SPF back-off delay

draft-ietf-rtgwg-spf-uloop-pb-statement-02

Stéphane Litkowski	Orange Business Service
Bruno Decraene	Orange
Martin Horneffer	Deutsche Telekom

draft-ietf-rtgwg-backoff-algo-01

Bruno Decraene	Orange
Stéphane Litkowski	Orange Business Service
Hannes Gredler	RtBrick
Acee Lindem	Cisco
Pierre François	
Chris Bowers	Juniper Networks

draft-ietf-rtgwg-spf-uloop-pb-statement-02

- Problem statement and problem analysis
- Distributed SPF computations requires that all routers:
  a. Perform the same computation (SPF)
  b. At the same time (t0)
  c. Using the same topology (Link State Database)
- Otherwise RIB are not consistent and micro forwarding loops may occur.
- Unfortunately, SPF delay algorithm has not been standardized and different vendors have different ones.

draft-ietf-rtgwg-backoff-algo

- Defines one standard SPF delay algorithm
- 3 draft revisions since IETF 93: -01 ␣-04
- Changes:
  - SPF delay algorithm re-specified using a Finite State Machine
  - New co-author: Chris Bowers
  - Many editorial improvements
    - Grammar, consistency in timer value and state names

Finite State Machine: 3 states / 3 delay

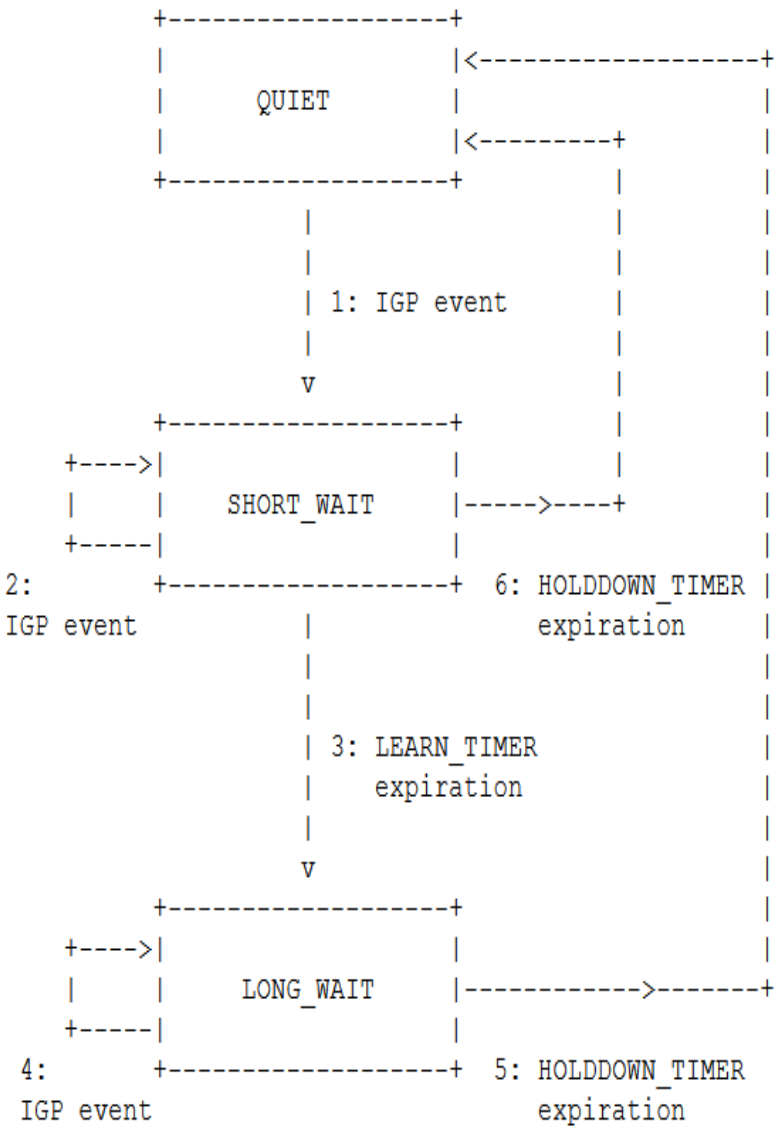

- QUIET state:
  - SPF delay:= INITIAL_SPF_DELAY
  - initial state / no IGP event since HOLDDOWN_INTERVAL
  - configured very low to react ASAP to the first event.
- SHORT_WAIT state:
  - SPF delay:= SHORT_SPF_DELAY
  - handle subsequent IGP events
  - designed for fast convergence, preserving stability
- LONG_WAIT state:
  - SPF delay set:= LONG_SPF_DELAY
  - designed for IGP unstable condition
    - set of IGP events for more than LEARN_TIMER

Implementations status

- 3 independent IS-IS implementations:
  - JunOS prototype
  - Quagga (Free Range Routing)
  - RtBrick
- 2 implementations tested: JunOS, Quagga
  - Tested all FSM states and transitions
  - IS-IS level 2

