Tales From the Bunker
Episode No. 6

Presented by: Gus Björklund, Dan Foreman, John Harlow

BRAVEPOINT

PROGRESS SOFTWARE


Chigwell
Stanford Rivers
A 113
Brentwood
Kelvedon Hatch
Industrial Estates
A 128
Secret Nuclear Bunker

Bunker tests: a brief history

- The bunker tests originated several years ago in an effort to see how much performance commodity priced X86 hardware could provide.
- In the interest of security we have almost never conducted a bunker test in the same location twice.
- Most bunker tests have resulted in improvements in process and performance.

Bunker History

- Bunker #1 - Florida 2002
- Bunker #2 - Atlanta 2002
- Bunker #3 – Nashua 2004
- Bunker #4 – Florida 2005
- Bunker #5 - Sept 2009 – in the Cloud

Tales From The Bunker

PUG Challenge Americas

Bunker 1-4 results

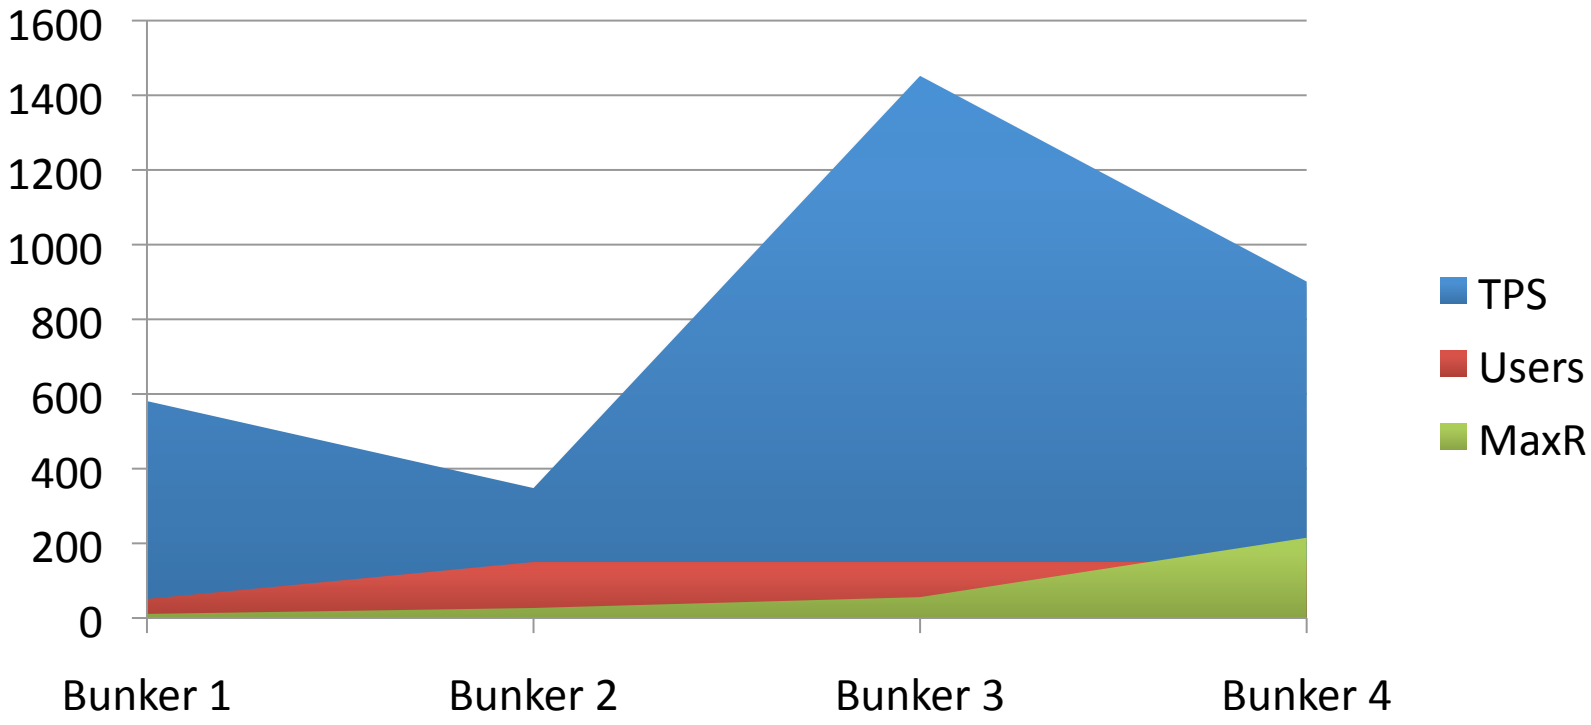

Bunker Team

- Gus Bjorklund (due to Bunker security protocols, we can't reveal his face; bodyguard in background)

