



Data Management Institute

BUILDING HIGH PERFORMANCE INFRASTRUCTURE FOR DATABASES

By

Jon Toigo

**Managing Principal, Toigo Partners International
Chairman, Data Management Institute**

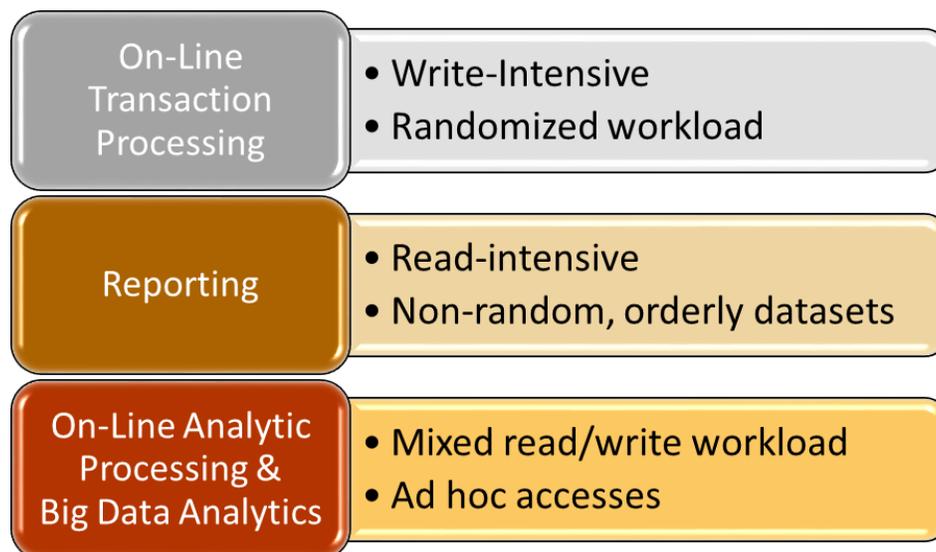
SUMMARY

Databases drive 60% of the applications in most companies. So, the performance of databases is extremely important to an organization's ability to process transactions as well as analyze data. Recently, many technology vendors have submitted all-flash products or other storage wares as "the right infrastructure" for databases and database-dependent applications. However, companies purchasing and deploying these high-performance servers and storage kits are still finding their databases to be underperforming. A slow database is often a reflection of latency brought on by a bottleneck in I/O processing inside the server: an issue that cannot be cost-effectively resolved by throwing more peripheral hardware (flash, disk or interconnects) at the problem. This paper looks at the impact of I/O latency on the performance of databases and their dependent applications and suggests an affordable option for surmounting the problems at a reasonable cost.

INTRODUCTION

Databases (and applications that depend on them) remain the mainstay of business computing. From traditional “systems of record” -- on-line transaction processing (OLTP) and reporting systems, and applications like Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), supply chain management and logistics, etc. -- to the latest “systems of insight” -- including on-line analytical processing (OLAP) and “Big Data” analytics, databases do most of the heavy lifting of IT services and generate much of the revenues for the organization.

Databases differ somewhat in the requirements that they make of underlying infrastructure, of course.

