ScaleMP™

Powering high-end x86 systems

Simplified HPC: SMP at Cluster Pricing

Use Cases

Aggregate. Scale. Simplify. Save.

ScaleMP – Overview

ScaleMP provides **virtualization for high-end computing**, delivering higher performance and lower Total Cost of Ownership (TCO) - allowing customers to **Aggregate, Scale, Simplify, and Save**

vSMP Foundation **aggregation** software creates a **virtual shared-memory multi-processor (SMP)** by aggregating multiple off-the-shelf **industry-standard x86** servers

New Server Virtualization Paradigm

Existing: Partitioning

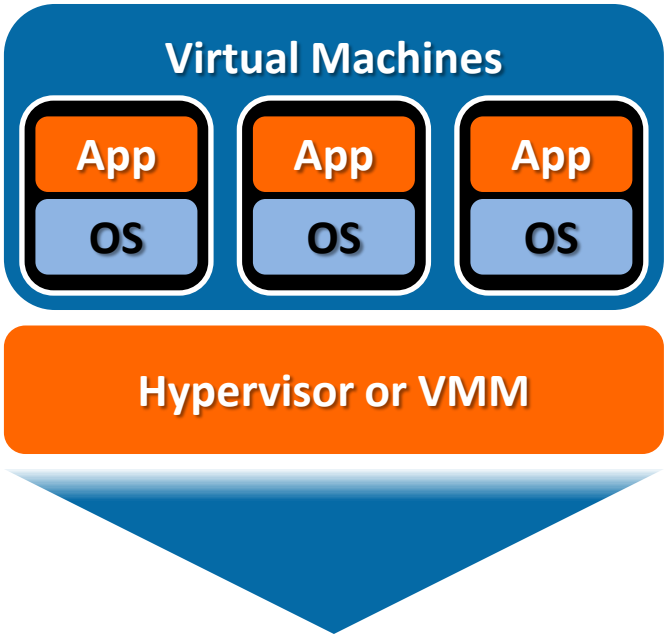

Applications requiring **fraction** of the physical server resources

New: Aggregation

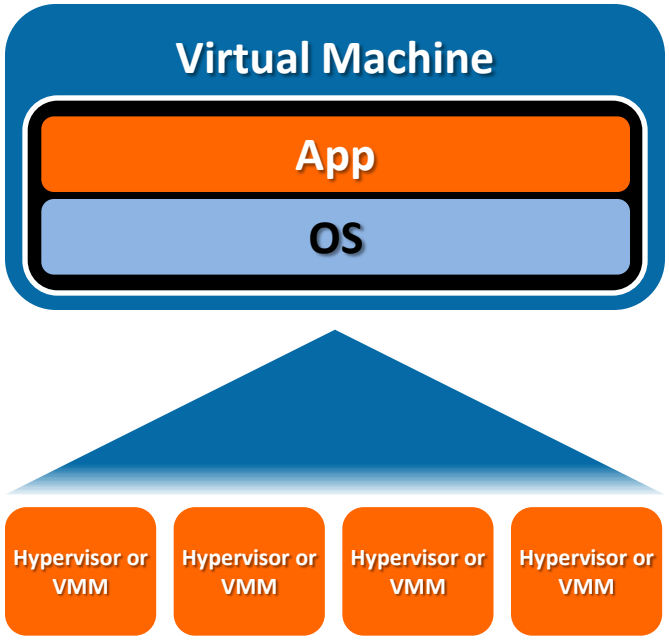

Applications requiring **superset** of the physical server resources

Existing HPC Deployment Models

Applications requiring **superset** of the physical server resources

Scale-Up

Fit the hardware to the problem size

Scale-Out

Break the problem to fit the hardware

Existing HPC Deployment Models

PROS AND CONS

Scale-Up

- Simplified IT infrastructure
  - Single system to manage
  - Consolidated I/O
- Proprietary hardware design
  - High cost
  - Vendor lock-in

Scale-Out

- High installation & management cost
  - Multiple operating systems
  - Cluster file system, etc.
- Leveraging industry standard servers
  - Low-cost
  - No vendor lock-in

Aggregation: Best of Both Worlds

Scale-Up

Scale-Out

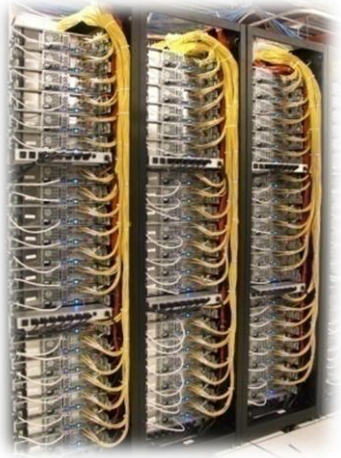

Aggregation

- Simplified IT infrastructure
  - Single system to manage
  - Consolidated I/O
- Leveraging industry standard servers
  - Low-cost
  - No vendor lock-in

vSMP Foundation™

HOW IT WORKS?