PUG Challenge Americas

Bunker Team

- Dan Foreman – revealing too much

PUG Challenge Americas

Bunker Team

- John Harlow – fallen comrade

PUG Challenge Americas

Supporting Cast

- Ty the Invisible Tech Guy

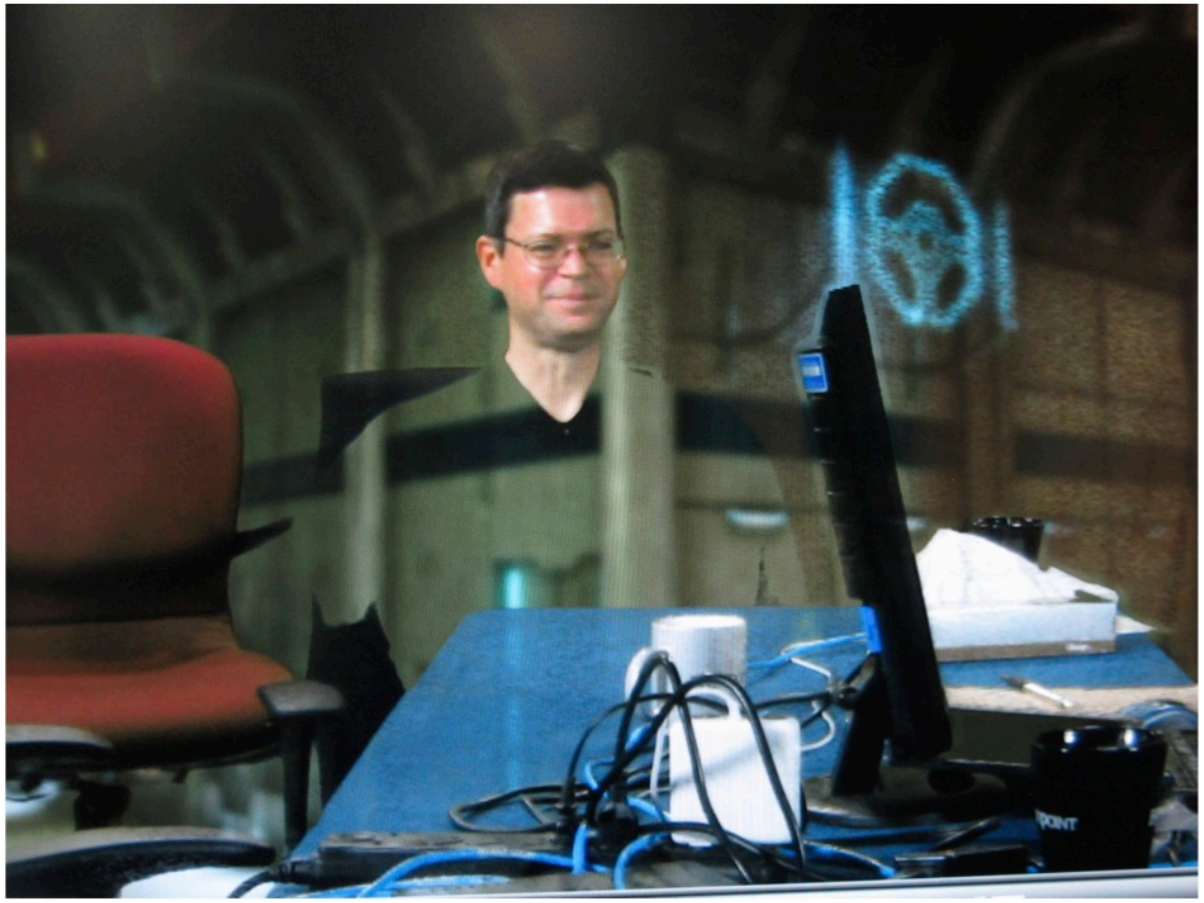

PUG Challenge Americas

Bunker #6 Location

PUG Challenge Americas

Bunker #6 Location

Tales From The Bunker

PUG Challenge Americas

Actual Work Area

PUG Challenge Americas

Virtual Work Area

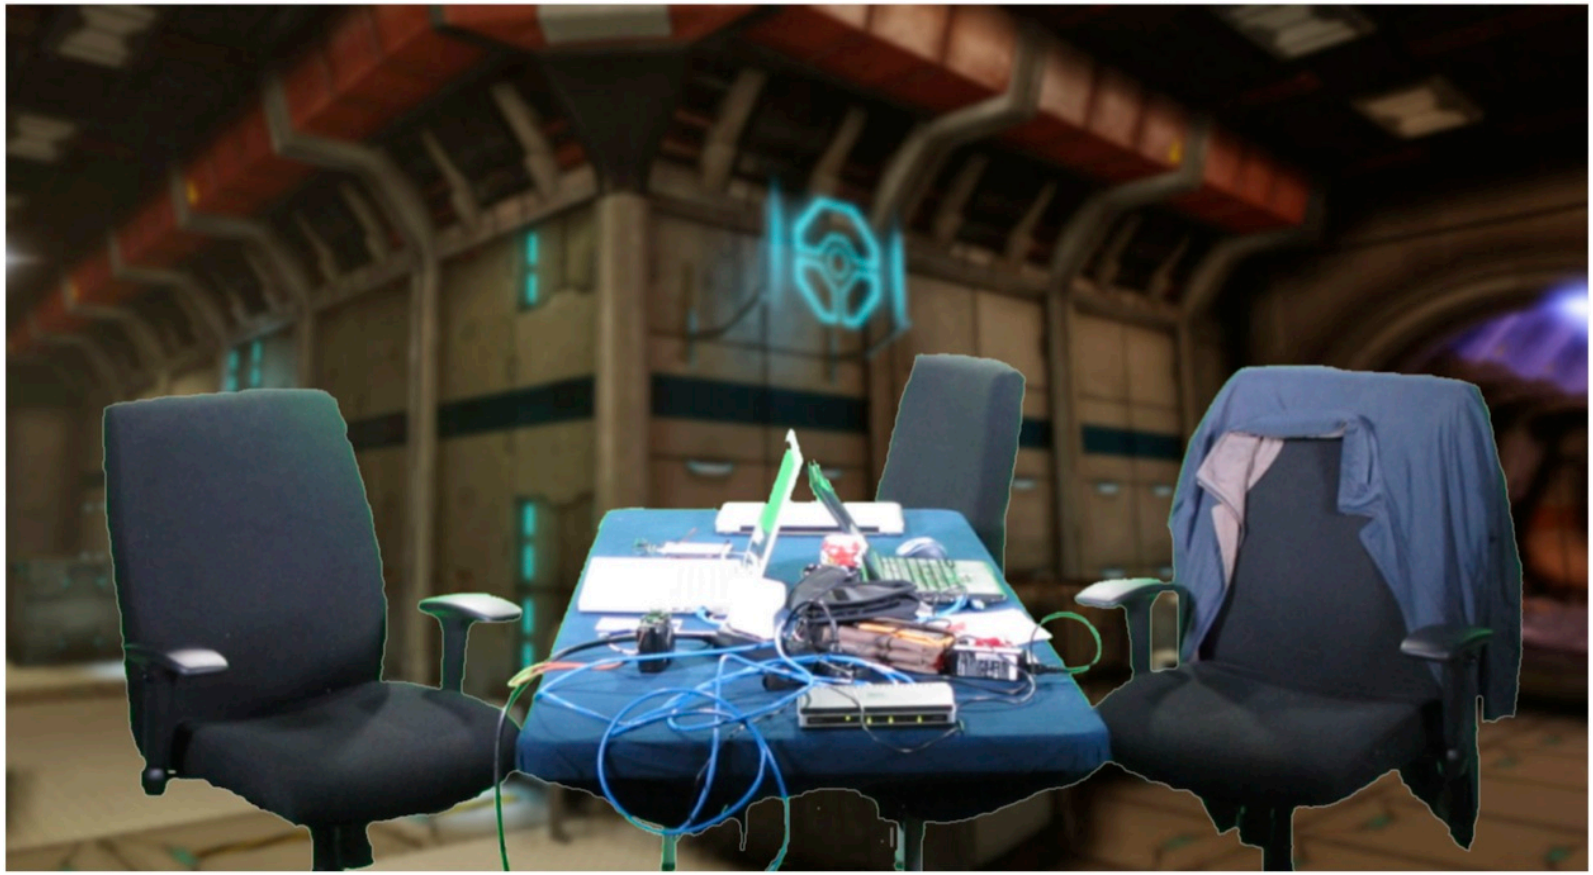

PUG Challenge Americas

Bunker Logo

PUG Challenge Americas

Server Info

- Dell R710 (two of them: bunker64 and bunker32)
  - 16 CPUs
  - 32 GB RAM

SAN Info

- EMC CX4-120 (Dan owns stock in EMC)
- Fabric: 4GB Fiber Channel
- 14 Disks + one hot swap spare
- 300 gb disks
- 15000 RPM
- Initially configured as RAID 5 !!!!!
- Later switched to RAID 10

Software Info

- VSphere Enterprise 4.1
- Progress V10.2B SP03
  - 64-bit
  - 32-bit
- Centos 5.5 (2.6.18-194.32.1.el5, if you must know)
  - 32 bit
  - 64 bit

Additional Equipment

PUG Challenge Americas

Test Protocols

- ATM Benchmark v 5.0
- Balanced Benchmark from BravePoint

About ATM ...

- Standard Secret Bunker Benchmark
  - baseline config always the same since Bunker#2
- Simulates ATM withdrawal transaction
- 150 concurrent users
  - execute as many transactions as possible in given time
