How to Optimize Capital Markets with Intel® Optane™ Technology

A Day in the Life of Financial IT

intel.

8:15 AM

*"Black. Always black.
No cream, thank you.
I need as much caffeine
per ounce as possible.*

*As soon as I walk in,
the gears are already cranking,
crunching after hours numbers.*

*But the real race begins
when that bell rings."*

Markets are volatile. The time to make money is when they spike or drop. IT infrastructure can capture millions of financial messages per day and must be able to handle these staggering volumes.[1]

In January of 2014, the average monthly financial message volumes including securities and payments were just over 20 million messages[2] and by June of 2021, they grew to over 40 million average monthly messages,[3] over a 90% increase in volume.

Every day, financial firms face a tsunami of real-time tick data hitting their infrastructures. The data needs to be instantaneously disseminated across all of their trading and analytics platforms.

And it's more than tick data. Today everyone trades on news. And there's more and more machines reading this news and these patterns—ingesting data, analyzing sentiment and making recommendations.

In addition to all of this public data, firms also have vast amounts of extremely valuable private data. Operational data for risk evaluation, for instance. The secret sauce is to mix them all, including unstructured, alternative sources like satellite news, social, sentiment, etc.—and do it at scale. Only then does a clear picture begin to form, and firms can be prescriptive around how to gain market advantage and how to deliver unique customer experiences.

The bottom line is that today's financial institutions are technology companies. They realize they need innovative platforms that enable them to revolutionize the way they create value and deliver services.
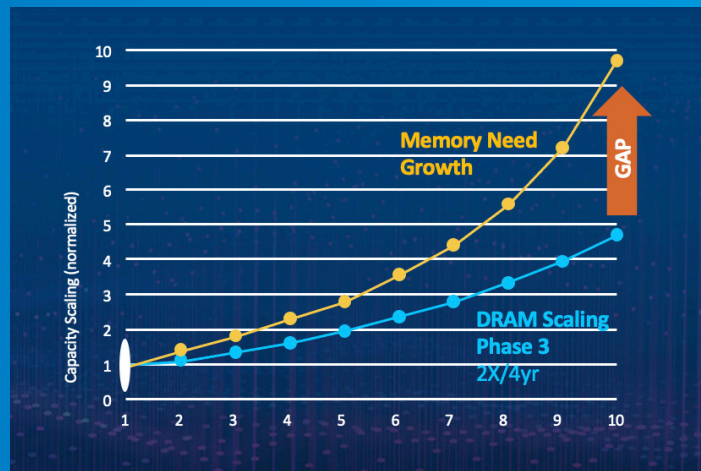
And it all stems from data. So now all your team has to do is manage this explosion of internal and external data, secure that data (and the infrastructure), and deliver new capabilities to the business units, all within a certain power, space, and cost envelope. There is no choice but to ensure the technology infrastructure keeps up with the needs of today and the growing needs of the future.

The bottom line is that today's
financial institutions are
technology companies.
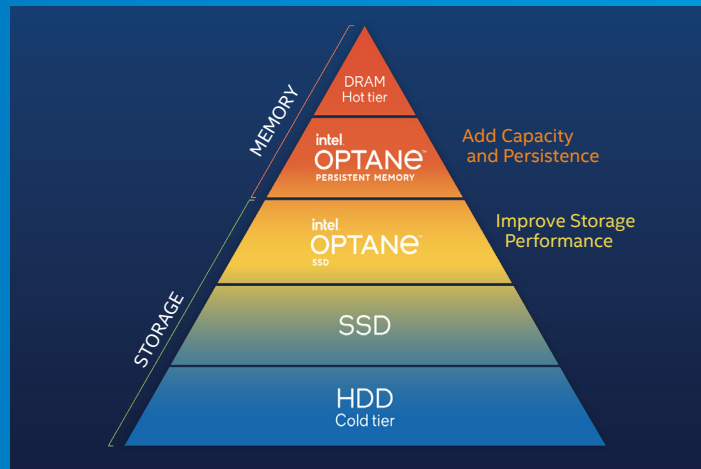And it all stems from data.

Gap Between Data and Memory is Increasing

Memory Need Growth

GAP

DRAM Scaling
Phase 3
2X/4yr

Capacity Scaling (normalized)

Source: Intel results may vary

Re-architecting the Memory/ Storage Hierarchy

MEMORY

DRAM
Hot tier

intel
OPTANE
PERSISTENT MEMORY

Add Capacity
and Persistence

intel
OPTANE
SSD

Improve Storage
Performance

STORAGE

SSD

HDD
Cold tier

10:00 AM

"Maybe we could get ahead if the quants weren't sitting around all of the time. If I hear one more time 'We're waiting on data,' I'll scream."

Traditionally, real time analytics were done solely on current data. The hot data lived in DRAM, while older data lived in storage.

