Internet Engineering Task Force	M. Scott, Ed.	
Internet-Draft	D. Wagner-Hall	
Intended status: Informational	J. Crowcroft	
Expires: April 21, 2011	University of Cambridge	
	October 18, 2010	

TOC

Addressing the Scalability of Ethernet with MOOSE
draft-malc-armd-moose-00

Abstract

Ethernet does not scale well to large networks. The flat MAC address space, whilst having obvious benefits for the user and administrator, is the primary cause of this poor scalability; other recent efforts to improve upon Ethernet's scalability have addressed symptoms, rather than this underlying cause. MOOSE, Multi-level Origin-Organised Scalable Ethernet, is an Ethernet switch architecture that performs in-place rewriting of MAC addresses in order to impose a hierarchy upon the address space without reconfiguration or modification of connected devices. This removes the need for switches to maintain large forwarding databases, is of direct use in implementing improved routing, and allows for a variety of other scalability and security innovations. MOOSE also includes a globally-scalable, distributed and resilient protocol for the automatic assignment of addresses to switches, and for detecting and cheaply resolving addressing conflicts.

Status of this Memo

Table of Contents

---

1. Introduction

TOC

Ethernet has lasted well since its inception in the '70s with Ethernet frame-structure and addressing remaining ubiquitous in the data centre environment as in many others. Alongside IP and IP-transported services such as iSCSI, it is now commonplace to see converged network services such as physical disk interfaces and cluster interconnects layered directly over Ethernet (e.g. ATA-over-Ethernet and variants of

Infiniband). However, Ethernet exhibits scalability issues on networks of more than a few thousand devices, such as costly and energy-dense address table logic and storms of broadcast traffic.

Aside from more physical devices, virtualised infrastructure further increases the density of Ethernet addresses in data centres. Widely-used layer-2 virtualisation (Clark, C. and others, “Live Migration of Virtual Machines,” 2005.) [Cl05] mandates a unique Ethernet address per virtual machine. This means that each physical machine in a data centre may represent many tens of Ethernet devices.

The traditional method of avoiding such problems is the artificial subdivision of a network, but this introduces an administrative burden, requires significant routing equipment and also precludes seamless migration--a necessity for virtualised infrastructure. While IP Mobility (Perkins, C., “IP Mobility Support for IPv4,” August 2002.) [RFC3344] addresses the problem of maintaining higher-layer connections when roaming between subnets, it requires client support that is neither ubiquitous or reliable. Common practice sees the provision of one physical Ethernet network covering an entire data centre, or even an entire WAN of data centres.

Our approach, Multi-level Origin-Organised Scalable Ethernet (MOOSE), provides all the advantages of an Ethernet network without the capital and running costs and administrative overhead of a IP router-based approach. MOOSE does this by providing a hierarchical addressing scheme without requiring host reconfiguration or modification.

Ethernet's scalability is limited firstly by the forwarding database that every switch in an Ethernet (IEEE, “802.1D: Standard for Local and Metropolitan Area Networks: Media Access Control (MAC),” 2004.) [802.1D] network must maintain. A switch's forwarding database contains one entry per source address seen in any frame passing through that switch, and stores that MAC address together with the learnt location of that address--the port on which packets from that address were last seen. This is later used to determine on which port to transmit frames destined for that address. Devices frequently broadcast frames throughout the network (e.g. ARP queries) so active devices on the network are listed in most switches' forwarding databases most of the time.

In modern switches the capacity of this database is generally of the order of 16,000 entries. (Higher-capacity forwarding databases exist but are currently constrained to very high-end switches.) On a moderately large network, full databases are a serious risk. If the database becomes full, entries will be discarded; frames for unknown addresses are flooded to all ports and the resulting traffic storm could cause major problems, especially in the presence of low-capacity edge links.

Traditionally the forwarding database has been stored in a content-addressable memory (CAM) as lookups must be very fast, particularly as 10 Gbit/s Ethernet becomes ubiquitous. As networks grow, the number of entries in a switch's forwarding database must naturally increase; however, increasing the capacity of CAMs without sacrificing speed

whilst constraining energy consumption is proving to be challenging. Cheaper switches use DRAM in place of a CAM, but this is likely to remain slower especially for large tables.

Secondly, Ethernet's inability to handle networks containing loops also presents a scalability problem. The Rapid Spanning Tree Protocol, RSTP, must remove loops by disabling any redundant links. On a dense mesh network, RSTP will disable a large proportion of links; this constrains frames to suboptimal routes and may introduce bottlenecks in the network, particularly around the root of the spanning tree. In a data centre environment, this potentially amounts to a very large proportion of capacity being wasted wherever redundant fibres are installed, e.g. between cabinet switches and between data centres.

Thirdly, not only does Ethernet flood frames destined for unknown hosts, but it also uses--and encourages higher-layer protocols to use--broadcast for control messages. For example, ARP (Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware," November 1982.) [RFC0826] performs address resolution via broadcast queries, and DHCP (Droms, R., "Dynamic Host Configuration Protocol," March 1997.) [RFC2131] uses broadcast messages for automatic configuration. It is impractical to replace these protocols entirely as this would require software upgrades to every device, but it would be desirable for the network to minimise the amount of broadcast traffic required to be forwarded.

