White Paper: Management of External Light Sources and Co-Packaged Optical Engines

OIF-MGT–Co-Packaging–ELSFP-01.0

November 1st, 2022

Abstract:

Co-packaged optical transceivers may use lasers in external light sources. OIF's Common Management Interface Specification (CMIS) can be used to manage the transceivers and light sources. The recommended management architecture is that the transceivers and the light sources are managed jointly by a host controller. In particular, OIF External Light Sources (ELSs) can provide optical power to OIF Co-packaged Optical Engines (CPO OEs), under the control of this management architecture.

Document Editor:

Eric Bernier	Huawei

Contributors:

Dominic Goodwill	Huawei
Jock Bovington	CISCO
Stefan Langenbach	CISCO
Todd Rope	Marvell
Jeff Hutchins	Ranovus

CONTENTS

Glossary

Application – An Application is a type of functional configuration that is characterized by specific signal propagation or signal processing between one or more host lanes and one or more media lanes, overall providing a well-defined signal or data transmission function to the host. For example, a module may support one 400Gbps Application that is characterized by a 400GAUI-8 host interface and a 400GBASE-DR4 media interface combination.

Color – A CW light source has a color attribute, which indicates what standard-compliant “color” it is. Illustrative example colors are: “400GBASE-DR4 gray”, “400GBASE-FR4 wavelength 2”.

CW light – An instance of continuous wave (non-modulated) light, delivered from an external light source to an optical engine.

Data Path – The specific host and media lanes of a module that are used to implement one instance of an Application, together with all required internal module resources, is called the Data Path of that Application instance.

ELS port – A physical fiber output from an external light source. An external light source port is used to deliver 1 or more CW light to 1 optical engine through the fiber.

External light source (ELS) – An optical module that provides light. An external light source provides optical power to an optical transceiver, for optical transceivers that do not have light source.

Host board – A board that contains OE and may contain ELS.

Shelf laser board – A board that contains ELS but does not contain OE.

OIF ELSFP – An ELS that complies with OIF ELSFP implementation agreement.

OIF OE – An OE that complies with OIF OE implementation agreement.

Optical engine (OE) – An optical communication device, the optical engines are expected to be in close proximity to the Host ASIC. The optical engine may include functions such as modulator, photodiodes, optical passives, control, driver, and receivers. The features of the engine depend on the final system architecture. An optical engine may have internal light source(s) (OE-with-laser) or may needlight from external light source(s) (OE-withoutlaser).

1 Introduction

Next generation datacenter switching networks and high-performance computing, such as machine-learning and artificial intelligence, are increasingly challenged by a combination of high power dissipation and the need for high bandwidth I/O escape from the ASICs enabling these applications. Scaling of current architectures suggest that next generation systems will challenge the cooling capabilities of these systems. New architectures and new technology implementations are required if the desired performance levels are to be achieved.

Co-packaging, where optical or electrical communications devices are attached on the same first-level substrate as the host ASIC (Figure 1), is expected to provide high bandwidth interconnects with significant power savings. By locating the optical engine in close proximity to the Host ASIC, the high-speed electrical channel losses and impedance discontinuities can be minimized, thus enabling the use of higher speed, lower power, off-chip I/O drivers.

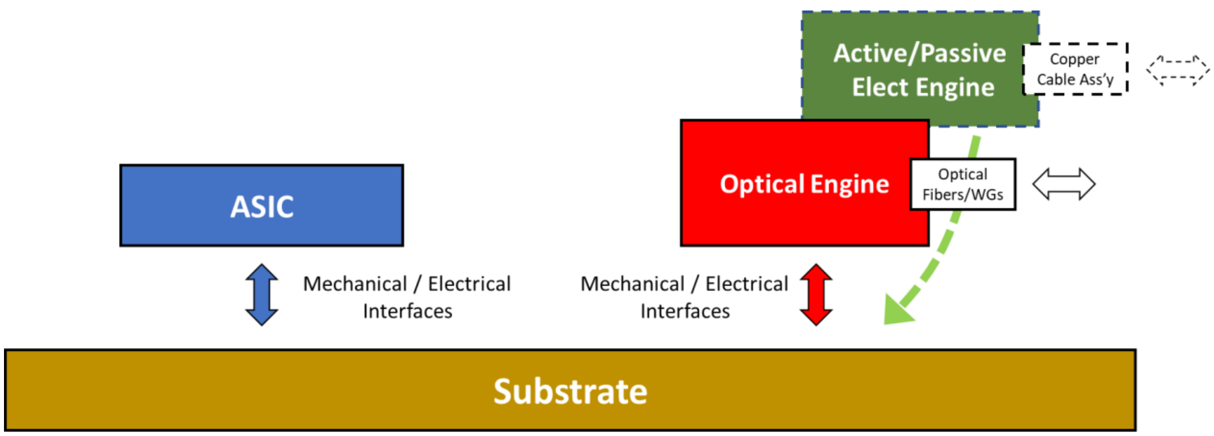

Figure 1: Co-Packaging implementation

In some cases, optical power is not generated within the optical engines themselves, as illustrated in Figure 2. Instead, optical power (also known as CW light) is created in external light sources.