- Highly update intensive
  - fetch 3 rows
  - update 3 rows
  - create 1 row with 1 index entry

About ATM ... database

account rows	80,000,000
teller rows	80,000
branch rows	8,000
data block size	4 k
database size	~ 12 gigabytes
maximum rows per block	64
allocation cluster size	512
data extents	6 @ 2 gigabytes
bi blocksize	16 kb
bi cluster size	16384

the "standard baseline" database setup

About ATM ... baseline config

```
-n 250          # maximum number of connections
-S 5108         # broker's connection port
-Ma 2           # max clients per server
-Mi 2           # min clients per server
-Mn 100         # max servers
-L 10240        # lock able entries
-Mm 16384       # max TCP message size
-maxAreas 20    # maximum storage areas
-B 64000        # primary buffer pool number of buffers
-spin 10000     # spinlock retries
-bibufs 32      # before image log buffers
```

About Balanced Benchmark ...

- BravePoint product
- Makes load & performance testing very easy for:
  - Servers
  - Operating systems
  - Databases
  - Applications
  - All the above

About Balanced Benchmark ...

- Easy Command & Control of a Progress Benchmark
- Can use code from various sources
  - Your application
  - Load simulation code generated by the BB based upon production DB activity
  - Custom
- Create different types of users (read/write, light/medium/ heavy, etc.)

