Database Dump/Load Performance Benchmark

By Nectar Daloglou

About the Speaker

Nectar Daloglou

- Principal Consultant at White Star Software
- Working with Progress and QAD for *almost* 20 years
- Performed specialized services at more than 80 Progress customer sites:
  - Progress Database Administration
  - Install/Upgrades/Migrations of Progress and QAD Applications
  - Technical Audits / Performance Tuning
  - Business Continuity Strategies

info@wss.com | wss.com

White Star software

What we will cover

- Dump/Load Process
- Benchmark Results

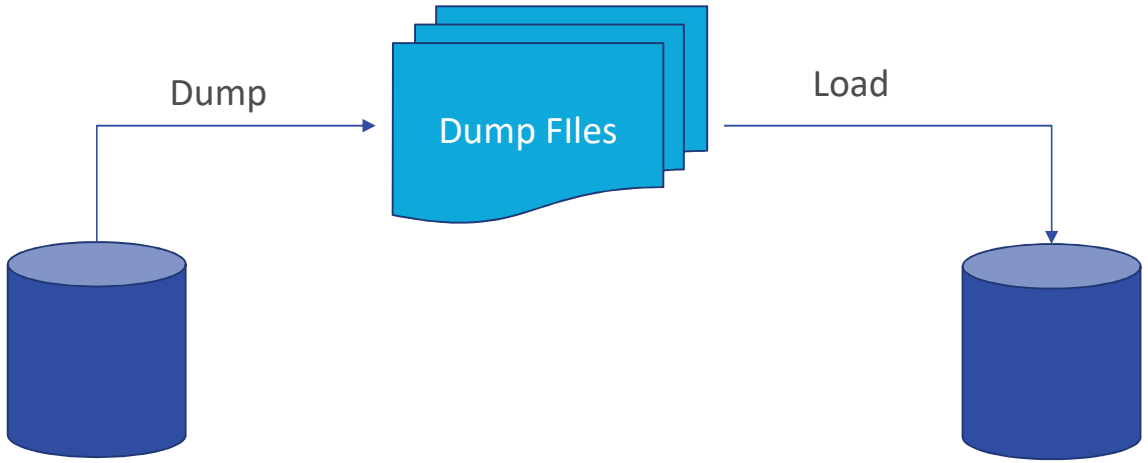

White Star
software

Your Mileage May Vary

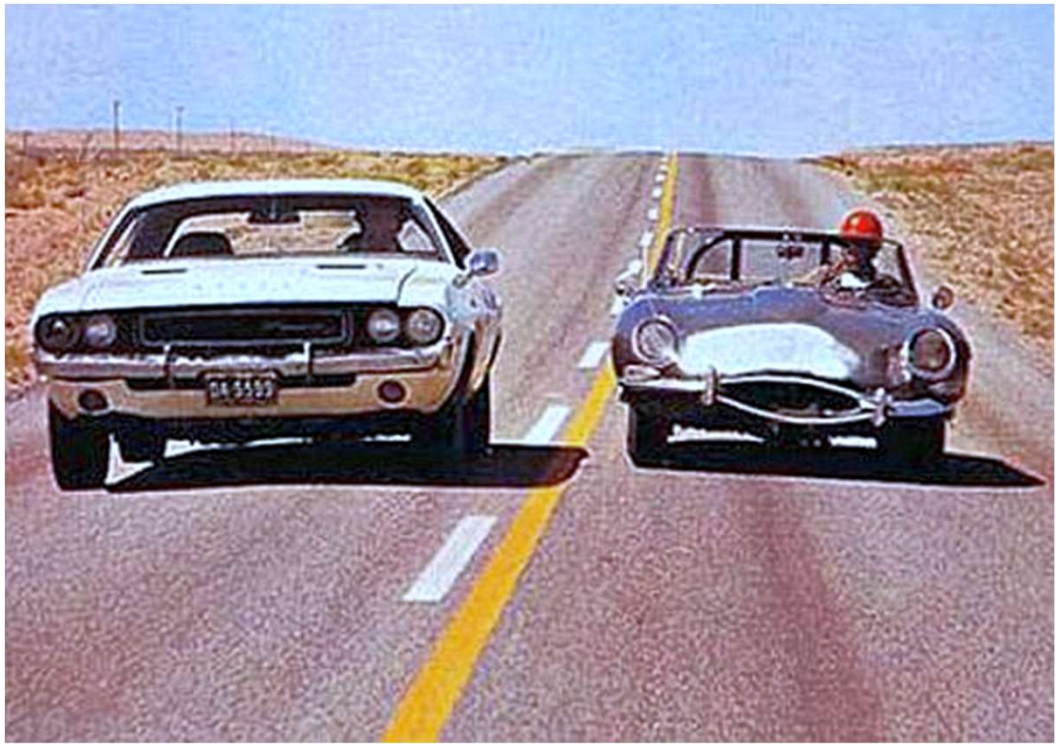

White Star
software

Benchmark Testing Hardware

- 8-Core Intel(R) Xeon(R) Platinum 8124M CPU @ 3.30GHz
- 16GB Memory
- Amazon Linux (CentOS 7.x equivalent)
- Different Disks:
  - 16k IOPS
  - 4k IOPS
  - 100 IOPS
- INTERESTING RESULTS! (and surprising)