Multiple off-the-shelf x86 boards, with processors and memory

Processors speed/amount or memory amount doesn't have to be same across all boards

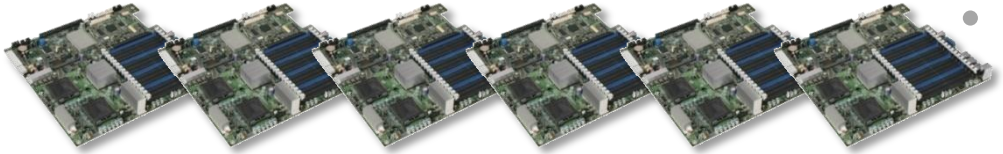

InfiniBand HCAs, cables and switch

vSMP Foundation™ Devices

The flash-devices plug into the boards and used as bootable device. vSMP Foundation is booted to present an aggregate coherent view to the OS

High-end x86 system, based on standard x86 components

vSMPowered™
scalable systems by ScaleMP

vSMP Foundation™

BEHIND THE SCENES

One System

- Software interception engine creates a uniform execution environment
- vSMP Foundation creates the relevant BIOS environment to present the OS (and the SW stack above it) as single coherent system

Seamless integration

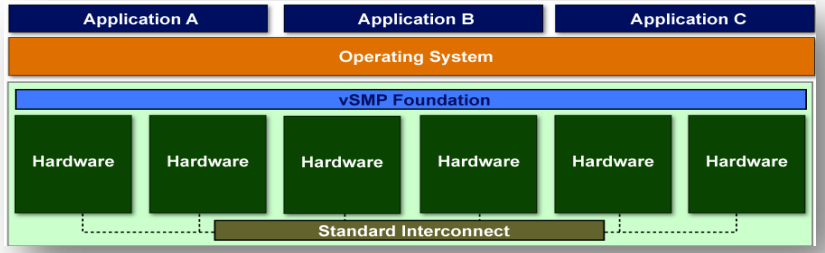

Coherent Memory

- vSMP Foundation maintains cache coherency between boards
- Multiple concurrent memory coherency mechanisms, on a per-block basis, based on real-time memory activity access pattern
- Leverage board local-memory for caching

Highest x86 SMP memory bandwidth!

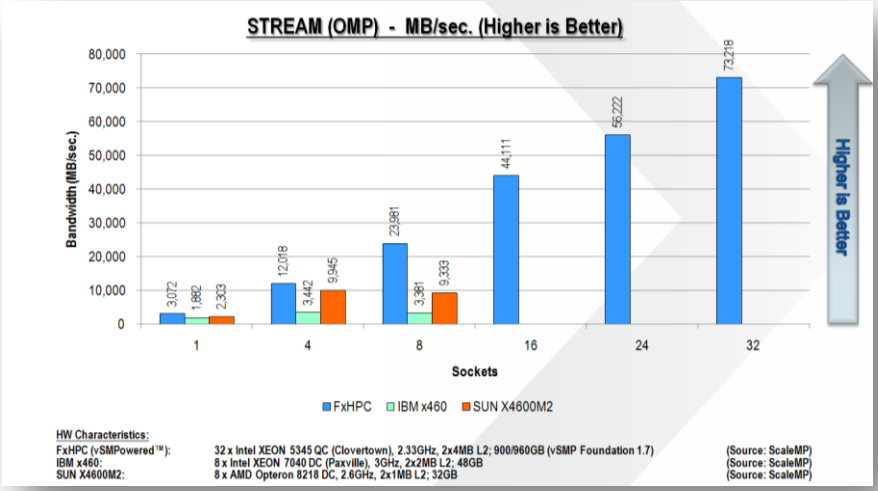

Shared I/O

- vSMP exposes all available I/O resources to the OS in a unified PCI hierarchy
- No need for cluster file systems

Why Aggregate?

