Azure Data Platform Performance Fun

How to run SQL Server @ Speed in the Cloud

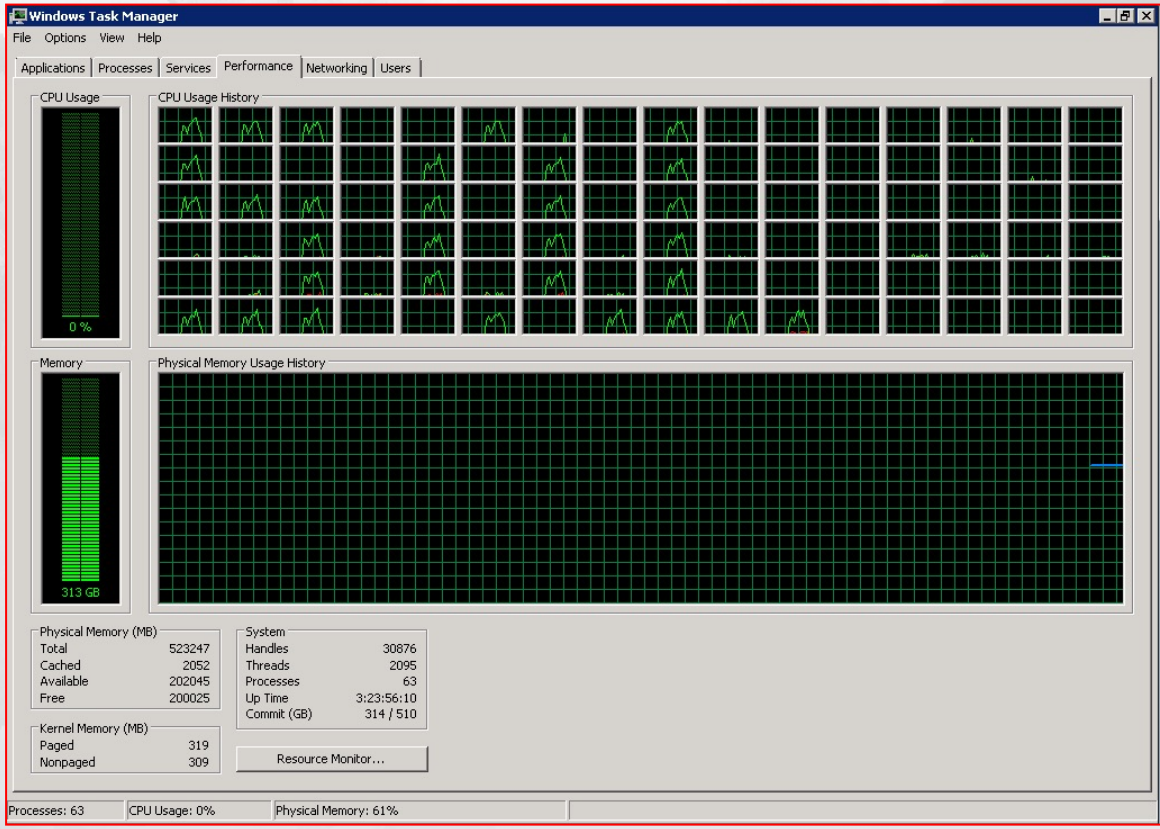

May 28th, 2021

henk@fulldata.nl

Henk van der Valk
@HenkvanderValk

FULLDATA
THE DATA PLATFORM SPECIALISTS

Agenda

Flashback!

- Recent SQL Benchmarking
  - Table scan speed from disk
  - Table scan speed In-Memory
  - Bulk Inserting from flat files!
- Backup / Restore
- Tip! Smart Bulk Copy
- Capacity Management
- Storage comparison
  - Azure Premium SSD (20 x 1.01 TB)
  - Azure Ultra disks (5x1TB)
  - Silk Cloud Data Platform Storage

FULLDATA
THE DATA PLATFORM SPECIALISTS

Speaker Introduction

- 15+ years active in SQL Pass community
- 10 years of Unisys-EMEA Performance Center
- 2002- Largest SQL DWH in the world (SQL2000)
  - **Project Real – (SQL 2005)**
- ETL WR - loading 1TB within 30 mins (SQL 2008)
- SQL Fasttrack & data loading performance guide whitepapers
- Perf Tips & tricks: www.henkvandervalk.com
- Schuberg Philis- 100% uptime for mission critical apps
- Since April 1st, 2011 – Microsoft SQL PDW/APS
- Since 2017- **Founder FullData!**

1999 - SQL2000 (SP1)

- 32Way / 64 GB RAM
- 32bit PIII - 550 MHz
- 64bit Itanium 800 MHz CPU's

```
sp_configure 'awe enabled', 1
RECONFIGURE
GO
```

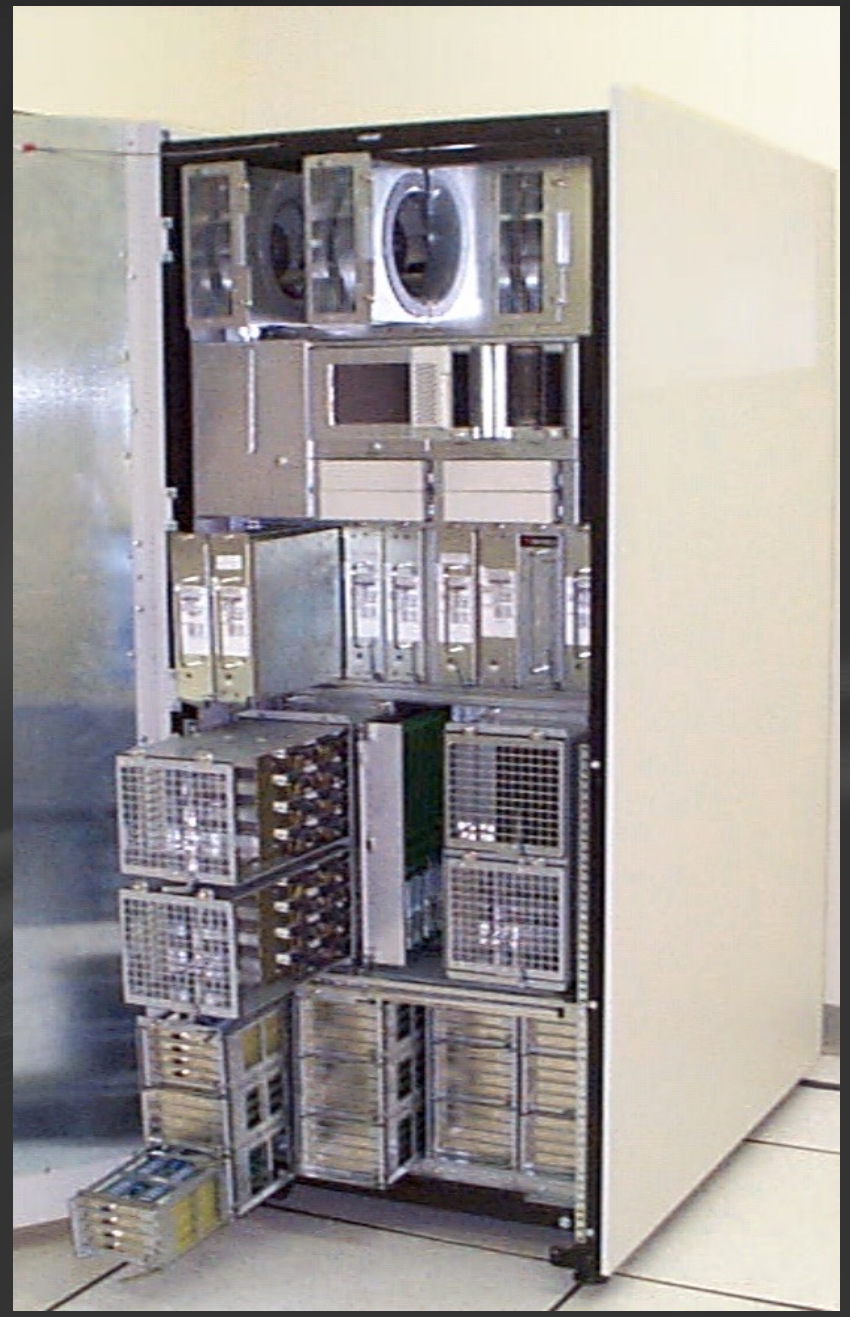

Big Windows / Wintel Mainframe Architecture

- 32 CPU's
- 64 GB RAM
- 96 PCI slots

Example :

Sequential Read throughput

- ES7000/100
- OS: Windows 2000 Datacenter
- 42 x 1Gbit HBA's
- EMC Clariion's with 250 physical drives

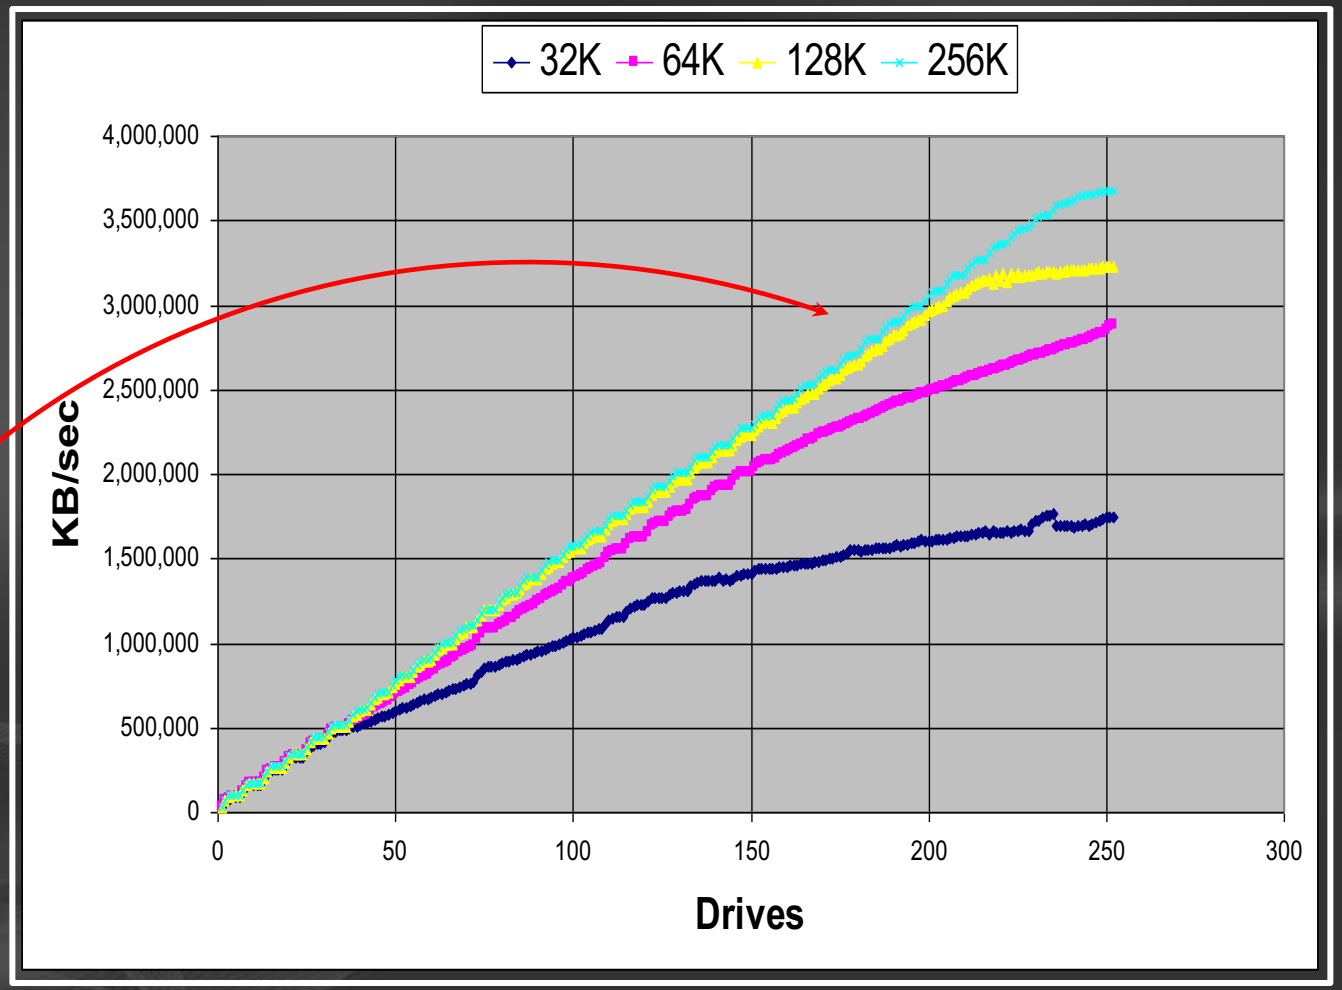

Partner Award winning solution on SQL2000/ **32-Bit**

Unisys digital Interactive Archive System (DIAS)

- 22+ Billion records
- 19 TB SQL Data + 19 TB images
- Redundant / Geo locations

Today:

- 30+ TB SQL Data
- 30+ TB images

W I N T E R C O R P

Database Size – All Environments – Data Warehousing

Company/ Organization	DB Size (GB)	Platform	DBMS	Architecture	DBMS Vendor	System Vendor	Storage Vendor
Yahoo!	100,386	Unix	Oracle	Centralized/SMP	Oracle	Fujitsu Siemens	EMC
AT&T	93,876	Unix	Daytona	Federated/SMP	AT&T	HP	HP
KT IT-Group	49,397	Unix	DB2	Centralized/Cluster	IBM	IBM	Hitachi
AT&T	26,713	Unix	Daytona	Federated/SMP	AT&T	Sun	Sun
LGR – Cingular Wireless	25,203	Unix	Oracle	Centralized/SMP	Oracle	HP	HP
Amazon.com	24,773	Linux	Oracle RAC	Centralized/Cluster	Oracle	HP	HP
Anonymous	19,654	Unix	DB2	Centralized/MPP	IBM	IBM	EMC
UPSS	19,467	Windows	SQL Server	Centralized/SMP	Microsoft	Unisys	EMC
Amazon.com	18,558	Linux	Oracle RAC	Centralized/Cluster	Oracle	HP	HP
Nielsen Media Research	17,685	Unix	Sybase IQ	Centralized/SMP	Sybase	Sun	EMC

