

Introduction to the IETF's Routing Area

IETF-93 : July 2015 : Prague

Adrian Farrel (adrian@olddog.co.uk)

Jeff Haas (jhaas@pfrc.org)

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 134 working groups
 - 42, 1, 20, 15, 23, 18, 15 (17%)
- IETF published 260 RFCs in year to June 30, 2015 [1]
 - {6+29+23}, 33, 32, 71, 21, 20 (27%) [2]

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

[2] The ART Area is formed from the previous APP and RAI Areas

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
 - The whole route is pre-determined
- 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
- These are largely in the APP Area
 - Content Delivery Networks Interconnection (CDNI)
 - INtermediary-safe SIP session ID (INSIPID)
 - Peer-to-Peer Session Initiation Protocol (P2PSIP)
 - Session Initiation Protocol Core (SIPCORE)
- Also the SEC and TSV Areas
 - Application Bridging for Federated Access Beyond web (ABFAB – SEC)
 - 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 23 working groups, a few of which ask for more than one meeting session
 - There are usually only 17 meeting slots
- 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 all 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)

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
 - At least two independent and interoperable implementations are needed before any protocol extension is advanced for publication as an RFC
- Essentially in “maintenance mode”, the working group works on protocol extensions intended to make the global routing system work more smoothly and scale better
 - Many suggestions come via the GROW WG
- Other protocol work comes from BESS
- A further important piece of work is BGP-LS
 - This allows the “export” of routing information (especially TE information) from a network to a management system (such a Path Computation Element) or to another network

PIM

Protocols for IP Multicast

- There used to be several competing protocols for multicast
 - Protocol Independent Multicast - Sparse Mode (PIM-SM)
“won”
 - Not widely deployed, but there are significant deployments
- Took over responsibility for IGMP and MLD
 - Used to be in INT Area
 - Puts all multicast expertise in one place
- Also a “maintenance mode” working group
 - Finalizing 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
- Each packet carries the full path that it should traverse
 - Compare with IP source route option
- Currently being worked on for MPLS and IPv6
- Ongoing work is architectural and protocol extensions where needed
- Routing protocol extensions (OSPF, IS-IS, BGP) happening in the respective working groups

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

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
- Work in the WG is notoriously slow and hotly debated
- Outstanding work items include...
 - DLEP : A protocol to report link characteristics to routers
 - AODVv2 : A distance vector protocol for MANETs
 - A number of extensions to OLSRv2 : A link state protocol
 - Enhanced security 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
 - 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

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
- Although moving to maintenance mode, the WG still generates at least 2 RFCs every month
- 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
- Potential future work includes completing flexi-grid, consideration of super-channels, routing with optical impairments
(One current and one recent AD chaired CCAMP)

L2TPEXT

Layer 2 Tunneling Protocol Extensions

- An almost dormant working group that exists to maintain and extend the Layer 2 Tunneling Protocol (L2TP) as necessary
- Likely to be folded into PALS in the near future

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
- Mainly a maintenance-mode WG
 - Expect a constant trickle of protocol refinements and pontificating Informational I-Ds

TRILL

Transparent Interconnection of Lots of Links

- Recently moved to RTG from the INT Area
- Originally conceived as just working on a new encapsulation
- Quickly became heavily involved with the application of IS-IS to this new encapsulation
 - About half of the work needed explicit review by IS-IS experts
- Currently working on...
 - Multicast
 - An implementation report
 - Security analysis and extensions
 - Applicability to data centers

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

Jeff co-chairs BFD

I2RS

Interface To the Routing System

- Software Defined Networking (SDN) has focused on the interface from the routing or control system to the physical forwarding components
- I2RS is at a higher level looking at the interface to the routing system
- Examples include
 - Installing routes into the Routing Information Base
 - Programming route admission policies into the BGP engine
- The WG has chosen YANG as its modeling language

Jeff co-chairs I2RS

PCE

Path Computation Element

- Originally conceived as an off-board tool for computing paths in multi-domain Traffic Engineered MPLS networks
- Now finding its place as an active network management tool
- The working group 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
- Future use cases and protocol work may arise from...
 - 6TiSCH working group in the INT Area
 - The DetNet BoF

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
- Close coordination with IDR for BGP extensions
- Coordination with...
 - MPLS for architectural considerations
 - NVO3 for data center VPNs
 - TRILL for EVPN interoperability

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
- Any form of Pseudowire service
 - IP, MPLS, L2TP
 - Pseudowire encapsulations
- Likely that L2TPEXT will be folded in soon

NVO3

Network Virtualization Overlays

- A relatively new working group
- Develop a set of protocols and/or protocol extensions that enable network virtualization within a data center
 - assumes an IP-based underlay
- Progressing slowly despite aggressive milestones
- Some work now off-loaded to BESS and PALS
- A lot of time focussing on new or proprietary encapsulations

SFC

Service Function Chaining

- Arguably not a classic routing problem
- Work concerns directing flows through service function nodes to apply features such as policing, access control, security, and load balancing
- Challenges are...
 - How to associate a packet with a flow
 - How to attach “metadata” to a packet or a flow
 - How to program the next service function (node) for a flow
- Work is mainly architecture and use cases
 - Protocol work is an encapsulation header intended to be layer-agnostic

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 work remains experimental while a fuller understanding of the impact is collected
- 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

- 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

Two of the current ADs used to chair RTGWG

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”
 - Deliberate “route hijacking”
- This working group 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

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 language
 - 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 least one YANG model
- RTGWG acts as a home for stray routing YANG models

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

BoFs

- There are two Birds of a Feather meetings related to routing at this IETF
 - Both aiming to form working groups
- Deterministic Networking (DetNet)
 - Looking to provide very predictable packet delivery
 - Well-know (and low) delay
 - Very small jitter
 - Particularly useful in high-density media environments
 - For example: video studios
- Simplified Use of Policy Abstractions (SUPA)
 - In the OPS Area
 - Looking to provide generic abstractions of “policy” for use in managing and operating a number of environments including routing

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/>

Work in Other Areas

- There is work directly related to RTG done in other Areas
- OPS Area
 - GROW
 - Global Routing Operations
 - L3SM
 - Layer 3 VPN Service Model
 - LIME
 - Layer Independent OAM Management in Multi-Layer Environment
 - MBONED
 - MBONE Deployment
- INT Area
 - HOMENET
 - Home Networking
 - 6TiSCH
 - IPv6 over the TSCH mode of IEEE 802.15.4e
 - HIP
 - Host Identity Protocol

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...
 - Delay-Tolerant Networking (DTNRG)
 - Global Access to the Internet for All (GAIA)
 - Network Function Virtualization (NFVRG)
 - Network Coding (NWCRG)
 - Software-Defined Networking (SDNRG)

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
 - Strange or wonderful 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>