



Rethinking Scale-Out with VESNIN for Data-Intensive Era

YADRO VESNIN Server White Paper




March 2018

Key Facts

Data explosion era drives an inextinguishable demand for data processing acceleration, proliferating new workload types that demand better performance than universal servers can provide. Innovative hardware architectures are essential to meet the modern and perspective performance requirements. VESNIN is optimized for in-memory and other data-intensive workloads, increases the value provided by HTAP deployments and helps to optimize TCO associated with high-performance data processing installations.



Memory



CPUs



Drives

8

TB DDR4 in 2U

4

**OpenPOWER
(RISC) Processors**

24

NVMe drives

460

gigabytes/s

48

SMT8 cores

60

gigabytes/s

16

risers

384

threads

12 mln

IOPS

128

modules

balanced

**distribution of
drives among CPUs**

Why Scale-Out Systems Demand Rethinking

We're deep in the age of data explosion — the blasting figures of annual data multiplication in our digital reality have become a common knowledge. In addition to growth of business and personal online activities, there is a rising tsunami of data generated by plethora of sensors and devices connected to the Internet of Things (IoT). Also, the world is facing the rapid development of public autonomous vehicles, and as their adoption spreads, the global pool of robotic automobiles is promising to become another enormous driver of data growth acceleration.

As data growth accelerates, the demand for data processing speed also skyrockets — to avoid turning unprocessed data to an inertial deadweight, the data streams increasingly demand to be processed and reacted upon in real time. Another powerful driver of craving for real-time data handling are users transferring expectations for instant system response from consumer apps to enterprise and public sector experience. One of the key challenges of real-time data processing and analytics is the substantial computational performance required for it.

A dominant approach for raising the data processing performance on the next level is the ongoing shift to in-memory operations. The family of in-memory technologies spans over plentiful software and some hardware development directions. The adoption of in-memory technologies is backed by ongoing price decrease for commodity RAM modules. As Gartner mentions in their Hype Cycle for In-Memory Computing Technology, 2017 report¹, it appears that the number of technologies are approaching widespread use and many more may reach it in the next several years to a decade — “While some IMC technologies are at the Innovation Trigger (thus reflecting a still-nascent state), a fair number of them have already progressed to the Trough of Disillusionment or Slope of Enlightenment. Several have already passed the Plateau of Productivity, signaling a degree of maturity and adoption of technologies typically favored by mainstream, risk-averse user organizations”.

Scale-out era commodity architectures are unable to meet the performance and capacity requirements imposed by modern types of workloads such as in-memory applications. In universal architectures not shaped for specific requirements of data-intensive workloads, it takes too long for data to reach the computing resources to be processed. The advancements in software solutions require new approaches for hardware design to leverage its potential.

Figure 2. Priority Matrix for In-Memory Computing Technology, 2017

benefit	years to mainstream adoption			
	less than 2 years	to 5 years	5 to 10 years	ore than 10 years
transformational	Solid -State Arrays	3D XPoint Event Stream Processing Operational In- Memory DBMS	IMC-Enabled Packaged ERP and F/SCPM Applications In-Memory Computing	
high		Analytical In-Memory DBMS Cloud IMDG Services (imgPaaS) In-Memory Data Grids Point-of-Decision HTAP STT-MRAM	IMC-Enabled Hyperscale Application Architecture In-Process HTAP Real-Time Analytics Scale-Up In-Memory Shared Accelerated Storage	
moderate		Hybrid DIMMs IMC-Enabled Application Infrastructure Solid-State DIMMs Spark	Cloud Event Stream Processing Services In-Memory Integrated Stack Next-Generation Memory	
low				

As of July 2017

© 2017 Gartner, Inc.

Source: Gartner, Hype Cycle for In-Memory Computing Technology, 2017, July 2017



IMC Market Improvement

As increasing number of vendors bring IMC products to maturity to meet the demand for greater performance, the hardware design approaches must adapt to support the technology adoption.