2005 Top Ten Program

SQL Server Integration Services Record ETL Performance

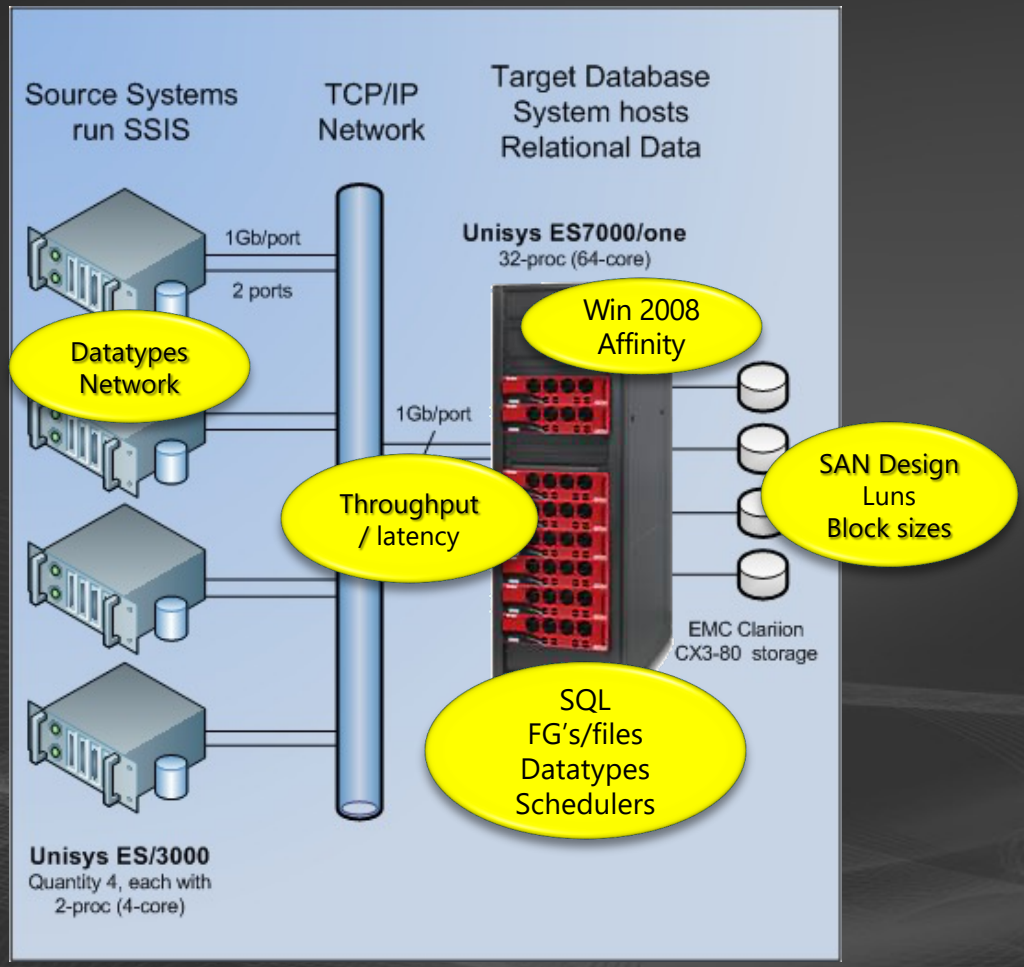

1.18 TB of data loaded in 29 min 54 sec

- Data model reflects a wholesale supplier data warehouse
- Data is read from text files, sent over network, and stored in a single database image
- **8.5+ Billion records**
- SSIS runs on source systems, reflecting a distributed environment
- Built upon
  - SQL Server 2008 Ent. Ed. February CTP
  - Windows Server 2008 Datacenter Ed.

Microsoft® SQL Server™ 2008

TPC-H Line item data Is most challenging to process

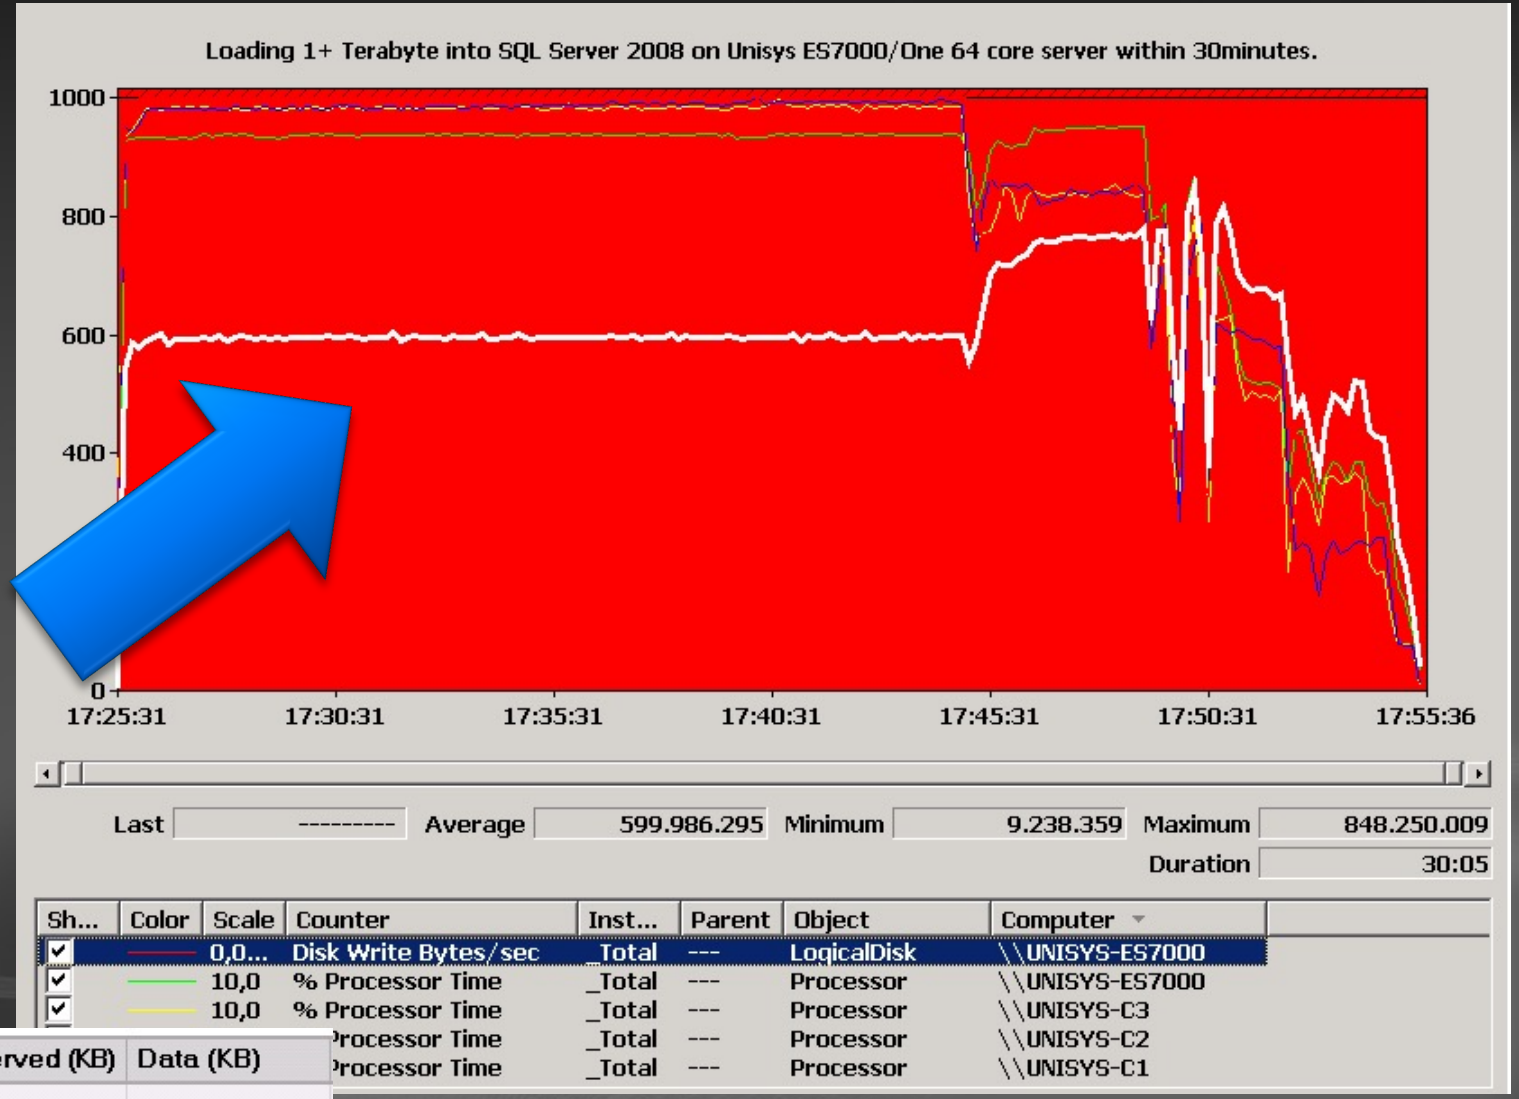

Table Name	# Records	Reserved (KB)	Data (KB)
dbo.LINEITEM	5,999,989,709	679,400,512	679,394,408
dbo.ORDERS	1,500,000,000	182,977,472	182,972,144
dbo.PARTSUPP	800,000,000	125,819,072	125,813,888
dbo.PART	200,000,000	29,367,808	29,361,840
dbo.CUSTOMER	150,000,000	26,157,248	26,152,152
dbo.SUPPLIER	10,000,000	1,598,336	1,592,608

Enterprise Class - High Speed Solid State Disk Storage

DSI3500

Solid State Disk System
Enterprise High Speed Disk

Quick Specs

- RAID3 Protection
- Up to 2TB Flash Storage
- 100,000 I/Os per second

- Up to 8 - 4Gb Fibre Channel Parts
- Hot Swap Architecture
- Chipkill Technology

DSI3600

Solid State Disk System
Enterprise High Speed Disk

Quick Specs

- Board Level RAID
- Up to 5TB Flash Storage
- 250,000 I/Os per second

- Up to 8 - 4Gb Fibre Channel Ports
- Hot Swap Architecture
- 3 GB/sec of bandwidth

Up to 8 Fibers each

SQL2008R2 on 96 Cores

Bulk Insert rate up from 4 million rows/sec (ETL WR)

