Link properties advertisement from modem to router

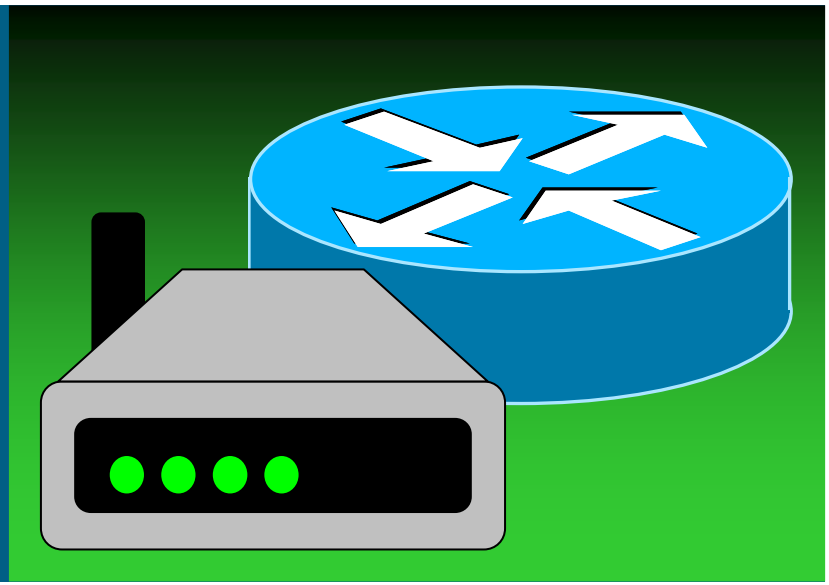

draft-wood-dna-link-properties-advertisement-01

Lloyd Wood, Rajiv Asati and Daniel Floreani
Cisco Systems

Detecting Network Attachments session
IETF 72, Dublin, July 2008.

Contents

- Outline of problem
- Alternatives to solving the problem

  leading to why we selected this approach
- How the selected approach works.
- Details of ‘rate block’ format.
- How block format is extensible.
- Questions raised by this approach.

