

Migrating to AKS: A Big Win for Big Data

We want to share some of the key lessons learned from our experience migrating the environment to Azure Kubernetes Service (AKS).

The main goal was to move off of dedicated servers and into AKS to gain more flexibility. With dedicated servers, our client was stuck overprovisioning to handle peak loads. And any time they needed to scale up or down, it required tedious manual work. They also ran into major availability issues whenever a single server went down.

Moving to AKS became a must.

One of the huge pluses we got with AKS was automatic scaling on demand. With big data loads, processing needs can swing wildly based on the volume and velocity of data. AKS allows us to scale up or down in real time based on actual workloads. This elasticity is perfect for handling the natural ebbs and flows of data processing.

AKS also provides built-in redundancy. If a node fails, workloads instantly shift to another node without missing a beat. In analytics, where uptime is critical, AKS ensures that applications always function flawlessly.

The Kubernetes architecture also improves resilience. By isolating services in containers, a failure is prevented from cascading through the system. This containment, along with rapid container restarts, avoids disruptions to data flows.

By moving to AKS, we gained agility, resilience, optimized infrastructure, and accelerated innovation. For any big data initiative, we'd strongly recommend switching to Kubernetes. The benefits we've experienced firsthand can take your analytics to the next level.



Dmitry Vishnyov
Co-founder & CBDO

OUR CONTACT

+357 25 059376

hello@itoutposts.com

www.itoutposts.com

TOP RATED DEVOPS COMPANY

50+

projects delivered remotely

90%

of certified engineers in the company

2 years

average client engagement duration

4.9/5

customer satisfaction score

OUR AWARDS

