Open Server Summit 2014
Santa Clara, CA
November 11-13, 2014

M. R. Pamidi, Ph. D.
Editor-in-Chief
IT Newswire

Executive Summary

Open Server and open 'everything'—software, storage, networking—are transforming the IT landscape. Smart vendors that are tactfully accepting the change will not only embrace it, but affect it. Others will go through Elisabeth Kübler-Ross's five stages of grief denial, anger, bargaining, depression, and acceptance, but it may be too late to guarantee their survival. This event covered more than just open servers and included storage, software, and networking, and probably should have more aptly been called *Open IT Summit*.

Computing and Storage

Data—structured and unstructured—is exploding, as you can see in the following graphic.

Figure 1. Data Explosion
Source: IBM

Sure, there is **data explosion**, but what we need is **knowledge deduction**. For instance, according to IBM, a four-engine, Boeing 747 flying between New York's JFK and London's Heathrow on an average collects about 350 TB of data about weather, altitude, fuel consumption, and crucial aspects of the aircraft's performance. This will get even more intense once GE, for instance, with its Industrial Internet initiative starts instrumenting **every one** of its jet-engine blades to monitor for possible failures—fatigue, thermal, or other structural problems.

The question arises: Should these massive amounts of crucial **data** be (transmitted possibly compressed) **close to computing** (on the ground) for analysis or should **computing** be **close to data** (for analysis in flight) for a quick response? The latter requires massive amounts of data-storage AND computing power in the aircraft, with the associated problems of weight, heating, and cooling. The crew of the fatal Asiana Airline Flight 214 between Incheon International Airport in Seoul and San Francisco in July 2013 had 1.2 seconds to

make a critical decision. Surely, one couldn't send terabytes of data to the ground, have it analyzed, and sent the results back to the aircraft in under 2 seconds.

There appears to be a battle between systems (server) vendors who want storage close to computing and storage vendors that desire computing close to storage. So, systems vendors are pushing SSD (still expensive, but prices falling rapidly) and flash, whereas storage vendors are bringing computing close to storage. For instance, EMC OEMs Lenovo servers (often called 'compute-storage servers') and earlier this acquired Andy Bechtolsheim's stealth startup DSSD developing technology for pooling server-based flash for high-performance data access. Similarly, Cisco acquired Whiptail (now christened Invicta) for its flash storage technology to bring storage close to computing.

So, should computing be close to storage or vice versa?

- It's cheaper to move computation to data.
- It takes 100X more energy to move a block of data than to perform computation on it.
- A computation requested by an application is much more efficient if it is executed near the data it operates on:
  - Minimizes network traffic
  - Increases throughput and performance of the system
- In-situ processing can significantly accelerate certain applications. It can also greatly increase energy efficiency at the system level.
- In-situ processing is a revolutionary approach that can fuel data analytics innovation and enables the seamless use of computational resources. Maybe the preceding two points is what driving Cisco to *Fog Computing*.
- In a 2009 ground-breaking paper on Cloud Computing, authors from UC Berkeley showed that it is faster to overnight a modern hard drive than it is to transfer its content over a network!
- *The best I/O is the one that doesn't happen.*

Big Data Explosion

That Big Data exploding is an understatement. By 2020, over 30 billion connected devices will be exchanging more than 44 Zettabytes of data.

- How do we manage this massive growth?
- Are the current hardware, software, and network architectures adequate to meet the new demands?

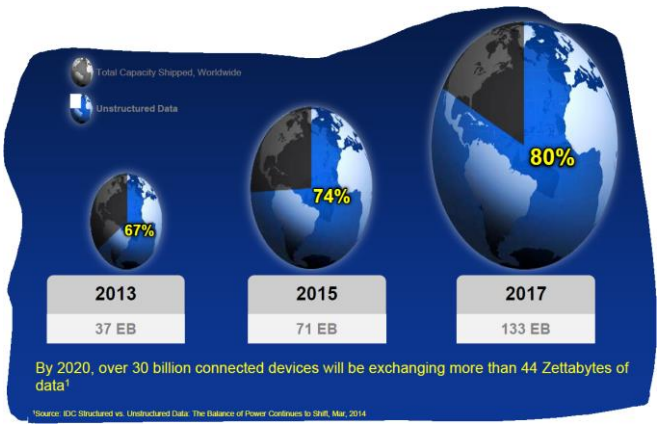

Figure 2. IoT and Data Explosion

Figure 3. Big Data Explosion
Source: IDC

StrongARM-ing the Industry

ARM was a fairly low-key vendor until a couple of years ago, supplying mostly to Apple's products, but the company has come on very strong in data centers and reportedly shipped over 10 billion processors in 2013. ARM is not going to displace AMD or Intel anytime soon, but is making a major push in to enterprise data centers, just as *n*Vidia is moving into the graphics-intensive and HPC spaces in data centers. ARM is becoming more relevant than AMD which is struggling.

