Introduction to the Routing Area

IETF 100 (November 2017, Singapore)

Joel Halpern (jmh@joelhalpern.com)

Sue Hares (shares@ndzh.com)

Martin Vigoureux (martin.vigoureux@nokia.com)

What We Want to Tell You

- We want to give you an overview of the breadth of work covered in the Routing Area
- We want to show how the work is divided between…
  - Support of core protocols without which the Internet would not operate
  - Applications of those protocols,
  - Specialist routing protocols for niche environments
  - Experimentation in new routing technologies
- We will do this by walking you through the list of working groups in the area

What We're Not Going to Tell You

- This is not a presentation about how routing works
- And it is not a discussion about how to design a routing protocol
- We have no plans to tell you whether OSPF is better or worse than IS-IS

History

- Routing has been recognized as a core division of the IETF's work from the beginning
  - In 1989 there were just 6 ADs
    - OSI co-existence (x2)
    - Internet Services
    - Network Management
    - Routing
    - Host-Based

Some Numbers

- IETF has seven Areas
  - ART, GEN, INT, OPS, RTG, SEC, TSV (14%)
- IETF has 15 Area Directors
  - 3, 1, 2, 2, 3, 2, 2 (20%)
  - Some ADs take responsibility for WGs in other Areas
- IETF has 130 working groups
  - 36, 1, 19, 16, 27, 16, 12 (21%)
- IETF published 250 RFCs in 1 year to Nov 7th, 2017 [1]
  - 218 WG: 47, 0, 30, 27, 68, 22, 24 (31%)
  - 32 non-WG (AD sponsored) [13%]

[1] http://www.rfc-editor.org/search

What is Routing?

- Hosts are not all directly connected to each other
- This means (IP) packets must be forwarded hop-by-hop across the Internet
- Routers receive packets on one interface and decide which interface to forward them out of
  - This is routing; the path followed by a packet is a route
- Routes are either known in a distributed fashion
  - Each router determines the next hop towards a destination from information about the network and an algorithm
- Or they are known in a programmed way (whole route predetermined)
- Routing protocols distribute information about the network or about pre-determined routes
- The Routing Area concerns itself with protocols and mechanisms to route packets, and with uses of those protocols

When is Routing Not Routing?

- There are problems in the IETF that are very similar to classic routing problems
  - Finding paths across a graph to deliver data
  - But they are not about delivering or routing packets
- Sample work in other Areas
  - Content Delivery Networks Interconnection (CDNI - ART)
  - INtermediary-safe SIP session ID (INSIPID - ART)
  - Session Initiation Protocol Core (SIPCORE - ART)
  - Application-Layer Traffic Optimization (ALTO –TSV)
  - Multipath TCP (MPTCP – TSV)
- RTG Area is largely not involved in this work
  - May do some informal review
  - Can give advice:
    - “problems already solved”
    - “things that may bite you later”

Why is the Routing Area so Hard to Schedule?

- There is a strong inter-relationship between many of the RTG WGs
  - Many routing technologies build on core routing protocols
  - Many routing protocols are complementary and need to work together
  - Some routing protocols address the same problem spaces
  - There is a relatively small core set of “routing experts”
  - There are 25 working groups, a few of which ask for more than one meeting session
  - There are usually less than 17 meeting slots (16 in SIN)
- Means that some meetings “conflict”
  - You have to choose where to go

The Nature of Routing Working Groups

- Two broad categories
  - Maintenance mode
    - Old WGs for long-established protocols
    - Usually plenty of new extensions, clarifications, bug-fixes
    - No indication that these will ever close!
  - New work
    - New ideas for specialist protocols or routing applications
    - Should be more “normal” as working groups
      - Deliver on charter and close down

Sub-Divisions in the Routing Area

- Core Routing Protocols
- Specialist Routing Protocols
- Sub-IP
- Routing Support and Operation
- Routing Services
- Experiments
- Closed but not forgotten!

