Correlation Node Types

Fifteen primitive node types are supplied with HP OpenView Event Correlation Services (ECS). Every correlation circuit must have one or more source nodes and one or more sink nodes.

Source Nodes. This is where events enter a correlation circuit. There can be more than one source node in a circuit. For a top-level circuit, source nodes are connected to the engine's input ports, where events are delivered by the HP OpenView DM postmaster.

Sink Nodes. This is where events leave a correlation circuit. For a top-level circuit, the events are returned to the HP OpenView DM postmaster for delivery to management entities that have registered to receive them.

It is necessary to be able to suppress unwanted events. In the circuit paradigm, events are filtered by preventing them from flowing through different paths in the circuit. This is done by filter nodes and unless nodes.

Filter Nodes. These suppress events based upon a configured expression which typically uses the incoming event as an argument.

Unless Nodes. These forward an event unless another event was created within a configured (positive or negative) time period relative to the creation time of the first event, and the configured (filtering) expression evaluates false.

Events can be delayed in the management network, especially when the network is stressed. This may result in delayed events and events arriving in a different order than the order in which they were created. Delay nodes can be used when correlation decisions depend upon events being processed in the strict event creation time order.

Delay Nodes. These hold an event until the creation time is a configured number of seconds before the current time. This has the effect of guaranteeing that the events are output from this node in creation time order.

Events may need to be stored for extended periods so that future correlation decisions can be made using the event history. Subsequently, it may be necessary to extract complete copies of events from the storage.

Table Nodes. These hold a logical copy of all events sent to the table, subject to configured retention parameters and conditions. Other nodes can examine the event list, stored in creation time order, and make processing decisions based upon the contents.

Extract Nodes. These search a table node and extract a copy of one or more events from the stored list, subject to a configured condition. The search is triggered by an event arriving at the input port of the extract node, and the extracted events are output as a composite event (see *Event Types*).

While most of the useful information will come from the event stream, data may need to be obtained from outside the correlation engine.

Annotate Nodes. These obtain data external to the engine and add it to an output composite event. The external data is now available in the event for subsequent use in the downstream circuit. The annotate request provides time for the request to be serviced, after which it will time out. Other events continue to be processed during this period.

One of the fundamental features of ECS is the ability to collect and consolidate discrete pieces of data from the event stream and from outside the engine to produce value-added information. Events need to be manipulated, including combining events into a single data unit, changing the structure of this unit, changing event data values, and creating new events.

Combine Nodes. At these nodes, two or more input event streams are combined into a single output stream, with each output event being a composite event containing an event from each input stream. Events on one stream can be held until events on other streams arrive.

Rearrange Nodes. These change the structure of a composite event, including pulling a single normal event out of a composite.

Modify Nodes. These change attribute values of incoming events to any values required. The values can be copied or calculated from any publicly available data. A copy of the original event is made, and the event copy is modified and output. The original event is not modified.

Create Nodes. These create a new event with a format controlled by a configured specification and attribute values set according to a configured specification. Event creation is triggered by an event arriving at the input port. The event's attribute values can be set from any publicly available data throughout the engine, including from the incoming event.

Some basic utility functions are provided by the following nodes.

Count Nodes. These count the events passing through the node.

Clock Nodes. These generate an empty event every configured time interval. This allows circuit logic to be triggered in the absence of any incoming events, enabling the absence of required events to be detected.

Rate Nodes. These calculate the rate at which events are passing through the node.