



European Organisation for Astronomical Research in the Southern Hemisphere

Programme: ELT

Project/WP: Core Integration Infrastructure / CCS

ELT CII - User Manual

Document Number: ESO-XXXXXX

Document Version: 1

Document Type: Manual (MAN)

Released on: -

Document Classification: Public

Owner: Schilling, Marcus

Validated by PM:

Validated by SE:

Validated by PE:

Approved by PGM:

Name



Authors

Name	Affiliation
Sekoranja, Matej, et al.	Cosylab
Schilling, Marcus, et al.	ESO/DOE/CSE

Change Record from previous Version

Affected Section(s)	Changes / Reason / Remarks
All	All sections updated

CONTENTS:

1	Introduction	7
1.1	Scope	8
1.2	Overview	9
2	MAL API	11
2.1	Overview	14
2.2	Installation	15
2.3	Introduction	16
2.4	CiiFactory	17
2.5	Publish-subscribe	20
2.6	Request-response	36
2.7	Entities	62
2.8	Appendix	64
3	MAL ICD Generation	66
3.1	Overview	69
3.2	Installation	70
3.3	Editing ICD file in Eclipse	71
3.4	ICD Building blocks	73
3.5	ICD Generation	83
3.6	Topic and interface definition document	84
4	MAL Mappings	92
4.1	Overview	95
4.2	Installation	96
4.3	C++	97
4.4	Java	105
4.5	Python	106
5	MAL Python Mappings	107
5.1	Introduction	109
5.2	Prerequisites	110
5.3	Using Python MAL API	111



6	Configuration	120
6.1	Overview	122
6.2	Includes and imports	123
6.3	API Components	124
6.4	Creation of Documents	125
6.5	Access to YAML source tree of a document	130
6.6	Document instance interface	131
6.7	Document Validation	135
6.8	Saving documents to YAML	137
6.9	Merging documents	139
6.10	Exceptions	141
6.11	YAML tags recognized by config-ng	142
6.12	List of built-in data types	144
7	Online Database	146
7.1	Overview	149
7.2	Introduction	150
7.3	Prerequisites	152
7.4	OLDB Library Usage	156
7.5	Advanced Topics	192
7.6	OLDB API LIBRARY	199
7.7	OLDB CLI	206
7.8	GUI	213
7.9	Data Point Value Type to Language Types Mapping	226
7.10	Default Metadata Instance Names	227
7.11	Metadata attribute constraints	228
7.12	OLDB API LISTING	229
8	Error Handling	238
8.1	Overview	240
8.2	Introduction	241
8.3	Installation	246
8.4	Usage	247
8.5	Using the QT Error Dialog Widget	261
8.6	Indexing service	262
8.7	API	266
8.8	Serialization code example	271
9	Logging	277
9.1	Overview	281
9.2	Includes and imports	282
9.3	Introduction	283
9.4	Prerequisites	288
9.5	Logging Library Usage	292
9.6	Logging services configuration	325
9.7	Logging User Tools	343



9.8 Tracing 359

9.9 CII Log Client API 377

9.10 CII Log Fields 396

10 Alarm 397

10.1 Overview 400

10.2 Introduction 401

10.3 Prerequisites 402

10.4 Usage 403

10.5 Advanced Topics 408

11 Telemetry 414

11.1 Overview 417

11.2 Introduction 418

11.3 Prerequisites 426

11.4 Data Capture Configuration 427

11.5 Telemetry Archiver API Usage 429

11.6 Telemetry CLI tools 457

11.7 Telemetry Subscription 463

11.8 Telemetry Archiver Deployment 469

11.9 Advanced Topics 472

11.10 Archive API 474

11.11 Data Capture Configuration Options 477

11.12 Telemetry Archiver Configuration 478

11.13 Big Data Storage Examples 480

12 Service Management 487

12.1 Overview 488

12.2 Introduction 489

12.3 cii-postinstall 490

12.4 cii-services 492

12.5 Example 495

13 Internal Config System 498

13.1 Overview 501

13.2 Introduction 502

13.3 Prerequisites 512

13.4 Basic usage 514

13.5 Advanced usage 531

13.6 Additional information 541

13.7 GUI 548

13.8 Config client API listing 568

13.9 Config point value type to language types mapping 574

13.10 Default metadata instance mapping 575

13.11 Class definition reserved words 576

13.12 JSON/YAML Schema 577



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 6 of 589

13.13 GUI configuration file	580
13.14 Configuration client settings	581
13.15 Configuration service settings	582
13.16 Code produced by the generators	584

**CHAPTER
ONE**

INTRODUCTION

The CII is the Common Integration Infrastructure of the ELT control Systemm, used within the Central Control System.



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 8 of 589

1.1 Scope

This document is the user manual for the ELT CII. The intended audience are ELT CCS team, IFW, WFC and AIV commissioning team.



1.2 Overview

List of CII user manuals:

- **Middleware Abstraction Layer:** The Middleware Abstraction Layer (MAL) is a part of Core Integration Infrastructure (CII) and serves to decouple the CII applications from the communication middleware, permitting transparent functioning over multiple middleware stacks as well easing the substitution of one middleware implementation for another.
 - *MAL API:* The applications do not need to know any specifics of the various middleware software, MAL hides all the specifics via its API. The MAL API is an API only; it does not contain any actual mapping or dependency to any communication middleware product. MAL supports two standardized communication patterns on top of which all applications shall be built. These are: - Data centric publish-subscribe: decoupled peers based on data producers and consumers. - Request-response: coupled, state-full, peer-to-peer communication.
 - *MAL Mappings:* Describes configuration options for the following MAL mappings: DDS, ZPB and OPC UA.
 - *MAL Python Mappings:* This document provides information on use of Python MAL mappings interface also referred to as Python MAL API. Python MAL API builds on MAL C++ API foundation. It provides programming interface that is very similar to the programming interface that is exposed by MAL C++ API.
 - *MAL ICD Generation:* Manual for developing Interface Control Documents, specifically Type Definition Documents holding ICD types.
- *Configuration:* Configuration provides the mechanisms for distributing configuration data to the control system application software and to allow for the collection of configuration data from these systems.
- *Online Database* The Core Integration Infrastructure (CII) Online Database (OLDB) provides distributed data publishing and access to actual or live data for user interface and control applications that do not have low-latency or real-time performance requirements. The term “online” refers to the fact that the database provides current and live values for the monitor points of the control system.
- *Error Handling:* Error Handling in CII provides the APIs and tools necessary to conveniently implement a common error handling mechanism across all EELT Control System applications and thus streamline the diagnosis of abnormal behavior of the system.
- *Logging:* CII Log provides the logging and tracing services of the Control System. Logging pertains to all aspects of CII services. Its purpose is to provide basic and detailed information of the different services and systems in the CII in order to have a better understanding of the behaviour of the distributed system, and in particular to troubleshoot the system in case of problems.
- *Alarm:* CII Alarm system uses Integrated Alarm System (IAS) [4] as its alarm system. IAS is a software whose main purpose is to get alarms and monitor points from different sources and



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 10 of 589

generate alarms to present to the users who can range from operators in the control room, up to engineers sitting at their desks. The alarm system is a message-passing facility that routes information about abnormal situations detected from hardware or software to the user.

- *Telemetry*: The purpose of the CII Telemetry Service is archiving the periodic, event-based and ad-hoc data that will be used for calculation of statistics, and the long-term trends and performance of the ELT telescope.

MAL API

Revision:	2.9
Status:	Released
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	ELT CII
Folder:	/trunk/deliverables/phase2
Document ID:	CSL-MAN-17-150198
File:	MAN-MAL_API.docx
Owner:	Matej Šekoranja, Cosylab Sweden
Last modification:	March 7, 2023
Created:	October 25, 2017

Prepared by	Reviewed by	Approved by
Matej Šekoranja, CSL SWE	Aljaž Podboršek Gregor Čuk	Gregor Čuk

Document History



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 12 of 589

Revision	Date	Changed/reviewed	Section(s)	Modification
0.1	2017-10-25	msekoranja	All	Created.
1.0	2017-11-06	msekoranja	All	Update on API, release.
1.1	2018-02-23	msekoranja	3.1, 3.3	Update on API changes, C++ MrvSubscriber example.
2.0	2019-02-22	msekoranja	All	Complete rewrite of a document.
2.1	2019-03-04	msekoranja	All	Example code revised.
2.2	2022-03-31	jpribosek	2.5.2	Example code revised.
2.3	2022-04-21	dkumar	2.6.2	Exception handling details
2.4	2022-12-14	mschilli	2.7 - 2.8	New section on API of ICD entity
2.5	2022-12-27	dkumar	2.7	Eigen example
2.6	2023-01-11	jrepinc	2.	Python subscription example update/note
2.7	2023-02-17	dkumar	Entites	Parent child relation when cloning
2.8	2023-02-17	dkumar	Subscriber	Subscriber's close guaranty
2.9	2023-03-07	dkumar	Subscriber	Subscriber's method call warning

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is a Middleware Abstraction Layer user manual document for the ELT Core Integration Infrastructure Software project.

Audience

Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 13 of 589

API	Application Programming Interface
MAL	Middleware Abstraction Layer
QoS	Quality of Service
CII	Core Integration Infrastructure
ICD	Interface Control Document
XML	Extensible Markup Language

References

1. Cosylab, Middleware Abstraction Layer Design document, CSL-DOC-17-147260 v1.5
2. Cosylab, ELT CII MAL Transfer document, CSL-DOC-18-168015, version 1.8
3. Cosylab, Data Addressing Specification, CSL-DOC-17-147264, version 1.2



2.1 Overview

This document is a user manual for MAL API. The document provides descriptions of the API and instructions to programmers how to use MAL API. All the details on the API and its functionality are covered by the MAL design document [1].



2.2 Installation

See ELT CII MAL Transfer document [2], Chapter 1.2 on how to build and install MAL.



2.3 Introduction

MAL is an abstraction API that provides complete independence from the communication middleware software. This is achieved by dynamically loadable MAL middleware implementation (mapping) that implements MAL API. In addition, data entities (classes, structures, interfaces) need to be independent from the communication-middleware too. This is achieved by using generation: user defines its data entities in ICD definition language (i.e. XML), the generator then produces middleware-agnostic entities that are used by the applications and middleware-specific code that needs to be generated. The specific code is dynamically loaded by the MAL mapping implementations when needed. This part is completely hidden from the programmer. The programmer must never use or reference any middle-ware specific code it only uses pure MAL API and agnostic data entities.

Middleware selection is done using Uniform Resource Identifier (URI) addressing. A URI is a string of characters that unambiguously identifies a particular resource. The scheme part of the URI (e.g. "dds.ps://" prefix) determines what MAL mapping to use, whereas the rest of the URI is completely middleware specific. For more info on the URIs please refer to the Data Addressing Specification ([3]).

MAL API supports two standardized communication patterns: data centric publish-subscribe (decoupled peers based on data publishers and subscribers) and request-response (coupled peer-to-peer client-server communication).

The following sections provide the description and examples of the API for both communication patterns. The examples demonstrate usage of the API, they do not do anything meaningful. Source code of the examples is available in `elt-mal/mal` module.



2.4 CiiFactory

CiiFactory is an entry-point class of the MAL API. Usually it is accessed as a singleton instance via `getInstance()` method. This provided a shared, thread-safe instance to be used among the threads in one process. However, in case of a more complex application where an insulation among different “sub-processes” is required, the API provides a `createInstance()` method that instantiates a completely new CiiFactory instance on every call of this method. Avoid using this method unless there is a good reason for it.

A CiiFactory instance has 2 groups of methods: methods for MAL mapping implementation (un-)registration and factory methods of entities (publisher, subscriber; client, server) for both communication patterns. The main task of CiiFactory is to delegate factory method calls to the appropriate registered MAL mapping implementation depending on the URI, given as a parameter to all the factory methods. CiiFactory is instantiated without any MAL mapping registered. A programmer needs to register MAL mappings. The registration is done by calling `registerMal` method that takes 2 parameters: a reference to the instance of MAL mapping implementation and an URI schema string to be handled by the implementation. CiiFactory takes no responsibilities on registered mappings lifecycle; it is responsibility of the programmer to take care of it, still CiiFactory provides a `close()` method that calls `close()` on all registered MAL mappings and unregisters them. A programmer has also an ability to unregister previously registered MAL mappings. Un-registration only removes a MAL instance from CiiFactory dictionary; any previously created entities (e.g. publishers) will remain active. The entities lifecycle is responsibility of the MAL mapping implementation, meaning releasing a MAL mapping implementation also releases all its entities.

Registration of a MAL mapping introduces a compile and runtime dependency to the mapping. `CiiFactory.loadMal` method dynamically loads the mapping and avoids any dependency on any MAL mapping. `CiiFactory.installMal` calls `CiiFactory.loadMal` and registers newly loaded mapping to the shared CiiFactory instance. Optional properties parameter can be passed as parameter to configure loaded MAL. Both methods accept MAL mapping (schema) name, e.g.: `dds`, `zpb`, `opcua`.

All factory methods accept 3 common parameters: a URI, a quality-of-service (QoS) parameters and a string name-value list of MAL mapping specific properties. MAL API defines a common set of QoS parameters per communication pattern (described in detail in the following sections). If none (null) are supplied defaults are used.

The list of MAL mapping specific properties allows specific/fine tuning of the middleware in a uniform way. The properties in a list are specific and documented per MAL mapping implementation, usually they are provided via some kind of configuration subsystem. If none (null) are supplied defaults are used.

A registration of another MAL instance with the same schema string will override any previous registration (un-registration no longer needed). Un-registration of non-registered scheme results in no operation action. A call to a factory method with a URI scheme that is not registered will result in `SchemeNotSupportedException` exception.

The following code examples show recommended initialization of MAL mapping with creation of a publisher.



2.4.1 Java

In Java it is highly recommended to use try-with-resources statements so that the code is safe and close() on the resource will be always called.

```
// install DDS MAL mapping
CiiFactory.installMal("dds");

// use try-with-resources statement to be safe that close()
// is always being called on factory and MAL closed
try (CiiFactory factory = CiiFactory.getInstance()) {

    // create URI
    URI uri = URI.create("dds.ps:///m1/CabinetTelemetry");

    // create publisher instance with default QoS and no specific properties
    // use try-with-resources statement to be safe that close()
    // is always being called on publisher
    try (Publisher<Sample> publisher =
        factory.getPublisher(uri, QoS.DEFAULT, null, Sample.class)) {
        // put publish code here ...
    }
}
}
```

2.4.2 C++

C++ implementation uses RAII idiom using smart pointers. Resources are closed automatically when their enclosing scope ends. To close resource explicitly call close() method.

```
try {

    // load DDS MAL mapping with (optional) properties
    std::shared_ptr<::elt::mal::Mal> ddsMal = mal::loadMal("dds",
↳mal::Mal::Properties{
    {"dds.domain", "100"}});
    ::elt::mal::CiiFactory &factory = mal::CiiFactory::getInstance();

    // register DDS MAL mapping with CiiFactory
    factory.registerMal("dds", ddsMal);

    // create URI
    mal::Uri uri("dds.ps:///m1/CabinetTelemetry");

    // create publisher instance with default QoS and no specific properties
    std::unique_ptr<ps::Publisher<mal::example::Sample>> publisher = factory.
↳getPublisher<mal::example::Sample>(
    uri,
```

(continues on next page)



(continued from previous page)

```
mal::ps::qos::QoS::DEFAULT,  
    {});  
    // put publisher code here...  
  
} catch (std::exception& exc) {  
    std::cerr << "EXCEPTION: " << exc.what() << std::endl;  
}
```

2.4.3 Python

In Python all generated entities are prefixed with a “Mod”<name of the ICD file in camelCase> string. The module also loads appropriate C++ binding classes and libraries. Submodule structure follows ICD package structure with capitalized name for each package.

Python MAL API supports context manager protocol to automatically release the resources after use. Closing resources explicitly with close() is also supported.

```
# Import MAL API.  
import elt.pymal as mal  
  
# Import Data Entity class from binding module.  
from ModManualExamples.Elt.Mal import Example  
  
uri = 'dds.ps:///ml/CabinetTelemetry'  
factory = CiiFactory.getInstance()  
  
# Load DDS MAL mapping.  
ddsMal = mal.loadMal("dds", {'dds.domain': '100'})  
factory.registerMal('dds', ddsMal)  
  
# Use context manager interface of the publisher  
# to ensure underlying instance of the publisher is  
# closed when context is destroyed.  
# Create publisher instance with default QoS and no  
# specific properties.  
with factory.getPublisher(uri,  
                          Example.Sample,  
                          qos=mal.ps.qos.DEFAULT) as publisher:  
    # Put publish code here...  
  
# Publisher is automatically closed at this point.
```



2.5 Publish-subscribe

Publish-subscribe is a data-centric decoupled communication pattern. Publishers (data producers) publish data to a topic, and subscribers (data consumers) subscribe to the topics they are interested in to receive the data published. This is not a coupled (connected) peer-to-peer communication – there can be many subscribers per one topic. MAL prescribes a limit of only one publisher per topic. Order of instantiation/startup also does not matter. Subscriber will miss any data sent to the channel prior their subscription. Subscribers can subscribe to a topic even if there is no publisher available. Once it became available the subscribers will start receiving the newly published data. The paradigm does not guarantee that the first n-messages are will be received by the subscribers, even if the subscribers are already running when the publisher starts. It takes some time to establish communication mechanism on startup of the publishers and subscribers leading to non-instant subscriptions to the topics and consequently missed data. In addition, MAL does not guarantee reliable communication (and queueing) – the data might be lost during the transmission or dropped if send/receive queues are full. MAL mapping specific properties can be used to set reliability QoS, if underlying communication middleware supports it.

MAL however does support 2 QoS settings:

- Deadline

The deadline QoS shall be the maximum age of a sample, i.e. $t(\text{now}) - t(\text{sample}) < \text{deadline_time}$, sample time is a time of creation of a sample. The sample time is given as parameter to the publish method call on publisher. If not given, current time is used. Note that this also implies that if a sample is left too long in the subscribers receive queue it might expire (and is removed from the queue).

If deadline QoS is not defined, it is assumed to be infinite.

- Latency

The latency QoS shall be the maximum time a sample may remain in-transit between the publisher and subscriber. That is, the arrival time minus sample time cannot be greater than latency QoS. If latency QoS is not defined, it is assumed to be infinite.

If QoS is violated a corresponding sample is discarded.

It is assumed that time among different network nodes where publishers and subscribers run is synchronized.

MAL also supports a concept of instances. Every structure (defined in ICD) can be assigned a key. A key is a set of structure fields (of primitive types) that uniquely identifies one instance. This allows programmers to publish/retrieve/filter data based on their instance.

To create data entities use `createDataEntity` methods available on both, publisher and subscriber.



2.5.1 Publisher

Publisher has only one task – to publish data within one topic. There are 2 publish methods available: one with and one without sample timestamp. In case of method without a source timestamp current time is used as a sample timestamp.

Timestamp is a double value representing seconds past 1.1.1970 UTC (to micro precision). There are utility classes that help conversion from/to standardized timestamp values/structures.

Publish methods accept timeout parameter. The method blocks until the message is sent. If the method execution time exceeds timeout time limit TimeoutException is thrown. A non-blocking behavior can be achieved by using timeout value of zero.

Depending on whether data structure contains a key there are two ways of publishing, as shown on the following examples.

This class is thread-safe.

Java

```
// install DDS MAL mapping
CiiFactory.installMal("dds");

try (CiiFactory factory = CiiFactory.getInstance()) {

    URI uri = URI.create("dds.ps:///m1/CabinetTelemetry");
    // create publisher with deadline and latency QoS
    try (Publisher<Sample> publisher =
        factory.getPublisher(uri,
            new QoS[] {
                new Deadline(1, TimeUnit.SECONDS),
                new Latency(100, TimeUnit.MILLISECONDS)
            }, null, Sample.class)) {

        //
        // publish code for keyless topic
        // and user provided source timestamp
        //
        Sample sample = publisher.createDataEntity();
        sample.setDaqId(0);
        sample.setValue(12.8);
        // MalUtil.timestampToDouble(long millis) from mal-common
        // might be helpful to convert Java timestamp from millis
        double timestamp = 1550791384.1d;
        publisher.publish(timestamp, sample,
            DEFAULT_TIMEOUT_SEC, TimeUnit.SECONDS);
    }
}
```

(continues on next page)



(continued from previous page)

```
//  
// publish code for keyed topic  
//  
Sample keySample = publisher.createDataEntity();  
// set key field(s)  
keySample.setDaqId(1);  
  
try (InstancePublisher<Sample> daq1Publisher =  
    publisher.createInstancePublisher(keySample)) {  
  
    Sample daq1Sample = daq1Publisher.createDataEntity();  
    // key is already set when createDataEntity is  
    // called on InstancePublisher  
    daq1Sample.setValue(Math.PI);  
    daq1Publisher.publish(daq1Sample, DEFAULT_TIMEOUT_SEC, TimeUnit.SECONDS);  
  
}  
// daq1Publisher is destroyed here,  
// since we are out of try-with-resource scope  
  
}  
  
}
```

C++

```
try {  
    std::shared_ptr<::elt::mal::Mal> ddsMal = mal::loadMal("dds",  
↳mal::Mal::Properties{  
    {"dds.domain", "100"}});  
    ::elt::mal::CiiFactory &factory = mal::CiiFactory::getInstance();  
    factory.registerMal("dds", ddsMal);  
  
    mal::Uri uri("dds.ps:///m1/CabinetTelemetry");  
    // Create publisher with deadline and latency QoS  
    std::unique_ptr<ps::Publisher<mal::example::Sample>> publisher = factory.  
↳getPublisher<mal::example::Sample>(  
        uri,  
        { std::make_shared<mal::ps::qos::Latency>(  
            std::chrono::milliseconds(100)),  
          std::make_shared<mal::ps::qos::Deadline>(  
            std::chrono::seconds(1)) }, {});  
  
    //  
    // publish code for keyless topic  
    // and user provided source timestamp
```

(continues on next page)



(continued from previous page)

```
//
std::shared_ptr<mal::example::Sample> sample = publisher->createDataEntity();
sample->setDaqId(0);
sample->setValue(12.8);
double timestamp = 1550791384.1;
publisher->publish(timestamp, *sample, std::chrono::seconds(3));

//
// publish code for keyed topic
//
std::shared_ptr<mal::example::Sample> keySample = publisher->
↪createDataEntity();
// set key field(s)
keySample->setDaqId(1);
{
    std::unique_ptr<::elt::mal::ps::InstancePublisher<mal::example::Sample>>
↪daqlPublisher =
        publisher->createInstancePublisher(*keySample);
    std::shared_ptr<mal::example::Sample> daqlSample = daqlPublisher->
↪createDataEntity();
    // daqId is already set as it is created by InstancePublisher
    daqlSample->setValue(3);
    daqlPublisher->publish(*daqlSample, std::chrono::seconds(3));
}
// daqlPublisher is destroyed here, since its enclosing scope ended
} catch (std::exception& exc) {
    std::cerr << "EXCEPTION: " << exc.what() << std::endl;
}
}
```

Python

Python does not support method overloading. Publisher provides `publish()` and `publishWithTimestamp()` methods. `publishWithTimestamp()` takes `timestamp` value as first argument. Type of the `timestamp` value must be float (as returned from `time.time()`).

Durations and timeouts must be specified as instances of standard `datetime.timedelta` class.

```
uri = 'dds.ps://m1/CabinetTelemetry'

ddsMal = mal.loadMal('dds', {'dds.domain': '100'})
factory = mal.CiiFactory.getInstance()
factory.registerMal('dds', ddsMal)

# Create publisher with deadline and latency QoS.
qosList = [mal.ps.qos.Deadline(datetime.timedelta(seconds=1)),
           mal.ps.qos.Latency(datetime.timedelta(milliseconds=100))]
with factory.getPublisher(uri, Example.Sample,
```

(continues on next page)



(continued from previous page)

```
                                qos=qosList) as publisher:
#
# Publish code for keyless topic
# and user provided source timestamp.
#

sample = publisher.createDataEntity()
sample.setDaqId(0)
sample.setValue(12.8)
timestamp = time.time()
publisher.publishWithTimestamp(sample,
                                timestamp,
                                datetime.timedelta(seconds=1))

#
# Publish code for keyed topic.
#

keySample = publisher.createDataEntity()

# Set key field(s).
keySample.setDaqId(1)

# Use Publisher as context manager.
with publisher.createInstancePublisher(keySample) as daq1Publisher:
    daq1Sample = daq1Publisher.createDataEntity()
    daq1Sample.setValue(3.14)
    daq1Publisher.publish(daq1Sample, DEFAULT_TIMEOUT)
# daq1Publisher is closed here.
```

2.5.2 Subscriber

Each subscriber instance has its own queue. The queue holds the events related to the subscribed topic ordered by the time of insertion into the queue (e.g. data arrival). The events are usually data events, they simply contain the data published, however there are also events that indicate e.g. that there is no publisher for given instance available. The events are represented using DataEvent interface.

The interface via its methods provides access the data (if available), its source timestamp, InstanceState and InstanceEvent enumerations. InstanceState always reports current (not the event was recorded) state of an instance: ALIVE or NOT_ALIVE. InstanceEvent describes the reason why event was created: CREATED (instance has been created/revived), UPDATED (updated data event), REMOVED (instance has been destroyed, i.e. no more publishers available for the instance). A programmer should never assume that data is always available, e.g. in case of REMOVED InstanceEvent there is no data. A programmer should always call hasValidData method to check whether the event holds (valid) data.



NOTE: REMOVED instance event is not guaranteed to be generated when instance is being destroyed (due to differences, or lack of notification, in middleware communication software), therefore software should not rely its logic on arrival of this event.

All the data received by subscriber instance (registered to one topic) is first processed by subscriptions on the subscriber. If any of the data matches a subscription a subscription callback is called and the data is provided in the callback, otherwise the data gets into the queue. A programmer has to take care of emptying the queue, otherwise it might get full and no new data can be received. In order to empty queue user need to invoke one of the read methods, or set deadline QoS property to instruct removal of (too) old data from the queue.

NOTE: You should never close the subscription within the subscription callback as that may cause deadlock, depending on the implementation. If you want your callback to stop processing but you don't want to close the subscription yet, you should pass in an atomic variable and early return from the callback depending on the state of the variable.

NOTE: You should never try to close the subscription or call any subscriber method within the subscription callback as that may cause a deadlock. See *one-shot subscriber* example in *icd-demo* module on how to handle this case.

NOTE: Subscription listener's methods are never invoked after a call to close the subscription has successfully completed.

The following text describes some of the use-cases shown in the example code below:

1. Read data from all instances. Maximum number of data is limited by `maxSamples` parameter. To read all the data (and clear the queue) use 0 as `maxSamples` parameter.
2. If there is no data, an empty (not null) list is returned. This method creates new entities and copies their data; avoid using this method when dealing with large arrays.
3. Read data from one instance. The instance to be read is given as parameter. To create an instance parameter call `createDataEntity` and set fields that form a key. If queue for specified instance is empty a null reference is returned. This method creates a new instance (there is an exception in case of C++ by-reference variation the method) and copies its data (in all method variations); avoid using it when dealing with large arrays.

Whenever there is a need to avoid any copies of the data (e.g. large arrays), there are load read methods available. If underlying middleware and mapping implementation support no-copy behavior this methods "loan" the data from the middleware and pass the reference to the data to the programmer. However, this is not always possible. If this is not possible (not supported by the implementation) a copy is made. Also be careful that some middleware might support this only for certain data types: if mapping from middleware data type matches MAL data type and no conversion is needed, then no-copy behavior is possible, otherwise a copy (during conversion) is made.

Synchronous read of events (not just data!) can be done using `readDataEvents` method. The method returns a count of added events to the list provided as an argument. The method blocks until one of the condition is met: number of events added to the list reaches `eventCountLimit` or `collectTime` time is passed. Zero `eventCountLimit` imposes no limit on number of events added, and zero `collectTime` turns method to non-blocking and only reads what is currently available in the queue (still limited by



eventCountLimit count limit). Do not forget that events might not always data event, check by calling hasValidData method if data is valid first.

For asynchronous subscriptions there are 2 methods: one that provides the data in the callback and the one without. The former method guarantees only that data is valid for the time of the callback. The implementation might avoid doing a copy in this case. The later method only notifies about the presence of new data, the data can be read by using one of synchronous read methods later. However, read call must not be invoke inside the callback. If you want to get the data, use the first method. The callback methods are being executed by some MAL mapping implementation or even middleware threads. Creating asynchronous subscription returns a Subscription instance that is used to cancel the subscription.

Lifetime of subscriptions is limited by the lifetime of a subscriber. If subscriber is being closed the subscriptions are also closed.

In addition there is a method that allows event-pump based subscription. Once subscription is being instantiated a poll method needs to be called in order to get the callback invoked. The callback is being invoked by the same thread that calls poll.

Subscriptions and readDataEvents methods support filtering. Filtering can be done on instance, InstanceState and InstanceEvent fields. Setting a filter on a field means selecting/limiting event to set value. Filter on multiple fields implies all fields must match their set value. Default filter allows all the events. A cached (shared) instance it accessed via DataEventFilter::all() method.

This class is thread-safe.

Java

```
public class SubscriberExample {  
  
    // user data event callback function  
    static void dataEventFunction(Subscriber<Sample> subscriber, DataEvent<Sample>_  
↪event) {  
        // NOTE: you should never close the subscription within the callback  
        // as that may cause deadlock.  
        if (event.hasValidData()) {  
            System.out.println(event.getData());  
        }  
    }  
  
    public static void main(String[] args) {  
  
        // install DDS MAL mapping  
        CiiFactory.installMal("dds");  
  
        try (CiiFactory factory = CiiFactory.getInstance()) {  
            URI uri = URI.create("dds.ps:///m1/CabinetTelemetry");  
            try (Subscriber<Sample> subscriber =  
                factory.getSubscriber(uri, QoS.DEFAULT, null, Sample.class)) {
```

(continues on next page)



(continued from previous page)

```
//  
// Read all alive instances w/o sample count limit.  
//  
List<Sample> allAliveInstances = subscriber.read(0);  
  
//  
// Read one instance, i.e. daqId == 1.  
//  
Sample keySample = subscriber.createDataEntity();  
keySample.setDaqId(1);  
Sample daq1Sample = subscriber.readInstance(keySample);  
  
//  
// Loaned (no-copy) example.  
// Loan is returned when close() method is called on LoanedDataEntity  
// or automatically by try-with-resources (as shown below).  
//  
try (LoanedDataEntity<Sample> loanedSample = subscriber.loanedRead()) {  
    Sample sampleData = loanedSample.getData();  
    if (sampleData != null) {  
        // use data here  
    }  
}  
  
//  
// Read data events, filtered by instance and instance state.  
// Block until maxEvents (1000) are read or timeout (3s) occurs.  
//  
final int maxEvents = 1000;  
List<DataEvent<Sample>> events = new ArrayList<DataEvent<Sample>>  
↪(maxEvents);  
DataEventFilter<Sample> filter = new DataEventFilter<Sample>();  
filter.setInstance(keySample);  
filter.setInstanceState(InstanceState.ALIVE);  
filter.setInstanceEvents(  
    EnumSet.of(InstanceEvent.CREATED, InstanceEvent.UPDATED)  
);  
int samplesRead = subscriber.readDataEvents(filter, events,  
    maxEvents, 3, TimeUnit.SECONDS);  
  
//  
// Asynchronous data subscription, no filtering.  
// 1 second sleep added to simulate some work.  
//  
try (Subscription subs = subscriber.subscribeAsync(DataEventFilter.all(),  
    SubscriberExample::dataEventFunction)) {  
    Thread.sleep(TimeUnit.SECONDS.toMillis(1));  
} catch (InterruptedException ie) {
```

(continues on next page)



(continued from previous page)

```
        // noop
    }
    // note that subscription is destroyed when leaving try-with-resources
↪scope.

    //
    // Asynchronous notification subscription, no filtering,
    // lambda callback implementation.
    // DO NOT CALL read methods in the callback!
    // 1 second sleep added to simulate some work.
    //
    try (Subscription subs =
        subscriber.subscribeAsync(DataEventFilter.all(),
            (Subscriber<Sample> sub) -> System.out.println("new data"))
        ) {
        Thread.sleep(TimeUnit.SECONDS.toMillis(1));
    } catch (InterruptedException ie) {
        // noop
    }

    //
    // Synchronous (aka event pump) subscription, no filtering.
    // Poll with 1000 events / 1 second limit, endless loop.
    //
    try (Subscription subs = subscriber.subscribe(
        DataEventFilter.all(),
        SubscriberExample::dataEventFunction)) {
        while (true) {
            subscriber.poll(0, 1, TimeUnit.SECONDS);
            // the call above blocks for max. 1 second or
            // until 1000 events are provided to the callback
        }
    }

    }
}
}
```

C++

```
std::shared_ptr<::elt::mal::Mal> ddsMal = mal::loadMal("dds",
↪mal::Mal::Properties{
    {"dds.domain", "100"}});
::elt::mal::CiiFactory &factory = mal::CiiFactory::getInstance();
factory.registerMal("dds", ddsMal);

mal::Uri uri("dds.ps:///m1/CabinetTelemetry");
```

(continues on next page)



(continued from previous page)

```
try {
    std::unique_ptr<ps::Subscriber<mal::example::Sample>> subscriber = factory.
↳getSubscriber<mal::example::Sample>(
        uri, mal::ps::qos::QoS::DEFAULT, {});

    //
    // Read all alive instances w/o sample count limit.
    //
    std::vector<std::shared_ptr<mal::example::Sample>> allAliveInstances =
↳subscriber->read(0);

    //
    // Read one instance, i.e. daqId == 1.
    //
    std::shared_ptr<mal::example::Sample> keySample = subscriber->
↳createDataEntity();
    keySample->setDaqId(1);
    std::shared_ptr<mal::example::Sample> daq1Sample = subscriber->
↳readInstance(*keySample);

    // ... or by providing sample by reference
    std::shared_ptr<mal::example::Sample> sample = subscriber->createDataEntity();
    bool sampleRead = subscriber->readInstance(*keySample, *sample);

    //
    // Loaned (no-copy) example.
    // Loan is returned when close() method is called on LoanedDataEntity
    // or automatically when scope containing loaned data entity
    // terminates.

    {
        std::shared_ptr<::elt::mal::ps::LoanedDataEntity<mal::example::Sample>>
↳loanedSample = subscriber->loanedRead();
        if (loanedSample) {
            mal::example::Sample &sampleData = loanedSample->getData();
            // use data here
        }
    }

    //
    // Read data events, filtered by instance and instance state.
    // Block until maxEvents are read or timeout occurs.
    //
    using Sample = mal::example::Sample;
    std::vector<std::shared_ptr<mal::ps::DataEvent<Sample>>> events(1000);
    mal::ps::DataEventFilter<Sample> filter;
    filter.setInstance(std::shared_ptr<Sample>(std::move(keySample)));
    filter.setInstanceState(mal::ps::InstanceState::ALIVE);
```

(continues on next page)



(continued from previous page)

```
filter.setInstanceEvents(
    std::set<mal::ps::InstanceEvent>{mal::ps::InstanceEvent::CREATED,
                                     mal::ps::InstanceEvent::UPDATED});
std::size_t samplesRead = subscriber->readDataEvents(
    filter,
    events, events.capacity(),
    std::chrono::seconds(3));

//
// Synchronous (aka event pump) subscription, no filtering.
// Poll for 1 second w/ no event count limit.
//
{
    std::unique_ptr<mal::ps::Subscription> subscription = subscriber->subscribe(
        mal::ps::DataEventFilter<Sample>::all(),
        [](mal::ps::Subscriber<Sample>& subscriber,
           const mal::ps::DataEvent<Sample>& event) -> void {
            // NOTE: you should never close subscription from
            // within a callback, since that may cause deadlock
            std::cout << event.getData() << std::endl;
        });
    subscriber->poll(10, std::chrono::seconds(1));
}
// Note that subscription is destroyed when leaving the scope.

//
// Asynchronous subscription, no filtering.
// 1 second sleep added to simulate some work.
//
{
    std::unique_ptr<mal::ps::Subscription> subscription = subscriber->
↪subscribeAsync(
    mal::ps::DataEventFilter<Sample>::all(),
    [](mal::ps::Subscriber<Sample>& subscriber,
       const mal::ps::DataEvent<Sample>& event) -> void {
        // NOTE: you should never close subscription from
        // within a callback, since that may cause deadlock
        std::cout << event.getData() << std::endl;
    });
    std::this_thread::sleep_for(std::chrono::seconds(1));
}

//
// Asynchronous notification subscription, no filtering.
// In callback sync read is called to read all alive samples.
//
{
    std::promise<void> eventPromise;
    std::future<void> eventFuture = eventPromise.get_future();
```

(continues on next page)



(continued from previous page)

```
std::unique_ptr<mal::ps::Subscription> subscription = subscriber->
↳subscribeAsync(
    mal::ps::DataEventFilter<Sample>::all(),
    [&eventPromise] (mal::ps::Subscriber<Sample>& subscriber) -> void {
        // NOTE: you should never close subscription from
        // within a callback, since that may cause deadlock
        eventPromise.set_value();
    });

std::future_status status = eventFuture.wait_for(std::chrono::seconds(1));
if (status == std::future_status::ready) {
    std::cout << subscriber->read(0).size() << std::endl;
} else {
    std::cout << "No event was received\n";
}
}
} catch (std::exception& exc) {
    std::cerr << "EXCEPTION: " << exc.what() << std::endl;
}
}
```

Python

Python subscriber offers (as described above) three versions of the subscribe method:

- subscribe()
 - Creates synchronous subscription. Received data events are passed to provided callback method via poll() call on the subscriber object synchronously.
- subscribeAsync()
 - Creates asynchronous subscription. Provided callback method is called from another thread as new data events arrive.
- subscribeAsyncNotifier()
 - Creates asynchronous subscription. Callback method receives change notifications. Method is called from another thread.

When creating data event filter, data entity class (not instance) must be provided as argument to factory function mal.ps.DataEventFilter.all().

```
DEFAULT_TIMEOUT = datetime.timedelta(seconds=3)
# Poll interval in seconds.
POLL_INTERVAL = datetime.timedelta(seconds=1)
uri = 'dds.ps://m1/CabinetTelemetry'
ddsMal = mal.loadMal('dds', {'dds.domain': '100'})
factory = mal.CiiFactory.getInstance()
```

(continues on next page)



(continued from previous page)

```
factory.registerMal('dds', ddsMal)
with factory.getSubscriber(uri,
                           Example.Sample,
                           qos=mal.ps.qos.DEFAULT) as subscriber:

    # Read all alive instances w/o sample count limit.

    allAliveInstances = subscriber.read(0)

    # Read one instance i.e. daqId == 1.
    keySample = subscriber.createDataEntity()
    keySample.setDaqId(1)
    daq1sample = subscriber.readInstance(keySample)

    # Loaned (no-copy) example.
    # Loan is returned when close() method is
    # called on LoanedDataEntity. Note: there is
    # no context management protocol available for
    # loaned samples. Use try, finally block.
    loanedSample = subscriber.loanedRead()
    try:
        sampleData = loanedSample.getData()
    finally:
        # Return loan
        loanedSample.close()

    # Read data events, filtered by instance and instance state.
    # Block until maxEvents (1000) are read or timeout (3s) occurs.

    maxEvents = 1000

    # Create new filter for Sample data entity.
    filter = mal.ps.DataEventFilter.all(Sample)
    filter.setInstance(keySample)
    filter.setInstanceState(mal.ps.InstanceState.ALIVE)
    filter.setInstanceEvents(set(mal.ps.InstanceEvent.CREATED,
                                 mal.ps.InstanceEvent.UPDATED))

    # List of samples is returned.
    samples = subscriber.readDataEvents(filter,
                                       maxEvents,
                                       timeout=DEFAULT_TIMEOUT)

    # Asynchronous data subscription, no filtering,
    # 1 second sleep added to simulate some work.

    def dataEventFunction(subscription, dataEvent):
        # NOTE: you should never close subscription from within
        # a callback, since that may cause a deadlock
```

(continues on next page)



(continued from previous page)

```
# Data event processing method
if dataEvent.isValidData():
    sample = dataEvent.getData()
    # use sample
else:
    # Note that in Python it is possible to create a subscriber
    # that expects samples with different data type that
    # are provided by publisher. This is especially difficult
    # to detect without this code (and the only way to find
    # out is to study the log files).
    # Usually subscriber receives non valid data when
    # publisher instance is not alive anymore, so in case
    # publisher instance is still alive, our subscription might
    # not be correct. The code below handles this case.
    instance_state = dataEvent.getInstanceState()
    if instance_state == mal.ps.InstanceState.ALIVE:
        print('WARNING: There is something wrong with data, please '
              ' check your publisher and subscriber data types')

with subscriber.subscribeAsync(
    mal.ps.DataEventFilter.all(Example.Sample),
                                dataEventFunction) as subscription:
    # Suspend main thread, dataEventFunction is
    # called from another thread.
    time.sleep(1)
# Note that subscription is destroyed once with block terminates.

# Asynchronous data subscription, no filtering,
# using lambda callback.
# DO NOT CALL read methods in the callback!
# 1 second sleep added to simulate some work.
with subscriber.subscribeAsync(mal.ps.DataEventFilter.all(),
                                lambda subscriber, dataEvent:
                                print(dataEvent.getData())) as subscription:
    # Suspend main thread, data is printed from another thread.
    time.sleep(1)

# Synchronous (aka event pump) subscription, no filtering.
# Poll with 1000 events / 1 seconds limit, endless loop.
with subscriber.subscribe(mal.ps.DataEventFilter.all(ExampleSample),
                                dataEventFunction) as subscription:
    while True:
        # dataEventFunction will be called synchronously
        # when new event arrives.
        subscriber.poll(maxEvents,
                        datetime.timedelta(seconds=1))
```

In Python it is possible to create subscriber that does not expect data with data type equal to data type of published event. This case is hard to detect by only checking whether event has valid data.



The situation is detected within subscription monitor and indicated in the event with publisher instance state set to ALIVE, so additional check can be done in the event processing method to handle this case. See Python example above.

2.5.3 MrvSubscriber

MrvSubscriber is a convenience variant of a subscriber intended to be used by GUIs. It has only one method that always blocks for a given amount of time. This might be used to implement a refresh loop. Use zero time to read samples from a queue and do not block for any more time. The method always returns a list of most-recent values for each instance. It does not return only changed instances over specified amount. If instance becomes unavailable it is returned in the list anymore. This method creates new entities and copies their data; avoid using this method when dealing with large arrays.

This class is thread-safe.

Java

```
// install DDS MAL mapping
CiiFactory.installMal("dds");

try (CiiFactory factory = CiiFactory.getInstance()) {

    URI uri = URI.create("dds.ps:///m1/CabinetTelemetry");
    try (MrvSubscriber<Sample> mrvSubscriber =
        factory.getMrvSubscriber(uri, QoS.DEFAULT,
            null, Sample.class)) {

        //
        // Every 1 second get most-recent-values of alive instances.
        //
        while (true) {
            List<Sample> mrvSamples = mrvSubscriber.readMostRecent(1, TimeUnit.
↪SECONDS);
            System.out.println(Arrays.toString(mrvSamples.toArray()));
        }
    }
}
```



C++

```
std::shared_ptr<::elt::mal::Mal> ddsMal = mal::loadMal("dds", {"dds.domain",  
↳ "100"});  
::elt::mal::CiiFactory &factory = mal::CiiFactory::getInstance();  
factory.registerMal("dds", ddsMal);  
  
mal::Uri uri("dds.ps:///m1/CabinetTelemetry");  
  
try {  
    std::unique_ptr<ps::MrvSubscriber<mal::example::Sample>> mrvSubscriber =  
        factory.getMrvSubscriber<mal::example::Sample>(uri, ↳  
↳ mal::ps::qos::QoS::DEFAULT, {});  
  
    //  
    // Every 1 second get most-recent-values of alive instance  
    //  
    while (true) {  
        std::vector<std::shared_ptr<mal::example::Sample>> mrvSamples =  
            mrvSubscriber->readMostRecent(std::chrono::seconds(1));  
        for (const auto &sample : mrvSamples) {  
            std::cout << "Sample " << sample->getDaqId() << " / " <<  
                sample->getValue() << std::endl;  
        }  
    }  
} catch (std::exception &exc) {  
    std::cerr << "EXCEPTION: " << exc.what() << std::endl;  
}
```

Python

```
uri = 'dds.ps:///m1/CabinetTelemetry'  
# Load DDS MAL mapping.  
ddsMal = mal.loadMal('dds', {'dds.domain': '100'})  
factory = mal.CiiFactory.getInstance()  
factory.registerMal('dds', ddsMal)  
  
with factory.getMrvSubscriber(uri,  
                               Example.Sample,  
                               qos=mal.ps.qos.DEFAULT) as mrvSubscriber:  
  
    # Every second get most-recent-value of alive instances.  
  
    timeout = datetime.timedelta(seconds=1)  
    while True:  
        mrvSamples = mrvSubscriber.readMostRecent(timeout)  
        if mrvSamples:  
            print(mrvSamples)
```



2.6 Request-response

Request-response is a coupled peer-to-peer client-server communication pattern. A server exposes to the network one or more instances of ICD-defined interfaces (services). A client connects to the server and invokes one or more methods on the services. MAL guarantees reliable QoS (i.e. request cannot be lost) and exactly-once delivery (i.e. for each method invocation exactly one request is made on the server-side; the request can neither be lost nor duplicated). MAL completely takes care of connection management. There is no need to explicitly issue a connect request, monitor connection status and do a reconnect. MAL takes care of this in the background. A programmer only needs to specify a timeout limit on the request. (Re-)connection time is included in the total execution time of the request). If connection cannot be (re-)established in that time or response is not delivered in time a `TimeoutException` is thrown. If response is delivered after the timeout expired the response is ignored. Any connection loss during the invocation of the method results in `DisconnectedException`, i.e. MAL does not try to invoke the method on the server side again after the reconnection.

In order to explicitly wait for connection to be established a connect method returns a future object that completes when the connection is established for the first time.

In cases where a programmer is interested in connection state it can register a connection listener on a client instance. The callbacks are called by MAL mapping or middleware thread.

When multiple responses are needed asynchronous method invocation (AMI) MAL support can be used. Basically, AMI is an asynchronous request with one or more asynchronous responses grouped as request. There is no limitation on number of responses nor time limit on lifetime. Lifetime is limited for the time of connection. Once connection is lost AMI request gets canceled (reported to the client as cancellation exception).

NOTE: Since the Migration to FastDDS, Request-Response communication became unavailable for MAL DDS C++, Java and Python languages.

2.6.1 Client

When an interface is defined in ICD, a synchronous and asynchronous client variants of the interfaces are generated. A programmer has a choice to select at client instantiation what interface to use. Invoking a method on synchronous interface will block and return its value via method return mechanism, or exception will be thrown. A request timeout can be set via `ReplyTime QoS` or via client instance created using `timeout` method. A `TimeoutException` is thrown in case of a timeout condition. Asynchronous methods always return a future object. A programmer then gets a response, or an exception, back via a future object. There can be more than one asynchronous requests issued/pending concurrently on one client instance.

Type of the interface on the client side does not affect (in any way) how a request is processed on the server side.

AMI methods can be recognized as methods that return `Ami<returnType>` object. The object is basically an iterator of future objects of `returnType` type. The future objects behave exactly the same as in case of asynchronous methods – a future object is a promise waiting for the next response. AMI call



is completed when the iterator runs out of elements or when exception is being thrown (also thrown in case of disconnection).

The example code below illustrates the usage of a client interface and using all three types of methods.

Java

```
public class ClientExample {

    public static void main(String[] args) {

        // install ZPB MAL mapping
        CiiFactory.installMal("zpb");

        try (CiiFactory factory = CiiFactory.getInstance()) {

            URI uri = URI.create("zpb.rr:///ml/Robot1");

            //
            // Synchronous client example w/ ReplyTime QoS set.
            // RobotControlSync interface is requested.
            //
            try (RobotControlSync robot =
                factory.getClient(uri,
                    new QoS[] {
                        new ReplyTime(3, TimeUnit.SECONDS)
                    },
                    null, RobotControlSync.class)) {

                //
                // Explicitly wait for connection to be established.
                // (This is usually not needed.)
                //
                CompletableFuture<Void> connectionFuture = robot.asyncConnect();
                connectionFuture.get(10, TimeUnit.SECONDS);

                //
                // Invoke synchronous method
                // TimeoutException is throw if default timeout is exceeded.
                //
                robot.command(Command.START_COMMAND);

                //
                // Some methods throw user-defined exceptions.
                // Handle them just like standard Java exceptions.
                //
                try {
                    float speed = robot.getSpeed();
                }
            }
        }
    }
}
```

(continues on next page)



(continued from previous page)

```
List<Float> speedValues = new ArrayList<Float>();
speedValues.add(66.6f);
speedValues.add(41.0f);
speedValues.add(90.5f);
boolean checker = robot.validSpeedValues(speedValues);
if (checker) {
    speed += 30.0f;
}
else {
    speed += 3.0f;
}

robot.setSpeed(speed);
}
catch (TooFast tf) {
    // handle too fast
}

//
// Create a new client instance w/ specific timeout.
// New instance is connected to the same service instance.
// Old instance is still valid.
//
RobotControlSync robotSlowCall = robot.timeout(1, TimeUnit.MINUTES);
robotSlowCall.command(Command.STOP_COMMAND);

//
// Multiple responses use-case using
// asynchronous method invocation (AMI).
//
try (Ami<String> systemCheckAmi = robot.systemCheck()) {
    for (CompletableFuture<String> future : systemCheckAmi) {
        String msg = future.get();
        System.out.println(msg);
    }
}

} catch (Throwable th) {
    // handle exceptions ...
}

//
// Asynchronous client example w/ reply timeout set.
// RobotControlAsync interface is requested.
//
try (RobotControlAsync robot =
    factory.getClient(uri,
        new QoS[] {
```

(continues on next page)



(continued from previous page)

```
        new ReplyTime(3, TimeUnit.SECONDS)
    },
    null, RobotControlAsync.class)) {

    //
    // Basic case, invoke and then wait, if necessary
    //
    CompletableFuture<Void> cfuture = robot.command(Command.START_COMMAND);
    // ... do some work here ...,
    // and get response (Void in this case), future will block if
    // the response is not yet received, or TimeoutException is throws in
↔reply-time
    // is exceeded
    cfuture.get();

    //
    // Chained CompletableFuture<> example:
    // async getSpeed() is called and on return setSpeed
    // with new value requested, exceptions are handled.
    // Wait for only 2 seconds to complete.
    //
    robot.getSpeed().thenAccept((Float speed) -> robot.setSpeed(speed + 10))
        .exceptionally((th) -> {
            /* handle exception */ return null;
        }).get(2, TimeUnit.SECONDS);

    //
    // Multiple responses use-case using
    // asynchronous method invocation (AMI).
    // Handling is the same as for sync. client.
    //
    try (Ami<String> systemCheckAmi = robot.systemCheck()) {
        for (CompletableFuture<String> future : systemCheckAmi) {
            String msg = future.get();
            System.out.println(msg);
        }
    }

    } catch (Throwable th) {
        // handle exceptions
    }
}
}
```

**C++**

```
std::shared_ptr<::elt::mal::Mal> zpbMal = mal::loadMal("zpb", {"zpb.domain",  
↵ "100"});  
::elt::mal::CiiFactory &factory = mal::CiiFactory::getInstance();  
  
factory.registerMal("zpb", zpbMal);  
  
mal::Uri uri("zpb.rr:///m1/Robot1");  
  
//  
// Synchronous client example w/ ReplyTime QoS set.  
// RobotControlSync interface is requested.  
//  
try {  
    std::unique_ptr<mal::example::RobotControlSync> robot = factory.getClient  
↵ <mal::example::RobotControlSync>(  
        uri,  
        {std::make_shared<mal::rr::qos::ReplyTime>(std::chrono::seconds(3))},  
        {});  
  
    //  
    // Explicitly wait for connection to be established.  
    // (This is usually not needed.)  
    //  
    ::elt::mal::future<void> connectionFuture = robot->asyncConnect();  
    std::future_status futureStatus =  
        connectionFuture.wait_for(::boost::chrono::seconds(10));  
    bool connected = (futureStatus == std::future_status::ready);  
  
    //  
    // Invoke synchronous methods  
    // TimeoutException is throw if default timeout is exceeded.  
    //  
  
    robot->command(mal::example::RobotControl::Command::START_COMMAND);  
  
    try {  
        float speed = robot->getSpeed();  
  
        mal::shared_vector<const float> speedArray = {66.6, 41.0, 90.5};  
        bool checker = robot->validSpeedValues(speedArray);  
        if(checker) {  
            speed += 30;  
        }  
        else {  
            speed += 3;  
        }  
  
        robot->setSpeed(speed);  
    }
```

(continues on next page)



(continued from previous page)

```
} catch (mal::example::RobotControl::TooFast& tf) {
    // handle too fast
}

//
// Multiple responses use-case using
// asynchronous method invocation (AMI).
//
std::shared_ptr<StrAmi> ami = robot->systemCheck();
for (auto future : *ami) {
    std::cout << future.get() << std::endl;
}

//
// Create a new client instance w/ specific timeout.
//
robot->timeout(std::chrono::seconds(10))
    ->command(mal::example::RobotControl::Command::STOP_COMMAND);
} catch (std::exception& exc) {
    std::cerr << "EXCEPTION: " << exc.what() << std::endl;
}

//
// Asynchronous client example w/ reply timeout set.
// RobotControlAsync interface is requested.
//
try {
    std::unique_ptr<mal::example::RobotControlAsync> robot = factory.getClient
    ↪<mal::example::RobotControlAsync>(
        uri,
        {std::make_shared<mal::rr::qos::ReplyTime>
         (std::chrono::seconds(3))},
        {});

    //
    // Basic case, invoke and then wait, if necessary
    //
    mal::future<void> cfuture = robot->command
        (mal::example::RobotControl::Command::START_COMMAND);
    // ... do some work here ...,
    // and get response (void in this case), future will block if
    // the response is not yet received, or TimeoutException is throws in
    // reply-time is exceeded
    cfuture.get();

    // Chained future<> example:
    // async getSpeed() is called and on return setSpeed
    // with new value requested, exceptions are handled.
    // Wait for 2 seconds to complete.
```

(continues on next page)



(continued from previous page)

```
try {
    mal::future<float> future = robot->
        getSpeed().then([&] (::mal::future<float> f)
            { return robot->setSpeed(f.get() + 0.1); }).get();

    if (future.wait_for(boost::chrono::seconds(2)) !=
        ::boost::future_status::ready) {
        // handle timeout
    }

} catch (mal::example::RobotControl::TooFast& tf) {
    // handle too fast
}

//
// Multiple responses use-case using
// asynchronous method invocation (AMI).
//
std::shared_ptr<StrAmi> ami = robot->systemCheck();
for (auto future : *ami) {
    std::cout << future.get() << std::endl;
}
} catch (std::exception& exc) {
    std::cerr << "EXCEPTION: " << exc.what() << std::endl;
}
}
```

Python

When using asynchronous client, result of the method call on the client reference is always a Future object. Actual result value must be obtained with get method on the result object.

To check state of the Future object, use its check method. Method takes no arguments and returns enumeration mal.rr.FutureStatus that can have the one of the following values:

- DEFERRED – requested method was not started yet.
- READY – requested method completed and result can be retrieved with call to get method.
- TIMEOUT – wait was performed on the Future, but timeout occurred.

```
import elt.pymal as mal
# Import synchronous client implementation
from ModManualExamples.Elt.Mal.Example.RobotControl import RobotControlSync
# Import asynchronous client implementation
from ModManualExamples.Elt.Mal.Example.RobotControl import RobotControlAsync
from ModManualExamples.Elt.Mal.Example import Command, TooFast

# 3 seconds.
```

(continues on next page)



(continued from previous page)

```
THREE_SECONDS = datetime.timedelta(seconds=3)

# Minute
MINUTE = datetime.timedelta(seconds=60)

def _main():
    """Main function implementation"""

    # sync example
    try:
        # Load ZPB MAL mapping.
        uri = 'zpb.rr://ml/Robot1'
        zpbMal = mal.loadMal('zpb', {'zpb.domain': '100'})
        factory = mal.CiiFactory.getInstance()
        factory.registerMal('zpb', zpbMal)

        # Synchronous client example w/ ReplyTime QoS timeout set.
        # RobotControlSync interface is requested.
        with factory.getClient(uri,
                               RobotControlSync,
                               qos=mal.rr.qos.ReplyTime(THREE_SECONDS)) as robot:

            # Explicitly wait for connection to be established.
            # (This is usually not needed).
            connectionFuture = robot.asyncConnect()
            connectionFuture.wait_for(THREE_SECONDS)

            # Invoke synchronous methods
            # TimeoutError is thrown if default timeout is exceeded.

            robot.command(Command.START_COMMAND)
            try:
                speed = robot.getSpeed()

                speedArray = mal.makeSharedVector('f', 3, 90.5)
                speedArray = speedArray.freeze()
                checker = robot.validSpeedValues(speedArray)
                if checker:
                    speed += 30
                else:
                    speed += 3

                robot.setSpeed(speed)
            except TooFast as e:
                # Handle too fast
                pass

            # Multiple responses use-case using
            # asynchronous method invocation (AMI).
```

(continues on next page)



(continued from previous page)

```
systemCheckAmi = robot.systemCheck()
# We are not using context manager protocol here.
# Once we are done with the AMI, it should be
# closed().
try:
    # Result set is iterable.
    for msg in systemCheckAmi:
        print(msg)
finally:
    systemCheckAmi.close()

# Asynchronous example w/reply timeout set.
# RobotControlAsync interface is requested.
with factory.getClient(uri,
                       RobotControlAsync,
                       qos=mal.rr.qos.ReplyTime(THREE_SECONDS)) as robot:

    # Basic case, invoke and then wait, if necessary
    future = robot.command(Command.START_COMMAND)

    # ... do some work here ...
    # and get response (None in this case). Future will block
    # it the response is not yet received or TimeoutError is
    # thrown when reply time is exceeded.

    future.get()

    # Multiple responses use-case using
    # asynchronous method invocation (AMI).
    # Handling is the same as for sync. client.

with robot.systemCheck() as systemCheckAmi:
    for future in systemCheckAmi:
        print('System Check: ', future.get())
except Exception as e:
    # Handle exceptions or
    raise
```

2.6.2 Server

Creating a server instance just creates a container where services can be registered. A service is a uniquely named instance reachable over the network. There can be multiple services of the same or different interface type registered to the server. Same service instance can also be registered multiple times using different names. The server will just propagate the request to a service by its name. Once a service is unregistered it is no longer reachable over the network. Any pending requests on unregistered service will complete normally.



When an interface is defined in ICD, a synchronous and asynchronous service variants of the interfaces are generated. A programmer needs to implement one of these interfaces. Only these interfaces can be registered as a service. An interface type dictates how requests are processed on the server side. Server processes requests of all the services using one thread. Once a request is received it decodes the request and calls a method on specified service. If the service implements a synchronous interface then invocation of its method will block. This includes processing of other requests of the same and other services. In case of asynchronous service the method will block just for a method to return a future object. The server will continue to process other pending requests. The server will send the response of asynchronous services back to the client immediately when the return value, or an exception, is given to the future object. This implies that asynchronous service responses might come back in a different order than their requests were issued.

When a C++ MAL service throws a known exception (i.e. declared in the MAL-ICD), it is passed through network to the client, where it is re-raised. However, if the service throws other exceptions: a MAL exception of a type that is not declared for this particular method, a standard exception, or throws a C++ object that is not derived from the std::exception, then MAL will behave as shown in Table 1.

Table with 4 columns: ICD, Exc. type, Server behaviour, Client behaviour. It details how the system handles various exception types like MAL exceptions, C++ exceptions, and non-MAL exceptions.

Table 1 MAL behaviour when an interface method throws an undeclared exception

Implementation of server-side AMI methods is the same for synchronous and asynchronous service interfaces. AMI methods declarations return the same Ami<returnType> as on the client side. However, the service needs to provide the implementation of the Ami<returnType>. The creation of it is done via ServerContextProvider singleton instance implemented as thread-safe-storage instance that provides access to server-side helper methods. A createAmi method will create a ServerAmi<returnType> instance (that implements Ami<returnType> interface). The implementation of the AMI method just needs to setup asynchronous execution of the method (the method should not block!) and return the ServerAmi<returnType> instance. Sending responses to the client is done by calling one of the three methods on the ServerAmi<returnType> instance: complete (non-last response), completeExceptionally (exception response, last response) and completed (last response). Once last response is sent the instance gets automatically released.



See the examples below on how to implement synchronous and asynchronous service.
The server has no responsibility over the lifecycle of registered services.

Java

```
public class ServerExample {

    private static void simulateWork() {
        try {
            Thread.sleep(100);
        } catch (InterruptedException ex) {
            // noop
        }
    }

    /**
     * Synchronous RobotControl implementation example,
     * must implement <ICD interface name> interface.
     */
    static class RobotControlImpl implements RobotControl {

        private final float maxSpeed;
        private List<Float> allSpeeds;
        private List<Float> tempSpeeds;

        public RobotControlImpl(float maxSpeed, List<Float> speeds) {
            this.maxSpeed = maxSpeed;
            this.tempSpeeds = speeds;
            this.allSpeeds = new ArrayList<Float>();
            this.allSpeeds.add(0.0f);
        }

        @Override
        public void command(Command com) {
            switch (com) {
                case START_COMMAND:
                    this.allSpeeds = new ArrayList<Float>(this.tempSpeeds);
                    break;

                case STOP_COMMAND:
                    this.allSpeeds = new ArrayList<Float>();
                    this.allSpeeds.add(0.0f);
                    break;

                default:
                    // undefined exceptions, i.e. exception not defined in ICD,
                    // translate to UnhandledRemoteException on the remote side
                    throw new IllegalArgumentException();
            }
        }
    }
}
```

(continues on next page)



(continued from previous page)

```
    }  
}  
  
@Override  
public float setSpeed(float speed) throws TooFast {  
    if (speed > this.maxSpeed) {  
        throw new TooFast();  
    }  
  
    this.allSpeeds.add(speed);  
    return this.allSpeeds.get(this.allSpeeds.size() - 1);  
}  
  
@Override  
public float getSpeed() {  
    return this.allSpeeds.get(this.allSpeeds.size() - 1);  
}  
  
@Override  
public boolean validSpeedValues(List<Float> values) {  
    boolean valid = true;  
    for (float val : values) {  
        if (val < 0 || val > this.maxSpeed) {  
            valid = false;  
            break;  
        }  
    }  
    return valid;  
}  
  
@Override  
public Ami<String> systemCheck() {  
  
    // request for ServerAmi instance  
    // to be returned by this method  
    ServerAmi<String> ami = ServerContextProvider.getInstance().  
↪createAmi(String.class);  
  
    // setup asynchronous execution  
    Executors.newCachedThreadPool().execute(() -> {  
  
        simulateWork();  
        // first response  
        ami.complete("Battery OK");  
  
        simulateWork();  
        // second response  
        ami.complete("Drivetrain OK");  
    });  
}
```

(continues on next page)



(continued from previous page)

```
        simulateWork();
        // final response
        // NOTE: completed vs complete method name
        ami.completed("Engine OK");
    });

    // return AMI instance
    return ami;
}

}

/**
 * Asynchronous RobotControl implementation example,
 * must implement Async<ICD interface name>.
 */
static class AsyncRobotControlImpl implements AsyncRobotControl {

    private final float maxSpeed;
    private List<Float> allSpeeds;
    private List<Float> tempSpeeds;

    public AsyncRobotControlImpl(float maxSpeed, List<Float> speeds) {
        this.maxSpeed = maxSpeed;
        this.tempSpeeds = speeds;
        this.allSpeeds = new ArrayList<Float>();
        this.allSpeeds.add(0.0f);
    }

    @Override
    public CompletableFuture<Void> command(Command com) {
        // create future
        CompletableFuture<Void> future = new CompletableFuture<>();

        // setup asynchronous execution
        Executors.newCachedThreadPool().execute(() -> {
            switch (com) {
                case START_COMMAND:
                    this.allSpeeds = new ArrayList<Float>(this.tempSpeeds);
                    // notify the completion w/ result
                    future.complete(null);
                    break;

                case STOP_COMMAND:
                    this.allSpeeds = new ArrayList<Float>();
                    this.allSpeeds.add(0.0f);
                    // notify the completion w/ result
                    future.complete(null);
                    break;
            }
        });
    }
}
```

(continues on next page)



(continued from previous page)

```
        default:
            // undefined exceptions, i.e. exception not defined in ICD,
            // translate to UnhandledRemoteException on the remote side
            // notify the completion w/ exception
            future.completeExceptionally(new IllegalArgumentException());
        }
    });

    // return future
    return future;
}

@Override
public CompletableFuture<Float> setSpeed(float speed) {

    // create future
    CompletableFuture<Float> future = new CompletableFuture<>();

    // setup asynchronous execution
    Executors.newCachedThreadPool().execute(() -> {
        if (speed > this.maxSpeed) {
            // notify the completion w/ exception
            future.completeExceptionally(new TooFast());
        }

        this.allSpeeds.add(speed);
        // notify the completion w/ result
        future.complete(speed);
    });

    // return future
    return future;
}

@Override
public CompletableFuture<Float> getSpeed() {
    // if the execution of the method is short
    // an already completed future can be returned
    return CompletableFuture.completedFuture(this.allSpeeds.get(this.allSpeeds.
↵size() - 1));
}

@Override
public CompletableFuture<boolean> validSpeedValues(List<Float> values) {

    // create future
    CompletableFuture<Float> future = new CompletableFuture<>();
```

(continues on next page)



(continued from previous page)

```
// setup asynchronous execution
Executors.newCachedThreadPool().execute(() -> {
    boolean valid = true;
    for (float val : values) {
        if (val < 0 || val > this.maxSpeed) {
            valid = false;
            break;
        }
    }
}

// notify the completion w/ result
future.complete(valid);
}

// return future
return future;
}

@Override
public Ami<String> systemCheck() {

    // request for ServerAmi instance
    // to be returned by this method
    ServerAmi<String> ami = ServerContextProvider.getInstance().
    ↪createAmi(String.class);

    // setup asynchronous execution
    Executors.newCachedThreadPool().execute(() -> {

        simulateWork();
        // first response
        ami.complete("Battery OK");

        simulateWork();
        // second response
        ami.complete("Drivetrain OK");

        simulateWork();
        // final response
        // NOTE: completed vs complete method name
        ami.completed("Engine OK");
    });

    // return AMI instance
    return ami;
}
}
```

(continues on next page)



(continued from previous page)

```
/**
 * Main entry-point.
 *
 * @param args command-line arguments.
 */
public static void main(String[] args) {

    URI uri = URI.create("zpb.rr:///m1");

    // install ZPB MAL mapping
    CiiFactory.installMal("zpb");

    try (CiiFactory factory = CiiFactory.getInstance()) {

        //
        // Create server, default QoS.
        // Currently there are no MAL QoS relevant for the Server.
        //
        try (Server server = factory.createServer(uri, QoS.DEFAULT, null)) {

            List<Float> speedsArray = new ArrayList<Float>();
            speedsArray.add(0.1f);
            speedsArray.add(49.4f);
            speedsArray.add(74.4f);

            //
            // Register three Robot services.
            // They are all separate separate instances.
            // First two instances have synchronous implementation,
            // the third has asynchronous.
            server.registerService("Robot1", new RobotControlImpl(120.0f, ↵
            ↵speedsArray));
            server.registerService("Robot2", new RobotControlImpl(80.0f, ↵
            ↵speedsArray));
            server.registerService("Robot3", new AsyncRobotControlImpl(20.0f, ↵
            ↵speedsArray));

            // Process commands until server.shutdown() is called.
            server.run();
        }
    }
}
}
```

**C++**

```
namespace mal = ::elt::mal;

/**
 * Simulate some work
 */
static void simulateWork() {
    std::this_thread::sleep_for(std::chrono::seconds(2));
}

/**
 * Synchronous RobotControl implementation example,
 * must implement RobotControl interface.
 */
class RobotControlImpl: public virtual mal::example::RobotControl {
public:

    RobotControlImpl(float max, mal::shared_vector<const float> speeds)
        : maxSpeed(max), tempSpeeds(speeds) {
        allSpeeds = {0.0};
    }

    void command(RobotControl::Command command) override {
        switch (command) {
            case RobotControl::Command::START_COMMAND:
                allSpeeds = tempSpeeds;
                break;

            case RobotControl::Command::STOP_COMMAND:
                allSpeeds = {0.0};
                break;

            default:
                throw std::invalid_argument("RobotControlImpl: Unknown command");
        }
    }

    float setSpeed(float newSpeed) override {
        if (speed > maxSpeed) {
            throw RobotControl::TooFast(
                "RobotControlImpl::setSpeed, speed too high");
        }

        ::elt::mal::shared_vector<float> mutableSpeeds = thaw(allSpeeds);
        mutableSpeeds.push_back(newSpeed);
        allSpeeds = freeze(mutableSpeeds);

        return allSpeeds.back()
    }
}
```

(continues on next page)



(continued from previous page)

```
float getSpeed() const override { return allSpeeds.back(); }

bool validSpeedValues(const ::elt::mal::shared_vector<const float>& values)
↳override {
    bool valid = true;
    for (const auto& val: values) {
        if (val < 0 || val > maxSpeed) {
            valid = false;
            break;
        }
    }
    return valid;
}

std::shared_ptr<RobotControl::StrAmi> systemCheck() override {
    // Obtain ptr to ServerAmi
    std::shared_ptr<RobotControl::StrAmi> ami = mal::rr::ServerContextProvider
        <mal::example::ServerContextImpl<std::string>>::
        getInstance().createAmi();

    auto future = boost::async(boost::launch::async, [=]() {
        simulateWork();
        ami->complete("Battery OK");
        simulateWork();
        ami->complete("Drivertrain OK");
        simulateWork();
        ami->completed("Engine OK");
    });

    // Return ptr to Ami
    return ami;
}

private:
    float maxSpeed;
    mal::shared_vector<const float> allSpeeds;
    mal::shared_vector<const float> tempSpeeds;
};

/**
 * Asynchronous RobotControl implementation example,
 * must implement AsyncRobotControl interface.
 */
class AsyncRobotControlImpl: public virtual mal::example::AsyncRobotControl {
public:

    AsyncRobotControlImpl(float max, mal::shared_vector<const float> speeds)
        : maxSpeed(max), tempSpeeds(speeds) {
```

(continues on next page)



(continued from previous page)

```
    allSpeeds = {0.0}
}

mal::future<void> command(mal::example::RobotControl::Command command)
override {
    std::function<void()> f = [this, command]() -> void {
        std::lock_guard<std::mutex> guard(speedMutex);
        switch (command) {
            case mal::example::RobotControl::Command::START_COMMAND:
                allSpeeds = tempSpeeds;
                break;
            case mal::example::RobotControl::Command::STOP_COMMAND:
                allSpeeds = {0.0};
                break;

            default:
                throw std::invalid_argument("RobotControlImpl: Unknown command");
        }
    };
    return boost::async(boost::launch::async, f);
}

mal::future<float> setSpeed(float newSpeed) override {
    return boost::async(boost::launch::async, [this, newSpeed]() -> float {
        std::lock_guard<std::mutex> guard(speedMutex);
        if (newSpeed > maxSpeed) {
            throw mal::example::RobotControl::TooFast(
                "RobotControlImpl::setSpeed, speed too high");
        }

        ::elt::mal::shared_vector<float> mutableSpeeds = thaw(allSpeeds);
        mutableSpeeds.push_back(newSpeed);
        allSpeeds = freeze(mutableSpeeds);

        return allSpeeds.back()
    });
}

mal::future<float> getSpeed() const override {
    boost::promise<float> promise;
    promise.set_value(allSpeeds.back());
    return promise.get_future();
}

mal::future<float> validSpeedValues(const ::elt::mal::shared_vector<const_
↵float>& values) override {
    boost::promise<bool> promise;

    bool valid = true;
```

(continues on next page)



(continued from previous page)

```
for (const auto& val: values) {
    if(val < 0 || val > maxSpeed) {
        valid = false;
        break;
    }
}

promise.set_value(valid);
return promise.get_future();
}

std::shared_ptr<AsyncRobotControl::StrAmi> systemCheck() override {
    // Obtain ptr to ServerAmi
    std::shared_ptr<AsyncRobotControl::StrAmi> ami =
mal::rr::ServerContextProvider
    <mal::example::ServerContextImpl<std::string>>::
    getInstance().createAmi();

    auto future = boost::async(::boost::launch::async, [=]() {
        simulateWork();
        ami->complete("Battery OK");
        simulateWork();
        ami->complete("Drivertrain OK");
        simulateWork();
        ami->completed("Engine OK");
    });

    // Return ptr to Ami
    return ami;
}

private:
    float maxSpeed;
    std::mutex speedMutex;
    mal::shared_vector<const float> allSpeeds;
    mal::shared_vector<const float> tempSpeeds;
};

/**
 * Server example.
 *
 * Main entry point.
 * @param argc number of command line arguments
 * @param argv command line arguments
 */
int main(int argc, char** argv) {

    // Load ZPB MAL mapping
    std::shared_ptr<::elt::mal::Mal> zpbMal = mal::loadMal("zpb", {"zpb.domain",
mal::"100"});
```

(continues on next page)



(continued from previous page)

```
// Register mapping with CiiFactory
::elt:: &factory = mal::CiiFactory::getInstance();
factory.registerMal("zpb", zpbMal);

mal::Uri uri("zpb.rr://m1");

try {
    //
    // Create server, default QoS.
    // Currently there are not MAL QoS relevant for the Server
    //
    std::unique_ptr<rr::Server> server = factory.createServer(uri,
↪mal::rr::qos::QoS::DEFAULT, {});

    mal::shared_vector<const float> speedsArray = {0.1, 49.4, 74.4};

    // Register three Robot services.
    // They are all separate instances.
    // First two instance have synchronous implementation,
    // third has asynchronous.
    std::shared_ptr<mal::rr::RrEntity> robot1(new RobotControlImpl(120.0,
↪speedsArray));
    std::shared_ptr<mal::rr::RrEntity> robot2(new RobotControlImpl(80.0,
↪speedsArray));
    std::shared_ptr<mal::rr::RrEntity> robot3(new AsyncRobotControlImpl(80.0,
↪speedsArray));

    server->registerService<mal::example::RobotControl, true>
        ("Robot1", robot1);
    server->registerService<mal::example::RobotControl, true>
        ("Robot2", robot2);
    server->registerService<mal::example::RobotControl, false>
        ("Robot3", robot3);

    // Process commands until server->shutdown() is called.
    server->run();
} catch (std::exception& exc) {
    std::cerr << "EXCEPTION: " << exc.what() << std::endl;
}
return 0;
}
```




Python

```
import datetime
import time
import threading
from concurrent.futures import ThreadPoolExecutor

# import top level Python MAL API module
import elt.pymal as mal
# import RobotControl binding
from ModManualExamples.Elt.Mal.Example import RobotControl
# import Command enumeration and TooFast exception
from ModManualExamples.Elt.Mal.Example import Command, TooFast

def simulateWork():
    time.sleep(1.0)

class RobotControlSyncImpl:
    """Synchronous RobotControl implementation example, must
    implement methods defined in the ICD definition file:
    command, setSpeed, getSpeed, systemCheck
    """
    def __init__(self, maxSpeed, speeds):
        """constructor"""
        self._allSpeeds = [0.0]
        self._maxSpeed = maxSpeed
        self._tempSpeeds = speeds

    def command(self, com):
        """Handle command"""
        if com == Command.START_COMMAND:
            self._allSpeeds = self._tempSpeeds
        elif com == Command.END_COMMAND:
            self._allSpeeds = [0.0]
        else:
            raise ValueError('Invalid command value')

    def setSpeed(self, speed):
        """Set new speed.
        Throws TooFast when speed exceeds maximum speed.
        """
        if speed > self._maxSpeed:
            raise TooFast()
        self._allSpeeds.append(speed)
        return self._allSpeeds[-1]

    def getSpeed(self):
        """Return current speed"""
        return self._allSpeeds[-1]
```

(continues on next page)



(continued from previous page)

```
def validSpeedValues(self, values):
    """Return Validation of entry Values"""
    result = True
    for val in values:
        if val < 0 or val > self._maxSpeed:
            result = False
            break

    return result

def systemCheck(self):
    """Perform system check"""
    # Obtain instance to the server context.
    # This call uses ServerContextProvider.getInstance()
    # method internally.
    context = RobotControl.ServerContextString.getInstance()
    # Obtain server AMI
    serverAmi = context.createAmi()
    # Setup asynchronous execution
    systemCheckThread = threading.Thread(self._systemCheck,
        args=(serverAmi,))
    systemCheckThread.start()
    return serverAmi

def _systemCheck(self, ami):
    """Actual systemCheck implementation"""
    simulateWork()
    # first response
    ami.complete('Battery OK')

    simulateWork()
    # second response
    ami.complete("Drivetrain OK")

    simulateWork()
    # final response, use completed method
    ami.completed("Engine OK")

class RobotControlAsyncImpl:
    """Asynchronous RobotControl implementation example
    all methods that don't return Ami, must return
    concurrent.futures.Future object
    """
    def __init__(self, maxSpeed, speeds):
        """constructor"""
        self._executor = ThreadPoolExecutor(max_workers=4)
        self._allSpeeds = [0.0]
        self._maxSpeed = maxSpeed
```

(continues on next page)



(continued from previous page)

```
self._tempSpeeds = speeds

def command(self, com):
    """Handle command"""
    return self._executor.submit(self, _command, com)

def _command(self, com):
    """Command implementation"""
    if com == Command.START_COMMAND:
        self._allSpeeds = self._tempSpeeds
    elif com == Command.END_COMMAND:
        self._allSpeeds = [0.0]
    else:
        raise ValueError('Invalid command value')

def setSpeed(self, speed):
    """Set speed"""
    self._executor.submit(self._setSpeed, speed)

def _setSpeed(self, speed):
    """setSpeed implementation
    Throws TooFast when speed exceeds maximum speed.
    """
    if speed > self._maxSpeed:
        raise TooFast()
    self._allSpeeds.append(speed)
    return self._allSpeeds[-1]

def getSpeed(self):
    """Get speed"""
    return self._executor.submit(lambda x: x, self._allSpeeds[-1])

def validSpeedValues(self, values):
    """Check if values are inside the limitations"""
    return self._executor.submit(self._validSpeedValues, values)

def _validSpeedValues(self, values):
    """validSpeedValues implementation
    Sanity check for list of values.
    """
    result = True
    for val in values:
        if val < 0 or val > self._maxSpeed:
            result = False
            break

    return result

def systemCheck(self):
```

(continues on next page)



(continued from previous page)

```
    """Perform system check"""
    # Obtain instance to the server context.
    # This call uses ServerContextProvider.getInstance()
    # method internally.
    context = RobotControl.ServerContextString.getInstance()
    # Obtain server AMI
    serverAmi = context.createAmi()
    # Setup asynchronous execution
    systemCheckThread = threading.Thread(self._systemCheck,
        args=(serverAmi,))
    systemCheckThread.start()
    return serverAmi

def _systemCheck(self, ami):
    """Actual systemCheck implementation"""
    simulateWork()
    # first response
    ami.complete('Battery OK')

    simulateWork()
    # second response
    ami.complete("Drivetrain OK")

    simulateWork()
    # final response, use completed method
    ami.completed("Engine OK")

def main():
    uri = 'zpb.rr:///m1'

    # Load ZPB MAL mapping
    zpbMal = mal.loadMal('zpb')
    factory = CiiFactory.getInstance()
    factory.registerMal('zpb', zpbMal)

    with factory.createServer(uri, mal.rr.qos.DEFAULT, {}) as server:

        speedsArray = [0.1, 49.4, 74.4]

        # Register three Robot services.
        # They are all separate instances.
        # First two have synchronous implementation,
        # the third has asynchronous implementation.
        server.registerService('Robot1',
            RobotControl.RobotControlSyncService,
            RobotControlSyncImpl(120.0, speedsArray))
        server.registerService('Robot2',
            RobotControl.RobotControlSyncService,
```

(continues on next page)



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 61 of 589

(continued from previous page)

```
        RobotControlSyncImpl(80.0, speedsArray))
server.registerService('Robot3',
        RobotControl.RobotControlAsyncService,
        RobotControlAsyncImpl(20.0, speedsArray))

server.run()
```



2.7 Entities

As a general rule, entities must be instantiated through factory methods provided by the MAL (and likewise provided by MAL Publishers). If the entity is complex, the factory method deep-creates also the nested entities inside. To populate the complex entity, you have to get the nested entities from it, and populate them.

