How To

Minimise the impact of unstable eBGP routes

Introduction

This document describes how to minimise the impact of unstable eBGP routes or flapping Border Gateway Protocol (BGP) routes, caused by recursive routing failure. Common symptoms of recursive routing failure in BGP are: constant deletion and re-insertion of BGP routes into the routing table and loss of connectivity towards destinations learned through BGP.

The solution to this problem is known as route flap damping. Route flap damping limits the impact and visibility of route flapping to a router's BGP peers, by suppressing routes that are considered too unstable to be used locally or advertised to any BGP peers.

What information will you find in this document?

This document provides information on:

- The problem known as route flapping
- The solution - route flap damping
- A description of the damping algorithm
- Useful configuration commands

What product and software version does this information apply to?

The information provided here applies to:

- Products: 8948, 9900-series, 9800-series
- Software version 2.6.6 and above

The problem

Route flapping

BGP-managed networks are more efficient and stable when routing update messages are kept to a minimum. Under some network conditions, BGP generates an excessive rate of update messages due to "route flapping", in which some routes frequently oscillate between being reachable and unreachable. Route flapping causes a ripple effect through the BGP network as the changes are propagated, and in extreme cases the network will not reach a stable, converged state for a substantial period of time.

The solution

Route flap damping

BGP route flap damping, as defined in RFC 2439, limits the impact and visibility of route flapping to a router's BGP peers. When a local BGP peer learns a route, it immediately adds it to its Routing Information Base (RIB), but may not immediately select or advertise it. It can only select or advertise the route once its internal BGP suppression engine considers that the route is sufficiently stable. A new route has no history of instability, so is immediately made available as normal. A route that has been previously learned but withdrawn, however, may be suppressed for a period of time, based on the severity of its previous instability. Therefore, BGP route flap damping suppresses routes that are considered too unstable to be used locally or advertised to any BGP Peers, until the route has remained stable for a sufficient period of time. Persistently unstable routes may be excluded from selection indefinitely. By taking into account the prior behaviour of that route, BGP can estimate the future stability of the route accurately enough to reduce router processing load without significantly impacting the convergence time for more stable routes.

Figure of Merit (FoM)

Route history

A route's history of instability is recorded via the maintenance of a statistic defined as a Figure of Merit (FoM) by RFC 2439. The FoM for a particular BGP route quantify that route's history of stability, or lack of stability. When the router learns a new route, it grants it an initial FoM of zero, indicating no history of instability. At this point, the BGP suppression engine is not interested in the route.

If an eBGP peer ever withdraws a route from the local device, it increments the FoM for that route by 1000. As soon as a route earns a non-zero instability metric, the BGP suppression engine begins to monitor it, but in most configurations takes no other action at this stage. If the route exhibits further instability, its FoM increases. Once the FoM exceeds a configurable suppression threshold, the BGP suppression engine begins to suppress it. At this stage, BGP no longer selects or advertises the route.

Each route's FoM is reduced over time, as the suppression engine gradually forgives the corresponding past transgression(s), so if a suppressed route remains stable for a sufficient period of time, its status is eventually downgraded to monitored. If a monitored route's FoM reaches zero, or a value close to zero, monitoring stops until further instability is observed.

The diagram below shows the FoM process of events.

