

Dramatic headlines on Lyft's AWS spending miss the point

Analysts - Owen Rogers

Publication date: Wednesday, March 6 2019

Last week, ride-hailing app Lyft submitted its S-1 filing to the US SEC as it begins the path to an IPO. Nestled among the details was a statement regarding its financial commitment to Amazon Web Services: between 2019 and 2021, Lyft is committed to pay AWS \$300m, about \$8.5m per month. In fact, this commitment has grown – Lyft extended the deal from \$150m in January. Those claiming this is too expensive are missing the bigger picture.

The 451 Take

Yes, Lyft is spending a lot on AWS, but this certainly doesn't mean cloud is expensive. For a start, just 5% of Lyft's revenue is being spent on AWS – for a company that makes its money from the reliable and performant operation of its ride-sharing app, this really isn't much. Perhaps it could have built its own huge private cloud of virtual machines at a lower price point than public cloud, but could it have built (and managed) a reliable infrastructure of serverless computing, storage, database, containers, streaming and other technologies for cheaper? We think not – especially when you factor in the capacity-planning challenge faced by Lyft, when the demand for its ride-sharing services is spontaneous and without reservation.

Some commentators claim Lyft's AWS outlay is evidence that cloud is poor value. But any assessment of value without a comparison figure or context of what that figure delivers makes little sense. If I spent a billion dollars on diamonds, there is no question that I've spent a significant amount of money. But if, for that billion dollars, I received a billion diamonds, few would argue that wasn't an incredibly good value.

In Lyft's S-1 filing, it states its net revenue in 2018 was \$2.2bn, about \$183m per month. This excludes the revenue it collects and returns to drivers (total bookings, which includes this figure, is \$8.1bn, which puts its average commission at 27%). This means it spends just $\$8.5m/\$183m = 5\%$ of its revenue on AWS.

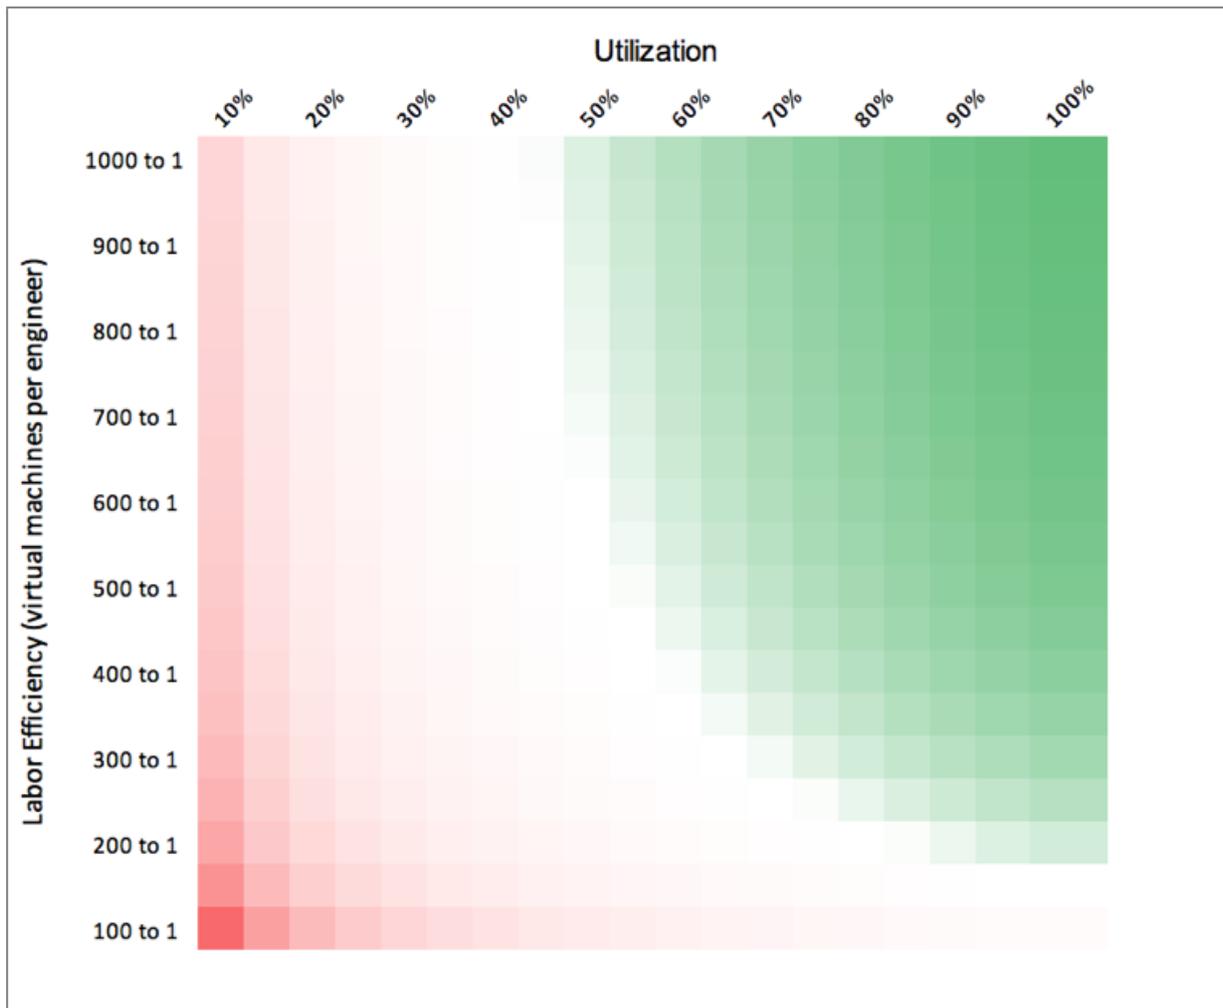
Now let's remember, Lyft is primarily a tech company. It owns no transportation assets; its only purpose is to deliver an application that puts passengers in touch with drivers (although it has recently expanded into scooter and bike rental.) If an IT department at a bank spent 5% of its budget on infrastructure, no one would bat an eyelid. For a ride-share app that wholly depends on

