



# Infrastructure Tuning

For SQL Server Performance

SQL PASS Performance Virtual Chapter

2014.07.24

## About David Klee



@kleegeek  
davidklee.net  
gplus.to/kleegeek  
linked.com/a/davidaklee

### Specialties / Focus Areas / Passions:

- Performance Tuning & Troubleshooting
- Virtualization
- Cloud Enablement
- Infrastructure Architecture
- Capacity Management
- High Availability
- Disaster Recovery
- Health Monitoring
- **PASS VC Chapters**
  - **Performance**
  - **Virtualization**
  - **HA & DR**



HX  
TECH

***“It’s running **slow** but  
nothing changed on my  
SQL Server!”***

.... becomes ...

***“How can I tell if my  
infrastructure **is**  
slowing it down?”***

© 2014 Heraflux Technologies

3

HX  
TECH

## Session Agenda

- Where to start?
- Bottleneck analysis
- Storage & interconnects
- Compute hardware
- Virtualization
- Operating system
- SQL Server instance and database
- Networking
- Application

© 2014 Heraflux Technologies

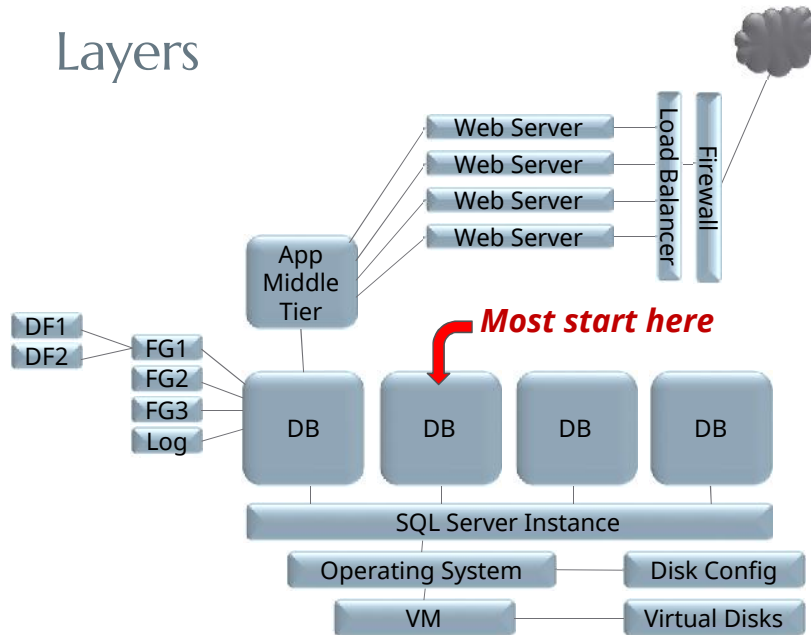
4

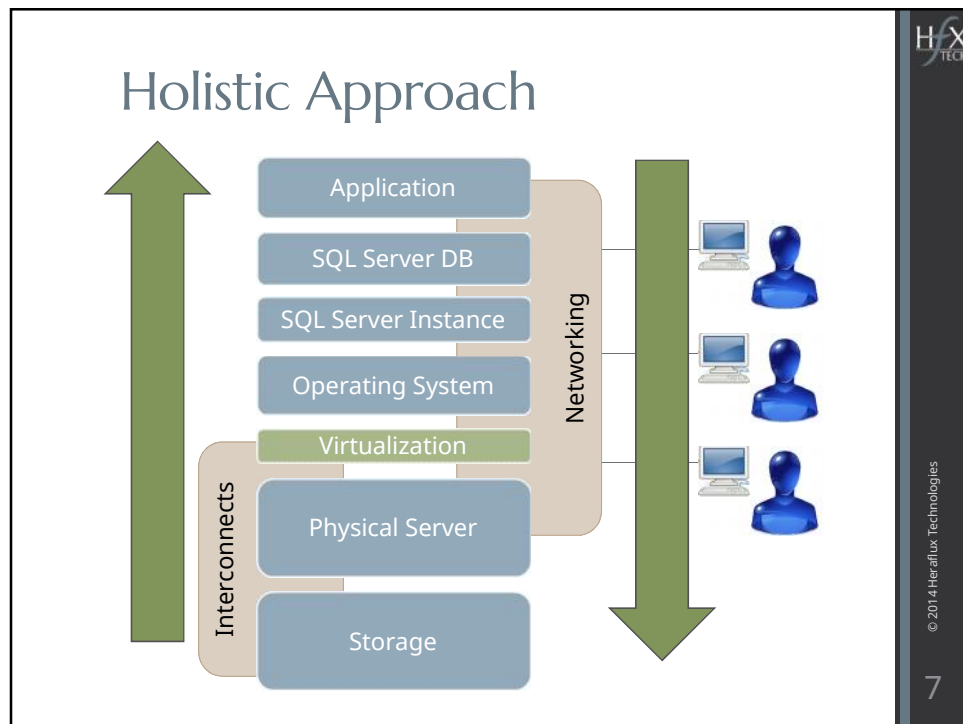
## Let's troubleshoot this!

