AMD DATACENTER SOLUTIONS

AMD and MongoDB Power the Intelligent Operational Data Platform

October 2018

High Performance: Cores, Memory, I/O

The AMD EPYC SoC brings a new balance to the datacenter. Utilizing an x86-architecture, the AMD EPYC processor, brings together high core counts, large memory capacity, ample memory bandwidth and massive I/O with the right ratios to help performance reach new heights.

Flexibly match core count to the workload

Match core count with application needs without compromising processor features. EPYC's balanced set of resources means more freedom to right-size the server configuration to the workload.

Dedicated Security Processor

AMD EPYC's dedicated security processor manages secure boot, memory encryption, and secure virtualization on the SoC itself. Encryption keys never leave the processor where they can be exposed to intruders.

Highly Scalable: Scale-up or Scale-out

Scale-up or scale-out, AMD and its ecosystem partners offer high-performance network connectivity options for applications at massive scale.

The Natural Way to Work with Data

MongoDB works with data in a natural intuitive way while providing ACID guarantees to ensure data integrity.

Intelligently Put Data Where You Need It

MongoDB delivers a globally resilient platform through sophisticated replication and failover that scales horizontally and allows operational and analytical workloads to run in the same cluster.

Freedom to Run Anywhere

MongoDB runs the same everywhere: on-premise, cloud, multi-cloud and hybrid.

AMD EPYC: Certified with MongoDB

The MongoDB Partner Program is designed to help reduce risk and implementation cost through joint testing and validation ensuring the optimal functioning of MongoDB products.

AMD EPYC™ Processor Meets the Needs of Today's Datacenter Workloads

Designed from the ground up for a new generation of solutions, AMD EPYC implements a philosophy of choice without restriction. Choose the number of cores and sockets that meet your needs without sacrificing key features like memory and I/O.

Each EPYC system-on-chip (SoC) can have from 8 to 32 cores with access to incredible amounts of I/O and memory regardless of the number of cores in use, including 128 PCIe lanes, and access to 2 TB of high-speed memory per socket.

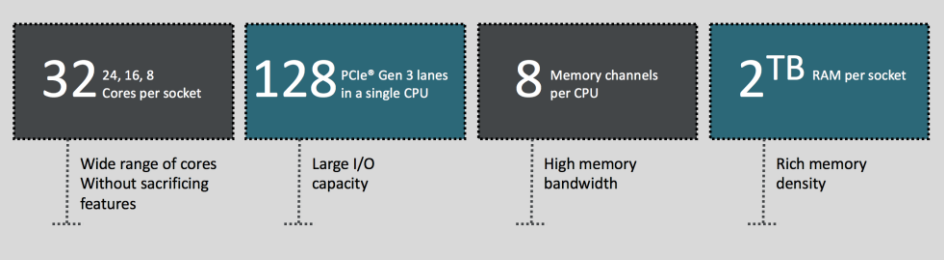

EPYC's innovative architecture translates to terrific performance at a low cost. More importantly, the performance you're paying for is appropriate to the performance you need.

Storage intensive workloads can utilize the plentiful I/O bandwidth with the right number of cores (avoiding overpaying for unneeded power), while compute-intensive workloads can make use of fully loaded core counts, dual sockets and plenty of memory.

MongoDB: The Intelligent Operational Data Platform

MongoDB is built around an intelligent distributed systems architecture that enables developers to place data where their apps and users need it. MongoDB can be run within and across geographically distributed data centers and cloud regions, providing levels of availability, workload isolation, scalability and data locality unmatched by relational databases.

MongoDB is designed to meet the demands of modern applications with a technology foundation that implements a versatile method to work with data, allows you to intelligently put the data where it's needed, and provides the freedom to run anywhere.

Better, Faster, Leaner

Organizations spend enormous amounts on databases and the resources required to maintain them. Technology does not stand still – the infrastructure, applications and even the data have evolved while the underlying data management tools, based on relational database management systems, have remained the same.

