

## Terminology for Describing Internet Connectivity

### Status of This Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2005).

### Abstract

As the Internet has evolved, many types of arrangements have been advertised and sold as "Internet connectivity". Because these may differ significantly in the capabilities they offer, the range of options, and the lack of any standard terminology, the effort to distinguish between these services has caused considerable consumer confusion. This document provides a list of terms and definitions that may be helpful to providers, consumers, and, potentially, regulators in clarifying the type and character of services being offered.

### Table of Contents

1.	Introduction . . . . .	2
1.1.	The Problem and the Requirement . . . . .	2
1.2.	Adoption and a Non-pejorative Terminology . . . . .	2
2.	General Terminology . . . . .	3
3.	Filtering or Security Issues and Terminology . . . . .	5
4.	Additional Terminology . . . . .	7
5.	Security Considerations . . . . .	9
6.	Acknowledgements . . . . .	9
7.	Informative References . . . . .	9

## 1. Introduction

### 1.1. The Problem and the Requirement

Different ISPs and other providers offer a wide variety of products that are identified as "Internet" or "Internet access". These products offer different types of functionality and, as a result, some may be appropriate for certain users and uses and not others. For example, a service that offers only access to the Web (in this context, the portion of the Internet that is accessible via the HTTP and HTTPS protocols) may be appropriate for someone who is exclusively interested in browsing and in Web-based email services. It will not be appropriate for someone who needs to download files or use email more frequently. And it is likely to be even less appropriate for someone who needs to operate servers for other users, who needs virtual private network (VPN) capabilities or other secured access to a remote office, or who needs to synchronize mail for offline use.

Recent and rapidly evolving changes to the Internet's email environment have led to additional restrictions on sending and retrieving email. These restrictions, most of them developed as part of well intentioned attempts to prevent or fight unsolicited mail, may be imposed independently of the service types described below and are discussed separately in Section 3.

This document describes only the functions provided or permitted by the service provider. It does not and cannot specify the functions that pass through and are supported by various user-provided equipment.

The terms SHOULD, MUST, or MAY are capitalized in this document, as defined in [1].

### 1.2. Adoption and a Non-pejorative Terminology

The definitions proposed here are of little value if service providers and vendors are not willing to adopt them. The terms proposed are intended not to be pejorative, despite the belief of some members of the IETF community that some of these connectivity models are simply "broken" or "not really an Internet service". The mention of a particular service or model in this document does not imply any endorsement of it, only recognition of something that exists or might exist in the marketplace. Thus, the Best Current Practice described in this document is about terminology and information that should be supplied to the user and not about the types of service that should be offered.

## 2. General Terminology

This section lists the primary IP service terms. It is hoped that service providers will adopt these terms, to better define the services to potential users or customers. The terms refer to the intent of the provider (ISP), as expressed in either technical measures or terms and conditions of service. It may be possible to work around particular implementations of these characteristic connectivity types, but that freedom is generally not the intent of the provider and is unlikely to be supported if the workarounds stop working.

The service terms are listed in order of ascending capability, to reach "full Internet connectivity".

- o Web connectivity.

This service provides connectivity to the Web, i.e., to services supported through a "Web browser" (such as Firefox, Internet Explorer, Mozilla, Netscape, Lynx, or Opera), particularly those services using the HTTP or HTTPS protocols. Other services are generally not supported. In particular, there may be no access to POP3 or IMAP4 email, encrypted tunnels or other VPN mechanisms.

The addresses used may be private and/or not globally reachable. They are generally dynamic (see the discussion of dynamic addresses in Section 3 for further discussion of this terminology and its implications) and relatively short-lived (hours or days rather than months or years). These addresses are often announced as "dynamic" to those who keep lists of dial-up or dynamic addresses. The provider may impose a filtering Web proxy on the connections; that proxy may change and redirect URLs to other sites than the one originally specified by the user or embedded link.

- o Client connectivity only, without a public address.

This service provides access to the Internet without support for servers or most peer-to-peer functions. The IP address assigned to the customer is dynamic and is characteristically assigned from non-public address space. Servers and peer-to-peer functions are generally not supported by the network address translation (NAT) systems that are required by the use of private addresses. (The more precise categorization of types of NATs given in [2] are somewhat orthogonal to this document, but they may be provided as additional terms, as described in Section 4.)

Filtering Web proxies are common with this type of service, and the provider SHOULD indicate whether or not one is present.

- o Client only, public address.

This service provides access to the Internet without support for servers or most peer-to-peer functions. The IP address assigned to the customer is in the public address space. It is usually nominally dynamic or otherwise subject to change, but it may not change for months at a time. Most VPN and similar connections will work with this service. The provider may prohibit the use of server functions by either legal (contractual) restrictions or by filtering incoming connection attempts.