- Our business critical application is running **80x** slower than yesterday, and our end users are *complaining*.
- **NOTHING** changed on my production SQL Server.
- *What do you do?*
- Most dive into queries.
- Start holistically before you dive into the queries...
- ...because....




## Layers





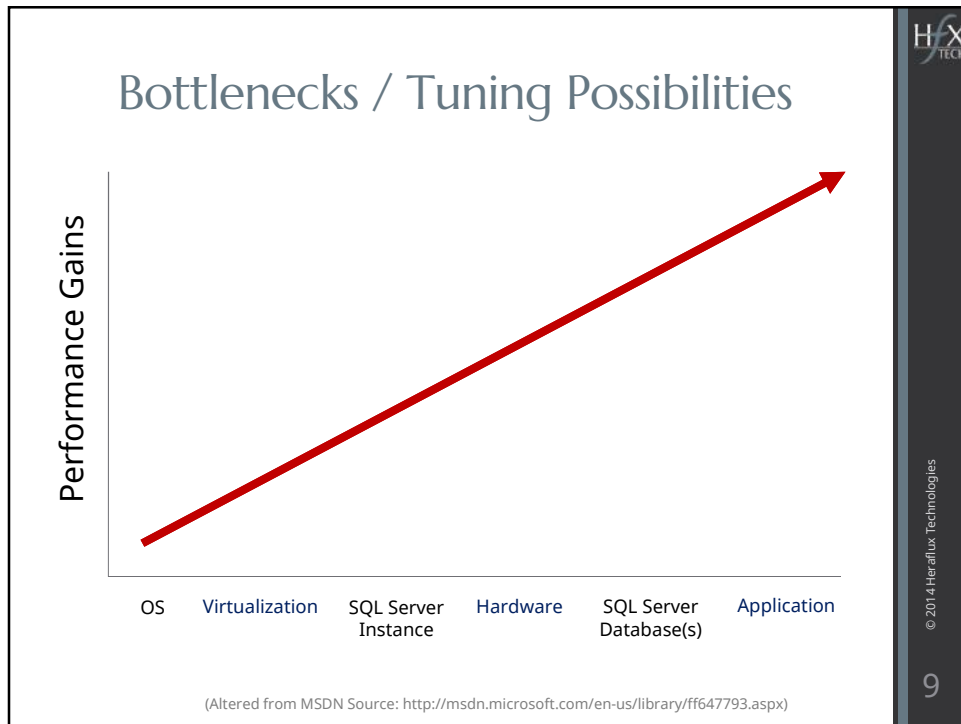
## Bottleneck Analysis

- SQL Server tells you what it's waiting on!
- Use SQL Server wait statistics as a guide
- Don't jump on any one statistic
- "Shift the bottleneck"
- Brent Ozar - sp\_Blitz@
  - [www.brentozar.com/blitz](http://www.brentozar.com/blitz)
- Glenn Berry diagnostic scripts available at:
  - [www.sqlskills.com/blogs/glenn](http://www.sqlskills.com/blogs/glenn)



© 2014 Heraflux Technologies

8



## Benchmarks & Baselines

- It always comes back to benchmarking and baselines...
- Methods:
  - Third party tools
  - Do it yourself solutions
- Infrastructure baselines:
  - Storage performance per test type
  - Network throughput
  - CPU test
- SQL Server baselines:
  - Long-running queries
  - Reports
  - ETL / backups / maintenance

