

Network Working Group
Internet Draft
Intended status: Informational

Expires: September 4, 2009

G. Montenegro
D. Thaler
Shyam Seshadri
Microsoft Corporation
March 4, 2009

Multiple Interfaces on Windows
draft-montenegro-mif-multihomeing-00

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 4, 2009.

Internet-Draft

Multiple Interfaces on Windows

March 2009

Copyright

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

Increasingly, hosts have more than one network interface active at any given point in time. Such multiplicity of interfaces leads to multiple and potentially conflicting (or overlapping) sets of configuration information and policies. How these are arbitrated and managed influence how the host resolves DNS queries, and-with respect to outgoing packets-how it selects a source address and an outgoing interface.

Table of Contents

1.	Introduction.....	2
1.1.	Requirements Language.....	3
1.2.	Terminology.....	3
2.	Default Router Selection.....	3
3.	Selecting the Proper Interface for Outgoing or Incoming Packets	4
4.	DNS Configuration.....	4
4.1.	Host-wide DNS configuration.....	4
4.2.	Interface-specific DNS configuration.....	5
5.	DNS Names and Dynamic DNS Behavior.....	5
6.	DNS Query Behavior.....	5
7.	Security Considerations.....	6
8.	IANA Considerations.....	6
9.	Acknowledgments.....	6
10.	References.....	6
10.1.	Normative References.....	6

1. Introduction

Increasingly, hosts have more than one network interface active at any given point in time. In such scenarios, it is important
Montenegro, et al. Expires August 4, 2009 [Page 2]

Internet-Draft

Multiple Interfaces on Windows

March 2009

to distinguish between host-wide and interface-specific networking configuration information and policies. In addition to these (hopefully) disjoint sets, such multiplicity of interfaces may lead to potentially conflicts or overlaps. The arbitration and management of these sets of policy or configuration influence how the host behaves: e.g., how it selects a source address and an outgoing interface with respect to outgoing packets, and how it resolves DNS queries.

This document explains how the Microsoft Windows family of operating systems handles multi-homing for hosts. Furthermore, it only discusses host behavior, not router behavior. It is offered to aid discussion with respect to the MIF BoF [[MIFPS](#)], and is not guaranteed to be correct nor complete. More detailed and authoritative information on Microsoft Corporation's protocol implementations can be found in other sources such as MSDN (<http://msdn.microsoft.com/en-us/default.aspx>), the MCPP (<http://www.microsoft.com/protocols/mcpp.msp>) or WSPP (<http://www.microsoft.com/protocols/wspp/wspp.msp>) Protocol Documentation series.

1.1. Requirements Language

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](#) [[RFC2119](#)].

1.2. Terminology

Preferred interface - This is the first interface bound to the TCP/IP stack during bootup, but this ordering can be changed programmatically or via user interface. The preferred interface has been used in previous versions of Windows for routing and is also used for DNS query purposes.

[2.](#) Default Router Selection

It is possible to configure default routes for different interfaces, leading to more than one default router. This is not desirable, as it may lead to inconsistent results, given that interfaces may connect to disjoint networks. If there are multiple such default routes, the usual routing precedence decides which one is used. First, the interface with the lowest metric is preferred (generally, the faster Montenegro, et al. Expires August 4, 2009 [Page 3]

Internet-Draft

Multiple Interfaces on Windows

March 2009

an interface, the lower its metric). If multiple interfaces share the same metric (e.g., because they have equal or very similar nominal speeds), then behavior changed starting with Windows Vista. Previously, the default route on the preferred interface is used.

With Windows Vista, such preferred interface is not used for routing. Instead, host-to-router load sharing [[RFC4311](#)] is used for both IPv4 and IPv6.

For more information, refer to [[DefGateway](#)].

[3.](#) Selecting the Proper Interface for Outgoing or Incoming Packets

For outgoing packets whose source address has not been determined previously by applications, the stack chooses from addresses assigned to its interfaces. Similarly, for incoming packets, valid destination addresses must match the address of one of its interfaces. Per host requirements [[RFC1122](#)], such choices depend on whether a host implements the strong or weak host model.

Windows has implemented the weak host model on IPv4 as follows: Windows 2000, Windows XP and Windows Server 2003. IPv6 has always implemented the strong host model (starting with Windows XP and Windows Server 2003). Starting with Windows Vista and Windows Server 2008, the strong host model became the default for IPv4 as well, although the weak model is still possible via per-interface configuration.

For IPv6, Windows follows [[RFC3484](#)]. Starting in Vista, Windows follows 3484 for IPv4 as well. Prior to Vista, IPv4 would choose the first address on the outgoing interface for all traffic where the application didn't specify a source address.

For more information, see [[StrongAndWeakHost](#)] and [[SrcDstSel](#)].

[4.](#) DNS Configuration

A Windows host has host-wide as well as interface-specific configuration items.

