

MIF Current Practices

draft-mrw-mif-current-practices-01.txt

Margaret Wasserman
mrw@sandstorm.net

Goals of the Document

- Collect information about how current implementations operate when they are configured with multiple physical or virtual interfaces
 - Still looking for additional information!
- Generalize or categorize current solutions
- Identify where there are gaps that require improvement
- Note: Not all of these goals have been realized in the current version

Systems Currently Included

- Nokia S60 3rd Edition, Feature Pack 2
- Microsoft Windows Mobile 2003, 2nd edition
- Blackberry
- Google Android
- Arena Connection Manager
- Microsoft Windows XP and Vista
- Linux and BSD Unix-based Operating Systems
- Apple Mac OS X

Two General Approaches

- Application (or user) chooses a “connection” among available choices
 - More common in mobile/cell phone environments
 - All configuration stored on a per-connection basis
 - After initial selection, all communication (including DNS, if any) uses the indicated connection
 - “Connection manager” may be provided
- DNS, address selection and interface selection are performed separately
 - More common in general-purpose operating systems
 - DNS invoked by application, remote address chosen from list returned, local address chosen to match, packets sent to default gateway
 - DNS server list, default gateway (& other routes, if any) stored on a per-system basis

Split-DNS

- Current Implementations fall into two groups:
 - DNS server list configured on per-host basis
 - manually, or by accepting list on primary interface via DHCP
 - DNS server list configured on a per-interface basis
 - manually, or by accepting lists on multiple interfaces via DHCP
- Different OSes use server addresses from lists in different orders/combinations
- In all cases, DNS queries are sent using a source address and outbound interface selected by the IP stack
 - The address and interface are not tied to the interface on which the DNS configuration was received
 - Typically consists of sending DNS queries to a default gateway

Interface Selection

- Selecting the outbound interface and next hop destination is a multi-step process
 - If the destination address is on a local link, address resolution (e.g. ARP) is performed on the corresponding interface, and the packet is sent directly to the destination
 - If not, the routing table is used to determine a longest-match route for the destination
 - Route will include next hop IP & interface
 - Will typically devolve to selecting from the default router list(s), as more specific routes are not typically configured on hosts
 - Selecting the default router is another complex process...

Interface Selection (Default Router)

- Default router list configurations fall into two groups:
 - Per-host default router lists
 - Overwritten by new config, or primary interface wins
 - Per-interface default router lists
- In general the default router with lowest metric is chosen for outbound traffic
- If multiple default routers have the same metric, hosts may:
 - Choose the first router in the list, or a “primary” default router
 - Perform load balancing, such as RFC 4311
- Packets are sent out the interface associated with the default router selected.
- The destination address isn't even considered in default router selection.

Address Selection

- Destination address selection
 - In almost all cases, the destination address will be chosen by the application based on DNS lookup results or some type of application-level configuration
- Source address selection
 - May be selected by the application
 - If not, implementations vary in what they choose:
 - First address on the outgoing interface
 - More complex selection mechanisms that involve choosing the best match for the destination address, such as RFC 3484

Going Forward

- Need information on more implementations
- Need to group implementations and provide more consistent information about them
- Common Solutions section needs to be expanded

If you are familiar with an implementation that is not included, please send text!

If you have more information about an included implementation, please let us know!