But now, quantitative analysts want the ability to instantly analyze yesterday—or last years'—data in relationship to today's. However, too often they're stuck with today's lightning fast data and yesterday's tortoise-slow data. Your IT team has this problem. And so does every bank, every hedge fund, every asset manager in the world. The core issue is that hot data sits in DRAM and is only nanoseconds away. Everything else is in storage, microseconds away. Not fast enough.

Data is at the core of all digital transformation, in which artificial intelligence (AI) and machine learning (ML) underpin business operations and insights. It represents an opportunity to change how a company operates. In fact, IDC reports that by 2025, the global datasphere will be 175 ZB with nearly 30 percent of it real time.[4] Financial services data is growing at a 26 percent compound annual growth rate (CAGR)![4]

Digital transformation means new apps and new capabilities—and embedded security to help protect all of this data. Data makes all of this possible, but old infrastructure and software aren't up to the task. Data also brings its own challenges, and new hardware and software can help you tackle these challenges.

As an increasing number of businesses engage in digital transformation, digital transformation fuels accelerating demand for compute, and with it exponential demand for the memory to support that compute.

Business data is also increasing exponentially, but traditional DRAM is not scaling to meet this demand. The growth rates for DRAM-density slow over time because it has become costly and complex to scale. And even as the growth in DRAM capacity slows, server memory size continues to increase as servers with more than 0.5 TB of memory are set to dominate the market. Simultaneous with these trends, the critical, data-intensive workloads that run on these servers need to marshal more hot data in memory (just as DRAM at scale grows expensive and limited in capacity).

These realities are converging into a need for a new memory tier to provide the large capacity and scale of storage with the speed and latency characteristics of DRAM, coupled with hardware encryption. The solution is Intel® Optane™ technology.

The first new memory and storage technology in 25 years, it comes in two flavors: Intel® Optane™ persistent memory (Intel® Optane™ PMem) and Intel® Optane™ Solid State Drives (Intel® Optane™ SSDs.) These technologies essentially add two layers of memory and storage to close the capacity, cost, and performance gaps between DRAM (nanoseconds) and NAND SSDs (microseconds). If we think of this in the view of a computing pyramid - Intel Optane PMem sits right below DRAM in memory, providing memory-like performance at a lower cost but with NAND-like persistence, while Intel Optane SSDs sits above cold storage, delivering faster data writes and information retrieval (reads and writes simultaneously), erasing bottlenecks even under heavy loads.

Quants make a lot of money and shouldn't be sitting around and waiting for yesterday's news. The closer they can get to more real-time and recent data, the more aggressive they can be trading. If they're waiting, they've lost the market advantage. At the core, this is why financial companies are becoming technology companies.

11:45 AM

"Another innovation meeting today. Still stuck on data. I mean, we keep track of every tick of every stock on every exchange. Globally. Forever.

That's our ever-growing data set. And I need as much of this as possible in memory, not storage—it has to be that fast. But with all of this data, we keep running into bottlenecks— and we fix it in one place and it moves to another. I think we need a rethink from the ground up, the whole enchilada.

Is it lunch time yet?"

Historically, financial services has been the leader in technology transitions. It's here that you consistently see new competitive trends and new innovation before they make their way to other verticals.

The hunger for innovation has led firms to completely rethink their memory and storage architecture. To paraphrase Einstein, the technologies that brought us to today will not solve the problems of tomorrow.

This pyramid approach to tiered data management with Intel Optane technology modernizes data center architecture by filling the gap between high-performing volatile memory and lower-performing non-volatile NAND storage (SSDs) and hard disk drives (HDDs). It augments DRAM with performance and capacity (Intel Optane PMem scales in TBs; DRAM in GBs.)

When real-time data hits DRAM, it needs to continuously flow out (as DRAM is scarce), but still needs to be memory bus accessible, so it's quickly "quant queryable." No longer does it trail down to slower SSDs. By placing data

closer to the compute source, you have a much faster, much more agile infrastructure.

Get faster, more consistent responses, more workloads per server — and better capital investment. For your IT team, this consistency will be key. As they add more data and workloads, it's stressing the system, and if this leads to degradation of any workload, it's unacceptable. It's critical to see the same response, for one—or thirty—workloads.

This gives your quants latitude: They can research the autonomous vehicle sector, for instance, and run both hot and historical market data across the board. Then they can layer on news data (including audio and video) and social sentiment from a wide variety of sources. All at the same time.

Intel Optane persistent memory and Intel Optane SSDs allow firms to bring in these kinds of massive, disparate data sets and efficiently access them in a predictable time window — pure gold for your team.

Intel Optane persistent memory and Intel Optane SSDs allow firms to bring in massive, disparate data sets and efficiently access them in a predictable time window.

