



BootStack your Big Data cloud

Big or small, gain cloud efficiencies and solution deployment in one sweet package.

Big Data was born in the cloud

The operational efficiencies of cloud computing for processing big data aren't just apparent, they're practically necessary just to get it done. And sometimes, you need to keep your data private. But building your own private cloud infrastructure can be a time consuming and daunting task, and there's no guarantee you'll get it right. So let the cloud experts at Canonical help you with that, and bootstrap a Canonical OpenStack cloud for you with BootStack.

BootStack is an offering from Canonical where we build, support and manage your OpenStack cloud, to an SLA, for just \$15 per server per day, plus hardware and hosting. Then, when you're ready, we'll train your team and smoothly hand over operational control. The up-front cost savings of a BootStack partnership with Canonical can be tremendous, and the ability to have a Canonical OpenStack private cloud up and running in days, not months, keeps you focused on your data.

Solution benefits beyond OpenStack infrastructure

One of the tools used to deploy and manage a BootStack infrastructure is a service modeling and deployment product called Juju. It lets us architect, design, deploy, repeat, and edit integrated and complex solutions like OpenStack; and Juju can do the exact same thing for your big data solutions.

Managing, manipulating, and processing big data is rarely a simple task. Most systems require multiple software solutions and often microservices. All of these pieces have relationships. Those relationships need to be visually represented, then technically associated with their respective components as resources. Once the design is right, then you go about deploying your solution.

Doing all of this manually, for each implementation, can take time, and be error-prone. It also can hinder the ability to easily scale, and if you use complex scripts for automation, that becomes another stack to manage.

As part of a BootStack Big Data Cloud, the BootStack team will make Juju available to you, so you're ready to benefit from it immediately. Juju makes deployments easy and powerful. It gives you a canvas to lay out your solution, it lets you define service relationships, and it handles the entire deployment, and potentially dynamically scaling, for you.

Juju uses concepts called Charms and Bundles to deploy solutions. There are already dozens of big data software and services available as Charms in Juju's software catalog. For example, you can start with a core framework like Hadoop, Spark, or Kafka. Then, extend with an analytics solution (Saiku, SpagoBI), or a search engine (SolR, Datafari, ElasticSearch), or instantiate a cluster of your favorite NoSQL store (Cassandra, DSE, MongoDB, etc.) You're ready to start designing and deploying your big data solutions in a BootStack environment as soon as it comes online.

Charms and bundles

Charms define the actions, configurations, and relationships for an application or service. A bundle is a group of applications in a solution that have the relationships, and possibly configurations, already defined for you. Charms make applications dynamic, easily connected, and scalable, while bundles create portability and reusability. Juju lets you manage it all, via CLI, API or GUI, from one place.

Focus on the data

BootStack brings you a lot more than just a managed private cloud. It provides you with an integrated solution that's ready to deploy your complex big data services almost immediately.

Spend your time on the data. Let us handle the infrastructure, and let Juju handle the solution deployment. Together, with the BootStack team at Canonical, your Big Data Cloud just got a whole lot easier.

Compare the cost with the BootStack calculator, and get started with BootStack today at ubuntu.com/bootstack