



Case Study: Aurea Software Goes Beyond the Limits of Amazon EBS to Run 200 Kubernetes Stateful Pods Per Host

CHALLENGES

- Create a single, multi-tenant Kubernetes platform capable of handling databases workloads for 80 different SaaS companies
- Traditional software-defined storage systems like Ceph and GlusterFS don't integrate well with Kubernetes or scale to the required levels
- Limited to only 40 EBS volumes per EC2 instance when using Amazon EBS for container storage, requiring them to overprovision VMs by 5x

SOLUTION

- Portworx provides a single, cloud native storage and data management layer for Kubernetes, no matter which stateful services Aurea's platform users require
- Portworx allowed Aurea to easily blow past the 40 EBS volume limit imposed by Linux-networking in order to run 200 container volumes per EC2 instance

RESULT

- Because of the densities enabled by Kubernetes and Portworx, Aurea saved 60-90% on compute costs
- Aurea saved roughly 5x on operations and development costs by deploying a single Kubernetes platform for 80 different SaaS applications

KEY TECHNOLOGIES DISCUSSED

Container Runtime – Docker

Scheduler – Kubernetes

Stateful Services – Cassandra, Solr, Memcache, ActiveMQ, Kafka, PostgreSQL

Infrastructure Provider – Amazon

We sat down with Sergey Pronin to learn about how Aurea Software is building a massive scale Kubernetes cluster with some very impressive density numbers. This is a transcript of that conversation.

Sergey Pronin

SaaS Ops Software
Engineering Manager at
Aurea Software



Can you tell me a little bit about what Aurea does as a company?

Aurea is the software engineering arm of ESW Capital. ESW Capital is a large investment fund which owns roughly 80 companies, mostly made up of SaaS and PaaS companies. Some of these companies provides SaaS, some provide on-prem software, etc. At Aurea, we write software, provide managed services, and professional services for these companies. All in all, the entire organization has about 4500 employees that make up the high-quality, fully remote teams (largely staffed by Crossover).

What is your role at Aurea?

I manage the container platform services for our organization. We call my group "Central" because we are the centralized Kubernetes service for 80 portfolio companies. We have built a Kubernetes platform that all of our internal customers, those 80 businesses, can consume. So instead of each of those SaaS companies needing to become the expert in running and operating Docker and Kubernetes, they can just consume a platform that's centrally provided by Aurea, and use it to host their applications.

Our company is based on the principle of economies of scale. We make a lot of acquisitions. We plan to acquire 50 more companies just this year. That's nearly one company per week. So we need to scale really fast. We don't want to have a SaaS Ops team for each company, we want to have a standardized approach for every company that we have, including, a single platform engineering team that can deliver value across all the companies. It's the same way we have Crossover legal and finance teams in place to support all our existing companies as well as new acquisitions.

In the end, the central platform that we build provides a standardized approach to IT Ops, so that every company can deploy their applications easily using standard tools, standard CI/CD pipelines, standard persistent storage, and standard monitoring.

Why did you settle on Kubernetes as the orchestration platform for containers?

We tested Swarm and ECS and they had their limitations and bugs that took too long to get fixed. Kubernetes has one of the largest open source communities, so we are confident that our issues are going to be solved fast and feature parity is keeping up. We run a lot of production apps on Kubernetes, including some really large clusters, so we're comfortable with the platform. We're also very involved in the community. We go to KubeCon and DockerCon and walk around talking to other companies to learn about their experiences and see how we can make better use of the platform.

The big benefits of Kubernetes are standardization and efficiency. We have all of these different companies, and they don't have to become experts in platforms. And because we have been able to standardize on a monitoring solution, a metric solution, a storage solution, and a scheduling solution, it's a lot faster for those teams to build applications and focus in on their specialty. The fact that they don't have to worry about running the platform and the specific apps on it, leads to big operations and development cost savings (roughly to the degree of 5x). We also realized that Kubernetes could help us drive massive opex cost efficiencies which is an important part of our overall business model.

“

We also realized that Kubernetes could help us drive massive opex cost efficiencies which is an important part of our overall business model.”