Filtering Web proxies are uncommon with this type of service, and the provider SHOULD indicate if one is present.

- o Firewallled Internet Connectivity.

This service provides access to the Internet and supports most servers and most peer-to-peer functions, with one or (usually) more static public addresses. It is similar to "Full Internet Connectivity", below, and all of the qualifications and restrictions described there apply. However, this service places a provider-managed "firewall" between the customer and the public Internet, typically at customer request and at extra cost compared to non-firewalled services. Typically by contractual arrangements with the customer, this may result in blocking of some services.

Other services may be intercepted by proxies, content-filtering arrangements, or application gateways. The provider SHOULD specify which services are blocked and which are intercepted or altered in other ways.

In most areas, this service arrangement is offered as an add-on, extra-cost, option with what would otherwise be Full Internet Connectivity. It is distinguished from the models above by the fact that any filtering or blocking services are ultimately performed at customer request, rather than being imposed as service restrictions.

- o Full Internet Connectivity.

This service provides the user full Internet connectivity, with one or more static public addresses. Dynamic addresses that are long-lived enough to make operating servers practical without highly dynamic DNS entries are possible, provided that they are not characterized as "dynamic" to third parties.

Filtering Web proxies, interception proxies, NAT, and other provider-imposed restrictions on inbound or outbound ports and traffic are incompatible with this type of service. Servers on a connected customer LAN are typically considered normal. The only compatible restrictions are bandwidth limitations and prohibitions against network abuse or illegal activities.

### 3. Filtering or Security Issues and Terminology

As mentioned in the Introduction, the effort to control or limit objectionable network traffic has led to additional restrictions on the behavior and capabilities of internet services. Such objectionable traffic may include unsolicited mail of various types (including "spam"), worms, viruses, and their impact, and in some cases, specific content.

In general, significant restrictions are most likely to be encountered with Web connectivity and non-public-address services, but some current recommendations would apply restrictions at all levels. Some of these mail restrictions may prevent sending outgoing mail (except through servers operated by the ISP for that purpose), may prevent use of return addresses of the user's choice, and may even prevent access to mail repositories (other than those supplied by the provider) by remote-access protocols such as POP3 or IMAP4. Because users may have legitimate reasons to access remote file services, remote mail submission servers (or, at least, to use their preferred email addresses from multiple locations), and to access remote mail repositories (again, a near-requirement if a single address is to be used), it is important that providers disclose the services they are making available and the filters and conditions they are imposing.

Several key issues in email filtering are of particular importance.

#### o Dynamic Addresses.

A number of systems, including several "blacklist" systems, are based on the assumption that most undesired email originates from systems with dynamic addresses, especially dialup and home broadband systems. Consequently, they attempt to prevent the addresses from being used to send mail, or perform some other services, except through provider systems designated for that purpose.

Different techniques are used to identify systems with dynamic addresses, including provider advertising of such addresses to blacklist operators, heuristics that utilize certain address ranges, and inspection of reverse-mapping domain names to see if

they contain telltale strings such as "dsl" or "dial". In some cases, the absence of a reverse-mapping DNS address is taken as an indication that the address is "dynamic". (Prohibition on connections based on the absence of a reverse-mapping DNS record was a technique developed for FTP servers many years ago; it was found to have fairly high rates of failure, both prohibiting legitimate connection attempts and failing to prevent illegitimate ones). Service providers SHOULD describe what they are doing in this area for both incoming and outgoing message traffic, and users should be aware that, if an address is advertised as "dynamic", it may be impossible to use it to send mail to an arbitrary system even if Full Internet Connectivity is otherwise provided.

- o Non-public addresses and NATs.

The NAT systems that are used to map between private and public address spaces may support connections to distant mail systems for outbound and inbound mail, but terms of service often prohibit the use of systems not supplied by the connectivity provider and prohibit the operation of "servers" (typically not precisely defined) on the client connection.

- o Outbound port filtering from the provider.

Another common technique involves blocking connections to servers outside the provider's control by blocking TCP "ports" that are commonly used for messaging functions. Different providers have different theories about this. Some prohibit their customers from accessing external SMTP servers for message submission, but they permit the use of the mail submission protocol ([3]) with sender authentication. Others try to block all outgoing messaging-related protocols, including remote mail retrieval protocols; however, this is less common with public-address services than those that are dependent on private addresses and NATs. If this type of filtering is present, especially with "Client only, public address" and "Full Internet Connectivity" services, the provider MUST indicate that fact (see also Section 4).

Still others may divert (reroute) outbound email traffic to their own servers, on the theory that this eliminates the need for reconfiguring portable machines as they connect from a different network location. Again, such diversion MUST be disclosed, especially since it can have significant security and privacy implications.

