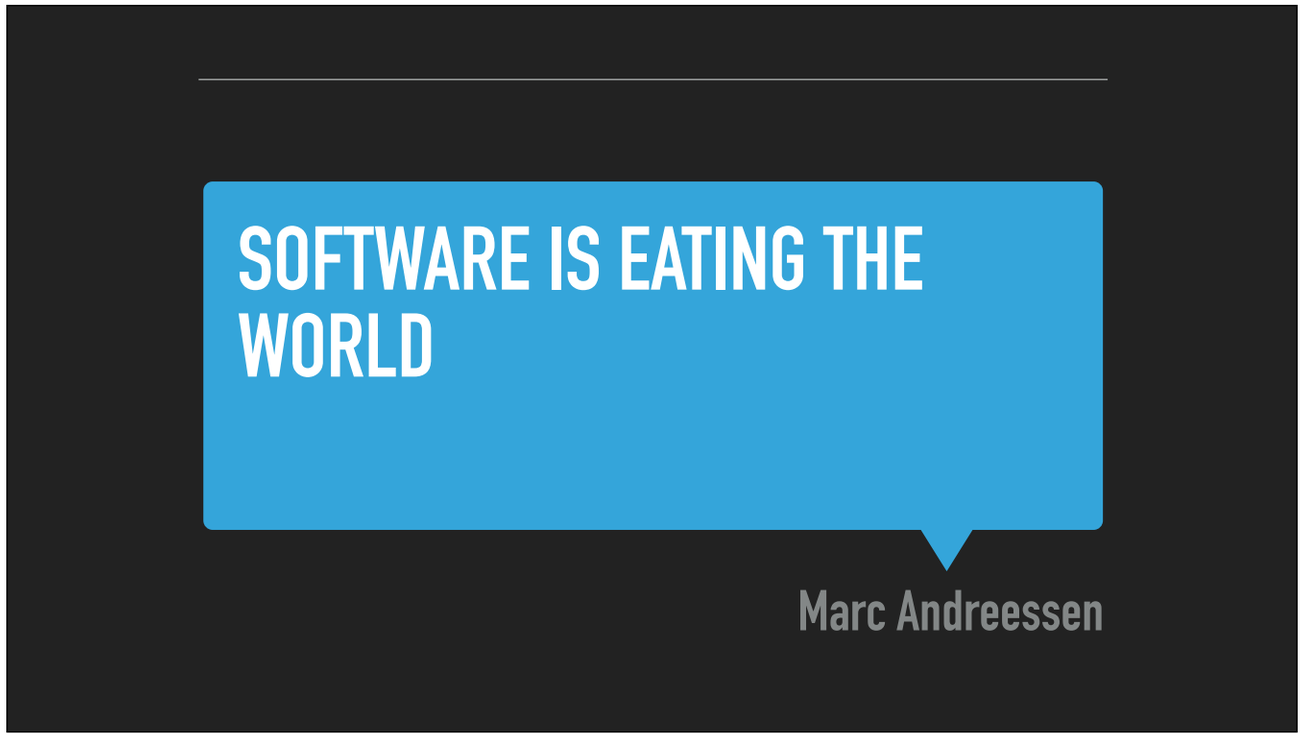

Software is eating the world.

Software is so pervasive that it is even turning some unexpected names into software shops. Let’s create our app!

