Network Working Group
Internet-Draft
Expires: January 11, 2005

J. Klensin

July 13, 2004

Terminology for Describing Internet Connectivity
draft-klensin-ip-service-terms-04.txt

Status of this Memo

Copyright Notice

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

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, but that does not support any other type of Internet services, may be entirely appropriate for someone who is exclusively interested in browsing and in web-based email services, but not for someone who requires access to download files or make more intense use of email. And it is likely to be even less appropriate for someone who requires the ability 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 of various types, may be imposed independently of the service types described below and are discussed separately in Section 3.

Of course, the 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.

[[Note in Draft: This paragraph to be removed by the RFC Editor if the document progresses that far.]] This document is a first attempt at establishing some definitions for these various services. It is hoped that the definitions will evolve into ones that can be standardized and adopted widely enough to be useful to users and consumers.

1.2 Adoption and a Non-pejorative Terminology

The definitions proposed here are clearly of little value if service providers and vendors are not willing to adopt them. Consequently, the terms proposed are intended to not be pejorative, despite the belief of some members of the IETF community that some of these connectively 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.

1.3 Definitional Terminology

When the terms SHOULD, MUST, or MAY are used, and capitalized, in this document, they are used as defined in [1].

2. General Terminology

The terms listed in this section are the primary "IP Service Terms" and it is hoped that service providers will adopt them in describing offerings to potential users or customers.

Terms are listed below more or less in order of ascending (to "full Internet") capability. In each case, the terminology refers 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 those flexibilities are generally not the intent of the provider and are unlikely to be supported if the workarounds stop working.

Web connectivity. This service provides connectivity to the web only. Other services are generally not supported. In particular, there may be no access to POP3 or IMAP email, encrypted tunnels or other VPN mechanisms. The addresses used may be private and/or not globally reachable. They are generally dynamic 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 (see Section 3). 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.

Client connectivity only, without a public address. This service provides access to the Internet without support for server or most peer to peer functions. The IP address assigned to the customer is dynamic and, as a distinguishing feature of this class, is 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 might 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.

Client only, public address. This service provides access to the Internet without support for server or most peer to peer functions. The IP address assigned to the customer is in public address space. It is usually nominally dynamic or otherwise subject to change, but 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 of incoming connection attempts. Filtering web proxies are uncommon with this type of service, and the provider SHOULD indicate if one is present.

Firewalled Internet Connectivity. This service provides access to the Internet and supports most server and most peer to peer functions with one or more (usually more) static public addresses. It is similar in most respects to "Full Internet Connectivity", below, and all of the qualifications and restrictions on limitations described there apply. However, a managed "firewall" is in place between the customer and the public Internet. This may result in blocking of some services, and others may be intercepted by proxies, content-filtering arrangements, or applications gateways (although the latter three are less common). 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.

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 and 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 including unsolicited mail of various types (including "spam"); worms, viruses, and their impact; and in some cases, specific content has led to additional restrictions on the behavior and capabilities of internet services. In general, significant restrictions are more likely to be encountered with web

connectivity and non-public-address services, but some current recommendations would apply them 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 depositories (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 depositories (again, a near-requirement if a single address is to be used), it is important that providers disclose the services, filters, and conditions they are making available or imposing.

Several key issues in email filtering are of particular importance:

Dynamic Addresses. A number of systems, including several "blacklists", are based on the assumption that most undesired email originated 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 both of 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.

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 as well as prohibiting the operation of "servers" (typically not precisely defined) on the client connection.

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 permit the use of the mail submission protocol ([3]) with sender authentication. Others try to block all outgoing messaging-related protocols, including the use of remote mail retrieval protocols (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 users of portable machines to reconfigure them as they connect from different network locations. Again, this MUST be disclosed, especially since it can have signficant 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, as listed as additional information that a service provider might choose to provide to complement those general definitions. Or a potential customer might use those that are relevant by, for example, constructing a list of specific questions to ask.

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

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.

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?

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

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?

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.

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, IMAP, or web access are provided and in what combinations and what types of authentication and privacy services are supported or required for each.

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 FTP PASV 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.

Use and Blocking of Inbound Applications Services. In addition to any issues raised by dynamic or private address space (when present), does the service take any other measures to 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?

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 UCE filtering for its mail services (if any), non-discretionary or "parental control" filtering of content, and so on.

Wiretapping and interception. The service should indicate whether traffic passing through it is subject to lawful intercept with or without notice? Is traffic data stored for possible use by law enforcement with or without notice?

5. Security Considerations

This document is about terminology, not protocols, and 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 definitions like the ones above for many years, that conversation convinced the author that it was finally time to get 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, Pekka Savola, and Vernon Schryver made very useful suggestions that were incorporated into subsequent versions. Susan Harris also gave the penultimate version an exceptional careful reading, which is greatly appreciated.

7. Disclaimers and Lawyers

[[Note to the IESG and in Draft: several of the people who have contributed to, or commented on, this document have observed that, if it is considered successful, sections of it could well end up in national or local regulations, other types of consumer protection provisions, or contractual terms and conditions. Given that concern, the IESG is requested, to consult legal counsel as to whether the normal disclaimers, which were designed somewhat more for protocol specifications, are adequate to prevent creating (quoting from one contributor), "the smallest atom of liability for the author, the IETF, the RFC Editor, ISOC, or anyone else within 10000 km" from liability. This section should then be removed and, if needed, replaced by text here or elsewhere in the document as appropriate.]]

8 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

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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