| Network Working Group | M. Blanchet | |
|---|---|---|
| Internet-Draft | Viagenie | |
| Intended status: Informational | P. Seite | |
| Expires: April 28, 2011 | France Telecom - Orange | |
| | October 25, 2010 | |

**Multiple Interfaces and Provisioning Domains Problem Statement**
**draft-ietf-mif-problem-statement-09.txt**

**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, others are local to the interface. Issues such as selecting the wrong interface to send trafic happen when conflicting node-scoped configuration objects are received and inappropriately used. Moreover, other issues are the result of simulatenous 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 single interface nodes situation.

**Status of this Memo**

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.
Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.
Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."
This Internet-Draft will expire on April 28, 2011.

**Copyright Notice**

**Table of Contents**

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, a 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] (Droms, R., "Dynamic Host Configuration Protocol," March 1997.), DHCPv6 [RFC3315] (Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," July 2003.), PPP [RFC1661] (Simpson, W., "The Point-to-Point Protocol (PPP)," July 1994.) and IPv6 Router Advertisements [RFC4861] (Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)," September 2007.). 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 servers IP addresses, and address selection policies, herein named "node-scoped".

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

Other issues are the result of simulatenous 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, the scope of this work, describe related work, list issues and then summarize the underlying problems.

A companion document [I-D.ietf-mif-current-practices] (Wasserman, M. and P. Seite, "Current Practices for Multiple Interface Hosts," October 2010.) 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] (Hares, S. and D. Katz, "Administrative Domains and Routing Domains: A model for routing in the Internet," December 1989.).

Provisioning domain

A set of consistent configuration information (e.g. Default router, Network prefixes, DNS,...). One administrative domain may have multiple provisioning domains.

Reference to IP version

When a protocol keyword such as IP, PPP, 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.  Below IP Interaction

Some types of interfaces have link layer characteristics which may be used in determining how multiple provisioning domain issues will be dealt with. For instance, link layers may have authentication and encryption characteristics which could be used as criteria for interface selection. However, network discovery and selection on lower layers as defined by [RFC5113] (Arkko, J., Aboba, B., Korhonen, J., and F. Bari, "Network Discovery and Selection Problem," January 2008.) 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 [802.21] (IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Part 21: Media Independent Handover Services, IEEE LAN/MAN Std 802.21-2008, January 2009.," 2010.), is also out of scope.
Some mechanisms (e.g., based on a logical IP interface)
allow sharing a single IP address over multiple  interfaces to networks with disparate access techologies. 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 current 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:

* A [RFC1122] (Braden, R., "Requirements for Internet Hosts - Communication Layers," October 1989.) IPv4 and/or [RFC4294] (Loughney, J., "IPv6 Node Requirements," April 2006.) IPv6 compliant node

* A MIF node is configured with more than one IP addresses (excluding loopback and link-local)

* A MIF node can attach to more than one provisioning domains, as presented to the IP stack.

* The interfaces may be logical, virtual or physical.

* Configuration objects come from one or more administrative domains.

* The IP addresses may be from the same or from different address families, such as IPv4 and IPv6.

* Communications using these IP addresses may happen simultaneously and independently.

* 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.

* While the MIF node may forward packets between its interfaces, forwarding packets is not taken into account in this definition and is out of scope for this document.

---

## 3.3.  Hosts Requirements

The requirements for Internet Hosts [RFC1122] (Braden, R., "Requirements for Internet Hosts - Communication Layers," October 1989.) describe 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.
The requirements states that The node maintains a route cache table where each entry contains the local IP address, the destination IP address, Type-of-Service and 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 these informations. [RFC1122] (Braden, R., "Requirements for Internet Hosts - Communication Layers," October 1989.) defines two host models:

> *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.

> *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 Type-of-Service. Under the weak model, the source IP address is not used, but only the destination IP address and the Type-of-Service.

---

### 3.4. Mobility and other IP protocols

The scope of this document is only about nodes implementing [RFC1122] (Braden, R., "Requirements for Internet Hosts - Communication Layers," October 1989.) for IPv4 and [RFC4294] (Loughney, J., "IPv6 Node Requirements," April 2006.) for IPv6 without additional features or special-purpose support for transport layers, mobility, multi-homing, 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] (Stewart, R., "Stream Control Transmission Protocol," September 2007.), [RFC5206] (Nikander, P., Henderson, T., Vogt, C., and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol," April 2008.), [RFC5533] (Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6," June 2009.), [RFC5648] (Wakikawa, R., Devarapalli, V., Tsirtsis, G., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration," October 2009.), [I-D.ietf-mptcp-architecture] (Ford, A., Raiciu, C., Handley, M., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development," October 2010.).
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 sessions on the new interface. The problem of the interface selection is within the MIF scope and may leverage specific node functions (Section 3.8 (Connection manager)).