A smart adaptive modem and a router

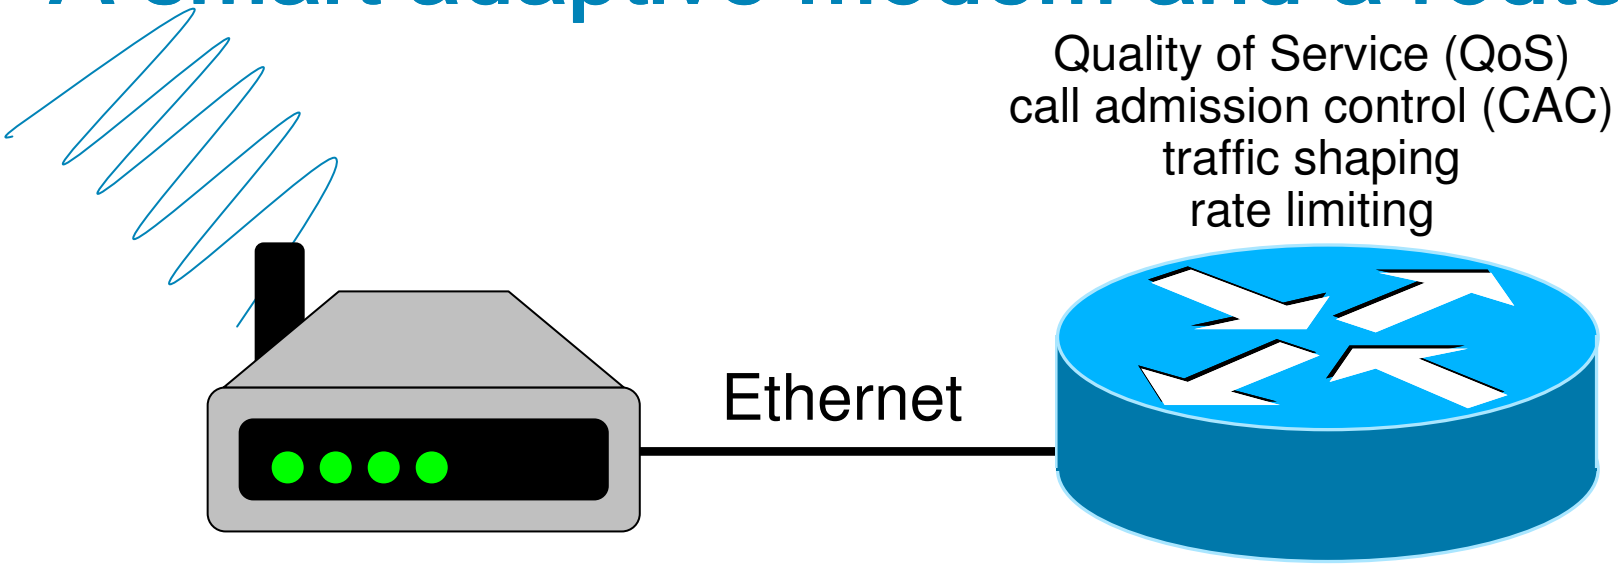

Connected together by Ethernet, because Ethernet is cheap.

Modem adapts its link to varying conditions – could be DVB ACM or VCM, ADSL, or any other smart adaptive link.

Router doesn't see the modem link speeds. Needs to set QoS and shaping the next hop along, and rate-limit over Ethernet to match the *current* modem speeds. How to do this? (Sending 100Mbps to the modem for it to drop is bad.)

Assumes router has the fuller traffic-munging featureset.

Alternatives for modem to describe link speed

- just clock a serial interface – but the world's gone Ethernet.
- RFC4938/draft-bberry-rfc4938bis-00 describes extensions to PPPoE to do this – fine ***if*** PPPoE is considered suitable.
- Some sort of tunnel, explicitly extending link to the router (PPPoE does this. Not keen.)
- an ICMP packet, using (obsolete) spare ICMP bits?
- Ethernet pause frames for flow control? Is an Ethernet-specific approach adding more control loops a good thing?

We wanted simple-as-possible method. So…

- Simple UDP packet, sent to link-local 'all routers' multicast address, advertising current modem interface rates.

