

White Paper:

The Versatile SMP™ (vSMP) Architecture and Solutions Based on vSMP Foundation™

Executive Summary

V100709

VSMP FOUNDATION EXECUTIVE SUMMARY

High performance and computational intensive applications have progressively demanded faster and larger systems to perform the work. In the late 80s, system-vendors started to offer multi-processor systems to accommodate such applications. Two types of multi-processor architectures evolved:

- **MPP – distributed memory systems.** MPP stands for Massively Parallel Processing systems, which required special programming techniques involved with message passing between the application fragments running on each processor.
- **SMP - shared memory systems.** SMP stands for Symmetric Multi-Processors, but later evolved to stand for any shared-memory system, even ones that use other memory architectures, as long as all system processors can access the entire memory address space.

SMP cost structure: In the late 80s, the SMP systems were the popular choice over MPP systems. The dominance of SMP was due to the ease of deployment, management and programming. However, these fully proprietary systems were expensive as they used custom processors, custom chipsets and ASICs to create high-speed backplanes as well as custom Operating Systems (OS).

MPP evolution: In the late 90s commodity processors became faster and offered a viable alternative for server solutions. This trend accelerated in 2002 with the adoption of loosely coupled compute clusters. These commodity servers interconnected with commodity high speed interconnects. Being similar in nature to MPP systems, it required more difficult programming model to allow the application to span across multiple systems. MPP systems deliver more performance than traditional SMPs, as the distributed architecture provides a bottleneck-free performance for share-nothing, or almost-nothing applications. Still, due to the need to implement and support multiple OS's and interconnect fabric technologies, these clusters are more difficult to deploy and manage.

Cluster complexity: Due to the complexity involved with cluster implementation and the cost associated with building SMP systems there is a need to provide a solution that can leverage the simplicity of SMP with the cost-structure of clusters. Leveraging x86 commodity components with software-based systems architecture can deliver the advantages of shared memory systems with the price point of cluster systems.

The Versatile SMP™ (vSMP) architecture: ScaleMP's vSMP architecture uses virtualization to abstract cluster deployment as large virtual SMP systems. The vSMP architecture eliminates the distributed nature of the cluster and enables unified compute resource optimized for capability, manageability, and flexibility. ScaleMP's vSMP Foundation virtualization platform implements the vSMP architecture. More about it can be found below.

THREE SOLUTIONS FOR THREE CUSTOMER REQUIREMENTS

vSMP Foundation for SMP: The vSMP Foundation for SMP solution targets customers seeking for large-scale x86 systems and rivals in performance traditional proprietary SMPs, yet at significantly lower cost. It is targeted at applications that benefit from large memory or high core-count such as simulation, analytics, and modeling. vSMP Foundation for SMP enables the aggregation of up to 16 servers into a single virtual machine, with up to 128 cores and 4TB or main memory. It supports the latest generation of Intel Xeon processors (Nehalem - 55xx series) and large memory as well as offering high availability features such as partitioning and redundant InfiniBand backplane.

vSMP Foundation for Cluster: The vSMP Foundation for Cluster solution provides a simplified compute architecture for high-performance clusters - it hides the InfiniBand fabric, offers built-in high-performance storage as cluster-filesystem replacement and reduces the number of operating systems to one, making it much easier to administer. This solution is ideally suited for smaller compute implementations in which management tools and skills may not be readily available. The target customers for this product are those with initial high-performance cluster implementations who are concerned with the complexity of creation and management of the cluster environment. vSMP Foundation for Cluster enables the aggregation of up to 16 servers into a single virtual machine with up to 128 cores and 512GB of RAM. It supports Intel Xeon 5500 (Nehalem) processors with speeds up to 2.4GHz.

vSMP Foundation for Cloud: The vSMP Foundation for Cloud solution enables on demand virtual SMP provisioning for enterprise cloud infrastructure. Tightly integrated with systems management and provisioning tools, vSMP Foundation for Cloud offers the flexibility of on-the-fly creation of virtual machines on a per job, per project, or per customer basis, eliminating the need to deploy dedicated shared memory systems. Using vSMP Foundation for Cloud allows for real-time deployment of large memory systems aggregated from cloud resources and redeployment back to the compute grid upon completion of application or batch jobs. Similar to the rest of our product lines, it support aggregation of up to 16 servers, utilizing Intel's latest generation of Intel Xeon processors (Nehalem - 55xx series) into a single virtual machine, with up to 128 cores and 4TB or main memory.

VSMP FOUNDATION SOLUTION SUMMARY

vSMP Foundation Key Advantages: ScaleMP's vSMP Foundation makes the next generation of affordable virtual high-end SMP systems possible. High-end SMPs can be created without investing tens or hundreds of millions of dollars in proprietary R&D, and losing valuable time to market. ScaleMP enables the creation of very affordable midrange to high-end SMP computers using commodity x86 server boards and standard interconnects that delivers the lowest overall Total Cost of Ownership by:

- Run demanding applications with best of breed performance for both cluster and SMP applications
- Use the latest generation of chips and interconnects to provide best performance at volume pricing
- Provide low management cost, by utilizing single point of management of SMP systems
- Maintain cost benefits of Clusters, minimizing the use of custom hardware and components

Functional Requirement	Business Value	Comments
Run applications that are designed to take advantage of SMPs	Leverage existing applications and programming methology	vSMP Foundation provides the versatility of being able to run different types of applications at equal to or better performance compared to both clusters and traditional SMP systems (multi threaded, multi process throughput, multi process cooperative, and single threaded, large memory applications. This gives customers the ability to deploy a single system to cover all their requirements.
Performance should be equal to or better than SMPs or Clusters	Provide solution with the highest levels of performance	vSMP Foundation optimizes memory locality, providing cluster performance for distributed applications and higher memory bandwidth than SMP for applications dependent on shared memory.
Leveraging the latest generation of chips and interconnects at any point in time	Faster time to market with higher performance lower cost technology	Solutions based on vSMP Foundation use software to replace chipset development, leveraging industry standard components, which speeds up time to market improve overall system performance and reduce system cost. System design and manufacturing cycle is reduced to less than 4 months, versus up to 3 years for a traditional SMP system or customers can chose to plug in vSMP Foundation software solution into standard systems.
Management costs should match SMP deployment models	Management flexibility and ease-of-use	Solutions based on vSMP Foundation provide a single point of management, thus reducing the on-going operational costs compared to clusters.
Acquisition should advantage of COTS cost take	Provide solutions that leverage low cost commodity systems for lowest TCO and fastest ROI	Solutions based on vSMP Foundation use off-the-shelf server systems and interconnects that are traditionally being used for clusters, delivering SMP ease of use at cluster pricing.