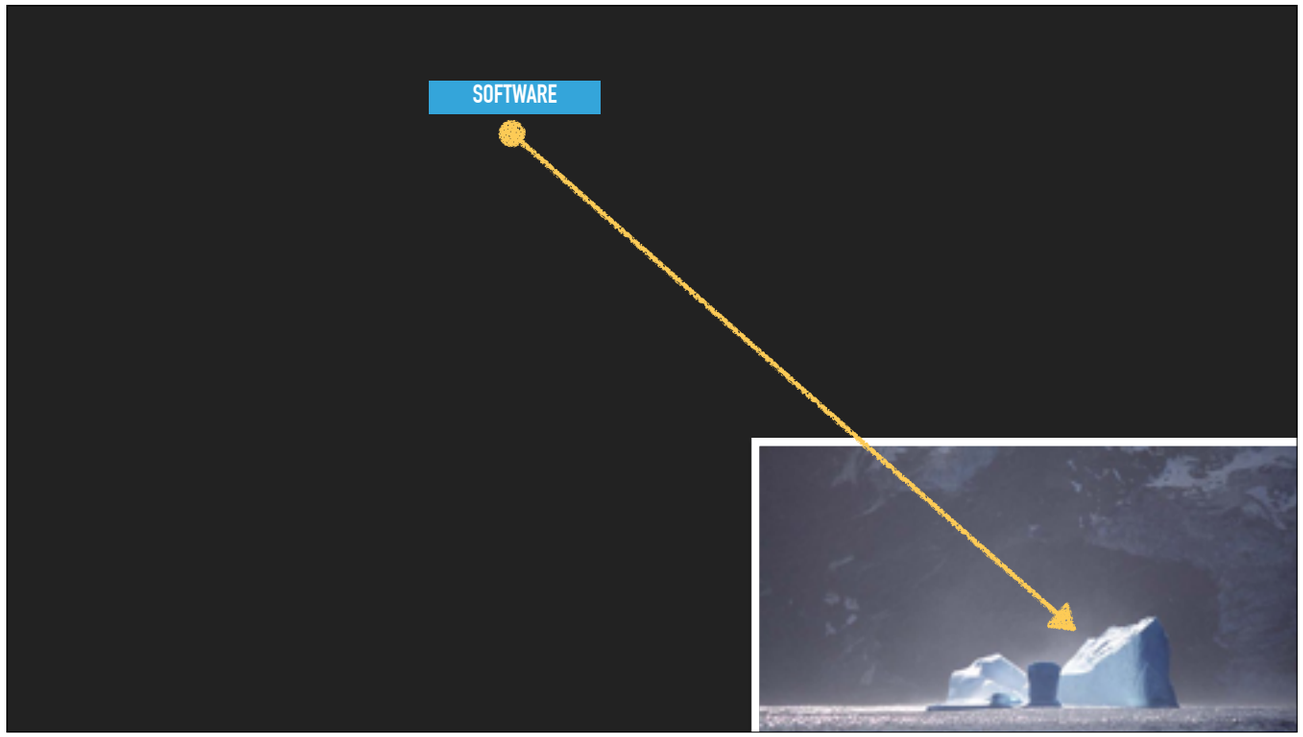

“Done — Houston, we have our app. We’re in the software business now and ready to sell our organically grown iceberg shards”