Caution: for some datatypes and middlewares, after getting the nested entity from the complex entity and populating it, you must (counter-intuitively) finally also call the setter on the complex entity. This is needed to trigger the marshalling of the data. We are seeing to improve this.

Caution: the C++ life-cycle of the child entities is bound to the life-cycle of its parent live-cycle. MAL implementation might use a middleware-specific structure to hold an entity data. The parent entity is the owner of the structure. Due to performance reasons (requirements), all the fields including child fields, are written directly to the structure without using any local storage. The middleware might not support ownership hierarchy, therefore the child entities are bound to the live-cycle of the specific structure owned by the parent. With the parent entity gone, the child entity would point to an invalid location of the previously located specific structure. If you need a child entity outside of the scope of a parent entity life-cycle, you need to clone the child entity, e.g. `parent->getChild()->clone()`. Do not assume that the smart pointers will take care of this limitation since the child entities do not hold a smart reference to the parent entities.

2.7.1 Array, Blob

To transfer ICD arrays, MAL uses the Shared Vector utility class, which is similar to the standard vector (or list) classes in the respective languages and additionally supports “freeze” and “thaw”. Only a “thawed” vector can be modified, only a “frozen” vector can be sent over MAL.

Python

Example: shared vector (used for sending/receiving arrays)

```
import elt.pymal as mal

# usage: makeSharedVector (<type>, <size>, <value>)
# <type>: float, double, string, uint8_t, int8_t (etc.)
vec = mal.makeSharedVector('int8_t', 10, 120)

print (vec)          # output: SharedVectorInt8[x, x, x, x, x, x, x, x, x, x]
print (vec[2:])     # output: SharedVectorInt8[x, x, x, x, x, x, x, x]

# before sending over MAL
vec_to_send = vec.freeze()

# vector to list
aslist = list (vec)
```

(continues on next page)



(continued from previous page)

```
print (aslist)    # output: [120, 120, 120, 120, 120, 120, 120, 120, 120, 120, 120]
# list to vector
vec2 = mal.SharedVectorInt8 (aslist)
```

Example: blob

```
complex_entity = publisher.createDataEntity()
nested_blob = complex_entity.getNested()
vec = mal.makeSharedVector('int8_t', 512, 0x55)
vec = vec.freeze()
nested_blob.setBlob (vec)
```

C++

Example: directly map one-dimensional arrays into your preferred library objects, e.g. for Eigen:

```
using Eigen;

elt::mal::shared_vector<double> values = ...

// square 3x3 matrix
Matrix3d mat1(values.data());

// more advanced mapping of 3x3 matrix
Map<Matrix<double, 3, 3, RowMajor> mat2(values.data());
```



2.8 Appendix

2.8.1 Building CII application with WAF

A CII application must be built using waf/wtools that has ICD generator (icd-gen tool) and all other supported middleware generators are integrated. All the artifacts are built into property formed libraries that are needed to be dynamically loaded by MAL mappings. Steps to build a CII application is the following:

- Create a project directory with a project wscript file, e.g.:

```
from wtools.project import declare_project

wtools.project.declare_project('icd-demo-folder', '0.1-dev-version', #
↳Project directory name and version
                                recurse='icd-app1 icd-app2', #
↳Name of all applications folders
                                requires='cxx boost java python protoc cii', #
↳Primary libraries to be linked
                                boost_libs='log log_setup thread system') #
↳Secondary libraries to be linked
                                                                    # ...
```

Project directory will contain as many applications necessary with the correspondent ICD XML files. Following the above example code 2 application will be implemented. So, after creating the applications directories:

- Inside bought application directory, 2 subdirectories should be created: 'icd' (contains ICD files for applications) and 'app' (where the application implementation is located). Names of those two subdirectories can be different.
- Parallel to this folders a wscript must be added with the following structure:

```
from wtools import package

# Declare external libraries to be used specifically by the respective
↳application
def configure(conf):
    """System installed libraries"""
    conf.check_cfg(package='opentracing_api', uselib_store='OPENTRACE', args='--
↳cflags --libs')

    """Created WAF libraries based on other applications"""
    conf.check_wdep(wdep_name='cii.mal-common', uselib_store='MALCOMMON')
    conf.check_wdep(wdep_name="client-api.config", uselib_store="CLIENTCONF")

package.declare_package(recurse='app icd') # Folders inside application
↳directory
```

Jumping in too the declared packages ('app' and 'icd'), each one must have a wscript file. Starting with



the 'app' subdirectories, is worth noting that depending on the developer language the WAF method call to create the correspondent executable file will change ('declare_cprogram' 'declare_pyprogram' 'declare_jar'):

```
from wtools.module import declare_cprogram

"""From pre-checked libraries in previous folder 'wscript', call the ones needed
↳for the respective application"""
declare_cprogram(target='executable-app-language', use='BOOST CLIENTCONF
↳MALCOMMON OPENTRACE')
```

And the 'icd' subdirectories with a wscript like:

```
from wtools.module import declare_malicd

declare_malicd() # Enables ICD artifacts generation
```

Remember that when building the project, the WAF Tool will only use files that exist inside a 'src' folder. So all necessary coding for the application to work as entended must be inside a source directory. In the end a developer should have a structure like this:

```
root
|-- icd-demo-folder
|   |-- icd-appl
|   |   |-- app
|   |   |   |-- src
|   |   |   |   |-- main-appl-file.xxx
|   |   |   |   |-- wscript
|   |   |-- icd
|   |   |   |-- src
|   |   |   |   |-- icd-appl-file.xml
|   |   |   |   |-- wscript
|   |   |-- wscript
|   |
|   |-- icd-app2
|   |   |-- app
|   |   |   |-- src
|   |   |   |   |-- main-app2-file.xxx
|   |   |   |   |-- wscript
|   |   |-- icd
|   |   |   |-- src
|   |   |   |   |-- icd-app2-file.xml
|   |   |   |   |-- wscript
|   |   |-- wscript
|   |
|   |-- wscript
```

See icd-demo project as example.

MAL ICD GENERATION

Revision:	1.6
Status:	Release
Repository:	https://gitlab.eso.org/cii/info/cii-docs
Project:	ELT CII
Folder:	/trunk/deliverables/phase5
Document ID:	CSL-MAN-19-170402
File:	MAN-MAL_icdgen.docx
Owner:	Matej Šekoranja, Cosylab Sweden
Last modification:	March 30, 2023
Created:	February 18, 2019

Prepared by	Reviewed by	Approved by
Damjan Kumar, CSL Matej Šekoranja, CSL SWE	Aljaž Podboršek Gregor Čuk	Gregor Čuk

Document History



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 67 of 589

Revision	Date	Changed/reviewed	Section(s)	Modification
0.1	2019-02-21	dkumar	All	Created.
1.0	2019-02-26	dkumar	All	Review, updated missing sections.
1.1	2019-04-02	mvitorovic	6	Added description for topics.
1.2	2019-04-10	msekoranja, dkumar	All	Review, updated ICD limitations
1.3	2019-04-24	dkumar	6.3, 6.4	Updated topics namespace names.
1.4	2019-06-06	msekoranja	4.9	Referencing types from other packages note added.
1.5	2023-03-15	mschilli	3.5	Info about icd-gen splitter
1.6	2023-03-30	rleo	3.4.8	arrays in different middlewares

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is an ICD development and ICD generation **user manual** document for the ELT Core Integration Infrastructure Software project.

Audience

Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms

API	Application Programmers Interface
CII	Core Integration Infrastructure
MAL	Middleware Abstraction Layer
ICD	Interface Control Document

References

1. Cosylab, ELT CII MAL Transfer document, CSL-DOC-18-168015, version 1.8
2. Cosylab, ELT CII MAL API User's Manual, CSL-DOC-17-150198, version 2.1
3. Cosylab, Interface Control Document, Specification, CSL-DOC-17-147262, version 1.6



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 68 of 589

4. Cosylab, Data Addressing Specification, CSL-DOC-17-147264, version 1.2



3.1 Overview

This document is a user manual for developing Interface Control Documents, specifically Type Definition Documents holding ICD types.



3.2 Installation

See ELT CII MAL Transfer document [1], Chapter 1.2 on how to build and install elt-mal.

3.2.1 Prerequisites

In order to use icd generation the following prerequisites must be fulfilled:

- CII MAL prerequisites, see ELT CII MAL Transfer document [1], Chapter 1.1.
- Installed elt-mal.



3.3 Editing ICD file in Eclipse

Developing Interface Control Documents involves writing XML documents. Developer is encouraged to use a XML editor that has support for XML content assist and support for a schema validation. Eclipse that is part of ESO ELT Development Environment has support for editing XML files with XML content assist and schema validation.

See more at <https://help.eclipse.org/oxygen/index.jsp?topic=%2Forg.eclipse.wst.xmleditor.doc.user%2Ftopics%2FworkXML.html>.

For the purposes of this user guide, complete the following steps which will enable support for XML content assist and schema validation for **icdDemo.xml** within Eclipse:

1. Open Eclipse (/opt/eclipse/oxygen/eclipse/eclipse)
2. Click **File > New > Project**.
3. Select **General > Project** and click **Next**.
4. Enter **icd-demo** as 'Project name' and click **Finish**.
5. To import the ICD XML schemas into the workbench, click **File > Import**.
6. Select **General > File System** and click **Next**.
7. Click **Browse** on the next page of the wizard to select the directories from which you would like to add the XML schemas (i.e. \$INTROOT/interface/schemas)
8. Select check box on 'schemas' which will import all schemas under the directory \$PREFIX/interface/schemas.
9. Make sure that option 'Create top-level folder' is selected and then click **Finish**.
10. Next step will create a Waf package named **icd** containing **icdDemo.xml**. Open the terminal and execute the following commands:

```
mkdir -p ~/icd-demo/icd/src

echo '<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">
</types>' > ~/icd-demo/icd/src/icdDemo.xml

echo 'from wtools.module import declare_malicd
declare_malicd()' > ~/icd-demo/icd/wscript
```

11. XML content assist and schema validation for **icdDemo.xml** will work within Eclipse only when the file **icdDemo.xml** is beside directory **schemas/** in the Eclipse project. A soft link needs to be created in `<user_eclipse_workspace>/icd-demo`, pointing to the **icdDemo.xml** that is part of Waf package. Open the terminal and execute the following commands (replace **user_eclipse_workspace** with a real path):



```
cd <user_eclipse_workspace>/icd-demo/  
ln -s ~/icd-demo/icd/src/icdDemo.xml .
```

12. Refreshing the Eclipse project (pressing F5) should show the icdDemo.xml as part of the icd-demo Eclipse project.

Support for the XML content assist should work by placing the cursor into the type's element body and pressing **Ctrl + Space**. The schema validity for icdDemo.xml is being checked as you type or it can be activated manually by Right-clicking on your file in the Navigator view and selecting Validate. Figure 1 is showing the XML content assist and a schema validation error (invalid package name) when editing the icdDemo.xml.

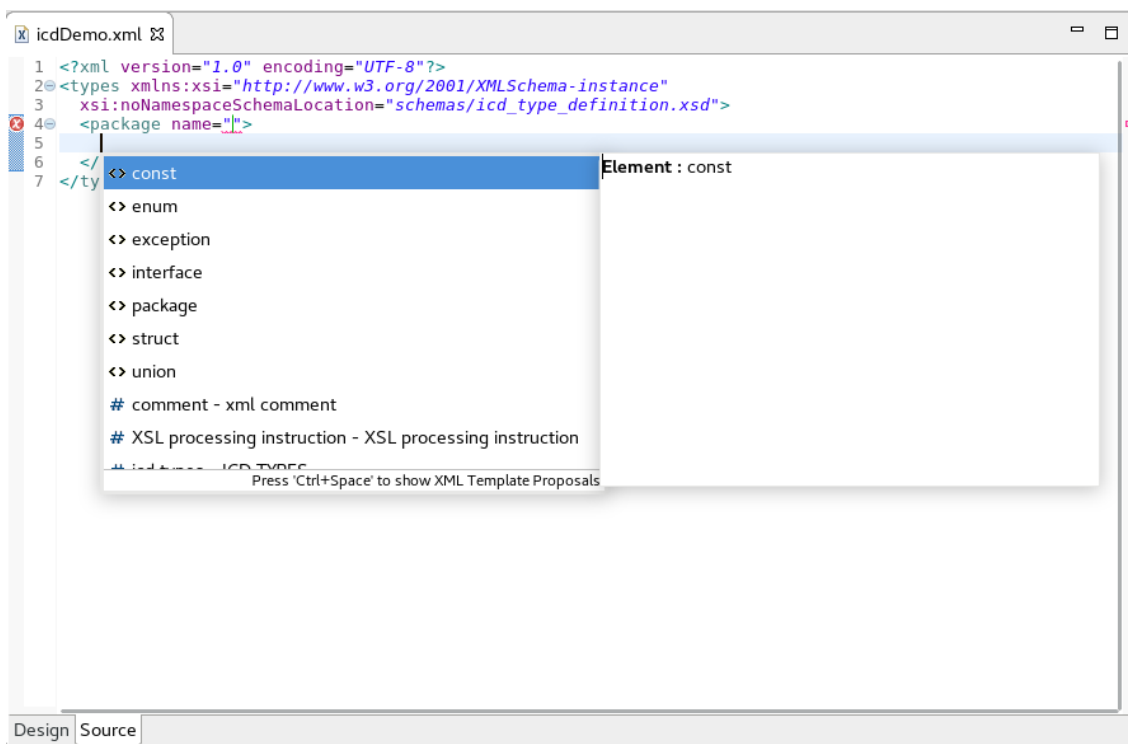


Figure 1: XML content assist and schema validation on icdDemo.xml



3.4 ICD Building blocks

Interface Control Document defines two types of XML documents [3], and this document describes **Type definition document** ([3], 1.1) that contains structured and interface types in chapters 4 and 5. **Topic and interface definition document** ([3], 1.2) is described in chapter 6.

ICD Type definition document is composed of:

1. **Packages** that group related types and prevent naming conflicts on similarly named types.
2. **Structures** that group data elements together under one name. These data elements, known as members, can be a primitive type or other complex types such as structures and unions.
3. **Basic types** are primitive programming language types (including string) and standard ICD types (blob, timestamp).
4. **Enumerations** that are distinct types whose value is restricted to a range of values that include several explicitly named constants (“enumerators”).
5. **Unions** are a special types that can hold only one of its data members at a time.
6. **Interfaces** describe related methods belonging to the same interface.
7. **Exceptions** define exceptional events that occur during the execution of interface methods.
8. **Constants** which are values used in the elements and attributes where a primitive value would be used.

Within one package types with duplicate names are not allowed. In addition some keywords are reserved (see [3], section 1.1.8.1)

3.4.1 Package

All ICD types belong to a package. There are no package-less types. Package has a non-empty unqualified name (i.e. must not contain other package names delimited with scope resolution operator ::). Package can include other packages but only the most inner package can hold types. ICD XML schema makes sure that the package constraints are maintained in the generation phase but also while the developer is creating the ICD in IDE (i.e. Eclipse).

Listing 1 and Listing 2 contain an example of a single and a nested package.

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">
  <package name="elt">
    <!-- type definition content -->
  </package>
</types>
```

Listing 1: Single package



```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">
    <package name="icdDemo">
      <!-- type definition content -->
    </package>
  </package>
</types>
```

Listing 2: Nested package

See Interface Control Document Specification for an example how package is mapped to Java and C++ ([3], 1.3.5).

3.4.2 Basic types

Basic types define:

1. Primitive types of members in structures and unions.
2. Types of union discriminators.
3. Define method arguments of primitive types or constant types.

Only basic types and enums can be keys in structures.

Basic types are:

- **boolean**, boolean primitive type holding true or false.
- **int8_t**, signed integer type with width of exactly 8 bits.
- **int16_t**, signed integer type with width of exactly 16 bits.
- **int32_t**, signed integer type with width of exactly 32 bits.
- **int64_t**, signed integer type with width of exactly 64 bits.
- **uint8_t**, unsigned integer type with width of exactly 8 bits.
- **uint16_t**, signed integer type with width of exactly 16 bits.
- **uint32_t**, unsigned integer type with width of exactly 32 bits.
- **uint64_t**, unsigned integer type with width of exactly 64 bits.
- **float**, single precision floating point type IEEE-754, 32 bit.
- **double**, double precision floating point type. IEEE-754 64 bit.
- **string**, string of characters. Length upper limit is 256 Bytes.
- **blob**, stream of bytes.



- **timestamp**, timestamp.

Listing 3 is showing a usage of basic types in different structured types and constants.

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">
  <package name="elt">

    <struct name="SimpleStruct">
      <member name="boolval" type="boolean"/>
      <member name="int8val" type="int8_t" key="true"/>
      <member name="uint8val" type="uint8_t" key="true"/>
    </struct>

    <const name="int8Const" type="int8_t" value="1"/>
    <const name="floatConst" type="float" value="6.4523424"/>

    <union name="SimpleUnion">
      <discriminator type="uint8_t"/>
      <case>
        <caseDiscriminator value="0"/>
        <member name="foo" type="string"/>
      </case>
      <case>
        <caseDiscriminator value="1"/>
        <member name="bar" type="int8_t"/>
      </case>
    </union>
  </package>
</types>
```

Listing 3: Usage of basic types for members in a structure, in constants and member of unions

3.4.3 Constants

Constants define values used in the elements and attributes where a primitive value would be used.

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">
    <const name="boolconst" type="boolean" value="true"/>
    <const name="int8const" type="int8_t" value="1"/>
    <const name="uint8const" type="uint8_t" value="2"/>
    <const name="int16const" type="int16_t" value="-3"/>
    <const name="uint16const" type="uint16_t" value="100"/>
    <const name="int32const" type="int32_t" value="-3200"/>
  </package>
</types>
```

(continues on next page)



(continued from previous page)

```
<const name="uint32const" type="uint32_t" value="100"/>  
<const name="floatconst" type="float" value="100.0"/>  
<const name="stringconst" type="string" value="batman"/>  
</package>  
</types>
```

Listing 4: Defining constants

3.4.4 Enumeration (Enum)

Enumeration is a distinct type whose value is restricted to a range of values explicitly named constants (enumerators). Enumeration has a non-empty unqualified name and at least one enumerator. The enumeration type is a building block for unions and can be included as a member of a structure and member of a union.

Listing 5 is showing an example of simple command modeled as an enumeration.

```
<?xml version="1.0" encoding="UTF-8"?>  
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">  
  
  <package name="elt">  
    <enum name="Command">  
      <enumerator name="START_COMMAND"/>  
      <enumerator name="STOP_COMMAND"/>  
    </enum>  
  </package>  
</types>
```

Listing 5: Enumerator Command

See Interface Control Document Specification for an example how enumeration is mapped to Java and C++ ([3], 1.3.1).

3.4.5 Union

Union is a special type that can hold only one of its data members at a time. The discriminator sets which member holds a value. Only numeric primitives and enums can be union discriminators. Union members can be: primitive types, unions, structures and enums. No arrays are allowed as union members.

Listing 6 and Listing 7 are showing examples of a union with a primitive type discriminator and a union with an enumerator as a union discriminator.

```
<?xml version="1.0" encoding="UTF-8"?>  
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
```

(continues on next page)



(continued from previous page)

```
    xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">
<package name="elt">
  <union name="SimpleUnion">
    <discriminator type="uint8_t"/>
    <case>
      <caseDiscriminator value="0"/>
      <member name="foo" type="string"/>
    </case>
    <case>
      <caseDiscriminator value="1"/>
      <member name="bar" type="int8_t"/>
    </case>
  </union>
</package>
</types>
```

Listing 6: Union with numeric primitive type discriminator

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">

    <enum name="Command">
      <enumerator name="START_COMMAND"/>
      <enumerator name="STOP_COMMAND"/>
    </enum>

    <union name="CommandData">
      <discriminator type="nonBasic" nonBasicTypeName="Command"/>
      <case>
        <caseDiscriminator value="START_COMMAND"/>
        <member name="startCommandData" type="string"/>
      </case>
      <case>
        <caseDiscriminator value="STOP_COMMAND"/>
        <member name="stopCommandData" type="int8_t"/>
      </case>
    </union>
  </package>
</types>
```

Listing 7: Union with enum discriminator

See Interface Control Document Specification for an example how union is mapped to Java and C++ ([3], 1.3.3).



3.4.6 Structure (struct)

Structure groups data elements (members) together under one name. Structure member can have a basic type or other structure type, union or enumeration type. Every structure can be assigned a key. A key is a set of fields (of primitive types) that uniquely identifies one instance. This allows programmers to publish/retrieve/filter data based on their instance. Only basic types and enums can be keys in structures. Structure can be extended. Extending structure means COPY (not inheritance!) from extended, including the keys.

Listing 8, Listing 9, Listing 10 and Listing 11 show typical type definitions of a structure.

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">

    <struct name="SimpleStruct">
      <member name="boolval" type="boolean"/>
      <member name="int8val" type="int8_t" key="true"/>
      <member name="uint8val" type="uint8_t" key="true"/>
      <member name="int16val" type="int16_t"/>
      <member name="uint16val" type="uint16_t"/>
      <member name="int32val" type="int32_t"/>
      <member name="uint32val" type="uint32_t"/>
    </struct>
  </package>
</types>
```

Listing 8: Structure with basic types as members

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">

    <enum name="Command">
      <enumerator name="START_COMMAND"/>
      <enumerator name="STOP_COMMAND"/>
    </enum>

    <struct name="ArrayStruct">
      <member name="boolvalArray" type="boolean" arrayDimensions="(16)"/>
      <member name="int8valArray" type="int8_t" arrayDimensions="(16)"/>
      <member name="CmdArray" type="nonBasic"
        nonBasicTypeName="Command" arrayDimensions="(20)"/>
    </struct>
  </package>
</types>
```

**Listing 9: Structure with arrays as members**

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">

    <struct name="SimpleStruct">
      <member name="boolval" type="boolean"/>
      <member name="int8val" type="int8_t" key="true"/>
    </struct>

    <struct name="ComplexStruct">
      <member name="boolval" type="boolean"/>
      <member name="int8val" type="int8_t"/>
      <member name="nestedStruct" type="nonBasic"
        nonBasicTypeName="SimpleStruct"/>
    </struct>
  </package>
</types>
```

Listing 10: Structure holding another structure as a member.

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">

    <struct name="BaseStruct">
      <member name="boolval" type="boolean"/>
      <member name="int8val" type="int8_t" key="true"/>
    </struct>

    <struct name="ExtendedStruct" baseType="BaseStruct">
      <member name="boolval" type="boolean"/>
      <member name="int8val" type="int8_t"/>
      <member name="nestedStruct" type="nonBasic"
        nonBasicTypeName="SimpleStruct"/>
    </struct>
  </package>
</types>
```

Listing 11: Extending a structure.

See Interface Control Document Specification for an example how structure is mapped to Java and C++ ([3], 1.3.2).



3.4.7 Interface and Exception

A server exposes to the network one or more instances of ICD-defined interface (a service). A client connects to the server and invokes one or more methods (method element) on that service. Methods defined in ICD Interface can throw (throws attribute) exceptions raising exceptional events during the execution of methods. Exception can have optional members with additional exception data. Exception type cannot publicly extend other exception types, while the interface can be extended (extends element). Extending interface means copying from the extended interface. Methods that are part of the interface can have zero or more arguments (argument attribute). Methods return results (returnType attribute). When multiple responses (results) are needed, asynchronous method invocation (AMI) MAL support can be used (methodType attribute). Declaring the AMI method with a void return-type is banned and it will produce a compilation error.

Listing 12 shows a typical service definition with different types of methods.

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">

    <struct name="SpeedArray" trace="true">
      <member name="values" type="float" arrayDimensions="(10)"/>
    </struct>

    <enum name="Command">
      <enumerator name="START_COMMAND"/>
      <enumerator name="STOP_COMMAND"/>
    </enum>

    <exception name="TooFast">
      <member name="speed" type="float"/>
    </exception>

    <interface name="Logger">
    <method name="setLevel" returnType="void">
      <argument name="level" type="uint8_t"/>
    </method>
    </interface>

    <interface name="RobotControl">

    <extends interfaceName="Logger"/>

    <method name="command" returnType="void">
      <argument name="com" type="nonBasic" nonBasicTypeName="Command"/>
    </method>

    <method name="setSpeed" returnType="float" throws="TooFast">
```

(continues on next page)



(continued from previous page)

```
<argument name="value" type="float" />
</method>
<method name="getSpeed" returnType="float" />
<method name="validSpeedValues" returnType="boolean">
  <argument name="values" type="float" arrayDimensions="(10)" />
</method>

<method name="systemCheck" returnType="string" methodType="ami"/>

<method name="setStrVec" returnType="void">
  <argument name="strvec" type="string" arrayDimensions="(5)" />
</method>

<method name="getStrVec" returnType="string" arrayDimensions="(5)" />

</interface>
</package>
</types>
```

Listing 12: RobotControl service definition

See Interface Control Document Specification for an example how an exception and an interface are mapped to Java and C++ ([3], 1.3.4 and 1.3.6).

3.4.8 Arrays

Arrays are specified with an attribute **arrayDimensions** on:

- Member of structured type (struct, union, exception): member of a structured type is an array.
- Method return type: return type of method is an array.
- Method parameter: method parameter is an array

Arrays are one of:

- BOUNDED: `arrayDimensions="(dim1, dim2, ..., dimN)"` where `dim1, dim2, ..., dimN > 0`
- FIXED: `arrayDimensions="[dim1, dim2, ..., dimN]"` where `dim1, dim2, ..., dimN > 0`

Fixed array is sent over the wire in full capacity, regardless how many elements are actually stored the array. On the other hand, in a bounded array, only currently stored elements will be sent, however the middleware might reserve entire capacity to avoid reallocation at runtime. Setting boundaries too big might cause in high memory usage.

It is also relevant to warn that some types of communication will not comply with the Bounded/Fixed array types and it's dimensional definitions. That is the case of ZPB communication method, because the actual array will be a repeated field in the protobuf. The following table shows how the array behave in a ZPD and DDS protocols:



Protocol/Pattern	4 Passed	2 Passed	1 Passed	Empty
ZPB RR Bounded (2)	4 Received	2 Received	1 Received	Empty
ZPB RR Fixed [2]	4 Received	2 Received	1 Received	Empty
ZPB PS Bounded (2)	4 Received	2 Received	1 Received	Empty
ZPB PS Fixed [2]	4 Received	2 Received	1 Received	Empty
DDS PS Bounded (2)	4 Received	2 Received	1 Received	Empty
DDS PS Fixed [2]	Exc on Publish	2 Received	Exc on Publish	Exc on Publish

Next is an ICD XML example file that shows how to declare an array:

```
<?xml version="1.0" encoding="UTF-8"?>
<types xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="schemas/icd_type_definition.xsd">

  <package name="elt">

    <struct name="SpeedArray" trace="true">
      <member name="values" type="float" arrayDimensions="(10)"/>
    </struct>

    <struct name="SampleMultidimension">
      <member name="matrixValue" type="string" arrayDimensions="[3,2]"/>
    </struct>

  </package>
</types>
```

Listing 13: Example of a Structure Array Specification



3.5 ICD Generation

When the ICD Type definition documents are written, MAL agnostic and MAL middlelayer specific types need to be generated based on the types defined in them. The ICD generator (`icd-gen` tool installed in `$PREFIX/bin`) must never be used directly by the user. ICD generation has been integrated into `waf/wtools`.

User needs to create a `waf` package composed of ICD Type definition document stored in `src/` directory and use `declare_maliced()` `wtools` support in package `wscript` file. During installation process all ICDs gets installed in `$PREFIX/interface/icd` directory so that they can be included by other ICDs.

During the generation process, the `icd` generator splits ICD files into smaller chunks. This was added as an optimisation for memory allocation and compilation duration, and is fully transparent to the users. Should there ever be a need, you can enable debug-logs from the splitter, set the environment variable `ICD_PYTHON_CPP_SPLITTER_DEBUG` to non-empty. To prevent the splitting, set the environment variable `ICD_WITHOUT_PYTHON_CPP_SPLITTER` to non-empty.

See ELT CII MAL API User's Manual [2], Chapter 7.1 on how to create a CII application using `waf/wtools` named `icd-demo` (containing `icdDemo.xml`) that was created in Chapter 3.



3.6 Topic and interface definition document

The ICD generation tool can also be used to parse the **Topic and interface definition document** and generate language specific source files that expose various parameters of the definition to the developer as constant values that can be used application source code.

This type of document defines the following elements per each subsystem:

1. **Messaging pattern** that may be publish/subscribe topics or method definitions for interfaces.
2. **Quality of Service** parameters for both of the above.
3. **Performance characteristics** for both of the above.
4. **Type and interface names**; these must be defined in the type definition document.
5. **Middleware mappings**; the document lists only the names of the protocols that apply.
6. **Address URI** as defined in [4].

In order to enable topic and interface definition generation use `-topic (-t)` command-line option when calling `icd-gen`.

3.6.1 Subsystems

Each document lists on or multiple subsystems. At the moment the only subsystem type supported is **device**.

```
<?xml version="1.0" encoding="UTF-8"?>
<topics xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
        xsi:noNamespaceSchemaLocation="schemas/icd_topic_definition.xsd">
  <device name="Dev1">
    <!-- subsystem definition content -->
  </device>
  <device name="Dev2">
    <!-- subsystem definition content -->
  </device>
</topics>
```

Listing 16: Subsystems (devices)

3.6.2 Messaging pattern

Each subsystem may contain one or multiple publish/subscribe topics or service interface definitions. The element for publish/subscribe topic is named **pubsub_topic**, and the element for the service interface definition is named **service**.

```
<?xml version="1.0" encoding="UTF-8"?>
<topics xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
```

(continues on next page)



(continued from previous page)

```
      xsi:noNamespaceSchemaLocation="schemas/icd_topic_definition.xsd">
<device name="Dev1">
  <pubsub_topic>
    <!-- topic definition content -->
  </pubsub_topic>
</device>
<device name="Dev2">
  <service name="MSA1" interface="MSAzimuth">
    <!-- subsystem definition content -->
  </service>
</device>
</topics>
```

Listing 17: Topic and service definition

3.6.3 Publish subscribe topic

The publish/subscribe topic definition contains the following required elements:

1. **topic_name** – this element defines the name of the topic.
2. **topic_type** – this element defines the type name of the structure defined in the type definition document (see chapter 4).
3. **address_uri** – this element contains the MAL address URI.
4. **qos** – this element defines the Quality of Service attributes.
5. **performance** - this element defines the performance attributes.
6. **mal** – this element defines provider specific elements, however it is not processed by the tool.

For a detailed description the elements and their attributes please see [3].

Based on the XML definition from Listing 18 the tool generates the code shown in Listing 19, Listing 20 and Listing 21.

```
<?xml version="1.0" encoding="UTF-8"?>
<topics xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="schemas/icd_topic_definition.xsd">
  <device name="Dev1">
    <pubsub_topic>
      <topic_name>Az:Pos_Actual</topic_name>
      <topic_type>Pos_Actual</topic_type>
      <address_uri>dds.ps://m1/Pos_Actual?p1=v1&amp;p2=v2</address_uri>
      <qos latency_ms="0.1" deadline_ms="20" />
      <performance rate_hz="0.2" latency_ms="25" synchronous="false" />
      <mal>
        <dds />
      </mal>
    </pubsub_topic>
  </device>
</topics>
```

(continues on next page)



(continued from previous page)

```
</pubsub_topic>  
</device>  
</topics>
```

Listing 18: Topic definition (full)

For all Java classes the package name is predefined. The device is a public class with static class members for all defined topics (and service interfaces). The name of the topic class is derived from the **topic_name** element with the **PubSub** suffix. The class member constants are derived from the other elements listed above.

```
package elt.icd.topicdefinition;  
  
public class Dev1 {  
  
    public static class AzPosActualPubSub {  
        public static String TYPE = "Pos_Actual";  
        public static String ADDRESS_URI = "dds.ps://m1/Pos_Actual?p1=v1&p2=v2";  
        public static double QOS_LATENCY_MS = 0.1;  
        public static double QOS_DEADLINE_MS = 20.0;  
        public static double PERFORMANCE_RATE_HZ = 0.2;  
        public static double PERFORMANCE_LATENCY_MS = 25;  
        public static boolean PERFORMANCE_SYNCHRONOUS = false;  
    } // AzPosActualPubSub  
  
} // Dev1
```

Listing 19: Topic Java code

For C++ the tool generates only the header file. The top 3 namespaces are predefined. The device is a namespace with namespace members for all defined topics (and service interfaces). The name of the topic namespace is derived from the **topic_name** element with the **PubSub** suffix. The constants are derived from the other elements listed above.

```
#ifndef __DEV_1_HPP  
#define __DEV_1_HPP  
  
namespace elt  
{  
    namespace icd  
    {  
        namespace topicdefinition  
        {  
  
            namespace Dev1  
            {  
  
                namespace AzPosActualPubSub
```

(continues on next page)



(continued from previous page)

```
{
    constexpr char[] type = "Pos_Actual";
    constexpr char[] addressUri = "dds.ps://m1/Pos_Actual?p1=v1&p2=v2";
    constexpr double qosLatencyMs = 0.1;
    constexpr double qosDeadlineMs = 20.0;
    constexpr double performanceRateHz = 0.2;
    constexpr double performanceLatencyMs = 25;
    constexpr bool performanceSynchronous = false;
} // namespace AzPosActualPubSub

} // namespace Dev1

} // namespace topicdefinition
} // namespace icd
} // namespace elt

#endif // __DEV_1_HPP
```

Listing 20: Topic C++ code

For Python the tool generates a class for each device with class members for all defined topics (and service interfaces). The name of the topic class is derived from the **topic_name** element with the **PubSub** suffix. The members are derived from the other elements listed above.

```
class Dev1:

    class AzPosActualPubSub:
        type = "Pos_Actual";
        addressUri = "dds.ps://m1/Pos_Actual?p1=v1&p2=v2";
        qosLatencyMs = 0.1;
        qosDeadlineMs = 20.0;
        performanceRateHz = 0.2;
        performanceLatencyMs = 25;
        performanceSynchronous = False;

    # end AzPosActualPubSub

# end Dev1
```

Listing 21: Topic Python code



3.6.4 Service interface

The attributes of each **service** definition are the service instance name and the interface type name. It contains the following required elements:

1. **address_uri** – this element contains the MAL address URI.
2. **performance** – this element defines the performance attributes.
3. **qos** – this element defines the Quality of Service attributes.
4. **is_configuration** – the configuration has no defined performance constraints and represents aspects of the Local Control System software which must be controllable externally from the CCS. This element is optional.
5. **mal** – this element defines provider specific elements, however it is not processed by the tool.
6. **methods** – this element contains all the method definitions for this service.

The **methods** element contains one or many **method** elements. Each element has a **name** attribute specifying the method name. Each method contains the following optional elements which override the values listed above:

1. **performance** – this element overrides the performance attributes.
2. **qos** – this element overrides the Quality of Service attributes.
3. **is_configuration** – this element overrides the configuration setting.

```
<?xml version="1.0" encoding="UTF-8"?>
<topics xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="schemas/icd_topic_definition.xsd">
  <device name="Dev2">
    <service name="MSA1" interface="MSAzimuth">
      <address_uri>zpb.rr://m4lsv.pl.eso.org:5002</address_uri>
      <performance rate_hz="0.2" latency_ms="25" synchronous="false" />
      <qos connection_timeout_sec="0.1" reply_timeout_sec="20" />
      <mal>
        <zpb />
      </mal>
      <methods>
        <method name="SomeMethodThatIsPartOfMSAzimuth">
          <performance rate_hz="10" latency_ms="2" synchronous="false" />
          <qos connection_timeout_sec="0.1" reply_timeout_sec="20" />
        </method>
        <method name="GetLogLevel">
          <is_configuration>true</is_configuration>
        </method>
      </methods>
    </service>
  </device>
</topics>
```




Listing 22: Service definition (full)

For all Java classes the package name is predefined. The device is a public class with static class members for all defined service interfaces (and topics). The name of the service class is derived from the service **name** attribute with the **Service** suffix. The **Methods** class name is also predefined. The class member constants are derived from the other elements listed above.

```
package elt.icd.topicdefinition;

public class Dev2 {

    public static class MSA1Service {
        public static String INTERFACE = "MSAzimuth";
        public static String ADDRESS_URI = "zpb.rr://m4lsv.pl.eso.org:5002";
        public static double PERFORMANCE_RATE_HZ = 0.2;
        public static double PERFORMANCE_LATENCY_MS = 25;
        public static boolean PERFORMANCE_SYNCHRONOUS = false;
        public static double QOS_CONNECTION_TIMEOUT_SEC = 0.1;
        public static double QOS_REPLY_TIMEOUT_SEC = 20.0;
        public static class Methods {
            public static class SomeMethodThatIsPartOfMSAzimuth {
                public static double QOS_CONNECTION_TIMEOUT_SEC = 0.1;
                public static double QOS_REPLY_TIMEOUT_SEC = 20.0;
                public static double PERFORMANCE_RATE_HZ = 10.0;
                public static double PERFORMANCE_LATENCY_MS = 2;
                public static boolean PERFORMANCE_SYNCHRONOUS = false;
            } // SomeMethodThatIsPartOfMSAzimuth
            public static class GetLogLevel {
                public static boolean IS_CONFIGURATION = true;
            } // GetLogLevel
        } // Methods
    } // MSA1Service

} // Dev2
```

Listing 23: Service Java code

For C++ the tool generates only the header file. The top 3 namespaces are predefined. The device is a namespace with namespace members for all defined service interfaces (and topics). The name of the service namespace is derived from the service **name** attribute with the **Service** suffix. The **Methods** namespace is also predefined. The constants are derived from the other elements listed above.

```
#ifndef __DEV_2_HPP
#define __DEV_2_HPP

namespace elt
{
```

(continues on next page)



(continued from previous page)

```
namespace icd
{
namespace topicdefinition
{
namespace Dev2
{
namespace MSA1Service
{
constexpr char[] interface = "MSAzimuth";
constexpr char[] addressUri = "zpb.rr://m4lsv.pl.eso.org:5002";
constexpr double performanceRateHz = 0.2;
constexpr double performanceLatencyMs = 25;
constexpr bool performanceSynchronous = false;
constexpr double qosConnectionTimeoutSec = 0.1;
constexpr double qosReplyTimeoutSec = 20.0;

namespace Methods
{
namespace SomeMethodThatIsPartOfMSAzimuth
{
constexpr double performanceRateHz = 10.0;
constexpr double performanceLatencyMs = 2;
constexpr bool performanceSynchronous = false;
constexpr double qosConnectionTimeoutSec = 0.1;
constexpr double qosReplyTimeoutSec = 20.0;
} // namespace SomeMethodThatIsPartOfMSAzimuth
namespace GetLogLevel
{
constexpr bool isConfiguration = true;
} // namespace GetLogLevel
} // namespace Methods
} // namespace MSA1Service
} // namespace Dev2

} // namespace topicdefinition
} // namespace icd
} // namespace elt

#endif // __DEV_2_HPP
```

Listing 24: Service C++ code

For Python the tool generates a class for each device with class members for all defined service interfaces (and topics). The name of the service class is derived from the service **name** attribute with the **Service** suffix. The **Methods** class name is also predefined. The members are derived from the other elements listed above.



```
class Dev2:

    class MSA1Service:
        interface = "MSAzimuth";
        addressUri = "zpb.rr://m4lsv.pl.eso.org:5002";
        performanceRateHz = 0.2;
        performanceLatencyMs = 25;
        performanceSynchronous = False;
        qosConnectionTimeoutSec = 0.1;
        qosReplyTimeoutSec = 20.0;

        class Methods:
            class SomeMethodThatIsPartOfMSAzimuth:
                performanceRateHz = 10.0;
                performanceLatencyMs = 2;
                performanceSynchronous = False;
                qosConnectionTimeoutSec = 0.1;
                qosReplyTimeoutSec = 20.0;

            class GetLogLevel:
                isConfiguration = True;

        # end Methods
    # end MSA1Service

# end Dev2
```

Listing 25: Topic Python code

MAL MAPPINGS

User's Manual

Revision:	2.7
Status:	Released
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	ELT CII
Folder:	/trunk/deliverables/phase3
Document ID:	CSL-MAN-18-155953
File:	MAN-MAL_mappings.docx
Owner:	Matej Šekoranja, Cosylab Sweden
Last modification:	June 30, 2022
Created:	July 21, 2018

Prepared by	Reviewed by	Approved by
Matej Šekoranja, CSL SWE	Aljaž Podboršek Gregor Čuk	Gregor Čuk

Document History



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 93 of 589

Re-vi-sion	Date	Changed/rev-iewed	Sec-tion(s)	Modificatio n
0.1	2018-07-21	mseko-ranja	All	Created.
0.9	2018-08-03	mseko-ranja	3	Added Python sections, properties.
1.0	2018-08-10	mseko-ranja	All	Minor updates, release.
2.0	2019-02-02	mseko-ranja	All	Updated property names, removed duplicated parts from MAL API manual and Transfer document.
2.1	2019-04-16	mseko-ranja	3 Refer-ences	Updated OPC UA C++ properties. References update
2.2	2020-03-24	mseko-ranja	3	Added ZPB C++/Java send queue size property.
2.3	2021-09-07	dkumar	3	Added OPCUA logging description for open62541
2.4	2021-09-24	dkumar	3	Added ZPB C++/Java disable type check property.
2.5	2022-04-26	dkumar	3	Added ZPB C++ enable ephemeral port check
2.6	2022-06-23	dkumar	3	Added MudPi C++ broadcast interface property
2.7	2022-06-30	mschilli	3	Added URI subsections

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is a Middleware Abstraction Layer **user manual** document for the ELT Core Integration Infrastructure Software project.

Audience

Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 94 of 589

API	Application Programmers Interface
CII	Core Integration Infrastructure
DDS	Data Distribution Service
MAL	Middleware Abstraction Layer
MUDPI	Multicast UDP Interface
URI	Uniform Resource Identifier
ZPB	ZeroMQ and Protocol Buffers
OPC UA	OPC (Object Linking and Embedding for Process Control) Unified Architecture

References

1. Cosylab, MAL Design document, CSL-DOC-17-147260 v1.5
2. Cosylab, ELT CII MAL Transfer document, CSL-DOC-18-168015, version 1.8
3. Cosylab, Python MAL Bindings Design document, CSL-DOC-18-152412, version 1.3
4. Cosylab, MAL API User Manual document, CSL-MAN-17-150198, version 2.1



4.1 Overview

This document describes configuration options for the following MAL mappings: DDS, ZPB and OPC UA. All the details on the MAL API and its functionality are covered by the MAL design document [1] and MAL User Manual [4].



4.2 Installation

See MAL Mappings Transfer document [2] on how to install MAL mappings.



4.3 C++

4.3.1 DDS

URIs

The underlying middleware is multicast. Discovery is topic-oriented. Publisher and Subscriber use the same URI to indicate which topic they both join. Request-Reply is not supported (any more).

- Publisher: *dds.ps:///DDS-TOPIC-NAME*
- Subscriber: same as Publisher
- Requester: Not supported

When IPv6 is enabled (through native configuration of the middleware), no changes to the URIs are needed.

Details

CII URI syntax is of the following form: *<scheme>:[//<authority>]</path>[?<query>][#fragment]*. The Authority, Query and fragment URI components are not used in DDS MAL URIs, therefore the DDS MAL URI is simplified: *<scheme>:///path*

Schema component of DDS MAL URI uses *dds.ps* string for specifying a DDS middleware publish subscribe profile. The string *dds.rr* for a DDS middleware request reply profile is no longer supported. A mandatory URI path component needs to hold the name of the topic (for publish subscribe profile).

Examples of MAL DDS URIs:

- *dds.ps:///Sample* - publisher for topicName = Sample
- *dds.ps:///Sample* - subscriber connected to topicName = Sample

Mal Specific Properties

The following table enumerates mapping specific properties:



Name	Type	Description	Used by	De-fault
dds.domain	integer	DDS domain ID.	MAL	0
dds.qos.profile.library	string	DDS QoS profile library.	MAL	(none)
dds.qos.profile.name	string	DDS QoS profile name.	MAL	(none)
dds.qos.profile.name.participant	string	DDS QoS profile name for participant.	MAL	(none)
dds.qos.profile.name.topic	string	DDS QoS profile name for topic.	MAL	(none)
dds.qos.profile.name.writer	string	DDS QoS profile name for writer.	MAL	(none)
dds.qos.profile.name.reader	string	DDS QoS profile name for reader.	MAL	(none)
dds.qos.profile.name.subscriber	string	DDS QoS profile name for subscriber.	MAL	(none)
dds.qos.profile.name.publisher	string	DDS QoS profile name for publisher.	MAL	(none)
dds.wait_for_ack_timeout_ms	integer	DDS wait for ack. before tearing down writer and publisher.	MAL	5000ms

Table 1 C++ DDS MAL mapping specific properties

Additional info for DDS QoS profile name selection:

Custom QoS profiles are enabled only when a user sets the “dds.qos.profile.library” mal property pointing to the file containing the QoS profiles. If the user sets only the “dds.qos.profile.name” mal property, then the same QoS profile name will be used for the creation of the following DDS entities: participant, topic, writer, publisher, reader, and subscriber. User has the option to override the default QoS profile name when creating a participant (“dds.qos.profile.name.participant”), a topic (“dds.qos.profile.name.topic”), a writer (“dds.qos.profile.name.writer”), a reader (“dds.qos.profile.name.reader”), a subscriber (“dds.qos.profile.name.subscriber”), and a publisher (“dds.qos.profile.name.publisher”). If the user does not override the profile name for a specific DDS entity, the default profile name is used.

Additional info for DDS Wait for acknowledgment timeout property:

If this property is not provided, a default timeout of 5000 ms shall be assumed as the timeout value. If the timeout value is larger than 0 then `writer->wait_for_acknowledgements()` and `publisher->wait_for_acknowledgements()` shall be called before tearing down the writers and the publisher. This means that developers can still suppress calling `wait_for_acknowledgements` if they specify the new property with a timeout value of 0, otherwise the call to `wait_for_acks` will always happen from now on.



4.3.2 ZPB

URIs

The underlying middleware is peer-to-peer. Discovery is based on host/port, plus name-filtering per port. A publisher exposes a topic-name on a local port, a subscriber needs to know where the publisher is. A server exposes one or more server-interfaces on a local port, a requester needs to know where the server is.

- Publisher: `zpb.ps://*:PORT/ZPB-TOPIC-NAME`
- Subscriber: `zpb.ps://PUBLISHER-HOST:PUBLISHER-PORT/TOPIC-NAME`
- Requester: `zpb.rr://PUBLISHER-HOST:PUBLISHER-PORT/SERVER-NAME/SERVER-IF-NAME`

ZMQ supports IPv6, but it needs to be enabled on the ZMQ socket as an option (currently not done by the ZPB-MAL). Once it is enabled on the sockets, IPv6 addresses need to be used in the URIs.

Details

ZPB MAL uses `zpb.ps` for publish-subscribe and `zpb.rr` for request-reply string for its URI schema component. A server-based authority component must be specified in the URI, without unsupported user-info element. When binding to the socket, i.e. in case of a publisher and request-reply server, an asterisk wildcard character is allowed to specify ANY network interface to bind to. A mandatory URI path component needs to hold the name of the topic or service. A path can be hierarchical. An empty URI path is allowed when creating a server: a path component is used as a service name prefix for all services registered to that server. Query and fragment URI components are not used.

There are helper methods available to create ZPB URIs, e.g. `ZpbUtil.createPsUri(String host, int port, String topicName)` and `ZpbUtil.createRrUri(String host, int port, String instanceName)`.

Examples of ZPB URIs:

- `zpb.ps://*:12345/Sample` - publisher bound on ANY interface on port 12345, topicName = Sample:
- `zpb.ps://eltmgr01:12345/Sample` - subscriber connected to eltmgr01, port = 12345, topicName = Sample:
- `zpb.rr://localhost:5555/RobotControl` - server bound to localhost interface, port = 5555, service name prefix = RobotControl:
- `zpb.rr://10.0.0.1:5555/RobotControl/Robot1` - client connected to 10.0.0.1, port = 5555, service-Name = RobotControl/Robot1:



Mal Specific Properties

The following table enumerates mapping specific properties:

Name	Type	Description	Used by	Default
zpb.ps.slowJoinerDelayMs	integer	Slow joiner delay. Time in milliseconds to wait after publisher has been instantiated. This mitigates slow joiner delay problem but does NOT guarantee that first N messages will not be lost.	Publisher	0 (no delay)
zpb.scheduledExecutorThreads	integer	Number of thread pool threads to handle internal timer based tasks (e.g. QoS deadline timeouts).	MAL	3
zpb.rr.callTimeoutSec (C++ only)	integer	Client asynchronous method call timeout in seconds.	Client	0 (no timeout)
zpb.ps.zmq_sndhwm	integer	Publisher send queue size. Use to limit memory usage when dealing with large messages. 0 stands for unlimited.	Publisher	1000 (ZMQ timeout)
zpb.ps.disableTypeHashCheck	boolean	Disable subscriber's check for type hash correctness	Subscribers	false
zpb.rr.zmqReconnectIntervalMs	integer	Set reconnection interval in ms	Client	100 ms
zpb.rr.zmqReconnectIntervalMaxMs	integer	Set maximum reconnection interval in ms	Client	0 ms
zpb.rr.enableEphemeralPortCheck	boolean	Throw an exception if ephemeral port is used, otherwise log warning	Server	false

Table 2 C++ ZPB MAL mapping specific properties



4.3.3 OPC UA

URIs

The underlying middleware is connection-oriented. Discovery via host-port. The URI encodes a mapping between ICD-entity and OPC-UA node, expressed as a semicolon-separated list of tuples (namespaceId, objectId), where objectId is either a nodeId or a methodId, e.g. “2,HelloWorld/ScalarTypes/Int64;2,HelloWorld/ScalarTypes/Double”. The order of the tuples must match the order of the fields in the ICD struct.

- Publisher: not applicable
- Subscriber: same as Requester
- Requester: *opcua.rr://SERVER-HOST:SERVER-PORT/NAMESPACE#MAPPING*

OPC-UA supports IPv6. To use it, the SERVER-HOST must be specified according to Format for Literal IPv6 Addresses.

Mal Specific Properties

The following table enumerates mapping specific properties:



Name	Type	Description	Used by	Default
opc.asyncLoopExecutionPeriodMs	integer	The sleep period of the asynchronous process loop.	MAL	50
opc.asyncCallSubmitTimeoutMs	integer	Timeout for submitting an asynchronous call.	Client	1000
opc.asyncCallRetryPeriodMs	integer	The retry period for failed asynchronous calls.	MAL	250
opc.log4plus.filename	string	log4plus file name.	MAL	undefined
opc.ps.pollingPeriodMs	integer	Multi-node subscriber polling period in milliseconds.	Subscriber, MrvSubscriber	1000
opc.ps.outstandingPublishRequests	integer	Number of outstanding publish requests.	Publisher, Subscriber, MrvSubscriber	10
opc.ps.usePollSubscriber	boolean	Force subscriber to poll for value updates instead of a monitor subscription	Subscriber, MrvSubscriber	false
opc.ps.requestedPublishingIntervalMs	integer	Server-side sampling interval of a subscription (i.e. an interval to sample/poll/monitor an item)	Subscriber, MrvSubscriber	5000

Table 3 C++ OPC UA MAL mapping specific properties

Data Access

The DA profile is implemented using special dedicated MAL request-response `DataAccess` interface. The interface is defined as:

```
public interface DataAccess extends RrEntity {  
    <T extends DataEntity<T>> T read(Class<T> dataClass);  
    <T extends DataEntity<T>> void write(T value);  
}
```

Synchronous (`DataAccessSync`) and asynchronous (`DataAccessAsync`) versions are available that directly issue OPC-UA DA read and write requests to the OPC-UA server. OPC UA nodes need to be specified at client instantiation time using URI fragment component using the following format: `opcua.rr://<host>:<port>/namespace#nsId1,nodeId1;nsId2,nodeId2;...;nsIdN,nodeIdN;`

IDs can be either a string or an integer ID, where `nsId` stands for namespace ID and `nodeId` for node ID. Node values are then mapped to the structure of type `T` at the read method call, or structure of type



T mapped to the nodes values at the write method call. Note that order and type of the node values and structure fields must match. In case of mismatch, an exception is thrown. In case of an error or timeout (specified as QoS), an exception is thrown. Connection management is handled by MAL transparently, as for every request-reply interface: connection process is initiated asynchronously and automatically and in case of disconnection a connection is restored automatically.

The OPC UA MAL library provides a predefined set of structures that hold a value of the supported types (see design document). User defined structures must be defined using ICD and their OPC UA specific structures generated by the icd-gen tooling.

In addition, MAL Publisher interface is also implemented as DA write operation, i.e. publish call actually calls DA write and is instantiated with the same URI semantic.

Logging

Additional info for opcua logging:

Open62541 library (client part) will log events via log4cplus malOpcuaOpen62541 logger. If a user does not provide a log4cplus configuration with the malOpcuaOpen62541 logger or a root logger, no log output will be generated from the open62541 library. Bellow is an example configuration pertaining to the open62541 library logging: malOpcuaOpen62541 logger configured for only ERROR and FATAL logging with additivity set to false (events won't be propagated to the root logger). With this configuration all other mal loggers will log with TRACE level. User needs to provide the location of the log4cplus configuration for mal-opcua using mal property opc.log4cplus.filename (see table 3).

```
log4cplus.rootLogger=TRACE, stdout
log4cplus.logger.malOpcuaOpen62541=ERROR, stdout
log4cplus.additivity.malOpcuaOpen62541=false
```

```
log4cplus.appender.stdout=log4cplus::ConsoleAppender
log4cplus.appender.stdout.layout=log4cplus::PatternLayout
log4cplus.appender.stdout.layout.ConversionPattern=%5p [%t] (%F:%L) - %m%n
```

4.3.4 MUDPI

URIs

The underlying middleware is packet-oriented, packets are addressed via host-port. The Publisher needs to know the address of the subscriber. The Publisher binds to a random local port. The publisher sends unicast or multicast depending on the address encoded in the URI (multicast typically: 224.0.0.0 - 239.255.255.255). In case of unicast, the subscriber must (obviously) run on the host described the URI. In case of multicast, the publisher can choose the outgoing_interface by passing a property to *getPublisher(...)*, see section Mal Specific Properties below.



- Publisher: *mudpi.ps://SUBSCRIBER-HOST:SUBSCRIBER-PORT/MUDPI-TOPIC-NAME*
- Subscriber: same as Publisher
- Requester: *mudpi.rr://SERVER-HOST:SERVER-PORT/MUDPI-SERVER-IF-NAME*

Examples of URIs

- *mudpi.ps:///127.0.0.1:12781/TELIF* - unicast on loopback interface
- *mudpi.ps://224.0.0.1:12000/TELIF* - multicast in subnet
- *mudpi.ps://134.171.3.160:12000/TELIF* - unicast to host 134.171.3.160

Mal Specific Properties

The following table enumerates mapping specific properties:

Name	Type	Description	Used by	Default
mudpi.ps. interfaceName	inter- string	Broadcast interface name for mudpi sender Calling setInterface	Pub- lisher	not set (empty string)



4.4 Java

4.4.1 DDS

Mal Specific Properties

For mapping specific properties refer to Table 1 C++ DDS MAL mapping specific properties.

4.4.2 ZPB

Mal Specific Properties

For mapping specific properties refer to Table 2 C++ ZPB MAL mapping specific properties.

4.4.3 OPC UA

Mal Specific Properties

The following table enumerates mapping specific properties:

Name	Type	Description	Used by	Default
opc.scheduledExecutorThreads	integer	Number of thread pool threads to handle internal timer based tasks (e.g. QoS deadline timeouts).	MAL	3
opc.ps.pollingPeriodMs	integer	Multi-node subscriber polling period in milliseconds.	Subscriber, MrvSubscriber	1000
opc.ps.reqPubRateMs	integer	OPC UA stack subscription requested publishing interval in milliseconds.	MAL	1000
opc.ps.queueSize	integer	OPC UA stack subscription queue size.	MAL	3

Table 4 Java OPC UA MAL mapping specific properties



4.5 Python

Python MAL API builds on MAL C++ API foundation. There is no native Python MAL mapping implementation, everything is provided via language mapping to C++.

Python applications use Python MAL API only and are not aware of the actual MAL implementation details.

Refer to Python Bindings Design [3] and MAL API User Manual [4] documents for detailed information.

MAL PYTHON MAPPINGS

Revision:	1.0
Status:	Released
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	E-ELT Core Integration Infrastructure Software
Folder:	/trunk/deliverables/phase3
Document ID:	CSL-DOC-18-162244
File:	MAN-Python_mappings.docx
Owner:	Matej Šekoranja, Cosylab Sweden
Last modification:	August 16, 2018
Created:	July 25, 2018

Prepared by	Reviewed by	Approved by
Matej Šekoranja (CSL SWE)	Aljaž Podboršek Gregor Čuk	Gregor Čuk

Document History

Revision	Date	Changed/reviewed	Section(s)	Modification
0.1	2018-07-25	msekoranja		Created.
1.0	2018-08-10	msekoranja		Minor reviewer changes, released.

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is a Middleware Abstraction Layer Python mappings user manual for the E-ELT Core Integration Infrastructure Software project.



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 108 of 589

Audience

All Cosylab and ESO employees involved with the E-ELT Core Integration Infrastructure Software.

Glossary of Terms

AMI	Asynchronous Method Invocation
API	Application Programmers Interface
CII	Core Integration Infrastructure
DDS	Data Distribution Service
GC	Garbage Collector
GIL	Global Interpreter Lock
MAL	Middleware Abstraction Layer
QoS	Quality Of Service
STL	Standard Template Library
URI	Uniform Resource Identifier
ZPB	ZeroMQ and Protocol Buffers

References

1. ESO, E-ELT Control System Architecture Concepts, ESO-283831_1
1. Cosylab, MAL ZPB Design Document, CSL-DOC-18-150609, version 1.0
2. Cosylab, MAL DDS Design Document, CSL-DOC-18-150607, version 1.0
3. <https://pybind11.readthedocs.io/>



5.1 Introduction

This document provides information on use of Python MAL mappings interface also referred to as Python MAL API.

Python MAL API builds on MAL C++ API foundation. It provides programming interface that is very similar to the programming interface that is exposed by MAL C++ API.

Python MAL API is compatible with Python 3.5 as provided with ESO E-ELT development environment v2.1.8.

Python applications use Python MAL API and are not aware of the actual MAL mappings details. The **CiiFactory** interface is the primary means to access MAL functionality.



5.2 Prerequisites

In order to use Python MAL API the following prerequisites must be fulfilled:

- ESO E-ELT Development Environment v2.1.8 must be properly set up.
- Directory hierarchy with installed MAL products must be established (commonly referred to as DATAROOT).
- MAL and MAL COMMON products must be installed under DATAROOT.
- Additional C++ MAL mappings must be installed (for example MAL DDS) under DATAROOT.
- Interface modules for specific data types, generated by ICD GENERATOR or created manually must be installed under DATAROOT.
- Linker path for dynamic loading of libraries must be properly set. On Linux this means that LD_LIBRARY_PATH environment variable should contain reference to the **lib64** subdirectory under DATAROOT. In more general terms, linker path must be setup in such a way that shared libraries required by Python modules and MAL mapping in use are within that path.
- PYTHONPATH environment variable must point to directory **lib/python3.5/site-packages** directory within DATAROOT hierarchy, for example on Linux:

```
# when using bash
```

```
::
```

```
# Define DATAROOT
```

```
export DATAROOT=~ /INTROOT
```

```
::
```

```
# Do other stuff ...
```

```
::
```

```
# Define LD_LIBRARY_PATH
```

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$DATAROOT/lib64
```

```
::
```

```
# Define PYTHONPATH
```

```
export PYTHONPATH=$PYTHONPATH:$DATAROOT/lib/python3.5/site-packages
```



5.3 Using Python MAL API

This chapter provides usage information for Python MAL API.

Python MAL API is implemented as language binding to C++ MAL mapping therefore it closely follows C++ API in its resemblance and operation.

5.3.1 Importing top level Python module, obtaining CiiFactory reference

Top level Python MAL API module is **elt.pymal**. It is recommended to be imported with alias import to mal.

Once module is imported, reference to CiiFactory can be obtained with call to **mal.CiiFactory.getInstance**. This call returns reference to the shared CiiFactory instance.

Alternatively, **malCiiFactory.createInstance** call creates new instance of the CiiFactory.

The following example provides basic steps:

```
#!/usr/bin/env python

# import top level Python MAL API module
import elt.pymal as mal

# use mal interface to obtain reference to shared CiiFactory instance
ciiFactory = mal.CiiFactory.getInstance()
```

elt.pymal module provides the following interfaces (classes or methods):

- CiiFactory, CII Factory class interface
- Exceptions, MalException, TimeoutException, SchemeNotSupportedException, IllegalUriException, IllegalStateExceptions
- Uri, wrapper for C++ `elt::mal::Uri` class implementation. Not needed by general applications, as methods within Python API that require uri as one of the parameters take a string as input.
- loadMal, method for loading MAL mappings
- PsQoSList, python class helper for creating list of Publish-Subscribe QoS options required by Publisher-Subscriber interface.
- RrQoSList, python helper for creating list of Request-Reply QoS options, required by Request-Reply interface.



5.3.2 Loading and registering MAL mappings

When imported, Python MAL API does not load any MAL mapping by default. Application must load required MAL mappings dynamically, by using **loadMal** method. **loadMal** method takes two required arguments:

- **malName**, string, name of the MAL mapping. Supported values are:
 - dds, for DDS MAL
 - zpb, for ZPB MAL
 - opcua, for OPCUA MAL
- **malProperties**, python dictionary with properties, used to initialize MAL mapping.

loadMal method returns opaque handle to the MAL mapping. In case MAL mapping could not be loaded, exception will be raised.

Opaque handle to the loaded MAL mapping must then be passed to the **CiiFactory.registerMal** class method as second argument. First argument to this method is name of the URI scheme that this particular MAL mapping covers.

Example that follows, illustrates loading and registering two separate MAL mappings:

```
# !/usr/bin/env python

# import top level python MAL API module
import elt.pymal as mal

# obtain reference to CiiFactory
ciiFactory = mal.CiiFactory.getInstance()

# Prepare MAL properties. Note that different MAL mappings
# require specific properties to be initially set. Refer to
# documentation for specific MAL mapping for details.
malProperties = {}

# Dynamically load DDS MAL mapping
ddsmal = mal.loadMal('dds', malProperties)

# Register DDS MAL for two schemes with CiiFactory
ciiFactory.registerMal('dds.ps', ddsmal)
ciiFactory.registerMal('dds.rr', ddsmal)

# Dynamically load ZPB MAL mapping
zpbmal = mal.loadMal('zpb', malProperties)

# Register ZPB MAL with CiiFactory
ciiFactory.registerMAL('zpb.ps', zpbmal)
```




5.3.3 Python MAL API relative time parameters (timeouts/durations)

Many API calls require relative time parameters, mostly for timeouts. Relative time parameters are also used within constructors of QoS parameters.

Python MAL API expects instance of **datetime.timedelta** class when relative time parameter is required. For example:

```
import datetime
# ...
publishTimeout = datetime.timedelta(milliseconds=500)
publisher.publish(sample, publishTimeout)
```

5.3.4 Using Publisher-Subscriber interface

Publisher-Subscriber interface relies on generated or handcrafted mapping module that provides Python interface for specific data entity type. Data entity type is typically represented as Python class within mapping module that must be imported. This Python class is a key to obtaining correct Subscriber or Publisher instances from the CiiFactory. Both **CiiFactory.getSubscriber** and **CiiFactory.getPublisher** methods expect reference to the data entity class (not instance) as second parameter. Both methods take URI as the first argument, properties dictionary as third and list of QoS parameters as fourth argument. References obtained with these methods should be checked to be valid (not None).