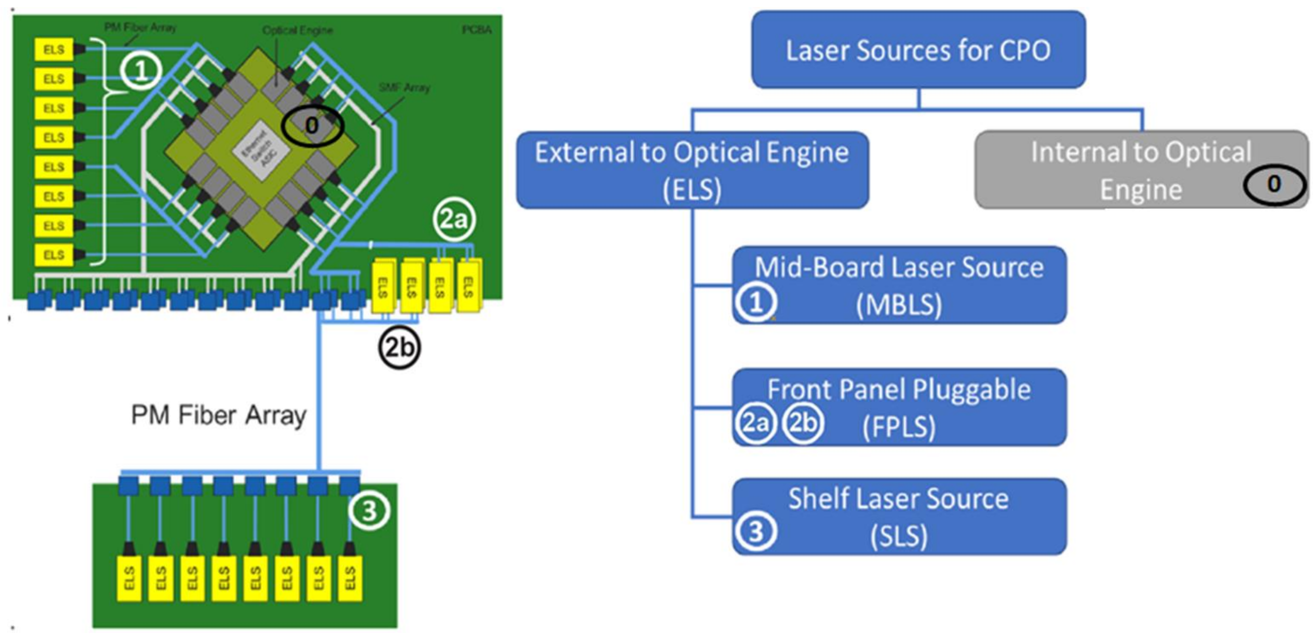

Figure 2: Laser light source use cases

To properly manage the optical engines and external light sources requires an efficient system management architecture. The recommended system management architecture is described in this White Paper. **The key feature is that the intelligence for controlling the CW light sources is in the host board controller, not in the optical modules.**

2 Scope

OIF has published a Framework for using co-packaged optical transceivers, which are also known as optical engines [1]. The Framework includes two fundamental solutions for providing optical power (Figure 2):

- lasers within the optical engine
- lasers in external light sources, outside the optical engine

OIF has published implementation agreements / specifications that may be used to implement the Framework:

- An implementation of a co-packaged optical engine (OIF OE) [2]
- An implementation of a pluggable external light source (OIF ELSFP) [3]
- An implementation of management communication (OIF Common Management Interface Specification CMIS) [4].

This White Paper describes the recommended system management architecture for the delivery of optical power to co-packaged optical engines. This system management architecture properly exploits the OIF implementation agreements of co-packaged optical engines and of external light sources. The management interfaces used in the architecture can be implemented using the OIF Common Management Interface Specification (CMIS). References to these OIF implementation agreements are listed in Section 10.

Optical modules that claim compliance to OIF implementation agreements must implement the management interfaces that are specified in those agreements.

However, system vendors may use the optical modules in any system management architecture of their choice. Therefore, this White Paper is informative only.

3 Managing optical engine with internal lasers

If a system uses only Optical Engines with internal lasers, then the lasers can be controlled in a traditional way. The controller inside the Optical Engine controls the lasers. This is shown in Figure 3.

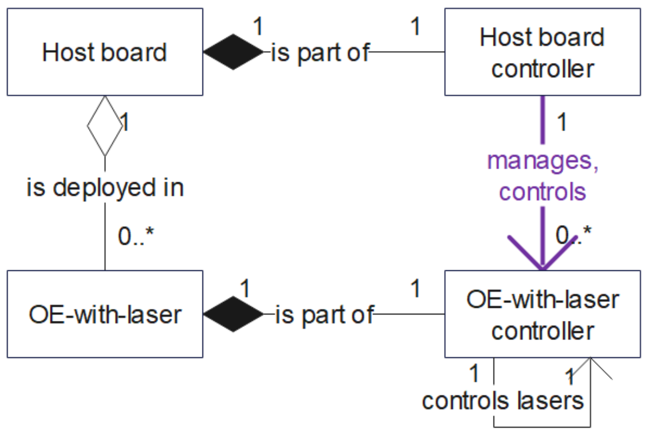

Figure 3: Management architecture for Optical Engines with internal lasers

This is architecturally identical to the management of traditional pluggable optical transceivers that have internal lasers. While there are minor CMIS enhancements specific to implementing co-packaged OIF OE, the overall system management concepts do not need to change.

4 Architectural requirements when using external light source

4.1 Overview

If the system uses ELS, then a change to the system management architecture is necessary.

A source of optical power (a laser in this case) is a necessary resource for an optical transceiver. If the lasers are in external light sources, then some of the necessary resources are no longer inside the optical transceiver. In particular, the sources of CW light are outside the optical transceiver. Thus, the traditional management architecture is not sufficient to control all the necessary resources.

A given optical module design may be used in a wide variety of systems. There are numerous potential ways to connect the CW light sources and the Optical Engines within a system and there are numerous deployment scenarios.

As a general rule, placing too much system context inside the optical modules would limit how they can be deployed. Therefore, the system management architecture must allow flexible deployments of simple optical modules. Hence the intelligence for managing CW light sources resides in the host board controller.

4.2 Relationships between ELS and OE

Figure 4 shows all the supported relationships between ELS and OE.

For example, an OIF 3.2Tb/s OE has either 4 or 8 CW light input ports. System vendors should choose an ELS with at least that number of fiber outputs and wavelengths with appropriate fiber management to deliver the CW light.

Similarly, an OIF ELSFP has up to 16 CW light output ports. System vendors may then choose any number from 1 to 16 OE to accept the light. Supporting for example 1 or more OEs.

In the short term, we envisage that most deployments would have an entire OE powered from a single ELS, and that a given ELS would power only 1 or 2 OEs. Nonetheless, the technologies will evolve and the system management architecture must remain flexible.

