Aspect Architecture Document

Serviceability

Author(s): John Logan
Save Date: 7/16/2010
Associated Product Requirements:

Feature-Level Variation Control Sequence:
EXCLUDE: <default>
INCLUDE:

Document: 48357-00
Preliminary Revision: A.1

TABLE OF CONTENTS

LIST OF FIGURES

Variation Point List

Definition:
Requirements Mapping:

REVISION HISTORY

Revision	Description of Changes
A.1	Initial version.

1. Element Functional Behavior Changes

The following subsections summarize the functional behavior changes between document revisions.

1.1. Revision A.1

Initial version – no changes apply.

2. Introduction

The centralized logging architecture describes a subsystem that unifies the processing of all *system events* within an OSA-based product, and maintains a *system event log* that records the system event history. A system event (also referred to simply as an “event” within this document) is an asynchronous message that conveys information about user-visible changes in system state, such as component failure state and object lifecycle state.

The system event log is distinct from *trace logs*, which maintain histories of debug-level events, and are extracted as components of the support bundle. See XREF for additional information on trace logs.

The architecture consists of the following software components:

- *Event sources* that generate system events. Event sources in a Unified Storage Platform (USP) include sources of B[illegible]DPL, NAS, platform hardware, serviceability, and management events.
- *Management clients* that access the *system event log*, and receive *alerts* derived from system events, including SNMP traps, email alerts, CIM indications, and Engenio Monitoring and Reporting Services (EMRS) data dumps. The principal management client is the embedded Amelia Element Management (EM) server, with which the Amelia Element Manager browser-based application and the Amelia Enterprise Console (EC) application communicate. Other management clients include the EMRS data warehouse, the user's email server, and third-party IT management frameworks.
- An *event router* that receives events from event sources, stores the events in the system event log, and sends notifications to management clients.

Figure 1 illustrates the relationship between event sources, the event router, and management clients in an OSA-based unified storage (file and block services) product.

The event routers in a dual-controller environment operates as a clustered application to ensure the proper processing of events despite controller failures.

Figure 1: Top-level functional model for central logging.

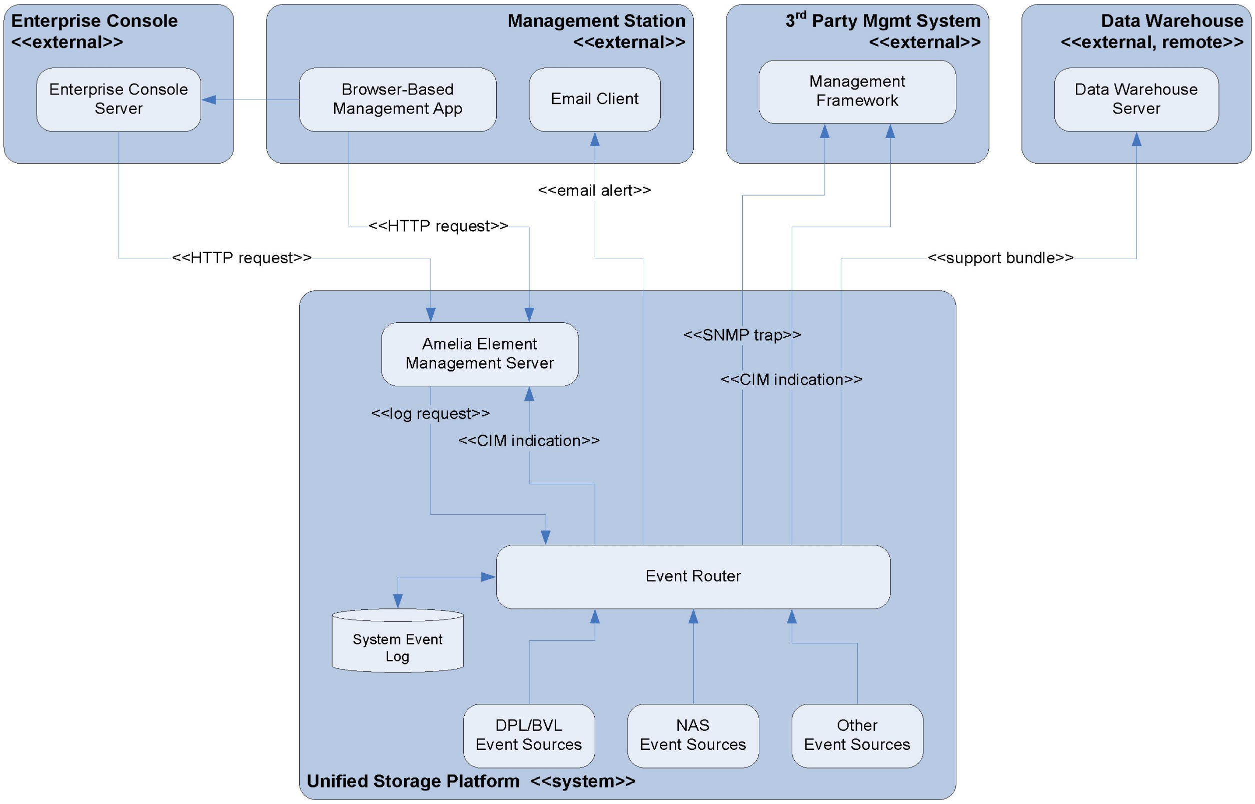
Enterprise Console
<<external>>
Enterprise Console Server
Management Station
<<external>>
Browser-Based Management App
Email Client
3rd Party Mgmt System
<<external>>
Management Framework
Data Warehouse
<<external, remote>>
Data Warehouse Server
<<email alert>>
<<HTTP request>>
<<HTTP request>>
<<support bundle>>
Amelia Element Management Server
<<SNMP trap>>
<<CIM indication>>
<<log request>>
<<CIM indication>>
Event Router
System Event Log
DPL/BVL Event Sources
NAS Event Sources
Other Event Sources
Unified Storage Platform <<system>>