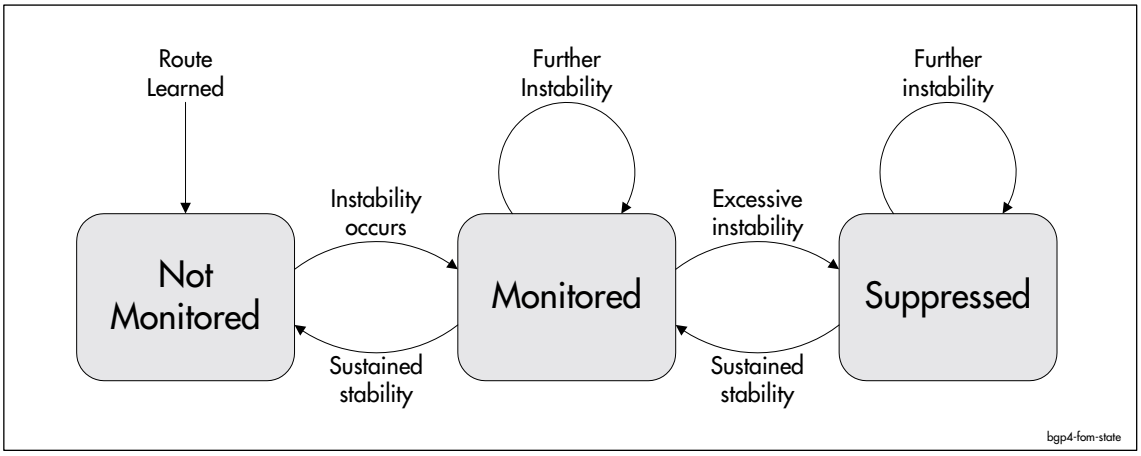

This next diagram shows how a route's FoM is maintained over time.

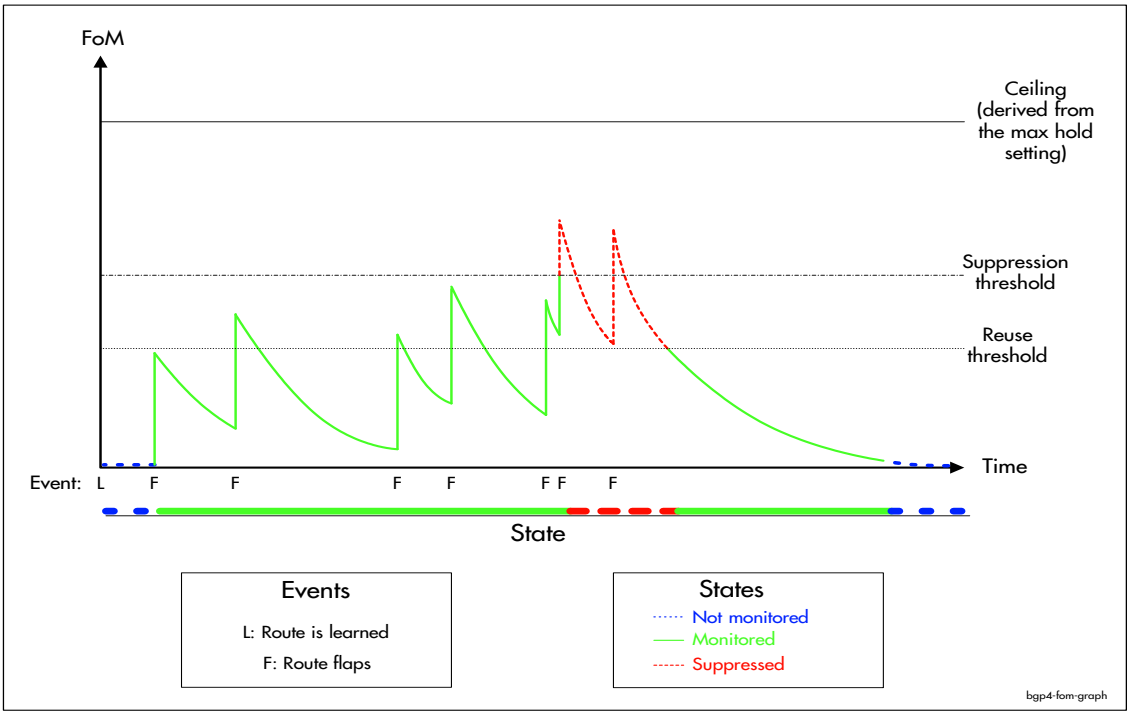

Configuring route flap damping

BGP route flap damping is disabled by default. You can enable it on all routes, or limit it to routes received by particular peers. You can use the default threshold settings, or specify different settings. The settings are captured by damping **parameter sets**, which are collections of four configuration parameters that determine the nature of the treatment received by relevant routes from the BGP suppression engine. The table below shows the parameters and the effect of increasing or decreasing each value.

