

Accelerating Analytics with KX and Intel Optane Persistent Memory

The onslaught of data continues. Predictions vary but some estimate the world will be generating 175 zettabytes annually by 2025 of which 90 would be from IoT alone, so drawing context and insights quickly is ever more important, and ever more challenging.

What's driving this demand is that organizations are deploying high-resolution sensors to improve the efficiency and effectiveness of their operations and products through analytics and automation.

Organizations are challenged with supporting analytics on higher data volumes cost-effectively. To address these challenges organizations had few options:

- Using more and larger infrastructure – more servers with a lot more DRAM to cache more data
- Use expensive high-performance storage and network infrastructure

Now we have a new option for consideration of performance demanding customers – Intel® Optane™ persistent memory (Intel® Optane™ PMem).

We partnered with Intel to test their persistent memory with our Streaming Analytics platform running mission critical industrial workloads.

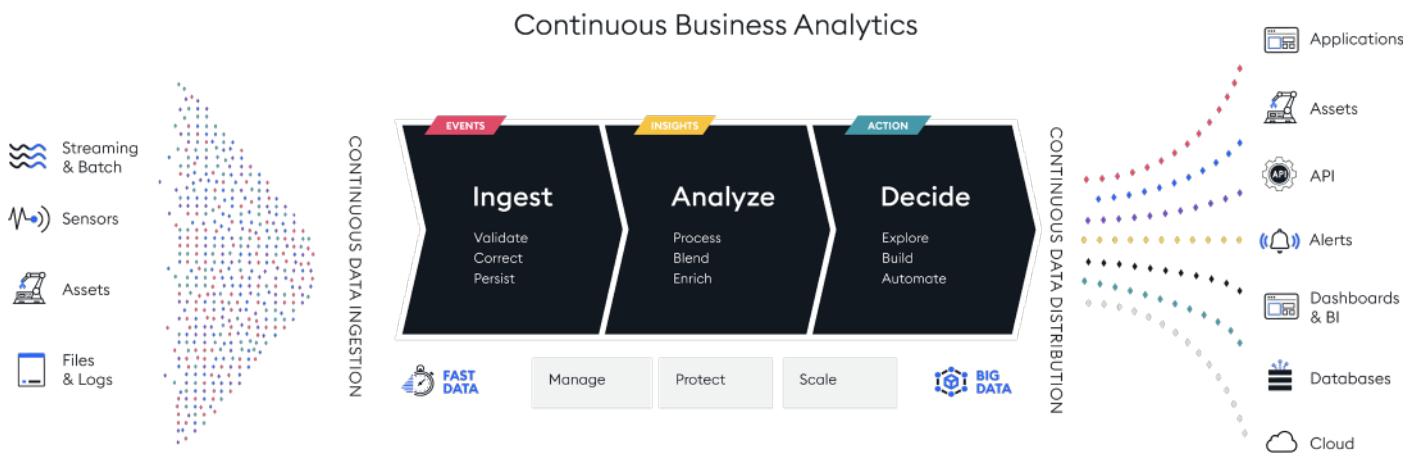
Key Takeaways

Use of Intel Optane persistent memory as block storage device with KX Streaming Analytics delivers 4x to 12x improved analytics performance over high performance NVMe storage, similar performance as DRAM for query workloads. We found that for key data processing workloads DRAM requirements were significantly reduced.

This enables organizations to support more demanding analytic workloads on more data with less infrastructure.

About KX Streaming Analytics

KX Streaming Analytics has been at the forefront of helping organizations ingest, process, analyze, and deliver actionable insights on massive amounts of real-time and historical data. KX Streaming Analytics is an integrated platform that runs in the cloud and/or on-premises to enable organizations to reduce costs and improve productivity through anomaly detection, analytics, predictions, corrective action and automation based on both real-time and historical data. It provides full life-cycle data ingestion, processing, analytics, and data management enabling mission-critical operations, 24x7, with no downtime and no data loss - all of your data, in any environment, instantly.



Source: KX

About Intel Optane persistent memory

Intel Optane persistent memory redefines traditional architectures, offering a large and persistent memory tier at an affordable cost. With breakthrough performance levels in memory intensive workloads, virtual machine density, and fast storage capacity Intel Optane persistent memory accelerates IT transformation to support the demands of the data era.

