

DNA Working Group
Internet-Draft
Expires: January 10, 2005

G. Daley
B. Pentland
Monash University CTIE
E. Nordmark
Sun Microsystems
July 12, 2004

Deterministic Fast Router Advertisement Configuration
draft-daley-dna-det-fastra-00.txt

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on January 10, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Neighbour Discovery forces routers responding to Router Solicitations to delay a random interval of 0-500 ms in order to prevent contention and bandwidth utilization by multiple responding routers. Routers receiving solicitations may need to quickly send responding Router Advertisements faster than allowed in IPv6 Neighbour Discovery.

In order to provide fast router response, Fast Router Advertisements

have been proposed, which do not randomly delay. This document describes an automatic arbitration and configuration mechanism to allow hosts to securely agree on the relative ordering and delay of their Fast Router Advertisements.

Table of Contents

1.	Introduction	4
2.	Terminology	4
3.	Protocol Operation	6
3.1	Deterministic Fast Router Advertisement Concepts	6
3.2	Router-to-Router Status Communication	7
3.3	Deterministic Fast Router Advertisement Options	8
4.	Message Formats	8
4.1	Router to Router ICMPv6 message	8
4.2	Deterministic Fast Router Advertisement Option Format	10
4.2.1	Rank	11
4.2.2	Fixed Delay	12
4.2.3	Separation	12
4.2.4	Relative Preference	12
5.	Sending Router-to-Router messages	12
5.1	Sending Router-to-Router Status-Requests	12
5.2	Sending Router-to-Router Status	13
6.	Ranking Calculation	13
6.1	Ranking Score Calculation Reasoning	14
7.	Becoming a Fast Router	14
8.	Receiving DetFRAO in ICMPv6 Router-to-Router	15
8.1	Dealing with Duplicate List Entries	15
8.2	Receiving Bootstrap Deterministic FastRA Options	16
8.3	Cleaning up old entries	16
8.4	Responding to Status-Requests with DetFRAO	17
9.	Sending DetFRAOs in ICMPv6 Router-to-Router	17
9.1	Sending RAs on Rank or Delay Change	17
9.2	Sending Advertisement Intervals	17
10.	Sending Fast Router Advertisements	17
10.1	Sending Unicast Fast RAs	18
10.2	Sending Multicast Fast RAs	18
11.	Interaction with Other Routers	19
11.1	Presence of Legacy Routers	19
11.2	Ceasing to be a Fast Router	19
11.3	Presence of Non Fast Routers	19
11.4	Presence of Incompatible Fast Routers	20
12.	Host interaction with DetFRAOs	20
12.1	Host Transmission of DetFRAOs	20
12.2	Reception of DetFRAOs in Router Solicitations	20
12.3	Transmission of DetFRAOs in Router Advertisements	20
12.4	Host Reception of DetFRAOs	20
13.	Configuration Parameters	21

14.	Protocol Constants	22
15.	IANA Considerations	23
16.	Security Considerations	23
17.	Acknowledgments	24
18.	References	24
18.1	Normative References	24
18.2	Informative References	24
	Authors' Addresses	25
A.	State Diagram	25
	Intellectual Property and Copyright Statements	27

1. Introduction

Existing standards require that hosts which respond to Router Solicitations introduce a random delay of 0-500ms before sending a Router Advertisement [2].

The Fast RA draft [6] introduces the concept of a Fast Router Advertisement. Fast Router advertisements provide the capability for a router to avoid performing the random delay before transmission, and send responses without the mean 250ms delay.

Additionally, routers may wish to allow a small number of multicast Router Advertisements to configuring devices with similar delays. Extensions to FastRA which support multicast Router Advertisement are considered in this document (possibly to move elsewhere later).

Fast Router Advertisement, as specified, only allows one host on a link to be a FastRA router. Unless FastRA specified this limit, responses from more than one FastRA router would result in the same MAC contention which [RFC 2461](#) sought to avoid.

Unfortunately, while Fast RA is an effective way of providing fast response, it has no defined automated configuration mechanism to determine which router is the fastest, nor any method to provide router reselection in case a router is no longer able to provide fast responses. This lack of a router reselection mechanism is seen a clear stumbling block to FastRA's deployment.

This document defines a secured mechanism relying on new Router-to-Router ICMPv6 messages to allow routers to make decisions about which router responds fastest, and additionally allows other routers to avoid random delays[5][2]. The new mechanism relies upon SEND-like security to determine trust-relationships between on-link routers, and test the reachability of existing routers.

It allows automated reconfiguration in the case that routers join or leave the Fast Router set, and ensures that Fast Routers do not collide with each other by providing negotiated fixed delays between responding advertisements.

2. Terminology

The following terms are used throughout the document

Bootstrapping Router: A router which has not yet determined its rank and delay for Fast Router Advertisement.