In this document we identify the relevant underlying problems in the design of Ethernet, review previous work and present the MOOSE switch architecture, which addresses inadequacies in the fundamental operation of Ethernet in a novel yet backwards-compatible way. By revisiting the addressing scheme itself, rather than simply addressing symptoms of the problem as many previous proposed solutions have done, we can go about solving all of the above scalability problems and more.

1.1. Requirements Language

TOC

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

---

2. Ethernet's Underlying Problem

TOC

The original Ethernet was a shared-medium network, where every frame was broadcast and no switching took place. Modern-day wired Ethernet-based networks instead consist almost entirely of point-to-point links; as a result of this, the distinction between unicast, broadcast and

multicast has become more important. 802.11 wireless LANs are the one remaining vestige of Ethernet operating over shared media, where one switch (access point) serves many hosts on the same radio channel. Ethernet's poor scalability arises in various guises, as outlined above. It would seem at first glance that these are entirely distinct and unrelated. However, there is a common underlying cause: that MAC addresses provide no location information.

Globally-unique MAC addresses are structured such that the first three bytes of a device's address contain an organisationally unique identifier (OUI) allocated to the device's manufacturer by the IEEE, with the remaining three bytes allocated by the manufacturer. This hierarchy exists solely for the purpose of allocating unique addresses in a decentralised fashion, and is of no use to Ethernet switches, which must treat the unicast address space as flat.

A flat address space has the advantage that no configuration of devices is required; a device can use its unique, manufacturer-assigned MAC address anywhere on any network. However, this leaves each switch with the task of discovering and storing the location of every addressable device.

If the MAC address space were not flat, but instead contained enough information to locate the device possessing the address, several advantages would be gained. Firstly, large forwarding databases would no longer have to be maintained on every switch. This location information could instead be distributed across the network so that frames are directed towards their destinations according to successive stages of a hierarchy.

Secondly, a hierarchical MAC address space would also make the addition of shortest-path routing considerably easier. Shortest-path routing is clearly a desirable property for a network, yet it is one that Ethernet does not provide. Flat addressing does not lend itself to easy routing: any address can be located anywhere on the network, which means either advertising every host's MAC address via the routing protocol--which scales very poorly--or providing some other location lookup service. The use of hierarchical addresses, with each switch handling a block of sequential addresses akin to an IP subnet, would reduce the routing problem to the one that routing protocols were designed to solve.

Thirdly, this would allow for reduction of broadcast traffic in a variety of different ways. Hierarchical MAC addresses could, for example, be mapped directly and deterministically onto the IP address space, if appropriate for the specific deployment. This would allow switches to respond directly and simply to DHCP and ARP queries, avoiding the need to forward the most common sources of broadcast frames. Alternatively, a distributed directory service can be used, which is less limiting and is thus our preferred approach as detailed below.

The facility for network administrators to assign locally administered addresses (LAAs) to devices has existed for as long as Ethernet. However, configuring and maintaining the LAA on every device based upon where they are connected would be a considerable and unwelcome

administrative overhead. We therefore present MOOSE, a system for applying hierarchical addressing to an Ethernet transparently and without any configuration to edge devices.

---

3. Related Work

TOC

It is well-known that traditional Ethernet scales poorly, and there have been various attempts in recent years to rectify this. The most widely-used of these in real-world networks is MPLS-VPLS (Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture," January 2001.) [RFC3031] (Multiprotocol Label Switching--Virtual Private LAN Service). This connects Ethernet islands together through tunnels across a MPLS cloud. MPLS works by adding one or more labels to the start of every frame, i.e. encapsulating the frame inside its own protocol.

In MPLS-VPLS, the label edge routers (LERs) must determine the frame's initial label(s) based upon the destination address via a lookup table. Frames follow prenegotiated label-switched paths (LSPs) that, unlike Ethernet, are not constrained to follow a spanning tree; LSPs are precomputed at connection setup time and the relevant next hop is stored in a lookup table on each intermediate switch. Each switch must hence use each frame's label to index into this lookup table to determine how to switch the frame.

The effect, once the connection has been negotiated, is to provide what appears to be one or more large Ethernet networks, transparently overlaid on the MPLS cloud. Whilst this solves effectively the problem of shortest-path routing across the MPLS cloud, the overlay Ethernets are still susceptible to the usual scalability problems--and in fact VPLS adds further large lookup tables on every switch that can in some configurations scale even worse than Ethernet's forwarding databases. LERs must map every MAC address to a LSP; label switch routers (LSRs) must store the next hop for every LSP in which they participate, which in the core of the network could scale as O(hosts^2).

A similar scheme is proposed by Hadzic (Hadzic, I., "Hierarchical MAC Address Space in Public Ethernet Networks," 2001.) [Ha01], with the difference that Ethernet-inside-Ethernet encapsulation is used rather than a new protocol. This has the advantage that less processing is required on intermediate switches in the backbone network. However, routes across the backbone are constrained to a spanning tree, and encapsulating switches must obtain a new destination address for every frame using a lookup table that--like Ethernet's forwarding database--must contain every transmitting MAC address. Due to its heavy basis on Ethernet, this shares many of Ethernet's scalability problems.