OVERCOMING LIMITATIONS OF EXISTING DEPLOYMENT MODELS

Fit the hardware to the problem size

- **Large memory** x86 resource
  - Enable larger workloads that cannot be run otherwise
  - Alternative to costly and proprietary RISC systems
- **High core-count** x86 shared-memory resource with high memory bandwidth
  - Allow threaded applications to benefit from shared-memory systems
  - Alternative to costly and proprietary RISC systems

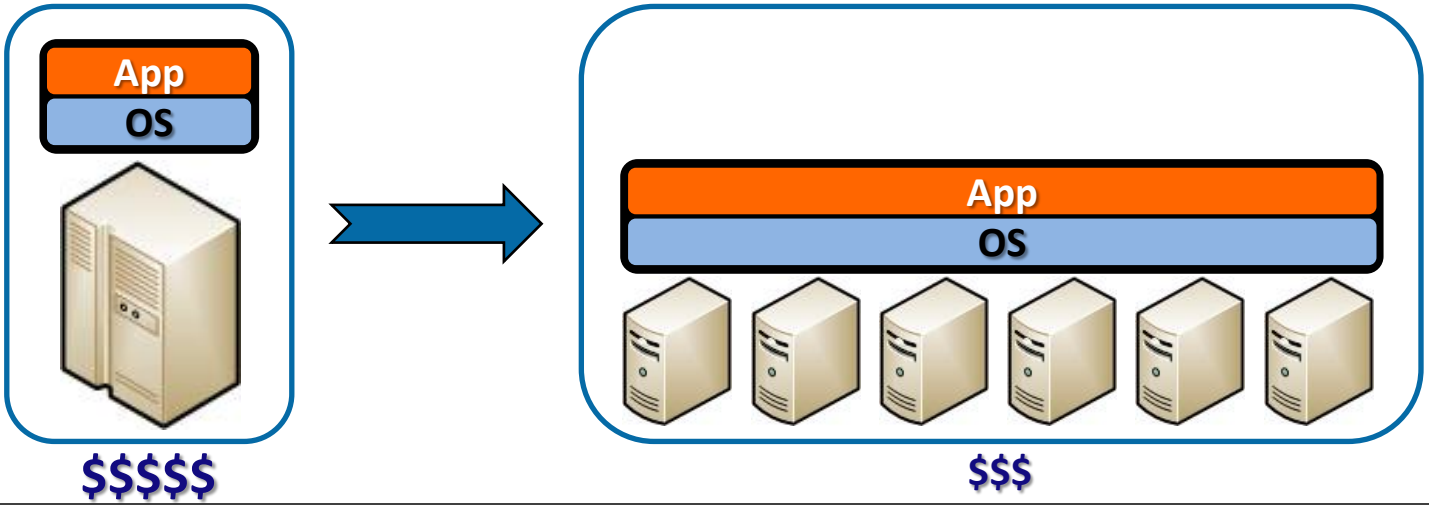

Why Aggregate?

OVERCOMING LIMITATIONS OF EXISTING DEPLOYMENT MODELS

Break the problem to fit the hardware

- **Ease of use**
  - One system to manage: fewer, larger nodes means less cluster management overhead
    - Single Operating System
    - Avoid cluster file systems
    - Hide InfiniBand complexities
  - Shared I/O