Deterministic Fast Router Advertisement Option (DetFRAO): An ICMPv6 option indicating the Fast Router's participation in this protocol, as well as its rank and delay. This option may appear in Router-to-Router Status-Requests and Status messages, as well as Router Solicitations and Advertisements. The option's format is defined in [Section 4.2](#).

Fast Router: A router participating in this protocol and exchanging Deterministic Fast Router Advertisement Options in Router-to-Router messages.

Fast Router Advertisement: A response to a Router Solicitation which is not randomly delayed. Please note that at a particular time, a Fast Advertising Router may advertise with a delay whose mean is slower than that defined by [2]. Even in this case, the response is still called a Fast Router Advertisement (or FastRA).

Fast Advertising Router List: A conceptual list where FastRA routers are sorted by a Ranking Score, and delays for various routers calculated.

Fastest Advertising Router: The router with Rank=1, which is able to transmit without delay.

Fixed Delay: The minimum amount of time which a Fast Advertising Router may delay before sending a Router Advertisement in response to a Router Solicitation. Typically this is the delay used for Router Advertisement transmission.

Legacy Router: A router which responds to router solicitations using the algorithm defined in [section 6.2.6](#) of the IPv6 Neighbour Discovery RFC[2].

Multicast Fast Router Advertisement A Fast Router Advertisement sent to the all-nodes multicast address. These advertisements have different parameters for their Token Bucket than a Unicast Fast Router Advertisement.

Rank: The position of the router in the advertising order. A router with a better (lower numbered) rank will always send responding advertisements faster than one with a worse rank.

Ranking Identifier: An value derived from the identity of the advertising router, used in ranking calculations.

Ranking Score: A score calculated from the router's preference and Ranking Identifier. This score is used to order the routers on the link and determine their Rank.

Router-to-Router message: A new ICMPv6 message which uses the same options format as IPv6 Neighbour Discovery, but is only used for communication between routers on a local link. It provides a means by which routers can advertise their configuration status and request that other routers do the same.

Separation: A period after the transmission of a previous RA that the router must wait. The separation time is defined by the preceding router.

Token Bucket A finite sized non-negative resource counter which increases at a set rate while the number of resource tokens in the bucket is not at maximum. When a resource is to be allocated, there must be a non-zero number of resource tokens in the bucket. As the token is allocated the resource counter decrements.

Unicast Fast Router Advertisement A Fast Router Advertisement sent to the unicast source address of a Router Solicitation. Unicast Fast RAs can be sent while there are tokens in the Unicast FastRA token bucket as defined in [6].

3. Protocol Operation

Routers wishing to provide Fast Router Advertisement need to check if other routers on the link are providing Fast RAs and agree on a relative ordering and delay before response. This negotiation takes place prior to Fast RA operation, when routers are first configured, start or change their preferences.

This allows all Fast Routers to send responses to Router Solicitations at the agreed delay, without introducing additional variable delays. Delays and ordering are therefore deterministic, and one responding Fast Advertising Router (the Fastest) will be able to transmit Router Advertisements in response to solicitation with no delay at all.

During transition periods where router ordering changes, router advertise their changed configuration to peer routers using ICMPv6 Router-to-Router messages, defined in this document.

3.1 Deterministic Fast Router Advertisement Concepts

To agree on when to send responses, Fast RA routers need to know which routers are Fast Routers, and must agree on their relative order for RS response. In this document, the ordinal position agreed to by a router is its Rank.

The lowest number provides the best Rank, and the fastest responding router has a Rank of 1. The ranking algorithm seeks to avoid ties, although from time to time, multiple routers will be seen to have the same Rank. Handling of this condition is specified in [Section 6](#).

Upon determining the advertisement order, each router needs to choose a delay exceeding that advertised by its better Ranked routers. To do this it inspects advertised values for other routers and advertises its own Fixed Delay value exceeding this. The Fixed Delay is the lowest value used by the router as delay for responding to Router Solicitations. In this fashion, a poorly ranked router needs only to inspect the immediately preceding routers' status messages to guarantee that it exceeds the expected values.

At this stage, routers introduce a Separation time, which is used to separate the Fast Router Advertisements. A worse Ranked router must select a delay greater than or equal to the sum of the advertised Fixed Delay and Separation of its immediately preceding router(s).

[3.2](#) Router-to-Router Status Communication

In order to accomplish Ranking and delay agreement between routers on a link, messages need to be exchanged between FastRA routers. These messages contain information about the router's current configuration status, and indicate that FastRA is in use. This document specifies the Router-to-Router(RtR) ICMPv6 message which allows configuration status to be requested from other routers while advertising the router's own status.

In negotiating with other routers on a link, trust of a router's identity and authorization needs to be established, and mechanisms to check these must be devised. The Router-to-Router message is able to use the existing message formats for Secure Neighbour Discovery, and uses the same signatures and keys which are used in Router Solicitations and Advertisements.