3. Operational Behavior

3.1. Event Content

3.1.1. Event Types

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDAR0BZB

A system event has an *event type* that classifies the event. All system events of the same type share the following user-visible fields:

- An integer identifying the event type.
- An integer indicating the severity level for events of the type. Valid severity levels conform to IETF RFC 5424:
    - EMERGENCY (0)
    - ALERT (1)
    - CRITICAL (2)
    - ERROR (3)
    - WARNING (4)
    - NOTICE (5)
    - INFORMATIONAL (6)
    - DEBUG (7)
- A boolean indicating whether the event is visible to the management clients accessing the system event log.
- A boolean indicating whether the event results in an email alert being sent to the user.
- A boolean indicating whether the event should generate an SNMP trap.
- A boolean indicating whether the event should generate an EMRS data dump.

An *event configuration table* (ECT) defines the default type-related information for the product.

An *alternate event configuration table* (AECT) permits OEMs to customize the severity, visibility, and alert attributes per event type for an OEM-specific product. Event type entries in the AECT override corresponding entries in the ECT. The event router loads the ECT and AECT at start-of-day.

3.1.2. Event Text

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDAYDCZB

Event text is a user-visible field associated with an event type; all system events of the same type share the same event text.

An LSI-defined *event text table* (ETT) maps (event type ID, ISO 639 locale) tuples to localized event text. The event text table contains a complete set of entries for the en-US locale, and may contain entries for some or all messages for other locales.

An *alternate event text table* (AETT) is an ETT that permits OEMs to customize event text for an OEM-specific product. Event text entries in the AETT override corresponding entries in the ETT.

The event router loads the ETT and AETT at start-of-day. Run-time modification of the ETT and AETT are not permitted.

3.1.3. Event-Specific Content

Topic ID: 2010-06-16T23:04:00Z-1514-8635-IDAIZLOC

System events contain the following user-visible fields:

- A monotonically increasing event sequence number (ESN).
- The type of the event.
- The UTC time at which the event occurred at the event source.
- The UTC time at which the event was written to the log.
- The array controller/NAS node containing the event source.
- The object ID for the component associated with the event.
- A variable-length component whose content is specific to the type of the event.

3.2. System Event Log

3.2.1. Centralized Logging

Topic ID: 2010-06-30T18:43:00Z-3203-18259-IDAFDE3B

The system event log appears to the user and to management clients as a single container of event messages for the controllers that constitute the OSA-based product.

3.2.2. Log Persistence

Topic ID: 2010-06-16T23:04:00Z-1514-8635-IDAP2LOC

Events written to the event log file persist and remain available following both controller swap and disk adoption procedures.

Events may be lost if the controller failure, controller swap, or disk adoption occurs while BVL volumes are unavailable.

3.2.3. Log Capacity

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDADJCZB

The system event log stores at least 131,072 events. The event router regularly purges the oldest system events from the system event log to prevent out-of-space conditions.