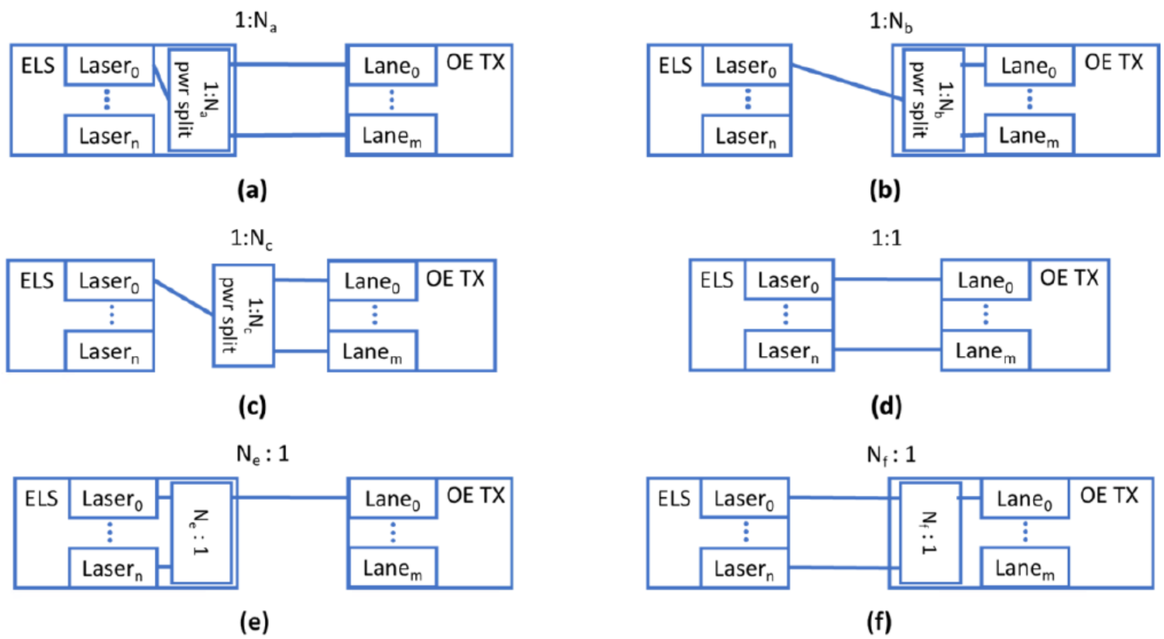

Figure 4: Relationships for delivering optical power from ELS to OE

4.3 System hierarchy

System hierarchies are shown in Figure 5. The three hierarchies in Figure 5 correspond to the cases of Figure 2 as follows:

(a) Scenario 0: OE with internal laser

(b) Scenarios 1, 2a, 2b: OE without internal laser, ELS on same host board

(c) Scenario 3: OE without internal laser, ELS on different board.

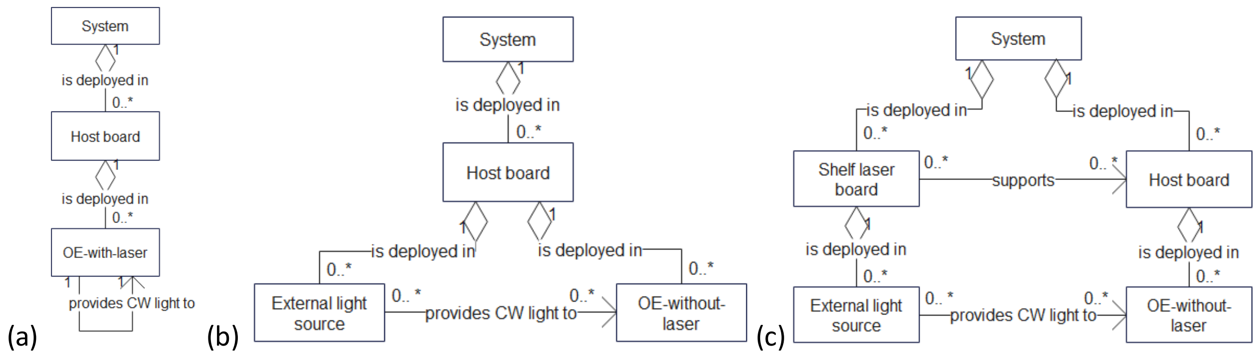

Figure 5: Hierarchy of system, boards, modules and CW light

4.4 CW Light

CW light and its relationships is illustrated in Figure 6 for a generic ELS, and Figure 7 for an OIF ELSP in particular. These figures are for properly connected and functioning systems. In principle, an ELS may deliver more than one CW light through each power delivery fiber, for example using WDM. In the case of an ELSFP, there is only one CW light through each power delivery fiber.

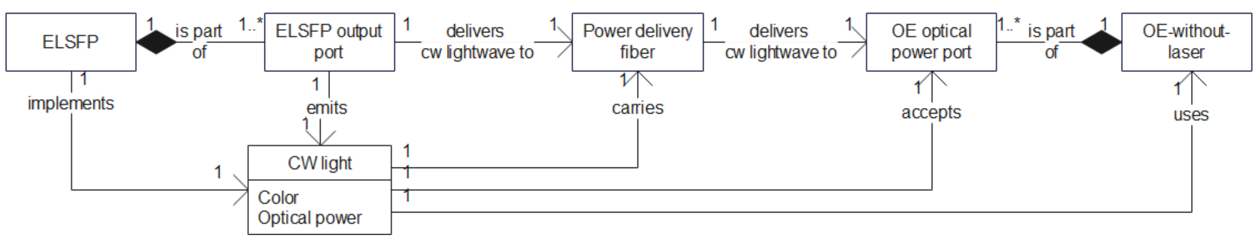

Figure 6: CW light relationships, for generic ELS

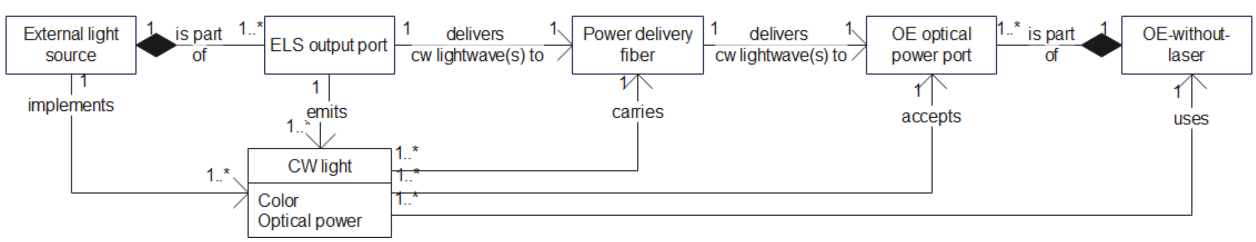

Figure 7: CW light relationships, for OIF ELSFP

A CW light is implemented by an ELS. Its attributes are color and optical power. It is emitted by an output port of the ELS. A fiber delivers the CW light to an OE.

A CW light has an identity and has a color and an optical power.

As shown in the OIF co-packaging Framework, the color attribute is static. In principle, the architecture could be extended to allow tunable lasers which can change the color, although this is currently out of scope.

The optical power attribute is dynamic. An OIF ELSFP is required to support setting of the optical power to within a specified range. The primary intended use case is that the OE advertises its currently-required optical power and the host board controller instructs the ELS to set the optical power to the required level. The goal is that the transmitted output of the OE complies with specified parameters such as average optical power and optical modulation amplitude (OMA).

