

Benchmarking Terminology for Firewall Performance

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Table of Contents

1. Introduction.....	2
2. Existing definitions.....	2
3. Term definitions.....	3
3.1 Allowed traffic.....	3
3.2 Application proxy.....	3
3.3 Authentication.....	4
3.4 Bit forwarding rate.....	5
3.5 Circuit proxy.....	6
3.6 Concurrent connections.....	6
3.7 Connection.....	7
3.8 Connection establishment.....	9
3.9 Connection establishment time.....	9
3.10 Connection maintenance.....	10
3.11 Connection overhead.....	11
3.12 Connection teardown.....	11
3.13 Connection teardown time.....	12
3.14 Data source.....	12
3.15 Demilitarized zone.....	13
3.16 Firewall.....	13
3.17 Goodput.....	14
3.18 Homed.....	15
3.19 Illegal traffic.....	15
3.20 Logging.....	16
3.21 Network address translation.....	16
3.22 Packet filtering.....	17
3.23 Policy.....	17
3.24 Protected network.....	18
3.25 Proxy.....	19
3.26 Rejected traffic.....	19

3.27 Rule set.....	20
3.28 Security association.....	20
3.29 Stateful packet filtering.....	21
3.30 Tri-homed.....	22
3.31 Unit of transfer.....	22
3.32 Unprotected network.....	23
3.33 User.....	23
4. Security considerations.....	24
5. References.....	25
6. Acknowledgments.....	25
7. Contact Information.....	25
8. Full Copyright Statement.....	26

1. Introduction

This document defines terms used in measuring the performance of firewalls. It extends the terminology already used for benchmarking routers and switches with definitions specific to firewalls.

Forwarding rate and connection-oriented measurements are the primary metrics used in this document.

Why do we need firewall performance measurements? First, despite the rapid rise in firewall deployment, there is no standard method of performance measurement. Second, implementations vary widely, making it difficult to do direct performance comparisons. Finally, more and more organizations are deploying firewalls on internal networks operating at relatively high speeds, while most firewall implementations remain optimized for use over relatively low-speed wide-area connections. As a result, users are often unsure whether the products they buy will stand up to relatively heavy loads.

2. Existing definitions

This document uses the conceptual framework established in RFCs 1242 and 2544 (for routers) and RFC 2285 (for switches). The router and switch documents contain discussions of several terms relevant to benchmarking the performance of firewalls. Readers should consult the router and switch documents before making use of this document.

This document uses the definition format described in RFC 1242, Section 2. The sections in each definition are: definition, discussion, measurement units (optional), issues (optional), and cross-references.

3. Term definitions

3.1 Allowed traffic

Definition:

Packets forwarded as a result of the rule set of the device under test/system under test (DUT/SUT).

Discussion:

Firewalls typically are configured to forward only those packets explicitly permitted in the rule set. Forwarded packets must be included in calculating the bit forwarding rate or maximum bit forwarding rate of the DUT/SUT. All other packets must not be included in bit forwarding rate calculations.

This document assumes 1:1 correspondence of allowed traffic offered to the DUT/SUT and forwarded by the DUT/SUT. There are cases where the DUT/SUT may forward more traffic than it is offered; for example, the DUT/SUT may act as a mail exploder or a multicast server. Any attempt to benchmark forwarding rates of such traffic must include a description of how much traffic the tester expects to be forwarded.

Unit of measurement:

not applicable

Issues:

See also:

policy
rule set

3.2 Application proxy

Definition:

A proxy service that is set up and torn down in response to a client request, rather than existing on a static basis.

Discussion:

Circuit proxies always forward packets containing a given port number if that port number is permitted by the rule set. Application proxies, in contrast, forward packets only once a connection has been established using some known protocol. When the connection closes, a firewall using application proxies rejects individual packets, even if they contain port numbers allowed by a rule set.

Unit of measurement:
not applicable

Issues:
circuit proxy
rule sets

See also:
allowed traffic
circuit proxy
proxy
rejected traffic
rule set

3.3 Authentication

Definition:
The process of verifying that a user requesting a network resource is who he, she, or it claims to be, and vice versa.

