



Border Router Discovery Protocol (BRDP) Based Routing

Exit routing for multi-homed networks



Teco Boot / 19 November 2008

Mobile Ad-hoc Networks (manet)



History and status

- ❖ Finding solution for scalable MANETs by using a backbone
- ❖ Connected MANETs have their problems:
 - Addressing (Autoconf)
 - Multi-homing (Autoconf & [MANET](#) & other WG)
- ❖ Autoconf I-D:
 - Current: [draft-boot-autoconf-brdp-01.txt](#)
 - Autoconf out-of-scope topics left out
- ❖ I-D on routing:
 - Current: [draft-boot-brdp-based-routing-01.txt](#)
 - Fixes problems with multi-homing and ingress filters
- ❖ More coming:
 - Source Address selection (related to RFC3484)
 - Tiny exit-routing IGP using routing headers

My illustrations are simplified




- Border Router is a router connecting the MANET to an ISP
- In practice, there is a Provider Edge (PE) and Customer Premises Equipment (CPE) router
- CPE – PE protocols are out of scope for this presentation
- PE may (should!) have ingress filter




My illustrations are simplified


- ISPs have their networks. Details on this are not important for BRDP
- Note the colors: Provider A is red, provider B is blue



Single-homed MANET: no problems 😊




Multi-homed MANET




Here, BRDP is born.
So far, so good. *But:*

In a MANET, nodes can move!




- MANET routing protocol converges to reflect the movement
- LFN is not aware of movement

Solution #1: Direct traffic to BR that owns SA prefix




Feature of BRDP Based Routing!

Solution #2: Adjust BR-b ingress filter for prefix BR-a



Not sure ISP is cooperative, so this is not the goal


Solution #3: Select SA that corresponds to BR used




LFN2 would use “metrics” of BRs

Session continuity problems,
doesn't work in a MANET
because of movements

Solution #4: Address translation on BRs



Goal: support for multi-homed edge networks with multi-homed nodes



Feature of BRDP and BRDP Based Routing!

Benefits depend on transport layer or application layer adjustments

Make-before-break!


Download your files twice as fast!

How does it work?

Border Router Discovery Protocol (BRDP):

- **Provides valid prefix information to interior routers**
- **Distributes Border Router Information Option (BRIO) via Neighbor Discovery Router Advertisements.**
- **BRIOs are distributed in the MANET**
- **BRIOs provide metrics for paths from and to DFZ (Default-Free Zone)**
- **BRDP uses link metrics from routing protocol or other source**
 - **Use link metrics for both directions**
- **BRIOs provide DHCP relay information**
- **BRDP is a Distance-Vector protocol**

BRIO flooding using Neighbor Discovery Router Advertisements




- Routers already send RA
- BRIOs piggy-backed on RA

Results:

1. All routers learn prefixes for BRs, with metric information
2. All routers learn DHCP addresses for relay or prefix delegation


BRIO flooding with scoping (needed in large MANETs)



- Reduction of forwarded BRIO set is acceptable, as only the best paths would be used

BRIO format

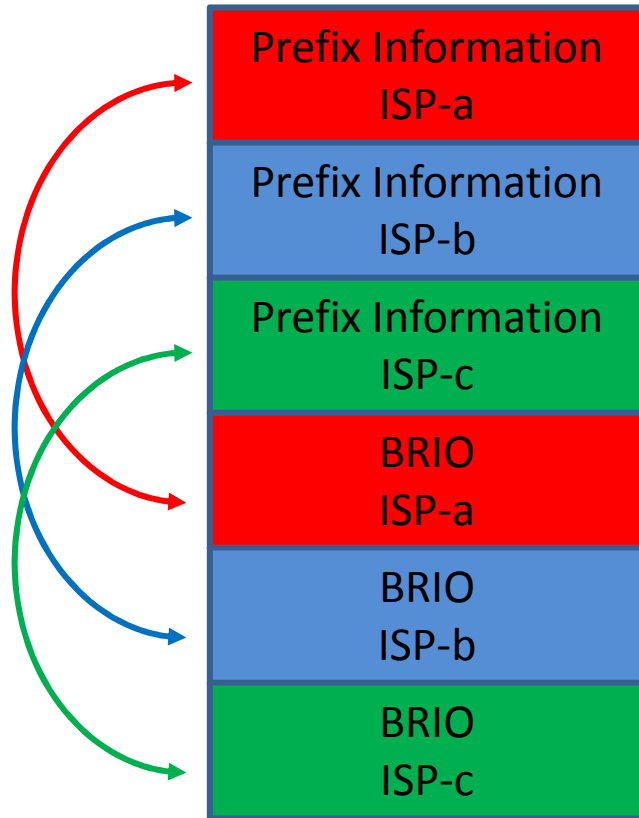
D=1: BR is DHCP server or relay



BR address and valid prefix for this BR

Metric for this prefix
Is a bidirectional metric between
DFZ and this BR

RA in MANET with 3 ISP uplinks



OK, the RA packet size increases.