Core Routing Protocols

- These are the protocols that are fundamental to how the Internet works today
- The working groups are mostly in “maintenance mode”
  - This does not mean that there is no new work
  - It does mean that the protocols are well-established and widely deployed
- New work is treated with a high degree of caution
  - We really do not want to break the Internet

OSPF

Open Shortest Path First IGP

- One of the two shortest path first (SPF) interior gateway protocols (IGPs) in wide use
- Work is on maintenance of OSPFv2 (for IPv4)
- Focus is moving to OSPFv3 (for IPv6 and IPv4)
- Extensions for a wide range of features
  - More routing metrics, Better scaling
  - More link/node characteristics
  - Support for other working groups (MPLS, CCAMP, SPRING, BIER)
  - Support for segment routing

ISIS

IS-IS for IP Internets

- Intermediate System to Intermediate System is an old ISO routing protocol
  - The IETF took over the specification of IS-IS for IP and published RFC 1195
- Much of the work mirrors that done in OSPF
  - Except that a new version was not needed to support IPv6
  - Extensions are also made for the same features and purposes
    - Sometimes sooner and sometimes later than for OSPF

IDR

Inter-Domain Routing

- The Border Gateway Protocol (BGP) is sometimes described as the glue that holds the internet together
  - The WG is probably the most conservative of all IETF WGs
  - Requires two independent and interoperable implementations before any protocol extension is published as an RFC
- Essentially in “maintenance mode”, WG works on protocol extensions to make the global routing system work more smoothly and scale better
  - GROW WG suggests additions,
  - BESS + Spring have protocol additions
- Two important change BGP-LS + Flow Specification tha
  - BGP-LS allows the “export” of routing information (TE) from nework to a management systems (for example PCE element)
  - Flow Specification controls flow within network

SIDR

Secure Inter-Domain Routing

- The Internet routing system depends on BGP
- The stability and resilience of routing tables used by BGP is under threat
  - Accidental “fat fingers” or Deliberate “route hijacking”
- This WG is tasked to develop a mechanism to sign route advertisements when they are originated
  - Requires a public key infrastructure
  - Requires a way to sign routes
  - Requires a way to distribute keys
- **WG has completed active work (just waiting for publication) - sidrops is working on operational issues**

PIM

Protocols for IP Multicast

- There used to be several competing protocols for multicast
  - Protocol Independent Multicast - Sparse Mode (PIM-SM) "won"
- Took over responsibility for IGMP and MLD
  - Used to be in INT Area
  - Puts all multicast expertise in one place
  - Very close collaboration with MBONED (OPS)
- Also a "maintenance mode" working group
  - Finalized work to advance PIM specification to Internet Standard
  - Improving authentication and scaling of PIM

SPRING

Source Packet Routing in Networking

- A new working group with a new look at an old concept
- Packet carries the waypoint that it should traverse
  - Compare with IP source route option
- Currently being worked on for MPLS and IPv6
  - Work on architecture and protocol extensions coming to an end.
- Complementary building blocks being worked on (conflict resolution, yang, …), and emergence of new applications for SR.
- Routing protocol extensions (OSPF, IS-IS, BGP) happening in the respective working groups
- Coordinates with MPLS and 6MAN (INT).

Specialist Routing Protocols

- Most routing protocols are general for IP in any environment
  - This has been part of the success of the Internet
- Some environments demand very specialized routing protocols
  - The devices may be exceptionally constrained
  - The cost of sending routing updates may be very high
- These specialist problems give rise to working groups targeted at niche environments

BABEL

Babel Routing Protocol

- Babel focuses on networks where some or all links have unstable metrics. For example, networks in where a wired and wireless mesh networks are combined
  - Babel is an augmented distance vector protocol
  - Currently specified in Experimental RFCs
- Babel WG is focused
  - Upgrading specification to standard track
  - Yang data models for management
- Babel is the mandatory to implement routing protocol for HOMENET WG (INT)