SmartBridge (Rodeheffer, T., Thekkath, C., and D. Anderson, "SmartBridge: A Scalable Bridge Architecture," 2000.) [Ro00] and RBridges (Perlman, R., "RBridges: Transparent Routing," March 2004.)

[Pe04] (TRILL (Touch, J. and R. Perlman, “Transparent Interconnection of Lots of Links (TRILL): Problem and Applicability Statement,” May 2009.) [RFC5556]) both encapsulate Ethernet frames in a new inter-switch protocol, and run a link-state routing protocol between switches. The link state graph includes the location of every MAC address--necessary because the address space remains flat and any address could appear anywhere--i.e. it again contains every host. Furthermore, switches must perform expensive computation to update routing tables whenever a MAC address joins or leaves the network. Myers et al (Myers, A., Ng, E., and H. Zhang, “Rethinking the Service Model: Scaling Ethernet to a Million Nodes,” November 2004.) [My04] suggest that Ethernet's main failing is its broadcast service, and propose a new architecture in which hosts make explicit use of directory services operated by switches rather than broadcasting queries. It is clear that switches' participation is necessary in order to deal with the broadcast problem; however the modifications to Ethernet suggested are not backwards-compatible and would require at least software modifications to all connected devices. Ethernet is, perhaps unfortunately, too widespread for this to be practical; transparent interception of broadcast frames and subsequent local handling or redirection via multicast or unicast remains the only practical solution. The use of hierarchical addressing is a useful stepping-stone to such a system, and our architecture includes a transparent directory service (ELK) for this purpose.
SEATTLE (Kim, C., Caesar, M., and J. Rexford, “Floodless in SEATTLE: A Scalable Ethernet Architecture for Large Enterprises,” 2008.) [Ki08] takes a more scalable approach. A routing protocol is operated between switches, but in contrast to the approaches described above and in common with MOOSE, the routing protocol only propagates switch location information, rather than every MAC address on the network. Flat MAC addresses are still used, and hence a mechanism is required to look up the switch to which a given address is connected. This is achieved by using a distributed hash table (DHT) operating on participating switches with local caching to alleviate load. This is certainly a step in the right direction but introduces considerable complexity to switches, since they now must maintain and update the DHT continually, and it is clear that a SEATTLE switch would have a significant software component in the data path. MOOSE alleviates some of the complexity of SEATTLE by a combination of hierarchical addresses and delegation to a separate directory service.

---

4. MOOSE Architecture

TOC

The basic operation of MOOSE is to assign a new hierarchical MAC address to each host on the network, assigned dynamically and automatically from the unicast LAA space. This dynamically-assigned

address is referred to as a MOOSE address to avoid confusion with hosts' static, manufacturer-assigned MAC addresses.
Every frame entering the network has its source address rewritten in-place to the sending host's MOOSE address by the first MOOSE-aware switch it traverses. The switch that performs address rewriting for a host--i.e. the closest MOOSE switch to that host--is the host's home switch and is responsible for assigning a MOOSE address to that host. (If non-MOOSE switches or hubs are in use, a host may have more than one "closest" MOOSE switch, in which case an RSTP-like protocol must be used to elect a switch to handle each edge segment.)
The destination address is left intact in the expectation that it already is a MOOSE address. Hosts' ARP caches will already contain the MOOSE addresses of any hosts being communicated with as any packet received will already have had its source address rewritten; a host's manufacturer-assigned MAC address is never seen outside of the segment containing that host. This is a crucial point since encapsulation-based technologies such as MPLS do not reveal to the destination host the address used for routing; as a result, switches must also convert destination as well as source addresses of frames entering the network. In other words, once again switches must maintain large tables of remote hosts on the network. The only destination rewriting that MOOSE switches perform, however, is of the destination addresses of frames destined for local hosts back to their manufacturer-assigned MAC addresses; this is simple as the required information is already known, and necessary because otherwise that host's network interface card would discard the frame as misaddressed.
A MOOSE address consists of a switch identifier followed by a host identifier. For our examples, we simply use a fixed three-byte switch identifier followed by a fixed three-byte host identifier:

Since these two identifiers when concatenated must form a unicast LAA, the settings of two bits in the first byte of the switch identifier are fixed: the least significant bit must be 0 to indicate a unicast address, and the second-least significant bit must be 1 to indicate a LAA. To cater for variable length switch identifiers, some means of

introducing separation between the switch and host identifiers is required. Two possible implementations would be for:

1. the first three bits of the address to indicate how many of the following 5-bit blocks make up the switch prefix;

2. some constant delimiter to appear between the switch identifier and host identifier, with switch identifiers not allowed to contain the delimiter.