How it works

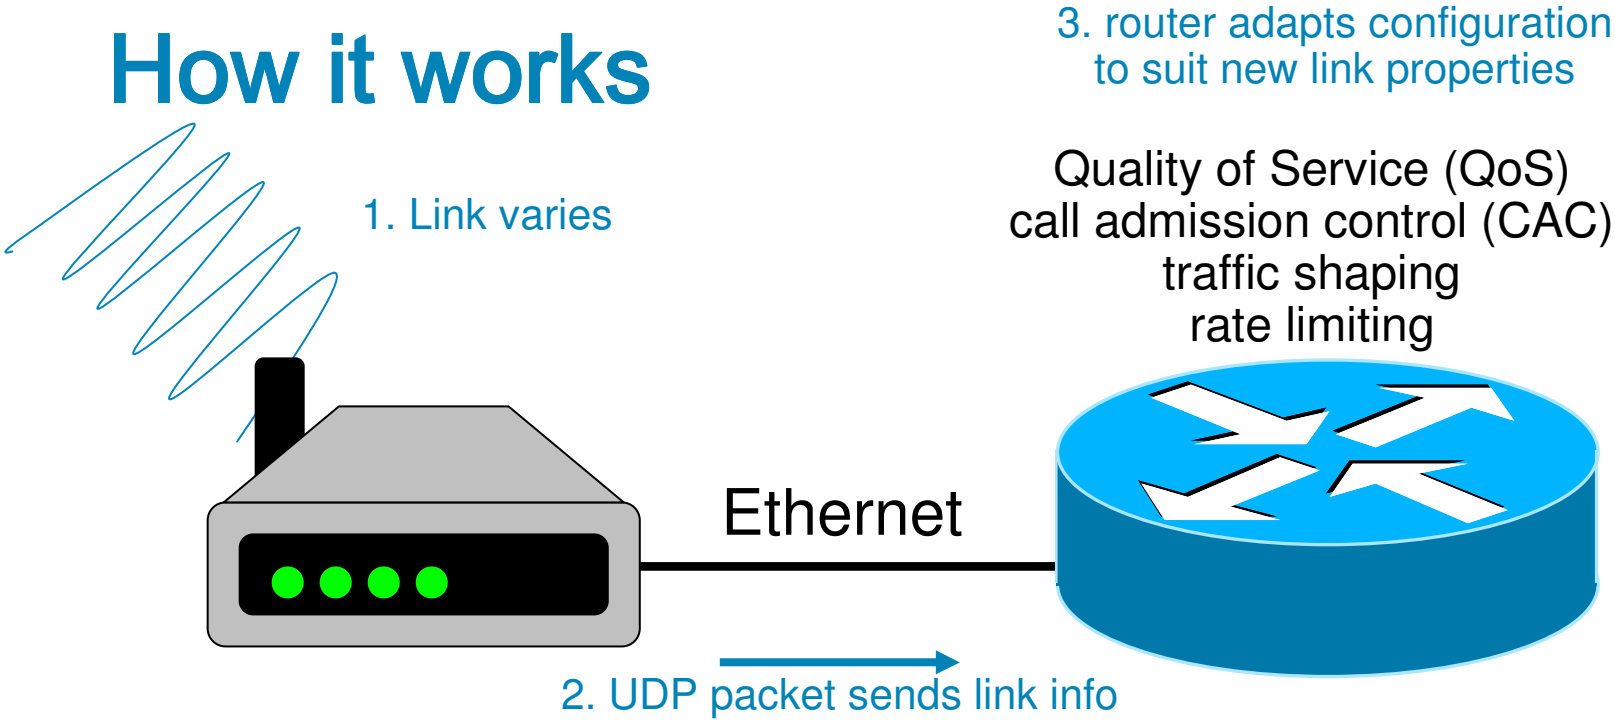

Modem sends UDP packet with information describing varying links to router when a link changes, and periodically.

Information sent on ‘all routers’ link-local multicast address on modem-router link, to avoid explicit configuration.

Router uses provided information to set traffic shaping of input and output, rate limiting of output to match modem’s link speed.

Send UDP packet with ‘rate block’

