CHAPTER 4

Causality Correlation and Root Cause Analysis

As mentioned in Chapter 5, "Understanding Fault Management," causality correlation is the process of relating an event to an existing alarm in a causality relationship. The root of the resulting causality hierarchy is called the *root cause*, and the correlation process that determines the root cause of a fault is referred to as root cause analysis.

Root cause analysis is performed on historic snapshots of the VNE model and forwarding information. These snapshots, maintained for ten minutes, enable the correlation logic to traverse the VNE network model at a time in the past, before the occurrence of the fault that is being analyzed. This is critical for root cause analysis in the presence of network faults.

The unique correlation capability of Cisco ANA is derived from this ability to learn the impact of the suspected fault by examining snapshots of the network before the occurrence of the fault.

These topics describe how Cisco ANA performs correlation logic decisions:

Attributes Used By the Event Correlation Algorithm

Table 4-1 describes the attributes that control the algorithm Cisco ANA uses to perform event correlation.

Table 4-1 Event Correlation Attributes

Attribute	Description
`correlate`	Determines whether the event should attempt to correlate and find a causing event. If `true`, it will correlate; if `false`, it will not correlate.
`activate-flow`	Determines if the new event will initiate a network correlation process (`true`) or only local correlation process (`false`).

Table 4-1 Event Correlation Attributes

Attribute	Description
Forwarding information	Event-specific flow information, such as the IP address of a BGP neighbor, used in network correlation. This information includes the following: • The component to start the correlation flow from. This might be the VNE component that identified the event (for example, an IP interface component in an *Interface status down* event) or any other component in the VNE. • Additional forwarding data used by the flow logic includes the flow direction and one of the following – The IP address used as the destination IP address of the flow. – The VC used as the entry point of the flow. – The Multiprotocol Label Switching (MPLS) label used as the entry point of the LSP. – The type of VNE component to stop at.
Correlation keys	A set of one or more unique identifiers that are typically derived from the event type, source, and parameters from the event message content, used in local and flow correlation.
`is-correlation-allowed`	If true, this event can be the possible cause for other events, and allow other alarms to correlate to it. An event with this attribute should also be ticketable.
`weight`	A positive integer which defines the relative weight of an event as a cause candidate in relation to other causing events. A new event can only correlate to an event which has a higher weight. The heavier the event, the more likely it will be chosen as the cause

Correlation Process

Figure 4-1 illustrates the correlation process in Cisco ANA:

Figure 4-1 Correlation Process in Cisco ANA

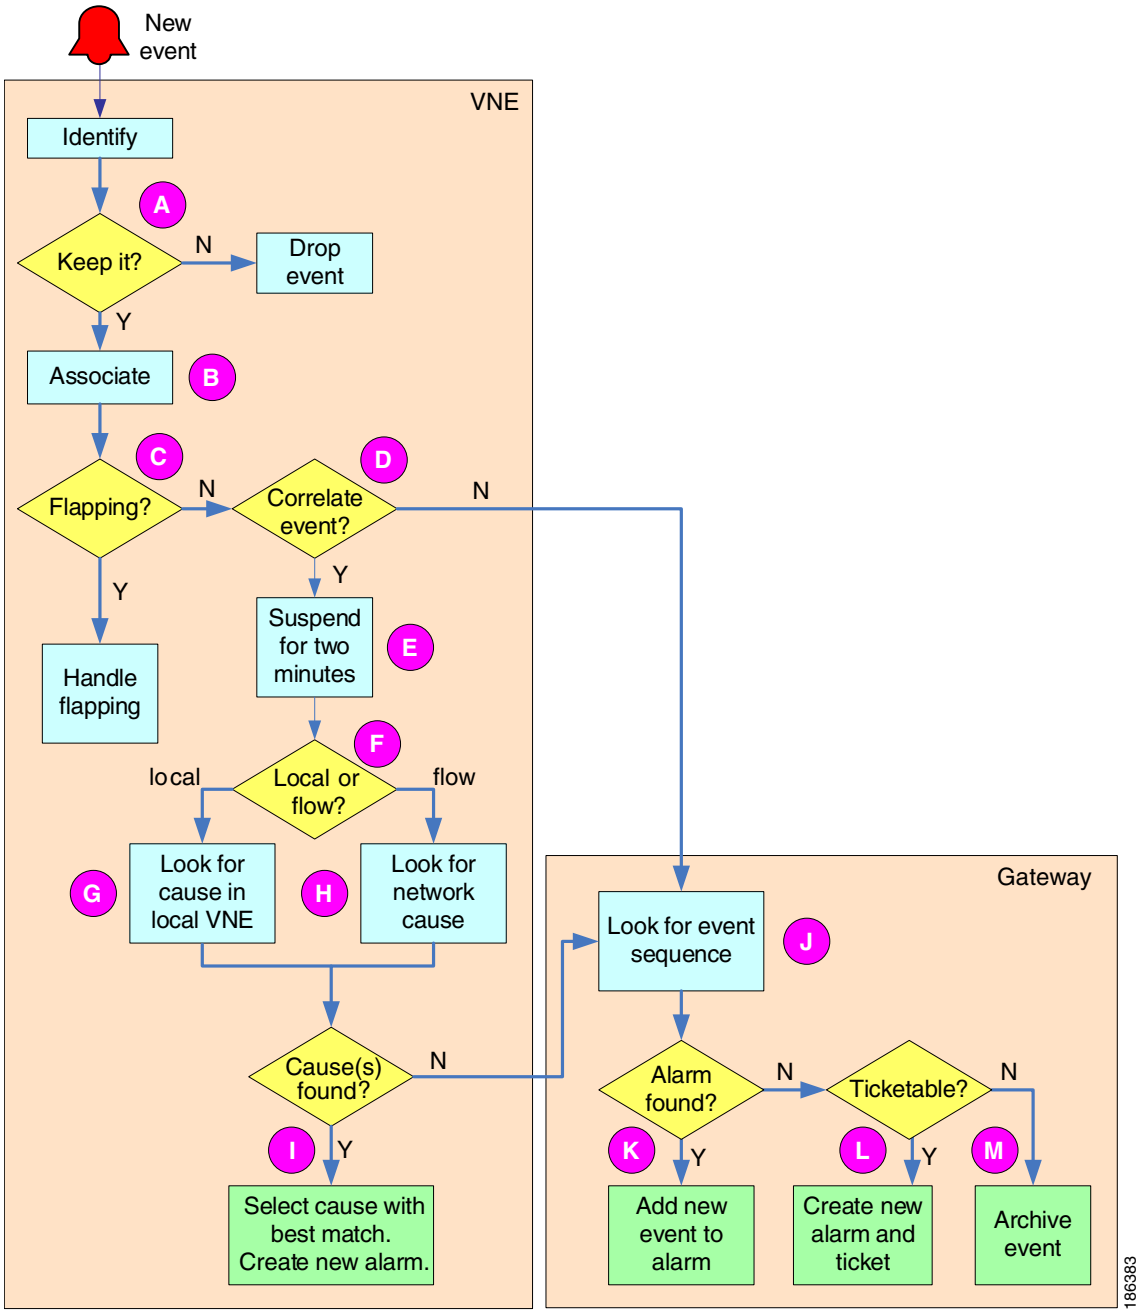

The correlation process kicks off when a new event is detected. A new event can be an incoming network event (for example, a trap event) or a generated event (for example, a service alarm). For more information on the correlate, activate-flow, and weight attributes, see Table 4-1 on page 4-1. The is-ticketable parameter is described in Event (Subtype) Configuration Parameters, page 18-2.

A The first step in the process is the identification, and after the trap has been identified and is not dropped by the VNE, its attributes are determined.

B Source association of the event, as described in Chapter 5, "Understanding Fault Management."

C Cisco ANA then examines the new event for flapping. If this event is part of a flapping sequence, it is suppressed as described in Chapter 5, "Understanding Fault Management."

D Cisco ANA then further examines the new event. If its `correlate` attribute is `true`, the new event is subjected to correlation. Most flagging events are subject to correlation; clearing events are not. If this event was configured to expedite polling, then the specific polling command is generated at this time as well, which might result in additional generated events.

If the correlate attribute is set to `false`, Cisco ANA tries to relate it to an event sequence. See steps J to M for additional details.

E The correlation process is now suspended for two minutes. This gives other events, possibly related in terms of cause and effect, a chance to be detected subsequently as well.

No further processing is performed on the event when it is suspended. Therefore the update to the database and Cisco ANA NetworkVision for this event is delayed by two minutes.

F After the suspension of two minutes, the next step is the examination of other events for possible causality. There are two different processes: *local correlation* and *flow correlation*. Which of the two processes is followed is subject to the value of the `activate-flow` attribute of the new event: `false` initiates local correlation (G) and `true` initiates flow correlation (H).

G Local correlation proceeds to examine and find other events related to the same (*local*) NE from which the new event originated; no network flow is initiated.

Local correlation is well-suited to the following network scenarios:

- The new event is in the device scope, for example, a module out event.
- The new event has a corresponding generated event message in Cisco ANA. It is desired that the event will correlate to its generated event message. For example, a Link Down syslog event tries to correlate to a local Link/Port Down generated event message.
- The new event does not contain information that can be used to perform correlation beyond the scope of the local device. Such information is used to initiate the correlation flow.

Most trap and syslog events use the local correlation process. The correlation logic looks for causing events in the VNE that:

- Have their is-correlation-allowed attribute set to True.
- Arrived within the last nine minutes (including the two-minute suspension time).
- Have a correlation key that matches at least one of the correlation keys of the new event.

H Flow or network correlation broadens the search for events to those coming from other NEs as well, including those that are several hops away.

Flow correlation is well-suited for scenarios such as the following:

- The event represents a failure in a connection or service that spans multiple devices. For example, an MPLS TE Tunnel Down event will try to correlate to faults on the path that the tunnel traverses.
- Logically, the new event can happen due to events that occurred in other devices. For example, a Device Unreachable event will try to find its root cause in other devices, by performing a flow to the management IP address.

Flow correlation uses historic snapshots of the VNE model to search the local VNE and other VNEs for causing events that meet the following criteria:

- Have their is-correlation-allowed attribute set to True.
- Arrived within the last seven minutes (including the two-minute suspension time).
- Exist on VNE components that appear on a flow path traversed according to the forwarding information of the new event.

I Often, the correlation processes (described previously) yield more than one candidate causing event. Additional, specific, rule-based filtering logic eliminates unlikely causing events, such as Cloud Problem, BGP Process Down, or LDP Neighbor Down. Finally, the causing event with the highest `weight` (or the closest in time, when there is more than one) is selected as the causing event.

The new event becomes the starting point of a new, correlated alarm.

J If the previously-described correlation process does not yield any causing events, or the new event was not subjected to correlation, it is possible that the new event simply relates to an existing event sequence (alarm). Cisco ANA searches for such an event sequence (with the same source and event type as the new event).

K If a matching event sequence is found, the new event is appended to it, and the alarm severity is updated accordingly.

L If no matching alarm is detected, the new event is potentially a root cause in its own right. This is determined by the `is-ticketable` attribute. If set to `true`, the new event is the basis for a new alarm that in turn causes the creation of a new ticket.

M If however, the new event is not ticketable (`is-ticketable` is set to `false`), it is now archived, and no longer involved in future correlations.