The former is simple and gives eight classes of switch identifier. Because the size of a MOOSE network is limited by the placement of IP routers, these classes should be sufficient. Additionally, because switches are free to change their identifiers, they may trivially switch to a larger class if they have too many attached hosts, or if a smaller class becomes full.
The latter removes the fixed classes, allowing for more flexibility with the sizes of switch identifiers, at the cost of complexity, and a reduction in the available address space.
Each switch can select for itself a unique switch identifier, as identifier conflict resolution is cheap (see below). When first joining the routing protocol, conflict should be very unlikely, as the switch will in the process gain an up-to-date list of in-use identifiers.
Depending on requirements, the switch identifier may itself be a hierarchical address--e.g. six bits to identify a network area followed by two bytes to identify a switch within that area--which could then be used to aid routing decisions.
Each host is assigned a host identifier by its home switch from the pool of identifiers available to that switch. Only a host's home switch ever bases a switching decision on the host identifier, so the detail of how these are allocated can vary from switch to switch. Suitable schemes include:

1. sequential assignment;

2. the port number followed by a sequential portion (to allow for multiple hosts connected to one port);

3. a hash of the host's real MAC address.

The latter two approaches are preferable to a simple sequential assignment, as they better isolate certain kinds of denial-of-service attack in which a malicious host attempts to use up all available host identifiers on the switch. They also require less state to be shared between ports. The third option has the further advantage that it is deterministic and hence can be recovered easily in the event of a crash.
It is hence possible to route frames through the network to remote hosts by simply inspecting the switch identifier in the destination

address, and ignoring the host identifier until the frame reaches the destination host's home switch. Switches no longer need to keep a table of all MAC addresses seen recently; they only need store the locations of other switches and of any directly-connected hosts.
As well as reducing the amount of data that must be consulted in order to make switching decisions, this provides extra resilience by making this data much more predictable. The number of MAC addresses in a network can increase unexpectedly in the event of an address flooding attack or even under normal operation if the network contains open wireless access points; relying on the MAC address list for forwarding leads to some of the vulnerabilities of Ethernet. The set of switch identifiers participating in MOOSE switching, on the other hand, is kept predictable and manageable by ensuring that neighbouring switches (discovered using LLDP (IEEE, “802.1AB: Station and Media Access Control Connectivity Discovery,” 2009.) [802.1AB]) are authenticated before they can participate in the routing protocol. This authentication can be achieved at layer 3 using the security features found in most popular routing protocols and/or at layer 2 (IEEE, “802.1X: Port Based Network Access Control,” 2004.) [802.1X]. As the switch identifier is the only address consulted for forwarding decisions, a MOOSE switch is likely to remain reliable in the face of attacks that could have brought down a traditional Ethernet. Furthermore, any attacks based upon MAC address spoofing cannot function on a MOOSE network as the user-provided MAC address is translated immediately.

4.1. Shortest Path Routing

TOC

As described so far, MOOSE switches must still forward frames along a spanning tree. As discussed above, this is an undesirable property of Ethernet as it can cause frames to take a highly suboptimal path through the network. The foundations are in place to do much better than this using shortest-path routing.
For the purpose of frame forwarding, a MOOSE switch can be considered akin to a layer 3 router; it has one locally-connected subnet--containing all addresses starting with its switch identifier--and delivers frames to other subnets by passing them to an appropriate neighbouring switch. Bearing this in mind, the switch can run a routing protocol of the kind normally used for IP, such as a variant of OSPF (Moy, J., “OSPF Version 2,” April 1998.) [RFC2328]. This allows frames to be routed along the shortest available path, rather than being constrained to a spanning tree. A multipath variant such as OSPF-OMP may be particularly desirable due to its ability to make use of multiple equal-cost routing paths in order to improve performance.

4.2. Address Selection and Conflict Resolution

TOC

For reasons akin to those of the flaws of Ethernet, it is undesirable to guarantee universally unique pre-determined MOOSE switch identifiers. Due to the reduced size of the switch ID space compared to the MAC address space, this would also be infeasible. We therefore propose that each switch selects an initial address for itself during startup. This could result in more than one switch claiming an address, which would be undesirable, so to mitigate the potential for MOOSE addresses to find themselves in conflict we additionally propose a simple and inexpensive conflict resolution protocol.
Suppose two switches each have the same identifer. We note that if these switches are on separate MOOSE networks (on disconnected networks, or separated by an IP router), this situation brings no issue. Should they be on the same MOOSE network, however, a conflict exists and must be resolved. Any routing protocol would require a switch to know which port other switches are connected to, for instance by OSPF neighbour lists, or simply by receiving frames and noting the switch port and source MOOSE address. When a switch receives a MOOSE frame, it looks up the source switch in its forwarding database, which is likely in fast Content Addressable Memory. If it finds that source switch to be on a port other than that which it recognises from its table, one of three situations may be possible:

1. the source switch may be the same as the known switch, and have physically moved, or a topology change has occurred;

2. the source switch may be a different one to the known switch, and they are in conflict;

3. the source switch may be the same as the known switch, but is sending frames down a different route to the last used route.

To avoid disruption to the network in the first case, and to give scope for switches to migrate within the network, the switch which detected the possible conflict should ascertain whether the known switch is still alive and present. The conflict-resolving switch thus attempts to send a unicast frame to the known switch, via the port stored in the forwarding database, asking whether it is there at a regular interval until a timeout. This will reach the known switch rather than the new switch if it is still present as other switches beyond that port must not have detected the conflict yet. The nature of the timeout we leave unspecified, and can be implementation specific. It may, for instance, be a pre-defined constant, or it may vary based on QoS information gathered if such capabilities are supported. When a MOOSE switch receives such a frame, it should promptly respond with an acknowledgement frame, showing that it is alive.