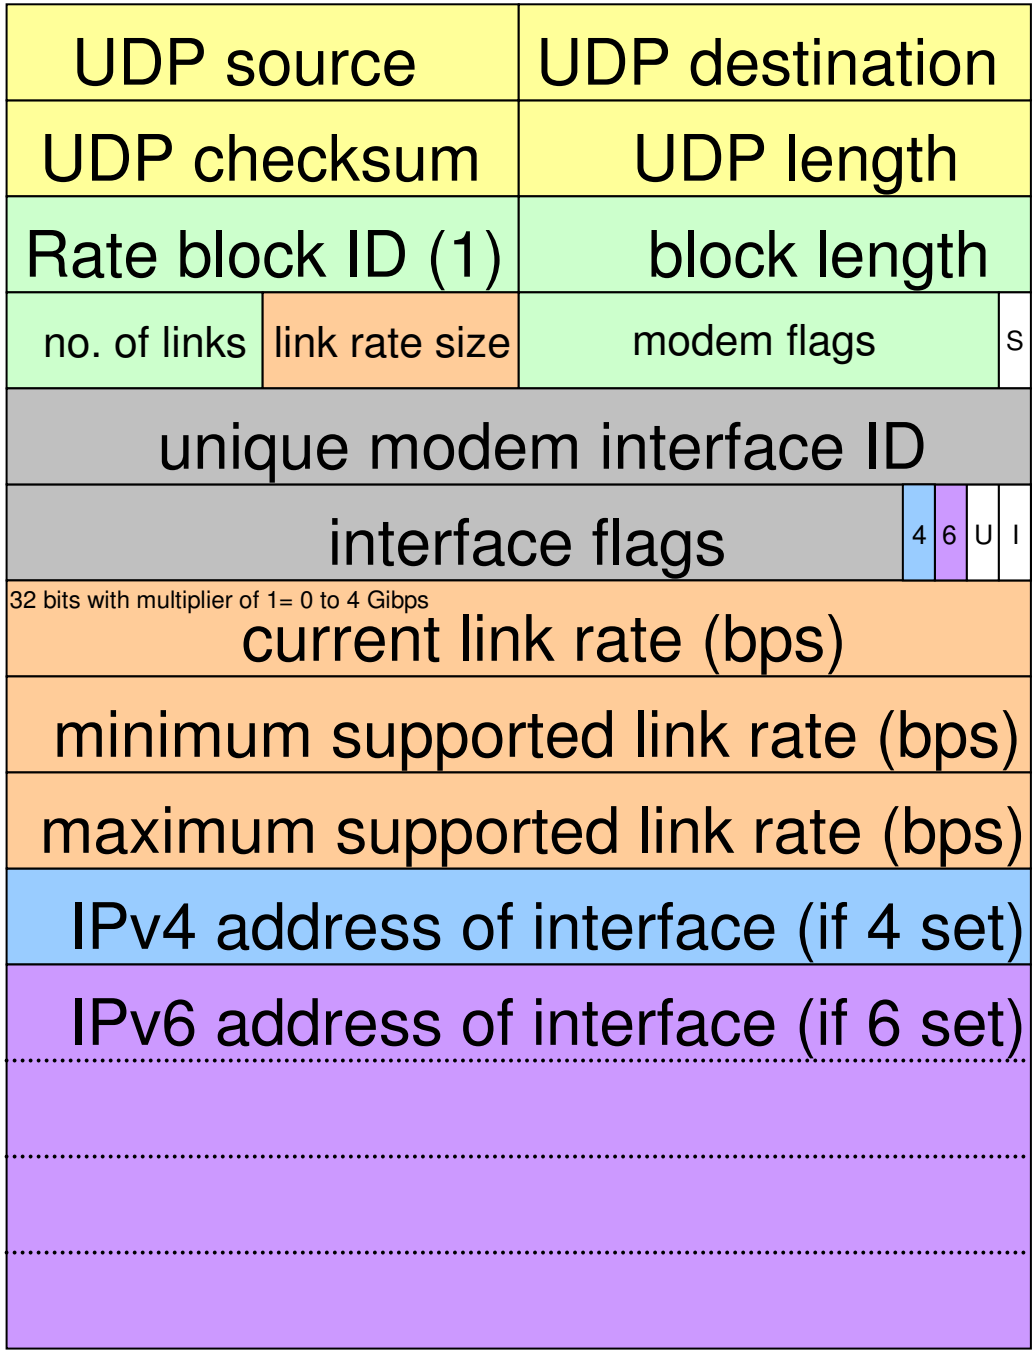

S – this block describes Some or All interfaces (handles exceeding MTU).

U – interface is Up or Down
I – incoming or outgoing; needed for asymmetrical links (different rates, or down in one direction) and unidirectional links.

repeated for each link interface: incoming to modem+router and outgoing from modem+router share ID

basic TLV structure (type/length/value). Sent to link-local ‘all routers’ multicast address