H X TECH  
© 2014 Heraflux Technologies  
10

# SQL Server Waits

- Start with SQL Server wait statistics analysis
- Paul Randal – “Tell me where it hurts” @ [bit.ly/1mibFXL](http://bit.ly/1mibFXL)
- Tom Davidson – Waits and queues whitepaper @ [bit.ly/1e1I38f](http://bit.ly/1e1I38f)
- Enrico van de Laar - Wait stat collector @ [bit.ly/1jVqHmF](http://bit.ly/1jVqHmF)
- Glenn Berry diag query – top waits



wait_type	wait_time_s	pct	running_pct
TRACEWRITE	610423.66	87.85	87.85
CXPACKET	20807.88	2.99	90.84
OLEDB	14058.50	2.02	92.86
BACKUPIO	10002.90	1.44	94.30
PAGEIOLATCH_SH	8806.83	1.27	95.57
ASYNC_IO_COMPLETION	7687.74	1.11	96.68
BACKUPTHREAD	4485.16	0.65	97.33
BACKUPBUFFER	4374.99	0.63	97.96
PREEMPTIVE_OS_WRITEFILE	2758.70	0.40	98.36
WRITELOG	2099.69	0.30	98.66
MSQL_XP	2118.21	0.30	98.96
BROKER_RECEIVE_WAITFOR	1800.00	0.26	99.22



# Storage

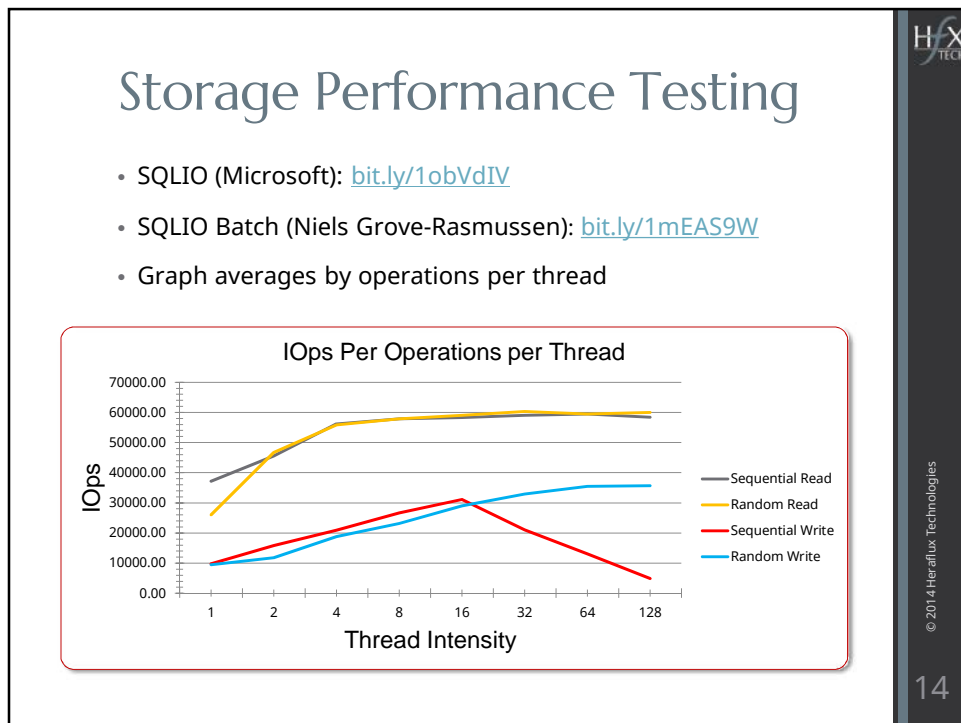
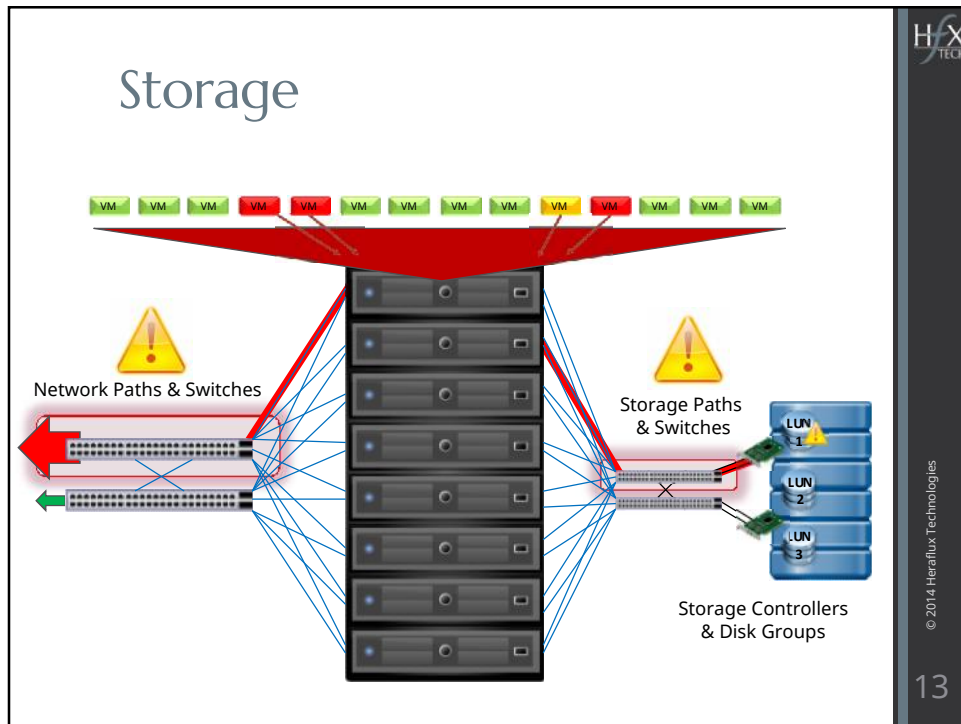
- Usually the slowest and least responsive part of the stack
- Perfmon – Physical Disk counters
  - Disk sec/Read & Write = latency in ms
  - Disk Reads/sec & Writes/sec = IOPs
- sys.dm\_io\_virtual\_file\_stats
  - Cumulative since the last service startup
- Check the error log

```

4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1140 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev14.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1833 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev18.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1833 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev18.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1846 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev17.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1927 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev16.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1940 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev18.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1944 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev12.nd] in database [tempdb] (2). The C
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1946 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev13.nd] in database [tempdb] (2). The C
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1983 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev1_1.nd] in database [tempdb] (2). Th
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1983 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev15.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1984 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev18.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1983 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev11.nd] in database [tempdb] (2). The
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1985 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev6.nd] in database [tempdb] (2). The D
4/15/2009 10:44:52 AM spid0s SQL Server has encountered 1985 occurrence(s) of I/O requests taking longer than 15 seconds to complete on file [C:\Writes\VD\tempdev6.nd] in database [tempdb] (2). The D
    
```

