

Internet Engineering Task Force (IETF)  
Request for Comments: 6418  
Category: Informational  
ISSN: 2070-1721

M. Blanchet  
Viagenie  
P. Seite  
France Telecom - Orange  
November 2011

## Multiple Interfaces and Provisioning Domains Problem Statement

### Abstract

This document describes issues encountered by a node attached to multiple provisioning domains. This node receives configuration information from each of its provisioning domains, where some configuration objects are global to the node and others are local to the interface. Issues such as selecting the wrong interface to send traffic happen when conflicting node-scoped configuration objects are received and inappropriately used. Moreover, other issues are the result of simultaneous attachment to multiple networks, such as domain selection or addressing and naming space overlaps, regardless of the provisioning mechanism. While multiple provisioning domains are typically seen on nodes with multiple interfaces, this document also discusses situations involving single-interface nodes.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6418>.

## Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction .....	3
2. Terminology .....	4
3. Scope and Existing Work .....	4
3.1. Interactions Below IP .....	4
3.2. MIF Node Characterization .....	5
3.3. Host Requirements .....	5
3.4. Mobility and Other IP Protocols .....	6
3.5. Address Selection .....	6
3.6. Finding and Sharing IP Addresses with Peers .....	7
3.7. Provisioning Domain Selection .....	7
3.8. Session Management .....	8
3.9. Sockets API .....	9
4. MIF Issues .....	9
4.1. DNS Resolution Issues .....	9
4.2. Node Routing .....	12
4.3. Conflicting Policies .....	13
4.4. Session Management .....	14
4.5. Single Interface on Multiple Provisioning Domains .....	14
5. Underlying Problems and Causes .....	15
6. Security Considerations .....	17
7. Contributors .....	18
8. Acknowledgements .....	18
9. Informative References .....	18

## 1. Introduction

A multihomed node may have multiple provisioning domains (via physical and/or virtual interfaces). For example, a node may be simultaneously connected to a wired Ethernet LAN, an 802.11 LAN, a 3G cell network, one or multiple VPN connections, or one or multiple tunnels (automatic or manual). Current laptops and smartphones typically have multiple access network interfaces and, thus, are often connected to different provisioning domains.

A multihomed node receives configuration information from each of its attached networks, through various mechanisms such as DHCPv4 [RFC2131], DHCPv6 [RFC3315], PPP [RFC1661], and IPv6 Router Advertisements [RFC4861]. Some received configuration objects are specific to an interface, such as the IP address and the link prefix. Others are typically considered by implementations as being global to the node, such as the routing information (e.g., default gateway), DNS server IP addresses, and address selection policies, herein referred to as "node-scoped".

When the received node-scoped configuration objects have different values from each provisioning domain, such as different DNS server IP addresses, different default gateways, or different address selection policies, the node has to decide which one to use or how it will merge them.

Other issues are the result of simultaneous attachment to multiple networks, such as addressing and naming space overlaps, regardless of the provisioning mechanism.

The following sections define the multiple interfaces (MIF) node and the scope of this work, describe related work, list issues, and then summarize the underlying problems.

A companion document, [RFC6419], discusses some current practices of various implementations dealing with MIF.

## 2. Terminology

### Administrative domain

A group of hosts, routers, and networks operated and managed by a single organization [RFC1136].

### Provisioning domain

A set of consistent configuration information (e.g., default router, network prefixes, DNS) and the corresponding interface. One administrative domain may have multiple provisioning domains. Successful attachment to the provisioning domain implies that the terminal attaches to the corresponding interface with appropriate configuration information.

### Reference to IP version

When a protocol keyword such as IP, PPP, or DHCP is used in this document without any reference to a specific IP version, then it implies both IPv4 and IPv6. A specific IP version keyword such as DHCPv4 or DHCPv6 is meant to be specific to that IP version.

## 3. Scope and Existing Work

This section describes existing related work and defines the scope of the problem.

