Understanding Architectural Priorities Inside Government Clouds

Not since the advent of the Internet has a technological concept achieved more attention than cloud computing. Driven by the promise of greater operational efficiency, lower costs and the potential to free up valuable resources for crucial mission goals, cloud computing is likely to dominate the tech sector for the next several years.

Research firms, such as IDC in Framingham, Mass., predict that sales of cloud computing products and services will grow at a compound annual rate of 25% per year, generating $56 billion in annual revenue in 2014. *(source: Worldwide Enterprise Server Cloud Computing 2010–2014 Forecast, IDC, Framingham, Mass.)*

Meanwhile, recent initiatives by public sector agencies, universities and other educational institutions around the world are seeking to reduce energy consumption, improve cyber security protections and migrate to more efficient IT platforms. In the U.S., federal agencies are incorporating cloud computing solutions to streamline IT operations, gain greater efficiency and meet a mandated goal to achieve $3 billion in cost savings by the end of fiscal 2012. Early in 2010, the Office of Management and Budget launched the Federal Data Center Consolidation Initiative (FDCCI) to accelerate consolidation and migration to cloud services. By the U.S. government's own estimates, data centers grew from 432 to 1,100 in the last decade, according to Federal CIO Vivek Kundra[1], who is working to migrate more of the federal government's IT infrastructure and services to private, public and hybrid cloud implementations. So far, U.S. government agencies and departments are striving to meet the mandates. An increasing number of examples of cloud computing are popping up at all levels of government. The cities of Los Angeles, Washington D.C., Carlsbad, Ca., and Miami, as well as GSA, the U.S. State Department, National Aeronautics and Space Administration, and Air Force are all using cloud implementations for everything from email services to online civic information services to global budget analysis at the State Department and a cloud system used by NASA for mission support, public education, data communications and storage.

At the same time, governments around the world are also investing in cloud-based solutions. The European Union (EU) is finalizing the implementation of its Resources and Services Virtualization without Barriers Project (RESERVOIR), which will deliver a cloud-based IT infrastructure to provide online services as utilities. Meanwhile, the Canadian government is hoping to leverage its cooler geographic climate to house cloud computing facilities, and Japan's Ministry of Internal Affairs and Communications (MIC) has launched a nationwide public cloud computing infrastructure plan currently known as the Kasumigaseki Cloud.

In the past, uncoordinated and decentralized IT asset and data center growth led to many government agencies owning and operating multiple data centers that generally met only a few functions, served a single part of the organization's needs and weren't interoperable. Today, virtualization and cloud computing are helping to drive greater consolidation and modernization in public sector data centers, allowing agencies to maximize their investments through greater hardware utilization, while lowering the total cost of ownership of IT infrastructures. For specific applications and services, early adopters have proven that provisioning servers in a cloud can dramatically shrink an organization's data center footprint, and offload at least some applications and services to public, private or hybrid cloud-based services.

Security and Privacy Concerns

Government organizations face more stringent regulatory requirements for protecting the privacy and security of information than most businesses in the private sector. From sensitive financial, health and personal data, to public safety services, national security, and other classified or top secret information, public sector IT operations are legally bound to protect such resources. This is why so many government IT officials are turning to private or hybrid cloud implementations, which grant the flexibility to gain many of the benefits of cloud computing, while still maintaining control over both infrastructure and IT operations.

While some less-sensitive public sector services and applications may still be good candidates for public cloud implementations, private cloud solutions may deliver greater control and security because the infrastructure is operated solely for one or more organizations' private use. This type of cloud may be managed by an organization or a designated representative (system integrator) and can exist on or off premise, depending on security requirements. Hybrid cloud implementations, meanwhile, leverage a combination of private and public cloud services for a government organization's applications and/or services. Balancing the decision to implement public versus private cloud services must be determined by the degree of security and control required. The greater the requirements for security, the more likely a customer will turn to a highly secured private cloud implementation.

AMD-V Features and Benefits

AMD-V technology includes:

FEATURES	BENEFITS
Virtualization extensions to the x86 instruction set	Enables software to more efficiently create virtual machines so that multiple operating systems and their applications can run simultaneously on the same computer.
Tagged TLB	Hardware features that facilitate efficient switching between virtual machines for better application responsiveness.
Rapid Virtualization Indexing (RVI)	Helps accelerate the performance of many virtualized applications by enabling hardware-based virtual machine memory management.
AMD-V™ Extended Migration	Hardware feature that helps virtualization software enable live migration of virtual machines between all available AMD Opteron™ processor generations.
I/O Virtualization	Enables direct device access by a virtual machine, bypassing the hypervisor for improved application performance and improved isolation of virtual machines for increased integrity and security.