Code to obtain reference to the subscriber generally looks like:

```
#!/usr/bin/env python
# import top level python MAL API module
import elt.pymal as mal

# Import data entity class. This usually resides in the interface module that
# is either handcrafted or generated.
# In this case TestSample class is our data entity class. Data entity classes
# must provide specific interface to the rest of the Python bindings,
# so they are usually generated automatically.
# Public interface provided by data entity class is specified by data entity
# definition for generator.

from TestSampleIf import TestSample

# obtain reference to CiiFactory
ciiFactory = mal.CiiFactory.getInstance()

# Prepare MAL properties
malProperties = {}

# load and register MAL with CiiFactory
zpbmal = mal.loadMal('zpb', malProperties)
```

(continues on next page)



(continued from previous page)

```
ciiFactory.registerMal('zpb.ps', malProperties)

# prepare subscriber specific MAL properties
subscriberMalProperties = {}

# obtain reference to Subscriber

subscriber = factory.getSubscriber(
    'zpb.ps://localhost:13455/test', TestSample, mal.ps.qos.DEFAULT,
    subscriberMalProperties)
if subscriber is None:
    throw RuntimeError('could not obtain subscriber')

# Read all available samples
for sample in subscriber.read(0):
    # Use interface specific to data entity type, to get values
    print(sample.getSampleId())
# Close the subscriber
subscriber.close()
```

Code to obtain reference to the publisher is similar:

```
#!/usr/bin/env python

import datetime
# import top level python MAL API module
import elt.pymal as mal

# Import data entity class. This usually resides in the interface module that
# is either handcrafted or generated.
# In this case TestSample class is our data entity class. Data entity classes
# must provide specific interface to the rest of the Python bindings,
# so they are usually generated automatically.
# Public interface provided by data entity class is specified by data entity
# definition for generator.

from TestSampleIf import TestSample

# obtain reference to CiiFactory
ciiFactory = mal.CiiFactory.getInstance()

# Prepare MAL properties
malProperties = {}

# load and register mAL with CiiFactory
zpbmal = mal.loadMal('zpb', malProperties)
ciiFactory.registerMal('zpb.ps', malProperties)

# prepare publisher specific MAL properties
```

(continues on next page)



(continued from previous page)

```
publisherMalProperties = {}  
  
# obtain reference to Publisher  
  
publisher = factory.getPublisher(  
    'zpb.ps://localhost:13455/test', TestSample, mal.ps.qos.DEFAULT,  
    publisherMalProperties)  
if publisher is None:  
    throw RuntimeError('Could not obtain publisher')  
  
# Use the publisher to create new data entity  
dataEntity = publisher.createDataEntity()  
  
# Use interface specific to data entity type, to set values  
dataEntity.setSampleId(100)  
dataEntity.setText('test')  
# Publish sample  
publisher.publish(dataEntity, datetime.timedelta(seconds=1))  
# Close the publisher  
publisher.close()
```

Using Subscriber/Publisher as a Context Manager

Subscriber and Publisher can be closed explicitly by calling **close** method or implicitly when they are destroyed by Python garbage collector. Alternative is to use them as Context Managers. In this case, underlying objects are automatically closed as soon as containing context is destroyed.

See the following example for illustration:

```
# ...setup...  
  
uri = 'zpb.ps://localhost:13455/test'  
properties = {}  
qos = mal.ps.qos.DEFAULT  
# Use subscriber as Context Manager  
with ciiFactory.getSubscriber(uri, TestSample, qos, properties) as subscriber:  
    samples = subscriber.read(0)  
# subscriber was automatically closed at the end of previous context
```



Constructing list of QoS parameters

Construction of Subscriber or Publisher object requires list of QoS parameters. Due to underlying implementation this list cannot be simple Python list.

In case list contains no elements, predefined symbol **mal.ps.qos.DEFAULT** should be used.

To construct non empty list, special helper class **mal.PsQoSList** has to be used. Constructor to this class takes either individual QoS parameters or Python list with individual QoS parameters:

```
import elt.pymal as mal

# ...

# Individual QoS parameters
deadline = mal.ps.qos.Deadline(datetime.timedelta(0, 30))
latency = mal.ps.qos.Latency(datetime.timedelta(0, 20))

# Construct QoS list from individual elements
qosList1 = mal.PsQoSList(deadline, latency)

# Construct QoS list from Python list
lst = [deadline, latency]
qosList2 = mal.PsQoSList(lst)
```

Constructing Data Event filters

Data Event filters are related to type of data entity that is to be filtered. Filter must be constructed for specific data entity class (not instance). To construct specific event filter, class methods **mal.ps.DataEventFilter.new** or **mal.ps.DataEventFilter.all** must be used with data entity class as argument as in following example:

```
import elt.pymal as mal

# Import data entity class definition
from TestSampleIf import TestSample

# ...

# Obtain reference to CiiFactory
ciiFactory = mal.CiiFactory.getInstance()
# Load and register MAL implementation to be used
# ...

# Create a subscriber
subscriber = factory.getSubscriber(
    'zpb.ps://localhost:13455/test', TestSample, mal.ps.qos.DEFAULT, {})

# create instance for filtering
```

(continues on next page)



(continued from previous page)

```
filterInstance = subscriber.createDataEntity()

# set parameters in filter instance, use API supported by specific data entity

# Create Data event filter. Default filter performs no filtering.
# Parameter is data entity class for which the filter is created.

eventFilter = mal.ps.DataEventFilter.new(TestSample)

# Set filter instance
eventFilter.setFilterInstance(filterInstance)

# Set filter events, input parameter is Python set of events
eventFilter.setInstanceEvents(set(mal.ps.InstanceEvent.UPDATED,
    mal.ps.InstanceEvent.REMOVED))

# use event filter...
```

Both methods return default filter for that data entity.

5.3.5 Using Request-Reply interface

As with Publisher-Subscriber, Request-Reply interface also relies on generated or handcrafted binding module that provides Python interface for specific Request-Reply client. Client can be implemented as either synchronous or asynchronous. Request-Reply client is typically represented as Python class within binding module that must be imported. This python class is key to obtaining correct Client instance from the CiiFactory. Method **CiiFactory.getClient** expects reference to the data Request-Reply client class (not instance) as second parameter. Method takes URI as the first argument, list of QoS parameters as third argument and dictionary of properties as the fourth argument.

The following example illustrates the concept:

```
#!/usr/bin/env python

# Import top level Python API module
import elt.pymal as mal

# from mapping module implementing specific Request-Reply interface,
# import client class
from MathModule import MathSyncClient

# obtain reference to CiiFactory, load MAL implementaton and register MAL with
# CiiFactory
ciiFactory = mal.CiiFactory.getInstance()
# ...

# Obtain synchronous client for Math RR interface
```

(continues on next page)



(continued from previous page)

```
client = factory.getClient(uri, MathSyncClient, mal.rr.qos.DEFAULT, {})  
  
# issue explicit connect request. This is asynchronous, therefore  
# we must wait on future returned from connect() method.  
  
client.connect().get()  
  
if client:  
    # issue explicit connect request. This is asynchronous, therefore  
    # we must wait on future returned from connect() method.  
  
    client.connect().get()  
  
    # This is synchronous client, client will wait for result to arrive  
    sum = client.sum(10.1, 15.10)  
    print(sum)  
    # once done, close the client  
    client.close()
```

Using Request-Reply client as Context Manager

Request-Reply clients also support Context Manager protocol to allow for closing client automatically on context exit.

```
# setup...  
#  
#  
qos = mal.rr.qos.DEFAULT  
  
with ciiFactory.getClient(uri, MathSyncClient, qos, {}) as client:  
    client.connect().get()  
    sum = client.sum(10.1, 15.10)  
  
# Client was automatically closed by now
```

Constructing list of QoS parameters

Construction of Client object requires list of QoS parameters. Due to underlying implementation this list cannot be simple Python list.

In case list contains no elements, predefined symbol **mal.rr.qos.DEFAULT** should be used.

To construct non empty list, special helper class **mal.RrQoSList** has to be used. Constructor to this class takes either individual QoS parameters or Python list with individual QoS parameters:



```
import elt.pymal as mal

# ...

# Individual QoS parameters
ct = mal.rr.qos.ConnectionTime(datetime.timedelta(0, 5))
rt = mal.rr.qos.ReplyTime(datetime.timedelta(0, 1))

# Construct QoS list from individual elements
qosList1 = mal.RsQoSlist(ct, rt)

# Construct QoS list from Python list
lst = [ct, rt]
qosList2 = mal.RrQoSlist(lst)
```

CONFIGURATION

Revision:	0.8
Status:	Final Draft
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	ELT CII
Document ID:	
File:	config-ng.rst
Owner:	Marcus Schilling
Last modification:	March 17, 2022

Document history

Revision	Date	Changed/ Reviewed	Section(s)	Modification
0.8	15.03.2022	Jure Repinc Marcus Schilling	All	Document created / reviewed

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is manual for the configuration system used by applications in the ELT Core Integration Infrastructure Software.

Audience

This document is aimed at those Cosylab and ESO employees involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 121 of 589

API	Application Programming Interface
CII	Core Integration Infrastructure
MDI	Metadata instance
YAML	YAML Ain't Markup Language



6.1 Overview

This is the user manual for the CII Configuration API (Config-ng API). All examples in this manual are also present in the `cii-demo` repository.

The Config-ng API is provided for C++ and Python. It provides support for parsing YAML documents (with support for additional information or directives via tags) into configuration documents. A Configuration document contains set of named instances with optional data type and metadata information.

A Configuration document allows access to its instances and supports additional operations like saving, merging with another configuration document and updating with programmatically prepared instance map.



6.2 Includes and imports

6.2.1 C++

To use the config-ng API from C++, the following directives are needed:

wscript (project)

```
cnf.check_wdep(wdep_name='config-ng.cpp.config-ng', uselib_store='config-ng')  
# Transitive Dependencies:  
cnf.check_cfg(package='log4cplus', uselib_store='log4cplus', args='--cflags --  
↪libs')
```

wscript (module)

```
use=[...,  
     'config-ng'  
     # Transitive Dependencies:  
     'log4cplus'
```

source

```
#include <config-ng/ciiConfigApi.hpp>
```

6.2.2 Python

To use the config-ng API from Python, the following directives are needed:

wscript (project)

```
cnf.check_wdep(wdep_name='config-ng.python.config-ng', uselib_store='config-ng')
```

wscript (module)

```
use=[..., 'config-ng' ]
```

source

```
import elt.configng
```



6.3 API Components

API provides the following basic components:

- Client
- Document
- Node
- Data type
- Metadata

Client is the entry point to the API. It provides basic operations like loading document from provided filename/URI, retrieval or modification of **config-ng** search path.

Document represents a configuration object that was either loaded from file/stream or created in memory. Operations like loading produce a new document instance. Document itself offers additional operations (like merging, check, update and save).

Node represents an item (also known as instance) within a configuration object. Generally, it can have a name and can represent a scalar (single value), a list of items or a map (dictionary) of items. Access to nodes is available through the document instance interface described later.

Data type describes a data type assigned to the configuration instance (Node) if any. Data type of instance is optional. There is a set of built-in data types. Defining custom data types is supported by using a predefined yaml syntax. A data type definition can contain its own metadata attributes. Some predefined attributes are understood by the **config-ng** (when running the Check() operation on a document):

default Provides default value of the configuration item when its value is not explicitly provided.

min Allowed minimum value of the scalar configuration item.

max Allowed maximum value of the scalar configuration item.

allowed_values Provides list of allowed scalar values for scalar configuration item.

Metadata is associated with Node and can be seen as a map of attributes assigned to item. These attributes have a name and can be scalars, lists, or maps. There is no data type assigned to the attributes of the item. Some attributes are automatically generated by the **config-ng** at the time of the parsing of the initial yaml document. There are **length** for lists and **rows** and **columns** for matrices.



6.4 Creation of Documents

Documents can be created by parsing YAML source either from input/output stream or provided URI. Additionally documents can also be constructed programmatically.

6.4.1 Parsing YAML source from input/output stream

Note that parsing document from YAML can cause exceptions either within **config-ng** or within the underlying YAML parser, therefore proper exception handling is needed (refer to **config-ng** examples).

C++

In C++, YAML source can be provided through a **std::istream** object as demonstrated below:

```
#include <iostream>
#include <sstream>
#include <config-ng/ciiConfigApi.hpp>

int main() {
    const char *document_source = R"DOC(
        a: 10
        b: this is a string
    )DOC";
    std::stringstream document_source_stream;
    document_source_stream << document_source;
    ::elt::configng::CiiConfigDocument document =
↪::elt::configng::CiiConfigClient::Load(document_source_stream);
    std::cout << document.Instances()["a"].As<std::int>() << ", "
        << document.Instances()["b"].As<std::string>()
        << std::endl;

    return 0;
}
```

Python

Loading from Python's standard **io** objects is supported:

```
import io
import elt.configng
document_source = """
    a: 10
    b: this is a string
"""
document = elt.configng.CiiConfigClient.load_from_stream(io.StringIO(document_
↪source))
```

(continues on next page)



(continued from previous page)

```
print(document.instances.a.as_value())
print(document.instances.b.as_value())
```

6.4.2 Loading YAML source from file/URI

Documents can be loaded from a local file by specifying a URI or a filename. The supported URI syntax is *cii.config://<authority>/<path>* where:

<authority> can be either **local**, **remote** or *****. Only **local** authority is currently supported.

<path> is path to the file. Note that path is always interpreted as relative regarding to the search path currently set in the client. Absolute paths are not supported.

In case a plain filename is provided, it is turned into a **local** URI:

```
my/path/filename.yaml -> cii.config://local/my/path/filename.yaml
```

Search Path

Client maintains a search path for files. The search path consists of a list of directories to be searched for a specific file. Directories are searched in the order they appear in the list. First match of the filename wins. From a user point of view, the search path is a string containing zero or more directories delimited by **:** (i.e. *"/path/to/dir1:/path/to/dir2:path/to/dir3"*).

The search path can be set to the empty string. In this case, search for files is limited to only relative to current working directory of the process.

Client initializes the search path from **CFGPATH** environment variable, if defined. To obtain the current search path, it provides method `GetSearchPath()` in C++ and `get_search_path()` in Python. When the search path was initialized from **CFGPATH**, the current working directory is prepended to the search path defined in **CFGPATH**.

To modify the search path `SetSearchPath(new_path)` is provided in C++ and `set_search_path(new_path)` in Python. The following examples demonstrate manipulation of the search path.

C++ example

```
// save current search path
std::string saved_search_path =
↳::elt::configng::CiiConfigClient::GetSearchPath();
std::string my_search_path = "/usr/local/conf:/local/conf";
// establish own search path
↳::elt::configng::CiiConfigClient::SetSearchPath(my_search_path);
// do something
...
```

(continues on next page)



(continued from previous page)

```
// restore search path  
::elt::configng::CiiConfigClient::SetSearchPath(saved_search_path);
```

Python example

```
# save current search path  
saved_search_path = elt.configng.CiiConfigClient.get_search_path()  
# establish own search path  
my_search_path = '/usr/local/conf:/local/conf'  
elt.configng.CiiConfigClient.set_search_path(my_search_path)  
# do something  
...  
# restore search path  
elt.configng.CiiConfigClient.set_search_path(saved_search_path)
```

6.4.3 Actual document loading

C++ Document is loaded by calling method `::elt::configng::CiiConfigClient::Load()` as demonstrated below. Note that parsing YAML and loading document can produce exceptions therefore exception handling is needed (refer to examples for more information):

```
// Load path/file.yaml using current search path  
::elt::configng::CiiConfigDocument document =  
↪::elt::configng::CiiConfigClient::Load("cii.config://local/path/file.yaml")
```

In the same manner as C++, Python version provides method `elt.configng.CiiConfigClient.load()`:

```
# Load path/file.yaml usign current search path  
::elt::configng::CiiConfigDocument document = elt.configng.CiiConfigClient.load(  
↪'cii.config://local/path/file.yaml')
```

6.4.4 Building document programmatically

To construct a document programatically, a special helper class **CiiConfigMapBuilder** must be used to construct the 'map' of the instances to be added to document. **CiiConfigMapBuilder** provides a simple interface to add elements to the map. The general workflow is like this:

- Create empty document instance
- Use map builder to prepare an instance map of instances to be added to the document
- Update document with prepared map. With C++, use method `SetInstancesFromMap()` on the document instance. With Python, use method `set_instances_from_map()` on the document instance.



Note that on the API level there are slight differences between Python and C++ API (due of language differences), but the general workflow is the same.

C++ Example

```
// See programmatic example for more comprehensive information
#include <vector>
#include <string>
#include <iostream>
#include <config-ng/ciiConfigApi.hpp>
// NOTE: additional includes needed
#include <config-ng/ciiConfigMapBuilder.hpp>
#include <config-ng/ciiConfigValueConverter.hpp>

int main() {
    // prepare empty document
    ::elt::configng::CiiConfigDocument document =
↪ ::elt::configng::CiiConfigDocument();
    // Create CiiConfigMapBuilder
    ::elt::configng::CiiConfigMapBuilder builder;
    // Builder provides operators to simplify construction of node hierarchy
    // Note that values cannot directly be assigned to the builder items,
    // factory method ::elt::configng::CiiConfigInstanceNode::From() must be used
    // to turn C++ value into correct node to be later assigned to document

    // Create untyped instance a and assign value 3 to it
    builder["a"] = ::elt::configng::CiiConfigInstanceNode::From(3);
    // Create untyped instance b and assign list 1,2,3 to it
    builder["b"] = ::elt::configng::CiiConfigInstanceNode::From(std::vector<int>
↪ {1,2,3});
    // Create type instance c, using builtin type vector_int32
    builder["c"] = ::elt::configng::CiiConfigInstanceNode::From(std::vector<int>
↪ {1,2,3},
        "!cfg.type:vector_int32");
    // Create map with untyped items a and b
    builder["map"]["a"] = ::elt::configng::CiiConfigInstanceNode::From(3);
    builder["map"]["b"] = ::elt::configng::CiiConfigInstanceNode::From("a string
↪");

    // Once map is constructed, document can be updated
    document.SetInstancesFromMap(builder);
    // Display instance a
    std::cout << document.Instances()["a"].as<int>();
    return 0;
}
```

Python Example

```
# See programmatic example for more comprehensive information
import elt.configng
# Prepare empty document
```

(continues on next page)



(continued from previous page)

```
document = elt.configng.CiiConfigDocument()
# Create CiiConfigMapBuilder
builder = elt.configng.CiiConfigMapBuilder()
# Create untyped instance a and assign value 3 to it
builder['a'] = 3
# Create untyped instance b and assign list [1,2,3] to it
builder['b'] = [1,2,3]
# Create typed instance c and assign list [1,2,3] to it.
# Whenever typed instance is required, use of elt.configng.CiiConfigNodeFactory.
↳make_from
# method must be used, as it takes additional parameters optional tag denoting
↳data type
# to be used when creating instance and optional converter method for custom
# value conversion when required
builder['c'] = elt.configng.CiiConfigNodeFactory.make_from([1,2,3], '!cfg.
↳type:vector_int32')
# Create map with untyped items a and b
builder['map']['a'] = 3
builder['map']['b'] = 'a string'
# Once map is constructed, document can be updated
document.set_instances_from_map(builder)
# Display instance a
print(document.instances.a.as_value())
```



6.5 Access to YAML source tree of a document

Once a document is loaded from YAML stream or file, the original YAML node tree is available to the user through `CiiConfigDocument::GetYamlSourceRootNode()` in C++ and `CiiConfigDocument.get_yaml_source_root_node()` in Python.

To operate on the returned YAML root node in C++, **yaml-cpp** API must be used.

To operate on the returned YAML root node in Python, **PyYAML** API must be used.

Note that subsequent changes on the returned YAML node tree do not reflect on the config document, and vice versa, subsequent changes on the config document do not reflect on the YAML node tree. It is recommended that when dealing directly with the YAML node tree, the best course is to save the modified tree as YAML into file or stream, and then load it into another config document. The same, analogously, applies to changes on the config document: save it to YAML stream/file, and then load it directly with a YAML parser.



6.6 Document instance interface

Once the document is created or loaded, access to its instances is available via the document instance

name Name of the instance (only valid with top level and Map member instances)

type Three user visible types are predefined: Scalar, Sequence, Map.

metadata Metadata attributes of the instance.

data type Optional config ng data type associated with the instance

origin Source file location (line, column) of the YAML element from which the instance has been constructed.

Sequence and Map instances provide non recursive (by default) and recursive iterators and methods to access the instance value.

6.6.1 Accessing document instance interface

C++

The document instance interface is available through method `::elt::configng::CiiConfigDocument::Instances()` method. i.e.:

```
::elt::configng::CiiConfigDocument document = ::elt::configng.  
↳CiiConfigClient::Load("cii.config://local/file.yaml");  
// For reading only  
const ::elt::configng::CiiConfigInstanceNamespace &instances = document.  
↳Instances();
```

Access to a specific instance is done through operator[]:

```
// Target type is known, instance value will be automatically  
// converted to target type (note exception is thrown if conversion not possible)  
double d = instances["theDouble"];  
// Target type is not known, conversion method As() must be used  
std::cout << instances["theDouble"].As<double>() << std::endl;  
// Assign new value to the instance  
instances["theDouble"].Assign(6.7);
```

Obtaining additional information:

```
::elt::configng::CiiConfigInstanceNamespace &theDouble = document.Instances() [  
↳"theDouble"];  
// Check the type  
if (theDouble.IsScalar()) {  
    ...
```

(continues on next page)



(continued from previous page)

```
} else if (theDouble.IsSequence()) {  
    // Tterate non recursively over instance  
    for (const auto &item: instance) {  
        ...  
    }  
    ...  
} else if (theDouble.IsMap()) {  
    ...  
}  
// Obtain the origin  
::elt::configng::CiiConfigItemOrigin origin = theDouble.GetOrigin();  
// Display the origin  
std::cout << "Filename: " << origin.source << " line: " << origin.line << " "  
↳column: "  
    << origin.column << std::endl;  
// obtain config-ng data type of the instance  
::elt::configng::CiiConfigDataType data_type = theDouble.GetDataType();  
// access metadata of the instance  
::elt::configng::CiiConfigMetadata &metadata = theDouble.GetMetadata();  
// and metadata properties  
if (metadata.Has("max")) {  
    std::cout << "theDouble has max property: " << metadata["max"].As<double>()  
        << std::endl;  
}
```

Python

The document instance interface is available through method `elt.config.CiiConfigDocument.get_instances()`, or through the attribute accessor **instances** on the document object, i.e.:

```
document = elt.config.CiiConfigClient.load('cii.config://local/file.yaml')  
# Use attribute accessor instances to access instance 'theDouble'  
theDouble = document.instances.theDouble  
# Or the classic way  
instances = document.get_instances()  
theDouble = instances['theDouble']  
# Print value of the instance  
print(theDouble.get_value())  
# Alternative to obtain the value  
# is to use as_value() method which takes optional converter argument  
# that is actually a function to call when instance value needs to be  
# converted to python value  
print(theDouble.as_value(lambda x: x.get_value() + 1))  
# Assign new value to the instance  
theDouble.assign(6.7)
```

Obtaining additional information:



```
theDouble = document.instances.theDouble
# Check the type
if theDouble.is_scalar():
    ...
elif theDouble.is_sequence():
    # Iterate non recursively over instance
    for item in theDouble:
        print(item.get_value())
elif theDouble.is_map():
    ...
# Obtain origin of the instance
origin = theDouble.get_origin()
# Display the origin, get_info() method returns tuple (filename, line, column)
# indicating origin of the instance
print('Origin filename %s, line: %s, column: %s' % origin.get_info())
# Obtain config-ng data type of the instance
data_type = theDouble.get_data_type()
# When data type is not assigned to the instance, None is returned
if data_type is None:
    print('Instance theDouble is untyped')
else:
    print('Data type of the theDouble instance is ', data_type.get_name())
# Access metadata of the instance
metadata = theDouble.get_metadata()
# And metadata properties ...
if metadata.has('max'):
    print('theDouble has max property: ', metadata['max'].get_value())
```

6.6.2 Iteration over document instances

The object returned by the document instance interface offers a recursive iterator over all document instances by default.

C++

```
const ::elt::configng::CiiConfigInstanceNamespace &instances = document.
↳Instances();
for (const auto &[name, node]: instances) {
    std::cout << "Name: " << name << ", Origin: " << node.GetOrigin() <<
↳std::cerr;
    // Obtain associated metadata
    ::elt::configng::CiiConfigMetadata &metadata = node.GetMetadata();
    // Obtain data type of the instance node from metadata
    ::elt::configng::CiiConfigDataType &data_type = metadata.GetDataType();
    std::cout << " ... data type: " << data_type.GetName()
        << " originating: " << data_type.GetOrigin() << std::endl;
```

(continues on next page)



(continued from previous page)

```
std::cout << " ... basic data type: "  
          << elt::configng::util::ToString(data_type.GetBasicDataType()) << "  
↪std::endl;  
    if (data_type.IsVector() || data_type.IsMatrix()) {  
        std::cout << ".... element data type: " << data_type.GetElementDataType().  
↪GetName();  
    }  
}
```

Python

```
document = elt.configng.CiiConfigLoad('filename.yaml')  
for name, item in document.instances:  
    # Data type can be None, indicating that config-ng data type is not assigned  
    # to the instance  
    data_type = item.get_data_type()  
    print(name, ' Data type: ', data_type)  
    print(name, ' Originating: ', item.get_origin())  
  
    if data_type is not None:  
        print('Basic data type: ', data_type.get_basic_data_type())  
        if data_type.is_matrix() or item.data_type.is_sequence():  
            print('..... element data type: ', data_type.get_element_data_type())
```



6.7 Document Validation

Metadata that is produced for each instance node is initially not validated, therefore it might not reflect the actual state of document.

To check the document validity against the schema (defined by elements with CII custom tags) and validity of the metadata, one must invoke check operation on the document. This validates the schema and all values of the document instances.

The Check operation returns an object describing the issues detected during document validation. In case this object does not contain any issues with error severity, the document is valid. This implies that the document is valid when only issues with warning severity were detected.

It is recommended to perform document validation after document loading or any operation that updates document like merge with another document or update with a programmatically built map of instances.

6.7.1 C++

The document validation operation is invoked by calling `::elt::configng::CiiConfigDocument::Check` on the document instance.

```
// Load document
::elt::configng::CiiConfigDocument document =
↳::elt::configng::CiiConfigClient::Load("cii.local://file.yaml");
// Perform document validation
::elt::configng::CiiConfigDocumentIssues issues = document.Check();
// The following test returns true in case there was at least one issue with
↳error severity
// was detected
if (issues) {
    std::cerr << "Document is not valid! The following issues were detected: " <<
↳std::endl;
    // One can iterate over issues. Each issue is instance
↳of::elt::configng::CiiConfigDocumentIssue
    // class.
    for (auto &issue: issues) {
        // Print out the issue
        std::cerr << " " << issue << std::endl;
        // Or examine it
        if (issue.IsError()) {
            std::cerr << "This is issue with code " << static_cast<int>(issue.
↳GetCode())
                << " and message: " << issue.GetMessage() << std::endl;
        }
    }
} else {
    // Document is valid at this point, but issues with warning severity can still
↳be present
```

(continues on next page)



(continued from previous page)

```
std::cout << "Document is valid." << std::endl;
if (issues.HasWarnings()) {
    std::cout << "The following warnings were issued during validation: " <<
↪std::endl;
    for (auto &issue: issues) {
        std::cout << " " << issue << std::endl;
    }
}
```

6.7.2 Python

Calling method `check()` on an `elt.configng.CiiConfigDocument` instance invokes the document validation operation.

```
# Load document
document = elt.configng.CiiConfigClient.load('cii.local://file.yaml')
# Perform document validation
issues = document.check()
if issues.has_errors():
    print('Document is not valid! The following issues were detected:')
    for issue in issues:
        # Print the issue
        print(' ', issue)
        # Or examine it
        if issue.is_error():
            print('This is issue with code ', issue.get_code(), ' with message: ',
                  issue.get_message())
else:
    # Document is valid, but issues with warning severity can still be present
    print('Document is valid.')
    if issues.has_warnings():
        print('The following warnings were issued during validation:')
        for issue in issues:
            print(' ', issue)
```




6.8 Saving documents to YAML

A document can be emitted in YAML format to local files or streams. It is recommended that the document is validated before emitting it in YAML format. The Save operation supports additional options that can prevent overwriting existing file or inlining includes (meaning not generating !cfg.include directives in the emitted YAML document).

6.8.1 C++

```
// Load document
::elt::configng::CiiConfigDocument document =
↳::elt::configng::CiiConfigClient::Load('cii.local://file.yaml');
// Validate document
::elt::configng::CiiConfigDocumentIssues issues = document.Check();
// Only save valid document
if (!issues) {
    // Save to file. Existing files are overwritten by default. Includes are not
↳inlined
    // Note that argument here is acutally file name not URI.
    document.Save('file1.yaml');
    // Save to file but prevent overwriting existing file.
    // In case file already exists, ::elt::configng::CiiConfigExistError is
↳thrown
    try {
        document.Save('file1.yaml,
            {::elt::configng::CiiConfigDocument::SaveOptions::NO_OVERWRITE});
    } catch (const ::elt::configng::CiiConfigExistsError &e) {
        std::cerr << "Could not overwrite the file (" << e.what() << ")" <<
↳std::endl;
    }
    // Save to stream, inline includes
    std::ostringstream stream;
    document.Save(stream,
        {::elt::configng::CiiConfigDocument::SaveOptions::INLINE_INCLUDES});
    // Save to stream, generate !cfg.include directives for included files
    std::ostringstream another_stream;
    document.Save(another_stream);
}
```



6.8.2 Python

```
# Load document
document = elt.configng.CiiConfigClient.load('cii.local://file.yaml');
# Validate document
issues = document.check()
# Only save valid document
if not issues.has_errors():
    # Save to file. Existing files are overwritten by default. Includes are not
    ↪inlined.
    # Note that argument here is actually file name not URI.
    document.save('file1.yaml')
    # Save to file but prevent overwriting existing file.
    # In case file already exists, elt.configng.CiiConfigExistsError is raised.
    try:
        document.save('file1.yaml',
                      (elt.configng.CiiConfigDocument.SaveOptions.NO_OVERWRITE,))
    except elt.configng.CiiConfigEistsError as e:
        print('Could not overwrite the file %s' % e)
# Save to stream, inline includes
stream = io.StringIO()
document.Save(stream,
              (elt.configng.CiiConfigDocument.SaveOptions.INLINE_INCLUDES,))
# Save to stream, generate !cfg.include directives for included files
another_stream = io.StringIO()
document.Save(another_stream)
```



6.9 Merging documents

One can request a merge of two documents via the merge operation. This operation modifies the target document in place.

The Merge operation does not modify target document in case it detects any issue that would prevent the merge from fully succeeding, unless explicitly permitted by the user.

It is recommended to validate both documents before attempting to execute merge operation. For clarity, validations are not performed in the following examples.

6.9.1 C++

```
// Load target document.
::elt::configng::CiiConfigDocument document =
↳::elt::configng::CiiConfigClient::Load("cii.config://local/document.yaml");
// Load document to merge
::elt::configng::CiiConfigDocument document_to_merge =
↳::elt::configng::CiiConfigClient::Load(
    "cii.config://local/document_to_merge.yaml");
// Try merge operation. List of issues that prevented merge is returned. In case
↳returned
// list is empty, merge succeeded.
std::vector<::elt::configng::CiiConfigDocument::MergeIssue> merge_issues =
↳document.Merge(document_to_merge);
if (merge_issues.empty()) {
    std::cout << "Merge succeeded" << std::endl;
} else {
    std::cerr << "Merge failed, the following issues were detected:" << std::endl;
    for (const auto &issue: merge_issues) {
        std::cerr << " " << issue << std::endl;
    }
    // Target document was not updated.
    // Perform partial merge, only update those instances that can be safely
↳merged.
    // Partial merge is requested by supplying additional boolean argument with
↳true value.
    // Merge issues for instances that could not be merged are still returned
    merge_issues = document.Merge(document_to_merge, true)
}
```



6.9.2 Python

```
# Load target document
document = elt.configng.CiiConfigClient.load('cii.config://local/document.yaml')
# Load document to merge
document_to_merge = elt.configng.CiiConfigClient.load(
    'cii.config://local/document_to_merge.yaml')
# Try merge operation. List of issues that prevented merge is returned. In case
↳returned
# list is empty, merge succeeded.
merge_issues = document.merge(document_to_merge)
if not merge_issues:
    print('Merge succeeded')
else:
    print('Merge failed, the following issues were detected:')
    for issue in merge_issues:
        print(' ', issue)
    # Target document was not updated.
    # Perform partial merge, only update thos unstances that can be safely
↳merged.
    # Partial merge is requested by supplying additional bool argument with True
↳value.
    # Merge issues for instances that could not be merged are still returned
    merge_issues = document.merge(document_to_merge, True)
```



6.10 Exceptions

Config-ng API calls can produce a number of exceptions. For details please refer to the API docs.

All exceptions generated by the config-ng API have common basic class (`::elt::configng::CiiConfigError` in C++ and `elt.configng.CiiConfigError` in Python).

List of config-ng exceptions that are defined in Python and C++:

CiiConfigError Base clas of config-ng exceptions.

CiiConfigBuildError Generated during process of building configuration document when major inconsistency is detected.

CiiConfigTypeError Indicates that argument with wrong data type was used.

CiiConfigValueError Indicates that argument with wrong value was used.

CiiConfigNotFoundError Item was not found

CiiConfigExistsError Item or file already exists.

CiiConfigNotImplementedError Feature was not implemented yet.

CiiConfigIterationError Inconsistency detected during iteration.

CiiConfigIllegalUriError Illegal URI was used.

An additional exception is defined in C++ implementation:

CiiConfigDocumentNotFromFileError Attempted to save document that was not loaded from file without providing file name.



6.11 YAML tags recognized by config-ng

Config-ng uses several specific tags in different context. All tags recognized by config-ng have prefix **cfg**.

6.11.1 `cfg.include`

Request inclusion of another YAML file.

Syntax: `!cfg.include <URI>: [NOERROR]`

Where `<URI>` is file name or URI of the file to include in the document. File is looked up according to the current search path of the client.

Optional `NOERROR` flag indicates that in case the file is not found, the operation should continue. Without this flag, an exception is thrown when search for a requested file failed.

Example:

```
# Include basic definitions
!cfg.include cii.config://local/basic_definitions.yaml:
# Include additional extensions, but do not fail if the file was not found
!cfg.include cii.config://local/extensions: NOERROR
```

6.11.2 `cfg.type`

Reference data type.

Syntax: `!cfg.type:<TYPENAME>`

`<TYPENAME>` is the name of a built-in data type or user-defined data type.

Example:

```
# Untyped instance
a: 10
# Typed instance, in this case b is of built in type int32
b: !cfg.type:int32 10
```

6.11.3 `cfg.typedef`

Define type alias or custom data type.

Syntax: `!cfg.typedef <TYPENAME>[(<BASE>)]`

`<TYPENAME>` is the name of the new data type. `<BASE>` is optional and must be the name of an existing built-in or user-defined data type. When specified, config-ng derives a new data type from the provided base data type.

**Example:**

```
# Type alias
!cfg.typedef myint(int32):

# Type alias with additional metadata
!cfg.typedef extint(int32):
    min: 10
    max: 100
    default: 25

# User defined data type (like struct)
!cfg.typedef Point:
    x: !cfg.type:double
    y: !cfg.type:double

# Derived, inherits the members of the basic data type,
# in this case x and y. Adds new member z.
!cfg.typedef Point3D(Point):
    z: !cfg.type:double
```

6.11.4 cfg.optional, cfg.required

Modifiers that can be used within user-defined data type member definition.

`cfg.optional` indicates that the value of this member needs not be present when initializing an instance of that data type.

`cfg.required` indicates that the value of this member must always be present when initializing an instance of this data type.

When neither of these modifiers are used, the value of the member is initialized from the default value, if not explicitly specified.

Example:

```
!cfg.typedef Point:
    # x must be always specified
    !cfg.required x: !cfg.type:double
    # y is initialized from default not specified
    y: !cfg.type:double
    # z is optional
    !cfg.optional z: !cfg.type:double

valid_point_1: !cfg.type:Point { x: 10 }
valid_point_2: !cfg.type:Point { x: 10, y: 55 }
valid_point_3: !cfg.type:Point { x: 10, z: 33 }
# The following instance is invalid, x is missing
invalid_point: !cfg.type:Point { y: 3, z: 3 }
```



6.12 List of built-in data types

List of built-in data types recognized by config-ng:

- int8
- uint8
- int16
- uint16
- int32
- uint32
- int64
- uint64
- single
- double
- boolean
- string
- binary
- vector_uint8
- vector_int8
- vector_uint16
- vector_int16
- vector_uint32
- vector_int32
- vector_uint64
- vector_int64
- vector_single
- vector_double
- vector_boolean
- vector_string
- matrix2d_uint8
- matrix2d_int8
- matrix2d_uint16



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 145 of 589

- matrix2d_int16
- matrix2d_uint32
- matrix2d_int32
- matrix2d_uint64
- matrix2d_int64
- matrix2d_single
- matrix2d_double

ONLINE DATABASE

Revision:	1.4
Status:	Released
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	ELT CII
Folder:	trunk/deliverables/phase8
Document ID:	CSL-DOC-19-176283
File:	MAN-ELT_OLDB.docx
Owner:	Jan Pribošek
Last modification:	April 04, 2022
Created:	July 5, 2019

Prepared by	Reviewed by	Approved by
Jan Pribošek (CSL SWE) Marko Novak (CSL) Jernej Strniša (CSL)	Miha Vitorovič (CSL) Borut Terpinc (CSL SWE)	Gregor Čuk (CSL)

Document History



ELT CII - User Manual

Re- vi- sion	Date	Changed/re- viewed	Section(s)	Modification
0.1	05.07.2019	prilosek, manovak	5	Document creation. GUI Section.
0.2	19.07.2019	prilosek, jstrnisa	1-4 Appendix A Appendix B	Added client API sections.
0.3	05.08.2019	torovic bter- pinc	All	Document review and fixes.
1.0	24.09.2019	prilosek	All	Final updates, release.
1.1	20.1.2020	jstrnisa	All	Major update of all sections after ESO topical review
1.2	13.7.2020	bter- pinc	3.1 4.1 5.6	Statement about pub/sub queue adjustment added. Note about the synchronous methods added. Additional External Redis serves explanation added.
1.3	16.7.2020	bter- pinc	7	CLI tools update
1.4	04.4.2022	schilli	7	Updated cpp py examples
1.5	09.5.2022	umar	All	HDFS to Remote FS, disable formula check when writing

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is manual for the Online Database system used by applications in the ELT Core Integration Infrastructure Software.

Audience

This document is aimed at those Cosylab and ESO employees involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 148 of 589

API	Application Programming Interface
CII	Core Integration Infrastructure
CLI	Command Line Interface
DP	Data Point
GUI	Graphical User Interface
ES	ElasticSearch
JSON	Javascript object notation
OLDB	Online Database
SVN	Subversion
YAML	YAML Ain't Markup Language

References

1. ESO, Core Integration Infrastructure Requirements Specification, ESO-192922 Version 6
2. Cosylab, ELT CII Online Database Transfer Document, CSL-DOC-19-176283, Version 1.4
3. Cosylab, CII Configuration User's Manual, CSL-DOC-19-176283, Version 1.4
4. Redis main page: <https://redis.io/>
5. Remote FS documentation: <https://docs.min.io/>
6. Octave documentation version 5.1.0: <https://octave.org/doc/v5.1.0/>
7. Redis INFO command <https://redis.io/commands/info>



7.1 Overview

This document is a user manual for usage of CII OLDB system. It explains how to use the OLDB Client API through Client API library and GUI Application to interact with the CII Online Database. All examples in this manual will be also presented in SVN code (project oldb-examples).

7.2 Introduction

The Core Integration Infrastructure (CII) Online Database (OLDB) provides distributed data publishing and access to actual or live data for user interface and control applications that do not have low-latency or real-time performance requirements. The term “online” refers to the fact that the database provides current and live values for the data points of the control system.

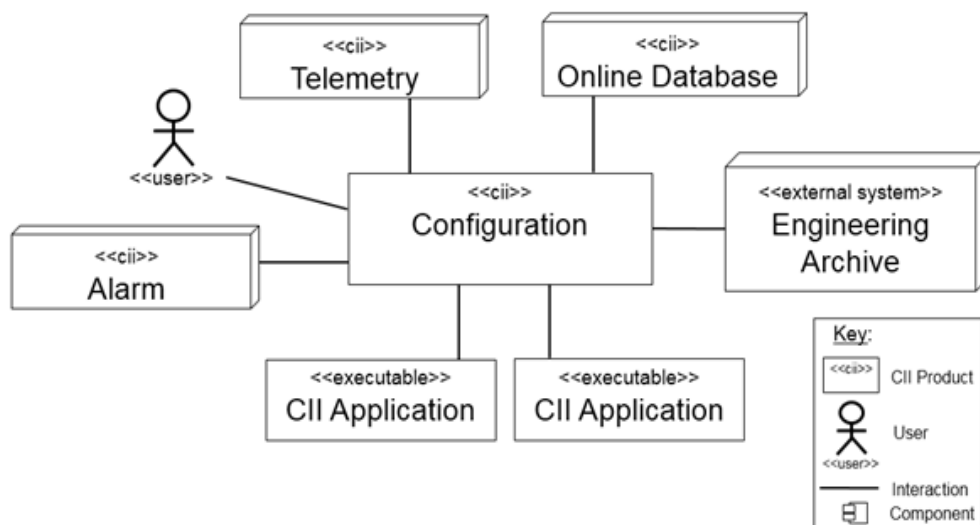


Figure 2-1: OLDB API interactions

A **data point** represents a single value in the OLDB and consists of a value, timestamp, quality, and metadata. Data points are identified by a unique URI (section 6.2.1).

Data point **value** is the actual data the data point is holding. It can be any of the predefined types (Described in Appendix A). The type of the data point can be a primitive type (e.g. integer, double, string) or a more complex type (vector or 2D matrix).

The **timestamp** is the timestamp of the last data point value modification.

The **quality** represents the state of the data point. It can be OK, SUSPECT and BAD. Only OK data points can be read and written to (except when correcting the quality to OK). The meaning of each quality state is as follows:

OK: the value of the Property is valid

BAD: the value is not OK and should not be used

SUSPECT: there is a value, but it is not clear if it is reliable. For example, the readout of a sensor is inconsistent with other readings and the sensor might be faulty.

Metadata is the meta information about the data point and consists of information like DP data type, formula, quality expression and other fields specific to DP data type (for a detailed description of metadata content see section 6.2.3). Metadata is maintained by CII Configuration service [3], but can also be manipulated through the OLDB Client.

Data points can define a calculation **formula** (included in metadata) that determines the value of the



data point. The formula can include URIs of other data points and so the value of the data point can depend on values of other data points. A data point with formula is called **calculated data point** and cannot be written to by users, since its value is determined by the formula.

OLDB **Client API** provides a means to manipulate data points (create, read, write and delete) through the OLDB client. Users can access and manipulate OLDB through the Client API library (srv-oldb) or OLDB GUI and CLI applications. The srv-oldb library is written in Java, C++ and Python. The inline code examples and class definitions in this document are written in Java. For example in C++ and Python see Chapter 4 where complete examples are listed in all three languages.

OLDB uses Redis key-value store [4] and Remote file storage [5] for storing the data. Redis is used for small data point values while large DP that exceed the limit defined in the OLDB configuration (section 3.3.1) are saved in Remote FS. The Redis storage is further divided into two parts:

default storage: This is the default location for data points. All small sized data points are stored here.

in-memory-only storage: This is a special storage where only predetermined data points (defined in data point configuration) are stored. It is intended for larger values that need quicker access than Remote FS can provide.

See [2] for a detailed description of data point configuration and how to manipulate it.

Both default and in-memory-only storages are horizontally scalable and can consist of any number of Redis servers either single instance or clusters.

The lifetime of the data points is not bounded by the lifetime of the OLDB client instance because their values and configurations are saved in the OLDB storage and Configuration service respectively. If the Redis storage is configured to be persistent, the data points also persist after Redis server shutdown.



7.3 Prerequisites

This section describes the prerequisites for using the OLDB library and applications. Note that the preparation of the environment (i.e. installation of required modules, configuration, running of services) is not in the scope of this document. Consult your administrator and the OLDB TRD document [2] for information on how to establish the environment for usage of the OLDB library and applications.

7.3.1 WAF Modules

For using the OLDB library and CLI application the following modules must be built and installed with WAF in the specified order:

1. elt-mal
2. client-api
3. srv-config
4. oldb-client
5. srv-oldb

For using the OLDB GUI application in addition to the above two more modules must be built and installed:

6. elt-qt-widgets
7. oldb-gui

Before building anything, Proper .pc files must be put in the “/eelt/System/pkgconfig” folder. These files are used by the WAF wscript to properly resolve the dependencies. Please refer to OLDB TRD for details [2].

7.3.2 Services

Redis

OLDB uses Redis [4] key-value store as its main storage of data point values. At a minimum the OLDB needs at least one default storage server, at least one in-memory-only (External Redis Server - 5.6) storage and exactly one pub/sub server. One Redis instance can serve as all three servers so minimally one Redis instance must be running for OLDB to work.

The number of default and in-memory-only storage servers is unbounded. Only one Redis server can be designated to be a pub/sub server and can be either a single instance or cluster.

See the OLDB TRD document [2] on how to setup the Redis server



Remote FS

Remote FS storage is used for large-sized data points (the ones that are not defined to be stored in in-memory-only storage). OLDB needs exactly one Remote FS storage server (currently S3 compatible minio) running to work properly. The installation and configuration of the Remote FS storage are out of the scope of this document. Guides for doing this can be found in [2].

Calculation Service

For the OLDB framework to work as intended, the **calculation service** (OLDB service) must be running. The calculation service is responsible for calculating the value and quality of calculated data points. The calculation service consists of any number of **calculation node** processes and a **scheduler process**. Each calculation node is responsible for the calculation of a subset of all existing data points and the scheduler process is responsible for distributing the calculated data points among the calculation nodes.

All the calculation nodes must be running before the scheduler process. The scheduler must be configured with addresses of existing calculation nodes so it can connect to them.

To add a calculation node, the scheduler configuration must be updated with the new node's address and the scheduler process must be restarted. The old nodes will retain their calculated data points and continue to recalculate their values. Note that during scheduler downtime newly created calculated data points will not be distributed and calculated. **They will be distributed however when the scheduler process restarts.**

If a calculation node shuts down the subset of calculated data points assigned to it will not be re-assigned to another calculation node and will not be calculated. For this to happen the scheduler process must be restarted.

The proper way to shut down the calculation service is to first stop the scheduler process and then the calculation nodes.

Configuration Service

OLDB uses CII Configuration Service [3] to store its configuration. The Configuration service must be accessible for the OLDB Client and Calculation Service to successfully initialize.

7.3.3 Configuration

Both the OLDB Client and the calculation service must be configured to work properly. The configuration storage is handled by the CII Configuration Service [3]. Both OLDB and Calculation Service configurations are stored in a versioned way.



OLDB Client Configuration

The OLDB client configuration is stored in the `CiiOldbConfigClass` object. The configuration is fetched from the CII Configuration Service by the OLDB client at start-up.

The configuration class contains the following fields:

1. `redisServers`: a `MdStringArray` class containing a list of strings that contains parameters for **default** storage servers. Each string corresponds to one server and is formatted as `<sequential-number>:<hostname>:<port>:<server-type>` where `server-type` is either `SINGLE` or `CLUSTER` and describes the type of the server (either single instance or cluster). The `<sequential-number>` should be the sequential number of servers in the list starting at 0.
8. `redisServersExternal`: a `MdStringArray` class containing a list of strings that contains parameters for **in-memory-only** storage servers. Each string corresponds to one server and is formatted as `<alias>:<hostname>:<port>:<server-type>` where `server-type` is either `SINGLE` or `CLUSTER` and describes the type of the server (either single instance or cluster).
9. `pubSubServerInfo`: a `MdString` class containing a string that contains parameters for the OLDB pub/sub server. The string is formatted in the same manner as strings in the `redisServers` field, except that the `<sequential-number>` has no significance since there is always one server.
10. `pubSubHandlerProcessingInterval`: a `MdInt32` containing an integer that represents the interval in milliseconds at which a pub/sub message handler should handle messages. To prevent the incorrect ordering of command execution, the incoming pub/sub messages are ordered by timestamp in unbounded queues and processed every `pubSubHandlerProcessingInterval` milliseconds. The performance of the pub-sub system will degrade with larger `pubSubHandlerProcessingInterval` value, but the correct order of message execution will be more guaranteed. The recommended value for this field is 10. The size of the pub/sub event queue is unbounded, and thus it is not adjustable. **Note: This reason for this setting is the proper timestamp ordering of the incoming messages.**
11. `valueSizeLimitRedisExt`: a `MdInt32` containing an integer that represents the value size limit in bytes for in-memory-only servers. A data point value written to these servers is sliced and distributed among several key-value pairs so that each chunk does not exceed the `valueSizeLimitRedisExt`. If this value is too small write operations on large data points will be slower. It should not be greater than 512 Mb (512 * 1024 * 1024 bytes) since this is the limit of one key-value pair on Redis. The recommended value for this field is 128 Mb (128 * 1024 * 1024 bytes). The value must be set in bytes (e.g., 134217728).
12. `valueSizeLimitHDFS`: a `MdInt32` containing an integer that represents the value size limit in bytes for default servers. If a data point value written to the default server exceeds this limit, it is redirected to the Remote FS storage. In this case, the data point information holds only the reference to its location. This value should not be greater than 512 Mb (512 * 1024 * 1024 bytes) since this is the limit of one key-value pair on Redis. The recommended value for this field is 100Kb (100 * 1024 bytes) The value must be set in bytes (e.g., 102400).
13. `extRedisKeyExpireTime`: a `MdInt32` containing an integer that represents the expire time in milliseconds of the keys on the in-memory-only servers. Whenever a data point is written to



in-memory-only storage, the old value keys are not rewritten (some client could at that time be reading them) but only set to expire (are deleted) after the `extRedisKeyExpireTime`. This value should always be bigger than a time of a write operation on any data point, else a read operation could return an error or malformed value. The recommended value is 5000.

14. `largeStorageLocation`: a `MdString` with the address of Remote FS storage.

15. `largeStorageEndpoint`: a `MdString` with the endpoint of Remote FS storage.

Note that the OLDB client configuration cannot be changed dynamically. If the configuration has been changed, the OLDB client must be restarted for the changes to take effect.

The OLDB configuration is stored in CII Config Service under the URI:

```
cii.config://remote/oldb/configurations/oldbclientconfig
```

For testing purposes, a default remote configuration can be deployed on the Configuration Service using the `oldb-initEs` script. This script will initialize testing OLDB configuration data. Every time the script is run, it will create a new version. The OLDB client always reads the last version of the configuration at start up.

```
oldb-initEs
```

Calculation Service Configuration

The scheduler process must be configured before it can be run. The configuration is done through the CII Configuration service [3] using the `CiiOldbSchedulerConfig`. The only field this configuration needs is a list of string URIs of calculation nodes that the scheduler should handle.

The calculation service configuration is stored in CII Config Service under the URI:

```
cii.config://remote/oldb/calculationnodes
```



7.4 OLDB Library Usage

This section explains through examples of how to use the OLDB API library in an application. The OLDB API library provides the user the functionality to create, read, write, query, subscribe to and delete data points in the CII OLDB. The API is distributed among two classes. The CiiOldb and CiiOldbDataPoint.

7.4.1 Includes/Imports

For basic usage of the OLDB library, the user needs the CiiOldb client class, the CiiOldbDataPoint and the CiiOldbDpValue class. The first two classes contain the API for manipulating data points and the last one is a container for the data point value triplet (i.e. value, quality, timestamp). Next the language specific URI class is needed. If data points are created with provided metadata, respective metadata classes need to be imported. See section 6.2.3 for the list of metadata classes.

Java

```
import java.net.URI;
import elt.olddb.client.CiiOldb;
import elt.olddb.client.CiiOldbDataPoint;
import elt.olddb.client.CiiOldbDpValue;
//Metadata classes
import elt.config.classes.meta.MdOldbNumber;
```

C++

```
#include <ciiOldbDpValue.hpp>
#include <ciiOldbFactory.hpp>
//Metadata classes
#include <meta/MdOldbNumber.hpp>
```

Python

```
import elt.olddb
import elt.config.Uri
```

For subscribing to data point changes, the CiiOldbSubscription interface or a class implementing this interface must be imported.



Java

```
import elt.olddb.client.CiiOldbDpSubscription;
```

C++

```
#include <ciiOldbSubscription.hpp>
```

Python

Included in elt.olddb module.

Since OLDB API also throws exceptions, these need to be imported to. For Java and C++ these exceptions can be found in elt.olddb.exceptions package and <ciiOldbExceptions.hpp> header. For Python these exceptions are included in elt.olddb module.

Linked library includes can be seen in code examples of 4.1.2, 4.1.3 and 4.1.4

Basic Example

The examples in this section can be found in the oldb-examples project (BasicExample.java, sample-app.cpp, cii-oldb-examples-sample-app-py.py).

In basic example (Listing 4-1, Listing 4-2, Listing 4-3) four constant (without formulas) data points of type DOUBLE, STRING, MATRIX2D_DOUBLE and VECTOR_INT32 are created. Multiple different creation methods are used to create these data points (see appendix C.1 on creation methods). These data points are then read to confirm their initial values. New values are then written to the data points.

Then the data points are queried using a glob expression and again with additional filter arguments (query for the specific type and value range).

Then the children of specific directory parent in the URI hierarchy are queried.

At last, the data points are deleted and OLDB client is closed.

Note: All OLDB client methods are synchronous towards the Redis server.

Listing 4-1: Java Basic Example

```
package sampleapp;  
  
import java.net.URI;  
import java.util.ArrayList;  
import java.util.Arrays;  
import java.util.List;
```

(continues on next page)



(continued from previous page)

```
import java.util.Map;

import elt.common.AttributeType;
import elt.config.classes.meta.MdOldbNumber;
import elt.config.client.CiiConfigClient;
import elt.config.exceptions.CiiConfigDeleteException;
import elt.config.exceptions.CiiConfigInitializationError;
import elt.config.exceptions.CiiConfigSaveException;
import elt.config.exceptions.CiiConfigWriteDisabledException;
import elt.error.CiiInvalidTypeException;
import elt.error.CiiInvalidURIException;
import elt.oldb.client.CiiOldb;
import elt.oldb.client.CiiOldbDataPoint;
import elt.oldb.client.CiiOldbDpSubscription;
import elt.oldb.client.CiiOldbDpValue;
import elt.oldb.client.impl.CiiOldbImp;
import elt.oldb.client.util.CiiOldbUriUtils;
import elt.oldb.exceptions.CiiOldbConfigError;
import elt.oldb.exceptions.CiiOldbCyclicDependencyException;
import elt.oldb.exceptions.CiiOldbDeletedException;
import elt.oldb.exceptions.CiiOldbDpConstraintsException;
import elt.oldb.exceptions.CiiOldbDpExistsException;
import elt.oldb.exceptions.CiiOldbDpQualityException;
import elt.oldb.exceptions.CiiOldbDpUndefinedException;
import elt.oldb.exceptions.CiiOldbExpressionException;
import elt.oldb.exceptions.CiiOldbIllegalOperationException;
import elt.oldb.exceptions.CiiOldbInitializationError;
import elt.oldb.exceptions.CiiOldbStorageError;
import elt.oldb.exceptions.CiiOldbURIException;
import elt.oldb.exceptions.CiiOldbUndefinedMetadataException;
import elt.oldb.exceptions.CiiOldbWriteDisabledException;

/**
 * @copyright (c) Copyright ESO 2019 All Rights Reserved ESO (eso.org) is an
 *             Intergovernmental Organisation, and therefore special legal
 *             conditions apply.
 * @ingroup srv-oldb-java-sample-app
 */
public class SampleApplication {

    public static void main(String[] args) {

        CiiOldb oldbClient = null;
        CiiConfigClient configClient = null;
        try {
            // Initialize OLDB and Config clients
            oldbClient = CiiOldb.getInstance();
            configClient = CiiConfigClient.getInstance();
        }
    }
}
```

(continues on next page)



(continued from previous page)

```
// Define URIs
URI doubleDpURI =
    URI.create("cii.oldb:/sampleroot/child/device/doubledp");
URI matrixDpURI = URI.create("cii.oldb:/sampleroot/child/matrixdp");
URI stringDpURI = URI.create("cii.oldb:/sampleroot/child/stringdp");
URI vectorDpURI = URI.create("cii.oldb:/sampleroot/child/vectordp");

oldbClient.deleteDataPoint(doubleDpURI);
oldbClient.deleteDataPoint(matrixDpURI);
oldbClient.deleteDataPoint(stringDpURI);
oldbClient.deleteDataPoint(vectorDpURI);

// CREATE DATA POINTS
// Create DOUBLE data point with provided metadata instance name.
// First we create the metadata and save it to Config.
MdOldbNumber<Double> doubleDpMeta =
    new MdOldbNumber<Double>("customDoubleDpMeta",
        "metadata of a constant DP", 0.0, null, null, 0.0, 10.0, null);
URI doubleDpMetaURI = URI
    .create(CiiOldbUriUtils.CONFIG_METADATA_URI +
        doubleDpMeta.getMetadataInstanceName());
configClient.saveMetadata(doubleDpMetaURI, doubleDpMeta);

CiiOldbDataPoint<Double> doubleDp = (CiiOldbDataPoint<Double>)
    oldbClient.createDataPoint(doubleDpURI,
        doubleDpMeta.getMetadataInstanceName(),
        doubleDpMeta.getClass().getSimpleName());

// Create MATRIX2D_DOUBLE data point with default matrix metadata
// OldbDoubleMatrixStd and provided initial value.
// The default metadata must exist on Config Service.
// Run oldb-initEs script to generate
// default metadata instances.
CiiOldbDataPoint<List<Double>> matrixDp =
    oldbClient.createDataPointByValue(matrixDpURI,
        new ArrayList<>(50), Double.class, true);

// Create STRING data point with default string metadata
// OldbStringStd and provided initial value.
CiiOldbDataPoint<String> stringDp =
    oldbClient.createDataPointByValue(stringDpURI, "ABCDEF");

// Create VECTOR_INT32 data point with provided metadata instance name.
// Metadata with this instance name must exist in Config Service.
CiiOldbDataPoint<List<Integer>> vectorDp =
    (CiiOldbDataPoint<List<Integer>>) oldbClient
        .createDataPoint(vectorDpURI, "OldbInt32ArrayStd", "MdOldbArray");

// Read data point values.
```

(continues on next page)



(continued from previous page)

```
CiiOldbDpValue<Double> constDpValue = doubleDp.readValue();
CiiOldbDpValue<List<Double>> matrixDpValue = matrixDp.readValue();
CiiOldbDpValue<String> stringDpValue = stringDp.readValue();
CiiOldbDpValue<List<Integer>> vectorDpValue = vectorDp.readValue();

// Write to data points
List<Double> newMatrixValue = new ArrayList<>();
for (int i = 0; i < 2500; i++) {
    newMatrixValue.add(new Double(i));
}

doubleDp.writeValue(2.0);
matrixDp.writeValue(newMatrixValue);
stringDp.writeValue("12345");
vectorDp.writeValue(new ArrayList<>(Arrays.asList(2, 4, 6)));

// Read the data point values.
constDpValue = doubleDp.readValue();
matrixDpValue = matrixDp.readValue();
stringDpValue = stringDp.readValue();
vectorDpValue = vectorDp.readValue();

// get Multiple data points satisfying an URI glob expression.
List<CiiOldbDataPoint<?>> searchResult = oldbClient
    .getDataPoints(new ArrayList<URI>(Arrays.asList(URI.create("cii.oldb:/
↪sampleroot/**"))));

// get Multiple DOUBLE data points with value between 1.0 and 3.0
List<CiiOldbDataPoint<Double>> filteredSearchResult =
    oldbClient.getDataPoints(
        new ArrayList<URI>(
            Arrays.asList(URI.create("cii.oldb:/sampleroot/**")),
            AttributeType.DOUBLE, 1.0, 3.0);

// get the children of directory URI cii.oldb:/sampleroot. Returns a map
// specifying whether
// a child is a data point or just a directory.
Map<URI, Boolean> children =
    oldbClient.getChildren(URI.create("cii.oldb:/sampleroot"));

// Delete data points
oldbClient.deleteDataPoint(doubleDpURI);
oldbClient.deleteDataPoint(matrixDpURI);
oldbClient.deleteDataPoint(stringDpURI);
oldbClient.deleteDataPoint(vectorDpURI);
configClient.deleteMetadata(
    URI.create(CiiOldbUriUtils.CONFIG_METADATA_URI +
        doubleDpMeta.getMetadataInstanceName()),
    doubleDpMeta.getMetadataInstanceName(), -1);
```

(continues on next page)



(continued from previous page)

```
// Close the clients
configClient.close();
oldbClient.close();

} catch (CiiConfigInitializationError | CiiOldbDpConstraintsException |
↪CiiInvalidTypeException
    | CiiInvalidURIException | CiiOldbDpExistsException |
↪CiiOldbCyclicDependencyException
    | CiiOldbExpressionException | CiiOldbDpQualityException |
↪CiiOldbUndefinedMetadataException
    | CiiOldbStorageError | CiiOldbConfigError | CiiOldbDpUndefinedException
↪| CiiOldbDeletedException
    | CiiOldbWriteDisabledException | CiiConfigWriteDisabledException |
↪CiiOldbInitializationError
    | CiiOldbURIException | CiiConfigDeleteException |
↪CiiConfigSaveException e) {
    e.printStackTrace();
    if (configClient != null) {
        configClient.close();
    }
    if (oldbClient != null) {
        oldbClient.close();
    }
}
}
```

Listing 4-2: C++ Basic Example

```
/**
 * @copyright (c) Copyright ESO 2019 All Rights Reserved
 * ESO (eso.org) is an Intergovernmental Organisation, and therefore special
↪legal conditions apply.
 * @ingroup oldb-examples
 */
#include <mal/utility/Uri.hpp>

#include <ciiOldbFactory.hpp>
#include <ciiOldbDpValue.hpp>
#include <ciiOldbExceptions.hpp>

int main(int ac, char *av[]) {
    try {
        // Initialize OLDB

        std::shared_ptr<elt::oldb::CiiOldb> oldb_client =
↪::elt::oldb::CiiOldbFactory::GetInstance();
```

(continues on next page)



(continued from previous page)

```
::elt::oldb::CiiOldbGlobal::SetWriteEnabled(true);
std::string RND_PREFIX = ::elt::oldb::CiiOldbUtil::NewUUID();

// Define URIs

::elt::mal::Uri double_dp_uri{"cii.oldb/" +RND_PREFIX+ "/sampleroot/child/
↪device/doubledp"};
::elt::mal::Uri matrix_dp_uri{"cii.oldb/" +RND_PREFIX+ "/sampleroot/child/
↪matrixdp"};
::elt::mal::Uri string_dp_uri{"cii.oldb/" +RND_PREFIX+ "/sampleroot/child/
↪stringdp"};
::elt::mal::Uri vector_dp_uri{"cii.oldb/" +RND_PREFIX+ "/sampleroot/child/
↪vectordp"};

// CREATE DATA POINTS

// Create DOUBLE data point with provided metadata instance.
// Metadata must first be created and saved to Config Service.

const std::string double_dp_meta_instance_name = RND_PREFIX +
↪"customDoubleDpMeta";
std::shared_ptr<elt::config::CiiDataPointMetadataBase> double_dp_meta =
↪CiiDataPointMetadataFactory::getNewMetadataInstance<
::elt::config::classes::meta::MdOldb<double>>(double_dp_meta_instance_
↪name);

double_dp_meta->setComment("metadata of a constant DP");
double_dp_meta->set_default_value(0.0);
double_dp_meta->set_min_limit(0.0);
double_dp_meta->set_max_limit(10.0);

// meta data needs to be created before calling CreateDataPoint methods

::elt::oldb::CiiOldbUtil::SaveOrUpdateMetadata(*double_dp_meta);

std::shared_ptr<::elt::oldb::CiiOldbDataPoint<double>> double_dp = oldb_
↪client->CreateDataPoint<double>(
double_dp_uri, double_dp_meta_instance_name);

// Create MATRIX2D_DOUBLE data point with default matrix metadata
↪OldbDoubleMatrixStd
// and provided initial value. The default metadata must exist on Config
↪Service.
// Run oldb-initEs script to generate default metadata instances.

std::vector<double> matrix_double =
{1.0, 20.0, 30.0, 40.0, 50.0, 4.4, 45.3, 34.4, 445.3, 301.3,
2.0, 33.3, 34.3, 33.3, 33.3, 5.5, 32.2, 33.4, 222.3, 203.1};
```

(continues on next page)



(continued from previous page)

```
std::shared_ptr<::elt::config::CiiConfigClient::CiiDataPointMetadataBase>
↳ mdi_matrix_double =
    ::elt::config::CiiConfigClient::RetrieveMetadata (
        ::elt::oldb::CiiOldbGlobal::GetTargetConfigStorage(),
        ::elt::oldb::CiiOldbGlobal::MDI_MATRIX_DOUBLE,
        ::elt::oldb::CiiOldbGlobal::CONFIG_VERSION);

bool is_matrix = true;
std::shared_ptr<::elt::oldb::CiiOldbDataPoint<std::vector<double>>> matrix_dp_
↳ =
    oldb_client->CreateDataPointByValue(matrix_dp_uri, matrix_double, is_
↳ matrix);

// Create STRING data point with default string metadata OldbStringStd and
// provided initial value.

std::shared_ptr<::elt::config::CiiConfigClient::CiiDataPointMetadataBase>
↳ mdi_string =
    ::elt::config::CiiConfigClient::RetrieveMetadata (
        ::elt::oldb::CiiOldbGlobal::GetTargetConfigStorage(),
        ::elt::oldb::CiiOldbGlobal::MDI_STRING,
        ::elt::oldb::CiiOldbGlobal::CONFIG_VERSION);

std::shared_ptr<::elt::oldb::CiiOldbDataPoint<std::string>> string_dp = oldb_
↳ client->CreateDataPointByValue (
    string_dp_uri, std::string("ABCDEF"));

// Create VECTOR_INT32 data point with provided metadata instance name.
// Metadata with this instance name must exist in Config Service.

std::shared_ptr<::elt::config::CiiConfigClient::CiiDataPointMetadataBase>
↳ mdi_vector_int32 =
    ::elt::config::CiiConfigClient::RetrieveMetadata (
        ::elt::oldb::CiiOldbGlobal::GetTargetConfigStorage(),
        ::elt::oldb::CiiOldbGlobal::MDI_VECTOR_INT32,
        ::elt::oldb::CiiOldbGlobal::CONFIG_VERSION);

std::cout << "Creating datapoints\n";

std::shared_ptr<::elt::oldb::CiiOldbTypedDataBase> vector_dp_base =
    oldb_client->CreateDataPoint(vector_dp_uri,
↳ ::elt::oldb::CiiOldbGlobal::MDI_VECTOR_INT32);
    ::elt::oldb::CiiOldbDataPoint<std::vector<std::int32_t>> vector_dp =
        std::dynamic_pointer_cast<::elt::oldb::CiiOldbDataPoint<std::vector
↳ <std::int32_t>>>(
            vector_dp_base);
    if (!vector_dp) {
        throw ::elt::oldb::CiiOldbException("Casting to data point of vector int32_
↳ type failed");
```

(continues on next page)



(continued from previous page)

```
}

// Read data point values.

std::shared_ptr<::elt::olddb::CiiOldbDpValue<double>> const_dp_value = double_
↪dp->ReadValue();
std::shared_ptr<::elt::olddb::CiiOldbDpValue<std::vector<double>>> matrix_dp_
↪value = matrix_dp->ReadValue();
std::shared_ptr<::elt::olddb::CiiOldbDpValue<std::string>> string_dp_value =
↪string_dp->ReadValue();

// vectors do not have default value when created, needs a write before read
vector_dp->WriteValue(std::vector<std::int32_t>{2, 4, 6});
std::shared_ptr<::elt::olddb::CiiOldbDpValue<std::vector<std::int32_t>>>
↪vector_dp_value = vector_dp->ReadValue();

// Write to data points

std::cout << "Writing datapoints\n";

std::vector<double> new_matrix_value(2500);
for (std::int32_t i = 0; i < 2500; ++i) {
    new_matrix_value.push_back(static_cast<double>(i));
}

double_dp->WriteValue(2.0);
matrix_dp->WriteValue(new_matrix_value);
string_dp->WriteValue("12345");
vector_dp->WriteValue(std::vector<std::int32_t>{2, 4, 6});

// Read the data point values.

std::cout << "Reading datapoints\n";

const_dp_value = double_dp->ReadValue();
matrix_dp_value = matrix_dp->ReadValue();
string_dp_value = string_dp->ReadValue();
vector_dp_value = vector_dp->ReadValue();

// Get multiple data points satisfying an URI glob expression

std::cout << "Multi-retrieval of datapoints\n";

std::vector<::elt::mal::Uri> uris = {::elt::mal::Uri("cii.olddb:/sampleroot/**
↪")};
std::vector<std::shared_ptr<::elt::olddb::CiiOldbTypedDataBase>> search_
↪result = oldb_client->GetDataPoints(uris);

// Get multiple DOUBLE data points with value between 1.0 and 3.0
```

(continues on next page)



(continued from previous page)

```
std::cout << "Filtered Search\n";

std::vector<std::shared_ptr<CiiOldbDataPoint<double>>> filtered_search_
↪result = oldb_client->GetDataPoints<double>(
    uris, ::elt::common::CiiBasicDataType::DOUBLE, 1.0, 3.0);

// Get the children of directory URI cii.oidb:/sampleroot. Returns a map_
↪specifying whether
// a child is a data point or just a directory

std::map<std::string, bool> children = oldb_client->
↪GetChildren(::elt::mal::Uri("cii.oidb:/sampleroot"));

// Delete data points

std::cout << "Deleting datapoints\n";

oidb_client->DeleteDataPoint(double_dp_uri);
oidb_client->DeleteDataPoint(matrix_dp_uri);
oidb_client->DeleteDataPoint(string_dp_uri);
oidb_client->DeleteDataPoint(vector_dp_uri);

return 0;
} catch (const ::elt::oidb::CiiOldbException& ex) {
    std::cerr << "CiiOldbException occured while executing sample code. What: "
        << ex.what() << '\n';
}
return -1;
}
```

Listing 4-3: Python Basic Example

```
#!/usr/bin/env python
"""
@copyright (c) Copyright ESO 2019 All Rights Reserved
ESO (eso.org) is an Intergovernmental organisation,
and therefore special legal conditions apply.
@ingroup client-apis-python
@author Cosylab
"""
#pylint: disable=E1101,C0103,R0914,C0330,W0612
#This script shows basic usage of OLDB API.

import sys
import traceback
import uuid
import elt.config
```

(continues on next page)



(continued from previous page)

```
import elt.olddb

# Define URIs

RND_PREFIX = str(uuid.uuid4())

double_dp_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/doubledp" %
↳RND_PREFIX)
string_dp_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/stringdp" %
↳RND_PREFIX)
vector_dp_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/vectordp" %
↳RND_PREFIX)
matrix_dp_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/matrixdp" %
↳RND_PREFIX)

# Initialize OLDB client
olddb_client = elt.olddb.CiiOldbFactory.get_instance()

def _main():
    """main method implementation"""
    # enable writing
    elt.olddb.CiiOldbGlobal.set_write_enabled(True)

    # create DOUBLE data point with provided metadata instance
    # Metadata must first be created and saved to Config Service.

    double_dp_meta = \
        elt.olddb.typesupport.DOUBLE.\
        get_new_number_metadata_instance("%scustomDoubleDpMeta" % RND_PREFIX)
    double_dp_meta.set_comment("metadata of a constant DP")
    double_dp_meta.set_min_limit(0.0)
    double_dp_meta.set_max_limit(10.0)
    double_dp_meta.set_default_value(0.0)

    # Create & Save Metadata
    elt.olddb.CiiOldbUtil.save_or_update_metadata(double_dp_meta)

    # Create DOUBLE data point
    double_dp = oldb_client.create_data_point(double_dp_uri, double_dp_meta.get_
↳instance_name())

    # Create MATRIX2D_DOUBLE data point with default matrix metadata,
↳OldbDoubleMatrixStd
    # and provided initial value. The default metadata must exist on Config,
↳Service.

    matrix_double = elt.olddb.VectorDOUBLE([
        1.0, 20.0, 30.0, 40.0, 50.0, 4.4, 45.3, 34.4, 445.3, 301.3,
        2.0, 33.3, 34.3, 33.3, 33.3, 5.5, 32.2, 33.4, 222.3, 203.1])
```

(continues on next page)



(continued from previous page)

```
mdi_matrix_double = elt.config.CiiConfigClient.retrieve_metadata(  
    elt.olddb.CiiOldbGlobal.get_target_config_storage(),  
    elt.olddb.CiiOldbGlobal.MDI_MATRIX_DOUBLE,  
    elt.olddb.CiiOldbGlobal.CONFIG_VERSION)  
  
is_matrix = True  
  
matrix_dp = oldb_client.create_data_point_by_value(matrix_dp_uri, matrix_  
↪double, is_matrix)  
  
# Create STRING data point with default string metadata OldbStringStd anf_  
↪provided initial value  
  
mdi_string = elt.config.CiiConfigClient.retrieve_metadata(  
    elt.olddb.CiiOldbGlobal.get_target_config_storage(),  
    elt.olddb.CiiOldbGlobal.MDI_STRING,  
    elt.olddb.CiiOldbGlobal.CONFIG_VERSION)  
  
string_dp = oldb_client.create_data_point_by_value(string_dp_uri, "ABCDEF")  
  
# Create VECTOR_INT32 data point with default vector metadata_  
↪OldbDoubleVectorStd  
  
mdi_vector32 = elt.config.CiiConfigClient.retrieve_metadata(  
    elt.olddb.CiiOldbGlobal.get_target_config_storage(),  
    elt.olddb.CiiOldbGlobal.MDI_VECTOR_INT32,  
    elt.olddb.CiiOldbGlobal.CONFIG_VERSION)  
  
vector_dp = oldb_client.create_data_point(vector_dp_uri,  
↪                               elt.olddb.CiiOldbGlobal.MDI_VECTOR_  
↪INT32)  
  
# Read values  
const_dp_value = double_dp.read_value()  
matrix_dp_value = matrix_dp.read_value()  
string_dp_value = string_dp.read_value()  
  
# vectors do not have default value when created, needs a write before read  
  
vector_dp.write_value([2, 4, 6])  
vector_dp_value = vector_dp.read_value()  
  
print('Initial Value of double_dp: ', const_dp_value.get_value())  
print('Initial Value of string_dp: ', string_dp_value.get_value())  
print('First Value of vector_dp: ', vector_dp_value.get_value())  
  
# Write new values to the Data points
```

(continues on next page)



(continued from previous page)

```
new_matrix_value = elt.olddb.VectorDOUBLE()
for x in range(250):
    new_matrix_value.append(x)

double_dp.write_value(2.7)
matrix_dp.write_value(new_matrix_value)
string_dp.write_value('New string value')
vector_dp.write_value([7, 8, 9, 10])

# Read values
const_dp_value = double_dp.read_value()
matrix_dp_value = matrix_dp.read_value()
string_dp_value = string_dp.read_value()
vector_dp_value = vector_dp.read_value()

print('New Value of double_dp: ', const_dp_value.get_value())
print('New value of matrix_dp: ', matrix_dp_value.get_value())
print('New Value of string_dp: ', string_dp_value.get_value())
print('New Value of vector_dp: ', vector_dp_value.get_value())

# Get Multiple data points satisfying URI glob expression

uris = [elt.config.Uri('cii.olddb:/%s/sampleroot/**' % RND_PREFIX)]

search_result = oldb_client.get_data_points(uris)
print('Search Result: ', search_result)

# Get Multiple DOUBLE data points with value between 1.0 and 3.0

filtered_search_result = oldb_client.get_data_points_double(
    uris, elt.config.CiiBasicDataType.DOUBLE, 1.0, 3.0)
print('Filtered search result: ', filtered_search_result)

# Get the children of directory URI cii.olddb:/sampleroot. Returns a map_
↪specifying whether
# a child is a data point or just a directory

children = oldb_client.get_children(elt.config.Uri('cii.olddb:/sampleroot'))

# Delete data points
oldb_client.delete_data_point(double_dp_uri)
oldb_client.delete_data_point(matrix_dp_uri)
oldb_client.delete_data_point(string_dp_uri)
oldb_client.delete_data_point(vector_dp_uri)

# Delete meta data
elt.config.CiiConfigClient.delete_metadata(
    elt.olddb.CiiOldbGlobal.get_target_config_storage(),
    double_dp_meta.get_instance_name())
```

(continues on next page)



(continued from previous page)

```
    return 0

def main():
    """main method wrapper"""
    result = 0
    try:
        result = _main()
        #pylint: disable=W0703
    except Exception as e:
        print(e)
        traceback.print_tb(sys.exc_info()[2])
        result = 5
    return result

if __name__ == '__main__':
    sys.exit(main())
```

Subscribing Example

The examples in this section can be found in the oldb-examples project (SubscriptionExample.java, subscription-sample-app.cpp, cii-oldb-examples-subscription-sample-app-py.py).