Discussion:
Trust is a critical concept in network security. Any network resource (such as a file server or printer) typically requires authentication before granting access.

Authentication takes many forms, including but not limited to IP addresses; TCP or UDP port numbers; passwords; external token authentication cards; and biometric identification such as signature, speech, or retina recognition systems.

The entity being authenticated might be the client machine (for example, by proving that a given IP source address really is that address, and not a rogue machine spoofing that address) or a user (by proving that the user really is who he, she, or it claims to be). Servers might also authenticate themselves to clients.

Testers should be aware that in an increasingly mobile society, authentication based on machine-specific criteria such as an IP address or port number is not equivalent to verifying that a given individual is making an access request. At this writing systems that verify the identity of users are typically external to the firewall, and may introduce additional latency to the overall SUT.

Unit of measurement:
not applicable

Issues:

See also:
user

3.4 Bit forwarding rate

Definition:

The number of bits per second of allowed traffic a DUT/SUT can be observed to transmit to the correct destination interface(s) in response to a specified offered load.

Discussion:

This definition differs substantially from section 3.17 of RFC 1242 and section 3.6.1 of RFC 2285.

Unlike both RFCs 1242 and 2285, this definition introduces the notion of different classes of traffic: allowed, illegal, and rejected (see definitions for each term). For benchmarking purposes, it is assumed that bit forwarding rate measurements include only allowed traffic.

Unlike RFC 1242, there is no reference to lost or retransmitted data. Forwarding rate is assumed to be a goodput measurement, in that only data successfully forwarded to the destination interface is measured. Bit forwarding rate must be measured in relation to the offered load. Bit forwarding rate may be measured with differed load levels, traffic orientation, and traffic distribution.

Unlike RFC 2285, this measurement counts bits per second rather than frames per second. Testers interested in frame (or frame-like) measurements should use units of transfer.

Unit of measurement:
bits per second

Issues:
Allowed traffic vs. rejected traffic

See also:
allowed traffic
goodput
illegal traffic
rejected traffic
unit of transfer

3.5 Circuit proxy

Definition:

A proxy service that statically defines which traffic will be forwarded.

Discussion:

The key difference between application and circuit proxies is that the latter are static and thus will always set up a connection if the DUT/SUT's rule set allows it. For example, if a firewall's rule set permits ftp connections, a circuit proxy will always forward traffic on TCP port 20 (ftp-data) even if no control connection was first established on TCP port 21 (ftp-control).

Unit of measurement:

not applicable

Issues:

application proxy
rule sets

See also:

allowed traffic
application proxy
proxy
rejected traffic
rule set

3.6 Concurrent connections

Definition:

The aggregate number of simultaneous connections between hosts across the DUT/SUT, or between hosts and the DUT/SUT.

Discussion:

The number of concurrent connections a firewall can support is just as important a metric for some users as maximum bit forwarding rate.

While "connection" describes only a state and not necessarily the transfer of data, concurrency assumes that all existing connections are in fact capable of transferring data. If a data cannot be sent over a connection, that connection should not be counted toward the number of concurrent connections.

Further, this definition assumes that the ability (or lack thereof) to transfer data on a given connection is solely the responsibility of the DUT/SUT. For example, a TCP connection that a DUT/SUT has

left in a FIN_WAIT_2 state clearly should not be counted. But another connection that has temporarily stopped transferring data because some external device has restricted the flow of data is not necessarily defunct. The tester should take measures to isolate changes in connection state to those effected by the DUT/SUT.

Unit of measurement:

Concurrent connections

Maximum number of concurrent connections

Issues:

See also:

connections

connection establishment time

connection overhead

3.7 Connection

Definition:

A state in which two hosts, or a host and the DUT/SUT, agree to exchange data using a known protocol.

Discussion:

A connection is an abstraction describing an agreement between two nodes: One agrees to send data and the other agrees to receive it.