Existing protocol messages such as Router Solicitation and Router Advertisement were explicitly designed for communication between routers and autoconfiguring hosts. While the options deployed in this document may have worked interoperably in existing Neighbour Discovery messages, existing assumptions about the roles of particular message senders (particularly as defined in [Appendix D](#) of [2]) would be violated in doing so.

The newly defined Router-to-Router message aims to be compatible with future protocols requiring router negotiation and agreement, and defines an extensible option format in the same manner as IPv6 Neighbour Discovery.

ICMP Fields:

Type	TBD (Suggest 191)
Code	0 - Status Request 1 - Status
Checksum	The ICMPv6 checksum.
Reserved	This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Valid Options:

Advertisement Interval

The maximum delay between unsolicited Router Advertisement messages. This option and its use are defined in Mobile IPv6.

Deterministic Fast Router Advertisement

The configuration status for FastRA, as defined in

[Section 4.2](#)

of this document.

CGA	An option providing information about the router's cryptographically generated address, as defined in SEND.
Nonce	This option is provided in Status-Request messages, and copied into Status messages when responding to a particular Status-Request. It is defined in SEND.
Timestamp	An indication of the time of a Status-Request or Status message, as perceived at the transmitter. This option aims to reduce the chance of messages being replayed, and is defined in SEND.
Signature	A digital signature over the message (including the CGA Type Tag, IPv6 Source and Destination addresses and the entire ICMPv6 message without Signature option). This option MUST be the last option in the message, if present, and is defined in SEND.

Future versions of this protocol may define new option types. Receivers MUST silently ignore any options they do not recognize and continue processing the message.

Mobile IPv6 is defined in [4], SEND is defined in [5].

The Router-to-Router message comes with two different codes. A Status-Request message asks peer routers which understand the message to report back with their configuration for the options contained in the message. A Status message either responds to a Status-Request or periodically updates its peers of the configuration.

Sending of Status-Request messages is defined in [Section 5.1](#).

Processing of Status-Request messages is performed as described for each received option type (unknown options being ignored). A router MUST respond to a Status-Request message with a Status message, even if it contains no configuration status information.

Transmission of Status messages is defined in [Section 5.2](#).

SEND options are ignored as regards requests for status by receiving routers.

SEND option processing is as detailed in [5] and particularly, all secured Router-to-Router messages MUST contain the Signature and Timestamp options. Status-Request messages act as if they are solicitations for the inclusion of CGA options and Nonce Options. CGA and Nonce Options MUST also be present in Status messages which respond to a Status-Request. Where the Status messages are sent otherwise they SHOULD include the CGA Option regularly to ensure that routers newly arrived on the link are able to verify their message.

[4.2](#) Deterministic Fast Router Advertisement Option Format

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Type										Length										Rank										Reserved									
Fixed Delay										Separation										Rel Pref																			

Fields:

Type	TBD (Suggest 13)
Length	The length of the option (including the type and length fields) in units of 8 octets. Routers compliant with this document MUST set the length to 1. If the option is received in a Router-to-Router message with a length greater than one the router MUST NOT read the option's contents, and MUST set the IncompatibleFastRouters flag on the interface until the advertising router disappears or begins advertising with an option length of 1.
Rank	An integer in the range 0-255 indicating the ordinal position of the Fast Router.
Reserved	This field MUST be set to zero and be MUST be ignored by receivers.
Fixed Delay	The minimum delay used by this router to send Router Advertisements in response to Router Solicitation. (Units: milliseconds)
Separation	A delay after the current router's Fixed Delay that worse ranked routers wait before transmitting. (Units: milliseconds)
Rel Pref	The Relative Preference of the router seeking to perform Fast Router Advertisement. This is set to the variable FastRARElPref.

This option may appear in Router-to-Router, Router Solicitation and Router Advertisement messages. Nodes receiving this option in other messages MUST ignore it, acting as if they didn't understand the option.

[4.2.1](#) Rank

The Router ranked 1 will be the fastest router, and SHOULD configure a Fixed Delay of 0. No router may adopt a rank (other than BOOTSTRAP_RANK) until it has undertaken router querying as defined in [Section 7](#).

4.2.2 Fixed Delay

The router determines its Fixed Delay by summing its immediately preceding router's Fixed Delay and Separation values.

The Fastest Router SHOULD set its Fixed Delay to 0.

4.2.3 Separation

This value for the Fast Router's Separation from subsequent Fast Routers exceeds the maximum additional delay over Fixed Delay required to transmit a Fast Router Advertisement.

This delay MUST allow for computation delays in forming the RA (such as incurred with SEND) [[5](#)].