However, if transfer of IP session is required, IP mobility mechanisms, such as [RFC3775], shall be used.

---

### 3.5. Address Selection

The Default Address Selection specification [RFC3484] (Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)," February 2003.) defines algorithms for source and destination IP address selections. It is mandatory to be implemented 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 being defined [I-D.ietf-6man-addr-select-sol] (Matsumoto, A., Fujisaki, T., and R. Hiromi, "Solution approaches for address-selection problems," March 2010.) to update the policy table.
Issues on using the Default Address Selection were found in [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," July 2008.) and [RFC5221] (Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Requirements for Address Selection Mechanisms," July 2008.) in the context of multiple prefixes on the same link.

---

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

Interactive Connectivity Establishment (ICE (Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols," April 2010.) [RFC5245]) is a technique for NAT traversal for UDP-based (and TCP) media streams established by the offer/answer model. The multiplicity of IP addresses, ports and transport in 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 itself 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 [I-D.carpenter-behave-referral-object] (Carpenter, B., Boucadair, M., Halpern, J., Jiang, S., and K. Moore, "A Generic Referral Object for Internet Entities," October 2009.) and

[I-D.ietf-shim6-app-refer] (Nordmark, E., "Shim6 Application Referral Issues," July 2005.).

---

### 3.7.  Interface and Provisioning domain selection

In a MIF context, the node may handle simultaneously 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 is difficult. For example, the interface selection mechanism defined in [TS23.234] (3GPP, "3GPP system to Wireless Local Area Network (WLAN) interworking; TS 23.234," 2009.) uses the following information:

*preferences provided by the user,

*policies provided by network operator,

*quality of the radio link,

*network resource considerations (e.g. available QoS, IP connectivity check,...),

*the application QoS requirements in order to map applications to the best interface

*and so on...

However, [TS23.234] (3GPP, "3GPP system to Wireless Local Area Network (WLAN) interworking; TS 23.234," 2009.) is designed for a specific multiple-interfaces use-case. A generic way to handle these characteristics is yet to be defined.

---

### 3.8.  Connection manager

Some implementations, specially in the mobile world, rely on higher-level connection managers to deal with issues brought by multiple provisioning domains. Typically, the connection manager manages the interface and provisioning domain selection on behalf of applications. As discussed previously in Section 3.7 (Interface and Provisioning

domain selection), the connection manager has the same issues to making decisions in the general way in front of multiple and diverse criteria. Connection managers usually leverage the link-layer interface to gather information and/or for control purpose. Such link-layer interface may not provide all required services to make a proper decision (e.g. interface selection). Some connection managers use specific MIF socket API (Section 3.9 (Socket API)) available in some vendor-specific platforms. The connection manager generic architecture and requirements are not currently standardized. Known connection managers behaviors have been documented in [I-D.ietf-mif-current-practices] (Wasserman, M. and P. Seite, "Current Practices for Multiple Interface Hosts," October 2010.).

## 3.9. Socket API

An Application Programming Interface (API) may expose objects that user applications, or connection managers, use for dealing with multiple interfaces. For example, [RFC3542] (Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "Advanced Sockets Application Program Interface (API) for IPv6," May 2003.) 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 similar issues to MIF.[RFC5014] (Nordmark, E., Chakrabarti, S., and J. Laganier, "IPv6 Socket API for Source Address Selection," September 2007.) defines an API to influence the default address selection mechanism by specifying attributes of the source addresses it prefers. [I-D.ietf-shim6-multihome-shim-api] (Komu, M., Bagnulo, M., Slavov, K., and S. Sugimoto, "Socket Application Program Interface (API) for Multihoming Shim," August 2010.) gives another example in a multihoming context, by defining a socket 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 results 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 server (S1) and another active interface (I2)
connected to a network (N2) which has its DNS server (S2). S1 serves
with 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. 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 and kept S2 address. M1 sends the forward
   DNS query for a.private.example.com to S2. S2 responds with an
   error for an non-existent domain (NXDOMAIN). The name is not
   resolved. This issue also arises when performing 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 and kept S2 address. M1 sends the DNS
   query for a.private.example.com to S2. S2 asks its upstream DNS
   and gets an IP address for a.private.example.com. However, the
   IP address is not the same one S1 would have given. Therefore,
   the application tries to connect to the wrong destination node,
   or to the wrong interface of the latter, 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] (Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot,
   G., and E. Lear, "Address Allocation for Private Internets,"
   February 1996.) address. Both N1 and N2 are [RFC1918] (Rekhter,
   Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear,
   "Address Allocation for Private Internets," February 1996.)