infrastructure for its revenue, spending just 5% of its budget on infrastructure is pretty reasonable. In a recent survey of 150 IT and procurement decision-makers conducted by 451 Research on behalf of Cloudability, 55% were spending more than 15% of their IT budgets on cloud.

Another claimed issue is that Lyft has committed to pay \$300m to AWS regardless of whether it consumes that value in resources. But this is how Lyft can squeeze its costs: we anecdotally hear of savings of 7-10% on the AWS Enterprise program for spending over \$1m per month. It's fair to assume that Lyft made far greater savings for its much larger commitment. But yes, \$300m is a lot of money – at Cloud Price Index benchmarks, it is equivalent of 250,000 virtual machines running for a year or the storage of an MP3 since Homo sapiens first wandered the Earth 300,000 years ago. But the fact that Lyft is willing to make such a massive commitment to AWS is a testament to its belief in the cloud provider. We expect that Lyft's spending over this period will be far greater than the commitment – it is unlikely Lyft would have made such a commitment if it didn't think it would easily beat it; on-demand provisioning means it can easily consume above this commitment on demand, albeit at a higher rate.

The question many have been asking is could Lyft have done it cheaper by hosting it on-premises? We think possibly, but probably not. The chart below shows where public and private cloud are the cheapest option, based on standard on-demand pricing, for different configurations of labor efficiency and server utilization (green reflects configurations where private cloud is cheaper; red shows where public cloud is cheaper). Obviously, this chart would be different for Lyft considering its ability to negotiate a lower price with cloud providers, but the general shape holds true – for Lyft to beat public cloud on price, it would need to have a high level of labor efficiency and utilization on its own private cloud.

Figure 1



Source: 451 Research, LLC

The question is, could Lyft beat these thresholds? We think it is possible that Lyft would be able to build and manage infrastructure at better than 1,000 virtual machines per engineer, as shown in the above diagram, and thus be in the 'goldilocks zone' where private cloud beats public cloud. But Lyft is a 'born in the cloud' company. It's unlikely to use just virtual machines. It probably uses serverless computing, different types of storage, NoSQL databases, and much newer, more scalable technology that it would also need to be executing at very high labor efficiencies, with expensive (and difficult to find) labor. Thus, the economic argument when you factor in technologies beyond virtual machines is certainly not clear cut. What we're saying is that Lyft wouldn't just have to be awesome at managing virtual machines – it would have to be expert at all the other underlying technologies it consumes from AWS.

Second, as a cloud-native company, it is highly likely Lyft's application is scalable to address variations in demand. Lyft is unlikely to have many applications that have constant capacity – the whole point of a ride-sharing app is that people can use it on demand, without booking. During sports events, for example, demand might spike hugely – how many people needed cabs at the end of the Superbowl in the US last month? How many British rugby fans needed taxis after Wales beat England last week (after all, Lyft has global aspirations)? This creates a capacity challenge. If Lyft has too much capacity to meet peak demand, then for much of that time, the capacity is wasted and sunk cost. Too little capacity, and performance and availability will suffer. As a result, using a public cloud means Lyft is not responsible for this capacity issue – it has outsourced it to AWS.

Dramatic headlines on Lyft's AWS spending miss the point

If you had \$300m to spend on cloud, would you blow it all on a hunch, or do your sums? We can be pretty certain that Lyft has analyzed in great detail the costs of the various options for its specific use case, as well as the risks of each, and has made an informed, strategic decision to commit to AWS.