White Star software

Benchmark Testing Hardware

Old HPUX

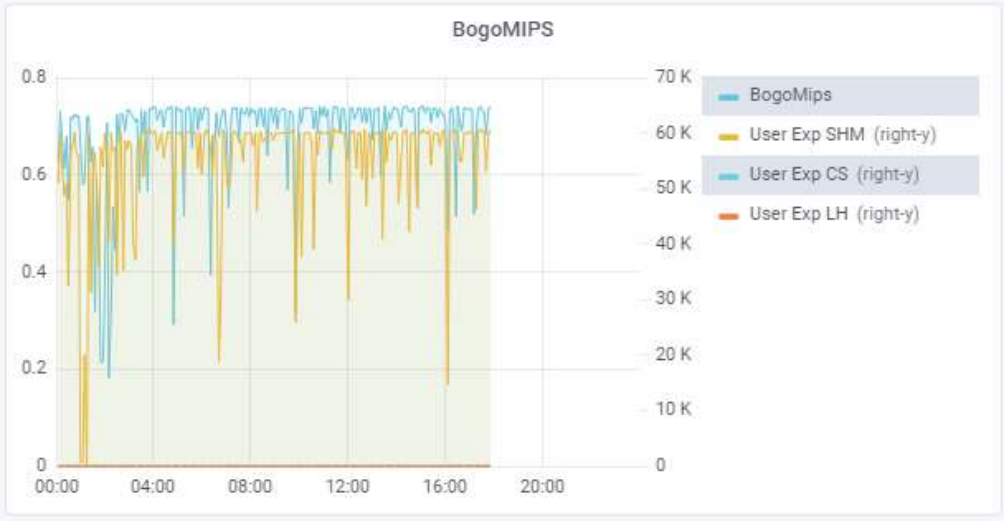

Amazon Instance

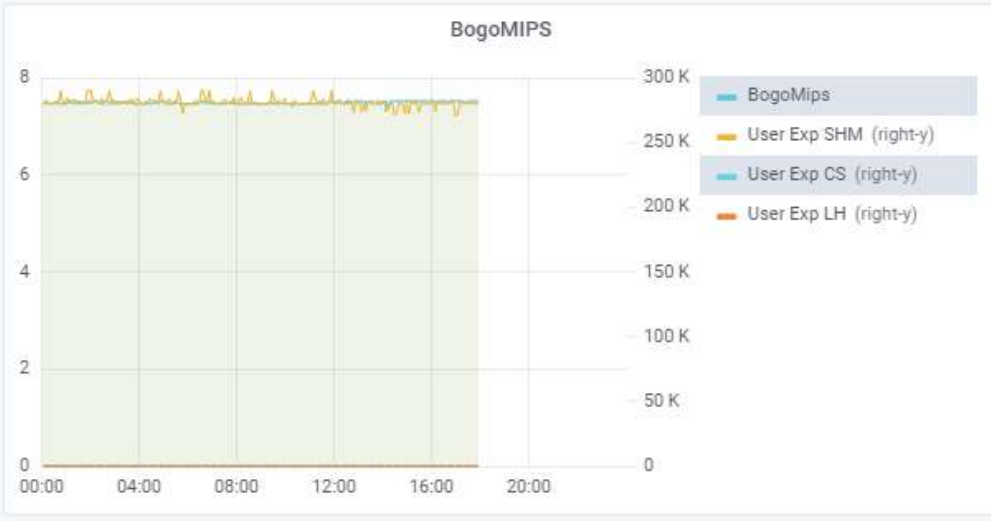

White Star software

Why or When to Dump/Load?

- Cross-platform Migration
- *Potential* Performance Increase
  - Reduce size after data purge
  - Reorganize to optimal RPB and BPC (Type 2)
  - Fix fragmented records
  - Eliminate physical scatter (Type 1)
  - Reduce logical scatter
  - Compact indexes (incidentally)
- Recover from corruption

White Star software

What to Dump/Load?

- Database schema (.df file)
- Data
- Sequence values
- Security (_user)
- SQL
  - User Permissions
  - Schema
  - Views
  - Etc..
- And more…

White Star software

Watch out!

- Check codepage
- Check DB/AI/BI Block size & BI cluster size
- Proutil Describe is your friend
- Validate loaded DB

White Star software

Binary D/L Process

- Preparation
  - Prepare scripts
  - Prepare new database
  - Draft checklist
- High level D/L process
  - Stop database/connections
  - Perform table analysis for record count
  - Dump database
  - Load new database
  - Rebuild Indexes

White Star software

Binary D/L Process

- Validation
  - Perform table analysis and compare record count
- Post D/L
  - Re-enable AI, OE Replication
  - Update Statistics
- Startup

White Star software

Why Benchmark D/L?

- Discover best possible configuration = Reduced outage
- Avoid surprises

White Star software

Benchmarking Tips

- Define goal
- Draft test scenarios
- Apply and measure one change at a time
- Automate
- Use tools to measure performance
- Document all results
- After each iteration:
  - Drop your cache not your CA$H (echo 3 > /proc/sys/vm/drop_caches)
  - Restart database
- On shared infrastructure: run test more than once

Well known Benchmarks

- Binary D/L faster than Data Dictionary D/L
- Multithreaded dump is faster
- Use of “no-integrity” -i or “raw”-r parameters will load faster

White Star software

Benchmark Database

- ERP database
- Type 1 SA
- Objects in Schema Area
- 14 GB
- 873 Tables
- 3140 Indexes
- Areas
  - 18 Table
  - 18 Index

Table over 1GB	Records	Size
gltmd_det	14,856,965	1.3GB
gltr_hist	5,145,602	1GB
glec_det	9,766,305	1GB
xxext_xref	18,535,331	3.7G

White Star software

Current Structure

- Objects in Schema Area

```
--------------------------------------------------------------------------------

Summary for AREA "Schema Area": 6

Subtotals:                                  20570303    4.5G      6  3201   236
20577868    1.0     2.3
```

- Type 1 Storage areas

```
d "TRANSACTION":7,64;1 .
d "TRANSACTION_IDX":8,32;1 .
d "STATIC":9,64;1 .
d "STATIC_IDX":10,32;1 .
d "HISTORY":11,64;1 .
d "HISTORY_IDX":12,32;1 .
```

