Host Status and Next Hop Overview

John Rushford - Comcast

Host Status

Parent selection consistent hash and round robin algorithms, along with host status determines route through the CDN

Host availability is changed with

transaction failures

- **transaction timeouts**
- **connection errors**

Tunable configuration parameters

- **failure thresholds**
- **timeout durations**
- **retries and retry windows.**

- **No external health checks**
- **No manual control for failed hosts or maintenance.**

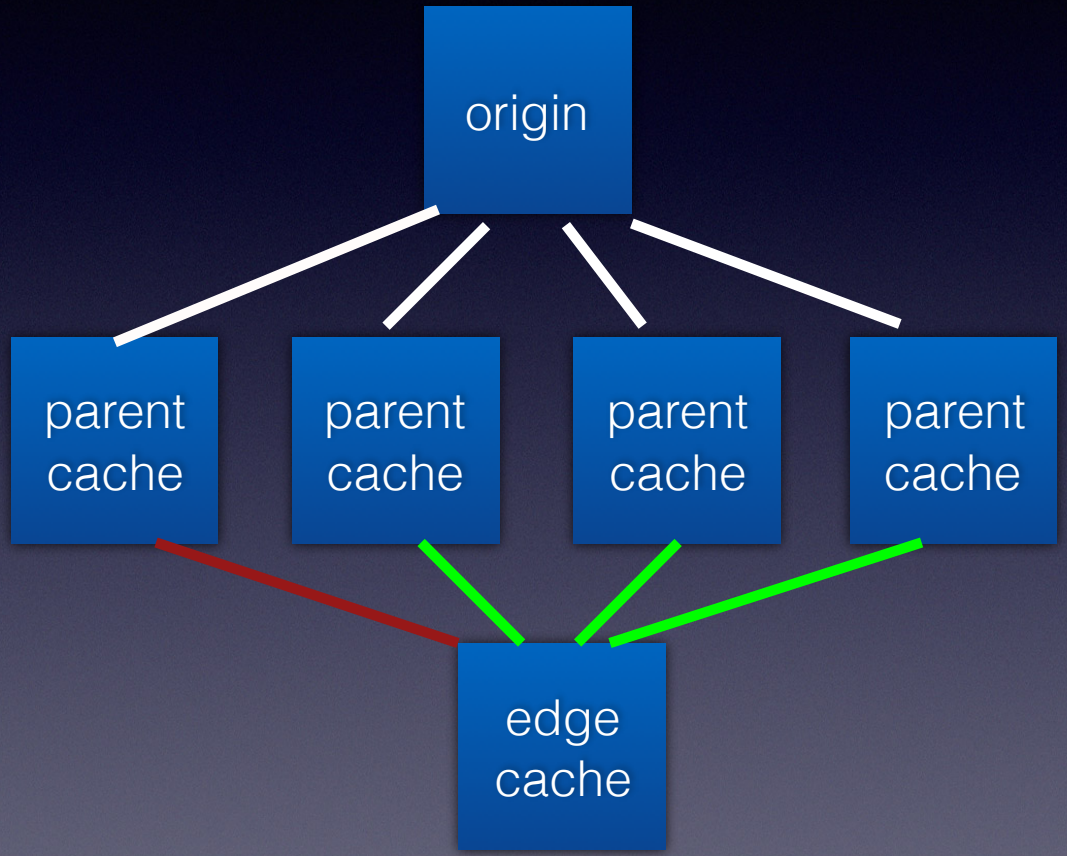

Host Status

John Rushford - Comcast
Vijay Mamidi - Apple

- A module used to check the availability of a upstream host for handling transactions
- Host Status updates are made by the operator or health checking tools when available.
- Parent Selection and HttpSM check HostStatus
- Completed - committed to master and in ATS 8.0

Host Status managed with traffic_ctl command

- Host sub-command added to traffic_ctl
- Retrieve a hosts status

```
# traffic_ctl host status HOSTNAME [HOSTNAME …]
```

- Mark a host up or down

```
# traffic_ctl host down –reason ‘manual | active | local’ –time seconds  HOSTNAME [HOSTNAME
…

# traffic_ctl host up –reason ‘manual | active | local’  HOSTNAME [HOSTNAME …]
```

Host Status

- When a parent proxy is marked down, next available parent proxy is chosen from the hash ring or round robin list.
- When an origin is marked down, objects are served stale from cache with a “Warning: 111” header indicating a revalidation as the origin is unavailable or a 503 is sent to the client on a cache miss.

proxy.config.http.cache.max_stale_age: 604800 (7 days by default)

Host status queries - traffic_ctl

 # traffic_ctl host status parent-proxy-cache-01.cdn.com

proxy.process.host_status.cache-1.cdn.com_active 1

proxy.process.host_status.cache-1.cdn.com_local 1

proxy.process.host_status.cache-cdn-cdn.com_manual 0

- Displays status by reason tag
- Reason tags are information
- **manual** - indicates the host status changed by the operator
- **active** - status changed by an active health check tool
- **local -** status changed by a health check plugin

Host status queries - astats

- You may obtain host availability details remotely through the “astats” or “stats_over_http” plugins

$ curl “http://edge-cache-01.cdn.com/_astats" | grep host_status

“proxy.process.host_status.parent-proxy-01.cdn.com_active”: “1”,

“proxy.process.host_status.parent-proxy-01.cdn.com_local”: “1”,

“proxy.process.host_status.parent-proxy-01.cdn.com_manual”: “0”,

traffic_ops - changes

- Traffic_ops and ORT writes and loads a new parent.config when parents are set to admin_down
- With ATS 8 - We should use traffic_ctl to mark parents down.