(Image from Henk van der Valk @ [bit.ly/1qA6Qwk](http://bit.ly/1qA6Qwk))



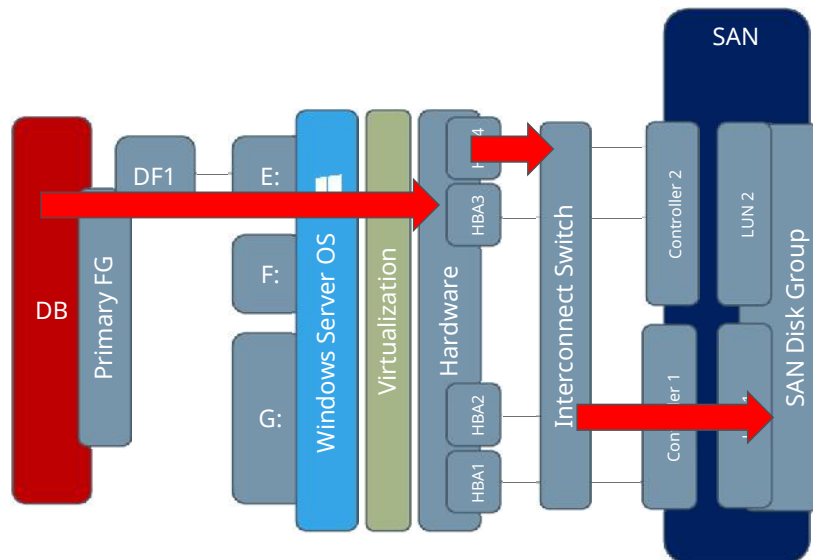


## Improve Storage Performance

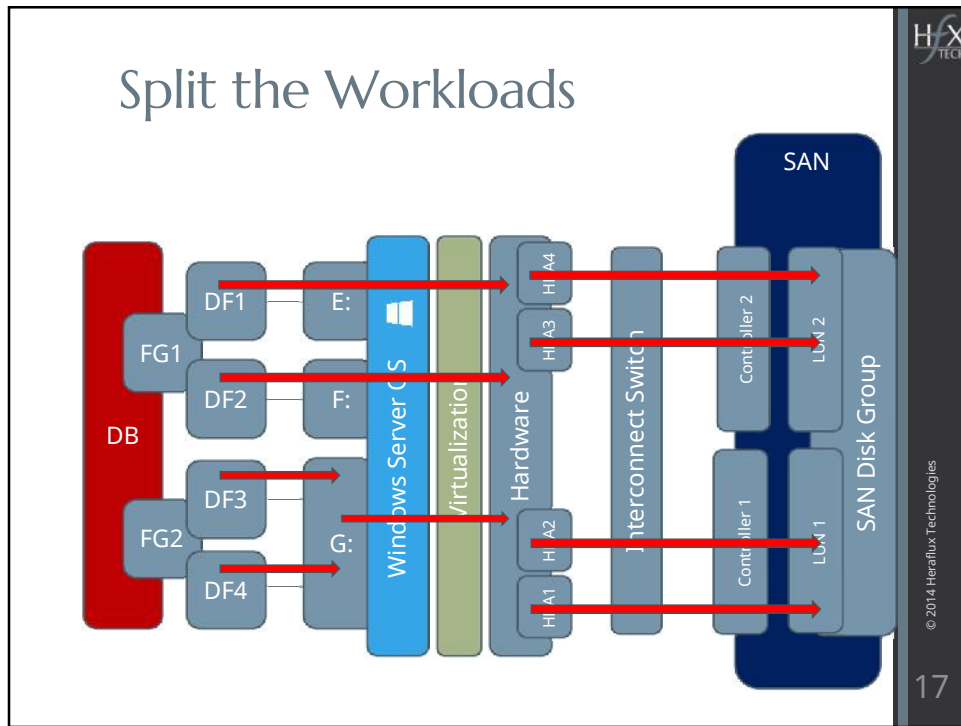
- Work with storage admins
- Add more SQL Server memory
- Use storage interconnect multipathing
- Faster individual interconnects
- Block size alignment with SAN / architecture
- HBA queue depth settings
- Improve disk configuration for performance
  - Faster RAID configuration
  - SSD caching / tier
  - Replace storage with faster array
- Spread out workload



## Split the Workloads







### Compute Hardware

- Have you overwhelmed the hardware itself?
- CPU?
- Memory?

The image shows a server rack on the left and a screenshot of the Windows Task Manager Performance tab on the right. The screenshot displays various system metrics:

- CPU Usage:** 100% (indicated by a full green bar chart).
- Memory:** 47.3 GB used.
- Physical Memory Usage History:** A line graph showing memory usage over time.
- Physical Memory Usage Error:** A line graph showing memory usage errors over time.
- System Statistics:**
  - Physical Memory (MB): Total 79161, Available 582, Free 22.
  - System: Handles 1,132, Threads 1871, Processes 135, Up Time 9:23:42+5, Commit (KB) 1,775b.
- Kernel Memory (MB):** Paged 199, Nonpaged 79.


At the bottom of the screenshot, it shows: Processes: 135, CPU Usage: 100%, Physical Memory: 90%.

© 2014 Heraflux Technologies

18

## High CPU Consumption?

- Waits - CXPacket, SOS\_Scheduler\_Yield, ThreadPool, LCK\_XX
- MaxDOP
- Cost Threshold for Parallelism
- Not enough CPUs?
- Is it really SQL Server?
- Virtualization CPU scheduling & background contention



## Out of Memory?

- Waits - Resource\_Semaphore
- Low page life expectancy & low buffer cache hit ratio
- Very high I/O as a result
- High CPU consumption
- Virtual machine memory ballooning
- RAM is cheap!
- Artificial licensing limitations are not...



