Death of the Data Center? Not According to the Analysts' Cloud Forecasts

Prepared by:

Jake Ring

Co-Founder and CEO

GIGA Data Centers LLC

1-800-301-7436 | www.gigadatacenters.com

Death of the Data Center? Not According to the Analysts' Cloud Forecasts

The spectacular growth of the public cloud continues and journalists constantly report that the number of enterprise data centers up for sale is greater than ever. So, are we really witnessing the death of the data center? Perhaps not and here's why.

As companies move workloads into public clouds, they may seek to sell their enterprise data centers and the spike in the sale-leaseback market is a good indicator that even older data centers are still valuable. The buyers in this market are colocation businesses and their actions underscore what leading industry analysts are predicting, e.g. we need more data centers and aren't building enough of them to satisfy current needs. Consider these market growth forecasts to support this assertion:

- IDC predicts global public cloud revenue growing from $180B in 2018 to $277B by 2021, or $97B.
- Forrester predicts global public cloud revenue growing 22% CAGR from $178B in 2018 to $323B by 2021, or $145B.
- Gartner's cloud prediction is $176B in 2018 to $278B by 2021, or $102B.

The cloud, whether public or private, is housed in data centers and these facilities must continue to evolve if they are going to support the aforementioned predictions. It's also important to note that evolution, in and of itself, does not ensure survival for all technologies. For data centers to properly evolve so that all types of organizations can benefit, they must also become more cost-effective and efficient to operate, be more flexible to support higher power configurations (such as hyper-converged), and scale faster.

There is no disputing that the adoption of public and private clouds as well as colocation hosting, are all growing at enormous rates—especially when it comes to large public cloud providers such as AWS and Microsoft. Case-in-point: AWS revenues have grown at a CAGR of over 60% since 2010, with the upward growth continuing through 2018. AWS revenues increased by over 48% through the first half of last year and by 49% year-over-year in Q2-2018. Revenue surged from Q3-2017 total of $3.66B to $6.11B in Q2-2018. Multiply that by four, and you get an annualized run rate of $24.4B, or $10B more than the prior year.

According to Microsoft's FY18 Q4 report, server and cloud services revenue increased $4.5B or 21%, driven by Azure and server products licensed on-premises revenue growth. Azure revenue growth was 91%, due to higher infrastructure-as-a-service and platform-as-a-service consumption-based and per-user-based services.

Server products licensed on-premises revenue increased 5%, mainly due to a higher mix of premium licenses for Windows Server and Microsoft SQL Server. In addition, Microsoft's operating expenses increased by $683M or 7%, driven by investments in commercial sales capacity and cloud engineering. Any way you slice it, Amazon, Microsoft and other providers are seeing enormous public cloud growth.

What does all this enormous cloud growth mean for infrastructure providers? One prominent analyst predicts the "death" of the enterprise-owned data center, as workloads migrate into hyper-scale data centers owned by the cloud providers—but they can't build enough capacity fast enough. Also, not all workloads are meant for a public cloud of shared resources, and some enterprises have investments in IT hardware that are strategically valuable for them to own and manage. These organizations just need to rid themselves of the cost and burden of managing their own data center facility - the prominent analyst now points to these facilities as competitive liabilities. So, the plausible answer to support the growth of cloud lies with the colocation market.

Privately hosted clouds represent a larger revenue amount than their public cloud counterparts, and the private cloud, almost by definition, within enterprise or colocation data centers—not in the shared resources that reside in hyper-scale facilities. 451 Research estimates that in 2018 private cloud revenue was about 43% greater than the public cloud revenue and growing only 2% less quickly. Also, while the hyper-scale companies continue to build sites, they don't have enough capacity to support this predicted growth on their own and have been increasingly turning to colocation providers for help.

This predictive model is great news to the colocation market and its supplier OEMs that provide mission-critical products. But just as the hyper-scale facilities need to adhere to the four best-practice principles (efficient, cost-effective to operate, more flexible to support higher or different power configurations, and scale faster), the same holds true for their colocation counterparts.

But how much capacity is needed if the analysts' growth predictions hold true? Using a few "back-of-the-envelope" calculations, we can begin to see how much data center infrastructure capacity growth may be needed. To find out, let's make a few basic assumptions:

- Assume a standard rack supplied at 8kW is used to support the servers & storage for hosting a cloud environment.
- Assume that each rack cabinet operating at that full load of 8kW, is capable of delivering $45,000 in monthly recurring revenue.

