

December 2005

*Geoff Huston*

ISOC has started reporting on IETF meetings in a new publication, the [IETF Journal](#). I had done a report on [routing-related activity at IETF-63](#), so I volunteered to write a similar report for IETF-64, held in November at Vancouver.

So here's the report I produced on activity within the Routing Area at IETF-64

## Routing IETF-64

The following is a rundown of current status of the working groups that either met at IETF-64 or whose status was reported at the Routing Area meeting during the week of IETF-64 in November 2005. This is of course a set of personal opinions and perspectives rather than any official report of the IETF.

### Routing Area Working Group (rtgwg)

Alex Zinen, one of the Area Directors of the Routing Area announced his intention to step down from this role in March 2006, at the expiration of his current term as AD. Alex has served for four years as an AD for the Routing Area of the IETF and has established a careful consultative style as an Area Director. I'd like to simply say here a personal thanks to Alex for his time and energy over the past four years.

RFC1264bis – a review of Routing Protocol Standardization Criteria. A number of changes are being proposed here, including turning the protocol analysis documentation, which was a mandatory requirement for Proposed or Draft Standard protocol specifications, into a chartered step if it is felt that such an analysis is a requirement for the protocol being developed. The experience with BGP was that this particular analysis was an exercise required for the IETF standards process, but was not felt to be a useful document in its own right. The current proposal is to either place the preparation of this document into the charter as an explicit Working Group deliverable, or do not prepare such an analysis. Also within scope of this review is a clarification of the independence of the 2 implementations from the proposed specification.

The area meeting also considered the manageability considerations proposal. This is a proposal for each routing protocol to have explicit consideration of manageability while designing the protocol. The discussion of this proposal highlighted the consideration that making this a required considerations section in a protocol specification may not necessarily be a lever to get folk to think about this topic, and it stands the risk of adding more boilerplate text of specification documents. On the other hand, thinking about manageability early in the process of protocol specification may be a useful exercise. However. There was no overwhelming push to make this a mandatory part of routing protocol specifications.

IP Fast Reroute - the microloop prevention specification has been updated, as have the base protocol and framework documentation.

---

## Common Control and Measurement Plane (ccamp)

There has been a collection of RFCs published recently published (RFCs 4201 through to RFC4210) on a common theme of Multi-Protocol Label Switching (MPLS) extensions and refinements, including six from the CCAMP Working Group on the topic of control and management extensions Generalized MPLS. A further eight documents are in the RFC Editor queue, nine documents have completed working group last call and seven are still being considered by the working group. Given this relatively high level of document generation, the pace of work in this working group has been quite intense in recent months. A revised charter for CCAMP reflects an intention to deliberately pace the next round of activity to match the capacity of the working group to carefully review material, but nothing dramatically different in terms of direction here. The meeting at IETF-64 had a relatively full agenda, including the following items. The group is working on an update of RFC3946 in an attempt to clear up a potential ambiguity, and in the way of many similar efforts, what was in the first instance a relatively straight forward minor task of altering a condition that was “greater than 1” to “greater than or equal to 1” has become infused with all kinds of complexities relating to already deployed implementations that have interpreted the existing text literally, while others have used a more liberal interpretation. It was reported that a resolution appears to be in sight, and it is expected that this will clarify some of the issues in interworking between the SONET and SDH systems. Other activity includes consideration of addressing in GMPLS networks, Traffic Engineering LS paths and the interaction with the RVSP protocol. Related work is on a Network-to-Network Interface specification (NNI) for GMPLS and an associated area of study of inter-domain GMPLS. One of the proposed work items I found interesting was that of virtual concatenation coupled with Link Capacity Adjustment within a GMPLS framework, which is proposed for a general inverse multiplexing technique that could be used across a number of transport technologies, including SONET, SDH, PDH and OTN. For those of us who have struggled with various forms of inverse multiplexing over the years in an effort to treat a number of parallel circuits as a single virtual circuit with a capacity equal to the sum of the multiplexed components, this news of a generalized approach is indeed promising news.

## Forwarding and Control Element Separation (forces)

It appears that this working group is relatively close to completion of its work. To recap from the charter of this working group, the emergence of off-the-shelf network processor devices that implement the fast path or forwarding plane in network devices such as routers, along with the appearance of a new generation of third party signalling, routing, and other router control plane software, has created the need for standard mechanisms to allow these components to be combined into functional systems. In other words ForCES is an effort to standardize a number of internal control interactions between the logical components of a routing engine. To continue from the charter, ForCES aims to define a framework and associated mechanisms for standardizing the exchange of information between the logically separate functionality of the control plane, including entities such as routing protocols, admission control, and signalling, and the forwarding plane, where per-packet activities such as packet forwarding, queuing, and header editing occur. At IETF-64 there was an interesting presentation of a ForCES router implementation, with a control element and a forwarding element linked by ForCES protocol messaging. It seems that we are nearing the completion of this effort.