2:00 PM

"This is the trick. If I can get more data closer to the CPU with Intel Optane PMem and Intel Optane SSDs, the quants can have more data faster!

And if I can deploy 50 workloads instead of 40 on a server, my performance— and budget(!)—is looking a lot better."

Intel® Optane™ technology is deployed in multiple form factors, including persistent memory in a DIMM form factor, or an NVMe SSD that sits on the PCIe bus.

Intel Optane persistent memory delivers a unique combination of affordable large capacity and support for data persistence. It acts as a bridge between the performance and high cost of DRAM and the capacity and low cost of storage. The advantage is that Intel Optane PMem provides a lower total cost of ownership (TCO) than large amounts of DRAM memory or additional servers.

Intel Optane PMem used in Memory Mode, enables applications to make use of it as expanded system memory, where the higher performance DRAM is used as a cache for the persistent memory. Using Intel Optane PMem in Memory Mode does not require applications to be modified. In today's work-from-home reality, Memory Mode is a perfect fit for virtual desktop infrastructure (VDI).

In App Direct Mode, Intel Optane PMem delivers Intel Optane technology as persistent memory for applications and operating systems that recognize two types of memory (volatile and non-volatile) and can direct data reads or writes to the memory resource most suited to the task. Operations like analytics with intelligent data placement that needs to be persistent can be routed to Intel Optane PMem, while those that require lower latency and don't need permanent storage can be executed on DRAM. The ability to have permanent data storage in memory is what makes this architecture unique. This gives developers a new way to think about the tradeoffs between capacity and data persistence or between memory and storage.

Intel Optane SSDs enable a new storage tier between Intel Optane PMem and traditional flash storage or NAND SSDs that offers fast caching or fast storage of hot and warm data. In contrast to traditional NAND-based SSDs, Intel Optane SSDs provide high random read/write performance. They offer low latency, higher drive writes per day, higher endurance and consistent responsiveness – even under heavy loads.[5] And, unlike NAND, Intel Optane SSDs can read and write simultaneously without performance degradation. This makes Intel Optane

technology unique: memory-like performance with NAND-like persistence, ideal for demanding storage environments.

Now the quants are happy. With this new architecture, you can efficiently tier data, so the hottest data is in DRAM and Intel Optane PMem, and frequently used data is on fast storage with Intel Optane SSDs. And Intel Optane PMem as a cache in front of storage delivers speed—and determinism, so you can reliably count on every transaction every single time. Most importantly, the hot and warm data is defined by the quants, not by the volume of DRAM. Now you can run all the analytics you desire on all of this new, instantly accessible, historical data.

Now the quants are happy. With this new architecture, you can efficiently tier data, so the hottest data is in DRAM and Intel Optane PMem, and frequently used data is on fast storage with Intel Optane SSDs.

Intel Optane technology benefits from a rich ecosystem of partners and open-source solutions to deliver even more accuracy, speed and scale.

3:30 PM

"Intel Optane sounds like incredible hardware but how does it benefit my analytics solutions and general purpose infrastructure?"

Intel Optane technology benefits from a rich ecosystem of partners and open-source solutions to deliver even more accuracy, speed and scale. Relevant for financial institutions are ecosystem partnerships such as VMware, MemVerge, KX, and Oracle.

VMware vSAN is an enterprise-class storage virtualization software that provides the easiest path to hyperconverged infrastructure (HCI) and hybrid cloud. vSAN is supercharged by Intel Optane SSDs and Intel Optane persistent memory. VMware requires memory to virtualize sessions, and often firms use less, since memory is expensive or because they have small capacity. **Intel Optane PMem has almost 8X (typical DRAM is 64GB; PMem goes to 512GB) more capacity than DRAM,** providing larger system memory, and an ideal ratio of memory to compute.

The result is faster responses, more consistency, and more VMs per server. The combination is ideal for a VMware VDI software solution; it delivers consistent performance under heavy write VDI workloads. With Intel Optane PMem powering the VMware VDI setup, **you can support 90% more remote virtual desktops. This can lower memory costs by almost a third, and TCO by over 15%.**[6]

MemVerge works with Intel Optane PMem to enable higher density, lower cost memory environments, by providing virtualized software-defined memory pools where data live. It allows **instant scaling by making 100% use of available memory capacity** while providing new operational capabilities to memory-centric workloads.

This makes more data more accessible and also enables MemVerge ZeroIO snapshot technology (ZeroIO) that takes small in-memory snapshots that can quickly recover terabytes of data from persistent memory for system re-start—orders of magnitude faster than from storage.

Now trading and financial market data analytics are not only easier to deploy, but you can train and infer from artificial intelligence (AI) / machine learning (ML) models faster, work with larger data sets in memory, complete more queries in less time and consistently replicate memory between servers.