More broadly, uses of the optical power attribute include:

- The ELS vendor configures the optical power to a default value and the host board controller verifies that this value meets the requirements of the OE.
- The host board controller sets the optical power during system initialization and then does not change it
- The host board controller varies the optical power, as the OE advertises its changing needs over life or over environmental conditions.

4.5 Manage use cases versus control use cases

For better understanding, we distinguish "manage" use cases from "control" use cases.

"Manage" use cases relate to the underlying needs of a module (OE or ELS) itself, such as updating firmware, and handling warnings and alarms.

"Control" use cases relate to behavior of CW light, such as advertising CW light requirements and capabilities, and control of the state and parameters of a CW light source.

4.6 Management architectures considered

Three management architectures were considered by OIF:

- Host controller manages the delivery of CW light from ELS to OE (chosen architecture)
- OE manages ELS ports that deliver CW light to it (rejected)
  - Given the large variety of deployment scenarios, the controller hardware within an OE likely has insufficient compute and memory resources to directly control ELS.
- ELS is managed like to electrical power supplies (rejected)
  - OE have more sophisticated relationships to ELS than to electrical power supplies.

5 Managing optical engine with external light source on host board

When the optical engine and external light source are both on the same board, the recommended system management architecture is shown in Figure 8. Enhancements to CMIS to support OE and ELS are specifically designed for use in this architecture.

In this scenario, the host board controller is responsible for "manage" use cases and "control" use cases.

The key features of the architecture are as follows:

- Control of the CW light is the responsibility only of the host board controller.
- The ELS-to-OE connectivity map is the responsibility only of the host board controller.
- OE does not know the identity or capability of ELS that provide its CW light.
- ELS does not know the identify or requirements of OE to which it provides CW light.
- ELS advertises to the host board controller its capabilities to deliver CW light. These capabilities are most likely static, although this is not mandatory.

- OE advertises to the host board controller its requirements for CW light. These may be time-varying with respect to required optical power, for example as the OE ages or its environmental conditions change.

Figure 8 shows the system management architecture for cases where the ELS is on the host board. These are cases 1, 2a, 2b of Figure 2, which are also illustrated as Figure 5b.

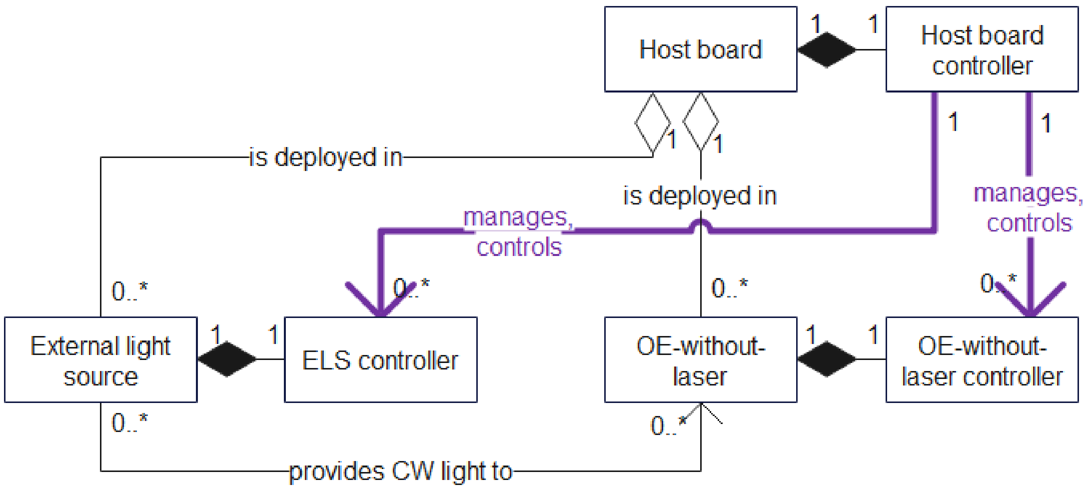

Figure 8: Management architecture for Optical Engines with external light sources on host board

The host board controller contains a map of the ELS/OE relationships, including the identity of ports and CW light. This map may be provisioned by the system designer or system installer. The map can be verified by the host board controller by, for example, varying the optical power from ELS on each CW light in turn, and verifying that optical power changes in the expected manner at the appropriate OE lanes. The signaling mechanisms of the CMIS interface support these capabilities. Alternatively, a similar mechanism can provide automated discovery of the ELS/OE relationship map by the host board controller.

The host board controller knows the advertised capabilities of the ELS and advertised requirements of the OE. These are advertised by the ELS and OE respectively, through the CMIS interface.

The host board controller thus has all the information it needs to control the delivery of CW light. Note however that the fibers and other passive components between ELS and OE are typically not equipped with sensors, actuators or controller and cannot be managed by the host controller.

Figure 8 illustrates the multiplicity in the CW light to OE relationship and can be understood as follows:

- Each CW light is a 1:1 relationship between 1 ELS and 1 OE.
- An ELS may provide separate CW light to multiple OE. This is illustrated as multiplicity "*" at the destination of the CW light relationship.
- An OE may be provided CW light from multiple ELS. This is illustrated as multiplicity "*" at the source of theCW light relationship.
- ELS or OE may be deployed without their companion, and thus there are no CW light provided. These situations are illustrated as multiplicity "0" at respectively the destination and source of the CW light relationship. For example, a user may install ELS or OE that are not immediately connected or placed into service, because they are intended for future growth.

6 Managing optical engine with external shelf laser board

In some cases, the ELS is not deployed on the same physical board as the OE. For example, there may be a shelf laser board containing ELS that provide CW light to OE on several host boards. This is case 3 of Figure 2, which is also illustrated as Figure 5c.

The system management architecture for this scenario is illustrated in Figure 9. The key differences versus Figure 8 are:

- The host board controller is responsible for control use cases.
- The host board controller is responsible for manage use cases of the OE.
- The shelf laser controller is responsible for manage use cases of the ELS.

The system management architecture is otherwise similar to that described in Section 5.

The multiplicity is important. For example, a shelf laser board may provide CW light to more than one host board. A host board may be provided with CW light from more than one shelf laser board. These decisions are at the discretion of the system implementer.

The system implementer should to pay careful attention to the implementation of the communication path between host board controllers and ELS. Factors to be considered include the reliability, latency and addressing of this communication path. Gateways are illustrated in Figure 9, only to assist discussion. The implementation of this communication path is outside the scope of OIF documents, except that the end-points at the ELS and OE must comply with OIF specifications if the modules are claimed to be OIF compliant.