### 3.1. Interactions Below IP

Some types of interfaces have link-layer characteristics that may be used in determining how multiple provisioning domain issues will be dealt with. For instance, link layers may have authentication and encryption characteristics that could be used as criteria for interface selection. However, network discovery and selection on lower layers as defined by [RFC5113] is out of scope of this document. Moreover, interoperability with lower-layer mechanisms such as services defined in IEEE 802.21, which aims at facilitating handover between heterogeneous networks [MIH], is also out of scope.

Some mechanisms (e.g., based on a virtual IP interface) allow sharing a single IP address over multiple interfaces to networks with disparate access technologies. From the IP-stack view on the node, there is only a single interface and single IP address. Therefore, this situation is out of scope of this problem statement. Furthermore, link aggregation done under IP where a single interface is shown to the IP stack is also out of scope.

### 3.2. MIF Node Characterization

A MIF node has the following characteristics:

- o A MIF node is an [RFC1122] IPv4- and/or [RFC4294] IPv6-compliant node.
- o A MIF node is configured with more than one IP address (excluding loopback and link-local).
- o A MIF node can attach to more than one provisioning domain, as presented to the IP stack.
- o The interfaces may be virtual or physical.
- o Configuration objects come from one or more administrative domains.
- o The IP addresses may be from the same or different address families, such as IPv4 and IPv6.
- o Communications using these IP addresses may happen simultaneously and independently.
- o Some communications using these IP addresses are possible on all the provisioning domains, while some are only possible on a smaller set of the provisioning domains.
- o While the MIF node may forward packets between its interfaces, the forwarding of packets is not taken into account in this definition and is out of scope for this document.

### 3.3. Host Requirements

"Requirements for Internet Hosts -- Communication Layers" [RFC1122] describes the multihomed node as if it has multiple IP addresses, which may be associated with one or more physical interfaces connected to the same or different networks.

Section 3.3.1.3 of [RFC1122] states that the node maintains a route cache table where each entry contains the local IP address, the destination IP address, Type(s) of Service (superseded by the Differentiated Services Code Point [RFC2474]), and the next-hop gateway IP address. The route cache entry would have data about the properties of the path, such as the average round-trip delay measured by a transport protocol. Nowadays, implementations are not caching this information.

[RFC1122] defines two host models:

- o The "strong" host model defines a multihomed host as a set of logical hosts within the same physical host. In this model, a packet must be sent on an interface that corresponds to the source address of that packet.
- o The "weak" host model describes a host that has some embedded gateway functionality. In the weak host model, the host can send and receive packets on any interface.

The multihomed node computes routes for outgoing datagrams differently, depending on the model. Under the strong model, the route is computed based on the source IP address, the destination IP address, and the Differentiated Services Code Point. Under the weak model, the source IP address is not used; only the destination IP address and the Differentiated Services Code Point are used.

### 3.4. Mobility and Other IP Protocols

The scope of this document is only about nodes implementing [RFC1122] for IPv4 and [RFC4294] for IPv6 without additional features or special-purpose support for transport layers, mobility, multihoming, or identifier-locator split mechanisms. Dealing with multiple interfaces with such mechanisms is related but considered as a separate problem and is under active study elsewhere in the IETF [RFC4960] [RFC5206] [RFC5533] [RFC5648] [RFC6182].

When an application is using one interface while another interface with better characteristics becomes available, the ongoing application session could be transferred to the newly enabled interface. However, in some cases, the ongoing session shall be kept on the current interface while initiating the new session on the new interface. The problem of interface selection is within the MIF scope and may leverage specific node functions (Section 3.8). However, if transfer of an IP session is required, IP mobility mechanisms, such as [RFC6275], shall be used.

### 3.5. Address Selection

"Default Address Selection for Internet Protocol version 6 (IPv6)" [RFC3484] defines algorithms for source and destination IP address selections. Default address selection as defined in [RFC3484] is mandatory to implement in IPv6 nodes, which also means dual-stack nodes. A node-scoped policy table managed by the IP stack is defined. Mechanisms to update the policy table are defined in [ADDR-SELECT-SOL].

Issues on using default address selection were found in [RFC5220] and [RFC5221] in the context of multiple prefixes on the same link.

