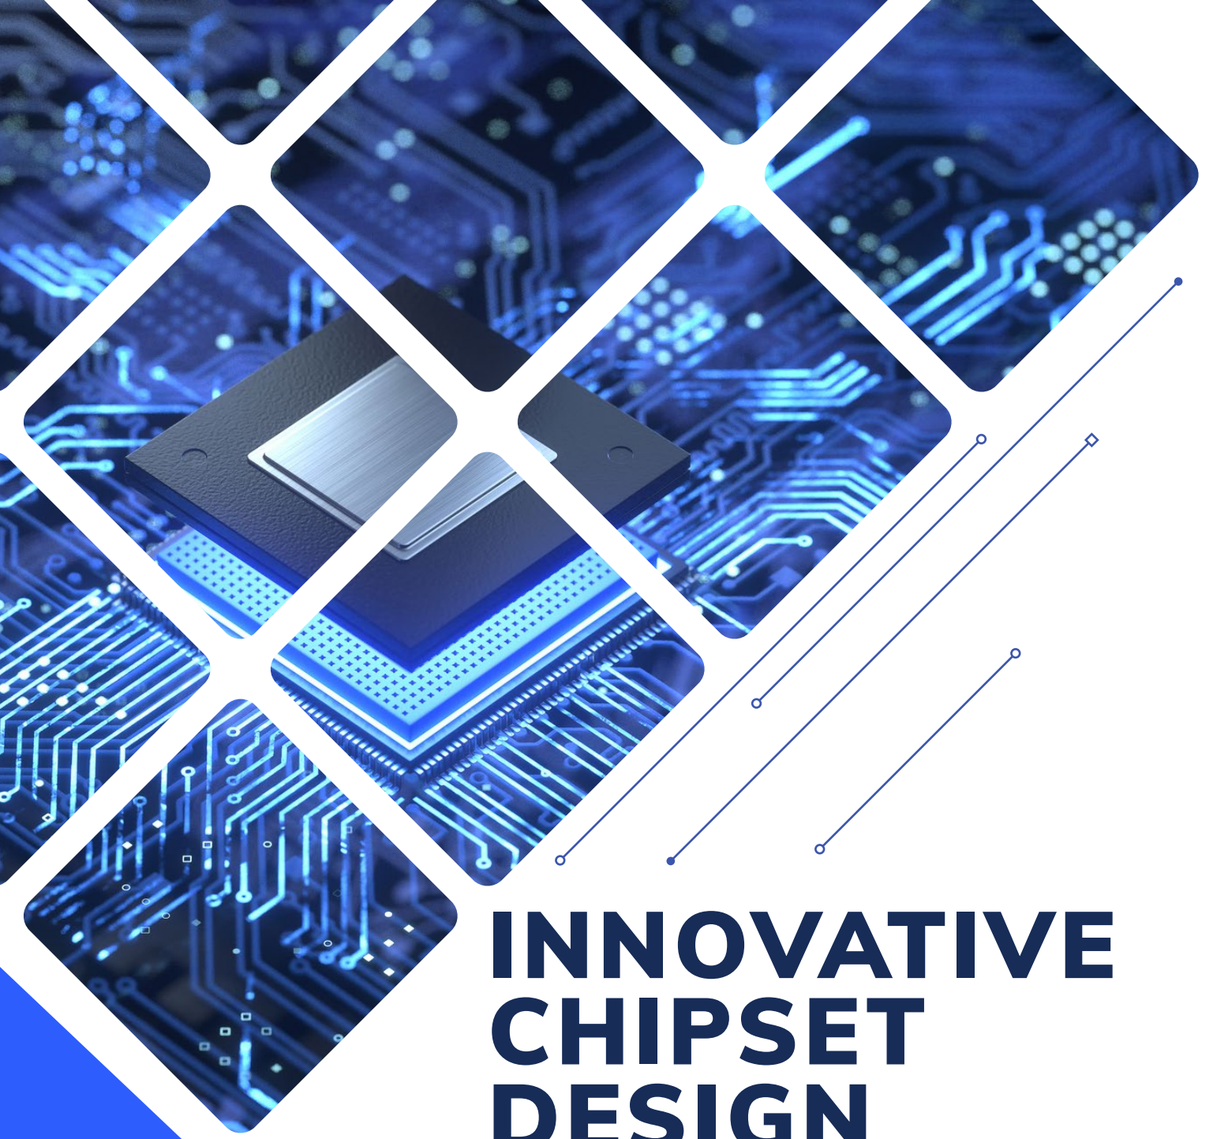

INNOVATIVE CHIPSET DESIGN

offers game-changing performance and security

One server can do the job of five in typical federal data centers with new line of advanced AMD microprocessors.

FedScoop Report

The decade-long effort by federal agencies to consolidate and optimize thousands of government-run data centers appears poised to take a promising turn, thanks to recent advances in microprocessor architecture and design.

While the speed and flexibility of the cloud have lured a growing share of government IT workloads away from federally run data centers in recent years, the fact remains the top 24 Cabinet-level agencies were still operating 2,441 purpose-built data centers at the end of fiscal 2019, according to a recent Government Accountability Office report.

That doesn't include more than 2,000 smaller-scale facilities operated governmentwide that the Office of Management and Budget decided in June of 2019 to stop tracking.

Those data centers continue to support essential services and sensitive data that, in most cases, make more sense to run in-house than in a leased cloud environment. Collectively, however, the cost to maintain all those data centers and servers is only growing more expensive.

A typical government data center accounts for 100 to 200 times the energy intensity of a commercial building, according to figures cited in the GAO report. And while server utilization rates have improved — from as low as 5% a decade ago across the federal government to within range of federal goals of 65% or more — a vast number of servers within those data centers today are built on outdated technology, says Greg Gibby, senior product manager for data center products at Advanced Micro Devices (AMD).

Changes at the core

It's easy to underestimate just how much microprocessors have advanced in the past few years, and why that has important operating and security implications for future investments, Gibby says.

"Most enterprise servers are five-plus years old. And when you look at when they were built, you're looking at core counts of 8, 10, or 12 [processor cores per CPU]. As you look at the performance levels available today, it just dwarfs them," says Gibby. "I can replace five servers with one, get the benefit of moving to a much more consolidated infrastructure — along with reduced maintenance, licensing and operating costs — and I get better performance to boot."