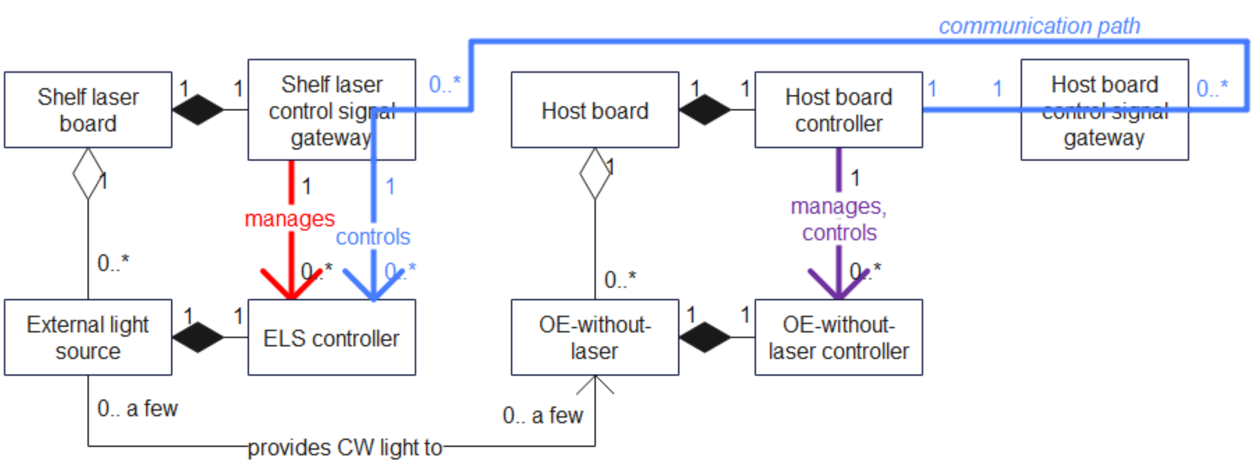

Figure 9: Management architecture for optical engines with external light source on shelf laser board[1]

[1] This diagram is a logical representation only, the control and manage signals may be carried from the "Self laser control signal gateway" to the "ELS controller" over a single physical interface, the control signals may be further processed by the "Self laser control signal gateway".

7 Data Path architecture

A Data Path is a key concept within CMIS. Data Paths represent the states, parameters and behaviors of transceivers.

For an optical engine with internal lasers, it is straightforward to understand how a Data Path is implemented by physical components. When an external light source is used, the implementation of a Data Path is not self-evident. This section describes how Data Paths are constructed, for these two cases.

7.1 Data Path concepts

CMIS section 6.2 defines Data Path as follows, by considering Applications [4]:

- *"An Application is a type of functional configuration that is characterized by specific signal propagation or signal processing between one or more host lanes and one or more media lanes, overall providing a well-defined signal or data transmission function to the host.*
- *An Application instance is one implementation of an Application by a module.*
- *An Application is typically characterized and specified by reference to a pair of industry standards, one for the host interface and one for the media interface, each comprising a number of lanes.*
- *The specific host and media lanes of a module that are used to implement one instance of an Application, together with all required internal module resources, is called the Data Path of that Application instance.*
- *For example, a module may support one 400Gbps Application that is characterized by a 400GAUI-8 host interface and a 400GBASE-DR4 media interface combination."*

For clarity, a complete end-to-end link needs compatible Applications and Data Paths on at least two modules.

The Data Path behavior is represented as a Data Path State Machine (DPSM), whose states are specified in CMIS [4].

CMIS describes the module behavior and Data Path State Machine as follows [4]:

- *"Fundamental power up, initialization, and reinitialization interactions between host and module are governed and described by state machine based behavioral models. Conceptually, these state machines are considered parts of the module.*
- *Note: These state machines are purely conceptual, precise models to specify required host and module interactions and behaviors associated with those interactions; they do not constrain software implementation.*
- *State machines describe both autonomous behavior (i.e. what the module does) and reactive behavior (i.e. how the module reacts to events caused by the host). In certain situations, the states of the state machines in the module cannot be observed by the host, but they still govern behavior and reactions of the module in these situations."*

Thus, when the host board controller wishes to act on an instance of an Application instance, it controls the module through the CMIS interface of the module. For example, the host board controller may wish to turn on or turn off the transmitter(s) of a Data Path. It would therefore instruct the module to change the state in the Data Path State Machine. The module then performs Reactive Behavior in response to the host action.

Alternatively, the Data Path may itself decide to act on a physical parameter within the module, to keep the transceiver performance within targets. The module performs Autonomous Behavior. For example, the Data Path may adjust the bias current of a laser or the settings of the equalizer taps in a receiver.

7.2 Optical engine with internal laser

In the case of an optical engine with internal laser, a Data Path is very similar to that of a traditional pluggable module.

The Data Path is owned and fully implemented by one optical engine. This is shown at a high level in Figure 10, and in more detail in Figure 11. An Application is therefore fully implemented within the optical engine. Therefore, there are no new architectural considerations, for this case.

In particular, consider a laser whose output power is determined by its bias current. If the Data Path needs to adjust the optical power of the laser (either autonomously or reactively), then the module directly sets the appropriate bias current within itself.