## Inter-Domain Routing (idr)

From a somewhat personal perspective, the good news from this meeting was the completion of this Working Group's efforts with preparing the 4-byte AS proposal. This has been parked in the working group for some time awaiting two independent implementations of the specification before being able to proceed as a Proposed Standard. With the recent implementation report on these implementations this draft is now on its way to being a Proposed Standard. A similar issue was associated with the AS Confederations specification, and with the recent implementation report this specification has also completed working group review, and is ready for publication as a Draft Standard. The Working Group has also been working on a revised base specification for the BGP protocol, and this group of documents is now with the RFC Editor.

The new work being introduced into this working include the use of an explicit AS Time To Live (TTL) for BGP advertisements. Currently it is possible to specify a TTL of 1, by specifying “NO EXPORT”, but not any values higher than 1, so advertisements are either highly constrained to immediate BGP peers, or completely global. Like similar previous efforts with the “NO PEER” community attribute, the TTL specification is an attempt to localize the propagation of a routing advertisement to a particular AS radius.

---

The IDR Working Group continues to see a wide variety of proposals for refinements to BGP, including outbound route filter grouping, aggregated withdrawals, dynamic AS renumbering, multicast signalling and explicit support for route tunnelling to support various forms of overlay configurations. The major criteria here for advancement of a proposal in the standards process is a writeup of two independent implementations of the proposed specification.

Of course there is also no shortage of proposals that appear to be on a continuous loop, and QoS routing, or in this context inter-domain QoS routing, is perhaps one of the best known of these proposals. Its not all that easy to identify precisely what has changed at each iteration of such proposals, and at each time the proposals tend to founder on one of the basic precepts of the Internet's inter-domain routing architecture, namely that routing is not a resource management system. The entire topic of how to manage a network's resources, and to how solve the associated feedback signalling mechanisms remain very resilient as outstanding problems in the routing space.

### IS-IS for IP Internets (isis)

As reported to the routing area, the IS-IS working group has now pushed most of its drafts through the process, including link attributes and router capability advertisements. Rechartering of this working group is the logical next step, and the decision at this point in time is whether to identify a number of work items relating to further IS-IS extensions (such as Layer 2 end point definition) and refinements (such as logical tunnel concentration) and recharter the group to work on these items, or to leave the working group dormant for a period while the current drafts complete their path through the publication process and await a critical mass of new work proposals for IS-IS before reactivating the working group with a new charter.

### Layer 1 Virtual Private Networks (l1vpn)

It is still early days for this particular working group, and this is their second meeting. Someone well versed in the 7-layer network model would see layer 1 as a media adaptation layer, primarily concerned with electrical voltages, plug and socket dimensions and encoding formats. This is not quite the case here. This form of VPN is based on a switched circuit-based network, that may be composed of optical cross-connects, time division cross-connects or fibre switches. The VPN control plane is used to provision a set of switching configurations to inter-connect Customer Edge (CE) devices in a specified topology. The working group is currently working on two documents, framework and applicability, and will shortly start looking at the solution aspects of this form of switch control. With a considerable level of interest in the research community in various form of light-path switched systems for very high speed point-to-point on demand circuits, this form of automation of control of the switching elements appears one promising way to handle on-demand high speed circuit provisioning.

### Mobile Ad-hoc Networks (manet)

To quote from the charter of the ad-hoc autoconfiguration working group, from the IP layer perspective, a MANET presents itself as a IP multi-hop network formed over a collection of links. Thus, each ad-hoc node in the MANET is, potentially, acting as a router in order to provide connectivity to other nodes within the MANET. Each ad-hoc node maintains host routes to other ad-hoc nodes within the MANET, in addition to potentially holding network routes to destinations outside the MANET. If connected to the Internet, MANETs are edge networks, i.e. their boundary is defined by their edge routers. Due to the nature of the links over which a MANET is formed, ad hoc nodes within a MANET do not share access to a single multicast-capable link for signalling. This implies that the usual delivery semantics of link-local multicast and broadcast are not preserved within a MANET.

The specification for this topic is now relatively well fleshed out and the working group is now calling for early implementation reports of the MANET protocols. A small number of drafts remain active in the working group, concerning dynamic source routing, on-demand routing, a link-state routing protocol and a simplified multicast forwarding protocol. It is likely that these documents will be completed in early 2006. Also the group is spinning off

---

activities in other areas, such as the autoconf working group in the Internet Area, and interest in a MANET research group to look at topics such as multicast, link metrics and the potential of QoS-related activity.

### Multiprotocol Label Switching (mpls)

As with IDR, OSPF and IS-IS, the MPLS working group is now one of the more venerable working groups in the routing area. Most of its chartered goals and milestones have been achieved, and the current work is focussed on a number of matters relating to ICMP handling, management considerations and OAM requirements and framework, failure detection and graceful restart mechanisms, point-to-multipoint paths. The decision point appears to be rapidly approaching whether to recharter MPLS, or to wind up with working group and charter more specific working groups on the basis of demonstrated interest in specific areas of further MPLS refinement.

