CONNECT-IP

draft-cms-masque-connect-ip-01

Alex Chernyakhovsky – achernya@google.com
Dallas McCall – dallasmccall@google.com
David Schinazi – dschinazi@google.com

CONNECT-IP

- Allows endpoints to set up an IP tunnel between one another

- This can be used to implement:
  - A consumer VPN
  - A point-to-point IP tunnel (e.g., overlay mesh)
  - A point-to-network IP tunnel (e.g., corporate VPN)
  - A network-to-network IP tunnel (e.g., Site-to-Site VPN)

- Use-cases detailed in draft-ietf-masque-ip-proxy-reqs-02

Datagrams

- IP Packets are conveyed using HTTP Datagrams
  - Either QUIC DATAGRAM frame
  - or DATAGRAM capsule

- Payload is full IP packet, from the IP Version field until the last byte of the IP Payload.
  - Extensions may add additional formats using REGISTER_DATAGRAM_CONTEXT

Aside: Forwarding packets

- CONNECT-IP endpoints act as IP routers
- When receiving an encapsulated packet
  - Parse IP header
  - Check local policy (e.g., source address validation)
  - Transmit using implementation-specific mechanism
- When receiving a CONNECT-IP tunnel-bound packet
  - Perform same forwarding checks as above
  - Transmit inside the tunnel

ADDRESS_REQUEST & ADDRESS_ASSIGN

- `ADDRESS_REQUEST`: used to request an IP address

```
ADDRESS_REQUEST Capsule {
  IP Version (8),
  IP Address (32..128),
  IP Prefix Length (8),
}
```

- `ADDRESS_ASSIGN`: used to assign an IP address to the peer

```
ADDRESS_ASSIGN Capsule {
  IP Version (8),
  IP Address (32..128),
  IP Prefix Length (8),
}
```

ROUTE_ADVERTISEMENT, ROUTE_REJECTION, ROUTE_RESET

- These allow an endpoint to communicate what routes they have to their peer
- When an endpoint receives it, it decides whether to use the route or not
- Can apply local policy

ROUTE_ADVERTISEMENT, ROUTE_REJECTION, ROUTE_RESET

- ROUTE_ADVERTISEMENT: Informs peer it's willing to route traffic to prefix.

```
ROUTE_ADVERTISEMENT Capsule {
  IP Version (8),
  IP Address (32..128),
  IP Prefix Length (8),
}
```

- ROUTE_REJECTION message allows an endpoint to communicate to its peer that it is not willing to route traffic to a given prefix. Message body is identical to ROUTE_ADVERTISEMENT

- ROUTE_RESET message allows an endpoint to cancel any routes it had previously advertised or rejected.

ATOMIC_START & ATOMIC_END

- Allows applying all configuration in one step, avoids leaking prefixes
- `ATOMIC_START`: used to begin a series of atomic set of messages.

```
ATOMIC_START Capsule {
}
```

- `ATOMIC_END`: used to end a series of atomic set of messages.

```
ATOMIC_END Capsule {
}
```

SHUTDOWN

- SHUTDOWN: used to indicate to peer that the endpoint is closing the CONNECT-IP stream, with a string explaining the reason for the shutdown.

```
SHUTDOWN Capsule {
  Reason Phrase (..),
}
```

CONNECT-IP Extensibility

- Three main extension mechanisms
  - HTTP Headers, Capsule types, Datagram context extensions
- Examples of possible extensions
  - Configure DNS server IPs as a Capsule
  - Authentication/Authorization data as an HTTP Header
  - Sending only IP Payload on the wire (compression of IP header)
    - HTTP Datagram context to map context ID to IP headers

```
STREAM(44): CAPSULE                 -------->
  Capsule Type = REGISTER_DATAGRAM_CONTEXT
  Context ID = 2
  Context Extension = {IP_COMPRESSION=tcp,192.0.2.6:9876,192.0.2.7:443}

DATAGRAM                            -------->
  Quarter Stream ID = 11
  Context ID = 2
  Payload = Compressed IP Packet
```

CONNECT-IP

draft-cms-masque-connect-ip-01