KX Systems kdb+ is an ultra-high-performance time series relational and columnar database designed for rapid analytics on large-scale datasets in motion and at rest. The columnar design offers greater speed and efficiency and its native support for time-series operations vastly improves both the speed and performance of queries, aggregation, and analysis of structured data. Kdb+ is unique in that it operates directly on the data in the database, removing the need to move it to other applications.

And since it employs virtual memory techniques, it pairs naturally with Intel Optane persistent memory. So, if you write a query, and if the data is sitting in DRAM, it's lightning fast. And if it's sitting in Intel Optane PMem, it's basically as fast. You can even run queries on the edge and its truly real-time, with no latency from data transitioning from an ingestion engine, then storing it in a file format, then writing your query. Now it's one step, with the data in DRAM or Intel Optane PMem— and instantly available.

In a recent whitepaper, KX reported they were able to reduce their DRAM consumption by **37% and speed up the queries 4-12X.**[7] That means your team can use machine learning or traditional regression techniques against a huge swath of data, because it's all right there. The result is accelerated queries, and also reduced cost of infrastructure running with fewer servers and DRAM to support data processing and analytic workloads.

Oracle Exadata is used by nine of the world's ten largest financial services firms to run databases for core banking applications and customer analytics. Exadata's industry-leading performance, scale, availability, and security allows them to accelerate their business, reduce risks, and cut costs.[8]

Oracle Exadata X9M takes these capabilities to the next level by integrating Intel® Optane™ persistent memory into the database. A persistent-memory enhanced scalable architecture with smart system software eliminates performance bottlenecks in crucial customer transaction processing and analytics workloads. X9M uses Intel Optane PMem as another data tier, with data exchange rates similar to direct random-access memory (DRAM) but with four times the capacity per module and the persistence capability.[9]

Oracle Exadata X9M capabilities deliver improved OLTP database performance and cost compared to Oracle Exadata X8M:

72% more IOPS: Enabling new transactional and hybrid database use cases.[10]

42% lower costs per IOPs: Making Exadata capabilities more affordable for many organizations.[10]

33% more cores: Increasing consolidation rates and resource utilization efficiency.[9]

5:00 PM

"At the end of the day, this solution can't just be compelling to IT. It needs to be compelling to the quants.

We have to be in tune with both the technology and the needs of our users."

Intel Optane technology is the key to handling today's data waves—and imperative for tomorrow's data tsunamis:

• Intel Optane technology delivers front office value by accelerating real-time analytics and quantitative insights for trading.

• Intel Optane technology delivers back office value by reducing expenditures through general purpose infrastructure efficiency and node consolidation.

The benefits affect a wide variety of technology—and business—issues:

• **Wealth Management:** Enables faster analytics for differentiated trading strategies, and new customer experiences with personally tailored, real-time management advisors and virtual assistants.

• **Risk Management:** Helps to optimize credit risk evaluation, loss scenarios, and standard deviation of financial portfolios for trading and enterprise risk management.

• **Compliance:** Allows the ability to create automated intelligent reporting for financial auditors; enables Anti-Money Laundering (AML), Know Your Customer (KYC), and stress testing.

• **Security:** Improves cyber security, fraud detection, malware, data leakage, and insider trading detection.

• **Operations:** Enables infrastructure anomaly detection and performance optimization using predicative analytics and machine learning, as well as application development and security vulnerability detection.

That's it. End of the day.

Now your team can head home knowing they can handle this data tsunami. They can optimize architecture. They can empower business units. And they can keep the quants happy.

Intel Optane technologies deliver
more access to valuable data
& insights, while optimizing
infrastructure efficiency and
reducing total cost of ownership.

Sources

[1] https://www.swift.com/about-us/swift-fin-traffic-figures

[2] https://www.swift.com/swift-resource/12366/download

[3] https://www.swift.com/swift-resource/251201/download

[4] IDC. "The Digitization of the World From Edge to Core." Sponsored by Seagate. November 2018.
http://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf.

[5] zhttps://www.intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/optane-ssd-better-option-
storage-cache-tier-paper.html

[6] https://www.techproviderzone.com/cloud-and-data-centers/boost-vdi-with-a-winning-combo-vmware-intel-optane-pmem

[7] https://kx.com/wp-content/uploads/2021/10/Intel_KX_whitepaper_accelerating-analytics-with-KX-and-Intel-Optane-PM.pdf

[8] https://www.oracle.com/a/ocom/docs/engineered-systems/exadata/exadata-in-financial-services-industry-brief.pdf

[9] https://www.oracle.com/a/ocom/docs/engineered-systems/exadata/exadata-x9m-2-ds.pdf

[10] https://www.oracle.com/a/ocom/docs/oracle-introduces-exadata-x9m-portfolio.pdf