Adding more memory to VM

- Start with 2 gb of RAM for the VM
- Double till we reach 16 gb

Adding memory to VM

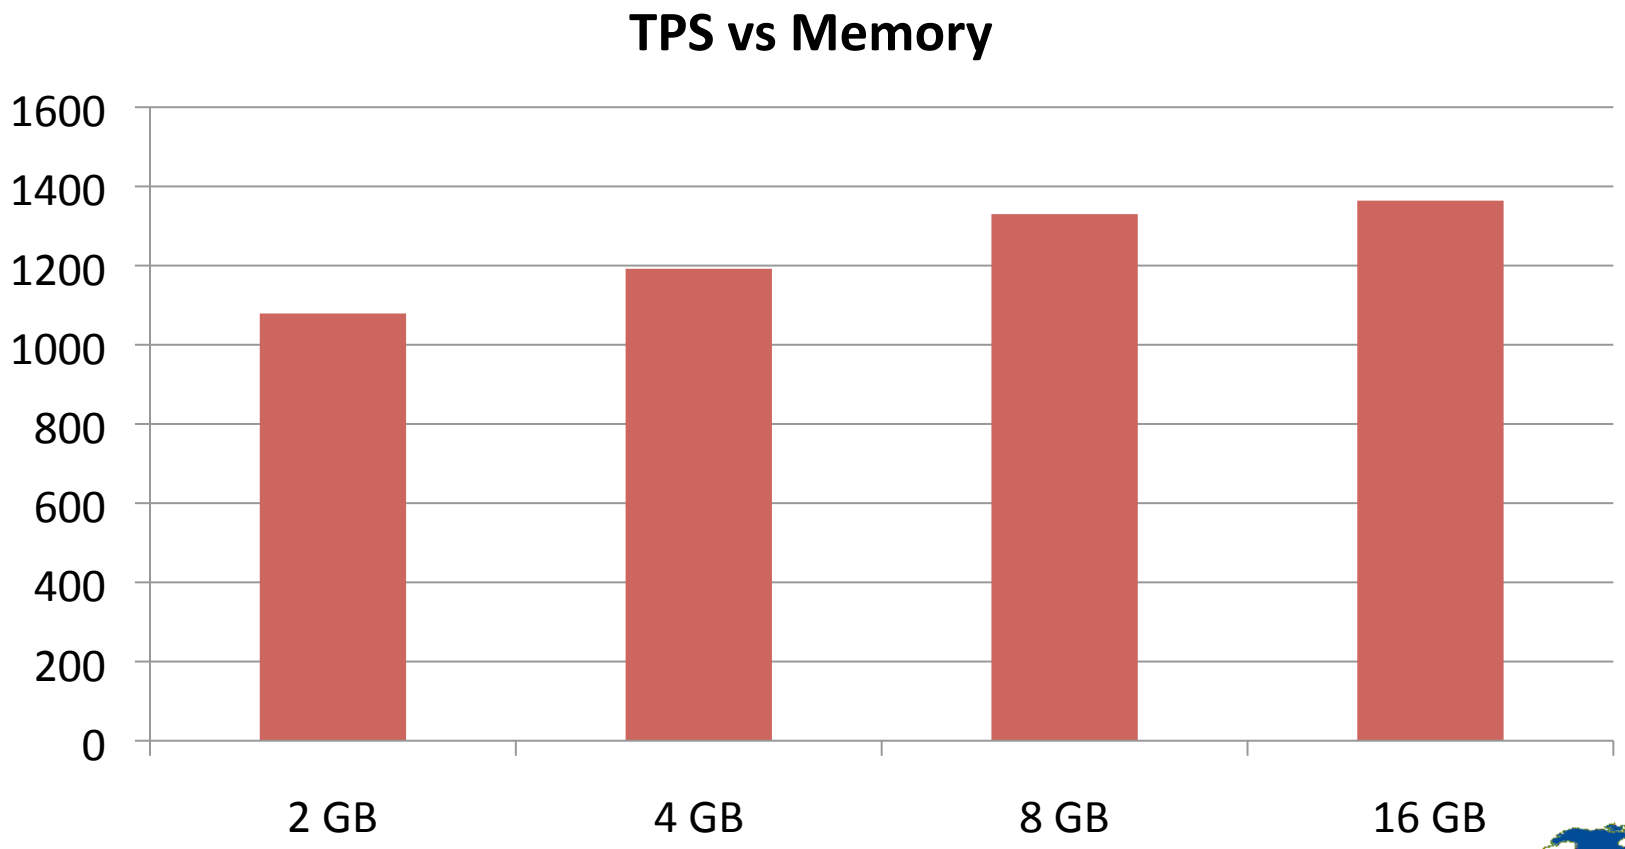

PUG Challenge Americas

QUIZ

Why is difference between 8 and 16 gb so small ?

Encryption (TDE)

- Done in 4 separate measurements
- Baseline with 8 GB of memory for VM
- Step 1: Create policy area and enable encryption
  "Encryption Policy Area":12,64;8 .
- Step 2: Define policies for tables and indexes.
  Only changed and new data will be encrypted
- Step 3: Run utility to encrypt all remaining unencrypted data

to define encryption policies

```
for T in account branch teller history1
do
   proutil atm -C epolicy manage table encrypt ${T}
done
for I in account.account branch.branch teller.teller \
         history1.histid
do
   proutil atm -C epolicy manage index encrypt ${I}
done
```

to encrypt all unencrypted data

```
for T in account branch teller history1
do
    proutil atm -C epolicy manage table update ${T}
done
for I in account.account branch.branch teller.teller \
         history1.histid
do
    proutil atm -C epolicy manage index update ${I}
done
```

