

POWER / PERFORMANCE DETERMINISM

MAXIMIZING PERFORMANCE WHILE DRIVING BETTER FLEXIBILITY & CHOICE

EXECUTIVE SUMMARY

Today's demanding workloads and complex environments leave businesses looking for every ounce of available performance from their systems. While CPU vendors strive to deliver the highest level of performance at all times, CPU manufacturing processes leave some variability that prevents vendors from capitalizing on every last drop of performance in an individual CPU. Because CPU vendors align their outputs around standardized curves, some individual performance upside could be sacrificed to meet a broader marketing model. Through power and performance determinism, AMD now enables IT to better align CPU choices to the exact needs of the workload, whether it requires a specific performance level or whether the desire is to capture every last ounce of performance from the CPU's available headroom.

TODAY'S CUSTOMER NEEDS

Flexibility is a highly desired capability for server platforms. To help businesses drive better economics, every 24-36 months companies will make standardization decisions that will guide their purchasing over the next few years. When these decisions are made, a business must consider both its immediate needs as well as what may be on the horizon, even if it is unsure of what that future holds. Uncertainty in today's IT directions makes this standardization process more difficult as emerging workloads and digital transformation leave less clarity around future requirements.

As businesses standardize, they need a high degree of flexibility in their platforms so that these can be deployed for a wider range of workloads. By limiting the number of standard platforms that procurement and IT can choose from, a business can help hold down its procurement costs, and ultimately its operational costs as there is less variability in the datacenter purchases and ongoing operation.

A key requirement of these next generation workloads is performance. For years, businesses ran servers at 10-20% utilization with a handful of cores and smaller memory footprints, because most applications were "bursty". As applications have changed and matured, we see higher utilization levels, a significantly higher core count, and massive memory footprints, all targeted at workloads that are now streaming and processing in a constant, always-on world. These constant streams of telemetry and data from IoT devices, along with densely populated cloud workloads, bring a

continuous need for high performance. This will only be accentuated by a future that will lean into machine learning, artificial intelligence, data analytics, and high performance compute (HPC) to make sense of all of that data being collected. These workloads operate best when they are maximizing their performance.

Just as businesses are trying to unlock the maximum potential of their compute investments, CPUs should be able to do the same, delivering the best value and price / performance for their hosted workloads.

MANUFACTURING VARIABILITY CAN UNLOCK POTENTIAL

CPU fabrication is variable by nature, with multiple layers, multiple process steps, and highly dense, complicated designs. Because of this complicated process, every die is unique. As a CPU manufacturer, planning for this variability is an important part of the process, and there is much time spent optimizing to a yield curve that delivers the best total output to drive lower cost for end customers. Over time, these standardized curves tend to scale upward as the process is tuned from learnings derived over time; this generally results in performance gains.

Towards the top of that yield curve there will be less variability and less headroom, but where the bulk of the CPUs are sold (mid bin) there will be more headroom and potentially more variability. This variability enables AMD to lean into the extra headroom above the rated bin speed and deliver some additional performance, provided the workload's power profile allows the frequency to boost up a little more, all while remaining within the technical specifications of the CPU.

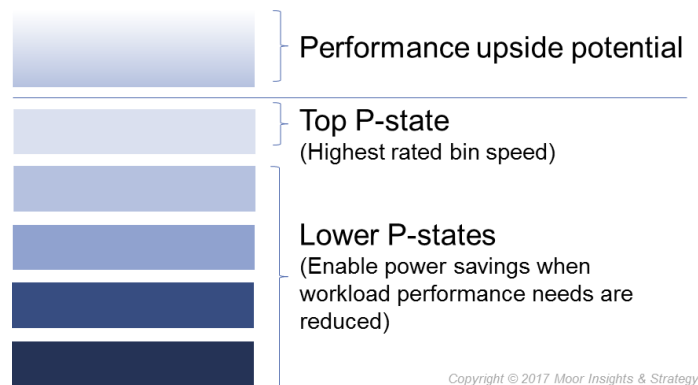
On the consumer side of the CPU world, one might assume that everyone would want extra performance upside. But for servers and enterprise applications, there might be some workloads where maintaining an exact and precise clock speed may be more advantageous than gaining a bit of performance upside. Consistency can be highly valued for some enterprise workloads, because it helps deliver better predictability, especially in very risk-sensitive environments.

For consumer CPUs, performance upside can be driven by overclocking the CPU (providing extra power above the rated power levels), but this practice limits CPU life and drives unpredictable performance. Overclocking is eschewed in the server world where the cost of an outage could negate any performance gains made over the life of the product. Additionally, overclocking can adversely stress the CPU, shortening its useful life—again, less of a concern for some consumers, but a much larger consideration when dealing with enterprise CPUs and platforms.

