Draft-campbell-dime-load-considerations-01

IETF 92

DIME Working Group Meeting

Dallas, Texas

What is Diameter Load

- Current state of a ***Non-overloaded*** Diameter Node
- Used in the peer selection process
  - Load balancing across set of peers
  - To try to avoid overload in the first place
  - What about *server* selection? (More to come)
- Different than Overload
  - Load is information that the recipient can use as needed
  - Overload is a request for action
  - High load does not necessarily imply overload
    - … but might predict it

Topology of Load

- The draft describes a set of topology use cases
  - To identify relevant use cases regarding use of load information and how is load information shared
- Alternatives
  - Between immediate peers
  - Between endpoints (clients and servers)
  - Hybrid – Both endpoint and peer load shared
  - Between any arbitrary nodes

Topology of Load

- Support for peer load is needed
  - Support for Diameter networks where server selection is done by last-hop agents
- Need to determine if there are scenarios where endpoint load is also needed
  - For scenarios where server selection is done by topology-aware Diameter nodes (endpoints or agents) that are not a peer to endpoints
- Open Question – Impact of redirect agent scenarios

Load Topology Preliminary Recommendations

- Support peer load reports
  - Any node can send load information
    - Agent can optionally aggregate upstream load information into their own load metrics
- Optionally add support for endpoint load reports
  - Adds complexity
    - Multiple load reports in each message
    - Agents must manage which load reports get passed on
- General Non-Adjacent node reports “out of scope”
  - Endpoints are only non-adjacent load reports considered

Scope of Load Information

- What does a load metric describe?
  - Load of an entire host? (i.e. Diameter-Identity)
    - Allows a simple metric
  - Load of an application at a host?
  - Load of a realm?
  - Load of a group of hosts?
- Some of these would require some additional metadata to describe the scope of a load metric
- Metadata could be explicit or implicit
  - Potentially more similar to DOIC

Scope of Load Information Preliminary Recommendations

- Do NOT address load of a realm or set of hosts in the initial specification
- Default load report applies to a nodes load for an individual application
- AVP in load report indicating the node to which the report applies
- Optionally add indicator in the load report saying that it applies to all applications at that host

Precedence between Load and Overload

- How do load and overload information interact?
- Preliminary Recommendations
  - Overload information takes precedence
    - Ignore load metric from overloaded host or use load information when routing non abated requests
    - Only for the “scope” of overload
  - Load does not imply overload
    - A 100% loaded node is not necessarily overloaded
    - A 0% loaded node is cannot necessarily be assumed to not be overloaded.
    - But load can be used as an input to predict and prevent overload.

Load Information Semantics

- Load value range alternatives
  - 0 to 10 (as suggested in I-D.tschofenig-dime-dlba)
  - 0 to 100 (as suggested I-D.korhonen-dime-ovl)
  - 0 to 65535 (as suggested in I-D.roach-dime-overload-ctrl and consistent with DNS SRV defined in RFC2782)
    - Gives greatest level of granularity
    - Consistent with other load balancing implementations

Negotiation

- Do we need to negotiate or declare support for Load?
  - May not be necessary if a non-supporting node can ignore it
  - … But if Load is strictly peer to peer, we need a way to make sure it doesn't leak across a non-supporting node
    - Could be done by negotiating support
    - Or by adding Diameter-Identity metadata to load metric

Negotiation
Preliminary Recommendations

- Do not specify Load negotiation mechanism
- Add SourceID AVP to load reports to identify the node that added the report
  - SourceID also needed if multiple reports are allowed in a single message
- Leaking load information addressed in topology hiding implementations

Transporting Load

- How should load be transported?
  - Piggy-backed on existing messages?
    - Consistent with DOIC
      - Should it be integrated with DOIC?
  - A dedicated Diameter application?
    - Easier to negotiate
      - Could use standard capabilities exchange if peer to peer
      - Could use subscription model if not peer to peer.
    - Creates additional traffic just to carry load info
  - Something else?
    - Something completely out of band?
    - (e.g. Web interfaces, SIP Events, etc.)

Transporting Load
Preliminary Recommendations

- Piggy-backed
- Not part of DOIC
  - It should be possible to use load without using DOIC

Frequency of Sending Load Information

- Alternatives:
  - Send load information in every message.
  - Send load information when it changes by some amount.

    For instance, only send a new load report when the load value has changed by some percentage.
  - Send load information every interval of time.

    With this approach, load information would be sent every some number of seconds.

Frequency of Sending Load Information

- Interacts with method of transporting information
- Some specification required if nodes are allowed to not send load information in every answer message
- Could put requirements on load capability announcement/negotiation

Frequency of sending Load Information

- Recommendations TBD

Summary of Preliminary Recommendations

- Load supports peer reports
  - TBD if Load also supports endpoint reports
- TBD on level of specification of upstream load aggregation
- Any node can send load
- SourceID in load reports to identify sender of load report
- Load and Overload are separate
- Load metric is scoped to an application at a host
  - TBD if there is mechanism to indicate scope is entire host
- Piggybacked on existing messages
- No explicit declaration of support
- Frequency of load information is TBD

Next Steps

- Add to working group charter
- Continue work on defining mechanism