encryption results

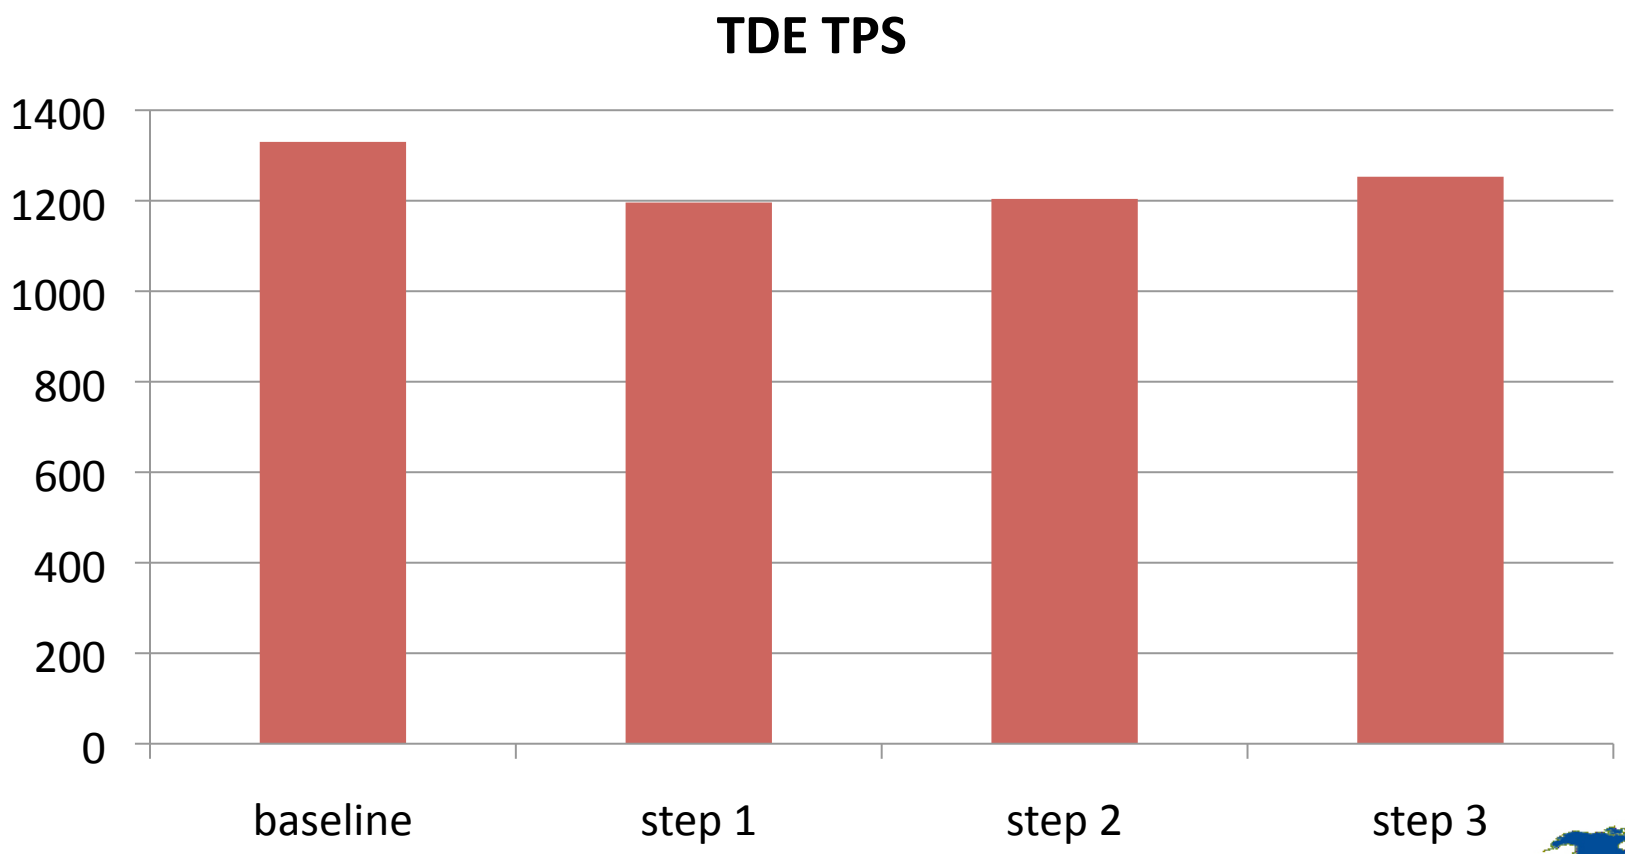

client-server vs self-serving

- Client on 32-bit OpenEdge 10.2B03
  on bunker32 VM
- Server on 64-bit OpenEdge 10.2B03
  on bunker64 VM

client-server results

PUG Challenge Americas

Surprises

- The default -napmax value of 5000 is too large
- More testing of this is required
- YMMV (your mileage may vary) !!!!
  - Transportation, meals, and accomodations not included

-napmax

- Upper bound on doubling sleep time when a shared-memory lock acquisition retry loop fails
- Starts at –nap milliseconds (10)
- 10, 20, 40, 80, 160, 320, 640, 1280, 2560, 5120
- Default value is 5000 milliseconds
  - Hypothesis: this value might be too large

