

Analysis of paths selection modes for Add-Paths

draft-vvds-add-paths-analysis-00

Virginie van den Schrieck
Olivier Bonaventure
Pierre François

Agenda

- Quick reminder
- Rule out some path selection modes ?
- Next

Add-Paths

- How to advertise multiple paths over a single iBGP session

Goal

- draft-ietf-idr-add-paths-01.txt
 - Encode multiple paths to same NLRI over a BGP session
 - Doesn't tell which paths to select for advertisement
 - Different applications lead to different selection modes
 - (maybe) No interop issues when different path selection modes are applied by speakers in an AS

Goal

- Informational draft
 - List the modes in one draft
 - Evaluate them (what it gives, what it costs)
 - Track potential interop issues

Modes

- All paths
- N paths (max)
- AS-Wide best paths
- Neighbor-AS group best paths
- Best Loc Pref / Second best Loc Pref paths
- Decisive step -I
- (put your favorite mode here)

	Path optimality	Backup availability / optimality	Control plane load and stress	DP Complexity	MED osc. avoidance
All	OK	OK	Max	EASIEST	OK
N	good	OK / ~OK	Bounded	Hard or Okay (MED)	good
LPI/LP2	OK	OK	~MAX	EASIER	OK
Decisive-I	OK	OK	~MAX	Easy but "complex"	OK
AS-Wide	OK	KO / ~OK	~MAX	EASY	OK
Group best	KO ...	KO	~MAX	?	OK

	Path optimality	Backup availability / optimality	Control plane load and stress	DP Complexity	MED osc. avoidance
All	OK	OK	Max	EASIEST	OK
N	good	OK / ~OK	Bounded	Hard or Okay (MED)	good
LPI/LP2	OK	OK	~MAX	EASIER	OK
Decisive-I	OK	OK	~MAX	Easy but "complex"	OK
AS-Wide	OK	KO / ~OK	~MAX	EASY	OK
Group best	KO ...	KO	~MAX	?	OK

	Path optimality	Backup availability / optimality	Control plane load and stress	DP Complexity	MED osc. avoidance
All	OK	OK	Max	EASIEST	OK
N	good	OK / ~OK	Bounded	Hard or Okay (MED)	good
LPI/LP2	OK	OK	~MAX	EASIER	OK
Decisive-I	OK	OK	~MAX	Easy but “complex”	OK
AS-Wide	OK	KO / ~OK	~MAX	EASY	OK
Group best	KO ...	KO	~MAX	?	OK

	Path optimality	Backup availability / optimality	Control plane load and stress	DP Complexity	MED osc. avoidance
All	OK	OK	Max	EASIEST	OK
N	good	OK / ~OK	Bounded	Hard or Okay (MED)	good
LPI/LP2	OK	OK	~MAX	EASIER	OK
Decisive-I	OK	OK	~MAX	Easy but "complex"	OK
AS-Wide	OK	KO / ~OK	~MAX	EASY	OK
Group best	KO ...	KO	~MAX	?	OK

	Path optimality	Backup availability / optimality	Control plane load and stress	DP Complexity	MED osc. avoidance
All	OK	OK	Max	EASIEST	OK
N	good	OK / ~OK	Bounded	Hard or Okay (MED)	good
LPI/LP2	OK	OK	~MAX	EASIER	OK
Decisive-I	OK	OK	~MAX	Easy but "complex"	OK
AS-Wide	OK	KO / ~OK	~MAX	EASY	OK
Group best	KO ...	KO	~MAX	?	OK

	Path optimality	Backup availability / optimality	Control plane load and stress	DP Complexity	MED osc. avoidance
All	OK	OK	Max	EASIEST	OK
N	good	OK / ~OK	Bounded	Hard or Okay (MED)	good
LPI/LP2	OK	OK	~MAX	EASIER	OK
Decisive-I	OK	OK	~MAX	Easy but "complex"	OK
AS-Wide	OK	KO / ~OK	~MAX	EASY	OK
Group best	KO	KO	~MAX	?	OK

	Path optimality	Backup availability / optimality	Control plane load and stress	DP Complexity	MED osc. avoidance
All	OK	OK	Max	EASIEST	OK
N	good	OK / ~OK	Bounded	Hard or Okay (MED)	good
LPI/LP2	OK	OK	~MAX	EASIER	OK
Decisive-I	OK	OK	~MAX	Easy but "complex"	OK
AS-Wide	OK	KO / ~OK	~MAX	EASY	OK
Group best	KO ...	KO	~MAX	?	OK

Add-N-Paths

- My preferred
 - Does what we're really looking for
 - Bounded, "easily" predictable overhead
- If you'll only implement one, pick this one
 - Hack N when things go wrong...
- Fails to provide some guarantees

Coming soon

Mix

- Different modes on different routers
 - Vendors making different choices
- Different modes for different sessions
 - One vendor allowing different modes for different “session types”

“Session type” dependent mode

- Do Add-All-Paths
 - over sessions with your RRs
 - over normal sessions
- Do Add-2-Paths
 - over sessions with your RR clients

Why ?

- No longer issue a show ip bgp on your ASBRs...
- Backup availability is ensured
- Fast BGP Add-Path Control Plane Convergence
 - RR -> RR-client updates only
- ASBRs loaded/stressed like with Add-2-Paths

?