- Intel Optane persistent memory provides an affordable alternative to expensive DRAM, that can deliver huge capacity and accommodate demanding workloads and emerging tools like in-memory databases (IMDB)
- Intel Optane persistent memory is propelling major infrastructure consolidation. The increase in memory size from Intel® Optane™ media provides the opportunity to consolidate workloads to concentrate them on fewer nodes, ultimately saving on more deployments and maximizing previously underutilized processors due to memory constraints
- As an entirely new memory tier that has new properties of performance and persistence, architects and developers are using Intel Optane persistent memory as a springboard for innovation taking advantage of new usages around restart and replication, ground-breaking performances from breaking system bottlenecks that used to severely constrain workloads and making use of what could be considered the world's tiniest, but fastest storage device sitting on a memory bus

Summary Results

Analytics

- Delivered query performance within 10% of DRAM involving data retrievals with table joins
- Delivered 4x to 12x read performance than 24 NVMe storage
- DRAM delivered 3x to 10x higher performance when performing single threaded calculations and aggregations on data

Data Processing and I/O operations

- Processed 1.6x more data per second than NVMe only storage where data was read from volumes residing on Intel Optane persistent memory and written to NVMe storage
- 2x to 10x faster reading data from files in parallel
- Similar performance of reading data as page cache (DRAM)
- Single threaded file write performance was within 10% in both configurations
- Multi-threaded file write performance was 42% slower

Infrastructure resources

- Required 37% less RAM to complete key I/O intensive data processing
- Required no page cache for querying or retrieving data stored volumes residing in Intel Optane persistent memory

Business benefits

- Collect and process more data with higher velocity sensors and assets
- Accelerate analytics and queries on recent data by 4x to 12x
- Reduce cost of infrastructure running with less servers and DRAM to support data processing and analytic workloads
- Align infrastructure more closely to the value of data by establishing a storage tier between DRAM and NVMe or SSD backed performance block storage

Organisations should consider Intel Optane persistent memory where there is a need to significantly accelerate analytic performance beyond what is available with NVMe or SSD storage.

Setup and Evaluation

Hardware Setup

We setup 2 different system configurations as follows:

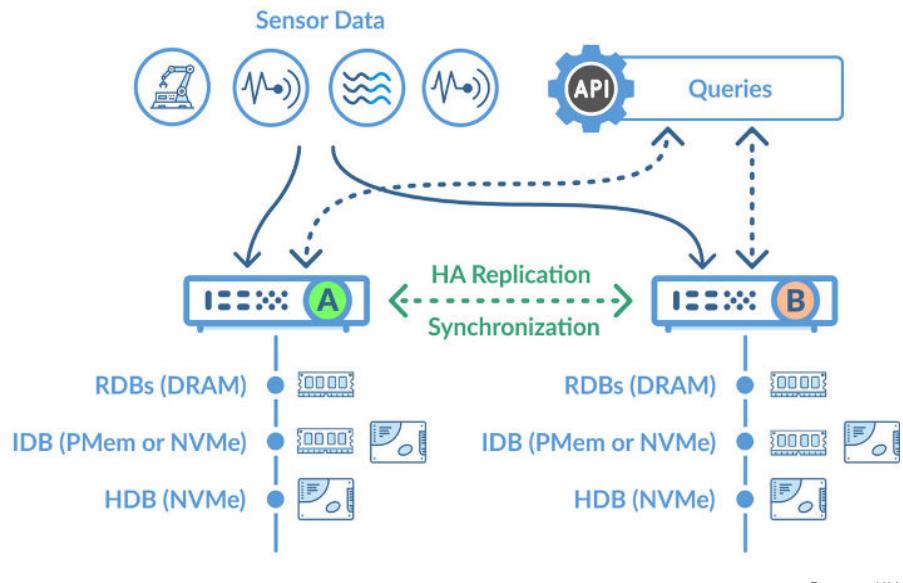
| Configuration | Baseline - Without Optane | With Optane |
|------------------------------|---|---|
| Server Operating System | Supermicro 2029U-TN24R4T Centos 8 | |
| RAM | 768 GB RAM (2666 MHz) LRDIMMs | |
| CPU | 36 physical cores: 2 x Intel® Xeon® Gold 6240L Gen.2 2.6 GHz CPU Hyper-threading turned on | |
| Optane Persistent Memory | N/A | 12 DIMMs x 512 GB NMA1XXD512GPS |
| Log | | Intel Optane ² persistent memory 3 TB EXT43 DAX |
| Intraday Database ≤ 24 Hrs | RAID 50 data volume 24x NVMe P4510 NVME 48 TB RAIDIX ERA ¹ Raid Software, XFS 6 RAID 5 Groups, 64 chunk | Intel Optane persistent memory 3 TB EXT4 DAX |
| Historical Database > 24 Hrs | | RAID 50 data volume Same as baseline configuration |

Environment Setup and Testing Approach

We configured a KX Streaming Analytics system operating in a high-availability (HA) cluster processing and analyzing semiconductor manufacturing data as follows. We performed a series of tests on both configurations for ingestion, processing, and analytics. Tests were done with the same data and duration.