to 15+ million rows/sec

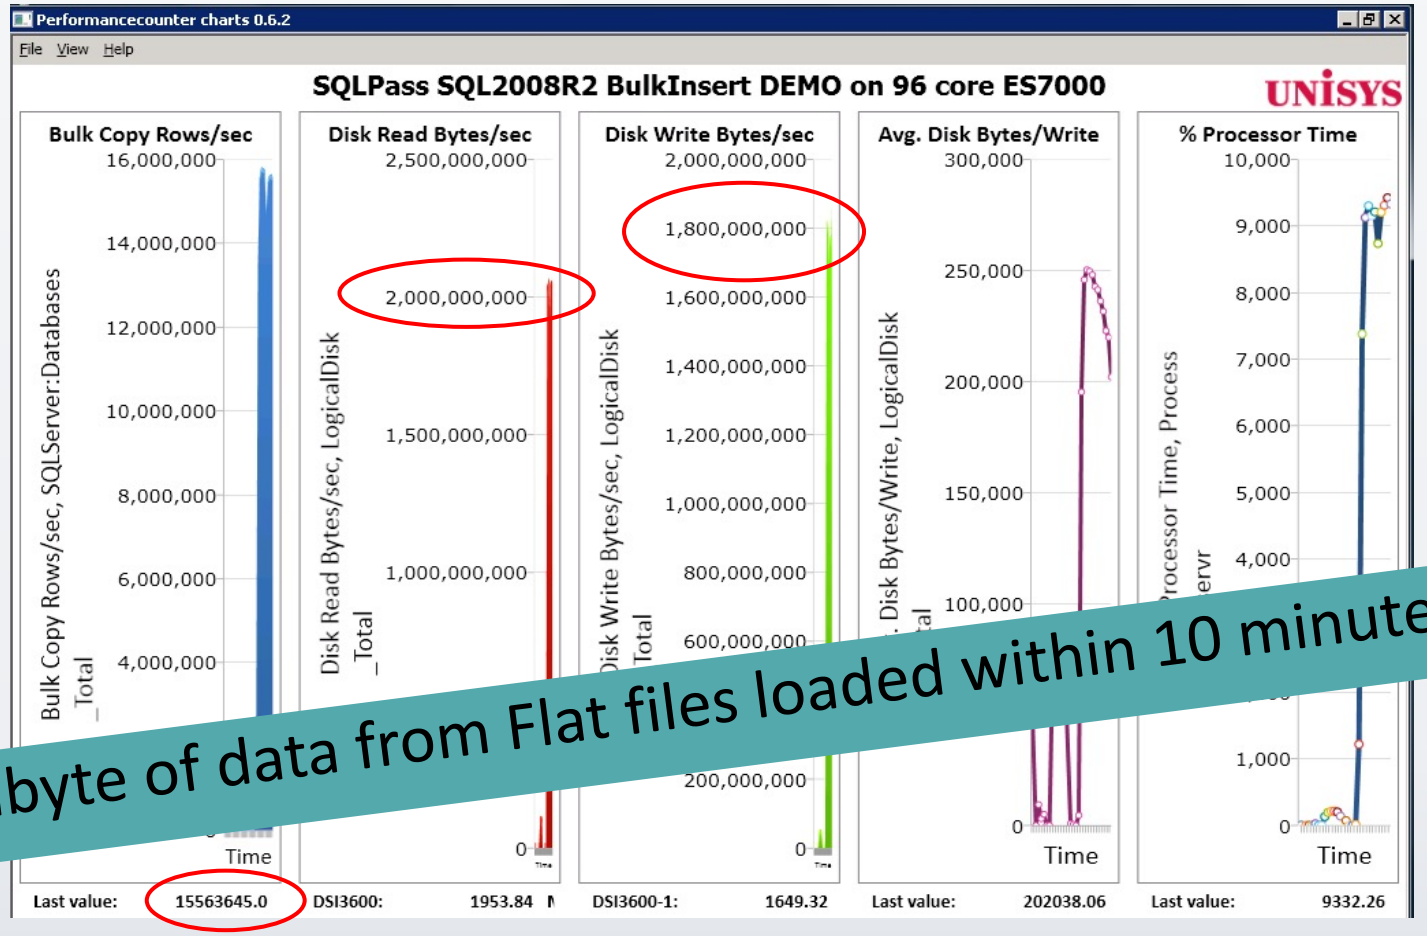

1 Terabyte of data from Flat files loaded within 10 minutes!

PDW: Querying 1 Petabyte of data in 1 second

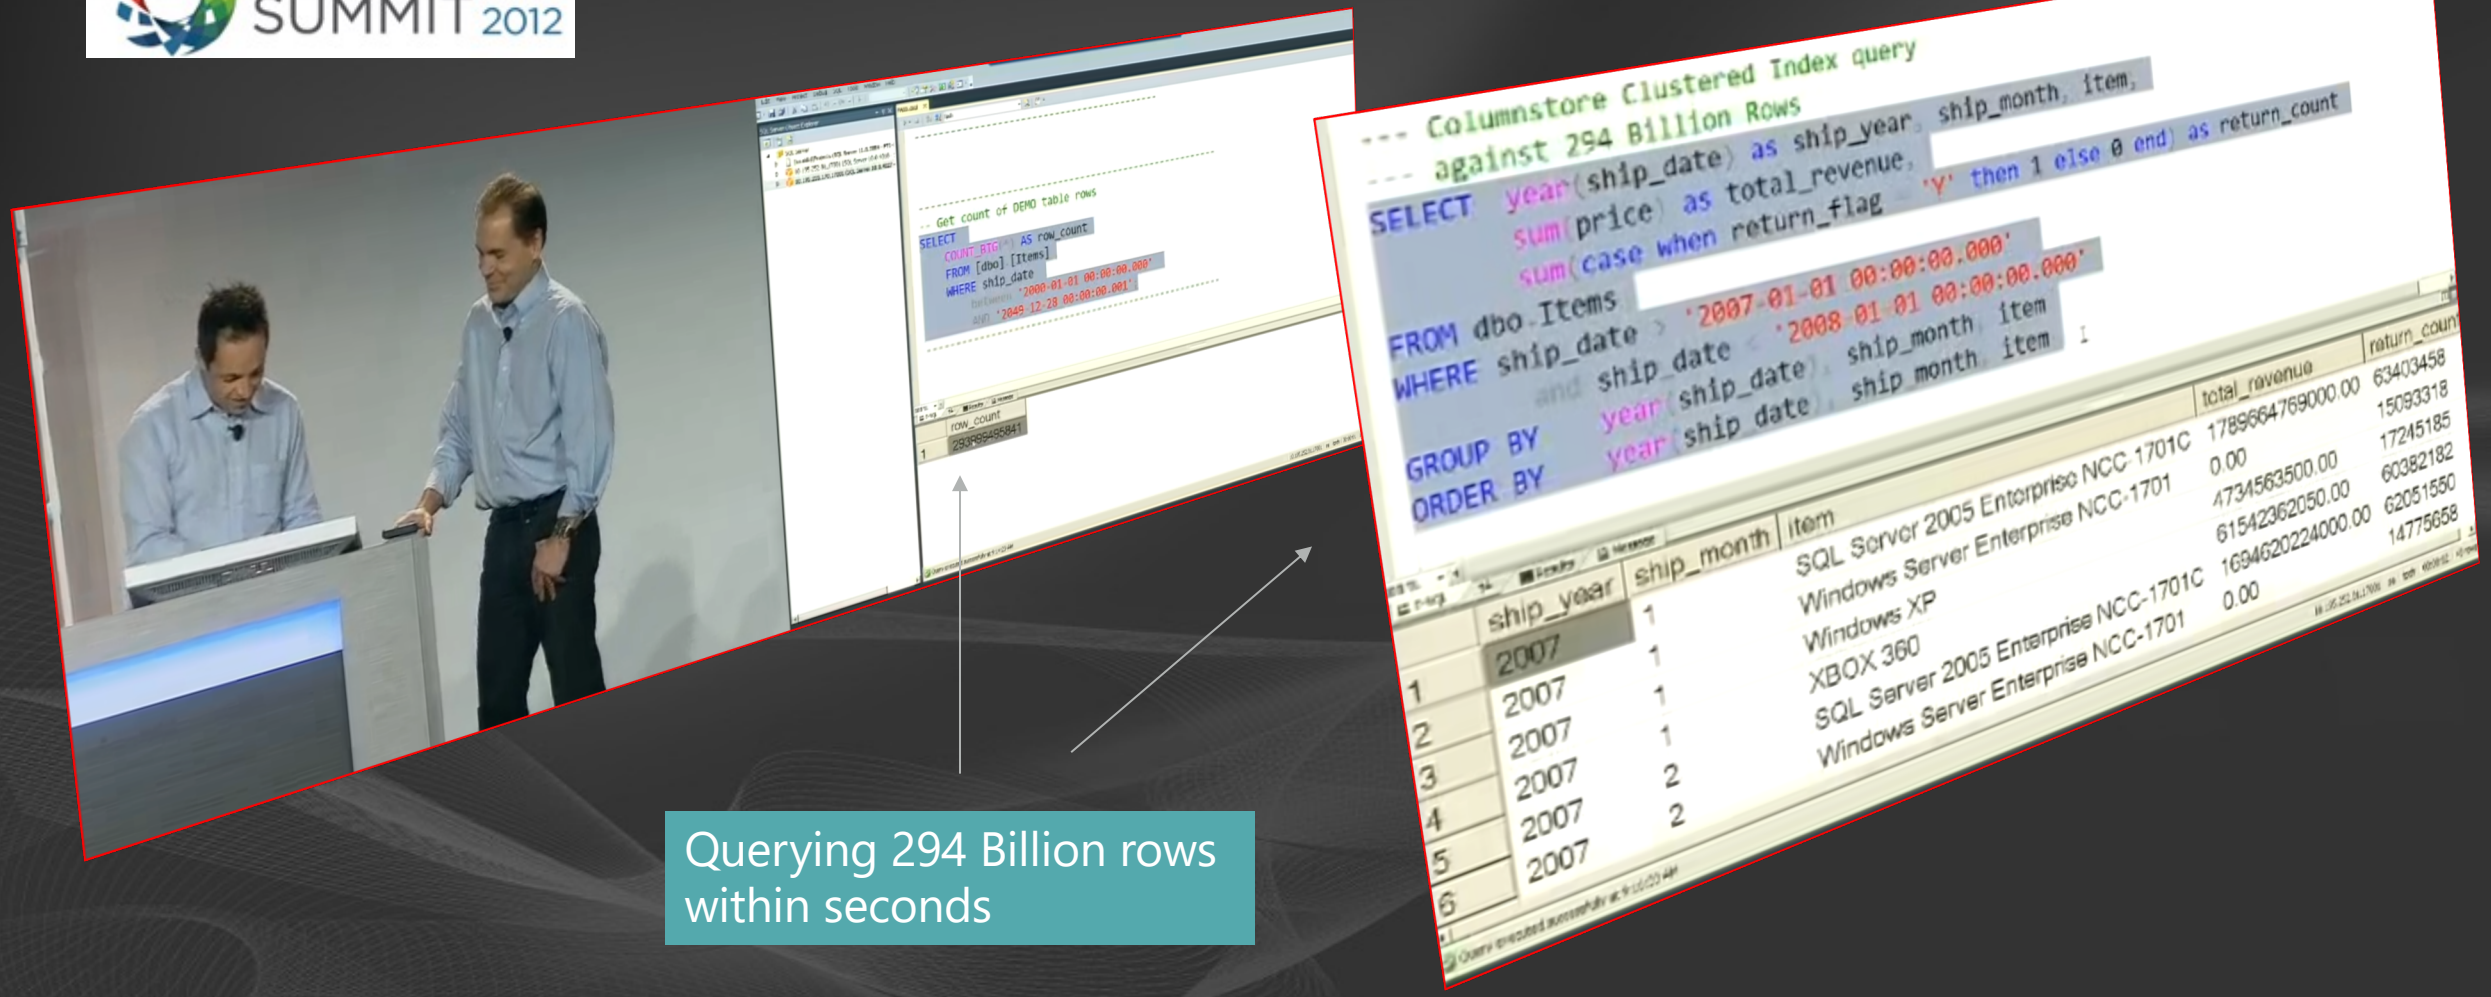

Querying 294 Billion rows within seconds

From Onprem into the Cloud!

FULLDATA
THE DATA PLATFORM SPECIALISTS

Azure Virtual Machines

- 8/16/32/64/96 Core Dxds_v4
- 128 Core M128

VM Size	Family	vCPUs	RAM (GiB)	Data disks	Max IOPS	Temp storage (GiB)	Premium disk	Cost/month
D64as_v4	General purpose	64	256	32	80000	512	Supported	$2,242.56
D64ds_v4	General purpose	64	256	32	80000	2400	Supported	$2,639.68
D64s_v4	General purpose	64	256	32	80000	0	Supported	$2,242.56
D96as_v4	General purpose	96	384	32	80000	768	Supported	$3,363.84

VM Size	Family	vCPUs	RAM (GiB)	Data disks	Max IOPS	Temp storage (GiB)	Premium disk	Cost/month
M64m	Memory optimized	64	1750	64	40000	8000	Supported	$7,546.01
M64ms	Memory optimized	64	1750	64	40000	2000	Supported	$7,546.01
M64s	Memory optimized	64	1000	64	40000	2000	Supported	$4,868.37
M128-64ms	Memory optimized	64	3800	64	80000	4000	Supported	$19,482.24
M128	Memory optimized	128	2000	64	80000	16000	Supported	$9,736.74
M128m	Memory optimized	128	3800	64	80000	16000	Supported	$19,482.24
M128s	Memory optimized	128	2000	64	80000	4000	Supported	$9,736.74