Directing packets to BR that correspond to source addresses

If an end host selected an upstream ISP, how can the *routing system* direct traffic to this ISP ?

- Packets must have an earmark for indicating the to-be-used Border Router
- Source Address or additional tag (routing header, encapsulation)
- BRDP Based Routing:
Edge networks, forwarding based on FIB and SA / BRIO cache

Legacy Internet Routing

RFC1812; Section 5.2.1.2 Unicast:

- (5) The forwarder determines the next hop IP address for the packet, **usually** by looking up the **packet's destination** in the router's routing table. This procedure is described in more detail in Section [5.2.4]. This procedure also decides which network interface should be used to send the packet.


RFC1812; 5.2.4.3 Next Hop Address

- (5) Default Route: This is a **route to all networks** for which there are **no explicit routes**. It is by definition the route whose prefix length is zero
- If a default route exists, legacy routing is used.

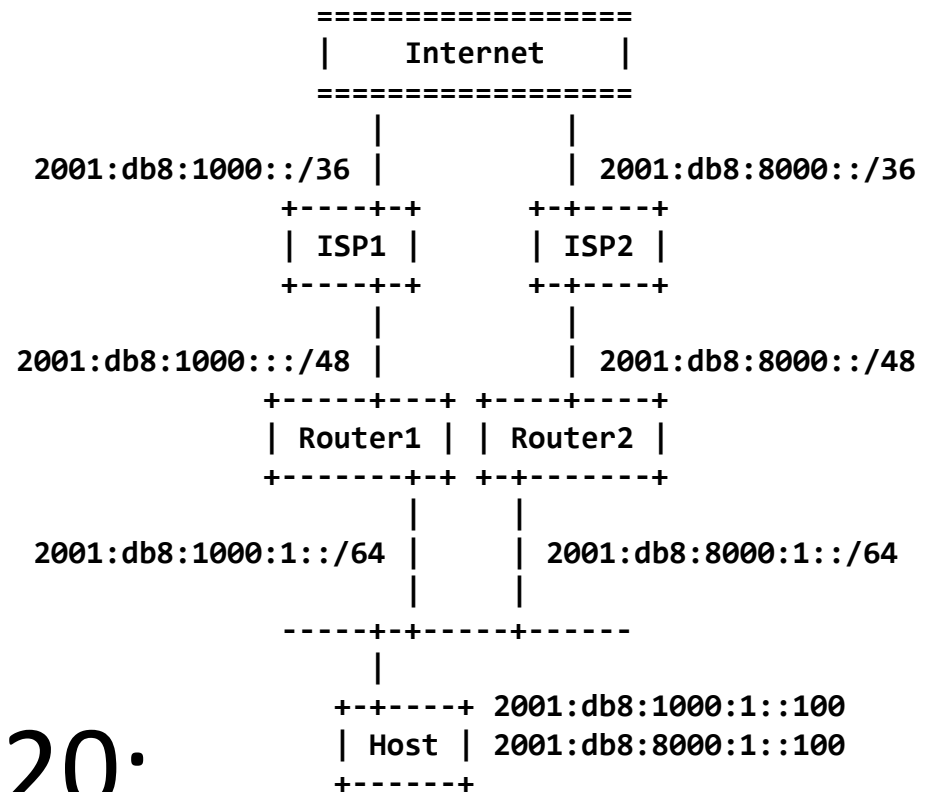
BRDP Based Routing

- With BRDP Based Routing, the default route is removed.
- New heuristic for finding a next hop, only used when first FIB lookup didn't find a nexthop:

(6) BRDP Route: This is a **route to all networks** for which there are **no explicit routes**, and a **default route is not used**. The nexthop IP address is found by means of a Border Router Information Cache (**BRIO-Cache**) **lookup** based on the source address and, if a matching BRIO-Cache entry is found, a subsequent **FIB lookup** based on the selected **Border Router address**.




