

CDNI Capacity Insights Capability Advertisement Extensions

Nov 2021

Capacity Insights

What draft are we discussing?

- CDNI Capacity Insights Capability Advertisement Extensions
 - draft-ryan-cdni-capacity-insights-extensions-01

Capacity Insights

What are the goals of the Capacity Interface?

- 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
- Bidirectional communication between delegation participants
- Leverage CDNi as much as possible, extend where needed

Capacity Insights

How does the communication work?

- ❑ FCI.CapacityLimits
 - ❑ Specify the limit of traffic that should be delegated in units, such as Bits per Second, called limit-types
 - ❑ References a specific Telemetry source which outlines current usage of a particular limit-type
- ❑ FCI.Telemetry
 - ❑ Allows the advertisement of what types of Telemetry sources are supported
 - ❑ Initially scoped to generic Telemetry sources, but paves the way for a formal Telemetry interface integration
- ❑ MI.RequestedCapacityLimits
 - ❑ Provides a vehicle for the uCDN to ask the dCDN for an update to established Limits

Capacity Insights – FCI.CapacityLimits

```
{
  "capabilities": [
    '...',
    {
      "capability-type": "FCI.CapacityLimits"
      "capability-value": {
        "total-limits": [
          {
            "limit-type": "egress",
            "maximum-hard": 500000000000,
            "maximum-soft": 250000000000,
            "telemetry-source": {
              "id": "capacity_metrics_region1",
              "metric": "egress_5m"
            }
          }
        ],
        "host-limits": [
          {
            "host": "serviceA.cdn.example.com",
            "limits": [
              "limit-type": "egress",
              "maximum-hard": 200000000000,
              "maximum-soft": 100000000000,
              "telemetry-source": {
                "id": "capacity_metrics_region1",
                "metric": "egress_service2_5m"
              }
            ]
          }
        ]
      }
    }
  ],
  '...',
  '...'
}
```