4.2.4 Relative Preference

This is an integer number providing the relative preference of the fast router for Rank determination. It is controlled by the variable FastRRelPref which MUST be configurable by the router's administrator. This value is prepended to the 24 bit Rank Identifier in order to provide a Ranking Score, as described in [Section 6](#).

A lower value will increase the preference of the router, and will outrank routers configured with the default preference value of DEF_REL_PREF.

Changes in a router's FastRA preference MAY be advertised immediately based on its newly calculated Rank, since routers should be checking if the Ranking Identifier associated with a particular router already exists in the Fast Router List, and will move the existing entry to its new location in the Fast Router List.

5. Sending Router-to-Router messages

Router-to-Router messages indicate the router's current configuration and may request a response from a peer router, with its own configuration status.

5.1 Sending Router-to-Router Status-Requests

When seeking to learn about other routers' status, a router may send Router-to-Router status Requests to its peers.

In doing so the router delays randomly for between 0 and MAX_RTR_STATUS_REQ_DELAY, and then sends up to MAX_RTR_STATUS_REQS requests separated by at least RTR_STATUS_REQ_INTERVAL seconds.

These messages are sent to the all-routers multicast group, and contain the Nonce, Signature, CGA and Timestamp options (at least) if CGA is used.

5.2 Sending Router-to-Router Status

Router-to-Router Status messages MAY be sent periodically to other routers to update the peers although the router's own reachability is readily confirmed by periodic unsolicited RA reception.

When changes in configuration occur, the router SHOULD send up to MAX_RTR_STATUS_REQS separated by MIN_DELAY_BETWEEN_RTR_STATUS after an initial random delay of between 0 and MAX_RTR_STATUS_DELAY.

Responses to Status-Request messages MUST be sent.

This responding Status MUST be sent after a random delay of 0 to MAX_RTR_STATUS_DELAY. Additionally, multicast Status Messages MUST NOT be sent more regularly than once every MIN_DELAY_BETWEEN_RTR_STATUS on average.

Responding Status messages MAY be sent to the unicast source address of the Status-Request. If MIN_DELAY_BETWEEN_RTR_STATUS has elapsed since the last multicast Status message, the response SHOULD be multicast to the all-routers group.

6. Ranking Calculation

When receiving Router-to-Router messages containing the DetFRA0, a router may maintain a list of Fast Routers and determines a relative order amongst them. Changes in the relative order trigger changes in Router Advertisement delays, so calculations for Rank need to be done simply, and reliably on information available in each Router-to-Router message with a Deterministic FastRA option.

Every Router-to-Router message is sent from a unicast link-local address uniquely owned (on the link) by the router. Additionally, every DetFRA Option contains a Rel Pref variable, which indicates the configured preference of this router over others.

When receiving an Router-to-Router message containing a DetFRA0, the Ranking Score is calculated by appending the Rank Identifier to the Rel Pref value received in the DetFRA0.

The Rank Identifier is created by taking the final 24 bits of the router's advertising link-local address and XOR'ing this value with 0xffffffff.

As an example, using 32 bit unsigned integers, Ranking Score computation is as follows:

```
RankID = ((ntohl(ll.s6_addr32[3]) & 0x00ffffff) ^ 0x00ffffff);  
RScore = RankID | (relpref << 24);
```

A discussion of the reasoning behind this algorithm is provided in [Section 6.1](#).

The router with the lowest Ranking Score assumes the Rank: 1, and the next lowest Rank: 2, etc.

In the situation where routers share the same Ranking Score and Identifier, comparison of the routers' complete IPv6 link-local address is made, in order to break ties.

The Ranking Identifier SHOULD be used to determine if a received router's preference has changed, by checking if the Ranking Identifier (as a lookup for the Source Address of the received message) is already present in the list, and has a different Ranking Score. In this case, the old entry is overridden by a more recent advertisement, and list entries consequently are re-Ranked.

Note that during ranking calculation the advertised option's Rank field is not consulted, although it may be checked subsequently.

[6.1](#) Ranking Score Calculation Reasoning

Selecting the low-order 24 bits of the IPv6 address allows selection of fastest router in the case of interface identifier creation from MAC-48 addresses, without reference to manufacturer's Organizationally Unique Identifier (OUI). Use of the XOR reverses the order of the IPv6 address suffix so that EUI-64 based addresses favour newer routers rather than older ones (unlike in [\[8\]](#) for the case where OUIs are the same), and the relative preference overrides all[\[7\]](#).

Maintaining this structure (RelPref,Suffix) allows routers to check not only the relative ordering of a router in the list on DetFRA0 reception, but allows even Router Advertisements which do not contain DetFRA0s to update the validity of the Fast Router's existing list entry by matching Rank Identifiers created from the RA's source address (if this is a unique match).

[7](#). Becoming a Fast Router

When a router wishes to be a Fast Router, it needs to check if anyone

else is acting as a Fast Router on this link. The router begins this bootstrap process sending Status-Request messages containing a DetFRA0, as defined in [Section 5.1](#).

