

AMD ROME- OPTIMIZED PLATFORMS

Selecting a new Server Platform
That is Workload-Optimized with
the 2nd Generation AMD EPYC
Processor with up to 64 Cores

An IDC InfoDoc, sponsored by Dell EMC | August 2019

Analyst: Ashish Nadkarni

Demanding and Unique Workload Profiles Drive the Choice of an IT Infrastructure Platform

The selection of a suitable infrastructure platform (i.e. server) must start with the processor and processor-centric optimizations (e.g., memory, I/O, connectivity and storage)

Attribute Category	What to look for?
Processor Sub-system Performance	<ul style="list-style-type: none"> • High Core and thread count per socket • High L2 and L3 cache per core • Higher TDP wattage • Memory Performance and Capacity
I/O Bandwidth and Latency	<ul style="list-style-type: none"> • PCIe capabilities for high bandwidth/low latency access to accelerators, connectivity interfaces and storage media • Number of PCIe lanes – flexibility in configuration for optimized performance
Storage Performance and Latency	<ul style="list-style-type: none"> • NVMe direct performance • PCIe lanes for non-direct I/O
Processor Sub-system Security	<ul style="list-style-type: none"> • Processor security to prevent hostile code execution attacks • Memory encryption • Workload and VM-aware isolation • Key management • Operating system and hypervisor support
Operating and Ownership Costs	<ul style="list-style-type: none"> • x86-based ISA eliminates the need to rewrite applications • Single-socket performance for reduced TCO in socket-based licensing • Dual-socket performance for improved ROI in mixed or specialized workload configuration • Based on industry-standard x86 ISA requiring no migration costs • Integration with existing management solution

The processor is a crucial building block for an infrastructure platform. The choice of a suitable platform must start with a critical examination of the processor sub-system. Many times, this is taken for granted, and can have serious repercussions on business outcomes when workload performance or scaling suffers due to lack of capabilities, availability, or resources. Or, in many cases, the infrastructure becomes vulnerable due to processor design shortcomings.

The AMD 2nd Generation EPYC Processor with up to 64 Cores Provides a Critical Building Block for an Optimized, Scalable and Secure IT Infrastructure

Advanced Micro Devices (AMD) introduced its AMD EPYC brand of x86-64 microprocessors in 2017. Based on the company's Zen microarchitecture, the server variant of these processors are specifically targeted for computationally and I/O-intensive datacenter and cloud workloads. With features such as higher core counts, more PCIe lanes, support for larger amounts of RAM, and larger cache memory, these processors support single and dual-socket system configurations. Naples is AMD's first generation EPYC server processor family followed by Rome.



AMD EPYC x86 server processors offer an integrated architecture with higher performance, improved bandwidth, balanced energy consumption, and enhanced security.

Workload Attribute	AMD EPYC capabilities	Naples	Rome
High compute parallelism	Higher core count with configurable TDP	<ul style="list-style-type: none"> • 32 cores • 64 threads • 2 x L3 cache/core (16MB per 4 cores) • 200W TDP 	<ul style="list-style-type: none"> • 64 cores • 128 threads • 2 x L3 cache/core (16MB per 4 cores) • Up to 240W
Faster in-memory workloads	RAM capacity and speed	<ul style="list-style-type: none"> • 2666 MT/s • Max. system memory 18TB 	<ul style="list-style-type: none"> • 3200 MT/s • Max. system memory 32TB
Massive I/O for data access and connectivity	More PCI lanes	<ul style="list-style-type: none"> • Up to 128 PCI Gen 3 lanes • xGMI 10GT/s 	<ul style="list-style-type: none"> • Up to 160 PCIe Gen4 lanes • xGMI-2 at 16 GT/s
Cryptographic isolation between hypervisor and VM	Secure Encryption Virtualization (SEV)	<ul style="list-style-type: none"> • Up to 16 SEV keys per hypervisor 	<ul style="list-style-type: none"> • Up to 509 SEV keys per hypervisor
Security for business-critical data	Secure Memory Encryption (SME)	<ul style="list-style-type: none"> • Not supported 	<ul style="list-style-type: none"> • Configurable

Advantages of the next-gen EPYC (Rome) processor platform

- Single I/O/memory die removes processor subsystem bottleneck, and provides 2x performance over Naples
- Uniform speed across cores, and high bandwidth and low latency network connectivity
- Separate I/O die from Zen2 chiplets* allows flexible core configurations
- Dedicated encryption offload sub-processor provides full stack encryption without any overhead
- Dedicated I/O plus memory die provides improvements for NUMA** workloads while also reducing latency for non-NUMA workloads
- Improved single socket and reimagined dual socket performance reduces TCO

* Chiplets are multiple smaller pieces of silicon that make up a processor vs. the older process of carving new processors from silicon as single chips

** Non-uniform memory access (NUMA) is a computer memory design used in multiprocessing, where the memory access time depends on the memory location relative to the processor. Under NUMA, a processor can access its own local memory faster than non-local memory (memory local to another processor or memory shared between processors). The benefits of NUMA are limited to workloads where the data is often associated strongly with certain tasks or users.

What to Look For in a “Rome-Optimized” Server Platform

A “Rome-Optimized” server platform is one that has been designed from the ground up to take full advantage of capabilities offered by the 2nd generation AMD EPYC Processor with up to 64 cores. Such server platforms offer investment protection as their capabilities are future-ready.



Did you know? Just because the features are on the AMD processor doesn’t mean the server has those features. Each server manufacturer decides which features and capabilities to use. Always check the server specs to ensure you are getting an AMD Rome-Optimized platform.