In the subscribing example (Listing 4-4, Listing 4-5, Listing 4-6) a data point of type DOUBLE is created and a subscription is added to it to detect value changes and deletion.

Listing 4-4: Java Subscribing Example

```
package sampleapp;

import java.net.URI;

import elt.oldb.client.CiiOldb;
import elt.oldb.client.CiiOldbDataPoint;
import elt.oldb.client.CiiOldbDpSubscription;
import elt.oldb.client.CiiOldbDpValue;
import elt.oldb.exceptions.CiiOldbConfigError;
import elt.oldb.exceptions.CiiOldbCyclicDependencyException;
import elt.oldb.exceptions.CiiOldbDeletedException;
import elt.oldb.exceptions.CiiOldbDpConstraintsException;
import elt.oldb.exceptions.CiiOldbDpExistsException;
import elt.oldb.exceptions.CiiOldbDpQualityException;
import elt.oldb.exceptions.CiiOldbExpressionException;
import elt.oldb.exceptions.CiiOldbIllegalOperationException;
import elt.oldb.exceptions.CiiOldbInitializationError;
import elt.oldb.exceptions.CiiOldbStorageError;
import elt.oldb.exceptions.CiiOldbURIException;
import elt.oldb.exceptions.CiiOldbUndefinedMetadataException;
import elt.oldb.exceptions.CiiOldbWriteDisabledException;
```

(continues on next page)



(continued from previous page)

```
public class SubscriptionExampleApplication {  
  
    public static void main(String[] args) {  
  
        CiiOldb oldbClient = null;  
  
        try {  
            // Initialize OLDB and Config clients  
            oldbClient = CiiOldb.getInstance();  
  
            // Define URIs  
            URI doubleDpURI = URI.create("cii.oldb:/sampleroot/child/device/doubledp");  
  
            CiiOldbDataPoint<Double> doubleDp = oldbClient.  
↪createDataPointByValue(doubleDpURI, 1.0);  
  
            // Subscribe to a data point value changes.  
            doubleDp.subscribe(new CiiOldbDpSubscription() {  
  
                @Override  
                public void newValue(CiiOldbDpValue<?> value, URI uri) {  
                    try {  
                        System.out.println("newValue: " + uri + " " + value.getValue());  
                    } catch (CiiOldbDpQualityException e) {  
                        // Should not happen  
                    }  
                }  
  
                @Override  
                public void dpRemoved(URI uri) {  
                    System.out.println("dpRemoved: " + uri);  
                }  
            });  
  
            // Write to data points  
            doubleDp.writeValue(2.0);  
  
            // Delete data points  
            oldbClient.deleteDataPoint(doubleDpURI);  
  
            //Close the clients  
            oldbClient.close();  
  
        } catch (CiiOldbInitializationError | CiiOldbDpExistsException | ↪  
↪CiiOldbCyclicDependencyException  
            | CiiOldbExpressionException | CiiOldbDpQualityException | ↪  
↪CiiOldbDpConstraintsException  
            | CiiOldbUndefinedMetadataException | CiiOldbStorageError | ↪  
↪CiiOldbConfigError | CiiOldbURIException
```

(continues on next page)



(continued from previous page)

```
        | CiiOldbWriteDisabledException | CiiOldbDeletedException e) {  
    e.printStackTrace();  
    if (oldbClient != null) {  
        oldbClient.close();  
    }  
}  
  
}  
  
}
```

Listing 4-5: C++ Subscribing Example

```
/**  
 * @copyright (c) Copyright ESO 2019 All Rights Reserved  
 * ESO (eso.org) is an Intergovernmental Organisation, and therefore special_  
↳ legal conditions apply.  
 * @ingroup oldb-examples  
 */  
#include <mal/utility/Uri.hpp>  
  
#include <ciiOldbFactory.hpp>  
#include <ciiOldbSubscription.hpp>  
#include <ciiOldbDpValue.hpp>  
#include <ciiOldbExceptions.hpp>  
  
namespace elt {  
namespace oldb {  
namespace app {  
  
class AppOldbDpSubscription : public CiiOldbDpSubscription<double> {  
public:  
    AppOldbDpSubscription():  
        CiiOldbDpSubscription<double> (::elt::common::CiiBasicDataType::DOUBLE) {}  
  
    void DpRemoved (::elt::mal::Uri uri) override {  
        std::cout << "dpRemoved:" << uri.string() << '\n';  
    }  
  
    void NewValue (std::shared_ptr<CiiOldbDpValue<double>> value, ::elt::mal::Uri_  
↳ uri) override {  
        std::cout << "newValue:" << uri.string() << " " << value->GetValue() << '\n';  
    }  
};  
  
} // namespace app  
} // namespace oldb  
} // namespace elt
```

(continues on next page)



(continued from previous page)

```
int main(int ac, char *av[]) {
    try {
        // Initialize OLDB
        std::shared_ptr<elt::olddb::CiiOldb> oldb_client =
↪::elt::olddb::CiiOldbFactory::GetInstance();
        ::elt::olddb::CiiOldbGlobal::SetWriteEnabled(true);

        // Define URIs
        std::string RND_PREFIX = ::elt::olddb::CiiOldbUtil::NewUUID();
        ::elt::mal::Uri double_dp_uri{"cii.olddb/" + RND_PREFIX + "sampleroot/child/
↪device/doubledp"};

        // CREATE DATA POINTS

        // meta data needs to be created before calling CreateDataPoint methods

        std::shared_ptr<::elt::config::CiiConfigClient::CiiDataPointMetadataBase>
↪mdi_double =
        ::elt::config::CiiConfigClient::RetrieveMetadata(
            ::elt::olddb::CiiOldbGlobal::GetTargetConfigStorage(),
            ::elt::olddb::CiiOldbGlobal::MDI_DOUBLE,
            ::elt::olddb::CiiOldbGlobal::CONFIG_VERSION);

        std::shared_ptr<::elt::olddb::CiiOldbDataPoint<double>> double_dp = oldb_
↪client->CreateDataPointByValue(double_dp_uri, 1.0);

        // Subscribe to a data point value changes.

        std::shared_ptr<::elt::olddb::app::AppOldbDpSubscription> subscription =
↪std::make_shared<::elt::olddb::app::AppOldbDpSubscription>();

        double_dp->Subscribe(subscription);

        // Subscription takes affect after some period after Subscribe

        std::this_thread::sleep_for(std::chrono::seconds(2));

        // Write to data points

        double_dp->WriteValue(2.0);

        // Takes time for subscription listener to be called

        std::this_thread::sleep_for(std::chrono::seconds(2));

        // Delete data points
```

(continues on next page)



(continued from previous page)

```
oldb_client->DeleteDataPoint(double_dp_uri);

std::this_thread::sleep_for(std::chrono::seconds(3));

return 0;
} catch (const ::elt::oldb::CiiOldbException& ex) {
    std::cerr << "CiiOldbException occured while executing sample code. What: "
                << ex.what() << '\n';
}
return -1;
}
```

Listing 4-6: Python Subscribing Example

```
#!/usr/bin/env python
"""
@copyright (c) Copyright ESO 2019 All Rights Reserved
ESO (eso.org) is an Intergovernmental organisation,
and therefore special legal conditions apply.
@author Cosylab
"""
#pylint: disable=E1101,C0103,R0914,C0330,W0612

#This script presents DP subscription example.

import sys
import traceback
import time
import threading
import uuid
import elt.config
import elt.oldb

RND_PREFIX = str(uuid.uuid4())

# Initialize OLDB client
oldb_client = elt.oldb.CiiOldbFactory.get_instance()

# Define URIs
double_dp_uri = elt.config.Uri("cii.oldb:/%s/sampleroot/child/device/double_dp"
↪ % RND_PREFIX)

class AppOldbDpSubscription:
    """Subscription listener implementation, must implement
    new_value and dp_removed methods
    """
    #pylint: disable=R0201
    def new_value(self, value, uri):
        """Handle DP value change event"""
```

(continues on next page)



(continued from previous page)

```
print('new_value, value=%s, uri=%s' % (value.get_value(), uri.string()))

def dp_removed(self, uri):
    """Handle DP remove event"""
    print('dp_removed, uri=%s' % (uri.string(),))

class WorkThread(threading.Thread):
    """Work thread implementation"""

    def run(self):
        """Modify DP then delete it"""
        dp = oldb_client.get_data_point(double_dp_uri)
        for x in range(1, 3):
            dp.write_value(x*.5)
            time.sleep(0.2)
        del dp
        oldb_client.delete_data_point(double_dp_uri)
        # Give subscription system time to pass the message around
        time.sleep(1.0)

def _main():
    """main method implementation"""
    # enable writing
    elt.oldb.CiiOldbGlobal.set_write_enabled(True)

    # create data points
    # metadata needs to be created before calling create_data_point method

    mdi_double = elt.config.CiiConfigClient.retrieve_metadata(
        elt.oldb.CiiOldbGlobal.get_target_config_storage(),
        elt.oldb.CiiOldbGlobal.MDI_DOUBLE,
        elt.oldb.CiiOldbGlobal.CONFIG_VERSION)

    double_dp = oldb_client.create_data_point_by_value(double_dp_uri, 1.0)

    # Subscribe to data point. Pass instance of SubscriptionListener() to the
    ↪subscription.
    listener = elt.oldb.typesupport.DOUBLE.get_new_subscription_
    ↪instance(AppOldbDpSubscription())
    double_dp.subscribe(listener)

    # Do not write to the data points from the same thread. Deadlock because of
    ↪GIL.
    # launch work thread that manipulates data
    workThread = WorkThread()
    workThread.start()
    workThread.join()
    return 0
```

(continues on next page)



(continued from previous page)

```
def main():
    """main method wrapper"""
    result = 0
    try:
        result = _main()
        #pylint: disable=W0703
    except Exception as e:
        print(e)
        traceback.print_tb(sys.exc_info()[2])
        result = 5
    return result

if __name__ == '__main__':
    sys.exit(main())
```

Advanced Example

The examples in this section can be found in the oldb-examples project (AdvancedExample.java, advanced-sample-app.cpp, cii-oldb-examples-advanced-sample-app-py.py).

The advanced example (Listing 4-8, Listing 4-9, Listing 4-10) covers the advanced topics of working with the OLDB library. Note that the calculation service must be running for this example to produce expected results. In the example the following actions are executed:

1. A constant DOUBLE data point is created (inputDubleDp),
2. a DOUBLE data point with quality expression is created (squareDoubleDp),
3. a constant VECTOR_DOUBLE data point is created (inputVectorDp),
4. a calculated DOUBLE data point is created (sumOfSquaresDp), It calculates the sum of squares of inputVectorDp elements. It uses a custom function which is defined in the \$ INT-ROOT/etc/oldb/octave/oldb_custom.m file (Listing 4-7),
5. a value that causes the quality expression evaluate to BAD is written to squareDoubleDp,
6. the value of squareDoubleDp is rewritten to change the quality back to OK,
7. the storage location of inputVectorDp value is changed to in-memory-only storage,
8. the metadata of squareDoubleDp is updated with formula so that squareDoubleDp becomes a calculated data point. It calculates the square of inputDubleDp value,
9. write operations are disabled and check is done that it is indeed impossible to write to data points,
10. data points are deleted.

Listing 4-7: Custom Formula Function.



```
function returnval = SUM_SQUARE(v)
    s = 0
    for i = v
        s = s + i^2
    endfor
    returnval = s
endfunction
```

Listing 4-8: Java Advanced Example

```
package sampleapp;

import java.net.URI;
import java.util.ArrayList;
import java.util.Arrays;
import java.util.List;

import elt.config.classes.CiiOldbDataPointInfo;
import elt.config.classes.meta.MdOldbNumber;
import elt.config.classes.meta.MdString;
import elt.config.client.CiiConfigClient;
import elt.config.exceptions.CiiConfigInitializationError;
import elt.config.exceptions.CiiConfigNoTcException;
import elt.config.exceptions.CiiConfigSaveException;
import elt.config.exceptions.CiiConfigUpdateException;
import elt.config.exceptions.CiiConfigWriteDisabledException;
import elt.error.CiiInvalidTypeException;
import elt.error.CiiInvalidURIException;
import elt.oldb.client.CiiOldb;
import elt.oldb.client.CiiOldbDataPoint;
import elt.oldb.client.CiiOldbDpValue;
import elt.oldb.client.impl.CiiOldbImp;
import elt.oldb.client.util.CiiOldbUriUtils;
import elt.oldb.exceptions.CiiOldbConfigError;
import elt.oldb.exceptions.CiiOldbCyclicDependencyException;
import elt.oldb.exceptions.CiiOldbDeletedException;
import elt.oldb.exceptions.CiiOldbDpConstraintsException;
import elt.oldb.exceptions.CiiOldbDpExistsException;
import elt.oldb.exceptions.CiiOldbDpQualityException;
import elt.oldb.exceptions.CiiOldbDpUndefinedException;
import elt.oldb.exceptions.CiiOldbExpressionException;
import elt.oldb.exceptions.CiiOldbIllegalOperationException;
import elt.oldb.exceptions.CiiOldbInitializationError;
import elt.oldb.exceptions.CiiOldbInvalidTypeException;
import elt.oldb.exceptions.CiiOldbStorageError;
import elt.oldb.exceptions.CiiOldbURIException;
import elt.oldb.exceptions.CiiOldbUndefinedMetadataException;
import elt.oldb.exceptions.CiiOldbWriteDisabledException;

public class AdvancedSampleApp {
```

(continues on next page)



(continued from previous page)

```
public static void main(String[] args) {

    CiiOldb oldbClient = null;
    CiiConfigClient configClient = null;
    URI inputDoubleUri = null;
    URI squareDoubleUri = null;
    URI inputVectorUri = null;
    URI sumofsquaresDoubleUri = null;

    try {
        // Initialize OLDB and Config clients

        oldbClient = CiiOldb.getInstance();
        configClient = CiiConfigClient.getInstance();

        inputDoubleUri = URI.create("cii.olddb:/sampleroot/inputdouble");
        squareDoubleUri = URI.create("cii.olddb:/sampleroot/squaredouble");
        inputVectorUri = URI.create("cii.olddb:/sampleroot/inputdector");
        sumofsquaresDoubleUri =
            URI.create("cii.olddb:/sampleroot/sumofsquaresdouble");

        oldbClient.deleteDataPoint(squareDoubleUri);
        oldbClient.deleteDataPoint(inputDoubleUri);
        oldbClient.deleteDataPoint(sumofsquaresDoubleUri);
        oldbClient.deleteDataPoint(inputVectorUri);

        // Create a constant DOUBLE data point
        CiiOldbDataPoint<Double> inputDubleDp =
            oldbClient.createDataPointByValue(inputDoubleUri, 1.1);

        // Create a DOUBLE data point with a quality expression
        MdOldbNumber<Double> squareMd =
            new MdOldbNumber<Double>("squareOfSumMd", "", 2.0, null,
                "if (CURR_VAL() > 1.0) SETQUAL(\"OK\") else SETQUAL(\"BAD\") endif",
                null, null, null);
        URI squareMdURI = URI.create(CiiOldbUriUtils.CONFIG_METADATA_URI
            + squareMd.getMetadataInstanceName());
        configClient.saveMetadata(squareMdURI, squareMd);

        CiiOldbDataPoint<Double> squareDoubleDp = (CiiOldbDataPoint<Double>)
            oldbClient.createDataPoint(squareDoubleUri,
                squareMd.getMetadataInstanceName(), squareMd.getClass().
                getSimpleName());

        // Create a VECTOR_DOUBLE data point
        CiiOldbDataPoint<List<Double>> inputVectorDp = oldbClient.
        createDataPointByValue(
            inputVectorUri,
```

(continues on next page)



(continued from previous page)

```
        new ArrayList<>(Arrays.asList(1.0, 2.0, 3.0)), Double.class);

        // Create a DOUBLE data point with formula
        MdOldbNumber<Double> sumOfSquaresMd = new MdOldbNumber<Double>(
↪ "sumOfSquares", "", 0.0,
        String.format("SUM_SQUARE(${%s})", inputVectorUri.toString()), null,
↪ null, null, null);
        URI squareOfSumMdURI = URI
            .create(CiiOldbUriUtils.CONFIG_METADATA_URI
                + sumOfSquaresMd.getMetadataInstanceName());
        configClient.saveMetadata(squareOfSumMdURI, sumOfSquaresMd);

        CiiOldbDataPoint<Double> sumOfSquaresDp = (CiiOldbDataPoint<Double>)
            oldbClient.createDataPoint(
                sumofsquaresDoubleUri, sumOfSquaresMd.getMetadataInstanceName(),
                sumOfSquaresMd.getClass().getSimpleName());

        try {
            Thread.sleep(100);
        } catch (InterruptedException e) {
            // TODO Auto-generated catch block
            e.printStackTrace();
        }

        // Read the values of data points to confirm the calculation have been
↪ done.
        CiiOldbDpValue<Double> squareDpValue = squareDoubleDp.readValue();
        System.out.println(squareDoubleUri);
        System.out.println("value :" + squareDpValue.getValue());
        System.out.println("quality :" + squareDpValue.getQuality());

        CiiOldbDpValue<Double> sumOfSquaresDpValue = sumOfSquaresDp.readValue();
        System.out.println(sumofsquaresDoubleUri);
        System.out.println("value :" + sumOfSquaresDpValue.getValue());
        System.out.println("quality :" + sumOfSquaresDpValue.getQuality());
        try {
            Thread.sleep(2000);
        } catch (InterruptedException e) {
            // TODO Auto-generated catch block
            e.printStackTrace();
        }

        // Write a value that causes the quality expression to evaluate quality of
↪ data
        // point to BAD
        squareDoubleDp.writeValue(0.1);
        try {
            Thread.sleep(2000);
        } catch (InterruptedException e) {
```

(continues on next page)



(continued from previous page)

```
// TODO Auto-generated catch block
e.printStackTrace();
}
// Check that the quality of data point is BAD.
squareDpValue = squareDoubleDp.readValue();
System.out.println(squareDoubleUri);
System.out.println("quality :" + squareDpValue.getQuality());
try {
    squareDpValue.getValue();
} catch (CiiOldbDpQualityException e) {
    System.out.println(
        "Cannot call getValue because the quality is bad."
        +" The value can be retrieved from exception: "
        + e.getValue());
}

// Write a value that causes the quality expression to evaluate quality of
↪data
// point to OK
squareDoubleDp.writeValue(1.3);
try {
    Thread.sleep(2000);
} catch (InterruptedException e) {
    // TODO Auto-generated catch block
    e.printStackTrace();
}
// Check that the quality of data point is back to OK.
squareDpValue = squareDoubleDp.readValue();
System.out.println(squareDoubleUri);
System.out.println("value :" + squareDpValue.getValue());
System.out.println("quality :" + squareDpValue.getQuality());

// Change the storage location of the data point to in-memory-only storage
URI inputVectorConfigURI = URI.create(CiiOldbUriUtils.CONFIG_INFO_URI_ROOT
    + inputVectorUri.getPath());
CiiOldbDataPointInfo infoConfig = configClient.
↪retrieveConfig(inputVectorConfigURI)
    .getData(CiiOldbDataPointInfo.class);
MdString serverAlias = new MdString("testExtRedisServer");
serverAlias.setVersion(1);
infoConfig.setServerAlias(serverAlias);
configClient.updateConfig(inputVectorConfigURI, infoConfig);

try {
    Thread.sleep(100);
} catch (InterruptedException e) {
    // TODO Auto-generated catch block
    e.printStackTrace();
}
}
```

(continues on next page)



(continued from previous page)

```
inputVectorDp = (CiiOldbDataPoint<List<Double>>) oldbClient.  
↳getDataPoint(inputVectorUri);  
inputVectorDp.writeValue(new ArrayList<>(Arrays.asList(2.0, 3.0, 4.0)));  
  
// Update a constant data point metadata to change it into calculated data_  
↳point  
// (add a formula to metadata)  
squareMd.setFormula(String.format("${%s}^2", inputDoubleUri.toString()));  
  
configClient.updateMetadata(  
    CiiOldbUriUtils.createMetaConfigUri(squareMd.  
↳getMetadataInstanceName()),  
    squareMd.getMetadataInstanceName(), -1, squareMd);  
squareDoubleDp.setMetadata(squareMd.getMetadataInstanceName());  
  
try {  
    Thread.sleep(1000);  
} catch (InterruptedException e) {  
    // TODO Auto-generated catch block  
    e.printStackTrace();  
}  
  
// Read the data point to confirm that the value is now calculated  
squareDpValue = squareDoubleDp.readValue();  
System.out.println(squareDoubleUri);  
System.out.println("value :" + squareDpValue.getValue());  
System.out.println("quality :" + squareDpValue.getQuality());  
  
// Disable write operations on data points  
oldbClient.setWriteEnabled(false);  
try {  
    inputDubleDp.writeValue(0.2);  
} catch (CiiOldbWriteDisabledException e) {  
    System.out.println("Write on DP not permitted");  
}  
oldbClient.setWriteEnabled(true);  
  
// Delete data points  
oldbClient.deleteDataPoint(squareDoubleUri);  
oldbClient.deleteDataPoint(inputDoubleUri);  
oldbClient.deleteDataPoint(sumofsquaresDoubleUri);  
oldbClient.deleteDataPoint(inputVectorUri);  
  
// close OLDB and Config clients  
configClient.close();  
oldbClient.close();  
  
} catch (CiiConfigInitializationError | CiiOldbInitializationError |  
↳CiiOldbDpExistsException
```

(continues on next page)



(continued from previous page)

```
        | CiiOldbCyclicDependencyException | CiiOldbExpressionException |  
↪CiiOldbDpQualityException  
        | CiiOldbDpConstraintsException | CiiOldbUndefinedMetadataException |  
↪CiiOldbStorageError  
        | CiiOldbConfigError | CiiOldbURIException |  
↪CiiOldbWriteDisabledException | CiiOldbDeletedException  
        | CiiInvalidTypeException | CiiConfigSaveException |  
↪CiiConfigWriteDisabledException  
        | CiiInvalidURIException | CiiOldbDpUndefinedException |  
↪CiiOldbInvalidTypeException  
        | CiiConfigNoTcException | CiiConfigUpdateException e) {  
    e.printStackTrace();  
    if (configClient != null) {  
        configClient.close();  
    }  
    if (oldbClient != null) {  
        oldbClient.close();  
    }  
}  
  
}
```

Listing 4-9: C++ Advanced Example

```
/**  
 * @copyright (c) Copyright ESO 2019 All Rights Reserved  
 * ESO (eso.org) is an Intergovernmental Organisation, and therefore special_  
↪legal conditions apply.  
 * @ingroup oldb-examples  
 */  
#include <mal/utility/Uri.hpp>  
  
#include <ciiOldbFactory.hpp>  
#include <ciiOldbExceptions.hpp>  
#include <meta/mdOldbNumber.hpp>  
#include <vector>  
int main(int ac, char *av[]) {  
    try {  
        // Initialize OLDB  
        std::shared_ptr<elt::oldb::CiiOldb> oldb_client =  
↪::elt::oldb::CiiOldbFactory::GetInstance();  
        ::elt::oldb::CiiOldbGlobal::SetWriteEnabled(true);  
        std::string RND_PREFIX = ::elt::oldb::CiiOldbUtil::NewUUID();  
  
        // Define URIs  
  
        ::elt::mal::Uri input_double_uri{
```

(continues on next page)



(continued from previous page)

```
"cii.oldb/" + RND_PREFIX + "/sampleroot/child/device/doubledp");
::elt::mal::Uri square_double_uri{"cii.oldb/" + RND_PREFIX + "/sampleroot/
↪squaredouble"};
::elt::mal::Uri input_vector_uri{"cii.oldb/" + RND_PREFIX + "/sampleroot/
↪inputvector"};
::elt::mal::Uri sum_of_squares_double_uri{
    "cii.oldb/" + RND_PREFIX + "/sampleroot/sumofsquaresdouble"};

// should not be any data points with this uri since uri has random part
// (consistent with java sample)

try {
    oldb_client->DeleteDataPoint(square_double_uri);
} catch (const ::elt::oldb::CiiOldbException &) {}

try {
    oldb_client->DeleteDataPoint(input_double_uri);
} catch (const ::elt::oldb::CiiOldbException &) {}

try {
    oldb_client->DeleteDataPoint(sum_of_squares_double_uri);
} catch (const ::elt::oldb::CiiOldbException &) {}

try {
    oldb_client->DeleteDataPoint(input_vector_uri);
} catch (const ::elt::oldb::CiiOldbException &) {}

// Create a constant DOUBLE data point

// meta data needs to be created before calling CreateDataPoint methods

std::shared_ptr<::elt::config::CiiConfigClient::CiiDataPointMetadataBase>_
↪mdi_double =
    ::elt::config::CiiConfigClient::RetrieveMetadata(
        ::elt::oldb::CiiOldbGlobal::GetTargetConfigStorage(),
        ::elt::oldb::CiiOldbGlobal::MDI_DOUBLE,
        ::elt::oldb::CiiOldbGlobal::CONFIG_VERSION);

std::shared_ptr<::elt::oldb::CiiOldbDataPoint<double>> input_double_dp = oldb_
↪client->CreateDataPointByValue(input_double_uri, 1.1);

// Create a DOUBLE data point with a quality expression

std::string square_md_instance_name = "squareOfSumMd";

std::shared_ptr<elt::config::CiiDataPointMetadataBase> square_md =_
↪CiiDataPointMetadataFactory::getNewMetadataInstance<
```

(continues on next page)



(continued from previous page)

```
        ::elt::config::classes::meta::MdOldbNumber<double>>(square_md_instance_
↪name);

square_md->set_default_value(2.0);
square_md->set_qualityExpression(
    "if (CURR_VAL() > 1.0) SETQUAL(\"OK\") else SETQUAL(\"BAD\") endif");

// meta data needs to be created before calling CreateDataPoint

::elt::oldb::CiiOldbUtil::SaveOrUpdateMetadata(*square_md);

std::shared_ptr<::elt::oldb::CiiOldbDataPoint<double>> square_double_dp =
↪oldb_client->CreateDataPoint<double>(
    square_double_uri, square_md_instance_name);

// Create a VECTOR_DOUBLE data point

std::vector<double> vector_double = { 1.0, 2.0, 3.0 };

std::shared_ptr<::elt::config::CiiConfigClient::CiiDataPointMetadataBase>
↪mdi_vector_double =
    ::elt::config::CiiConfigClient::RetrieveMetadata(
        ::elt::oldb::CiiOldbGlobal::GetTargetConfigStorage(),
        ::elt::oldb::CiiOldbGlobal::MDI_VECTOR_DOUBLE,
        ::elt::oldb::CiiOldbGlobal::CONFIG_VERSION);

bool is_matrix = false;
std::shared_ptr<::elt::oldb::CiiOldbDataPoint<std::vector<double>>> input_
↪vector_dp = oldb_client->CreateDataPointByValue(
    input_vector_uri, vector_double, is_matrix);

// Create a DOUBLE data point with formula

std::string sum_of_squares_md_instance_name = "sumOfSquares";

std::shared_ptr<elt::config::CiiDataPointMetadataBase> sum_of_squares_md =
↪CiiDataPointMetadataFactory::getNewMetadataInstance<
    ::elt::config::classes::meta::MdOldbNumber<double>>(sum_of_squares_md_
↪instance_name);

std::string formula = "SUM_SQUARE(${ " + input_vector_uri.string() + "})";
sum_of_squares_md->set_default_value(0.0);
sum_of_squares_md->set_formula(formula);

::elt::oldb::CiiOldbUtil::SaveOrUpdateMetadata(*sum_of_squares_md);

std::shared_ptr<::elt::oldb::CiiOldbDataPoint<double>> sum_of_squares_dp =
↪oldb_client->CreateDataPoint<double>(
    sum_of_squares_double_uri, sum_of_squares_md_instance_name);
```

(continues on next page)



(continued from previous page)

```
std::this_thread::sleep_for(std::chrono::milliseconds(100));

// Read the values of data points to confirm the calculation have been done.

std::shared_ptr<::elt::olddb::CiiOldbDpValue<double>> square_dp_value =
↳square_double_dp->ReadValue();
std::cout << square_double_dp->string() << '\n';
std::cout << "value: " << square_dp_value->GetValue() << '\n';
std::cout << "quality:"
    << ::elt::olddb::CiiOldbUtil::ToQualityStr(square_dp_value->GetQuality()) <
↳< "\n\n";

std::shared_ptr<::elt::olddb::CiiOldbDpValue<double>> sum_of_squares_dp_value
↳= sum_of_squares_dp->ReadValue();
std::cout << sum_of_squares_double_dp->string() << '\n';
std::cout << "value: " << sum_of_squares_dp_value->GetValue() << '\n';
std::cout << "quality:"
    << ::elt::olddb::CiiOldbUtil::ToQualityStr(sum_of_squares_dp_value->
↳GetQuality()) << "\n\n";

std::this_thread::sleep_for(std::chrono::milliseconds(2000));

// Write a value that causes the quality expression to evaluate quality of
↳data point to BAD

square_double_dp->WriteValue(0.1);

std::this_thread::sleep_for(std::chrono::milliseconds(2000));

// Check that the quality of data point is BAD.

try {
    square_dp_value = square_double_dp->ReadValue();
    std::cerr << "Reading bad value should throw an exception.\n ";
} catch(const ::elt::olddb::CiiOldbDpQualityWithValueException<double>& ex) {
    std::cout << "Quality: BAD\n";
}

// Write a value that causes the quality expression to evaluate quality of
↳data point to OK

square_double_dp->WriteValue(1.3);

std::this_thread::sleep_for(std::chrono::milliseconds(2000));

// Check that the quality of data point is back to OK.

square_dp_value = square_double_dp->ReadValue();
```

(continues on next page)



(continued from previous page)

```
std::cout << square_double_uri.string() << '\n';
std::cout << "value: " << square_dp_value->GetValue() << '\n';
std::cout << "quality:"
    << ::elt::oldb::CiiOldbUtil::ToQualityStr(square_dp_value->GetQuality()) <
↳< "\n\n";

// Change the storage location of the data point to in-memory-only storage

::elt::mal::Uri input_vector_config_uri(
    ::elt::oldb::CiiOldbGlobal::GetTargetConfigDataPointsUriPrefix()
    + input_vector_uri.path().to_string());

std::shared_ptr<::elt::config::classes::CiiOldbDataPointInfo> info_config =
    elt::config::CiiConfigClient::GetConfigData
↳<::elt::config::classes::CiiOldbDataPointInfo>(
    input_vector_config_uri, ::elt::oldb::CiiOldbGlobal::CONFIG_VERSION);

info_config->set_serverAlias("testExtRedisServer");

::elt::config::CiiConfigClient::SaveTargetConfig(input_vector_config_uri,
↳*info_config);

std::this_thread::sleep_for(std::chrono::milliseconds(100));

input_vector_dp = oldb_client->GetDataPoint<std::vector<double>>(input_
↳vector_uri);

input_vector_dp->WriteValue(std::vector<double> { 2.0, 3.0, 4.0 });

// Update a constant data point metadata to change it into calculated data_
↳point
// (add a formula to metadata)

formula = "${" + input_double_uri.string() + "}^2";
square_md->set_formula(formula);
square_double_dp->SetMetadata(square_md_instance_name);

std::this_thread::sleep_for(std::chrono::milliseconds(1000));

// Read the data point to confirm that the value is now calculated

square_dp_value = square_double_dp->ReadValue();
std::cout << square_double_uri.string() << '\n';
std::cout << "value: " << square_dp_value->GetValue() << '\n';
std::cout << "quality:"
    << ::elt::oldb::CiiOldbUtil::ToQualityStr(square_dp_value->GetQuality()) <
↳< "\n\n";
```

(continues on next page)



(continued from previous page)

```
// Disable write operations on data points

::elt::olddb::CiiOldbGlobal::SetWriteEnabled(false);

try {
    input_double_dp->WriteValue(0.2);
} catch(const ::elt::olddb::CiiOldbWriteDisabledException&) {
    std::cerr << "Write on DP not permitted\n";
}

::elt::olddb::CiiOldbGlobal::SetWriteEnabled(true);

// Delete data points

olddb_client->DeleteDataPoint(square_double_uri);
olddb_client->DeleteDataPoint(input_double_uri);
olddb_client->DeleteDataPoint(sum_of_squares_double_uri);
olddb_client->DeleteDataPoint(input_vector_uri);

return 0;
} catch (const ::elt::olddb::CiiOldbException& ex) {
    std::cerr << "CiiOldbException occured while executing sample code. What: "
        << ex.what() << '\n';
} catch (const std::exception& ex) {
    std::cerr << "Unknown exception occured while executing sample code. What: "
        << ex.what() << '\n';
}
return -1;
}
```

Listing 4-10: Python Advanced Example

```
#!/usr/bin/env python
"""
@copyright (c) Copyright ESO 2019 All Rights Reserved
ESO (eso.org) is an Intergovernmental organisation,
and therefore special legal conditions apply.
@ingroup client-apis-python
@author Cosylab
"""
#pylint: disable=E1101,C0103,R0914,C0330

# Make sure OLDB services are running before executing this script.

import sys
import traceback
import time
import uuid
import elt.config
```

(continues on next page)



(continued from previous page)

```
import elt.olddb

RND_PREFIX = str(uuid.uuid4())

# Define URIs
input_double_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/doubledp
↳" % RND_PREFIX)
square_double_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/
↳squaredouble"
                                   % RND_PREFIX)
input_vector_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/
↳inputvector" % RND_PREFIX)
sum_of_squares_uri = elt.config.Uri("cii.olddb:/s/sampleroot/child/device/
↳sumofsquaredouble"
                                     % RND_PREFIX)

def _main():
    """main method implementaton"""
    # Initialize OLDB
    #pylint: disable=R0915
    elt.olddb.CiiOldbGlobal.set_write_enabled(True)
    elt.olddb.CiiOldbGlobal.set_calc_enabled(True)

    # Initialize OLDB client
    oldb_client = elt.olddb.CiiOldbFactory.get_instance()

    # Remove Data points if they already exist
    # pylint: disable=W0703,W0702
    try:
        oldb_client.delete_data_point(input_double_uri)
    except:
        pass
    try:
        oldb_client.delete_data_point(square_double_uri)
    except:
        pass
    try:
        oldb_client.delete_data_point(input_vector_uri)
    except:
        pass
    try:
        oldb_client.delete_data_point(sum_of_squares_uri)
    except:
        pass

    # Create a constant DOUBLE data point

    # meta data needs to be created before calling create_data_point method
```

(continues on next page)



(continued from previous page)

```
#pylint: disable=W0612
mdi_double = elt.config.CiiConfigClient.retrieve_metadata(
    elt.olddb.CiiOldbGlobal.get_target_config_storage(),
    elt.olddb.CiiOldbGlobal.MDI_DOUBLE,
    elt.olddb.CiiOldbGlobal.CONFIG_VERSION)

input_double_dp = oldb_client.create_data_point_by_value(input_double_uri, 1.
↪1)

# Create a DOUBLE data point with quality expression

square_md = elt.olddb.typesupport.DOUBLE.get_new_number_metadata_instance(
↪'squareOfSumMd')
square_md.set_default_value(2.0)
square_md.set_quality_expression(
    'if (CURR_VAL() > 1.0) SETQUAL("OK") else SETQUAL("BAD") endif')

# meta data needs to be created before calling create_data_point

elt.olddb.CiiOldbUtil.save_or_update_metadata(square_md)

square_double_dp = oldb_client.create_data_point(square_double_uri,
square_md.get_instance_
↪name())

# Create VECTOR_DOUBLE data point
vector_double = elt.olddb.VectorDOUBLE([1.0, 2.0, 3.0])

mdi_vector_double = elt.config.CiiConfigClient.retrieve_metadata(
    elt.olddb.CiiOldbGlobal.get_target_config_storage(),
    elt.olddb.CiiOldbGlobal.MDI_VECTOR_DOUBLE,
    elt.olddb.CiiOldbGlobal.CONFIG_VERSION)

input_vector_dp = oldb_client.create_data_point_by_value(input_vector_uri, ↪
↪vector_double, False)

# Create DOUBLE data point with formula

sum_of_squares_md = elt.olddb.typesupport.DOUBLE.get_new_number_metadata_
↪instance('sumOfSquares')

formula = 'SUM_SQUARE(${s})' % (input_vector_uri.string(),)
sum_of_squares_md.set_default_value(0.0)
sum_of_squares_md.set_formula(formula)

# Create & Save Metadata
elt.olddb.CiiOldbUtil.save_or_update_metadata(sum_of_squares_md)

sum_of_squares_dp = oldb_client.create_data_point(sum_of_squares_uri,
```

(continues on next page)



(continued from previous page)

```
sum_of_squares_md.get_
↪instance_name())

    # Wait for calculations to finish
    time.sleep(3.0)

    square_dp_value = square_double_dp.read_value()
    print('%s, value: %s, quality: %s' % (square_double_uri.string(), square_dp_
↪value.get_value(),
        square_dp_value.get_quality()))

    sum_of_squares_dp_value = sum_of_squares_dp.read_value()
    print('%s, value: %s, quality: %s' % (sum_of_squares_uri.string(),
        sum_of_squares_dp_value.get_value(),
        sum_of_squares_dp_value.get_quality()))

    # Write a value that causes the quality expression to evaluate quality of
↪data point to BAD

    square_double_dp.write_value(0.1)

    time.sleep(2.0)

    try:
        square_dp_value = square_double_dp.read_value()
        print('%s, quality: %s' % (square_double_uri, square_dp_value.get_
↪quality()))
        square_dp_value.get_value()
    except elt.oldb.CiiOldbDpQualityException as e:
        print("Cannot call get_value because quality is bad.")

    # Write a value that causes the quality expression to evaluate quality of
↪data point back to OK

    square_double_dp.write_value(1.3)

    time.sleep(2.0)

    # Check that quality of data point is back to OK

    square_dp_value = square_double_dp.read_value()
    print('%s, value: %s, quality: %s' % (square_double_uri, square_dp_value.get_
↪value(),
        square_dp_value.get_quality()))

    # Change the storage location of the data point to in-memory-only storage

    input_vector_config_uri = elt.config.Uri(
```

(continues on next page)



(continued from previous page)

```
elt.olddb.CiiOldbGlobal.get_target_config_data_points_uri_prefix() +
input_vector_uri.path()

info_config = elt.config.CiiConfigClient.retrieve_config(input_vector_config_
↪uri,
                                                    elt.olddb.
↪CiiOldbGlobal.CONFIG_VERSION)

info_config.set_server_alias("testExtRedisServer")
elt.config.CiiConfigClient.save_target_config(input_vector_config_uri, info_
↪config)

time.sleep(0.1)

input_vector_dp = oldb_client.get_data_point(input_vector_uri)

input_vector_dp.write_value([2.0, 3.0, 4.0])

# Update a constant data point metadata to change it into calculated data_
↪point
# (add formula to metadata)

formula = '${%s}^2' % (input_double_uri.string(),)
square_md.set_formula(formula)
square_double_dp.set_metadata(square_md.get_instance_name())

time.sleep(2.0)

# Read the data point to confirm that the value is now calculated

square_dp_value = square_double_dp.read_value()
print('%s, value: %s, quality: %s' % (square_double_uri.string(), square_dp_
↪value.get_value(),
                                square_dp_value.get_quality()))

# Disable write operations on data points
elt.olddb.CiiOldbGlobal.set_write_enabled(False)

try:
    input_double_dp.write_value(7.7)
except elt.olddb.CiiOldbWriteDisabledException as e:
    print('Write on DP not permitted')

# Enable write operations
elt.olddb.CiiOldbGlobal.set_write_enabled(True)

# Delete data points
oldb_client.delete_data_point(square_double_uri)
oldb_client.delete_data_point(input_double_uri)
```

(continues on next page)



(continued from previous page)

```
oldb_client.delete_data_point(sum_of_squares_uri)
oldb_client.delete_data_point(input_vector_uri)

# Delete meta data
elt.config.CiiConfigClient.delete_metadata(elt.oldb.CiiOldbGlobal.get_target_
↪config_storage(),
                                           square_md.get_instance_name())
elt.config.CiiConfigClient.delete_metadata(elt.oldb.CiiOldbGlobal.get_target_
↪config_storage(),
                                           sum_of_squares_md.get_instance_
↪name())
    return 0

def main():
    """main method wrapper"""
    result = 0
    try:
        result = _main()
    #pylint: disable=W0703
    except Exception as e:
        print(e)
        traceback.print_tb(sys.exc_info()[2])
        result = 5
    return result

if __name__ == '__main__':
    sys.exit(main())
```



7.5 Advanced Topics

This section describes in detail the advanced functionalities of OLDB API library.

7.5.1 OLDB Statistics

Basic traffic statistics can be obtained for the OLDB Client and Calculation service.

OLDB Client Statistics

For OLDB client the basic operation statistics can be obtained directly from Redis by using the redis-cli command line client application. Note that Redis must be installed on the machine so that the redis-cli application is available. See [2] for how to install Redis. Redis client app can be run in the terminal. It can be given options -h (for hostname) and -p (for the port) to connect to specific Redis server. By default it will connect to localhost:6379.

```
redis-cli -h 10.71.0.247  
10.71.0.247:6379>
```

To get basic statistics about the read and write operations the INFO command with commandstats option should be run (Listing 5-1).

Listing 5-1: INFO commandstats result

```
10.71.0.247:6379> INFO commandstats  
# Commandstats  
cmdstat_expire:calls=54,usec=754,usec_per_call=13.96  
cmdstat_subscribe:calls=690,usec=9138,usec_per_call=13.24  
cmdstat_exists:calls=296,usec=2683,usec_per_call=9.06  
cmdstat_get:calls=948,usec=8600,usec_per_call=9.07  
cmdstat_mget:calls=11,usec=88,usec_per_call=8.00  
cmdstat_mset:calls=59,usec=662,usec_per_call=11.22  
cmdstat_del:calls=190,usec=1708,usec_per_call=8.99  
cmdstat_command:calls=3,usec=3138,usec_per_call=1046.00  
cmdstat_set:calls=29601,usec=411882,usec_per_call=13.91  
cmdstat_publish:calls=29992,usec=318727,usec_per_call=10.63  
cmdstat_info:calls=3,usec=273,usec_per_call=91.00
```

The most important fields in the (Listing 5-1) are:

cmdstat_get:calls – number of read operations.

cmdstat_set:calls – number of write operations.

cmdstat_del:calls – number of delete operations.

See [7] for full explanation of the INFO command and how to get other information about the Redis server (e.g. master – slave statistics for cluster servers).



Calculation Service Statistics

Calculation service exposes basic statistic using the Java MBeans JMX interface. Statistics can be obtained using various JMX tools. The example shows usage of jconsole. The user should start the jconsole with:

```
jconsole&
```

Calculation service can be chosen from the list of processes and connected to by clicking on the Connect button (Figure 5-1). In this example, a calculation node process is chosen. The same statistics can be found for the Scheduler process.

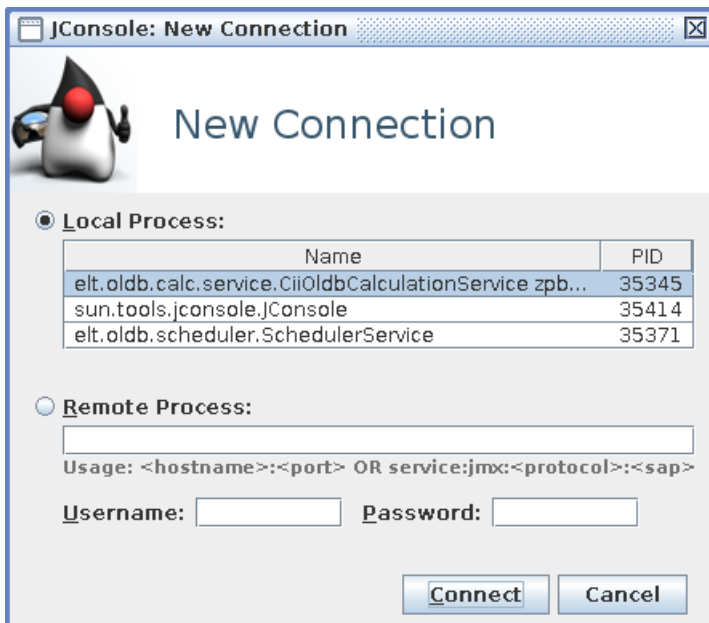


Figure 5-1: Connect to Calculation Service with Jconsole

After the connection is made, the console with statistics is displayed. The console displays heap usage, CPU usage, number of threads, and similar java statistics. The specific Calculation service statistics are exposed under the `elt.olddb.calc.service.stat` namespace.

To obtain statistics, the user should click on the MBeans tab and browse the `elt.olddb.calc.service.stat` namespace (Figure 5-2). The statistics show up-time of the calculation node and number of data points calculated by this calculation node (the number of calculation operations made).

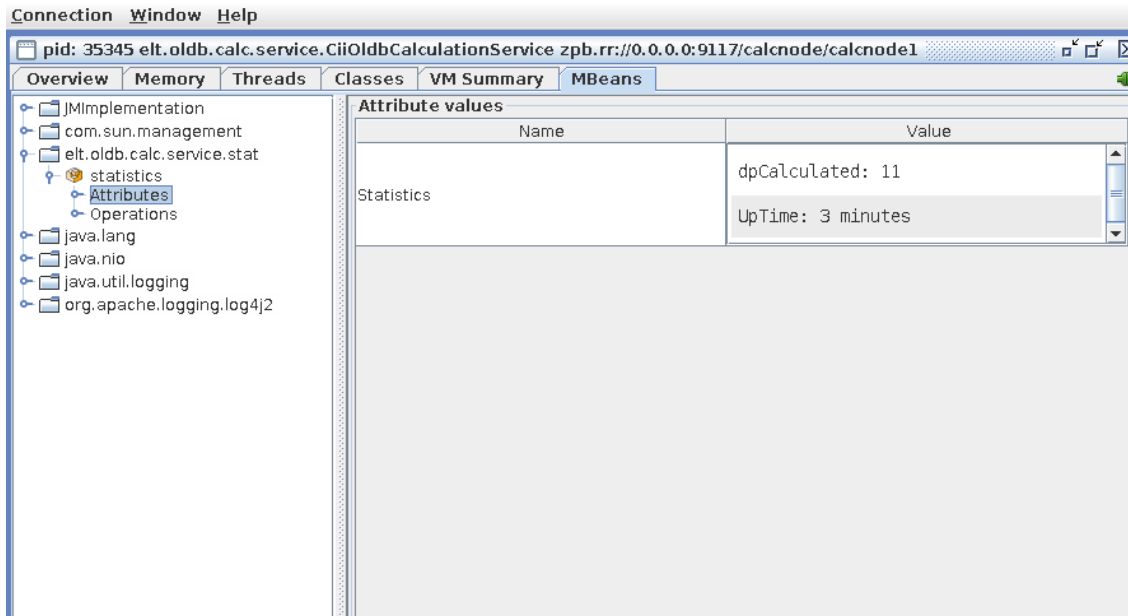


Figure 5-2: Calculation Service Statistics

7.5.2 Disabling Writes to Data Points

Write operations on data points can be disabled or enabled by calling the `setWriteEnabled` method on the OLDB client. If disabled, all calls of `writeValue` on data points will throw a `CiiOldbWriteDisabledException`. Note that the write disabling only effects the same process environment in which it was done. Two OLDB clients on different machines (or in different processes on the same machine) can have different settings regarding the permission of write operations. Write operations are by default enabled.

```
//Disable write operations on data points
oldbClient.setWriteEnabled(false);
try {
    inputDubleDp.writeValue(0.2);
} catch (CiiOldbWriteDisabledException e) {
    System.out.println("Write on DP not permitted");
}
oldbClient.setWriteEnabled(true);
```

When writes to the data points are enabled, a user can also enable an additional check for a formula existence when writing. If the data point formula is not empty and the write is not done by a calculation engine, write will fail with an exception `CiiOldbIllegalOperationException`. User can enable this check by calling the method `SetCheckFormulaIfWriteEnabled` in the OLDB client API.



7.5.3 Custom functions in Data Point Formulas

Formulas of calculated data points can include any functions from the Octave language, but also user defined functions. The user can include their custom Octave functions in the \$INT-ROOT/etc/olddb/octave/olddb_custom.m file.

```
function returnval = SUM_SQUARE(v)
    s = 0
    for i = v
        s = s + i^2
    endfor
    returnval = s
endfunction
```

These functions must be included in the file before the Calculation Service is started. Changes to oldb_custom.m file will not take effect if the Calculation Service is already running.

The custom functions can then be included in data point formulas.

```
//Create a DOUBLE data point with formula
MdOldbNumber<Double> sumOfSquaresMd = new MdOldbNumber<Double>("sumOfSquares", "
↳", 0.0, String.format("SUM_SQUARE(${s})", inputVectorUri.toString()), null,
↳null, null, null);
URI squareOfSumMdURI = URI.create(CiiOldbImp.CONFIG_METADATA_URI +
↳sumOfSquaresMd.getMetadataInstanceName());
configClient.saveMetadata(squareOfSumMdURI, sumOfSquaresMd);

CiiOldbDataPoint<Double> sumOfSquaresDp = oldbClient.
↳createDataPoint(sumofsquaresDoubleUri, sumOfSquaresMd);
```

To see how to write Octave functions see [6].

7.5.4 Manipulating Data Point Metadata

Metadata of the data point can be updated dynamically. Certain metadata attributes (e.g. limits or formula) can be changed or the data point's metadata instance can be replaced altogether (change of the metadata instance name). This can be done by calling the setMetadata method on the data point which accepts a string metadata instance name. The data point metadata can be changed in two ways:

1. The attributes of the metadata of this data point have been changed (the provided metadata instance name is the same as the one the data point already has). These changes must be saved to Configuration Service before calling this method for it to have any effect.
2. The metadata instance of the data point has been replaced (the provided metadata instance name is different than the one the data point already has). The new metadata must already exist and must be of the same type as the old one. This change will be propagated to the data point configuration (Section 5.5).



The change will be propagated to the calculation service so that if the change is relevant to it (e.g. a data point has become a calculated data point), it can take proper actions (e.g. add it to a calculation node).

```
//Update a constant data point metadata to change it into calculated data point
//(add a formula to metadata)
squareMd.setFormula(String.format("${%s}^2", inputDoubleUri.toString()));

configClient.updateMetadata (
CiiOldbUriUtils.createMetaConfigUri (squareMd.getMetadataInstanceName()),
squareMd.getMetadataInstanceName(), -1, squareMd);

squareDoubleDp.setMetadata (squareMd);
```

Note that the type of data point cannot be changed. Trying to update metadata with a different type of metadata will result in `CiiOldbInvalidTypeException`. That is, the class of the metadata cannot change once it has been defined in the data point at creation. If there is a need for data point type change the data point must be deleted and created again with the updated metadata.

The change will however not be propagated to other instances of the data point. For the change to take effect on these data points must be fetched again from the OLDB.

7.5.5 Data Point Configuration

Each data point has a configuration in Configuration Service which is created automatically at data point creation. It is primarily used for internal management so that the user needs not to concern themselves with it except in a case where the storage location needs to be changed.

The data point configuration is stored in the `CiiOldbDataPointInfo` object. This configuration class contains the following fields.

1. `uri`: the `MdString` class containing the URI of the data point as the string
2. `metadataInstanceName`: the `MdString` class containing the data point's metadata instance name as a string.
3. `metadataClassName`: the `MdString` class containing the data point's metadata class name as a string (See Table 6-1 for a list of OLDB metadata classes).
4. `serverAlias`: The `MdString` class containing the alias for external Redis server (in-memory-only data storage) on which this data point value is stored. All newly created data points have this field *null* which means that their value is stored in default storage (default Redis server or Remote FS).

The only field relevant to the user is the `serverAlias`, for it can be used to change the storage location of the data point (Section 5.7).



7.5.6 External Redis Servers

To provide the ability for low latency high throughput, the OLDB client has an additional setting that enables the usage of multiple independent Redis servers to run in parallel.

By default, all newly created data points are stored in the default storage location (default Redis servers). This location is defined in the data point configuration. For low latency high throughput configuration, it is possible to also define in-memory-only storage locations (external Redis servers) and transfer data point values there.

External Redis servers are used for data points with larger values that need quicker access than an Remote FS storage is able to provide. They are regular Redis servers (can be single instances or clusters), but data point values sliced and saved under multiple key-value pairs. The number of slices depends on the data size limit of each slice which is defined with the `valueSizeLimitRedisExt` field in the OLDB configuration (Section 3.3.1). The list of the external Redis servers that OLDB client will connect to is written in the OLDB configuration. This setting is read at OLDB client start up (See `redisServersExternal` filed in Section 3.3.1). There each external server is formatted as `<alias>:<hostname>:<port>:<server-type>` where `server-type` is either `SINGLE` or `CLUSTER` and describes the type of the server. This alias, which is defined by the administrator can be used to assign a data point to the external server. A single specific data point can be assigned to specific Redis External Server. The data point setting `serverAlias` is used to control this behavior (5.5).

As the in-memory-only servers are intended to be used for high throughput, the properly tuned Redis setting need to be applied. The setting can be set to single or cluster Redis, but when low latency is required a cluster can cause additional delay. By design, Redis is single-threaded, which means that every write or read to the server, will block it. In case the cluster is used, the time to copy the data between the cluster parts need to be considered. If large blocks of data are used, these times can be very long.

The data points that have value written to the external Redis server have also a key-value field on the default Redis storage. This key-value doesn't represent the actual data point value, but it is only the address of the external server (alias of the external Redis).

7.5.7 Changing the Storage Location of the Data Point

To change the Storage location of the data point, the user must update the data point configuration and change the `serverAlias` field. If this field is set to null the data point value will be moved to default Redis servers, otherwise, it will be moved to external Redis server with this alias. In code, this can be done using the `ConfigClient` API from the Configuration Service [3]. Example of changing data point storage location can be seen in Listing 5-2.

Listing 5-2: Example of changing data point storage location

```
//Change the storage location of the data point to in-memory-only storage
URI inputVectorConfigURI = URI.create(CiiOldbImp.CONFIG_INFO_URI_ROOT +
↳inputVectorUri.getPath());
```

(continues on next page)



(continued from previous page)

```
CiiOldbDataPointInfo infoConfig = configClient.  
↪ retrieveConfig(inputVectorConfigURI).getData(CiiOldbDataPointInfo.class);  
  
MdString serverAlias = new MdString("testExtRedisServer");  
serverAlias.setVersion(1);  
infoConfig.setServerAlias(serverAlias);  
configClient.updateTargetConfig(inputVectorConfigURI, infoConfig);  
  
try {  
    Thread.sleep(100);  
} catch (InterruptedException e) {  
    // TODO Auto-generated catch block  
    e.printStackTrace();  
}  
  
//The data point needs to be reread and written to so that the change takes  
↪ effect  
inputVectorDp = (CiiOldbDataPoint<List<Double>>) oldbClient.  
↪ getDataPoint(inputVectorUri);  
inputVectorDp.writeValue(new ArrayList<>(Arrays.asList(2.0, 3.0, 4.0)));
```

Note: the in-code configuration change does not alert data points. For the configuration change to take effect in the OLDB storages, the data point must be fetched again from the OLDB and a value needs to be written to it. After that also other instances of data points will read and write to the right storage location.

It is possible to alert all clients using a data point instance of the config change with Config GUI or CLI tool.

1. If the configuration is changed through the Config GUI [3], there is a “notify oldb” button to send a notification to all data points.
16. The CLI tool config-oldb-update notifies the clients of data point config changes. It takes one parameter (the config URI under which the data point configuration is saved in the configuration service). The data point configuration URI can be constructed from the data point URI

Example:

If the data point URI is `cii.oldb:/sampleroot/inputvector`, the URI of configuration is `cii.config:/remote/oldb/datapoints/sampleroot/inputvector`.

```
./config-oldb-update cii.config://remote/oldb/datapoints/sampleroot/inputvector
```



7.6 OLDB API LIBRARY

This section explains the details of the srv-oldb library. The OLDB API is split between two classes. The CiiOldb and CiiOldbDataPoint.

7.6.1 CiiOldb

The core OLDB Client API is the singleton CiiOldb class. The CiiOldb class contains methods for creating retrieving and deleting data points. See Appendix C for a description of CiiOldb API.

7.6.2 CiiOldbDataPoint

The other part of the OLDB Client API is the CiiOldbDataPoint class (Figure 3). It represents the data point in the OLDB and allows reading from and writing to it. The CiiOldbDataPoint class includes three main fields:

URI: the unique data point identifier.

CiiOldbDpValue: the value class containing the data point data value, timestamp and quality. It is parameterized with the type of the data value.

CiiOldbMetaData: the metadata class that contains data point type, formula, quality expression and other data type specific information.

The CiiOldbDataPoint is a generic class parameterized with the programming language type of the data value (for Java and C++).

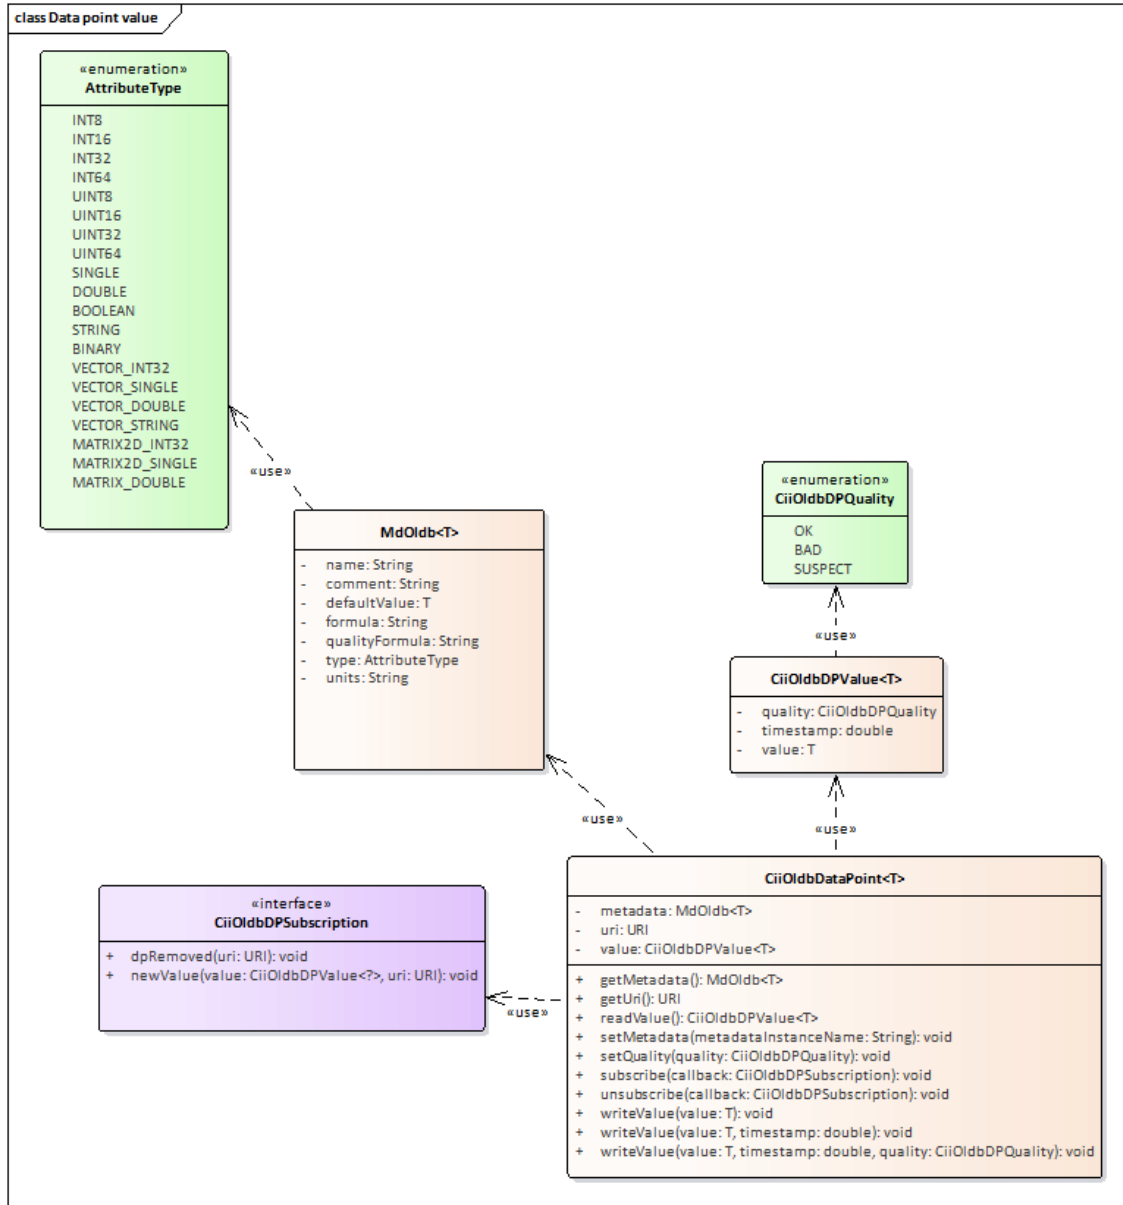


Figure 3: CiiOldbDataPoint class diagram

Data Point URI

Data point URI is the unique hierarchy identifier of the data point and contains OLDB schema `cii.oidb:///` and the path to the data point (e.g. `cii.oidb:///root/child/dp1`).

Note: To make OLDB hierarchy working as expected, the data point URI should not contain any triple underscore (“___”) character. Because internally all “/” characters are converted to triple underscores, a triple underscores in a level name would result in an additional level.



Note: For the same reason, do not use path components that end with an underscore. For example `cii.olddb:///a/b_c` will translate to a sequence of 4 underscores which would wrongly translate back to `cii.olddb:///a/b/_c`.

Note: Due to a limitation of the config service, the path component of the data point URI should not contain a sole number, i.e. `cii.olddb:///a/12/dp1` is not allowed. Creating a data point with a number in the URI path will throw an exception clearly stating that a number in URI path is not allowed.

```
URI uri = point.getUri();
```

Note: The OLDB is case insensitive so two URIs differing only in case are considered the same. All URIs are internally stored in lower case.

Note: It is valid that a data point URI path is a sub-path of another data point URI path since the Configuration Service does not hold the URIs in a hierarchy structure. So for example, if there a data point with URI `cii.olddb:///some/path/a`, a data point with `cii.olddb:///some/path/a/b` can be created.

Data Point Value

The `CiiOldbDpValue` contains three pieces of information that dynamically change during the data point lifetime:

Value: The actual data value of the data point. Data types of values are listed in Appendix A.

Quality: The `CiiOldbDpQuality` enumeration that represents the quality state of the data point. Possible values are OK, SUSPECT or BAD.

Timestamp: The moment in time when the data point was last modified.

The meanings of the quality state are as follows:

OK: the value of the Property is valid

BAD: the value is not OK and should not be used

SUSPECT: there is a value, but it is not clear if it is reliable. For example, the readout of a sensor is inconsistent with other readings and the sensor might be faulty.

If the data point quality is not OK, data point value cannot be read and written to. Note that the difference between BAD and SUSPECT is just semantic. The OLDB library treats these two states the same. If anything goes wrong with the data point internally (e.g. error in the data point formula) the quality is set to BAD. It is up to the user to set the quality to SUSPECT if appropriate.

The `CiiOldbDpValue` is a generic class parameterized with the programming language type of the data value (for Java and C++, but not Python because it is not strongly typed).

```
CiiOldbDpValue<T> value = point.readValue();
```



Metadata

The MdOldb is the class containing all the metadata information about the data point. It is not expected that this information would change often but it is possible to manipulate it through the OLDB API (Section 5.4). It is a subtype of MdBase from the Cii Config Service [3] so it contains all the fields contained in the MdBase. The OLDB relevant fields of the MdBase are genType (the type of the data point value) and defaultValue (the default value of the data point). The OLDB specific fields contained in the base class MdOldb are data point formula and qualityExpression.

For other fields in MdBase class consult the [3] document.

The CiiOldbDataPoint class contains a subtype of MdOldb specific to the data point type (Table 6-1). The user of the CiiOldbDataPoint must know the type of the data point and the metadata class corresponding to it so they can retrieve the metadata object, as seen in the following example:

```
MdOldbString metadata = point.getMetadata(MdOldbString.class);
```

If at runtime the type of the data point is not known, the user can determine it by retrieving the metadata as the MdOldb base type, retrieving the generic type from it and casting the metadata in the right class.

```
MdOldb metadata = point.getMetadata(MdOldb.class);  
AttributeType type = metadata.getGenType();  
if (type == AttributeType.STRING) {  
    metadata = (MOldbString) metadata;  
}
```

All the built-in OLDB metadata classes are listed in Table 6-1. All classes derive from MdOldb (which derives from config MdBase) class which defines basic OLDB fields (quality and formula). The metadata class associated with the data type can be read in Appendix B.

Note that it is not meant to extend the existing metadata classes and define new “types” of data points since the data types defined by Core Integration Infrastructure are covered with existing metadata. It is in theory possible to extend OLDB metadata classes, but it means extending the OLDB design and coding of new extended classes.

Table 6-1: Built-in metadata classes



Class	Description	Specific Fields
MdOldNumber	Metadata for all scalar number types. Parameterized with the language mapping of the type	minLimit (data point value type): the lower limit of the data point value. maxLimit (data point value type): the upper limit of the data point value. units (string): the units of the data point.
MdOldString	Metadata for string data types. Not parameterized.	allowedValues (list): all the values the data point value is allowed to have. If the list is empty or <i>null</i> , all values are allowed. size (integer): the length of the string value. This is not a limit just informational filed.
MdOldbArray	Metadata for all vector data points. Parameterized with the type of the elements.	
MdOldMatrix	Metadata for all matrix data points. Parameterized with the type of the elements. This is a subtype of MdOldbArray*	width (integer): The width of the matrix (the number of elements in each matrix row).
MdOldBinary	Metadata for byte array data points.	

*Note: The matrix type data point actually holds the matrix flattened in to a vector. The metadata then holds the width of the matrix (number of elements in each row) so the user can reconstruct the matrix (e.g. if data point value is a vector of length 25 and the width property of the metadata is 5, the data point represents a 5x5 matrix). The length of the vector value must be divisible by the width otherwise the CiiConfigDataLimitsException is thrown.

Data Point Type

The Core Integration Infrastructure defines the types of all data values. All data point values are one of these predefined types. The predefined type effectively defines the type of the data point. Programmatically the data point type is defined in the AttributeType enumeration class (elt-common). See Appendix A for table of types relevant for data point values and their mappings to specific programming language types. The data point type is stored in the data point metadata in the genType field. Data point type cannot change during the lifetime of a data point. Data point with specific URI must be deleted and created again if the type needs to be changed.

```
AttributeType type = metadata.getGenType();
```



Data Point Formula

Data point formula is a mathematical expression by which the value of the data point can be calculated. If data point defines a formula, it is a calculated data point. Calculated data points cannot be directly written to (trying to do so will result in the `CiiOldbIllegalOperationException`). They are handled by the Calculation Service. The data point formula is defined in the data point metadata.

```
String formula = metadata.getFormula();
```

Since the Calculation Service uses the Octave language, the data point formula must be a valid Octave expression [6] and must return a value of the data point type. With that satisfied, the formula can contain all Octave language functions, Calculation Service build-in functions and user defined functions.

See [6] for definitions of Octave language functions.

Calculation Service defines its own built-in functions for accessing the data point's current and previous values and previous and current timestamp:

`PREV_VAL()`

`CURR_VAL()`

`PREV_TIMESTAMP()`

`CURR_TIMESTAMP()`

See Section 5.3 for how to define custom Octave functions to use in the data point formulas.

Data point formula can contain references to other data points. This makes the data point dependent on other data points. The dependent data point value will be recalculated when one of its dependencies change. A valid reference to another data point is the data point URI enclosed in curly brackets with a preceding dollar sign (e.g. `#{...}`). The engine accepts either the full URI or only a relative path of the data point.

Example:

```
#{/root/child/device/dp1} + 3 * sqrt(#{cii.oldb:/root/child/device/dp2})
```

Data point formula must not establish cyclic dependencies. Doing that will result in a `CiiOldbCyclicDependencyException` at the data point creation or metadata update.

Any syntax errors in the data point formula will result in data point quality to be set to BAD by the Calculation Service.

Note: Octave numbers are by definition float values so if a formula of a data point is a constant number and the user expects a different type it should be converted by an Octave conversion function (e.g. `int32(100)`).



Data Point Quality Expression

The quality expression is an expression which defines the data point quality. The data point quality represents the quality state of the data point and can have values OK, SUSPECT or BAD.

Data point quality evaluation is done by the Calculation Service (Section 3.2.3) and must be a valid Octave expression. The data point quality expression must not return any value but must set the quality of the data point by using the Calculation Service built-in function SETQUAL(quality) with the parameter either OK, SUSPECT or BAD.