The Deterministic FastRA option in this case MUST contain the values:

```
Rank          = BOOTSTRAP_RANK
Fixed Delay= BOOTSTRAP_DELAY
Separation = FastRASep
Rel Pref      = FastRARElPref
```

Receivers will see this option and recognise that the requester is a bootstrapping router. As defined in [Section 8.4](#), the routers MUST send a Status message responding to the request containing the DetFRA option. This informs the requester of their own identity, Rank and delays.

While collecting information about other routers, the bootstrapping router MUST send Router-to-Router messages with the DetFRA option. It MUST delay when sending responses to Router Solicitations by 0 to MAX_RA_DELAY_TIME, and ensure that MIN_DELAY_BETWEEN_RAS separates multicast advertisements.

After collecting this information, the new Fast Router calculates its Rank and begins advertising as described in [Section 9.1](#).

8. Receiving DetFRA0 in ICMPv6 Router-to-Router

When receiving Router-to-Router messages with the DetFRA0 option, the router determines whether the transmitting router has been seen before, and whether the transmitted Rank, Fixed Delay or Separation have changed. If either the router is previously unseen or the ranking or delay parameters have changed, the entry is inserted into the list at the position indicated by the router's Ranking Scores.

If the delay or rank of the receiving router have changed, it advertises its changed configuration as indicated in [Section 9.1](#).

It SHOULD also send a unicast Status message to the transmitter of a Status-Request message if change occurred due to the option in that request.

8.1 Dealing with Duplicate List Entries

Where Router-to-Router messages are received with a DetFRA Option indicating the the same rank as another preceding entry, then the receiving router MUST configure its Rank to be the number of elements

preceding it in the Fast Routers List plus one, and the router's own fixed delay MUST be configured to the maximum immediately preceding routers' fixed delays plus both of the other routers' received Separation values.

This ensures the router is will not then collide with a router which subsequently increases its Rank.

When this change occurs, advertisement follows the procedures in [Section 9.1](#).

[8.2](#) Receiving Bootstrap Deterministic FastRA Options

Deterministic FastRA routers or bootstrapping routers may receive Router-to-Router messages containing a DetFRAO from a bootstrapping router.

Routers SHOULD add such routers into their Fast Router List, in anticipation of the router's arrival as a fully functioning Fast Router.

Calculations for the Rank and Fixed delay of the bootstrapping Router MUST NOT be made based on the values in the received Option, but on the Ranking Score generated as described in [Section 6](#). The Fixed Delay calculated for the bootstrap router should be the sum of the preceding router's Fixed Delay and Separation.

[8.3](#) Cleaning up old entries

If a Fast Router fails to receive multiple expected Router Advertisement packets from a peer router, it SHOULD check if the peer router is dead using either router or neighbour discovery . Such mechanisms may be invoked upon non-reception of advertisements in multiple retransmission intervals as advertised by the peer router, or non-response to previous Router-to-Router Status-Request [[4](#)].

If the peer is dead, the local router removes its entry from the list, and re-advertises its own preference and distance as defined in [Section 9.1](#), if any change has occurred.

If one of a router's preceding routers reduces their Rank, so that it conflicts with another router in the list, a router MAY send a Router-to-Router Status-Request message containing DetFRAO after a random delay between 0 and MAX_RTR_STATUS_REQ_DELAY, to establish whether routers are still reachable.

8.4 Responding to Status-Requests with DetFRAO

A router MUST send a Deterministic FastRA option to a router which sends a Router-to-Router Status-Request to it, if the router is trusted. Delays to Status message are described in [Section 5.2](#).

9. Sending DetFRAOs in ICMPv6 Router-to-Router

9.1 Sending RAs on Rank or Delay Change

In the case that a router's Rank, Fixed Delay or Separation change, it MUST transmit a Router-to-Router Status message after a delay of between 0 and MAX_RTR_STATUS_DELAY seconds and then MAX_INITIAL_RTR_STATUS-1 messages each separated by MIN_DELAY_BETWEEN_RTR_STATUS seconds (as described in [Section 5.2](#)).

If subsequent changes occur within this interval, it extends such that three multicast Router-to-Router Status messages are sent with the new configuration. This allows all peer routers to be updated in a configuration interval of less than 12.5 seconds, if one set of changes occurs.

9.2 Sending Advertisement Intervals

When a router advertises Deterministic Fast RA Options in Router-to-Router messages, it MAY also indicate the frequency of its periodic Router Advertisements using Advertisement Interval Options in the message if there is room.

Alternatively, a router may advertise the interval in its Router Advertisement messages as defined in [\[4\]](#).

This option allows receiving routers to know how often to expect these Router Advertisements, so that they can check that the advertising router is active if expected packets are not received (as in [Section 8.3](#)).