addressed networks. 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 an FQDN to locally valid IP address when
   connected to N1. If the node looses 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 AAAA record from a DNS server on a network that
   uses protocol translators and DNS64 [I-D.ietf-behave-dns64]
   (Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum,
   "DNS64: DNS extensions for Network Address Translation from
   IPv6 Clients to IPv4 Servers," October 2010.). If the M1
   receives synthesized AAAA record, it is guaranteed to be valid
   only on the network it was learned from. If the M1 uses
   synthesized AAAA on any other network interface, traffic may be
   lost, dropped or forwarded to the wrong network.

---

### 4.2.  Node Routing

A MIF node (M1) has 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, 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 N2
   path has better characteristics, such as better round-trip
   time, least cost, better bandwidth, etc.... These preferences
   could be defined by user, provisioned by the network operator,
   or else. M1 stack uses R1 and tries to send through I1. IP1 is
   reached but the service would be better by 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

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, M1 stack
sends through I2 and 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
decision regarding which interface to use. For instance, a user on such
multihomed node might want a local policy to influence which interface
will be used based on various conditions. Some SDOs have defined
policy-based selection mechanisms. For instance, the Access Network
Discovery and Selection Function (ANDSF) [TS23.402] (3GPP,
"Architecture enhancements for non- 3GPP accesses; TS 23.402," 2010.)
provides selection policies to terminals with both a 3GPP and non-3GPP
interfaces. However, the selection may still be difficult, due to
disjoint criteria as discussed in Section 3.8 (Connection manager).
Moreover, information required to make the right decision may not be
available. For instance, interfaces to lower layer may not provide all
required hints to the selection (e.g. information on interface
quality).
A node usually has a node-scoped routing table. 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 reconciliate 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]
(Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating
Denial of Service Attacks which employ IP Source Address Spoofing,"
May 2000.),[RFC3704] (Baker, F. and P. Savola, "Ingress Filtering for
Multihomed Networks," March 2004.)) and reply packets may never come
back to the source.

---

### 4.3.  Address Selection Policy

The distribution of address selection policy to end nodes is being
discussedSection 3.5 (Address Selection). Depending on the final
solution, Such mechanism may bring issues in a multiple provisioning
domains context. 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). When the user or the application tries to reach an IP address (IP1), the following situations may occur:

> 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 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.

---

## 4.4.  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 as described above 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 domain. 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, which 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) connection management and API. They are shown as below:

1. Configuration. In a MIF context, configuration information specific to a provisioning domain may be ignored because:

   1. Configuration objects (e.g. DNS servers, NTP servers, ...) are node-scoped. So the IP stack is not able to maintain the mapping between information and corresponding provisioning domain.

2. Same configuration objects (e.g. DNS server addresses, NTP
   server addresses, ..) received from multiple provisioning
   domains may be overwritten.

3. Host implementations usually do not keep separate network
   configuration (such as DNS server addresses) per
   provisioning domain.

2. DNS resolution

1. Some FQDN can be resolvable only by sending queries to the
   right server (e.g. intranet services). However, DNS query
   could be sent to the wrong interface because DNS server
   addresses may be node-scoped.

2. 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 from.

3. Routing

1. In the MIF context, routing information could be specific
   to each interface. This could lead to routing issue
   because, in current node implementations, routing tables
   are node-scoped.

2. Current node implementations do not take into account the
   Type-of-Service or path characteristics in the routing
   table.

3. Even if implementations take into account path
   characteristics, the node has no way to properly merge or
   reconciliate the provisioning domain preferences.

4. a node attached to multiple provisioning domain 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 reconciliate multiple selection
   policies.

