

Benchmarking Terminology Working Group

H.Berkowitz

Internet Draft

S.Hares

[draft-ietf-bmwg-conterm-00.txt](#)

A.Retana

Expires February 2002

P.Krishnaswamy

M.Lepp

July 2001

Terminology for Benchmarking External Routing Convergence Measurements

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

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>.

Abstract

This draft establishes terminology to standardize the description of benchmarking methodology for measuring eBGP convergence in the control plane of a single router. Future documents will address iBGP convergence, the initiation of forwarding based on converged control plane information and internet-wide convergence.

Conventions used in this document

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](#) [2].

Table of Contents

1	Introduction	3
1.1	Overview and Roadmap	3
1.2	Definition Format	4
2	Constituent elements of a router or network of routers.	5
2.1	BGP Peer	5
2.2	Default Route, Default Free Table, and Full Table	5
2.2.1	Default Route	6
	2.2.2 Default Free Routing Table	6
2.2.3	Full Default Free Table	7
2.2.4	Full Provider Internal Table	7
2.3	Classes of BGP-Speaking Routers	8
2.3.1	Provider Edge Router	8
2.3.2	Subscriber Edge Router	8
2.3.3	Interprovider Border Router	9
2.3.4	Intraprovider Core Router	9
3	Routing Data Structures	9
3.1	Routing Information Base (RIB)	9
3.1.1	Adj-RIB-In and Adj-RIB-Out	10
3.1.2	Loc-RIB	10
3.2	Policy	11
3.3	Policy Information Base	11
3.4	The Forwarding Information Base (FIB)	12
4	Components and characteristics of Routing information	13
4.1	Prefix	13
4.2	Route	13
4.3	BGP Route	14
4.4	Route Instance	14
4.5	Active Route	15
4.6	Unique Route	15
4.7	Non-Unique Route	15
4.8	Route Packing	16
4.9	Update Train	16
4.10	Route Flap	16
5	Route Changes and Convergence	17
5.1	Route Change Events	17
5.2	Convergence	18
6	Factors that impact the performance of the convergence process	19
6.1	General factors affecting BGP convergence	19
6.1.1	Number of peers	19
6.1.2	Number of routes per peer	20
6.1.3	Policy processing/reconfiguration	20
6.1.4	Interactions with other protocols.	20
6.1.5	Flap Damping	20
6.1.6	Churn	21
6.2	Implementation-specific and other factors affecting BGP convergence	21
6.2.1	Forwarded traffic	21
6.2.2	Timers	22
6.2.3	TCP parameters underlying BGP transport	22

6.2.4	Authentication	22
7	Security Considerations	22
8	References	22
9	Acknowledgments	22
10	Author's Addresses	23

[1](#) Introduction

This document defines terminology for use in characterizing the convergence performance of BGP processes in routers or other devices that instantiate BGP functionality. It is the first part of a two document series, of which the subsequent document will contain the associated tests and methodology.

The following observations underly the approach adopted in this, and the companion document:

- The principal objective is to derive methodologies to standardize conducting and reporting convergence-related measurements for BGP.
- It is necessary to remove ambiguity from many frequently used terms that arise in the context of such measurements.
- As convergence characterization is a complex process, it is desirable to restrict the initial focus in this set of documents to specifying how to take basic control plane measurements towards characterizing BGP convergence. For path vector protocols such as BGP, the primary initial focus will therefore be on network and system control-plane activity consisting of the arrival, processing, and propagation of routing information

Subsequent drafts will explore the more intricate aspects of convergence measurement, such as the impacts of the presence of policy processing, simultaneous traffic on the control and data paths within the DUT, and other realistic performance modifiers. Convergence of Interior Gateway Protocols will also be considered in separate drafts.

[1.1](#) Overview and Roadmap

A characterization of the BGP convergence performance of a device must take into account all distinct stages and aspects of BGP functionality. This requires that the relevant terms and metrics be as specifically defined as possible. Such definition is the goal of this document.