White Star software

Type 1 vs Type 2

- TRANSACTION AREA in Type I

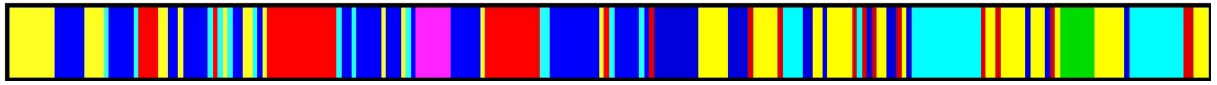

- TRANSACTION AREA in Type II

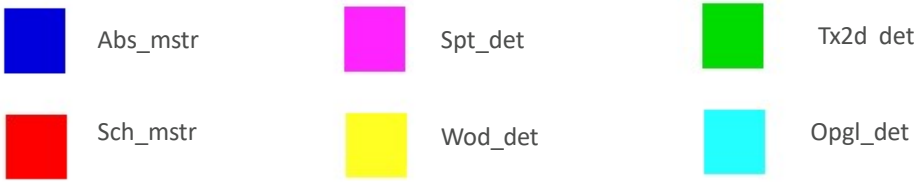

White Star software

High Scatter Factor

RECORD BLOCK SUMMARY FOR AREA "TRANSACTION" : 7

Table	Records	Size	-Record Size (B)- Min	Max	Mean	---Fragments--- Count	Factor	Scatter Factor
PUB.absc_det	12210	727.4K	61	61	61	12210	1.0	4.0
PUB.absd_det	100449	11.2M	79	154	116	100457	1.0	3.8
PUB.abs_mstr	84520	26.2M	195	641	324	84953	1.0	5.2
PUB.acd_det	6108	584.1K	82	125	97	6108	1.0	4.3
PUB.ap_mstr	18405	3.0M	154	194	168	18465	1.0	5.4
PUB.ard_det	83500	9.9M	114	157	123	83506	1.0	3.8

White Star software

D/L Preparation

- Move objects out of Schema Area
- Separate large tables & indexes
- Group common Records per Block areas
- Create new structure file
- Create new data definition file (.df) or
- Move objects to new areas

White Star software

ProTop is Your Friend

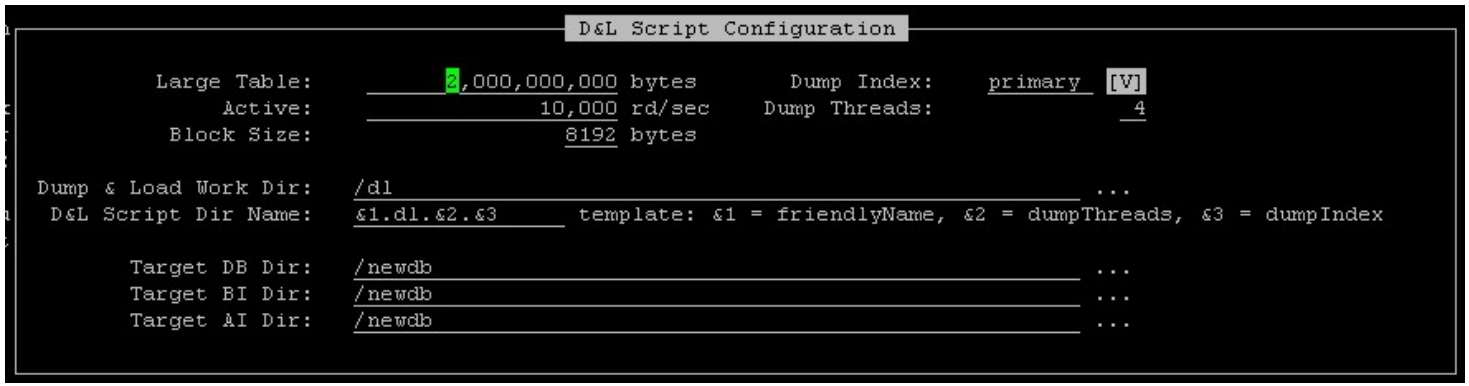

- Create new structure file
  - mfgprod.st
- Create new data definition file (.df) or
  - mfgprod.df.new
- Move objects to new areas
  - mfgprod.tblmv.sh