Connections might use TCP, but they don't have to. Other protocols such as ATM also might be used, either instead of or in addition to TCP connections.

What constitutes a connection depends on the application. For a native ATM application, connections and virtual circuits may be synonymous. For TCP/IP applications on ATM networks (where multiple TCP connections may ride over a single ATM virtual circuit), the number of TCP connections may be the most important consideration.

Additionally, in some cases firewalls may handle a mixture of native TCP and native ATM connections. In this situation, the wrappers around user data will differ. The most meaningful metric describes what an end-user will see.

Data connections describe state, not data transfer. The existence of a connection does not imply that data travels on that connection at any given time, although if data cannot be forwarded on a previously established connection that connection should not be considered in any aggregate connection count (see concurrent connections).

A firewall's architecture dictates where a connection terminates. In the case of application or circuit proxy firewalls, a connection terminates at the DUT/SUT. But firewalls using packet filtering or stateful packet filtering designs act only as passthrough devices, in that they reside between two connection endpoints. Regardless of firewall architecture, the number of data connections is still relevant, since all firewalls perform some form of connection maintenance; at the very least, all check connection requests against their rule sets.

Further, note that connection is not an atomic unit of measurement in that it does not describe the various steps involved in connection setup, maintenance, and teardown. Testers may wish to take separate measurements of each of these components.

When benchmarking firewall performance, it's important to identify the connection establishment and teardown procedures, as these must not be included when measuring steady-state forwarding rates. Further, forwarding rates must be measured only after any security associations have been established.

Though it seems paradoxical, connectionless protocols such as UDP may also involve connections, at least for the purposes of firewall performance measurement. For example, one host may send UDP packets to another across a firewall. If the destination host is listening on the correct UDP port, it receives the UDP packets. For the purposes of firewall performance measurement, this is considered a connection.

Unit of measurement:

- concurrent connections
- connection
- connection establishment time
- maximum number of concurrent connections
- connection teardown time

Issues:

- application proxy vs. stateful packet filtering
- TCP/IP vs. ATM

- connection-oriented vs. connectionless

See also:

- data source
- concurrent connections
- connection establishment

- connection establishment time
- connection teardown
- connection teardown time

3.8 Connection establishment

Definition:

The data exchanged between hosts, or between a host and the DUT/SUT, to initiate a connection.

Discussion:

Connection-oriented protocols like TCP have a proscribed handshaking procedure when launching a connection. When benchmarking firewall performance, it is import to identify this handshaking procedure so that it is not included in measurements of bit forwarding rate or UOTs per second.

Testers may also be interested in measurements of connection establishment time through or with a given DUT/SUT.

Unit of measurement:

not applicable

See also:

- connection
- connection establishment time
- connection maintenance
- connection teardown

Issues:

not applicable

3.9 Connection establishment time

Definition:

The length of time needed for two hosts, or a host and the DUT/SUT, to agree to set up a connection using a known protocol.

Discussion:

Each connection-oriented protocol has its own defined mechanisms for setting up a connection. For purposes of benchmarking firewall performance, this shall be the interval between receipt of the first bit of the first octet of the packet carrying a connection establishment request on a DUT/SUT interface until transmission of the last bit of the last octet of the last packet of the connection setup traffic headed in the opposite direction.

This definition applies only to connection-oriented protocols such as TCP. For connectionless protocols such as UDP, the notion of connection establishment time is not meaningful.

Unit of measurement:

Connection establishment time

Issues:

See also:

concurrent connections
connection
connection maintenance

3.10 Connection maintenance

Definition:

The data exchanged between hosts, or between a host and the DUT/SUT, to ensure a connection is kept alive.

Discussion:

Some implementations of TCP and other connection-oriented protocols use "keep-alive" data to maintain a connection during periods where no user data is exchanged.

When benchmarking firewall performance, it is useful to identify connection maintenance traffic as distinct from UOTs per second. Given that maintenance traffic may be characterized by short bursts at periodical intervals, it may not be possible to describe a steady-state forwarding rate for maintenance traffic. One possible approach is to identify the quantity of maintenance traffic, in bytes or bits, over a given interval, and divide through to derive a measurement of maintenance traffic forwarding rate.