The necessary definitions are classified into two separate categories:

- . Descriptions of the constituent elements of a network or a router that is undergoing convergence
- . Descriptions of factors that impact convergence processes

[1.2](#) Definition Format

The definition format is equivalent to that defined in [[RFC1812](#)], and is repeated here for convenience:

X.x Term to be defined. (e.g., Latency)

Definition:

The specific definition for the term.

Discussion:

A brief discussion about the term, its application and any restrictions on measurement procedures.

Measurement units:

The units used to report measurements of this term, if applicable.

Issues:

List of issues or conditions that affect this term.

See Also:

List of other terms that are relevant to the discussion of this term.

[2](#) Constituent elements of a router or network of routers.

Many terms included in this list of definitions were described originally in previous standards or papers. They are included here because of their pertinence to this discussion. Where relevant, reference is made to these sources. An effort has been made to keep this list complete with regard to the necessary concepts without over definition.

[2.1](#) BGP Peer

Definition:

A BGP peer is another BGP instance to which the the Device Under Test (DUT) has established a TCP connection over which a BGP session is active. In the test scenarios in the methodology discussion that will follow this draft, peers send BGP advertisements to the DUT and receive

DUT-originated advertisements.

Discussion:

This is a protocol-specific definition, not to be confused with another frequent usage, which refers to the business/economic definition for the exchange of routes without financial compensation.

It is worth noting that a BGP peer, by this definition is associated with a BGP peering session, and there may be more than one such active session on a router or on a tester. The peering sessions referred to here may exist between various classes of BGP routers (see [section 2.3](#))

.

Measurement units: number of BGP peers

Issues:

See Also:

[2.2](#) Default Route, Default Free Table, and Full Table

An individual router's routing table may not necessarily contain a default route. Not having a default route, however, is not synonymous with having a full default-free table.

It should be noted that the references to number of routes in this section are to routes installed in the loc-RIB, not route instances, and that the total number of route instances may be 4 to 10 times the number of routes.

MPLS speaking routers are outside the scope of this document

[2.2.1](#) Default Route

Definition:

A Default Route is a route entry that can match any prefix. If a router does not have a route for a particular packet's destination address, it forwards this packet to the next hop in the default route entry, provided its Forwarding Table (Forwarding Information Base (FIB)) contains one. The notation for a default route is 0.0.0.0/0

Discussion:

Measurement units: N.A.

Issues:

See Also: default free routing table, route, route instance

[2.2.2](#) Default Free Routing Table

Definition:

A routing table with no default routes, as typically seen in routers at the core or top tier.

Discussion:

The term originates from the concept that routers at the core or top tier of the Internet will not be configured with a default route (Notation 0.0.0.0/0). Thus they will forward every prefix to a specific nexthop based on the longest match on the IP addresses.

Default free routing table size is commonly used as an indicator of the magnitude of reachable Internet address space. However, default free routing tables may also include routes internal to the router's AS.

Measurements: The number of routes

See Also: Full Default Free

[2.2.3](#) Full Default Free Table

Definition:

A set of BGP routes generally accepted to be the complete set of BGP routes announced by all autonomous systems to the public Internet. Due to the dynamic nature of the Internet, the exact size and composition of this table may vary slightly depending where it is observed.

Discussion:

Several investigators [Bates, Smith] measure this on a weekly basis; June 2001 measurements put the table at approximately 105,000 routes, growing exponentially.

It is generally accepted that a full table, in this usage, does not contain the infrastructure routes or individual subaggregates of routes that are otherwise aggregated by the provider before announcement to other autonomous systems.

Measurement Units: number of routes

Issues:

See Also: Routes, Route Instances, Default Route

[2.2.4](#) Full Provider Internal Table

Definition:

A superset of the full routing table that contains infrastructure and non-aggregated routes.

Discussion:

Experience has shown that this table can contain 1.3 to 1.5 times the number of routes in the externally visible full table. Tables of this size, therefore, are a real-world requirement for key internal provider routers.

Measurement Units: number of routes

Issues:

See Also: Routes, Route Instances, Default Route

[2.3](#) Classes of BGP-Speaking Routers

A given router may perform more than one of the following functions, based on its logical location in the network.