Questions?

- Host Status

Next Hop Overview

A new layer 7 framework used to provide an upstream connection.

The following modules:

- Flexible routing architecture configuration
- Upstream Generator
- Modular routing logic
- Shared data storage

Next Hop Working Group

- Alan Carroll - Oath
- Aaron Canary - Oath
- Vijay Mamidi - Apple
- John Rushford - Comcast

Upstream generator

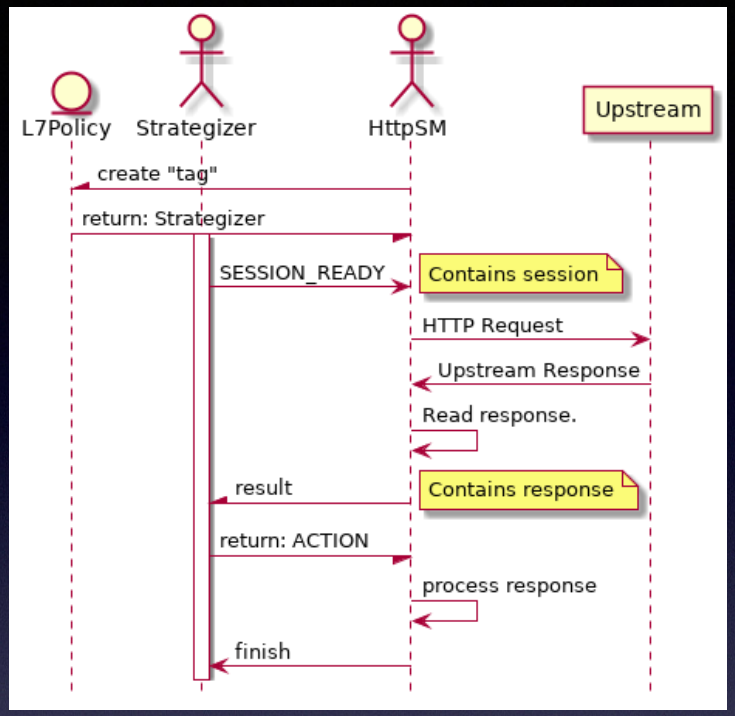

Generic transaction

Alan Carroll analysis notes - https://solidwallofcode.github.io/notes/L7R-Denver.en.html

- **Instantiated by the HttpSM per remap**
- **Provides transaction ready connections to the HttpSm**
- **Handles connection and timeout errors providing a new connection to the HttpSM as necessary.**

Shared data storage container

Aaron Canary - **https://cwiki.apache.org/confluence/display/TS/Presentations+-+2018?preview=/75958509/94798211/NextHop%20-%20Fall%202018.pptx**

- One in memory data storage container - for use with parent selection, HostStatus, HostDB and other modules.
- C++ Template Library
- Extendable - systems may extend the data schema with additional attributes as needed.
- One container to manage thread safety, indexing, hashing, and performance optimization.
- Work in progress with code committed to master - tscore/AcidPtr.h, tscore/Extendible.h

NextHop Routing Config

Analysis notes - https://github.com/apache/trafficserver/pull/3870

- YAML config file
- Define upstream host addresses and protocols
- Define upstream selection strategies and routing
- Define health check attributes
- Define transaction error and result handling.
- Some work has been done on a C++ yaml parser
- Upcoming change to use YAML for all ATS 9 config files - changes required in traffic_ops

parent hosts.yaml

```
# hosts.yaml
#
hosts:
  p1: &p1
    host: p1-cache.foo.com
    protocol:
      - http: 80
      - https: 443
    healthcheck:
      url: tcp://192.168.1.1:80
  p2: &p2
    host: p2-cache.foo.com
    protocol:
      - http: 8080
      - https: 8443
    healthcheck:
      url: tcp://192.168.1.2:80
  s1: &s1
    host: s1-cache.bar.com
    protocol:
      - http: 80
      - https: 443
    healthcheck:
      url: tcp://192.168.2.1:80
  s2: &s2
    host: s2-cache.bar.com
    protocol:
      - http: 8080
      - https: 8443
    healthcheck:
      url: tcp://192.168.2.2:80
groups:
  - g1: &g1
    - weight: 1.0
      <<: *p1
    - weight: 2.0
      <<: *p2
  - g2: &g2
    - weight: 0.1
      <<: *s1
```

strategies.yaml

```
#include unit-tests/hosts.yaml
#
strategy:
          policy: consistent_hash
          hash_key: path+query
          groups:
            - *g1
            - *g2
          protocol: http
          failover:
              ring_mode: exhaust_rings
              response_codes:
                  - 404
                  - 503
         health_check:
             - passive
```

Questions?

- NextHop