gRPC on Network Devices

IETF 101, London UK

robjs@google.com

Why?

- Available APIs for configuration and telemetry from network devices are fragmented, proprietary or overly complex.

- Build on modern RPC framework to define a standard set of services for device interaction.
  - Setting configuration from a management system.
  - Receiving device telemetry - **streaming telemetry** in a scalable way.
  - Running operational commands.
  - Injecting routing entries.

- Publish specifications, tooling and reference code to allow consumption of these APIs.

Service Landscape.

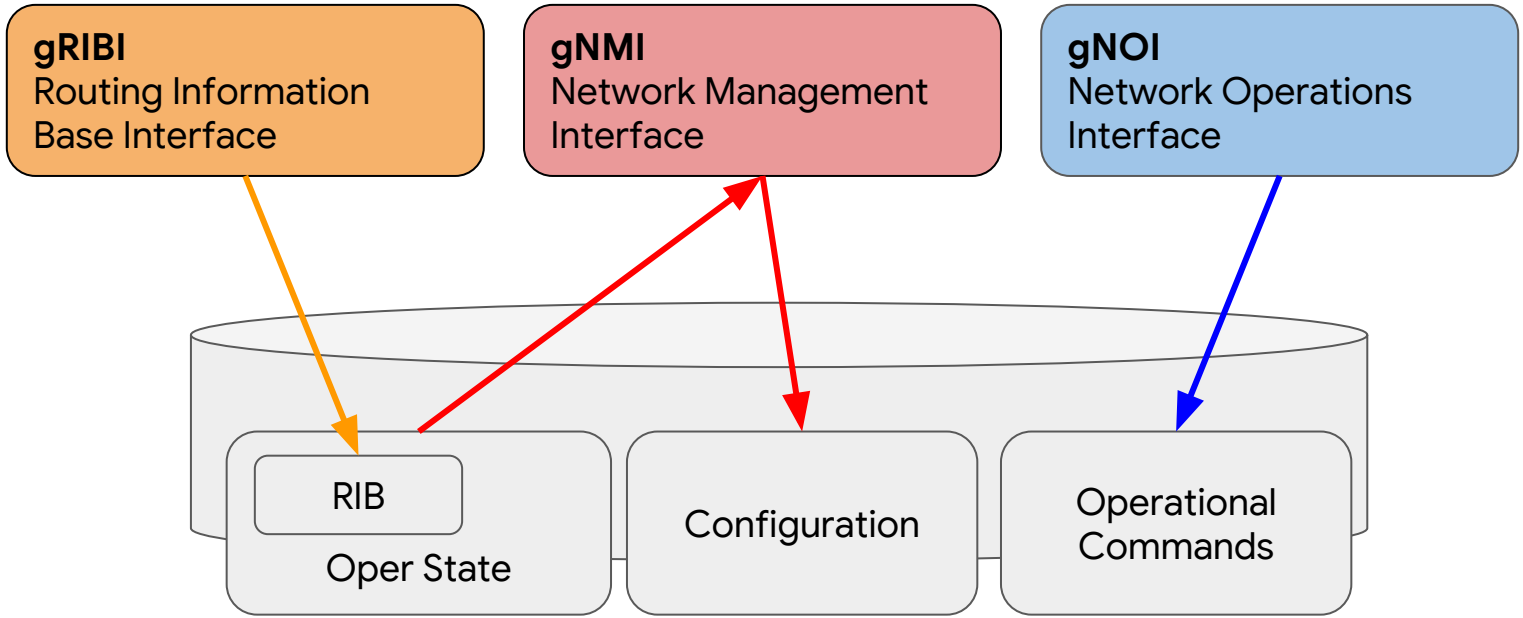

gNMI

- Single service for state manipulation and retrieval of a data tree.
  - Could be YANG, but doesn't have to be.
- Set RPC used for manipulating device state.
  - One RPC = one transaction - no long-lived candidates.
  - Encoding can be ASCII, JSON, or Protobuf.
- Subscribe RPC used for telemetry.
  - Pushed by the device - **streaming telemetry** - in some modes (STREAM)
  - Or requested by client (ONCE, POLL)
  - Streams can carry event-driven or sampled path updates.
  - Efficient on the wire encoding approaches to support high-volume.
  - Target timestamped.

gNOI

- Suite of microservices - each corresponding to a set of operations.
  - Allows adoption of only the services that the device supports.
  - Reflection service (in gRPC library) can be used to discover which services a device supports.

- Growing coverage, today:
  - BGP, Certificate management, MPLS, interface, layer 2, system (ping, traceroute etc.)

- Natively described in protobuf.
  - No YANG model for operations contents.
  - Path within data tree used to relate to other state on the device.

gRIBI

- Single service used to inject entries into the RIB of a network device.
  - Bi-directional streaming RPC over long-lived channel to modify (add, modify, delete) entries on the device.

- Motivation is for a simple interface to add entries to device's RIB without overloading existing protocol semantics.
  - gRIBI service on device essentially becomes another protocol RIB client.
  - Coexists with other protocols - gRIBI entry can resolve via another protocol's entry.

- Schema within the protocol is OpenConfig "abstract forwarding table".
  - Payload is binary encoded protobuf.
  - YANG model is machine translated into a protobuf for use within the protocol.

Implementations

- gNMI running in production - at least 6 vendors have implementations.
- gNOI and gRIBI in engineering code from multiple implementations.

- Open source reference gNMI implementation, telemetry collector on GitHub.
- Stratum project in ONF builds on gNMI and gNOI alongside P4 runtime - will publish reference code.

- Adoption of gRPC internally to multiple implementations - simplifying implementation.