Unit of measurement:

maintenance traffic
forwarding rate

See also:

connection
connection establishment time
connection teardown
connection teardown time

Issues:

not applicable

3.11 Connection overhead

Definition:

The degradation in bit forwarding rate, if any, observed as a result of the addition of one connection between two hosts through the DUT/SUT, or the addition of one connection from a host to the DUT/SUT.

Discussion:

The memory cost of connection establishment and maintenance is highly implementation-specific. This metric is intended to describe that cost in a method visible outside the firewall.

It may also be desirable to invert this metric to show the performance improvement as a result of tearing down one connection.

Unit of measurement:

bit forwarding rate

Issues:

3.12 Connection teardown

Definition:

The data exchanged between hosts, or between a host and the DUT/SUT, to close a connection.

Discussion:

Connection-oriented protocols like TCP follow a stated procedure when ending a connection. When benchmarking firewall performance, it is important to identify the teardown procedure so that it is not included in measurements of bit forwarding rate or UOTs per second.

Testers may also be interested in measurements of connection teardown time through or with a given DUT/SUT.

Unit of measurement:

not applicable

See also:

connection teardown time

Issues:

not applicable

3.13 Connection teardown time

Definition:

The length of time needed for two hosts, or a host and the DUT/SUT, to agree to tear down a connection using a known protocol.

Discussion:

Each connection-oriented protocol has its own defined mechanisms for dropping a connection. For purposes of benchmarking firewall performance, this shall be the interval between receipt of the first bit of the first octet of the packet carrying a connection teardown request on a DUT/SUT interface until transmission of the last bit of the last octet of the last packet of the connection teardown traffic headed in the opposite direction.

This definition applies only to connection-oriented protocols such as TCP. For connectionless protocols such as UDP, the notion of connection teardown time is not meaningful.

Unit of measurement:

Connection teardown time

Issues:

See also:

- concurrent connections
- connection
- connection maintenance

3.14 Data source

Definition:

A host capable of generating traffic to the DUT/SUT.

Discussion:

One data source may emulate multiple users or hosts. In addition, one data source may offer traffic to multiple network interfaces on the DUT/SUT.

The term "data source" is deliberately independent of any number of users. It is useful to think of data sources simply as traffic generators, without any correlation to any given number of users.

Unit of measurement:

not applicable

Issues:

- user

See also:
connection
user

3.15 Demilitarized zone

Definition:

A network segment or segments located between protected and unprotected networks.

Discussion:

As an extra security measure, networks may be designed such that protected and unprotected segments are never directly connected. Instead, firewalls (and possibly public resources such as HTTP or FTP servers) reside on a so-called DMZ network.

DMZ networks are sometimes called perimeter networks.

Unit of measurement:
not applicable

Issues:
Homed

See also:
protected network
unprotected network

3.16 Firewall

Definition:

A device or group of devices that enforces an access control policy between networks.

Discussion:

While there are many different ways to accomplish it, all firewalls do the same thing: control access between networks.

The most common configuration involves a firewall connecting two segments (one protected and one unprotected), but this is not the only possible configuration. Many firewalls support tri-homing, allowing use of a DMZ network. It is possible for a firewall to accommodate more than three interfaces, each attached to a different network segment.

The criteria by which access are controlled are not specified here. Typically this has been done using network- or transport-layer criteria (such as IP subnet or TCP port number), but there is no

reason this must always be so. A growing number of firewalls are controlling access at the application layer, using user identification as the criterion. And firewalls for ATM networks may control access based on data link-layer criteria.

Unit of measurement:
not applicable

Issues:

See also:
DMZ
tri-homed
user

3.17 Goodput

Definition:
The number of bits per unit of time forwarded to the correct destination interface of the DUT/SUT, minus any bits lost or retransmitted.