New Structure

```
d "misc8_dat":100,8;64 /data-fastest/newdb/mfgprod_100.d1
#
d "misc8_idx":101,1;64 /data-fastest/newdb/mfgprod_101.d1
#
d "misc16_dat":102,16;64 /data-fastest/newdb/mfgprod_102.d1
#
d "misc16_idx":103,1;64 /data-fastest/newdb/mfgprod_103.d1
#
d "misc32_dat":104,32;64 /data-fastest/newdb/mfgprod_104.d1
#
d "misc32_idx":105,1;64 /data-fastest/newdb/mfgprod_105.d1
#
d "misc64_dat":106,64;64 /data-fastest/newdb/mfgprod_106.d1
#
d "misc64_idx":107,1;64 /data-fastest/newdb/mfgprod_107.d1
#
d "misc128_dat":108,128;64 /data-fastest/newdb/mfgprod_108.d1
#
d "misc128_idx":109,1;64 /data-fastest/newdb/mfgprod_109.d1
#
d "misc256_dat":110,256;64 /data-fastest/newdb/mfgprod_110.d1
#
d "misc256_idx":111,1;64 /data-fastest/newdb/mfgprod_111.d1
#
d "gltr_hist":112,64;512 /data-fastest/newdb/mfgprod_112.d1
#
d "gltr_hist_idx":113,1;64 /data-fastest/newdb/mfgprod_113.d1
```

Here We Go!

```
$ ./load.sh

============================================================
shut down /data-fastest/newdb/mfgprod and load from scratch? y
shutting down mfgprod
Shutdown is executing. (1613)
Shutdown complete. (1614)

============================================================
clean old logs? y

============================================================
clean dump & stage directories? y
```

White Star software

Ready to Dump Source Data

```
This is a full backup of /data-fastest/newdb/mfgprod.db. (6759)
This backup was taken Sat Oct  5 03:34:41 2019. (6760)
The blocksize is 8192. (6994)
mfgprod already exists.
Do you want to over write it? [y/n]:  Start of extending target DB to
needed size... (9432)
It will require a minimum of 248798 blocks to restore. (6763)
Start of restoring the target DB... (9433)
Read 131110 db blocks in 00:00:03


============================================================
dump source data? y
```

White Star software

Dump Benchmark Results

CPU Cores	IOPS	Time	Improvement
2	300	6m40s	
4	300	6m2s	10% with more cores
8	300	6m2s	None
2	4000	3m8s	215% with more IOPS
4	4000	3m8s	None
8	4000	3m6s	3% with more cores
2	16000	1m17s	244% with more IOPS
4	16000	1m15s	2% with more cores
8	16000	1m2s	20% with more cores

White Star software

Dump Slow?

- Dump using smaller index
  - proutil sports –C dump customer . -index 5
- Dump -index 0
  - proutil sports -C dump customer . -index 0
  - Does not address logical scatter
  - Only on Type 2
- Dumpspecified
  - proutil "dbname" -C dumpspecified "[owner.]table.index_field" "operator" "value" AND "operator" "value" "directory" -preferidx index-name

White Star software

Load

```
      Stage: /data-fastest/dl/stage
       Load: /data-fastest/dl/load
    Archive: /data-fastest/dl/arc
       Logs: /data-fastest/dl/log

 Target  DB: /data-fastest/newdb/mfgprod

   # Tables:         873
```