-napmax results

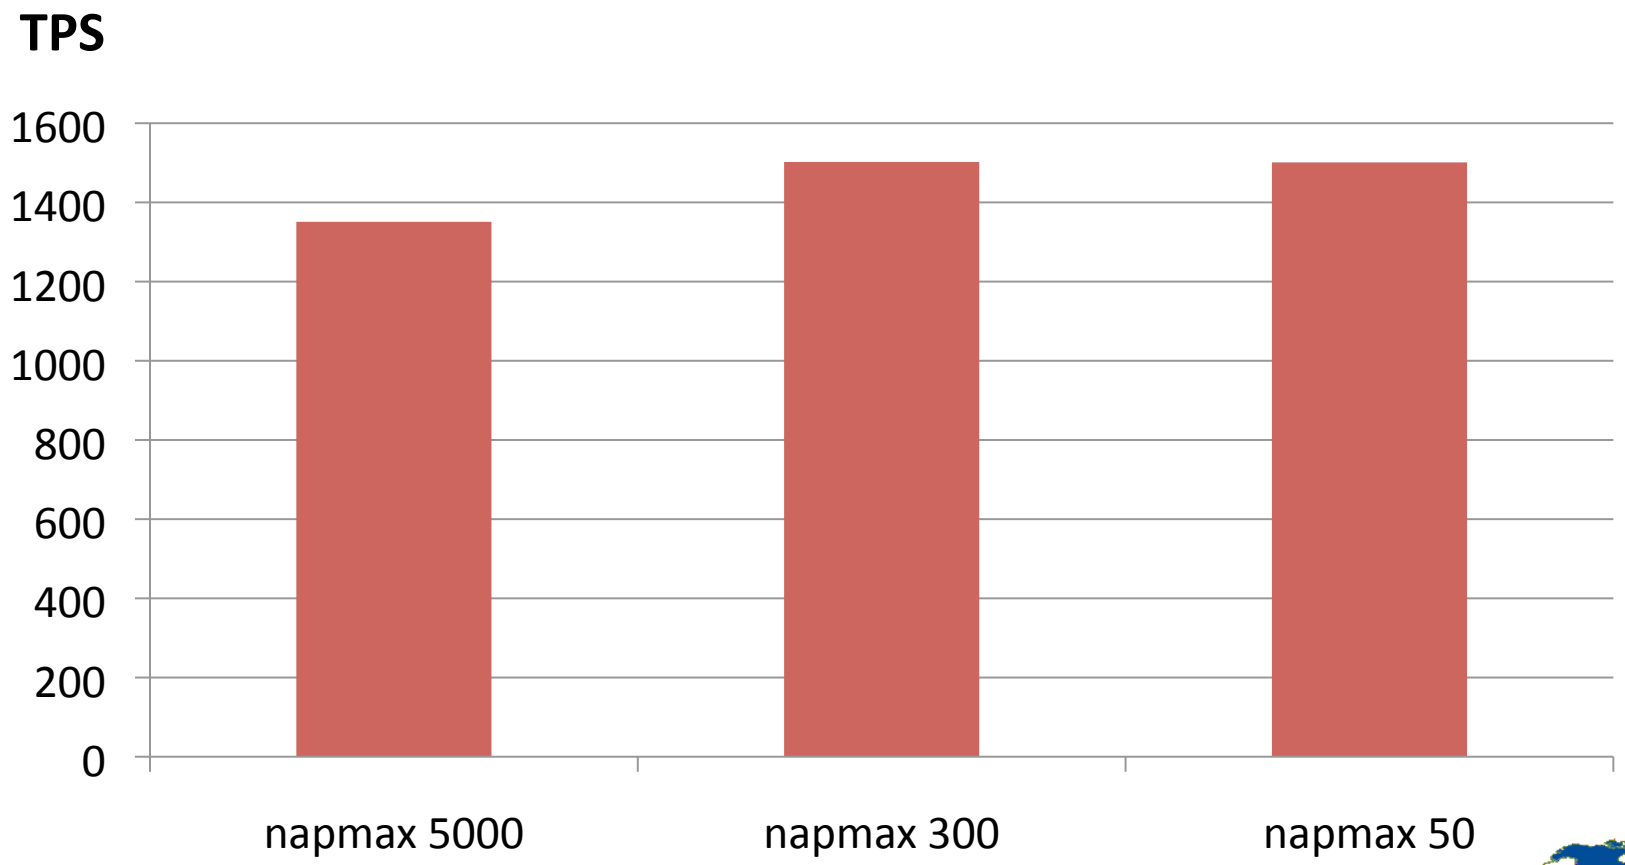

PUG Challenge Americas

Other Surprises

- Adding –directio was a very good thing
  - This is contrary to older Bunker testing
  - Improved TPS (~3000 > 3700+)
  - Reduced Max Response time
  - Huge number of Buffers Flushed (see next slide) using 4 APWs & 16mb Cluster Size)
  - To reduce BF to a reasonable number required a 128mb BI Cluster size & 8 APWs

Surprises - directio

Ckpt No.	Time	Len	Freq	------ Database Writes ------ Dirty	CPT Q	Scan	APW Q	Flushes
92	17:12:05	2267	0	63689	0	2217	76962	0
91	17:11:36	29	29	63138	0	0	35772	14991
90	17:11:08	28	28	63364	0	0	30548	21384
89	17:10:39	29	29	62945	0	0	35525	15534
88	17:10:11	28	28	63167	0	0	31926	19737
87	17:09:42	29	29	62835	0	0	34992	16059
86	17:09:14	28	28	62833	0	0	32153	19588
85	17:08:44	30	30	62569	0	0	35052	16061

