The Motley Fool

Why ARM Will Fail in the Data Center

By Steve Heller | More Articles
December 28, 2012 | Comments (3)

Over the last two years, data has seen explosive growth and now 2.5 quintillion bytes are created daily. This number is so large that it represents 2,220 petabytes, where each petabyte represents 1,024 terabytes. It's essentially an incomprehensible amount of data that's created each day. An easier way to understand it is that it's so much data that 90% of the world's data was created in the last two years alone.

Running an ever-increasing army of processors to power and ultimately analyze that data has its fair share of associated costs. Data centers that house this information have become challenged keep a lid on energy costs. According to *The New York Times*, data centers consumed 76 billion kilowatt-hours of electricity in 2010, or about 2% of all U.S. electricity consumption. Since data hasn't stopped growing, it's likely data center energy needs will grow in lockstep. The evidence to support this claim lies within the rise of extreme low-energy servers, which **Intel** (NASDAQ: INTC) estimates to represent 6% to 10% of the server market. **ARM Holdings** (NASDAQ: ARMH) and its proven power-efficient architecture is currently being courted for the job, but many challenges lie ahead before this low-energy dream can become a reality. However, if ARM can succeed, there's a sizable opportunity for it to encroach on Intel's stronghold in the server market.

The rise of the cluster

No, I'm not talking about your favorite cereal, I'm talking about server farms that utilize tens of thousands low-power processors to collectively take on larger tasks by breaking them into byte-sized, manageable pieces. This style of data processing would be like using an ant farm instead of the Incredible Hulk to get the job done. It's compelling because it works really well for cloud-based services like **Facebook**, **Amazon**, and **Google**. Given ARM's current undisputed energy advantage in mobile computing, efforts have been ramped up to develop the same pedigree for server applications.

Viva Barcelona!

The Barcelona Supercomputer Center is building Mont-Blanc, the world's most energy efficient supercomputer by using **NVIDIA's** (NASDAQ: NVDA) Tegra 3 mobile processor over the typical Intel Xeon found in supercomputers. For comparison, the typical Xeon processor consumes 50 to 100 watts of energy, where each Tegra will be tuned to consume just 4 watts. Although the project is using off-the-shelf materials, it's not an off-the-shelf initiative. Practically every piece of code written for Intel x86-based architecture needs to be repurposed in order to take advantage of Tegra 3's ARM-based architecture. For data centers, this is great news, because whatever is learned and developed during the process will be openly shared with the community.

First movers

NVIDIA, **Dell** (NASDAQ: DELL) , and **Hewlett-Packard** (NYSE: HPQ) all have their sights set on developing products for the ARM server revolution. NVIDIA is developing Project Denver, a custom 64-bit ARM chip design for servers, which is important because servers rely on 64-bit code, and this will be one of the first 64-bit ARM designs available sometime in 2014.

Dell has begun shipping its Copper servers to select customers so they can begin testing and validating ARM architecture for 32-bit standards since 64-bit designs won't be shipping until 2014.

Like a good competitor, HP has announced similar efforts under its Project Moonshot, except that it decided to ditch ARM for Intel's newest Atom chip before shipping the goods. That's because the Atom chip delivers 64-bit support, the same power consumption, a broader ecosystem (and user base), and generally less operational challenges. If HP's decision is any indication, there's far greater demand for Intel's out-of-the-box compatibility over the academic promise ARM hopes to deliver.

Challenges and uncertainty

Between exploiting a massive number of low-power processors, sustaining performance, and programming every aspect of the initiative, it's safe to say the challenges are plentiful for ARM to transition into the server space by 2014. But the greatest challenge of them all is Intel. By the time ARM is ready for prime time, Intel's Atom chip will be fabricated on 14-nanometer transistors, essentially closing the power gap between itself and ARM's latest and greatest. At that time, ARM chips are expected to be fabricated on a larger, less power efficient 20-nanometer node.

If low-power server clusters become the de facto standard for cloud-based data centers, Intel has a huge lead on ARM, as nearly everything in the data center has been developed in favor of Intel's architecture – which commands a 90% market share in the server market. When you also consider Intel's fabrication process will be a solid generation ahead of ARM's fabrication partner, **Taiwan Semiconductor**, there won't be much reason left to convince businesses to deploy ARM server farms. Despite the efforts being made, the practicality of implementing ARM-based architecture in the data center doesn't appear to be all that practical.

When it comes to dominating markets, it doesn't get much better than Intel's position in the PC microprocessor arena. However, that market is maturing, and Intel finds itself in a precarious situation longer term if it doesn't find new avenues for growth. In this premium research report on Intel, our analyst runs through all of the key topics investors should understand about the chip giant. Better yet, you'll continue to receive updates for an entire year. Click here now to learn more.

Fool contributor Steve Heller owns shares of Intel and Google. The Motley Fool owns shares of Amazon.com, Facebook, Google, and Intel and has the following options: long JAN 2014 $20.00 calls on Facebook. Motley Fool newsletter services recommend Amazon.com, Facebook, Google, Intel, and NVIDIA. Try any of our Foolish newsletter services free for 30 days. We Fools may not all hold the same opinions, but we all believe that considering a diverse range of insights makes us better investors. The Motley Fool has a disclosure policy.

Comments from our Foolish Readers

On December 30, 2012, at 8:11 AM, applefan1 wrote:

> If the OS and apps are written for which ever processor that proves to be the fastest, lowest cost, most reliable and lowest power consumption, what does it matter if it an ARM, X86, or XYZ?
>
> Should it matter?

On December 30, 2012, at 11:30 PM, pigz wrote:

> Thanks for the analysis, Steve. Take advantage of the current PC bashing to pick up INTC on the cheap.

On December 31, 2012, at 2:44 AM, NOTvuffett wrote:

> If I am not mistaken, the ARMH chips utilize a x86 architecture. They also have the distinct disadvantage of not having their own fab facilities. Building one can cost hundreds of millions of dollars.
>
> On a positive note, I will say that ARMH chips have gained a strong foothold in small devices and especially micro-controllers. I recently read that they will soon sell one in an 8 pin DIP package for only 39 cents in bulk.
>
> I am so tired of hearing the 'death of the PC' bullshiite. They are here to stay. Not so many years ago, AMD had superior chips so I bought those. Now INTC has the best ones at the top end of the spectrum, but things could change. I just wish Intel would quit wasting time trying to develop on-board graphics. Even the crappiest of video cards will beat on-board graphics.
>
> As far as supercomputers go, they are used mostly for modelling the real world. It may be super-sonic flow around a turbine blade, folding of an organic molecule, the physics of the femto-seconds before the detonation of a thermo-nuclear bomb, etc., etc. Most of these tasks mainly require the processing of large numbers vectors and scalars. So I would have to with the dark horse here and suggest NVDA. The computer that I am on right now has an NVDA card. Sure, it sucks in the juice, it has 3 fans to keep it cool, heat-sinks the size of a medium sized rodent... but it also has over 1300 cores.