

CDNI CAPACITY CAPABILITY ADVERTISEMENT EXTENSIONS

Andy Ryan

Ben Rosenblum

IETF 115

11/11/2022

OBJECTIVES OF CDNI CAPACITY CAPABILITY ADVERTISEMENT DRAFT

- ▶ Provide framework for information exchange to facilitate traffic delegation decisions
- ▶ Establish limits that are specific to delegation relationship
- ▶ Define limits in unambiguous, mutually understood units

WHERE WE ARE?

- ▶ How could we define a limit specific to the type of traffic being delegated and keeping in mind, not all traffic is created equal. For example,
 - ▶ Game download traffic might be high bps low rps
 - ▶ Whereas, Low Latency HLS could be high rps and high bps.
- ▶ These different traffic profiles could have different impact on the underlying CDN
- ▶ Therefore, we want to allow for some additional granularity which would leverage an association of traffic type with the delivery host (or an identifier which maps to a configuration)

WHAT ARE WE THINKING?

- ▶ Footprint provides the geographical/network boundary, which is consistent with FCI, but within the CapacityLimits FCI payload, we would leverage a scope object which would further define the limit to apply to not only the footprint but also this additional scope-type

CAPABILITY ADVERTISEMENT

```
{
  "capabilities":[
    {
      "capability-type": "FCI.CapacityLimits",
      "capability-value":{
        "limits":[
          {
            "id": "capacity_limit_region1",
            "limit-type": "egress",
            "maximum-hard": 50000000000,
            "maximum-soft": 40000000000,
            "telemetry-source":{
              "id": "capacity_metrics_region1",
              "metric": "egress_5m"
            }
          }
        ],
      },
      {
        "id": "capacity_limit_serviceA_region1",
        "scope":{
          "type": "published-host",
          "values":[
            "serviceA.cdn.example.com"
          ]
        },
        "limit-type": "egress",
        "maximum-hard": 30000000000,
        "maximum-soft": 20000000000,
        "telemetry-source":{
          "id": "capacity_metrics_serviceA_region1",
          "metric": "egress_5m"
        }
      }
    ]
  ],
},
```

```
{
  "capability-type": "FCI.CapacityLimits",
  "capability-value":{
    "limits":[
      {
        "id": "capacity_limit_region2",
        "limit-type": "egress",
        "maximum-hard": 20000000000,
        "maximum-soft": 10000000000,
        "telemetry-source":{
          "id": "capacity_metrics_region2",
          "metric": "egress_5m"
        }
      }
    ]
  },
  "footprints":[
    {
      "footprint-type": "ipv4cidr",
      "footprint-value": ["10.0.10.0/24"]
    }
  ]
}
}
```

IN SUMMARY

- ▶ The following is a summary of what the FCI.CapacityLimits payload specifies represented in a hierarchical manner of the IPv4 CIDR ranges.
 - ▶ 10.0.0.0/8
 - ▶ ALL traffic \leq 50000000000
 - ▶ serviceA.cdn.example.com \leq 30000000000
 - ▶ 10.0.10/24
 - ▶ ALL traffic \leq 20000000000
- ▶ In a scenario a uCDN is considering how to delegate traffic for serviceA.cdn.example.com towards 10.0.10/24, the following conditions need to be met:
 - ▶ traffic within 10.0.10/24 must stay under 20000000000 as per capacity_metrics_region2
 - ▶ traffic for serviceA.cdn.example.com within 10.0.0.0/8 must stay under 30000000000 as per capacity_metrics_serviceA_region1

Ideally, the scoping would be fully contained within the Footprint object but in lieu of that, we implemented sub-scope specific to the CapacityLimits object

SO, WHAT'S NEXT?

Welcome any feedback on the topic

THANK YOU