## Virtualization

- Background activity on the same host can cause performance variations

40 CPU  
512GB RAM

© 2014 Heraflux Technologies

21

## VM Layer Statistics

- Get access! Host and in-guest-level metrics. *Read-only!*

CPU / Real-time, 7/8/2014 8:19:17 PM - 7/8/2014 9:19:17 PM | Clear | Overview

Graph refreshes every 20 seconds

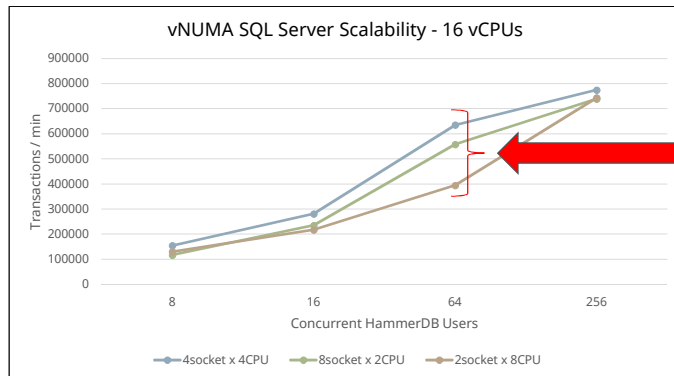
Status bar: CPU

© 2014 Heraflux Technologies

22

## vNUMA Configuration

- Validate your VM-level CPU configuration
- Generate a HammerDB or sample workload that best represents your average SQL Server workload
- Adjust vNUMA config / vCPU counts & retest



**31%  
Difference!**

## Operating System

- Background tasks
  - Antivirus real-time scanning
  - Sysadmin tools
- OS-level disk and I/O queueing
  - More disks / drive letters / mount points
  - Spread out the workload
- Licensing limitations
  - Artificial limits on CPU & memory utilization



## SQL Server Instance

- Reduce I/O
  - Backup compression
  - Instant file initialization
- Multiple tempdb data files
- Large memory pages
- Min and max memory
- SQL Server 2014 Buffer Pool Extensions
- CPU Affinity
- MaxDOP & Cost Threshold for Parallelism
- Stagger job timings



## SQL Server Database

- Spread out the workload
  - File groups, data files, partitions
- Slow I/O? Reduce footprint with table & index compression
- In-memory constructs such as In-Memory OLTP and updatable Columnstore indexing
- Workload analysis for Cost Threshold value
  - @sqlpoolboy script at [bit.ly/1nXrxin](http://bit.ly/1nXrxin)
- SSD read caching



## Application & Networking

- Waits - ASYNC\_NETWORK\_IO
- Application might not be able to consume data quickly
- OR
- Application has to communicate with the DB server
- End-user application has to communicate to the app server or DB
- Validate network throughput and stability between application and database servers



## Network Performance Test

- Verify and test end-to-end performance
- Use *iperf3* to test network throughput
  - Download @ [bit.ly/1zuyRy](http://bit.ly/1zuyRy)
  - How-To Guide: [tinyurl.com/ost9gll](http://tinyurl.com/ost9gll)



```
iperf2 on
File View VM
Mikael@iperf2:~$ iperf -s
Server listening on TCP port 5001
TCP window size: 65.5 KByte (default)
[ 0] local 192.168.31.113 port 5001 connected with 192.168.31.113
[ ID] Interval      Transfer     Bandwidth
[ 0] 0.0-10.0 sec  1.10 GBytes  941 Mbits/sec
Mikael@iperf2:~$




iperf3 on
File View VM
Mikael@iperf3:~$ iperf -c 192.168.31.113 -l 10
Client connecting to 192.168.31.113, TCP port 5001
TCP window size: 65.5 KByte (default)
[ 0] local 192.168.31.113 port 5001 connected with 192.168.31.113 port 5001
[ ID] Interval      Transfer     Bandwidth
[ 0] 0.0-10.0 sec  1.10 GBytes  941 Mbits/sec
Mikael@iperf3:~$
```

## Conclusions

- SQL Server performance *might* not be a code or query problem
- Opportunities to tune the infrastructure are all over
- Use the tools that you have available to you to find the infrastructure bottlenecks to pursue
- Shift the bottleneck until it's acceptable



## Questions?

 @kleegeek  
 davidklee.net  
 gplus.to/kleegeek  
 linked.com/a/davidaklee