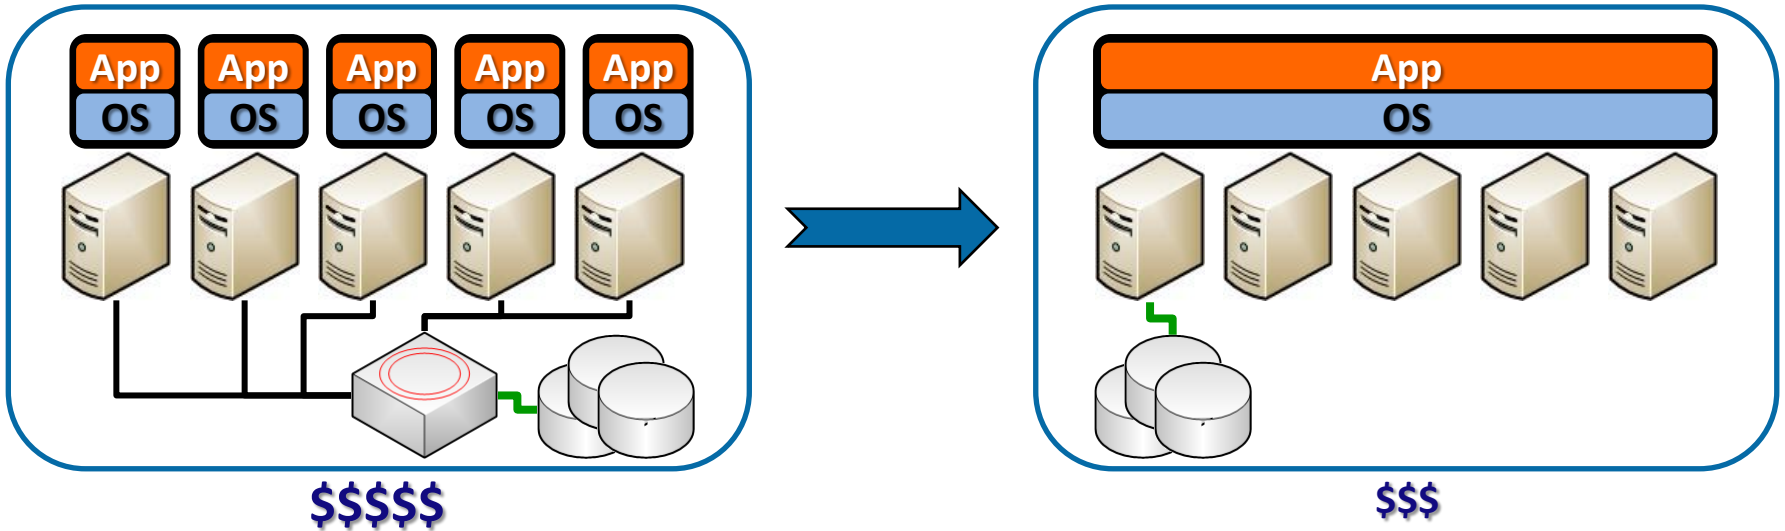

Why Aggregate?

OVERCOMING LIMITATIONS OF EXISTING DEPLOYMENT MODELS

Break the problem to fit the hardware

- **Improved utilization**
  - Fewer, larger nodes reduce resource fragmentation
  - More throughput without increase of hardware resources