### Open Shortest Path First IGP (ospf)

There has been some good progress on some long-standing work items in this Working Group, with work on refresh and flooding in stable networks, graceful restart and prioritization and congestion avoidance all being published as RFCs. The working group is currently completing work on IANA Considerations to create a number of IANA registries for OSPF types and options, as well as traffic engineering extensions for OSPF v3. Current activity includes consideration of multi-topology routing, where a number of basic approaches including reuse of the IPv4 Type of Service (TOS) bits with altered semantics in the context of OSPF v2, or use of OSPF v3 with separate instances of OSPF for each topology instance, or the use of tagging OSPF protocol elements with Type Length Value (TLV) headers to allow a number of routing contexts to co-exist in one OSPF environment. The OSPF working group is also looking at the integration of the Mobile Ad-hoc network (MANET) requirements into OSPF v3, looking, in particular, at how to manage potential flooding instances and reduction in the level of formed adjancies. Rechartering of the OSPF Working Group also appears to be a near term option.

### Path Computation Element (pce)

The PCE Working Group is chartered to specify a Path Computation Element (PCE) based architecture for the computation of paths for MPLS and GMPLS Traffic Engineering LSPs. In this architecture path computation does not occur on the head-end label switching device, but on some other entity that may physically not be located on the head-end device. As reported to the Routing Area, this working group is evidently making good progress, with the architecture description at a mature state and the requirements document also close to completion. The group is intending to complete these documents before heading into the protocol specification phase of their work. At this stage the working group is looking at candidate path computation communications protocols, and protocols of the discovery of path computation elements.

### Routing Protocol Security Requirements (rpsec)

This group did not meeting at IETF-64. As reported to the routing area meeting, the main work item at present is the security requirements document for BGP. This document is supported by reasonable agreement on most aspects, but there remain a small number of strongly contested items remain, and there is no clear way forward at this stage to resolve this. There has evidently been some discussion in the working group on starting a work item on Interior Routing security requirements at this stage, and defer the resolution of the remaining BGP items for the moment.

### Source-Specific Multicast (ssm)

The SSM architecture document has been approved by the IESG. As this was the last remaining work item for the Working Group, it may be that the Working Group has now completed all its work!

---

## BOF Sessions

One way to charter new work in the IETF is via the BOF, which is a more informal session designed to assess the level of interest in the work, and see what related issues may be exposed when considering a particular topic. Two BOFs were held in the IETF-64 within the Routing Area.

### Secure Inter-Domain Routing (sidr)

The BOF reviewed the current status of RPSEC, and the current state of design activity in the area of secure inter-domain frameworks. The proposition was advanced that while RPSEC has not concluded as yet, there is sufficient impetus to commence work on infrastructure and protocol support mechanisms intended to address aspects of securing inter-domain routing. The specific area where there has been clear agreement in the requirements specification activity is that of authentication of route origination.

The proposed work would include consideration of the relevant certificate infrastructure to support information validation. It was noted that the outcomes of this activity should be capable of supporting hierarchical rooted PKI models as well as decentralized "web of trust" models if at all possible, as the intended scope of application of this framework encompasses a broad diversity of deployment environments.

There was support from the BOF attendees for the aspects of the work where there is clear agreement on requirements, concerning authentication of route origination information and use of associated certificate frameworks, to be undertaken immediately. The question of charter scope was considered and the rough consensus in the BOF was to support a charter that encompassed a more comprehensive security framework for inter-domain routing, but with a caveat that commencement on any particular component of the work would be conditional on clear agreement on requirements from the RPSEC Working Group.

### GMPLS-controlled Ethernet Label Switching (gels)

When all you have is a hammer, then everything looks like a nail, or so goes the saying. So when all you have is Generalized Multi-Protocol Label Switching (GMPLS), then everything looks like a collection of potential label switching devices! (Although some folk have been heard to comment that when all you have is GMPLS then, unfortunately, everything still looks like a nail!) This session was to see if there was interest in applying GMPLS to Ethernet switches in support of point-to-point label switched paths. This is very close to the existing effort in CCAMP, but with the addition of wanting to place label information into the Ethernet frame and then coordinate the switches via a GMPLS superstructure. The proposed work was to include definition of protocol-independent attributes for describing links and paths that are required for routing and signalling Ethernet switched point-to-point paths, and specification of routing protocol extensions (OSPF, ISIS) and signalling protocol extensions (RSVP-TE) required for Ethernet switched point-to-point path establishment. If you are looking for a clean delineation between layers 2 and 3 of the OSI protocol stack model in this work you are probably not going to see it! This a blurring of the original protocol model that attempts to create logical point-to-point circuits between Ethernet switching devices, where the circuits are constructed using a label path across label switching devices using some form of routing mechanism to determine edge-to-edge paths. Not all BOFs become chartered as working groups in the IETF, and there was evidently little support in this case to continue with this work in the IETF.

---

### Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

### About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and is currently the Senior Internet Researcher at APNIC, the Regional Internet Registry serving the Asia Pacific region. He was a member of the Internet Architecture Board from 1999 until 2005.