The event router does not permit management clients to purge the system event log, nor does it purge events following EMRS support bundle collection. It is the responsibility of the management client to implement features that allow users to "virtually" trim the event log.

3.2.4. Log Access

3.2.4.1. Event Retrieval

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDA3KCZB

Management clients may incrementally browse the event log by retrieving ~~retrieve~~ ranges of events from the event router. Management clients request ranges based on ESN, event log time, or event occurrence time. Ranges may be specified by a start value and an end value, or by a start value and a number of events to retrieve.

3.2.4.2. CIM Interface Support

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDA1LCZB

A CIM object manager, acting as a management client, can access the event router and present a CIM interface that conforms to the DMTF Record Log Profile as described in DMTF Document Number DSP1010, version 2.0.0.

3.2.4.3. Event Filtering

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDA1MCZB

Management clients may provide a filter specification when retrieving ranges of events. The filter specification allows filtering by any event field, event type field, or combination thereof, using comparison and substring match operators.

The system event log API supports the SELECT capability exposed through the Amelia management CLI described in XREF.

Architecture Note: Deferring on “dependent event filtering/hiding” to a later iteration when we have a decent definition of the desired behavior. This requires identifying all cases where there are dependencies between DPL, BVL, and NAS events that result in event filtering.

3.2.4.4. Localized Event Text

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDAHOCZB

Management clients may specify a locale string in the form of an ISO 639 language code when making event requests to retrieve localized versions of event text. The event router retrieves US English (en-US) event text when clients do not specify a locale string, or the locale string is unrecognized for an event.

3.2.4.5. Timezone Support

Topic ID: 2010-07-07T17:16:00Z-2185-12456-IDAKPCZB

The event router returns all time values as UTC timestamps. Management clients are responsible for converting time values to local time.

3.3. Relationship with Fault Management

Topic ID: 2010-07-15T17:27:00Z-2356-13431-IDAX4XQD

The embedded EM server retrieves fault data from major subsystems (such as BVL, DPL, and NAS) via subsystem-specific interfaces. The Recovery Guru feature of the EM server presents fault-specific recovery actions to the user.

When a subsystem detects a persistent fault, it sends an event with CRITICAL severity.

When recovery actions result in the clearing of a persistent fault, the responsible subsystem sends an event with CRITICAL severity.

3.4. Management Client Alerts

3.4.1. Alert Generation

Topic ID: 2010-07-15T17:27:00Z-2356-13431-IDAJBYQD

With the exception of delayed alerts (see Section Section 3.4.2. Delayed Alerts), events with CRITICAL severity generate the following alerts to management clients:

- A CIM indication to all clients registered to receive the indication.
- An SNMP trap to all registered trap receivers.
- An email message to all registered email receivers. The message contains a subject line that includes the event text, and a message body that includes all event fields.
- An EMRS data dump, sent to the EMRS data warehouse. XREF describes the EMRS subsystem architecture.

3.4.2. Delayed Alerts

Topic ID: 2010-07-08T21:32:00Z-2325-13253-IDA3QAQC

A delayed alert is an alert that is dispatched for persistent fault conditions that may resolve themselves without user intervention within a specified time period. The event router logs the event associated with the delayed alert immediately, but does not dispatch alerts for the event. The embedded EM server detects the associated fault condition (see Section), and logs a separate event with CRITICAL severity if the fault remains after the fault-specific time period, and the event router dispatches alerts for this event.

3.5. Reliable Operation

3.5.1. Reliable Delivery

Topic ID: 2010-06-16T23:04:00Z-1514-8635-IDAHWLOC

In the absence of controller failures and event router failures, the event router delivers a sequence of events sent by an event source such that the system event log and management clients receive the sequence in order, without omission or duplication of events.

3.5.2. Controller Failure Handling

Topic ID: 2010-06-30T19:18:00Z-3122-17802-IDAVQLHC

Controller failure may result in brief disruption (no greater than 5 seconds) to system event log availability. Log access requests pending at the time of failure time out and should be retried by management clients.

Events being processed by an event router on an unexpectedly failing controller are not guaranteed to generate alerts or be logged. Trace logs at the event source (see Section Section 3.5.7. Traceability After Failure) capture such events.

Event logging and alert dispatch continue to function normally for event sources on the surviving controller.

Event sequence numbering continues without discontinuity despite controller failure.

