

New Hardware = Worse Performance?

By Nectar Daloglou



White Star
software

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



Picture this!

- New server purchased (or new VM provisioned)
- More CPU, more memory, SSD disks
- Vendor assured you performance will be better.....WHY TEST?

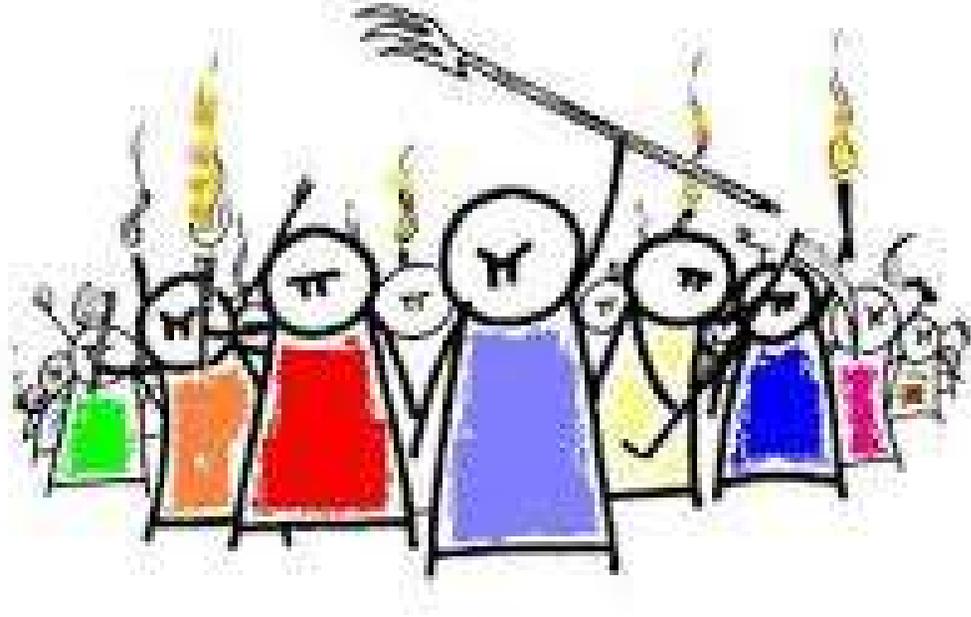
But...

- After migrating, performance got WORSE!

info@wss.com | wss.com



Users are not very happy



[in](#)

Management wants to “talk”



Agenda

- Virtualization
- Disk Subsystem
- Network
- CPUs

Case #1: Performance Virtually SUCKS!

- Progress-based ERP system
- AppServer and Webspeed brokers
- CPU often spiking to 99%
- Users complained about performance
- Server was often rebooted to resolve issue



BogoMIPS

- A ProTop FREE metric
- A CPU measure from an OpenEdge application's perspective



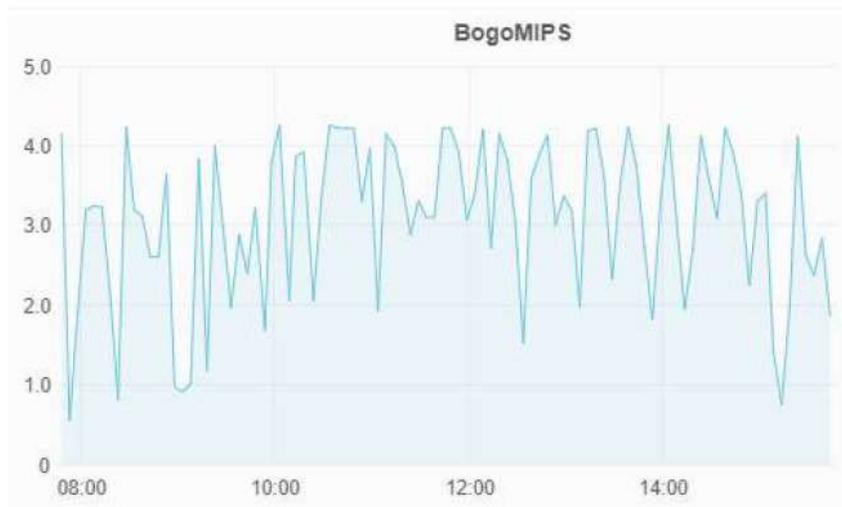
Findings

- Virtual Host running Hypervisor
- Windows 2 X Quad Core 2.5Ghz CPU (8 cores)
- OpenEdge VM had 2 cores allocated
- VM Host was provisioning a total of 13 cores
- Wait what??? $13 > 8$

BogoMIPS

- Problem resolved after reducing cores from 13 to 7:

Before



After



Case #2: The SAN Scam

- Customer migrating from HP-UX to Windows VM
- Opted for performance testing
- All tests were normal except...



