

Twitter Thread by Corey Quinn



Corey Quinn

@QuinnyPig



So there's a lot of confusion about what Parler being kicked off of Amazon Web Services (@awscloud) means. Let me do a quick thread to explain it to folks who aren't deep in the technical weeds...

You use an app (in this case, Parler). There's a web site you can use, and apps you can download onto your phone via Apple and Google. Those three versions of the app all talk to servers (big computers) behind the scenes.

In the Olden Days, getting those servers took months. Then you had to sign deals with companies to host those servers--they take massive amounts of power, they run hot so air conditioning is a big deal, and they need a *lot* of bandwidth.

About 15 years ago or so, Amazon (the online bookstore) had a revelation and started AWS. The idea being that instead of needing tens of thousands of dollars to get those servers built up, and spending months on it, that they would rent you these servers by the hour.

"A credit card and a few clicks of a mouse later" you were basically able to spin up huge numbers of these servers.

This powers Netflix! This is what Capital One does! It's not just random social media apps. Today AWS sees around \$50 billion a year in revenue. It's a real business.

Getting booted off of AWS is virtually unheard of; when people leave intentionally the planning takes months; execution can take years. It's a lot harder than you think for a few reasons.

It takes time to move all of that data (there's a reason downloads take a while), but that's just the beginning.

AWS doesn't just sell "big empty computers." They offer higher level services. Preconfigured databases, automatic video streaming, etc.

These services aren't really directly compatible with other companies' offerings. Making what you built on AWS's systems work elsewhere is super challenging.

Parler claims they didn't use these higher level services. Taking that as true, there are still problems. The way AWS's services work--how you create them, how long that creation process takes, how you get data onto them? They behave differently in AWS's world than elsewhere.

A lot of assumptions about how the servers behave are "baked in" to how Parler (and any AWS hosted application) are built. A lot of companies don't realize that those assumptions are there until they try to move.

That's why migrations take months / years. Parler has 30 hours.

A lot of folks in the cloud computing world make noises about this, about why "deploying on multiple clouds at once" is imperative. I have and do disagree with this notion--it takes a LOT of work and time that could be used to make the product better.

But if you're going to get booted off of basically every cloud provider out there because you can't follow their very, very, very permissive rules? It's no longer optional.

I will now field questions!

"What other providers will tolerate this?"

None of the household names. *Maybe* AliCloud? Anyone who tries for it is going to inevitably end up with a lot of FBI inquiries and a pile of bad press.

And they'll of course be blamed for any Parler outages. <https://t.co/CEOzSAQG6u>

Where can they even go? We can joke about Oracle or IBM, but they have to be persona non grata at this point, yeah?

— Andrew Nelson (@holophrastic_an) [January 10, 2021](#)

"Containers" is an argument against this stuff. It's a bit technical, but the way those containers all talk to each other? How the servers underneath the containers get built? How the networking is configured? How they find one another? There are provider-specific assumptions. <https://t.co/yYrbG5C90R>

What if they are highly containerized? Isn't that supposed to facilitate hybrid clouds where all your eggs aren't in one cloud provider's basket?

— Kurt Hoyt (@kurtinatlanta) [January 10, 2021](#)