Living Inside the Cloud

While the migration to cloud computing services seems inevitable, crucial obstacles, including security, privacy, a lack of standards, a fear of vendor lock-in and performance concerns have slowed the pace of public sector cloud adoption. Another important consideration as government organizations start to migrate toward predominantly private cloud-based services – the ongoing, strong requirement to acquire technology and server resources to run cloud-based implementations. In its worldwide cloud computing server forecast, IDC predicted public and private cloud computing will shape all server sales through 2014. Servers used for cloud computing will grow from $3.8 billion in revenue in 2010, representing over 600,000 units shipped, to a $6.4 billion market in 2014, with over 1.3 million units shipped. *(source: Worldwide Enterprise Server Cloud Computing 2010–2014 Forecast, IDC, Framingham, Mass.)*

What many public sector organizations may not yet realize is that no matter which type of cloud-based solutions are selected, organizations must still pay close attention to the underlying hardware architecture to be deployed for any cloud computing environments. It's crucial for government IT officials to understand the primary architectural characteristics of the server resources planned for use. Without a doubt, government agencies should be seeking the best balance of scalability with low power consumption, leveraging the most cost-effective technologies. Raw performance isn't as critical as balanced performance and scalability, because cloud applications are spread over many hundreds or thousands of nodes and cores, not constrained to a single server or a limited set of cores. In these environments, multicore processors are required, with higher core counts being preferred over raw clock speed.

In the server hardware realm, differentiators today revolve around how many users and/or applications can be supported per server – as well as the actual power draw (at the wall) for each server. Also, public sector IT officials must weigh how much each server or virtual machine will cost over the course of its useful life. In cloud computing, total cost of ownership (TCO) should drive server purchase decisions. And only one supplier delivers the ease of management agencies require, along with an extremely long server lifespan.

AMD Stands Apart

AMD's solutions stand out in three primary ways: low power consumption; core density/scalability; and integrated virtualization features. This is how AMD's server platform technology is uniquely tuned to address critical government requirements for reducing costs, simplifying operations and raising energy efficiency in server price/performance. "Using AMD Opteron™ processors to power our HP ProLiant servers has been the perfect complement to Windows Server® 2008 R2. The combination delivers optimal real-world performance and increased server responsiveness and reliability. [AMD] exceeded our expectations – both in terms of performance and power savings," said Sumeeth Evans, Director of Information Technology, Collegiate Housing Services. *(Source: http://sites.amd.com/us/Documents/48753A_Collegiate_Case_Study_WEB.pdf)*

AMD Opteron™ 4100 and 6100 Series processors currently offer up to 12 cores per processor – up to a 120% increase in performance from the previous generation of AMD models.[2] This means that servers can be designed that offer up to 48 cores in the popular 4P, 2U platform form factor with an extremely efficient power envelope, with some processor models using as little as six watts per core. AMD also incorporates a suite of power

features that allow IT managers to control and monitor power consumption. Looking ahead, the next generation of AMD Opteron processors will be based on a brand new, even more efficient core architecture that will feature up to 16 cores per processor in the same server footprint from a power and heat generation standpoint as the 2010 models that featured up to 12 cores – anticipating a delivery of another 50% increase in throughput.

With more cores in the same power and thermal envelope, customers gain greater utilization from virtualized applications and cloud computing, using the same or less space and power. Also, as government organizations grow cloud implementations to meet the ever-increasing demand for services, AMD customers can rest assured servers featuring the next generation AMD Opteron 4200 and 6200 Series processors, with consistent platform technology, and consistent system image, will help ease management now and in the future.

Meanwhile, AMD Virtualization™ (AMD-V™) technology can help public sector organizations do more with fewer systems. By allowing multiple operating systems and applications to run simultaneously on a single computing system, virtualization delivers dramatic operational benefits that can help provide a cost savings while increasing the value and capabilities of any public sector IT investment. Virtualization also helps enable cloud computing and its CPU and memory intensive applications while maintaining power efficiency. By utilizing virtualization as a core element of a cloud deployment, customers increase both agility and manageability by spreading the cloud load over a large number of systems and processor cores.

AMD Virtualization (AMD-V™) technology is a set of hardware extensions to the x86 system architecture that allow cloud operators to optimize resources, making servers, clients and data centers more effective.

Addressing Security in the Cloud

Another advantage of AMD-based servers is strong, hardware-based security. Security forms the basis for protecting sensitive data, IT networks and critical infrastructure from unauthorized access, malicious intrusions and other breaches. Most security solutions today tend to focus on software applications to address identity management, multi-tenant support in virtualized environments, network governance, storage isolation/encryption, as well as regulatory and compliance issues (SAS70, PCI DSS, HIPAA, etc.), among other concerns. For maximum levels of protection, true end-to-end security requires the integration of protection mechanisms at all layers of the IT stack, beginning with the hardware layer.

When it comes to hardware-based protection, security begins at the bedrock of computing – the processor. Comparable to all x86 solutions, AMD offers competitive security features that address information-sharing and control at the platform level.

Energy Efficiency and Total Cost of Ownership

Energy efficiency and total cost of ownership (TCO) benefits make AMD's processor technology a great fit for cloud computing data centers, where efficient servers allow for increased density through more complete rack utilization – without sacrificing performance or functionality.