[2.3.1](#) Provider Edge Router

Definition:

A router at the edge of a provider's network, configured to speak BGP, which peers with a BGP speaking router operated by the end-user. The traffic that transits this router may be destined to, or originate from non-contiguous autonomous systems.

Discussion:

Such a router will always speak eBGP and may speak iBGP.

Measurement units:

Issues:

See Also:

[2.3.2](#) Subscriber Edge Router

Definition:

A BGP-speaking router belonging to an end user organization that may be multi-homed, which carries traffic only to and from that end user AS.

Discussion:

Such a router will always speak eBGP and may speak iBGP.

Measurement units:

Issues:

See Also:

[2.3.3](#) Interprovider Border Router

Definition:

A BGP speaking router that maintains BGP sessions with another BGP speaking router in another provider AS. Traffic transiting this router may be directed to or from another AS that has no direct connectivity with this provider's AS.

Discussion:

Such a router will always speak eBGP and may speak iBGP.

Measurement units:

Issues:

See Also:

[2.3.4](#) Intraprovider Core Router

Definition: A provider router speaking iBGP to the provider's edge routers, other intraprovider core routers, or the provider's interprovider border routers.

Discussion:

Such a router will always speak iBGP and may speak eBGP.

Measurement units:

Issues:

MPLS speaking routers are outside the scope of this document. It is entirely likely, however, that core Label Switched Routers, especially in the P router role of [RFC 2547](#), may contain little or no BGP information.

See Also:

[3](#) Routing Data Structures

[3.1](#) Routing Information Base (RIB)

The RIB collectively consists of a set of logically (not necessarily

literally) distinct databases, each of which is enumerated below. The RIB contains all destination prefixes to which the router may forward, and one or more currently reachable next hop addresses for them. Routes included in this set potentially have been selected from several sources of information, including hardware status, interior routing protocols, and exterior routing protocols. [RFC 1812](#) contains a basic set of route selection criteria relevant in an all-source context. Many implementations impose additional criteria. A common implementation-specific criterion is the preference given to different routing information sources.

[3.1.1](#) Adj-RIB-In and Adj-RIB-Out

Definition:

Adj-RIB-In and Adj-RIB-Out are "views" of routing information from the perspective of individual peer routers. The Adj-RIB-In contains information advertised to the DUT by a specific peer. The Adj-RIB-Out contains the information the DUT will advertise to the peer.

See [RFC 1771](#)

Discussion:

Issues:

Measurement Units: Number of route instances

See Also: Route, BGP Route, Route Instance, Loc-RIB, FIB

[3.1.2](#) Loc-RIB

Definition:

The Loc-RIB contains the set of best routes selected from the various Adj-RIBs, after applying local policies and the BGP route selection algorithm.

Discussion:

The separation implied between the various RIBs is logical. It does not necessarily follow that these RIBs are distinct and separate entities in any given implementation.

Issues:

Specifying the RIB is important because the types and relative proportions of routes in it can affect the convergence efficiency. Types of routes can include internal BGP, external BGP, interface, static and IGP routes.

Measurement Units: Number of route instances.

See Also: Route, BGP Route, Route Instance, Adj-RIB-in, Adj-RIB-out,

FIB

[3.2](#) Policy

Definition:

Policy is "the ability to define conditions for accepting, rejecting, and modifying routes received in advertisements"[16]

Discussion:

[RFC 1771](#) [[RFC1771](#)] further constrains policy to be within the hop-by-hop routing paradigm.

Policy is implemented using filters and associated policy actions.

Measurement Units:

Issues:

See Also: Policy Information Base.

[3.3](#) Policy Information Base

Definition:

A policy information base is the set of incoming and outgoing policies.

Discussion:

All references to the phase of the BGP selection process here are made with respect to [RFC 1771](#) [[RFC1771](#)] definition of these phases.