Routers MAY send DetFRAOs occasionally in their periodic unsolicited Router Advertisements, in order to show hosts their FastRA configuration.

10. Sending Fast Router Advertisements

Fast Router Advertisements are sent in response to Router Solicitations. Where Deterministic Fast Router Advertisement Options have been exchanged, and the routers know their fixed delays, Router Advertisements are transmitted to the solicitor after delaying for

the Fixed Delay.

The number of FastRAs which may be sent at any time are determined by trading off the reasonable arrival rate of Router Solicitations, and the bandwidth and delay consumption caused by having all of these packets transmitted successively[6].

Determining good values for these outstanding FastRA bucket sizes is not well described and may require further work. The values defined in this document are approximations thought to be relatively harmless based on other protocol defaults, at the time of writing.

10.1 Sending Unicast Fast RAs

The same resource DoS protection mechanisms for unicast FastRAs used in the FastRA Draft are used in this document. Particularly, the maximum token bucket size is limited to MaxFastRAs, which defaults to MAXFASTRAS (10) [6].

The FastRA draft places up to MaxFastRAs tokens into the bucket each multicast Router Advertisement interval.

In order to provide more flexible replenishment of FastRA resources, a router MAY place unicast FastRA tokens into the bucket at a rate of one every FastRARefreshIval ms, where this defaults to FASTRAREFRESHIVAL.

10.2 Sending Multicast Fast RAs

Multicast Fast RAs are not supported in the original FastRA draft [6]. It does make sense for routers to provide fast multicast responses to Router Solicitations, where such solicitations do not create sufficient neighbour cache state to allow immediate unicast response.

Also, if a router has multiple Unicast FastRAs delayed before transmission, it may be possible to send a multicast FastRA instead at the earliest of the delay times, in order to reduce the number of responses required.

Multicast Fast RA transmission is managed separately to unicast transmission, and has a token bucket with a size of MaxMCFastRAs which defaults to MAXMCFASTRAS.

Tokens are placed into this bucket at a rate of one every MIN_DELAY_BETWEEN_RAS.

11. Interaction with Other Routers

Not all routers will understand Deterministic FastRA options or want to be in a Fast Router List. This section describes interactions between Fast Routers and other routers.

11.1 Presence of Legacy Routers

Deterministic Fast Routers ignore the presence of Legacy Routers when building their Fast Router List. Fast Routers rely upon these routers undertaking random delays according to IPv6 Neighbour Discovery[2].

11.2 Ceasing to be a Fast Router

A router which no longer wishes to undertake FastRA may stop making fast responses at any time, but SHOULD delay by a random value between 0 and MAX_RTR_STATUS_DELAY, before sending MAX_INITIAL_RTR_STATUS successive multicast Router-to-Router Status messages. These messages MUST contain the DetFRA option with Rank set to NOT_FAST_RTR_RANK and an advertised Fixed Delay of NOT_FAST_RTR_DELAY. These multicast Router Advertisements SHOULD be sent once every MIN_DELAY_BETWEEN_RTR_STATUS to the all-routers group.

While a router advertises its Fixed Delay as NOT_FAST_RTR_DELAY, it in fact reverts to the procedures defined in IPv6 Neighbour Discovery [2].

Routers receiving this option in a Router-to-Router message SHOULD remove the router from the Fast Routers List, and recalculate subsequent Ranks and delays of routers remaining in the list.

This may lead to peer routers' re-advertisement of new Ranks and delays as described in [Section 9.1](#).

11.3 Presence of Non Fast Routers

While a router may be able to understand and process DetFRAOs, it may not wish to be a Fast Router. Routers which have completed advertising their leaving of the Fast Routers List fall into this category.

These routers SHOULD act like legacy routers, and ignore Deterministic FastRA options as if they didn't understand them.

11.4 Presence of Incompatible Fast Routers

When routers wishing to undertake FastRA exist on the link but are not trusted for inclusion in the Fast Routers List, two distinct sets of routers may appear.

Each set may not trust another, and may instead have its own router at each Rank. In this case, the `IncompatibleFastRouters` flag SHOULD be set on the interface until the untrusted routers become trusted, or cease being Fast Routers.

Each Fast Router which has the `IncompatibleFastRouters` flag MUST attempt to inform its administrator about the misconfiguration.

12. Host interaction with DetFRAOs

While the principle use of Deterministic FastRA options is for router to router configuration, the options may be used to pass information to hosts.

12.1 Host Transmission of DetFRAOs

A host may transmit a Deterministic FastRA option in a Router Solicitation. This option MUST be sent with a Rank of `NOT_FAST_RTR_RANK` and a Fixed Delay of `NOT_FAST_RTR_DELAY`.

This option in a solicitation requests that the responding Router Advertisement contains a Deterministic FastRA option.

