

Parent Aggregate Change

- Parent Aggregate Change, on page 1

Parent Aggregate Change

This alarm detects an unexpected supernet or threshold violation.

A network operator is usually aware of their immediate supernet prefix (aggregate or summary) of their advertised prefixes, other aggregated higher order supernets, and their origin AS. The user must configure at least one set of expected IPv4 and IPv4 supernets by specifying their Classless inter-domain routing (CIDR) prefix-lengths. The user can also enforce that the observed aggregates originate from a list of allowed Origin ASs.

Note It is useful to know which of your peers may be doing something wrong (leaking route information or having some type of misconfiguration) so that you can address the problem right away. A **My Peers** rule is available for this alarm with certain Crosswork Cloud subscriptions. The **My Peers** option follows BGP updates *only* from your peers, whereas **All Peers** follow BGP updates from your peers *and* global peers. To configure this option, see Add Crosswork Cloud Network Insights policies.

Possible Problem Detected

This alarm can help identify the accidental withdrawal or route-leak of summary prefixes.

Relevant Alarm Rule Configurations

The following options must be configured when adding this alarm rule to a Prefix policy configuration (**External Routing Analysis > Configure > Policies > Add Policy > Prefix Policy > Add Rule > Parent Aggregate Change**):

- Thresholds (per advertised aggregate)
- Allowed Origin ASNs (optional)
- Allowed IPv4/IPv6 supernets

Example

You create a Prefix policy with the **Parent Aggregate Change** alarm rule and is linked to prefix 8.8.0.0/24. The policy is configured with allowed IPv4 aggregate prefix lengths [22, 9] and an Allowed Origin AS 3356. The following events will trigger an alarm:

- An expected supernet of 8.8.0.0/22 is hijacked (prefix originates from an unexpected origin AS).
- An aggregate, prefix 8.8.0.0/20, is advertised and is identified as a potential leak.

The alarms are cleared when either the leak or hijack is resolved, or the user changes the alarm configuration to indicate that these aggregate advertisements are legitimate.