### 3.6. Finding and Sharing IP Addresses with Peers

Interactive Connectivity Establishment (ICE) [RFC5245] is a technique for NAT traversal for UDP-based (and TCP-based) media streams established by the offer/answer model. The multiplicity of IP addresses, ports, and transport mechanisms in Session Description Protocol (SDP) offers are tested for connectivity by peer-to-peer connectivity checks. The result is candidate IP addresses and ports for establishing a connection with the other peer. However, ICE does not solve issues when incompatible configuration objects are received on different interfaces.

Some application protocols do referrals of IP addresses, port numbers, and transport for further exchanges. For instance, applications can provide reachability information to themselves or to a third party. The general problem of referrals is related to the multiple-interface problem, since, in this context, referrals must provide consistent information depending on which provisioning domain is used. Referrals are discussed in [REFERRAL-PS] and [SHIM6-APP-REFER].

### 3.7. Provisioning Domain Selection

In a MIF context, the node may simultaneously handle multiple domains with disparate characteristics, especially when supporting multiple access technologies. Selection is simple if the application is restricted to one specific provisioning domain: the application must start on the default provisioning domain if available; otherwise, the application does not start. However, if the application can be run on several provisioning domains, the selection problem can be difficult.

There is no standard method for selecting a provisioning domain, but some recommendations exist while restricting the scope to the interface selection problem. For example, [TS23.234] proposes a default mechanism for the interface selection. This method uses the following information (non-exhaustive list):

- o preferences provided by the user
- o policies provided by the network operator
- o quality of the radio link

- o network resource considerations (e.g., available Quality of Service (QoS), IP connectivity check)
- o the application QoS requirements in order to map applications to the best interface

However, [TS23.234] is designed for a specific multiple-interfaces use case. A generic way to handle these characteristics is yet to be defined.

### 3.8. Session Management

Some implementations, especially in the mobile world, rely on a higher-level session manager, also called a connection manager, to deal with issues brought by simultaneous attachment to multiple provisioning domains. Typically, the session manager may deal with the selection of the interface, and/or the provisioning domain, on behalf of the applications, or tackle complex issues such as how to resolve conflicting policies (Section 4.3). As discussed in Section 3.7, the session manager may encounter difficulties because of multiple and diverse criteria.

Session managers usually leverage the link-layer interface to gather information (e.g., lower-layer authentication and encryption methods; see Section 3.1) and/or for control purposes. Such a link-layer interface may not provide all required services to make a proper decision (e.g., interface selection). Some Oses or terminals already implement session managers [RFC6419], and vendor-specific platforms sometimes provide a specific sockets API (Section 3.9) that a session manager can use. However, the generic architecture of a session manager and its associated API are not currently standardized, so session manager behavior may differ between Oses and platforms.

Management of multiple interfaces sometimes relies on a virtual interface. For instance, a virtual interface allows support of multihoming, inter-technology handovers, and IP flow mobility in a Proxy Mobile IPv6 network [LOGICAL-IF-SUPPORT]. This virtual interface allows a multiple-interface node sharing a set of IP addresses on multiple physical interfaces and can also add benefits to multi-access scenarios such as Third Generation Partnership Project (3GPP) Multi Access Packet Data Network (PDN) Connectivity [TS23.402]. In most cases, the virtual interface will map several physical network interfaces, and the session manager should control the configuration of each one of these virtual and physical interfaces, as well as the mapping between the virtual and sub-interfaces.



In a situation involving multiple interfaces, active application sessions should survive path failures. Here, the session manager may come into play but only relying on existing mechanisms to manage multipath TCP (MPTCP) [RFC6182] or failover (Mobile IPv6 (MIPv6) [RFC6275], Shim6 [RFC5533]). A description of the interaction between these mechanisms and the session manager is out of scope of this document.

### 3.9. Sockets API

An Application Programming Interface (API) may expose objects that user applications or session managers use for dealing with multiple interfaces. For example, [RFC3542] defines how an application using the advanced sockets API specifies the interface or the source IP address through a simple `bind()` operation or with the `IPV6_PKTINFO` socket option.