At the heart of those performance advances is an innovative, multi-generation line of microprocessors, designed and produced by AMD, that provide up to 64 cores of computing power in a single-socket processor. The architecture that facilitates single-socket designs is part of a larger processor development roadmap at AMD that delivers greater flexibility and lower total operating costs now, according to Gibby, and will continue to do so over the long run.

Multi-socket servers and large symmetric multiprocessor systems are becoming dinosaurs in the digital age, argues Robert Hormuth, technology fellow and vice president of Dell Technologies' Infrastructure Solutions

Group. They were designed when a variety of factors prevented the notion of large-scale systems and applications that are common today, he says.

As core counts grew, so did the demand for memory and power, creating data center hot spots that often forced IT departments to leave racks half empty. Single-socket servers, on the other hand, eliminate the bottlenecks that occur when multiple processor cores share memory, input-output bandwidth and other resources between sockets. That newer, more innovative design from AMD reduces power density, improves performance and lowers total operating costs.

Flexible deployments

The underlying chiplet-based Zen architecture of AMD's current and future generations of chipsets is not only delivering greater performance for the money than the competition — "We have 170-plus world records and counting," says Gibby — the design is also giving customers more flexibility in how they can utilize newer servers. Chiplets are independent components which make up a large chip using multiple smaller dies.

> ***EVERY MAJOR CLOUD PROVIDER HAS INVESTED IN AMD AND IS CONTINUING TO ROLL OUT AMD PLATFORMS AS PART OF THEIR PORTFOLIO.***
>
> **Greg Gibby, AMD**

That flexibility and lower operating costs are among the reasons AMD's EPYC family of microprocessors has become the microprocessor of choice to power AWS Elastic Compute Cloud and Microsoft Azure.

