

BSR issues:
Maintaining group-to-RP mappings
IPv6 and scopes

Stig Venaas
venaas@uninett.no

Overview

- Slides at <http://domen.uninett.no/~venaas/bsrietf61.pdf>
- Basically three approaches for how to withdraw RPs from a group range
- IWNH
 - Implicit by not including RP in list
 - Holdtimes ignored
- EW
 - Expiry using holdtimes
 - Explicit by using holdtime of 0 for the RP
- IW
 - Combination of the two above
 - Implicit but also using holdtimes
 - Holdtime of 0 also removes the RP

Objective

- We believe the BSR specs are ambiguous
- Previous specs seems to be IW or IWNH, but not clear to us
- In draft version 04 we tried to specify EW
- Implementations of IW and EW; any of IWNH?
- Want to get consensus for which to specify
- What criteria for choosing?
 - Simplicity, backwards compatibility, others?

Comparing simplicity

- In document sent to mailing list we tried to do this by comparing how they cope with different events
- Fragmented BSMs, lost fragments
- RP disappearing and possibly replaced
- E-BSR disappearing, possibly replaced
- Change in group ranges mapping to RP
- Not enough time for all details here

Network partitioned, new BSR 1/2

- Assume that some RPs are now unreachable. Should be removed ASAP and new C-RPs should be used
- How can new E-BSR know which RPs to remove and which are present
 - Wait at least 60s to get C-RP advs and assume those missing gone? How to know which are missing? Or should old RP mappings stay on routers until e.g. holdtime expiry

Network partitioned, new BSR 2/2

- Should new BSR use information from previous BSR?
- Some way to expire RPs and group ranges is needed. Holdtime is one option
- How to make sure live RPs are not lost during switchover. Takes longer than default holdtimes
- Special treatment of empty BSM from new BSR? Send C-RP-adv immediately? Make routers keep old mappings longer; at least until non-empty BSM received?

Backwards compatibility

- Depends on which we choose
- If choose EW
 - IW routers ok with EW E-BSR
 - IWNH not. Can possibly send additional BSMs omitting the RP to remove. Kludge
- If choose IW
- EW routers ok if IW E-BSR uses holdtime 0, not if omitting RPs to be removed
- IWNH works with latter, not former

Making a choice

- How do we proceed making a choice?
- We have at least three alternatives. Which are worth exploring further?
- Major effort tweaking each of the approaches to satisfy all known issues
- What do current implementations do?
- Which alternative do you prefer and why?

IPv6 and scopes 1/2

- Scopes 1, 2, 5, 8, E etc.
- For e.g. scope 5 we have FFx5::/16
- 16 possible ranges for each scope
 - FF05::/16, FF15::/16, FF25::/16, ...
- A BSM can contain an admin scope group range, but only a single range
- Don't want to send 16 BSMs for each scope value

IPv6 and scopes 2/2

- In effect we also have separate BSR election for each of the ranges
- Is a simple 4 bit scope identifier sufficient?
 - How to encode?
- Encoding swapping flag and scope nibbles?
- Allow an admin scope set consisting of union of ranges? Several scope ranges in one BSM
- Possible issues with arbitrary mix of unions of ranges and partly overlapping admin scope sets