AMD EPYC Presents Opportunity To Save On Software Licensing Costs

Business Selection Of Processor Should Factor In Software Costs

Executive Summary

Software licensing models for many server applications have long been tied to the hardware infrastructure resources on which they are run, particularly server processor resources. As server processor architectures have evolved, software licensing models have evolved with them.

The software licensing model evolution has varied widely by software vendor, both in their types of changes and in each vendor's timing of changes after the introduction of a significant change in processor architecture. This diversity in license models, and the timing lag of their evolution, presents businesses with an opportunity to save on licensing costs when a new processor architecture substantially changes the price/performance relationship between hardware and software. In these circumstances, buyers should assess the combined costs of hardware and software when making hardware infrastructure purchases.

The new AMD EPYC processor includes significant architectural change from past x86 processor generations in a way that presents exactly this type of opportunity for cost savings on software licensing. AMD offers a set of EPYC processor options for optimizing hardware and software licensing cost and performance across the number of processor sockets and cores per processor utilized per server. EPYC processors present compelling price/performance value for many applications on a pure hardware basis. Factoring in software cost savings can increase that value and make EPYC the higher value choice for more applications.

Cost Savings Opportunity

To understand why the opportunity for saving on software licensing costs occurs as processor architectures change, it helps to understand how software licensing models are determined.

The aim of licensing models in applying typical pricing principles is to charge customers at a level corresponding to the value they derive from the software. For some server software, there are strong proxies for value based on the scale of use in the number of users. For example, an application with a set of users in an organization can be licensed according to the maximum number of connected users or the number of users with access to the application as is common in virtual desktop, mail, and file server application licensing. That number of client users is straightforward to track and audit.

For many types of server software, however, usage metrics are not as easily tracked and audited. For example, a database server would be difficult to license and audit based on a dynamically scaling connection count or database storage size. Therefore, the licensing model for the majority of server software is mapped to the hardware resources it is run on as a proxy for the scale, and therefore value, of use.

The server processor is the hardware resource most frequently chosen as the best proxy for the infrastructure resources applied overall across compute, network, and storage. For many applications, the number of processors scales (e.g. from 1 to 4) in strong correlation to the level of application workload served, and this can be easily tracked and audited.

In the simplest situation for software licensed on processor count, new processor generations historically would offer higher performance through improvements in speed (i.e. clock speed) and efficiency (i.e. instructions processed per clock cycle). That simple relationship lasted until the introduction of dual-core processors, particularly x86 dual-core processors in 2007. Over the decade that followed, as two cores became four cores and beyond, customers could serve higher scale application needs from fewer servers and processors. This motivated many software vendors to switch to a per-core licensing model. Some vendors, for example Oracle, made the switch quickly. Others, for example Microsoft, moved in steps over a longer period with Windows Server licensing not switched until 2016.

In a minority of cases, application vendors leave the benefit of higher compute density with multi-core processors to their customers with licensing remaining at the per processor level. VMware is a top example of this model with vSphere still licensed per processor.

This varied software licensing approach across vendors in the context of evolving processor architectures presents two types of cost saving opportunities:

- For applications licensed per core that are performance-bound by memory and/or I/O resources instead of processing unit resources, reducing processor core

count offers savings in processor and software costs – if it can be accomplished without impacting memory and/or I/O resources.
- For applications licensed per processor and performance-bound by processing unit resources, a doubling of cores per processor can enable consolidation from four processors to two, or two processors to one, for savings in processor and software costs.

While there are processor and software savings in both of those scenarios, common enterprise software costs can easily dwarf the processor cost being considered. Here are three per-core savings examples to consider at different levels of the software stack, when the necessary memory and I/O performance can be delivered using 8 fewer cores per processor (e.g. 8 cores vs 16 cores, 16 cores vs 24 cores, or 24 cores vs 32 cores):

- **Microsoft** for operating system software. Windows Server 2016 DataCenter Edition Open NL list pricing for the server-side cost is currently just over $6,000 per 16 processor cores. If able to run with 8 fewer cores per processor, that is a potential **savings of $3,000 per processor**.
- **Oracle** for application server software. Current Oracle WebLogic list pricing ranges by edition from $10,000 (Standard) to $45,000 (Suite) per pair of x86 processor cores. If able to run with 8 fewer cores, that is a potential **savings of $40,000 to $180,000 per processor**.
- **MicroStrategy** for business intelligence & analytics application software. MicroStrategy's server application current list pricing ranges by edition from $300,000 per core (Web/Mobile) to $600,000 per core (Server). If able to run with 8 fewer cores, that is a **savings opportunity of $2,400,000 to $4,800,000 per processor**.