Add other information in other blocks later

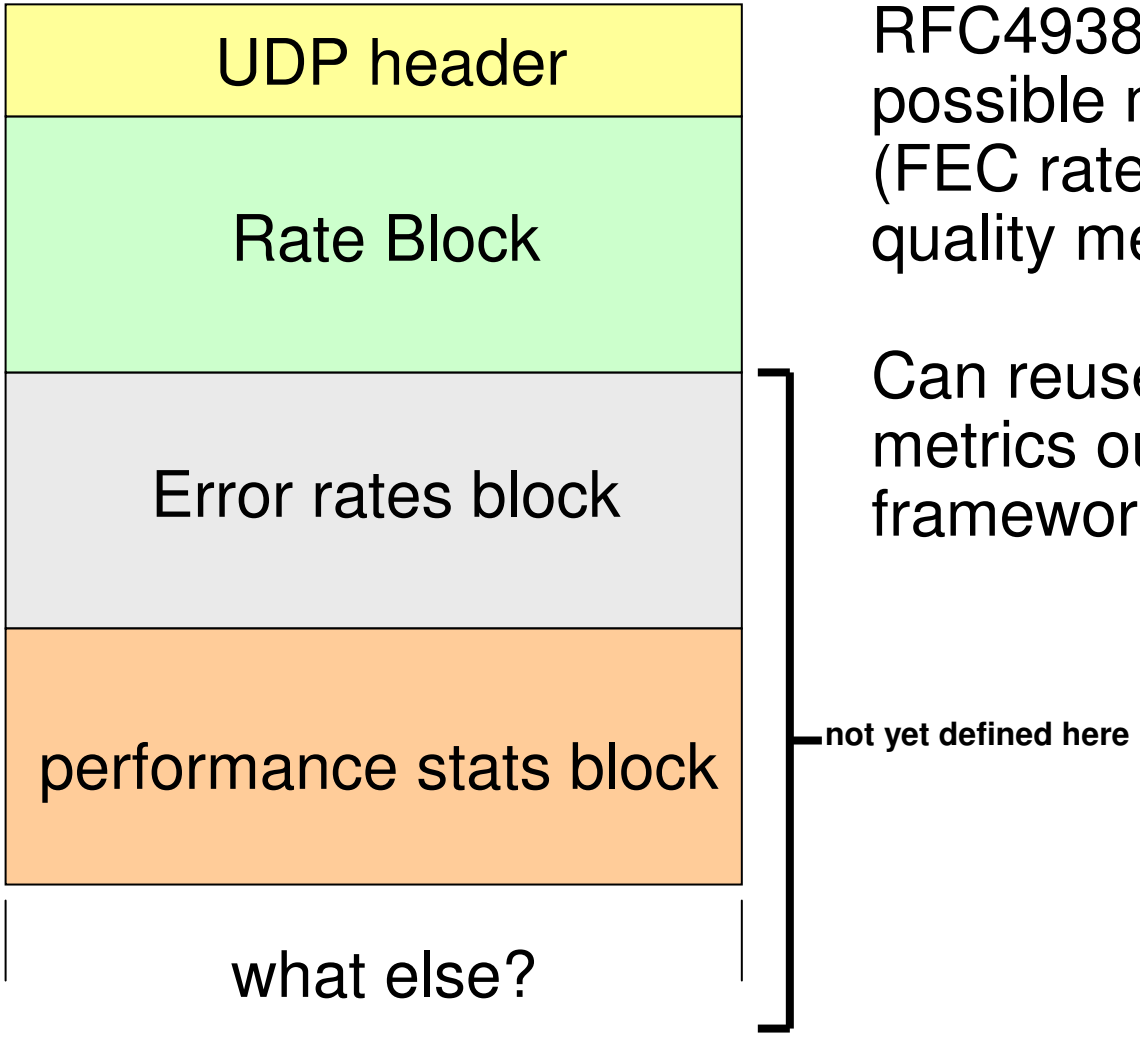

RFC4938 work has a handle on possible metrics for other blocks (FEC rates, bit error rates, link quality metrics…)

Can reuse these thought-out metrics outside the PPPoE framework.

Questions raised by this approach #1

- Where is this applicable?

  All sorts of wireless modems, cable modems…
  will ADSL always be using PPPoE?

- Is providing a current rate to describe an interface sufficient, or is explicit flow control needed?
- If a rate is enough, what granularity indication is appropriate? (favour exact bps. Is that misleading?)
- Should hysteresis be introduced to prevent too-frequent updates? If so, should it be in the modem, or in the router?

Questions raised by this approach #2

- How to handle a truly shared modem MAC that isn't point-to-point? How to describe separate rates to each peer modem? How do we leverage a peer discovery mechanism and tie this to routing?
- Security – how to prevent spoofing?
- Header compression happens on the modem, and can't be factored into shaping on the router. Could shaping consider potential compression savings? Is this even worthwhile to consider?

The all-important questions

- Is this appropriate work for the (rechartered) DNA workgroup?
- If not, should this work be somewhere else? (another group/BOF?)

draft-wood-dna-link-properties-advertisement-01

More discussion is needed.
thankyou.