SyncIO Test

```
=== proutil sports -C bigrow 2 -zextendSyncIO ===  
OpenEdge Release 11.7.2 as of Tue Oct 24 19:02:01 EDT 2017  
Execution time: 48.829 s  
9 seconds = 10MB/sec -- anything longer and your disk subsystem is junk  
5 seconds = 20MB/sec -- not horrible  
3 seconds = 30MB/sec -- good  
1 second = 100MB/sec -- excellent!
```

- How fast can we grow 100MB unbuffered
- Storage on a RAID 5 NetApp FAS2552 SAN with SSD



How fast is your Ferrari in traffic?



Why SANs do not perform well for databases

- Storage is external
 - Longer distance from server to disk
- Storage is shared
 - Other I/O intensive operations may affect performance
- An accounting solution
 - Savings in mind not performance

Minivan vs. sports car

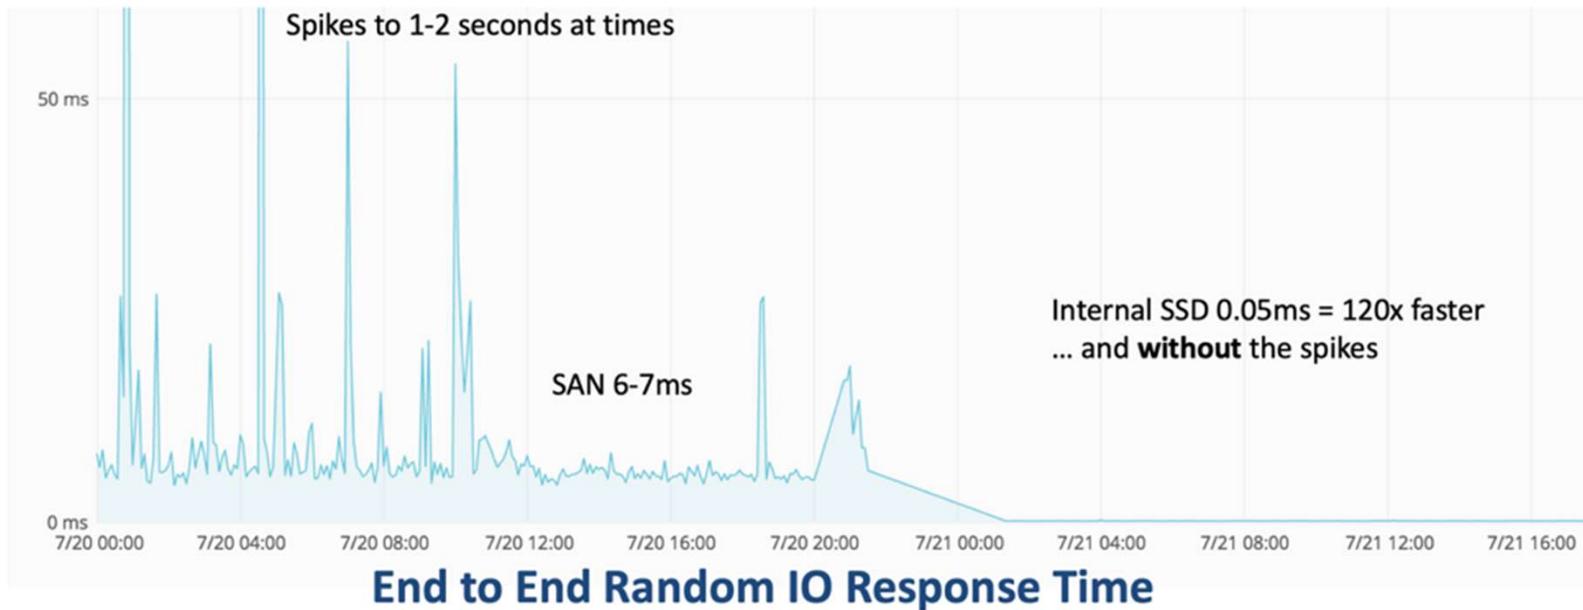


Read performance on memory vs. disk

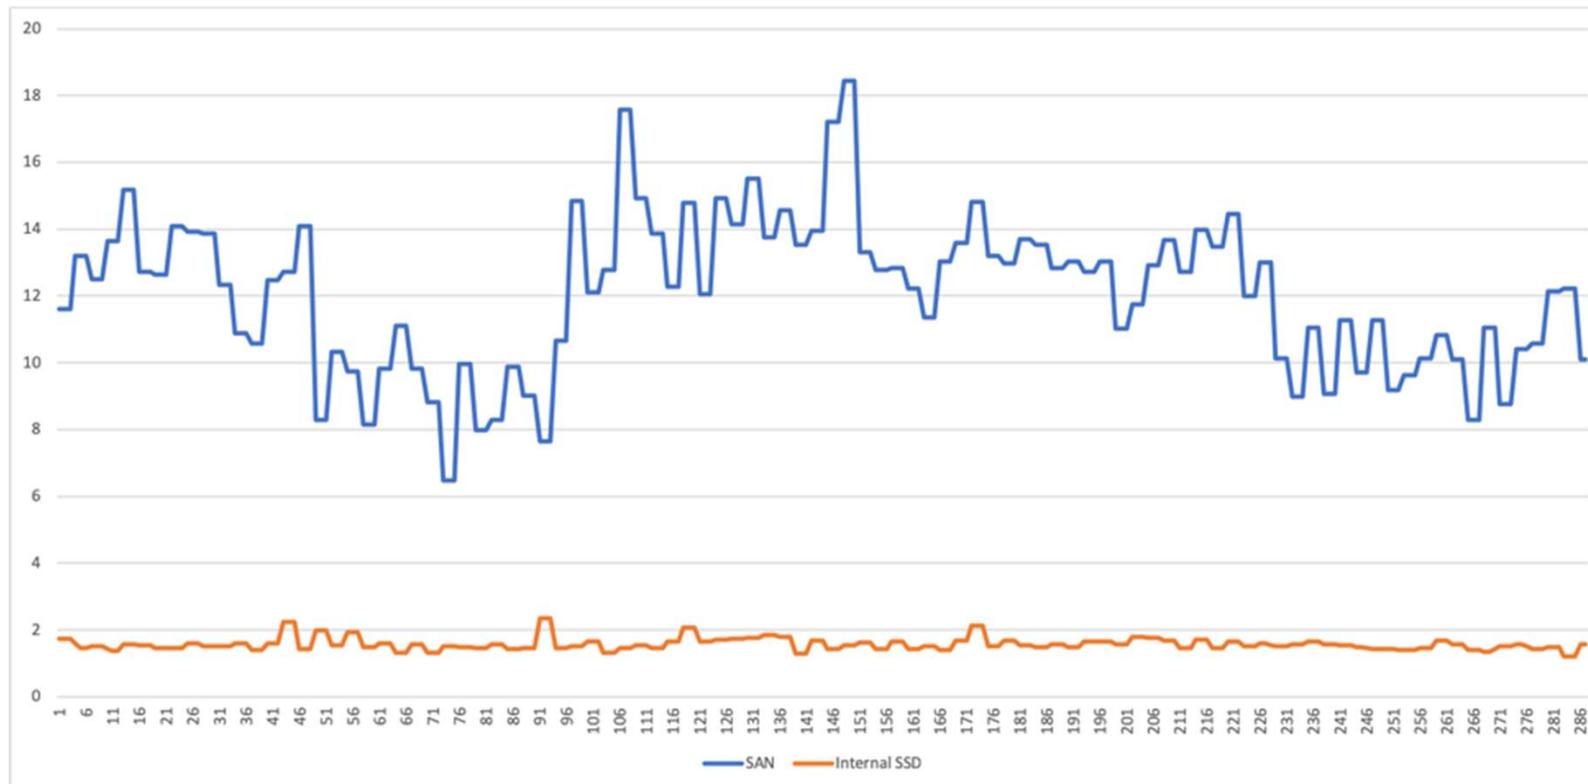
Layer	Time	# of Recs	# of Ops	Cost per Op	Relative
Progress to -B	0.96	100,000	203,473	0.000005	1
-B to FS Cache	10.24	100,000	26,711	0.000383	75
FS Cache to SAN	5.93	100,000	26,711	0.000222	45
-B to SAN Cache	11.17	100,000	26,711	0.000605	120
SAN Cache to Disk	200.35	100,000	26,711	0.007500	1500
-B to Disk	211.52	100,000	26,711	0.007919	1585

(Approximately 4 records per read op in non -B cases.)