Azure Storage

- 20 x 1TB Premium SSD LUNs
- 4 x Ultra disks

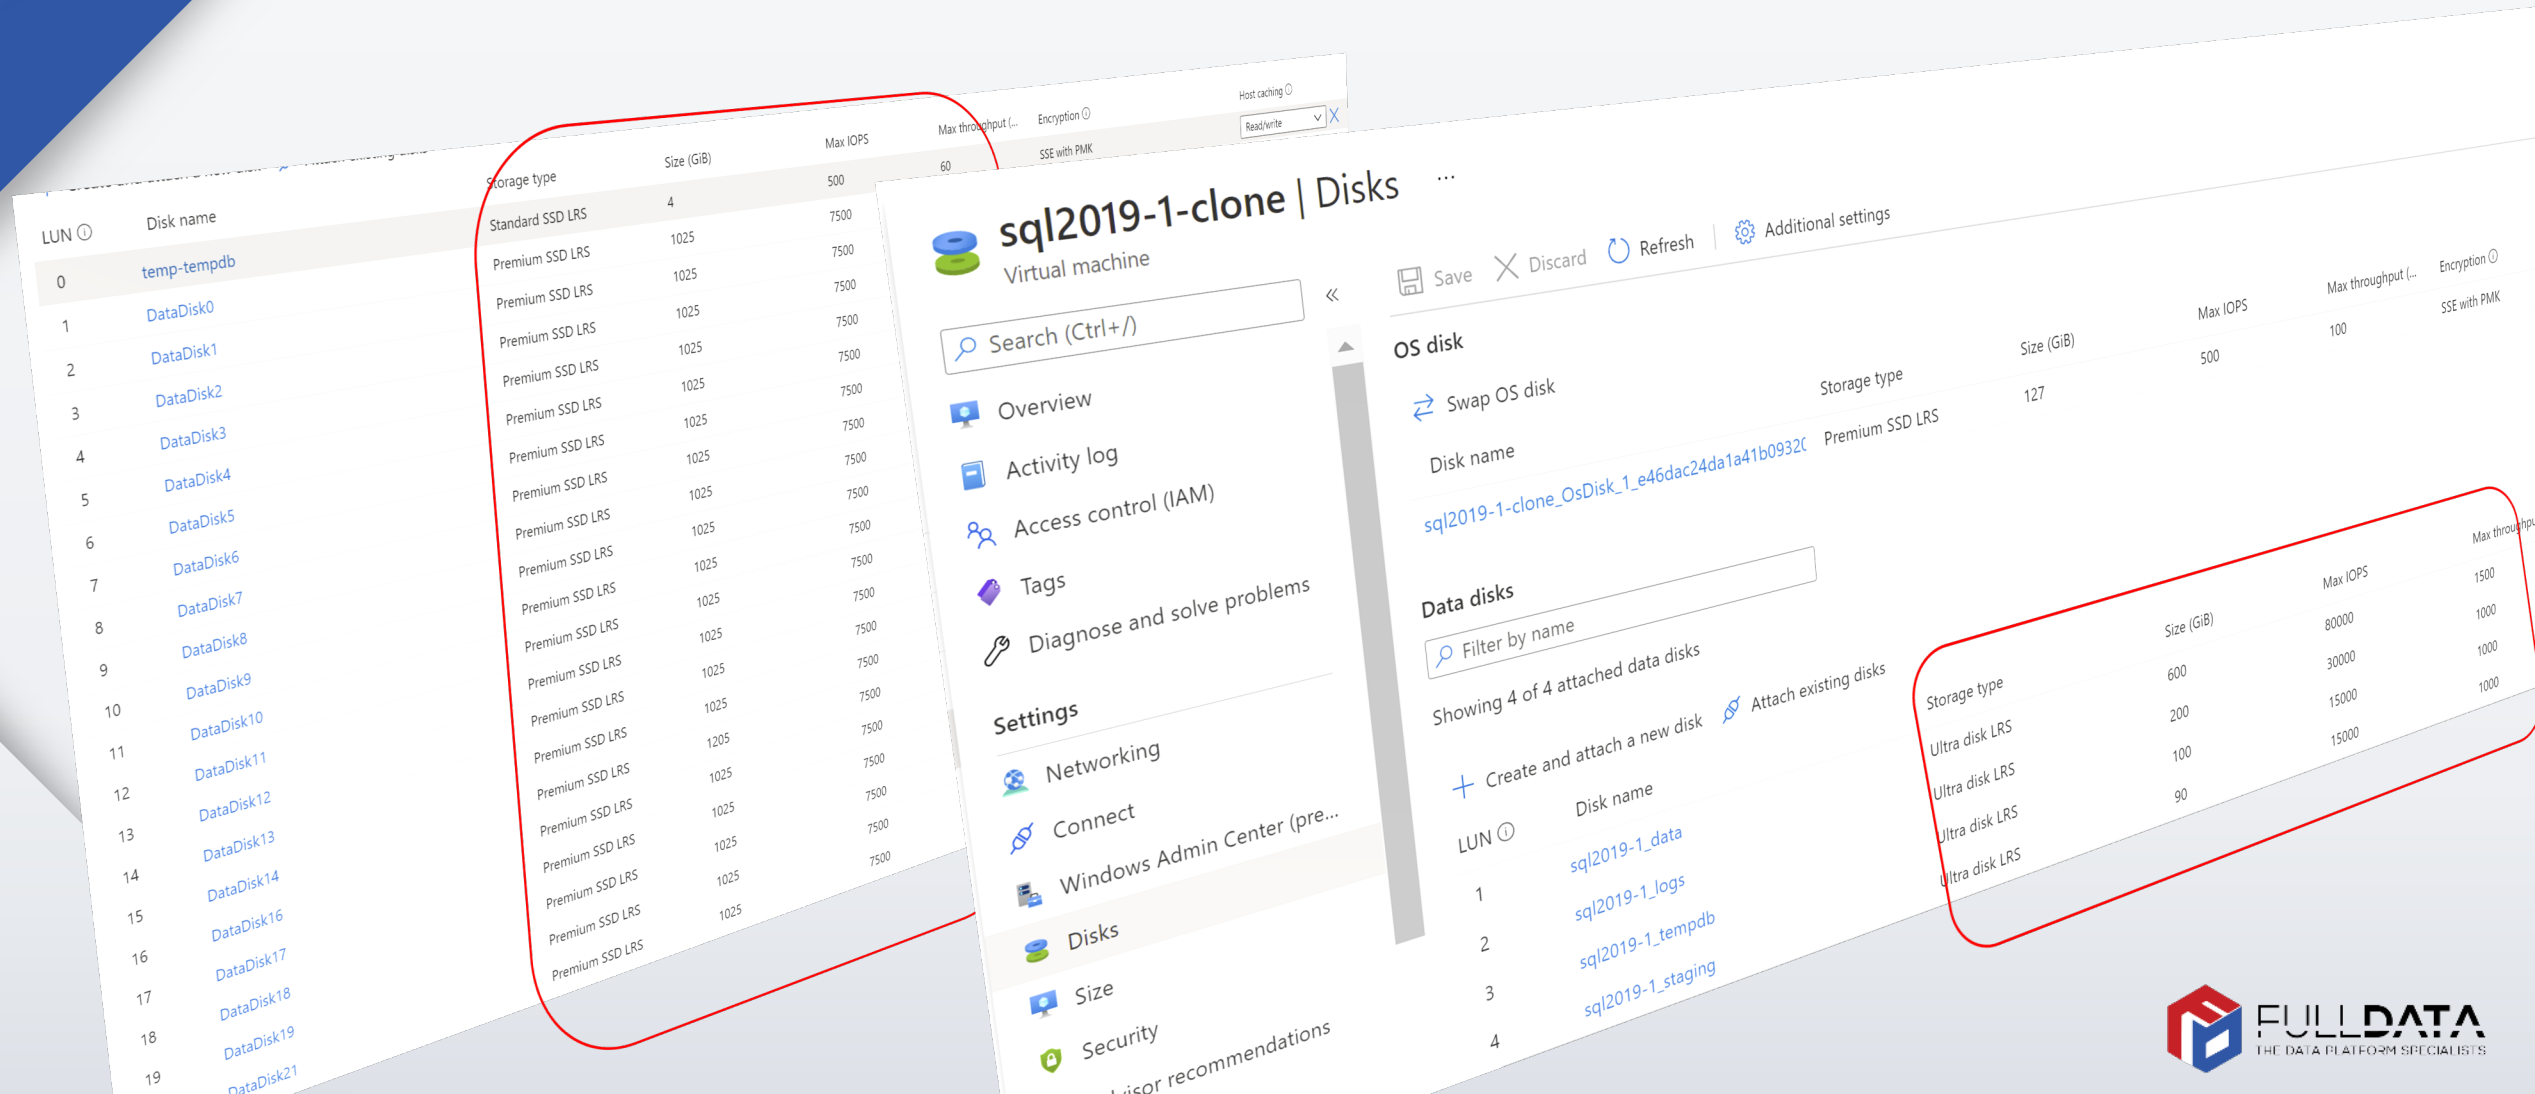

FULLDATA
THE DATA PLATFORM SPECIALISTS

Synthetic test – Premium SSDs

Disk Management

File Action View Help

Volume	Layout	Type	File System	Status	Capacity
Windows (C:)	Simple	Basic	NTFS	Healthy (Boot, Crash Dump...	126.51 GB
Temporary Storage (D:)	Simple	Basic	NTFS	Healthy (Page File, Primary ...	2399.87 GB
tempdb (T:)	Simple	Basic	NTFS	Healthy (Primary Partition)	100.00 GB
temp tempdb	Simple	Basic	NTFS	Healthy (Primary Partition)	4.00 GB
System Reserved	Simple	Basic	NTFS	Healthy (System, Active, Pri...	500 MB
Staging_Stripe4_PDisks (R:)	Striped	Dynamic	NTFS	Healthy	4099.99 GB
staging (S:)	Simple	Basic	NTFS	Healthy (Primary Partition)	80.00 GB
PSSD_Stripe12 (...	Striped	Dynamic	NTFS	Healthy	12299.96 GB
PSDD_Stripe4 (J:)	Striped	Dynamic	NTFS	Healthy	4099.99 GB
Logs (L:)	Simple	Basic	NTFS	Healthy (Primary Partition)	200.00 GB
DATA (E:)	Simple	Basic	NTFS	Healthy (Primary Partition)	800.00 GB

```
throughput metrics:
IOs/sec: 81605.43
MBs/sec:   637.54
latency metrics:
Min_Latency(ms): 1
Avg_Latency(ms): 7
Max_Latency(ms): 77
histogram:
ms: 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24+
%:  0 31 18  5  4  3  3  2  2  2  2  2  2  3  3  3  2  2  2  1  1  1  1  1  6
```

```
c:\Toolbox-SQL\Utils>sqlio -kR -s10 -fsequential  -o64 -b256 -LS -Fparam-P.txt
sqlio v1.5.SG
using system counter for latency timings, 10000000 counts per second
parameter file used: param-P.txt
        file p:\test\testfile.dat with 2 threads (0-1) using mask 0x0 (0)
        file p:\test\testfile1.dat with 2 threads (2-3) using mask 0x0 (0)
        file p:\test\testfile2.dat with 2 threads (4-5) using mask 0x0 (0)
        file p:\test\testfile3.dat with 2 threads (6-7) using mask 0x0 (0)
        file p:\test\testfile8.dat with 2 threads (8-9) using mask 0x0 (0)
10 threads reading for 10 secs from files p:\test\testfile.dat, p:\test\testfile1.dat, p:\test\test
t
        using 256KB sequential IOs
        enabling multiple I/Os per thread with 64 outstanding
using specified size: 2000 MB for file: p:\test\testfile.dat
size of file p:\test\testfile1.dat needs to be: 2097152000 bytes
current file size:      1048576000 bytes
need to expand by:      1048576000 bytes
expanding p:\test\testfile1.dat ... done.
using specified size: 2000 MB for file: p:\test\testfile1.dat
using specified size: 2000 MB for file: p:\test\testfile2.dat
using specified size: 2000 MB for file: p:\test\testfile3.dat
size of file p:\test\testfile8.dat needs to be: 2097152000 bytes
current file size:      1048576000 bytes
need to expand by:      1048576000 bytes
expanding p:\test\testfile8.dat ... done.
using specified size: 2000 MB for file: p:\test\testfile8.dat
initialization done
CUMULATIVE DATA:
throughput metrics:
IOs/sec:  4695.42
MBs/sec:  1173.85
latency metrics:
Min_Latency(ms): 14
Avg_Latency(ms): 135
Max_Latency(ms): 1647
histogram:
ms: 3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24+
%:  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0  0 100
```

sqlio -kR -s30 -fsequential -o64 –b8 -LS -Fparam-E.txt

FULLDATA
THE DATA PLATFORM SPECIALISTS

Synthetic test - Ultra Disk

- ~81K IOPS , ~2 GByte/sec <u>in Total, VM bound</u>

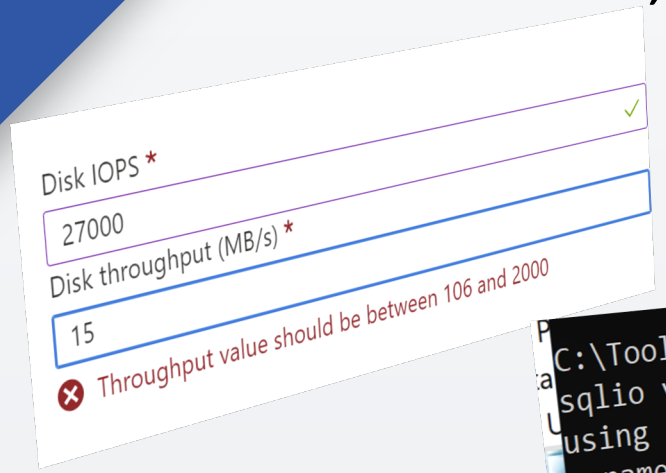

sqlio -kR -s30 -fsequential -o64 –b8 -LS -Fparam-E.txt

```
C:\Toolbox-SQL\Utils>sqlio -kR -s30 -fsequential  -o64 -b8 -
sqlio v1.5.SG
using system counter for latency timings, 10000000 counts pe
parameter file used: param2.txt
        file e:\test\testfile.dat with 2 threads (0-1) using
        file e:\test\testfile1.dat with 2 threads (2-3) using
        file e:\test\testfile2.dat with 2 threads (4-5) using
6 threads reading for 30 secs from files e:\test\testfile.dat
        using 8KB sequential IOs
        enabling multiple I/Os per thread with 64 outstanding
using specified size: 2000 MB for file: e:\test\testfile.dat
using specified size: 2000 MB for file: e:\test\testfile1.dat
using specified size: 2000 MB for file: e:\test\testfile2.dat
initialization done
CUMULATIVE DATA:
throughput metrics:
IOs/sec:  80401.05
MBs/sec:    628.13
latency metrics:
```

```
2.dat and s:\test\testfile3.dat
        using 256KB sequential IOs
        enabling multiple I/Os per thread with 8 outstanding
using specified size: 2000 MB for file: e:\data\testfile.dat
using specified size: 2000 MB for file: e:\data\testfile1.dat
using specified size: 2000 MB for file: e:\data\testfile2.dat
using specified size: 2000 MB for file: l:\logs\testfile5.dat
using specified size: 2000 MB for file: l:\logs\testfile6.dat
using specified size: 2000 MB for file: l:\logs\testfile7.dat
using specified size: 2000 MB for file: t:\test\testfile.dat
using specified size: 2000 MB for file: t:\test\testfile1.dat
using specified size: 2000 MB for file: t:\test\testfile2.dat
using specified size: 2000 MB for file: s:\test\testfile3.dat
initialization done
CUMULATIVE DATA:
throughput metrics:
IOs/sec:  5211.85
MBs/sec:  1302.96
latency metrics:
Min_Latency(ms): 1
Avg_Latency(ms): 30
Max_Latency(ms): 629
```

SQL Test - Reading data from disk as fast as we can

SELECT COUNT (*) from 600 Million Row table

Premium SSD – the Read cache effect

Storage Configuration	Duration (Seconds)	MByte/sec	Remarks
4x1025GB PremiumSSD	67	993	ReadAhead Cache Enabled
8x1025GB PremiumSSD	34	1956	ReadAhead Cache Enabled
12x1025GB PremiumSSD	23	2892	ReadAhead Cache Enabled
12x1025GB PremiumSSD	59	1127	NoCache
Ultra disk	57	1167	

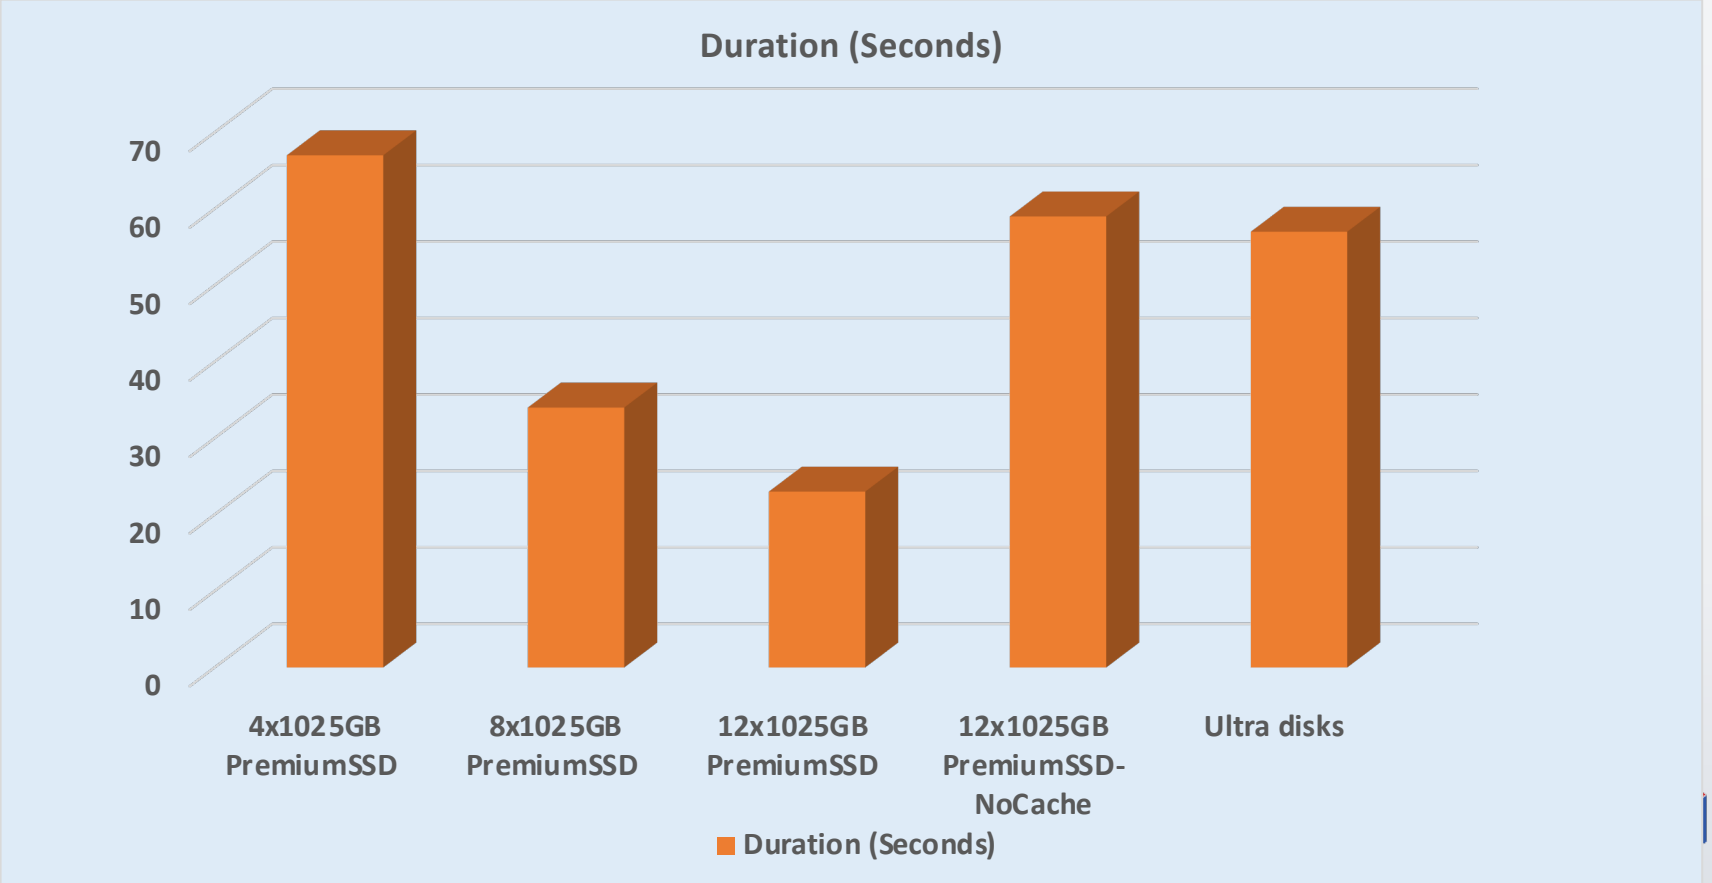

How to beat a 10 year old data loading record

FULLDATA
THE DATA PLATFORM SPECIALISTS

The Database Superchargers!