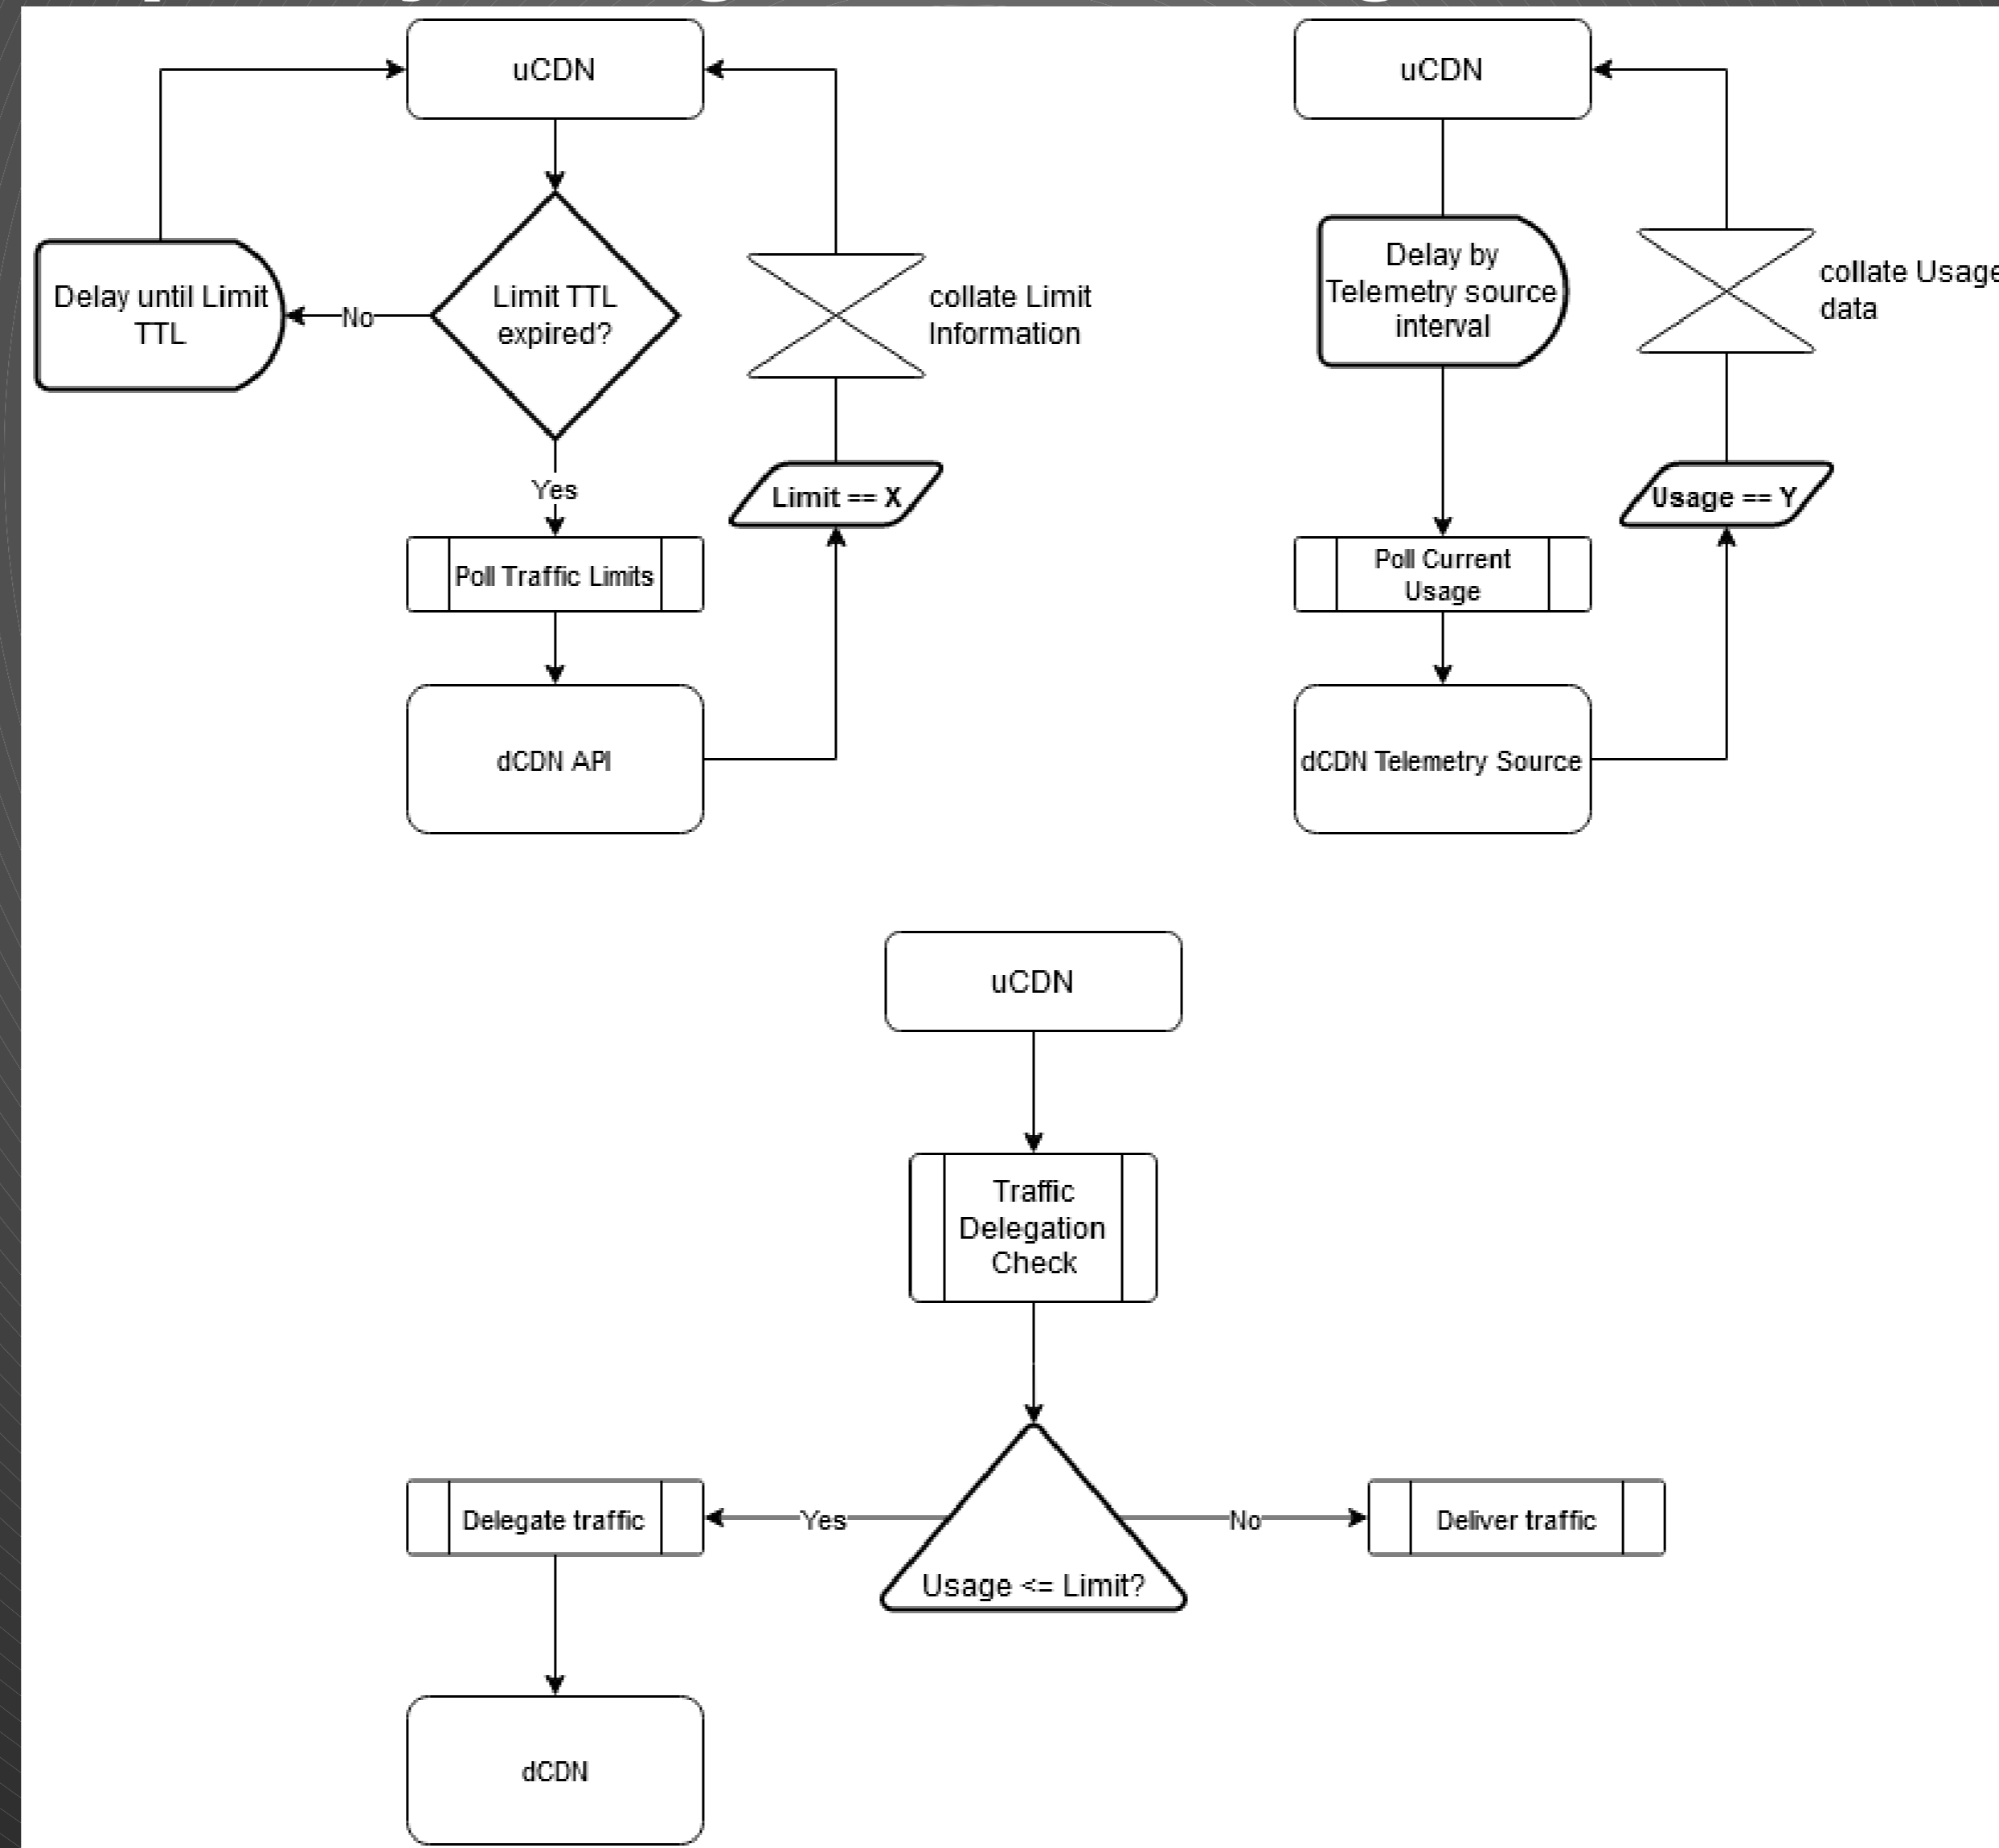
- Total-limits: limit for all traffic associated with a uCDN/dCDN relationship
- Host-limits: limit for traffic for a specific CDN-Domain (published host)
- Each limit is scoped to a footprint (as per the CDNi FCI specification)
- Each limit references a Telemetry source which provides data with the same scope as the defined limit (i.e Bits per Second for Host X within Footprint Y)
- GOAL: How much can you send, and how can you track how much you are sending