Of course, the balance of server energy efficiency, density and performance required to optimize TCO can vary dramatically depending on each organization's requirements. AMD Opteron processors, for example, offer the lowest power per core of any server processor[4]. AMD Opteron processors improve power efficiency by leveraging key technologies, including:

- **Enhanced AMD PowerNow!™ Technology** – integrated power monitoring and management that allows the processor to dynamically adjust power based on utilization loads
- **Independent Dynamic Cores** – allows each core to vary its frequency based on each system's specific needs. This enables more precise power management to reduce data center energy consumption and TCO.
- **Dual Dynamic Power Management** – permits each processor to maximize the power-saving benefits of Enhanced AMD PowerNow! technology without compromising available performance. Dual Dynamic Power Management can reduce idle power consumption and allow for per-processor power management in multi-socket systems to decrease power consumption.
- **AMD CoolCore™ Technology** – reduces power consumption within each core by evaluating the parts needed to support currently running applications. It can cut power to unused transistor areas to reduce power consumption and lower heat generation.
- **AMD Smart Fetch Technology** – helps reduce power consumption by allowing idle cores to enter a 'halt' state, to draw even less power during processing idle times.

Public sector organizations can check out estimated savings using AMD's power saving estimator tool listed here: *http://sites.amd.com/us/mod/tools/Pages/server-power-estimator.aspx*

This includes:

- Dynamic Root of Trust technology designed to ensure a secure state before launching an operating system in an unknown state. This technology protects system integrity by alleviating potential vulnerabilities associated with BIOS, BIOS extensions, bootstrap loaders and other system applications.
- Enhanced Virus Protection technology, a feature that sets portions of system memory aside as 'data only.' Any code resident in these areas may not be executed, only read from or written to. AMD's Enhanced Virus Protection, working in concert with OS support, acts as a preventative measure, localizing the virus where it will be short-lived, non-contagious and eventually flushed entirely from system memory. As part of a comprehensive security program, AMD strongly recommends enabling Enhanced Virus Protection (EVP) and using up-to-date third party anti-virus software.[3]
- AMD-V also includes key extensions to enable the establishment of a 'root of trust' within an initially 'untrusted' operating mode. This technology enables pre-authentication of a hypervisor or virtual machine image before users are able to decrypt and load them.

Establishing trust and transparency in a remote and highly dynamic shared cloud environment, similar to what most organizations already possess in IT infrastructures they operate today, remains a difficult challenge. Achieving that trust requires platform integrity monitoring and reporting capabilities. This is why AMD supports the ongoing work of the Trusted Computing Group (TCG), an international standards body that's developing and defining elements that promote open, vendor-neutral, industry standards for trusted computing across multiple platforms. AMD is a founding member of TCG. More information is available at *www.trustedcomputinggroup.org*.

Systems integrators, meanwhile, should also investigate AMD's server components to help deliver cloud-based solutions for government audiences as cost effectively as possible. AMD provides the low-cost cloud computing server alternative, featuring easy to manage virtualized servers that deliver exceptional energy efficiency and low operational costs.

This is likely why many government agencies, universities and schools turn to OEMs that provide AMD-based platforms for efficient, cost-effective and secured server hardware technologies. By selecting AMD as the underlying server platform technology, government agencies around the world may achieve significant performance gains compared to available alternatives, along with a reduction in ongoing operational costs, and greater efficiencies in power consumption and systems management.

In the coming year, it will likely become clear that not all cloud platforms function alike. What goes into each cloud's internal infrastructure, no matter whether an agency builds or buys a cloud-based service, will impact a customer's energy efficiency, IT security and systems management for the foreseeable future. To meet growing public sector energy, cost and performance mandates and derive superior performance for low prices, public sector audiences can rely on AMD.

1 – Source: Keynote address by Federal Chief Information Officer Vivek Kundra at a Brookings Institute conference on April 7, 2010, called the "The Economic Gains of Cloud Computing."
http://www.brookings.edu/~/media/Files/events/2010/0407_cloud_computing/0407_cloud_computing_kundra_presentation.pdf

2 – Based on Six-Core AMD Opteron™ processor Model 2435 vs. AMD Opteron™ processor Model 6174 running SPECfp®_rate2006. The results stated above reflect the best performing two-socket servers using AMD Opteron™ processor Models 2435 and 6174. The result for AMD Opteron™ processor Model 6174 is based upon data submitted to Standard Performance Evaluation Corporation as of March 17, 2010. For the latest SPEC results, visit *http://www.spec.org/cpu2006/results*. SPEC and SPECfp are registered trademarks of the Standard Performance Evaluation Corporation

3 – As part of a comprehensive security program, AMD strongly recommends enabling Enhanced Virus Protection (EVP) and using up to date third party anti-virus software

4 – As of October 8, 2010, AMD Opteron™ processor Models 4162 EE /4164 EE have the lowest known power per core of any server processor, at 5.83W (35W/6 = 5.83W/core). Intel's L5609 is 10W/core (40W/4 cores). See *http://www.intel.com/p/en_US/products/server/processor/xeon5000/specifications*.