The data point quality expression can contain Octave language functions, Calculation Service built-in functions listed above (with the addition of SETQUAL function) and custom defined functions (Section 5.3).

Example:

```
if (CURR_VAL() > 0) SETQUAL(OK) else SETQUAL(BAD) endif
```

The data point quality expression must not include references to any other data points. Any syntax error in the quality expression will result in data point quality to be set to BAD by the Calculation Service.



7.7 OLDB CLI

This section describes how to use OLDB command line application. The oldb-cli program can be used for read, write and subscribe operations and with data points that have textual outputs. The output of the tool can be parsed and used in other scripts.

7.7.1 Read operations

The tool reads the current value of the data point at the given data point URI and displays it on the standard output. Depending on the options used it can also display data point's quality expression and formula as well. Arrays and matrices are displayed in the standard JSON format. It is also possible to write the value to a file (Section 7.1.3).

The syntax for read operations is:

```
oldb-cli read <data point URI> [-f <output file>|-e|-vf|-time|-quality|-value]
```

The -h argument prints help description to standard output.

Standard Usage

To read a value of a data point, use the client together with an URI that points to the desired data point. If the data type of value is BINARY or its string representation exceeds 500 characters, the value won't be displayed on the standard output and you will have to use the output file option to retrieve it (Section 7.1.3), otherwise an error will be thrown. To output just a single part of the (timestamp/quality/value) group, the options can be used (-time, -quality, -value).

Example:

```
$ oldb-cli read cii.oldb:///glob/root/child/testintdp  
Timestamp: 2020-01-13T14:49:54.702Z  
Quality: OK  
Value: 13
```

Retrieving Value Formula and Quality Expression

To also retrieve value formula or quality expression of a data point use the options -vf (for value formula) and/or -e (for quality expression).

Example:

```
$ oldb-cli read cii.oldb:///glob/root/child/testintdp -vf -e  
Quality expression: if (CURR_VAL() > 0) SETQUAL(OK) else SETQUAL(BAD) endif  
Value formula: ${/root/child/device/DP1} + 3 * ${/root/child/device/DP2}  
Timestamp: 2020-01-13T14:49:54.702Z
```

(continues on next page)



(continued from previous page)

```
Quality: OK  
Value: 13
```

Saving value to a File

To save the data point's value to a file use the `-f` option, followed by a path to file to write to. Unless the datatype of value is BINARY or its string representation exceeds 500 characters, the value will still get displayed on standard output as well.

Example:

Saving to file:

```
$ oldb-cli read cii.oldb:/glob/root/child/teststringdp -f testStringDpValue  
Timestamp: 2020-01-07T19:09:13.117Z  
Quality: OK  
Value: ABCDEF
```

Reading the value from file:

```
$ cat testStringDpValue  
ABCDEF
```

7.7.2 Write operation

The tool gets the data point at the given URI and writes a new value to it. Data point's quality, timestamp, formula and quality expression can all be changed through the tool with the right options. Arrays and matrices must be given in the standard JSON format. Values can be given through a file too.

The syntax for writing is:

```
oldb-cli write <data point uri> <value> [-f] [-h] [-q <quality>] [-e <quality_↵  
↵expression>] [-t <timestamp>] [-vf]
```

The `-h` argument prints help description to standard output.



Standard Usage

To write a value to data point, use the client together with an URI that points to the desired data point and value to be written. If the datatype of the value to be written is BINARY, the value must be given through the -f option (Section 7.2.5).

Example:

```
$ oldb-cli write cii.oldb:/glob/root/child/testintdp 1234  
Timestamp: 2020-01-14T15:34:52.392Z  
Quality: OK  
Value: 1234
```

Setting Quality and Timestamp

When setting a data point's value, the quality and timestamp can also be set.

To set the quality use the -q option followed by the desired quality (OK/SUSPECT/BAD).

Example:

```
$ oldb-cli write cii.oldb:/glob/root/child/testintdp 1234 -q BAD
```

```
Timestamp: 2020-01-14T15:34:52.392Z
```

```
Quality: OK  
Value: 1234
```

To set a timestamp to the data point use the -t option, followed by the timestamp in ISO 8601 extended date-time format.

Example:

```
$ oldb-cli write cii.oldb:/glob/root/child/testintdp 1234 -t 2020-01-07T19:09:13.  
↪117Z  
Timestamp: 2020-01-07T19:09:13.117Z  
Quality: OK  
Value: 1234
```

Note: all three data point properties (value, quality, timestamp) can be set in one operation using all the above options.



Setting Formula

If you want to set a formula to a data point, replace the value argument with the value formula, and use the -vf option. The functionality will first read the metadata instance name that belongs to the given data point and then writes the formula to the metadata instance.

Note: Formula is a part of the data point metadata instance. As one metadata instance can belong to many datapoints, this change can affect multiple data points.

Example:

```
olddb-cli write cii.olddb:/glob/root/child/testintdp3 '${cii.olddb:/glob/root/child/↵  
↵testintdp2} + 3 * sqrt(${cii.olddb:/glob/root/child/testintdp})' -vf
```

Setting Quality Expression

If instead of setting a fixed quality, you want to set a quality expression, you can do so by using the -e option followed by the desired quality expression. In this case, the value argument can be omitted.

Example:

```
olddb-cli write cii.olddb:/glob/root/child/testintdp2 -e 'if (CURR_VAL() > 0) ↵  
↵SETQUAL(OK) else SETQUAL(BAD) endif'
```

Writing Value from a File

To write a value from a file, replace the value argument with a path to the file containing the value and use the -f option.

Example:

Writing value from a file:

```
$ oldb-cli write cii.olddb:/glob/root/child/testintdp /home/esodev/values/↵  
↵testIntDpValue -f  
Timestamp: 2020-01-14T15:40:34.783Z  
Quality: OK  
Value: 25
```



Writing Value directly from the formula.

Value to specific data point can be written directly from the expression formula. To use the formula, the value must start with the “=” character.

Example:

Writing values from a formula:

```
$ oldb-cli write cii.oldb:/glob/root/child/testdoubledp '=e(3) '
Timestamp: 2020-07-17T08:48:22.399Z
Quality: OK
Value: 20.085536923187668

$ oldb-cli write cii.oldb:/glob/root/child/testintdp '=${cii.oldb:/glob/root/
↪child/testintdp2} + ${cii.oldb:/glob/root/child/testintdp3}'

Timestamp: 2020-07-17T08:50:00.445Z
Quality: OK
Value: 4435
```

7.7.3 Subscribe operations

The tool retrieves the data point at the given URI, subscribes to it and displays value changes to the standard output until user interrupts it or the data point is deleted. Arrays and matrices are displayed in the standard JSON format.

The syntax for subscribe functionality is:

```
$ oldb-cli subscribe <data point uri> [-f <output file prefix>] [-h]
```

Standard Usage

To subscribe to a data point, use the client together with an URI that points to the desired data point. If the data type of value is BINARY or its string representation exceeds 500 characters, the value won't be displayed on the standard output and you will have to use the output file prefix option (Section 7.3.2) to have changes written to files.

Example:

```
$ oldb-cli subscribe cii.oldb:/glob/root/child/testintdp
Subscribed to cii.oldb:/glob/root/child/testintdp
To stop the subscription interrupt the program with Ctrl-C.
Timestamp: 2020-01-10T13:04:52.788Z
Quality: OK
Value: 30
```

(continues on next page)



(continued from previous page)

```
Timestamp: 2020-01-10T13:05:11.474Z  
Quality: OK  
Value: 25  
Timestamp: 2020-01-10T13:05:19.885Z  
Quality: OK  
Value: 20  
Timestamp: 2020-01-10T13:05:27.098Z  
Quality: OK  
Value: 15  
Timestamp: 2020-01-10T13:05:34.134Z  
Quality: OK  
Value: 10  
(base) [esodev@localhost values]$
```

Output to Files

To write the received values to files, use the `-f` option with a file name prefix. The files will then be placed in the current directory and named as `<prefix_timestamp>` (for each value update a new file is created).

Example:

```
$ oldb-cli subscribe cii.oldb:/glob/root/child/testintdp -f testIntDpValues  
Subscribed to cii.oldb:/glob/root/child/testintdp  
To stop the subscription interrupt the program with Ctrl-C.  
Timestamp: 2020-01-10T13:04:52.788Z  
Quality: OK  
Value: 30  
Timestamp: 2020-01-10T13:05:11.474Z  
Quality: OK  
Value: 25  
Timestamp: 2020-01-10T13:05:19.885Z  
Quality: OK  
Value: 20  
Timestamp: 2020-01-10T13:05:27.098Z  
Quality: OK  
Value: 15  
Timestamp: 2020-01-10T13:05:34.134Z  
Quality: OK  
Value: 10  
^C(base) [esodev@localhost values]$ ls  
testIntDpValues_2020-01-10T13:04:52.788Z testIntDpValues_2020-01-10T13:05:11.  
↪474Z testIntDpValues_2020-01-10T13:05:19.885Z testIntDpValues_2020-01-  
↪10T13:05:27.098Z testIntDpValues_2020-01-10T13:05:34.134Z
```

The values are stored inside the files.

Example:



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 212 of 589

```
[esodev@localhost values]$ cat testIntDpValues_2020-01-10T13:05:19.885Z  
20  
[esodev@localhost values]$
```



7.8 GUI

This section describes the components of OLDB GUI. OLDB GUI has functionalities to read, write, create, delete and update data points. The GUI can also be used to read metadata of data points, but changing of metadata is not supported. Config GUI [3] must be used to change the metadata instances, which the OLDB uses.

7.8.1 General

OLDB GUI is a graphical application that can be used for managing the data points and displaying their values. The data points are retrieved from the remote OLDB service.

After the OLDB GUI application has been successfully installed (See [2] on how to install OLDB GUI application), it can be executed by issuing the following command from the shell:

```
olddb-gui [path_to_application_config_file] [path_to_log_config_file]
```

The `olddb_gui` executable takes 2 optional arguments:

- `path_to_application_config_file`: a path to the file that contains the settings for the OLDB GUI application. An example value for the attribute: `../config/app-config.ini`. A detailed description of the application configuration file can be found in section 8.2.
- `path_to_log_config_file`: a path to the file that contains the settings for the log4cplus logging library. If it is not provided, the basic log4cplus configuration is used. Usually, the configuration file for the log4cplus library is named `log4cplus.properties`. Example value for the attribute: `../log4cplus.properties`.

7.8.2 Configuration File

The configuration file for the OLDB GUI application contains the information needed to connect to the remote configuration service (i.e. the data source). The configuration file contains 3 sections:

- **DataSource**: this section contains all the data needed to connect to the primary configuration service. This service is used for retrieving the metadata about every selected data point. The **hostname** attribute provides the address of the service (in a form of an IP address or a host name) and the **port** attribute provides a port on which the service listens for incoming connections.
- **DataSourceBackup**: this section contains the same parameters as **DataSource**, only the values are set to the ones of the backup configuration service. In case you don't have a backup configuration service, insert the values of the primary configuration service into this section.
- **MatrixEdit**: this section determines the max number of rows (`maxHeight`) and columns (`maxWidth`) that a given vector or matrix can have in order for its data to be displayed and edited directly in the OLDB GUI (for vectors, the `maxWidth` defines the size limit). Vectors and



matrices that are larger than the MatrixEdit values can only be downloaded to the local disk and uploaded from the local disk.

If “path_to_application_config_file” is not provided to the application, default values will be assigned to respective parameters:

- in DataSource section, the hostname attribute will get the **ciiconfservicehost** value assigned and the port attribute will get the **9116** value assigned,
- in DataSourceBackup section, the hostname attribute will get the **ciiconfservicehost2** value assigned and the port attribute will get the **9116** value assigned,
- in MatrixEdit section, the maxWidth and MaxHeight attributes will both get value of 100.

Note: If “path_to_application_config_file” is not provided, both **ciiconfservicehost** and **ciiconfservicehost2** must be set in order to run the GUI. These hostnames must be set in the /etc/hosts file.

7.8.3 Functionality

Main Window

The main window is displayed when we execute the OLDB GUI. The application screenshot is shown in Figure 8-1. As we can see from the figure, the functionality of the OLDB GUI is related to the management of data points.

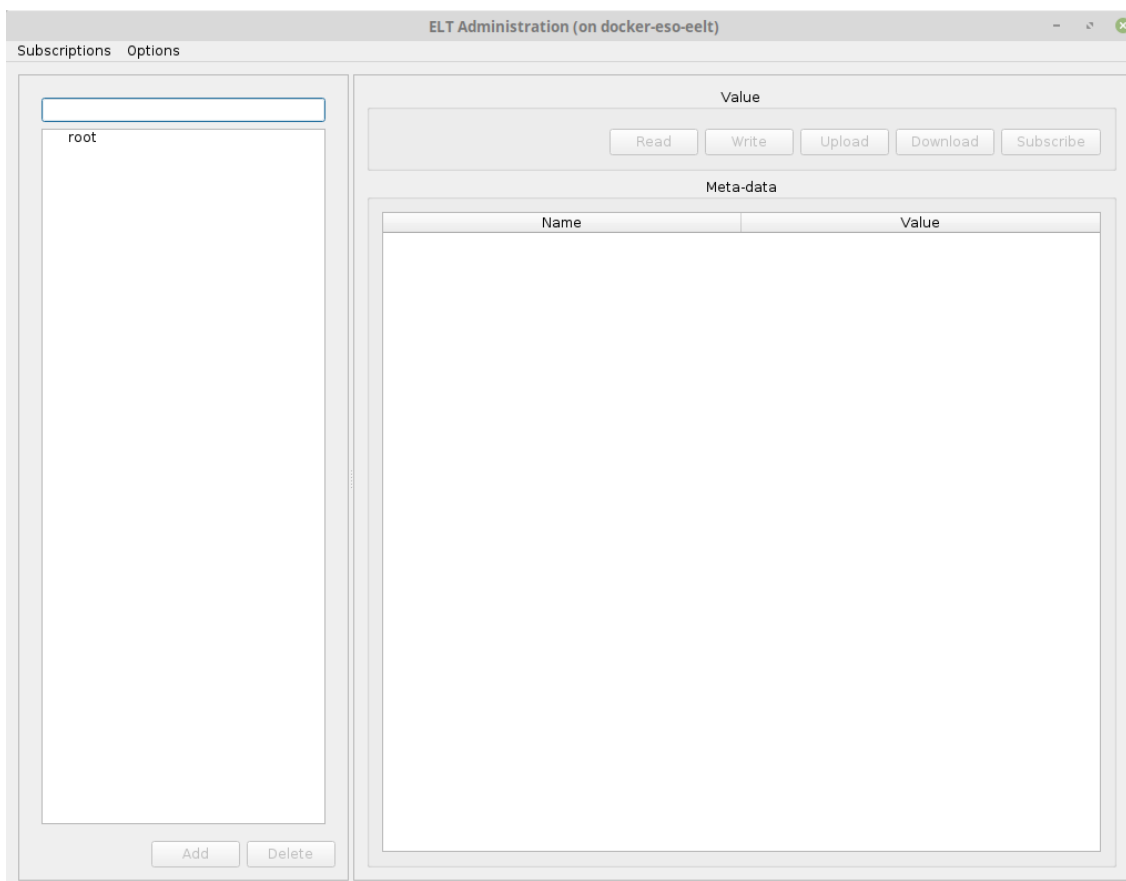


Figure 8-1: Main window

The GUI is divided into 2 parts. The left part displays the list of folders and data points retrieved from the remote service, search bar and the buttons for adding and deleting data points. The user can select the folders and data points by clicking them.

The right part is used for displaying the data about the selected data point: its value and metadata.

Menu

The OLDB GUI application only contains two menu entries, the **Subscriptions** (Figure 8-2) and **Options**. The Subscription menu entry contains the 3 items for managing data point subscriptions (i.e. the windows for continuously displaying the latest data point values):

- **Load Subscriptions:** loads the list of data point subscriptions from a file on disk (through file explorer) and opens them in windows in saved positions. All subscriptions are started (the value updates will be visible). If a subscription window is already displayed for a data point with a given URI, a new subscription window won't be displayed once more.
- **Save Subscriptions:** saves the list of currently displayed data point subscription windows (and their position on the screen) into a CSV file on disk. A file explorer window is opened and the user can enter the file name and location.



- **Close All Subscriptions:** closes all the data point subscription windows that are currently displayed (unsubscribes from the data points).

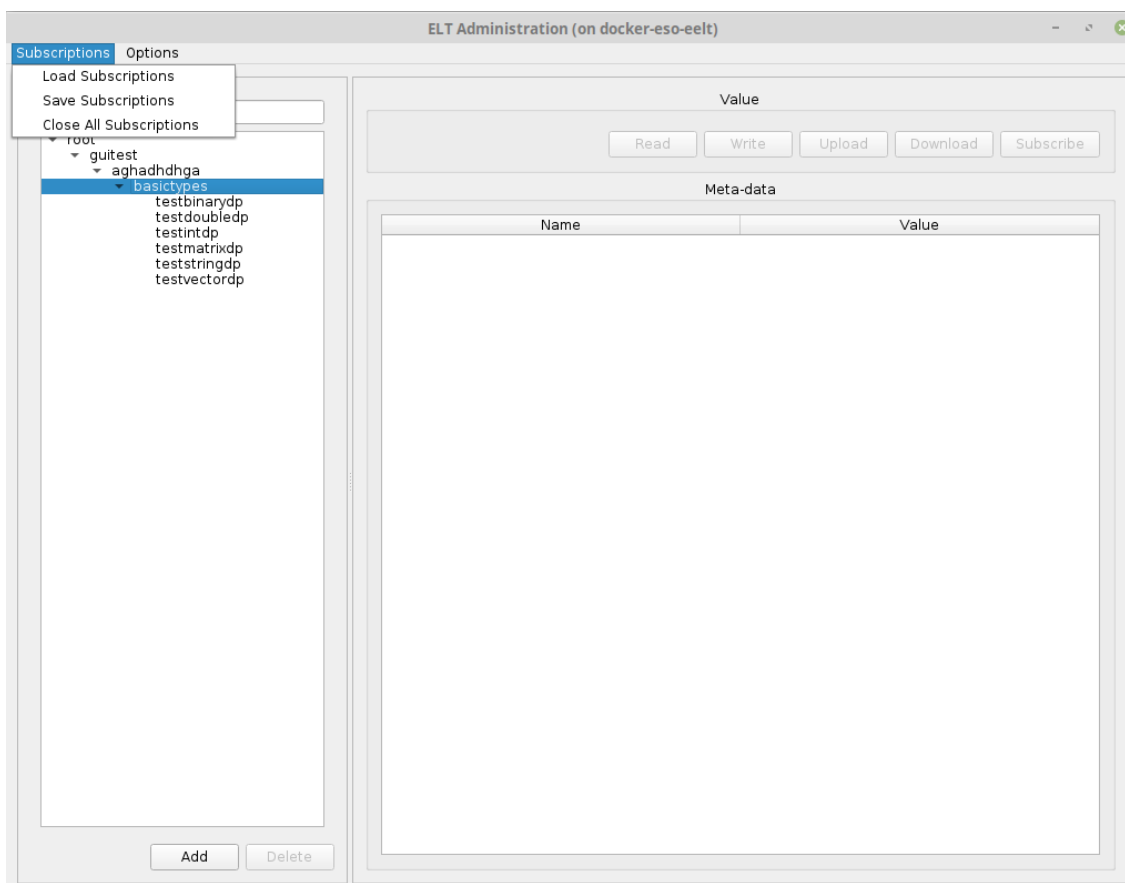


Figure 8-2: Application menu

The Options menu contains a checkbox item “write enabled”. This checkbox must be checked to enable write operations on data points. If this checkbox is not checked, writing to data points will not be possible. It is not checked by default.

7.8.4 Hierarchy view

The hierarchical view displays the hierarchical structure of folders and data points. It enables the user to do the following actions:

- Browse and view the data points that are stored in the remote service.
- Add new data points.
- Delete existing data points.
- Trigger the loading of data point values.

1. Selecting Folder



The user selects a folder by clicking the inner node in the data point hierarchy tree (Figure 8-3). If the leaf tree node is clicked, the data point is selected. The folders are created automatically, based on the data points' URIs.

When a folder is selected, all the data in the data point frame is cleared. The folders don't have any values or metadata, they are only used for organizing data points. When a folder is selected the user can create a new data point by clicking on the Add button below the hierarchy tree.




Figure 8-3: Selecting folder

2. Selecting Data Point

The user selects a data point by clicking the leaf node in the data point hierarchy tree (Figure 8-4). The data points are organized in the tree structure based on their URIs.

For the easier searching of data points, a search filter bar is located at the top of the hierarchy tree. The user can input a part of the URI and when he/she hits the Enter key, all the data points whose name doesn't contain the string input are filtered out (the string comparison is case insensitive). In order to display all the data points again, the user needs to clear the filtering text and hit the Enter key again.

When a configuration is selected, the following data is loaded in the data point frame of the GUI:

- data point value, quality and timestamp: the most recent value, quality and timestamp of the



selected data point.

- the list of data point's metadata attributes and their values.

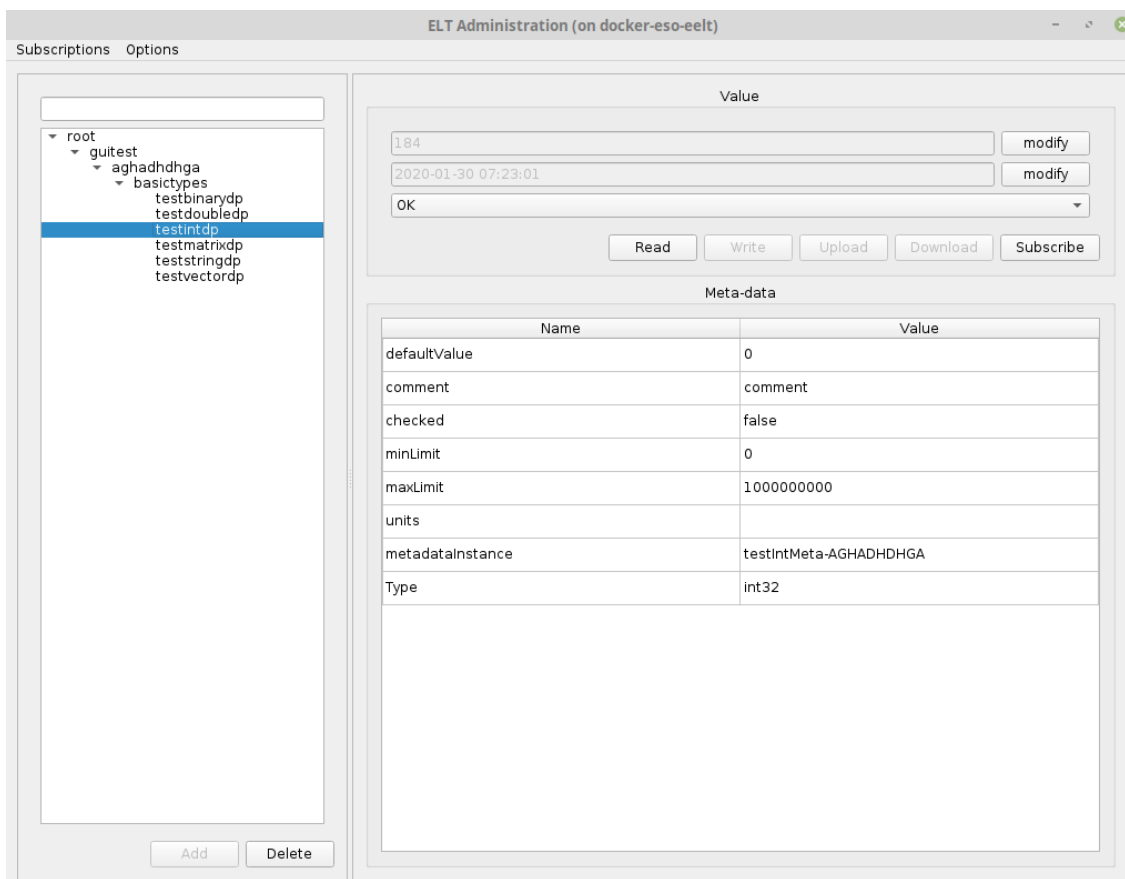


Figure 8-4: Data point selected

Note that it is not possible to change data point metadata through OLDB GUI application. The user should use the Config GUI application to change the metadata.

The data point URI is at the moment not showed anywhere in the data point frame. The user can reconstruct the data point URI from the tree path.

3. Adding a New Data Point

The adding of new data point is initiated by clicking the Add button below the hierarchy tree. The button is only enabled if a folder is selected.

When the Add button is clicked, a dialog for inserting a data point is displayed (Figure 8-5). The dialog contains the following fields:

- URI prefix: a read-only field that contains the URI prefix
- URI suffix: the input field where the user can input the URI suffix or the name of the data point. The final URI of the configuration is formed by concatenating the URI prefix and suffix.



- Metadata instance: here, the user selects the metadata instance for the newly created data point.

After the user clicks the OK button, the dialog is closed and a new data point is created on the remote server.

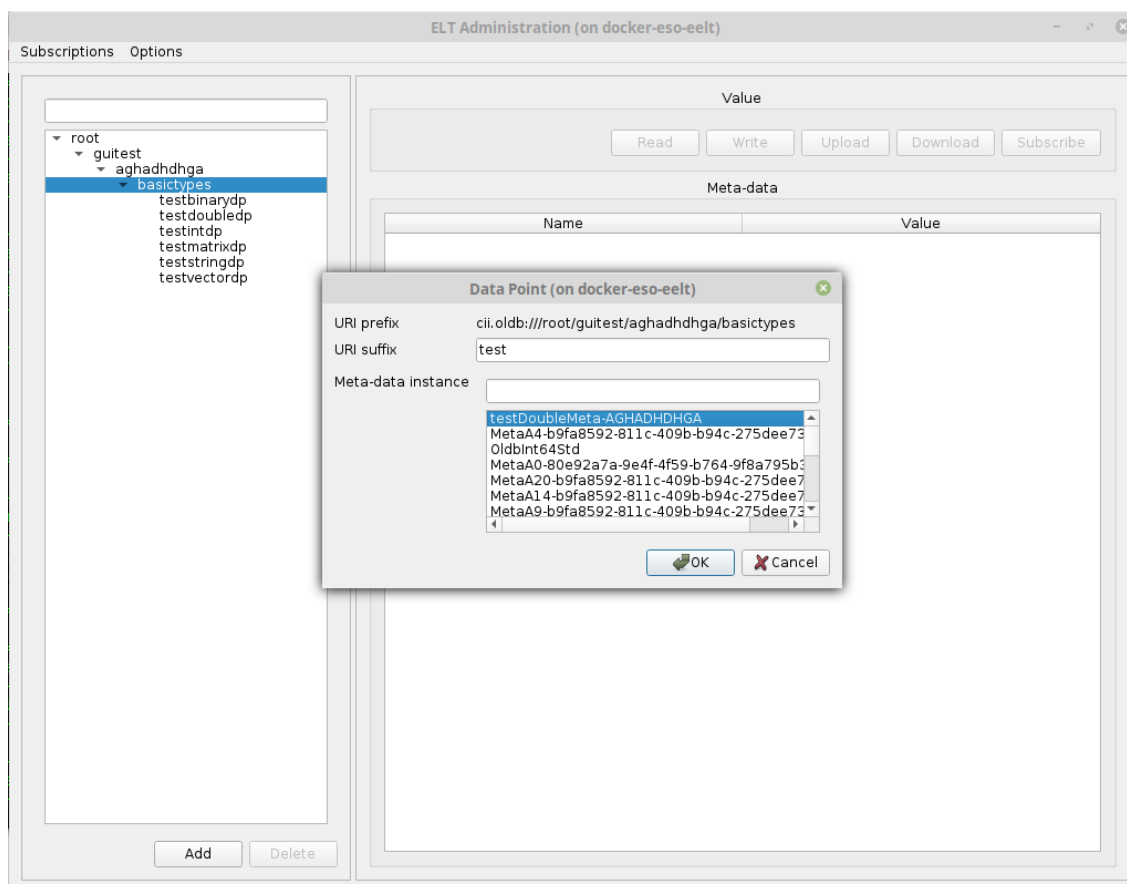


Figure 8-5: Adding new data point

4. Deleting an Existing Data Point

The deleting of existing data point is initiated by clicking the Delete button below the hierarchy tree. The button is only enabled if a data point is selected.

When the Delete button is clicked, a confirmation dialog is displayed where the user can confirm or cancel the deletion. On user confirmation, the data point is deleted on the remote server.

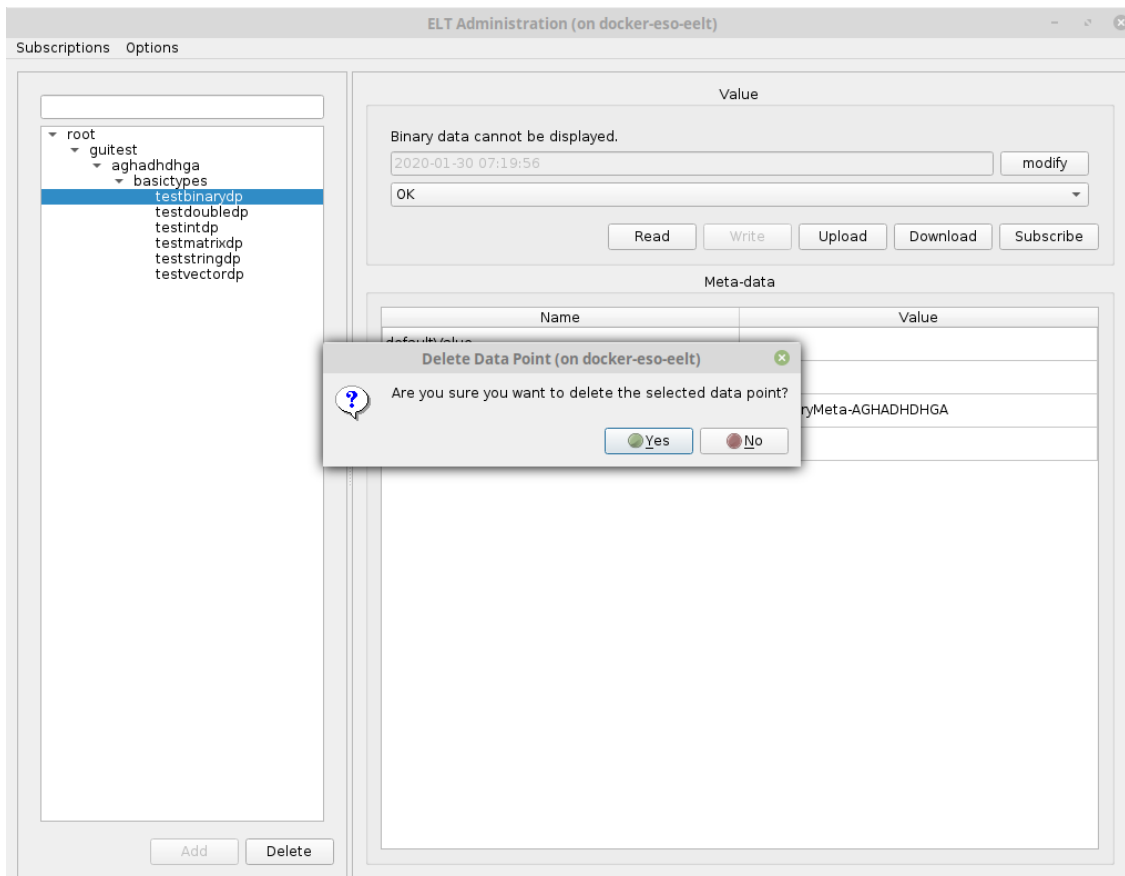


Figure 8-6: Data point deletion confirmation dialog

1. Data Point Value and Metadata Frame

The data point value and metadata frame occupy the right side of the application window. It displays the value and all the metadata values for the selected data point. It enables the user to do the following actions:

- Refresh data point values.
- Update data point values.
- Subscribe to data point value changes.

5. Manually Refreshing Data Point Value

The manual refreshing of the value for the selected data point is triggered by clicking the Read button just below the area for displaying the data point value. After doing this, a request is sent to a remote server, the most recent data point value is retrieved and displayed in the GUI.

6. Setting the Data Point Value

There are 3 groups of data point values:

- primitive values: the setting of the primitive value for the data point is performed using the input data dialog (Figure 8-7) that is displayed when clicking the modify button next to the text field



for displaying the data point value. The dialog contains a simple input field where the user can input the data value. After the value is entered and the dialog is closed (by clicking OK), the Write button must be clicked for value to be written to the OLDB storage.

- vector & matrix values: they are displayed in a tabular form (Figure 8-8). The table values are set by clicking different table cells. After doing this, the same dialog displays as with the primitive values. Here, the user inputs the value for the table cell. In order for the modified values to be sent to the remote service, the Write button needs to be clicked.
- binary data values: the binary data values don't get displayed in the OLDB GUI. It must be stored to or loaded from the local disk.

Whenever the user modifies the data point value in the GUI the timestamp field updates to "now", so that when Write button is clicked the timestamp written to the OLDB is of the last change. The user can also modify the timestamp field in the same manner than the value.

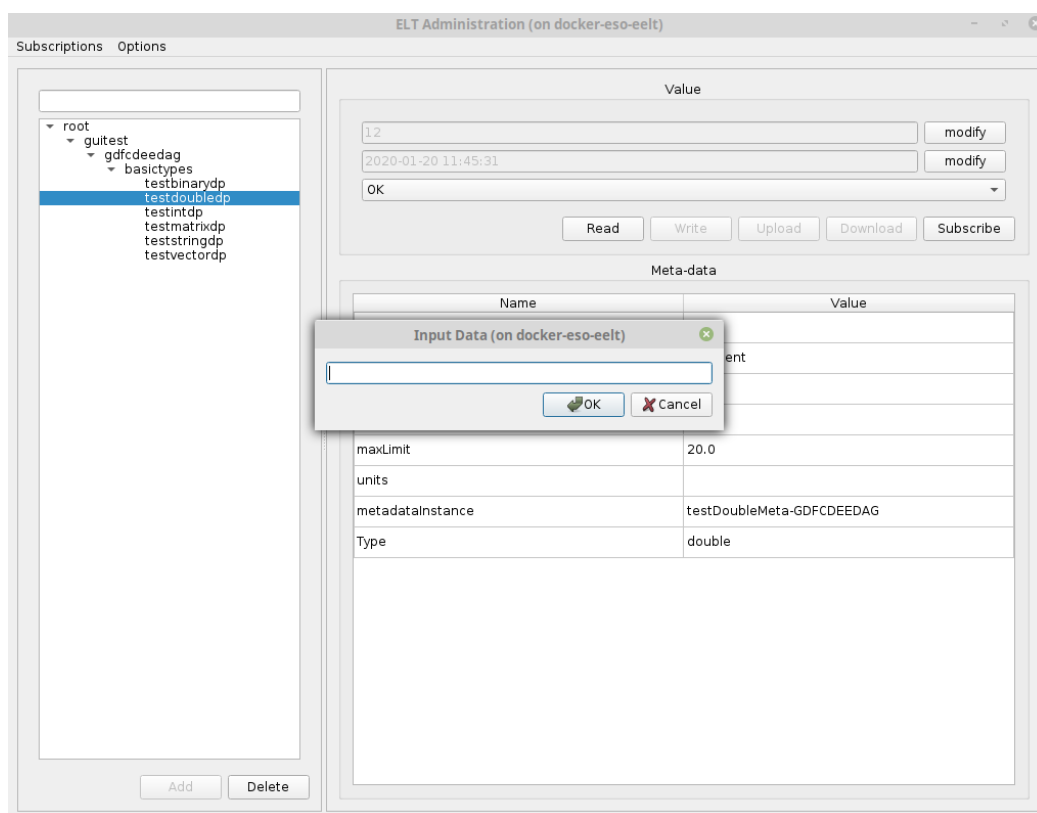


Figure 8-7: Data input dialog

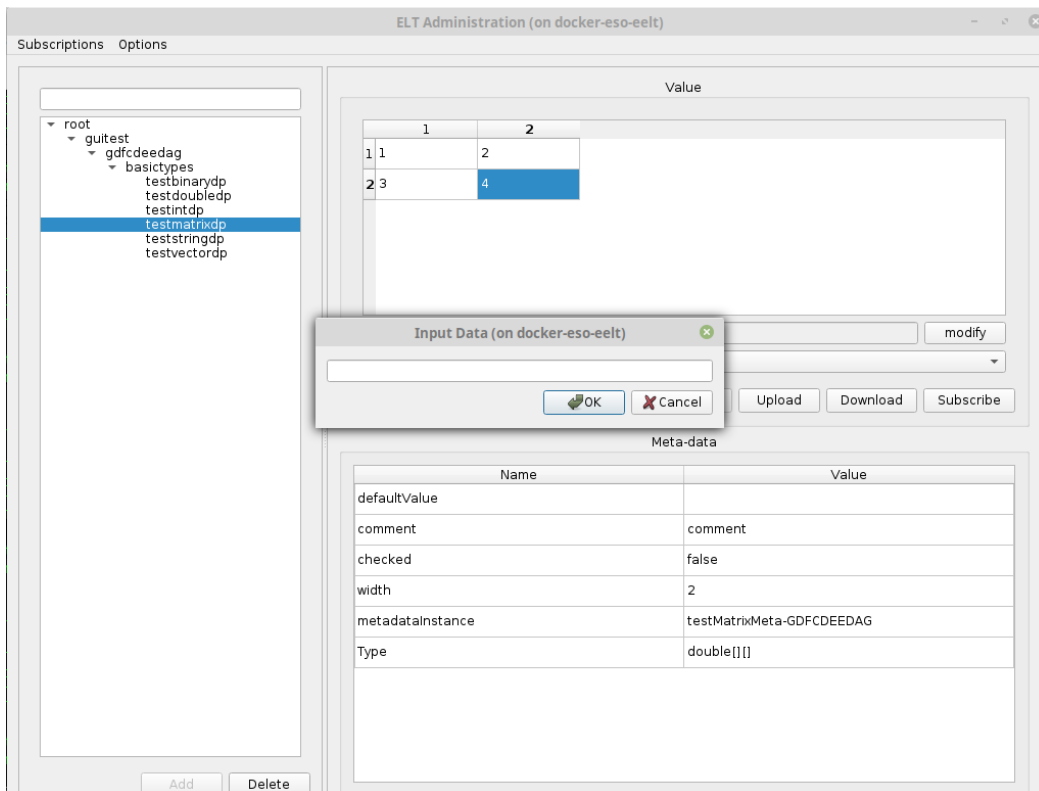


Figure 8-8: Vector & matrix data point value

The vector & matrix values as well as binary data values can also be loaded from a local disk and be stored to a local disk:

- the loading from a local disk can be started by clicking the Upload button. In the dialog that gets shown, the user selects the file on local disk, from which the data point values will be loaded. After the file has been loaded it can be sent to the remote service by clicking the Write button.
- the storing to the local disk can be started by clicking the Download button. In the dialog that gets shown, the user selects the file on local disk, to which the data point values will be stored.

For the vector & matrix data the user can also change the number of rows and columns. The row-and-column-setting dialog (Figure 8-10) is displayed by right-clicking the table and selecting the Resize table item from the context menu (Figure 8-9). If the matrix size is made smaller, any values already entered will be deleted for the removed columns or rows. If the matrix is made larger, the additional cells are empty.

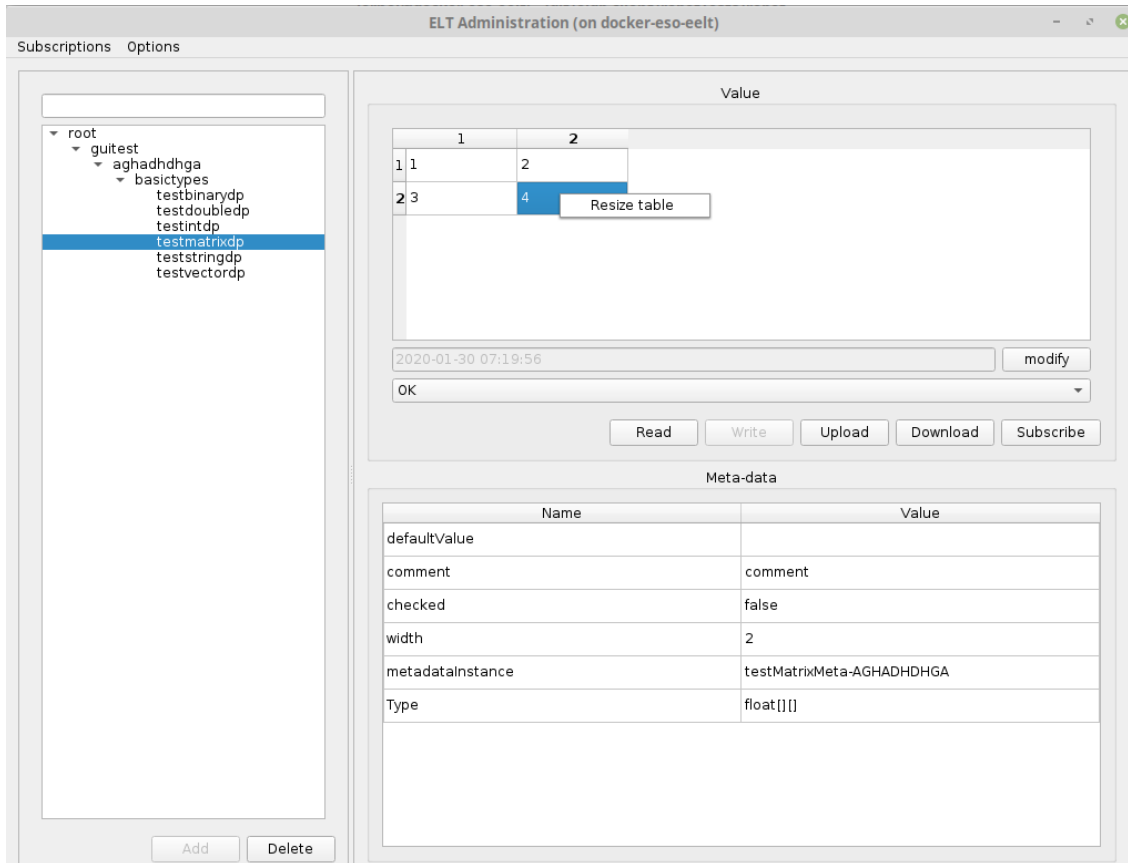


Figure 8-9: “Resize table” context menu

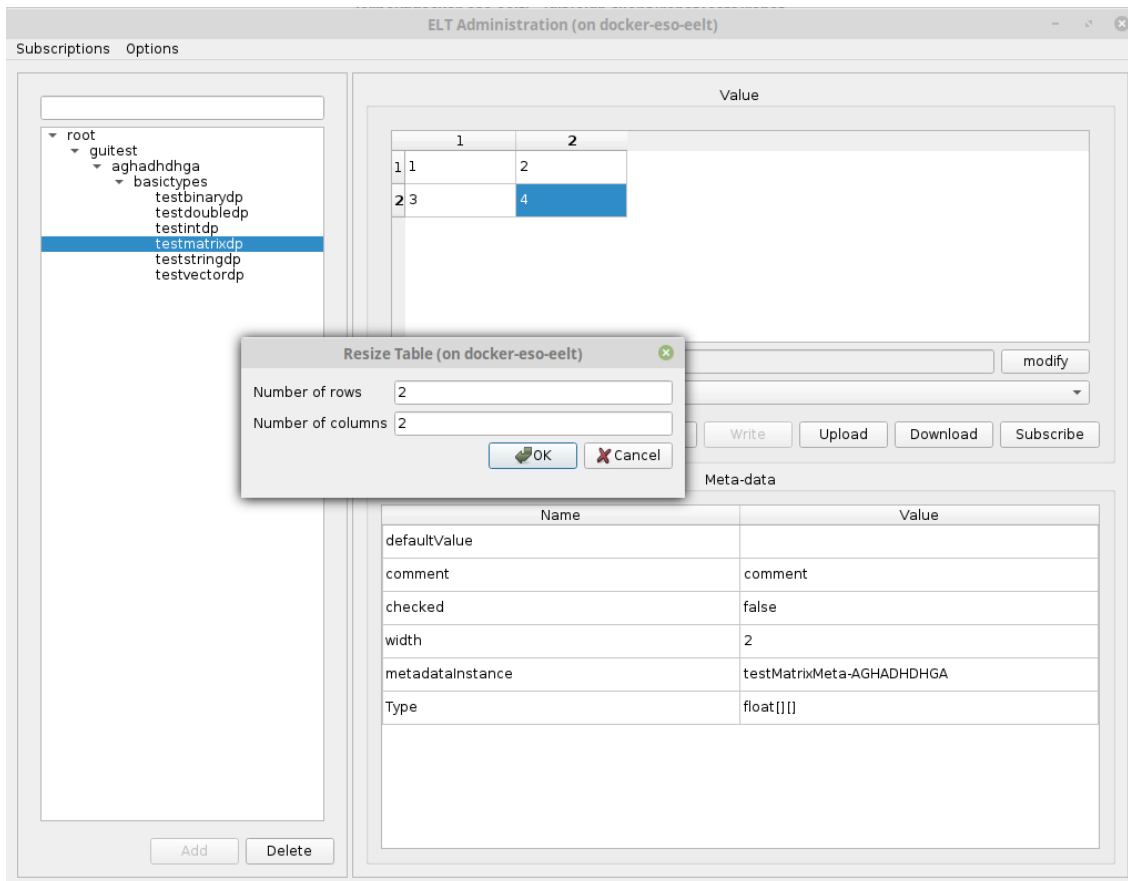


Figure 8-10: Dialog for setting the number of rows and columns.

7. Subscribing to Data Point Value Changes

Subscribing to the value changes of the selected data point is performed by clicking the Subscribe button that is located under the widget for displaying data point value. At this point the new window dialog is displayed that displays the most recent value of the given data point. Whenever the data point value changes on the server it is automatically updated in the subscription dialog. Multiple subscription dialogs can be displayed at the same time, meaning that the user can monitor the values of multiple data points. For a given data point, however, only a single subscription window can be displayed, having duplicate subscription windows for the same data point is not possible.

The user can organize the subscription windows in any way they like, they can even store the position of the subscription windows into a file. Later, the user can use the file to restore the windows. This enables the user to have a pre-defined sets of data points that they can monitor (see section 8.3.2 on how this is done).

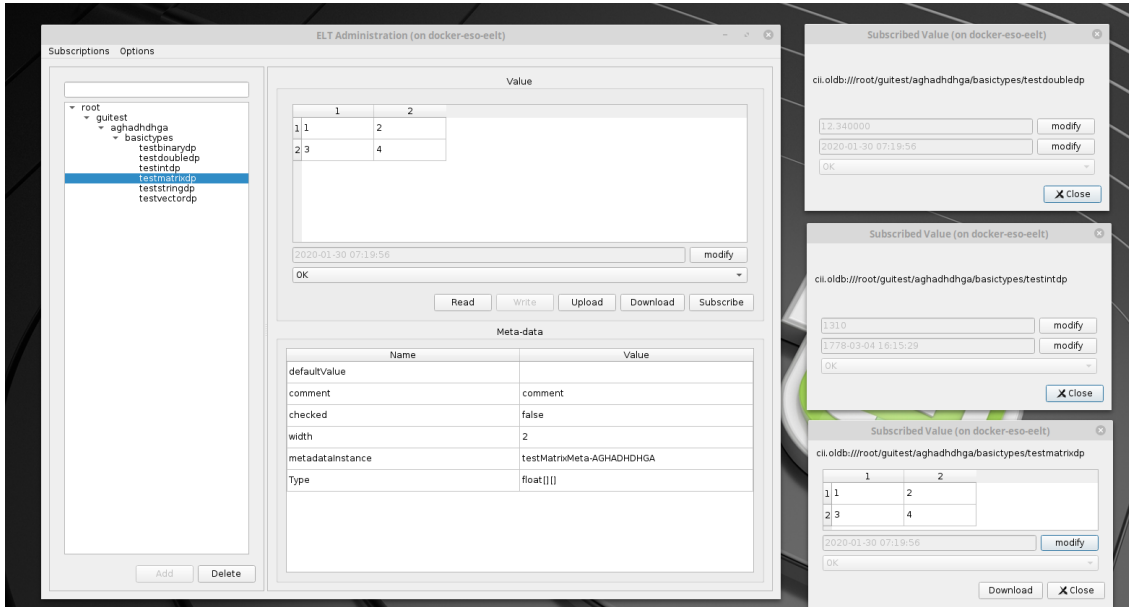


Figure 8-11: Arrangement of multiple subscription windows on the screen



7.9 Data Point Value Type to Language Types Mapping

Below is the mapping table between data point types and data types of all programming languages that the OLDB can use.

Table 8-1: Data point types mapping table

ELT CII BASIC TYPE	Java	C++	Python
INT8	Byte	std::int8_t	int
INT16	Short	std::int16_t	int
INT32	Integer	std::int32_t	int
INT64	Long	std::int64_t	int
UINT8	Byte	std::uint8_t	int
UINT16	Short	std::uint16_t	int
UINT32	Integer	std::uint32_t	int
UINT64	Long	std::uint64_t	int
SINGLE	Float	float	float
DOUBLE	Double	double	float
BOOLEAN	Boolean	bool	bool
STRING	String	std::string	str
BINARY	java.nio.ByteBuffer	std::vector<std::uint8_t>	bytes
VECTOR_INT32	List<Integer>	std::vector<std::int32_t>	list
VECTOR_SINGLE	List<Float>	std::vector<float>	list
VECTOR_DOUBLE	List<Double>	std::vector<double>	list
VECTOR_STRING	List<String>	std::vector<std::string>	List
MATRIX2D_INT32	List<Integer>	std::vector<std::int32_t>	list
MATRIX2D_SINGLE	List<Float>	std::vector<float>	list
MATRIX2D_DOUBLE	List<Double>	std::vector<double>	List



7.10 Default Metadata Instance Names

ELT CII BASIC TYPE	Metadata Instance Name	Metadata Class name
INT8	OldbInt8Std	MdOldbNumber
INT16	OldbInt16Std	MdOldbNumber
INT32	OldbInt32Std	MdOldbNumber
INT64	OldbInt64Std	MdOldbNumber
UINT8	OldbUInt8Std	MdOldbNumber
UINT16	OldbUInt16Std	MdOldbNumber
UINT32	OldbUInt32Std	MdOldbNumber
UINT64	OldbUInt64Std	MdOldbNumber
SINGLE	OldbSingleStd	MdOldbNumber
DOUBLE	OldbDoubleStd	MdOldbNumber
STRING	OldbStringStd	MdOldbString
BINARY	OldbBinaryStd	MdOldbBinary
VECTOR_INT32	OldbInt32ArrayStd	MdOldbArray
VECTOR_SINGLE	OldbFloatArrayStd	MdOldbArray
VECTOR_DOUBLE	OldbDoubleArrayStd	MdOldbArray
VECTOR_STRING	OldbStringArrayStd	MdOldbArray
MATRIX2D_INT32	OldbInt32MatrixStd	MdOldbMatrix
MATRIX2D_SINGLE	OldbFloatMatrixStd	MdOldbMatrix
MATRIX2D_DOUBLE	OldbDoubleMatrixStd	MdOldbMatrix



7.11 Metadata attribute constraints

METADATA ATTRIB.	CONSTRAINT	CONSTRAINT VALUE
Comment	Max size in characters	80 chars



7.12 OLDB API LISTING

This a listing of OLDB API methods with short description for both CiiOldb and CiiOldbDataPoint classes. The method signatures are written in Java. The signatures are equal (with some exceptions in some of createDataPoint methods) in all three languages except for the native data types of each language (e.g. for string, double, integer). For data type mapping between languages see Table 8-1. For URI variables java uses java.net.URI, C++ uses <mal/utility/Uri.hpp> and Python uses elt.config.Uri.

Note that all OLDB API methods are synchronous (blocking) methods.

For detail explanation of exceptions thrown by these methods, consult the oldb-client source code.

7.12.1 CiiOldb

<code>static CiiOldb getInstance()</code>
The method returns the singleton instance implementation of the OLDB client CiiOldb.

<code>CiiOldbDataPoint<?> createDataPoint(URI uri, String metadataInstanceName, String metadataClassName)</code>
The method creates a data point at the defined URI with defined metadata. The data point data type and other restrictions are defined by its metadata which must already exist. Once the data point is created its data type cannot be changed. Note: This method signature is different in C++ and Python where the String metadataClassName parameter is omitted. This is because of Java language specifics.

<code>CiiOldbDataPoint<String> createDataPointByValue(URI uri, String value)</code>
Creates The method creates a data point of type STRING at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "OldbStringStd". The metadata must already exist.

<code>CiiOldbDataPoint<Byte> createDataPointByValue(URI uri, Byte value)</code>
Creates The method creates a data point of type INT8 at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "OldbInt32Std". The metadata must already exist.

<code>CiiOldbDataPoint<Short> createDataPointByValue(URI uri, Short value)</code>
Creates The method creates a data point of type INT16 at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "OldbInt16Std". The metadata must already exist.



CiiOldbDataPoint<Integer> createDataPointByValue(Uri uri, Integer value)

Creates The method creates a data point of type INT32 at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "OldbInt32Std". The metadata must already exist.

CiiOldbDataPoint<Float> createDataPointByValue(Uri uri, Float value)

Creates The method creates a data point of type SINGLE at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "OldbSingleStd". The metadata must already exist.

CiiOldbDataPoint<Double> createDataPointByValue(Uri uri, Short value)

Creates The method creates a data point of type DOUBLE at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "OldbDoubleStd". The metadata must already exist.

CiiOldbDataPoint<ByteBuffer> createDataPointByValue(Uri uri, ByteBuffer value)

Creates The method creates a data point of type BINARY at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "OldbBinaryStd". The metadata must already exist.

CiiOldbDataPoint<List<T>> createDataPointByValue(Uri uri, List<T> value, Class<T> clazz, boolean isMatrix)

Creates The method creates a data point of matrix or vector type at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "Oldb<element-type>ArrayStd". The metadata must already exist. The <element-type> is "Int32", "Single", "Double" or "String" (vector only). The isMatrix parameter defines whether a matrix or vector data point is created. The clazz parameter defines the class of the elements of matrix/vector.

Note: The Class<T> clazz parameter is only needed in Java version of the method because of language specifics.

CiiOldbDataPoint<List<T>> createDataPointByValue(Uri uri, List<T> value, Class<T> clazz)

Creates The method creates a data point of vector type at the defined URI with defined value. It uses the default OLDB string metadata with metadata instance name "MdO<element-type>ArrayStd". The metadata must already exist. The <element-type> is "Int32", "Single", "Double" or "String". The clazz parameter defines the class of the elements of vector.

Note: The Class<T> clazz parameter is only needed in Java version of the method because of language specifics.

void deleteDataPoint(Uri uri)

The method deletes a data point referenced by the URI.



Map<URI, Boolean> getChildren(Uri uri)

Finds all direct children of the given URI. Returns the map of all found URI with boolean value indicating whether the URI represents a data point (true if data point, false if only folder). Since OLDB allows to create a data point with an URI that is a sub-path of another URI, there can be both folder and data point URIs present. In that case the boolean value is true.

Note: Because OLDB is case insensitive all URIs returned will be in lower case.

Example

Consider the following OLDB URIs:

```
cii.olddb:/root/child/device/dp1  
cii.olddb:/root/child/dp2  
cii.olddb:/root/child/dp2/dp3
```

A call to the getChildren method with URI = cii.olddb:/root/child will return the map in Table 8-2.

Table 8-2: Example of getChildren call result

Key	Value
cii.olddb:/root/child/device	FALSE
cii.olddb:/root/child/dp2	TRUE

Note that the path cii.olddb:/root/child/dp2 hierarchically represents both data point and a folder. In this case the getChildren returns TRUE for this path. This means that a TRUE for an URI does not mean that the URI doesn't also represent a folder.

CiiOldbDataPoint<?> getDataPoint(Uri uri)

Returns a data point at a specified URI.

Note that the caller of the getDataPoint method must know the data point type to cast the received object to the right type (only required for the strongly typed languages like Java and C++). The type of the data point can be obtained from metadata (Section 6.2.3).

Note: this method does not provide transactional safety. There is no guarantee that the data point is not deleted the moment after it was retrieved (data point object of a deleted data point is not usable).



List<CiiOldbDataPoint<?>> getDataPoints(List<URI> uris)

The method returns a list of data points matching the specified URIs.

The URIs in the list can be written as a simple glob pattern using three special characters.

* for any number of characters restricted to the URI hierarchy level

** for any number of characters spanning multiple levels

^ for exactly one character

Example:

If we have a simple OLDB that contains data points with the following URIs:

cii.oidb:/root/child/device/dp1

cii.oidb:/ root/child/device/dp21

cii.oidb:/ root/child/device/dp23

cii.oidb:/ root/child1/motor/dp2

Then calls to getDataPoints with these glob patterns would result in following data points:

cii.oidb:/root/child/device/dp2* -> dp21, dp23

cii.oidb:/root/child/** -> dp1, dp21, dp23

cii.oidb:/root/child/device/dp^ -> dp1

Note: this method does not provide transactional safety. There is no guarantee that a data point is not deleted the moment after it was retrieved (data point object of a deleted data point is not usable).

List<CiiOldbDataPoint<T>> getDataPoints(List<URI> uris, AttributeType dataType, T minValue, T maxValue)

The method returns a list of data points matching the specified URIs, filtered by type and min and max limits. As with the getDataPoints above, this method supports the simple glob patterns.

Note: this method does not provide transactional safety. There is no guarantee that a data point is not deleted the moment after it was retrieved (data point object of a deleted data point is not usable).

void setWriteEnabled(boolean writeEnabled)

Enables or disables all writes to data points.

void close();

Closes connections and cleans up resources.

Below are the Java and C++ CiiOldb API listings. The Python client uses Python bindings and has therefore the same API as the C++ client.

Listing 8-1: Java CiiOldb API

```
static CiiOldb getInstance() throws CiiOldbInitializationError {  
    return CiiOldbImp.getInstance();  
}
```

(continues on next page)



(continued from previous page)

```
CiiOldbDataPoint<?> createDataPoint(Uri uri, String metadataInstanceName, String  
↳ metadataClassName) throws CiiOldbDpExistsException, ↳  
↳ CiiOldbCyclicDependencyException, CiiOldbExpressionException,  
CiiOldbDpConstraintsException, CiiOldbUndefinedMetadataException, ↳  
↳ CiiOldbStorageError, CiiOldbConfigError, CiiOldbUriException;
```

```
CiiOldbDataPoint<String> createDataPointByValue(Uri uri, String value) throws ↳  
↳ CiiOldbDpExistsException, CiiOldbCyclicDependencyException, ↳  
↳ CiiOldbExpressionException, CiiOldbDpConstraintsException,  
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError, ↳  
↳ CiiOldbUriException;
```

```
CiiOldbDataPoint<Byte> createDataPointByValue(Uri uri, Byte value) throws ↳  
↳ CiiOldbDpExistsException, CiiOldbCyclicDependencyException, ↳  
↳ CiiOldbExpressionException, CiiOldbDpConstraintsException,  
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError, ↳  
↳ CiiOldbUriException;
```

```
CiiOldbDataPoint<Short> createDataPointByValue(Uri uri, Short value) throws ↳  
↳ CiiOldbDpExistsException, CiiOldbCyclicDependencyException, ↳  
↳ CiiOldbExpressionException, CiiOldbDpConstraintsException,  
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError, ↳  
↳ CiiOldbUriException;
```

```
CiiOldbDataPoint<Integer> createDataPointByValue(Uri uri, Integer value) throws ↳  
↳ CiiOldbDpExistsException, CiiOldbCyclicDependencyException, ↳  
↳ CiiOldbExpressionException, CiiOldbDpConstraintsException,  
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError, ↳  
↳ CiiOldbUriException;
```

```
CiiOldbDataPoint<Long> createDataPointByValue(Uri uri, Long value) throws ↳  
↳ CiiOldbDpExistsException, CiiOldbCyclicDependencyException, ↳  
↳ CiiOldbExpressionException, CiiOldbDpConstraintsException,  
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError, ↳  
↳ CiiOldbUriException;
```

```
CiiOldbDataPoint<Float> createDataPointByValue(Uri uri, Float value) throws ↳  
↳ CiiOldbDpExistsException, CiiOldbCyclicDependencyException, ↳  
↳ CiiOldbExpressionException, CiiOldbDpConstraintsException,  
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError, ↳  
↳ CiiOldbUriException;
```

```
CiiOldbDataPoint<Double> createDataPointByValue(Uri uri, Double value) throws ↳  
↳ CiiOldbDpExistsException, CiiOldbCyclicDependencyException, ↳  
↳ CiiOldbExpressionException, CiiOldbDpConstraintsException,  
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError, ↳  
↳ CiiOldbUriException;
```

(continues on next page)



(continued from previous page)

```
CiiOldbDataPoint<ByteBuffer> createDataPointByValue (URI uri, ByteBuffer value)
↳throws CiiOldbDpExistsException, CiiOldbCyclicDependencyException,
↳CiiOldbExpressionException, CiiOldbDpConstraintsException,
CiiOldbUndefinedMetadataException, CiiOldbStorageError, CiiOldbConfigError,
↳CiiOldbURIException;

<T> CiiOldbDataPoint<List<T>> createDataPointByValue (URI uri, List<T> value,
↳Class<T> clazz, boolean isMatrix) throws CiiOldbDpExistsException,
↳CiiOldbCyclicDependencyException, CiiOldbExpressionException,
CiiOldbDpConstraintsException, CiiOldbUndefinedMetadataException,
↳CiiOldbStorageError, CiiOldbConfigError, CiiOldbURIException,
↳CiiInvalidTypeException;

<T> CiiOldbDataPoint<List<T>> createDataPointByValue (URI uri, List<T> value,
↳Class<T> clazz) throws CiiOldbDpExistsException,
↳CiiOldbCyclicDependencyException, CiiOldbExpressionException,
CiiOldbDpConstraintsException, CiiOldbUndefinedMetadataException,
↳CiiOldbStorageError, CiiOldbConfigError, CiiOldbURIException,
↳CiiInvalidTypeException;

void deleteDataPoint (URI uri) throws CiiOldbStorageError, CiiOldbURIException,
↳CiiOldbConfigError;

Map<URI, Boolean> getChildren (URI uri) throws CiiOldbDpUndefinedException,
↳CiiOldbURIException;

CiiOldbDataPoint<?> getDataPoint (URI uri) throws CiiOldbURIException,
↳CiiOldbDpUndefinedException, CiiOldbUndefinedMetadataException,
↳CiiOldbStorageError;

List<CiiOldbDataPoint<?>> getDataPoints (List<URI> uris)
    throws CiiOldbUndefinedMetadataException;

<T> List<CiiOldbDataPoint<T>> getDataPoints (List<URI> uris, AttributeType
↳dataType, T minValue, T maxValue) throws CiiOldbDpConstraintsException,
↳CiiOldbDpUndefinedException, CiiOldbDpQualityException,
CiiOldbUndefinedMetadataException;

void setWriteEnabled (boolean writeEnabled);

void close();
```

Listing 8-2: C++ CiiOldb API

```
std::shared_ptr<CiiOldbTypedDataBase> CreateDataPoint (const ::elt::mal::Uri& uri,
↳ const std::string& meta_data_instance_name);

template<typename T>
    std::shared_ptr<CiiOldbDataPoint<T>> CreateDataPointByValue (
```

(continues on next page)



(continued from previous page)

```
const ::elt::mal::Uri& uri, const T& value);

template<typename T>
std::shared_ptr<CiiOldbDataPoint<std::vector<T>>> CreateDataPointByValue (
const ::elt::mal::Uri& uri, const std::vector<T>& value, bool is_matrix);

std::shared_ptr<CiiOldbTypedDataBase> GetDataPoint(const ::elt::mal::Uri& uri)
↳const;

template<typename T> std::shared_ptr<CiiOldbDataPoint<T>> GetDataPoint (
const ::elt::mal::Uri& uri, ::elt::common::CiiBasicDataType data_type) const;

std::vector<std::shared_ptr<CiiOldbTypedDataBase>> GetDataPoints (
const std::vector<::elt::mal::Uri>& uris) const;

template<typename T> std::vector<std::shared_ptr<CiiOldbDataPoint<T>>>
↳GetDataPoints(const std::vector<::elt::mal::Uri>& uris,
::elt::common::CiiBasicDataType data_type, T min_value, T max_value) const;

std::map<::elt::mal::Uri, bool> GetChildren(const ::elt::mal::Uri& uri) const;

static void SetWriteEnabled(bool enabled) noexcept;

static bool IsWriteEnabled() noexcept;

void DeleteDataPoint(const ::elt::mal::Uri& uri);
```

7.12.2 CiiOldbDataPoint

CiiOldbDpValue<T> readValue()

Returns a copy of this Data Point value.

Every call to readValue retrieves an up to date value from the OLDB, but the CiiOldbDpValue object is a snapshot at the time of retrieval and does not update automatically. The user should call readValue to get the updated value.

void writeValue(T value)

Writes the value to the data point. The quality is unchanged and timestamp is now.

writeValue(T value, Instant timestamp)

Writes the value and the timestamp to the data point. The quality is unchanged.

void writeValue(T value, Instant timestamp, CiiOldbDpQuality quality)

Writes all three fields of the data point value to the data point.



<code>void subscribe(CiiOldbDpSubscription callback)</code>
Subscribes a callback to the data point to listen to value changes.

<code>void unsubscribe(CiiOldbDpSubscription callback)</code>
Unsubscribes a callback from the data point.

<code>void setQuality(CiiOldbDpQuality quality)</code>
Sets the data point quality.

<code><U extends MdOldb<T>> U getMetadata(Class<U> clazz)</code>
Returns the metadata of the data point. Caller must provide the metadata class (extends MdOldb).

<code><U extends MdOldb> void setMetadata(String metadataInstanceName)</code>
Updates the data point metadata with the metadata which instance name is provided. The changes made in the metadata must be saved to config service otherwise this call will have no effect. If the instance name is different from the one this data point already has (new metadata instance name has been assigned to the data point), the new metadata must respect the data point type.

Below are the Java and C++ CiiOldbDataPoint API listings. The Python client uses Python bindings and has therefore the same API as the C++ client.

Listing 8-3: Java CiiOldbDatapoint API

```
public abstract CiiOldbDpValue<T> readValue() throws CiiOldbDpUndefinedException,  
↳ CiiOldbDeletedException, CiiOldbStorageError;  
  
public abstract void writeValue(T value) throws CiiOldbDpQualityException, ↳  
↳ CiiOldbDpConstraintsException, CiiOldbWriteDisabledException, ↳  
↳ CiiOldbDeletedException, CiiOldbStorageError;  
  
public abstract void writeValue(T value, Instant timestamp) throws ↳  
↳ CiiOldbDpConstraintsException, CiiOldbWriteDisabledException, ↳  
↳ CiiOldbDpQualityException, CiiOldbDeletedException, CiiOldbStorageError;  
  
public abstract void writeValue(T value, Instant timestamp, CiiOldbDpQuality ↳  
↳ quality) throws CiiOldbWriteDisabledException, ↳  
↳ CiiOldbIllegalOperationException, CiiOldbDeletedException,  
CiiOldbStorageError, CiiOldbDpConstraintsException;  
  
public abstract void subscribe(CiiOldbDpSubscription callback) throws ↳  
↳ CiiOldbDeletedException;  
  
public abstract void unsubscribe(CiiOldbDpSubscription callback) throws ↳  
↳ CiiOldbDeletedException;
```

(continues on next page)



(continued from previous page)

```
public abstract void setQuality(CiiOldbDpQuality quality) throws ↳
↳CiiOldbDeletedException, CiiOldbStorageError, CiiOldbWriteDisabledException, ↳
↳CiiOldbIllegalOperationException;

public <U extends MdOldb<T>> U getMetadata(Class<U> clazz);

public abstract <U extends MdOldb> void setMetadata(String metadataInstanceName)
throws CiiOldbInvalidTypeException, CiiOldbCyclicDependencyException, ↳
↳CiiOldbExpressionException, CiiOldbUndefinedMetadataException, ↳
↳CiiOldbInitializationError, CiiOldbDeletedException, CiiOldbURIException;

public URI getUri();
```

Listing 8-4: C++ CiiOldbDataPoint API

```
std::shared_ptr<CiiOldbDpValue<T>> ReadValue(bool check_bad_quality = true);

void WriteValue(const T& value, int64_t timestamp = CiiOldbUtil::Now());

void WriteValue(const T& value, int64_t timestamp, CiiOldbDpQuality quality,
const std::string *overridden_server_alias = nullptr,
const std::size_t *overridden_value_size_limit_hdfs = nullptr);

void WriteValue(std::istream& value, std::uint64_t size, int64_t timestamp = ↳
↳CiiOldbUtil::Now());

void WriteValue(std::istream& value, std::uint64_t size, int64_t timestamp,
CiiOldbDpQuality quality);

void Subscribe(const std::shared_ptr<CiiOldbDpSubscription<T>>& listener);

void Unsubscribe(const std::shared_ptr<CiiOldbDpSubscription<T>>& listener);

std::shared_ptr<::elt::config::classes::meta::MdOldb<T>> GetMetadata() const;

void SetMetadata(const std::string& meta_data_instance_name);

void SetQuality(CiiOldbDpQuality quality) override;

elt::common::CiiBasicDataType GetType() const override;
```

ERROR HANDLING

Revision:	1.2
Status:	Released
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	ELT CII
Folder:	/trunk/deliverables/phase7
Document ID:	CSL-MAN-19-172698
File:	MAN-ELT-ErrorHandling
Owner:	Borut Terpinc
Last modification:	November 23, 2020
Created:	April 2, 2019

Prepared by	Reviewed by	Approved by
Borut Terpinc (CSL SWE) Jan Pribošek (CSL SWE)	Miha Vitorovič (CSL)	Gregor Čuk (CSL)

Document History

Revision	Date	Changed/reviewed	Section(s)	Modification
0.9	9.4.2019	bterpinc	All	Created.
1.0	16.9.2019	bterpinc	All	Updated after internal review.
1.1	10.2.2020	bterpinc	All	Major update of all sections after Topical review comments.
1.2	20.11.2020	bterpinc	2.2	Added explanation about the error severity.

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.



Scope

This document is an error handling manual for the ELT Core Integration Infrastructure Software project.

Audience

This document is aimed at those Cosylab and ESO employees involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms

API	Application Programmers Interface
CII	Core Integration Infrastructure
GUI	Graphical User Interface
ICD	Interface Control Document
IDE	Integrated development environment
MAL	Middleware Abstraction Layer
SVN	Subversion version control system

References

1. Cosylab, Interface Control Document, Specification, CSL-DOC-17-147262, version 1.3
2. Cosylab, Error handling design document, CSL-DOC-17-167386, version 1.2
3. ELT Config, Error Handling Transfer document, CSL-DOC-19-172644, version 1.5
4. Elasticsearch Query DSL, <https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>
5. Elasticsearch, Standard Analyzer, <https://www.elastic.co/guide/en/elasticsearch/reference/6.8/analysis-standard-analyzer.html>



8.1 Overview

This document is a user manual for CII Error handling system.

8.2 Introduction

Error Handling in CII provides the APIs and tools necessary to conveniently implement a common error handling mechanism across all EELT Control System applications and thus streamline the diagnosis of abnormal behavior of the system.

