

Standardized back-off SPF delay algorithm

draft-decraene-rtgwg-backoff-
algo-00

Bruno Decraene

Orange

Recap 1: Distributed SPF

- Distributed SPF computations require that all routers:
 - perform the same computation (SPF)
 - at the same time (T0)
 - using the same data (Link State DataBase)
- Otherwise RIB are not consistent and micro forwarding loops may occur.
 - cf <http://tools.ietf.org/html/draft-litkowski-rtgwg-spf-uloop-pb-statement-00>

Recap 2: Back off SPF delay

- Goal 1: first SPF should be quick for Fast Convergence
 - SPF delay should be low for the first event (e.g. 10ms)
- Goal 2: if the IGP gets unstable, favor stability over speed
 - Playing safe, avoid multiplying micro-loops
 - SPF delay should be high for multiple consecutive events/SPF (e.g. 2s)
- Between those 2 extremes, lots of free space and different vendors have proposed different algorithm.

Back off SPF delay is not a local consideration

- Back off SPF algorithm/delay is **not** a local consideration.
- Using different algorithm/delay breaks the distributed SPF requirements
 - compute RIB at the same time (T0)
 - using the same data (Link State DataBase)
 - between T0(R1) and T0(R2) the topology may change (LSP/LSA).
Especially since we are in an unstable situation.

Example

IGP Events
LSP/LSA



Node A
SPF

Algo: N
rapid run



Node B
SPF

Algo:
exponential
backof



Delta time



Inconsistent RIBs
using different LSDB
(LSP4 used by B but not by A)

How often - how long

- How often: roughly each time the SPF algo delay is used
- How long: hundred of ms up to seconds → definitely significant

Iterations	N Rapid Run	Exponential	Delta
1	50	50	0
2	50	150	-100
3	50	250	-200
4	2000	450	1550
5	2000	850	1150
6	2000	1650	350
7	2000	2000	0

Iterations	N Rapid Run	Exponential	Delta
1	50	50	0
2	50	150	-100
3	50	250	-200
4	50	450	-400
5	50	850	-800
6	4000	1650	2350
7	4000	3250	750
8	4000	4000	0

	Rapid Run	Exponential
Initial	50	50
N	3	
Increment		50
Max	2000	2000

	Rapid Run	Exponential
Initial	50	50
N	5	
Increment		50
Max	4000	4000

Very easy to solve

- We only need to standardize one algorithm.
- Vendors would still be free to implement a proprietary algo.
- IETF would still be free to specify another algo in some future.

Micro-loops

- **Does not remove all micro-loops**, because others factors/delta are involved.
 - mostly different time to update the FIB
- **Do**
 - solves the SPF delay factor
 - avoid adding micro-loops for no reason.

Summary & next step

- draft
 1. calls for a standardized SPF back-off algorithm, for interoperability purpose
 2. proposes an algorithm

- Calling for feedback
 - Short term: can we agree on “1” ?
 - Medium term: discuss the choice of the algorithm
 - v00 proposes the exponential algo for the only reason that this is the most implemented/deployed.
 - Choice of the algorithm should be discussed on the mailing list.

Thank you