3.5.3. Event Router Failure Handling

Topic ID: 2010-06-30T19:02:00Z-3102-17682-IDAYRF3B

Failure of the event router may result in brief disruption (no greater than 5 seconds) to system event log availability. Log access requests pending at the time of failure time out and should be retried by management clients.

Temporary unavailability of the event router is tolerated such that event sources need not handle outage-related errors. A compile-time parameter determines the maximum amount of time that constitutes a temporary outage. Outages that exceed this limit trigger a controller reset.

Events being processed by a failing event router are not guaranteed to generate alerts or be logged. Trace logs at the event source (see Section Section 3.5.7. Traceability After Failure) capture such events.

Event sequence numbering continues without discontinuity despite event router component failure.

3.5.4. Service Mode Handling

Topic ID: 2010-07-15T17:27:00Z-2356-13431-IDAXLYQD

Placing a controller in service mode, or removing the controller from service mode, may result in brief disruption (no greater than 5 seconds) to system event log availability. Log access requests pending at the time of failure time out and should be retried by management clients.

Events being processed by an event router are not lost when a controller enters or leaves service mode.

Event logging and alert dispatch continue to function normally for event sources on a controller that is in service mode.

Events sequence numbering continues without discontinuity when a controller enters or leaves service mode.

3.5.5. Lockdown Handling

Topic ID: 2010-07-15T17:27:00Z-2356-13431-IDA5MYQD

A brief disruption (no greater than 5 seconds) to system event log availability may result when one controller of a pair enters or leaves the lockdown state. Log access requests pending at the time of failure time out and should be retried by management clients.

The event log is not accessible, and event routing ceases to operate, when both controllers in a pair are locked down. In this state, the event router returns error responses to both event log access requests and submissions from event sources.

Events being processed by an event router are not lost when a single controller enters or leaves the lockdown state.

Events logging and alert dispatch continue to function normally for event

Events sequence numbering continues without discontinuity when a controller enters or leaves the lockdown state.

3.5.6. Solicited Controller Resets

Topic ID: 2010-07-15T17:27:00Z-2356-13431-IDA3NYQD

Event sources may request confirmation that event router processing completes for an event so that a diagnostic event can be logged prior to a solicited controller reset.

3.5.7. Traceability After Failure

Topic ID: 2010-06-30T18:43:00Z-3203-18259-IDA5TD3B

To facilitate debugging of controller and event router component failures, the event router ensures that events are synchronously written to a trace log at the event source. The EMRS capability (see XREF) is responsible for extracting support bundles, including trace logs, following such failures.

3.6. Unified and Block Only Product Support

Topic ID: 2010-06-18T22:09:00Z-1911-10896-IDA3MLRD

The event router operates in both unified (NAS and block) and block-only product configurations.

3.7. Single-Controller and Dual-Controller Product Support

Topic ID: 2010-06-30T19:18:00Z-3122-17802-IDAZNLHC

The event router operates in both single-controller (simplex) and dual-controller product configurations.

4. Administrative and Configuration Interfaces

Architecture Note: TODO: SNMP trap receiver configuration interface, email configuration interface. EMRS document specifies warehouse configuration.

5. Error Handling and Event Notification

5.1. Event Router Restart Event

Topic ID: 2010-07-15T17:27:00Z-2356-13431-IDAAUYQD

EVENT SYNOPSIS: The event router restarted a failed component.
MEL AFFECTED COMPONENT: Controller
FAILURE TYPE NAME: N/A
RECOVERY PROCEDURE: None
ADDITIONAL EVENT DETAILS:

6. Serviceability

6.1. System Event Log Serviceability

Topic ID: 2010-06-30T16:46:00Z-3197-18224-IDARFXOC

A system event log dump utility extracts a human-readable copy of the system event log from a root shell or support bundle script on a running system.

A system event log development library allows offline tools to analyze system event logs obtained from support bundles.

7. Compatibility and Migration

7.1. No Software Upgrade to USP

Topic ID: 2009-11-05T03:46:00Z-830-4737-IDACQCO

There is no software upgrade path from prior products to the USP product.

7.2. Event Log Schema Migration

Topic ID: 2010-07-02T22:37:00Z-2581-14714-IDA0WDND

The data fields associated with events and event types may change between different firmware revisions. Following firmware upgrades and (where applicable) downgrades, the event router automatically converts the event log to conform to the current event log schema.

8. Restrictions and Limits