

Case Study



A subsidiary of SMART Global Holdings, Inc.

MEMX Brings Innovation to Trading With a Deterministic Exchange Using Penguin Computing® and Intel® HPC Technologies

The Penguin Computing® Tundra® Extreme Scale (ES) platform's extreme density and Intel HPC technology performance minimize latency and jitter and deliver an entire exchange in a single rack

Penguin Computing's Tundra ES Platform

- **2nd Gen Intel® Xeon® Gold Scalable Processor**
- **Intel Server Boards**
- **Intel Optane™ persistent memory**
- **Intel SSD DC Drives**



Executive Summary

MEMX is a new members-owned stock exchange that was initially founded in 2019 by nine companies: Bank of America, Merrill Lynch, Fidelity Investments, Morgan Stanley, UBS, Charles Schwab, TD Ameritrade, Virtu Financial, E*TRADE, and Citadel Securities. According to Jonathan Kellner, MEMX CEO, MEMX was established to bring new competition into the market to drive three effects: lower fees, give its members a voice in market dialogue and decisions, and foster innovation.

Entering a market where established exchanges are large and pervasive requires a strong competitive advantage. MEMX has developed advantage through sound engineering, business acumen, and application of advanced Intel HPC and other technologies. These technologies include a performance-optimized exchange platform provided by Penguin Computing using 2nd Gen Intel Xeon Gold processors, Intel Server boards, Intel SSDs, and Intel Optane persistent memory (Intel Optane PMem).

Challenge

The idea that ‘time is money’ is expressed most urgently in the business of equity and commodity trading. With trading initiated by people and conducted by computers, software, and networks, value can be measured in nanoseconds that traders have little control over. Thus, traders are concerned with the predictability of the behaviors of the exchanges with which they interact.

Several characteristics can affect this predictability—or determinism. Main concerns include:

- Latency and jitter of hardware, including the performance of the processors, memory, storage, and network.
- Software performance related to order entry, matching and market data distribution.
- Operational practices of the exchange business itself.

The Need For Determinism

Deterministic behavior is a relatively recent concern, according to Dom Paniscotti, CTO of MEMX.

“Ten to fifteen years ago,” Paniscotti said, “trading systems executed orders in hundreds of microseconds to tens of milliseconds. At the time, little emphasis was given to jitter or determinism as the overall time to execute orders far exceeded system jitter. Then, as infrastructure and software latencies reduced in magnitude, the jitter became meaningful. Traders found themselves placing orders before their counterparts, only find that those orders would execute later than expected. And so, industry focus shifted from latency reduction to jitter reduction.”

According to Paniscotti, technologies were continually introduced to reduce latency and jitter. Faster, low latency, cut through switches were deployed. Network architectures changed to reduce the number of switches used. Storage evolved from

spinning drives to SSDs. Server network stacks were moved out of kernel space and into user space. Computation was moved out of software and into hardware. And all the while, order traffic and market data bandwidth increased.

Determinism is key to building trust between exchanges and traders. While seeking out “best execution”, brokers seek exchange relationships that not only provide fast execution, but also consistency of execution.

“Our end customers want the most deterministic system they can get,” Paniscotti concluded.

MEMX endeavors to bring predictable trading through its exchange with streamlined order handling and a new trading technology platform designed for determinism.

Solution

While architecture and technology are a critical component of MEMX’s solution, maximizing efficiency of operations also affects determinism. As with any business operation, processes are designed around technologies and other requirements at the time. When technologies advance, processes often lag in evolution or they need to be backward-compatible to maintain requirements across the organization or industry. MEMX addressed such challenges early on.

“We looked at the business challenges that exist in this industry that actually create non-determinism,” Thomas Fay, MEMX COO explained. “We asked ‘why’ a lot. Why is this process or that exchange rule the way it is? Is it because of market structure, a regulatory apparatus, or is it just an artifact of some technical debt being carried? That allowed us to evaluate everything in the design of the business and technology stack and trim the operational aspects quite significantly to eliminate functions that cause delay.”

In the end, MEMX built an efficient, responsive business. But their technology choices formed the foundation of a deterministic exchange.

Building an Exchange Platform From the Ground Up

The financial services industry (FSI) is often an early adopter of new technologies. Companies upgrade critical parts of their exchanges with faster processors, memory, and networks to reduce latency and jitter. But, many FSI systems are very large and complex; global technology upgrades are not feasible, leaving older technologies in place. This technical debt can build layer on layer and impact innovation and operational efficiency.

“Lately, the technology industry has been in a fast innovation cycle, offering new paradigms and capabilities that can significantly transform IT services,” CEO Kellner commented. “Simultaneously, focus on balanced architecture and design strategies has shifted from compute to data.”

Data drives decisions and actions, and the exchanges process a lot of data in real time. As exchanges adopt new technologies and architectures through upgrades, however, fine tuning systems to gain the most out of the latest technologies becomes a challenge due to technical debt.

“We chose Penguin Computing because they are an **agile organization**,” Paniscotti explained. “We needed to work with someone who could not only build a **highly custom platform** but would be a **true technology partner** that would help us with the many technology, engineering and logistics challenges that we faced along the way.”