Incoming policies are applied in Phase 1 of the BGP selection process [[RFC1771](#)] to the Adj-RIB-In routes to set the metric for the Phase 2 decision process. Outgoing Policies are applied in Phase 3 of the BGP process to the Adj-RIB-Out routes preceding route (prefix and path attribute tuple) announcements to a specific peer.

Policies in the Policy Information Base have matching and action conditions. Common information to match include route prefixes, AS paths, communities, etc. The action on match may be to drop the update and not pass it to the Loc-RIB, or to modify the update in some way, such as changing local preference (on input) or MED (on output), adding or deleting communities, prepending AS in the AS path,

etc.

The amount of policy processing (both in terms of route maps and filter/access lists) will impact the convergence time and properties of the distributed BGP algorithm. The amount of policy processing may vary from a simple policy which accepts all routes and sends all routes to complex policy with a substantial fraction of the prefixes being filtered by filter/access lists.

[3.4](#) The Forwarding Information Base (FIB)

Definition:

The Forwarding Information Base is generated from the RIB. The FIB is referred to in [[RFC1771](#)] as well as [[RFC1812](#)] but not defined in either. For the purposes of this document, the FIB is the subset of the RIB used by the forwarding plane to make per-packet forwarding decisions.

Discussion:

Most current implementations have full, non-cached FIBs per router interface. All the route computation and convergence occurs before a route is downloaded into a FIB.

Measurement Units: N.A.

Issues:

See Also: Route, RIB

[4](#) Components and characteristics of Routing information

[4.1](#) Prefix

Definition:

A destination address in CIDR format. Expressed as prefix/length. The definition in [[RFC1812](#)] is "A network prefix is à a contiguous set of bits at the more significant end of the address that defines a set of systems; host numbers select among those systems."

Discussion:

A prefix is expressed as a portion of an IP address followed by the associated mask such as 10/8.

Measurement Units: N.A.

Issues:

See Also

[4.2](#) Route

Definition:

In general, a 'route' is the n-tuple <prefix, nexthop [,other non-routing protocol attributes] >. A route is not end-to-end, but is defined with respect to a specific next hop that will move traffic closer to the destination defined by the prefix. In this usage, a route is the basic unit of information about a target destination distilled from routing protocols.

Discussion:

This term refers to the concept of a route common to all routing protocols. With reference to the definition above, typical non-routing-protocol attributes would be associated with diffserv or traffic engineering.

Measurement Units: N.A.

Issues: None.

See Also: BGP route

[4.3](#) BGP Route

Definition:

The n-tuple <prefix, nexthop, aspath [, other BGP attributes]>.

Discussion:

Nexthop is one type of attribute. Attributes are defined in [RFC 1771](#) [[RFC1771](#)], and are the qualifying data that accompanies a prefix in a BGP route. From [RFC 1771](#): " For purposes of this protocol a route is defined as a unit of information that pairs a destination with the attributes of a path to that destination... A variable length sequence of path attributes is present in every UPDATE. Each path attribute is a triple <attribute type, attribute length, attribute value> of variable length."

Measurement Units:N.A.

Issues:

See Also: Route, prefix, Adj-RIB-in.

[4.4](#) Route Instance

Definition:

A single occurrence of a route sent by a BGP Peer for a particular prefix. When a router has multiple peers from which it accepts routes, routes to the same prefix may be received from several peers. This is then an example of multiple route instances.

Discussion:

Each route instance is associated with a specific peer. A specific route instance may be rejected by the BGP selection algorithm due to local policy.

Measurement Units: number of route instances

Issues:

The number of route instances in the Adj-RIB-in bases will vary based on the function to be performed by a router. An interprovider router, located in the default free zone will likely receive more route instances than an provider edge router, located closer to the end-users of the network.

See Also:

[4.5](#) Active Route

Definition:

Route for which there is a FIB entry corresponding to a RIB entry.

Discussion:

Measurement Units:

Issues:

See also: RIB.

[4.6](#) Unique Route

Definition:

A unique route is a prefix for which there is just one route instance across all Adj-Ribs-In.

Discussion:

Measurement Units: N.A.

