

Multi Cloud API Routing and Policing

Challenge

Public cloud infrastructure has fundamentally changed enterprise IT strategies. The level of automation made available by Amazon Web Services (AWS), Microsoft Azure and Google Cloud have made it possible for many enterprises to retire data centers and migrate all applications to the public cloud region. Unfortunately, using a public cloud does not ensure your applications automatically become highly available. Many enterprises would prefer to not be do not want to be locked into one vendor's cloud offerings and may continue to use a private data center.

To address these high-availability and vendor-proofing requirements, many enterprises choose to deploy application stacks across multiple cloud regions, multiple cloud providers and data centers. In doing so, enterprises are encountering a performance impact; modern applications tend to rely on API or microservices gateways to carry out user authentication and policing functions. Gateways are usually packaged as virtual machines and deployed in the public subnet of a VPC in one public cloud region (or in the DMZ in a data center). If an application has microservices deployed across multiple cloud regions, end user traffic will enter the application environment through the gateway, and then get routed to the right public cloud region or data center where the relevant microservice has been deployed. This "traffic tromboning" leads to added network latency and impacts overall application response times.

Instead, what if the all intra-application routing and policing decisions could be carried out closer to endpoints or end users, allowing application owners to curtail traffic tromboning and reduce application response times? Many enterprises would prefer to adopt such a strategy, but lack access to distributed infrastructure and tools to internally develop such distributed API routing and policing services.

Solution

Rafay enables enterprises to rapidly deploy API and microservices routing and policing policies across a network of distributed edge clusters - without any DevOps effort. Called the Programmable Edge™, Rafay's platform delivers powerful yet easy-to-use APIs for deploying microservices worldwide in minutes. Enterprises can leverage Rafay's multi-tenant network to gain instant access to global locations close to end users or deploy Programmable Edge technology in their own private edge application network.

The platform applies a variety of patent-pending algorithms to optimize microservice placement decisions based on configured policies. The platform also provides a suite of application lifecycle management tools to help developers deploy and run their applications across a number of edges, so you don't have to.

By leveraging Rafay's Programmable Edge platform, developers can apply API and microservices routing and policing strategies closer to endpoints and end users. As a result, application owners can effectively curtail traffic tromboning and reduce application response times, resulting in improved user engagement.

Benefits

- **On-Demand Expansion of Your Application's Global Footprint:** All traffic from endpoints is handled by your containerized microservices, which are dynamically placed close to endpoints.
- **Highly Improved End User Experience – Globally:** Distributed API routing, global load balancing, end point authentication and policing ensure that end users experience a consistently high performing application.
- **Developer-Friendly Application Lifecycle Management Tools:** A suite of application life cycle features empower developers to deploy microservices to the large number of edge clusters without worrying about internally developing code and artifact deployment, log aggregation, and a variety of other tools.