"Every major cloud provider has invested in AMD and is continuing to roll out AMD platforms as part of their portfolio," says Gibby.

Google also recently announced the availability of AMD powered N2D machine types as well as Google Cloud Confidential Computing with Confidential VM's and Confidential GKE Nodes, based on AMD EPYC powered servers. Additionally, VMware just announced support for SEV-ES in their vSphere 7.0 update 1 release.

Federal agencies now have the ability to provide their internal and external-facing customers a range of more-flexible, higher-performance computing options, comparable to what AWS offers, such as:

- **Optimized compute** — ideal for compute-intensive workloads that require high performance processors, including scientific modeling, batch processing, distributed analytics, machine learning, log analytics and web applications.
- **General purpose compute** — built for business and enterprise applications and application development environments where a balance of compute, memory and networking resources are required.
- **Large L3 cache, memory capacity** — designed for workloads that process large data sets in memory and require fast performance, such as real-time big data analysis workload, distributed web scale in-memory caches and other enterprise applications.

Federal agencies looking to refresh and optimize their server fleets can now capitalize on the same long-term operating advantages AMD's processor are delivering to the big cloud providers, says Rick Indyke, who manages AMD's public sector business development.

"By improving the density, you begin to deliver greater virtual machine counts and reduce the total cost of ownership," he says,