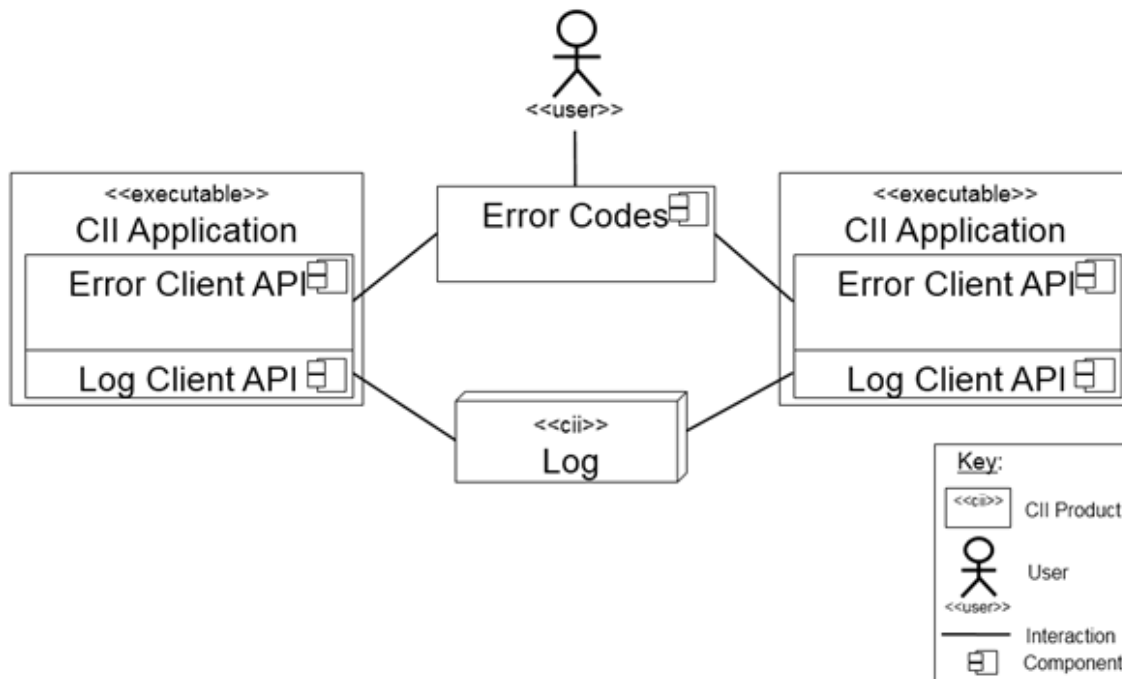


Figure 2-1 Interaction diagram of error handling

In general, error handling covers the following error management activities:

Error detection: language-specific try/catch mechanisms provide error detection.

Error handling: errors are based on exceptions. The name of the exception with the namespace defines the error type and error severity. Catching errors is based on the try/catch mechanism of each language. Error grouping is supported by exception inheritance. Error message and description is a part of exception class.

Propagation of error information: Error information is contained inside the exception entities. This information is passed between the different program parts. Passing information over the network is realized with ICD-based exceptions.

Administration and collection of all-important information: Definition of the error information is contained inside the exception class. New classes can be added and changed and it is also possible to change existing exception classes. All error information can be logged to the logging system.

Error messages will be displayed to the user: A QT GUI error display widget provides a GUI component for displaying the error messages. In the case of an error, the error message is displayed in a pop-up dialog.

All components of the error handling are provided for 3 languages: Java, CPP and Python.

8.2.1 Standard workflow

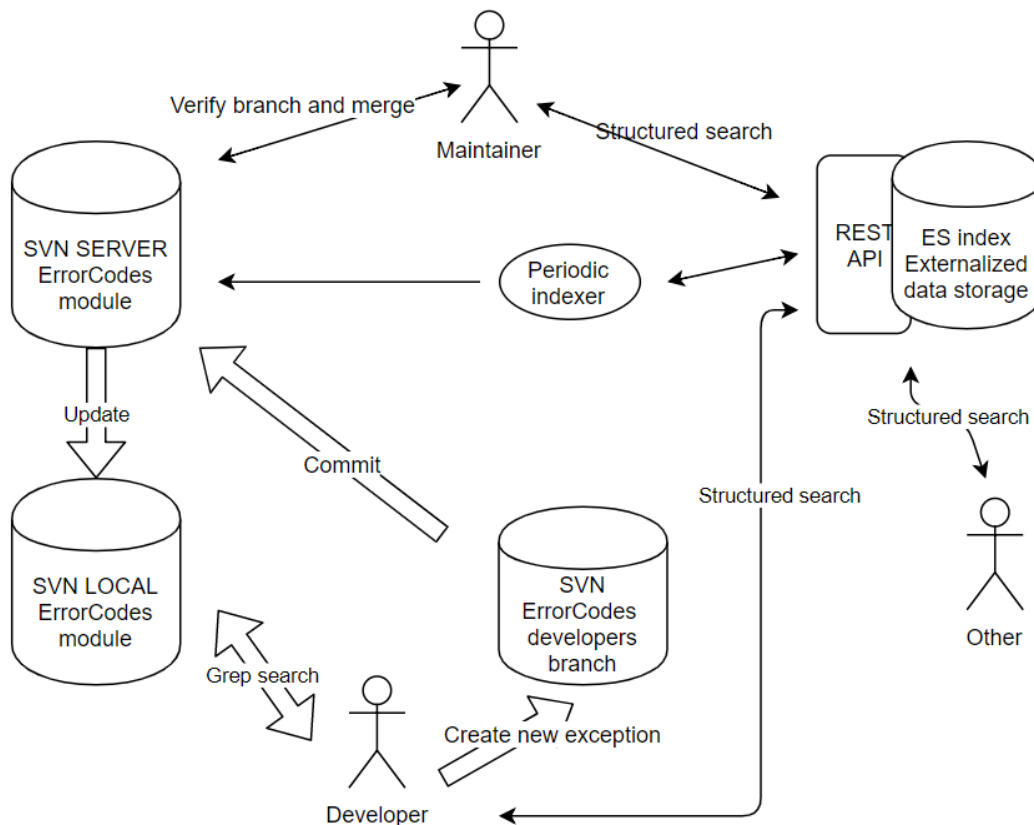


Figure 2-2 Error code repository and workflow

Figure 2-2 presents the typical workflow for storing the exception information in a central repository. It is advised to have one or more general maintainers (lead developers) of the error code system. The maintainer is responsible for the error code system integrity and has an overview of all the exceptions in the system (or a specific section).

Whenever the developers need a specific exception, they first checkout the latest version of the SVN module where needed exception is located. They then use *grep*, *find* or any other file text search tool to find the desired exception. They can also browse the folder tree to find the exceptions that belong to specific groups. A helper search script *searchExceptions.sh* is provided to aid the developer when searching. If a developer cannot find the needed exception, one creates a new exception that is stored in a separate development branch. The new exception can be used immediately in the development branch. Then the developer commits the proposed new exception to central storage in the separate branch. The developer informs the maintainer about the new exception. The maintainer reviews the new exception and either merges the exception into the trunk or proposes the use of an existing exception from the trunk to the developer. When a new exception is merged into the system, the



maintainer sends the message to the developer. The developer checks out the latest version of the *trunk* and uses the new exception from the *trunk*.

If a change to the exception information is needed, the same process as for creating an exception is used; the only difference is that the existing exception is modified and stored to the development branch. Subsequent steps are identical.

8.2.2 CII Exceptions

The CII error handling mechanism is based on CiiException, which is the base exception from which all other exceptions are derived. The CiiException derives from the special ICD defined CiiSerializedException, which provides serialization (2.4). Derived from the CiiException, two additional exceptions are defined as the part of the error handling system: CiiError and CiiBaseException. These two exceptions correspond to the error severity:

- Normal - CiiBaseException – these exceptions cover the errors which can be recovered, and can be handled at higher software layers.
- Error - CiiError – these exceptions cover the unrecoverable errors. When unrecoverable error happens, the program shall exist with the corresponding error message.

The severity of the error is defined by the exception name and the parent exception from which all errors of the selected severity are defined. The exception and error severities are pre-defined with exception names. All exceptions defined for severity »error« must contain the word error in the exception class name. All exceptions with »normal« severity must contain the word exception in the class name. Exceptions that define error severity are based on CiiError exception and exceptions with normal severity must derive from CiiBaseException.

All user defined exceptions must derive from the CiiBaseException or from CiiError.

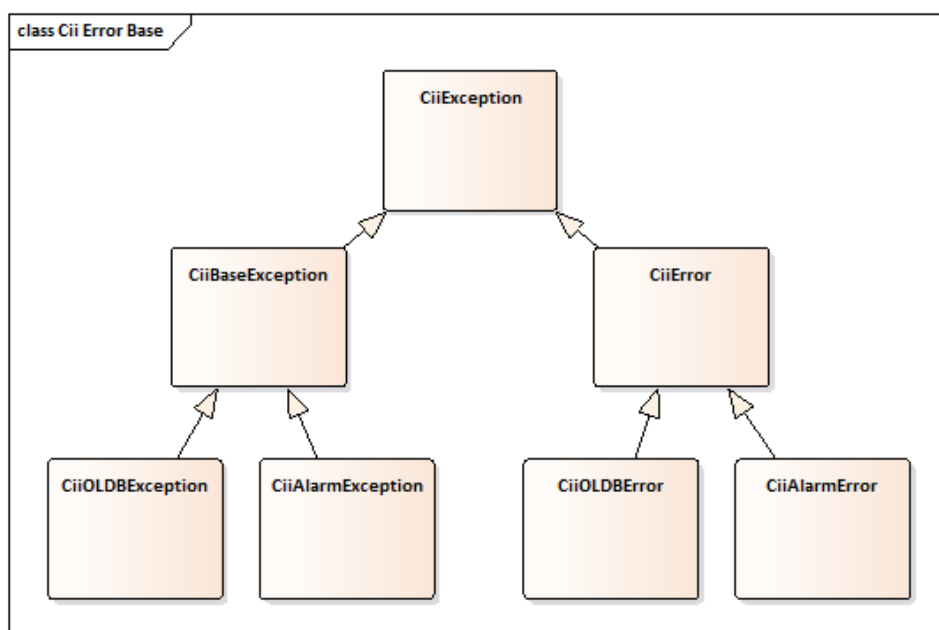




Figure 2-3 CII base exceptions and errors

Figure 2-3 shows the class inheritance diagram of the base errors and exceptions. The exceptions defined by the particular CII service all derive from the `CiiBaseException` for »normal« severity and from `CiiError` for »error« severity. Every CII service also defines root exceptions for the two severities (e.g., `CiiOLDBException`, `CiiOLDBError`). From these two exceptions, all the other exceptions for a particular system are derived.

Error stack

Single exception instance can nest exceptions and create an error stack. The stack is created by construction exceptions with the included information about the nested exception. A hierarchy of nested exceptions constructs the exception stack. Every nested exception presents an entry to the exception error stack. Exceptions (entries to the error stack) can only be added to the error stack, they cannot be changed or deleted. The error stack is automatically removed when reference to the exception with all the nested exceptions is lost.

In C++, a special throw mechanism must be used to create a stack of errors.

Whenever there is a need for the stack to be logged, an explicit entry to the log system must be programmed by the developer.

The examples of creating the exception stack can be found in 4.6.1.

Indexing Service

The Indexing service is a Java-based program that collects information about the exceptions and stores them in the Elasticsearch. The Indexing Service will traverse the SVN directory where exceptions are stored and crawl over the source code to extract the following information from the exceptions:

- error type,
- error message,
- error description,
- author.

Extraction uses the following rules:

The error type is extracted from the exception name and folder path.

The error description is extracted from the class comment. The description starts with the `@desc` tag.

The author of the exception is extracted from the class comment. It starts with the `@author` tag.

The error message is extracted from the exception class member variable named `message`. The variable is of type string.

The indexing service expects all information other than the description to be stored in one line.



To provide namespace information for the error type, the directory structure information is used. As the exceptions must be properly structured in the named folders that reflect the namespace, this information will provide the proper information for the error type namespace.

8.2.3 Conditions API

Error handling also provides conditions API. The purpose of the API is to provide a convenience API to enable checking of pre and post conditions, and method invariants. The API also supports enabling and disabling conditions and fail-fast checks. API provides static methods for verifying arguments, null values, states and valid index positions in lists and arrays. The methods, in general, accept one or more arguments for checking the condition and an argument that defines the exception type. The API interacts with the error handling system.

Whenever the condition is not met, the API will throw an exception. API contains an option for setting the fail-fast mode. When the fail-fast mode is enabled, the language-specific root exception will be thrown. This will cause the application to exit with the defined error message.

The error data is centrally stored in SVN. Chapter A.3 contains a description of Conditions API methods.

8.2.4 Serialization

Transportation of the exceptions over the distributed processes, written in different programming languages is supported with exception serialization. Serialized exceptions are transferred over MAL, which provides serialization of exceptions into plain objects. Therefore, exceptions are passed over the network with information serialized to string.

MAL uses ICD schema to define the exception structure. Custom user-defined ICD exceptions with desired information can be created. For these exceptions, there is no mechanism to implicitly transfer the data from the CiiException based exceptions to the ICD exception. Exception data can be read from the CiiException based exceptions and used to fill the ICD user-defined member fields. When an ICD based exception is transferred over the network, it is the developer's responsibility to catch the exception and manage the data transfer and pass the exception to the higher levels of error handling.

For simplicity of use, the ICD CiiSerializedException is defined. This is a special ICD exception with hard-coded additions. This exception has predefined member fields, which carry the data of the exception. CiiSerializedException is also a base exception for all CiiExceptions. This special exception simplifies transferring exceptions over the network from the sender side, as no data copying between exceptions is needed. All CiiExceptions are derived from the CiiSerializedException, this means that CiiSerializedException will automatically catch all user-defined CiiExceptions. All exceptions that use CiiException will be automatically transferred over the network as CiiSerializedException.



8.3 Installation

Error Handling provides the following components that need to be installed:

- Error indexer
- Error handling dialog widget
- Conditions API

See Configuration and Error Handling Transfer Document [3] for instructions to install the error handling system.

8.3.1 Prerequisites and waf modules

Error Handling uses SVN as central storage for exceptions. Error Handling also uses Elasticsearch as an engine for more structured searching of the exceptions. The two products must be installed and accessible from the local hosts to use the Error Handling:

- SVN server,
- Elasticsearch.

Error handling uses the following CII modules:

- elt-mal: MAL, MAL ZPB.
- client-api: Conditions API, Error Handling API.
- srv-error-indexer: Error indexer, search script.
- elt-qt-widgets: Error handling dialog widget.



8.4 Usage

The example in this manual is based on CLI SVN interface. Standard language IDE's (Eclipse, IntelliJ Idea, QT Creator) can be used for writing exception definitions and to communicate with the SVN server. Examples of IDE use are not part of this manual.

8.4.1 Includes/Imports

For basic usage of the Error Handling, the following programming language imports must be added:

Java

```
import elt.error.CiiError;  
import elt.error.CiiBaseException;  
import elt.error.CiiConditions;
```

CPP

```
#include <CiiException.hpp>  
#include <CiiConditions.hpp>
```

Python

```
import elt.error
```

QT

```
#include <CiiException.hpp>  
#include <CiiExceptionDialog.hpp>
```

8.4.2 Creating an SVN repository

To use the error codes system an SVN repository must be created, which will act as a storage for all CII Exceptions. It is expected that this repository is already created by the lead developer or repository maintainer. The developer should consult the responsible person for the information about the location.

In this manual we will create a repository in an SVN server hosted at *https://svnServer*, at */pathToErrorCodes/errorCodes* path. The repository can be created with the following command:



```
svn mkdir -m "Create error codes directory" https://svnServer/pathToErrorCodes/  
↩errorCodes --username myusername
```

In the *errorCodes* repository, the basic error codes system structure must be created. Folders can be created in the same manner as the root folder, or a local copy of structure can be prepared and committed to the server. The following structure must be created:

errorCodes/trunk/java

errorCodes/trunk/cpp

errorCodes/trunk/python

errorCodes/branches

Note: The root folder in SVN for the exceptions can be changed to the specific path (e.g. *p8/trunk/CONTROL-SYSTEMS/CentralControlSystem/CII/CODE/srv-error-indexer/sampleErrorCodeTestRepo/*). However, the structure of folders after the root folder is fixed and cannot be changed in order for Indexing service (6) to properly collect the error codes from SVN repository. Indexing service expects *trunk* folder and subfolders *java*, *cpp* and *python*. The names of the subfolder's are a part of indexing service source code. A small change in Indexing service code must be made to use the different names for the folder structure.

8.4.3 Creating the exceptions

Exceptions must be created in the folder reflecting the namespace and with a file name reflecting the exception class name. Multiple exceptions with the same name are allowed, if they are put in separate namespaces. Exception name together with the namespace defines the error type and severity.

The workflow of error codes is described in section 2.1. Shortly, a developer must create exceptions in his own branch, which is then checked by the maintainer and merged into the trunk.

Note: It is expected that the development environment was properly set and that the variable *\$INT-ROOT* points to the root location where all the CII components are installed (command "*module load introot eeltdev*" was run).

If the SVN *errorCodes* folder is not present on local developer's computer, we first need to check-out the *errorCodes* directory from the SVN.

```
cd $INTROOT/sources/  
svn co https://svnServer/pathToErrorCodes/errorCodes --username myusername
```

New exceptions, defined by a developer, should be created in their own branch. It is suggested to create this branch with a folder named after developer username. For each language a separate folder is required (java, cpp, python).

```
cd branches/  
mkdir developerName
```

(continues on next page)



(continued from previous page)

```
cd developerName  
mkdir java cpp python
```

Exceptions source code must be structured as shown in the following examples. Exception files must be placed in the folders, whose names reflect the namespace. Each single exception is placed in own file with the filename named the same as the exception class.

In the example below we create `CiiFileHandlingException`, which is in the `elt.error.common` package. We create this exception for all 3 languages. Folders for this package must be created:

```
mkdir -p java/src/elt/error/common  
mkdir -p python/src/elt/error/common  
mkdir -p cpp/elt/error/src/include
```

Information from the exceptions is parsed by the indexing service. In order for the indexing service to properly parse the exceptions, the exceptions must be placed in the described folders. Also, the structure of exceptions source code must follow the structure presented in the following chapters (4.3.1, 4.3.2, 4.3.3). Information about the author and description of the exception is parsed from the source code comments. The information about the exception message is parsed from the message string.

Examples of the exceptions SVN repository can be found in the `srv-error-indexer` module under the `sampleErrorCodesRepo` directory. The following sections show the programming language specific examples.

Java

To create a Java exception, create `CiiFileHandlingException.java` file inside the `java/src/elt/error/common` folder with the desired exception specifics. The following example shows the Java exception structure:

```
package elt.error.common  
  
/**  
 * @desc This is the description of file handling error.  
 * @author Somebody@eso.org  
 */  
  
public class CiiFileHandlingException extends CiiBaseException {  
    public static final String MESSAGE = "Error while handling file: %s";  
  
    public CiiFileHandlingException(String path) {  
        super(String.format(MESSAGE, path));  
    }  
}
```



CPP

For CPP exceptions to be browsable, they must follow the Java style for exception filenames. Each exception is created in its own .hpp file, with the file named the same as the exception class name. Exception details (description and author) must be added inline in the header code. CPP exceptions are header only based, no .cpp files shall be created.

To create a CPP exception, create CiiFileHandlingException.hpp file inside the *cpp/elt/error/common/src/include* folder with the desired exception specifics. The following example shows the CPP exception structure:

```
/**
 * @desc This is the description of file handling error.
 * @author Somebody@eso.org
 */

namespace elt {
namespace error {
namespace common {

class CiiFileHandlingException : public CiiBaseException {

const std::string getMessage(){
    static const std::string& message = "Error while handling file: %s";
    return message;
    //other option for message definition
    // return std::string("Bad version number %i");
}

public:
    CiiFileHandlingException(const char* fileName):
        CiiException(getMessage(), fileName){}
};

//other option std::string
// if the argument passed to the CII Exception is string it has to be converted_
↳to const char*
/*
public:
    CiiFileHandlingException(const std::string& fileName):
        CiiException(getMessage(), fileName.c_str()){}
}; */
}
}
}
```



Python

For Python exceptions to be browsable, they must follow the Java style for exception filenames. Each exception is created in its own .py file, with the file named the same as the exception class name.

To create a Python exception, create CiiFileHandlingException.py file inside the *python/src/elt/error/common* folder with the desired exception specifics. The following example shows the Python exception structure:

```
"""  
@desc This is the description of file handling error.  
@author Somebody@eso.org  
"""  
  
class CiiFileHandlingException(CiiBaseException):  
    _message = "Error while handling file: %s"  
    def __init__(self, arg1):  
        super().__init__(_message % arg1)
```

After the exceptions are created in a separate branch, the developer commits the exceptions to the SVN server and informs the exceptions maintainer about the existence of new exceptions. The maintainer then merges the exception into the trunk.

It is the maintainer's responsibility to check that the exception structure (folders and namespace) conforms to the convention.

8.4.4 Wrapping native exceptions

As native language exceptions follow their own inheritance tree, they cannot be directly caught by the CiiException. For this reason, all the native exceptions must be caught and rethrown with the details passed to the variant of the CiiException. This way information provided by the native exceptions can be part of the CII Exception.

The following example shows how to pass the native exception information to CII Exceptions in Java:

```
try {  
    testList.get(arrayIndex);  
} catch (IndexOutOfBoundsException ex) {  
    throw new CiiBadIndexException(arrayIndex, ex);  
}
```

Same mechanism should be used in Python and CPP.



8.4.5 Search script

Search script **searchExceptions.sh** is a simple bash script providing capabilities of automatic SVN update, local search and remote search. The script is a wrapper script around the Linux *find* and *grep* commands. The script automatically checks out the latest version of *errorCodes* from the SVN server and executes a search. It also provides search for indexed exceptions in Elasticsearch. While the functionality of the search in Elasticsearch is practically the same as for other *find* and *grep* style searches, it provides more structured output with more detailed information about the exceptions.

The script is intended to be used locally by developers. All the settings are embedded directly in the script.

Configuring the script

To configure the script, open it in a text editor:

```
nano $INTROOT/bin/searchExceptions.sh
```

Change the following settings:

- **REMOTE_REPO**: location of the root of remote SVN repository
- **SVN_USERNAME**: username with SVN permissions
- **SVN_PASSWORD**: password for the SVN user
- **LOCAL_REPO**: location of the root of local SVN repository

Example:

```
REMOTE_REPO="https://svnhq8.hq.eso.org/p8/trunk/CONTROL-SYSTEMS/  
↪CentralControlSystem/CII/CODE/srv-error-indexer/sampleErrorCodesRepo"  
SVN_USERNAME="myusername"  
SVN_PASSWORD="mypassword"  
LOCAL_REPO="myLocalRepoRoot/sampleErrorCodesRepo"
```

Usage of the search script

A search without a special flag argument will perform a *grep* style search:

```
./searchExceptions.sh "Error while handling file:"
```

The search with the **-f** argument will execute a *find* style search:

```
./searchExceptions.sh -f "CiiFileHandlingException.java"
```

Direct Elasticsearch index searches can be executed with the **-r** argument. The following example shows how to search the exception index using the Elasticsearch query:



```
searchExceptions.sh -r "file error"
```

This searches for the exceptions that contain either word *file* or word *error* in the following fields of the Elasticsearch index:

- Name
- Folder
- Namespace
- Author
- Description
- Message
- Language
- Comment

A more relevant match (i.e. all or many search words are present) appear sooner in the returned list of results.

To see all functions supported by the script, the -h option can be used:

```
searchExceptions.sh -h
```

8.4.6 Exceptions examples

The examples show usage of the CII error handling system for Java, CPP and Python. As the majority of native language exception handling is similar across all three languages, a general example is shown only for Java:

```
public class ErrorHandlerExample {  
  
    //initialize logger  
    private static Logger logger =  
        CiiLogManager.getLogger("ErrorHandlerUsageExample");  
  
    public static void main(String[] args) {  
  
        /* Main method is the on the highest level (first) of exception handling.  
        * All unhandled exception will end the application.  
        */  
  
        //initialize example object  
        ErrorHandlerExample ex = new ErrorHandlerExample();  
  
        /* Handle the exception from the runProductionArrayElementCode() and
```

(continues on next page)



(continued from previous page)

```
* log the current exception to log (Level.WARNING).
*/
try {
    ex.runProductionArrayElementCode();
} catch (CiiException e) {
    /* Exception is treated as warning. Current exception stack
    * is logged on a warning.
    * Exception is not passed to higher level, so the stack is cleared.
    */
    logger.warn(CiiLogManager.formatException(e));
}

/* Error will not be logged at this level.
* Stack will also contain entries from the ex.runProductionFileOpen();
*/
try {
    ex.runProductionFileOpen();
} catch (CiiException e) {

    /*
    * Access individual information
    */
    System.out.println(e.getTypeName()); //errorTypeID
    System.out.println(e.getCreationDate()); //detectionTimestamp
    System.out.println(e.getStackTraceAsString()); //localization information
    System.out.println(e.getMessage()); //message

    /* This will exit the application with CiiException as the exception
    * is thrown from the main method.
    */

    throw e;
}

}

/**
* Methods on ErrorHandlerExample present the 2 level of exception
* handling.
*/
public void runProductionArrayElementCode() {

    /*
    * Error from lower level is handled as non-critical and logged
    * as warning.
    * The super-type CiiException also catches CiiArrayOutOfRangeException
    */
    ProductionCode pcHandled = new ProductionCode();
```

(continues on next page)



(continued from previous page)

```
try {
    pcHandled.codePartThatGetsArrayElement();
} catch (CiiException ex) {
    logger.warn(CiiLogManager.formatException(ex));
}

/*
 * Error from lower level is handled but passed also to higher level.
 */
ProductionCode pc = new ProductionCode();
try {
    pc.codePartThatGetsArrayElement();
} catch (CiiBadIndexException ex) {
    /*
     * Stack will contain details form CiiBadIndexException
     * Error will be propagated to upper level
     */
    throw ex;
}
}

public void runProductionFileOpen() {
    /*
     * Non handled, error will be handled in the higher level.
     */
    ProductionCode pcNonHandled = new ProductionCode();
    pcNonHandled.codePartThatTriesToOpenFile();
}

private class ProductionCode {

    /**
     * Application level 3 error handling
     */
    public void codePartThatGetsArrayElement() {

        final ArrayList<Integer> testList =
            new ArrayList(Arrays.asList(1, 2, 3, 4, 5));
        final int arrayIndex = 9;

        /*
         * Example of array out of range error, where java native exception
         * is wrapped into
         * CiiArrayOutOfRangeException. The error is propagated to higher level.
         */

        try {
            testList.get(arrayIndex);
        } catch (IndexOutOfBoundsException ex) {
```

(continues on next page)



(continued from previous page)

```
    /*
     * Internal exception will be caught, but nothing will be passed.
     * IndexOutOfBoundsException is native java exception.
     */
}

try {
    testList.get(arrayIndex);
} catch (IndexOutOfBoundsException ex) {
    /* Internal exception will be caught, and passed
     * as a CiiBadIndexException
     *
     * Error message example: "Index 9 is out of bounds!
     * Error will be propagated to higher level.
     */
    throw new CiiBadIndexException(arrayIndex, ex);
}
}

public void codePartThatTriesToOpenFile() {

    /*
     * Non existing file is logged as info.
     */
    final String path = "data.txt";
    final File myDatafile = new File(path);
    if (!myDatafile.exists()) {
        /*
         * We treat this as non - error, we only log the information
         * to logger.
         */
        logger.log(Level.INFO, "File %s could not be opened", path);
    }

    /*
     * Non existing file is threaded as error and passed to higher level.
     */
    if (!myDatafile.exists()) {
        /*
         * Pass the exception to higher level.
         * Exception will have path added to the message.
         */
        throw new CiiFileHandlingException(path);
    }
}
}
```

Presented example will produce the following log messages:



```
14:53:41.056 [main] INFO ErrorHandlingUsageExample - Index 9 is out of bounds!,  
↳ErrorTypeName=elt.error.CiiBadIndexException, ErrorDateTime=1579445621042,  
↳ErrorMessage=Index 9 is out of bounds!, ErrorTrace=elt.error.  
↳CiiBadIndexException  
    at elt.error.ErrorHandlingExample$ProductionCode.  
↳codePartThatGetsArrayElement (ErrorHandlingExample.java:127)  
    at elt.error.ErrorHandlingExample.  
↳runProductionArrayElementCode (ErrorHandlingExample.java:75)  
    at elt.error.ErrorHandlingExample.main (ErrorHandlingExample.java:41)  
Caused by: java.lang.IndexOutOfBoundsException: Index: 9, Size: 5  
    at java.util.ArrayList.rangeCheck (ArrayList.java:657)  
    at java.util.ArrayList.get (ArrayList.java:433)  
    at elt.error.ErrorHandlingExample$ProductionCode.  
↳codePartThatGetsArrayElement (ErrorHandlingExample.java:120)  
    ... 2 more  
,  
14:53:41.069 [main] WARN ErrorHandlingUsageExample - File data.txt could not be  
↳opened  
Exception in thread "main" elt.error.CiiFileHandlingException: Error while  
↳handling: data.txt  
    at elt.error.ErrorHandlingExample$ProductionCode.  
↳codePartThatTriesToOpenFile (ErrorHandlingExample.java:154)  
    at elt.error.ErrorHandlingExample.  
↳runProductionFileOpen (ErrorHandlingExample.java:90)  
    at elt.error.ErrorHandlingExample.main (ErrorHandlingExample.java:49)
```

Source code of this example can be found in the client-api module under `elt-common/java/common/test/elt/error/ErrorHandlingExample.java`

Building an Error Stack

The following examples show sample code of building the exception stack in different languages. The exception stack is built by nesting the exceptions:

Building Error Stack in Java

```
try {  
    try {  
        throw CiiOldDbDpExistsException("testDPUri");  
    } catch (CiiOldDbDpExistsException ex) {  
        throw (CiiBadIndexException(10, ex));  
    }  
} catch (CiiBadIndexException ex) {  
    ex.getCiiExceptionStackAsString();  
}
```



Building Error Stack in Python

```
try:
    try:
        raise CiiInvalidTypeException("")
    except CiiException as x:
        raise CiiBadIndexException(x, 10)
except CiiException as cix:
    print(cix.get_stack_trace())
```

Building Error Stack in CPP

```
try
{
    try
    {
        throw elt::oldb::CiiOldbDpExistsException("testDPUri");
    } catch (const elt::error::CiiOldbException& e) {
        // build exception stack of nested exceptions
        // CiiOldbDpExistsException is nested in CiiBadIndexException
        CII_THROW_WITH_NESTED(elt::error::CiiBadIndexException, e, 10);
    }
} catch (const elt::error::CiiException& e) {
    // flatten all nested exceptions and dump to string
    std::string nestedDetails = e.dumpWithNested();
}
```

Throwing a CPP Exception with Extra Information

CPP API provides additional macros which add line number and code filename to the error messages. Section A.2 in Appendix A contains the details.

The following example shows throwing a CPP exception with details (line number, function, file name, etc.). Standard CPP exception what() method (inherited from std::exception) shall be used to obtain the exception message. CPP API also provides the extra Dump() method (see details in section A.1).

```
try
{
    CII_THROW(elt::error::CiiFileHandlingException, "myfileName");
}
catch(const elt::error::CiiException& e)
{
    std::cout << e.What(); //error message
    std::cout << e.GetDetails(); //error details
    std::cout << e.Dump(); //error dump
}
```



Exception Serialization from CPP to Java

The following code shows an example of transferring the exception from CPP to Java. All the exceptions are serialized using the MAL CiiSerializableException.

```
//CPP method with CiiException on the sender side
std::string methodWithException(const std::string &par1) override
{
    throw elt::error::CiiInvalidDataTypeException(par1);
    return par1;
}

//Java - Receiving exception on the reciever side
try {
    final String unexpectedReply = serialization.methodWithException("Test");
} catch (CiiSerializableException e) {
    logger.info(String.format("Type: %s", e.getCiiType()));
    logger.info(String.format("Host name: %s", e.getCiiHostName()));
    logger.info(String.format("File name: %s", e.getCiiFileName()));
    logger.info(String.format("Class name: %s", e.getCiiClassName()));
    logger.info(String.format("Method name: %s", e.getCiiMethodName()));
    logger.info(String.format("Message: %s", e.getCiiMessage()));
    logger.info(String.format("Creation Date: %d", e.getCiiCreationDate()));
    logger.info(String.format("Stack trace: %s", E.getCiiStackTrace()));
    logger.info(String.format("Exception stack: %s", e.getCiiExceptionStack()));
    logger.info(String.format("Details: %s", e.getCiiDetails()));
}
}
```

8.4.7 Using the Condition API

The following example shows the use of the condition API. Since the use is similar for Java, CPP and Python, only a Java example is presented.

```
public static void main(String[] args) {
    // Enable condition API for the process.
    CiiConditions.setEnabled(true);

    ErrorConditionsExample ex = new ErrorConditionsExample();

    try{
        ex.setRotation(250);
    } catch (CiiException x){
        System.out.println(x.getMessage());
    }
    ex.getListValue(10);
}

// get value from array of float
```

(continues on next page)



(continued from previous page)

```
public float getListValue(int index){
    // check is element is in range.
    CiiConditions.checkElementIndexRange(index, list.length,
        new CiiBadIndexException(index));
    return list[index];
}

public void setRotation(int degrees){
    //pre condition check -- check if rotation in smaller than 360 deg.
    CiiConditions.checkArgument(degrees <= 360, new CiiOverRotatedException());
    axis.setRotation(degrees)
    axis.rotate()
    final int readRotation = 200;
    // post condition check -- check if axis is properly rotated.
    CiiConditions.checkArgument(readRotation == degrees,
        new CiiNotProperlyRotatedException());
}
```



8.5 Using the QT Error Dialog Widget

To use the error dialog widget, CiiExceptionDialog.hpp must be included into the source code. The dialog is located under the elt::error::gui namespace. The CiiExceptionDialog.hpp derives from the QDialog. The method exec() is therefore directly inherited from this QDialog class.

Exception dialog constructor takes a CiiException as an input parameter. All exception details are automatically parsed into the dialog, based on the CiiException type.

To display the widget, exec() method must be called. The following code shows the typical usage:

```
#include <CiiExceptionDialog.hpp>

try {
    throw elt::error::CiiInvalidTypeException();
} catch (const elt::error::CiiException& e) {
    gui::CiiExceptionDialog w(e);
    w.exec();
}
```

Figure 5-1 shows an example of the Error widget dialog screen output. The widget displays all the important error information: Exception class (error type), exception message, time of the occurrence, thread ID, process ID, host, line number and stack trace.

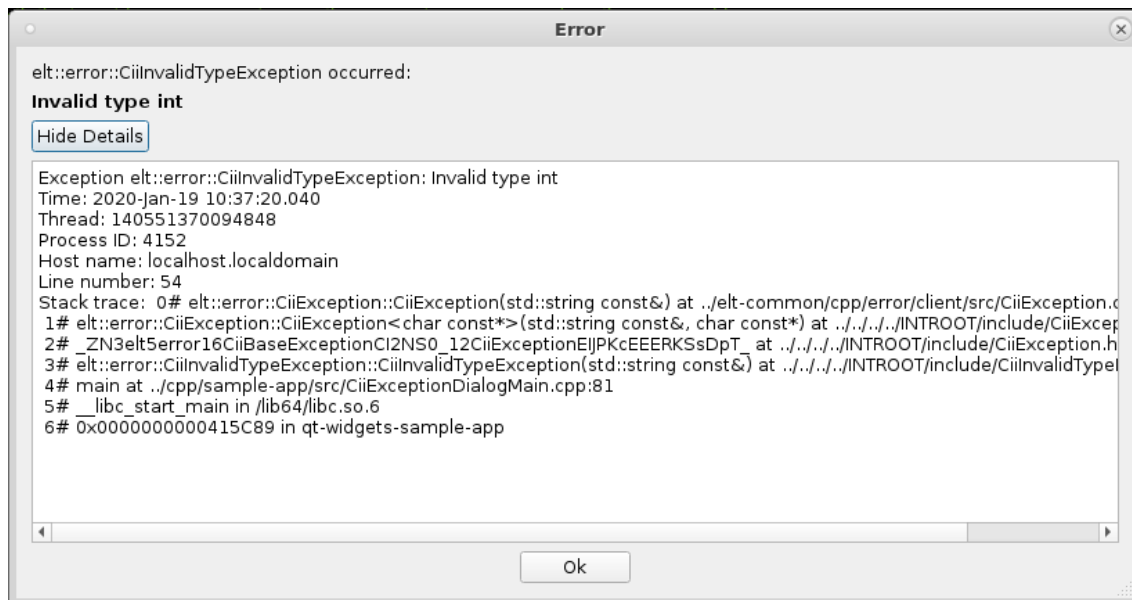


Figure 5-1 Error widget dialog screen output



8.6 Indexing service

Error indexer (2.2.2) is a tool that collects the data of exceptions and stores the exceptions data to Elasticsearch index. Using Elasticsearch for searching the exceptions provides structured and powerful search. Data provided by indexing service is used by the search script (4.5), when using the ElasticSearch option. As ElasticSearch returns data in JSON format, the data returned by the queries can be used for fast integration with other services.

8.6.1 Configuring indexing service

Error indexer uses configuration localDB for setting up the needed configuration. The configuration instance file must be created and deployed for error indexer to work properly. Use the following template to create the config file:

```
config:
  instance:
    __comment__: Configuration for error-indexer
  data:
    "@type": CiiErrorIndexerConfigClass
  fields:
    elasticSearchServiceAddress:
      "@type": MdString
      metadataInstance: ConfStringStd
      value: "localhost"
    elasticSearchIndexName:
      "@type": MdString
      metadataInstance: ConfStringStd
      value: "errors"
    centralErrorCodesRepositoryHost:
      "@type": MdString
      metadataInstance: ConfStringStd
      value: "https://svnhq8.hq.eso.org"
    centralErrorCodesRepositoryPaths:
      "@type": MdStringArray
      metadataInstance: ConfStringArrayStd
      value:
        - "p8/trunk/CONTROL-SYSTEMS/CentralControlSystem/CII/CODE/srv-error-
↔indexer/sampleErrorCodesRepo/"
    lastProcessedRevision:
      "@type": MdInt32
      metadataInstance: ConfInt32Std
      value: 0
    centralErrorCodesPassword:
      "@type": MdString
      metadataInstance: ConfStringStd
      value: "myPass"
    centralErrorCodesUserName:
      "@type": MdString
```

(continues on next page)



(continued from previous page)

```
metadataInstance: ConfStringStd  
value: "myUsername"
```

After the file is created (we assume it is saved under `ErrorIndexerConfig.cdi.yaml`), it has to be deployed to the localDB. If localDB is not initialized, it should be initialized first.

```
config-tool init #LocalDb initialization.  
config-tool deploy -i localDbConfiguration/ErrorIndexerConfig.cdi.yaml -a cii.  
↪config://services/CiiErrorIndexer
```

Following properties are set:

- `elasticSearchServiceAddress`: Address of the elastic search server.
- `elasticSearchIndexName`: Index name, which will hold the indexed data.
- `centralErrorCodesRepositoryHost`: SVN hostname of central repository.
- `centralErrorCodesRepositoryPaths`: Array of paths in the repository. Indexer will look for exception code under these subfolders.
- `lastProcessedRevision`: Internally, last processed revision is stored. The setting should initially be set to 0.
- `centralErrorCodesPassword`: password credential for SVN.
- `centralErrorCodesUserName`: user name credential for SVN.

8.6.2 Running the indexing service

To manually run the indexer, execute the following command:

```
error-indexer
```

It is proposed to run the indexer every 30 minutes. A Linux time-based utility *cron* shall be used for this. To define the schedule to use the *cron*, put the following line into the `/etc/crontab` (change the `{INTROOT}` to proper location of the CII root installation folder – variable `$INTOOT` cannot be used here):

```
30 * * * * {INTROOT}/bin/error-indexer
```

Another option (not documented in this manual) is to hook running of indexer on every SVN update.



8.6.3 Accessing the indexed data

Elasticsearch provides more powerful search capabilities of the exception definitions and more structured output searched data. As the data is returned in JSON format, it is more suitable for integration with other services, for example logging. Elasticsearch uses REST API to access the data. Any REST client can be used to access the errors data. For the following examples we will use *curl*.

Displaying all the errors in the index:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty'
```

Elasticsearch provides Query DSL language [4], with powerful search capabilities. The following fields can be used, when searching the error index:

- folder,
- author,
- namespace,
- description,
- language,
- message,
- version,
- revision,
- timestamp.

Elasticsearch supports different types of searches. When using the *wildcard* type of search, only lower case is supported (ElasticSearch standard analyzer is used [5]). Data in index can be in camel case, but the search will only work if the lower-case words are used in the search query. The search will not distinguish between lower and upper case.

Getting the error data based on the exception name:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty' -d '{"query":{"match":{"name": " CiiFileHandlingException"}}}' -H "Content-Type: application/json"
```

Wildcard search based on exception name:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty' -d '{"query":{"wildcard":{"name": "ciifile*"}}}' -H "Content-Type: application/json"
```

Searching the errors with the desired description:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty' -d '{"query":{"match":{"description": "This is the description"}}}' -H "Content-Type: application/json"
```




Searching the errors with the desired description and name:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty' -d '{"query": {  
  ↪ "bool": {"should": [{ "match": { "description": "This is the description"}}, {  
  ↪ "match": { "message": "Error while"}}]}}}' -H "Content-Type: application/json"
```

Wildcard search with the desired description and name:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty' -d '{"query": {  
  ↪ "bool": {"should": [{ "wildcard": { "description": "* is the description"}}, {  
  ↪ "wildcard": { "message": "*while"}}]}}}' -H "Content-Type: application/json"
```

Search with the desired language and name:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty' -d '{"query": {  
  ↪ "bool": {"must": [{ "match": { "language": "JAVA"}}, { "match": { "name": "CiiFileHandlingException"}}]}}}' -H "Content-Type: application/json"
```

Return the last record of wildcard search sorted by timestamp:

```
curl -XPOST 'http://ciielastichost:9200/errors/_search?pretty' -d '{"query": {  
  ↪ "wildcard": { "name": "ciifile*"}, "sort": [ { "timestamp": { "order": "desc" }  
  ↪ } ], "size": 1}' -H "Content-Type: application/json"
```



8.7 API

8.7.1 CII Exception methods

CII-specific fields and methods are added to CiiException to provide the required functionalities:

stack trace: (execution stack trace - back trace) is provided by language-specific trace mechanisms in Java and Python. CPP uses the *boost libbacktrace* module. The `getStackTraceAsString` method provides a stack trace in a string format.

creation date: creation date/time is implicitly set by CiiException. The date/time is stored in milliseconds since Epoch. The `getCreationDate` method will return the number.

type name: type name is implicitly set by the exception definition. The `getCiiType` method will provide the type of the exception class.

The CPP CiiException class provides additional functionalities:

The following fields are provided as member fields with getter and setter methods:

threadId

lineNumber

functionName

filename

processId

hostname

details: additional exception trace details can be added to CPP exceptions. This attribute should be explicitly set by the developer. The `getDetails` and `setDetails` methods should be used to retrieve and store the information.

The two additional methods are added:

dump: method that returns string formatted information about all exception details and values.

dumpWithNested: method that returns string formatted information about all nested exceptions details and values.

8.7.2 CPP macros

CPP API provides macros which add details (file, function and line number) to the error messages. If these macros are not used, the exceptions will not contain the mentioned fields. The following macros are provided:



`CII_THROW(exceptionType_t, ...)`

Throws the desired exception with additional information. First parameter is the exception type. All other parameters are variadic, and depend on the number of parameters the exception class defines.

`CII_THROW_WITH_NESTED(exceptionType_t, nested_exception, ...)`

Throws the desired exception with additional information and nested exception. The first parameter is the exception type and the second parameter is the exception to be nested. All other parameters are variadic, and depend on the number of parameters the exception class defines.

8.7.3 Conditions API

The conditions API provides static methods for verifying arguments, null values, states and valid index positions in lists and arrays. It provides the following methods:

`setEnabled(bool enabled)`
`getEnabled()`

Methods for enabling and disabling the Conditions API. If Conditions API is disabled, checks will not be executed. The `getEnabled` method will return the current state of the enabled flag.

`setFailFastMode(bool enabled)`
`getFailFastMode()`

Methods for enabling and disabling the fail-fast mode. If fail-fast is enabled, failing the condition will cause the application to exit with an error. The `getFailFastMode` method will return the current state of the enabled flag.

`checkArgument(bool expr, Exception ex)`

Checks that the parameter `expr` is true. It can be used for validating arguments in methods. If the expression evaluates to false, the given exception is thrown.

`checkNotNull(T value, Exception ex)`

Checks that the value is not null. It returns the value directly, so you can use `checkNotNull(value, ex)` inline. In case of null, the given type exception is thrown.



`checkState(boolean value, Exception ex)`

Checks some state of the object, not dependent on the method arguments. For example, an iterator might use this to check that method `next()` has been called before any call to `remove()` method. If the expression evaluates to false, the given exception is thrown.

`checkElementIndexRange(int index, int size, Exception ex)`

Checks that `index` is a valid element index in a list, string, or array with the specified size. An element index may range from 0 inclusive to size exclusive. Instead of passing the whole list, string, or array, only its size is passed as parameter. If the index exceeds the size, the given exception is thrown.

The condition API for specific languages (Java, CPP and Python) is shown in the following sections.

Java

```
public class CiiConditions {
    public static boolean getEnabled() {...}

    public static void setEnabled(boolean enabled) {...}

    public static boolean getFailFastMode(){...}

    public static void setFailFastMode(boolean failFastMode) {...}

    public static <E extends Exception> void checkArgument(boolean expr, E ex)
↳throws E {...}

    public static <T, E extends Exception> void checkNotNull(T value, E ex)
↳throws E {...}

    public static <E extends Exception> void checkState(boolean value, E ex)
↳throws E {...}

    public static <E extends Exception> void checkElementIndexRange(int index,
↳int size, E ex) throws E {...}
}
```



CPP

```
class CiiConditions {
public:

    static bool GetEnabled();

    static void SetEnabled(bool enabled);

    static bool GetFailFastMode();

    static void SetFailFastMode(bool failFast);

    template<typename E, class T, typename ... Args>
    static void CheckArgument(bool expr, Args ... args);

    template<typename E, class T, typename ... Args>
    static void CheckNotNull(T value, Args ... args)

    template<typename E, class T, typename ... Args>
    static void CheckState(T value, Args ... args)

    template<typename E, class T, typename ... Args>
    static void CheckElementIndexRange(size_t index,
        size_t size, Args ... args);

};
```

Python

```
class CiiConditions:

    @staticmethod
    def get_enabled():
        pass

    @staticmethod
    def set_enabled(enabled):
        pass

    @staticmethod
    def get_fail_fast_mode():
        pass

    @staticmethod
    def set_fail_fast_mode(failFastMode):
        pass
```

(continues on next page)



(continued from previous page)

```
@staticmethod
def check_argument(expr, ciiException):
    pass

@staticmethod
def check_not_null(value, ciiException):
    pass

@staticmethod
def check_state(expr, ciiException):
    pass

@staticmethod
def check_element_index_range(index, size, ciiException):
    pass
```



8.8 Serialization code example

The following code presents example of serialization of exceptions over MAL. The example consists of MAL CPP server and Java client. Examples for other languages can be found in the client-api module source code:

- Java: client-api/elt-common/java/errorNonUnitTests
- CPP: client-api/elt-common/cpp/error/errorSerialCInt, errorSerialSrv
- Python: client-api/elt-common/runScripts/testZpbClientPythonException, testZpbServerPythonException

8.8.1 CPP serialization server code

```
#include <CiiSerializableException.hpp>
#include <CiiSerializationTest.hpp>
#include <CiiInvalidDataTypeException.hpp>

#include <mal/rr/qos/QoS.hpp>
#include <mal/Cii.hpp>
#include <mal/Mal.hpp>
#include <mal/utility/LoadMal.hpp>
#include <mal/rr/ServerContextProvider.hpp>
#include <mal/rr/ServerContext.hpp>
#include <mal/rr/ServerAmi.hpp>
#include <mal/rr/RrEntity.hpp>

#include <stdexcept>
#include <memory>
#include <iostream>
#include <boost/thread/future.hpp>

namespace {

const std::string EXCEPTION_CALL = "error";

/**
 * Simulate some work
 */
static void simulateWork() {
    std::this_thread::sleep_for(std::chrono::seconds(2));
}

class CiiSerializationTestImpl : public virtual
    elt::error::icdTest::CiiSerializationTest
{
public:
    std::string methodWithException(const std::string &par1) override
    {
```

(continues on next page)



(continued from previous page)

```
std::cout << "Got exception call with parameter: " << par1 << std::endl;
if (par1 == EXCEPTION_CALL)
{
    elt::mal::throw_exception(elt::error::CiiInvalidDataTypeException(par1));
}
return par1;
}

std::shared_ptr<:elt::mal::rr::Ami<std::string>>
amiMethodWithException(const std::string& par1)
{
    // Obtain ptr to ServerAmi
    std::shared_ptr<:elt::mal::rr::Ami<std::string>> ami =
↪elt::mal::rr::ServerContextProvider::
    getInstance<elt::mal::rr::ServerContext<std::string>>().createAmi();

    std::cout << "Ami call with param: " << par1 << std::endl;

    if (par1 != EXCEPTION_CALL)
    {
        auto future = boost:::async(boost:::launch:::async, [=]() {
            simulateWork();
            // first response
            ami->complete("ONE");

            simulateWork();
            // second response
            ami->complete("TWO");

            simulateWork();
            // third response
            ami->completed("THREE");
        });
    }
    else
    {
        auto future = boost:::async(boost:::launch:::async, [=]() {
            simulateWork();
            // first response
            ami->complete("ONE");

            simulateWork();
            // second response
            :elt::error::CiiInvalidDataTypeException e =
↪elt::error::CiiInvalidDataTypeException(par1);
            e.setFileName(boost:::filesystem::path(__FILE__).filename().string());
            e.setFunctionName(__FUNCTION__);
            e.setLineNumber(__LINE__);
            e.setClassName(boost:::core::demangle(typeid(*this).name()));
```

(continues on next page)



(continued from previous page)

```
ami->completeExceptionally(e);

simulateWork();
// third response
ami->completed("THREE");
});
}
return ami;
}
};

} // namespace

int main(int ac, char* av[]) {
    std::cout << "CPP ZPB server test start." << std::endl;

    // Load DDS MAL mapping
    std::shared_ptr<elt::mal::Mal> ddsMal = elt::mal::loadMal("zpb", {"dds.
↪domain", "100"});
    ::elt::mal::CiiFactory &factory = elt::mal::CiiFactory::getInstance();
    factory.registerMal("zpb", ddsMal);

    elt::mal::Uri uri("zpb.rr://0.0.0.0:12081/CiiSerializationTest");

    try {
        // Create server, default QoS.
        std::unique_ptr<elt::mal::rr::Server> server = factory.createServer(uri,
↪elt::mal::rr::qos::QoS::DEFAULT, {});

        // Register the service.
        std::shared_ptr<elt::mal::rr::RrEntity> exSer(new
↪CiiSerializationTestImpl());
        server->registerService<elt::error::icdTest::CiiSerializationTest, true>(
↪"inst_1", exSer);

        // Process commands until server->shutdown() is called.
        server->run();
    } catch (std::exception& exc) {
        std::cerr << "EXCEPTION: " << exc.what() << std::endl;
    }

    std::cout << "Server stopped." << std::endl;
    return 0;
}
```



8.8.2 Java serialization client code

```
import java.net.URI;
import java.util.Properties;
import java.util.concurrent.CompletableFuture;
import java.util.concurrent.ExecutionException;
import java.util.concurrent.TimeUnit;
import org.apache.logging.log4j.Logger;
import elt.error.icd.CiiSerializableException;
import elt.error.icdTest.CiiSerializationTestSync;
import elt.log.CiiLogManager;
import elt.mal.CiiFactory;
import elt.mal.Mal;
import elt.mal.rr.Ami;
import elt.mal.rr.qos.QoS;
import elt.mal.rr.qos.ReplyTime;
import elt.mal.zpb.ZpbMal;

public class ExceptionSerializationClientTest {

    private static final Logger logger = CiiLogManager.
↳getLogger(ExceptionSerializationClientTest.class.getName());
    private static final String NO_EX_CALL = "first";
    private static StringBuilder outData = new StringBuilder();

    public static void main(String[] args) {
        logger.debug("Java test zpb client started.");

        final CiiFactory factory = CiiFactory.createInstance();

        final URI uri = URI.create("zpb.rr://127.0.0.1:12081/CiiSerializationTest/
↳inst_1");
        logger.debug("URI: " + uri.toString());

        Mal mal = new ZpbMal();
        factory.registerMal("zpb", mal);

        try (CiiSerializationTestSync serialization =
            factory.getClient(uri, new QoS[] {new ReplyTime(3, TimeUnit.SECONDS)},
↳new Properties(),
            CiiSerializationTestSync.class)) {

            logger.info("Perfomrming simple call test.");
            try {
                final String reply = serialization.methodWithException(NO_EX_CALL);
                if (!reply.equals(NO_EX_CALL)) {
                    logger.error("Error!!!! Did not get expected string.");
                    System.exit(1);
                }
            }
            final String unexpectedReply =
```

(continues on next page)



(continued from previous page)

```
        serialization.methodWithException(ExceptionSerializationServerTest.  
↳EXCEPTION_CALL);  
        logger.error("Error!!!! Should not be here.");  
    } catch (CiiSerializableException e) {  
        final CiiException ex = CiiException.fromSerializable(e);  
        logger.debug("Got CiiSerializableException.");  
        logger.info(String.format("Type: %s", ex.getCiiType()));  
        logger.info(String.format("Host name: %s", ex.getCiiHostName()));  
        logger.info(String.format("File name: %s", ex.getCiiFileName()));  
        logger.info(String.format("Class name: %s", ex.getCiiClassName()));  
        logger.info(String.format("Method name: %s", ex.getCiiMethodName()));  
        logger.info(String.format("Message: %s", ex.getCiiMessage()));  
        logger.info(String.format("Creation Date: %d", ex.getCiiCreationDate()));  
        logger.info(String.format("Stack trace: %s", ex.getCiiStackTrace()));  
        logger.info(String.format("Exception stack: %s", ex.  
↳getCiiExceptionStackAsString()));  
        logger.info(String.format("Details: %s", ex.getCiiDetails()));  
  
        //the data to be checked by the integration test  
        outData.append("Type -- ").append(e.getCiiType()).append(";");  
        outData.append("Message -- ").append(e.getCiiMessage()).append(";");  
        outData.append("Exception stack -- ").append(ex.  
↳getCiiExceptionStackAsString());  
    }  
  
    logger.info("Perfomrming Ami call test.");  
    try {  
        logger.info("OK call.");  
        final Ami<String> okTasks = serialization.amiMethodWithException(NO_EX_  
↳CALL);  
        for (CompletableFuture<String> reply : okTasks) {  
            final String result = reply.get();  
            logger.info("Reply: " + result);  
            logger.info("Task complete: " + okTasks.isDone());  
        }  
  
        logger.info("Error call.");  
        final Ami<String> errorTasks = serialization.  
↳amiMethodWithException(ExceptionSerializationServerTest.EXCEPTION_CALL);  
        for (CompletableFuture<String> reply : errorTasks) {  
            final String result = reply.get();  
            logger.info("Reply: " + result);  
            logger.info("Task complete: " + errorTasks.isDone());  
        }  
        logger.error("(Ami) Error!!!! Should not be here.");  
    } catch (InterruptedException e) {  
        logger.error("Unexpected error: ", e);  
    } catch (ExecutionException e) {
```

(continues on next page)



(continued from previous page)

```
final Throwable th = e.getCause();
if (th instanceof CiiSerializableException) {
    final CiiSerializableException se = (CiiSerializableException) th;
    final CiiException ex = CiiException.fromSerializable(se);
    logger.debug("Got CiiSerializableException.");
    logger.info(String.format("Type: %s", ex.getCiiType()));
    logger.info(String.format("Host name: %s", ex.getCiiHostName()));
    logger.info(String.format("File name: %s", ex.getCiiFileName()));
    logger.info(String.format("Class name: %s", ex.getCiiClassName()));
    logger.info(String.format("Method name: %s", ex.getCiiMethodName()));
    logger.info(String.format("Message: %s", ex.getCiiMessage()));
    logger.info(String.format("Creation Date: %d", ex.
↪getCiiCreationDate()));
    logger.info(String.format("Stack trace: %s", ex.getCiiStackTrace()));
    logger.info(String.format("Exception stack: %s", ex.
↪getCiiExceptionStackAsString()));
    logger.info(String.format("Details: %s", ex.getCiiDetails()));
    } else {
        logger.error("Unexpected error: ", e);
    }
} catch (Throwable e) {
    logger.error("Unexpected throwable: ", e);
}

} finally {
    mal.close();
    factory.close();
}
logger.debug("Client done.");
System.out.print(outData.toString());
System.out.print("PASS");
}
}
```

LOGGING

Revision:	1.10
Status:	Released
Repository:	https://gitlab.eso.org/cii/info/cii-docs
Project:	ELT CII
Folder:	userManual/ciiman/src/docs
Document ID:	CSL-DOC-20-181435
File:	log.rst
Owner:	Marcus Schilling
Last modification:	Apr 27, 2023
Created:	July 5, 2019

Prepared by	Reviewed by	Approved by
Rodrigo Fernández Rodríguez (CSL)	Borut Terpinč (CSL SWE) Jan Pribošek (CSL SWE)	Gregor Čuk (CSL)

Document History



ELT CII - User Manual

Re- vi- sion	Date	Changed/re- viewed	Section(s)	Modificatio n
0.1	25.02.2020	Man- dez	All	Document creation.
1.0	06.05.2020	Man- dez/ bterpinc	All	Updated all sections and added tracing. Final review.
1.1	11.09.2020	Mar	9.2.3 9.2.2.2 9.2.3	Tracing changes. New heading 9.2.3.
1.2	28.09.2020	Man- dez	All	Updates after LAT topical review
1.3	28.09.2020	Mar	9.2.4 9.2.3	OLDB Context propagation . Mal pub/sub examples.
1.4	29.09.2020	Man- dez	4.8 5	Rearranged TOC structure. Moved String formatting and CLI tool sections
1.5	15.10.2020	Man- dez	2, 8.1 6.1.1 8.2.1.1, 8.2.1.2	Added “separation ” between logging and tracing. Added description of Filebeat multiline options. New chapters added.
1.6	01.12.2020	Man- dez	4	Fixed inconsisten cies in C++ example executables names
1.7	27.01.2021	Man- dez	2, 4, 6 , 7.1	Replaced usages of INTROOT with CIISRV_ROOT
1.8	14.03.2021	Schilli	5, 6, 7	logSend(5) now part of Tools(now 6) Added ciiLogViewer
1.9	02.02.2021	Schilli	all	corrected logsink: /var/log/elt and \$CII_LOGS
1.10	27.04.2021	Schilli	2	new section “Includes and Imports”

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is manual for the CII Logging system.

Audience

This document is aimed at those Cosylab and ESO employees involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms



API	Application Programming Interface
CII	Core Integration Infrastructure
MAL	Middleware Abstraction Layer
OLDB	Online Database
EA	Engineering Archive
CLI	Command Line Interface
GUI	Graphical User Interface
ES	Elasticsearch
FB	Filebeat
LS	Logstash
JSON	JavaScript object notation
SVN	Subversion
XML	Extensible Markup Language
YAML	YAML Ain't Markup Language
TTL	Time-To-Live

References

1. ESO, Core Integration Infrastructure Requirements Specification, ESO-192922 Version 7
2. Cosylab, ELT CII Log Transfer Document, CSL-DOC-20-181435, Version X.X
3. Cosylab, ELT CII Error handling Manual Document, CSL-DOC-20-181435, Version X.X
4. Log4j2, <https://logging.apache.org/log4j/2.x/>
5. Log4cplus, <https://sourceforge.net/p/log4cplus/wiki/Home/>
6. Python logging, <https://docs.python.org/3/library/logging.html>
7. Logstash Monitoring APIs: <https://www.elastic.co/guide/en/logstash/6.8/monitoring.html>
8. Filebeat HTTP Monitoring endpoint: <https://www.elastic.co/guide/en/beats/filebeat/6.8/http-endpoint.html>
9. Logstash-to-Logstash Communication, <https://www.elastic.co/guide/en/logstash/6.8/ls-to-ls.html>
10. Logstash Settings File, <https://www.elastic.co/guide/en/logstash/6.8/logstash-settings-file.html>
11. Log4j2 Async Logging, <https://logging.apache.org/log4j/2.x/manual/async.html>
12. Logstash Reference, <https://www.elastic.co/guide/en/logstash/6.8/index.html>
13. Filebeat Reference, <https://www.elastic.co/guide/en/beats/filebeat/6.8/index.html>
14. Elasticsearch Reference, <https://www.elastic.co/guide/en/elasticsearch/reference/6.6/index.html>
15. Kibana User Guide, <https://www.elastic.co/guide/en/kibana/6.6/index.html>
16. Logtrail, <https://github.com/sivasamyk/logtrail>



17. Checker Framework, <https://checkerframework.org/>
18. Kibana Index Patterns, <https://www.elastic.co/guide/en/kibana/6.8/index-patterns.html>
19. Lucene Query Syntax, <https://www.elastic.co/guide/en/kibana/6.8/lucene-query.html>
20. Log4cplus logging macros, https://log4cplus.sourceforge.io/docs/html/loggingmacros_8h_source.html
21. Elasticsearch cross-cluster search, <https://www.elastic.co/guide/en/kibana/6.6/management-cross-cluster-search.html>
22. OpenTracing API docs, <https://opentelemetry.io/docs/>
23. Jaeger Tracing docs, <https://www.jaegertracing.io/docs/1.17/>
24. Jaeger Client Libraries, <https://www.jaegertracing.io/docs/1.17/client-libraries/>
25. Jaeger Trace Adjusters, <https://pkg.go.dev/github.com/jaegertracing/jaeger/model/adjuster>



9.1 Overview

This document is a user manual for usage of CII Logging and tracing systems. It explains how to use the Logging Client API through ELT CII common library and GUI Applications to interact with it. It also explains how to use the different services that comprise the distributed logging CII system.

All examples in this manual are also presented in SVN code (project logging-examples).



9.2 Includes and imports

9.2.1 C++

To use the logging API from C++, the following directives are needed:

wscript (project)

```
cnf.check_wdep(wdep_name='client-api.el-common.cpp.log', uselib_store='cii-  
↳logging')  
# Transitive Dependencies:  
cnf.check_cfg(package='log4cplus', uselib_store='log4cplus', args='--cflags --  
↳libs')
```

wscript (module)

```
use=[...,  
    'cii-logging'  
    # Transitive Dependencies:  
    'log4cplus'
```

source

```
#include <ciiLogManager.hpp>
```

9.2.2 Python

To use the logging API from Python, the following directives are needed:

wscript (project)

```
(nothing to be done, just require 'cxx python cii')
```

wscript (module)

```
(nothing to be done)
```

source

```
from elt.log import CiiLogManager
```

9.3 Introduction

CII Log provides the logging and tracing services of the Control System. Logging pertains to all aspects of CII services. Its purpose is to provide basic and detailed information of the different services and systems in the CII in order to have a better understanding of the behaviour of the distributed system, and in particular to troubleshoot the system in case of problems. Figure 2-1 shows how the CII Log Service interacts with other elements of the ELT Core Integration Infrastructure. Logging is considered a basic service and doesn't use any other CII service.

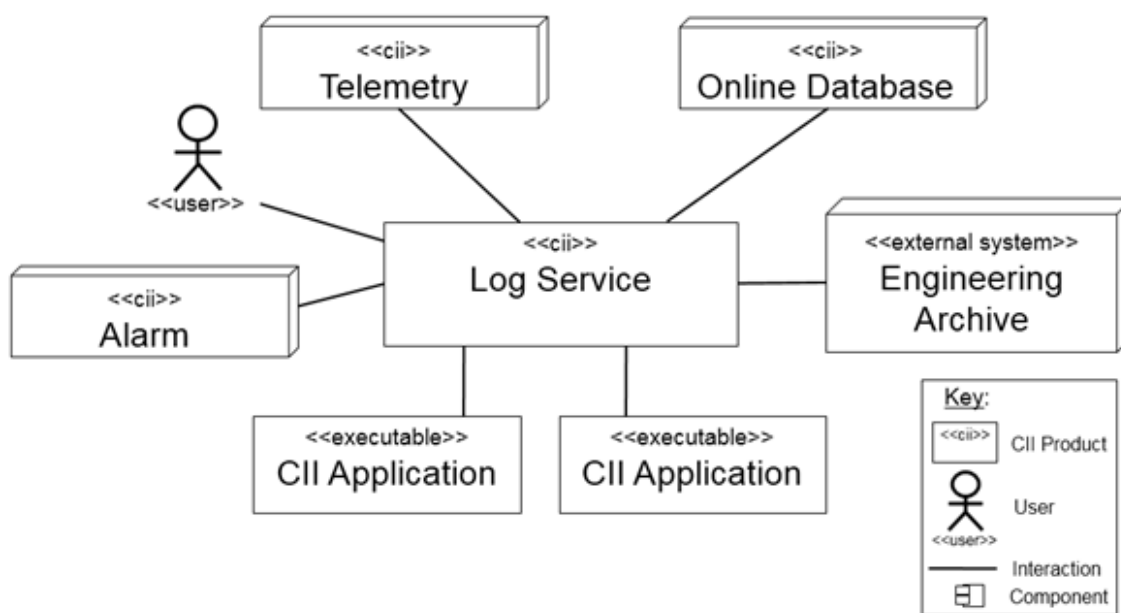


Figure 2-1 Log interaction

CII Log system components can be divided into four subcomponents according to their responsibilities:

- Client APIs
- Transmission and Management
- Persistence
- Data visualization

Figure 2-2 Standard workflow of CII Log system shows the standard workflow between the different parts of the CII Log system, and the elements that comprise each of the sub-parts.

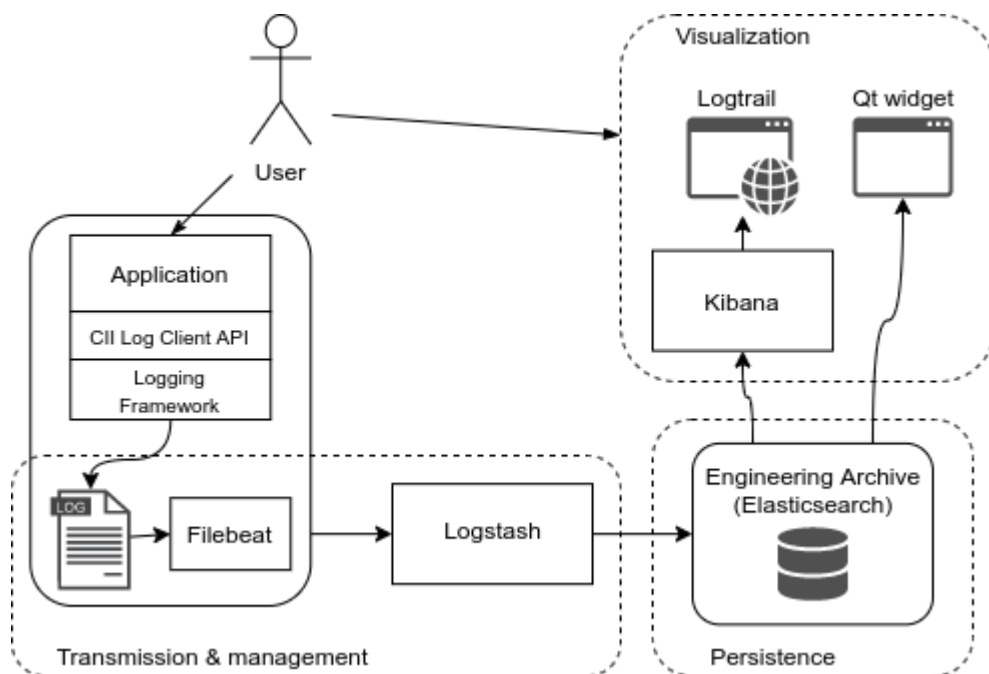


Figure 2-2 Standard workflow of CII Log system

The CII Log Client API enables users to produce log messages from their applications using Logger objects (see 2.2). These messages are produced as strings and follow a predefined format defined in the CII Logging library (see 2.3). The API leverages on existing logging frameworks and extends them in order to provide all CII functionalities. The logging frameworks it leverages on are:

- Log4j (Java) [4]
- Log4cplus (C++) [5]
- Python logging (Python) [6]

Filebeat [13] and Logstash [12] are used as transmission and management services. A Filebeat instance is responsible for harvesting CII logs from a single host (see 2.1). This log message is then wrapped as JSON format and sent to a Logstash instance.

The serialized message is not parsed to the individual log fields yet, as this is the responsibility of the Logstash instance that first receives the unparsed log event.

Apart from extracting all the CII Log fields (see Appendix B), Logstash is responsible for persisting log events to an Elasticsearch cluster, forwarding them to another Logstash instance, and/or outputting them to a file or console. In addition to this, it allows for further management and routing of the serialized log events.

Elasticsearch [14] is used as persistence backend, allowing for remote storage of log events. It is part of the same software stack as Filebeat and Logstash, therefore it provides a seamless integration of the Engineering archive with the CII Log transmission and management subsystem.

Kibana [15] with Logtrail plugin [16], as well as a Qt GUI and widget are used as visualization tools for querying and browsing CII logs.



Tracing is separated from logging, as it serves a different purpose. Tracing is explained in detail in section 8.

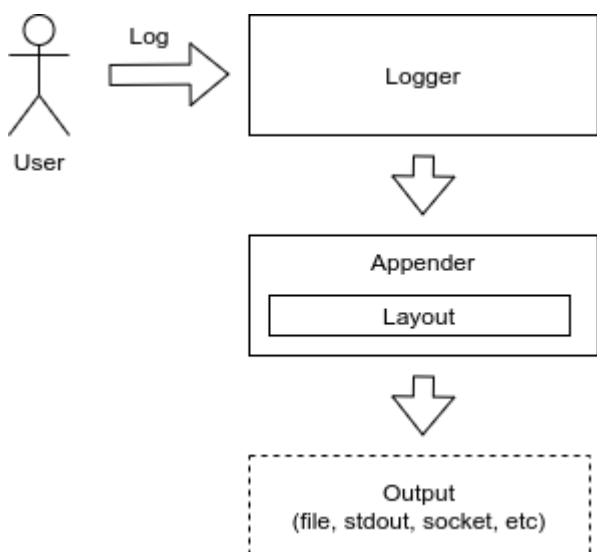
9.3.1 Logsink directory

In order to enable the CII Logging transmission and management subsystem to collect logs produced at a host using the CII Log Client API, a logsink directory is used. This directory is defined by the **\$CII_LOGS** which is a per-user setting. It is recommended to set the variable to `/var/log/elt`, since log files created here will be harvested by a Filebeat instance, if running on the host.

It is the responsibility of the system administrator to ensure that old unused log files are deleted regularly from the logsink directory. A cron job that periodically deletes log files that haven't been modified for a certain amount of time can be used for this purpose.

9.3.2 Loggers and Appenders

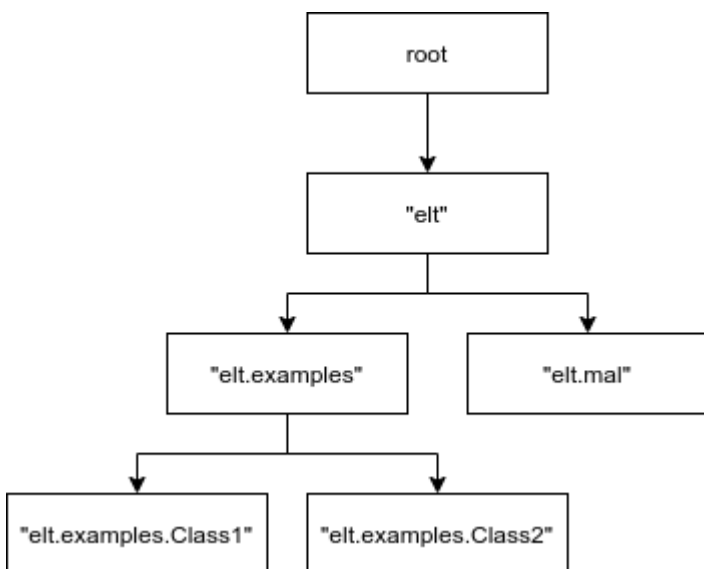
The architecture of the three used logging frameworks is relatively similar. The main elements of this architecture are the loggers, the appenders, and the layouts.



While the loggers are the elements responsible for exposing the API to create log events, appenders are responsible for writing those events to the appropriate outputs, whether it is stdout, a file, a database, a socket, or any other.

Loggers are organized in a named hierarchical structure in a parent-child relationship fashion. Each logger has a name with dot-separated components, which indicates the position of the logger in the hierarchy. An example hierarchy is shown in Figure 2-4.

There must always be a root logger with an empty name.



Appenders are attached to loggers. Each logger can have one or more appenders attached, and the same appender can be attached to multiple loggers.

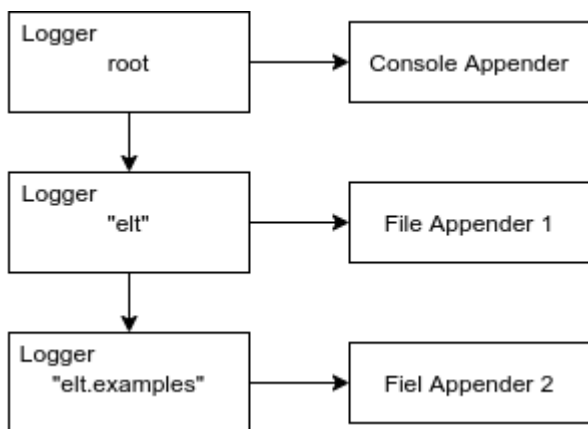


Figure 2-5 Loggers and Appenders example

Another possible scenario is having multiple child loggers with no appenders attached, having the appenders attached only to the root/parent logger. This can be useful to handle logs from different classes or packages through the child loggers but leveraging the output of those log messages to the parent/root logger and its appenders.



9.3.3 Log Layouts

Log layouts are components of the logging libraries responsible for serializing a log event into the format that will be output by the appenders, usually a string.

CII Logging library provides two different layouts that produce plain strings according to the CII Log format:

- CiiLayout
- CiiSimpleLayout

The CiiLayout is intended both for human and machine consumption, and **it must be used for logs written to the logsink directory** (see 2.1). It contains the fields described in Appendix B, in the same order as shown there.

The CiiSimpleLayout is a simplified version and is intended for human consumption. It is composed of the date and time of the log event, the log type, a log ID, the message and all the optional fields described in Appendix B.

Additional details and usage instructions can be found in section 4.4.

Example log messages using both the CiiLayout and the CiiSimpleLayout are included in section 4.5 as well as examples how to use it in configurations are included in section 4.3.

Appendix B describes the fields of which a CII log is composed. All optional fields can be added to the log message using the CII Log Client API. See section 4.7 for further details.

TraceID and SpanID are used to log information about the current tracing context (see section 8).

ErrorTypeName, ErrorDateTime, ErrorMessage, and ErrorTrace are used to log a formatted CiiException.

9.3.4 Postmortem buffer

CII Log Client API offers a logs postmortem buffer facility. It enables users to obtain on demand logging information that otherwise would have been lost due to filtering by the active logging configuration.

The postmortem buffers stores produced log events for all logging levels, regardless of the filtering configurations of the loggers and appenders.

This buffer is time-bound. Events are evicted after a defined time is elapsed. The time-to-live (TTL) for the buffered log events can be configured by the user.

The flushing of this buffer is triggered on demand through the CII Log Client API. When this happens, the buffered events are flushed to the appropriate destinations.

By default, this buffer is disabled. Configuration and usage instructions can be found in section 4.8.



9.4 Prerequisites

This section describes the prerequisites for using the CII Logging libraries and applications. Note that the preparation of the environment (i.e. installation of required modules, configuration, running of services) is not in the scope of this document. Consult your administrator and the Log TRD document for information on how to establish the environment for usage of CII Logging libraries and applications.

9.4.1 WAF Modules

For using the CII Logging and tracing libraries and CLI tools, the following modules must be built and installed with WAF in the specified order:

1. MAL, (elt-mal)
2. Client APIs (client-api)

For using the Logging Qt GUI application and Logging Qt widgets, in addition to the above, the following module must be built and installed:

3. Qt Widgets library (elt-qt-widgets)

9.4.2 Services

This section describes the services used by logging. These services should be managed and maintained by system administrators. For details on how to use the CII Log Client API, skip to section 3.3.

Logging Services

The services described below are used for transmission, management, persistence, and querying of distributed CII logs. It must be noted that they are not necessary in order to use the Logging Client API.

Distributed CII Logging management is handled using the Elastic stack, which is comprised of the following components:

- Filebeat [13]
- Logstash [12]
- Elasticsearch [14]
- Kibana [15]



Filebeat

Filebeat is a lightweight log harvester and shipper service. It is used by CII Logging system to collect logs produced by applications using the CII Log Client API. Each host producing logs that need to be collected from logsink directory requires one properly configured instance of Filebeat running.

See the Log TRD document on how to set up the Filebeat service. Configuration instructions can be found in section 6.1.

Logstash

Logstash is a server-side data processing pipeline service. It is used by the CII Logging system to ingest and parse the log events sent by one or many Filebeat instances, or to receive the log events forwarded by another Logstash instance.

These log events can then be processed by any number of output handlers, such as handlers for persistence to an `ref::Elasticsearch` instance, forwarding to another Logstash instance, writing to a file, or to stdout. How the log events are passed to the Logstash outputs can be managed based on the log event properties, such as the level or the audience.

At least one instance of Logstash is necessary to process and persist the log events collected by the Filebeat instances.

See the Log TRD document on how to set up the Logstash service. Configuration instructions can be found in section 6.2.

:: Elasticsearch:

Elasticsearch is a RESTful distributed document-oriented database and a search and analytics engine. It is used by CII Logging system as a persistence database for logs produced using the CII Log Client API in the distributed system.

It is required to have at least one Elasticsearch instance to be able to persist the log events.

See the Log TRD document on how to set up the Elasticsearch service.

Kibana

Kibana is a frontend application that provides search and data visualization capabilities for data indexed in Elasticsearch. Kibana is used by CII Logging system, together with the Logtrail plugin [16], in order to allow users to query and browse persisted CII logs.

At least one Kibana instance must be available for the user to be able to browse the persisted logs. Each Kibana instance is configured to use one Elasticsearch instance or cluster. It is also possible to configure a Kibana instance to query to other remote Elasticsearch clusters.



See the Log TRD document on how to set up the Kibana service and the Logtrail plugin. Configuration instructions for Kibana and the Logtrail plugin can be found in section 7.2.

9.4.3 Includes/Imports

For basic usage of the Logging library, the user needs the CiiLogManager class. This class contains the API for interacting with the underlying Logging framework using the CII Log extensions, whether it's Log4j (Java), Log4cplus (C++), or Python logging.

It provides methods for configuring the logging system, getting a Logger object, or flushing the post-mortem buffer. Detailed API per language is provided in Appendix A. Additional details can be found in section 4.1.

In addition, the language-specific Logger and Level classes shall be imported. Level class is not necessary if convenience log methods for the different levels are used, such as debug("message") instead of log(Level.DEBUG, "message").

C++ also needs to include ciiException class, since the ciiLogManager configuration methods can throw an exception in case of an error during configuration initialization.

For more advanced use, it is necessary include the CiiLogMessageBuilder class. The CiiLogAudience enum must be imported as well if Audience field value is provided. Also, the appropriate CiiException class must be imported.

These includes, in addition to linked libraries, are shown in the example code in logging-examples module.

Java

