APM and the Incredible Shrinking Computer

Mark Preston, First Data Corporation

Application Performance Management (APM) is undergoing significant change, which is highly disruptive to the established field of computer instrumentation and management. Proactive organizations are moving from vertical monitoring and analysis of a few tiers of service to horizontal transaction monitoring across all tiers of a business transaction. The convergence of three disruptive technologies - manufacturing process densities, packaging and distributed computing - has given us the incredible shrinking computer. While horizontal APM and distributed high-density computing can be viewed in isolation, they are also complementary, eventually providing us with a better IT solution.

Several years ago, while evaluating the performance of next-generation storage solutions, I remarked at the incredible change in the makeup of a server. Staring down at the abyss of the hardware enclosure there appeared few remaining discrete components as manufacturing processes has enabled ever increasing integration of functionality. The processor, memory and I/O real estate were dominated by power and thermal management. Sure, there were many support components for legacy I/O, bios, boot and management that consumed real estate and power but the impact of Moore’s law was obvious. Component integration opened up new opportunities in server packaging, initially leveraged as inefficient and hot pizza boxes.

Enter the blade server in the early 2000’s arguably as a disruptive server technology. Initially targeting the web tier the early blade designs were based on low power processors while utilizing traditional motherboard designs. Further refinement introduced shared I/O and integration of fabric switches. But even with blades the actual compute card has remained relatively complex.

Distributed computing also emerged as a way of dealing with complexity, the high cost of scale-up computing, and the explosion of data. Instead of buying big iron, scale out with low cost computing that communicate over a network. Push the work out to several lower power nodes to leverage parallelism vs. pulling the work into a single high-powered node. This strategy has also being employed for Big Data as Hadoop and its spin-offs (Hive, Pig, ZooKeeper) enable structured processing to be pushed out to distributed storage vs. the traditional strategy of hub-and-spoke that pulls data into the compute node. The recently announced Microsoft SQL Server 2012 will leverage Hadoop for scale-out. How will the next disruptive technology impact instrumentation and management? We will get to that shortly.

A brief revisit of history is needed to set the architectural landscape for the four quadrants of computing; CPU, Memory, Storage and Networking. In the late 90's AMD design work x86-ized the Alpha architecture by moving to a 64-bit NUMA (Non-Uniform Memory Access) x86, integrating the memory controller and introducing Alpha Processor Inc's Lightning Data Transport (LDT) point-to-point processor interconnect, later to be called Direct Connect Architecture, now HyperTransport. With the Nehalem core Intel also moved to the NUMA architecture with an Integrated Memory Controller and Common System Interface, later named Quick Path Interconnect (QPI). Recently announced, Intel's Sandy Bridge family adds PCIe 3.0 to QPI. Tight integration of the CPU and Memory complex is clear although there remain distinct differences between Intel and AMD relative to I/O integration. Embedded design are clearly integrating all four quadrants; CPU, Memory, Storage, and Network. Enterprise data center's on the other hand continue to require I/O flexibility, moving towards a CPU/Memory and Storage/Networking architecture.

A recent development in compute density is the technology introduced by SeaMicro. Viewed by the industry as a disruptive technology, SeaMicro's Technology Overview list CPU I/O virtualization, a super-computer fabric and CPU management and load balancing as primary technical innovations. [1] The key technology, named Freedom Fabric and embodied in an ASIC (Application-Specific

Integrated Circuit), is the placement of a fabric between CPU/Memory and Network/Storage. According to press releases this has allowed Seamicro to remove 90 percent of components from the traditional definition of the motherboard. The Seamicro ‘microserver’ is comprised of a single-socket, multi-core CPU, memory and Freedom Fabric ASIC. The microserver is packaged as a credit card size FRU (Field Replaceable Unit) module. Multiple microserver~~’~~s are interconnected with a multi-dimensional torus fabric. Network and Storage is virtualized with I/O-virtualization technology (IOVT). CPU/Memory FRU management allows dynamic creation of application compute pools and allocation of workload.

As with early blades the initial SeaMicro servers were based on the low power adequate computing Atom processor. Prior to AMD’s acquisition of SeaMicro Sandy-bridge high-density servers were announced. Intel is also developing a fabric although details have not been released. [2] So while this has been characterized as a “micro-server” play, the same technologies could be applied to other server segments.

But this article is not about the choices Intel and AMD will make with respect to system architectures. At a high level that path is clear, further integration, low latency fabric between CPU/Memory and Storage/Networking, simplification of functional modules (FRU’s), high density computing frames and hardware integration of dynamic virtualization technology.

As recognition of this disruptive technology IDC, beginning with the 4Q11 Tracker, added the category Density Optimized systems. IDC characterized this category as *“Density Optimized servers are designed for large scale datacenter environments where parallelized workloads are prevalent. The form factor serves the unique needs of these datacenters with streamlined system designs that focus on performance, energy efficiency, and density.”* [3]

Density Optimized systems could be defined as hundreds of processors in a xU rack space, packaged as field replaceable modules (FRU’s) within the frame. As with SeaMicro, density-optimized systems offer dynamic provisioning of Compute+Memory, Network and Storage FRU's.

Today the typical IT organization is structured around a load-balanced multi-server multi-site model for performance and availability. Notification of outages sometimes is the phone-call or virtual environment event notification. In far too many cases resolution is characterized as all-hands-on-deck and mean-time-to-innocence as disciplines attempt to point fingers at others. Outage resolution can take minutes to days as a very manual intervention strategy plays out. Pinpointing the source of the outage or degradation can be very painful and costly as remediation must be immediate.

It is conceivable that new IT models comprised of highly parallelized applications, hosted by density optimized platforms will move towards an over-provisioned hardware with an aggressive alerting, component isolation and replacement model that takes advantage of density computing. This IT model will rely on enhanced transaction monitoring and application performance management (APM) to effect changes on provisioning and service levels. Return-To-Service needs will most likely drive a proactive strategy embodied in integration of APM and density computing vs. today's traditional reactive all hands on board strategy.

Integration of APM and Density server introduces some interesting possibilities. Similar to Solid State Storage performance management, can a density server be over-provisioned with microservers for both workload changes and failure remediation? As we see with the latest product announcements microservers can be dynamically brought on-line and allowed to hibernate. This goes beyond today's load-balanced redundant server and site availability solutions that are large, always-on, power consuming and data center facilities hungry.

If we move to a FRU strategy how deep do monitoring tools need to go? Does instrumentation need to capture each sample or can we move to an event-driven, alerting model? Is the density server monitored with Bayesian analytics such that microservers are "removed from the pool" when they fall outside the parameters of normal observation? Do we need to resolve a server failure or simply decommission/replace a FRU for future replacement? How many failed FRU's can exist in a frame before it needs attention? Is the answer one or many? If the cost of a microserver is low then an organization will most likely evaluate the cost of a fully populated frame vs. the cost of immediate remediation. Maybe simply fully populating a high density microserver-based frame is the better choice. For Return-to-Service issues driven by hardware failure or site outages and Capacity Planning for new customer workload the APM+FRM module is most likely the computing model of choice.

On the other hand for deeper dive issues today's tool set will still be required but as history has shown us, when two solution domains become tightly integrated new tool opportunities arise and the need for multiple tools is either diminished or is no longer needed. Is the next disruptive technology the tight integration of Application Performance Management (APM) and IDC's new density-computing category?

1. SeaMicro Technology Overview
Anil Rao, May 2010
www.seamicro.com

2. Intel Dismisses AMD Acquisition of SeaMicro Deal
eweek
IT & Network Infrastructure News & Reviews

3. IDC - Press Release
28 Feb 2012