

Quick Take

AWS New Bare Metal Instance: Courting the HPC Crowd by Combining Bare Metal Performance with Virtualized Cloud Functionality

Bob Sorensen, Earl Joseph, Steve Conway, Alex Norton
January 2018

HYPERION RESEARCH OPINION

This Quick Take looks at Amazon's recent announcement of a new EC2 Bare Metal offering that provides users with direct, non-virtualized access to a processor, memory, storage and related networking instance. Amazon uses custom hardware in an effort to wring out the highest possible performance from the basic hardware set up while still offering a full complement of cloud-based software support. This development is a positive one for HPC users looking to migrate workloads to public clouds as Amazon is addressing one of the most vexing hurdles of HPC in the cloud: the performance overhead of running applications in a virtualized, performance limiting environment.

SITUATION OVERVIEW

[Amazon](#) recently announced the availability of a new EC2 Bare Metal offering that provides users with direct, non-virtualized access to a processor, memory, storage and related networking instance. Amazon claims that the new capability is targeting users who want the performance of running on bare metal while still needing access to all the benefits of a cloud infrastructure. Specifically, AWS announced that they were launching a public preview for users to test their Bare Metal i.3 instance, as the first in a series of EC2 instances, with the following specifications:

- Processing - Two Intel Xeon E5-2686 v4 processors running at 2.3 Ghz., with a total of 36 hyperthreaded cores
- Memory - 512 GB
- Storage - 15.2 terabytes of local, SSD-based NVME storage
- Network - 25 Gbps of ENA-based enhanced networking, a custom AWS network interface designed to deliver high throughput and low latencies.

Amazon also states that these new Bare Metal instances will provide the full functionality of the EC2 family and can access the full suite of AWS services including their data analytics products and machine learning development infrastructure as well as support for database, IoT, mobile, and security services. To date there is no information about the pricing scheme for either the i.3 instances or expected subsequent AWS Bare Metal offerings.

- These Bare Metal instances are based on AWS's Nitro system, a collection of AWS-built custom hardware offload and hardware security components that AWS says deliver high performance, high availability, and improved security for critical applications.

- Amazon's growing interest in developing their own hardware was publicly highlighted in 2015 when Amazon acquired Annapurna Labs - an Israeli-based chip design firm - to design a wide range of components to meet Amazon's specific hardware requirements.

FUTURE OUTLOOK

Amazon is not the first cloud service provider to offer bare metal capabilities. At least two other notable cloud service providers, IBM and Rackspace, already have similar solutions. What is different here is that Amazon has designed and uses its own custom hardware in an effort to wring out the highest possible performance from the basic hardware setup while still offering a full complement of cloud-based software support.

- It remains to be seen if other bare metal players or perhaps even new aspirants to the sector will follow suit with their own custom enhancements to their bare metal stable. Lacking such capabilities, these vendors could have difficulty competing with Amazon from a pure performance perspective.

This recent development is a positive one for HPC users looking to migrate at least some of their workloads to public cloud environments. Although performance claims have yet to be verified, Amazon is moving to address one of the larger hurdles in migrating HPC workloads to the clouds: the overhead of running applications in a virtualized, and hence performance limiting, environment.

- Interested HPC users should monitor closely any performance results made publicly available during or after this test phase.

In addition, Amazon has hinted that there may be a wide range of additional bare metal instances on the horizon, offering HPC users the option to pick configurations that most closely match their price/performance requirements.

Finally, and perhaps most important, Amazon has raised the bar on bare metal instance performance that could drive other providers to look for even more innovative ways to drive the price/performance capabilities to the benefit of HPC users everywhere.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.hyperionresearch.com

Copyright Notice

Copyright 2018 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.hyperionresearch.com to learn more. Please contact 612.812.5798 and/or email ejoseph@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.