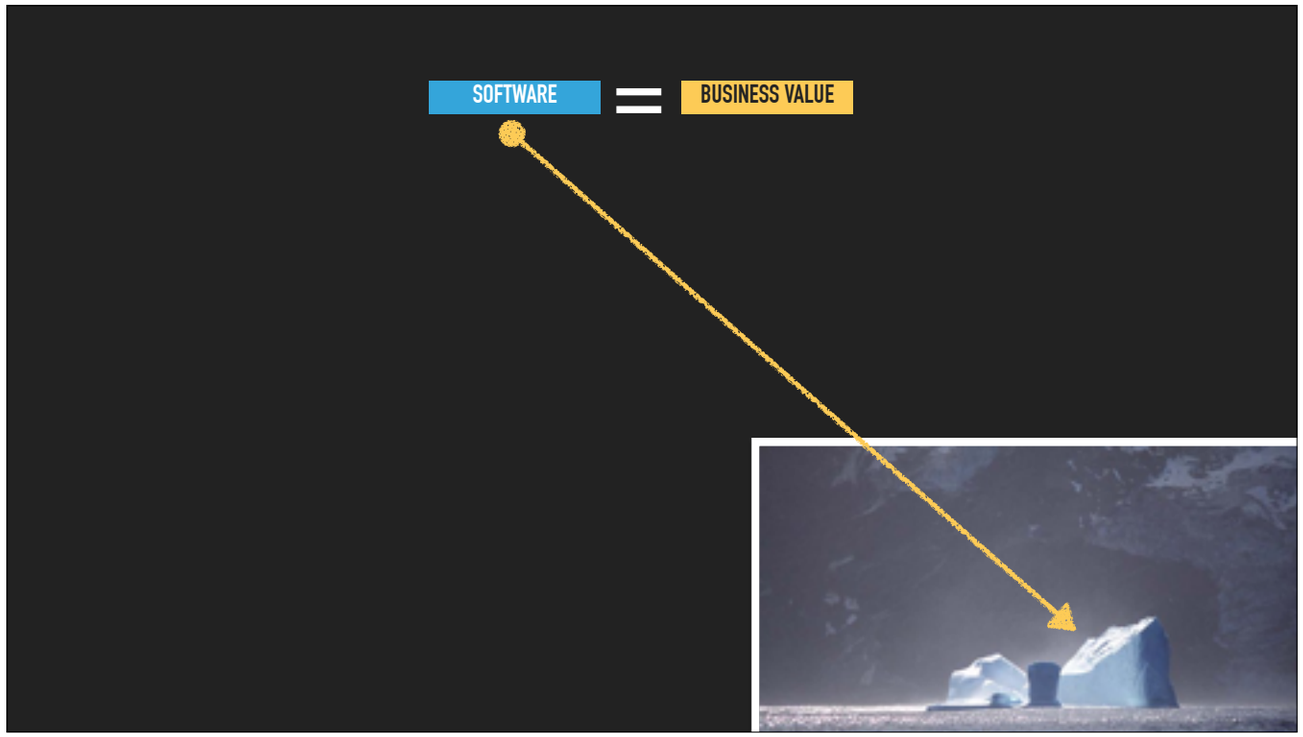

There is a catch though…

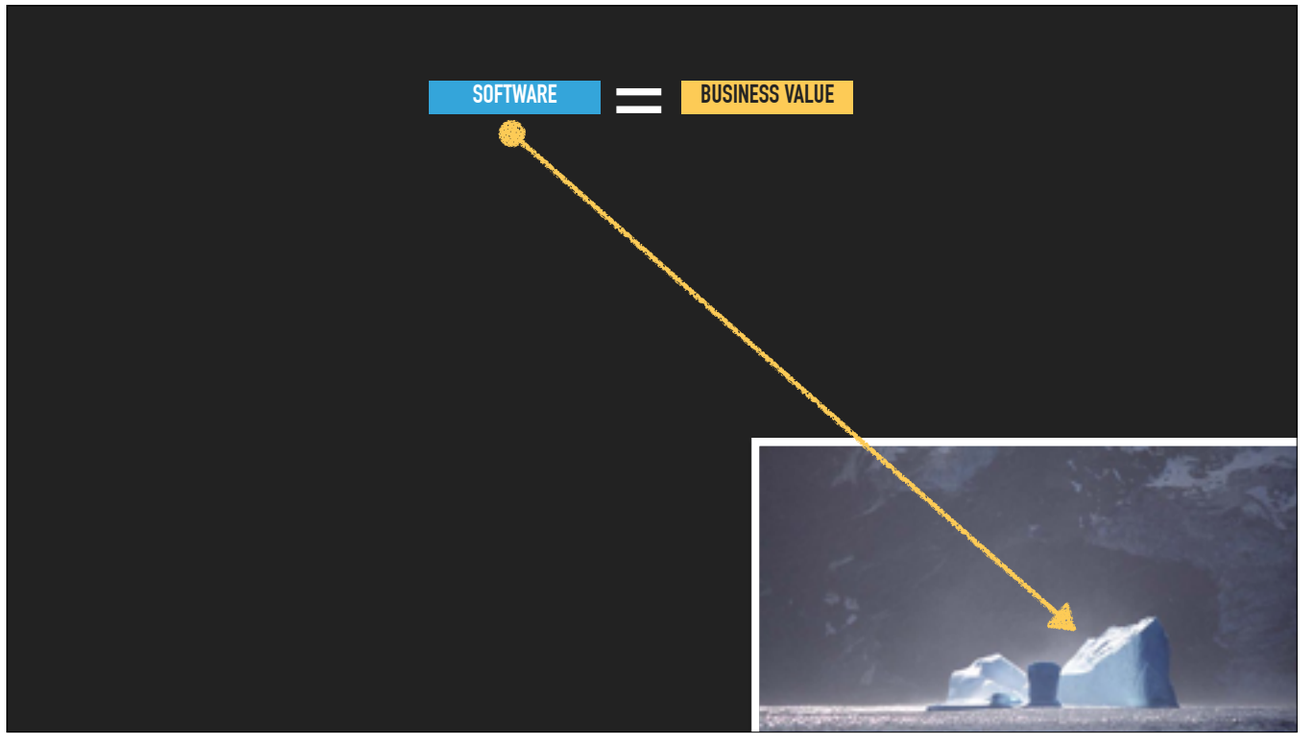

Businesses that are getting into the software game must have infrastructure. That is the real computers to run "workloads".

