opposing
Example: RAID 5 – Data protected but performance and material cost for extra overhead and extra disk space
- If the users can't get at the data (poor performance) then they don't care about other things you did right.



Build the DB for Your Environment

- Use the best database block size
 - 8k most places
 - 4k Windows Server
- Use Type II Areas (Always best)
- Use the right BI cluster size – default is wrong
- Use the right network message size (-Mm 8192)
- Do you want every area to end with a variable extent or do you want all fixed.(*)

* You will **crash** if you fill an area with no variable extents



OpenEdge Startup Options: Self-Healing/Self-Protection

- The Watchdog process
 - Finds local users that have died abnormally and disconnects them
 - Remote Users are determined via a different mechanism
- Database Integrity Check Parameters
 - -DbCheck
 - -AreaCheck
 - -TableCheck
 - -IndexCheck



Database Integrity Check

- -DbCheck – Physical consistency check of all record and index blocks (except blobs) for the entire database
- -AreaCheck – Physical consistency check of all record and index blocks (except blobs) for an area
- -TableCheck – Enables consistency checks for a given table. Record block check only.
- -IndexCheck – Enables consistency checks for a given index. Index block check only.



Block-Level Integrity Check Rules

- Each option can only be used once.
- If `-DbCheck` is used then it takes precedence and consistency checking for the whole database will be applied
- `-AreaCheck`, `-IndexCheck` and `-TableCheck` can be used at the same time.
- These options can be enabled online via `promon`
 - `promon <dbname>`
 - R&D
 - Option 4 (Administrative Functions)
 - Option 8 (Enable/Disable block level consistency check)



Database Maintenance Options

- dbtool
- aiverify
- prostrct list
- prostrct –validate option
- prorest –vp and –vf options
- Update SQL Statistics (if necessary)



dbtool

- Option 1 – Report SQL width and date violations
- Option 2 – Report and correct SQL width violations
- Option 3 – Compare defined vs. physical records
- Option 4 – Validate record schema versioning
- Option 5 – Validate database block headers
- Option 6 – Scan records for possible corruption
- Option 7 – Schema validation (_StorageObject)
- Option 9 – Enable/disable redirection of tool output

Yes, There is no 8



Checking After Image files: aiverify

- Validates after image files before they are applied

- Partial verification

```
rfutil <dbname> -C aiverify partial -a <ai_file_name>
```

- AI block verification
- rlnote header verification

- Full verification

```
rfutil <dbname> -C aiverify full -a <ai_file_name>
```

- AI block verification
- rlnote header verification
- rlnote data verification



Maintenance – prostrct list

- Provides the current structure of the database
- Should be done after all structure maintenance
- I recommend running it as a scheduled task (cron) daily and putting the current copy in a separate directory (separate machine) for restore purposes



Maintenance – prostrct validate option

- Should be used in conjunction any prostrct add or create
- Determines if the .st file is valid
 - No blank lines
 - Correct starting character(s)
 - Extent definition info is valid and quoted
 - If it has area number, rec. per block, blocks per cluster they are correct
 - Correct delimiters
 - Colon followed by area number
 - Comma followed by records per block
 - Semicolon followed by blocks per cluster
 - Plus a bunch of other stuff



Maintenance: Testing backups –vp and –vf options for prorest

- probkup is the only way to backup your database truly online
- -vp – Is a partial verification pass
 - Does a CRC check of the blocks to check validity
 - Does not require the database to be down
 - No additional disk/tape needed as check is done in memory
- -vf – Full verification pass
 - Block for block check of backup against source database
 - Database must be down to use this option (or use copy)




Maintenance – update SQL statistics

- OpenEdge is SQL/92 compliant
- SQL, not ABL supports weight-based queries
- Weight-based queries are only as good as the data on which they are based.
- Active tables may warrant weekly statistic updates
- Generally, monthly is fine



Questions



Still have questions?

Please feel free to contact me directly.

Adam Backman
White Star Software
(603)897-1010
adam@wss.com



Thank you for your time!