Publishing and ingestion

- Publish and ingest over 2.25M sensor readings in 894 messages per second, 2.5 TB per day
- Ingesting sensor trace, aggregation, event, and lot data using 4 publishing clients from a semiconductor front-end processing environment



Source: KX

Analytics

- 81 queries per second spanning real-time data, intraday data (< 24 hours), and historical data
- 100 queries at a time targeting the real-time database (DRAM), intra-day database (on Intel Optane PMem), and historical database (on NVMe storage)

Data Processing

- Performing a data intensive process involving reading and writing all of the data ingested per day

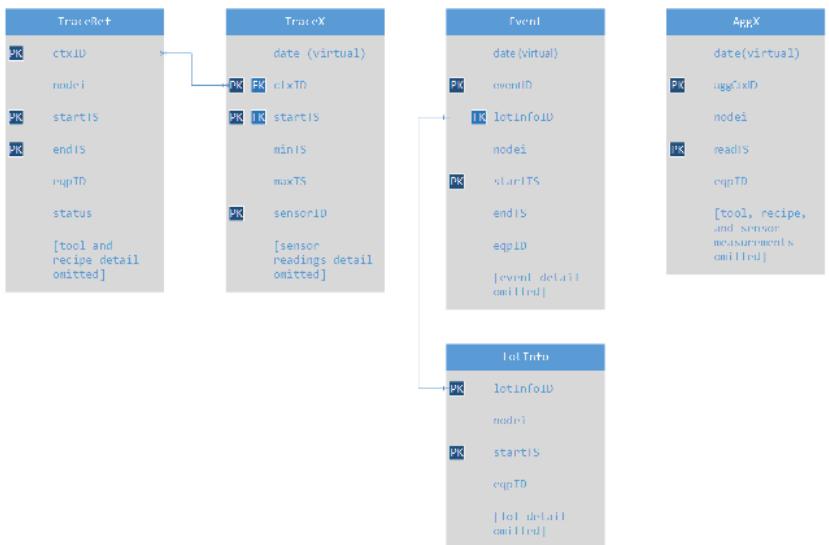
High Availability and Replication

- System running 24x7 with real-time replication to secondary node
- Logging of all data ingested to support data protection and recovery
- Data fed to two nodes with data mediated to ensure no data loss in event disruptions to the primary system

Data model and ingestion

Data model used for workload involved multiple tables representing reference or master data, sensor reading, event and aggregation data. This relational mode is used in fulfilling in streaming analytics and queries spanning real-time and historical data. KX ingests raw data streams, processes and persists data into the following structure. For efficient queries and analytics, KX batches and stores data for similar time range together using one or more sensor or streaming data loaders.

The tables and fields used in our configuration are illustrated to the right.



Source: KX

Test Results Details

Reading and writing to disk

We used the kdb+ nano I/O benchmark for reading and writing data to a file system backed by block storage. Nano benchmark calculates basic raw I/O capability of non-volatile storage, as measured from kdb+ perspective. Note the cache is cleared after each test, unless specified.

Information about the benchmarking tools is available here: <https://github.com/KXSystems/nano>

Please note that with most block storage devices data is read into page cache so that data can be used by the application. However, reads and writes to Intel Optane persistent memory configured as block storage bypass page cache (in DRAM). This improves overall system performances and reduces demands on Linux kernel in moving data in/out of page cache and overall memory management.

Read performance

(Intel Optane persistent memory as block device vs. NVMe storage)

- 2x to 9x faster reading data from 36 different files in parallel.
- Comparable to retrieving data from page cache (near DRAM performance)
- 41x better for reading a file in a single thread.