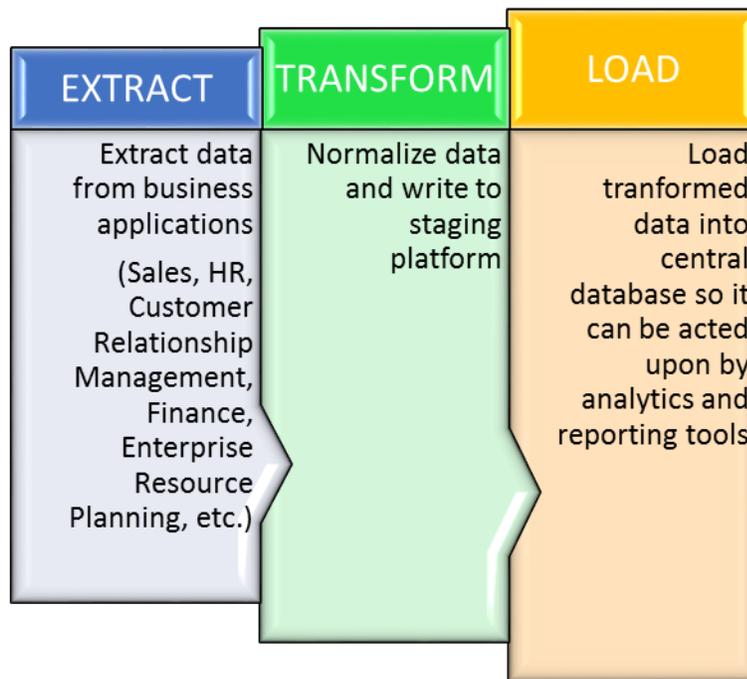


- OLTP system, as some analysts have observed, tend to present a mixture of read and write operations. Plus, they are latency sensitive as a rule, because they tend to be “consumer facing” and concurrent in operation (routinely processing many requests at the same time). These workload characteristics create a set of infrastructure support requirements that must be met to make the system perform responsively. In the case of OLTP databases, I/O activity is mostly randomized and very write intensive, so an infrastructure must be provided that can accommodate these workload characteristics.
- Reporting systems, by contrast, tend to be read intensive, requiring throughput in order to deliver information often within a narrow processing window to impact decisions. This represents a very different set of requirements from a workload and infrastructure standpoint.

- OLAP and analytics systems are thought to be ad hoc in operation, more challenging to characterize in terms of infrastructure and workload requirements. Some queries may require mostly reads to process, but Big Data analysis is often presenting a need for companies to write data, and then read it, in order to derive just-in-time insights about the transaction or the customer or both. Typically, OLAP and analytics platforms are designed to satisfy “boundary conditions” – the demands exerted by the processing of the largest or most complex queries because “typical” schemas are hard to nail down.

Despite these differences, many vendors seek to advance a “one size fits most” approach to database acceleration, usually predicated on substituting faster storage I/O devices for slower ones in the hopes of achieving an increase in overall speed. “In memory databases” – placing entire databases into system memory during operation – are the “shiny new things” in discussions of business analytics. They also tend to carry the highest price tags.

Other vendors focus on improving the processes for extracting, transforming and loading data (ETL) from OLTP systems and databases into data warehouses for real-time data validation or analysis using reporting or analytics platforms like Hadoop or SPARK. These efforts break the problem into small slices and distribute it over many servers in hopes of speeding up the process – but they may not provide a comprehensive fix for performance problems.



Simply put, an effective solution to the challenge of database performance optimization likely resides between the two approaches of infrastructure modification and ETL process streamlining. Regardless of the database involved, three basic requirements must be addressed.

REQUIREMENTS FOR LATENCY REDUCTION IN DATABASES

First, database data needs to be laid out in a manner that expedites data access – that is, reads and writes to data tables. Many companies have been deploying all-flash storage or other “in-memory” storage infrastructure at the recommendation of their vendors with the expectation that this will produce huge performance advantages. Often, the only real value of this infrastructure strategy comes in the form of revenue to the vendor.

Databases, whether deployed on disk or in memory, tend to operate quite efficiently at first, but they become noticeably slower as data amasses in table structures over time – from days to weeks, depending on the volume and complexity of transactions. While storing database contents in silicon reduces the rotational latency accrued to disk-based storage, it is an imperfect solution for achieving performance. Certain types of memory, such as flash, that are being pressed into service as database storage are optimized for data reads, but not data writes. This can make them less efficient at processing workloads that are write intensive. System RAM (DRAM) is better for reads and writes, but it is volatile and requires power backup to hedge against catastrophic loss.

Something to keep in mind when considering I/O performance is latency. Latency accrues not only to storage I/O, but also to raw I/O – where it is a function of the way that concurrent I/O requests (databases are usually packed with concurrent operations) are handled by essentially sequential processes. Most flash memory-based in-memory database solutions do nothing about sequential processing and flash is inherently less well suited to writes than it is to reads. By contrast, DataCore Software’s Adaptive Parallel I/O technology delivers parallel I/O processing of concurrent raw I/O operations, delivering significant improvements in data reads and writes into in-memory databases and caches. This has been validated by two recently audited Storage Performance Council™ SPC-1™ benchmarks submitted by DataCore in late 2015 and mid 2016.¹ Notably, DataCore achieved a record high on performance – at 5,120,098.98 SPC-1 IOPS™ with latency at 0.28 milliseconds – using its Parallel I/O technology. It is now the fastest performance and response time combination ever measured on the SPC-1 enterprise workload benchmark. For reference, the next fastest systems, which were either multi-million dollar systems or all-flash arrays, reported latency times that were more than 3 to 18 times slower.

