ScaleMP
CREATING THE POWER OF ONE

Scale-up With AMD-EPYC to Achieve Breakthrough Performance Using vSMP ServerONE

The Need for Larger and Faster Systems

Large enterprises, governments and leading research institutions demand faster and larger systems for processing and analyzing very large data sets and for computational-intensive applications. Some of those applications can utilize a collection of computers (cluster) to handle a specific workload, while others require a shared-memory system (SMP) to execute because the program needs access to very large memory, many CPUs or both.
SMP systems have great advantages: They are easy to deploy, manage and program, and they are versatile They can run any application, including those demanding applications that needs hundreds of cores or more or dozens of terabytes of RAM. However, SMP systems are traditionally proprietary systems and are very expensive. SMPs sometimes use custom processors and operating systems, and they always require the development of custom chipsets and ASICs to create high-speed backplanes.

AMD EPYC SMPs With Virtualization for High-End Computing

With the introduction of **EPYC** server processors, AMD has bypassed the competition in server processor specification. **EPYC** offers 14 percent more cores per chip, 33 percent more memory channels and capacity, and 33 percent more PCIe lanes per dual-socket platform. The only thing missing is an SMP capability to outmaneuver the glue-less 8-socket Xeon and the limited-scale 32-socket SMPs with node controllers. For this, AMD has partnered with ScaleMP: ScaleMP's patented **Versatile SMP (vSMP) architecture** leverages server systems based on AMD's **EPYC** processors along with industry-standard interconnects to create a virtual SMP system with superior capabilities to proprietary legacy SMP systems, while maintaining an affordable cost structure. ScaleMP's **vSMP ServerONE** software aggregates up to 128 **EPYC**-based servers into a single virtual machine. This translates to 256 **EPYC** processors, up to 16,384 CPUs, and up to 512 TB of main memory.
By its software-defined nature, it also makes the following advantages available over legacy SMPs:

1. Composable infrastructure: Organizations can transform a (part of a) cluster into a shared-memory system on-demand and, exactly at the size that fits their workload requirements.
2. Future-proof: It supports and will continue to support the latest AMD processor generations.
3. RAS: It provides features rarely found elsewhere, such as redundant SMP backplane, fast detection and isolation of failed hardware components, and partitioning support.

As a result, SMP users now have a broader choice of vendors and can procure better products with a lower overall cost.

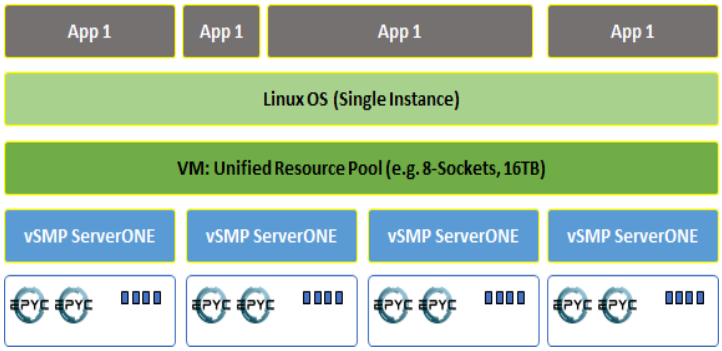

ScaleMP™
CREATING THE POWER OF ONE

IT Challenges Driving the Adoption of SMP Systems

Computing Power: Where Dual-Socket SMP Doesn't Have Enough Cores

AMD **EPYC** processors offer up to 32 cores and 64 computing threads – higher than the competition. Still, some workloads, especially multi-threaded applications in the high-performance and technical computing domain (science/research, computer aided engineering, computational chemistry, data analytics, physical simulations) may require hundreds or even thousands of cores in order to complete execution within a reasonable timeframe. This is especially true in computing domains where parallelization frameworks such as OpenMP and OpenCL gain traction.

EPYC-based SMP is an easy win for anyone looking into 8 sockets or 32 sockets offered by the competition: It actually offers these customers a proven path to grow to much larger single-system images with linear cost structure.

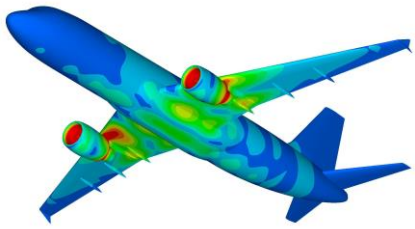

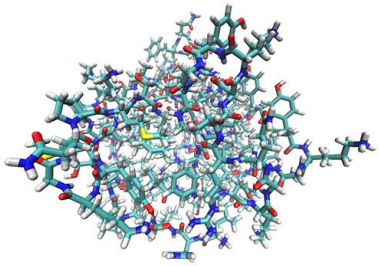

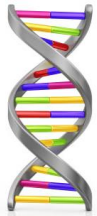

For users in the HPC domain, there is an additional bonus: If they source clusters based on **EPYC** and InfiniBand, they can deploy SMP services on the very same hardware, **on demand**. That composable-computing approach eliminates the need to provision proprietary SMPs or fat nodes to service specific workloads.

The SMP is composed in an automated manner by the cluster provisioning and management through the job-scheduling software. SMPs are created for a specific workload out of available cluster nodes and torn down once the SMP workload finishes execution. This greatly simplifies HPC IT environments, as all workloads can be served by the same computing infrastructure, reducing TCO and increasing utilization.