Capacity Insights – FCI.Telemetry

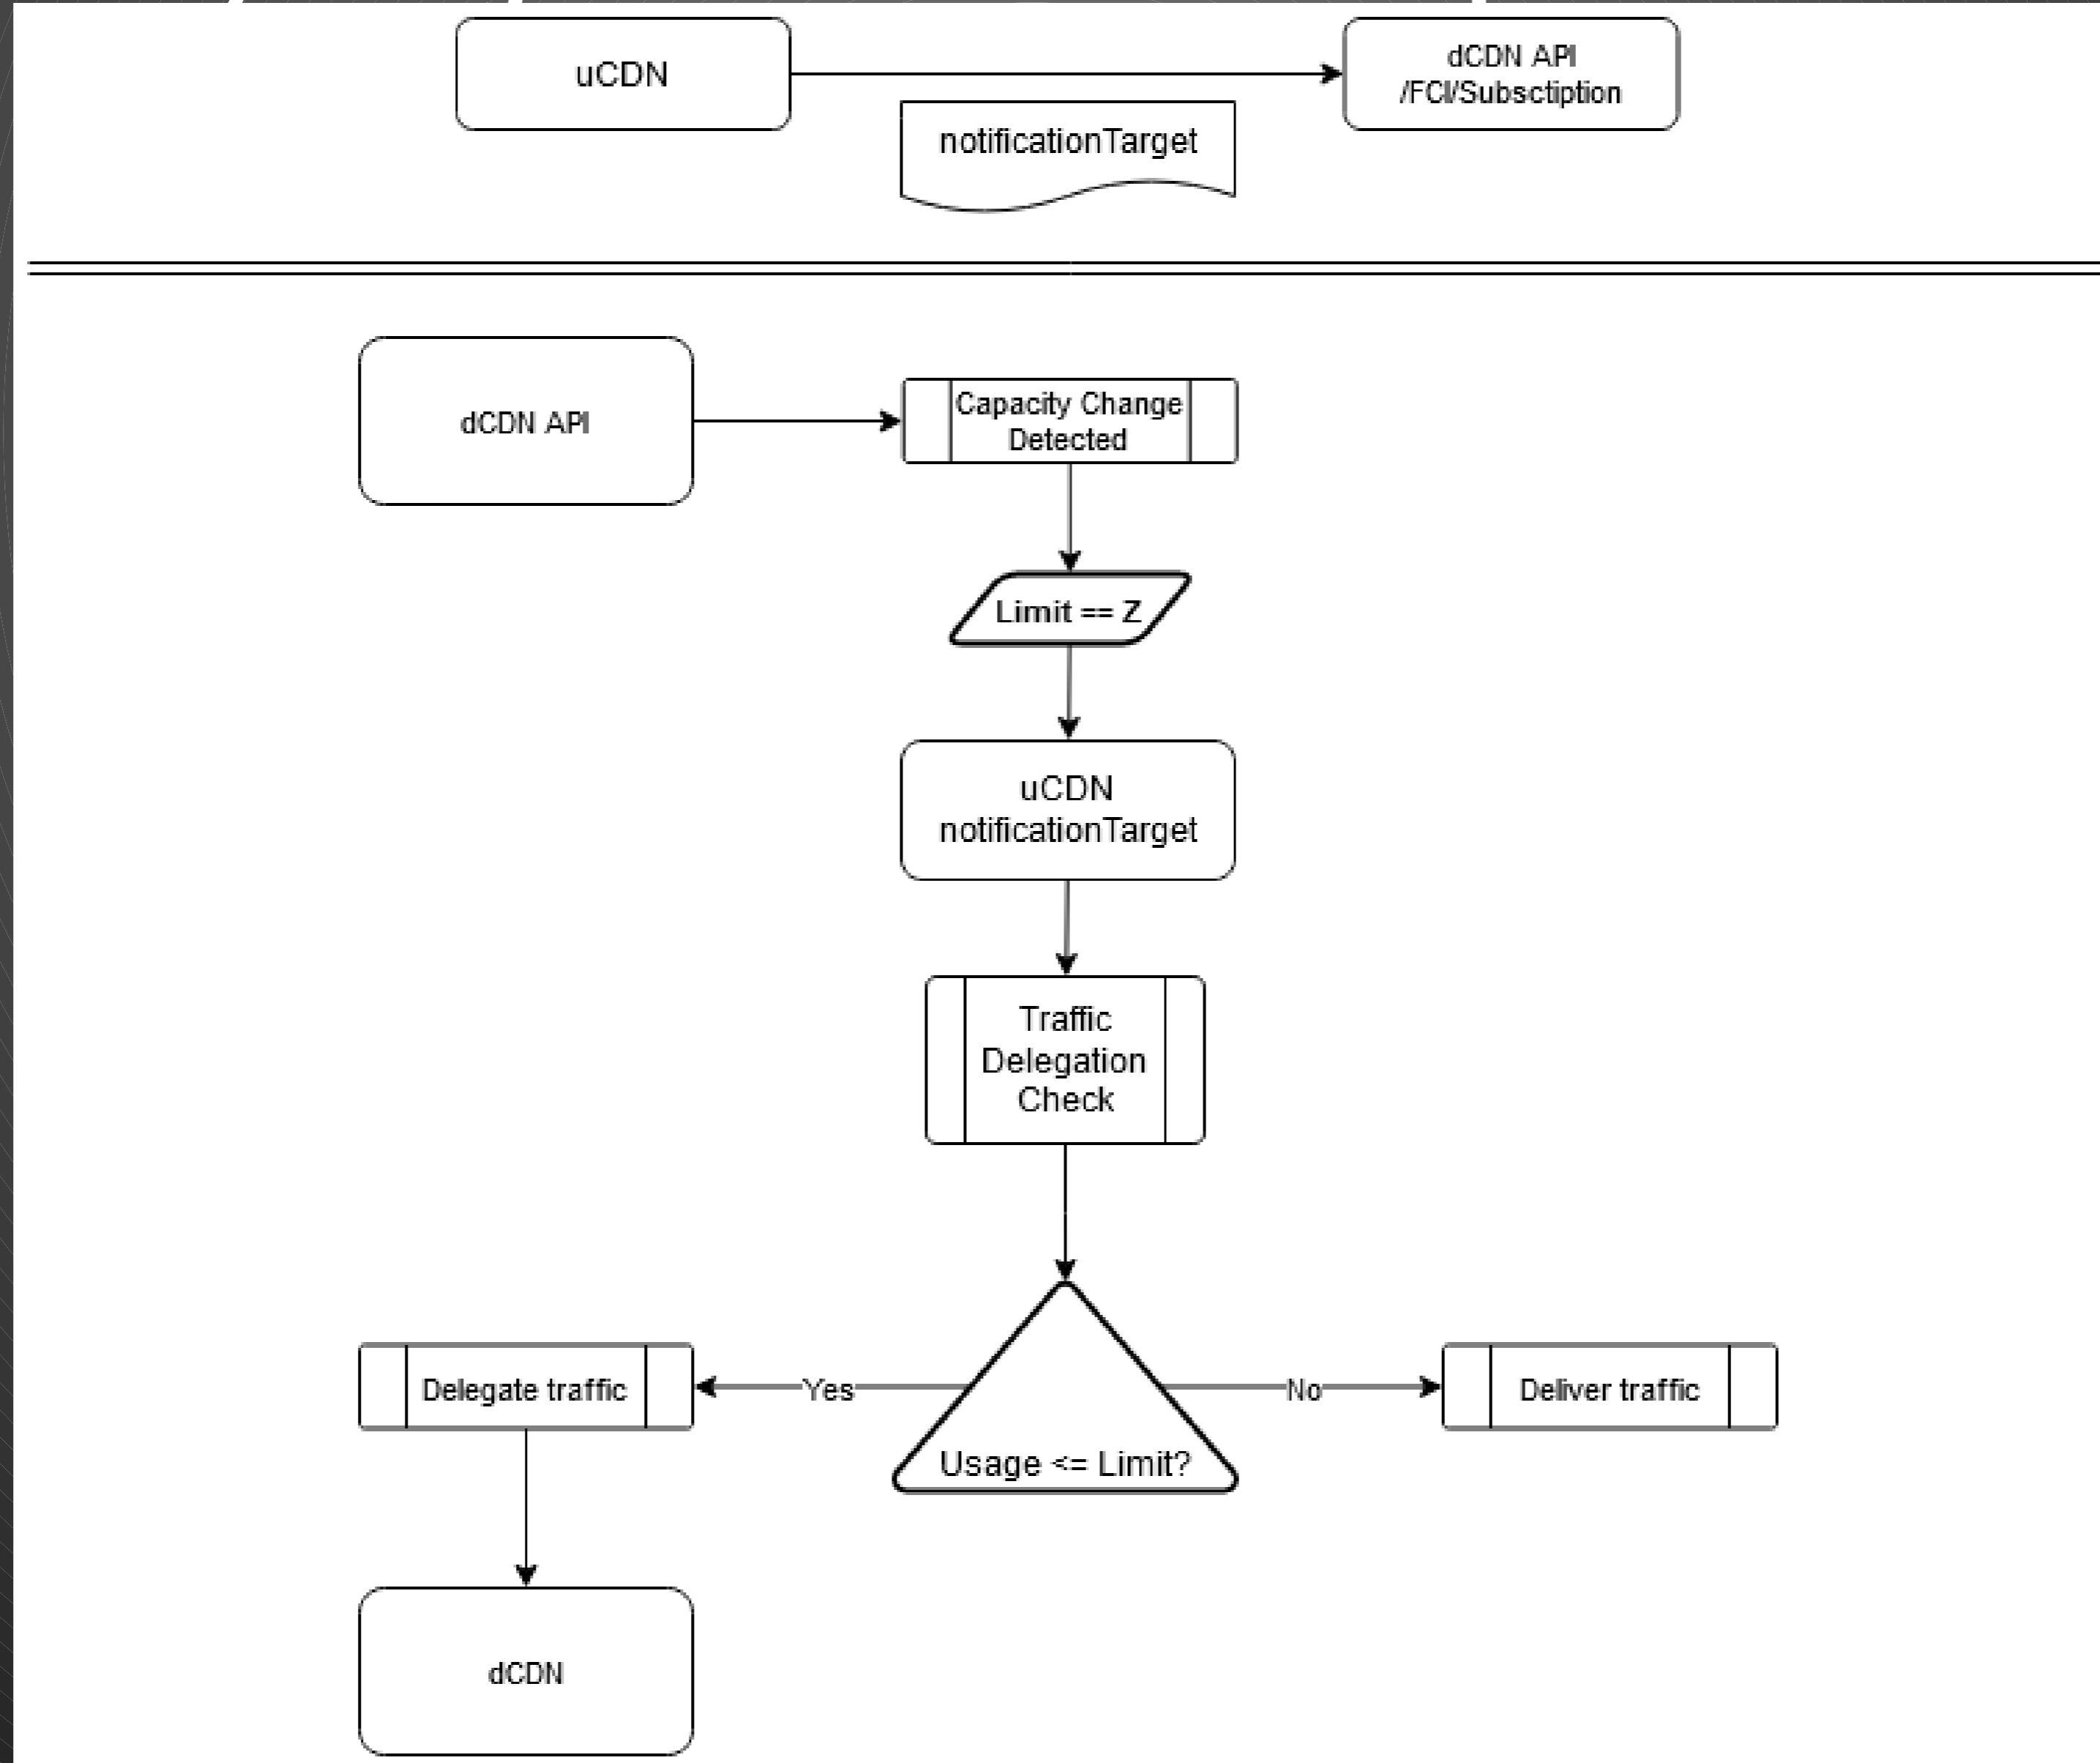
```
"capabilities": [  
  {  
    "capability-type": "FCI.Telemetry",  
    "capability-value": {  
      "sources": [  
        {  
          "id": "capacity_metrics_region1",  
          "type": "generic",  
          "metrics": [  
            {  
              "name": "egress_5m",  
              "time-granularity": 300,  
              "data-percentile": 50,  
              "latency": 1500  
            },  
            {  
              "name": "requests_5m",  
            },  
            ...  
          ],  
          "configuration": { ... optional ... }  
        }  
      ]  
    }  
  }  
]
```