VESNIN: Power of Scale-Up Balanced in Scale-Out Footprint

Aiming for advanced workloads and searching for a new optimal architecture, YADRO created the VESNIN server. The machine is designed with workload-centric approach to offer significant advantages in high-impact application environments. With VESNIN server, YADRO reaches customers who are constrained by limitations of traditional architectures in modern era applications and thus consider adopting innovative architectures and deploying alternative platforms from emerging vendors.

It provides a large memory volume to handle larger data chunks, augments it with high memory bandwidth an incredible local storage performance to speed up the data transfer operations, and incorporates a group of four RISC processors optimized for intensive parallel processing and enterprise workloads.


VESNIN provides great level of flexibility for various scenarios — customers can start from small config on single node and scale up as their applications demand growth, and if the applications require more resources VESNIN can scale out and provide benefit from commodity price/performance ratio.

VESNIN architecture is optimized to balance capacity and performance of all major server resources for data-intensive workloads.

Outstanding architecture of YADRO VESNIN offers impressive performance and capabilities for data-intensive workloads and cost-effective solution for more general enterprise tasks.

**Up to
8 TB RAM**

**Up to
12 mln IOPS /
60 GB/s**



Designed for Data-Intensive Applications

In-Memory Application Infrastructure Platforms

IMC products are rapidly maturing, and market of IMC technologies grows with intensifying competition. As Gartner Market Guide² mentions, “digital business is generating an inexhaustible demand for faster performance, greater scalability and deeper real-time insight, which is boosting the market for IMC technologies, expected to reach almost \$15 billion by 2021 from \$6.8 billion in 2016”. From the application infrastructure platforms side, main market traction is within the segments of In-Memory Data Grid products (IMDG) and in-memory database management systems (IMDBMS). The same Market Guide by Gartner² informs that “in the key IMDBMS segment, competition has further intensified”.

IMDBMS store databases in server RAM and offer both row-store and column-store options. They are usually classified for analytical and operational purposes, although the second segment also tends to include

With VESNIN, you:

Process more data faster

Simplify large-memory installations

Minimize overhead costs

increasing number of analytical capabilities. The main market driver in this segment is SAP HANA, followed by Redis. IMDG products are more lightweight object-oriented options of in-memory DBMS designed for distributed data storage. Most notable IMDG systems in the market are GridGain, Hazelcast, Aerospike, XAP and InsightEdge by GigaSpaces Technologies. Some ISVs offer products for both segments of in-memory application platforms, e.g. Couchbase or TerracottaDB by Software AG.

Bigger memory and higher processing speed — these are the fundamental requirements for all in-memory application infrastructure platforms. VESNIN outstanding memory capacity and bandwidths enables bigger in-memory databases for more complex customer systems. For IMDGs it also helps in handling available data with fewer systems to minimize overhead costs for hardware and software associated with clusterization.



Capacity and performance

With 8 TB RAM and balanced NVMe local storage featuring up to 12 mln IOPS and 60 GB/s VESNIN has become the leading server for data-intensive workloads.

IMC-Enabled Real-time Analytics

Digital business is generating an insatiable demand for faster and deeper knowledge derived in real-time from various data flows. In-memory analytics (BI) software brings data to memory for faster executions of queries run over large datasets. The most advanced systems of that class (HTAP) can handle live streams of data and process operational data as it arrives. Eliminating the need for precalculated datacubes, these systems provide answers in real-time but require a lot of memory.

VESNIN is designed to host significantly more memory than other servers with scale-out footprint — enabling BI and HTAP software run analytics over bigger datasets. From customer perspective, it means more data gets analyzed and deeper

With VESNIN, you:

Analyze more data at once




Achieve deeper real-time insights

Speed up the value extraction from data

real-time insights are achieved. At the end of the day, running BI&HTAP on VESNIN improves enablement of superior business flexibility practices like dynamic pricing and builds the competitive advantage for customer business.

Record-breakingly fast and balanced NVMe local storage of VESNIN brings data up to memory with blazing speed, so time to get the value from data is minimized and analytics are performed with maximum efficiency.

Any PPC64-compatible in-memory analytics software like SAP HANA, Polymatica, etc. will make more value with VESNIN.



