

Mike Wolfe's Blog : IT Innovation in the "Cloud"

Designing world-class semiconductors, platforms and cores for our customers and OEM partners around the globe is one of AMD's most complex and demanding engineering tasks. From an IT perspective, it means we have to meet the massive demand of computing power and the need for a flexible infrastructure while finding a way to deliver these resources to our engineers across multiple geographies on an as-needed basis. It's been a daunting challenge but one that we have embraced.

Our motto was "**Compute Anywhere**" - inspiring our IT team to create the high-performance private AMD Cloud, an enterprise-wide grid that now powers all R&D at our company. This multi-year initiative received considerable attention and praise when AMD was recently recognized as one of [CIO Magazine's 100](#) most innovative organizations.

At AMD, we have spent a lot of time thinking about the cloud. Our recent study on "[Adoption, Attitudes and Approaches to Cloud Computing.](#)" revealed that 60% of the organizations that use cloud solutions already see business value. Moreover, the success of our private cloud demonstrates we are [walking our own talk](#) and embracing cloud computing's greatest benefits – **flexibility, increased efficiency, and hardware cost savings**. If your business is thinking about the cloud, you may want to consider what we have been able to accomplish.

More and more companies need to be able to function fully anywhere, anytime. This is particularly essential when your business relies on technology to deliver technology. The AMD Cloud runs on more than 115,000 AMD CPU cores and over 4 PB of storage around the world. It supports thousands of engineers worldwide with approximately 40 million jobs executed in a month. **It is an integral part of how AMD is able to deliver true R&D innovation to our customers.** As IDC's Senior Vice President & Chief Analyst, Frank Gens states: "[this year, the market will recognize "the cloud model as a fundamental part of IT service delivery."](#)

We have stretched our IT muscles through the cloud and are reducing our spending on hardware and data center costs while responding dynamically to the project needs of our engineering teams. By consolidating the compute infrastructure into a logical cloud of a few large geographically diverse physical cloud clusters with common standards, we have achieved a sustained utilization rate of greater than 90%.

Our flexibility has also dramatically expanded what we are able to do. Consider that projects are no longer bound by geography due to the location of the data that they require. And, over half of our storage was once consumed by huge data sets, and had to be replicated between sites. This used to take weeks to carry out. Today, we are leveraging available engineers, regardless of their location, to engage on any project using consolidated data sets.

By way of example, in 2010, we needed to shift massive amounts of compute power to a single project. The IT team dynamically reallocated capacity in one night, giving the project approximately 45% of the cloud - up from their typical consumption of only 4%. Consequently, the project team was able to get two months of testing done in only five days – an outcome simply not possible without the cloud.

So next Monday morning, while you're heading to your office, think about what task you and your team can achieve before Friday ends. Five well spent days can result in a dramatic impact on time to market for you and save millions of dollars.

Additionally, you may be asking other questions, such as, what does *flexibility* mean to your business? How about improved customer service, 24/7? What about being able to sell and deliver products online to a family in Beijing? How should you assess the value of investing in the cloud model?

Our strategic approach to the AMD Cloud has included a strong focus on standardization across platforms resulting in reduced investment costs for equipment upgrades, and basing the cloud entirely on AMD technology. Specifically, our servers based on [AMD Opteron™](#) chips offer a unique capability to swap out processors and achieve in-place upgrades, allowing us to double our capacity without buying a single new server. Over \$6 million in savings were realized with just an in-place upgrade from dual to quad-core chips, demonstrating a direct ROI that comes with the right IT investment.

The AMD Cloud demonstrates daily how IT innovation can be used to create more innovation. We are now better able to respond to dynamic needs, enable resources across the globe, and maintain high levels of service delivery. Our virtualization-ready technology –[cloud clusters of servers equipped with AMD Opteron™ processors, running AMD-Virtualization™ technology](#) – has helped us drive and materialize new ideas that will turn into enhanced solutions for our clients.

Our team at AMD is now able to fuel our extremely compute-heavy R&D activities – and to do so with the knowledge that we have the resources we need to embark on any new research opportunity. We, therefore look at our future with even greater optimism, as we contemplate what we will be able to do with this new greatly expanded capability.

The sky is the limit, right there among the clouds.