MANET

Mobile Ad-hoc Networks

- A MANET includes routers and hosts that may be mobile and that may come and go
  - Consider battle-field environments, emergency response radio systems, or the Internet in the developing world
- MANET protocols are used in niche environments including community networks across Europe
- Outstanding work items include…
  - DLEP : A protocol to report link characteristics to routers
  - A number of extensions to OLSRv2 : A link state protocol
  - Enhanced security and manageability for MANETs

ROLL

Routing Over Low-power and Lossy networks

- The Internet of Things (IoT) poses a new set of routing problems
  - Networks may be ad-hoc as in MANET
  - But devices may be extremely constrained in CPU, Power availability, Memory, etc.
  - Additionally, links may be subject to high degrees of interference
- The WG developed a new protocol called RPL
- Work now focused on special cases…
  - Multicast
  - Compression of routing information
  - Deployment and implementation advice for different environments
    - Factory
    - Domestic
    - Public space
    - Office

DetNet

DETerministic NETworking

- Focuses on deterministic data paths
  - Over Layer 2 bridged and Layer 3 routed segments
  - Provide bounds on latency, loss, and packet delay variation (jitter); and high reliability
  - Data plane will be compatible with the work done in IEEE802.1 Time Sensitive Networking (TSN)
  - Will use IP and/or MPLS to support a method of data plane flow identification and packet forwarding
- Use cases include
  - Pro audio and video
  - Electrical Utilities
  - Building Automation Systems
  - Industrial M2M and Wireless
  - Cellular Radio Front, Mid and Back Haul

Sub-IP

- Sub-IP was, for a short time, a sub-area with its own Area Director
- Covers routing and signaling protocols for forwarding technologies that lie below IP
  - MPLS
  - Layer 2
  - Optical technologies

MPLS

Multiprotocol Label Switching

- One of the largest and most prolific working groups
- MPLS is now almost as successful as IP and Ethernet
  - Nearly all IP traffic traverses an MPLS network somewhere along its path
- The working group has progressed key technologies
  - Label Distribution Protocol (LDP)
  - Resource Reservation Protocol for Traffic Engineering (RSVP-TE)
  - Extensions to OSPF and IS-IS for Traffic Engineering
  - MPLS Transport Profile (MPLS-TP)
  - MPLS OAM
- Generic extensions to RSVP-TE, OSPF-TE, and IS-IS-TE have now moved to TEAS
- While certain aspects of the technology are in “maintenance mode”, the WG still tackles new work (e.g. flow identification) and generates at least 2-3 RFCs per meeting cycle
- Possible new work includes refinements for OAM, security, forwarding plane protection mechanisms

CCAMP

Common Control and Measurement Plane

- Responsible for Generalized Multiprotocol Label Switching (GMPLS)
  - Extensions and generalizations to RSVP-TE and OSPF-TE for non-MPLS uses
  - Largely thought of as signaling and routing for optical technologies
    - Lambda switching, TDM, OTN, flexi-grid
    - Also covers Ethernet and MPLS
- Generic extensions to RSVP-TE, OSPF-TE, and IS-IS-TE have now moved to TEAS
  - Leaves CCAMP with technology-specific work
- The current work includes GMPLS extensions to B100 OTN, FlexE, and Yang models for non-packet technology-specific networks (OTN, WSON, Flexigrid, Microwave).

L2TPEXT

Layer 2 Tunneling Protocol Extensions

- A seasonal working group with active and dormant times, that exists to extend and maintain the Layer 2 Tunneling Protocol (L2TP) as necessary
- Recent RFCs include S-BFD for L2TPv3, and Keyed IPv6 Tunnel
- Currently working on YANG models for Keyed IPv6 Tunnel

TEAS

Traffic Engineering Architecture and Signaling

- A new working group formed to off-load some of the work from MPLS and to coordinate the work of MPLS and CCAMP
- Handles high level architectural views of TE
- Produces generic extensions to TE protocols
  - RSVP-TE, OSPF-TE, and IS-IS-TE
- Has oversight of protocol work from MPLS and CCAMP to see whether it should be generalized
- Ongoing work includes:
  - Modeling TE specific data (Topology, Tunnels, RSVP-TE)
  - SDN Control of TE Networks: Abstraction and Control of TE Networks
  - Segment-Routing and RSVP-TE Co-existence
  - Protocol refinements (“maintenance mode” work)