Monitoring IO Response Time in ProTop



Monitoring Unbuffered Writes in ProTop: SyncIO



Resolution

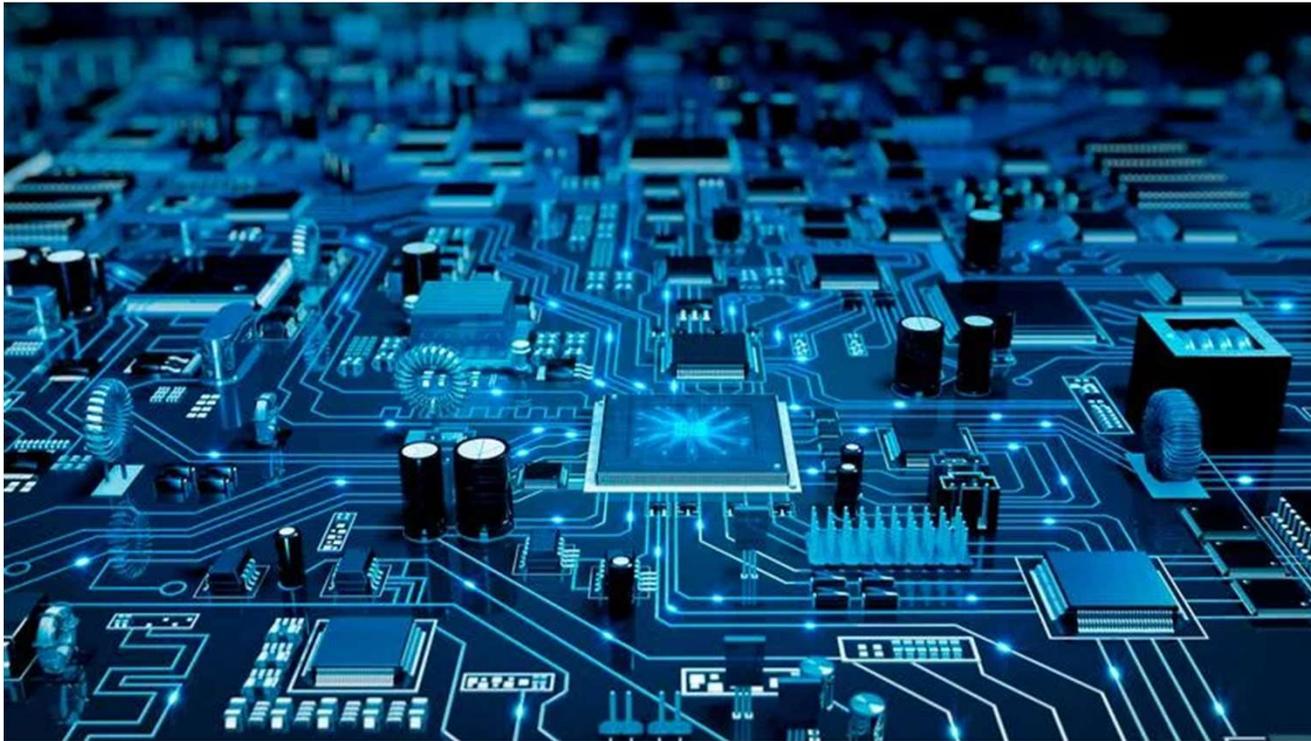
- Firmware update on the SAN brought SyncIO down to 9 seconds

Case #3: The “Switch”

- Users complained of performance issues
- Problem mysteriously surfaced on August 1st



What could go wrong?