5. The problem of first hop selection could not be solved via
   configuration (Section 3.7 (Interface and Provisioning
   domain selection)), and may leverage on sophisticated and
   specific mechanisms (Section 3.8 (Connection manager)).

4. Address selection

   1. Default Address Selection policies may be specific to
      their corresponding provisioning domain. However, a MIF
      node may not be able to manage per-provisioning domain
      address selection policies because default Address
      Selection policy is node-scoped.

   2. 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.

   3. Source address or address selection policies could be
      specified by applications. However, there is no advanced
      APIs to allow applications realizing such operations.

5. Connection management and API

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

   2. Connection managers usually leverage on interface to link
      layer to gather information and/or for control purpose.
      However, such link layer interface may not provide all
      required services (e.g. may not provide all information
      allowing to make a proper interface selection).

   3. A MIF node can support different connection managers,
      which may have contradictory ways to solve the MIF issues.
      For instance, because of different selection algorithms,
      two different connection managers could select different
      domains in a same context. Or, when dealing with different
      domain selection policies, a connection manager may give
      precedence to user policy while another could favor mobile
      operator policy.

## 6.  Security Considerations

The problems discussed in this document have security implications, such as when the 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 bring 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.  IANA Considerations                                    [TOC]

This document has no actions for IANA.

---

## 8.  Authors                                                [TOC]

This document is a joint effort with authors of the MIF requirements draft [I-D.yang-mif-req] (Yang, P., Seite, P., Williams, C., and J. Qin, "Requirements on multiple Interface (MIF) of simple IP," March 2009.). The authors of this document, in alphabetical order, include: Marc Blanchet, Jacqni Qin, Pierrick Seite, Carl Williams and Peny Yang.

---

## 9.  Acknowledgements                                       [TOC]

The initial Internet-Drafts prior to the MIF working group and the discussions during the MIF BOF meeting and on the mailing list around the MIF charter scope on the mailing list brought very good input to the problem statement. This draft steals a lot of text from these discussions and initial drafts (e.g. [I-D.yang-mif-req] (Yang, P., Seite, P., Williams, C., and J. Qin, "Requirements on multiple Interface (MIF) of simple IP," March 2009.),

[I-D.hui-ip-multiple-connections-ps] (Hui, M. and H. Deng, "Problem Statement and Requirement of Simple IP Multi-homing of the Host," March 2009.), [I-D.savolainen-mif-dns-server-selection] (Savolainen, T. and J. Kato, "Improved DNS Server Selection for Multi-Homed Nodes," September 2010.)). Therefore, the editor would like to acknowledge the following people (in no specific order), from which some text has been taken from: 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, Rémi Denis-Courmont, Alexandru Petrescu, Zhen Cao. Sorry if some contributors have not been named.

## 10. Informative References

| [802.21] | IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Part 21: Media Independent Handover Services, IEEE LAN/MAN Std 802.21-2008, January 2009.," 2010. |
| --- | --- |
| [I-D.carpenter-behave-referral-object] | Carpenter, B., Boucadair, M., Halpern, J., Jiang, S., and K. Moore, "A Generic Referral Object for Internet Entities," draft-carpenter-behave-referral-object-01 (work in progress), October 2009 (TXT). |
| [I-D.hui-ip-multiple-connections-ps] | Hui, M. and H. Deng, "Problem Statement and Requirement of Simple IP Multi-homing of the Host," draft-hui-ip-multiple-connections-ps-02 (work in progress), March 2009 (TXT). |
| [I-D.ietf-6man-addr-select-sol] | Matsumoto, A., Fujisaki, T., and R. Hiromi, "Solution approaches for address-selection problems," draft-ietf-6man-addr-select-sol-03 (work in progress), March 2010 (TXT). |
| [I-D.ietf-behave-dns64] | Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers," draft-ietf-behave-dns64-11 (work in progress), October 2010 (TXT). |
| [I-D.ietf-mif-current-practices] | Wasserman, M. and P. Seite, "Current Practices for Multiple Interface Hosts," draft-ietf-mif-current-practices-05 (work in progress), October 2010 (TXT). |
| [I-D.ietf-mptcp-architecture] | Ford, A., Raiciu, C., Handley, M., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development," draft-ietf-mptcp-architecture-02 (work in progress), October 2010 (TXT). |

