

TuneIn lowers cloud costs and increases visibility into cloud spending with Flexera

The challenge

TuneIn is an audio streaming service that delivers live news, radio, sports, music and podcasts to more than 75 million active monthly users. After experiencing some outages in their data centers several years ago, TuneIn's engineering leadership decided to migrate their applications and services to the AWS cloud to stop managing their own physical infrastructure.

After the migration, the cost of cloud operations was higher than prior hosting costs and drawing the finance team's attention, who wanted to reduce the bills for Kubernetes and cloud infrastructure-related services. Although they were hoping to save on costs with AWS Reserved Instances, it was not a reliable option—siload teams were running their own EC2, Autoscaling Groups (ASGs), ECS and Kubernetes/EKS stacks with little consistency in instance families and types across workloads made it difficult to commit to reserved capacity.

Engineering teams also needed more visibility into spending data as they were moving more applications to Kubernetes. But the task of calculating these metrics took significant time and effort, and with no dedicated FinOps team to track it, costs were quickly growing.



» AT A GLANCE

Industry: Media

Products: Flexera Ocean, Flexera One Cloud Commitment Management

TuneIn, the leading live streaming and on-demand audio service, brings together live sports, news, music, podcasts and radio from around the globe. With 75 million monthly active users, TuneIn is one of the most widely used streaming audio platforms in the world. TuneIn broadcasts over 100,000 owned and operated and partner radio stations and boasts more than 5.7 million podcasts. With premiere distribution across 200 platforms and connected devices, TuneIn empowers listeners to 'hear' what they love wherever 'here' might be. TuneIn premium subscribers also unlock exclusive access to commercial-free news from top networks like CNBC, CNN, FOX News Radio and MSNBC, as well as commercial-free music channels for every mood and activity.

The solution

Flexera offered TunelN a solution for FinOps that would give them the savings they were looking for by using spot instances, the reliability they needed with an enterprise-grade SLA and in-depth visibility into the cost of workloads running in their Kubernetes clusters.

“We were never willing to use spot instances,” said Ryan White, Senior Director of Engineering Operations at TunelN. The inherent volatility of spot instances makes them risky, but TunelN trusted Flexera to run their Kubernetes workloads on spot instances after a successful trial of Ocean* and coordination with developers to adjust their workloads to make them more resilient to the nature of spot disruptions. With Flexera, TunelN saw immediate benefits in cost savings, cost analysis capabilities and high availability for their applications.

The results

Using Flexera Ocean and Flexera One Cloud Commitment Management to leverage spot and reserved instances for maximum cost savings

Running spot instances is inherently cheaper—up to 90% in some cases—but they come with the risk that the cloud provider can take back the spot instances applications are running on if they’re needed elsewhere. With millions of users streaming TunelN content, using their website and apps, and sending requests for content from hundreds of types of devices like Alexa and Google Home, losing service was not an option. Instead of hiring someone to analyze their AWS usage, manage instances and ensure availability, TunelN uses Ocean to leverage spot instances and Cloud Commitment Management to manage reserved instances (RIs) for workloads running

on Kubernetes. TunelN is able to save more than 25% of their entire AWS bill (including storage, networking, etc.) while also mitigating the risk of service disruption by utilizing spot capacity across a range of instance families and sizes.

The savings don’t stop there. Cloud Commitment Management provided TunelN with savings of over 40% on RIs. Additionally, the product’s flexibility allows TunelN to maximize savings while avoiding long-term commitment lock-in. In one instance, TunelN needed to temporarily increase capacity to support a testing project. Cloud Commitment Management cost-effectively increased TunelN’s RIs by more than 340% over 11 days. Then, two weeks later when testing was complete, it reduced their total RIs by an equivalent amount in only three days.

“We don’t even have to think about provisioning cloud infrastructure, Spot* just handles it all for us. The beauty of Spot is that we just set it and forget it,” said White.

High availability for mixed instance types

For TunelN, highly available systems are critical to maintaining quality customer experiences. In addition to Ocean’s predictive capabilities that anticipate disruptions to the spot market, TunelN can also leverage Ocean’s unique ability to use different instance types and sizes to ensure even more reliability when using spot instances. By leveraging Kubernetes affinity and anti-affinity rules, which can be applied by node, availability zone, and even instance type/size, TunelN is able to strategically declare rules for instance types to keep infrastructure from going dark when there is an outage or disruption in the spare capacity market.

» HIGHLIGHT

More 25% than savings of the entire AWS bill. TunelN uses Flexera Ocean to leverage spot instances and Flexera One Cloud Commitment Management to manage RI’s for workloads running on Kubernetes.

Kubernetes cost allocation and showback helps make business decisions

As their footprint in the cloud grew, so did Tuneln's cloud bill, especially after they built their in-house real-time transcoding service, which converts upstream feeds (i.e., from radio stations, podcasts, live sporting events etc.) into the right audio format. This application, running at scale on top of Kubernetes, is critical to the business, ensuring that Tuneln's streams could be hosted in-house and highly available.

The teams managing Tuneln's API's and streaming services needed to know how much they were spending per application/service and stream, along with other deployments, to understand profit and loss margins. Using Ocean, they've seen how much they are spending per Kubernetes namespace, deployment and even pod, and can easily calculate the cost per stream. This granular visibility into spending data isn't possible by looking at only an AWS bill, and any estimate based on just that would have been an educated guess at best.

"Being able to pull out the cost data from the Ocean UI to determine how much we're spending by service and by stream has been hugely beneficial to multiple teams," said White.

* Flexera acquired Spot in March 2025.


About Flexera

Flexera helps organizations understand and maximize the value of their technology, including the rising costs and risks introduced by AI, saving billions of dollars in wasted spend. Our Flexera One platform connects the dots between what technology you have, how it is used, what it costs, and where it creates risk, helping teams take control of the increasingly complex IT estate across cloud, SaaS and on-premises. We are leading the way to unify IT asset management, FinOps and SaaS management with high fidelity data from Technopedia, our proprietary reference library of technology asset data, and intelligent automation fueled by AI. That's why thousands of global organizations rely on the Flexera One platform and Technopedia. Learn more at flexera.com

“

We don't even have to think about provisioning cloud infrastructure, Spot* just handles it all for us. The beauty of Spot is that we just set it and forget it”

Ryan White
Senior Director of Engineering Operations



» NEXT STEPS

Understand what your Kubernetes workloads really cost—and where you can save with Flexera Ocean

[Get your savings analysis →](#)