

Improving Rate Limit Experience for Developers in Istio

Zufar Dhiyaulhaq



#IstioCon

Why Rate Limit?

Essential component of
Security

Reduce throttling errors

Reduce resource consumption

Rate limits is implemented via EnvoyFilter object.

```
metadata:
  name: filter-ratelimit-svc
  namespace: istio-system
spec:
  workloadSelector:
    labels:
      istio: ingressgateway
  configPatches:
    - applyTo: VIRTUAL_HOST
      match:
        context: GATEWAY
        routeConfiguration:
          vhost:
            name: ""
            route:
              action: ANY
      patch:
        operation: MERGE
        # Applies the rate limit rules.
        value:
          rate_limits:
            - actions: # any actions in here
              - request_headers:
                  header_name: ":path"
                  descriptor_key: "PATH"
EOF
```

How Ratelimit Implemented In Istio Today

Problem Statement

Developers need to understand very complex configurations within the mesh

Hard to maintain and guard trail EnvoyFilter, since it's breaking glass API

Prone to human error

Slow down Istio upgrade

Requirements

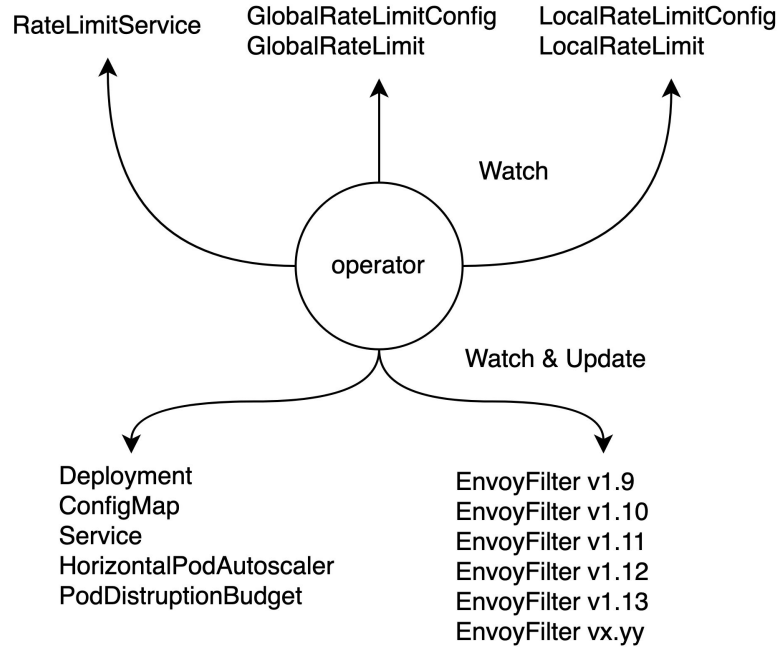
Abstracts the complex EnvoyFilter configurations

Automatically create rate limit configuration for every Istio version

Guard trail the configuration for every Istio version

Support Global & Local Rate Limit

Approach



RateLimitService

```
apiVersion: ratelimit.zufardhiyaulhaq.com/v1alpha1
kind: RateLimitService
metadata:
  name: internal-gateway-rls
  namespace: istio-system
spec:
  kubernetes:
    replica_count: 3
    auto_scaling:
      max_replicas: 15
      min_replicas: 3
    resources:
      limits:
        cpu: "2"
        memory: "2048Mi"
      requests:
        cpu: "1"
        memory: "512Mi"
  backend:
    redis:
      type: "single"
      url: "10.121.211.209:6379"
  monitoring:
    statsd:
      enabled: true
      spec:
        host: statsd-exporter.infrastructure.svc.cluster.local
        port: 8125
```

GlobalRateLimitConfig

```
apiVersion: ratelimit.zufardhiyaulhaq.com/v1alpha1
kind: GlobalRateLimitConfig
metadata:
  name: internal-gateway-grlc
  namespace: istio-system
spec:
  type: "gateway"
  selector:
    labels:
      istio: internal-gateway
  istio_version:
    - "1.10"
    - "1.11"
    - "1.12"
  ratelimit:
    spec:
      domain: "internal-gateway"
      failure_mode_deny: false
      timeout: "25"
    service:
      type: "service"
      name: "internal-gateway-rls"
```

