ITIL® Poster Series #51

Event Correlation

Introduction

The Service Operation process of Event Management describes how events generated by configuration items are gathered by monitoring tools and presented to operational staff so that they may respond quickly to any issues detected. Event management is one of the main activities of IT operations. It would be very difficult, or even impossible to manage complex modern infrastructures without the use of these tools. They enable the monitoring of large numbers of configuration items simultaneously, identifying any issues as soon as they arise and notifying technical management staff.

The huge number of events generated by monitoring tools need to be managed and understood. This is the job of the correlation engine. In this article we examine the role of this software, and its importance to achieving effective event management.

Event-generating software and hardware enables the management of large numbers of configuration items. Events are usually in a standard Simple Network Management Protocol (SNMP) format. These events are trapped and forwarded by the event management tool to an operator console.

EVENTS MAY BE GENERATED BY:

- Operating system logs
- Servers
- Application logs
- Firewalls and VPN gateways
- Network routers and switches.

The more complex the infrastructure, the more events are generated. The volume of events generated means that operators would be overwhelmed with data, and unable to focus on what is important. To prevent this, event management tools filter and correlate events

Remember, the purpose of event management is to detect events, understand what they mean, and take action if necessary. Monitoring enables the first part – it detects the events. Filtering and correlation are concerned with the understanding of the significance of the event.

Filtering prevents the event management system from being overwhelmed by discarding notifications of events that have no significance to the organization. We may still want these events to be created and stored, to show an audit trail, for example, but we do not need to be notified that they have occurred.

As events are detected, the event management system must interpret and make decisions about how to handle them. This is done by software known as a correlation engine. This a software application that can be programmed to understands the relationships between CIs. It uses predefined business rules to determine the significance of any warning or exception events. It can then be programmed to decide the appropriate next steps.

Using a correlation engine will enable the system to determine the significance of each event and also to determine whether there is any predefined response to an event. Patterns of events are defined and programmed into correlation tools for future recognition. The diagram shown here shows how this works.

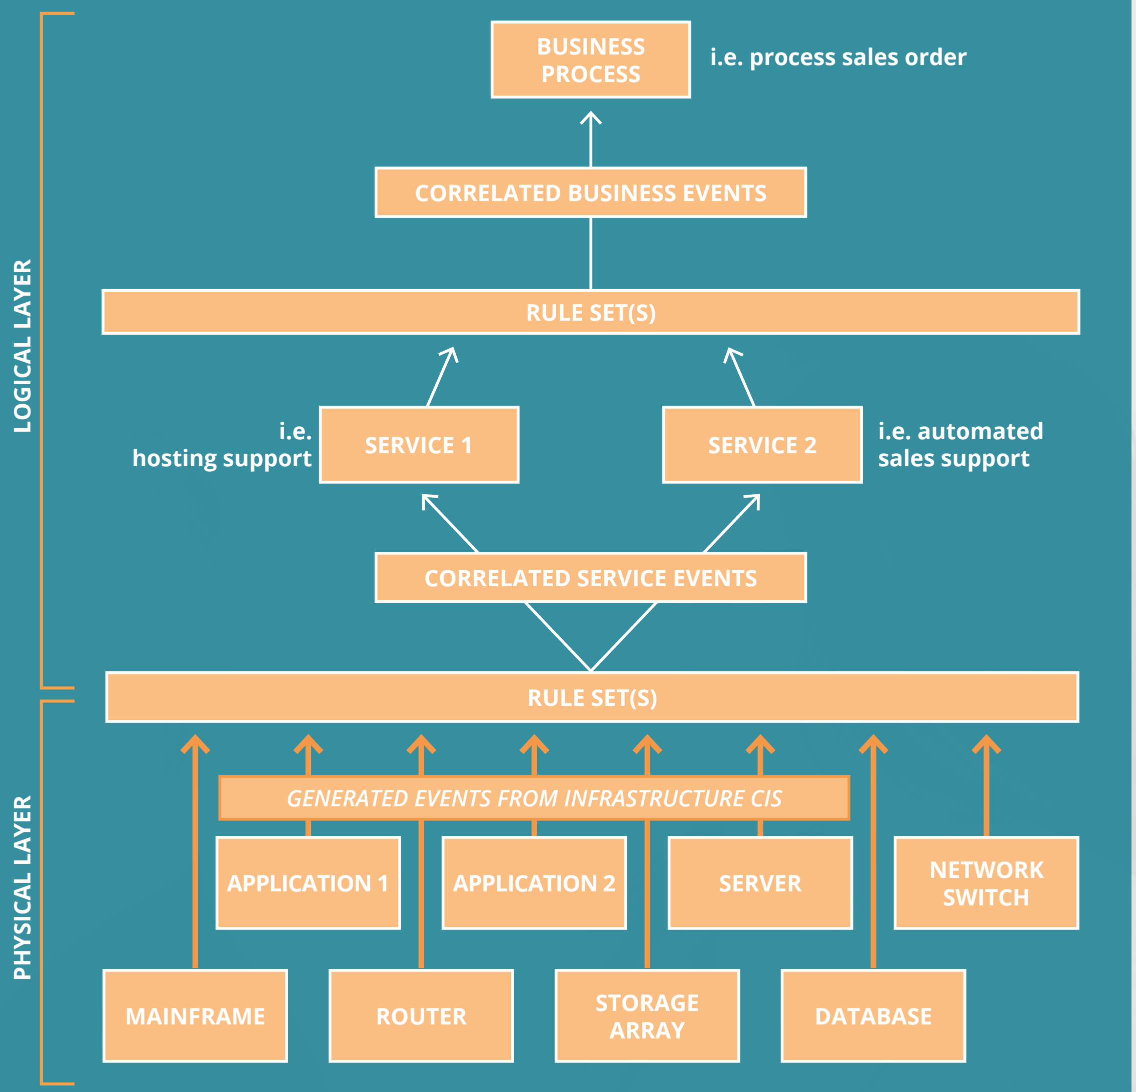

1. CORRELATION ENGINE

The correlation engine rules can associate events both from a technical perspective, and a business perspective. The correlation engine translates component-level events into service impacts and business impacts. This will enable an accurate assessment of the business impact to be made, and the incident to be correctly prioritised.

EXAMPLES OF EVENT CORRELATION FROM A TECHNICAL PERSPECTIVE INCLUDE:

- Compression which reports multiple occurrences of the same event, them as a single event.
  So rather than receiving 1000 "route failed" alerts a single alert that says "route failed 1,000 times" is generated.
- Suppression is when alarms are not reported for low-priority events if a higher-priority event has already occurred.
- Generalization correlates multiple events on the same device such as connection failures to multiple ports on the same switch. The rules determine that it is the switch that has failed, and so the individual alarms need not be sent, just one high-priority event for the whole device.