More generally, filters that block some or all mail being sent to (or submitted to) remote systems (other than via provider-supported servers), or that attempt to divert that traffic to their own servers, are, as discussed above, becoming common and SHOULD be disclosed.

#### 4. Additional Terminology

These additional terms, while not as basic to understanding a service offering as the ones identified above, are listed as additional information that a service provider might choose to provide to complement those general definitions. A potential customer might use those that are relevant to construct a list of specific questions to ask, for example.

- o Version support.

Does the service include IPv4 support only, both IPv4 and IPv6 support, or IPv6 support only?

- o Authentication support.

Which technical mechanism(s) are used by the service to establish and possibly authenticate connections? Examples might include unauthenticated DHCP, PPP, RADIUS, or HTTP interception.

- o VPNs and Tunnels.

Is IPSec blocked or permitted? Are other tunneling techniques at the IP layer or below, such as L2TP, permitted? Is there any attempt to block applications-layer tunnel mechanisms such as SSH?

- o Multicast support

Does the user machine have access to multicast packets and services?

- o DNS support.

Are users required to utilize DNS servers provided by the service provider, or are DNS queries permitted to reach arbitrary servers?

- o IP-related services.

Are ICMP messages to and from end user sites generally blocked or permitted? Are specific functions such as ping and traceroute blocked and, if so, at what point in the network?

- o Roaming support.

Does the service intentionally include support for IP roaming and, if so, how is this defined? For "broadband" connections, is some dialup arrangement provided for either backup or customer travel? If present, does that arrangement have full access to mailboxes, etc.

- o Applications services provided.

Are email services and/or Web hosting provided as part of the service, and on what basis? An email services listing should identify whether POP3, IMAP4, or Web access are provided and in what combinations, and what types of authentication and privacy services are supported or required for each.

- o Use and Blocking of Outbound Applications Services.

Does the service block use of SMTP or mail submission to other than its own servers or intercept such submissions and route them to its servers? Do its servers restrict the user to use of its domain names on outbound email? (For email specifically, also see Section 3 above.) Is the FTP PASV command supported or blocked? Are blocks or intercepts imposed on other file sharing or file transfer mechanisms, on conferencing applications, or on private applications services?

More generally, the provider should identify any actions of the service to block, restrict, or alter the destination of, the outbound use (i.e., the use of services not supplied by the provider or on the provider's network) of applications services.

- o Blocking of Inbound Applications Services.

In addition to issues raised by dynamic or private address space (when present), does the service take any other measures that specifically restrict the connections that can be made to equipment operated by the customer? Specifically, are inbound SMTP, HTTP or HTTPS, FTP, or various peer-to-peer or other connections (possibly including applications not specifically recognized by the provider) prohibited and, if so, which ones?

- o Application Content Filtering.

The service should declare whether it provides filtering or protection against worms or denial of service attacks against its customers, virus and spam filtering for its mail services (if



any), non-discretionary or "parental control" filtering of content, and so on.

- o Wiretapping and interception.

The service SHOULD indicate whether traffic passing through it is subject to lawful intercept, and whether the provider will make a proactive attempt to inform the user of such an intercept when such notice is legal. Analogous questions can be asked for traffic data that is stored for possible use by law enforcement.

## 5. Security Considerations

This document is about terminology, not protocols, so it does not raise any particular security issues. However, if the type of terminology that is proposed is widely adopted, it may become easier to identify security-related expectations of particular hosts, LANs, and types of connections.

## 6. Acknowledgements

This document was inspired by an email conversation with Vernon Schryver, Paul Vixie, and Nathaniel Bornstein. While there have been proposals to produce such definitions for many years, that conversation convinced the author that it was finally time to put a strawman on the table to see if the IETF could actually carry it forward. Harald Alvestrand, Brian Carpenter, George Michaelson, Vernon Schryver, and others made several suggestions on the initial draft that resulted in clarifications to the second one and Stephane Bortzmeyer, Brian Carpenter, Tony Finch, Susan Harris, David Kessens, Pekka Savola, and Vernon Schryver made very useful suggestions that were incorporated into subsequent versions. Susan Harris also gave the penultimate version an exceptionally careful reading, which is greatly appreciated, as are editorial suggestions by the RFC Editor.

## 7. Informative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.
- [3] Gellens, R. and J. Klensin, "Message Submission", RFC 2476, December 1998.

Author's Address

John C Klensin  
1770 Massachusetts Ave, #322  
Cambridge, MA 02140  
USA

Phone: +1 617 491 5735  
EMail: john-ietf@jck.com

## Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the IETF's procedures with respect to rights in IETF Documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.