Write performance

- 42% slower than NVMe devices, due to striping only across 6 DIMM devices vs 24 NVME drives.
- Similar single threaded write performance across the two configurations.

| | Before | | After | | Comparison ⁴ | |
|---|--------|---------|---------|---------|-------------------------|---------------|
| | NVMe | NVMe | PMem | PMem | PMem vs. NVMe | PMem vs. NVMe |
| Threads | 1 | 36 | 1 | 36 | 1 | 36 |
| Total Write Rate (sync) | 1,256 | 5,112 | 1,137 | 2,952 | 0.91 | 0.58 |
| Total create list rate | 3,297 | 40,284 | 4,059 | 30,240 | 1.23 | 0.75 |
| Streaming Read (mapped) | 1,501 | 12,702 | 61,670 | 118,502 | 41.08 | 9.33 |
| Walking List Rate | 2,139 | 9,269 | 3,557 | 28,657 | 1.66 | 3.09 |
| Streaming ReRead (mapped) Rate (from DRAM for NVMe) | 35,434 | 499,842 | 101,415 | 479,194 | 2.86 | 0.96 |
| random1m | 828 | 12,050 | 1,762 | 24,700 | 2.13 | 2.05 |
| random64k | 627 | 8,631 | 1,905 | 36,970 | 3.04 | 4.28 |
| random1mu | 607 | 10,216 | 1,099 | 14,679 | 1.81 | 1.44 |
| random64ku | 489 | 6,618 | 1,065 | 8,786 | 2.18 | 1.33 |

Query performance

We tested query performance by targeting data that would be cached in DRAM, on Intel Optane PMem, and NVMe drives, with parallel execution of each query using multiple threads where possible.

Each query involved retrieving trace data with a range of parameters including equipment, chamber, lot, process plan, recipe, sequence, part, sensor, time range, columns of data requested. The parameters were randomized for time range of 10 minutes.

Query response times using Intel Optane persistent memory were **similar to DRAM** and **3.8x to 12x faster** than NVMe.

| | DRAM RDB 2 query processes | PMem IDB 8 query processes | NVMe HDB 8 query processes | PMem vs. DRAM ⁵ | PMem vs. NVMe |
|------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|------------------|
| 1 query at a time | | | | | |
| Mean response time (ms) | 23 | 26 | 319 | 1.17 | 12.10 |
| Mean Payload Size (KB) | 778 | 778 | 668 | 1 | 1 |
| 100 queries at a time | | | | | |
| Mean response Time (ms) | 100 | 82 | 310 | 0.82 | 3.77 |
| Mean Payload Size (KB) | 440 | 440 | 525 | 1 | 1 |

Two real-time database query processes were configured matching typical configuration, with each process maintain a copy of the recent data in DRAM. Additional real-time processes may be added to improve performance with higher query volumes at the cost of additional DRAM.

Data processing performance

KX Streaming Analytics enables organizations to develop and execute data and storage I/O intensive processes. We illustrate the performance of leveraging a mix of Intel Optane persistent memory with NVMe storage vs NVMe only storage configuration when reading significant volume of data from the intra-day database and persisting it to the historical database on NVME storage.

Optane improved data processing time by 1.67x and reducing the amount of RAM required by 37% by reading data from Intel Optane persistent memory and writing to NVMe backed storage.

| | Before NVMe Only No Optane | After PMem & NVMe | PMem vs NVMe Only ⁶ |
|-----------------------------------|---|------------------------------------|--------------------------------|
| | | | |
| Data processed (GB) | 2,200 | 3,140 | 1.43 |
| Processing Time (minutes) | 24.97 | 21.40 | 0.86 |
| Processing time GB/s | 1.47 | 2.45 | 1.67 |
| Max DRAM Utilisation ⁷ | 56% | 35% | 0.63 |



Notices

¹ We used software RAID from RAIDIX to deliver lower latency and higher throughput for reads and writes over and above VROC and MDRAID. More information is available at: <https://www.raidix.com/blog/kx-streaming-analytics-platform-raises-its-performance-with-raidx-era/>

² Intel Optane persistent memory configured in App Direct Mode as EXT4 volume single block device.

³ In our testing we found EXT4 performed significantly better than XFS, with EXT4 performing 1.5x to 12x better than XFS

⁴ Higher number is better. Factor of 1 = same performance. Factor of 2 = 200% faster than comparator.

⁵ Higher number is better. Factor of 1 = same performance. Factor of 2 = 200% faster than comparator.

⁶ Higher number is better. Factor of 1 = same performance. Factor of 2 = 200% faster than comparator.

⁷ Maximum DRAM utilisation as measured by the operating system during the process and is primarily a function of amount of data that needed to be maintained in RAM for query access. The faster the completion of the process the less RAM that is required on the system.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations, and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services, and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications, and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel, the Intel logo, Intel Optane, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

3D XPoint is a trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© 2020 Intel Corporation Printed in USA XXX/XXX/XXX/XXX Please Recycle XXXXXX-XXX-XXXX