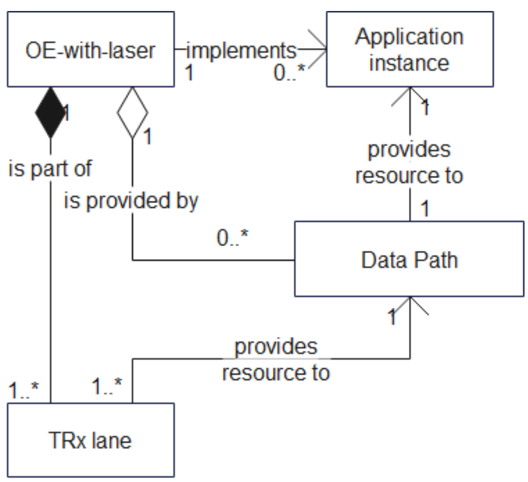

Figure 10: Data Path, using optical engine with internal laser

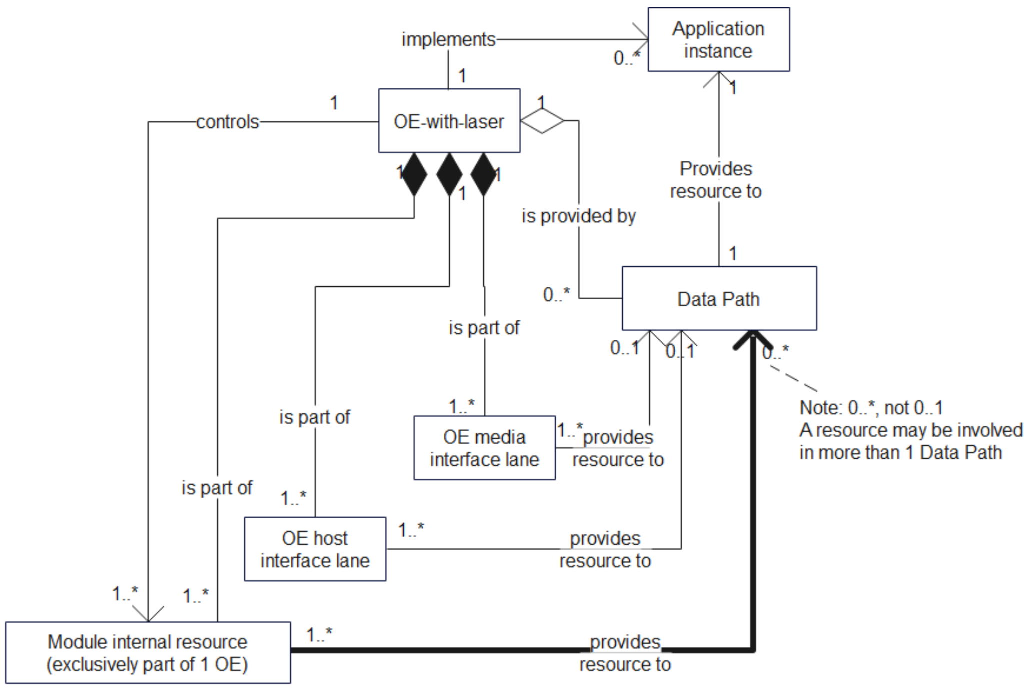

Figure 11: Data Path, using optical engine with internal laser – detailed view

7.3 Optical engine with external light source

In the case of an optical engine with external laser, the situation is more subtle. Now, the laser resources needed by the Data Path are outside the optical engine.

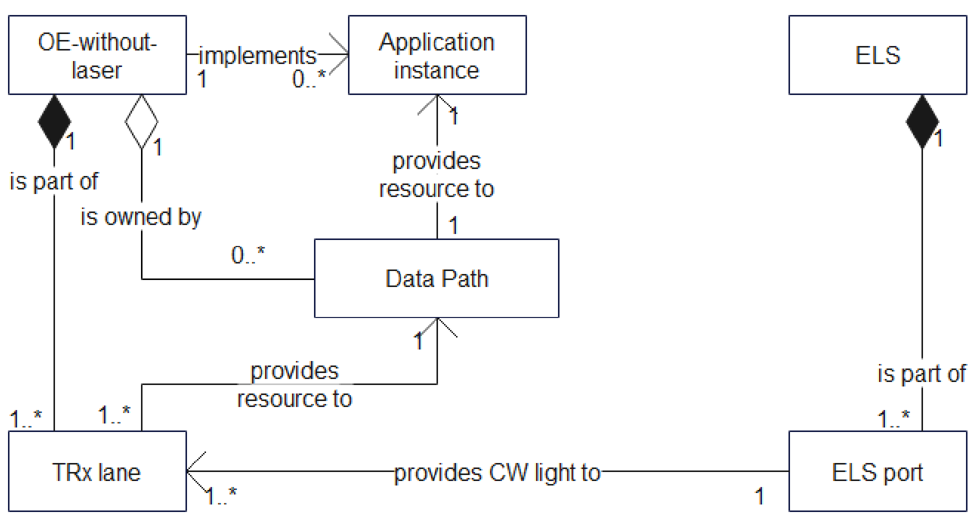

Figure 12: Creation of a Data Path, using optical engine and external light source

OIF considered two potential solutions:

- Optical engine implements Application and owns Data Path. Laser is not part of Data Path. (chosen architecture)
    - This architecture was chosen because it is consistent with traditional practice for everything except the lasers.
- Host board controller owns Application and Data Path (rejected architecture)
    - This architecture was rejected because it would move the Data Path State Machine out of the module, which would be a major change to traditional practice.

The chosen system architecture is shown in **Error! Reference source not found.**.

8 Example behaviors when using external light source

When using an external light source, the majority of the resources of a Data Path are still in the optical engine – for example: electrical data I/O, optical modulator driver, optical modulator, photodetector, optical receiver. Thus, most of the Data Path behavior implementation is within the module. Hence the majority of management behavior is not changed from traditional optical modules.

The special situation is how CW light instances are controlled. We illustrate the recommended system management architecture by means of examples.

These examples show high-level behavior, to illustrate overall concepts. The figures are conceptual and does not show details of communication or implementation. The CMIS specification provides more detailed flows, specifies which capabilities are mandatory / optional, and specifies default behaviors.

We emphasize that the OE does not instruct the ELS. Instead, the OE makes information available to a higher-up management entity (the host board controller) that specifies what the OE needs, and the host board controller coordinates external resources, such as required laser power level, to meet that need if possible.

8.1 Example use case 1: Optical engine Tx optical output is becoming too low

An OE is operating normally, with a given Data Path in a DataPathActivated [5] state. The OE then observes that its Tx optical signal level on a lane is approaching an action threshold, while the CW light power monitored at the OE remains within its target. This may occur due to aging of the OE or due to environmental conditions at the OE. The threshold is vendor-specific, and is chosen by the vendor to prevent the optical Tx performance from falling outside OIF-compliant values.

The OE therefore decides that it wants more CW light power on that lane, so that it can move the Tx optical output closer to nominal levels. The activity diagram is illustrated at a high level in Figure 13.

To effect a change in the CW light (typically a change in optical power), the optical engine advertises its requirements on its CMIS interface. The host board control reacts to this updated advertisement by:

- Identifying the related ELS resources (external light source instance(s) and external light source port instance(s)) that will be needed to handle the OE advertised requirement.

- Verifying that the advertised capabilities of the related ELS resources can meet the advertised requirements of the OE.
- Instructing the related ELS to set the required CW light parameters (typically, the optical power).