Traditional relational database management systems have struggled to adapt to the rapid pace of change:

Data volumes have grown. Ever chart showing data growth looks the same: up and to the right at a very steep angle.

Storage and compute costs have declined. The rise of commodity servers and storage has driven down infrastructure costs, and new data management software has been developed that takes advantage of commodity infrastructure.

The EPYC Advantage: Performance - The AMD EPYC SoC brings new balance to the datacenter. The highest core count yet in an AMD x86-architecture server processor, large memory capacity, bandwidth and I/O density are all brought together with the right ratios to help performance reach new heights.

Cloud is mainstream. Traditional on-premise architectures are being supplemented, and in some cases supplanted, by private and public cloud infrastructures that provide significant benefits in cost-effectiveness, elasticity and time-to-market.

People, and things, are always online. Ubiquitous mobile devices and internet connectivity have made continuous online presence and activity the norm. The number of simultaneous users has exploded and users want access to increasing volumes and types of information, and they expect seamless, high-quality experiences across all devices, all the time.

New types of applications are now a requirement. Traditional business applications are no longer sufficient. Mobile, social and real-time analytical applications are now a requirement to stay relevant.

Data types continue to change. Applications now require a wide variety of data: rich text, user comments, images, geo-spatial data, video and social media. All are used to create a seamless user experience.

AMD has partnered with leading software and server vendors to create certified reference architectures for customers to demonstrate strong performance, lower risk and to help reduce implementation cost.

Why MongoDB?

MongoDB is designed to meet the demands of modern applications with a technology foundation that implements a versatile method to work with data, allows you to intelligently put the data where it's needed, and provides the freedom to run anywhere.

Data Versatility

Traditional relational databases struggle in key areas that are now commonplace with modern systems. They require that the data, data types and the relationships between the data be clearly understood before any data is loaded or applications developed. Thus, these systems are rigid, and difficult to enhance or change as business requirements evolve.

Less obvious is the disconnect between how the data is stored and how it used in applications. Relational databases use a tabular data model, storing data across many tables – hundreds to thousands of tables is not uncommon. Applications, on the other hand, perform their function by creating objects, including their associated data,

Data Versatility (continued)

that resemble real-world entities like customers, orders or products.

The impact is two-fold: application development is slower as developers have to translate between the data model and the objects that use data; and performance is slower as many tables have to be accessed to build the objects the application needs to perform its function.

MongoDB solves these problems by using a document data model. Documents are a much more natural way to describe data. They present a single data structure with related data embedded as sub-documents and arrays. This means how the data is stored is closely aligned to how it is used making application development simpler and faster.

MongoDB is a distributed system that handles the complexity of maintaining multiple copies of the data via replication. Data consistency is tunable (down to the level of individual queries) based on the application's requirements.

MongoDB also supports multi-document ACID transactions providing a consistent view of data, and enforcing all-or-nothing execution.

Intelligent Data Placement

MongoDB is built around an intelligent distributed systems architecture that enables developers to place data where their apps and users need it. MongoDB can be run within and across geographically distributed data centers and cloud regions providing high availability, workload isolation, scalability and data locality.

Self-healing replica sets are used to maintain multiple copies of the data. Failover and recovery are fully automated.

These replicas sets serve as the foundation for workload isolation enabling a single cluster to handle many different workloads each operating on its own copy of the data. In this way, operational applications can be isolated from analytics, ad hoc queries, machine learning modeling, etc.

MongoDB provides horizontal scale-out on commodity hardware or cloud infrastructure using a technique called sharding. Sharding automatically partitions and distributes data across multiple physical instances called shards in a completely transparent manner.

The EPYC Advantage: Outstanding performance and scalability – both scale-up and scale-out. AMD and its ecosystem partners offer high-performance network connectivity options for massive scale-out systems.

Freedom to Run Anywhere

MongoDB can be deployed anywhere: on-premise, private cloud, public cloud or any combination. Application development is completely unaffected by the deployment model. Installation, maintenance and optimization all use a unified set of tools.