| [I-D.ietf-shim6-app-refer] | Nordmark, E., "Shim6 Application Referral Issues," draft-ietf-shim6-app-refer-00 (work in progress), July 2005 (TXT). |
|---|---|
| [I-D.ietf-shim6-multihome-shim-api] | Komu, M., Bagnulo, M., Slavov, K., and S. Sugimoto, "Socket Application Program Interface (API) for Multihoming Shim," draft-ietf-shim6-multihome-shim-api-14 (work in progress), August 2010 (TXT). |
| [I-D.savolainen-mif-dns-server-selection] | Savolainen, T. and J. Kato, "Improved DNS Server Selection for Multi-Homed Nodes," draft-savolainen-mif-dns-server-selection-04 (work in progress), September 2010 (TXT). |
| [I-D.yang-mif-req] | Yang, P., Seite, P., Williams, C., and J. Qin, "Requirements on multiple Interface (MIF) of simple IP," draft-yang-mif-req-00 (work in progress), March 2009 (TXT). |
| [RFC1122] | Braden, R., "Requirements for Internet Hosts - Communication Layers," STD 3, RFC 1122, October 1989 (TXT). |
| [RFC1136] | Hares, S. and D. Katz, "Administrative Domains and Routing Domains: A model for routing in the Internet," RFC 1136, December 1989 (TXT). |
| [RFC1661] | Simpson, W., "The Point-to-Point Protocol (PPP)," STD 51, RFC 1661, July 1994 (TXT). |
| [RFC1918] | Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets," BCP 5, RFC 1918, February 1996 (TXT). |
| [RFC2131] | Droms, R., "Dynamic Host Configuration Protocol," RFC 2131, March 1997 (TXT, HTML, XML). |
| [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 (TXT). |
| [RFC3315] | Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," RFC 3315, July 2003 (TXT). |
| [RFC3484] | Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)," RFC 3484, February 2003 (TXT). |
| [RFC3542] | Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "Advanced Sockets Application Program Interface (API) for IPv6," RFC 3542, May 2003 (TXT). |
| [RFC3704] | Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks," BCP 84, RFC 3704, March 2004 (TXT). |

| [RFC4294] | Loughney, J., "IPv6 Node Requirements," RFC 4294, April 2006 (TXT). |
|---|---|
| [RFC4477] | Chown, T., Venaas, S., and C. Strauf, "Dynamic Host Configuration Protocol (DHCP): IPv4 and IPv6 Dual-Stack Issues," RFC 4477, May 2006 (TXT). |
| [RFC4861] | Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)," RFC 4861, September 2007 (TXT). |
| [RFC4960] | Stewart, R., "Stream Control Transmission Protocol," RFC 4960, September 2007 (TXT). |
| [RFC5014] | Nordmark, E., Chakrabarti, S., and J. Laganier, "IPv6 Socket API for Source Address Selection," RFC 5014, September 2007 (TXT). |
| [RFC5113] | Arkko, J., Aboba, B., Korhonen, J., and F. Bari, "Network Discovery and Selection Problem," RFC 5113, January 2008 (TXT). |
| [RFC5206] | Nikander, P., Henderson, T., Vogt, C., and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol," RFC 5206, April 2008 (TXT). |
| [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 (TXT). |
| [RFC5221] | Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Requirements for Address Selection Mechanisms," RFC 5221, July 2008 (TXT). |
| [RFC5245] | Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/ Answer Protocols," RFC 5245, April 2010 (TXT). |
| [RFC5533] | Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6," RFC 5533, June 2009 (TXT). |
| [RFC5648] | Wakikawa, R., Devarapalli, V., Tsirtsis, G., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration," RFC 5648, October 2009 (TXT). |
| [TS23.234] | 3GPP, "3GPP system to Wireless Local Area Network (WLAN) interworking; TS 23.234," 2009. |
| [TS23.402] | 3GPP, "Architecture enhancements for non- 3GPP accesses; TS 23.402," 2010. |

## Authors' Addresses

| | |
|---|---|
| | Marc Blanchet |
| | Viagenie |
| | 2875 boul. Laurier, suite D2-630 |
| | Quebec, QC G1V 2M2 |
| | Canada |
| Email: | Marc.Blanchet@viagenie.ca |
| URI: | http://viagenie.ca |
| | |
| | Pierrick Seite |
| | France Telecom - Orange |
| | 4, rue du Clos Courtel, BP 91226 |
| | Cesson-Sevigne 35512 |
| | France |
| Email: | pierrick.seite@orange-ftgroup.com |