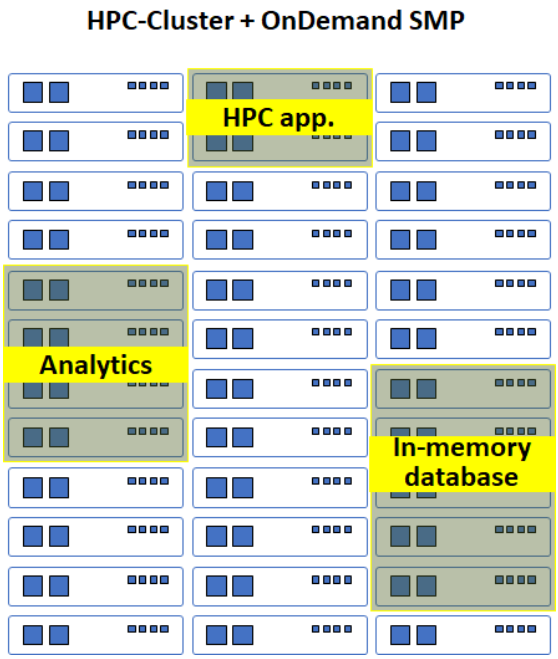

ScaleMP™
CREATING THE POWER OF ONE

Larger Memory for In-Memory Databases

AMD **EPYC** processors are a great building block for database machines. More memory and higher bandwidth for both IO and memory compared to competing solutions, **EPYC** provides plenty of the key requirements of the database engine.
In recent years, the industry has turned to in-memory databases in order to deliver the performance that end users demand, with very large datasets being managed. For some customers and use cases, the 4TB RAM limitation is simply not high enough. With the introduction of **EPYC**-based vSMP ServerONE systems, customers can scale their total system memory to the previously unimaginable 512 terabytes of DRAM.

With this, customers can rest assured that whatever their database scale, they can fit it into memory. Users of in-memory database implementations from Oracle, Microsoft, SAP and others can take advantage of **EPYC**-based SMP to reduce the TCO and expand well beyond competing solutions.

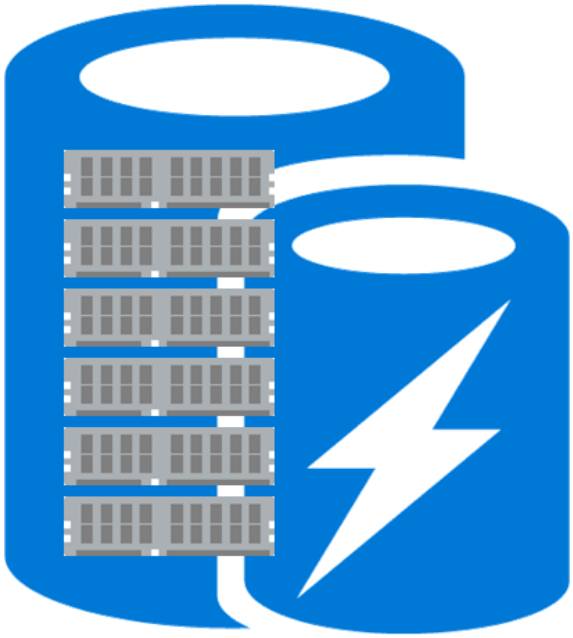

The Versatile SMP (vSMP) Architecture

Patented Versatile SMP (vSMP) architecture enables users to create high-end SMP systems. The vSMP architecture replaces the functionality of custom and proprietary backplane chipsets with software and industry-standard high-performance interconnects such as InfiniBand or PCI switches. It utilizes only a fraction of the system's CPUs and RAM to provide chipset-level services, without sacrificing system performance.
vSMP Foundation is ScaleMP's implementation of the vSMP architecture. vSMP ServerONE is a vSMP Foundation product that aggregates multiple **EPYC**-based server systems into one larger SMP system, allowing system vendors and value-added resellers to create high-end server solutions using industry-standard components that are superior in performance and price.

The vSMP Architecture: Software-Based SMP

The vSMP architecture uses industry-standard components: It does not require any custom parts. Its key value is the use of software to provide the functionality that is otherwise provided by a chipset found in traditional multi-processor systems.
vSMP Foundation provides cache coherency, shared I/O and the system interfaces (BIOS, ACPI) that the OS requires. The vSMP architecture is implemented completely transparently; users need no device drivers, no OS and no application modifications.

One System

Once loaded into the memory of each of the system boards, vSMP Foundation aggregates the compute, memory and I/O capabilities of each system and presents a unified virtual system to both the OS and the applications running above the OS.

Coherent Memory

vSMP Foundation maintains cache coherency between the individual boards using multiple advanced memory management and coherency algorithms, and leverages the system's local-memory and caching algorithms to minimize the effect of interconnect latencies.

Shared I/O

vSMP Foundation aggregates I/O resources across all boards into a unified PCI hierarchy and presents them as a common pool of I/O resources to the OS and the application. The OS can utilize all the system storage and networking controllers, providing high-I/O system capabilities.

Requirements for an EPYC-Based SMP Solution

The building blocks for the solution are simple, and system vendors provide them as turnkey:

- **EPYC**-based servers
- Fast private interconnect
- Software-defined stack

Users get to choose their favorite Linux OS, along with system components and size — storage, networking, etc. — just as they would with any other server.

AMD **EPYC**-Based Software-Defined SMP	
Supported server models	HPE DL385 Gen10 Dell PowerEdge R7425
Supported operating systems	Linux® 64-bit OS RHEL 7, CentOS 7, SLES 15
Max. sockets/cores/threads	256 / 8,192 / 16,384
Max. memory	512TB RAM in a single SMP

Summary

Customers are now able to benefit from the excellent performance and cost efficiency of AMD EPYC processors, while being able to scale up the compute and memory capacity within a single system image (a.k.a. shared-memory server).
The ability to aggregate up to 128 EPYC-based servers into a single virtual machine translates to 256 EPYC processors, up to 16,384 CPUs, and up to 512TB of main memory.