Issues:

See Also: route, route instance

[4.7](#) Non-Unique Route

Definition:

A Non-unique route is a prefix for which there is at least one other route in a set including more than one Adj-RIB-in.

Discussion:

Measurement Units: N.A.

Issues:

See Also: route, route instance, unique active route.

[4.8](#) Route Packing

Definition:

Number of route prefixes that exist in a single Routing Protocol Update or Withdraw Message.

Discussion:

In general, a routing protocol update MAY contain more than one prefix. In BGP, a single update MAY contain multiple prefixes with identical attributes.

Protocols that do not support such a concept implicitly have a Route Packing of 1. "

[4.9](#) Update Train

Definition:

A set of updates, containing one or more route prefixes, which an external router desires to send to the DUT. When there is more than one prefix in the set, the multiple updates (including withdrawals) may be sent as individual BGP UPDATE packets, or as one or more BGP packets with multiple routes packed (q.v.) into them.

Discussion:

The more individual update packets that are sent, the more TCP and BGP header processing will be imposed on the receiving router that is the DUT.

An update train may be caused by a variety of network conditions. For example, an update train could be caused by an influx of UPDATES from different peers that have been received and moved to RIB-out or caused by a peer coming up and advertising its routes, or by a local or remote peer flapping a link. Other causes are, of course, possible.

Measurement units: Number of prefixes and update packets in the train.

Issues:

See Also:

[4.10](#) Route Flap

Definition: RIPE 210 [RIPE]define a route flap as "the announcement and withdrawal of prefixes." For our purposes we define a route flap as the rapid withdrawal/announcement or announcement/withdrawal/ of a prefix in the Adj-RIB-in.

A route flap is not a problem until a route is flapped several times in close succession. This causes negative repercussions throughout the internet.

Discussion:

Route flapping can be considered a special and pathological case of update trains.

A practical interpretation of what may be considered excessively rapid is the RIPE recommendation of "four flaps in a row". See [Section 6.1.5](#) on flap damping for further discussion.

Measurement units

Flapping events per unit time.

Issues:

Specific Flap events can be found in [Section 5.1](#) Route Change Events. A bench-marker should use a mixture of different route change events in testing.

See Also: Route change events, flap damping, packet train

[5](#) Route Changes and Convergence

The following two definitions are central to the benchmarking of external routing convergence, and so are singled out for more extensive discussion.

[5.1](#) Route Change Events

A taxonomy characterizing routing information changes seen in operational networks is proposed in [Ahuja et al] as well as Labovitz et al[4]. These papers describe BGP protocol-centric events, and event sequences in the course of an analysis of network behavior. The terminology in the two papers addresses similar but slightly different behaviors with some overlap. We would like to apply these taxonomies to categorize the tests under definition where possible, because these tests must tie in to phenomena that arise in actual networks. We avail of, or may extend, this terminology as necessary for this purpose.

A route can be changed implicitly by replacing it with another route or explicitly by withdrawal followed by the introduction of a new route. In either case the change may be an actual change, no change, or a duplicate. The notation and definition of individual categorizable route change events is adopted from [Labovitz et al] and given below.

- a) AADiff: Implicit withdrawal of a route and replacement by a route different in some path attribute.
- b) AADup: Implicit withdrawal of a route and replacement by route that is identical in all path attributes.
- c) WADiff: Explicit withdrawal of a route and replacement by a different route.
- d) WADup: Explicit withdrawal of a route and replacement by a route that is identical in all path attributes.

To apply this taxonomy in the benchmarking context, we need both terms to describe the sequence of events from the update train perspective, as listed above, as well as event indications in the time domain so as to be able to measure activity from the perspective of the DUT.

With this in mind, we incorporate and extend the definitions of [4] to the following:

- a) Tup (TRx): Route advertised to the DUT by Test Router x
- b) Tdown: Route being withdrawn
- c) Tupinit: The initial announcement of a route to a unique prefix
- d) TWF: Route fail over after an explicit withdrawal.

The basic criterion for selecting a "better" route is the final tiebreaker defined in [RFC1771](#), the router ID. As a consequence, this memorandum uses the following descriptor events.