White Star software

Load

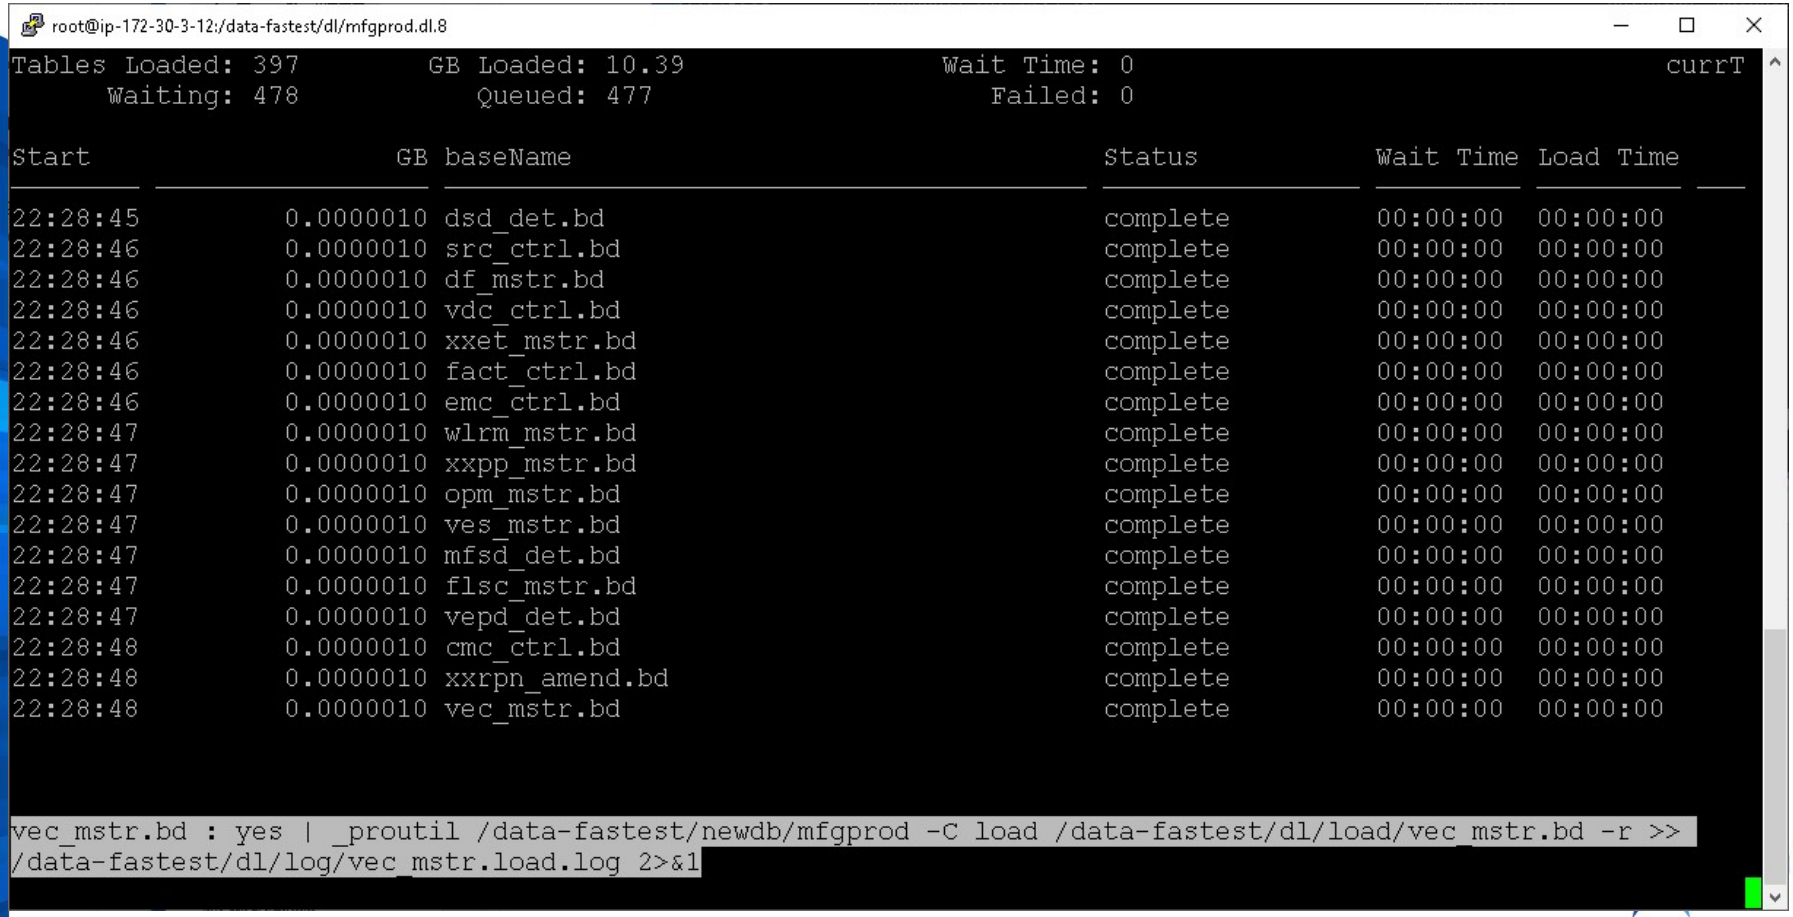
root@ip-172-30-3-12:/data-fastest/dl/mfgprod.dl.8
Tables Loaded: 397 GB Loaded: 10.39 Wait Time: 0 currT
Waiting: 478 Queued: 477 Failed: 0
Start GB baseName Status Wait Time Load Time
22:28:45 0.0000010 dsd_det.bd complete 00:00:00 00:00:00
22:28:46 0.0000010 src_ctrl.bd complete 00:00:00 00:00:00
22:28:46 0.0000010 df_mstr.bd complete 00:00:00 00:00:00
22:28:46 0.0000010 vdc_ctrl.bd complete 00:00:00 00:00:00
22:28:46 0.0000010 xxet_mstr.bd complete 00:00:00 00:00:00
22:28:46 0.0000010 fact_ctrl.bd complete 00:00:00 00:00:00
22:28:46 0.0000010 emc_ctrl.bd complete 00:00:00 00:00:00
22:28:47 0.0000010 wlrm_mstr.bd complete 00:00:00 00:00:00
22:28:47 0.0000010 xxpp_mstr.bd complete 00:00:00 00:00:00
22:28:47 0.0000010 opm_mstr.bd complete 00:00:00 00:00:00
22:28:47 0.0000010 ves_mstr.bd complete 00:00:00 00:00:00
22:28:47 0.0000010 mfsd_det.bd complete 00:00:00 00:00:00
22:28:47 0.0000010 flsc_mstr.bd complete 00:00:00 00:00:00
22:28:47 0.0000010 vepd_det.bd complete 00:00:00 00:00:00
22:28:48 0.0000010 cmc_ctrl.bd complete 00:00:00 00:00:00
22:28:48 0.0000010 xxrpn_amend.bd complete 00:00:00 00:00:00
22:28:48 0.0000010 vec_mstr.bd complete 00:00:00 00:00:00
vec_mstr.bd : yes | _proutil /data-fastest/newdb/mfgprod -C load /data-fastest/dl/load/vec_mstr.bd -r >>
/data-fastest/dl/log/vec_mstr.load.log 2>&1