Other APIs have been defined to solve issues similar to MIF. For instance, [RFC5014] defines an API to influence the default address selection mechanism by specifying attributes of the source addresses it prefers. [RFC6316] gives another example, in a multihoming context, by defining a sockets API enabling interactions between applications and the multihoming shim layer for advanced locator management, and access to information about failure detection and path exploration.

## 4. MIF Issues

This section describes the various issues when using a MIF node that has already received configuration objects from its various provisioning domains, or when multiple interfaces are used and result in wrong domain selection, addressing, or naming space overlaps. They occur, for example, when:

1. one interface is on the Internet and one is on a corporate private network. The latter may be through VPN.
2. one interface is on one access network (i.e., WiFi) and the other one is on another access network (3G) with specific services.

### 4.1. DNS Resolution Issues

A MIF node (M1) has an active interface (I1) connected to a network (N1), which has its DNS servers (S1 as primary DNS server) and another active interface (I2) connected to a network (N2), which has its DNS servers (S2 as primary DNS server). S1 serves some private

namespace, "private.example.com". The user or the application uses a name "a.private.example.com", which is within the private namespace of S1 and only resolvable by S1. Any of the following situations may occur:

1. The M1 stack, based on its routing table, uses I2 to reach S1 to resolve "a.private.example.com". M1 never reaches S1. The name is not resolved.
2. M1 keeps only one set of DNS server addresses from the received configuration objects. Let us assume that M1 keeps S2's address as the primary DNS server. M1 sends the forward DNS query for a.private.example.com to S2. S2 responds with an error for a nonexistent domain (NXDOMAIN). The name is not resolved. This issue also arises when performing a reverse DNS lookup. In the same situation, the reverse DNS query fails.
3. M1 keeps only one set of DNS server addresses from the received configuration objects. Let us assume that M1 keeps S2's address. M1 sends the DNS query for a.private.example.com to S2. S2 queries its upstream DNS and gets an IP address for a.private.example.com. However, the IP address is not the same one that S1 would have given. Therefore, the application tries to connect to the wrong destination node, or to the wrong interface, which may imply security issues or result in lack of service.
4. S1 or S2 has been used to resolve "a.private.example.com" to an [RFC1918] address. Both N1 and N2 are [RFC1918]-addressed networks. If addresses overlap, traffic may be sent using the wrong interface. This issue is not related to receiving multiple configuration objects, but to an address overlap between interfaces or attaching networks.
5. M1 has resolved a Fully Qualified Domain Name (FQDN) to a locally valid IP address when connected to N1. If the node loses connection to N1, the node may try to connect, via N2, to the same IP address as earlier, but as the address was only locally valid, connection setup fails. Similarly, M1 may have received NXDOMAIN for an FQDN when connected to N1. After detachment from N1, the node should not assume the FQDN continues to be nonexistent on N2.

6. M1 requests a AAAA record from a DNS server on a network that uses protocol translators and DNS64 [RFC6147]. If M1 receives a synthesized AAAA record, it is guaranteed to be valid only on the network from which it was learned. If M1 uses synthesized AAAA on any other network interface, traffic may be lost, dropped, or forwarded to the wrong network.

Some networks require the user to authenticate on a captive web portal before providing Internet connectivity. If this redirection is achieved by modifying the DNS reply, specific issues may occur. Consider a MIF node (M1) with an active interface (I1) connected to a network (N1), which has its DNS server (S1), and another active interface (I2) connected to a network (N2), which has its DNS server (S2). Until the user has not authenticated, S1 is configured to respond to any A or AAAA record query with the IP address of a captive portal, so as to redirect web browsers to an access control portal web page. This captive portal can be reached only via I1. When the user has authenticated to the captive portal, M1 can resolve an FQDN when connected to N1. However, if the address is only locally valid on N1, any of the issues described above may occur. When the user has not authenticated, any of the following situations may occur:

1. M1 keeps only one set of DNS server addresses from the received configuration objects and kept S2 address. M1 sends the forward DNS query for a.example.com to S2. S2 responds with the correct answer, R1. M1 attempts to contact R1 by way of I1. The connection fails. Or, the connection succeeds, bypassing the security policy on N1, possibly exposing the owner of M1 to prosecution.
2. M1 keeps only one set of DNS server addresses from the received configuration objects and kept S1 address. M1 sends the DNS query for a.example.com to S1. S1 provides the address of its captive portal. M1 attempts to contact this IP address using I1. The application fails to connect, resulting in lack of service. Or, the application succeeds in connecting but connects to the captive portal rather than the intended destination, resulting in lack of service (i.e., an IP connectivity check issue, as described in Section 4.4).

#### 4.2. Node Routing

Consider a MIF node (M1) with an active interface (I1) connected to a network (N1) and another active interface (I2) connected to a network (N2). The user or the application is trying to reach an IP address (IP1). Any of the following situations may occur:

1. For IP1, M1 has one default route (R1) via network (N1). To reach IP1, the M1 stack uses R1 and sends through I1. If IP1 is only reachable by N2, IP1 is never reached or is not the right target.
2. For the IP1 address family, M1 has one default route (R1, R2) per network (N1, N2). IP1 is reachable by both networks, but the N2 path has better characteristics, such as better round-trip time, least cost, better bandwidth, etc. These preferences could be defined by the user, provisioned by the network operator, or otherwise appropriately configured. The M1 stack uses R1 and tries to send through I1. IP1 is reached, but the service would be better via I2.
3. For the IP1 address family, M1 has a default route (R1), a specific X.0.0.0/8 route R1B (for example, but not restricted to an [RFC1918] prefix) to N1, and a default route (R2) to N2. IP1 is reachable by N2 only, but the prefix (X.0.0.0/8) is used in both networks. Because of the most specific route R1B, the M1 stack sends packets through I2, and those packets never reach the target.

A MIF node may have multiple routes to a destination. However, by default, it does not have any hint concerning which interface would be the best to use for that destination. The first-hop selection may leverage on local routing policy, allowing some actors (e.g., network operator or service provider) to influence the routing table, i.e., make a decision regarding which interface to use. For instance, a user on such a multihomed node might want a local policy to influence which interface will be used based on various conditions. Some Standards Development Organizations (SDOs) have defined policy-based routing selection mechanisms. For instance, the Access Network Discovery and Selection Function (ANDSF) [TS23.402] provides inter-system routing policies to terminals with both a 3GPP interface and non-3GPP interfaces. However, the routing selection may still be difficult, due to disjoint criteria as discussed in Section 3.8. Moreover, information required to make the right decision may not be available. For instance, interfaces to a lower layer may not provide all required hints concerning the selection (e.g., information on interface quality).

A node usually has a node-scoped routing table. However, a MIF node is connected to multiple provisioning domains; if each of these domains pushes routing policies to the node, then conflicts between policies may happen, and the node has no easy way to merge or reconcile them.

On a MIF node, some source addresses are not valid if used on some interfaces. For example, an [RFC1918] source address might be appropriate on the VPN interface but not on the public interface of the MIF node. If the source address is not chosen appropriately, then packets may be filtered in the path if source address filtering is in place ([RFC2827], [RFC3704]), and reply packets may never come back to the source.

#### 4.3. Conflicting Policies

The distribution of configuration policies (e.g., address selection, routing, DNS selection) to end nodes is being discussed (e.g., ANDSF in [TS23.402], [DHCPv6-ROUTE-OPTIONS]). If implemented in multiple provisioning domains, such mechanisms may conflict and create issues for the multihomed node. Considering a MIF node (M1) with an active interface (I1) connected to a network (N1) and another active interface (I2) connected to a network (N2), the following conflicts may occur:

1. M1 receives from both networks (N1 and N2) an update of its default address selection policy. However, the policies are specific to each network. The policies are merged by the M1 stack. Based on the merged policy, the chosen source address is from N1, but packets are sent to N2. The source address is not reachable from N2; therefore, the return packet is lost. Merging address selection policies may have important impacts on routing.
2. A node usually has a node-scoped routing table. However, each of the connected provisioning domains (N1 and N2) may push routing policies to the node; conflicts between policies may then happen, and the node has no easy way to merge or reconcile them.
3. M1 receives from one of the networks an update of its access selection policy, e.g., via the 3GPP/ANDSF [TS23.402]. However, the policy is in conflict with the local policy (e.g., user-defined or default OS policy). Assuming that the network provides a list of overloaded access networks, if the policy sent by the network is ignored, the packet may be sent to an access network with poor quality of communication.