Build your SAN in the Cloud

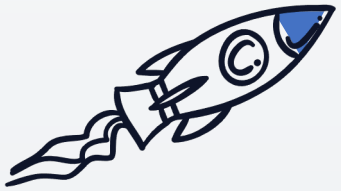

Ultra-high performance

Run 10x faster and shorten run-time

Performance Consistency

Avoid inconsistency and performance degradations

Dynamic Performance

Support constantly changing workloads

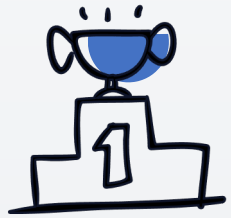

Shared performance

Avoid waste of IOPS provisioning

Silk Architecture (Azure)

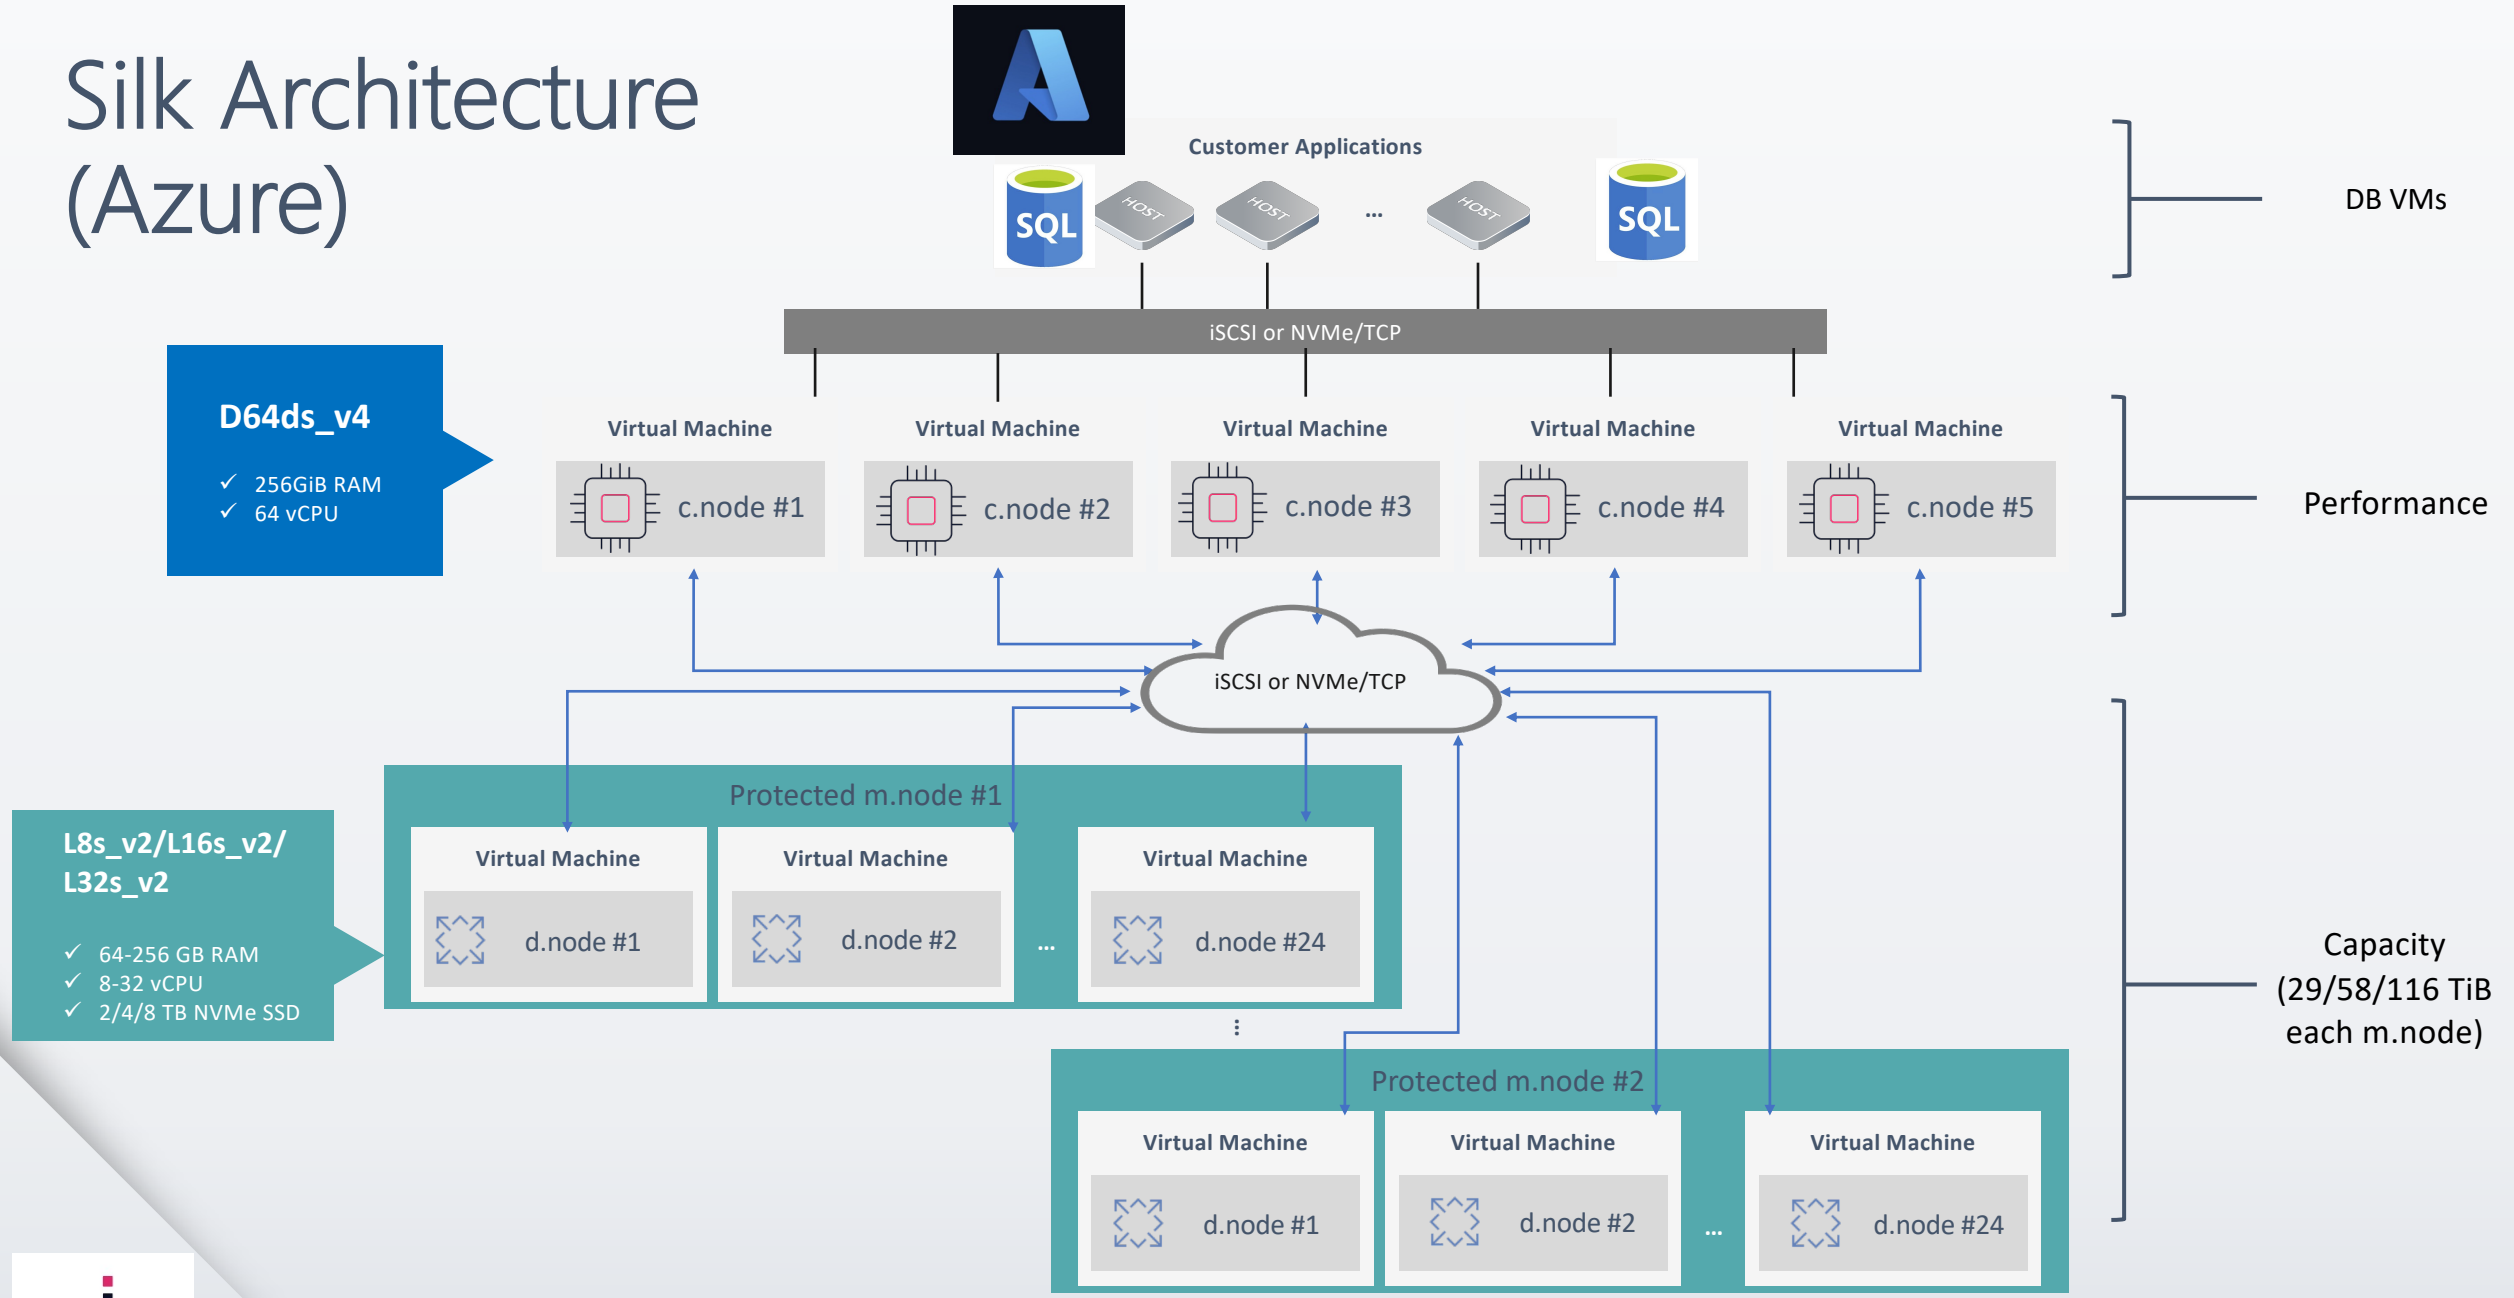

Synthetic test - Silk

iSCSI! - throughput capped by 50 Gbit network bandwidth

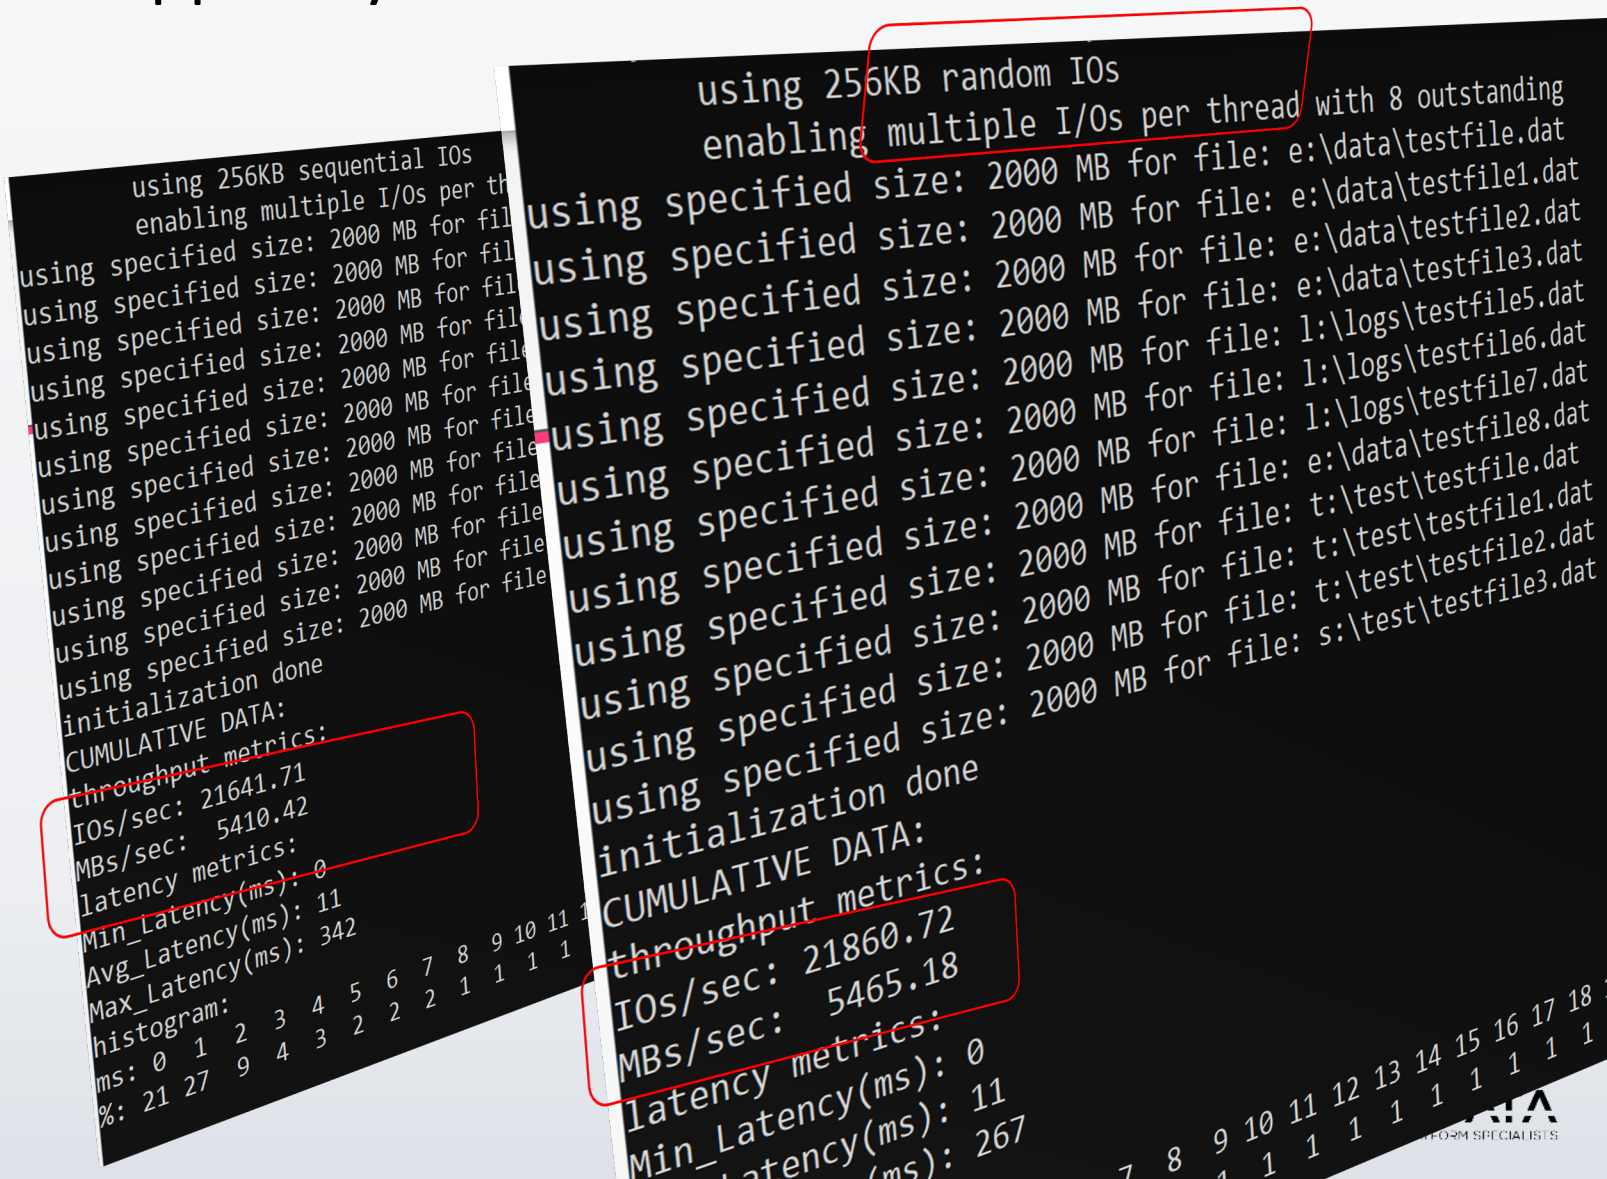

Synthetic test Premium Disk vs Silk

```
sqlio -kR -s30 -fsequential  -o64 -b256 -LS -Fparam-E.txt
```