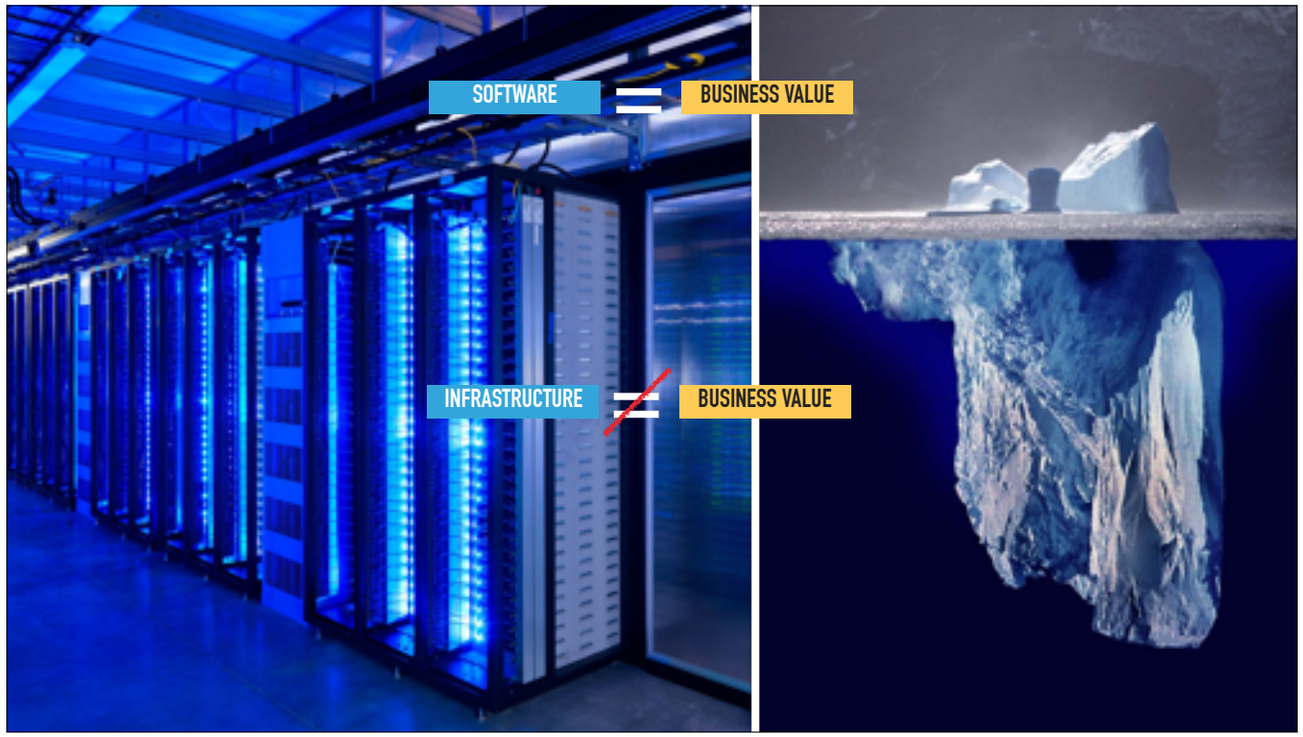

These are racked and stacked in large and costly facilities called data centers, wired to the electrical grid, with onsite power generation that will last for days in a Godzilla event. The meter is always running…