Discussion:
Firewalls are generally insensitive to packet loss in the network. As such, measurements of gross bit forwarding rates are not meaningful since (in the case of proxy-based and stateful packet filtering firewalls) a receiving endpoint directly attached to a DUT/SUT would not receive any data dropped by the DUT/SUT.

The type of traffic lost or retransmitted is protocol-dependent. TCP and ATM, for example, request different types of retransmissions. Testers must observe retransmitted data for the protocol in use, and subtract this quantity from measurements of gross bit forwarding rate.

Unit of measurement:
bits per second

Issues:
allowed vs. rejected traffic

See also:
allowed traffic
bit forwarding rate
rejected traffic

3.18 Homed

Definition:

The number of logical interfaces a DUT/SUT contains.

Discussion:

Firewalls typically contain at least two logical interfaces. In network topologies where a DMZ is used, the firewall usually contains at least three interfaces and is said to be tri-homed. Additional interfaces would make a firewall quad-homed, quint-homed, and so on.

It is theoretically possible for a firewall to contain one physical interface and multiple logical interfaces. This configuration is discouraged for testing purposes because of the difficulty in verifying that no leakage occurs between protected and unprotected segments.

Unit of measurement:

not applicable

Issues:

See also:

tri-homed

3.19 Illegal traffic

Definition:

Packets specified for rejection in the rule set of the DUT/SUT.

Discussion:

A buggy or misconfigured firewall might forward packets even though its rule set specifies that these packets be dropped. Illegal traffic differs from rejected traffic in that it describes all traffic specified for rejection by the rule set, while rejected traffic specifies only those packets actually dropped by the DUT/SUT.

Unit of measurement:

not applicable

Issues:

See also:

- accepted traffic
- policy
- rejected traffic
- rule set

3.20 Logging

Definition:

The recording of user requests made to the firewall.

Discussion:

Firewalls typically log all requests they handle, both allowed and rejected. For many firewall designs, logging requires a significant amount of processing overhead, especially when complex rule sets are in use.

The type and amount of data logged varies by implementation. Testers may find it desirable to log equivalent data when comparing different DUT/SUTs.

Some systems allow logging to take place on systems other than the DUT/SUT.

Unit of measurement:

not applicable

Issues:

rule sets

See also:

- allowed traffic
- connection
- rejected traffic

3.21 Network address translation

Definition:

A method of mapping one or more private, reserved IP addresses to one or more public IP addresses.

Discussion:

In the interest of conserving the IPv4 address space, RFC 1918 proposed the use of certain private (reserved) blocks of IP addresses. Connections to public networks are made by use of a device that translates one or more RFC 1918 addresses to one or more public addresses--a network address translator (NAT).

The use of private addressing also introduces a security benefit in that RFC 1918 addresses are not visible to hosts on the public Internet.

Some NAT implementations are computationally intensive, and may affect bit forwarding rate.

Unit of measurement:
not applicable

Issues:

See also:

3.22 Packet filtering

Definition:

The process of controlling access by examining packets based on the content of packet headers.

Discussion:

Packet-filtering devices forward or deny packets based on information in each packet's header, such as IP address or TCP port number. A packet-filtering firewall uses a rule set to determine which traffic should be forwarded and which should be blocked.

Unit of measurement:
not applicable

Issues:
static vs. stateful packet filtering

See also:
application proxy
circuit proxy
proxy
rule set
stateful packet filtering

3.23 Policy

Definition:

A document defining acceptable access to protected, DMZ, and unprotected networks.

Discussion:

Security policies generally do not spell out specific configurations for firewalls; rather, they set general guidelines for what is and is not acceptable network access.

The actual mechanism for controlling access is usually the rule set implemented in the DUT/SUT.

Unit of measurement:

not applicable

Issues:**See also:**

rule set

3.24 Protected network

Definition:

A network segment or segments to which access is controlled by the DUT/SUT.

Discussion:

Firewalls are intended to prevent unauthorized access either to or from the protected network. Depending on the configuration specified by the policy and rule set, the DUT/SUT may allow hosts on the protected segment to act as clients for servers on either the DMZ or the unprotected network, or both.

