Wood Harbinger

3009 112th Ave NE, Suite 100
Bellevue, WA 98004
425.628.6000
woodharbinger.com

Key Trends Impacting Data Center Design

What to Expect as Demand for Data Grows

By Tom Leonidas, Jr., P.E.

The next time you "like" your friend's new band or comment on a photo of your baby nephew's first steps, consider this: Facebook is now used by one in every 13 people on earth. More than 700 billion minutes a month are spent on Facebook, 20 million applications are installed per day and more than 250 million people interact on Facebook from outside the official website on a monthly basis, across 2 million websites. In a 20 minute timeframe, over 1 million links are shared, 2 million friend requests are accepted and some 3 million messages are sent.

Not only is Facebook using up more and more of your free time, it is requiring an unprecedented demand for data. And Facebook isn't the only data-hungry technology in town. These are staggering statistics from only one of many lenses into the growing world of information technology. The increasing demand for data continues to grow at staggering rate, with worldwide IP traffic expected to quadruple by the year 2015. The impacts on the data centers that must support the exponential growth of information are immense. Whether it be social media, educational, financial or healthcare; the need to deliver more data at a faster rate will continue to grow at a staggering pace.

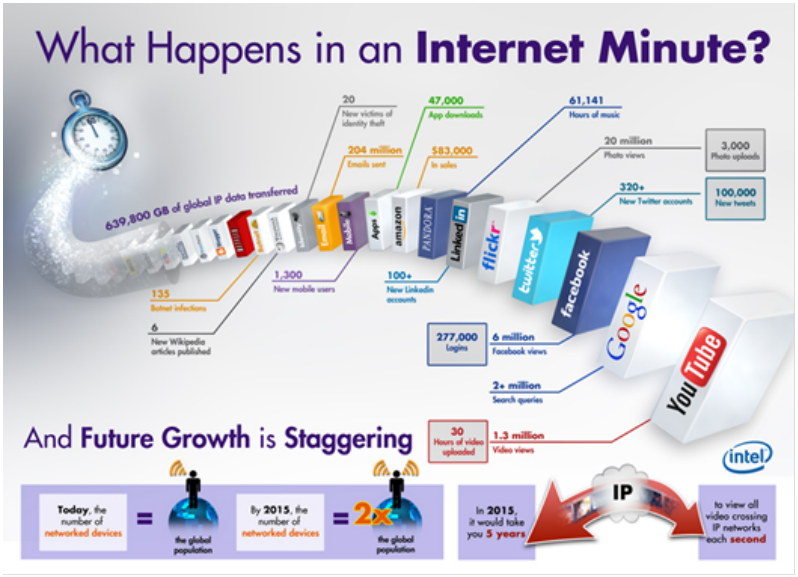

Figure 1 – An Internet Minute *(source: Intel)*

There are several key trends and drivers for this growth and those trends impact planning for new and existing data centers. Among them:

Mobile Computing

The introduction of the Apple iPad made the tablet computer a highly available, low-cost handheld computing device. The tablet has become an approved computing device in most business industries, from finance to healthcare. In addition to its convenience, tablets allow the user access to huge amounts of information at their fingertips; anytime, anywhere, anyplace.

The growing popularity of the tablet means that more users are moving information in and out of their hands. Data centers that support the storage and transmission of this data must increase in terms of capacity and performance in order to deliver the experience expected by users. This will continue to create a need for data center capacity; in many cases within limitations of physical space. This will drive innovation in terms of virtualization, which in turn will drive design toward highly energy efficient data centers to balance the increased power and cooling needed to move more data.

Virtualization

Server virtualization is the practice of enabling more computing ability out of a single physical server. This is done by using one server as a "host" for many "virtual" servers that ride on top of the host server. This means a single physical server hardware platform can accommodate a large number of virtual servers and, in turn, applications. While the quantity of the virtual servers depends on the physical memory and processor power of the host server, it is not uncommon to have 30 or more virtual server instances on one physical server. This is quite different from the abilities of servers in the past, which were capable of running just a single application and the background computing services to support that application.