Disk Management

File Action View Help

Volume	Layout	Type	File System	Status	Capacity
Windows (C:)	Simple	Basic	NTFS	Healthy (Boot, Crash Dump...	126.51 GB
Temporary Storage (D:)	Simple	Basic	NTFS	Healthy (Page File, Primary ...	2399.87 GB
tempdb (T:)	Simple	Basic	NTFS	Healthy (Primary Partition)	100.00 GB
temp tempdb	Simple	Basic	NTFS	Healthy (Primary Partition)	4.00 GB
System Reserved	Simple	Basic	NTFS	Healthy (System, Active, Pri...	500 MB
Staging_Stripe4_PDisks (R:)	Striped	Dynamic	NTFS	Healthy	4099.99 GB
staging (S:)	Simple	Basic	NTFS	Healthy (Primary Partition)	80.00 GB
PSSD_Stripe12 (...	Striped	Dynamic	NTFS	Healthy	12299.96 GB
PSDD_Stripe4 (J:)	Striped	Dynamic	NTFS	Healthy	4099.99 GB
Logs (L:)	Simple	Basic	NTFS	Healthy (Primary Partition)	200.00 GB
DATA (E:)	Simple	Basic	NTFS	Healthy (Primary Partition)	800.00 GB

Silk

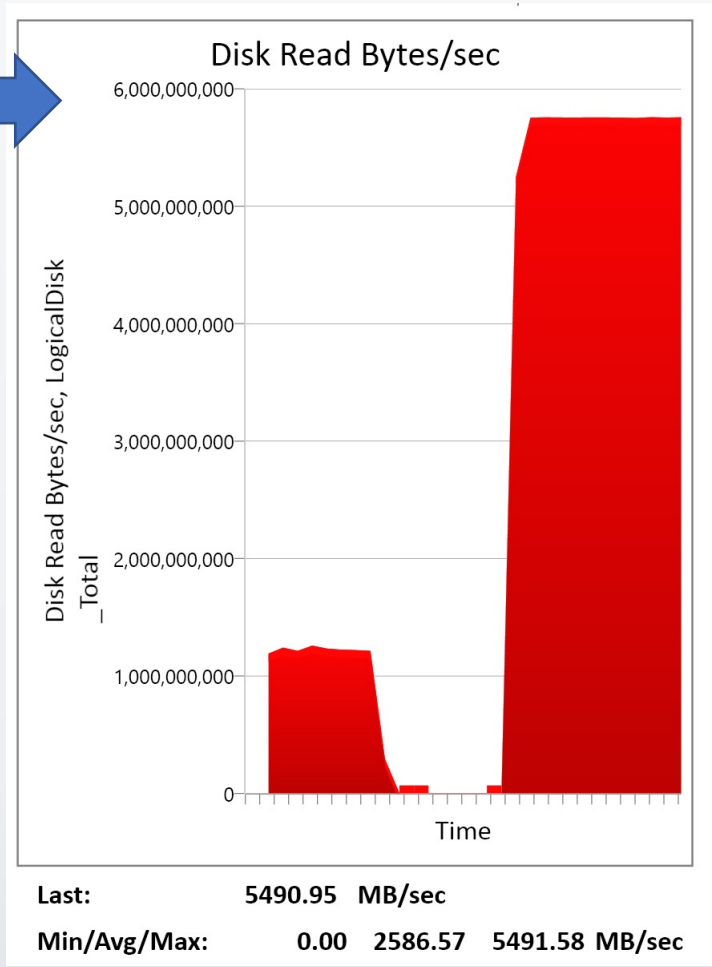

12x LUN

Premium SSD

Windows Stripe Volume

FULLDATA
THE DATA PLATFORM SPECIALISTS

Table Scanning within SQL Server

- Select count (*) from 73GByte Line Item table

- Silk
- Premium SSD

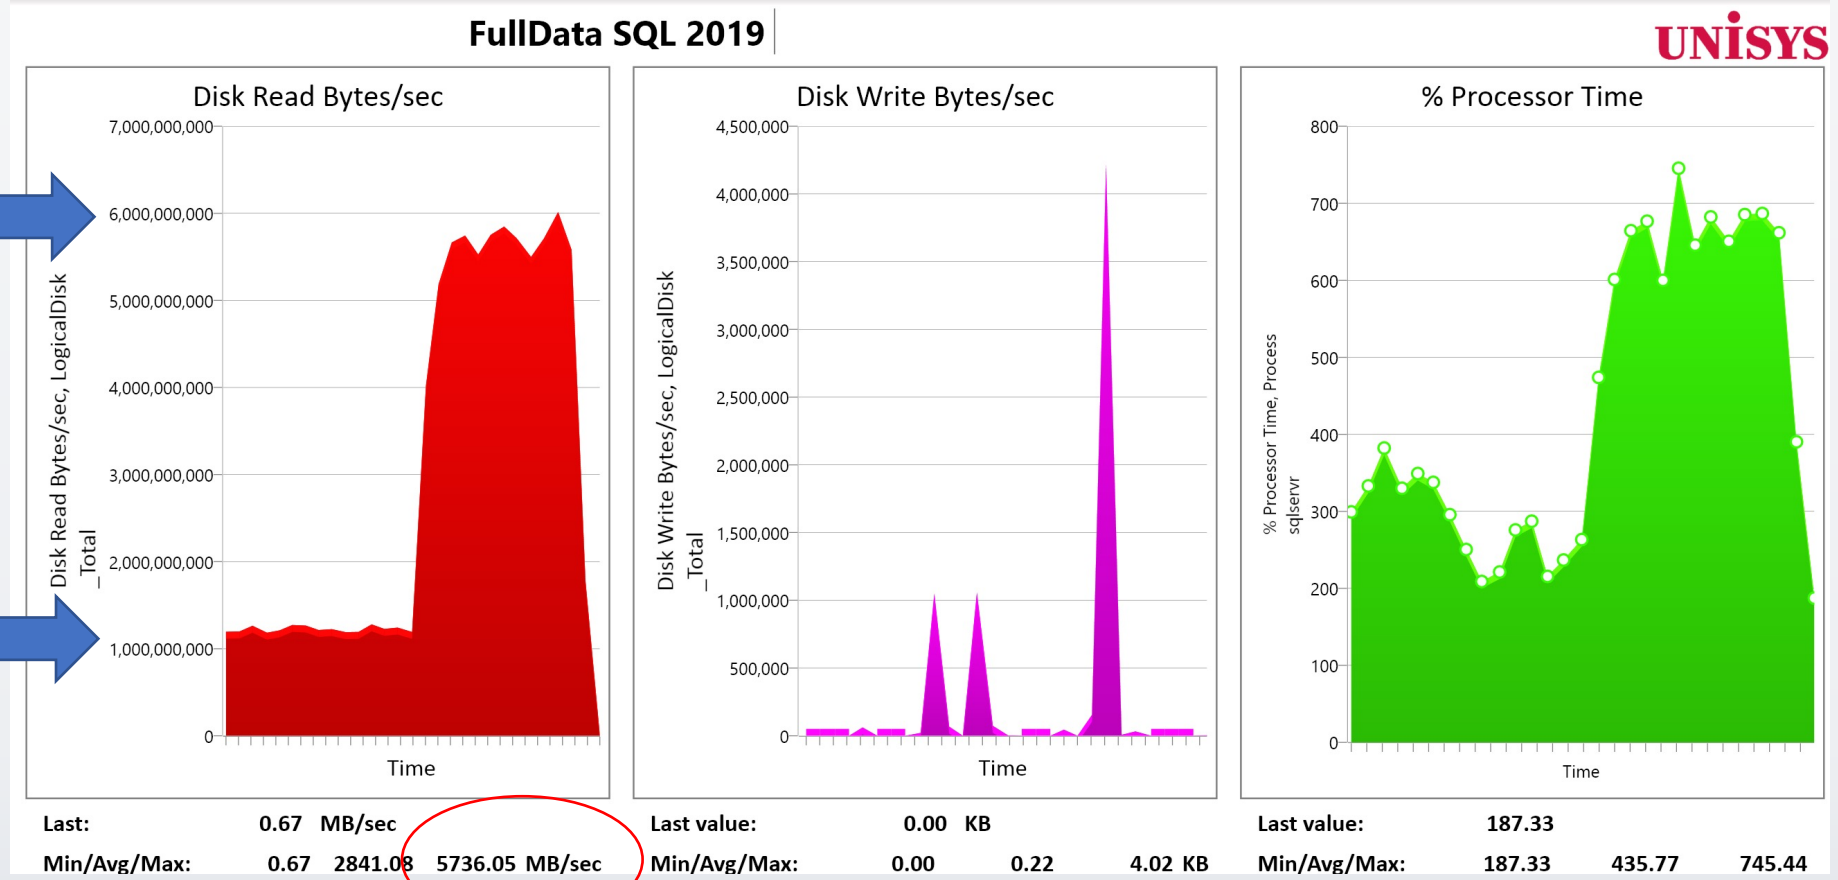

Bulk Inserting

Number to Beat:
15 Million rows/Second

Setting the scene

- 128 separate files,
- 730 Mbyte each (CDROM+ size) source files
- TPC-H Line item data
- 768 Millions rows

- Standard M128 (128 vcpus, 2048 GiB memory)

Impact of sharpening datatypes

```
11  CREATE TABLE [dbo].[LINEITEM_NVarchar](
12      [L_SHIPDATE] [nvarchar](max),
13      [L_ORDERKEY] [nvarchar](max),
14      [L_DISCOUNT] [nvarchar](max),
15      [L_EXTENDEDPRICE] [nvarchar](max),
16      [L_SUPPKEY] [nvarchar](max),
17      [L_QUANTITY] [nvarchar](max),
18      [L_RETURNFLAG] [nvarchar](max),
19      [L_PARTKEY] [nvarchar](max),
20      [L_LINESTATUS] [nvarchar](max),
21      [L_TAX] [nvarchar](max),
22      [L_COMMITDATE] [nvarchar](max),
23      [L_RECEIPTDATE] [nvarchar](max),
24      [L_SHIPMODE] [nvarchar](max),
25      [L_LINENUMBER] [nvarchar](max),
26      [L_SHIPINSTRUCT] [nvarchar](max),
27      [L_COMMENT] [nvarchar](max)
28  )
29  GO
```

```
2   CREATE TABLE [dbo].[LINEITEM_0](
3       [L_SHIPDATE] [date] NOT NULL,
4       [L_ORDERKEY] [bigint] NOT NULL,
5       [L_DISCOUNT] [smallmoney] NOT NULL,
6       [L_EXTENDEDPRICE] [money] NOT NULL,
7       [L_SUPPKEY] [int] NOT NULL,
8       [L_QUANTITY] [smallint] NOT NULL,
9       [L_RETURNFLAG] [char](1) NOT NULL,
10      [L_PARTKEY] [int] NOT NULL,
11      [L_LINESTATUS] [char](1) NOT NULL,
12      [L_TAX] [smallmoney] NOT NULL,
13      [L_COMMITDATE] [date] NOT NULL,
14      [L_RECEIPTDATE] [date] NOT NULL,
15      [L_SHIPMODE] [varchar](10) NOT NULL,
16      [L_LINENUMBER] [int] NOT NULL,
17      [L_SHIPINSTRUCT] [varchar](25) NOT NULL,
18      [L_COMMENT] [varchar](44) NOT NULL
19  )
20  GO
```

Datatype impact when Bulk inserting

- 730 MByte flat file

Scenario	Duration (Sec)	Throughput (MByte/Sec)	MBytes written to DB File	Mbytes written to Log file	Total MB
Bulkinsert into NVARCHAR(MAX)	67	11	1390	17.5	1408
Bulkinsert into VARCHAR(MAX)	59	12	904	10.2	914
Datatypes sharpened	18	41	646	7.2	653

3.5 x faster!