Protected networks are often called "internal networks." That term is not used here because firewalls increasingly are deployed within an organization, where all segments are by definition internal.

Unit of measurement:

not applicable

Issues:**See also:**

demilitarized zone (DMZ)
unprotected network
policy
rule set
unprotected network

3.25 Proxy

Definition:

A request for a connection made on behalf of a host.

Discussion:

Proxy-based firewalls do not allow direct connections between hosts. Instead, two connections are established: one between the client host and the DUT/SUT, and another between the DUT/SUT and server host.

As with packet-filtering firewalls, proxy-based devices use a rule set to determine which traffic should be forwarded and which should be rejected.

There are two types of proxies: application proxies and circuit proxies.

Unit of measurement:

not applicable

Issues:

application

See also:

application proxy
circuit proxy
packet filtering
stateful packet filtering

3.26 Rejected traffic

Definition:

Packets dropped as a result of the rule set of the DUT/SUT.

Discussion:

For purposes of benchmarking firewall performance, it is expected that firewalls will reject all traffic not explicitly permitted in the rule set. Dropped packets must not be included in calculating the bit forwarding rate or maximum bit forwarding rate of the DUT/SUT.

Unit of measurement:

not applicable

Issues:

See also:

- allowed traffic
- illegal traffic
- policy
- rule set

3.27 Rule set

Definition:

The collection of access control rules that determines which packets the DUT/SUT will forward and which it will reject.

Discussion:

Rule sets control access to and from the network interfaces of the

DUT/SUT. By definition, rule sets do not apply equally to all network interfaces; otherwise there would be no need for the firewall. For benchmarking purposes, a specific rule set is typically applied to each network interface in the DUT/SUT.

The tester must describe the complete contents of the rule set of each DUT/SUT.

To ensure measurements reflect only traffic forwarded by the DUT/SUT, testers are encouraged to include a rule denying all access except for those packets allowed by the rule set.

Unit of measurement:

not applicable

Issues:

See also:

- allowed traffic
- demilitarized zone (DMZ)
- illegal traffic
- policy
- protected network
- rejected traffic
- unprotected network

3.28 Security association

Definition:

The set of security information relating to a given network connection or set of connections.

Discussion:

This definition covers the relationship between policy and connections. Security associations (SAs) are typically set up during connection establishment, and they may be reiterated or revoked during a connection.

For purposes of benchmarking firewall performance, measurements of bit forwarding rate or UOTs per second must be taken after all security associations have been established.

Unit of measurement:

not applicable

See also:

- connection
- connection establishment
- policy
- rule set

3.29 Stateful packet filtering

Definition:

The process of forwarding or rejecting traffic based on the contents of a state table maintained by a firewall.

Discussion:

Packet filtering and proxy firewalls are essentially static, in that they always forward or reject packets based on the contents of the rule set.

In contrast, devices using stateful packet filtering will only forward packets if they correspond with state information maintained by the device about each connection. For example, a stateful packet filtering device will reject a packet on port 20 (ftp-data) if no connection has been established over the ftp control port (usually port 21).

Unit of measurement:

not applicable

Issues:**See also:**

- application proxy
- packet filtering
- proxy

3.30 Tri-homed

Definition:

A firewall with three network interfaces.

Discussion:

Tri-homed firewalls connect three network segments with different network addresses. Typically, these would be protected, DMZ, and unprotected segments.

A tri-homed firewall may offer some security advantages over firewalls with two interfaces. An attacker on an unprotected network may compromise hosts on the DMZ but still not reach any hosts on the protected network.

Unit of measurement:

not applicable

Issues:

Usually the differentiator between one segment and another is its IP address. However, firewalls may connect different networks of other types, such as ATM or Netware segments.

See also:

homed

3.31 Unit of transfer

Definition:

A discrete collection of bytes comprising at least one header and optional user data.

Discussion:

This metric is intended for use in describing steady-state forwarding rate of the DUT/SUT.

