



Market Insight Report Reprint

NTT Docomo, NEC achieve energy-efficiency gains running 5GC software on AWS Graviton2

November 18 2022

by **Brian Partridge, William Fellows, Melanie Posey**

Working toward ESG and net-zero targets is already commonplace in the telecom sector. Add in regional conflicts that are driving up energy costs, and we have a perfect storm that calls for more drastic changes in network architecture with the goal of increasing efficiency.

451 Research

S&P Global

Market Intelligence

This report, licensed to AWS, developed and as provided by S&P Global Market Intelligence (S&P), was published as part of S&P's syndicated market insight subscription service. It shall be owned in its entirety by S&P. This report is solely intended for use by the recipient and may not be reproduced or re-posted, in whole or in part, by the recipient without express permission from S&P.

Introduction

The mobile telecom and public cloud industries are today closer than ever as hyperscalers (headlined by Amazon Web Services, Google Cloud Platform and Microsoft Corp.'s Azure) are increasingly viable suppliers of the underlying core and edge digital infrastructure needed to deliver in-demand telecom services such as 5G and multi-access edge computing. Hyperscalers can also be solid partners for joint business-to-business/consumer initiatives that combine network services and a highly distributed physical topology with their ever-expanding array of cloud infrastructure, platform and value-added services.

For its part, AWS is touting another interesting advantage versus legacy on-premises virtualization architectures that could trump them all: energy efficiency. A three-way engineering proof of concept (POC) combining AWS-designed Arm chips (Graviton2) with 5GSA core (5GC) software from NEC Corp., and hosted by Japan-based mobile giant NTT Docomo, has delivered impressive results.

THE TAKE

ESG and net-zero mandates are now ubiquitous C-level initiatives in the telecom space as regional conflicts drive global energy costs higher. Against this backdrop, global telcos are desperate for technology services that can materially reduce carbon footprint and energy bills. The move to telecom network cloudification that is well underway (the disaggregation of core and edge network appliances to run as software on general IT infrastructure and ultimately via cloud-native technologies and deployment patterns) has finally catalyzed a generational disruption event in how telecom networks are built and run. This has led to a flood of new participants in what had been a closed and proprietary club of telecom original equipment manufacturers and supporting telecom operational software shops.

Executing on this complex and difficult transformation offers telcos the agility and cost benefits that global enterprises have enjoyed for years by shifting their own workloads to the cloud. The complexity and disruption of refactoring core and edge applications to be truly cloud native is the greatest inertia holding back such transformation at scale. We believe the efficiency results approaching the level of what NTT Docomo has achieved with AWS/Graviton2 and NEC, along with a relatively light engineering effort, will pique interest and accelerate transformation for some operators. Over time, efficiency could end up being one of the industry's most potent digital transformation accelerators and a catalyst for AWS, trumping the near-term promises of revenue growth and agility that has been lackluster for 5G services thus far.

Details

The POC results recently shared by NTT Docomo and NEC include a claimed 72% average reduction in power consumption using AWS Graviton2 processors vs. Intel Corp.'s second-generation Xeon Scalable Processor-based infrastructure for 5G stand-alone core control plane workloads. In speaking with AWS, we learned that the joint engineering effort to effectively optimize the NEC 5GC software for porting to AWS/Graviton2 took about a couple of months from start to finish.

While the general perception of “lifting and shifting” of NFV software from X86 to Arm architectures is that it is complex and costly, in this case NEC reported a relatively smooth port while taking advantage of several features of the Graviton2 chip design such as a high-capacity memory architecture. AWS also noted that mapping one vCPU core to one physical core positively impacted both efficiency and performance results. According to the vendor, it is reasonable to assume that it has similar efforts underway with other 5GC software ecosystem leaders and AWS core network partners.

AWS offered its engineering support to NEC and NTT as it believes that investing in such a partnership accrues material benefits for the company as well as sets the stage for long-term growth of its cloud services within NTT and its peers. While the press release detailing this effort was focused solely on the sustainability results, we also learned that the raw performance results achieved parity or increased for the 5GC workloads vs. X86-based service. AWS indicated that while using Graviton2/3 as part of Outposts deployment (a fully managed edge computing stack operated by AWS within an operator-controlled datacenter) is possible, the most likely optimized path for 5GC deployments in Japan will be hosting in nearby AWS regions or local zones, which brings theoretically limitless capacity for traffic spikes.

The company's early successes and continued innovation (it is now on the third iteration of Graviton) in creating its own chipset designs using foundational Arm intellectual property hasn't gone unnoticed by its rivals. Both Azure and GCP have announced a similar intention to release their own Arm-based server designs in the coming months.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2022 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.