If, within the timeout period, the conflict resolver finds the known host not to be alive, no conflict exists, so the switch updates its view of the network by removing the old entry from its forwarding database and triggering a routing protocol refresh.
If, on the other hand, the host is found to be alive, a conflict exists. The conflict resolver then sends a frame to the more recently found switch indicating that it is in conflict and should change its address. That switch, upon receiving this frame, changes its address and sends a gratuitous ARP for each of its connected hosts, so that the rest of the network is aware of the change. To mitigate the risks of a denial of service attack, or faulty equipment sending out conflict frames, an exponential backoff algorithm should be used when receiving conflict notification frames.
A switch should have a timer, and counter influencing the maximum value of the timer, both initialised to 0. When a conflict notification frame is received, the counter is incremented (subject to a saturation value to avoid excessive timeouts). After a conflict has been resolved--i.e. the switch has changed its address--a timer starts counting down from some time exponential in that counter; subsequently the switch will only change its address if the timer has returned to 0 by the time the conflict frame is received. The counter should be reset to 0 when the timer reaches 0. Using this scheme the event of true conflict is handled quickly, even in the unlikely case that the newly acquired address is also in conflict. Any node emitting malicious or erroneous conflict notifications, however, is rate-limited enough that their damage potential is much restricted, subject to a sufficient timer being chosen.
Pseudocode: Conflict resolution backoff:

```
    if timer > 0:
        if counter < counter_max:
            counter = counter + 1
        # Discard conflict notification frame
    else:
        timer = k^counter
        change_address()
```

Pseudocode: Conflict resolution timer:

```
    foreach clock tick do:
        if timer > 0:
            timer = timer - 1
        else:
            counter = 0
```

This could be further enhanced by detecting repeated conflicts involving the same switch or switches, in a manner similar to BGP Route Flap Damping (Villamizar, C., Chandra, R., and R. Govindan, “BGP Route Flap Damping,” November 1998.) [RFC2439], and performing more

aggressive steps to avoid further conflicts--for example using a significantly increased timeout, and/or having *both* switches in conflict select new addresses.
The conflict resolution algorithm brings a marked improvement on the equivilent vulnerability of Ethernet, that MAC addresses can be spoofed. We build in a flexible, well-defined system of recovery. The decentralised nature of the system makes it much less open to denial of service attack than any centralised directory may be. Having every MOOSE switch acting as a barrier to the propagation of packets from addresses in conflict provides a strong separation between recently bridged networks with conflicting addresses, so that communication within the individual networks may continue without modification, until bridge-crossing traffic appears, at which point resolution quickly happens. We also remove the possibility for forwarding databases to frequenty have to switch their entry for a conflicted address, which can happen with MAC conflicts in traditional Ethernet. Additionally, in the case of a switch identifier spoofing attack, the conflict resolver acts as a hard boundary for the effects of such an attack.
It is possible that the switch performing conflict resolution could send a suggested replacement switch address to the switch in conflict, known by the conflict resolver to have a low probability of being present on the network (because it is not present in its forwarding database). This would reduce the chance of repeated collisions, and potentially allow for longer backoff periods, but may be premature optimisation.
Because multi-path routing is often desirable, we could introduce an extra datum during the source address rewriting performed by MOOSE switches. When an ingress MOOSE switch rewrites the source address of an Ethernet frame to a MOOSE address, it could also prepend some hash of its manufacturer-assigned MAC address to the data field, and increment the length field as necessary. The egress switch, when rewriting the MOOSE destination address to a host's MAC address, then strips out this added datum. This allows the conflict resolver to check whether conflicts actually exist by local lookup, rather than probing other switches, at the cost of added memory requirements in every switch. This may push the frame to be larger than Ethernet's maximum, so may require fragmenting the packet into two, at small added cost. Alternatively, assuming jumbo frames are permitted by the hardware, the maximum frame size could be marginally reduced to allow for this in the same manner as for 802.1Q VLAN tags.
From the cheapness of conflict resolution, certain other address management tasks become simple. A switch is free to choose its address when it joins the network however it wishes--attempting to re-use its last-used address, from a list of preferred addresses, or by generating an address entirely at random. More intricate addressing schemes may be used on managed networks if desired, perhaps encapsulating deeper layers of hierarchy.

4.3. Broadcast and Multicast

TOC

Since Ethernet does still need to support arbitrary broadcast frames, these must still be forwarded along a spanning tree in order that they reach each host exactly once. An explicit spanning tree protocol is not required however, as the tree can be deduced from the routing table via reverse path forwarding in a similar manner to Protocol-Independent Multicast (PIM) (Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)," January 2005.) [RFC3973]. In other words, broadcast packets are routed as if they had been sent to the all-hosts multicast group. More general multicast groups can be implemented using a combination of IGMP snooping (Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches," May 2006.) [RFC4541] as used by modern Ethernet switches, and participation of the MOOSE switches in PIM routing.