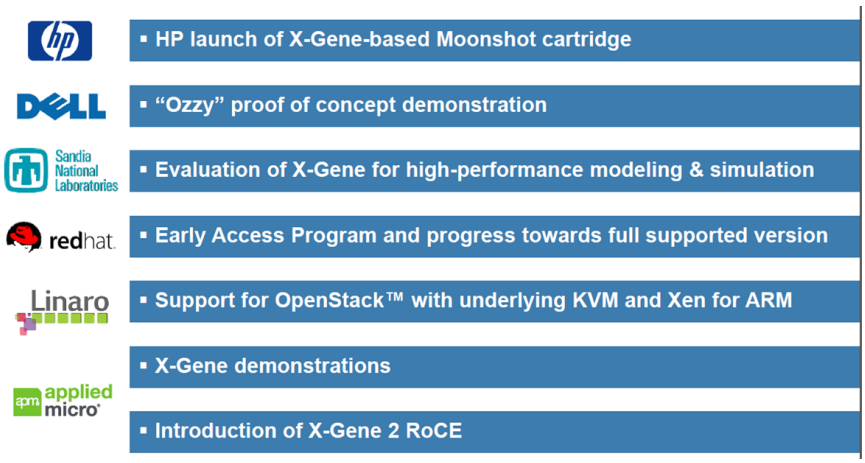

Figure 4. Vendors adopting ARM's X-Gene

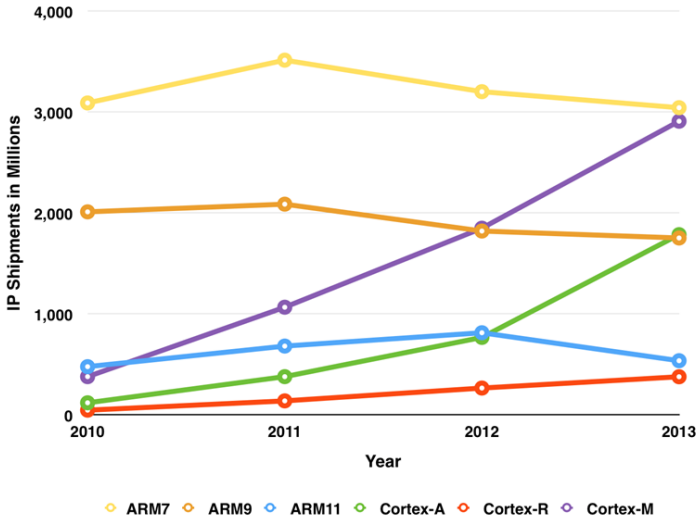

Figure 4. ARM Shipments by Year

Hyper-scale Mega Data Centers

There were some interesting discussions about Facebook's Open Compute Project (OCP), white-box servers, and open everything. Even early skeptics and critics of the OCP have now embraced it. But, is OCP the cure-all for IT's headaches? We don't believe so. Sure, OCP delivers cheap...er, inexpensive COTS components with compelling CapEx and OpEx. If a server fails, it is just disposed of and replaced by another cheap one. As one panelist noted, "What's Google's HA (high-availability) solution? Discard the failed server and replace it with another inexpensive box." Surely, if you are doing high-frequency, sub-millisecond transactions on Wall Street, you aren't going to bet on white boxes imported from Red countries. Although Facebook drove the OCP initiative, even Mark Zuckerberg is quoted to have said (and we are paraphrasing), "I want to enable people, and not computers, to communicate." Also, are major enterprises such as banks and financial institutions going to embrace OCP and open-source everything? A CIO of one of the Top 5 U. S. banks told us, "I want to run a bank, not manage servers."

What is even more interesting is that the same Asian companies that are manufacturing for U. S. companies are also building white boxes and selling them cheaper to the mega data centers. We expect in three years the leading PC (and possibly server) vendors will be all Asian—Lenovo, Acer, and Asus. Lenovo is already the world's leading PC vendor and just acquired most of IBM's hardware, except the P series, and what is left of its mainframe business. What are U. S. vendors doing about it?

Conclusions

Open 'everything'—servers, software, networking, and storage—is revolutionizing the IT industry and has traditional vendors on the ropes. Server hardware has already been commoditized and it may soon spread to networking gear. Sadly, most established vendors look to their **left, right**, and **above** for competitors, but will have their lunch eaten by hungry, nimble vendors **below**. (For a related fascinating lecture by Clayton Christensen visit http://nicolaiwadstrom.com/blog/2011/12/05/reinventing-it-clayton-m-christensen-gartner-symposium-2011.) Software and Services are still safe, with gross margins in the ~80% and 60% range, respectively. But open source software will threaten software margins and most companies won't easily morph to Services vendors. It took IBM over 20 years to do so and now Big Blue seems to be transforming into the next-unknown avatar.