← My example



RFC 5220:

Figure 1



Example


FIB R3:

<code>2001:DB8:A::/64</code>	-> FE80::1	# BR-a
<code>2001:DB8:A:1::/64</code>	-> FE80::1	
<code>2001:DB8:A:2::/64</code>	-> FE80::2	
<code>2001:DB8:A:3::/64</code>	-> local	
<code>2001:DB8:B::/64</code>	-> FE80::2	# BR-b
<code>2001:DB8:B:1::/64</code>	-> FE80::1	
<code>2001:DB8:B:2::/64</code>	-> FE80::2	
<code>2001:DB8:B:3::/64</code>	-> local	
<code>2001:DB8:B:4::/64</code>	-> FE80::2	

BRIO Cache R3:

<code>2001:DB8:A::A/48</code>	<= FE80::1 metric 100
<code>2001:DB8:A::A/48</code>	<= FE80::2 metric 250
<code>2001:DB8:B::B/48</code>	<= FE80::2 metric 150
<code>2001:DB8:B::B/48</code>	<= FE80::1 metric 200

Example - FIB



All internal routes are found in FIB
 Forwarding based on DA
 No default route !!


FIB R3:

2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200

Example BRIO-Cache



Border Router information,
with sender info and metrics

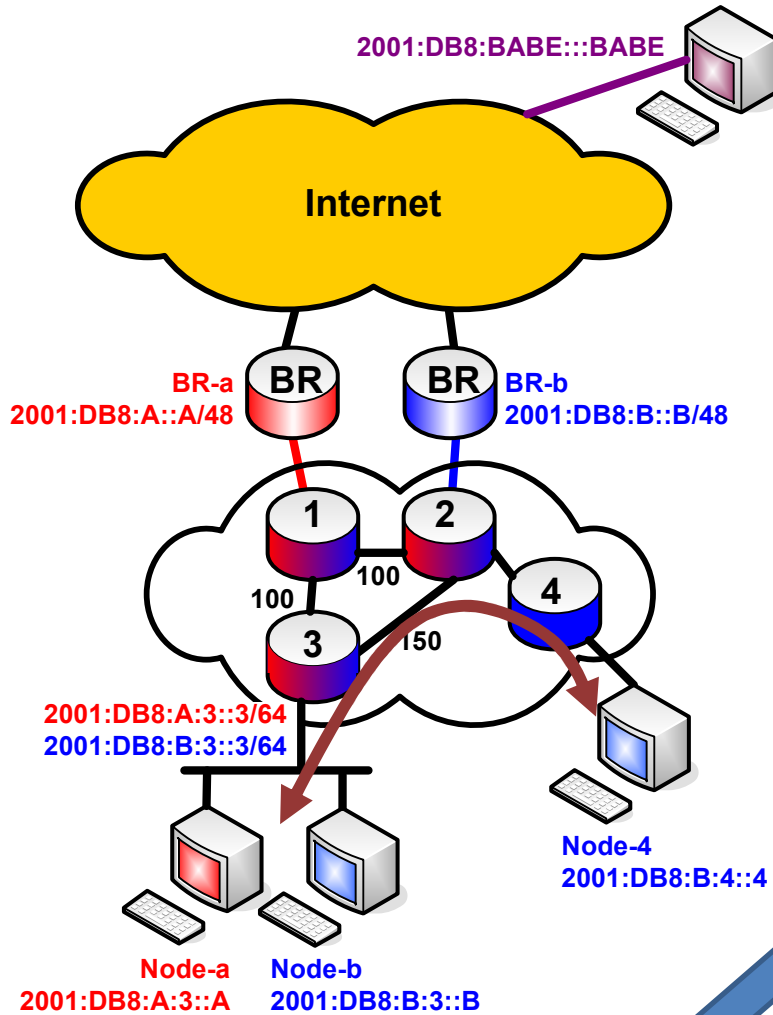


FIB R3:

2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200



Node-a or Node-b sends to Node-4

FIB R3:


2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200

1) 2001:DB8:B:4::4 in FIB?

Yes



Node-a sends to Babe

FIB R3:


2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200

1) 2001:DB8:BABE::BABE in FIB?

NO



Node-a sends to Babe

FIB R3:


2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200

2) 2001:DB8:A:3::A match in BRIO Cache?