If the advertised capabilities of the related ELS resource cannot meet the advertised requirements of the OE, the host board controller should inform that OE that its advertised requirements cannot be met. The OE can then react in a similar manner as would an OE-with-laser that had a fault in an internal laser. For example, the OE may disable a Data Path and raise an alarm.

In this example, the OE remains in the DataPathActivated [5] state during the change. Alternatively, implementors may choose to implement a conservative approach wherein the Data Path is deactivated during the change in optical power.

OIF

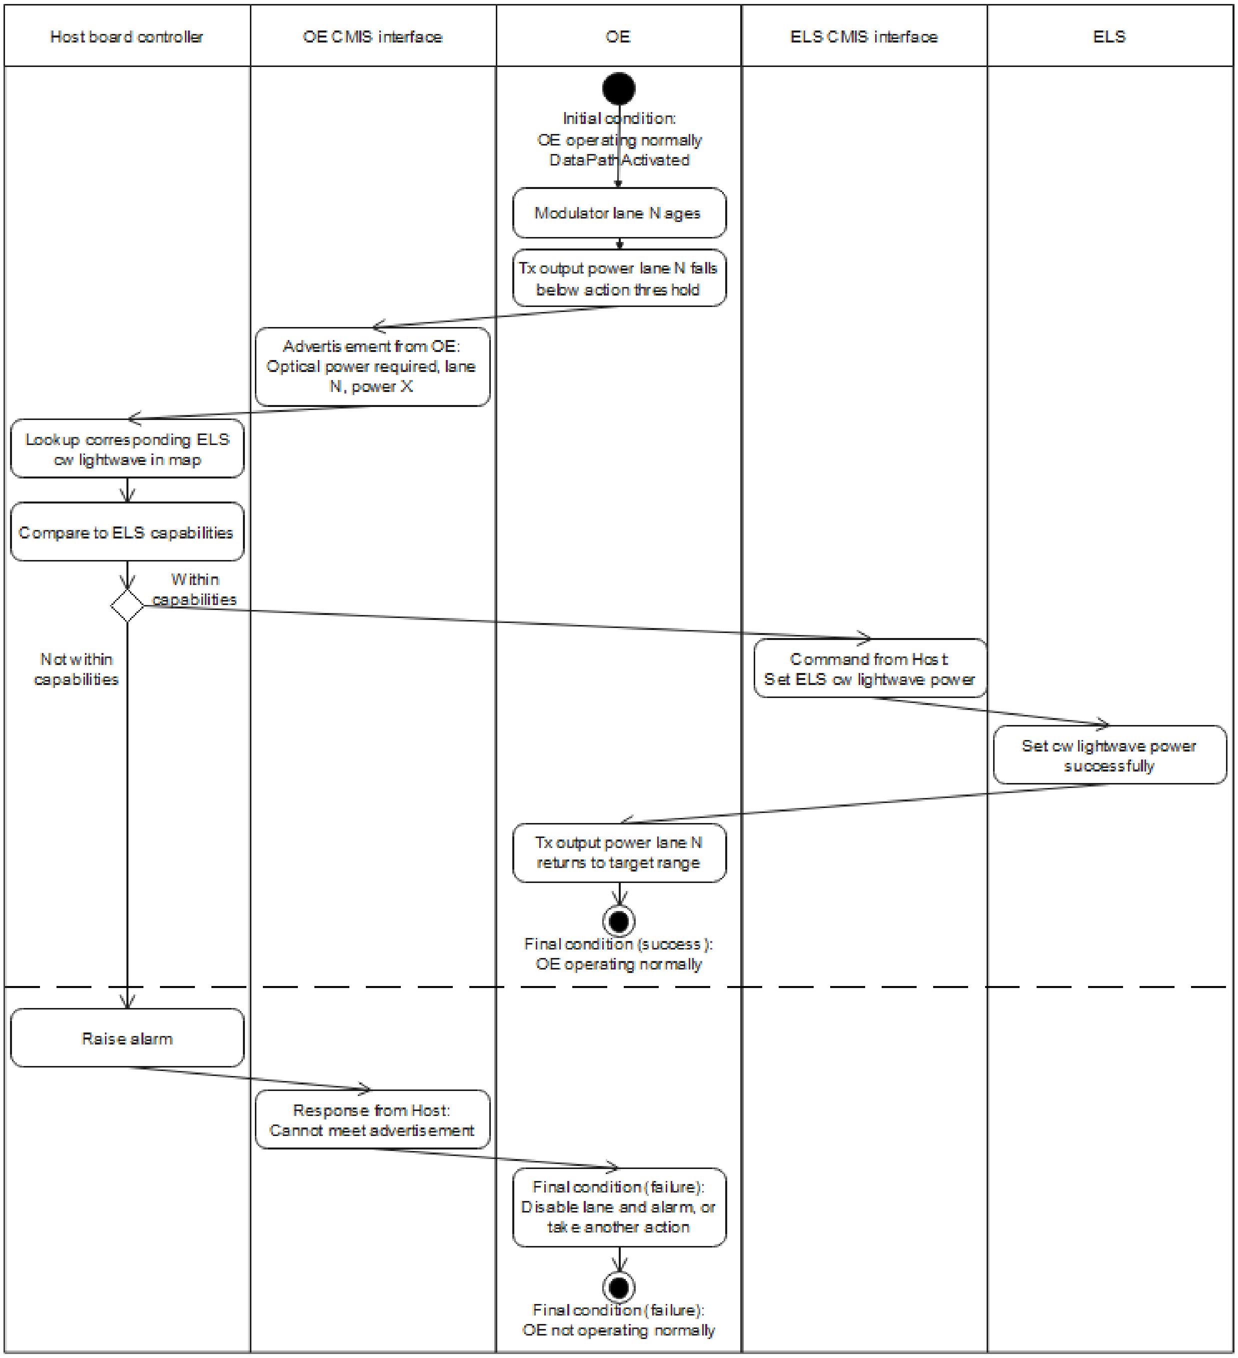

Figure 13: Example use case: OE needs more CW light power due to OE aging

8.2 Example use case: Fault in power delivery fiber

Failure path scenarios can be handled in a similar manner.

For example: The laser output power degrades due to a fault in the power delivery fiber. Such a fault is not generally detectable by the ELS. The optical engine, by means of its input laser power monitor, observes that the laser power delivered to it is not within its previously-advertised requirements. The optical engine raises an alarm to be read by the host board controller. The host board controller then decides what to do. For example, it may instruct the Data Path and the related ELS resources to be turned off, and raise alarms to trigger action by maintenance personnel.