Our clusters are highly dense, meaning we run a lot of containers per host. On AWS, we use huge instances, like x1.32xlarge, as much as possible. The recommendation from Kubernetes is 100 pods per VM. Already, we're running 200-300 pods per host. Also, since most of the apps that we run are stateful, we can easily have 200-300 volumes per host as well. And we're working to push these limits even further. Right now, we have a test project that runs on dedicated worker nodes. They run 500 pods with a persistent volume each. Because of these densities enabled by Kubernetes and Portworx, we're easily saving 60-90% on our compute costs. Portworx itself was between 30-50% cheaper than any other storage solution we tested.

“

Because of these densities enabled by Kubernetes and Portworx, we're easily saving 60-90% on our compute costs. Portworx itself was between 30-50% cheaper than any other storage solution we tested.”

What were some of the challenges you had to overcome to run stateful services on Kubernetes?

First, on Amazon you simply cannot use EBS for Kubernetes, because you are limited to only 40 or so volumes per host, and remember, we are trying to get to 500 volumes per host. That was one of the main reasons we had to go out and find a cloud native storage solutions designed for Kubernetes.

Second, we run our Kubernetes clusters multi-tenant, which means that we have 80 different companies running on the same cluster, all simultaneously. One of the biggest challenges is to isolate resources within the Kubernetes cluster, so that they do not kill each other. This is especially important with resource intensive databases.

Additionally, and another challenge for us is that we don't really know what is going to be deployed to our cluster. For example when we acquire a new company, they will take their app, containerize it and then run it on our platform. We don't know what stateful services they run. They could be running Cassandra, or Solr, or Kafka, or maybe a huge SMTP server with billions and billions of files. We have no idea ahead of time, so we need a consistent way to manage those stateful apps without being an operational expert in each. For us, running stateful containers on our cluster is not just optional. We have to do it.

“

For us, running stateful containers on our cluster is not just optional. We have to do it.”

And for that reason, we tested seven or eight different solutions before landing on Portworx. We looked at Amazon EBS, Amazon EFS, GlusterFS, and Ceph running on EBS and instance storage among others.

Since we're on Amazon, the most obvious solution was EBS, but as I mentioned, that wasn't possible because we could only run 40 volumes per host, when our goal is a minimum of 200-300 pods. Right there, our compute savings would have vaporized.

From there we started looking at various software-defined options. But we had issues with performance with almost everyone. Ceph, on the other hand gave us good performance, but Red Hat told us don't run Ceph over EBS because you're basically running Ceph over Ceph, and you will face issues if you do that. We also found issues with Ceph's Kubernetes integration.

“

Ceph, on the other hand gave us good performance, but Red Hat told us don't run Ceph over EBS because you're basically running Ceph over Ceph, and you will face issues if you do that. We also found issues with Ceph's Kubernetes integration.”

Ceph is a good solution for storage when you create a block device now and then, and you use it, and you just store data there. But when you are running containers, and you need to create a new volume say every second because you are scaling one app, and at the same time removing a bunch of volumes because you're scaling down another app, Ceph is really slow. It is not a good solution for dynamic container environments.

What advice would you give someone who is thinking about building a container platform and needs to support stateful applications?

People need to be ready to face the challenges of limiting the resources consumed by apps in the clusters. Just because we have Cgroups, doesn't mean the problem is solved. This is especially the case in large multi-tenant Kubernetes clusters where you are allowing business to essentially deploy anything, which if your not careful can kill your cluster. You need to understand how to segregate the resources, how to separate one container from another. This is the hardest part, and where every container-as-a-service provider struggles as well.

LEARN MORE

Portworx is the cloud native storage company that enterprises depend on to reduce the cost and complexity of rapidly deploying containerized applications across multiple clouds and on-prem environments.

With Portworx, you can manage any database or stateful service on any infrastructure using any container scheduler. You get a single data management layer for all of your stateful services, no matter where they run. Portworx thrives in multi-cloud environments.

Contact us to find out how Portworx's unique storage solution can deliver the availability, performance, and features necessary to minimize stateful application operations costs and improve revenue growth -

<https://portworx.com/request-a-demo>

portworx.com

www.aurea.com

 @portwx

 linkedin.com/company/portworx



4940 El Camino Real, Ste 200, Los Altos, CA 94022
Tel: 650-241-3222 | info@portworx.com | portworx.com