[4.1.](#) Host-wide DNS configuration

Host-wide DNS configuration is input either via static configuration or via Microsoft's Group Policy. The latter, however, is available only within deployments that make use of Active Directory (e.g., within corporate or enterprise environments). Upon joining an Active Directory domain, clients may receive such configuration.

The host-wide DNS configuration comprises the following:

Montenegro, et al. Expires August 4, 2009 [Page 4]

Internet-Draft

Multiple Interfaces on Windows

March 2009

- . Primary DNS suffix - This is typically set to the domain name of the Active Directory in environments in which the host joins a domain. Otherwise, this can be set via static configuration. This can be set to something other than the Active Directory domain, and is sometimes unavoidable (e.g., upon corporate mergers), but is not recommended.
- . Host-wide suffix list - Host-wide list of suffixes that can be appended to names being queried (see "DNS Query Behavior" below). Prior to Windows Vista, this would be appended to any unqualified multi-label query (e.g., to names not ending in a "."). As of Vista, this is by default only appended to single-label queries (although the previous behavior is still configurable).

Before Windows Vista and Windows Server 2008 there used to be a global DNS servers configuration which took precedence over the per-interface DNS servers (see below). This global value is now deprecated.

[4.2.](#) Interface-specific DNS configuration

Interface-specific DNS configuration is input either via static configuration or via DHCP.

- . Interface-specific suffix list
- . List of DNS server IP addresses - The first one is the "primary" with all others being secondary.

[5.](#) DNS Names and Dynamic DNS Behavior

TBD

[6.](#) DNS Query Behavior

The DNS Client service queries the DNS servers in the following order:

1. The DNS Client service sends the name query to the first DNS server on the preferred interface's list of DNS servers and waits one second for a response.
2. If the DNS Client service does not receive a response from the first DNS server within one second, it sends the name query to the

Montenegro, et al.

Expires August 4, 2009

[Page 5]

Internet-Draft

Multiple Interfaces on Windows

March 2009

first DNS servers on all interfaces that are still under consideration and waits two seconds for a response.

3. If the DNS Client service does not receive a response from any DNS server within two seconds, the DNS Client service sends the query to all DNS servers on all interfaces that are still under consideration and waits another two seconds for a response.
4. If the DNS Client service still does not receive a response from any DNS server, it sends the name query to all DNS servers on all interfaces that are still under consideration and waits four seconds for a response.
5. If it the DNS Client service does not receive a response from any DNS server, the DNS client sends the query to all DNS servers on all interfaces that are still under consideration and waits eight seconds for a response.

More information is available at sources such as [[DNSWorks](#)] (See, for example, "Overview of DNS Query Process" or [[DNSClient](#)]).

7. Security Considerations

TBD.

8. IANA Considerations

This document has no IANA considerations.

9. Acknowledgments

The authors would like to acknowledge the following people for their feedback and interesting discussions.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

Montenegro, et al.

Expires August 4, 2009

[Page 6]

Internet-Draft

Multiple Interfaces on Windows

March 2009

10.2. Informative References

[DefGateway] Default Gateway Behavior for Windows TCP/IP:

<http://technet.microsoft.com/en-us/library/bb878104.aspx>

[DNSClient] Domain Name System Client Behavior in Windows Vista:

<http://technet.microsoft.com/en-us/library/bb727035.aspx>

[DNSWorks] How DNS Works: <http://technet.microsoft.com/en-us/library/cc772774.aspx>

[MIFPS] Blanchet, M., "Multiple Interfaces Problem Statement", [draft-blanchet-mif-problem-statement-00.txt](#) (work in progress), December 2008.

[RFC1122] Braden, R., "Requirements for Internet Hosts-

communication Layers", STD 3, [RFC 1122](#), October 1989.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC 3484](#), February 2003.

[RFC4191] R. Draves, D. Thaler, "Default Router Preferences and More-Specific Routes", [RFC 4191](#), November 2005.

[RFC4311] R. Hinden, D. Thaler, "IPv6 Host-to-Router Load Sharing", [RFC 4311](#), November 2005.

[SrcDstSel] Source and Destination Address Selection for IPv6:
<http://microsoft.com/technet/community/columns/cableguy/cg0206.aspx>

[StrongAndWeakHost] Strong and Weak Host Models
<http://technet.microsoft.com/en-us/magazine/2007.09.cableguy.aspx>

Montenegro, et al.

Expires August 4, 2009

[Page 7]

Internet-Draft

Multiple Interfaces on Windows

March 2009

Author's Addresses

Gabriel Montenegro
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA

Phone:

Email: gabriel.montenegro@microsoft.com

Dave Thaler

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA

Phone:
Email: dthaler@microsoft.com

Shyam Seshadri
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA

Phone:
Email: sseshad@microsoft.com