Tests results

Test case				Theorethical			JunOS measured		Quagga measured					
Computed					*Computed*	*Computed*					*Computed*			
FSM State before event	**event #**	**Wait after this event**	**FSM event tested (from current to next state)**	**FSM State after event**	***FSM Timer***	***Next SPF***	**State**	***Next SPF***	**Event**	**Next SPF**	***Delay***	**LSP version**	**Old State**	**New State**
Quiet	1	12,000	1 & 6	Quiet	*00,100*	*00,100*	Quiet	*00,100*	15:48:34,520	15:48:34,620	*00,100*	0F	Quiet	Short, Long, Quiet
Quiet	2	02,000	1	Short	*00,100*	*00,100*	Quiet	*00,100*	15:48:46,735	15:48:46,835	*00,100*	10	Quiet	Short
Short	3	02,000	2	Short	*00,200*	*00,200*	Short	*00,200*	15:48:48,748	15:48:48,948	*00,200*	11	Short	Short
Short	4	12,000	6	Quiet	*00,200*	*00,200*	Short	*00,200*	15:48:50,756	15:48:50,956	*00,200*	12	Short	Short, Long, Quiet
Quiet	5	06,000	1 & 3	Long	*00,100*	*00,100*	Quiet	*00,100*	15:49:02,754	15:49:02,854	*00,100*	13	Quiet	Short, Long
Long	6	05,000	4	Long	*10,000*	*10,000*	Long	*10,000*	15:49:08,754	15:49:18,754	*10,000*	14	Long	Long
Long	7	05,000	4	Long	*10,000*	*05,000*	Long	*05,000*	15:49:13,000	15:49:18,754	*05,754*	15	*Long*	*Long*
Long	8	12,000	5	Quiet	*10,000*	*10,000*	Long	*10,000*	15:49:18,759	15:49:28,759	*10,000*	16	Long	Long, Quiet
Quiet	9		1		*00,100*	*00,100*	Quiet	*00,100*	15:49:30,760	15:49:30,860	*00,100*	17	Quiet	Short, Long, Quiet

FSM state	delay (s)	JunOS	Qugga
Long	10,000	long-spf-delay	long-delay
Quiet	00,100	initial-spf-delay	init-delay
Short	00,200	short-spf-delay	short-delay
	10,000		
	05,000	time-to-learn-interval	time-to_learn
	10,000	holddown-interval	holddown

Focus on test 1: show logs

JunOS:

```
[... event 1]
Sep 28 17:22:47.371917 L2 Hi-Prio SPF trigger: Updated LSP ca1a.00-00
Sep 28 17:22:47.371937 task_timer_uset: timer IS-IS_learn_timer <Touched> set to offset 5 at 17:22:52.371917
Sep 28 17:22:47.371955 task_timer_uset: timer IS-IS_holddown_timer <Touched> set to offset 10 at 17:22:57.371937
Sep 28 17:22:47.371973 task_timer_uset: timer IS-IS_spf_timer <Touched SpawnJob> set to offset 0.100000 at
17:22:47.471955
Sep 28 17:22:47.371986 Shared SPF trigger sent to FSM isis L2 shared fsm: old state=quiet, new state=short-wait
Sep 28 17:22:47.371999 New SPF scheduled in 0.100000s
```

Quagga:

```
==·Event·1¶
2016/10/03·17:48:34·ISIS:·ISIS-Upd·(1):·Rcvd·L2·LSP·on·eth3,·cirType·L2,·cirID·3¶
2016/10/03·17:48:34·ISIS:·ISIS-Upd·(1):·Rcvd·L2·LSP·0000.2a00.0003.00-00,·seq·0x0000000f,·cksum·0x5b98,·lifetime·65498s,·len·76,·on·eth3¶
2016/10/03·17:48:34·ISIS:·ISIS-Snp·(1):·Compare·LSP·0000.2a00.0003.00-00·seq·0x0000000f,·cksum·0x5b98,·lifetime·65498s¶
2016/10/03·17:48:34·ISIS:·ISIS-Snp·(1):·····is·newer·than·ours·seq·0x0000000e,·cksum·0x12ce,·lifetime·64889s¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·SPF·F0·schedule·called·using·IETF·algorithm,·ttl·thread·:·(nil),·holddown·thread·:·(nil)¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·:·disabling·timetolearn·(thread·(nil))¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·launch·TIMETOLEARN·timer¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·:·disabling·holddown·(thread·(nil))¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·launch·HOLDDOWN·timer¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·disabling·SPF·timer·(nil)¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·move·from·state·0·(QUIET)·to·state·1·(SHORT_WAIT)¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·:·next·SPF·scheduled·at·GMT·15:48:34.620570·(wait·time·:·100,·current·time·:·GMT·15:48:34.520570)¶
2016/10/03·17:48:34·ISIS:·ISIS-Spf·(1)·L2·SPF·needed,·periodic·SPF¶
2016/10/03·17:48:39·ISIS:·ISIS-Spf·(1)·L2·IETF·SPF·delay·TIMETOLEARN·expired·moved·to·state·2·(LONG_WAIT)·family·0¶
2016/10/03·17:48:44·ISIS:·ISIS-Spf·(1)·L2·F0·IETF·SPF·delay·:·disabling·timetolearn·(thread·(nil))¶
2016/10/03·17:48:44·ISIS:·ISIS-Spf·(1)·L2·IETF·SPF·delay·HOLDDOWN·expired·moved·to·state·0·(QUIET)·family·0¶
¶
```

Summary & next steps

- Simple solution to reduce micro-loops duration & probability
- 3 implementations so far
- Ran tests on 2 implementations, covering all FSM states and transitions:
  - Conformant with the IETF draft
  - Same results
- Asking for WG Last Call:
  - draft-ietf-rtgwg-spf-uloop-pb-statement
  - draft-ietf-rtgwg-backoff-algo

Thank you