GlobalRateLimit

```
apiVersion: ratelimit.zufardhiyaulhaq.com/v1alpha1
kind: GlobalRateLimit
metadata:
  name: foo-gopay-com-default-route
  namespace: istio-system
spec:
  config: "internal-gateway-grlc"
  selector:
    vhost: "foo.gopay.com:80"
    route: "default-route"
  matcher:
    - request_headers:
        header_name: ":method"
        descriptor_key: "method"
    - request_headers:
        header_name: ":path"
        descriptor_key: "path"
    - generic_key:
        descriptor_value: "default-route"
        descriptor_key: "route"
    - generic_key:
        descriptor_value: "foo.gopay.com"
        descriptor_key: "domain"
  limit:
    unit: hour
    requests_per_unit: 1
```



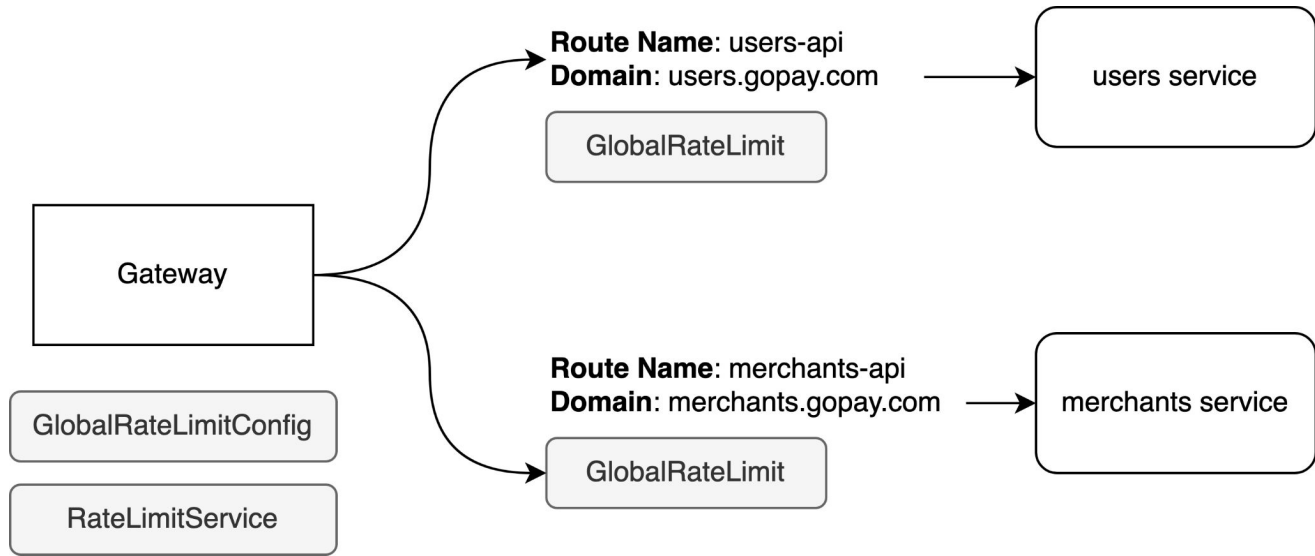
GlobalRateLimit

```
apiVersion: ratelimit.zufardhiyaulhaq.com/v1alpha1
kind: GlobalRateLimit
metadata:
  name: foo-gopay-com-default-route
  namespace: istio-system
spec:
  config: "internal-gateway-rls-grlc"
  selector:
    vhost: "foo.gopay.com:80"
    route: "default-route"
  matcher:
    request_headers:
      header_name: ":method"
      descriptor_key: "method"
    - request_headers:
        header_name: ":path"
        descriptor_key: "path"
    - generic_key:
        descriptor_value: "default-route"
        descriptor_key: "route"
    - generic_key:
        descriptor_value: "foo.gopay.com"
        descriptor_key: "domain"
  limit:
    unit: hour
    requests_per_unit: 1
```

VirtualService

```
apiVersion: networking.istio.io/v1beta1
kind: VirtualService
metadata:
  name: foo-gopay-com-vs
  namespace: istio-system
spec:
  hosts:
    - "foo.gopay.com"
  gateways:
    - istio-system/internal-gateway
  http:
    - name: "default-route"
      route:
        - destination:
            host: foo.default.svc.cluster.local
            port:
              number: 8080
```





The background consists of a repeating pattern of overlapping circles. The circles are arranged in a grid-like fashion, with each circle overlapping its neighbors. There are two shades of blue: a darker, royal blue and a lighter, cyan blue. The circles are arranged such that the darker blue circles are slightly offset from the lighter blue circles, creating a sense of depth and movement.

Demo

Production ready Istio Ratelimit Operator.

Integration with gopay.sh, our SuperApp for Engineer.



**Next
Milestones**

Thank you!

@zufardhiyaulhaq

<https://zufardhiyaulhaq.com>

#IstioCon