4.4. Example

TOC

To illustrate the basic behaviour of MOOSE switches, before we go on to describe further features, we will offer a simple example. We will describe the steps involved in forwarding a broadcast frame containing a query in some higher-layer IPv4-based protocol, and subsequent unicast frame containing the response, between two hosts A and B via three MOOSE switches 02:11:11, 02:22:22 and 02:33:33.

4.4.1. Query

TOC

1. Host A transmits the broadcast query frame as it would on any Ethernet network, with its own manufacturer-assigned MAC address in the Ethernet header's source field and the broadcast address (FF:FF:FF:FF:FF:FF) in the destination field.

2. The frame is received by switch 02:11:11, which observes the non-MOOSE address in the frame's source field, and rewrites the source field into a MOOSE address containing the switch identifier and the appropriate host identifier. As this is Host A's first frame, the switch must allocate a host identifier (in this case 00:00:01, making Host A's complete MOOSE address 02:11:11:00:00:01).

3. The three switches broadcast the frame using reverse path forwarding away from Host A.

4. The frame is received by Host B (and any other hosts on the network) in its current form; no further rewriting is performed.

4.4.2. Response

TOC

1. Host B looks up Host A's IP address in its ARP cache to determine a suitable destination address for the response frame. Since the rewritten query frame arrived at Host B with the source field containing the MOOSE address 02:11:11:00:00:01, this is the address returned by the cache lookup.

2. As above, switch 02:33:33 assigns a MOOSE address to Host B (02:33:33:00:00:01) and rewrites the source address of the frame.

3. The frame is now routed through the network based solely on the destination switch identifier--the host identifier is ignored for now. The routing table is consulted for the location of switch 02:11:11 and the frame is forwarded accordingly.

4. On receiving the frame, switch 02:11:11 observes that it is destined for a directly-connected host (02:11:11:00:00:01). It prepares the frame for transmission along its final hop by rewriting the destination address to Host A's manufacturer-assigned MAC address. The source field of the frame is again left as the MOOSE address of Host B in order that this address is used for any further communication with Host B.

4.5. Directory Service

TOC

A directory service, Enhanced Lookup (ELK), runs in conjunction with the basic MOOSE switch described so far. ELK exists to handle ARP and DHCP queries in a broadcast-free manner by learning mappings from IP addresses to MOOSE addresses. The master ELK directory is served by one or multiple systems for resilience and is reached using an anycast MOOSE address; the layer-2 anycast feature is a convenient side-effect of running a routing protocol. Slave copies of the directory can be

held nearer the edge of the network in order to take load away from the masters; slaves can be reached for lookups via a separate anycast address, and the entire herd of ELK can be kept synchronised via the masters using a combination of multicast and unicast.
MOOSE switches intercept ARP and DHCP packets broadcast by hosts and convert them into anycast ELK queries to the nearest slave (for ARP) or master (for DHCP). (DHCP handling could make use of the protocol's existing DHCP relay mechanism.) The ELK slave answers ARP queries directly using information in the directory; as it does so, if the query is from a host not in the directory, it learns the sender's IP address to MOOSE address mapping. The ELK master can also act as a DHCP server, populating the ELK directory as it grants IP address leases to clients.
The one case in which the ELK directory will not contain the answer to a query is when answering an ARP request for a host that is not configured to use DHCP and that has not yet itself sent an ARP packet (i.e. has not yet communicated via IP). This must be dealt with by flooding the query to every active switch port, in a manner akin to current Ethernet switches, and caching the result in the ELK directory. Although this is not ideal, it is necessary in order to deal with this scenario in a compatible manner, and is unlikely to happen frequently.

4.6. Mobility

TOC

A consequence of introducing location-based hierarchy into MAC addresses is the need to explicitly handle host mobility. In a traditional Ethernet, hosts can migrate between switches as the switches will learn the host's new location as soon as it sends a frame. With MOOSE, if a host relocates to a new switch its address changes and any ARP cache entries on other hosts pertaining to the migrated host become incorrect; frames will continue to be sent to the host's old location for a while. There are two strategies for dealing with this, which can be used separately or in conjunction:

1. The previous home switch of the migrated host can forward frames sent to the host's old address until outdated ARP cache entries expire. This is similar to IP Mobility: the previous home switch essentially becomes a care-of agent for the host. However, unlike IP Mobility, it requires no host support. A handover protocol is necessary for the old and new home switches to set up such forwarding: on the arrival of a new host at a switch, that switch would ask all other switches (via multicast) whether any had seen this host before, identifying it using its manufacturer-assigned MAC address, and would instruct such switches to redirect frames.

2. A broadcast ARP announcement (or "gratuitous ARP") can be sent by the new home switch to immediately update remote ARP caches and the ELK directory with the new MOOSE address. This is the technique used by Xen when migrating live virtual machines. Unlike the previous approach, this works even if the previous switch is no longer reachable, for example if this host migration was as a result of a switch failure. This is a simpler approach as a handover protocol is not required, but results in additional broadcast traffic.

