

Home > Blog >

Architect's Corner: Qomplx leverages automation to run stateful services in containers for the financial and cybersecurity industries

Architect's Corner: Qomplx leverages automation to run stateful services in containers for the financial and cybersecurity industries

September 20, 2018 | 5 mins | Architect's Corner



DEVOPS REALIZED: QOMPLX LEVERAGES AUTOMATION TO RUN STATEFUL SERVICES IN CONTAINERS AT MASSIVE SCALE



In today's [Architect's Corner](#), we speak with Sunil Pentapati, Director of Platform Operations at [Qomplx](#), who manages the DevOps & QA teams. Sunil speaks with us about the crucial need for real-time automation in today's modern data centers, especially while running stateful services in containers at scale. You can also find a PDF of this case study [here](#).

QOMPLX:

Can you tell us a little bit about Qomplx?

Qomplx is a 120+ person software company that applies artificial intelligence to solve complex, real-world problems at scale with our next-gen Data Analytics-as-a-Service platform: Qomplx OS. We have spent years thoughtfully combining foundational data handling, analytics, and automation with cutting-edge simulation modeling and deep learning to offer decision-making support that guides today's businesses toward total optimization.

specifically, the company's founders recognized from the beginning that the most challenging problems in that space could only be addressed with a unified, cloud-based platform that ingests, integrates, and correlates data from every available source in real time. The distributed computing power Qomplx OS provides the immediate context needed to understand what's happening on the network when it's happening, and what to do about it in order to minimize your risk exposure.

Tell us about your role at Qomplx?

I joined Qomplx last year as Director of Platform Operations. The mission and vision for the team is to create a multi-cloud enabled production SaaS platform and to be a differentiator. I work with a very talented team that tackles a lot of interesting use cases, including developing frameworks, tools, and processes that range from Infrastructure Automation, CI/CD, and Quality Engineering to the day-to-day operational management of very large, distributed analytic systems that allow Qomplx to run these cutting-edge services at cloud scale.

How are you using containers at Qomplx?

To build this integrated platform, we knew we wanted to use a container orchestration system because our software platform needs to be able to run in multiple environments including our own multi-tenant SaaS infrastructure, or a customer's data center, or their VPC in Amazon, Google, or Azure. We have to be flexible enough to support whatever deployment works best for the customer. If you're a financial services company that operates in a sensitive market or vertical, for instance, you may not want to run on a public cloud. But you still need the cybersecurity monitoring and response system that our platform provides. We want to make it easy for the customer to pick the model that is right for them, but that means we will have to run in multiple environments, and containers are the only way to do that easily and effectively.



The biggest problem we faced was that we have a set of stateful services and a set of stateless services, each operating differently.



Most container orchestration platforms are built for stateless services, and can scale them elastically. But we did not find a good solution for stateful services until we started using Portworx.

We tried everything from RexRay, mounted volumes, local volumes, etc. We quickly learned that, like our unified cybersecurity platform, we needed a unified container storage solution to solve the container data management problem. We realized that each of the solutions that we looked at solved one small problem, but left a lot for us to still do ourselves. Because of this gap, we initially achieved stability by moving all our stateful services back into running on VMs. But that was never going to work long-term because a VM-based solution wouldn't work across clouds and on-prem data centers without significant re-work for each deployment.

So that's when we started looking into real enterprise solutions. I come from security & systems background, having worked for companies like IBM, EMC, and RSA, and I understand the importance of having a stable and scalable infrastructure platform to build a stable SaaS product.

The biggest challenge is that every stateful service is a snowflake. Each is unique and has its own way of doing things in regards to how it functions and scales. Each has its own operational challenges related to things like performance, backups, recovery, encryption, etc.

manage it in production, in an automated way, when infrastructure inevitably fails. For any of the stateful services we are using—all the databases, Flink, or other distributed services—running them in production at scale and operationalizing them with respect to HA, backups, and disaster recovery was a challenge that we needed to overcome.

In other words, we didn't just need persistent storage to run a stateful app in containers. We needed data management capabilities to manage a stateful app in production. There are lots of persistent storage solutions for containers, but few true container data management options. The reason is scale.



We not only need to run stateful services in containers, we need to be able to run them at scale. And when we say scale, it's massive scale. We need to maintain millisecond query times on petabyte-sized databases.

The biggest challenge was ensuring that we have a solution that works not only for one stateful service but rather for all stateful services in a common way so that we can scale this platform without a lot of tech debt and without a lot of manual effort.

You emphasize automation being so important. How does automation factor into your decision about data management solutions?

I strongly believe that firefighting kills innovation. I truly believe that to the core.





The more time engineering teams and operations teams spend fighting fires or fixing yesterday's problems, the less time they spend solving real customer problems before they happen.

And working for so many Fortune 500 companies, I saw first-hand how constant firefighting puts some of these companies at a disadvantage compared to some of the newer companies. You're constantly in reactive mode battling issues, and that's not allowing you to innovate enough to compete with the newer companies which are small and nimble.

One of the biggest factors in these newer companies being so much more agile is automation, i.e., the mindset that says "treat everything as code." And so, for us, we went back to our initial core principle of needing a unified platform that ingests, integrates, and correlates data. For that to happen at scale, we cannot do things manually. It has to be in real time.



For high-volume analytics, real time and manual don't go together. For us to really scale this platform in real time, we need solid automation, not just for the compute layer, but also for the data layer. That's where Portworx comes in.

What advice would you give to another architect who wants to run stateful services in containers?

Whatever solution you're using, make sure it solves a majority of your data management needs. In other words, your production operations problems as well.

is a reason why they do that. Fewer variations make easier for maintenance purposes. The mindset of those companies is: you have fewer moving parts that you need to worry about. The same concept is valid for software and how you run things on the IT side.



The fewer number of things that you deal with, the fewer potential problems you'll need to deal with. Rather than going for a siloed solution, look for a unified solution that solves the majority of your problems.

And understand that no matter what you do, there will always be use cases that require some tweaking or customization. You'll need to solve that anyway. But I would advise customers to look for a solution that helps solve the majority of their problems.

We tried out all kinds of solutions. And one of the reasons we chose Portworx is that it solves a majority of our use cases and problems in a fully automated way. It allowed us to standardize on one solution to scale faster and better.

Share



[Subscribe for Updates](#)



Submit





Let's Talk

Ready for a demo, or just want to talk? Reach out to our cloud experts today.

Call: (800) 379-7873

Contact Us

© 2026 | Portworx

Products

- Portworx Enterprise
- Portworx Backup
- Portworx Disaster Recovery
- Free Trial

Resources

- Resources
- Install & Docs
- News & Events
- Cookie Policy
- Privacy Notice

Solutions

- Kubernetes Storage
- Modern Virtualization
- Kubernetes Backup
- Kubernetes Disaster Recovery
- Developer Platform
- AI/ML
- Financial Services
- FlashArray + Portworx

Company

- About Us
- Careers
- Partners
- Security and Trust Center
- Contact Sales



Portworx by Pure
Storage
2555 Augustine Dr.
Santa Clara, CA 95054
hi@portworx.com