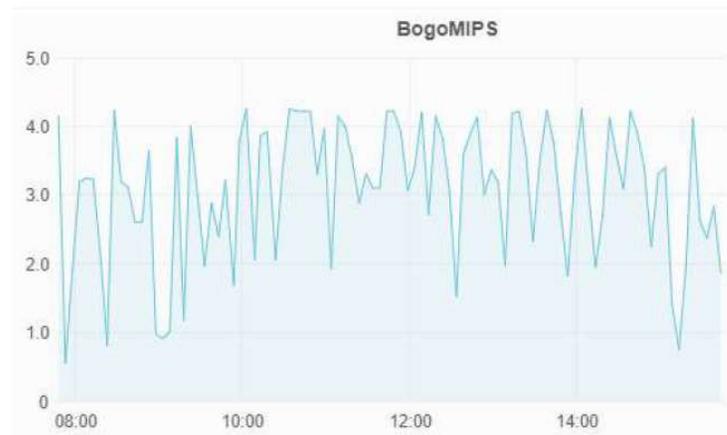
Client/Server User Experience Decline

- A network switch was replaced



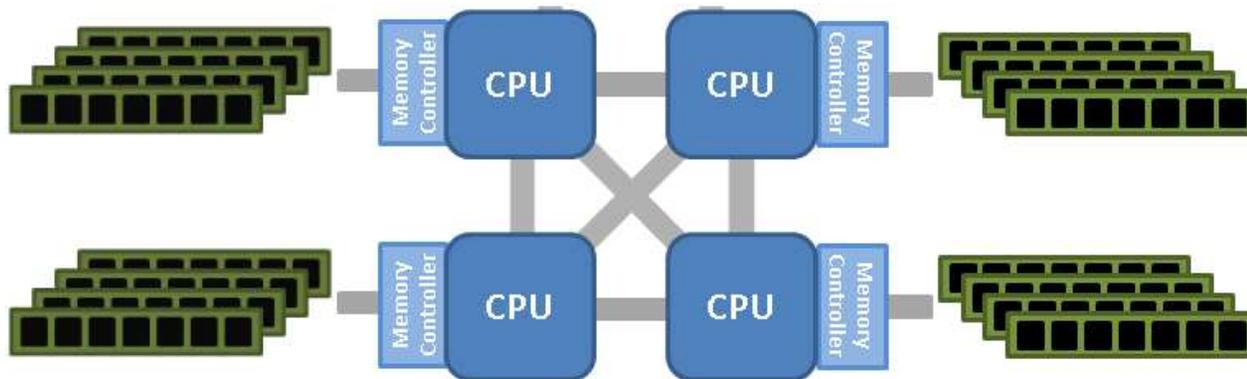
Case #4: NUMA

- Variable performance reported by customer
- Turns out they had migrated on a NUMA machine and planning on purchasing another.
- Whoa!



What is NUMA?

- Non-Uniform Memory Access
- Like a cluster of tightly coupled machines into one



Avoid Issues: Benchmark test on new hardware

- Discover Best Possible Configuration
- Avoid Surprises



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
 - Restart database

Benchmarking Tools

- syncio.sh for disk writes (ProTop Free)
- ATM for tx throughput, with AI enabled
- ReadProbe for single-threaded rapid readers (ProTop Free)
- Spawn to simulate users logging in to the server and connecting to the DB (ProTop Free)
- Leave ProTop running on the future prod box for a while, get a feeling for bogoMIPS and IOResponse over time (ProTop \$\$\$)



Readprobe

```
# of Sessions:      100   After:  50   Jump by:  10
```

```
CPU info:          8 x   3320 MHz
```

```
Comment:  edit rp.sh to pre-load comment >>
```

```
Scenario: >>
```

```
uname -a:  Linux ip-172-30-3-12.ec2.internal 4.14.138-114.102.amzn2.x86
```

```
Ext Sys: >>
```

```
Contributor:  root
```

```
    e-mail:  nd@wss.com
```

```
Progress Version:  12.0
```



Readprobe

Loop:	86	Best Single:	1,742,913
Sessions:	18	Best Rec/sec:	2,638,380
Time:	5.0010	Best Users:	14
Latch Waits:	561	Worst Latch Waits:	561
LRU Waits:	66	Worst LRU Waits:	66
LRU2 Waits:	0	Worst LRU2 Waits:	0
Rec/sec:	2542664	%usr:	97.53
Rec/User:	141259	%sys:	0.85
Blk/sec:	3412923	%idle:	1.62
User Exper:	14	%Entitled:	0.00
		Load Avg:	11.85



Conclusion

- Monitor environment for infrastructure and performance changes
- Perform benchmark tests on new hardware before migrating
 - Avoid surprises
- ProTop can help
 - Free performance tools
 - Ongoing performance monitoring and alerting



Q&A
Tell us your
story.



White Star
software

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





THE BEST OPENEDGE PERFORMANCE, MONITORING, AND ALERTING TOOL IN THE GALAXY! | [WSS.COM/PROTOP](https://wss.com/protop)



Detect and correct issues before they affect your critical business processes