http://topwalls.net/server-data-center-lights-neon-computer-server-blue/

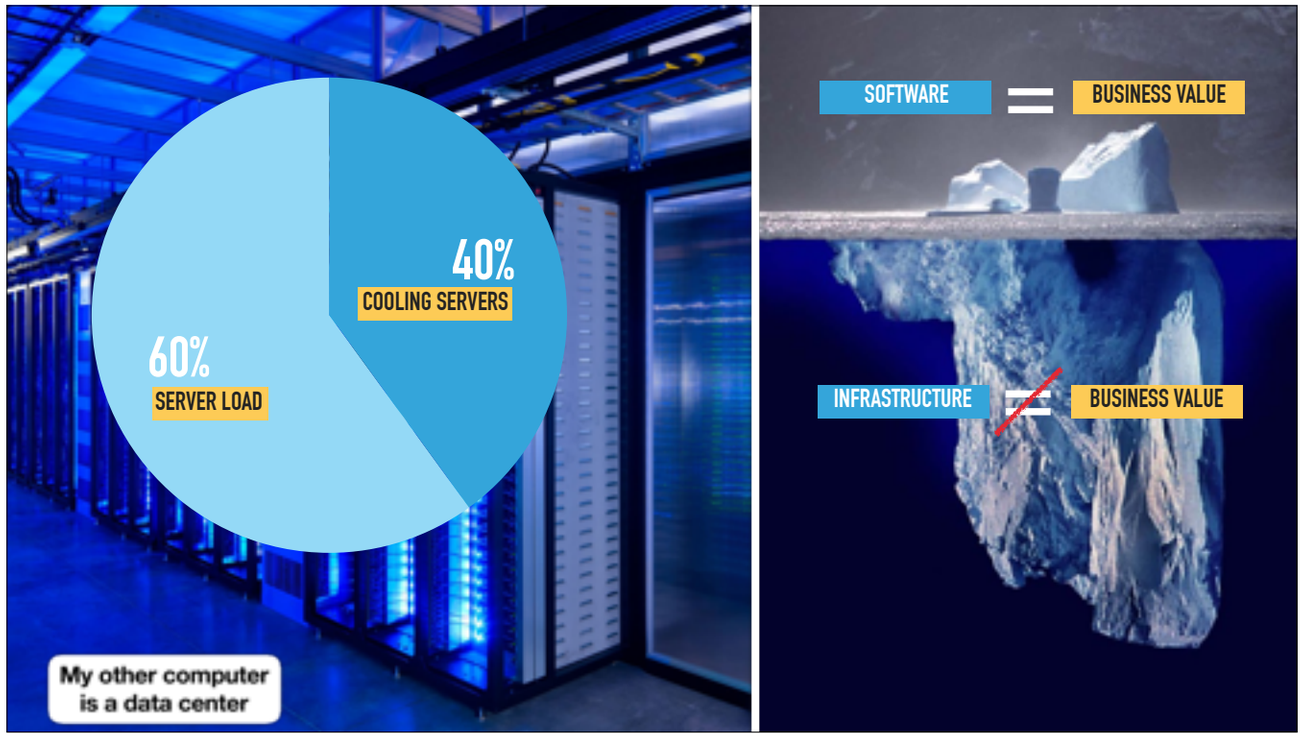

Everything in a data center is planned around watt usage. Over half of the wattage goes to server load, “compute” as it is called, and most of the rest to cooling equipment. Electricity is a large ongoing cost for a datacenter.

http://www.nytimes.com/2012/09/23/technology/data-centers-waste-vast-amounts-of-energy-belying-industry-image.html
http://www.nytimes.com/2011/09/09/technology/google-details-and-defends-its-use-of-electricity.html?_r=0

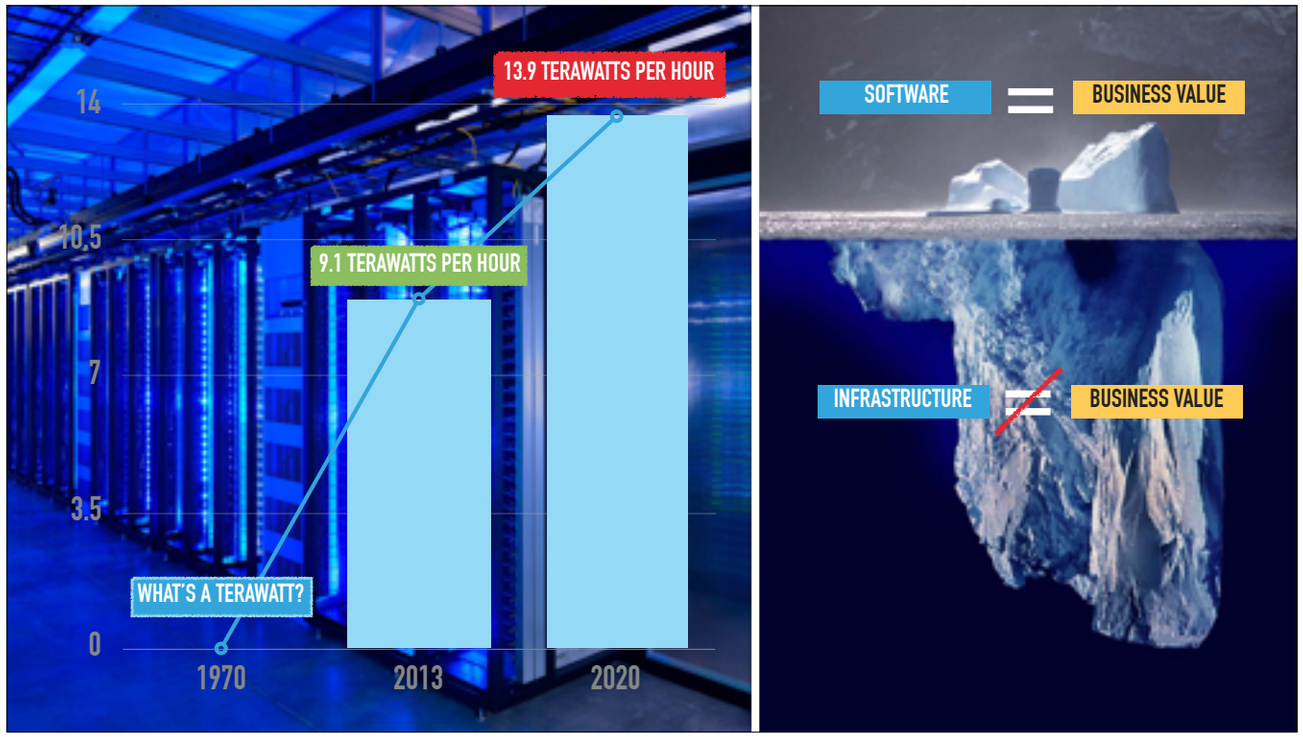

Lots and lots of electricity is used in data centers throughout the United States on the order of about 9.1 Terawatts per hour! With a projected rise to 13.9 Terawatts per hour by 2020.