12.2 Reception of DetFRAOs in Router Solicitations

Routers receiving a DetFRAO option in a Router Solicitation sends a Router Advertisement responding to the solicitation if it is able to do so. In responding to the solicitation with this option, the Router Advertisement MAY contain the router's own configuration in a DetFRAO.

12.3 Transmission of DetFRAOs in Router Advertisements

Routers MAY inform hosts of their status by including their Router-to-Router status in some unsolicited Router Advertisements. If a router does this, it SHOULD inform them with updated options after configuration status change.

12.4 Host Reception of DetFRAOs

If a host which configures a router receives a DetFRAO from this router in a Router Advertisement, it can determine the routers' Ranks

and delays.

If when receiving a Router Advertisement in response to its solicitation it receives a DetFRA0 with a Rank which matches an existing configured router, it may suspect change has occurred.

Where no prefixes exist in common with the previously received advertisements, reception of a FastRA from a router within the reception interval defined by another router either implies a new router joining the link, or a link change.

Since new routers joining a Fast Router Advertisement set are unlikely to occur often, a host MAY assume that the new router is on a new link, and begin configuration operations.

Where a single router has joined a link, the next best Fast Router will then be the previously configured router, if the host remains on the same link. Therefore, a host MAY delay until after the sum of the old and newly received routers' Separation have elapsed before concluding that link change has occurred. This additional cost is likely to be less than that incurred with Legacy Router Advertisements, but provides slightly more safety than the previous heuristic.

13. Configuration Parameters

The following parameters are defined in this document:

FastRRelPref	<p>A parameter which controls the ranking of Fast Routers. It sets the advertised Rel Pref field in the DetFRA0. Setting this value lower betters the preference of the router.</p> <p>Minimum: 0 Maximum: 255 Default: DEF_REL_PREF</p>
FastRASep	<p>A parameter controlling the delay between scheduling of Fast Router Advertisements on adjacent routers in the Fast Routers List. (Units: milliseconds)</p> <p>Minimum: 1 Maximum: 255 Default: DEF_SEP</p>
MaxFastRAs	<p>The maximum bucket size for Unicast FastRAs</p>
[6]	<p>Minimum: 0 (not Fast RA) Default: MAXFASTRAS</p>
MaxMCFastRAs	<p>The maximum bucket size for Multicast FastRAs.</p> <p>Minimum: 0 (no Multicast Fast RA) Default: MAXMCFASTRAS</p>

[14.](#) Protocol Constants

BOOTSTRAP_DELAY	500ms
BOOTSTRAP_RANK	254
DEF_SEP	50ms
DEF_REL_PREF	127
MAXMCFASTRAS	5
MAX_INITIAL_RTR_STATUS	3
MAX_RTR_STATUS_REQ_DELAY	1000ms
MAX_RTR_STATUS_REQS	3
MAX_RTR_STATUS_DELAY	500ms
MIN_DELAY_BETWEEN_RTR_STATUS	3 seconds
NOT_FAST_RTR_DELAY	500ms
NOT_FAST_RTR_RANK	255
RTR_STATUS_REQ_INTERVAL	4 seconds
UCASTFASTRAREFRESHIVAL	400ms

15. IANA Considerations

The new ICMPv6 message type - Router to Router (with two codes) and a new ICMPv6 'Deterministic Fast Router Advertisement Option' are defined in this document.

The Router-to-Router message is suggested to be an Informational ICMPv6 message and is defined in [Section 4.1](#).

(DetFRA0) is defined in [Section 4.2](#) of this document. It is suggested that the option number 13 is used for the Type of this option.

16. Security Considerations

The Router-to-Router message's use of the Neighbour Discovery option format allows SEND options to be used to check the authenticity of the messages sent from the peer router.

The authorization of a node to be a router is described by a delegation chain advertised in a succession of Delegation Chain Advertisement messages. When using SEND, routers MUST check that the router from which they receive a DetFRA0 is an authorized router from that link, and that the router trusts the delegation service used to sign this authorization [[5](#)].

Where the Router-to-Router message is not trusted, but contains a Deterministic FastRA Option, the router MUST NOT include the router in its Fast Router List, but SHOULD set the IncompatibleFastRouters flag on that interface. This will turn attempt to inform the administrator of the configuration problem.

Where disjoint sets of routers each undertake FastRA (but don't trust each other), they can choose the same delay timers for sending FastRA. In the worst case this means that every FastRA sent will collide with another RA. Administrative action is currently required to fix this issue, but further work may establish if automated robustness to this issue is desirable.

17. Acknowledgments

This work is based on and expands the FastRA internet-draft [6].

