

CUSTOMER:
AMD
www.amd.com

IN PARTNERSHIP WITH:
AMD IT

INDUSTRY
Semiconductor design

CHALLENGES

- Compute power in isolated pockets
- Software non-standardized
- Resources not flexible and accessible

SOLUTION

- Develop a private cloud for AMD

The AMD IT organization transforms the AMD infrastructure into a private cloud.

RESULTS

- Grew base compute resource capacity over 20%
- Achieved sustained utilization rate greater than 90%
- Saved over \$6 million through in-place upgrades

AMD TECHNOLOGY AT A GLANCE
AMD Opteron™ Processors

"We've never seen increases in productivity, collaboration and efficiency like we've seen with the cloud. Projects that used to take weeks can take days. We believe that the AMD cloud represents a significant business advantage."

Bob Luong
Director for Systems Engineering
AMD



AMD creates a private cloud to increase its own productivity, improve accessibility, lower costs and maintain its position as industry innovator

AMD is a semiconductor design firm with a history of innovations that have shaped the modern computing landscape. With the emergence of cloud computing, AMD again assumed the role of technological leader and pioneered a generation of power-efficient, multi-core processors to serve as the computational bedrock of cloud computing installations around the world.

AMD Faced the Challenges of Complexity
Designing semiconductors is a complex engineering task that requires massive computing power. For AMD, that power must be accessible to a global cadre of engineers in multiple disciplines. AMD built numerous clusters of systems to meet that need, deploying over 15,000 servers. The servers themselves had numerous releases of operating systems and file systems, as well as differing versions of storage management and cluster software. And with infrastructure growing at 20-30% annually, upgrading both hardware and software was a near-constant activity.

Moreover, the challenge of complexity was not limited to hardware and OSs. Managing and moving data were equally problematic. Because the global installations were not connected, huge data sets had

to be replicated between sites. As a result, it could take weeks to establish usable project data in a given location.

Once a project was underway, still newer obstacles would arise. Testing semiconductor designs requires access to tens of thousands of cores. However, to avoid wasteful and costly over-provisioning, no single AMD installation had excess available servers on that scale. So, to complete testing cycles that meet AMD's rigorous standards, engineers would often have to run tests on available machines in sequence, a time-intensive approach. While AMD had more than enough servers in total to accelerate testing, there was no way to shift the compute power to a particular project team that needed it for only a brief time. In short, the computing power AMD had created was massive. But it did not meet AMD's need for flexibility, manageability, accessibility or cost-effectiveness.

AMD 

AMD IT Migrates Infrastructure to the Cloud

In April of 2009, the AMD IT organization initiated a project to transform the isolated pockets of servers and storage AMD had deployed worldwide into a private cloud. The overarching strategy, internally referred to as “Compute Anywhere,” rapidly took shape. The tactical plan was to centralize data centers, provide appropriate connectivity tools, and thereby establish a private AMD cloud to fuel all AMD design projects worldwide.

In an early phase, AMD IT replaced aging processors in servers throughout the organization with newer, more power-efficient AMD *Opteron*™ processor technology. The AMD cloud benefited from AMD’s unique capability to swap out processors and achieve in-place upgrades. As a result, AMD IT was able to double AMD’s computing capacity while leveraging its existing infrastructure, resulting in over \$6 million in savings.

In parallel, AMD IT initiated a project to standardize software and versions of software throughout the emerging AMD cloud. The Red Hat Enterprise Linux operating system was selected as the foundation based on its multi-core support and scalability. Following a study of usage patterns and software capabilities, AMD IT also standardized on Platform Computing for workload management and electronic design automation (EDA) software from Cadence, Mentor Graphics, and Synopsys for design and simulation tools.

The AMD Cloud Takes Shape

Throughout the migration, engineering projects continued to execute to tight schedules even as the base capacity of the compute resource was grown at over 20%. Today, the newly-created AMD cloud infrastructure runs on more than 115,000 AMD CPU cores and over 4 Petabytes of storage. It is an integral part of AMD’s ability to deliver the R&D innovation AMD customers expect. All engineering projects at AMD reside in the private cloud.

The AMD Cloud Provides Flexibility and Speed

At the same time, AMD IT is able to respond dynamically to the project needs of engineering teams. Projects are no longer bound by geography due to the location of the data that they require. Resources can be re-prioritized in a matter of hours. To cite one example, a project team working on the Bobcat Project in August of 2010 needed to speed up the design process to adapt to new market needs. The team asked for a Simulation Rush. AMD IT was able to dynamically reallocate capacity from the AMD cloud in one night, and give the project team 42,000 CPU cores (approximately 45% of computing resources) for five days. As a point of comparison, a typical project team normally utilizes only 4% of the cloud. With the added computing power, the project team was able to get two months of testing done in five days. This directly correlated to faster time to market and saved millions of dollars for AMD.

“The AMD cloud project, ‘Compute Anywhere,’ hinged on bringing existing resources together and into the cloud, without disrupting productivity. The single greatest advantage we had was that our infrastructure was AMD Opteron™ processor-based, allowing us to do rapid, in-place upgrades with absolutely minimal downtime.”

Tom Painter

Corporate Vice President, IT, AMD

The AMD Cloud Enhances Manageability

To ensure maximum efficiency, AMD’s Engineering Enablement Group currently oversees the allocation of the cloud resources. Engineering Enablement consists of senior members of AMD’s engineering leadership team, who have a deep understanding of ever-changing customer dynamics and can set project priorities accordingly. The actual allocation of resources is then done at the operational level by the Grid Maintenance Services group within IT. With standardization and flexibility, AMD can better utilize hardware spend and reduce data center costs associated with supporting hardware installations.

The AMD Cloud Provides Wide Access to Resources

With computing resource flexibility optimized, AMD IT deployed thin client technology from Wyse. Thin clients based on AMD Fusion APU technology provide high speed access from remote sites to the AMD cloud.

The AMD Cloud Lowers Storage Costs

Because data is now available everywhere, redundant local copies of data sets are no longer required. AMD is able to leverage available engineers, regardless of their location, to engage on any project as needed, using consolidated data sets. Consolidation also allows AMD to leverage existing, high-end design software that once sat idle in distributed sites.

The AMD Cloud Drives Increased Productivity

Currently, thirty major projects reside in the AMD cloud, including design work related to AMD Accelerated Processing Units (APUs) and the new “Bulldozer” core. Some 3,000 engineers worldwide execute approximately 40 million jobs per month on AMD cloud, which is able to achieve sustained utilization rate, both in terms of the consumption of available servers and in terms of the average workload per core, of greater than 90%.

The AMD Cloud Enables Superior Forecasting

The AMD cloud allows AMD IT to respond to dynamic needs, enabling resources across the globe, and maintain high levels of service delivery. Flexibility and standardization allow any engineering project to be easily moved across the cloud. In close partnership with senior engineering, AMD IT delivers the most efficient and reliable service possible to AMD’s engineering organization. As a result, AMD is able to fuel its extremely compute heavy R&D activities – and do so with the knowledge that it has the resources it needs to embark on any new R&D challenge.

For more information about AMD’s innovative technologies in cloud computing, go to: <http://www.amd.com/cloud>