Transaction processing

Business and IT transaction streams are critical and intense for any big business. Fast and efficient transaction processing is crucial for business flexibility and competitive strength. It is important for meeting the compliance requirements, attracting new customers and promoting their loyalty.

The high volume of memory in VESNIN stores more transactions for processing. Also, the server can accelerate the operations with hosting more transactional database instances like Redis, PostgreSQL or MariaDB.

With VESNIN, you:

Store more transactions in memory

Accelerate processing by using multiple instances with high bandwidth to drives

Use fewer servers and cut down TCO

Moreover, on VESNIN all instances are having high bandwidth to drives thanks to balanced NVMe subsystem. With VESNIN, customers can use fewer servers and infrastructure to save money.

High-speed VESNIN memory boosts processing performance, while fast local storage accelerates snapshot operations and improves safety of transactional data.



High-speed memory & fast storage

High-speed VESNIN memory boosts processing performance, while fast local storage accelerates snapshot operations and improves safety of transactional data

Workload-independent Benefits VESNIN Offers

Best Density for Best TCO

Bigger memory and higher processing speed — these are the fundamental requirements for all in-memory application infrastructure platforms. VESNIN outstanding memory capacity and bandwidths enables bigger in-memory databases for more complex customer systems. For IMDGs it also helps in handling available data with fewer systems to minimize overhead costs for hardware and software associated with clusterization.

Best memory density in 2U enterprise server market makes VESNIN a TCO leader. Whenever a business needs a high RAM space, few VESNIN servers offer the same volume as a large group of commodity x86 servers. For example, only 6 VESNIN servers with full 8 TB RAM comprise a pool of 48 TB — the same as 32 servers with 1.5 TB each. And VESNIN pool comes with even higher memory bandwidth.

High Memory and Storage Bandwidth Per Core

VESNIN features full OpenPOWER memory bandwidth of 460 GB/s per socket. Available CPU options range from 8 to 12 cores, slicing each core a large segment of memory bandwidth. Balanced local NVMe storage scores outstanding performance of up to 12 mln IOPS and 60 GB/s of bandwidth. Each CPU core gets a big share of IOPS and bandwidth to drives for loading and saving data. Combined, these two features make a boost for data-intensive workloads with appetite for fast processing of large data streams and sets.

Learn more about VESNIN

Find out technical details on VESNIN architecture and more product information at yadro.com/vesnin.


To learn more about how YADRO VESNIN can boost your business, please email our sales team at sales@yadro.com, or complete our demo request form to have us contact you.

About YADRO

YADRO is a technology company committed to influencing the way enterprises address their IT transformation challenges. We design and offer game-changing products with outstanding quality. YADRO offers server and storage products with outstanding features, focused on enterprise workloads, primarily in-memory computing, deep learning, hyperscale datacenter and large volume storage.

We are an OpenPOWER Gold Member and the first vendor in foundation with 4-Socket design. Member of SNIA, Linux Foundation, Gen-Z Consortium, PCI-SIG, OpenCAPI, etc. YADRO established partnership with Avago, Marvell, PLX, LSI, Molex, Echostreams, PMCS, Micron, NVIDIA, Wistron, MSI, HGST, and Seagate.

Key focus of our R&D team lays in new type of memory driven systems area. We prototype innovative high density low-latency disaggregated compute/memory fabrics for ultra-scalable in-memory installations. YADRO is a member of Gen-Z and OpenCAPI consortiums developing an open systems interconnect designed to provide memory semantic access to data and devices via direct-attached, switched or fabric topologies. Next generation products in company roadmap will enable full power of memory class storage technologies for outstanding applications performance.





Visit our website
yadro.com

General inquiries
info@yadro.com

Media
media@yadro.com

Rev. Mar 2018

COPYRIGHT AND TRADEMARK INFORMATION

© Copyright 2018 YADRO. The information contained herein is subject to change without notice. The only warranties for YADRO products are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. YADRO shall not be liable for technical or editorial errors or omissions contained herein. VESNIN is a trademark of YADRO. All other third-party trademarks are property of their respective owners.