TRILL

Transparent Interconnection of Lots of Links

- Moved to RTG from the INT Area
- A protocol transparent to end stations operating above bridging but below IP routing that provides load spreading and optimal unicast forwarding
  - Supports multi-destination traffic and active-active connection
  - Supports arbitrary topologies and link technologies
  - Uses IS-IS as control plane
- Currently working on…
  - Multi-topology, multi-topology,
  - Data Center Extensions, Directory service additions
- This set of work to be completed by March 2018

Routing Support and Operation

- In order that routing protocols can work well they need support from operational and management tools
- Operations, Management, and Administration (OAM) is a set of tools that monitor and report on the behavior of traffic flows, connections, and links
- Other management tools enable configuration and operation of the routing system through…
  - Reading information about the network
  - Injecting information into the routing system
  - Programming the routing system to behave in specific ways

BFD

Bidirectional Forwarding Detection

- “This will be a short-lived working group lasting only around nine months”
- BFD is a liveness monitoring OAM tool
  - Are my packets getting through?
  - Is my link / tunnel up?
- Closely coordinated with the MPLS WG
- Also some interaction with the core routing protocol working groups
- Current focus on…
  - Multicast
  - Seamless BFD for end-to-end monitoring

I2RS

Interface To the Routing System

- Software Defined Networking (SDN) and Data Center automation have focused on the interface from the routing to the physical forwarding components
- I2RS is at a higher level interface to the routing system.
- Examples include:
  - Installing routes into the Routing Information Base
  - Tracking network topologies,
  - Programming route admission policies for forwarding or BGP engine
- The WG has chosen YANG as its modeling language with the Revised Data Store Concepts.
- WG has Yang models

PCE

Path Computation Element

- Originally conceived as an off-board tool for computing paths in multi-domain Traffic Engineered MPLS networks
- Now extending its scope to be active network management tool
- The WG mainly works on extensions to the PCE protocol (PCEP)
  - Handling sophisticated computation requirements
    - Multiple protection paths
    - Complex constraints (such as for optical networks)
  - Reporting network events
  - Supplying unsolicited updates to previously requested paths
  - Requesting new paths to be set up
  - Encompassing segment routing
- Future use cases and protocol may come from:
  - Proposal in TEAS WG to use PCE as central controller
  - 6TiSCH WG in INT + DetNet WG in RTG Area

Routing Services

- Many WGs in RTG focus on the use of existing protocols to enable new services
- Historically this has been seen in…
  - Layer 3 VPN
  - Layer 2 VPN
  - Pseudowires
- There is a recent increase in the number of new ideas in this area
- There has also been some recent consolidation of WGs

BESS

BGP Enabled ServiceS

- Formed from parts of the L3VPN and L2VPN WGs
- Any service (but especially a VPN) achieved using BGP
  - Major focus is on EVPN, though MVPN still generates work
- Close coordination with IDR for BGP extensions
- Coordination with…
  - MPLS for architectural considerations
  - NVO3 for data center VPNs

PALS

Pseudowire and LDP-enabled Services

- Formed partly from L2VPN WG and partly from PWE3 WG
- Any service enabled by LDP including…
  - Layer 2 VPNs including data center VPNs
  - Pseudowire services (transporting Layer 1 and 2 services over an IP and/or MPLS network)
- Any form of Pseudowire service
  - IP, MPLS, L2TP
  - Pseudowire encapsulations

NVO3

Network Virtualization Overlays

- Develop protocols/protocol extensions that enable network virtualization over IP within a data center
- Progress was slow, so to expedite it the working group has pioneered new meeting formats
  - Round table discussion/debate
- A lot of time focusing on new or proprietary encapsulations
- Security and control plane are also hot topics
  - Some distributed control plane work off-loaded to BESS

