

MOOSE

Multi-level Origin-Organised Scalable Ethernet

draft-malc-armd-moose-00

Malcolm Scott

University of Cambridge Computer Laboratory




Aim: Hierarchical MAC address space

- **Current Ethernet:** manufacturer-assigned MAC address valid anywhere on any network
 - But every switch must store the location of every host
- **Hierarchical MAC addresses:** address depends on location *e.g. [switch ID].[port ID].[host ID]*
 - Route frames according to hierarchy
 - Small forwarding databases
 - Run a routing protocol between switches
 - One “subnet” per switch – e.g. “02:11:11:00:00:00/24”
 - Don’t advertise individual MAC addresses (cf. TRILL Rbridges)
- **LAAs?** High administrative overhead. So, instead....:

MOOSE

- “NAT for Ethernet”

- Dynamically allocate hosts hierarchical addresses
- Perform source **MAC address rewriting** on ingress
- No encapsulation: no costly rewriting of dest address
- Looks like Ethernet from outside: *transparent to hosts*
- We have an OpenFlow implementation



Beyond simple protocols

- Some protocols must be rewritten by switches
 - Anything which puts MAC address in payload
 - ARP, DHCP: trivial for switches to deal with
- Broadcast: unfortunate legacy
 - Propagate broadcast traffic using reverse path forwarding (PIM): no explicit spanning tree protocol
- Multicast and anycast for free
 - (if we use a suitable routing protocol)
 - May be able to convert broadcast into multicast by inferring groups (e.g. DHCP servers) – see SEATTLE

This is ongoing research;
comments very welcome

*This was a very brief overview: much more detail in
draft-malc-armd-moose-00*

Malcolm Scott


Malcolm.Scott@cl.cam.ac.uk

<http://www.cl.cam.ac.uk/~mas90/MOOSE/>

- Spare slides follow

Mobility

- If a host moves, it is allocated a new MAC address by its new switch
- Other hosts may have the old address in ARP caches
 1. **Forward frames,**
IP Mobility style
(new switch discovers host's old location by querying other switches for its real MAC address)
 2. **Gratuitous ARP,**
Xen VM migration style



Allocation of host identifiers

- Only the switch which allocates a host ID ever uses it for forwarding
 - More distant switches just use the switch ID
- Therefore the detail of how host IDs are allocated can vary between switches
 - Sequential assignment
 - Port number and sequential portion
(reduces address exhaustion attacks)
 - Hash of manufacturer-assigned MAC address
(deterministic: recoverable after crash)

Security and isolation benefits

- The number of switch IDs is more predictable by the network admin than the number of MAC addresses
 - Address flooding attacks are ineffective
- Host-specified MAC address is not used for switching
 - Spoofing is ineffective