Copying tables @ Top speed

- Smartbulkcopy utility

```
4  select %%PhysLoc%%
5  from LINEITEM
6
```

121 %

Results | Messages

	(No column name)
1	0x98C90D0001000000
2	0x98C90D0001000100
3	0x98C90D0001000200
4	0x98C90D0001000300
5	0x98C90D0001000400
6	0x98C90D0001000500
7	0x98C90D0001000600
8	0x98C90D0001000700
9	0x98C90D0001000800
10	0x98C90D0001000900
11	0x98C90D0001000A00
12	0x98C90D0001000B00
13	0x98C90D0001000C00
14	0x98C90D0001000D00
15	0x98C90D0001000E00
16	0x98C90D0001000F00
17	0x98C90D0001001000
18	0x98C90D0001001100
19	0x98C90D0001001200

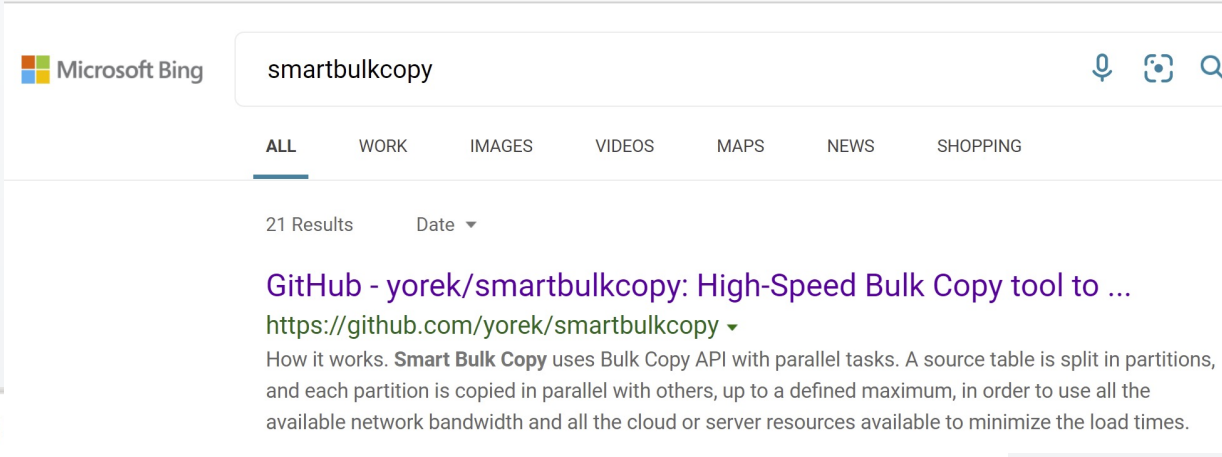

smartbulkcopy.config -

File Edit Format View Help

```
{
"source": {
"connection-string": "Server=.;Initial Catalog=Knuts;Integrated security=SSPI; Packet Size=16384;
Connection Timeout=300;",
"csv-path": "C:\\FullData\\DBGen\\Input"
},
"destination":{
"connection-string": "Server=.;Initial Catalog=Knuts_copy;Integrated security=SSPI; Packet
Size=16384;Connection Timeout=300;"
},
"tables": ["dbo.lineitem"],
"options": {
"tasks": 6,
"logical-partitions": 8,
"inputcsv": false,
"fieldseparator" : "|",
"file-extension" : "tbl",
"readOnly": true,
"batch-size": 100000,
"truncate-tables": true,
"safe-check": "none",
"retry-connection": {
"delay-increment": 1,
"max-attempt": 5
}
}
```

How to lose your data within 1 second ☺

Workqueue for bulk insert tasks / 2 check files locally + ‘hara-kiri support’

Results | Messages

	work_id	group	priority	assignment	partition	target	is_active	is_done	start_date	end_date	time_sec	logical_reads	reads	writes	cpu_time
1	1	BULK	0	C:\Toolbox-SQL\Utils\lineitem.tbl	NULL	LINEITEM_01	0	1	2021-03-10 08:58:41.143	2021-03-10 08:59:09.297	28	123276	0	81993	27000
2	2	BULK	0	D:\lineitem.tbl	NULL	LINEITEM_02	0	1	2021-03-10 08:58:41.173	2021-03-10 08:59:06.660	25	123312	0	81971	25432
3	3	BULK	0	T:\FullData\lineitem.tbl.11	NULL	LINEITEM_03	0	1	2021-03-10 08:58:41.190	2021-03-10 08:59:15.667	34	123329	0	81981	29772
4	4	BULK	0	T:\FullData\lineitem.tbl.12	NULL	LINEITEM_04	0	1	2021-03-10 08:58:41.203	2021-03-10 08:59:14.480	33	123307	0	81966	28157
5	5	BULK	0	T:\FullData\lineitem.tbl.13	NULL	LINEITEM_05	0	1	2021-03-10 08:58:41.220	2021-03-10 08:59:16.280	35	123329	0	81979	29945
6	6	BULK	0	T:\FullData\lineitem.tbl.14	NULL	LINEITEM_06	0	1	2021-03-10 08:58:41.237	2021-03-10 08:59:14.833	33	123288	0	81965	28235
7	7	BULK	0	T:\FullData\lineitem.tbl.15	NULL	LINEITEM_07	0	1	2021-03-10 08:58:41.250	2021-03-10 08:59:15.157	33	123318	0	81973	28661
8	8	BULK	0	T:\FullData\lineitem.tbl.16	NULL	LINEITEM_08	0	1	2021-03-10 08:58:41.253	2021-03-10 08:59:14.140	32	123307	0	81968	27963
9	9	BULK	0	T:\FullData\lineitem.tbl.17	NULL	LINEITEM_09	0	1	2021-03-10 08:58:41.267	2021-03-10 08:59:13.880	32	123313	0	81978	27554
10	10	BULK	0	T:\FullData\lineitem.tbl.18	NULL	LINEITEM_10	0	1	2021-03-10 08:58:41.270	2021-03-10 08:59:13.527	32	123331	0	81974	27873
11	11	BULK	0	T:\FullData\lineitem.tbl.19	NULL	LINEITEM_11	0	1	2021-03-10 08:58:41.270	2021-03-10 08:59:16.053	34	123293	0	81965	29614
12	12	BULK	0	T:\FullData\lineitem.tbl.2	NULL	LINEITEM_12	0	1	2021-03-10 08:58:41.283	2021-03-10 08:59:14.243	32	123306	0	81976	27786
13	13	BULK	0	T:\FullData\lineitem.tbl.20	NULL	LINEITEM_13	0	1	2021-03-10 08:58:41.287	2021-03-10 08:59:15.613	34	123308	0	81979	29495
14	14	BULK	0	T:\FullData\lineitem.tbl.21	NULL	LINEITEM_14	0	1	2021-03-10 08:58:41.287	2021-03-10 08:59:15.973	34	123319	0	81962	29947
15	15	BULK	0	T:\FullData\lineitem.tbl.22	NULL	LINEITEM_15	0	1	2021-03-10 08:58:41.827	2021-03-10 08:59:15.663	33	123306	0	81978	27749
16	16	BULK	0	T:\FullData\lineitem.tbl.23	NULL	LINEITEM_16	0	1	2021-03-10 08:58:41.830	2021-03-10 08:59:16.317	34	123309	0	81972	29333
17	17	BULK	0	T:\FullData\lineitem.tbl.24	NULL	LINEITEM_17	0	1	2021-03-10 08:58:41.867	2021-03-10 08:59:15.337	33	123310	0	81968	28026
18	18	BULK	0	T:\FullData\lineitem.tbl.25	NULL	LINEITEM_18	0	1	2021-03-10 08:58:41.897	2021-03-10 08:59:15.687	33	123333	0	81972	28348
19	19	BULK	0	T:\FullData\lineitem.tbl.26	NULL	LINEITEM_19	0	1	2021-03-10 08:58:41.907	2021-03-10 08:59:15.283	33	123297	0	81978	28501
20	20	BULK	0	T:\FullData\lineitem.tbl.27	NULL	LINEITEM_20	0	1	2021-03-10 08:58:41.910	2021-03-10 08:59:16.140	34	123345	0	81965	28810
21	21	BULK	0	T:\FullData\lineitem.tbl.28	NULL	LINEITEM_21	0	1	2021-03-10 08:58:41.913	2021-03-10 08:59:16.317	34	123303	0	81978	29135
22	22	BULK	0	T:\FullData\lineitem.tbl.29	NULL	LINEITEM_22	0	1	2021-03-10 08:58:41.913	2021-03-10 08:59:15.970	34	123329	0	81972	29345
23	23	BULK	0	T:\FullData\lineitem.tbl.3	NULL	LINEITEM_23	0	1	2021-03-10 08:58:41.990	2021-03-10 08:59:16.840	34	123316	0	81978	29601
24	24	BULK	0	T:\FullData\lineitem.tbl.30	NULL	LINEITEM_24	0	1	2021-03-10 08:58:42.000	2021-03-10 08:59:15.977	33	123318	0	81975	28521
25	25	BULK	0	T:\FullData\lineitem.tbl.31	NULL	LINEITEM_25	0	1	2021-03-10 08:58:42.007	2021-03-10 08:59:15.970	33	123300	0	81976	29073
26	26	BULK	0	T:\FullData\lineitem.tbl.32	NULL	LINEITEM_26	0	1	2021-03-10 08:58:42.033	2021-03-10 08:59:15.630	33	123335	0	81974	28590
27	27	BULK	0	T:\FullData\lineitem.tbl.0	NULL	LINEITEM_27	0	1	2021-03-10 08:58:42.050	2021-03-10 08:59:16.343	34	123319	0	81968	29694
28	28	BULK	0	T:\FullData\lineitem.tbl.4	NULL	LINEITEM_28	0	1	2021-03-10 08:58:42.050	2021-03-10 08:59:15.633	33	123341	0	81974	28138
29	29	BULK	0	T:\FullData\lineitem.tbl.5	NULL	LINEITEM_29	0	1	2021-03-10 08:58:42.080	2021-03-10 08:59:16.727	34	123289	0	81973	29758
30	30	BULK	0	T:\FullData\lineitem.tbl.6	NULL	LINEITEM_30	0	1	2021-03-10 08:58:42.087	2021-03-10 08:59:15.470	33	123332	0	81976	28194
31	31	BULK	0	T:\FullData\lineitem.tbl.7	NULL	LINEITEM_31	0	1	2021-03-10 08:58:44.843	2021-03-10 08:59:17.510	32	123311	0	81965	28169
32	32	BULK	0	T:\FullData\lineitem.tbl.8	NULL	LINEITEM_32	0	1	2021-03-10 08:59:06.663	2021-03-10 08:59:26.917	20	123301	0	81951	19444

Best result with 4x Ultra disks:

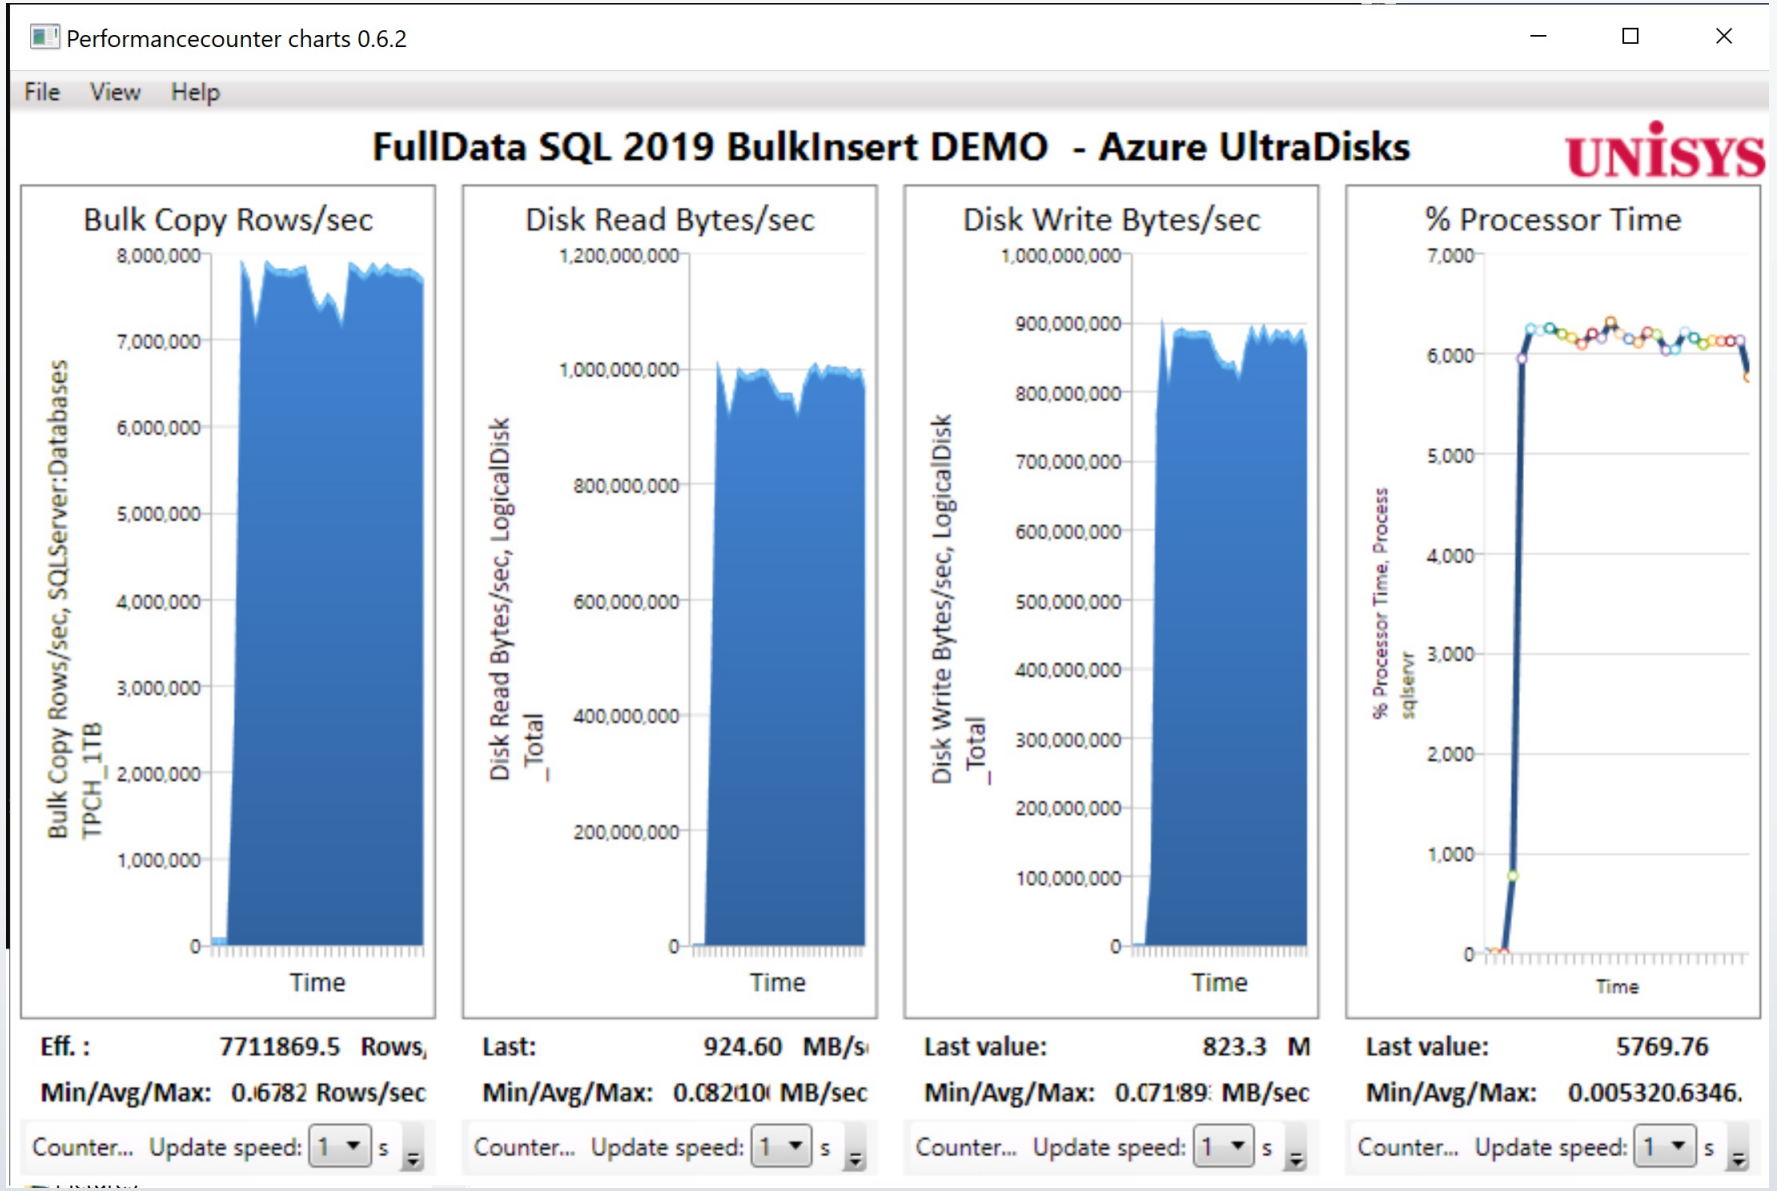

FULLDATA
THE DATA PLATFORM SPECIALISTS

The new record!