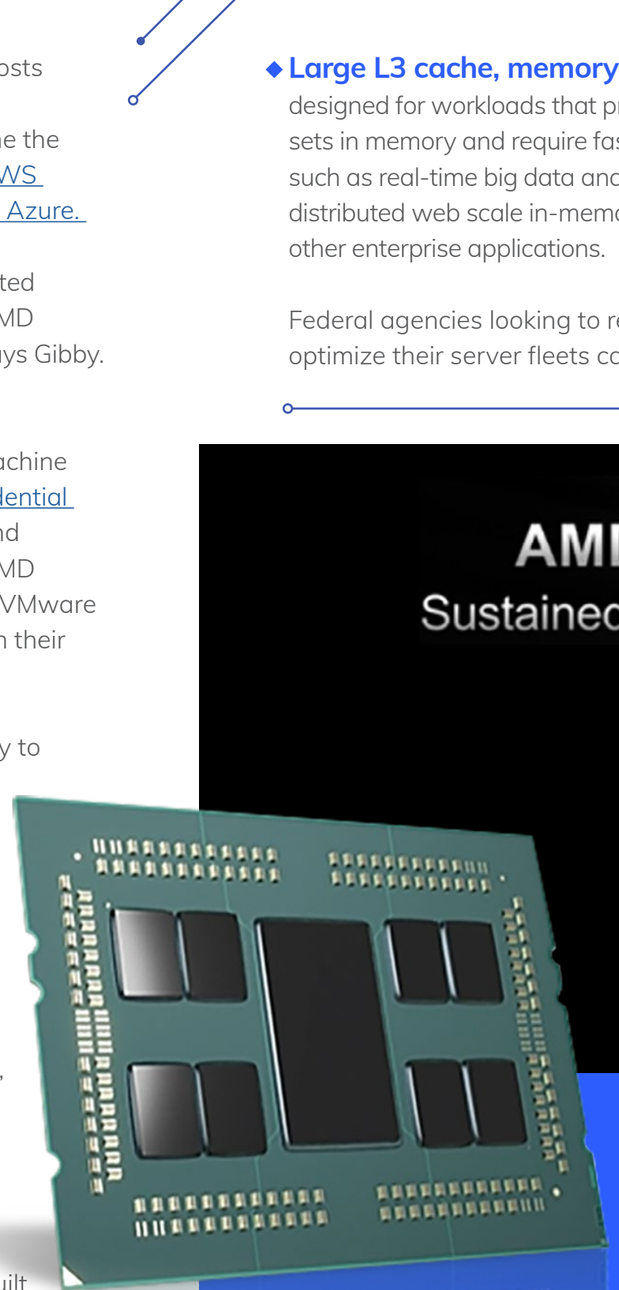

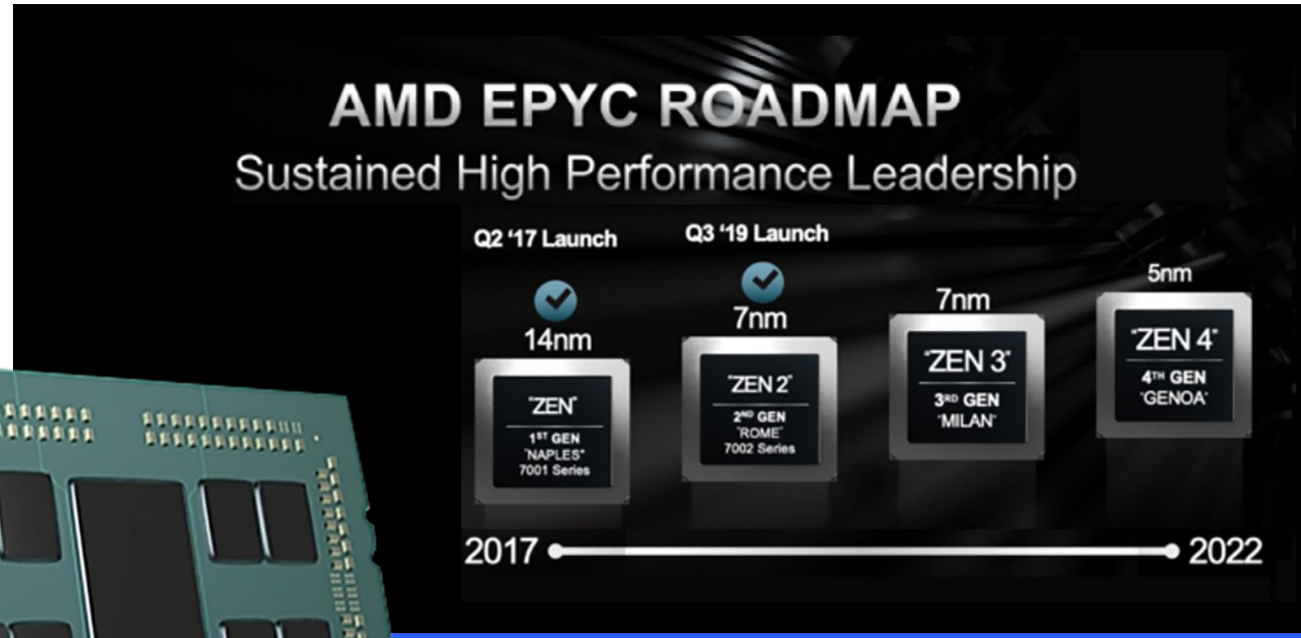

AMD's innovative, high-performance features — including the ability to deliver up 64 cores in a single socket processor — and AMD's EPYC design "Roadmap" assure users of long-term flexibility and security, and lower total operating costs.

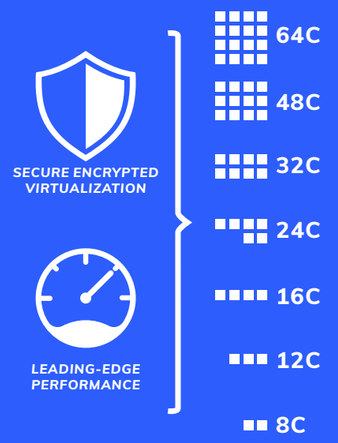

> ***BY IMPROVING DENSITY, YOU BEGIN TO DELIVER GREATER VIRTUAL MACHINE COUNTS AND REDUCE TOTAL COST OF OWNERSHIP.***
>
> ***Rick Indyke, AMD***

"and reduce the support and maintenance services required, because you're taking care of fewer machines."