#### 4.4. Session Management

Consider that a node has selected an interface and managed to configure it (i.e., the node obtained a valid IP address from the network). However, Internet connectivity is not available. The problem could be due to the following reasons:

1. The network requires a web-based authentication (e.g., the access network is a WiFi hot spot). In this case, the user can only access a captive portal. For instance, the network may perform HTTP redirection or modify DNS behavior (Section 4.1) until the user has not authenticated.
2. The IP interface is configured as active, but Layer 2 is so poor (e.g., poor radio condition) that no Layer 3 traffic can succeed.

In this situation, the session manager should be able to perform IP connectivity checks before selecting an interface.

Session issues may also arise when the node discovers a new provisioning domain. Consider a MIF node (M1) with an active interface (I1) connected to a network (N1) where an application is running a TCP session. A new network (N2) becomes available. If N2 is selected (e.g., because of better quality of communication), M1 gets IP connectivity to N2 and updates the routing table priority. So, if no specific route to the correspondent node is in place, and if the node implements the weak host model [RFC1122], the TCP connection breaks as the next hop changes. In order to continue communicating with the correspondent node, M1 should try to reconnect to the server via N2. In some situations, it could be preferable to maintain current sessions on N1 while new sessions start on N2.

#### 4.5. Single Interface on Multiple Provisioning Domains

When a node using a single interface is connected to multiple networks, such as different default routers, similar issues to those described above will happen. Even with a single interface, a node may wish to connect to more than one provisioning domain: that node may use more than one IP source address and may have more than one default router. The node may want to access services that can only be reached using one of the provisioning domains. In this case, it needs to use the right outgoing source address and default gateway to reach that service. In this situation, that node may also need to use different DNS servers to get domain names in those different provisioning domains.

## 5. Underlying Problems and Causes

This section lists the underlying problems, and their causes, that lead to the issues discussed in the previous section. The problems can be divided into five categories: 1) configuration, 2) DNS resolution, 3) routing, 4) address selection, and 5) session management and APIs. They are shown below:

1. Configuration. In a MIF context, configuration information specific to a provisioning domain may be ignored because:
  - A. Configuration objects (e.g., DNS servers, NTP servers) are node-scoped. So, the IP stack is not able to maintain the mapping between configuration information and the corresponding provisioning domain.
  - B. The same configuration objects (e.g., DNS server addresses, NTP server addresses) received from multiple provisioning domains may be overwritten.
  - C. Host implementations usually do not keep separate network configurations (such as DNS server addresses) per provisioning domain.
2. DNS resolution
  - A. Some FQDNs can be resolvable only by sending queries to the right server (e.g., intranet services). However, a DNS query could be sent to the wrong interface because DNS server addresses may be node-scoped.
  - B. A DNS answer may be only valid on a specific provisioning domain, but applications may not be aware of that mapping because DNS answers may not be kept with the provisioning from which the answer comes.
3. Routing
  - A. In the MIF context, routing information could be specific to each interface. This could lead to routing issues because, in current node implementations, routing tables are node-scoped.
  - B. Current node implementations do not take into account the Differentiated Services Code Point or path characteristics in the routing table.

- C. Even if implementations take into account path characteristics, the node has no way to properly merge or reconcile the provisioning domain preferences.
- D. A node attached to multiple provisioning domains could be provided with incompatible selection policies. If the different actors (e.g., user and network operator) are allowed to provide their own policies, the node has no way to properly merge or reconcile multiple selection policies.
- E. The problem of first-hop selection could not be solved via configuration (Section 3.7), and may leverage on sophisticated and specific mechanisms (Section 3.8).