Parameter	Meaning of parameter	Change	Effect of change
suppression	**Suppression** is a FoM value. When a route's FoM exceeds this threshold, the route is suppressed.	Raised	Increases the number of times the route can become unreachable before it is suppressed.
		Lowered	Decreases the number of times the route can become unreachable before it is suppressed.
reuse	**Reuse** is a FoM value. Once a route is suppressed, it remains suppressed until its FoM falls below this threshold.	Raised	Increases the minimum time that the route is suppressed for.
		Lowered	Decreases the minimum time that the route is suppressed for.
halflife	**Halflife** is the time interval within which the route's FoM will halve, if the route remains stable. For example, if the halflife is 15, the FoM of a stable route reduces by 50% over a 15 minute period, 75% over a 30 minute period, and so on.	Lengthened	Lowers the FoM more slowly, so increases the time the route is suppressed for.
		Shortened	Lowers the FoM more quickly, so decreases the time the route is suppressed for.
maxhold	**Maxhold** multiplied by **halflife** is the maximum period of time that a route must remain stable in order to become unsuppressed. For example, if **halflife** is 15 and **maxhold** is 4, the route is unsuppressed after 60 min. of stability even if its FoM still exceeds **reuse**.	Increased	Increases the time that a severely unstable route must be stable for, before it is unsuppressed.
		Reduced	Decreases the time that a severely unstable route must be stable for, before it is unsuppressed.

Enabling route flap damping

BGP route flap damping is disabled by default. To apply route flap damping to all incoming routes, use the command:

```
ENABLE BGP DAMPING
```

Damping all routes

To apply route flap damping to all incoming routes, using the default parameter settings, use one of the commands:

```
enable bgp damping
```

```
enable bgp damping parameterset=0
```

If you do not want to use the default parameter settings, change the settings for the default parameter set 0, using the command:

```
set bgp damping parameterset=0 ...
```

When you enable route flap damping globally, the router examines the instability history of every route it receives from every remote peer, and suppresses the route if appropriate.

Damping for a certain Peer

To enable damping for a certain Peer, the following steps are required:

1. Create a new route flap damping parameter set with the desired characteristics, and enable **only** that parameter set.

```
CREATE BGP DAMPING PARAMETERSET=1..100 [DESCRIPTION=description]
   [SUPPRESSION={DEFAULT|1..20000}] [REUSE={DEFAULT|1..20000}]
   [HALFLIFE={DEFAULT|1..45}] [MAXHOLD={DEFAULT|1..8}]

ENABLE BGP DAMPING PARAMETERSET=1
```

2. Create an IP routemap that applies the given parameter set to matched routes. In this example, all routes with a next hop of 192.168.1.1 will be damped using parameter set 1. However, the match clause may be changed or even omitted. If no match clause is provided, the route map will match all routes, and so all routes to which this route map is applied will be damped using parameter set 1.

```
ADD IP ROUTEMAP=dampRouteMap ENTRY=1 MATCH NEXTHOP=192.168.1.1

ADD IP ROUTEMAP=dampRouteMap ENTRY=1 SET BGPDAMPID=1
```

3. Configure the relevant Peer to use the given routemap as its inroutemap.

```
SET BGP PEER=peerIp INROUTEMAP=dampRouteMap
```

Viewing a complete list of monitored and suppressed routes

To view a complete list of all monitored and suppressed routes, including their current FoM and the period of time which they must remain stable in order to progress from suppressed to monitored, or monitored to not monitored, use the command:

```
SHOW BGP DAMPING ROUTES
```