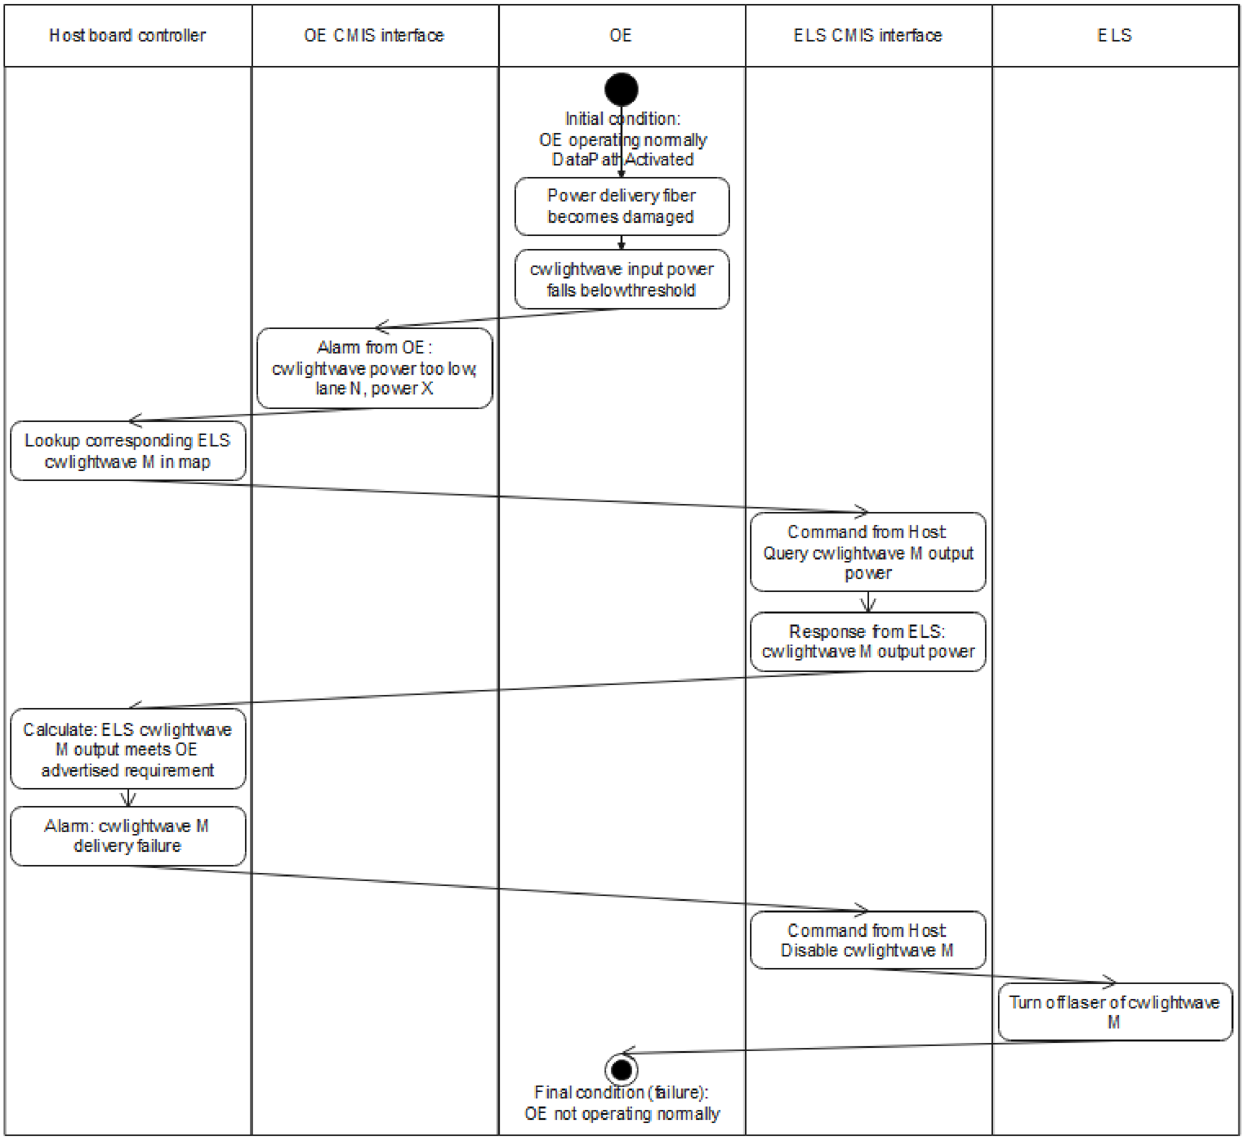

Figure 14: Example use case: Fault in power delivery fiber

9 Summary

This White Paper has presented a system management architecture that can conveniently manage optical engines, both those with internal lasers, and those that need external light sources. Key to external laser source (ELS) management is that the intelligence for controlling the CW light sources is in the host board controller, not in the optical modules.

OIF implementation agreements and specifications are optimized for use with this system management architecture. The architecture minimizes the management complexity that must be embedded within optical modules, while allowing a great variety of system deployment possibilities for future-proofing.

10 References to OIF implementation agreements

[1] OIF-Co-Packaging-FD-01.0 – Co-Packaging Framework Document, February 2022.

[2] <reference OIF OE IA when published>

[3] <reference OIF ELSFP IA when published>

[4] <reference OIF CMIS 5.3 when published>

[5] Common Management Interface Specification (CMIS) Revision 5.2, April 27, 2022, Section 6.3.3.9

11 Acknowledgments

Some diagrams in this White Paper use the Universal Modeling Language (UML®) notation. UML® is a registered trademark of the Object Management Group, Inc. The Object Management Group® Standards Development Organization (OMG® SDO) is an international (27 countries), membership-driven (230+ organizations) and not-for-profit consortium.

The symbols used in this document are using the class diagram UML representation. The notation is described in Figure 15.

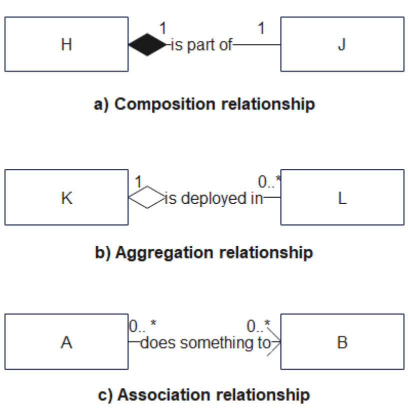

Figure 15 Legend for UML diagram used in this document