YES: 2001:DB8:A::A



Node-a sends to Babe

FIB R3:

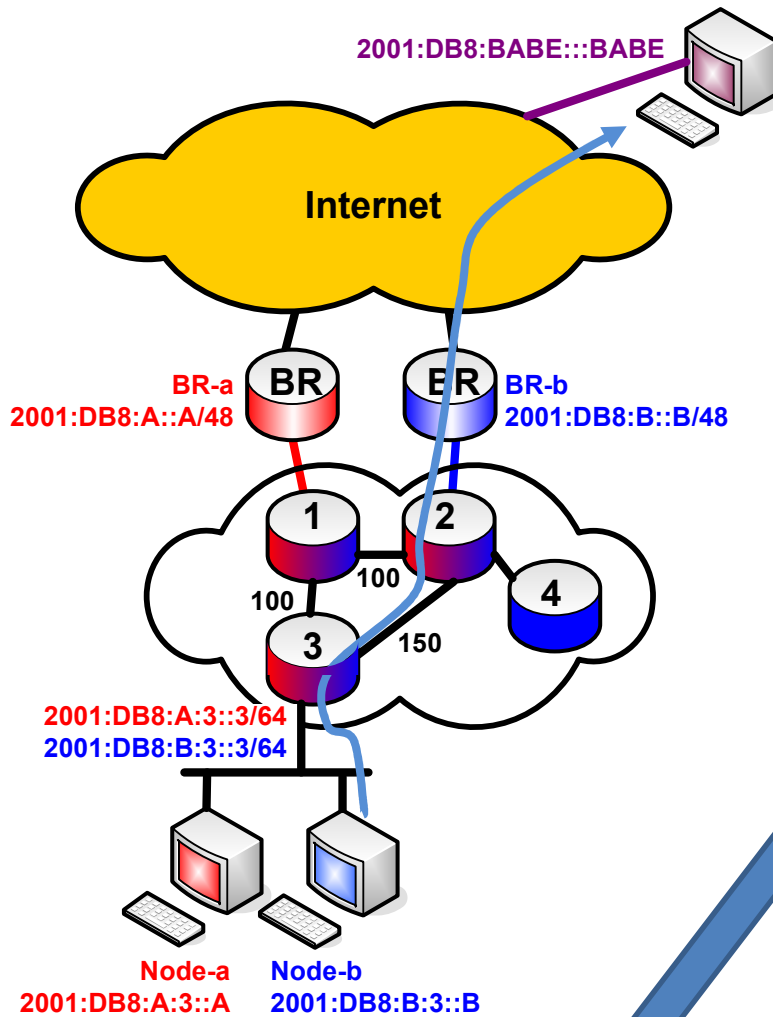
2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200

3) 2001:DB8:A::A in FIB?

YES: 2001:DB8:A::/64 -> FE80::1



Node-b sends to Babe

FIB R3:


2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200

- 1: 2001:DB8:BABE::BABE in FIB: **NO**
- 2: 2001:DB8:B:3::B match in BRIO Cache :
Yes, 2001:DB8:B::B
- 3: 2001:DB8:B::B in FIB: Yes, FE80::2

Bad sends to Babe



1: 2001:DB8:BABE::BABE in FIB: **NO**

2: 2001:DB8:BAD::BAD match in BRIO Cache : **NO**

3: **DROP**

FIB R3:

2001:DB8:A::/64	-> FE80::1	# BR-a
2001:DB8:A:1::/64	-> FE80::1	
2001:DB8:A:2::/64	-> FE80::2	
2001:DB8:A:3::/64	-> local	
2001:DB8:B::/64	-> FE80::2	# BR-b
2001:DB8:B:1::/64	-> FE80::1	
2001:DB8:B:2::/64	-> FE80::2	
2001:DB8:B:3::/64	-> local	
2001:DB8:B:4::/64	-> FE80::2	

BRIO Cache R3:

2001:DB8:A::A/48	<= FE80::1 metric 100
2001:DB8:A::A/48	<= FE80::2 metric 250
2001:DB8:B::B/48	<= FE80::2 metric 150
2001:DB8:B::B/48	<= FE80::1 metric 200

BRDP Based Routing

Introduction

Analysis of the problem

How does it work?

Next steps

Next steps

- Support for IPv4
- Finish BRDP implementation (Linux, Opnet)
- Implement BRDP Based Routing (Linux, Opnet)
- Continue research on Border Router Routing Header
- Website:
<http://www.inf-net.nl/brdp.html>
- Any help is welcome !
- Any comment is welcome !

Thanks for your attention !