- Lists the supported Telemetry Sources which include an id, type and metrics
- Telemetry data is uniquely identified within a source by id and metrics.name
- FCI.CapacityLimits reference Telemetry sources by id and metrics.name
- Each metric outlines key characteristics of telemetry data such as data latency, granularity, etc..
- Currently type is limited to 'Generic' and is meant to allow for advertisement of ad-hoc/preexisting data sources
- Types will eventually expand to include reference to a formal Telemetry interface
- While this is being defined within the Capacity specification, this is meant to provide a general way to advertise Telemetry data

Capacity Insights – Polling Workflow



Capacity Insights – Incremental Update Workflow

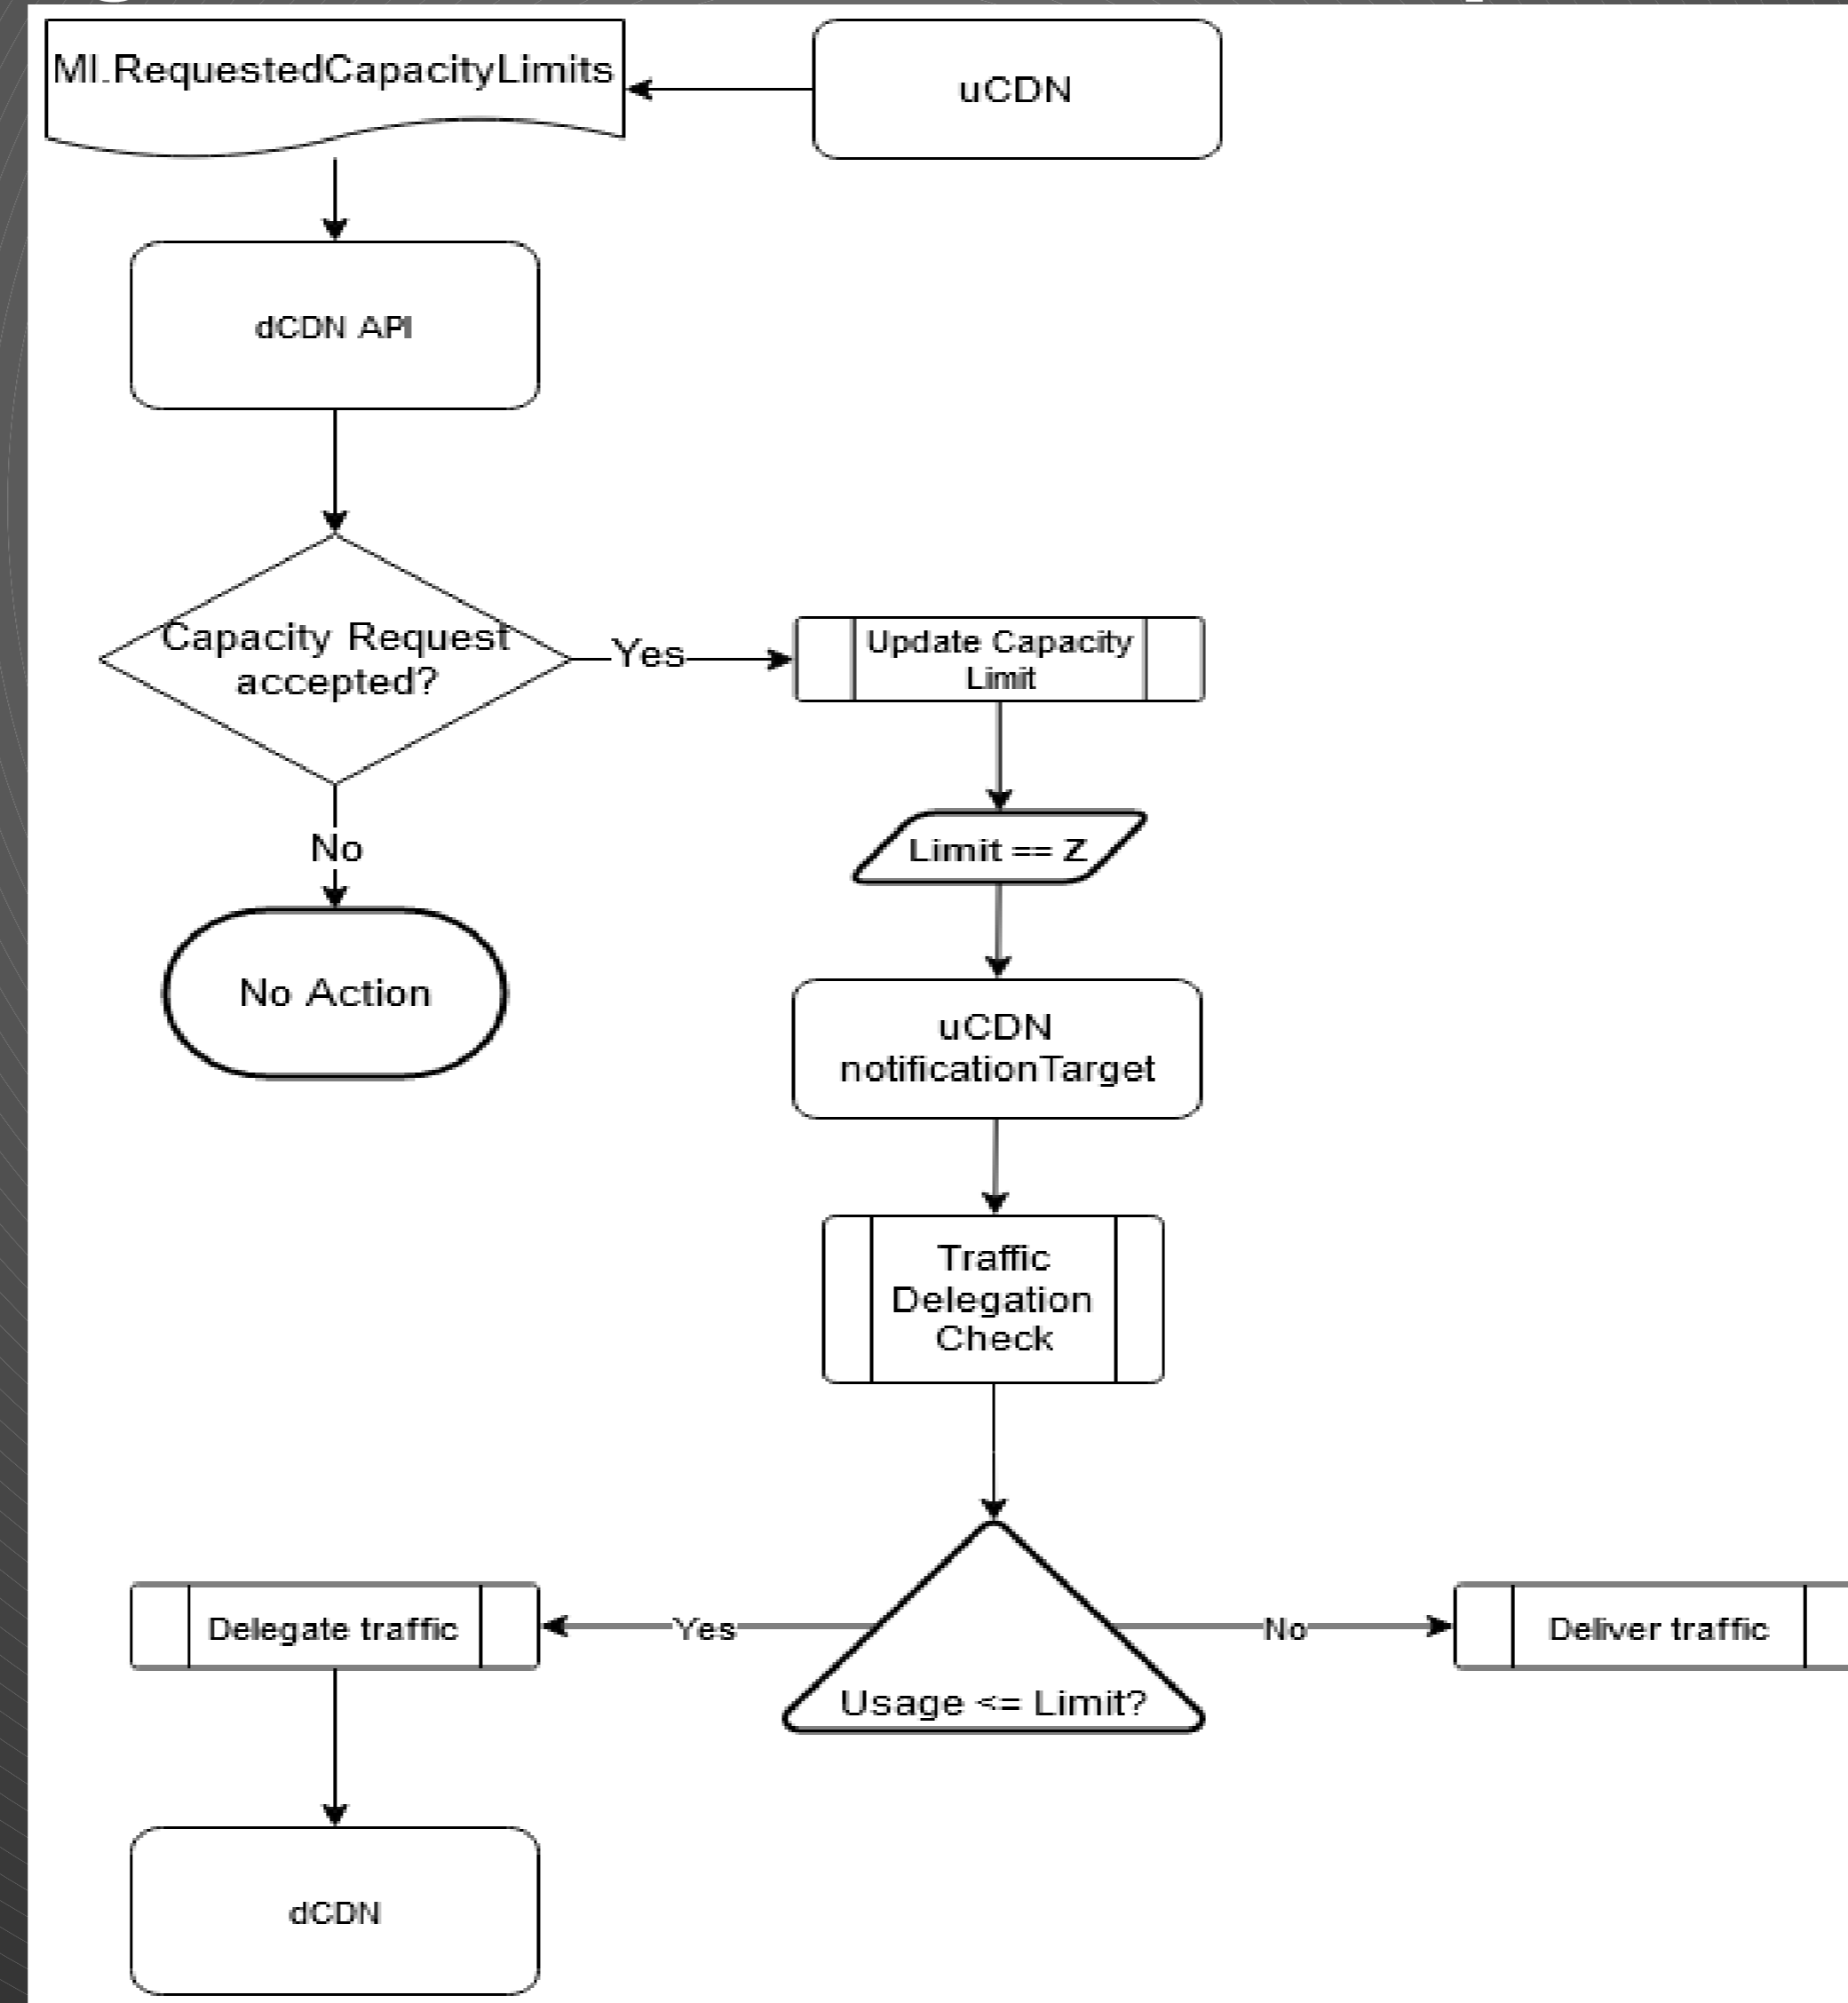


Capacity Insights – MI.RequestedCapacityLimits

```
{  
  "host": "serviceA.cdn.example.com",  
  "host-metadata": {  
    "metadata": [  
      ...,  
      {  
        "generic-metadata-type": "MI.RequestedCapacityLimits",  
        "generic-metadata-value": {  
          "requested-limits": [  
            {  
              "limit-type": "egress",  
              "limit-value": 8000000000000,  
              "footprints": [{  
                "footprint-type": "ipv4cidr",  
                "footprint-value": ["192.0.2.0/24",  
"198.51.100.0/24"]  
              }]  
            },  
            {  
              ...  
            }  
          ]  
        }  
      ]  
    }  
  }  
}
```

- Allows uCDN to ask dCDN to reconsider current CapacityLimit with suggested values
- Solicitations are specific to a CDN-Domain/Published host
- uCDN does a PUT to the dCDN Configuration endpoint
- dCDN considers request, and if it accepts sends an updated FCI.CapacityLimit to the uCDN
- If dCDN rejects, no action is taken

Capacity Insights – Solicitation for Update Workflow



Thank you!