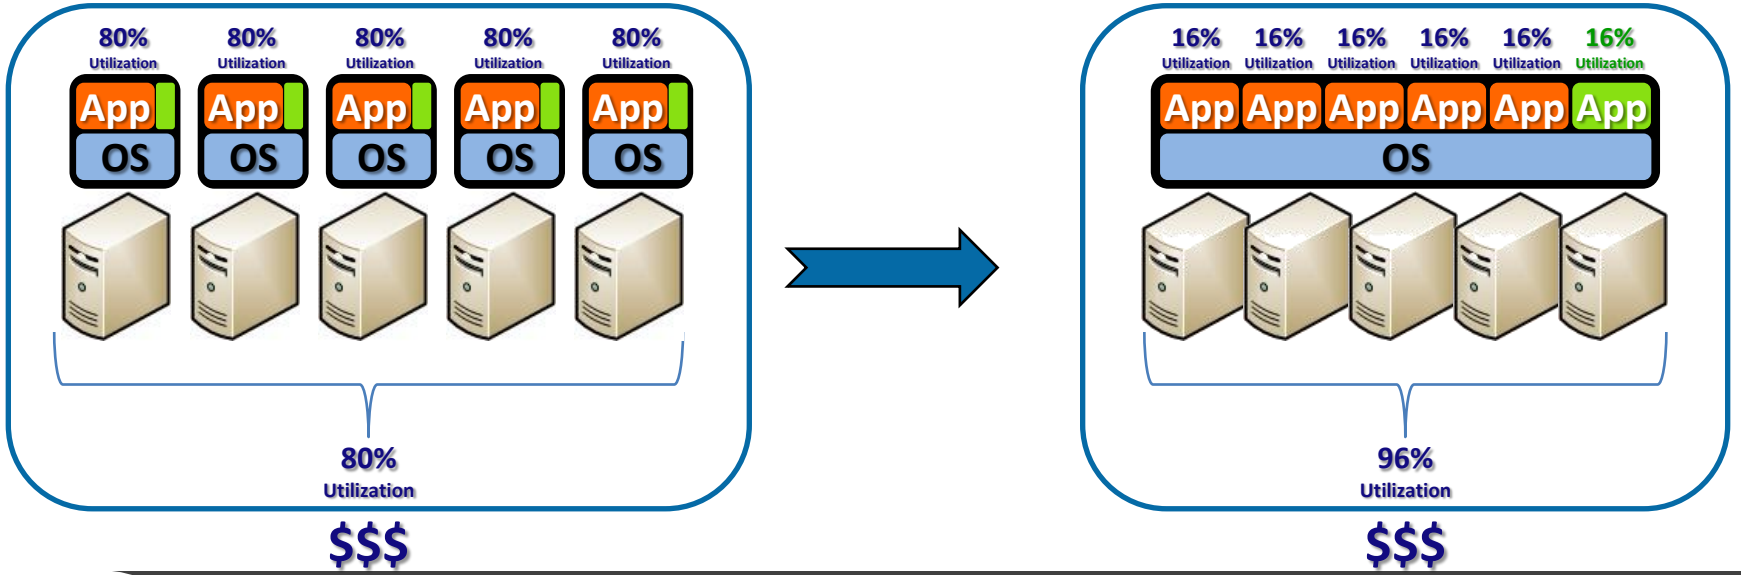

Customer Use Case #1

FORMULA1 TEAM

- **Customer:** Formula1 team
- **Current platform:** Large-memory SGI Itanium-based system
- **Problems:**
  - Need to generate large mesh as part of pre-processing of whole-car simulation (FLUENT TGrid)
  - Mesh requirements are ~200GB in size
  - Expect to grow significantly within 12 months after initial deployment
  - Would like to standardize on x86 architecture due to lower costs and open standards
- **Solution:**
  - 12 Intel dual-processor Xeon systems to provide 384GB RAM single virtual system running Linux with vSMP Foundation
- **Benefits:**
  - **Better performance:** Solution evaluated and found to be faster than alternative systems (x86 and non-x86)
  - **Cost:** Significant savings compared to alternative system
  - **Versatility:** Also being used to run FLUENT (MPI) as part of large cluster
  - **Investment protection:** Solution can grow to 640GB RAM with additional 4 Intel dual-processor systems with 64GB each

SCALEUP AT SCALEOUT PRICING

Customer Use Case #2

ENGINEERING SERVICES COMPANY

- **Customer:** Mid-size Engineering Services Company
- **Current platform:** Multiple 2-socket workstation
- **Problems:**
  - Existing models (Abaqus) grow fast and can't fit the engineers workstation
  - Interested in running apps in batch at night
  - No in-house skills to run x86 InfiniBand cluster (although the application running nicely on InfiniBand cluster)
  - Can't afford RISC systems
- **Solution:**
  - 4 Intel dual-processor Xeon systems to provide 128GB RAM, 8 sockets (16 cores) single virtual system running Linux with vSMP Foundation
- **Benefits:**
  - **Performance:** Solution evaluated and found to be significantly faster than existing workstations. Performance found to be comparable to cluster performance (using vendor benchmarks).
  - **Low OPEX:** No IT required for day-to-day operation
  - **Versatility:** Night time executed in batch mode. Most day time jobs are executed on the system while using the workstation for display only. Multi-user environment provides perfect scaling – allowing sharing without performance degradation.
  - **Investment protection:** Expected to expand the system by adding additional 4 nodes (to a total of 256GB RAM, 32 cores)

ENABLE INNOVATION WITHOUT THE COMPLEXITY

Customer Use Case #3

MULTINATIONAL ENERGY COMPANY

- **Customer:** Multinational Energy Company
- **Current platform:** x86 grid
- **Problems:**
  - Using in-house single-threaded simulation tools in throughput mode. Memory footprint of each simulation has grown over the years and sometimes (10%) exceeds 32GB.
  - Application runs on x86 only
  - Used to reschedule failed runs on large-memory systems
- **Solution:**
  - 6 Intel dual-processor Xeon systems to provide 192GB RAM, 12 sockets (48 cores) single virtual system running Linux with vSMP Foundation
- **Benefits:**
  - **Versatility:** Both large and small workloads used concurrently on the same system
  - **Utilization:** Higher utilization compared to grid due to lower infrastructure fragmentation
  - **Investment protection:** Solution expanded by 100% since initial installation

SINGLE INFRASTRUCTURE FOR HORIZONTAL AND VERTICAL APPLICATION SCALING – PLUG & PLAY

ScaleMP™

Powering high-end x86 systems

Shai Fultheim
Founder and President

Shai@ScaleMP.com, +1 (408) 480 1612

Aggregate. Scale. Simplify. Save.

Certified Hardware

Discrete Systems

- Up to 2 systems connected
- Direct connection – no IB switch
- Most affordable four-socket solution

• SR1560SF • SR1560SFHS	• 6015TW-INF	• X3450 • DX340 (iDataPlex) • DX360 (iDataPlex)	• PowerEdge 1950 III

Blade Servers

- Scale to full chassis capabilities
- Integrated InfiniBand required
- Shared-memory supercomputer

• **SuperBlade:**
  • SBI-7125W-S6
	• **BladeCenter:**
	  • HS21XM
		• **M1000E:**
		  • M600