The unit of transfer (UOT) definition is deliberately left open to interpretation, allowing the broadest possible application. Examples of UOTs include TCP segments, IP packets, Ethernet frames, and ATM cells.

While the definition is deliberately broad, its interpretation must not be. The tester must describe what type of UOT will be offered to the DUT/SUT, and must offer these UOTs at a consistent rate. Traffic measurement must begin after all connection establishment routines complete and before any connection completion routine begins. Further, measurements must begin after any security associations (SAs) are established and before any SA is revoked.

Testers also must compare only like UOTs. It is not appropriate, for example, to compare forwarding rates by offering 1,500-byte Ethernet UOTs to one DUT/SUT and 53-byte ATM cells to another.

Unit of measurement:

Units of transfer

Units of transfer per second

Issues:

See also:

bit forwarding rate
connection

3.32 Unprotected network

Definition:

A network segment or segments to which access is not controlled by the DUT/SUT.

Discussion:

Firewalls are deployed between protected and unprotected segments. The unprotected network is not protected by the DUT/SUT.

Note that a DUT/SUT's policy may specify hosts on an unprotected network. For example, a user on a protected network may be permitted to access an FTP server on an unprotected network. But the DUT/SUT cannot control access between hosts on the unprotected network.

Unit of measurement:

not applicable

Issues:

See also:

demilitarized zone (DMZ)
policy
protected network
rule set

3.33 User

Definition:

A person or process requesting access to resources protected by the DUT/SUT.

Discussion:

"User" is a problematic term in the context of firewall performance testing, for several reasons. First, a user may in fact be a process or processes requesting services through the DUT/SUT. Second, different "user" requests may require radically different amounts of DUT/SUT resources. Third, traffic profiles vary widely from one organization to another, making it difficult to characterize the load offered by a typical user.

For these reasons, testers should not attempt to measure DUT/SUT performance in terms of users supported. Instead, testers should describe performance in terms of maximum bit forwarding rate and maximum number of connections sustained. Further, testers should use the term "data source" rather than user to describe traffic generator(s).

Unit of measurement:
not applicable

Issues:

See also:
data source

4. Security Considerations

The primary goal of this memo is to describe terms used in benchmarking firewall performance. However, readers should be aware that there is some overlap between performance and security issues. Specifically, the optimal configuration for firewall performance may not be the most secure, and vice-versa.

Further, certain forms of attack may degrade performance. One common form of denial-of-service (DoS) attack bombards a firewall with so much rejected traffic that it cannot forward allowed traffic. DoS attacks do not always involve heavy loads; by definition, DoS describes any state in which a firewall is offered rejected traffic that prohibits it from forwarding some or all allowed traffic. Even a small amount of traffic may significantly degrade firewall performance, or stop the firewall altogether. Further, the safeguards in firewalls to guard against such attacks may have a significant negative impact on performance.

Since the library of attacks is constantly expanding, no attempt is made here to define specific attacks that may affect performance. Nonetheless, any reasonable performance benchmark should take into

consideration safeguards against such attacks. Specifically, the same safeguards should be in place when comparing performance of different firewall implementations.

5. References

Bradner, S., Ed., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.

Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.

Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.

Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

6. Acknowledgments

The author wishes to thank the IETF Benchmarking Working Group for agreeing to review this document. Several other persons offered valuable contributions and critiques during this project: Ted Doty (Internet Security Systems), Kevin Dubray (Ironbridge Networks), Helen Holzbaaur, Dale Lancaster, Robert Mandeville, Brent Melson (NSTL), Steve Platt (NSTL), Marcus Ranum (Network Flight Recorder), Greg Shannon, Christoph Schuba (Sun Microsystems), Rick Siebenaler, and Greg Smith (Check Point Software Technologies).

7. Contact Information

David Newman
Data Communications magazine
3 Park Ave.
31st Floor
New York, NY 10016
USA

Phone: 212-592-8256
Fax: 212-592-8265
EMail: dnewman@data.com

8. Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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