SFC

Service Function Chaining

- Arguably not a classic routing problem
- Work concerns directing traffic flows through service function nodes to apply features
  - policing, access control, security, load balancing
  - Where applicable, TCP proxies, transcoders, ...
- Produced RFC 7665 on the desired architecture
- Network Service Header (NSH) work near completion.
- Now working on associated topics such as improved security and OAM mechanisms

Successful Experiments

- Sometimes in routing we act a bit cautiously
- New ideas need to be given space for experimentation, but we don't want to qualify them as Proposed Standards until we know how they behave
  - PIM is a good example of a successful experiment that was moved onto the Standards Track
- There are currently two working groups in RTG tasked with producing Experimental RFCs

BIER

Bit Indexed Explicit Replication

- A new take on an old idea
  - Give every node in the network a bit in a bitmask
  - Indicate on each packet the intended recipients
  - Use routing protocols to build next-hop trees
  - Replicate packets as necessary
  - (Of course, it is a little more complicated than that)
- One challenge is whether this can be achieved without replacing all of the routers in the Internet
- This is a new and enthusiastic working group
  - Architectures and protocols are under discussion

LISP

Locator/ID Separation Protocol

- Relatively old work coming out of the Internet Research Task Force (IRTF)
- Originally conceived to handle the explosive growth of the global routing table
- Now looks at a large number of “layering” or “overlay” scenarios best typified by VPNs
- The working group is close to producing Standards Track documents for the overlay uses of LISP.
- The work has an enthusiastic core of supporters

Catch-All and Specialist Work

- There is important work in the RTG Area that does not fit into any of the WGs just described
- Some of this work is advanced under the care of the AD
  - Published as AD-sponsored RFCs
  - Open discussion on the routing-discussion mailing list

RTGWG

Routing Working Group

- Looking at overall architectural approach to DC, 5G, potentially MEC, SD-WAN, ...
- Some pieces of routing work don't fit comfortably into any existing WG
  - But they may be too small to justify a new working group
- Other pieces of work are highly technical but don't require the development of a new routing protocol
  - They describe how routers can behave to improve routing success
- The Routing (Area) Working Group is the catch-all for these
  - Do not confuse this on your agenda with the Routing Area Open Meeting
- RTGWG also acts as a venue for "mini-BoFs"
  - Proponents can float new ideas in a skilled and critical environment
    - Just a 20 or 30 minute slot

A Word About YANG

- ***Everyone*** seems to be talking about YANG models
- There are around 120 active I-Ds with the term "YANG" in their titles or filenames [1]
  - Although some of these may belong to Chinese authors ☺
- YANG and NETCONF have replaced ASN.1 and SNMP as the configuration mechanisms of choice in the IETF
  - A more parsable modeling langague
  - A more flexible protocol
- Riding on the back of a lot of OpenSource SDN work
- I2RS focuses specifically on YANG models
- Every other working group has at last one YANG model
- RTGWG acts as a home for stray routing YANG models

[1] http://datatracker.ietf.org

BoFs

- There is one Birds of a Feather meetings related to routing at this IETF

- Data Center Routing (dcrouting)

  Data Centers, because of their topologies (traditional and emerging), traffic patterns, need for fast restoration and low human intervention, among other things, are driving a set of routing solutions specific to them – in this case, one size probably doesn't fit all.
  The focus of this effort is on new potential solutions: ones that may require a standalone effort.