Built-in security

Indyke also says that AMD's current line of EYPC "Rome" chips, along with its rollout of EPYC "Milan" 7 nanometer chips later this year and EPYC "Genoa" 5 nanometer chips that will follow "Milan" — offer federal data centers another advantage: Long-term security built into every chipset that "gives agencies a platform to help future-proof their servers."

AMD's Secure Encrypted Virtualization (SEV) gives users "the ability to isolate your virtual machines from each other and from the hypervisor itself, so it helps protect against both internal and external threats," explains Gibby. "And it encrypts the full memory space within that virtual machine to provide an excellent solution for data protection," using built-in security technologies such as Secure Root of Trust, Secure Memory Encryption and Secure Encrypted Virtualization. Those technologies are managed by a dedicated security processor called the AMD Secure Processor.

"AMD has designed protection checks for speculative execution requests. If certain protection checks fail, [the AMD chips] don't deliver data or any of the downstream requests for that data," Gibby says.

SEV — currently available for machines running on Linux and VMware — do not require any changes to applications, making it easy to deploy.

"With most servers being deployed for five-plus years — if you think about buying a new architecture that does not have this capability, that is a long time not to reap the benefits of AMD's security features, so think about enabling your hardware ecosystem now," recommends Gibby.

AMD's partnership with Dell Technologies, along with Dell's proven supply chain experience in delivering certified servers and engineering support to federal agencies, provides even greater assurance to agencies, Indyke says.

That partnership has led to deep engineering engagements over the past few years and resulted in Dell bringing a full portfolio of AMD-based PowerEdge server platforms — including several single socket servers — to the federal market.

Dell has also brought federal customers a growing array of AMD-based solutions and appliances in recent years, such as vSAN ReadyNodes and VxRail E665, that provide agencies with industry-leading performance and that are optimized right out of the box.

"If you invest in AMD, this is something that not just pays off today. It pays off in the future as well," says Gibby.

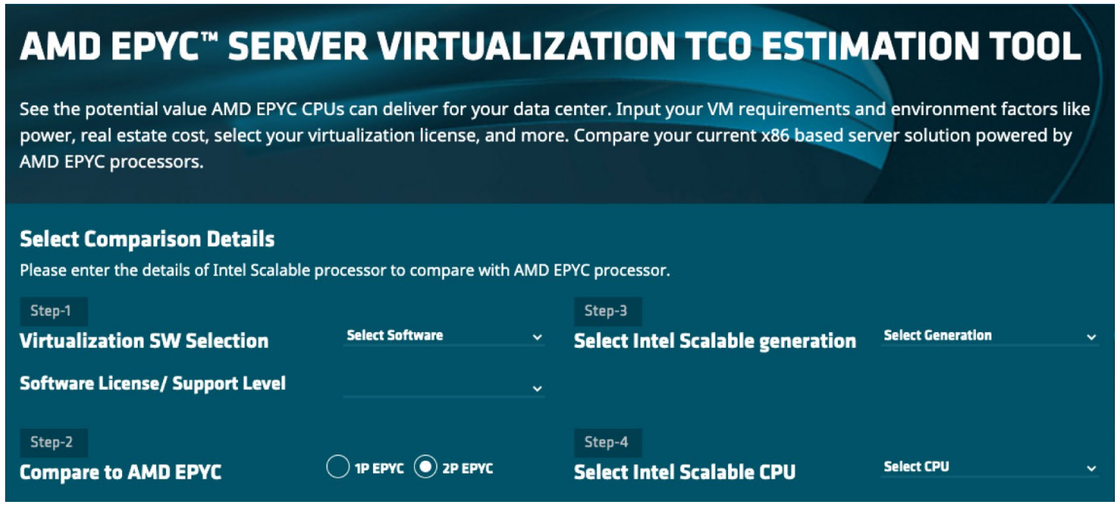

By delivering superior performance, security and scalability, AMD's EPYC CPUs — and Dell Technologies' latest servers — can lower your total cost of operations. Check out AMD's "Virtualization TCO Estimation Tool.

Learn more about AMD's microprocessor innovations and Dell Technologies' latest line up of high-performance servers.

fedscoop DELL Technologies AMD