These per-core license model savings can be compared to processor cost savings of around $1,000 to $2,000 per processor when dropping by 8 cores. *The software licensing cost reduction on common types of licensed server software across the software stack from operating system to application can be much larger than the hardware cost savings opportunity and should be a major consideration in overall price/performance in processor selection.*

Likewise, this is also applicable in per processor licensing models where **VMware** provides a commonly used example to consider with vSphere Enterprise list pricing per processor currently starting just above $4,000. Being able to decrease processor-bound workloads running virtualized on VMware from two processors at 16 cores each to one processor at 32 cores would cost $2,000 to $3,000 more per processor, but save

$4,000 in licensing costs across a fleet of virtualized server hosts — commonly around 10-50 times the scale of high performance, single workload back-end database, analytics, etc. servers. The software cost savings overrides the hardware cost increase by $1,000-$2,000 for a potential **savings of $10,000 to $100,000 per set of 10-50 servers**.

TABLE 1: EXAMPLES OF SOFTWARE COST SAVINGS OPPORTUNITY

Software	Licensing Model	Processor Architecture Advantage for Minimizing Licensing Cost	Potential Cost Savings
VMware vSphere Enterprise	Per processor	High core count per processor allows consolidating from 2 processors to 1 processor	$1,000 to $2,000 per server
Microsoft Windows Server 2016 DataCenter Edition	Per core	Decoupling of memory & I/O performance from number of processor cores allows reducing processor core count	$3,000 per processor
Oracle WebLogic	Per core	Decoupling of memory & I/O performance from number of processor cores allows reducing processor core count	$40,000 to $180,000 per processor
Microstrategy Web/Mobile & Server editions	Per core	Decoupling of memory & I/O performance from number of processor cores allows reducing processor core count	$2,400,000 to $4,800,000 per processor

Source: Moor Insights & Strategy

INTRODUCING THE AMD EPYC PROCESSOR

The new AMD EPYC processor architecture introduces advances in processor and server platform architecture that provide unique flexibility in available combinations of processor cores versus memory and I/O bandwidth to address varied application workload requirements. Architectural advances start with significant per core Instructions Per Cycle (IPC) improvement, including use of a 2MB L3 cache per core.

The cores are structured in 8-core building blocks, allowing high scale processor performance options from 8 to 32 cores.

What is most new and unique to x86 processor architecture is EPYC's de-coupling of memory and I/O bandwidth from processor core count. Each processor provides 8 DDR4 channels for up to 170GB/s memory bandwidth with 2TB of memory and 128 lanes of PCI Express. In a 2-socket server this allows for scaling to 64 cores and 340GB/s memory bandwidth with 4TB of memory.

All of that memory and I/O bandwidth is available whether using an 8-core, 16-core, 24-core, or 32-core EPYC processor. This enables memory and I/O performance-sensitive applications to benefit without overpaying for more cores, and therefore more per-core software licensing cost.

Examples of application workloads where this is most advantageous are:

- **High-density virtualization** running horizontally scaled application tier virtual machines (VMs) of general-purpose web and enterprise applications, which typically consume memory and I/O highly disproportionate to the amount of processor resources. VMware vSphere virtualization serving these types of applications is an example.
- **I/O-intensive workloads like application serving, analytics, databases, and HPC** where memory and storage I/O performance are often critical, while processor utilization varies widely by application. Oracle WebLogic and MicroStrategy Analytics Server are good examples of application serving and analytics applications. Database and HPC workloads can be memory and I/O-bound as well, particularly NoSQL and in-memory databases like Oracle's NoSQL Enterprise and TimesTen, which are currently priced in a similar range to Oracle WebLogic at $10,000 and $23,000 per pair of x86 cores, respectively.

These types of workloads that benefit most in performance from the EPYC architecture provide the best targets for potential software savings through reduced core count or processor socket count.

CALL TO ACTION

Software licensing cost per hardware resource presents a significant opportunity for cost savings when choosing hardware infrastructure, particularly when a new processor architecture introduces new options in resources purchased across processor core count, memory bandwidth, and I/O bandwidth. AMD EPYC's unique de-coupling of processor cores from memory and I/O bandwidth introduces this type of opportunity.

Moor Insights & Strategy recommends that businesses consider total hardware and software cost when making processor purchase decisions, as software licensing cost can weigh heavily in price/performance consideration. In addition, businesses should consider reviewing their application portfolio for applications licensed per core or per socket that are memory and/or I/O-bound. These provide opportunities to apply EPYC-based servers to achieve software cost savings in future license renewals greater than the cost of new hardware.

MOOR INSIGHTS & STRATEGY

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR

Rhett Dillingham, Senior Analyst In-Residence 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

DISCLOSURES

This paper was commissioned by Advanced Micro Devices, Inc. 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.