To make better use of the physical hardware, each computer is “virtualized”.

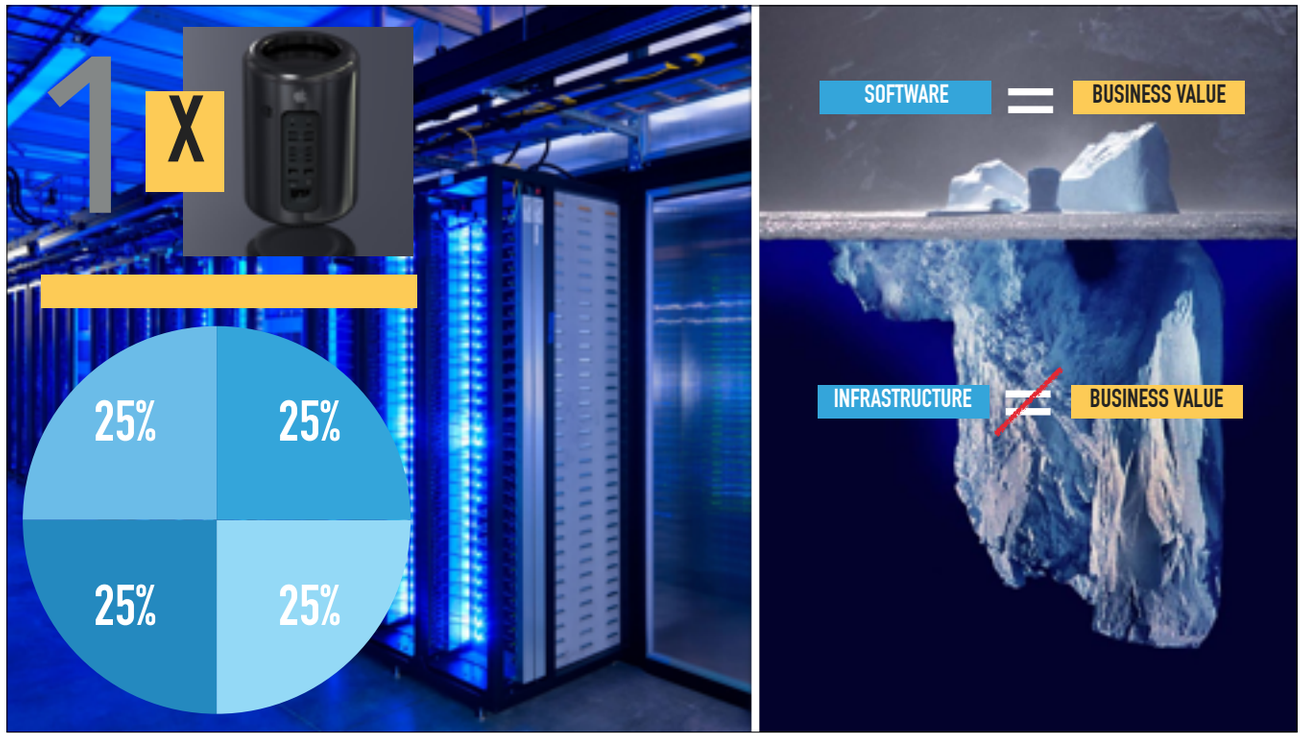

So that many virtual computers can run at the same time on a single physical computer. Little instances that do one thing well… like run a web server with our application on it.