If a CPU vendor can drive additional upside to the manufacturing process and enable an end customer to unlock a potential for higher performance, there is an opportunity for a selective upgrade with zero risk to the stability of the application or the longevity of the platform. In the case of enterprise servers, to meet the exacting needs of businesses, all variances that are managed will rely on a tradeoff between clock speed predictability and maximum performance, not lifespan / stability versus maximum performance.

All CPUs have multiple performance states (P-states) that allow the CPU to drop the effective clock speed to reduce power consumption at times of lower utilization. There is a maximum rated clock speed, and that is the “marked” speed on the CPU, effectively a governed speed limit for the CPU. But beyond the top bin speed and the lower P-states, AMD’s EPYC CPUs also have some degree of performance headroom that the CPU can tap into above the highest rated clock speed, depending on the workload.

FIGURE 1: PROCESSOR P-STATES & PERFORMANCE HEADROOM



Source: Moor Insights & Strategy

WORKLOAD VARIANCES DRIVE PROCESSOR PROFILES

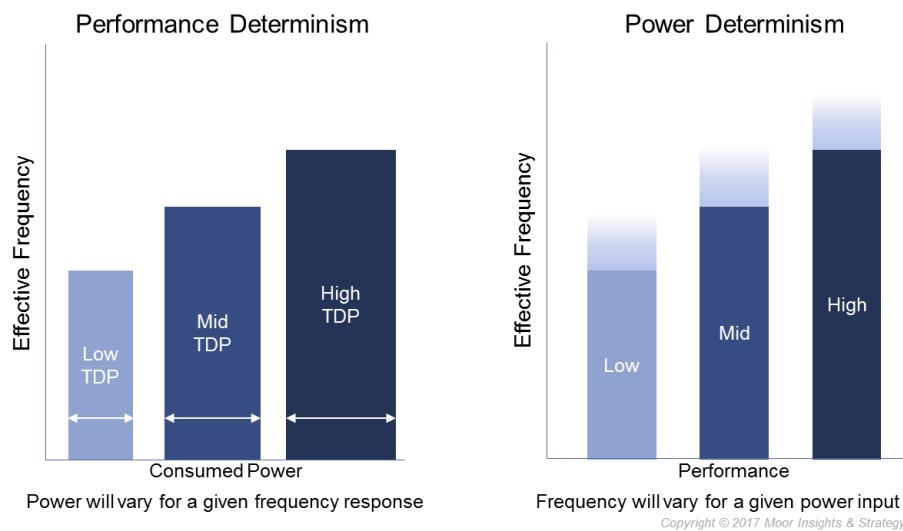
To understand how processor headroom can be turned into additional performance, it is first important to understand workload dynamics with respect to power, because not every workload will have the same power profile. Additionally, not every workload benefits from incremental performance increases in the same way. This may be counterintuitive in a world that is always buying “the next fastest thing”, but there are cases where additional performance may buy very little overall benefit.

Performance for some workloads needs to be consistent, with a bias towards continuous predictability over fluctuating performance increases. Luckily, most IT organizations will test, optimize, and characterize workloads so they have a better understanding of what each requires and how each performs. Some workloads, such as

highly dependent or tightly coupled applications may have less impact from an unanticipated performance increase. When workloads must wait on past instruction results to execute future instructions, acceleration can be less relevant because that wait can offset forward progress. Additionally, some embedded applications are designed around equipment that must have exacting performance levels to ensure the consistent experience on each run. In such cases predictability surpasses performance.

To address both the workloads that crave every last drop of performance and those that demand stringent adherence to pre-determined outcomes, AMD has developed two processor operating modes that enable an IT organization to optimize around either outcome. These two modes, power deterministic mode and performance deterministic mode, are enabled through a selection in the BIOS of the server.

FIGURE 2: THE PERFORMANCE & POWER MODES FOR AMD EPYC



Source: Moor Insights & Strategy

The first of the two modes for the EPYC processor is the **power deterministic** mode. Most enterprise applications will run in this mode, as it delivers the highest performance possible, based on the power budget of the application. Single server enterprise applications will almost always want a power deterministic mode to maximize performance. In some HPC / cloud / clustered workloads, if the work is highly decentralized and not node-dependent, then the power deterministic setting could generate greater aggregate performance across the cluster, deriving outcomes faster.

In power determinism mode, the CPU will run as fast as it can for a given power input / thermal design power (TDP). There will be variability from CPU to CPU as to what that

maximum / effective frequency will be, ultimately resulting in some small degree of variability for the application from server to server, but always boosting performance safely to rated clock speeds.