MEMX did not have historical technology debt to hinder their decisions; they started with a greenfield design. The company’s IT architects were able to analyze afresh how exchanges and the trading markets work. They looked at data flow and location, traffic patterns and densities, processes, and the like, and then leveraged the best and most advanced technologies available. They identified and addressed challenges to efficiency and data availability through innovations from leading technology companies, such as Intel and Penguin Computing. They took advantage of modern capabilities, such as secure cloud infrastructure. Thus, they were able to create a highly predictable trading platform that is also scalable and flexible.

“We chose Penguin Computing because they are an agile organization,” Paniscotti explained. “We needed to work with someone who could not only build a highly custom platform but would be a true technology partner that would help us with the many technology, engineering and logistics challenges that we faced along the way. We have powerful relationships with Intel, Penguin Computing, and others. They are all good at providing a deep partnership with insight into their roadmaps, which enables us to deploy a truly custom platform optimized for our workloads that delivers true value to our clients.”

The MEMX Platform

The MEMX exchange is built on Penguin Computing’s next-generation Tundra ES platform. The platform features Intel server boards, 2nd Generation Intel Xeon Gold processors, Intel SSD DC series drives, and Intel Optane persistent memory. Based on the Open Compute Project (OCP), Tundra ES racks offer extreme high density with more than 75 nodes in a single OCP rack.

“We asked for the latest and most advanced technologies in a high-density platform,” said Paniscotti. “With this design we have a density of five to ten times what we could have built with one-U

servers. So, with OCP, we can run the entire exchange in a single rack.”

Running in one high-efficiency rack means fewer switches and hops between end points, reducing points of latency and points of failure. The system runs on a streamlined version of the Linux operating system that ensures highly stable and deterministic behavior.

“We are always evaluating the speed, jitter, space and power of our systems and seeking solutions at the sweet spot where speed is maximized while those others factors are minimized,” Paniscotti stated. “We work with senior Intel engineers who offer their suggestions about new technologies we should consider. Intel Optane persistent memory is an example.”

Optimizing the Exchange for Performance

One of the key challenges in an exchange is the latency introduced when trades must wait to complete while data is copied from system memory to permanent storage. With today’s message rates, spinning drives and even SSDs create bottlenecks, because of information crossing from the memory bus to the PCIe bus or storage network. This costs precious time, introducing latency and jitter.

MEMX’s exchange integrates high-capacity, high-performance Intel Optane PMem configured in App Direct Mode. With App Direct Mode, applications are able to use PMem as extremely fast, persistent storage. Trades are executed and safely written to PMem at memory bus speed—protecting against data and activity losses in the event of outages—before being copied to disk storage. Speed and persistence eliminate delays in the trading process.

“Intel Optane persistent memory is the next evolution of storage,” Paniscotti commented. “At MEMX, we are able to take advantage of its performance to reduce latency of the entire infrastructure.”

The choice of processors optimizes performance for characteristics of the trading workloads. The design

leverages high-clock-frequency Intel Xeon processors for time-sensitive activities and large-core-count Intel Xeon processors for highly parallelized workloads.

Result

Innovating a new exchange on a greenfield design with the latest and most advanced technologies available today, MEMX enters the market with a strong competitive position. It will drive lower fees and attract customers due to its deterministic platform. By choosing Penguin Computing’s extreme high-density Tundra ES platform and standardizing on Intel architecture, MEMX is able to run a high-performance system at lower cost than with other traditional approaches to an exchange infrastructure. Having a homogenous computing architecture, recovery from failovers are seamless, and service replacements are quick. The approach creates an infrastructure with no single point of failure; one that is designed to lower the cost of ownership by reducing footprint, power, and cooling consumption, and improve availability. Further, the single rack design, along with the use of a closed loop cooling solution, significantly reduced the required networking gear, cabling, and power requirements, resulting in the savings of tens of kVA in power.

By balancing the architecture across a streamlined OS, intelligent choice of processors, volatile and persistent memories, and high-performance data center SSDs in a high-density platform, MEMX can consistently provide customers with a predictable platform on which they can execute their trades at the best price possible.

“The entire system is optimized and balanced for speed and stability across the entire architecture—from quoting to the matching engine to executing trades and eliminating latency when moving data across the platform and to the cloud,” Kellner concluded.

Solution Summary

Members Exchange was engineered to maximize determinism and minimize latency and jitter. With hardware developed by Penguin Computing on its Tundra ES platform, the exchange integrates advanced technologies from Intel and other vendors. These technologies include Intel server boards, 2nd Gen Intel Xeon Gold processors, Intel Optane PMem, Intel SSD DC series drives, and high-speed, low-latency networking.

Where to get more information

- ▶ Learn more about [Penguin Computing's solutions](#).
- ▶ Find out more about MEMX, the Member's Exchange at memx.com.
- ▶ Explore the capabilities of the [2nd Generation Intel Xeon Scalable processors](#) with integrated Intel Deep Learning Boost capabilities for accelerated AI inferencing.

About Penguin Computing

Penguin Computing, a subsidiary of SMART Global Holdings, is expanding the world's vision of what is possible in Analytics/Artificial Intelligence, HPC, Data, and Cloud. We leverage decades of deep and broad experience and expertise to deliver fully packaged, tested, and supported platforms that simplify the emerging technologies required for today's most critical and demanding workloads.

Notices & Disclaimers

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.