Unless the frequency of host migrations is very high, the additional load introduced by either mobility approach is expected to be negligible.
Illustration of the two ways to handle a host A roaming onto another switch whilst maintaining communication with another host B:

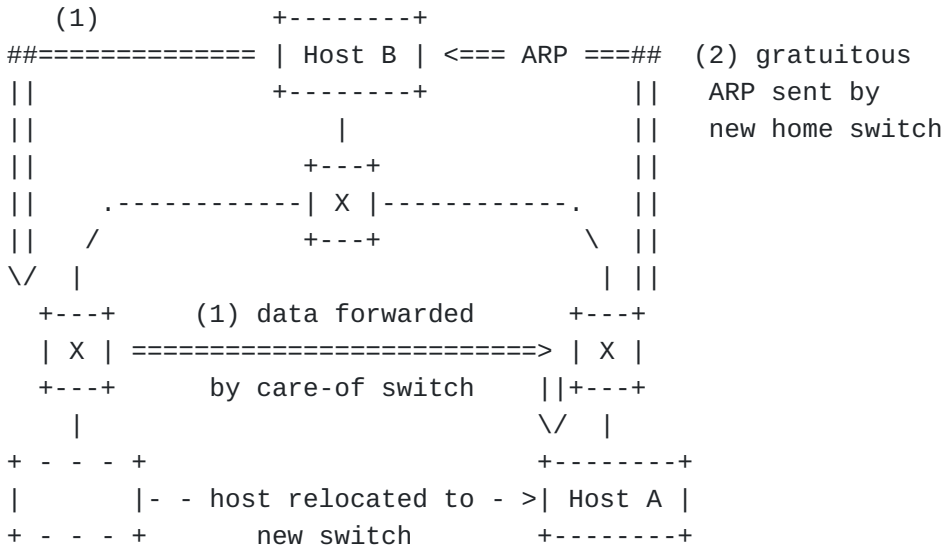

5. Interoperability Considerations

TOC

5.1. Layer-violating Protocols

TOC

In an ideal world, free from layering violations, all layer 3 protocols would operate correctly on top of MOOSE in exactly the same way that they currently operate on top of Ethernet, with no protocol-specific handling necessary in the switch. In reality, however, protocols abound which use hosts' MAC addresses for purposes other than layer 2

addressing or which place MAC addresses in the frame payload. DHCP and ARP have already been mentioned as such protocols which must be specifically handled by edge switches in order to operate; luckily, the rewriting required for these important protocols is simple.
Of particular concern are recent standards for layering on top of Ethernet protocols which were previously used solely on dedicated hardware interconnects, such as Fibre Channel over Ethernet (FCoE (T11 FC-BB-5 working group, "Fibre Channel Backbone - 5," June 2009.) [FC-BB-5]). In order to support FCoE and similar protocols on a MOOSE network, each edge switch will need to be able to interpret and rewrite individual protocols that are in use. A production MOOSE switch would, therefore, need to be implemented such that it is possible to add rewriting support for additional protocols after manufacture, for example by loading an additional software or FPGA configuration module. Ultimately, in the general case, this problem could be addressed more satisfactorily by extending the Ethernet standard to provide a protocol-agnostic method for a layer 2 network to inform hosts of their own addresses; LLDP (IEEE, "802.1AB: Station and Media Access Control Connectivity Discovery," 2009.) [802.1AB] would make a good basis for this extension. This would allow the use of network-assigned MAC addresses for any protocol, with some rewriting performed either partially (within the frame payload) or fully by the host itself, and furthermore would allow higher-layer protocols to respond to changes of the host's network-assigned address (e.g. due to mobility). Such a mechanism could be deployed incrementally as needed, with switches able to perform address rewriting for hosts which are not able to do this themselves. This is, however, a very long-term solution, and protocol-specific rewriting on the switch is likely to be required for the foreseeable future.
FCoE in particular is unusual, however, as it already does its own dynamic allocation of MAC address to devices. It is conceivable that an extension to FCoE could be developed which allows a network-wide dynamic address assignment scheme such as MOOSE to be exploited to provide addresses directly to fibre channel devices.

5.2. Edge Virtual Bridging

TOC

The rise of virtualisation has caused an unanticipated proliferation of software switches, usually in the host operating system or hypervisor which provides network connectivity to multiple virtual machines. Since software switches are almost always neither fast nor centrally manageable in the same way as hardware switches, there is ongoing work to standardise--by Cisco as Port Extension and by the IEEE as Edge Virtual Bridging (Jeffree, A., Congdon, P., and J. Pelissier, "P802.1Qbg: Edge Virtual Bridging," September 2009.) [P802.1Qbg]--a means of making these software switches act merely as additional ports

which are logically part of a more central hardware switch. This reduces the work required by a virtual edge switch: frames from local virtual edge ports can be forwarded straight out via the uplink to a physical switch without consideration, and frames from the uplink will arrive simply tagged with a virtual edge port identifier.
(The scope of Port Extension in particular is greater than this, and allows for physical port extenders to exist in place of switches where a large number of ports but a small amount of processing is required, but virtualisation is likely to be the most significant use case.)
Edge Virtual Bridging and Port Extension require very little adaptation to be implemented on a MOOSE switch. It is unlikely, although too early in the standardisation process to say for certain, that the virtual bridge will need to be MOOSE-aware. A virtual-bridging-aware physical MOOSE switch will thus simply need to take into account the possibility that one physical port may hide a large number of virtual ports when allocating host identifiers, as it would if it had an Ethernet switch connected on that port. If, however, the virtual bridge is made MOOSE-aware, the hierarchical addressing of MOOSE could be exploited to allow the virtual bridge to allocate host identifiers itself, given that it is likely to be aware of the exact number and nature of virtual edge ports. The parent MOOSE switch would accordingly allocate an address prefix to each child virtual bridge, and hosts' full MOOSE addresses could be formed as:

```
     SWITCH ID   :   CHILD ID   :   HOST ID
      (parent)      (allocated     (allocated
                     by parent)     by child)
```