Based on these assumptions, we can calculate the incremental cloud revenue growth and what it can mean for the infrastructure needed to support the expansion.

If we look at the median forecast and use Gartner's incremental growth figures of $102B over the 36 months between 2018 and 2021, and use the aforementioned assumptions, it translates to new incremental revenue of $2.83B per month. And adding our assumption of $45K in MRC per rack cabinet, this translates to just over 62,888 new racks of equipment that will be needed each month. At 8kW per rack, this equates to 503,104 kW or 503.1 megawatts of new data center infrastructure globally, each month. The infrastructure support can come from new facilities or retrofitted data centers, but it would be an incremental infrastructure necessity to some degree.

A common figure in the industry is that it takes an average of $8-$10M to build a megawatt of capacity. But let's assume costs have been improving, so it only requires $8M in capital expenses per megawatt. This means it will require over $4.0B in capital spending—each month—to build the required global capacity to support the forecasted revenue growth of the cloud.

It's also worth noting IDC's estimate that 60% of the global cloud capacity is built in the United States. So, by extension, this means that approximately $2.42B in CapEx is needed—each month—to add the incremental infrastructure of over 300 MW per month required in the U.S. to support the growth of cloud systems. Is this real or are forecasters being over-exuberant?

Based on the reported data—it appears to be very real—since each analyst firm revised their 2018 estimates to be higher than their initial forecasts. How will the colocation industry be able to support this growth? Even if we take the lower estimate from IDC, it indicates that $2.3B will be needed—each month—to build the incremental infrastructure to support the cloud growth in the U.S. Cut the estimate in half, and the prediction is still over 140 MW per month! However, if we can construct these facilities in a quick and cost-effective manner, the needed data center support can be realized. Yet how to do it?

Looking at hyper-scale providers, one consistent method that is used by each is the prefabricated modular approach. Just like car manufacturers want their assembly plants to be efficient, cost-effective, and support high-quality processes, hyper-scalers use modular designs in their data centers as it delivers those benefits of lower cost, higher efficiency, and better flexibility and scalability.

The Department of Energy's National Energy Technology Lab (DOE-NETL) HPC data center build provides another solid example of a present-day, cost-effective facility and also gives us some necessary guidance toward achieving the goal of delivering hyper-scale performance to a broader market.

The DOE-NETL facility is a modular 1 MW data center that can flexibly support up to 35kW per rack. The facility uses adiabatic evaporative cooling that delivers the required ASHRAE standard humidity and temperature ranges without the higher cost of a chilled water system and raised floor with CRAC units. This reduces the cost of the build to under $6M per megawatt, yet the data center has been averaging a PUE of 1.06 since it was commissioned back in 2012. The manager estimates the higher efficiency has delivered an operational cost savings of $450,000 per year. Even better, it was deployed on site in two weeks to help realize those savings faster.

Regardless of which forecast you use, a huge capacity of colocation data centers is needed over the next three years. In order to ensure the infrastructure growth is capable of meeting the analysts' predicted requirements for *all* workload processing needs, future data center retrofits and new builds *must* be low-cost, highly efficient, support higher power densities and have the ability to rapidly scale. Modular designs and methods, such as used at the DOE-NETL and many hyper-scale facilities, can deliver colocation data centers that are capable of supporting these requirements. And most importantly, these new facilities can be deployed quickly! So as you evaluate your next colocation provider, explore how they stack up compared to a data center using modular methods. The DOE-NETL figures do sound like data center Nirvana, but I assure you, they are indeed real and being replicated right now. Please visit GIGA Data Centers to learn more.

###

About GIGA Data Centers

Headquartered in Atlanta, Georgia, GIGA Data Centers (GIGA) takes a modern approach to building & operating data centers with innovative, modular technology that is proven to be more efficient than the decades-old tradition of raised floors facilities with over-provisioned cooling systems. GIGA is motivated by the belief that all companies deserve colocation data centers that provide hyper-scale efficiency and flexibility at a reasonable price. The result is a guaranteed PUE of 1.15, flexible rack power densities from 5kW to 50kW per rack-cabinet, and an immensely scalable design delivered at very competitive rates.