

CASE STUDY

Regional Electoral Courts – TRE (Tribunal Regional Eleitoral)

Regional Electoral Court protects necessary electoral data and ease administration of multi-cluster environment with CloudCasa

The Challenge

As a government institution in Pará Brazil, regional electoral courts are mandated to have backup copies of their information, stored at an off-site location. Because the electoral courts are responsible for the control and inspection of the entire electoral process, as well as production of reports and electoral maps when counting votes, this requirement protects the essential electoral data from being lost, and/or tampered with.

For most courts, the backup requirement has prevented developers from taking advantage of the benefits of containerization and Kubernetes, because of Kubernetes shortcomings when it comes to data protection. Although Kubernetes is designed for efficiency, flexibility, scalability, and code portability, these benefits oftentimes lead to complex, unorganized, and unregulated environments that are difficult to manage.

Overview

Industry - Government

Headquarters - Belém
of Pará, Brazil

Partner – SUSE Rancher
on Longhorn storage

Solution Requirements

- ❑ Consider government regulations while protecting their Kubernetes applications.
- ❑ Besides offsite backup, needed a backup solution for cloud-native Kubernetes data.
- ❑ Scalable solution as Kubernetes footprint keeps growing.



Off-Site Backup and Other Government Requirements

As application developers at this court began to experiment with Kubernetes, it was an immediate concern that the data must adhere to the government requirements and have a full backup of the data protected at an off-site location. Natively, this process is an incredibly difficult task, but CloudCasa makes this easy.

Using CloudCasa, Kubernetes clusters and resources on SUSE Rancher using Longhorn storage are snapshotted and copied to an Amazon S3 bucket, managed by CloudCasa. These backup jobs include large Postgres databases that are prepared for backup using pre-built application hooks provided in the CloudCasa web-based GUI. Once protected, the clusters and applications can be recovered to the original location, or migrated to different clusters for disaster recovery, prod/staging testing, infrastructure upgrades, etc.

Challenges

- ❑ Ensuring data protection of Electoral data due to government mandate
- ❑ Limited resources for software infrastructure
- ❑ Application-consistent data protection of Kubernetes-hosted postgres databases
- ❑ Management and administration of multi-cluster environment

Besides the offsite backup, government organizations have several other regulations or requirements that a Kubernetes backup solution had to abide by. The solution had to be offered “as-a-service.” With limited resources when it came to infrastructure and storage, this court was able to register their clusters, and backup their data to the cloud without any software or hardware installation, and without purchasing and managing their own cloud object storage.

“CloudCasa has a rich and yet simple to use cloud dashboard, with the possibility to manage multiple clusters and perform backups and restores among them all with a single license.”

Manoel Frederico Beltrão, Judicial Analyst at Tribunal Regional Eleitoral do Pará

Solution

CloudCasa makes the complex task of managing backup and restores of multi-cluster Kubernetes environments a simple process.

Benefits

- ❑ Easy-to-use web interface and snapshot catalog
- ❑ Administration of multi-cluster environment made easier.
- ❑ Capacity-based pricing model.
- ❑ Seamless Rancher and Longhorn integration and upgrades.

Administration Benefits

The primary reason for utilizing a backup solution for their Kubernetes data was simply because it was required. However, CloudCasa provides the electoral court with several additional benefits that make administration of this multi-cluster environment much easier. Registration of new clusters and agent installation on those clusters is simple and straightforward — create an account and organization at cloudcasa.io/signup, and then register clusters with a simple process that adds the cluster info and installs the CloudCasa agent. For users with Kubernetes hosted in the cloud, this is even easier, as the clusters are auto-discovered after adding the cloud account to CloudCasa.

Once the cluster is registered, the CloudCasa agent can be installed via YAML files using a quick kubectl command. For SUSE Rancher customers, this is even easier, as installation of the CloudCasa agent can be done from a Helm chart in the Rancher marketplace. This makes it easy to quickly add newly created Kubernetes clusters to your existing backup policies in CloudCasa, without having to perform any manual, error-prone administration.

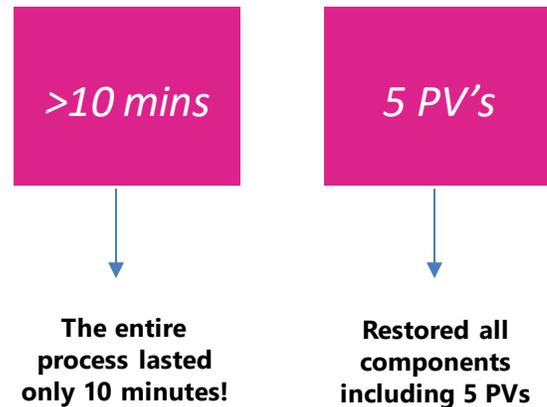
Finally, one of the biggest benefits for the Electoral court, was how CloudCasa assisted with seamlessly performing Rancher and Longhorn upgrades without developers temporarily losing access to their data. In this case, the Kubernetes administrator was able to upgrade SUSE Rancher from v2.5.11 to v2.7.1 without the developers even noticing. Instead of performing the upgrade in-place and interrupting the Kubernetes workloads, they simply took a snapshot and backup of the cluster and application data running on the old version of Rancher, created a new cluster using the new version of Rancher, and performed a cross-cluster recovery to migrate the data and configuration from the old to the new cluster.

Best Data Protection Service

Although the intention of using CloudCasa for data protection of their Kubernetes environment was to fulfill a requirement, this electoral court has found that CloudCasa offers so much more than data protection. Along with its many administration benefits, and its easy-to-use web interface and snapshot catalog, CloudCasa makes the complex task of managing backup and restores of multi-cluster Kubernetes environments a simple process.

Results

These were the results when restoring an application from a TRE-PA cluster to a TRE-ES cluster, with Longhorn on source cluster, and vSphere CSI on destiny.



Pricing Model

The backup solution also needed to have a consistent pricing model. Most Kubernetes backup providers license their solutions based on number of worker nodes and/or number of clusters. Because this court is progressively growing its Kubernetes footprint, it is hard to predict how many worker nodes and how many clusters the environment will consist of a year from now, or three years from now. CloudCasa uses simple capacity pricing tiers based on the amount of data backed up. This allows for flexibility and growth over a multi-year subscription, without having to worry about the number of worker nodes or clusters being used over time.

To learn more about CloudCasa and/or to sign up, visit [CloudCasa.io](https://cloudcasa.io). Please feel free to [get in touch](#) with any feedback on CloudCasa and what your requirements for cloud native backup are.