While virtualization reduces the number of physical host servers needed, it does not decrease the amount of power needed in a data center. The impact of virtualization on data centers is to dramatically increase the amount of power density per server rack to power more and smaller physical host servers. Since these host servers are now smaller - typically 1U or 2U in size - each rack can accommodate a higher density of physical servers. The end result is a need for higher amounts of power. This is a trend that will continue, along with a need for creative data center cooling solutions to address higher density heat loads.

Expect to see more virtual desktop computers crop up in the workplace. Virtual desktops are powerful in that they can be easily customized and easily deployed using a standardized and stored virtual hardware platform. Users only need a

network or internet connection for their own personal virtual desktop to follow them anywhere they go, even on their tablet device. The compelling reason for virtual desktops is standardization, quick deployment and lower IT costs, along with user ability to have ubiquitous access to their desktop.

Virtualization allows data centers to get more use out of the existing data center footprint rather than investing significant dollars to build more data center space. It also can allow for a more efficient data center operation with the proper cooling strategies in place.

IT Power Consumption

This increased need for power is also driving an effort to create sustainable data centers in terms of energy. Data center operators are constantly comparing a data center's infrastructure to its existing IT load and trying to achieve the best energy output for the lowest cost. As data centers expand vertically and become denser, expect the trend toward high efficiency data center design to continue. Cooling strategies that focus at the rack level rather than room level are, and will, continue to be the best practice model as they more efficiently bring the cooling to where it is needed. Innovative strategies such as those that take advantage of outside air (where the climate conditions are appropriate) will be used to cool the data center using less power.

Savvy data center managers are also exploring chargeback rate structures in their quest to lower energy costs. Hosted data center providers as well as private organizations that act as data center hosts have, and will, continue to implement tiered rate structures that create higher rates for higher per rack power density usage. These rate structures will force data center users to use the most efficient hardware to lower their power consumption and decrease their monthly IT costs.

Patterns and Analytics

There is expected to be an 800 percent growth in data over the next five years with 80 percent of that growth in the form of unstructured data. Accompanying this growth will be an increase in measuring patterns and analytics on how data is stored, accessed and distributed. The data collected will be dynamic and will be used to optimize the computing environment and reduce energy. Detailed and dynamic data on power consumption at the rack and server level can be measured using sophisticated Data Center Infrastructure Management (DCIM) tools. These will continue to evolve into Data Center Predictive Modeling (DCPM) tools that data center operators can use to predict how best to utilize and

optimize the data center and model their IT assets to the most energy efficient design; as well as allow them to change it dynamically to adjust and scale with growth.

In line with the sharp growth in the analysis of data handling, expect an increased focus on optimizing storage utilization with an eye toward reducing physical floor space as well as energy usage. Solid state drives (SSDs) will play center stage here as larger capacity drives come available at a lower cost. As the data growth continues, so too will the need for highly reliable data recovery strategies to improve business continuity. Key technologies that will be important in addition to SSDs are data de-duplication and automated tiering and indexing of data for faster access.

Fabric Computing

The growth of data will also require faster processing power and faster networks to deliver that data at high speeds. Fabric computing is a consolidated, high-performance computing platform created from loosely coupled processors, storage and networking; linked by high bandwidth interconnects to make up a grid that looks like a weave of fabric when looked at collectively.

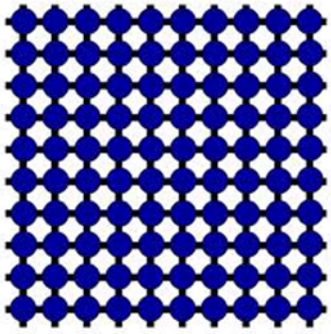

Figure 2 – Connected Fabric Computing Nodes

The fundamental components of the fabric system are the nodes and the links between nodes that together create a unified computing system. Fabric computing is projected to grow exponentially and will become the very foundation of the next generation IT infrastructure for data centers and the organizations they serve. Expect to hear more about fabric data centers that are comprised of highly efficient processing power connected with core-to-edge networking and high bandwidth.