```
import org.apache.logging.log4j.Logger;  
import org.apache.logging.log4j.Level;  
import elt.log.CiiLogManager;  
import elt.log.CiiLogMessageBuilder;  
import elt.log.CiiLogAudience;  
import elt.log.layout.CiiSimpleLayout;  
import elt.log.layout.CiiLayout;
```

C++

```
#include <log4cplus/loggingmacros.h>  
#include <ciiLogManager.hpp>  
#include <ciiLogMessageBuilder.hpp>  
#include <ciiLogConfigurator.hpp>  
#include <ciiLogAudience.hpp>  
#include <ciiException.hpp>
```



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 291 of 589

Python

```
import logging
from elt.log import CiiLogManager
from elt.log import CiiLogMessageBuilder
from elt.log import CiiLogAudience
```



9.5 Logging Library Usage

This section describes the CII Logging Client API, and explains through examples how to configure and use the CII Logging API library in an application.

The CII Log Client API is split between two classes. The `CiiLogManager` and the `CiiLogMessageBuilder`.

Because of the fact that Cii Log Client API leverages on different libraries for Java (Log4j), C++ (Log4cplus) and Python (Python logging), and the constraints this imposes due to the different implementations and APIs, the CII Log client API methods signatures vary across languages and some methods are only available for certain languages.

A complete description of the CII Log Client API for Java, C++, and Python can be found in Appendix A.

9.5.1 CiiLogManager

This is the main class of the CII Log Client API. It provides static methods for configuring and initializing the application's logging, retrieving, creation of logger objects used to log messages, or flushing of the post-mortem appender.

The loggers retrieved using the `CiiLogManager` are used according to the APIs of the specific logging frameworks for each language [4][5][6].

9.5.2 CiiLogMessageBuilder

The `CiiLogMessageBuilder` is aimed at helping users to create correctly formatted log messages with extended information, such as:

- Audience
- Tracing context
- CII Error messages
- User-defined log ID

See section 4.7 for usage details and examples.



9.5.3 Logging Configuration

Each of the three logging libraries CII Log Client API provides different ways of configuration. The following sections explain how logging is configured for Java, C++, and Python.

If logging is initialized without providing any custom configuration, a default configuration will be used. This configuration will have:

- A root logger configured to an ERROR level with two appenders attached.
- An stdout appender using a CiiSimpleLayout layout.
- A rolling file appender to logsink directory using a CiiLayout. This appender logs to a file with the name eeltdev-output-{timestamp}.log and rotated files with file names following the pattern eeltdev-output-{timestamp}.%i.log, where %i is the index of the rotated file.

If using default configuration, logs will be output to console and logsink directory only if \$CII_LOGS environment variable is defined, otherwise logs will be only written to console and warning will be issued to the user.

Examples of default configuration files for Log4j, Log4cplus, and Python are shown in sections 4.3.1, 4.3.2, and 4.3.2.

It is relevant to notice the default path for any file creation or search is the compilation folder generated when the 'waf build' command is called. So any developer string containing a specific path file should be relative to the project build directory.

Java

Due to extensions to the Log4j library, only a subset of the configuration options is supported when using CII Log Client API. There is no support for configuring logging using Log4j's ConfigurationBuilder.

The following file configuration options are available:

- XML
- Properties
- JSON
- YAML

By default, Log4j will look for a file called log4j2.xml in the classpath. If this file is not found, it will fall back to the default CII logging configuration. If there's any issue creating the default CII configuration, it will fall back to the default Log4j configuration, which consists of a root logger with an ERROR level and a console appender attached to it.

It is possible to specify which configuration file will be loaded by Log4j by setting the system property log4j.configurationFile to the path of the configuration file.

It can be passed as a command line argument as:



```
-Dlog4j.configurationFile=path/to/log4j2configFile.xml
```

Or set at runtime before Log4j has been initialized:

```
System.setProperty("log4j.configurationFile", "path/to/log4jconfigFile.xml")
```

It is also possible to initialize a custom configuration dynamically using the Log4j Configurator class:

```
Configurator.initialize("configuration_name", "path/to/log4jconfigFile.xml")
```

Below it is shown how an XML with a configuration equivalent to the default CII configuration would like.

Listing 4-1 Example Java default configuration

```
<?xml version="1.0" ?>
<Configuration name="1584706371399" status="WARN">
  <Appenders>
    <Console name="ConsoleAppender" target="SYSTEM_OUT">
      <CiiSimpleLayout/>
    </Console>
    <RollingFile name="FilesinkFileAppender" append="true"
createOnDemand="true"
fileName="/var/log/elt/eeltdev-output-1584706371399.log"
filePattern=
  "/var/log/elt/eeltdev-output-1584706371399.%i.log">
      <Policies>
        <SizeBasedTriggeringPolicy/>
      </Policies>
      <CiiLayout/>
    </RollingFile>
  </Appenders>
  <Loggers>
    <Root level="ERROR" includeLocation="true">
      <AppenderRef ref="ConsoleAppender"/>
      <AppenderRef ref="LogsinkFileAppender"/>
    </Root>
  </Loggers>
</Configuration>
```

Note the use of CiiLayout and CiiSimpleLayout. These are loaded as plugins by Log4j on initialization. It is required for this to work as intended that class CiiLogManager is loaded by the classloader before Log4j is initialized. CII Log layouts are described in more detail in section 4.4.

To troubleshoot the initialization, it is possible to enable the debug level on the internal Log4j Status-Logger by setting system property log4j2.debug. If this property is set, Log4j will output its internal logs to console.

```
-Dlog4j2.debug=true
```



C++

In addition to the default configuration, the CII Log Client API provides support for custom configuration of the Log4cplus logging framework. Two configuration options are available:

- Properties object
- Properties file

The CiiLogManager class offers methods that can take either an object or a file containing Log4cplus configuration options. It also provides a method without arguments that uses the default CII configuration. Listing 4-2 shows an example properties configuration equivalent to the default one.

Listing 4-2 Example C++ default configuration

```
log4cplus.appender.ConsoleAppender=log4cplus::ConsoleAppender
log4cplus.appender.ConsoleAppender.layout= elt::log::layout::CiiSimpleLayout
log4cplus.appender.LogsinkFileAppender=log4cplus::RollingFileAppender
log4cplus.appender.LogsinkFileAppender.File=
    /var/log/elt/eeltdev-output-1584706371399.log
log4cplus.appender.LogsinkFileAppender.layout=elt::log::layout::CiiLayout
log4cplus.rootLogger=ERROR, ConsoleAppender, LogsinkFileAppender
```

Detailed CiiLogManager API can be found in Appendix A.

Python

Three configuration modes are supported:

- Default CII Log configuration
- Custom dictionary-based configuration
- Custom file configuration

If a file configuration is used, the file should contain the equivalent to a dictionary based configuration.

In Listing 4-3 an example of a configuration equivalent to the default configuration can be seen.

Listing 4-3 Example Python default configuration

```
'version': 1,
'cii_log_buffer': {
    'enabled': False,
    'ttl_ms': 60000
},
'handlers': {
    'cii-logsink': {
        'class': 'logging.handlers.RotatingFileHandler',
        'filename': /var/log/elt/eeltdev-output-1584706371399.log,
        'maxBytes': 10 * 1024 * 1024,
```

(continues on next page)



(continued from previous page)

```
        'backupCount': 5,  
        'formatter': 'cii'  
    },  
    'cii-console': {  
        'class': 'logging.StreamHandler',  
        'formatter': 'cii-brief',  
        'stream': 'ext://sys.stdout',  
        'level': 'ERROR'  
    }  
},  
'root': {  
    'level': 'ERROR',  
    'handlers': ['cii-logsink', 'cii-console']  
}
```

Detailed API can be found in Appendix A.

9.5.4 Log Layouts

Two CII Log Layouts are provided:

- CiiLayout
- CiiSimpleLayout

These layouts are described in section 2.3.

Examples on how to use the CII Log Layouts can be found in the configurations shown in section 4.3.

Java

The CII Log Layouts are provided as Log4j plugins, implemented by the classes CiiLayout and CiiSimpleLayout. They can be used as any other Log4j layout plugin. Usage example can be found in Listing 4-1.

C++

The CII Log Layouts are provided as Log4cplus plugins, implemented by the classes CiiLayout and CiiSimpleLayout. They can be used as any other Log4j layout plugin. Usage example can be found in Listing 4-2.



Python

To use the CII Log Layouts, the `cii` and `cii-brief` formatters can be used to enable the `CiiLogLayout` and the `CiiSimpleLayout` respectively. The definition of those formatters is injected in the configuration by the `CiiLogManager` configuration methods. Usage example can be found in Listing 4-3.

9.5.5 Basic usage example with default configuration

The following examples show a basic usage workflow of the CII Logging API to produce log messages of different levels using a default CII configuration, as described in 4.3.

Java

The example shown in Listing 4-4 can be run from the root directory of the `logging-examples` module with the `run-logging-example-default-config` provided script.

Listing 4-4 Java basic example

```
public class LoggingBasicCustomExample {

    public static void main(String[] args) {

        /* Automatically initialize the logging framework and get the root
        * logger.
        * No configuration is provided so a default configuration is used.
        * This default configuration has two appenders: one that outputs to
        * stdout using CiiSimpleLogLayout, and another one that outputs to a
        * rolling file in /var/log/elt directory
        * Only log events of levels ERROR and FATAL will be output to the
        * appenders.
        */
        Logger rootlogger = CiiLogManager.getLogger("");

        /* Send INFO level message to root logger
        * This log will be discarded since the minimum log level for the root
        * logger is ERROR
        */
        rootlogger.log(Level.INFO, "First root logger message");

        /* Send ERROR level message to root logger
        * This log will be passed to the appenders for processing, since the
        * minimum log level for the root logger is ERROR
        */
        rootlogger.log(Level.ERROR, "Second root logger message");

        /* Get a logger with name 'elt.log.test'. There's no logger configured
        * with that name, therefore it will be created inheriting the
```

(continues on next page)



(continued from previous page)

```
* configuration of it's nearest ancestor in the loggers named
* hierarchy. In this case, since there is only root logger
* configured, the configuration of the root logger will be used.
*/
Logger logger = CiiLogManager.getLogger("elt.log.test");

/* Send INFO level message to elt.log.test logger
* This log will be discarded since the minimum log level is ERROR
*/
logger.log(Level.INFO, "First elt.log.test logger message");

/* Send ERROR level message to elt.log.test logger
* This log will be passed to the appenders for processing, since the
* minimum log level is ERROR
*/
logger.log(Level.ERROR, "Second elt.log.test logger message");
}
}
```

The logs output produced by this example would be the following in stdout:

```
2020-03-19T20:05:03.104+0100, ERROR, elt.LoggingBasicDefaultExample/main, Second
↪root logger message
2020-03-19T20:05:12.571+0100, ERROR, elt.LoggingBasicDefaultExample/main, Second
↪elt.log.test logger message
```

And in log file in logsink directory:

```
Cref=LoggingBasicDefaultExample.java-31, Date=2020-03-19T20:05:03.104+0100,
↪HostName=docker-eso-eelt, LogType=ERROR, SourceID=elt.
↪LoggingBasicDefaultExample/main, LogID=elt.log.CiiLogManager.main-31,
↪Message=Second root logger message
Cref=LoggingBasicDefaultExample.java-49, Date=2020-03-19T20:05:12.571+0100,
↪HostName=docker-eso-eelt, LogType=ERROR, SourceID=elt.
↪LoggingBasicDefaultExample/main, LogID=elt.log.test.main-49, Message=Second
↪elt.log.test logger message
```

C++

This example shown in Listing 4-5 can be run with the \$CIISRV_ROOT/bin/logging-example-app executable file installed by the logging-examples module.

Listing 4-5 C++ basic example

```
int main(int ac, char *av[]) {
    try {
        /* Automatically initialize the logging framework and get the root logger.
```

(continues on next page)



(continued from previous page)

```
* No configuration is provided so a default configuration is used. This
* default configuration has two appenders: one that outputs to stdout
* using CiiSimpleLogLayout, and another one that outputs to a rolling
* file in /var/log/elt directory
* Only log events of levels ERROR and FATAL will be output to the
* appenders.
*/
log4cplus::Logger root_logger = ::elt::log::CiiLogManager::GetLogger();

std::vector<log4cplus::SharedAppenderPtr> appenders = root_logger.
↪getAllAppenders();
for (const auto& appender: appenders) {
    std::cout << "Appender:" << appender->getName() << '\n';
}

/* Send INFO level message to root logger
* This log will be discarded since the minimum log level for the root
* logger is ERROR
*/
root_logger.log(log4cplus::INFO_LOG_LEVEL, "First root logger message");

/* Send ERROR level message to root logger
* This log will be passed to the appenders for processing, since the
* minimum log level for the root logger is ERROR
*/
root_logger.log(log4cplus::ERROR_LOG_LEVEL, "Second root logger message");

/* Get a logger with name 'elt.log.test'. There's no logger configured
* with that name, therefore it will be created inheriting the
* configuration of it's nearest ancestor in the loggers named hierarchy.
* In this case, since there is only root logger configured, the
* configuration of the root logger will be used.
*/
log4cplus::Logger logger = ::elt::log::CiiLogManager::GetLogger("elt.log.test
↪");

/* Send INFO level message to elt.log.test logger
* This log will be discarded since the minimum log level is ERROR
*/
root_logger.log(log4cplus::INFO_LOG_LEVEL,
                "First elt.log.test logger message");

/* Send ERROR level message to elt.log.test logger
* This log will be passed to the appenders for processing, since the
* minimum log level is ERROR
*/
root_logger.log(log4cplus::ERROR_LOG_LEVEL,
                "Second elt.log.test logger message");
```

(continues on next page)



(continued from previous page)

```
return 0;
} catch (const ::elt::error::CiiException& ex) {
    std::cerr << "CiiException occured while executing sample code. What: "
        << ex.what() << '\n';
}
return -1;
}
```

This would produce the following output on stdout:

```
2020-04-19T21:49:41.894+0200, ERROR, root/140400261358272, Second root logger_
↪message
2020-04-19T21:49:41.894+0200, ERROR, root/140400261358272, Second elt.log.test_
↪logger message
```

And this one on the log file in logsink directory:

```
Cref=./cpp/logging-app/src/logging-app.cpp-35, Date=2020-04-19T21:49:41.
↪894+0200, HostName=docker-eso-eelt, LogType=ERROR, SourceID=140400261358272,
↪LogID=root.main-35, Message=Second root logger message
Cref=./cpp/logging-app/src/logging-app.cpp-53, Date=2020-04-19T21:49:41.
↪894+0200, HostName=docker-eso-eelt, LogType=ERROR, SourceID=140400261358272,
↪LogID=root.main-53, Message=Second elt.log.test logger message
```

Python

This example shown in Listing 4-6 can be run with the \$CIISRV_ROOT/bin/cii-logging-example-app-py executable file installed by the logging-examples module.

Listing 4-6 Python basic example

```
def main():
    """
    Main application code
    @return status code, int
    """
    result = 0
    try:
        # Automatically initialize the logging framework and get the root
        # logger.
        # No configuration is provided so a default configuration is used.
        # This default configuration has two handlers: one that outputs to
        # stdout using elt.log.CiiLogConfigurator.CII_SIMPLE_LOG_PATTERN and
        # another one that outputs to a rolling file in /var/log/elt
        # directory.
        # Only log events of levels ERROR and FATAL will be output to the
        # handlers.
```

(continues on next page)



(continued from previous page)

```
root_logger = CiiLogManager.get_logger()

# Send INFO level message to root logger
# This log will be discarded since the minimum log level for the root
# logger is ERROR
root_logger.info("First root logger message")

# Send ERROR level message to root logger
# This log will be passed to the handlers for processing, since the
# minimum log level for the root logger is ERROR
root_logger.error("Second root logger message")

# Get a logger with name 'elt.log.test'. There's no logger configured
# with that name, therefore it will be created inheriting the
# configuration of it's nearest ancestor in the loggers hierarchy. In
# this case, since there is only root logger configured,
# the configuration of the root logger will be used
logger = CiiLogManager.get_logger("elt.log.test")

# Send INFO level message to elt.log.test logger.
# This log will be discarded since the minimum log level is ERROR
logger.info("First elt.log.test logger message")

# Send ERROR level message to elt.log.test logger
# This log will be passed to the handlers for processings, since the
# minimum log level is ERROR
logger.error("Second elt.log.test logger message")
except Exception as e:
    print('Exception ocured while executing sample code: %s', e)
    result = 5
return result
if __name__ == '__main__':
    sys.exit(main())
```

This would produce the following output on stdout:

```
2020-04-19T21:49:41.894+0200, ERROR, cii-logging-example-app-py/MainThread, ↵
↵Second root logger message
2020-04-19T21:49:41.894+0200, ERROR, cii-logging-example-app-py/MainThread, ↵
↵Second elt.log.test logger message
```

And this one on the log file in logsink directory:

```
Cref cii-logging-example-app-py.main-35, Date=2020-04-19T21:49:41.894+0200, ↵
↵HostName=docker-eso-eelt, LogType=ERROR, SourceID= cii-logging-example-app-py/
↵MainThread, LogID=main-35, Message=Second root logger message
Cref= cii-logging-example-app-py.main-50, Date=2020-04-19T21:49:41.894+0200, ↵
↵HostName=docker-eso-eelt, LogType=ERROR, SourceID= cii-logging-example-app-py/
↵MainThread, LogID=main-50, Message=Second elt.log.test logger message
```



9.5.6 Basic usage example with custom configuration

This section shows how the use a custom configuration. In this case, we have a configuration with two different loggers, and each logger has a different appender, as described in Table 4-1.

Table 4-1 Example loggers and appenders

Logger name	Logger level	Appender type
"" (root)	WARN	File
"elt.log.test"	TRACE	Console

This configuration can be useful if, for example, we want to send all logs of severity WARN or higher to the logsink directory, but we are trying to troubleshoot a bug in package elt.log.test, therefore we want logs of severity as low as TRACE to be output as well for that specific package.

In certain cases, it is necessary or useful for a developer to be able to change the existing logging configuration at runtime. In the C++ and Python cases, the respective log4cplus and Python logging APIs are unaffected by the CII Log extensions of those frameworks, and appropriate documentation on runtime configuration of the frameworks is available in their respective APIs documentation [4][5][6]. In the case of Java, due to the implementation details of the postmortem buffer, it is not possible to change levels or filters of the loggers or appenders if it is enabled. If the postmortem buffer is not enabled the Log4j API is unaffected and can be modified according to the Log4j API documentation [4]. More details on the limitations on the usage of the Log4 API are described in section 4.9.1.

The example code and configurations for each language and outputs produced by the examples from section 4.4 with these configurations are shown below. These examples can be also found in logging-examples module.

Java

The example shown in Listing 4-7 uses the configuration shown in Listing 4-8.

It can be run from the root directory of the logging-examples module with the run-logging-example-custom-config provided script.

Listing 4-7 Java example app with custom configuration

```
public class LoggingBasicExampleCustomConfig {  
  
    public static void main(String[] args) {  
  
        /* Logging is initialized with custom configuration  
        * custom-example-log4j2-config.xml located in resources directory.  
        * This custom configuration has two loggers:  
        * - root: With a INFO level and a rolling file appender to  
        *   /var/log/elt directory  
        * - "elt.log.test": With a TRACE level, a console appender, and
```

(continues on next page)



(continued from previous page)

```
*      additivity enabled
*
* It is also possible to provide the path of the configuration file
* as system property log4j.configurationFile. This can be achieved
* passing it as an argument when running the application:
*   $ java -Dlog4j.configurationFile=<path/to/config/file>
* Or setting it programmatically before Log4j is initialized:
*   System.setProperty(
*       "log4j.configurationFile", "path/to/config/file");
*/
Configurator.initialize("custom-config",
    "java/src/resources/custom-example-log4j2-config.xml");

/* Get the root logger.
*/
Logger rootlogger = CiiLogManager.getLogger("");

/* Send INFO level message to root logger
* This log will be output to the log file in logsink directory
*/
rootlogger.log(Level.INFO, "First root logger message");

/* Send TRACE level message to root logger. This log message will be
* discarded due to the root logger having a threshold level of WARN
*/
rootlogger.log(Level.TRACE, "Second root logger message");

/* Get a logger with name 'elt.log.test'
*/
Logger logger = CiiLogManager.getLogger("elt.log.test");

/* Send TRACE level message to elt.log.test logger
* This log will be output to console since the minimum log level is
* TRACE.
* It will be also passed to its parent logger (root) since additivity
* property of this logger is set to true. The root logger will accept
* it also, despite of its level being set to INFO. This is due to the
* fact that the logger level threshold is only checked at the logger
* where the message is created. If a threshold level would be set to
* on the appender attached to the root logger, that one would be
* applied.
*/
logger.log(Level.TRACE, "First elt.log.test logger message");

/* Send ERROR level message to elt.log.test logger
* This log will be output to console since the minimum log level is
* TRACE. It will be also passed to its parent logger (root) since
* additivity property of this logger is set to true. The root logger
* will also pass it to the logsink file appender.
```

(continues on next page)



(continued from previous page)

```
        */  
        logger.log(Level.ERROR, "Second elt.log.test logger message");  
    }  
}
```

Listing 4-8 Log4j custom configuration

```
<Configuration status="WARN">  
  <Appenders>  
    <Console name="Console" target="SYSTEM_OUT">  
      <CiiSimpleLayout/>  
    </Console>  
    <RollingFile name="FileLogsink" append="true"  
      createOnDemand="true"  
      fileName="/var/log/elt/custom-example.log"  
      filePattern="/var/log/elt/custom-example.%i.log">  
      <Policies>  
        <SizeBasedTriggeringPolicy/>  
      </Policies>  
      <CiiLayout/>  
    </RollingFile>  
  </Appenders>  
  <Loggers>  
    <Logger name="elt.log.test" level="TRACE" additivity="true">  
      <AppenderRef ref="Console"/>  
    </Logger>  
    <Root level="INFO">  
      <AppenderRef ref="FileLogsink"/>  
    </Root>  
  </Loggers>  
</Configuration>
```

This configuration would produce the following output on console:

```
2020-04-19T23:32:38.600+0200, TRACE, elt.log.examples.  
↪LoggingBasicExampleCustomConfig/main, First elt.log.test logger message  
2020-04-19T23:32:38.600+0200, ERROR, elt.log.examples.  
↪LoggingBasicExampleCustomConfig/main, Second elt.log.test logger message
```

And the following output on the the custom-example.log file in the logsink directory:

```
Cref=LoggingBasicExampleCustomConfig.java-38, Date=2020-04-19T23:32:38.597+0200, ↵  
↪HostName=docker-eso-eelt, LogType=INFO, SourceID=elt.log.examples.  
↪LoggingBasicExampleCustomConfig/main, LogID=.main-38, Message=First root ↵  
↪logger message  
Cref=LoggingBasicExampleCustomConfig.java-58, Date=2020-04-19T23:32:38.600+0200, ↵  
↪HostName=docker-eso-eelt, LogType=TRACE, SourceID=elt.log.examples.  
↪LoggingBasicExampleCustomConfig/main, LogID=elt.log.test.main-58, ↵  
↪Message=First elt.log.test logger message
```

(continues on next page)



(continued from previous page)

```
Cref=LoggingBasicExampleCustomConfig.java-66, Date=2020-04-19T23:32:38.600+0200, ↵  
↵HostName=docker-eso-eelt, LogType=ERROR, SourceID=elt.log.examples.  
↵LoggingBasicExampleCustomConfig/main, LogID=elt.log.test.main-66, ↵  
↵Message=Second elt.log.test logger message
```

C++

This example shown in Listing 4-9 can be run with the \$CIISRV_ROOT/bin/logging-example-customconfig-app executable file installed by the logging-examples module.

Listing 4-9 C++ example app with custom configuration

```
int main(int ac, char *av[]) {  
    try {  
        /*  
        * Logging is initialized with custom configuration.  
        *  
        * This custom configuration has two loggers:  
        *   - root: With a INFO level and a rolling file appender to  
        *           /var/log/elt directory  
        *   - "elt.log.test": With a TRACE level, a console appender, and  
        *                       additivity enabled  
        */  
        ::elt::log::CiiLogManager::Configure("log4cplus-test.prop");  
  
        /* Get the root logger.*/  
        log4cplus::Logger root_logger = ::elt::log::CiiLogManager::GetLogger();  
  
        /* Send INFO level message to root logger  
        * This log will be output to the log file in logsink directory  
        */  
        root_logger.log(log4cplus::INFO_LOG_LEVEL, "First root logger message");  
  
        /* Send TRACE level message to root logger. This log message will be  
        * discarded due to the root logger having a threshold level of INFO  
        */  
        root_logger.log(log4cplus::TRACE_LOG_LEVEL, "Second root logger message");  
  
        /* Get a logger with name 'elt.log.test' */  
        log4cplus::Logger logger = ::elt::log::CiiLogManager::GetLogger("elt.log.test  
↵");  
  
        /*  
        * Send TRACE level message to elt.log.test logger  
        * This log will be outputted to console since the minimum log level is  
        * TRACE.  
        * It will be also passed to its parent logger (root) since additivity  
        * property of this logger is set to true. The root logger will accept it
```

(continues on next page)



(continued from previous page)

```
* also, despite of its level being set to INFO. This is due to the fact
* that the logger level threshold is only checked at the logger where the
* message is created.
* If a threshold level would be set to on the appender attached to the
* root logger, that one would be applied.
*/
logger.log(log4cplus::TRACE_LOG_LEVEL,
           "First elt.log.test logger message");

/*
 * Send ERROR level message to elt.log.test logger
 * This log will be output to console since the minimum log level is
 * TRACE.
 * It will be also passed to its parent logger (root) since additivity
 * property of this logger is set to true. The root logger will also pass
 * it to the logsink file appender.
*/
logger.log(log4cplus::ERROR_LOG_LEVEL,
           "Second elt.log.test logger message");

return 0;
} catch (const ::elt::error::CiiException& ex) {
    std::cerr << "CiiException occured while executing sample code. What: "
               << ex.what() << '\n';
}
return -1;
```

Listing 4-10 Log4cplus custom configuration

```
log4cplus.rootLogger=INFO, ROLLINGFILE
log4cplus.logger.elt.log.test=TRACE, STDOUT
log4cplus.additivity.elt.log.test=TRUE
log4cplus.appender.STDOUT=log4cplus::ConsoleAppender
log4cplus.appender.STDOUT.layout=elt::log::layout::CiiSimpleLayout
log4cplus.appender.ROLLINGFILE=log4cplus::RollingFileAppender
log4cplus.appender.ROLLINGFILE.File=/home/rfernandez/INTROOT/logsink/custom-
↪example.log
log4cplus.appender.ROLLINGFILE.MaxFileSize=16MB
log4cplus.appender.ROLLINGFILE.MaxBackupIndex=1
log4cplus.appender.ROLLINGFILE.layout=elt::log::layout::CiiLayout
```

This configuration would produce the following output on console:

```
2020-04-19T23:32:38.600+0200, TRACE, elt.log.test/139781441696448, First elt.log.
↪test logger message
2020-04-19T23:32:38.600+0200, ERROR, elt.log.test/139781441696448, Second elt.
↪log.test logger message
```

And the following output on the the custom-example.log file in the logsink directory:



```
Cref=../cpp/logging-customconfig-app/src/logging-customconfig-app.cpp-105,
↪Date=2020-04-19T23:32:38.597+0200, HostName=docker-eso-eelt, LogType=INFO,
↪SourceID=139781441696448, LogID=root.main-105, Message=First root logger
↪message
Cref=../cpp/logging-customconfig-app/src/logging-customconfig-app.cpp-125,
↪Date=2020-04-19T23:32:38.600+0200, HostName=docker-eso-eelt, LogType=TRACE,
↪SourceID=139781441696448, LogID=elt.log.test.main-125, Message=First elt.log.
↪test logger message
Cref=../cpp/logging-customconfig-app/src/logging-customconfig-app.cpp-134,
↪Date=2020-04-19T23:32:38.600+0200, HostName=docker-eso-eelt, LogType=ERROR,
↪SourceID=139781441696448, LogID=elt.log.test.main-134, Message=Second elt.log.
↪test logger message
```

Python

This example shown in Listing 4-11 can be run with the `$CIISRV_ROOT/bin/cii-logging-example-customconfig-app-py` executable file installed by the logging-examples module.

Listing 4-11 Python example app with custom configuration

```
def main():
    """
    Main application code
    @return status code, int
    """
    result = 0
    try:
        # Logging is initialize with custom configuration
        #
        # This custom configuration has two loggers:
        # - root: With a INFO level and rolling file handler to
        #         /var/log/elt directory
        # - "elt.log.test": With a TRACE level, a console appender and
        #                   propagation enabled
        CiiLogManager.configure(CONFIG_FILE_NAME)

        # Get the root logger
        root_logger = CiiLogManager.get_logger()

        # Send INFO level message to root logger
        # This log will be forwarded to the lof file in logsink directory
        root_logger.info("First root logger message")

        # Send TRACE level message to root logger.
        # This log message will be discarded due to the
        # root logger having ad threshold level of INFO
        root_logger.trace("Second root logger message")
```

(continues on next page)



(continued from previous page)

```
# Get a logger with name 'elt.log.test'
logger = CiiLogManager.get_logger("elt.log.test")

# Send a TRACE level message to elt.log.test logger
# This log will be forwarded to console since the minimum log level is
# TRACE.
# It will be also passed to its parent logger (root) since propagate
# property of this logger is set to True (by default). The root logger
# will also pass it to the logsink file handler. This is due to the
# fact that the logger level threshold is only checked at the logger
# where the message is created.
logger.trace("First elt.log.test logger message")

# Send ERROR level message to elt.log.test logger
# This log will be displayed on console since the minimum log level is
# TRACE.
# It will be also passed to its parent logger (root) since propagate
# property of this logger is set to True (by default). The root logger
# will also pass it to the logsink file handler.
logger.error("Second elt.log.test logger message")
except Exception as e:
    print('Exception occured while executing sample code: %s', e)
    result = 5
return result
if __name__ == '__main__':
    sys.exit(main())
```

Listing 4-12 Python logging custom configuration

```
{
  "version": 1,
  "handlers": {
    "logsink": {
      "class": "logging.handlers.RotatingFileHandler",
      "backupCount": 5,
      "filename": "CII_LOGSINK_FILENAME",
      "formatter": "cii",
      "maxBytes": 10485760
    },
    "console": {
      "class": "logging.StreamHandler",
      "formatter": "cii-brief",
      "stream": "ext://sys.stdout"
    }
  },
  "root": {
    "handlers": [
      "logsink"
    ]
  }
}
```

(continues on next page)



(continued from previous page)

```
    "level": "INFO"
  },
  "loggers": {
    "elt.log.test": {
      "handlers": [
        "console"
      ],
      "level": "TRACE"
    }
  }
}
```

This configuration would produce the following output on console:

```
2020-04-19T23:32:38.600+0200, TRACE, cii-logging-example-customconfig-app-py/
↪MainThread, First elt.log.test logger message
2020-04-19T23:32:38.600+0200, ERROR, cii-logging-example-customconfig-app-py/
↪MainThread, Second elt.log.test logger message
```

And the following output on the eelt-dev-output.log file in the logsink directory:

```
Cref= cii-logging-example-customconfig-app-py.main-92, Date=2020-04-19T23:32:38.
↪597+0200, HostName=docker-eso-eelt, LogType=INFO, SourceID= cii-logging-
↪example-customconfig-app-py/MainThread, LogID=main-92, Message=First root_
↪logger message
Cref= cii-logging-example-customconfig-app-py.main-108, Date=2020-04-19T23:32:38.
↪600+0200, HostName=docker-eso-eelt, LogType=TRACE, SourceID= cii-logging-
↪example-customconfig-app-py/MainThread, LogID=main-108, Message=First elt.log.
↪test logger message
Cref= cii-logging-example-customconfig-app-py.main-115, Date=2020-04-19T23:32:38.
↪600+0200, HostName=docker-eso-eelt, LogType=ERROR, SourceID= cii-logging-
↪example-customconfig-app-py/MainThread, LogID=main-115, Message=Second elt.log.
↪test logger message
```

9.5.7 Adding optional information to log messages

Some of the fields in a CII log message are optional and have to be explicitly added to a log event by the user. As explained in section 4.2, the CII Log Client API, through the CiiLogMessageBuilder, provides a mechanism to create log messages with this additional information using a builder pattern.

This optional information can be added to a log message:

- Audience: The intended audience of the log message
- UserLogID: A log ID provided by the user.
- Formatted CiiException: A CII Exception [3] provided by the user that will be formatted into the fields:



- ErrorMessage
- ErrorTypeName
- ErrorDateTime
- ErrorTrace
- Tracing info: The builder is responsible for extracting the following data from an existing active Tracing context:
 - TraceID
 - SpanID

Once the optional parameters are defined, the user can call the `build()` method on the builder to produce the formatted string with the selected optional fields.

The builder also provides convenience factory methods for the most common logging use cases: setting a log's audience and logging a formatted CII Error.

Support for tracing information is currently not available for C++ and Python. It will be added in following releases.

Complete API can be found in Appendix A.

Sections 4.7.1, 4.7.2, 4.7.3 show usage examples for the `CiiLogMessageBuilder` on Java, C++, and Python. All examples follow the same workflow.

First, the language-specific logging frameworks are initialized with the default CII configuration and a logger named `myLogger` is retrieved.

(How to enable tracing to be added in following versions)

Then, the following logs with optional information are produced:

- A log with a message, a custom user ID, and an audience value.
- A log with a message and a formatted CII exception using the builder pattern.
- A log with a message and a formatted CII exception using the builder's convenience method.
- A log with a message and an audience value using the builder's convenience method.
- A log without message and a formatted `CiiException`. Since no message is passed, the `CiiException` error message will be used as log message field.
- A log with current tracing context information. Note that these examples will not output any tracing information, since tracing is not enabled. These examples are meant to illustrate the usage of the `CiiLogMessageBuilder`.



Java

The example shown in Listing 4-13 can be run from the root directory of the logging-examples module with the run-logging-example-message-builder-java provided script.

Listing 4-13 Adding optional info to logs (Java)

```
package elt.log.examples;

import elt.error.CiiException;
import elt.error.CiiInvalidTypeException;
import elt.log.CiiLogAudience;
import elt.log.CiiLogMessageBuilder;
import org.apache.logging.log4j.Logger;
import elt.log.CiiLogManager;

public class LoggingMessageBuilderExample {

    public static void main(String[] args) {

        //initialize Log4j with default CII configuration and retrieve
        // "myLogger" Logger
        Logger logger = CiiLogManager.getLogger("myLogger");

        //Build and log message with message value, user ID and DEVELOPER
        //audience
        logger.error(CiiLogMessageBuilder.create("Message value")
            .withUserLogId("user_log_id")
            .withAudience(CiiLogAudience.DEVELOPER)
            .build());

        CiiException e =
            new CiiInvalidTypeException("Exception error message");

        //Build and log message with formatted CiiException using builder
        //pattern
        logger.error(CiiLogMessageBuilder.create()
            .withMessage("This is and exception")
            .withCiiException(e)
            .build());

        //Build and log message with formatted CiiException using the
        //convenience method
        logger.error(CiiLogMessageBuilder.createAndBuildWithException(
            "This is an exception", e));

        //Build and log message with DEVELOPER audience
        logger.error(CiiLogMessageBuilder.createAndBuildWithAudience(
            "Message for DEVELOPER audience", CiiLogAudience.DEVELOPER));
    }
}
```

(continues on next page)



(continued from previous page)

```
//Build and log message with formatted CiiException using the
//convenience method without passing a message field value
logger.error(
    CiiLogMessageBuilder.createAndBuildWithException(null, e));

//Build and log a message with tracing context information
logger.error(CiiLogMessageBuilder.create(
    "This is a log message with tracing info")
    .withTracingInfo()
    .build());
}
}
```

The log output from this code, using the CiiSimpleLayout, would be the following:

Listing 4-14 Output CiiLogMessageBuilder usage example Java

```
2020-03-30T10:05:46.377+0200, ERROR, elt.log.examples.
↳LoggingMessageBuilderExample/main, Message value, Audience=DEVELOPER,
↳UserLogID=user_log_id

2020-03-30T10:05:46.401+0200, ERROR, elt.log.examples.
↳LoggingMessageBuilderExample/main, This is and exception, ErrorTypeName=elt.
↳error.CiiInvalidTypeException, ErrorDateTime=1585523146401,
↳ErrorMessage=Invalid type detected Exception error message, ErrorTrace=elt.
↳error.CiiInvalidTypeException
    at elt.log.examples.LoggingMessageBuilderExample.
↳main(LoggingMessageBuilderExample.java:23)

2020-03-30T10:05:46.402+0200, ERROR, elt.log.examples.
↳LoggingMessageBuilderExample/main, This is an exception, ErrorTypeName=elt.
↳error.CiiInvalidTypeException, ErrorDateTime=1585523146401,
↳ErrorMessage=Invalid type detected Exception error message, ErrorTrace=elt.
↳error.CiiInvalidTypeException
    at elt.log.examples.LoggingMessageBuilderExample.
↳main(LoggingMessageBuilderExample.java:23)

2020-03-30T10:05:46.402+0200, ERROR, elt.log.examples.
↳LoggingMessageBuilderExample/main, Message for DEVELOPER audience,
↳Audience=DEVELOPER

2020-03-30T10:05:46.402+0200, ERROR, elt.log.examples.
↳LoggingMessageBuilderExample/main, Invalid type detected Exception error
↳message, ErrorTypeName=elt.error.CiiInvalidTypeException,
↳ErrorDateTime=1585523146401, ErrorMessage=Invalid type detected Exception
↳error message, ErrorTrace=elt.error.CiiInvalidTypeException
    at elt.log.examples.LoggingMessageBuilderExample.
↳main(LoggingMessageBuilderExample.java:23)
```

(continues on next page)



(continued from previous page)

```
2020-03-30T10:05:46.403+0200, ERROR, elt.log.examples.  
↔LoggingMessageBuilderExample/main, This is a log message with tracing info
```

C++

This example can be run with the \$CIISRV_ROOT/bin/logging-example-messagebuilder executable file installed by the logging-examples module.

Listing 4-15 Adding optional info to logs (C++)

```
int main(int ac, char *av[]) {  
    try {  
        log4cplus::Logger logger = ::elt::log::CiiLogManager::GetLogger("myLogger");  
  
        //Build and log message with message value, user ID and DEVELOPER audience  
        logger.log(log4cplus::ERROR_LOG_LEVEL,  
            ::elt::log::CiiLogMessageBuilder::Create().WithMessage("Message value")  
                .WithUserLogId("user_log_id")  
                .WithAudience(::elt::log::CiiLogAudience::DEVELOPER)  
                .Build());  
  
        try {  
            CII_THROW(::elt::error::CiiException, "TEST EXCEPTION");  
        } catch (const elt::error::CiiException& e) {  
            //Build and log message with formatted CiiException using builder  
            //pattern  
            logger.log(log4cplus::ERROR_LOG_LEVEL,  
                ::elt::log::CiiLogMessageBuilder::Create()  
                    .WithMessage("This is and exception")  
                        .WithException(e)  
                        .Build());  
  
            //Build and log message with formatted CiiException using the  
            //convenience method  
            logger.log(log4cplus::ERROR_LOG_LEVEL,  
                ::elt::log::CiiLogMessageBuilder::CreateAndBuildWithException(  
                    "This is an exception", e));  
        }  
  
        //Build and log message with DEVELOPER audience  
        logger.log(log4cplus::ERROR_LOG_LEVEL,  
            ::elt::log::CiiLogMessageBuilder::CreateAndBuildWithAudience(  
                "Message for DEVELOPER audience",  
                ::elt::log::CiiLogAudience::DEVELOPER));  
  
        //Build and log a message with tracing context information  
        logger.log(log4cplus::ERROR_LOG_LEVEL,  
            ::elt::log::CiiLogMessageBuilder::Create(  
                "Message for DEVELOPER audience",  
                ::elt::log::CiiLogAudience::DEVELOPER,  
                "Tracing context information")
```

(continues on next page)



(continued from previous page)

```
        "This is a log message with tracing info")
        .WithTracingInfo(true).Build();

    return 0;
} catch (const ::elt::error::CiiException& ex) {
    std::cerr << "CiiException occured while executing sample code. What: "
        << ex.what() << '\n';
}
return -1;
}
```

The output of the code in Listing 4-15 using the CiiSimpleLayout would be the following:

Listing 4-16 Output CiiLogMessageBuilder usage example C++

```
2020-04-19T21:56:56.130+0200, ERROR, myLogger/140218499314368, Message value,
↳Audience=DEVELOPER, UserLogID=user_log_id
2020-04-19T21:56:56.303, ERROR, myLogger/140218499314368, This is and exception,
↳ErrorTypeName=elt::error::CiiException, ErrorDateTime=1587333416,
↳ErrorMessage=TEST EXCEPTION, ErrorTrace= 0#
↳elt::error::CiiException::CiiException(std::string const&) at ../elt-common/
↳cpp/error/client/src/ciiException.cpp:29
1# main at ../cpp/logging-messagebuilder-app/src/logging-messagebuilder-app.
↳cpp:25
2# __libc_start_main in /lib64/libc.so.6
3# 0x0000000000414039 in cii-logging-example-messagebuilder-app

2020-04-19T21:56:56.303+0200, ERROR, myLogger/140218499314368, This is an
↳exception, ErrorTypeName=elt::error::CiiException, ErrorDateTime=1587333416,
↳ErrorMessage=TEST EXCEPTION, ErrorTrace= 0#
↳elt::error::CiiException::CiiException(std::string const&) at ../elt-common/
↳cpp/error/client/src/ciiException.cpp:29
1# main at ../cpp/logging-messagebuilder-app/src/logging-messagebuilder-app.
↳cpp:25
2# __libc_start_main in /lib64/libc.so.6
3# 0x0000000000414039 in cii-logging-example-messagebuilder-app

2020-04-19T21:56:56.303+0200, ERROR, myLogger/140218499314368, Message for
↳DEVELOPER audience, Audience=DEVELOPER
2020-04-19 21:56:56.303, ERROR, myLogger/140218499314368, This is a log message
↳with tracing info
```



Python

This example can be run with the `$CIISRV_ROOT/bin/cii-logging-example-messagebuilder-py` executable file installed by the logging-examples module.

Listing 4-17 Adding optional info to logs (Python)

```
def main():
    """
    Main application code
    @return status code, int
    """
    result = 0
    try:
        logger = CiiLogManager.get_logger("myLogger")

        # Build and log message with message value, user ID and DEVELOPER
        # audience
        logger.error(CiiLogMessageBuilder()
                    .with_message("Message value")
                    .with_user_log_id("user_log_id")
                    .with_audience(CiiLogAudience.DEVELOPER)
                    .build())

    try:
        raise CiiException("TEST EXCEPTION")
    except CiiException as e:
        # Build and log message with formatted exception using builder
        # pattern
        logger.error(CiiLogMessageBuilder.create()
                    .with_message("This is an exception")
                    .with_exception(e)
                    .build())

        # Build and log message with formatted exception using convenience
        # method
        logger.error(CiiLogMessageBuilder
                    .create_and_build_with_exception("This is an excaption", e))

        # Build and log message with DEVELOPER audience
        logger.error(CiiLogMessageBuilder
                    .create_and_build_with_audience("Message for DEVELOPER_
↪ audience", CiiLogAudience.DEVELOPER))

        # Build and log a message with tracing context information
        logger.error(CiiLogMessageBuilder
                    .create("This is a log message with tracing info")
                    .with_tracing_info(True)
                    .build())

        # Build and log a message with keyword args
        logger.error(CiiLogMessageBuilder
                    .create_and_build_with_args(message='This is a message',
                                                user_log_id='USER_LOG_ID',
                                                audience=CiiLogAudience.DEVELOPER))
```

(continues on next page)



(continued from previous page)

```
except Exception as e:
    print('Exception occurred while executing sample code: %s', e)
    result = 5
return result

if __name__ == '__main__':
    sys.exit(main())
```

The output of the code in Listing 4-17 using the CiiSimpleLayout would be the following:

Listing 4-18 Output CiiLogMessageBuilder usage example Python

```
2020-04-19T21:56:56.130+0200, ERROR, cii-logging-example-messagebuilder-app-py/
↳MainThread, Message value, Audience=DEVELOPER, UserLogID=user_log_id
2020-04-19T21:56:56.303+0200, ERROR, cii-logging-example-messagebuilder-app-py/
↳MainThread, This is an exception, ErrorTypeName=elt.error.Exceptions.
↳CiiException, ErrorDateTime=1588727172250, ErrorMessage=TEST EXCEPTION,
↳ErrorTrace=Traceback (most recent call last):
  File "/home/rfernandez/INTROOT/bin/cii-logging-example-messagebuilder-app-py",
↳line 30, in main
    raise CiiException("TEST EXCEPTION")
elt.error.Exceptions.CiiException: TEST EXCEPTION

2020-04-19T21:56:56.303+0200, cii-logging-example-messagebuilder-app-py/
↳MainThread, This is an exception, ErrorTypeName=elt.error.Exceptions.
↳CiiException, ErrorDateTime=1588727172250, ErrorMessage=TEST EXCEPTION,
↳ErrorTrace=Traceback (most recent call last):
  File "/home/rfernandez/INTROOT/bin/cii-logging-example-messagebuilder-app-py",
↳line 30, in main
    raise CiiException("TEST EXCEPTION")
elt.error.Exceptions.CiiException: TEST EXCEPTION

2020-04-19T21:56:56.303+0200, ERROR, cii-logging-example-messagebuilder-app-py/
↳MainThread, Message for DEVELOPER audience, Audience=DEVELOPER
2020-04-19T21:56:56.303+0200, ERROR, cii-logging-example-messagebuilder-app-py/
↳MainThread, This is a log message with tracing info
2020-04-19T21:56:56.303+0200, ERROR, cii-logging-example-messagebuilder-app-py/
↳MainThread, This is a message, Audience=DEVELOPER, UserLogID=USER_LOG_ID
```

9.5.8 String formatting

Usage of formatted strings in the printf style is supported in both Java, C++, and Python.

In Java, the String.format method can be used to construct the required message. Also, the Logger instances returned by CiiLogManager.getLogger support string formatting.

The following examples would produce a log with the message “Number: 100”.



```
logger.debug(String.format("Number: %d", 100));  
logger.debug("Number: %d", 100);
```

In C++ it is supported through the LOG4CPLUS_<LEVEL>_FMT macros [20], as shown in the following example:

```
LOG4CPLUS_WARN_FMT(root, "Number: %d", 100);
```

In Python it is also supported using the language's regular string formatting constructs.

```
logger.error("Number: %d", 100)  
logger.error("Number: %d" % 100)
```

Compile-time validation

Compile-time validation of the formatted strings is supported for Java and for C++.

In the C++ case, compile time validation is automatically done if the LOG4CPLUS_<LEVEL>_FMT macros [20] are used.

In Java, this is supported through the Checker Framework [17]. This solution doesn't only apply to formatted Strings used in log messages, but to any formatted string in the source code.

The Format String Checker prevents the use of incorrect format strings in format methods such as System.out.printf, String.format, or the Logger log methods. It warns the developer if an invalid format string is written, and it warns if the other arguments are not consistent with the format string (in number of arguments or in their types).

Checker Framework library is installed in the system as part of srv-support-libs when client-apis module is installed. Once installed it can be used during compilation as shown in the example below that compiles class TestUnitLogStringCheckerTest.java:

```
export CHECKERFRAMEWORK=$CIISRV_ROOT/lib/srv-support-libs/checker-framework-3.2  
  
javac -J-Xbootclasspath/p:$CHECKERFRAMEWORK/javac.jar \  
-Xbootclasspath/p:$CHECKERFRAMEWORK/jdk8.jar \  
-cp $CHECKERFRAMEWORK/checker-qual.jar:$CIISRV_ROOT/lib/srv-support-libs/*:  
↪$CIISRV_ROOT/lib/ \  
-processorpath $CHECKERFRAMEWORK/checker.jar \  
-processor org.checkerframework.checker.formatter.FormatterChecker \  
-Astubs=$CHECKERFRAMEWORK/Logger.astub \  
-source 8 -target 8 TestUnitLogStringCheckerTest.java
```

In Python this is not supported, since it is not compiled.



9.5.9 Using the Postmortem Buffer

This section describes the usage of the postmortem buffer. By default, this buffer is disabled and it is the responsibility of the user to enable it.

The postmortem buffer is configured through two environment variables:

- `CII_LOG_BUFFER_ON`: If set to true, the postmortem buffer will be enabled. If not set, or set to any other value, the postmortem buffer will be disabled.
- `CII_LOG_BUFFER_TTL_MS`: Time-To-Live in milliseconds of the buffered log events.

Changes to these environment variables, once the logging has been initialized, will not have any effect on the current configuration. In order for it to be applied, logging must be shut down and initialized again.

Programatic and file-based configurations of the postmortem buffer will take precedence over these environment variables.

Due to the differences in implementations and APIs of the logging frameworks CII Log Client API leverages on, the implementation and usage of the postmortem buffer vary from language to language. Configuration and usage instructions are described for each of the three languages in the following subsections.

Note: Enabling the postmortem buffer can make significant performance impact. The impact depends on the number of produced logs. When it is, enabled, all the logging events (regardless of its level) will gather all the log context information (hostname, line number, file name, etc...). Gathering this information will take significant time, when a lot of logs are produced.

Java

In addition to the configuration of the postmortem buffer through environment variables, it is possible to configure it using system properties.

- `cii.log.buffer.on`: This property defines if the buffer is enabled or not. By default, it is disabled. Must be set to true to enable the buffer.
- `cii.log.buffer.ttl.ms`: This property defines the TTL value in milliseconds for entries in the buffer. Default is 60000 ms (1 min).

These properties must be set before initialization of the Log4j framework in order for them to have effect.

If any of these system properties are defined, it will override the environment variable for that configuration parameter.

Once the logging framework is initialized, it is not possible to enable or disable the postmortem buffer.

If the postmortem buffer is enabled, **once the configuration is initialized it is not possible to perform any changes to the active configuration using the Log4j API**. The `Logger.getLevel()` method or any method to retrieve levels or filters configuration at runtime can return an incorrect



value and should not be relied upon.

The example in Listing 4-19 describes a standard workflow using the postmortem buffer.

In the first place, the `cii.log.buffer.on` property is set in order to enable the postmortem buffer. Also, the `cii.log.buffer.ttl.ms` property is set to a custom value of 30 seconds, setting the events TTL in the buffer to that value.

Then, a logger object is obtained with the name of the current class. Since no custom configuration has been defined, it will automatically initialize the logging framework to the default CII configuration. Default configuration only has the root logger configured; therefore, the retrieved object will use the root logger configuration, which filters out logs from levels below ERROR.

After, a log message is produced for every log level, from TRACE to FATAL. Only logs of levels ERROR and FATAL will be logged by the appenders.

Finally, the `CiiLogManager.flushBufferedAppender()` method is called. This will output all the buffered logs to all appenders that originally rejected them, whether it was filtered out by the logger or by the appender itself. Listing 4-20 shows the console output of the execution of the example. As can be seen, first, the ERROR and FATAL messages are printed to console, and then, when the postmortem buffer is flushed, the rest of the log messages are printed to the console.

The example shown in Listing 4-19 can be run from the root directory of the logging-examples module with the `run-logging-example-postmortem-buffer-java` provided script.

Listing 4-19 Example postmortem appender (Java)

```
package elt.log.examples;

import elt.log.CiiLogManager;
import org.apache.logging.log4j.Level;
import org.apache.logging.log4j.Logger;

public class PostmortemAppenderExample {

    public static void main(String[] args) {

        // Enable the postmortem appender
        // This must be done before initializing Log4j
        System.setProperty("cii.log.buffer.on", "true");
        // Set the time-to-live of the log events in the buffer to 30s
        System.setProperty("cii.log.buffer.ttl.ms", "30000");

        // Automatically initialize the logging framework with console and
        // file appender and get the logger with the name
        // elt.log.examples.PostmortemAppenderExample
        Logger logger = CiiLogManager.getLogger(
            PostmortemAppenderExample.class.getCanonicalName());

        // Create a log message of each level. Only messages of levels ERROR
```

(continues on next page)



(continued from previous page)

```
// and FATAL will be output to the console and file
logger.log(Level.TRACE, "Trace message");
logger.log(Level.DEBUG, "Debug message");
logger.log(Level.INFO, "Info message");
logger.log(Level.WARN, "Warn message");
logger.log(Level.ERROR, "Error message");
logger.log(Level.FATAL, "Fatal message");

try {
    Thread.sleep(500);
} catch (InterruptedException e) {
    e.printStackTrace();
}

// Flush the postmortem appender. Logs of levels TRACE, DEBUG, INFO,
// and WARN will be now output to the console and file
CiiLogManager.flushPostmortemBuffer();
}
```

This example produces the following output:

Listing 4-20 Console output from postmortem example (Java)

```
2020-03-12T22:33:14.524+0100, ERROR, elt.log.examples.PostmortemAppenderExample/
↳main, Error message
2020-03-12T22:33:14.524+0100, FATAL, elt.log.examples.PostmortemAppenderExample/
↳main, Fatal message
2020-03-12T22:33:14.524+0100, TRACE, elt.log.examples.PostmortemAppenderExample/
↳main, Trace message
2020-03-12T22:33:14.524+0100, DEBUG, elt.log.examples.PostmortemAppenderExample/
↳main, Debug message
2020-03-12T22:33:14.524+0100, INFO, elt.log.examples.PostmortemAppenderExample/
↳main, Info message
2020-03-12T22:33:14.524+0100, WARN, elt.log.examples.PostmortemAppenderExample/
↳main, Warn message
```

Async Logging

Log4j supports asynchronous logging, as described in [11]. It's important to note that when using custom configurations, **the includeLocation property of loggers' configuration must be explicitly set to true**. Otherwise, information as the source code file and line references will be lost. It is also necessary for the postmortem buffer to work as intended since it relies on a log event's context data, which would otherwise be lost.

This also needs to be taken into consideration when flushing the buffer, since it's possible that some of the last produced log events might have not been processed yet by the buffer at the time of flushing when asynchronous logging is used.



C++

In C++, the postmortem buffer can be configured using the environment variables mentioned earlier, as well as programmatically. The CiiLogConfigurator class provides methods to initialize a configuration with the postmortem buffer enabled, as well as a method to enable the postmortem buffer once the configuration has been already initialized.

When the postmortem buffered is flushed, log messages indicating the START and END of the flushed logs will be injected, as can be seen in Listing 4-22, helping the user identify the logs that were buffered by the postmortem facility.

More details on how to configure the postmortem buffer can be found in Appendix A.2.

A usage example of the postmortem buffer in C++ is shown in Listing 4-21. In this example, the postmortem buffer is enabled and a log TTL of 30 seconds is set. Then, logging is initialized with the default configuration. A series of log messages of different levels are produced, and finally, the postmortem buffer is flushed to all the configured appenders.

This example can be run with the \$CIISRV_ROOT/bin/logging-example-postmortem-buffer-app executable file installed by the logging-examples module.

Listing 4-21 Example postmortem appender (C++)

```
int main(int ac, char *av[]) {
    try {
        //Enable postmortem appender and set the time-to-live of the log events in
        //the buffer to 30s
        std::string postmortem_enable_env = std::string(
            ::elt::log::CiiLogManager::POSTMORTEM_APPENDER_ENABLED_ENV_VAR) +
            "=true";
        putenv(const_cast<char*>(postmortem_enable_env.c_str()));

        std::string postmortem_ttl_env = std::string(
            ::elt::log::CiiLogManager::POSTMORTEM_APPENDER_ENTRYTTL_MS_ENV_VAR) +
            "=30000";
        putenv(const_cast<char*>(postmortem_ttl_env.c_str()));

        // Automatically initialize the logging framework with console and file
        // appender and get the logger with the name
        // elt.log.examples.PostmortemAppenderExample
        log4cplus::Logger logger = ::elt::log::CiiLogManager::GetLogger(
            "PostmortemAppenderExample");

        // Create a log message of each level
        // Only messages of levels ERROR and FATAL will be output to the console
        // and file
        logger.log(log4cplus::TRACE_LOG_LEVEL, "Trace message");
        logger.log(log4cplus::DEBUG_LOG_LEVEL, "Debug message");
        logger.log(log4cplus::INFO_LOG_LEVEL, "Info message");
        logger.log(log4cplus::WARN_LOG_LEVEL, "Warn message");
        logger.log(log4cplus::ERROR_LOG_LEVEL, "Error message");
    }
}
```

(continues on next page)



(continued from previous page)

```
logger.log(log4cplus::FATAL_LOG_LEVEL, "Fatal message");
std::this_thread::sleep_for(std::chrono::seconds(1));

// Flush the postmortem appender.
// Logs of levels TRACE, DEBUG, INFO, and WARN will be now output to the
// console and file
::elt::log::CiiLogManager::FlushPostmortemBuffer();

return 0;
} catch (const ::elt::error::CiiException& ex) {
    std::cerr << "CiiException occured while executing sample code. What: "
        << ex.what() << '\n';
}
return -1;
}
```

This example produces the following output:

Listing 4-22 Console output from postmortem example (C++)

```
2020-04-19T22:27:50.415+0100, ERROR, PostmortemAppenderExample/140577659003584, ↵
↵Error message
2020-04-19T22:27:50.415+0100, FATAL, PostmortemAppenderExample/140577659003584, ↵
↵Fatal message
2020-04-19T22:27:51.416+0100, TRACE, root/140577659003584, [START] DUMPING ↵
↵POSTMORTEM LOGS
2020-04-19T22:27:50.415+0100, TRACE, PostmortemAppenderExample/140577659003584, ↵
↵Trace message
2020-04-19T22:27:50.415+0100, DEBUG, PostmortemAppenderExample/140577659003584, ↵
↵Debug message
2020-04-19T22:27:50.415+0100, INFO, PostmortemAppenderExample/140577659003584, ↵
↵Info message
2020-04-19T22:27:50.415+0100, WARN, PostmortemAppenderExample/140577659003584, ↵
↵Warn message
2020-04-19T22:27:51.416+0100, TRACE, root/140577659003584, [END] DUMPING ↵
↵POSTMORTEM LOGS
```

Python

The postmortem buffer can be configured by one of the methods provided by the `elt.log.CiiLogManager` module. Note that once the buffered logger is configured into the system, it can not be disabled.

The standard Python Logging API is unaffected by the use of the postmortem buffer.

The example in Listing 4-23 shows an example usage of the postmortem buffer in Python.

More details on how to configure the postmortem buffer can be found in Appendix A.3.

An usage example of the postmortem buffer in Python is shown in Listing 4-23. In this example, the



postmortem buffer is enabled and a log TTL of 30 seconds is set. Then, logging is initialized with the default configuration. A series of log messages of different level are produced, and finally the postmortem buffer is flushed to all the configured handlers.

This example can be run with the `$CIISRV_ROOT/bin/cii-logging-example-postmortem-buffer-app-py` executable file installed by the `logging-examples` module.

Listing 4-23 Example postmortem buffer (Python)

```
def main():
    """
    Main application code
    @return status code, int
    """
    result = 0
    try:
        # Enable postmortem functionality and set time-to-live of the log events
        # in the buffer to 30s
        os.environ[CiiLogManager.POSTMORTEM_BUFFER_ON_ENV_VAR] = 'true'
        os.environ[CiiLogManager.POSTMORTEM_BUFFER_TTL_MS_ENV_VAR] = '30000'

        # Automatically initialize the logging framework with console and file
        # handler and get the logger with name 'PostmortemBufferExample'
        logger = CiiLogManager.get_logger("PostmortemBufferExample")

        # Create a log message of each level
        # Only messages of levels ERROR and FATAL will be outoput to the console
        # and file

        logger.trace("Trace message")
        logger.debug("Debug message")
        logger.info("Info message")
        logger.warning("Warning message")
        logger.error("Error message")
        logger.fatal("Fatal message")

        # Wait a little
        time.sleep(1)

        print('Flushing postmortem buffer')
        # Flush postmortem buffer, trace, debug and info messages should be seen
        # on stdout
        CiiLogManager.flush_postmortem_buffer()

    except Exception as e:
        print('Exception occured while executing sample code: %s', e)
        result = 5
    return result

if __name__ == '__main__':
    sys.exit(main())
```



This example produces the following output:

```
2020-05-06T11:45:02.407+0100, ERROR, cii-logging-example-postmortem-buffer-app-  
↳py/MainThread, Error message  
2020-05-06T11:45:02.407+0100, CRITICAL, cii-logging-example-postmortem-buffer-  
↳app-py/MainThread, Fatal message  
2020-05-06T11:45:02.407+0100, TRACE, cii-logging-example-postmortem-buffer-app-  
↳py/MainThread, Trace message  
2020-05-06T11:45:02.407+0100, DEBUG, cii-logging-example-postmortem-buffer-app-  
↳py/MainThread, Debug message  
2020-05-06T11:45:02.407+0100, INFO, cii-logging-example-postmortem-buffer-app-py/  
↳MainThread, Info message  
2020-05-06T11:45:02.407+0100, WARNING, cii-logging-example-postmortem-buffer-app-  
↳py/MainThread, Warning message
```



9.6 Logging services configuration

These sections describe how to configure the Filebeat and Logstash services according to the CII log transmission & management use cases.

It is necessary to deploy a Filebeat instance on each host that produces logs that need to be persisted to Elasticsearch. Said instance can handle logs produced by any number of applications running on the host if appropriate configuration is provided.

A single instance of Logstash is necessary, which can be running on the same host or on a remote host. Multiple Logstash instances can be deployed, if needed, in different hosts for performance, infrastructure management, or any other reason. This allows for scalability of the transmission and management subsystem.

9.6.1 Filebeat

The Filebeat service needs to be configured to harvest logs from the logsink directory and then send those to a Logstash instance. An introduction to Filebeat configuration and a basic example configuration for the CII Log system, are described below.

Basic configuration example

In order to create a Filebeat configuration, it is necessary to define **inputs** and **outputs**. The **inputs** describe which files are harvested for logs and how, and the outputs describe where these harvested log events are sent. In Listing 6-1 an example serialized log event is shown. The CII Log message string is not parsed at this stage, and its raw value is stored in the message field of the JSON object output by Filebeat.

Listing 6-1 Example Filebeat log event output

```
{
  "@timestamp": "2020-01-17T10:57:07.656Z",
  "@metadata": {
    "beat": "filebeat",
    "type": "doc",
    "version": "6.8.6"
  },
  "offset": 0,
  "log": {
    "file": {
      "path": "/home/eelt-dev/INTROOT/logsink/test_logs_1.log "
    }
  },
  "prospector": {
    "type": "log"
  },
}
```

(continues on next page)



(continued from previous page)

```
"input": {
  "type": "log"
},
"beat": {
  "name": "docker-eso-eelt",
  "hostname": "docker-eso-eelt",
  "version": "6.8.6"
},
"host": {
  "name": "docker-eso-eelt"
},
"message": "Cref=LogProducer.java-16, Date=2020-01-28T17:51:12.906+0100, ↵
↵HostName=docker-eso-eelt, LogType=TRACE, SourceID=elt.log.testing.LogProducer/
↵main, LogID=elt.log.testing.LogProducer.main-16, Message=Test TRACE log message
↵",
"source": "/home/eelt-dev/INTROOT/logsink/test_logs_1.log"
}
```

In addition to inputs and outputs, it also allows for further Filebeat options customization.

A basic example configuration that harvests log files from the logsink directory and sends them to a Logstash instance running in host logstashhost on port 5044 is shown in Listing 6-2.

This example includes the following sections:

- Queue
- Inputs
- Outputs
- Logging level
- HTTP metrics endpoint

Filebeat supports buffering of events between the harvesting and shipping stages. This allows for batch processing of events by the output, being able to do bulk operations to send batches of logs in a single transaction. It is configured in the **queue** section.

In the inputs section, multiline options can be defined as well. This allows for harvesting log events that span across multiple lines as a single log event. In the CII Log case, this is used for log events that include stack traces. Specific multiline configuration for CII log events is shown in Listing 6-2. For further details on multiline options, refer to the Filebeat documentation [13].

It is possible to define the log level for the logs produced by the Filebeat process by setting the property logging.level to the desired level. If this property is not set, the default value used is INFO.

The http.enabled property is used to enable the Filebeat HTTP endpoint. Usage of this endpoint to obtain Filebeat metrics is described in section 6.3.1.

By default, the internal Filebeat logs are output in rotating files in /var/log/filebeat directory (see Table 6-1).



Listing 6-2 Basic Filebeat configuration

```
##### General #####
# Internal queue configuration for buffering events to be published.
queue:
  # Queue type by name (default 'mem')
  # The memory queue will present all available events (up to the outputs
  # bulk_max_size) to the output, the moment the output is ready to process
  # another batch of events.
mem:
  # Max number of events the queue can buffer.
  events: 4096

  # Hints the minimum number of events stored in the queue,
  # before providing a batch of events to the outputs.
  # The default value is set to 2048.
  # A value of 0 ensures events are immediately available
  # to be sent to the outputs.
  flush.min_events: 500

  # Maximum duration after which events are available to the outputs,
  # if the number of events stored in the queue is < min_flush_events.
  flush.timeout: 1s

##### Filebeat inputs #####
# This input defines which files will be harvested and the rules to do so
filebeat.inputs:

- type: log

  # Set to true to enable this input configuration.
  enabled: true

  # Paths that should be crawled and fetched. Glob based paths.
  paths:
    - /var/log/elt/*.log

  ### Multiline options

  # Multiline can be used for log messages spanning multiple lines, like
  # CII log messages including an error trace
  # This pattern matches any line starting with Cref
  multiline.pattern: ^Cref

  # Defines if the pattern set under pattern should be negated or not.
  multiline.negate: true

  # Match is set to "after" to tell logstash to append lines that don't
  # match the pattern to the last matched line
  multiline.match: after
```

(continues on next page)



(continued from previous page)

```
# This value defines the time after which new lines not matching the
# pattern will stop being appended to the last event and will
# start being processed as new events
multiline.timeout: 1s

#==== Outputs =====
# Configuration of Logstash output to use when sending the data collected.
output.logstash:
  # The Logstash hosts
  hosts: ["logstashhost:5044"]

#==== Logging =====
# Sets log level. The default log level is info.
# Available log levels are: error, warning, info, debug
logging.level: debug