20 Million Rows/Second
8+ Terabyte/hour

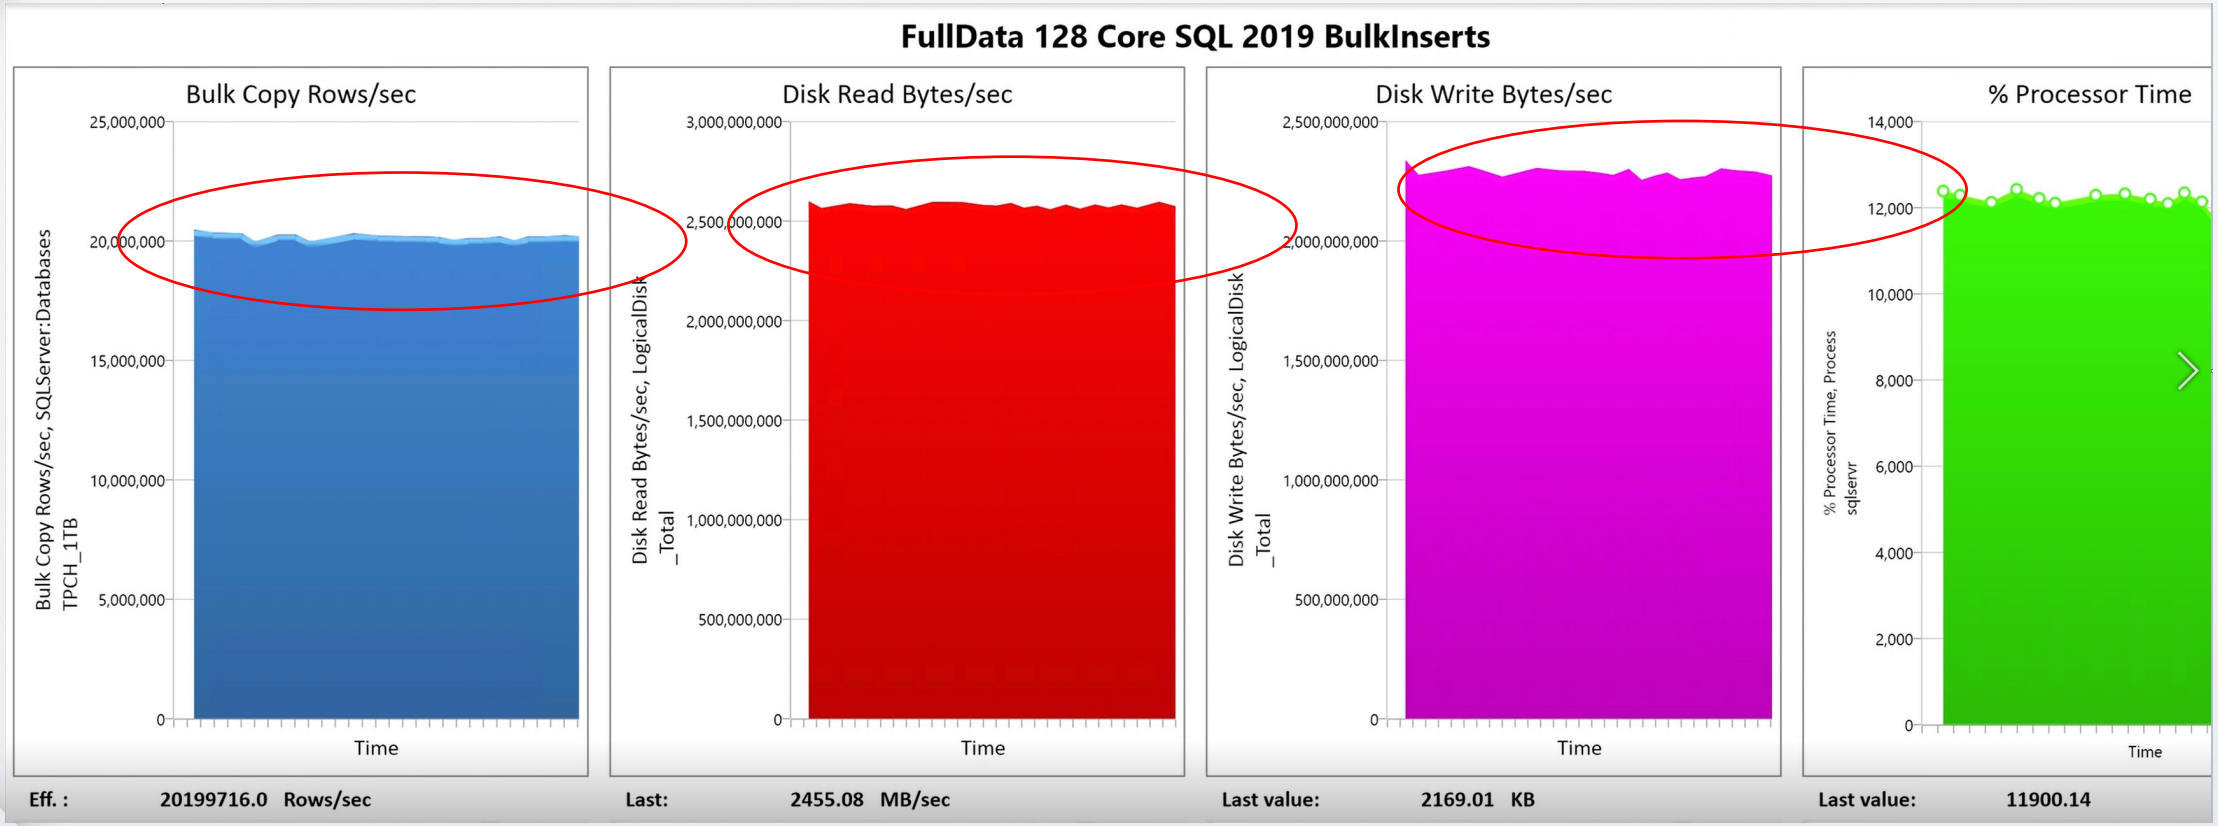

128 Core M-Series Azure VM with 3 cNode Silk Storage

SQL Capacity Management

Analyze SQL workload over time

www.fulldata.nl

FULLDATA
THE DATA PLATFORM SPECIALISTS

Cosmic | Onprem | HV Henk van der Valk

- Dashboard
- General Information
- Backup Information
- File Statistics
- Session Information
- Memory Statistics
- Principal Information
- Log Data

SQL Workload Overview

Disk Activity / Period	Apr vs May 2020	May vs Jun 2020	Jun vs Jul 2020	Jul vs Aug 2020	Aug vs Sep 2020	Sep vs Oct 2020	Oct vs Nov 2020	Nov vs Dec 2020	Dec 2020 vs Jan 2021	Jan vs Feb 2021	Feb vs Mar 2021	Mar vs Apr 2021	Apr vs May 2021
Total User Data Read	Baseline 824.11 TB	↓ -32.6% 555.09 TB	→ 0.7% 558.77 TB	→ 0.3% 560.3 TB	→ 4.2% 583.7 TB	→ -2.4% 569.44 TB	↗ 31.5% 748.91 TB	↘ -21.4% 588.85 TB	↑ 114.7% 1264.1 TB	↓ -57.0% 543.21 TB	↗ 38.4% 751.97 TB	↘ -19.7% 604.02 TB	→ -4.6% 576.09 TB
Total User Data Write	Baseline 227.52 TB	↓ -57.9% 95.75 TB	↑ 174.9% 263.18 TB	↓ -58.4% 109.41 TB	↗ 23.0% 134.59 TB	↗ 19.0% 160.22 TB	↑ 60.2% 256.65 TB	↓ -49.6% 129.44 TB	↘ -16.8% 107.71 TB	↑ 55.1% 167.09 TB	↑ 51.1% 252.52 TB	↓ -38.4% 155.49 TB	↗ 22.8% 190.94 TB

FULLDATA
THE DATA PLATFORM SPECIALISTS

ADF Capacity Management

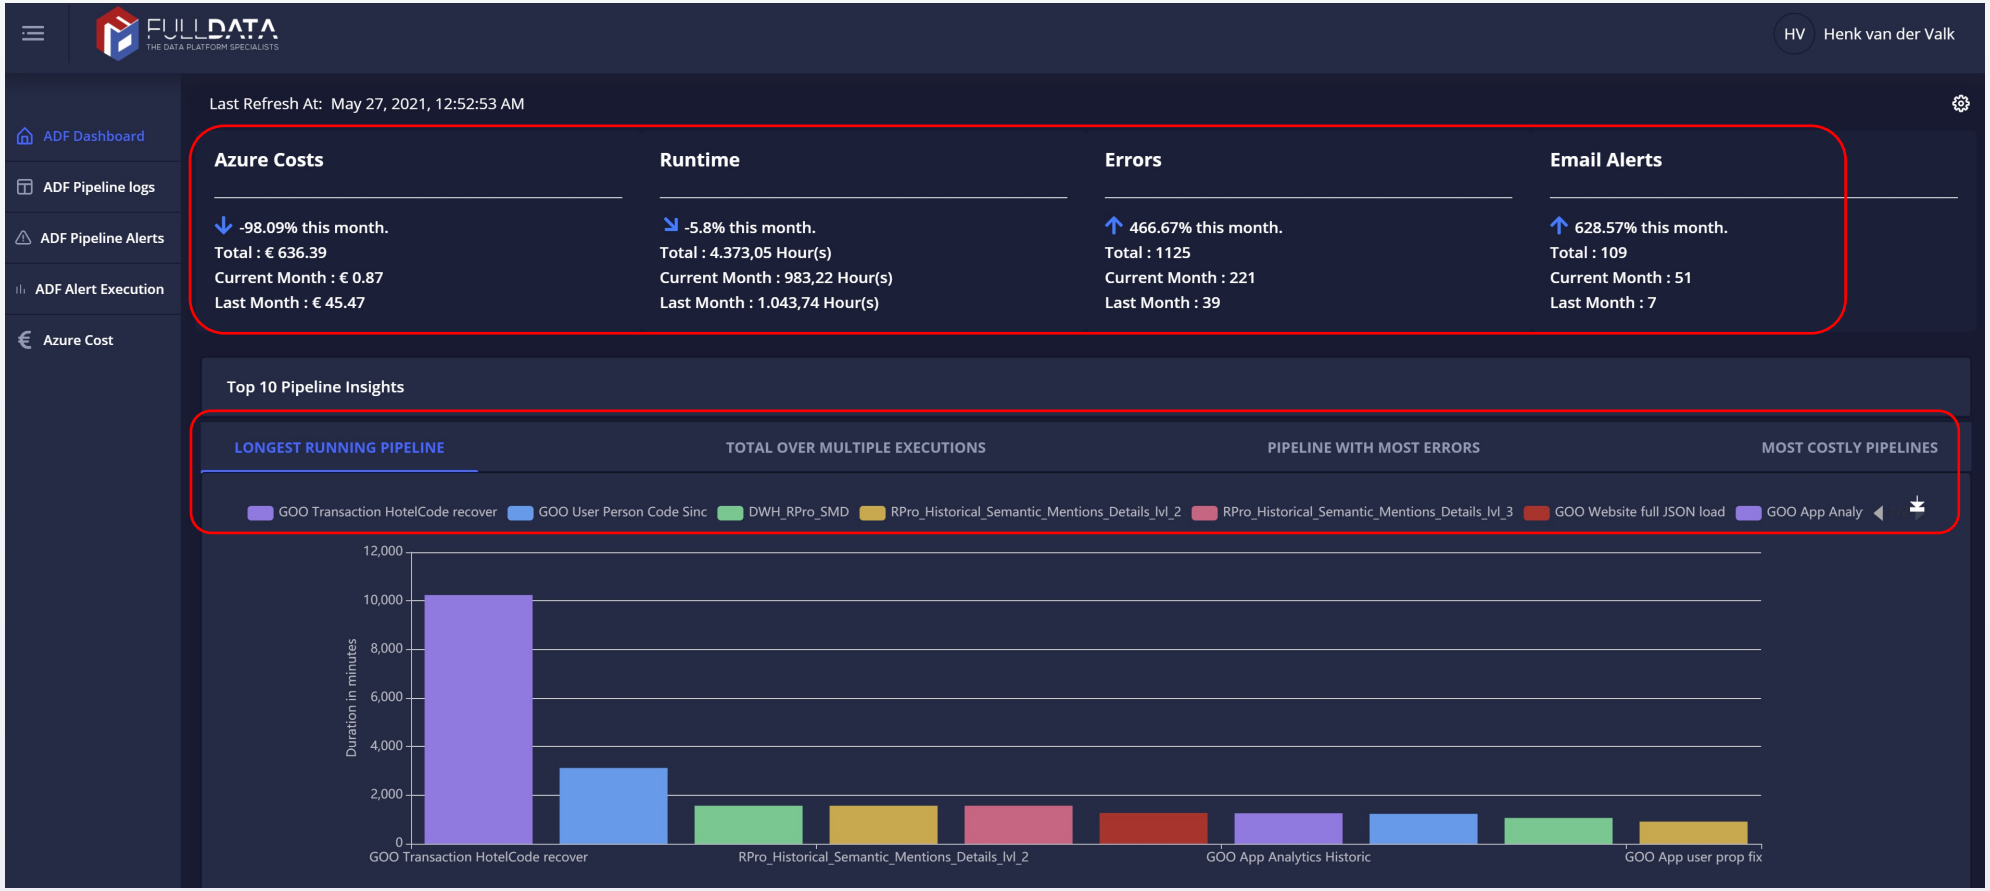

FULLDATA
THE DATA PLATFORM SPECIALISTS

Wrap-up

- Running SQL in Azure @ Speed works great!

Hope you got a bit better understanding of

- How-to size & test
- VM throughput / capacity / in-memory

- Best performance with Silk Storage via iSCSI