Cloud Computing

Cloud computing is the use of computing resources consisting of hardware and software that is delivered as a service over the internet. Cloud computing allows users to “rent” computing infrastructure from cloud server providers such as Amazon, Google and Microsoft, where they can create virtual servers and virtual desktop computers.

The compelling business strategy for using cloud services comes down to economics and ability to have applications and documents follow users anywhere and on any device, especially mobile devices such as tablets. The economics play is that organizations don't have the capital or labor investment expense in maintaining their own on-site computing infrastructure. The ability of ubiquitous access to information at a low cost makes cloud computing compelling for end users.

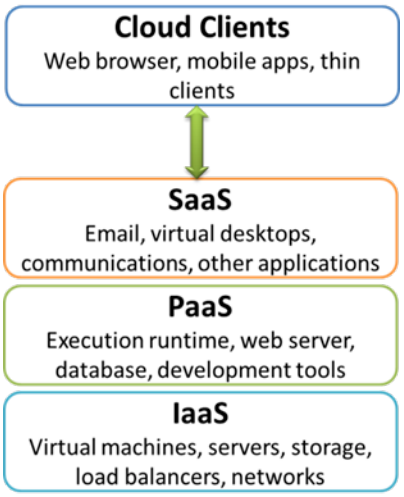

Figure 2 – Cloud Computing Elements

Information technology research and advisory firm Gartner estimates that, over the course of the next five years, enterprises will spend $112 billion cumulatively on cloud-related software, platforms and infrastructure. Worldwide cloud services revenue is expected to top $148 billion by the year 2015. In the healthcare industry alone, cloud services revenue is expected to grow by 20 percent per year over the next five years, making it a $5.7 billion industry by the year 2017.

One of the key issues that has held organizations back on full use of cloud computing is related to security. For this reason, organizations should keep their most sensitive and critical use data on-site and off the cloud. Expect firms to continue to use cloud computing with hybrid solutions where their most sensitive or critical data and applications live on site, but less sensitive data and applications live in the cloud.

Social Networking and IT

Social networking sites account for more than 25 percent of the most active web URLs. In 2011, there were 155 million daily "tweets" versus 55 million in 2010. Social networking is growing and here to stay and is, along with instant messaging and text messaging, displacing email communication.

Companies are using social networking to collect opinions that drive their product and service innovations. Employees walk in the door and bring their own networks with them, creating a fabric of social networks together to bring clients and like-minded individuals closer together. The impact on IT and the data center is that we are headed toward more unified communication environments where social networking, instant messaging, video messaging, voice mail and email are converged together into a common platform. There are

several negatives that come with social networking. First is security and identity theft potential. Company networks will need to be more robust in terms of capacity as well as security. Second is the impact on the social work environment. Social networking can become an addiction and it can be a platform for employees to provide negative information about their employer or their services. Companies cannot stop social networking, but they can and should manage it with prudent policies regarding the use of the company's IT infrastructure.

As demand for data in the hands of users grows at an exponential pace, so does the demand for higher capacity data centers. Data center growth is a challenge in terms of physical space and power needs. Much of the growth drivers in the IT industry such as social networking are not profitable business models, but nonetheless must be accommodated. This will drive innovative strategies to first maximize the available space by expanding vertically using virtualization.

Organizations will look to lower their IT costs by taking advantage of cloud space computing that is created by connected nodes of fabric data centers that integrate processing power, end-to-end networking and high bandwidth in a single provider that provides shared resources at a lower cost model than can be achieved internally.

The ultimate goal? Driving down energy cost, deployment and maintenance cost and providing highly accessible, highly secure data to users in a growing mobile computing environment on a 24/7 basis.

Tom Leonidas, Jr., P.E. is Executive Vice President of the Bellevue, Washington based consulting firm, Wood Harbinger.