#### 4. Address selection

- A. Default address selection policies may be specific to their corresponding provisioning domain. However, a MIF node may not be able to manage address selection policies per provisioning domain, because default address selection policies are node-scoped.
- B. On a MIF node, some source addresses are not valid if used on some interfaces or even on some default routers on the same interface. In this situation, the source address should be taken into account in the routing table, but current node implementations do not support such a feature.
- C. Source address or address selection policies could be specified by applications. However, there are no advanced APIs that support such applications.

#### 5. Session management and APIs

- A. Some implementations, especially in the mobile world, have higher-level APIs and/or session managers (aka connection managers) to address MIF issues. These mechanisms are not standardized and do not necessarily behave the same way across different OSes and/or platforms in the presence of MIF problems. This lack of consistency is an issue for the user and operator, who could experience different session manager behaviors, depending on the terminal.



- B. Session managers usually leverage on an interface to the link layer to gather information (e.g., lower-layer authentication and encryption methods) and/or for control purposes. However, such a link-layer interface may not provide all required services (e.g., may not provide all information that would allow a proper interface selection).
- C. A MIF node can support different session managers, which may have contradictory ways of solving MIF issues. For instance, because of different selection algorithms, two different session managers could select different domains in the same context. Or, when dealing with different domain selection policies, one session manager may give precedence to user policy while another could favor mobile operator policy.
- D. When host routing is updated and if the weak host model is supported, ongoing TCP sessions may break if routes change for these sessions. When TCP sessions should be bound to the interface, the strong host model should be used.
- E. When provided by different actors (e.g., user, network, default OS), policies may conflict and, thus, need to be reconciled at the host level. Policy conflict resolution may impact other functions (e.g., naming, routing).
- F. Even if the node has managed to configure an interface, Internet connectivity could be unavailable. This could be due to an access control function coming into play above Layer 3, or because of poor Layer 2 conditions. An IP connectivity check should be performed before selecting an interface.

## 6. Security Considerations

The problems discussed in this document have security implications, such as when packets sent on the wrong interface might be leaking some confidential information. Configuration parameters from one provisioning domain could cause a denial of service on another provisioning domain (e.g., DNS issues). Moreover, the undetermined behavior of IP stacks in the multihomed context brings additional threats where an interface on a multihomed node might be used to conduct attacks targeted to the networks connected by the other interfaces. Corrupted provisioning domain selection policy may induce a node to make decisions causing certain traffic to be forwarded to the attacker.

Additional security concerns are raised by possible future mechanisms that provide additional information to the node so that it can make a more intelligent decision with regards to the issues discussed in this document. Such future mechanisms may themselves be vulnerable and may not be easy to protect in the general case.

## 7. Contributors

This document is a joint effort with the authors of the MIF requirements document [MIF-REQ]. This includes, in alphabetical order: Jacni Qin, Carl Williams, and Peng Yang.

## 8. Acknowledgements

The documents written prior to the existence of the MIF working group, and the discussions during the MIF Birds of a Feather (BOF) meeting and around the MIF charter scope on the mailing list, brought very good input to the problem statement. This document steals a lot of text from these discussions and initial documents (e.g., [MIF-REQ], [IP-MULTIPLE-CONN], [MIF-DNS-SERVER-SELECT]). Therefore, the authors would like to acknowledge the following people (in no specific order), from whom some text has been taken: Jari Arkko, Keith Moore, Sam Hartman, George Tsirtsis, Scott Brim, Ted Lemon, Bernie Volz, Giyeong Son, Gabriel Montenegro, Julien Laganier, Teemu Savolainen, Christian Vogt, Lars Eggert, Margaret Wasserman, Hui Deng, Ralph Droms, Ted Hardie, Christian Huitema, Remi Denis-Courmont, Alexandru Petrescu, Zhen Cao, Gaetan Feige, Telemaco Melia, and Juan-Carlos Zuniga. Apologies to any contributors who have inadvertently not been named.