Bunker 1-4 & 6 TPS results

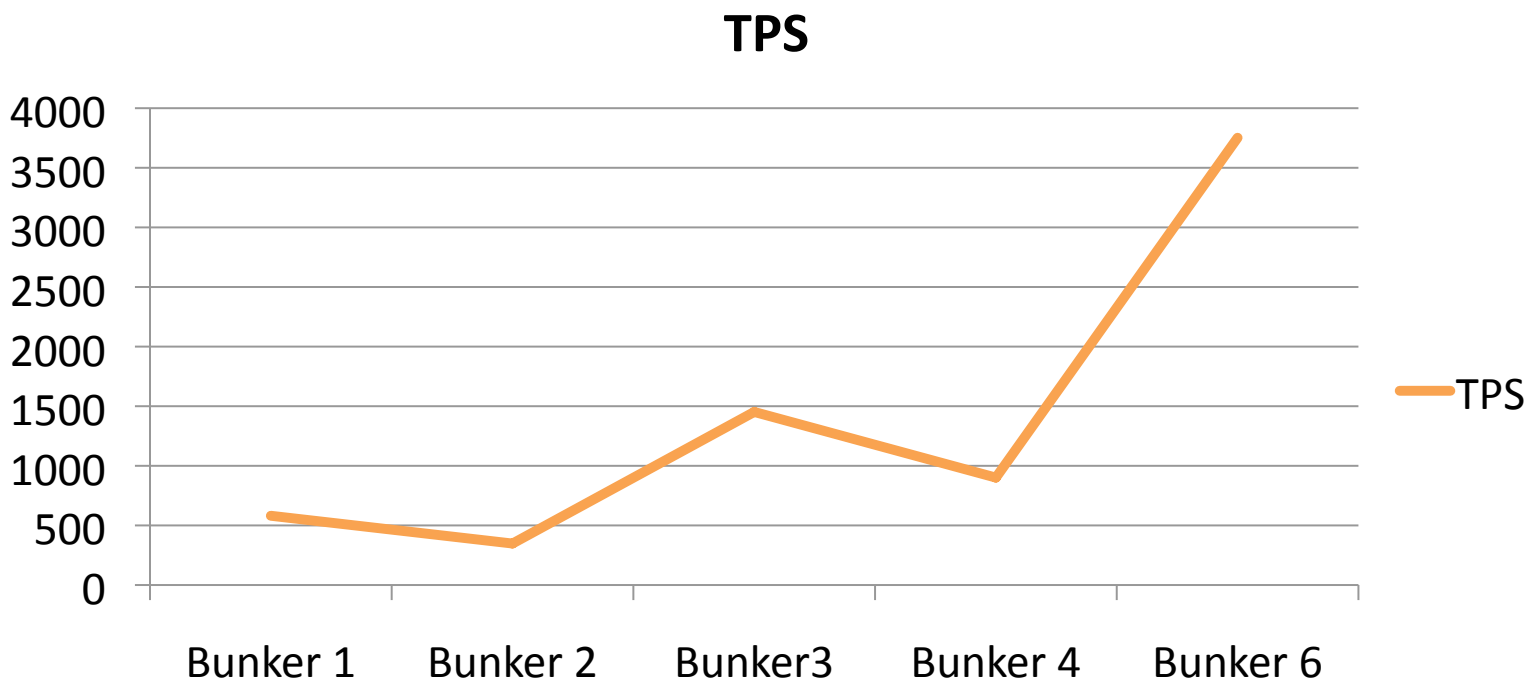

Bunker 1-4 & 6 Max-R results

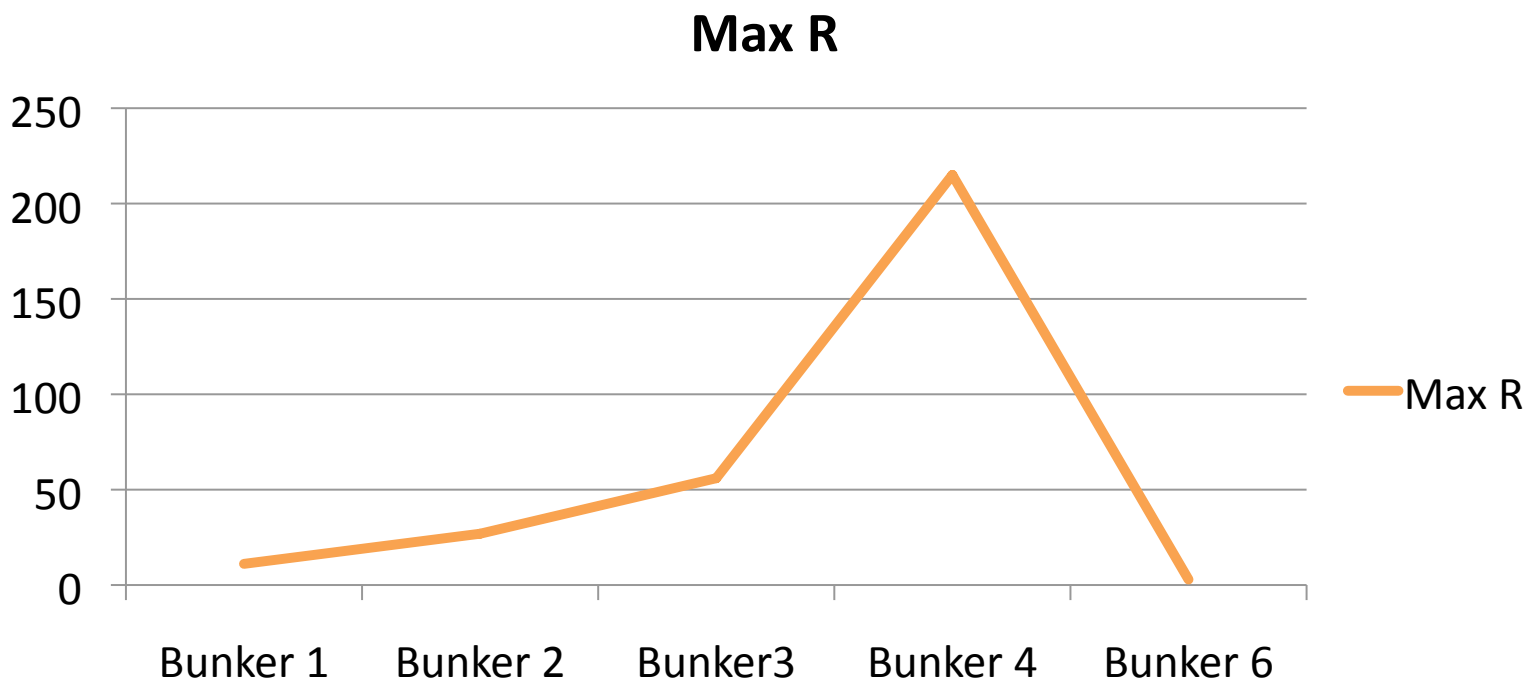

Surprises

- No Integrity (-i) > No Improvement in Performance

No Surprise

- Converting from RAID 5 to RAID 10 improved TPS dramatically
  - John was *slightly* surprised by the magnitude of improvement
  - Gus was offended that we even used RAID 5 since we know it is evil
  - Best RAID 5 TPS was 1502
  - It allowed Dan to set the Bunker TPS record (3766 TPS)
- Deadline scheduler versus CFQ
  - Deadline is still slightly better for database i/o

VM Specific Stuff

- The additional complexity of adding a VM can be astronomical
- There are a huge # of buttons, knobs, and switches
- Example: Caching
  - VMware
  - SAN
  - Centos
  - Progress DB Buffers

