

BIER WG Interim

Day 2

Sept 28, 2016

Minute taker, nranns@cisco.com

Minutes

alia: does using a specific address allow to express different forwarding

shep: mcast address are already an abstraction

tony: problems with changing the dst addr in flight

shep: changes are only within the BIER domain, not the overlay.

dino: a new higher order byte in the address implies diff forwarding, so ok for diff behaviour.

greg: not changing address -it's address switching. e.g. NAT, MPLS.

dino: hop-by-hop options.

tony: v6 arch, hbh is last resort.

tony: pcp38

pierre: pcp38 will filter bad source address.

- out of the domain, host cannot use BIER prefix.

- in domain, host will have one address with BIER bit. this host can only use this address to send

tony: implies /128 for each host

pierre: /64 is not on the link.

tony: when sending, use BIER addr or other

pierre: BIER address is routeable.

- /128s are at the edge???

- BIER in the domain.

pierre: if you want to reach BIER domain from outside, you advertise.

tony: AH header,

dino: sounds complex. separate ethertype. new I3 proto.

dino: is a new forwarding paradigm;

shep: i disagree.

ice: new paradigm. but not complicated.

alia: just cos it's exp, doesn't mean we can affect v6 arch.

pierre: how do we prove it can be used, if we can't do ucmnet/try it?

alia: we need demonstrate a use case to change v6 arch

ice: server-to-server. no IGMP, DR/DF election.

- prototype, demo it works.

tony: UDP checksum. error in flight. or recompute on each hop.

dino: 3rd forward

dino: new etherte

ice: if it's a new ehtertype, new header, what about applications

dino, tony: raw sockets, provide library.

toerless: raw sockets require priviledge

shep: app writers, put them in the way, or make them oblivious.

dino: make it compeling, they will come.

ice, toerless: IGMPv3 took decades for adoption.

alia: we have layers, but they can be arbitrarily stack. do we want to solve only apps for L3, or any layer.

dino: one receiver => unicast. BIER mcast nad ucast same. It's a new forwrding paradigm.

alia: if theres an app sending, e.g. NSH over transport, would you want BIER hdr to be at same layer as NSH, or only at transport.

dino: at transport. it's an underlay. BUT it can can over VXLAN, it just won't be as effective. it's a peer proto to MPLS and IP.

pierre: we could have new/one encap - which we have. v6 brings as weel, capability to send through internet + hosts that are not BIER aware to send BIER packets.

dino: other protos can transpoirt mcast over internet, but not without added encap.

tony: securitiy problems for domain being hit by BIER preifixes.

tony: storage/paxos app, use case for BIER

dino: receiver discovery in app;

shep: yes. rx/tx discover out of scope.

tony: is it easy for app to write v6 headers or new encap headers???

all: MPLS encap is best.

shep: are there other use cases that require another encap.

ice, pierre: app opens socket on BIER /128 address.

dino: ok.

tony: IPSEC b0rked.

tony: 64 bit is limited. Arch says must support 256.

ice; divide by 64 still a great saving.

ice: change MPLS encoding draft, to say MPLS enc must do 256. BIER arch does not specify.

shep: MPLS encap good; done. do we have other use cases? is there an application that is not addressed by MPLS.

greg: if there's no MPLS DP.

shep: operational. often non-technical.

dino: network must upgrade all boxes. not the same as depending as on existing MPLS boxes. not teaching ops MPLS.

shep: we are not running MPLS.

toerless: homenet as use case?

tony: signalling in BABEL

pierre: HNCP. need to prop ampping between bit and ip addr.

pierre: src addr from DHCP. dst addr from service discovery (of sort sort).

- allows host to be non-BIER supporting.

tony, neale: tunnel through non-BIER nodes in the same way.

LUNCH.

alia: comments from individuals...

linda: questions re. ethernet BIER address

sandy: still questions on v6.

pierre: mpls/ethernet good. need to do BIER without knowing MPLS. there are other link layers. we should look at v6 as an encap (but not necessary Pierre's proposal). Hosts

neale: mpls&ethernet combined great. sufficient use cases for continuing investigation of v6.

ice: mpls&ethernet combined great. important to keep BIER forwarding in unicast FIB. v6 purists might disagree, but there are enough advantages to prototype.

shep: v6, change inevitable. resistance expected. who do we need to engage with about should and should not do.

dino: one encap only. message if there an arch change, and ietf accepted, not necessarily uptake from operators. Need operator in the room, and an app coder, understand from them, reqs.

greg: good converging to single BIER header o different transport - pursue and complete. if req. expose qos ttl into BIER layer, then encap must be capable. Would result in 'generic' encap.

- use IP next proto for getting generic

ice, toerless; lose control.

alia: see draft for defining next proto field.

toerless: best chance for common encap changin existing and not calling it MPLS.

neale: when the encap is not explicit that it's a label, do you set the EOS bit in when tunneling over MPLS.

tony: 8847 simple reason gets us there fast. selling may be harder. new ethertype later. v6 enticing. current proposal up high sell - low rate success. BIER o UDP - something usefull for the host.

(impromptu white board session, tony brainstorming)

shep: volunteers for host-to-host BIER arch;

- jeff, toerless, ice, tony.

greg: encap transport and payload agnostic

jeff: keep 8847. transport and payload agnostic. IGMP MLD on host, better than host-to-host BIER.

alia: common encap good direction. needs to be single encap. fine to explore host-to-host. include homenet group.

shep: host-to-host not homenet, it's DC use-case.

pierre: homenet needs multicast. not necessarily BIER.

alia: ack. BIER better than PIM.