- a) Tbest -- The current best path.
- b) Tbetter -- Advertise a path that is better than Tbest.
- c) Tworse -- Advertise a path that is worse than Tbest.

[5.2](#) Convergence

Definition:

A router is said to have converged onto a route advertised to it, given that the route is the best route instance for a prefix (if multiple choices exist for that prefix), when this route is advertised to its downstream peers.

Discussion:

The convergence process can be subdivided into three distinct phases:

- convergence across the entire Internet,
- convergence within an Autonomous System,
- convergence with respect to a single router.

Convergence with respect to a single router can be

- convergence with regard to the routing process(es), the focus of this document
- convergence with regard to data forwarding process(es).

It is of key importance to benchmark the performance of each phase of convergence separately before proceeding

to a composite

characterization of routing convergence, where implementation-specific dependencies are allowed to interact.

The preferred route instance must be unambiguous during test setup/definition.

Measurement Units: N.A.

Issues:

See Also:

[6](#) Factors that impact the performance of the convergence process

While this is not a complete list, all of the items discussed below have a significant affect on BGP convergence. Not all of them can be addressed in the baseline measurements described in this document.

[6.1](#) General factors affecting BGP convergence

These factors are conditions of testing external to the router Device Under Test (DUT).

[6.1.1](#) Number of peers

As the number of peers increases, the BGP route selection algorithm is increasingly exercised. In addition, the phasing and frequency of updates from the various peers will have an increasingly marked effect on the convergence process on a router as the number of peers grows.

[6.1.2](#) Number of routes per peer

The number of routes per BGP peer is an obvious stressor to the convergence process. The number, and relative proportion, of multiple route instances and distinct routes being added or withdrawn by each peer will affect the convergence process, as will the mix of overlapping route instances, and IGP routes.

[6.1.3](#) Policy processing/reconfiguration

The number of routes and attributes being filtered, and set, as a fraction of the target route table size is another parameter that will affect BGP convergence.

Extreme examples are

-Minimal Policy: receive all, send all,

-Extensive policy: up to 100% of the total routes have applicable policy.

[6.1.4](#) Interactions with other protocols.

There are interactions in the form of precedence, synchronization, duplication and the addition of timers, and route selection criteria. Ultimately, understanding BGP4 convergence must include understanding of the interactions with both the IGP and the protocols associated with the physical media, such as Ethernet, SONET, DWDM.

[6.1.5](#) Flap Damping

A router can use flap damping to respond to route flapping. Use of flap damping is not mandatory, so the decision to enable the feature, and to change parameters associated with it, can be considered a matter of routing policy.

The timers are defined by [RFC 2439](#) [[RFC 2439](#)] and discussed in RIPE 210 [17].

If this feature is in effect, it requires that the router keep additional state to carry out the damping, which can have a direct impact on the control plane due to increased processing. In addition, flap damping may delay the arrival of real changes in a route, and affect convergence times

[6.1.6](#) Churn

In theory, a BGP router could receive a set of updates that completely defined the Internet, and could remain in a steady state, only sending appropriate KeepAlives. In practice, the Internet will always be changing.

Churn refers to control plane processor activity caused by announcements received and sent by the router. It does not include KeepAlives.

Churn is caused by both normal and pathological events. For example, if an interface of the local router goes down and the associated prefix is withdrawn, that withdrawal is a normal activity, although it contributes to churn. If the local router receives a withdrawal of a route it already advertises, or an announcement of a route it did not previously know, and readvertises this information, again these are normal constituents of churn. Routine updates can range from single announcement or withdrawals, to announcements of an entire default-free table. The latter is completely reasonable as an initialization condition.

Flapping routes are a pathological contributor to churn, as is MED oscillation [medosc]. The goal of flap damping is to reduce the contribution of flapping to churn.

The effect of churn on overall convergence depends on the processing power available to the control plane, and whether the same processor(s) are used for forwarding and for control.