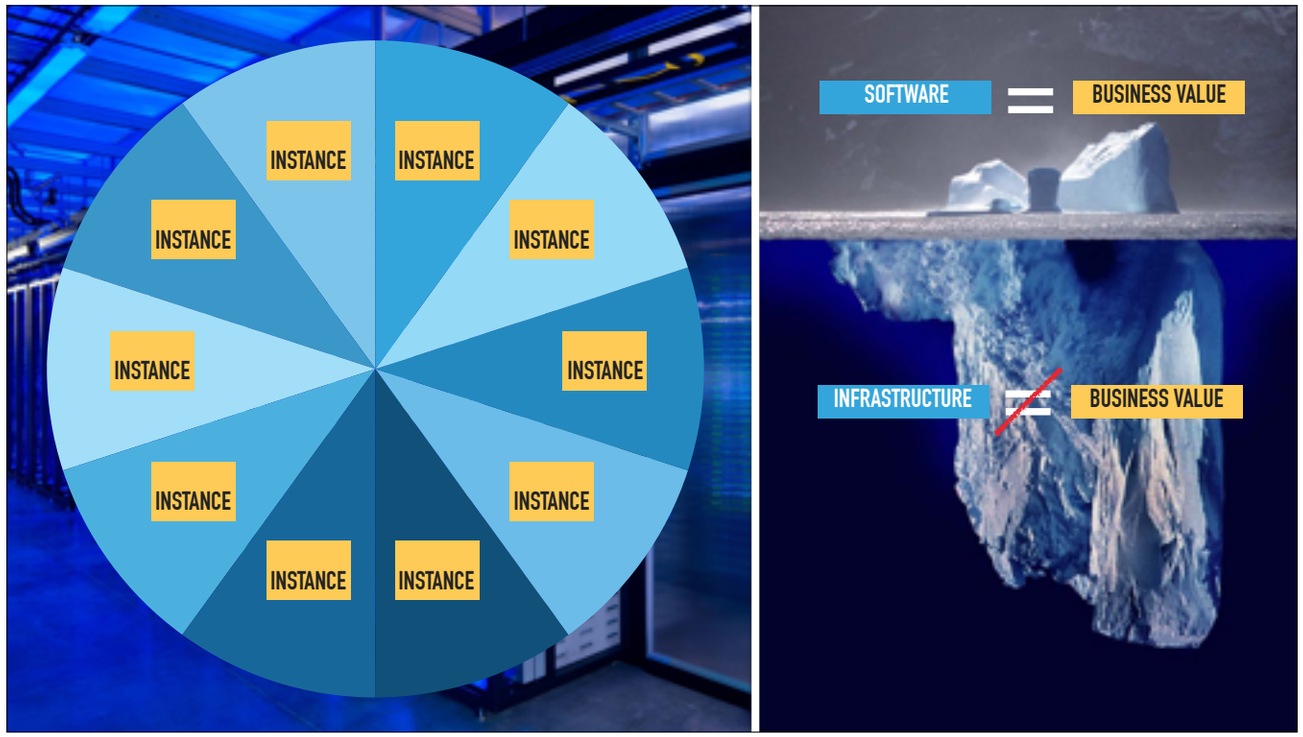

These instances allow us to run as many copies of our app as we need. So when our app is under-siege with the growing demand for fresh organically grown iceberg shards — we're ready!

But how much “compute” horsepower do we need?

How many instances should we have spun-up to meet demand?

The problem is that demand for our app is highly variable. On a cold day or at night, no one really wants to buy organically grown iceberg shards. If it’s summer or someones birthday party… we have serious demand.

What if we could better understand demand as it happens and respond to it by instantly “scaling” OUT our app to add more containers giving us more “compute”.

And, when we no longer need it, we can simply “scale” IN to just enough compute to handle the current demand.

Elastic compute, knowing when to automatically scale resources IN or OUT on demand is a great question to answer with data science.