## 9. Informative References

### [ADDR-SELECT-SOL]

Matsumoto, A., Fujisaki, T., and R. Hiromi, "Solution approaches for address-selection problems", Work in Progress, March 2010.

### [DHCPv6-ROUTE-OPTIONS]

Dec, W., Ed., Mrugalski, T., Sun, T., and B. Sarikaya, "DHCPv6 Route Options", Work in Progress, September 2011.

### [IP-MULTIPLE-CONN]

Hui, M. and H. Deng, "Problem Statement and Requirement of Simple IP Multi-homing of the Host", Work in Progress, March 2009.

- [LOGICAL-IF-SUPPORT]  
Melia, T., Ed., and S. Gundavelli, Ed., "Logical Interface Support for multi-mode IP Hosts", Work in Progress, October 2011.
- [MIF-DNS-SERVER-SELECT]  
Savolainen, T., Kato, J., and T. Lemon, "Improved DNS Server Selection for Multi-Interfaced Nodes", Work in Progress, October 2011.
- [MIF-REQ] Yang, P., Seite, P., Williams, C., and J. Qin, "Requirements on multiple Interface (MIF) of simple IP", Work in Progress, February 2009.
- [MIH] IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Part 21: Media Independent Handover Services", IEEE LAN/MAN Std. 802.21-2008, January 2009.
- [REFERRAL-PS]  
Carpenter, B., Jiang, S., and Z. Cao, "Problem Statement for Referral", Work in Progress, February 2011.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.
- [RFC1136] Hares, S. and D. Katz, "Administrative Domains and Routing Domains: A model for routing in the Internet", RFC 1136, December 1989.
- [RFC1661] Simpson, W., Ed., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC3542] Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "Advanced Sockets Application Program Interface (API) for IPv6", RFC 3542, May 2003.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC4294] Loughney, J., Ed., "IPv6 Node Requirements", RFC 4294, April 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [RFC5014] Nordmark, E., Chakrabarti, S., and J. Laganier, "IPv6 Socket API for Source Address Selection", RFC 5014, September 2007.
- [RFC5113] Arkko, J., Aboba, B., Korhonen, J., Ed., and F. Bari, "Network Discovery and Selection Problem", RFC 5113, January 2008.
- [RFC5206] Nikander, P., Henderson, T., Ed., Vogt, C., and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol", RFC 5206, April 2008.
- [RFC5220] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules", RFC 5220, July 2008.
- [RFC5221] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Requirements for Address Selection Mechanisms", RFC 5221, July 2008.

- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.
- [RFC5648] Wakikawa, R., Ed., Devarapalli, V., Tsirtsis, G., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", RFC 5648, October 2009.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, April 2011.
- [RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, March 2011.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.
- [RFC6316] Komu, M., Bagnulo, M., Slavov, K., and S. Sugimoto, Ed., "Sockets Application Program Interface (API) for Multihoming Shim", RFC 6316, July 2011.
- [RFC6419] Wasserman, M. and P. Seite, "Current Practices for Multiple-Interface Hosts", RFC 6419, November 2011.
- [SHIM6-APP-REFER]  
Nordmark, E., "Shim6 Application Referral Issues", Work in Progress, July 2005.
- [TS23.234]  
3GPP, "3GPP system to Wireless Local Area Network (WLAN) interworking", TS 23.234, December 2009.
- [TS23.402]  
3GPP, "Architecture enhancements for non-3GPP accesses", TS 23.402, December 2010.

## Authors' Addresses

Marc Blanchet  
Viagenie  
2875 boul. Laurier, suite D2-630  
Quebec, QC G1V 2M2  
Canada

EEmail: [Marc.Blanchet@viagenie.ca](mailto:Marc.Blanchet@viagenie.ca)  
URI: <http://viagenie.ca>

Pierrick Seite  
France Telecom - Orange  
4, rue du Clos Courtel, BP 91226  
Cesson-Sevigne 35512  
France

EEmail: [pierrick.seite@orange.com](mailto:pierrick.seite@orange.com)