Using MongoDB, organizations can adopt cloud at their own pace by moving select workloads as needed. For example, they may run the same workload in a hybrid environment to manage sudden peaks in demand, or use the cloud for services in regions where they don't have a physical datacenter.

Security

Securing customer data is a must-have from both a compliance and basic data access perspective. MongoDB features extensive capabilities to defend, detect and control access to data including Kerberos, x.509, LDAP and Active Directory for authentication. Role Based Access Control is configurable at a granular level, and MongoDB's native audit log can be used for regulatory compliance and auditing. Lastly, MongoDB can be encrypted on disk, on the wire and in backups.

The EPYC Advantage: Security - AMD EPYC features the industry's first dedicated security processor embedded in an x86-architecture server SoC. The processor manages secure boot, memory encryption, and secure virtualization on the SoC itself. Encryption keys never leave the system where they can be exposed to intruders.

Reference Architecture for AMD EPYC and MongoDB

The reference architectures for AMD EPYC processors and MongoDB provide options for the performance and scalability requirements needed to maximize an organization's investment.

AMD EPYC's ability to provide a no compromise single-socket solution ensures you are only paying for the processing power the application needs. A single-socket server has all of the I/O and memory bandwidth available in a dual-socket server without the extra cost

For high-performance applications, AMD EPYC-based dual-socket servers offer phenomenal core density.

The EPYC Advantage: Flexibility - Match core count with application needs without compromising processor features. EPYC's balanced set of resources means more freedom to right-size the server configuration to the workload.

These reference architectures offer the foundation for creating custom configurations that meet unique application demands with the various storage and connectivity options to choose from our leading server OEM partners.

No Compromise Single-Socket		Versatile Dual-Socket	
MongoDB Nodes: scalable to hundreds		**MongoDB Nodes**: scalable to hundreds	
CPUs	1 x AMD EPYC 7601, 7551, or 7501*	CPUs	2 x AMD EPYC 7451 or 7401*
Cores/Threads	32/64	Cores/Threads	24/48 per CPU, 48/96 total
Memory	128/256 GB (supports up to 2 TB)	Memory	128/256 GB (supports up to 4 TB)
Network	10/25/40/100 Gigabit Ethernet	Network	10/25/40/100 Gigabit Ethernet
Storage: OS	2 x SSDs or 2 x HDDs	Storage: OS	2 x SSDs or 2 x HDDs
Storage: Data High Performance Performance Capacity	 4-8 x NVMe drives or 8-12 x SAS/SATA SSDs 24 x SAS/SATA HDDs (SFF) 12-16 x SAS/SATA HDDs (LFF)	Storage: Data High Performance Performance Capacity	 4-8xNVMe drives or 8-12 SAS/SATA SSDs 24 x SAS/SATA HDDs (SFF) 12-16 x SAS/SATA HDDs (LFF)

*other AMD EPYC 7000 series processors are also supported

Conclusion

Versatility and agility are among the most important requirements in modern datacenters. The AMD EPYC system-on-a-chip (SoC) enables organizations to deploy systems that precisely meet today's needs while positioning themselves for tomorrow's requirements.

MongoDB is designed to meet the demands of modern applications with a technology foundation that implements a versatile method to work with data, allows you to intelligently put the data where it's needed, and provides the freedom to run anywhere. It delivers a globally resilient platform through sophisticated replication and failover that scales horizontally and allows operational and analytical workloads to run in the same cluster.

Together, AMD and MongoDB empower the development of modern data applications that implement solutions to a diverse set of business problems with high-performance processing and cost-effective solutions that are perfectly sized for current needs and easily scalable as your business grows.

For more information about AMD's EPYC line of processors visit: http://www.amd.com/epyc

For more information about MongoDB visit: http://www.mongodb.com/

For more information about the MongoDB Partner Program visit: https://www.mongodb.com/partners/partner-program