Rome-Optimized Platforms	Design capabilities that enable optimization
Are built with a high-performance processor sub-system, reduce latency, and deliver fast response times	<ul style="list-style-type: none"> • Support for full processor stack – up to 240W TDP • 20% faster memory speed at 3200MT/s
Enable faster data access, transfer speeds for storage workloads, and faster connectivity for networking and storage workloads	<ul style="list-style-type: none"> • 25% more PCIe lanes and 2x PCIe speed - up to 160 PCIe Gen4 • 60% faster interconnect fabric with xGMI-2 at 16GT/s (Up to 24 GT/s with a supported peripheral)
Are designed to make full use of processor capabilities for integrated end-to-end security	<ul style="list-style-type: none"> • Secure data with AMD Secure Memory Encryption (SME) • Workload VM to hypervisor isolation with AMD Secure Encrypted Virtualization (SEV) - 509 unique keys per hypervisor

Additional configuration enabled by vendor’s design and engineering

- Design makes use of higher sustained TDP of 240W per processor
- Risers optimized to balance workloads (flexible configuration)
- Flexible use of PCIe lanes for network and management access
- Socket direct networking for balanced network I/O
- Direct NVMe configuration eliminates switches or bridges
- Extended speed mode ready for faster peripherals operating at 24 GT/s when available
- BIOS supported AMD Secure Memory Encryption
- Liquid cooling options via partners
- Reduction in deployment time through automation capabilities and system management solutions
- Improvement in datacenter cooling power utilization efficiency (cooling PUE)
- Easy BIOS tuning with workload-optimized server configuration profiles
- Verifiable secure root of trust that is etched at the factory


A Trusted Vendor Delivers a Best-in-Class “Rome-Optimized” Server Platform

Organizations must consider the role of IT vendors as partners when selecting a suitable Rome-Optimized platform. It is crucial for organizations to partner with vendors that are investing in design and engineering capabilities to produce such a platform at scale.

Equally imperative is the level of trust that firms must place on the vendor’s ability to build and deliver the platform.

Key Capabilities of Trusted Vendors

(top 6 showing)



Three Important Factors in Building Trust

Maintaining a secure supply chain. This includes verifying the authenticity of components or parts, procuring components or parts from trusted suppliers, and physically securing the build environment, the system build process, and the process of shipping the system to the customer.

Building security into every design step. This includes incorporating security features in the hardware to prevent malicious attacks as well as developing, testing, and verifying the integrity of the firmware prior to installing it.

Providing proactive updates and system recovery. This includes proactively patching firmware bugs or vulnerabilities and ensuring that updates are delivered and applied in a timely fashion. This also includes providing a restore to a known good state when a system recovery is necessary.



Why wait? Trusted and reputable vendors often take their time to design, build, and deliver a quality product. They focus on getting things right vs. being first to market. It is worth investigating what such vendors have planned for release when investing in new infrastructure.

How to Benefit From Investing in “Right-Fit” Infrastructure

Essential Guidance

- **“Right-fit” IT infrastructure provides a solid foundation** on which firms can accelerate their digital journey. IT can enable their businesses to produce consistent outcomes by making timely investments in a workload-optimized and innovation-packed infrastructure platform, from a trusted vendor.
- Additionally, IDC studies consistently find that **infrastructure operating costs continue to spiral upwards** if businesses do not upgrade in a timely manner.
- **One-size-fits-all approaches to infrastructure platforms can potentially slow down the pace of business innovation** due to performance and scalability. At the same time, an aging infrastructure can quickly become a liability for an organization.
- **Older or generic infrastructure platforms are vulnerable to data breaches.** Recent disclosures have revealed that vulnerabilities whose root lies in the processor design are difficult to patch without incurring heavy penalties.
- **Organizations must therefore discard the “one-size-fits-all” dictum and replace it with a “workload-optimized approach”** to making infrastructure decisions – which are based on technology and business economics. They must rely on a trusted vendor as a partner to deliver a fully-optimized and secure platform. Only then can they enable business outcomes at the pace of innovation.



Breaking the “good enough” barrier: Investing in infrastructure is a business decision

Server operating costs in years 4–6 of deployed life **are more than 10x higher than the initial acquisition cost of the server.**

Compared with continuing to operate installed servers, refreshed servers deliver:

59% lower three-year cost of operations

\$4.66M additional revenue per year per organization (\$123,400 per server)

Potential benefits of a regular, faster refresh cycle (two three-year life cycles versus one six-year life cycle):

33% lower net cash flow

\$14.6M cash flow over six years per 300 servers

<1 YEAR return on investment

IDC Methodology

This IDC InfoDoc provides a summary of an extensive validation process performed by IDC in collaboration with the vendor's teams. IDC relied on the data from the vendor and its own independent research to make statements in this document.

This document is meant to provide a quick set of inferences and insights for IT professionals and business decision makers seeking to perform further due diligence on the capabilities of the product and/or services that have been evaluated in this InfoDoc. However, the goal of this InfoDoc is not to supply detailed hands-on test plans and validation jobs. It is not meant to replace the evaluation process that most businesses will conduct before making any decision to purchase the product and/or services. It is for this reason that this InfoDoc is not designed to be an all-inclusive document on all the capabilities of the product, but rather as a concise document that highlights features/functions of products, their relative performance with respect to a traditional environment and the value these features bring to businesses looking to solving certain problems for Hadoop workloads.

Finally, even though this InfoDoc is a sponsored document, it is not meant to be an IDC endorsement of the product, service or the sponsoring supplier. IDC's opinions are its own and not influenced by the production of this document.