SPC-1 Numbers	IOPS	Price/ Performance dollars/IOPS	Av response in milliseconds at 100% load	Supplier provided total Price
DataCore Parallel Server 2-node	5,120,098.98	0.10	0.28	\$506,525.24
Huawei OceanStor 18800 V3	3,010,007.37	0.79	0.92	\$2,370,764.00
Hitachi VSP G1000 HAF	2,004,941.89	1.00	0.96	\$2,003,803.00
DataCore Parallel Server 1-node	1,510,090.52	0.09	0.10	\$136,759.00
Kaminario K2 v4 Flash	1,239,898.00	0.80	2.95	\$997,348.00
DataCore SANsymphony HA FC	1,201,961.83	0.10	0.22	\$115,142.76
Huawei OceanStor 18800	1,005,893.43	2.78	5.17	\$2,794,971.80
IBM Power 780 Server with SSDs	780,081.02	4.56	18.9	\$3,557,709.00
NetApp FAS8080 EX all-flash	685,281.71	2.77	1.23	\$1,897,999.00
Huawei OceanStor 6800 v3	650,987.88	2.29	3.36	\$1,488,036.50
Hitachi VSP all-flash	602,019.47	2.86	0.72	\$1,721,620.95

SOURCE: *The Register*, 15 June 2016ⁱⁱ

Even with latency-reducing parallel I/O handling, data will accrue to a database over time, slowing its performance. That is when it may help to “de-stage” data from the database. De-staging is a process by which older, unchanging, data sets may be moved out of the primary database, replaced with a checkpoint or stub. Often, de-staging is a time-consuming process, requiring that the database be quiesced until the datasets to be de-staged are isolated and transferred to new storage targets. DataCore™ Adaptive Parallel I/O technology expedites de-staging by allowing the de-staging process to be completed in an “interrupt-free” environment from system memory to disk or flash. DataCore has recently certified its platform to work on SAP HANA and has thousands of customers running their business on Oracle and Microsoft SQL based environments. In the most recent certification testing with a major analytics database vendor, the vendor remarked about how “insanely fast” DataCore was in de-staging then re-enabling the vendor’s database.

This same I/O acceleration effect accrued to DataCore’s Adaptive Parallel I/O technology also applies to database reorganization tasks and to backups. Database re-orgs are part of the maintenance of any database. However, the maintenance window for completing reorgs is getting much shorter in most firms. In much the same way that DataCore’s Parallel I/O technology enables fast data transfers from system memory to storage with respect to de-staging, it can also expedite database reorgs.

Used in conjunction with DataCore’s own snapshot feature, Adaptive Parallel I/O can be harnessed to dramatically reduce the time required for database backups, which also addresses shrinking maintenance windows. With DataCore’s snapshot technology, it is possible to create a mirror image of a database that is consistent and complete. Adaptive Parallel I/O expedites the replication of the data by accelerating raw I/O handling. That means that the database can go on about its production business with the snapshot being used to generate the backup of the database at a given point in time.

CONCLUSION

By reducing the latency of raw I/O processing upstream of the peripheral storage devices, DataCore's Adaptive Parallel I/O technology can help your databases perform at the highest possible speed. That advantage is achieved using the infrastructure you already own and with little or no additional investment required in all-flash or high-end disk storage gear. DataCore technology will also reduce the latency of ETL processes, including data de-staging, maintenance reorganizations, and even database backups, with the additional capabilities of the DataCore Software-Defined Storage platform, such as snapshots.

In the final analysis, database efficiency is a function of raw I/O handling both in the production environment and in the ETL support chain. The good news is that DataCore's Adaptive Parallel I/O technology can deliver an affordable and effective solution for optimizing the performance of your mission-critical databases.

ⁱ The SPC-1™ benchmark is itself a database workload, providing a better insight than do other testing regimes into the benefits that DataCore's Adaptive Parallel I/O technology can bring to database workload. See: http://www.storageperformance.org/benchmark_results_files/SPC-1/DataCore/A00179_DataCore_Parallel-Server_DN-FC-SAN/a00179_DataCore_Parallel-Server_DN-FC-SAN_SPC-1_executive-summary_r1.pdf

ⁱⁱ http://www.theregister.co.uk/2016/06/15/datacore_drops_spc1_bombshell/