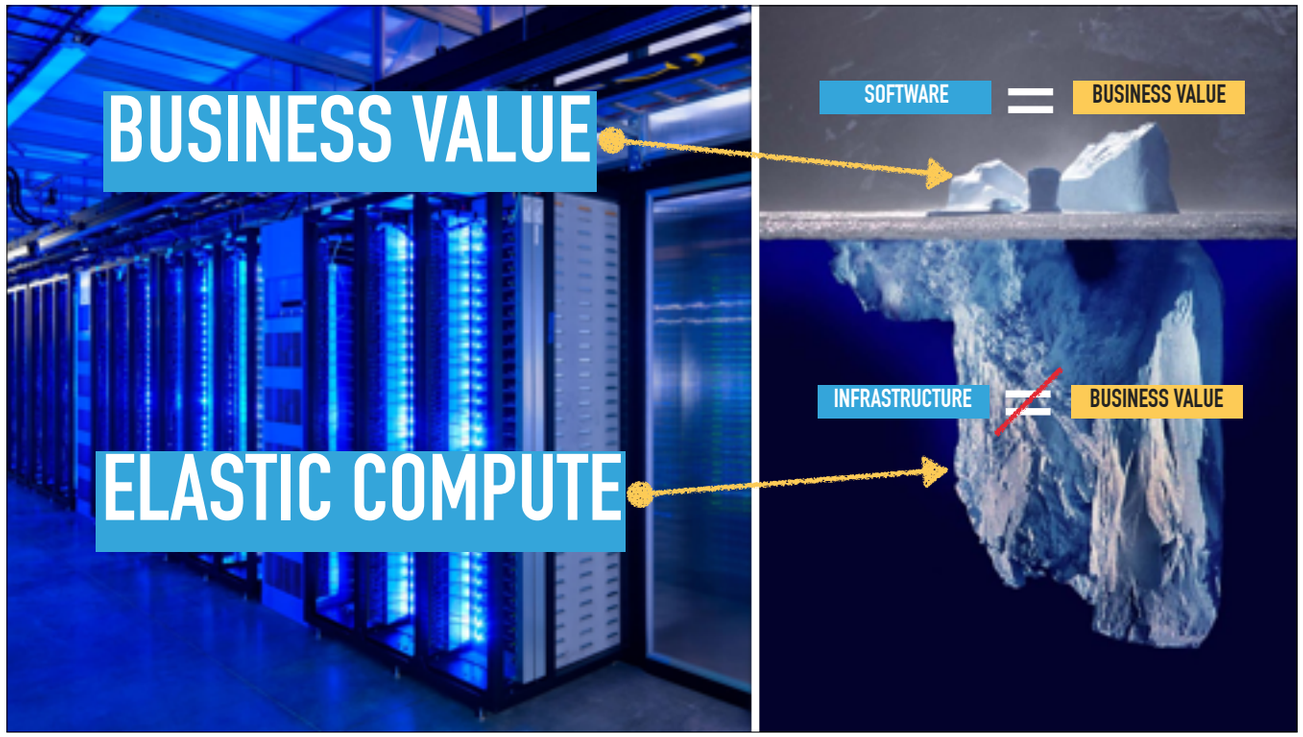

In the Elastic Compute project we explore whether we can begin reducing waste and "right-sizing" energy consumption for applications by linking them more directly to actual demand using data science techniques. We invite you to see the results of our research!

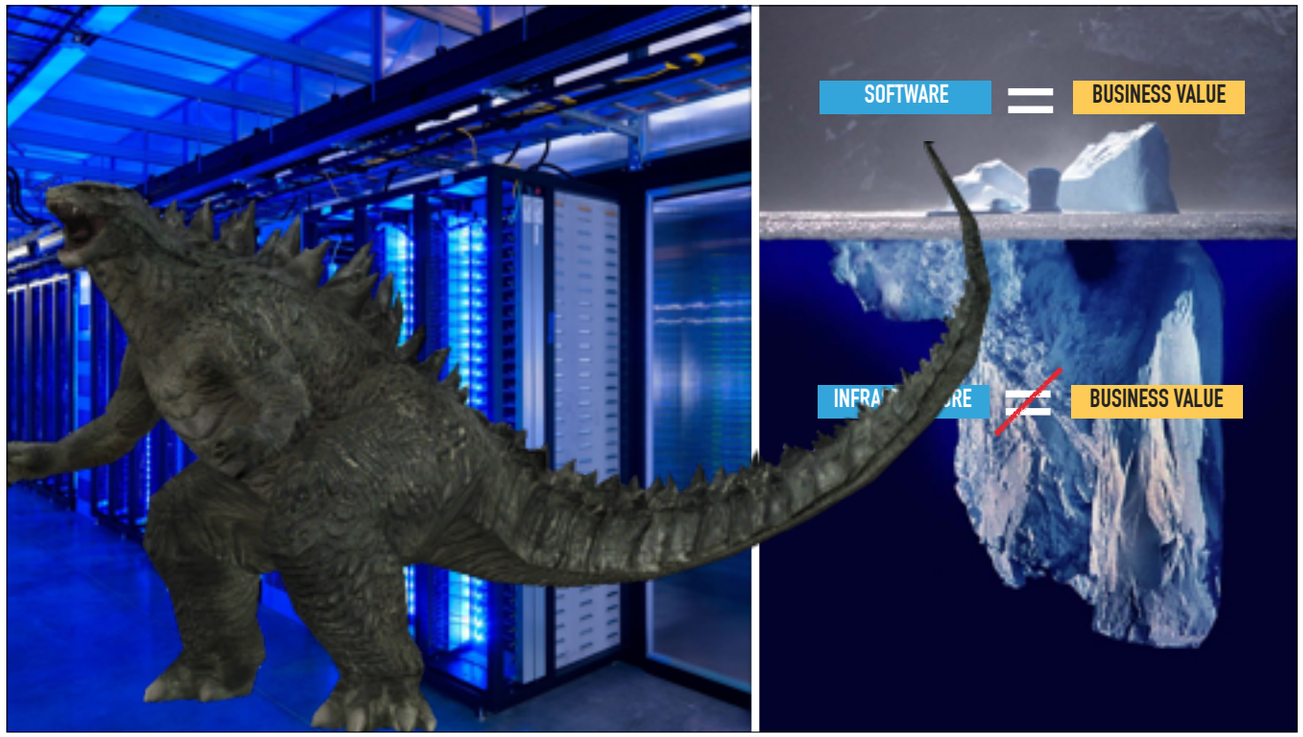

There he is… at least our datacenter was ready for it.

http://fantendo.wikia.com/wiki/File:Godzilla_transparent.png