Closed Working Groups

- When a working group is closed it means it has finished its work
  - It ***does not*** the protocol it developed is dead or pointless
    - Although sometimes it does!
  - A working group should aim to close: this is good!
- Notable examples include…
  - Routing Information Protocol (RIP and RIPv2)
  - Virtual Router Redundancy Protocol (VRRP)
  - Forwarding and Control Element Separation (ForCES)
- Look at the very long list at…
  http://datatracker.ietf.org/group/concluded/

Routing Directorate

- Panel of routing area experts
  - appointed by the ADs
  - 46 current members
- Purpose of the directorate
  - Review routing area drafts as they pass through IETF last call
    - Guardians of Quality – last chance to spot issues
  - Review other routing-related drafts at IETF last call
    - Do not aim to review all IETF drafts, unlike some other directorates e.g. security
  - “Early” review of any routing-area WG document before WG last call
    - Aim to improve and debug documents early in their lifetime

How the Routing Directorate works

- The RTGDIR coordinators (Jon Hardwick and Amy Ye) manage the directorate's operation
  - Contact details on the wiki: https://trac.ietf.org/trac/rtg/wiki/RtgDir
- An AD or WG chair can request a routing directorate review of a draft at any time using a button on the draft's datatracker page
  - ADs – usually at IETF last call
  - WG chairs – when looking for an expert pair of eyes on a WG draft
- Reviews assigned by round robin by default
  - If specific expertise is needed, this can be flagged in the review request
- Review turnaround is usually ~2 weeks

Many Ways to contribute

- Pick your favorite working group, write a draft and discuss it on the list or live at a meeting
- But you can also
  - Review drafts and share your comments on the list
  - Sit at the front and take minutes, sit near the microphone and relay Jabber
  - Volunteer to become a WG secretary
  - Volunteer to shepherd documents through the standardization process

Work in Other Areas

- OPS Area
  - Global Routing Operations (GROW)
  - Layer 2 VPN Service Model (L2SM)
  - Layer Independent OAM Management in Multi-Layer Environment (LIME)
  - MBONE Deployment (MBONED)
  - SIDR Operations (SIDROPS)
- INT Area
  - Home Networking (HOMENET)
  - IPv6 over the TSCH mode of IEEE 802.15.4e (6TISCH)
  - IPv6 over Networks of Resource-constrained Nodes (6LO)
  - Host Identity Protocol (HIP)
- TSV Area
  - IP Performance Measurement (ippm)

IRTF

- The Internet Research Task Force has always done work of importance to RTG
  - For years the Routing Research Group (RRG) was a key place for discussion of the next steps in routing
- Current RGs of interest are…
  - Global Access to the Internet for All (GAIA)
  - Network Function Virtualization (NFVRG)
  - Network Coding (NWCRG)
  - Path Aware Networking Proposed RG (PANRG)
  - Thing-to-Thing (T2TRG)

Independent Stream

- A number of routing protocols are published as RFCs on the independent Stream
- These are not the work of the IETF
  - The only IETF review they receive is to check that they do not directly conflict with IETF work
- There is a variety of such work…
  - Proprietary protocols published so that people can implement and interoperate
  - Academic or other experiments
  - Failed ideas published for the record
  - Work that the IETF was not interested to pursue
- Sometimes Independent Stream work gains traction and is brought back into the IETF for more work

Resources

- Datatracker for information about all working groups and documents
  - http://datatracker.ietf.org
- BoF wiki for details of all BoF meetings
  - http://trac.tools.ietf.org/bof/trac/
- The Routing Area wiki
  - http://trac.tools.ietf.org/area/rtg/trac
- The Routing Area Directorate's wiki pages
  - http://trac.tools.ietf.org/area/rtg/trac/wiki/RtgDir

Resources

- General Routing discussion list
  - https://www.ietf.org/mailman/listinfo/routing-discussion
- Routing and Open Source discussion list
  - https://www.ietf.org/mailman/listinfo/rtg-open-source

Roll of Thanks

- original authors: Adrian Farrell, Jeff Haas
- update authors and presenters:

  Joel Halpern, Susan Hares, Martin Vigoureux
- Routing Area Directors: Alia Atlas, Deborah Brungard, Alvaro Retana
- WG chairs in Routing Area
  - https://datatracker.ietf.org/wg
- Routing Area Directorate
  - https://trac.ietf.org/trac/rtg/wiki/RtgDir