In power determinism mode, parts will naturally boost, based on the power consumption of the application, as manufacturing variability may allow the CPU to run at a higher effective frequency for a given power input. The determinism is maximized for a pre-determined power load. For instance, a workload that consumes a lot of power on one platform will generally act the same across all platforms (although this may not necessarily be exact), platforms may act differently because of the unique properties of each CPU. For instance, if one server can boost a certain workload by 50-100MHz on one platform, that does not imply that the workload will boost by 50-100MHz under all similar platforms in this mode, but it will generally be close to that range.

The second mode is the **performance deterministic** mode. This mode would be favored for highly dependent clustered applications where the performance deterministic mode will help ensure every SoC behaves the same way for a given workload to keep the processing in sync across many nodes. Every processor within a given TDP will behave with the same predictable performance under performance determinism. However, in this mode, with the exact same performance per node, the variances in the CPUs may lead to slightly different power consumption per CPU. In embedded applications where speed might need to be locked down, performance deterministic mode will help drive more a tightly controlled user experience.

Performance determinism will ensure that all applications react the same, since each processor will be set to deliver a consistent performance profile—regardless of the power required. For instance some parts may draw 175W to reach an effective maximum frequency, while others may require only 170W—but both parts will have the exact same performance profile under performance determinism.

Businesses should understand the differences in their workloads and should be able to set the profile based on the workload at a per-server basis, putting control back into the hands of the customer with this new feature that was not available in the past.

DETERMINISM MODES CAN DRIVE DATACENTER FLEXIBILITY

By default, all AMD EPYC processors are “boost enabled”, and there are already established power and performance states for every CPU. The power determinism mode may safely add an additional bump on top of that rated clock speed.

Each EPYC SoC (system on a chip) includes four CPU dies, all of which are connected and can communicate with each other to optimize for stability and product longevity. While the frequencies are fixed at the part number level (each part has an associated maximum clock speed), each of the dies could fluctuate under power determinism to drive additional performance upside.

The base frequency for the CPU will always be the marked speed. Customers will never receive less than the performance rated on the CPU unless they were to intentionally edit the BIOS to turn off particular P-states, as some have done in the past to maximize a power budget at the rack level.

While the small incremental increases may not seem like a lot, when applied across a large population in an HPC cluster, that performance difference could shave hours or possibly days of processing off some larger jobs. However, with a single database server, cutting latency by a small percentage may not register large incremental savings, meaning that IT should model both modes to determine the best performance and power profiles for each workload with respect to their overall operations.

Where these two determinism modes can truly benefit a business is in the flexibility that they can deliver. With businesses standardizing their purchasing decisions, this flexibility enables them to standardize on a smaller set of CPU options and still have a decent range of performance. This helps hold down both acquisition and management costs over time by tightening up the options and variances.

MI&S PERSPECTIVE ON POWER DETERMINISM

Businesses are looking for more control, so any feature that increases control should be seen as a positive step forward. Each feature becomes one more tool that IT can use to help drive better business outcomes based on aligning infrastructure to workloads.

Like density and power savings, most of the impact will be felt at scale. Large HPC clusters and cloud deployments are areas where even a small performance increase can help drive better business outcomes. Even if the headroom is small, when that is multiplied across nodes at scale, value quickly increases. A few percentage points of performance may not seem like a lot, but in a cluster with hundreds of nodes, this may mean the savings of several systems, multiplied by not only the system cost and the operational cost (maintenance, power, and cooling) but also the rack space which can equate to possibly the savings of an entire rack, depending on the deployment.

This is clearly the beginning of the next layer of flexibility at the CPU level. As technology continues to mature and IT gains more control, we believe features like these will help drive the agility that businesses seek, especially in an age of digital transformation. Between software definition and application orchestration, businesses are getting a higher level of optimization for their environments, and this feature brings some of that capability down to the hardware level.

By including this as a universal feature and not branding it as a premium SKU, AMD is taking a very customer-centric approach that greatly assists IT's ability to standardize.

CALL TO ACTION

Power and performance are important variables in any datacenter plan, and it is important for businesses to maximize these variables as they chart out their datacenter strategies. The ability for IT to better manage these variabilities brings a business more control in a world full of unknowns. Infrastructure control is a primary underpinning of driving better operating expense, not just a capital variable. Businesses would be wise to maximize their control wherever possible.

MI&S recognizes the power and performance determinism capabilities of AMD EPYC CPUs and believes that businesses interested in better managing their environment should consider platforms based on these processors.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR

John Fruehe, Senior Analyst at [Moor Insights & Strategy](#)

PUBLISHER

Patrick Moorhead, Founder, President, & Principal Analyst at [Moor Insights & Strategy](#)

EDITOR & DESIGN

Scott McCutcheon, Director of Research at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

This paper was commissioned by Advanced Micro Devices. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

© 2017 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.