Load Complete

```
Load complete! Sun Oct 6 22:30:01 UTC 2019

Running tabanalys in background on source db...
```

White Star software

Load Benchmark Results

CPU Cores	IOPS	Time	Improvement
2	300	13m55s	
4	300	13m54	
8	300	13m44s	
2	4000	4m46	Almost x3 faster
4	4000	4m42	
8	4000	4m40s	
2	16000	5m3s	Slightly faster
4	16000	4m38s	
8	16000	4m45s	

White Star software

Onto Index Rebuild

```
Starting index build...
Use:  tail -f /data-fastest/dl/log/idxbuild.2019.10.06.22.30.log  from another
window to monitor the output (if desired).
This session will quietly wait for the idxbuild to finish.
2019.10.06.22.30
2019.10.06.22.39 idxbuild finished
3270 indexes were rebuilt. (11465)
Index rebuild complete.  0 error(s) encountered.  Elapsed time: 496.537

Resource usage: CPU user 647.110769, system 18.091232
Resource usage: DISK reads: 11907578 KB at 23 MB/sec, writes: 5590976 KB at 11
MB/sec

============================================================
```

White Star software

Index Rebuild Benchmark Results

idxbuild all -rusage –z -B 512 -SG 64 -TB 64 -TM 32 -TMB 4096 -TF 80 -datascanthreads 8 -mergethreads 8

CPU Cores	IOPS	Time	Improvement
2	300	11m11s	
4	300	12m	
8	300	12m20s	
2	4000	4m4s	
4	4000	3m47	Slightly faster than 2,8 cores
8	4000	4m02s	
2	16000	4m56s	
4	16000	3m53s	Slightly faster than 2,8 cores
8	16000	4m01s	

White Star software

Index Rebuild Parameters Benchmark Results

8 cores 16k IOPS TEST: idxbuild all -rusage -z <PARAMS>:

-B	-TB	-TM	-TMB	-TF	-SG	-thread	-threadnum	-mergethreads	-datascanthreads	Time (s)	Time (m)
512	64	32	512	80	64	1	8	4	8	352	5.9
512	64	32	**4096**	80	64	1	8	4	8	386	6.4
512	64	32	512	80	64	1	8	**2**	8	278	4.6
512	64	32	512	80	64	1	**2**	**8**	8	276	4.6
512	64	32	512	80	64	1	**4**	**4**	8	278	4.6
512	64	32	**4096**	80	64	1	**8**	**2**	8	386	6.4
512	64	32	**64**	80	64	1	8	2	8	249	4.2
512	32	32	**32**	80	64	1	8	2	8	246	4.1
512	16	32	**16**	80	64	1	8	2	8	247	4.1
512	16	32	16	80	64	1	**2**	**8**	8	247	4.1
512	**32**	32	**32**	80	64	1	2	8	8	243	4.0
512	**64**	32	**64**	80	64	1	2	8	8	244	4.1
512	**16**	32	**16**	80	64	1	**4**	**4**	8	245	4.1
512	**32**	32	**32**	80	64	1	4	4	8	245	4.1
512	**64**	32	**64**	80	64	1	4	4	8	242	4.0

Index Rebuild Parameters Benchmark Results

```
-B 512 -TB 64 -TM 32 -TMB 64 -TF 80 -SG 64 -thread 1
-threadnum 4 -mergethreads 4 -datascanthreads 8 -z -rusage
```

- Best Time: 4 minutes
- Smaller -threadnum and larger -mergethreads is better (product < 2 X #cpu)
- Smaller –TMB is better
- Smaller –TMB affected datascan time; not only merge time

White Star software

And Finally Record Count Check

```
execute tabanalys on /data-fastest/newdb/mfgprod? y

Running tabanalys on /data-fastest/newdb/mfgprod...
output will be in  /data-fastest/dl/log/mfgprod.new.tba  and  /data-
fastest/dl/log/mfgprod.new.tbx to verify record counts run:

  sum /data-fastest/dl/log/*.tbx

once both the old & new tabanalys are complete, if the checksums match then the
record counts are the same

waiting for tabanalys to complete

45882    14 /data-fastest/dl/log/mfgprod.new.tbx
45882    14 /data-fastest/dl/log/mfgprod.old.tbx
```

White Star software

End-to-end Dump/Load Benchmark Results

- Parallel dump/load faster with higher IOPS disks

CPU Cores	IOPS	Dump	Load	Combined Dump/Load	Parallel Dump/Load	Combined vs. Parallel	Index Rebuild	End-to-end Parallel
8	300	6m2s	13m44s	19m46s	15m27	28% Slower	12m20s	36m10s
8	4000	3m6s	4m40s	7m46s	5m11s	50% Faster	4m02s	13m04s
8	16000	1m2s	4m45s	5m47	5m10s	12% Faster	4m01	10m59s

White Star software

Other Tricks

- OpenEdge Replication
  - Use the replication target to off-load some work
  - Specifically move the 8 largest tables to the replication server
  - Dump time reduced to 2hr 31min
  - Total time: 5hr 30min

- Use CDC to reduce outage window
  - Establish baseline
  - Enable CDC
  - Restore new DB
  - Apply changes to new DB
  - Switch to new DB once changes have caught up

White Star software

Conclusion

- Disk I/O can have the most impact on timing
- Benchmarking and testing is essential
  - Reduce outage window
  - Ensure integrity
- ProTop can help
  - Dump/Load toolkit
  - Performance metrics

White Star software

Q&A
Tell us your story.

About White Star Software

- The Oldest and Most Respected Independent Progress OpenEdge Consulting Firm
- 5 of the top OpenEdge DBAs in the world: Adam Backman, Tom Bascom, Dan Foreman, Paul Koufalis and Nectarios Daloglou
- 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

White Star
software

protop

THE BEST OPENEDGE PERFORMANCE, MONITORING, AND ALERTING TOOL IN THE GALAXY! | WSS.COM/PROTOP

Detect and correct
issues before they affect
your critical business processes

White Star software