This has nothing to do with our topic

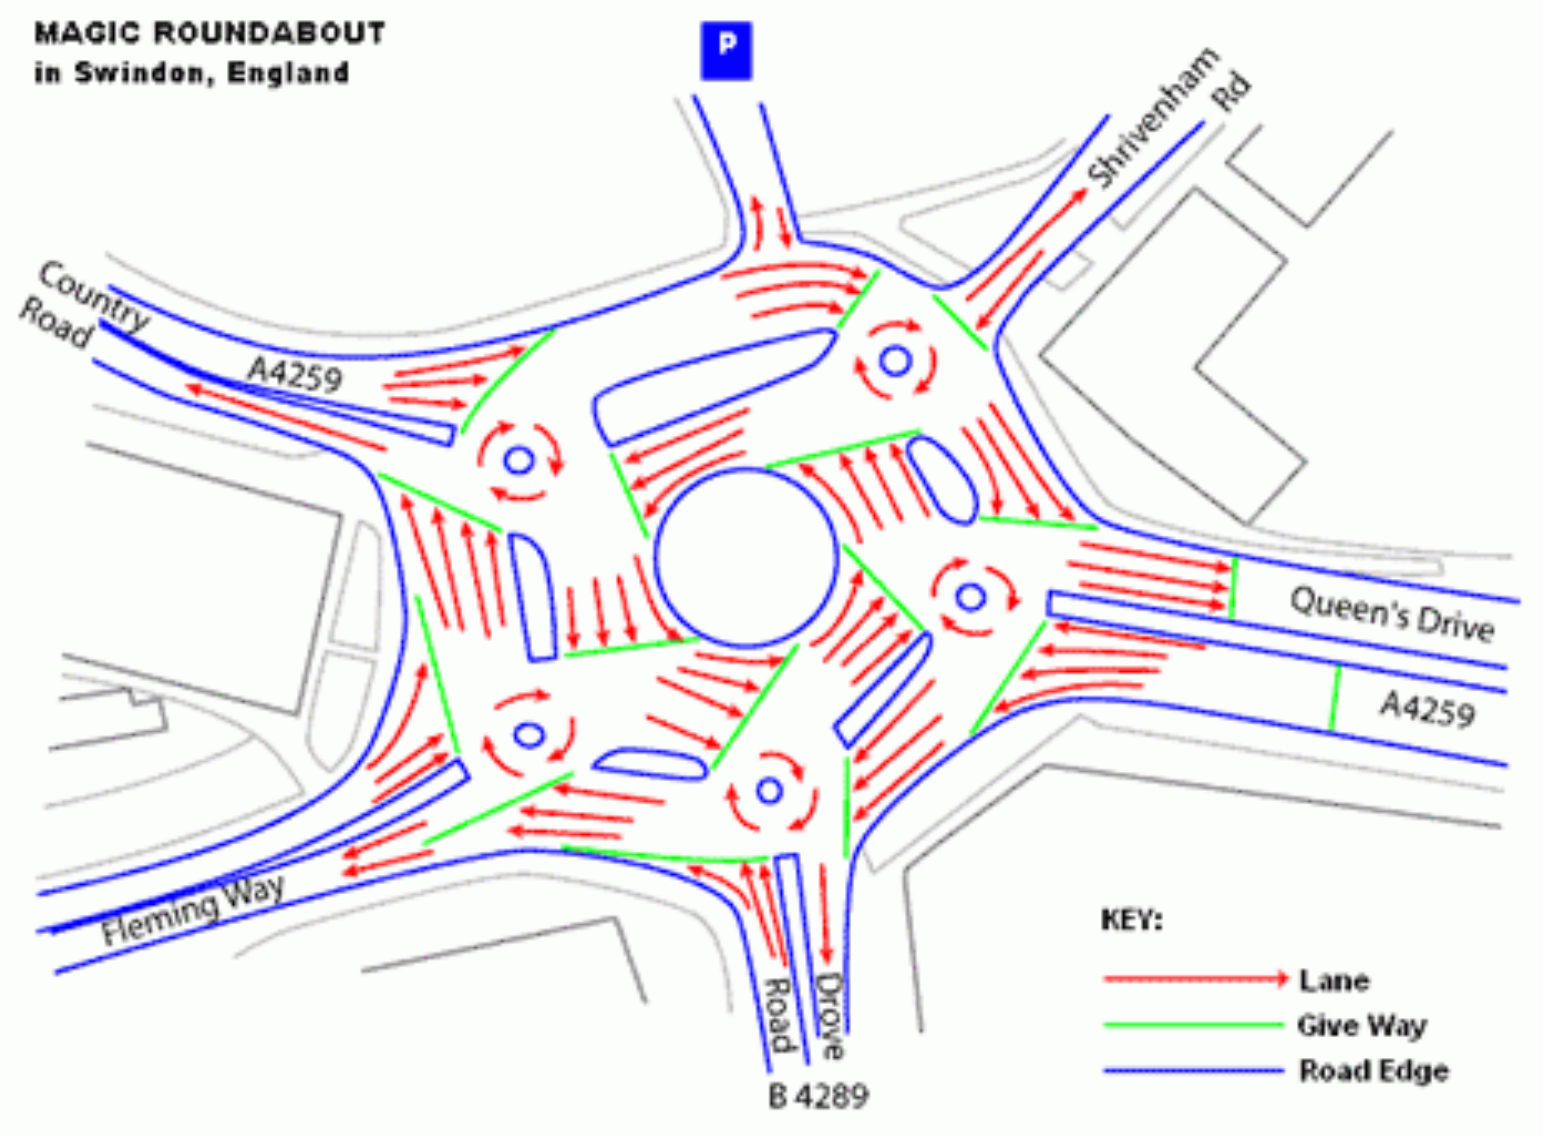

VM Specific Stuff

- Client/Server connections between two VMs on the same box were **much** faster than a Client/Server connection between VMs on separate physical servers
- There was no benefit from changing the VM Swap location
- There was no benefit to using a Private Network
- VM Snap during a test cut TPS almost in half

VM Specific Stuff

- RAM: 2gb > 4gb > 8gb > 16gb
  - No significant change in VM overhead
- A note on this benchmark and CPU intensive tasks
  - The ATM is a database intensive task
  - It doesn't stress the CPUs the way a real world app with java servlets, web servers, file transfers, etc… do

What we didn’t get to

- Auditing
- Alternate Buffer Cache (-B2)
- “Bare metal” versus VmWare
- VMotion overhead

Conclusions

- When you introduce a VM, performance troubleshooting becomes **exponentially** more difficult
- It's like introducing a whole new operating system
- More beer is required

Bunker Links

- http://www.facebook.com/pages/Progress-Bunker-6/204029846278025
- http://www.johnharlow.com/page2/page2.html
- http://www.johnharlow.com/bunkers/Bunker6.html
- Other Secret Bunkers: www.secretbunker.co.uk
- BAARF: www.baarf.com
  - Don't go here if you are a fan of RAID 5
- Awesome Japanese Food: www.shoyaatlanta.com
  - Endorsed by Dan who lived in Japan for 5 years

any questions

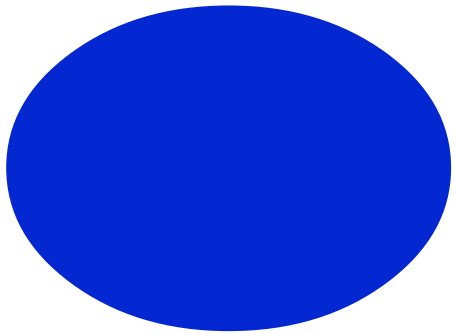

PUG Challenge Americas