18. References

18.1 Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", [RFC 2461](#), December 1998.
- [3] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 2463](#), December 1998.
- [4] Johnson, D., Perkins, C. and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [5] Arkko, J., Kempf, J., Sommerfeld, B., Zill, B. and P. Nikander, "SEcure Neighbor Discovery (SEND)", [draft-ietf-send-ndopt-05](#) (work in progress), April 2004.
- [6] Kempf, J., Khalil, M. and B. Pentland, "IPv6 Fast Router Advertisement", [draft-mkhalil-ipv6-fastra-02](#) (work in progress), October 2002.

18.2 Informative References

- [7] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", [RFC 3513](#), April 2003.
- [8] "IEEE standard for local and metropolitan area networks - common specifications - Media access control (MAC) Bridges", ISO/IEC IEEE Std 802.1d, 1998.

Authors' Addresses

Greg Daley
Centre for Telecommunications and Information Engineering
Department of Electrical and Computer Systems Engineering
Monash University
Clayton, Victoria 3800
Australia

Phone: +61 3 9905 4655
EMail: greg.daley@eng.monash.edu.au

Brett Pentland
Centre for Telecommunications and Information Engineering
Department of Electrical and Computer Systems Engineering
Monash University
Clayton, Victoria 3800
Australia

Phone: +61 3 9905 5245
EMail: brett.pentland@eng.monash.edu.au

Erik Nordmark
Sun Microsystems, Inc.
17 Network Circle
Mountain View, CA
USA

Phone: +1 650 786 2921
EMail: erik.nordmark@sun.com

[Appendix A](#). State Diagram

Below is a state diagram for Fast Router Advertisement.

Legend:

FD : Fixed Delay

TmrSet : Set Timer value(millisecons)

DetFRA0: Deterministic Fast RA Option

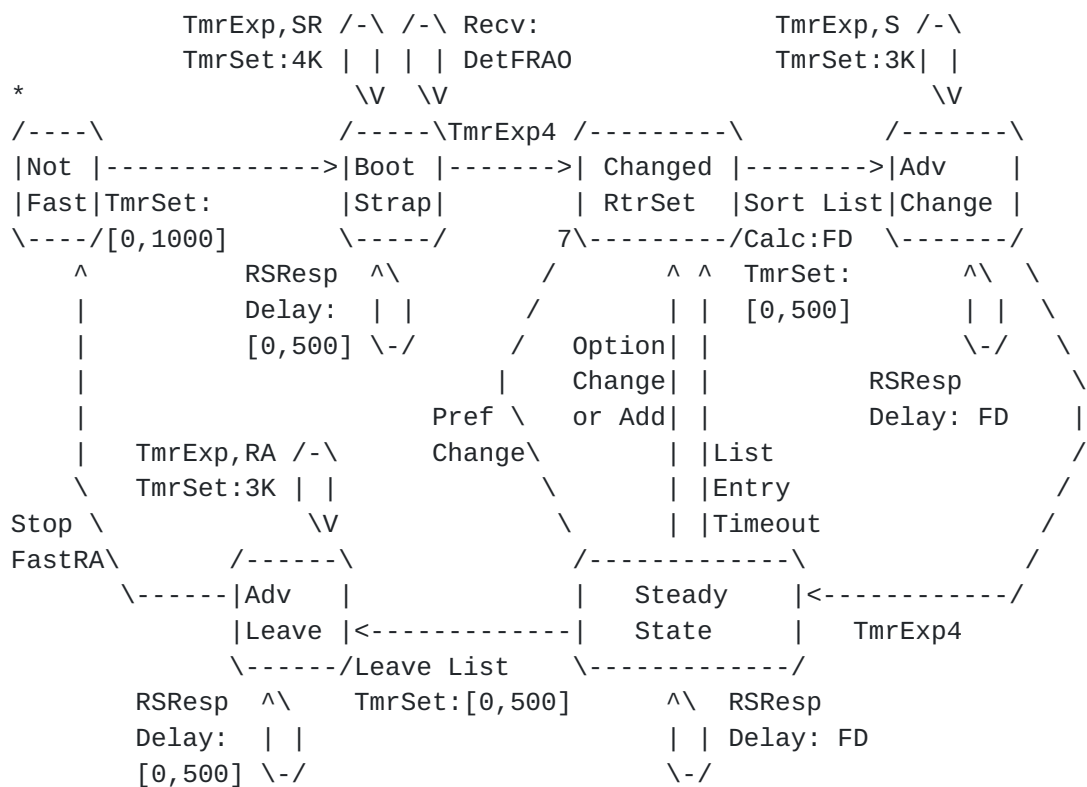
RSResp : Response to Router Solicitation

S : Router-to-Router Status

SR: : Router-to-Router Status-Request

TmrExp : A timer expiry

TmrExp4: Fourth Timer Expiry(resets counter)



Intellectual Property Statement

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 procedures with respect to rights in RFC 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.

Disclaimer of Validity

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.

Copyright Statement

Copyright (C) The Internet Society (2004). 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.

Acknowledgment

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