[6.2](#) Implementation-specific and other factors affecting BGP convergence
These factors are conditions of testing internal to the router Device Under Test (DUT), although they may affect its interactions with test devices.

[6.2.1](#) Forwarded traffic

The presence of actual traffic in the router may stress the control path in some fashion if both the offered load due to data and the control traffic (FIB updates and downloads as a consequence of flaps) are excessive. The addition of data traffic presents a more accurate reflection of realistic operating scenarios than if only control traffic is present

[6.2.2](#) Timers

Settings of delay and hold-down timers at the link level as well as for BGP4, can introduce or ameliorate delays. As part of a test report, all relevant timers should be reported if they use non-default value. Also, any variation in standard behavior, such as overriding TCP slow start, should be documented.

[6.2.3](#) TCP parameters underlying BGP transport

Since all BGP traffic and interactions occur over TCP, all relevant parameters characterizing the TCP sessions should be provided: eg Max window size, Maximum segment size, timers.

[6.2.4](#) Authentication

Authentication in BGP is currently done using the TCP MD5 Signature Option [Heff]. The processing of the MD5 hash, particularly in routers with a large number of BGP peers and a large amount of update traffic can have an impact on the control plane of the router.

[7](#) Security Considerations

The document explicitly considers authentication as a performance-affecting feature, but does not consider the overall security of the routing system.

[8](#) References

- [1](#). [[RFC 2026](#)]Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996
- [2](#). [[RFC 2119](#)]Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997

3. [[RFC 1771](#)] "A Border Gateway Protocol 4 (BGP-4)", [RFC 1771](#), Y. Rekhter, T. Li. March 1995.
4. [[RFC 2539](#)] ""BGP Route Flap Damping", [RFC 2439](#), C. Villamizar, R. Chandra, R. Govindan. November 1998."
5. [[RFC 1812](#)] "Requirements for IP Version 4 Routers", [RFC 1812](#), F. Baker. June 1995.
6. [Ahuja et al] "An Experimental Study of Delayed Internet Routing Convergence." Abha Ahuja, Farnam Jahanian, Abhijit Bose, Craig Labovitz, RIPE 37 - Routing WG.
7. [Labovitz et al] "Origins of Internet Routing Instability," Infocom 99 Craig Labovitz, G. Robert Malan, Farnam Jahanian],
8. [RPSL] Routing Policy Specification Language (RPSL), [RFC 2622](#), " C.Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T.Bates, D. Karrenberg, M. Terpstra. June 1999.
9. [RIPE] RIPE 210, "RIPE Routing-WG Recommendation for coordinated route-flap damping parameters, Tony Barber, Sean Doran, Daniel Karrenberg, Christian Panigl, Joachim Schmitz
10. [Heff] "Protection of BGP Sessions via the TCP MD5 Signature Option", [RFC 2385](#), A. Heffernan. August 1998.
11. [medosc] <[draft-ietf-idr-route-oscillation-00.txt](#)>, BGP Persistent Route Oscillation Condition, McPhersonm, Gill, Walton, Retana

9 Acknowledgments

Thanks to Francis Ovenden and Elwyn Davies for review and Abha Ahuja for encouragement. Much appreciation to Jeff Haas, Matt Richardson, and Shane Wright at Nexthop for comments and input. Debby Stopp and Nick Ambrose contributed the concept of route packing.

10 Author's Addresses

Howard Berkowitz
Nortel Networks
5012 S. 25th St
Arlington VA 22206
Phone: +1 703 998-5819 (ESN 451-5819)
Fax: +1 703 998-5058

EMail: hberkowi@nortelnetworks.com

Susan Hares
Nexthop Technologies
517 W. William
Ann Arbor, Mi 48103
Phone:
Email: skh@nexthop.com

Padma Krishnaswamy
Nexthop Technologies
517 W William
Ann Arbor, Mi 48103
Phone: 734 973 2200
Email: kri@nexthop.com

Marianne Lepp
Juniper Networks
51 Sawyer Road
Waltham, MA 02453
Phone: 617 645 9019
Email: mlepp@juniper.net

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
Email: aretana@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2000). 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.