---

6. Prototype Implementation

TOC

We have implemented a MOOSE switch in OpenFlow and NOX, which can be run on off-the-shelf switches. Details can be found in our paper (Wagner-Hall, D., “A Prototype Implementation of MOOSE on a NetFPGA/OpenFlow/NOX Stack,” September 2010.) [Wa10].

---

7. Conclusions

TOC

Ethernet remains popular due to its simplicity and ubiquity, but is showing its age and exhibits serious scalability issues in large deployments. Previously-proposed improvements address either a few of the problems in a simple way, or most of the problems in a highly complex or backwards-incompatible way. We have demonstrated a simple, novel and easily-implementable approach for significantly boosting the

scalability of Ethernet, which has a working prototype switch firmware implementation.

---

8. IANA Considerations

TOC

This memo includes no request to IANA.

---

9. Security Considerations

TOC

Security will be considered in a later revision of this document.

---

10. Informative References

TOC

[802.1AB]	IEEE, “802.1AB: Station and Media Access Control Connectivity Discovery,” 2009.
[802.1D]	IEEE, “802.1D: Standard for Local and Metropolitan Area Networks: Media Access Control (MAC),” 2004.
[802.1X]	IEEE, “802.1X: Port Based Network Access Control,” 2004.
[Cl05]	Clark, C. and others, “Live Migration of Virtual Machines,” USENIX NSDI 2005, 2005.
[FC-BB-5]	T11 FC-BB-5 working group, “Fibre Channel Backbone - 5,” June 2009.
[Ha01]	Hadzic, I., “Hierarchical MAC Address Space in Public Ethernet Networks,” IEEE GLOBECOM vol 3, 2001, 2001.
[Ki08]	Kim, C., Caesar, M., and J. Rexford, “Floodless in SEATTLE: A Scalable Ethernet Architecture for Large Enterprises,” ACM SIGCOMM 2008, 2008.
[My04]	Myers, A., Ng, E., and H. Zhang, “Rethinking the Service Model: Scaling Ethernet to a Million Nodes,” ACM SIGCOMM Workshop on Hot Topics in Networking 2004, November 2004.
[P802.1Qbg]	Jeffree, A., Congdon, P., and J. Pelissier, “P802.1Qbg: Edge Virtual Bridging,” September 2009.
[Pe04]	Perlman, R., “RBridges: Transparent Routing,” Proc. INFOCOM vol 2, 2005, March 2004.
[RFC0826]	Plummer, D., “Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware,” STD 37, RFC 826, November 1982 (TXT).
[RFC2131]	

	Droms, R., "Dynamic Host Configuration Protocol," RFC 2131, March 1997 (TXT, HTML, XML).
[RFC2328]	Moy, J., "OSPF Version 2," STD 54, RFC 2328, April 1998 (TXT, HTML, XML).
[RFC2439]	Villamizar, C., Chandra, R., and R. Govindan, "BGP Route Flap Damping," RFC 2439, November 1998 (TXT, HTML, XML).
[RFC3031]	Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture," RFC 3031, January 2001 (TXT).
[RFC3344]	Perkins, C., "IP Mobility Support for IPv4," RFC 3344, August 2002 (TXT).
[RFC3973]	Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)," RFC 3973, January 2005 (TXT).
[RFC4541]	Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches," RFC 4541, May 2006 (TXT).
[RFC5556]	Touch, J. and R. Perlman, "Transparent Interconnection of Lots of Links (TRILL): Problem and Applicability Statement," RFC 5556, May 2009 (TXT).
[Ro00]	Rodeheffer, T., Thekkath, C., and D. Anderson, "SmartBridge: A Scalable Bridge Architecture," ACM SIGCOMM 2000, 2000.
[Wa10]	Wagner-Hall, D., "A Prototype Implementation of MOOSE on a NetFPGA/OpenFlow/NOX Stack," First European NetFPGA Developers' Workshop Cambridge, September 2010.

Authors' Addresses

TOC

	Malcolm Scott (editor)
	University of Cambridge
	15 JJ Thomson Ave
	Cambridge, CB3 0FD
	UK
Phone:	+44 1223 763500
Fax:	+44 1223 334678
Email:	Malcolm.Scott@cl.cam.ac.uk
URI:	http://www.cl.cam.ac.uk/~mas90/MOOSE/
	Daniel Wagner-Hall
	University of Cambridge
Email:	dwh@cantab.net
	Jon Crowcroft

	University of Cambridge
	15 JJ Thomson Ave
	Cambridge, CB3 0FD
	UK
Phone:	+44 1223 763500
Fax:	+44 1223 334678
Email:	Jon.Crowcroft@cl.cam.ac.uk