#==== HTTP Endpoint =====
# Defines if the HTTP endpoint is enabled.
http.enabled: true
```

Directory structure

Filebeat uses different directories to run. These directories follow the layout described in Table 6-1.

Table 6-1 Filebeat directory structure

Table with 4 columns: Dir, Description, Default location, Config option. Rows include home, bin, config, data, and logs.

These default locations are applied when Filebeat is started as a systemd service. If Filebeat is started as a regular command-line process it will use the installation directory /usr/share/filebeat as path.home unless provided in the configuration options.

Path configuration options can be defined as command-line arguments or in the Filebeat configuration file.

If any of the other path configuration options are not provided, the default values shown in Table 6-2 will be used.

Table 6-2 Filebeat default path properties values



Dir	Default	Config option
home		path.home
bin	{path.home}/bin	
config	{path.home}/config	path.config
data	{path.home}/data	path.data
logs	{path.home}/logs	path.logs

By default, Filebeat configuration file is located in the {path.home}/config directory with filename filebeat.yml.

The location of the Filebeat configuration file can be also defined in command line using the -c / -c option, e.g:

```
$ filebeat -c /path/to/my/filebeatconfig.yml
```

This is an example command line input to start Filebeat defining custom values for all path configuration options. It will look for configuration files in the /home/eelt-dev/filebeat_config directory

```
$ filebeat -path.home /usr/share/Filebeat \  
> --path.configs /home/eelt-dev/filebeat_config \  
> --path.data /home/eelt-dev/filebeat_data \  
> --path.logs /home/eelt-dev/filebeat_logs
```

9.6.2 Logstash

Logstash processes log events as a pipeline. The Logstash pipeline architecture is shown in Figure 6-1 and consists of these stages:

- **Inputs.** These elements receive log events. In the CII Logging architecture (Figure 2-2), events can be received from a Filebeat instance or from another Logstash instance.
- **Filters.** Are used to process events. Here is where raw CII Log events sent by Filebeat are parsed and their individual fields extracted.
- **Outputs.** In this stage, the processed events can be sent to any number of outputs

Codecs are used on the inputs and outputs to handle different types of serialization, such a JSON, single-line string, or multi-line string. Unless stated otherwise, all inputs and outputs use a JSON codec by default.






Figure 6-1 Logstash pipeline architecture

A default setup for Logstash in the CII Logging system consists of multiple Filebeat instances collecting logs and sending them to a Logstash instance, which then persists the received log events to Engineering Archive as depicted on Figure 6-2.




Figure 6-2 Typical log events data flow

A configuration for Logstash for this default CII configuration is shown in Listing 6-4. Example events before and after being processed by Logstash are shown in Listing 6-1 and Listing 6-3 respectively

Listing 6-3 Example Logstash parsed log event output

```
{
  "source" => "/home/eelt-dev/INTROOT/logsink/test_logs_1.log ",
  "beat" => {
    "name" => "docker-eso-eelt",
    "version" => "6.8.6",
    "hostname" => "docker-eso-eelt"
  },
  "tags" => [
    [0] "beats_input_codec_plain_applied"
  ],
  "prospector" => {
    "type" => "log"
  },
  "@timestamp" => 2020-01-28T17:51:12.906+0100,
  "log" => {
    "file" => {
      "path" => "/home/eelt-dev/INTROOT/logsink/test_logs_1.log "
    }
  },
  "message" => "Test TRACE log message",
  "@version" => "1",
  "host" => {
    "name" => "docker-eso-eelt"
  },
  "offset" => 0,
  "input" => {
    "type" => "log"
  },
}
```

(continues on next page)



(continued from previous page)

```
"cii" => {  
  "sourceid" => "elt.log.testing.LogProducer/main",  
  "date" => "2020-01-28T17:51:12.906+0100",  
  "logtype" => "TRACE",  
  "logid" => "elt.log.testing.LogProducer.main-16",  
  "cref" => "LogProducer.java-16",  
  "message" => "Test TRACE log message",  
  "hostname" => "docker-eso-eelt"  
}
```

In the **input** section, a single beats input plugin is configured. This input will use the 5044 port to receive log events from the Filebeat instances.

The **filter** section of this example is composed of three filter plugins: grok, mutate, and date.

As mentioned earlier, Filebeat will serialize the logs as JSON objects, but the collected raw logs from the logsink directory will be passed to Logstash as a whole string in the message field of the JSON object. It is the responsibility of Logstash to parse that string and extract the appropriate fields. To do so, a CILOG pattern installed in `#{CIISRV_ROOT}/etc/log/logstash/patterns` matching the CII Log format is used by the grok filter plugin. Since it's a custom pattern, the path to the directory containing the file with the CILOG pattern definition must be passed to the grok filter by setting the `patterns_dir` option to the appropriate path.

If the string in the event's message field fails to match the pattern, a `_grokparsefailure` tag will be added to the tags field of the log event and no fields will be extracted.

After parsing the log, if the tag `_grokparsefailure` is not found in the tags field, the event is further processed, otherwise is passed directly to the outputs. In the example configuration provided in Listing 6-4, these log messages would be persisted to the same index as the rest of the log events, but only containing the "message" field, which wraps the whole unparsed log event. This behaviour is completely configurable. An additional output could be configured to a specific ES index for failed parsing log messages.

The `cii.message` field is copied to the message field with the mutate filter, replacing the original unparsed log string.

Next, the `cii.date` field is parsed and copied to the `@timestamp` field of the event, which by default is set to the time the event gets processed by the Filebeat instance, rather than the original time of the event.

Finally, the **output** section contains a single Elasticsearch output plugin. It has two parameters defined:

- `hosts` – Defines the Elasticsearch instances to which it will write the log events. In this case, `ciilogelastichost`.
- `index` – The index to which the log events will be written. This configuration uses a pattern that includes the current date, e.g. `cii_log-2020.06.10`. This is a recommended good practice in order to avoid log indices from growing indefinitely. Usage of these patterns is discussed in



more detail in section 6.2.2.

Listing 6-4 Default Logstash CII configuration

```
# Default Logstash CII configuration
# Filebeat -> Logstash -> Elasticsearch.

input {
  beats {
    port => 5044
  }
}

filter {
  grok {
    patterns_dir => ["${CIISRV_ROOT}/etc/log/logstash/patterns"]
    match => { "message" => "%{CIILOG}" }
  }
  if "_grokparsefailure" not in [tags] {
    mutate {
      copy => {"[cii][message]" => "message"}
    }
    date {
      match => [ "[cii][date]", "ISO8601" ]
    }
  }
}

output {
  elasticsearch {
    hosts => ["ciilogelastichost"]
    index => "cii_log-%{+YYYY.MM.dd}"
  }
}
```

Forwarding log events

Another typical use case consists of a Logstash instance that forwards log events to another Logstash instance, as depicted in Figure 6-3.

In this use case, both Logstash instances receive events from different Filebeat instances, but instance 1 forwards the processed events to instance 2, which persists all received events to the Engineering Archive.

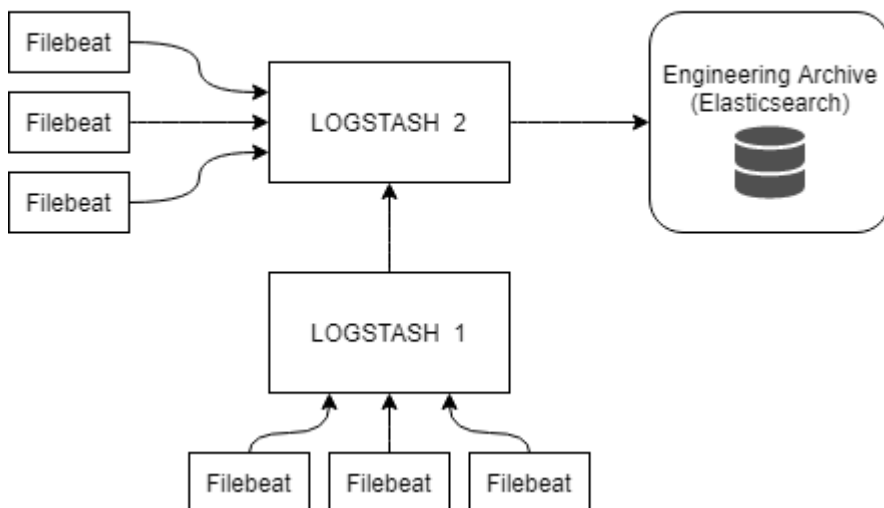


Figure 6-3 Logstash forwarding events

Configuration for Logstash instance 1 is shown in Listing 6-5 Logstash upstream forwarding configuration. The only difference with the default configuration seen in the previous section (Listing 6-4) is the outputs configuration, which is coloured in orange.

In this case, a lumberjack output plugin is used with the following configuration:

- `hosts` – The Logstash instance on `ciilogstashhost` hostname where events will be forwarded to (instance 2).
- `ssl_certificate` – Path to SSL certificate to use to connect to the downstream Logstash (instance 2). Forwarding between Logstash instances requires using a secured SSL connection. More information on the SSL configuration can be found in the Logstash documentation [9].
- `port` – TCP port on the downstream Logstash instance where events will be sent to.

Listing 6-5 Logstash upstream forwarding configuration

```
# Logstash configuration for creating a simple
# Filebeat -> Logstash -> Logstash pipeline.

input {
  beats {
    port => 5044
  }
}

filter {
  grok {
    patterns_dir => ["${CIISRV_ROOT}/etc/log/logstash/patterns"]
    match => { "message" => "%{CIILog}" }
    overwrite => ["message"]
  }
}
```

(continues on next page)



(continued from previous page)

```
}  
if "_grokparsefailure" not in [tags] {  
  mutate {  
    copy => {"[cii][message]" => "message"}  
  }  
  date {  
    match => [ "[cii][date]", "ISO8601" ]  
  }  
}  
}  
  
output {  
  lumberjack {  
    codec => json  
    hosts => "ciilogstashhost"  
    ssl_certificate => "/path/to/cert/lumberjack.cert"  
    port => 5045  
  }  
}
```

The configuration for the Logstash instance 2 is shown in Listing 6-6. In this case there are two differences with the default Logstash configuration example from Listing 6-4.

In the **inputs** section, two beats input plugins are configured. The first one to receive events from Filebeat instances, and the second one to receive events forwarded from other Logstash instances using the lumberjack protocol. SSL is enabled and fields for SSL certificate and key files are provided.

The second difference is in the **filters** section. Events that have been forwarded by another Logstash instance have already been processed. Therefore, it is not necessary to process them again. To avoid so, it checks if `cii.logtype` field is defined. If it's not, it means that the event hasn't been processed yet. In this case it is then passed to the filter plugins, otherwise it is sent directly to the outputs.

Listing 6-6 Logstash downstream forwarding configuration

```
# Logstash configuration for creating a  
# Filebeat/Logstash -> Logstash -> Elasticsearch pipeline.  
  
input {  
  # Input for Filebeat  
  beats {  
    port => 5044  
  }  
  # Input for Logstash forwarding  
  beats {  
    codec => json  
    port => 5045  
    ssl => true  
    ssl_certificate => "/path/to/cert/lumberjack.cert"  
    ssl_key => "/path/to/cert/lumberjack.key"
```

(continues on next page)



(continued from previous page)

```
}  
}  
  
filter {  
  if ! [cii][message]  
    grok {  
      patterns_dir => ["${CIISRV_ROOT}/etc/log/logstash/patterns"]  
      match => { "message" => "%{CIILog}" }  
      overwrite => ["message"]  
    }  
    if "_grokparsefailure" not in [tags] {  
      mutate {  
        copy => {"[cii][message]" => "message"}  
      }  
      date {  
        match => [ "[cii][date]", "ISO8601" ]  
      }  
    }  
  }  
}  
  
output {  
  elasticsearch {  
    hosts => ["ciilogelastichost"]  
    index => "cii_log-%{+YYYY.MM.dd}"  
  }  
}
```

Persisting Log events

As mentioned earlier, it is a good practice to define time-based indices for our logs, in the same fashion regular logging rotates log files using time-based policies.

For example, the index in a Logstash output plugin configuration can be configured as `cii_log-%{+YYYY.MM.dd}`. Which will produce indices with the current date appended to the name of the index (e.g. `cii_log-2020.06.10`).

This feature can be used to further manage log events in indices based on their audience. Listing 6-7 shows an example output configuration that uses a pattern including the `cii.audience` field value to persist the event in different audience-based indices. Events with an audience field defined will be stored in indices following the pattern `cii_log-%{[cii][audience]}-%{+YYYY.MM.dd}` (e.g. `cii_log-DEVELOPER-2020.06.10` or `cii_log-OPERATOR-2020.06.10`). Events with no audience defined will be stored in indices following the pattern `cii_log-ALL-%{+YYYY.MM.dd}` (e.g. `cii_log-ALL-2020.06.10`).

A similar approach can be taken with other CII Log fields, such as `cii.logtype`.

Listing 6-7 Audience-based output configuration



```
output {
  if ! [cii][audience] {
    elasticsearch {
      hosts => ["ciilogelastichost"]
      index => "log-ALL-%{+YYYY.MM.dd}"
    }
  } else {
    elasticsearch {
      hosts => ["ciilogelastichost"]
      index => "cii_log-%{[cii][audience]}-%{+YYYY.MM.dd}"
    }
  }
}
```

Handling Log Events

It is possible to handle events based on the value of any of their fields. One example use case of event handling would be a situation where we want logs of levels TRACE and DEBUG to be saved to a local Elasticsearch instance, and events with a level above DEBUG to be persisted to a remote Elasticsearch instance (e.g. Engineering Archive). An output configuration for this use case is shown in Listing 6-8 Logtype-based output configuration.

Listing 6-8 Logtype-based output configuration

```
output {
  if [cii][logtype] in ["TRACE", "DEBUG"] {
    elasticsearch {
      hosts => ["localhost"]
      index => "cii_log-%{+YYYY.MM.dd}"
    }
  } else {
    elasticsearch {
      hosts => ["ciilogelastichost"]
      index => "cii_log-%{+YYYY.MM.dd}"
    }
  }
}
```

Configuring Logstash

The directory layout of a Logstash installation is shown in Table 6-3.

By default, when configuring pipelines, Logstash will try to load all the files with .conf extension from the conf directory. This can be overridden when running Logstash from the command line using the -f flag with the path of the configuration file we want to load:



```
$ logstash -f <path_to_config_file>
```

The settings directory contains the logstash.yml file. This file is used to define configuration flags for Logstash. It allows defining batch sizes for the pipelines, queue type, and settings, or paths to the different configuration directories described in Table 6-3.

Additional information on the logstash.yml file can be found in the Logstash documentation

Table 6-3 Logstash directory structure

Type	Description	Default Location	Setting
home	Home directory of the Logstash installation	/usr/share/logstash	
bin	Logstash executables directory	/usr/share/logstash/bin	
settings	Configuration files, including logstash.yml, jvm.options, and startup.options	/etc/logstash	path.settings
conf	Logstash pipeline configuration files	/etc/logstash/conf.d/*.conf	Defined in /etc/logstash/pipelines.yml
logs	Log files generated by Logstash	/var/log/logstash	path.logs
plugins	Local plugin files	/usr/share/logstash/plugins	path.plugins
data	Data files used by Logstash and its plugins for any persistence needs.	/var/lib/logstash	path.data

Using multiple pipelines

Logstash supports using multiple pipelines. This allows for modularization and encapsulation of responsibilities. Pipelines can work independently or they can be connected to each other, allowing to effectively split the processing of a log event in different pipelines.

In Figure 6-4 an example setup with multiple pipelines is shown. This example consists of three pipelines:

- logsink-input: This pipeline is responsible for receiving log events from the Filebeat instances.
- forwarded-input: This pipeline is responsible for receiving log events forwarded by another Logstash instance.
- ea-output: This pipeline is responsible for persisting to the Engineering Archive log events passed on by any of the input pipelines.

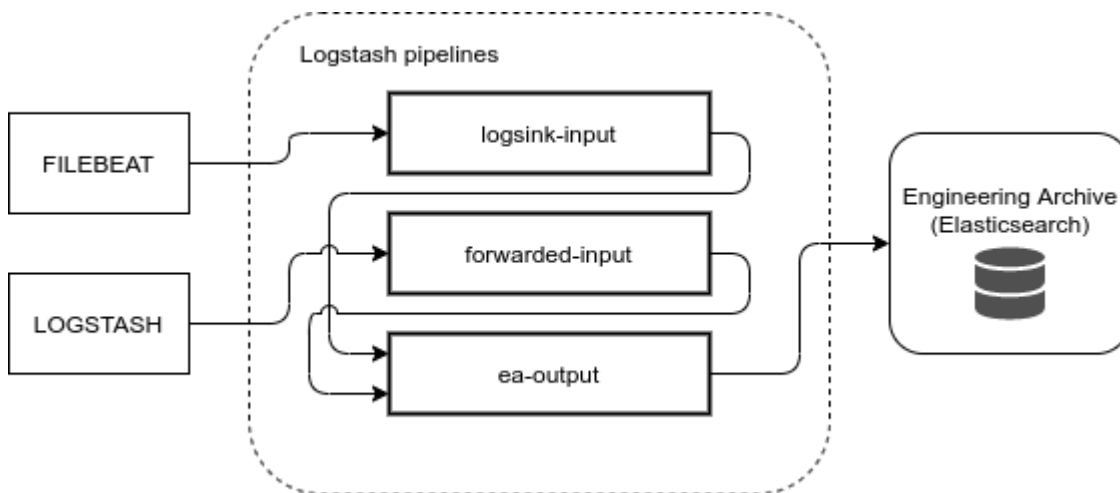


Figure 6-4 Multiple pipeline Logstash setup

Configuration of Logstash pipelines is done through the pipelines.yml file. This file must be placed in the path.settings directory. Configuration for each pipeline can be provided as a string in the pipelines.yml configuration, as shown in Listing 6-9. It is also possible to define the paths to the individual pipeline, as shown in Listing 6-10.

Listing 6-9 Logstash multiple pipelines string configuration

```
# config/pipelines.yml
- pipeline.id: logsink-input
  config.string: |
    input {
      beats {
        port => 5044
      }
    }
    filter {
      grok {
        patterns_dir => ["${CIISRV_ROOT}/etc/log/logstash/patterns"]
        match => { "message" => "%{CIILOG}" }
        overwrite => ["message"]
      }
      if "_grokparsefailure" not in [tags] {
        mutate {
          copy => {"[cii][message]" => "message"}
        }
        date {
          match => [ "[cii][date]", "ISO8601" ]
        }
      }
    }
    output {
      pipeline { send_to => ea-output }
    }
  }
```

(continues on next page)



(continued from previous page)

```
- pipeline.id: forwarded-input
  config.string: |
    input {
      beats {
        codec => json
        port => 5045
        ssl => true
        ssl_certificate => "/path/to/cert/lumberjack.cert"
        ssl_key => "/path/to/cert/lumberjack.key"
      }
    }
    output {
      pipeline { send_to => ea-output }
    }
- pipeline.id: ea-output
  config.string: |
    input { pipeline { address => ea-output } }
    output {
      elasticsearch {
        hosts => [es_cluster_b_host]
      }
    }
}
```

Listing 6-10 Logstash multiple pipelines path configuration

```
# config/pipelines.yml
- pipeline.id: logsink-input
  path.config: "/path/to/config/logsink-input.conf"
- pipeline.id: forwarded-input
  path.config: "/path/to/config/forwarded-input.conf"
- pipeline.id: ea-output
  path.config: "/path/to/config/ea-output.conf"
```

Persistent queues

Logstash supports buffering of events in queues. These queues are located between the input and the filter stages, as shown in Figure 6-1. By default, Logstash uses an in-memory queue, but it can be configured to use persistent queues, which will buffer the received events in a file.

Using persistent queues has two main benefits. It helps absorbing bursts of events that can not be processed at the same rate as they are receive; and helps avoiding loss of events in case of an abnormal termination of the Logstash service, since those buffered events will be processed once the Logstash service is started again.

Configuration of persistent queues is defined in the `logstash.yml` settings file. An example configuration is shown below. This configuration defines a persistent queue with a limit of 2000 buffered events.



```
queue.type: persisted  
queue.max_events: 2000
```

It is possible to define other parameters, like the maximum size in bytes of the buffer, or the location of the queue data files.

Additional information on persistent queues can be found in the Logstash documentation [12].

9.6.3 Logging Services Statistics

It is possible to obtain internal metrics data for the Filebeat and the Logstash instances by configuring and querying their Metrics HTTP endpoints. The responses are encoded as JSON objects.

Filebeat Metrics HTTP Endpoint

Filebeat exposes internal metrics through an HTTP endpoint.

By default, the Metrics endpoint is disabled, but can be enabled and configured in the Filebeat configuration through the setting described in Table 6-4.

Table 6-4 Filebeat HTTP Metrics endpoint properties

Property	Description	Default value
http.enabled	Enables the HTTP endpoint	false
http.host	Binds to this hostname	localhost
http.port	The HTTP endpoint will bind to this port	5066

Two types of queries can be performed:

- Basic info
- Advanced stats

The following example queries show how the metrics can be queried. First for the basic info:

```
$ curl -XGET 'localhost:5066/?pretty'
```

And then for the advanced stats

```
$ curl -XGET 'localhost:5066/stats?pretty'
```

Additional information and output examples can be found in the Filebeat documentation [8].



Logstash Monitoring APIs

The metrics provided by Logstash are more exhaustive than the ones provided by Filebeat. While Filebeat only offers one endpoint, Logstash offers one endpoint for each of the different Monitoring APIs it offers, as shown in Table 6-5.

Table 6-5 Logstash monitoring APIs

Metrics API	Description	Endpoint name
Node Info	Retrieves information about the node. Supports the following node types: <ul style="list-style-type: none">• pipelines• os• jvm	/_node/<types>
Plugins Info	Retrieves information about the currently installed Logstash plugins.	/_node/plugins
Node Stats	Retrieves runtime information about the Logstash process. Supports the following node types: <ul style="list-style-type: none">• jvm• process• events• pipelines• reloads• os	/_node/stats/<types>
Hot Threads	Retrieves information about the current hot threads for Logstash. A hot thread is a Java thread that has high CPU usage and executes for longer than a normal period of time.	/_node/hot_threads

In addition to these resources, the root endpoint can be queried to obtain general information about the Logstash instance, as is shown in the query and response below. The ?pretty parameter is optional for any of the queries, and if used the JSON response object from Logstash is pretty-printed, instead of being returned as a one-line JSON.



```
curl -XGET logstashhost:9600/?pretty
```

Configuring Monitoring APIs port

By default, the Logstash Monitoring API is bound to the port 9600, but it is possible to override this value and set it to a custom TCP port number with the command line flag `-http.port`.

Additional information on the Monitoring APIs can be found in the Logstash documentation [7].



9.7 Logging User Tools

This section describes the different user interface tools that enable users to create, query, and browse CII Log events.

9.7.1 logSend

This section describes how to use the logSend command line tool. This tool allows a user to create a log entry in the logsink directory, where it can be collected by the CII Logging transmission and management subsystem. The tool will create a file with the name eltdev-log-output-logsend-<current_date>.log and append the log message to it.

This tool takes the Message field value and accepts parameters for all the rest of the CII Log format fields as described in Appendix B. For the required fields, the "N/A" value is used if no value is provided, except for the LogType field, for which the INFO level will be used, and for the HostName field, for which the current hostname will be used. The Date field is populated with the current time. Table 5-1 shows the mappings between the sendLog command flags and the CII Log fields.

Table 6-1 sendLog CLI tool flags mappings

Table with 2 columns: Flag and Field name. It lists various CLI flags like -h, -help, -lt, -logtype, etc., and their corresponding field names like LogType, Cref, SourceID, etc.

A basic usage example can be seen below:

```
$ logSend "This is a log message"
```

Which will create the following log:

```
Cref=N/A, Date=2020-03-27T19:44:57.872+0200, HostName=docker-eso-eelt,
LogType=INFO, SourceID=N/A, LogID=N/A, Message=This is a log message
```

This is a more advanced example, where several of the log fields are passed as arguments:



```
$ logSend "This is a log message" -lt ERROR -cr example_code_reference -si \
example_source_id -li example_log_id -au DEVELOPER
```

Which will produce the following log:

```
Cref=example_code_reference, Date=2020-03-27T19:50:28.739+0200, HostName=docker-
eso-eelt, LogType=ERROR, SourceID=example_source_id, LogID=example_log_id,
Message=This is a log message, Audience=DEVELOPER
```

9.7.2 Log Viewer

This section describes usage of the ciiLogViewer application, which is a graphical viewer application for logs.

In "follow" mode, it detects all newly created log messages on the local host, found in the local log sink directory (defined by \$CII_LOGS), and displays them in an automatically refreshing table.

The tool is new in version 2.0 of CII Services, and in this version only supports a limited set of features.

Table 6-2 log viewer command line options

Table with 2 columns: Flag, Description. Rows include -h, -help; -f, -follow.

9.7.3 Log Viewer Widget (Qt)

This is a Qt widget, meant for integration into an ELT application, allowing to display and filter ELT log data. Log data can be stored locally (i.e. in log files) or remotely (i.e. in Elasticsearch server). It also allows the user to create a log entry in the local logsink directory.

To start a demo application containing the widget, run the qt-widgets-log-dialog-sample-app command from a terminal. When the widget is displayed by an application, it opens the main window in local mode. Initially two files are open: the most recently modified file from \$CII_LOGS directory and the default output file for current date in the same directory. The default output file is used as default destination by the widget when an ad-hoc log entry is created. This file name follows the pattern default-eeltdev-output-{YYYYMMDDHHMMSS}.log.



Main Window

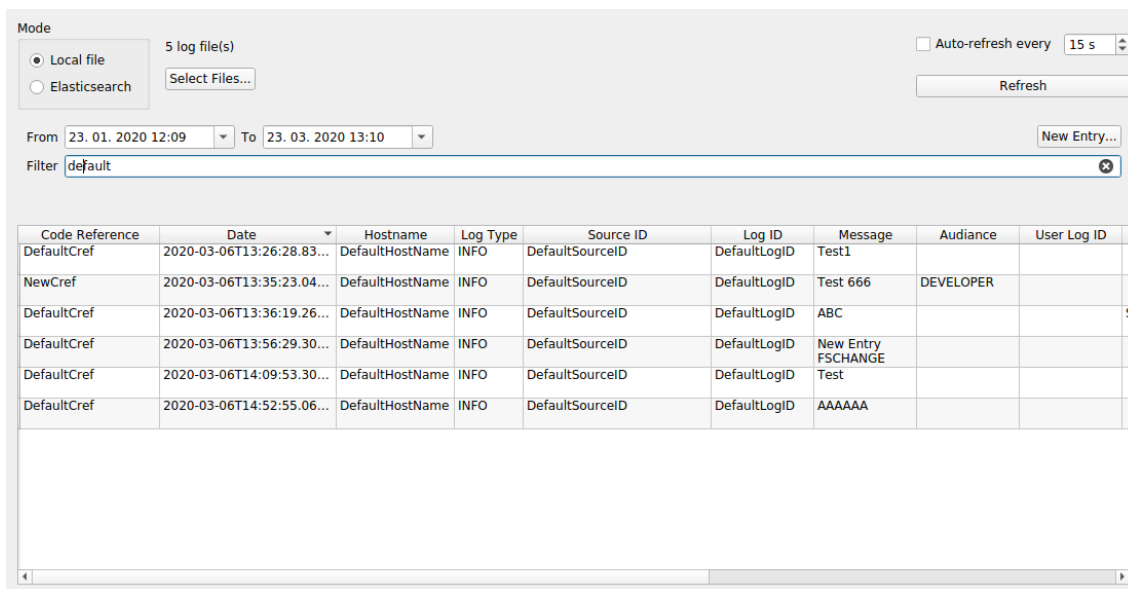


Figure 6-1 LogViewer main window (Local)

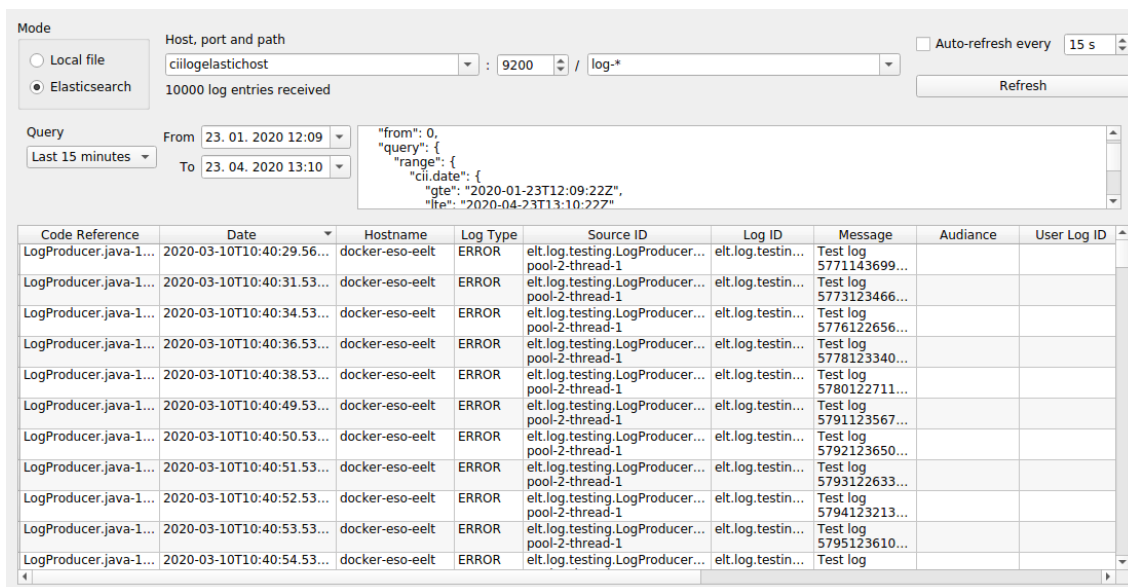


Figure 6-2 LogViewer main window (Remote)

The main window contains four major parts:

- The Mode toggle control is located in the upper-left corner. It allows switching between a local file and Elasticsearch mode. In Local file mode, the log entries are read from local log files. In Elasticsearch mode, log entries are retrieved from an Elasticsearch server (local or remote). It is possible to use a search query to filter the browsed log events.



- To the right of the mode selection controls, some information is displayed about the current mode. Mode-specific settings are displayed as well next to the mode selection controls.
 - In local mode, you can see the number of open log files (if you hover the mouse over the label you can see the list of file paths in the tooltip). To select log files that are open, click the Select Files button.
 - In Elasticsearch mode this part displays host name/address, port, and path for the server. You can edit those fields and when a query succeeds new host and path are added to the drop-down list for quick future selection. The label below server settings displays the basic status of query and possible network errors.




Figure 6-3 Log Widget remote mode settings

- The Refresh controls are located in the upper-right corner. The Refresh button can be clicked to manually retrieve the log data from an appropriate source and re-display it in the log entries table. To enable periodic automatic refresh of the log data, tick the Auto-refresh option, and enter the desired time period in seconds. This time will be used as the refresh frequency for the logs data. When this option is enabled in local mode, it also enables refreshing of log's data each time a displayed log file is changed on the hard disk (when some external tool or application writes to it).




Figure 6-4 Log Widget refresh controls

- In the middle part, Filtering and query controls are located. These differ depending on the current mode and are described in section 7.1.2.
- At the bottom of the window, log entry data is shown in a table. By default, log entries are sorted by date ascending, but sorting can be customized. To sort by any of the log properties click on the different column headers. To resize the width of a column to the longest value in that column, double-click on a header separator at the right edge of the column.




Filter and query controls

When the application is in local mode these are the controls displayed:

Figure 6-5 Local mode filter controls

Use From/To date and time picker to set the time period of displayed log entries. Time is displayed in UTC. You can change time/date by entering any part directly using the keyboard, or using mouse scrolling to increase/decrease the value of the selected part. The date can also be selected with a date drop down menu, which you can display by clicking the arrow button.

To filter display of log entries by desired text, enter the desired text into Filter field. Filtering is not case-sensitive. This field also accepts grep-like regular expressions. When filtering by specific log entry property, the original names are used. You can see the original name in a tooltip, displayed when you hover the mouse cursor over the column header. You can also see the list of all original property names in the context help for the Filter field. To display the context help, press Shift+F1 with the field focused or use the question-mark button  from window title bar, if available. Note that not all windows managers are able to display this button.

To add a new entry into a local log file in the \$CII_LOGS directory click on the New Entry button and a new dialog window will be opened. For details about it see the section below.

In global mode the controls are a bit different:

Figure 6-6 Remote mode query controls

The Query dropdown list enables to select a predefined Elasticsearch query. When desired query is selected, it is copied to the query text area on the right. You can edit it, before submitting it by clicking the Refresh button.

The From/To date and time fields behave the same way as in the local mode. When the date or time values are modified, the query text area on the right will also reflect these changes by updating the cii.date field range constraint.

The query text area on the right and enables to edit the entire Elasticsearch query as needed.



Adding new local log file entry to logsink

When New Entry button is clicked in local mode, a dialog window opens. The log message you would like to append to a local log file is entered into a text field in the top half of the dialog. The message is mandatory and the Send button will only be enabled if you enter some message text.

File: /home/jlp/INTROOT/logsink/default-eeltdev-output.log

Message

My Message

Additional information

Cref=DefaultCref, HostName=DefaultHostName, LogType=INFO, SourceID=DefaultSourceID, LogID=DefaultLogID

Figure 6-7 Send log message dialog

An additional information text box can be used to set other log entry property fields. These must be set in the form of “property=value”. Note that some property fields (e.g. Cref) are mandatory. These are filled out in advance in this text field with their default values. It is recommended you change them as needed. For information about all available property fields, you can consult the contextual tooltip help for this field and also Appendix B. To access it use the question-mark/button at the window title bar.

When you are finished with the message and additional information, click the Send button to append the log entry to the end of the local log file, or the Cancel button to close the dialog without appending any entry.

9.7.4 Logtrail

Logtrail can be accessed in a browser by navigating to the appropriate Kibana URL (eg: <http://ciikibanahost:5601>). Access to the Logtrail app can be found in the left-side menu of the Kibana GUI. It is also possible to navigate directly to the URL app (e.g.: <http://ciikibanahost:5601/app/logtrail>)

To use Logtrail log browser, two configurations need to be in place:

- Kibana index patterns



- Logtrail configuration

Kibana index patterns

Index patterns [18] are used to define the Elasticsearch log indices that are available for exploring. Index patterns settings can be found in Kibana Management app, available on the left-side menu of the Kibana GUI.

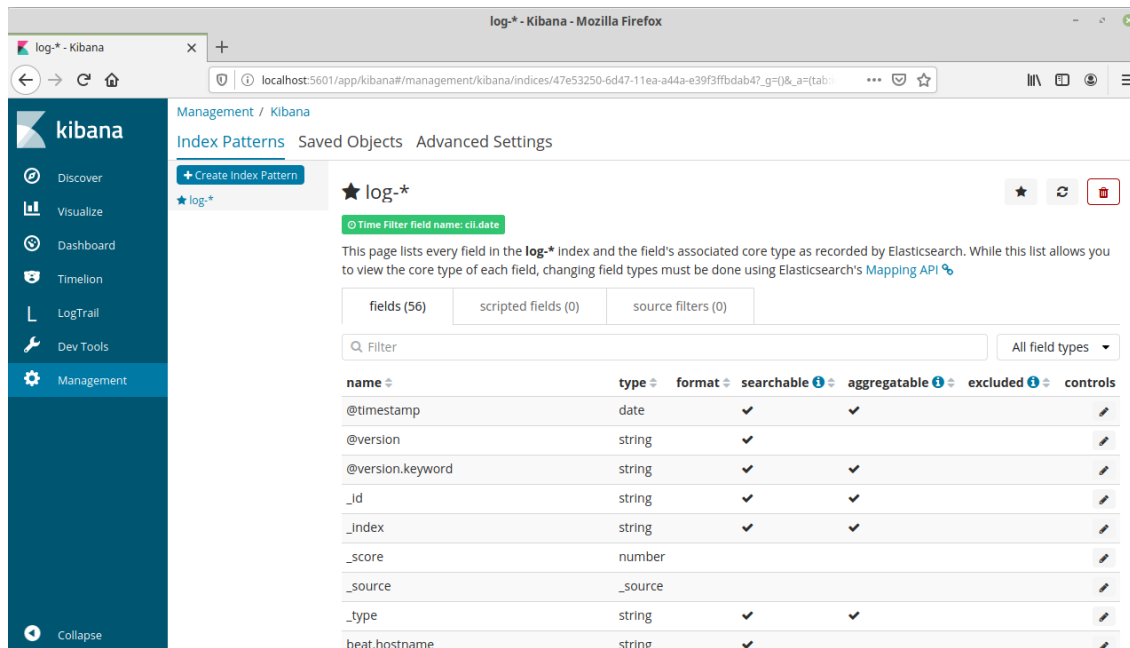


Figure 6-8 Kibana Index Patterns view

When creating an index pattern for CII logs it is important to select `cii.date` as the time filtering field, otherwise, it will use the default `@timestamp` field. This setting is primarily used to order all the log events being queried from the different Elasticsearch indices matched by the pattern.

For example, if log index patterns are defined as shown in section 7.2.3. following the pattern `log-%{AUDIENCE}-%{+YYYY.MM.dd}`, it can be useful to define patterns like:

- `cii_log-*` : Provides access to all log indices in the Elasticsearch instance
- `cii_log-DEVELOPER-*` : Provides access to all log indices for DEVELOPER audience.
- `remoteciilogelastichost:cii_log-` : Provides access to all logs in remoteciilogelastichost remote Elasticsearch cluster.

Additional information can be found in the Kibana documentation [18].



Cross-cluster search

Kibana supports searches across remote Elasticsearch clusters. To enable cross-cluster search, the appropriate remote clusters must be configured in the local Elasticsearch cluster.

Remote clusters can be configured as persistent or transient. If the remote cluster is configured as persistent, the configuration will be persisted between restarts of the Elasticsearch cluster, while a transient remote cluster configuration won't be available after a full cluster restart.

For each remote cluster, a name is defined, and at least a seed must be provided. These seeds are hosts belonging to the remote cluster being configured. For each seed, a host must be provided, as well as the Elasticsearch transport port for that host. The transport port is used for communication between Elasticsearch nodes and is set to port 9300 by default. It must not be confused with the Elasticsearch HTTP API port, which is set by default to port 9200.

In the example below, a transient remote cluster with name remoteciilogelastichost on host ciilogelastichost is added to a local cluster.

Listing 6-1 Configuring a remote cluster for cross-cluster search

```
curl -X PUT "localhost:9200/_cluster/settings?pretty" \
-H 'Content-Type: application/json' -d'
{
  "transient": {
    "search": {
      "remote": {
        "remoteciilogelastichost": {
          "seeds": [
            "ciilogelastichost:9300"
          ]
        }
      }
    }
  }
}
```

To remove a remote cluster configuration, the seeds must be set to null, as shown in the example . If the remote cluster configuration contains any optional settings, these settings must be set to null as well

Listing 6-2 Removing a cross-cluster search remote cluster configuration

```
curl -X PUT "localhost:9200/_cluster/settings?pretty" \
-H 'Content-Type: application/json' -d'
{
  "transient": {
    "search": {
      "remote": {
        "remoteciilogelastichost": {
```

(continues on next page)



(continued from previous page)

```
        "seeds": null
      }
    }
  }
}
```

Additional information on cross-cluster search can be found in the Kibana documentation [21].

Logtrail configuration

Logtrail can be configured through `logtrail.json` file located in Logtrail plugin directory `/usr/share/kibana/plugins/logtrail/`. It also supports configuration through Elasticsearch.

At Kibana start-up, Logtrail will automatically look for a configuration. First, it will look in the `.logtrail` index in the Elasticsearch. If a valid configuration is found, it will be used. Otherwise, the one from the `logtrail.json` file will be used.

Changes in `logtrail.json` file configuration require Kibana service to be restarted. If Elasticsearch configuration is used, it will be reloaded automatically by accessing or refreshing the Logtrail app page.

To write a configuration to the Elasticsearch, the command shown in Listing 7-3 can be used. In that command, `elasticsearch_host` should be replaced with the correct Elasticsearch instance hostname or IP, and `path_to_logtrail.json_file` with the path to the file containing the Logtrail configuration to be saved to Elasticsearch.

Listing 6-3 Saving Logtrail config to ES command

```
curl -XPUT '<elasticsearch_host>:9200/.logtrail/config/1?pretty' \  
-H 'Content-Type: application/json' \  
-d@<path_to_logtrail.json_file>
```

An example Logtrail configuration for CII logs can be found in Listing 7-4. This configuration is defined for index pattern `log-*`.

It is possible to define configuration for multiple index patterns. Each index pattern will have its own separate configuration. Switching between different index-pattern configurations is performed through the Logtrail app. The example in Listing 7-5 shows a configuration for index patterns `log-*` and `log-DEVELOPER-*`.

Mappings for Logtrail fields `timestamp`, `hostname`, `program` and `message` are required. From these fields, the ones mapped as `timestamp`, `hostname`, and `program` will be displayed first for each displayed log event on the Logtrail app. The fields following those are configured through the `message_format` field. This field allows to define a pattern for the fields displayed for the log events.

It is also possible to define colours for the message field based on the values of another field. In the



provided examples, the `cii.logtype` field is used to apply different colours to the message field based on the severity of the log event.

Listing 6-4 Logtrail configuration example

```
{
  "version" : 2,
  "index_patterns" : [
    {
      "es": {
        "default_index": "log-*"
      },
      "tail_interval_in_seconds": 10,
      "es_index_time_offset_in_seconds": 0,
      "display_timezone": "local",
      "display_timestamp_format": "YYYY-MM-DDTHH:mm:ss.SSSZ",
      "max_buckets": 500,
      "default_time_range_in_days" : 0,
      "max_hosts": 100,
      "max_events_to_keep_in_viewer": 5000,
      "default_search": "",
      "fields" : {
        "mapping" : {
          "timestamp" : "cii.date",
          "hostname" : "cii.hostname",
          "program" : "cii.sourceid",
          "message" : "cii.message"
        },
        "message_format": "{{{cii.cref}}} | {{{cii.logid}}} | {{{cii.logtype}}}_
↩| {{{cii.message}}} | {{{cii.audience}}} | {{{cii.userlogid}}} | {{{cii.
↩traceid}}} | {{{cii.spanid}}} | {{{cii.error.typename}}} | {{{cii.error.
↩datetime}}} | {{{cii.error.message}}} | {{{cii.error.trace}}}",
        "keyword_suffix" : "keyword"
      },
      "color_mapping" : {
        "field": "cii.logtype",
        "mapping": {
          "FATAL": "#FF0000",
          "ERROR": "#FF9933",
          "WARN": "#FFFF66",
          "INFO": "#B5E7A0",
          "DEBUG": "#FFFFFF",
          "TRACE": "#CFE0E8"
        }
      }
    }
  ]
}
```

Listing 6-5 Multiple index patterns Logtrail configuration example



```
{
  "version" : 2,
  "index_patterns" : [
    {
      "es": {
        "default_index": "log-*"
      },
      "tail_interval_in_seconds": 10,
      "es_index_time_offset_in_seconds": 0,
      "display_timezone": "local",
      "display_timestamp_format": "YYYY-MM-DDTHH:mm:ss.SSSZ",
      "max_buckets": 500,
      "default_time_range_in_days" : 0,
      "max_hosts": 100,
      "max_events_to_keep_in_viewer": 5000,
      "default_search": "",
      "fields" : {
        "mapping" : {
          "timestamp" : "cii.date",
          "hostname" : "cii.hostname",
          "program": "cii.sourceid",
          "message": "cii.message"
        },
        "message_format": "{{{cii.cref}}} | {{{cii.logid}}} | {{{cii.logtype}}}_
↔| {{{cii.message}}} | {{{cii.audience}}} | {{{cii.userlogid}}} | {{{cii.
↔traceid}}} | {{{cii.spanid}}} | {{{cii.error.typename}}} | {{{cii.error.
↔datetime}}} | {{{cii.error.message}}} | {{{cii.error.trace}}}",
        "keyword_suffix" : "keyword"
      },
      "color_mapping" : {
        "field": "cii.logtype",
        "mapping": {
          "FATAL": "#FF0000",
          "ERROR": "#FF9933",
          "WARN": "#FFFF66",
          "INFO": "#B5E7A0",
          "DEBUG": "#FFFFFF",
          "TRACE": "#CFE0E8"
        }
      }
    },
    {
      "es": {
        "default_index": "log-DEVELOPER-*"
      },
      "tail_interval_in_seconds": 10,
      "es_index_time_offset_in_seconds": 0,
      "display_timezone": "local",
      "display_timestamp_format": "YYYY-MM-DDTHH:mm:ss.SSSZ",
      "max_buckets": 500,

```

(continues on next page)



(continued from previous page)

```
"default_time_range_in_days" : 0,  
"max_hosts": 100,  
"max_events_to_keep_in_viewer": 5000,  
"default_search": "",  
"fields" : {  
  "mapping" : {  
    "timestamp" : "cii.date",  
    "hostname" : "cii.hostname",  
    "program" : "cii.sourceid",  
    "message" : "cii.message"  
  },  
  "message_format": "{{{cii.cref}}} | {{{cii.logid}}} | {{{cii.logtype}}}  
↪ | {{{cii.message}}} | {{{cii.audience}}} | {{{cii.userlogid}}} | {{{cii.  
↪ traceid}}} | {{{cii.spanid}}} | {{{cii.error.typename}}} | {{{cii.error.  
↪ datetime}}} | {{{cii.error.message}}} | {{{cii.error.trace}}}",  
  "keyword_suffix" : "keyword"  
},  
"color_mapping" : {  
"field": "cii.logtype",  
  "mapping": {  
    "FATAL": "#FF0000",  
    "ERROR": "#FF9933",  
    "WARN": "#FFFF66",  
    "INFO": "#B5E7A0",  
    "DEBUG": "#FFFFFF",  
    "TRACE": "#CFE0E8"  
  }  
}  
}  
]  
}
```

Main Window

The main window is presented when the Logtrail app is accessed is composed of three main parts:

- The Kibana Menu is located on the left side of the window and provides access to the different Kibana applications. On the bottom part it has a Collapse button that toggles the menu mode between icon with names and only icons.
- The log filtering and navigation controls are located at the bottom of the screen. They allow a user to apply different types of filters to the logs, navigate to the logs at a specific point in time, and to enable or disable live tailing of events.
- The logs viewer is the main element of the window. It provides a scrollable view of the currently browsed logs according to the format defined in the Logtrail configuration.



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
 Doc. Version: 1
 Released on: -
 Page: 355 of 589




Figure 6-9 Logtrail main window

Filtering logs

The first element from the left on the bottom controls is the filtering field (Figure 7-10). It is possible to filter logs using the Lucene Query Syntax [19].



Figure 6-10 Logtrail filtering field

For example, the following would filter logs of `cii.logtype` ERROR and `cii.hostname` including substring "elt".

```
cii.logtype: "ERROR" AND cii.hostname: "*elt*"
```

To apply the filter, press enter, or click the Search button.

Time-range queries also support relative values. For example, the following example query will display all the log events of the last 24 hours:

```
cii.date: [NOW-1DAY TO NOW]
```

When no field name is provided it will search in the field mapped as message in the Logtrail configuration, `cii.message` in this case. The matched substring will be highlighted on the message field.

Clicking on one of the fields of a log event will automatically filter logs according to the value of the clicked field. For example, clicking on the `cii.logtype` field of a log event with value ERROR will filter the logs and display only the ones with `cii.logtype` value ERROR.



Selecting a hostname

The All Systems button is located to the right of the Search button, and when clicked it will display a popup containing a list of all the hostnames found in the current log index pattern. The number of hosts displayed is limited and can be configured with the max_hosts property in the configuration. When a hostname is selected, only the logs for that hostname will be displayed, and the button will change its label to reflect the selected hostname.

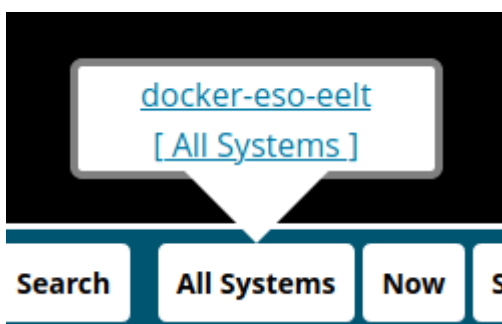


Figure 6-11 Logtrail hostname filter

Navigating to a point in time

The “seek to” button can be found to the right of the hostname filtering button (Figure 7-12). The button displays as its label the current seek to point in time, which by default is Now. When clicked, it opens a popup containing a field to enter a point in time. The user must define a point in time and click the Seek to button. When applied, Logtrail will seek to the logs at the selected point in time.

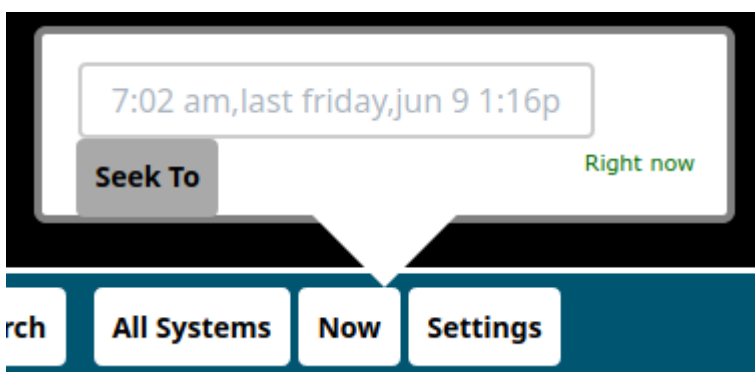


Figure 6-12 Logtrail seek-to control



Selecting an index pattern

The Settings button (Figure 7-13) is located to the right of the “seek to” control. This button allows a user to switch between the different index patterns Logtrail configurations. When clicked, it will display a dropdown menu with the list of available index pattern configurations where the desired pattern can be selected by clicking on it. Once an index pattern is selected, only logs from indices matching this pattern will be displayed.

It must be noted that the appropriate Logtrail configuration must be provided for an index pattern to be browsable with Logtrail. It is not enough to only define that index pattern in Kibana.

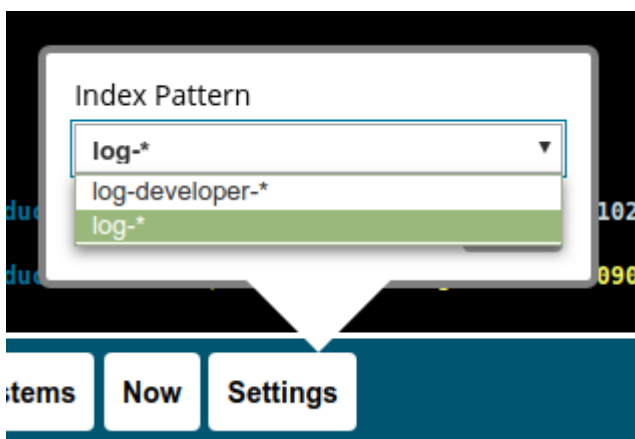


Figure 7-13 Logtrail index pattern selection

Live tailing of log events

Logtrail supports live tailing of events. The refresh rate when on live tail mode is defined on the `tail_interval_in_seconds` property of the Logtrail configuration.

The control for live tail mode is located in the right-bottom corner of the window, and it can be on three different states, as described in Table 6-3.

Table 6-3 Logtrail live tail control states

Icon	State description
	When this icon is displayed, it means that the currently displayed log events are not the latest available from the last performed search or last refresh. When clicked, Logtrail will seek to the latest available log events and start polling periodically the Elasticsearch live tailing the new results.
	This icon is displayed when Logtrail is in live tailing mode. When clicked it will stop the live tailing
	This icon means we are at the latest available log from the last search or refresh and live tailing of log events is paused. Clicking it will start the live tailing again.



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 358 of 589

It must be noted that the Logtrail version used in the current release has a bug, which makes these **live tailing control buttons not render in certain web browsers**, such as Google Chrome.

9.8 Tracing

This section describes the CII Tracing system, including application instrumentation examples as well as the different services used to manage and browse the tracing data produced by the CII distributed systems.

9.8.1 Introduction

Tracing purpose is to aid developers and engineers in the understanding of interactions between components in a distributed system.

It must be noted that logging and tracing are handled separately and there is no automatic connection between logs and traces. This logical connection must be created manually by the user.

It is supported in CII with the OpenTracing API [22], and Jaeger [23] is used as the OpenTracing implementation.

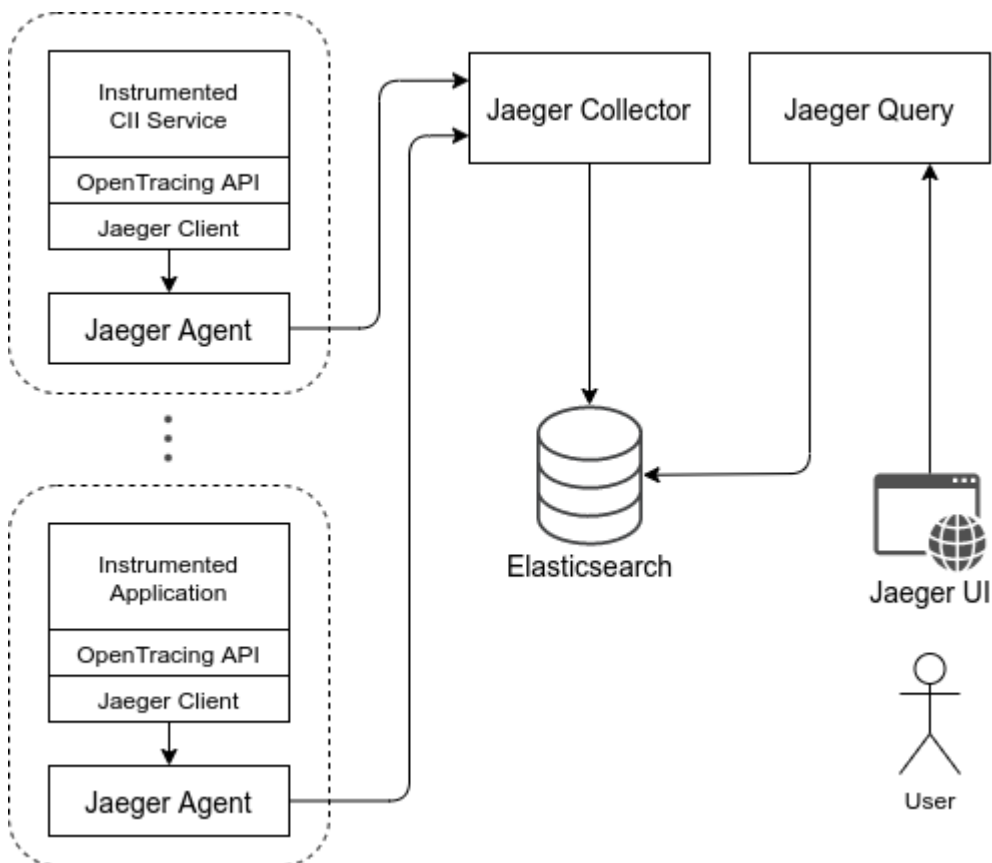


Figure 8-1 CII Tracing architecture

The architecture of the tracing system is shown in the figure above.

It can be divided in:



- Instrumented application using OpenTracing API and Jaeger Client libraries
- Jaeger services
- Elasticsearch persistence backend

The Jaeger services used are:

- Jaeger Agent
- Jaeger Collector
- Jaeger Query

A standard workflow of the tracing system is also shown in Figure 8-1.

Applications and services instrumented with the OpenTracing API produce tracing data, that is sent to a Jaeger Agent by the Jaeger Client library.

The Jaeger Agent collects tracing data from multiple applications and sends batches of data to the Jaeger Collector.

The Jaeger Collector is responsible for persisting the tracing data received from the Jaeger Agents to an Elasticsearch instance.

Finally, the Jaeger Query service provides access to the persisted tracing data, as well as the Jaeger UI, which can be used to browse the tracing data.

Traces and Spans

Distributed tracing relies on traces and spans. A trace is the complete processing of a request. The trace represents the whole execution path of a request as it moves through all of the services or components of a distributed system. All trace events generated by a request share a trace ID that tools use to organize, filter, and search for specific traces.

Each trace is comprised of several spans. A span is an activity or operation that takes place within individual services or components of the distributed system. Each span is another step in the total processing of the overall request. Spans are typically named and timed operations. Spans each carry a unique span ID, and can also carry a reference to their parent span, metadata, or other annotations.

A trace can be seen as a directed acyclic graph, where spans are the nodes and references between spans are the edges.

For example, the following figure illustrates a trace made up of 6 spans:

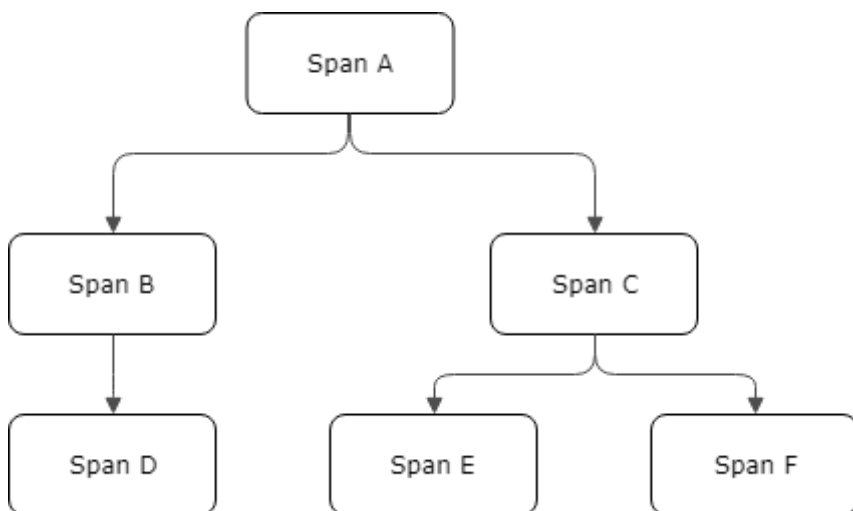


Figure 8-2 Causal relationships between spans in a trace

While this representation allows us to visualize the relationships between the different spans in a trace, it can also be useful to visualize traces in a timeline:




Figure 8-3 Temporal relationships between spans in a trace

Each trace is identified by a unique Trace ID. A traced transaction that spans across multiple hosts, processes, or threads shall have a tracing context propagated between them.

This trace context must contain the following information to be able to trace the call flow of a transaction between services:

- Trace ID: Identifier for a specific operation spanning across one or more distributed systems. When this operation is initiated, a Trace ID is generated and passed to other services that are called as part of the operation.

A trace is composed of one or more spans in one or more distributed hosts. Each span describes a logical unit of work, and contains the following information:

- Operation name: String describing the unit of work represented by the span, e.g. ConfigService::readConfig.
- Span ID: Identifies the span. It must be unique within a trace.



- Parent ID: Identifier that references the ID of the span that triggered the current span.
- Start time: Timestamp indicating when a span was started.
- End time: Timestamp indicating when a span was ended.

The flow of this tracing context information across different services through operation can be seen in Figure 8-4.

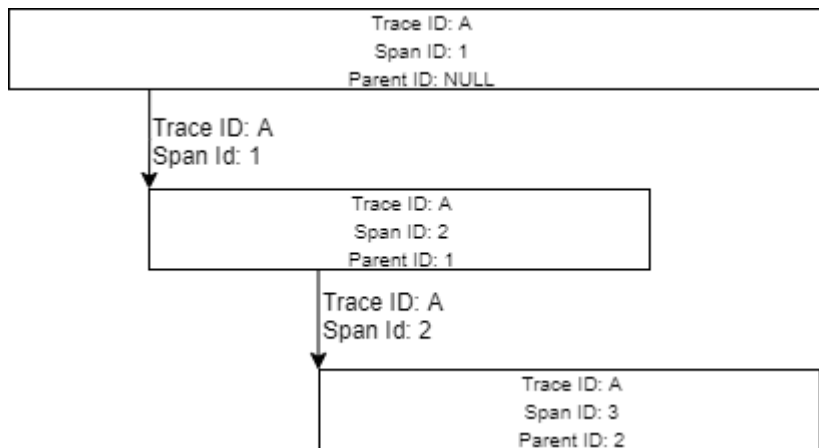


Figure 8-4 Tracing context information flow

In this case, Service 1 would create a tracing context with Trace ID = A, Span ID = 1, and Parent ID = NULL. Together with the call to Service 2, Trace ID and Span ID of the caller are passed. Then Service 2 will create a trace context with Parent ID equal to the Span ID from Service 1 and the Trace ID to the one passed by Service 1. This newly created context will be set as the current context for Service 2. The same procedure applies to the interaction between Service 2 and Service 3 and to subsequent calls to any other service.

9.8.2 Tracing Library Usage

The following sections describe basic tracing instrumentation usage examples.

Advanced usage and configuration of the OpenTracing API and Jaeger are out of the scope of this document. For further information, see [22] and [23].

Tracing request-reply communication

Every request-reply action is instrumented.

When a new request is being issued by the client a new span is created. The MAL middleware implementation must check for existence of an active span using `GlobalTracer.get().scopeManager().activeSpan()` method. If there is a span active a newly created span Parent Id is set to Span Id of the active span, otherwise Parent Id is to 0. Operation name of the span is set to `<instance name>::<method name>`, e.g. `ConfigService::readConfig`. Trace Id and Span Id of



the span are propagated to the server side among other request data. Upon completion or cancelation of the request the `Span.finish()` method is called on the created span instance.

On the server side a new span is created when request is received. The span shares the same operation name and Trace Id. It is created as a child span of the request span, i.e. Span Id of the span is set to Parent Id the of the request span. Upon completion or cancelation of the reply the `Span.finish()` method is called on the created span instance. The server-side implementation of a CII request-reply interface can create a child span of the span that CII created using mal tracing utilities by calling `elt::mal::util::tracing::startSpan`, which will create a child span of an active span if an active span exists, otherwise a childless span will be created.

Tracing publish-subscribe communication

Publish-subscribe communication is not instrumented by default. It needs to be enabled per data entity in ICD, i.e. by setting structure attribute `trace` to "true". This will instruct ICD generator to declare agnostic data entity structure as traceable by extending `Traceable` interface. The interface provides access to Trace Id and Span Id.

```
interface Traceable {  
    int64_getTraceId();  
    uint64 getSpanId();  
};
```

Figure 8-5 Definition of `Traceable` interface in pseudo-IDL.

On publisher side the process is the same as on the client side of request-reply communication pattern with exception that the operation name equals "publish(<data entity name>)" and the span finishes as soon as it is published (enqueued).

On the subscriber side it is programmer's responsibility to create a new span, mal mapping will not create any spans automatically. User needs to extract a tracing context (i.e. `sample->getTrace()`) and create a span using mal tracing utilities. This decision has been made mostly due to the fact that multiple data entities can be read at once. The Logging API will provide helper method to create and activate a new span from received data entity that implements `Traceable`, e.g. `TracingUtils#createSpan(methodName, traceable)`.

Includes/Imports

Java

```
import elt.utils.TracingUtils;  
import io.opentracing.Span;  
import io.opentracing.SpanContext;  
import io.opentracing.Tracer;  
import io.opentracing.util.GlobalTracer;
```



C++

```
#include <ciiTrace.hpp>
```

Python

```
import elt.trace
```

Java class path dependencies

To support the CII Tracing, the Opentracing and Jaeger libraries need to be added to the Java classpath. It is also important to set the Jaeger service name and sampling type and rate.

Listing 8-1 Java tracing classpath libraries

```
LIBDIR="$CIISRV_ROOT/lib"  
SRV_SUPPORT_LIBS="$LIBDIR/srv-support-libs/*"  
JAEGER_LIBS="$LIBDIR/srv-support-libs/jaeger/*"  
OPENTRACING_LIBS="$MAL_ROOT/lib/opentracing/*"  
MALCOMMON_LIB="$MAL_ROOT/lib/mal-common.jar"  
export JAEGER_SERVICE_NAME=tracing-example  
  
java -DJAEGER_SAMPLER_TYPE=const -DJAEGER_SAMPLER_PARAM=1 -DJAEGER_SERVICE_NAME=  
↪$JAEGER_SERVICE_NAME -cp "$MALCOMMON_LIB:$OPENTRACING_LIBS:$JAEGER_LIBS:$SRV_  
↪SUPPORT_LIBS:$LIBDIR/*" elt.trace.examples.TracingBasicExample
```

Basic usage example

In these examples, tracing is initialized using the provided `TracingUtils.initializeTracing()` helper method, which can be used to initialize the tracing framework and register a global tracer for the current application.

Then a root span is created, and a second span with the root span as its parent.

Finally, after all spans are finished, the tracer is closed.



Java

This example can be run from the root directory of the logging-examples module with the run-tracing-example-java provided script. Tracing enabled examples for the request/reply and publish/subscribe are in icd-demo in directory app/java.

Listing 8-2 Tracing basic example (Java)

```
/**
 * CII tracing utility usage
 */
public class TracingBasicExample {

    public static void main(String[] args) {

        /* Initialize tracing framework and register global tracer.
         */
        Tracer tracer = TracingUtils.initializeTracing();

        /* Once the global tracer is registered, it is also
         * possible to obtain an instance from GlobalTracer.
         */
        tracer = GlobalTracer.get();

        /* Start a parent span */
        Span parentSpan = tracer.buildSpan("method1")
            .start();

        /* Start a child span */
        String str;
        Span span = tracer.buildSpan("method2")
            .asChildOf(parentSpan)
            .start();

        try {
            Thread.sleep(2000);
        } catch (InterruptedException e) {
            e.printStackTrace();
        }

        /* close the child span */
        span.finish();

        try {
            Thread.sleep(3000);
        } catch (InterruptedException e) {
            e.printStackTrace();
        }

        /* close the parentSpan */
    }
}
```

(continues on next page)



(continued from previous page)

```
parentSpan.finish();  
  
/* close global tracer */  
tracer.close();  
}  
}
```

C++

This example shown in Listing 8-3 can be run with the `$CIISRV_ROOT/bin/cii-tracing-example-app` executable file installed by the logging-examples module. Tracing enabled examples for a request/reply and a publish/subscribe are in `icd-demo` in directory `app/cpp`.

Listing 8-3 Tracing basic example (C++)

```
int main(int ac, char *av[]) {  
    try {  
        /* Initialize tracing framework and register global tracer.*/  
        ::elt::trace::InitializeTracing("demo-service-name");  
  
        /* Obtain a tracer (registered in previous step) */  
        std::shared_ptr<opentracing::Tracer> tracer = opentracing::Tracer::Global();  
        if (!tracer) {  
            throw ::elt::error::CiiException("No tracer registered");  
        }  
  
        /* Start a parent span */  
        std::unique_ptr<opentracing::Span> parentSpan = tracer->StartSpan("method1");  
  
        /* Start a child span */  
        std::unique_ptr<opentracing::Span> span = opentracing::Tracer::Global()->  
↪StartSpan(  
            "method2", { opentracing::ChildOf(&parentSpan->context()) });  
        std::this_thread::sleep_for(std::chrono::seconds(2));  
  
        /* close the child span */  
        span->Finish();  
        std::this_thread::sleep_for(std::chrono::seconds(3));  
  
        /* close the parentSpan */  
        parentSpan->Finish();  
  
        /* close global tracer */  
        tracer->Close();  
        return 0;  
    } catch (const ::elt::error::CiiException& ex) {  
        std::cerr << "CiiException occured while executing sample code. What: "  

```

(continues on next page)



(continued from previous page)

```
<< ex.what() << '\n';  
}  
return -1;  
}
```

Python

In this example, the use of a span and the tracer as context manager is shown as well.

The example shown in Listing 8-3 can be run with the \$CIISRV_ROOT/bin/cii-tracing-example-app-py executable file installed by the logging-examples module.

Listing 8-4 Tracing basic example (Python)

```
def main():  
    """  
    Main application code  
    @return status code, int  
    """  
    result = 0  
    try:  
        # Initialize tracing framework and register global tracer  
        elt.trace.initialize_tracing('demo-service-name')  
  
        # Obtain a tracer (registered in previous step)  
        tracer = elt.trace.opentracing.Tracer.get_global()  
        if tracer is None:  
            raise RuntimeError('No tracer registered')  
  
        # Start a parent span  
        parent_span = tracer.start_span('method 1')  
  
        # Start a child span  
        span = tracer.start_span("method2",  
                                [elt.trace.opentracing.child_of(parent_span.context())])  
  
        time.sleep(2)  
  
        # Close the child span  
        span.finish()  
  
        # Span as context manager  
        with tracer.start_span("method3",  
                              [elt.trace.opentracing.child_of(parent_span.context())]):  
            time.sleep(1)  
        # span.finish() was called on context exit  
  
        time.sleep(1)
```

(continues on next page)



(continued from previous page)

```
# Close the parent span
parent_span.finish()

# Use tracer as context manager that closes itself automatically
with tracer:
    # make another span and use it as context manager
    with tracer.start_span('another span'):
        time.sleep(1)
# global tracer is closed a this point
# to explicitly close the tracer, use: tracer.close()

except Exception as e:
    print('Exception ocured while executing sample code: %s', e)
    traceback.print_exc()
    result = 5
return result

if __name__ == '__main__':
    sys.exit(main())
```

Mal pub/sub and request/response tracing basic examples

Mal tracing examples are part of icd-demo project and are listed in Table 2.

Table 2: MAL tracing basic examples in icd-demo

Table with 2 columns: Example type, Example location. Rows include Mal pub/sub java tracing example, Mal pub/sub c++ tracing example, Mal pub/sub python tracing example, Mal request/response java tracing example, Mal request/response c++ tracing example, and Mal request/response python tracing example.

Basic examples for mal pub/sub tracing (based on examples in icd-demo):

Listing 8-5 Tracing mal pub/sub example (C++)

```
/**
 * @brief Initialize tracing (implementation specific to the user
```

(continues on next page)



(continued from previous page)

```
* environment, mal is using only opentracing API to obtain the tracer)
*
* In this example we are using jaeger implementation of opentracing
* standard.
*
* @param service_name opentracing service name
*
*/
void InitializeTracing(const std::string& service_name) {
    const bool log_span = true;
    const bool disabled = false;
    const double always_sample = 1.0;

    /* Configuration specific to jaeger (opentracing API implementation) */

    jaegertracing::Config config = jaegertracing::Config(
        disabled,
        jaegertracing::samplers::Config(
            CONST_SAMPLING_STRATEGY,
            always_sample),
        jaegertracing::reporters::Config(
            jaegertracing::reporters::Config::kDefaultQueueSize,
            jaegertracing::reporters::Config::defaultBufferFlushInterval(),
            log_span,
            ::jaegertracing::reporters::Config::kDefaultLocalAgentHostPort));

    std::shared_ptr<opentracing::Tracer> tracer = jaegertracing::Tracer::make(
        service_name, config, jaegertracing::logging::consoleLogger());

    /* Register jaeger tracer as global opentracing tracer,
    MAL will call opentracing API to obtain this tracer*/
    opentracing::Tracer::InitGlobal(
        std::static_pointer_cast<opentracing::Tracer>(tracer));
}

InitializeTracing("Some service");

/* Create a subscriber and a publisher in usual way and start publishing */

::elt::mal::Uri uri("zpb.ps://127.0.0.1:12404/test");
std::unique_ptr<ps::Subscriber<elt::SampleDataEntity>> subscriber =
    mal::CiiFactory::GetInstance().getSubscriber<elt::SampleDataEntity>(
        uri, m_standardQoS, m_malProperties);

std::unique_ptr<ps::Publisher<elt::SampleDataEntity>> publisher =
    mal::CiiFactory::GetInstance().getPublisher<elt::SampleDataEntity>(
        uri, m_standardQoS, m_malProperties);
```

(continues on next page)



(continued from previous page)

```
std::shared_ptr<elt::SampleDataEntity> dataEntity = publisher->
↳createDataEntity();

dataEntity->setText("Initial");
dataEntity->setSampleId(0);

const std::size_t SAMPLE_COUNT = 3;
for (std::size_t i = 0; i < SAMPLE_COUNT; i++) {
    std::stringstream ss;
    ss << "Some text" << i;
    dataEntity->setText(ss.str());
    dataEntity->setSampleId(i);

    /* Since the global tracer is set (InitializeTracing) and
       elt::SampleDataEntity has struct attribute trace="true",
       MAL will start a new span and finish it immediately and propagate the span
       context using sample->setTrace() */

    publisher->publish(*dataEntity, std::chrono::milliseconds(100));
}

std::this_thread::sleep_for(std::chrono::milliseconds(100));

std::vector<std::shared_ptr<elt::SampleDataEntity>> samples = subscriber->
↳read(0);

std::shared_ptr<elt::SampleDataEntity> sample = samples.at(0);

/* Obtain tracer using opentracing API */
std::shared_ptr<opentracing::Tracer> tracer = opentracing::Tracer::Global();

/* Obtain the span context from sample and create a child span */
std::stringstream contextStrm(sample->getTrace());

std::unique_ptr<::opentracing::SpanContext> spanContext =
↳::elt::mal::util::tracing::deserializeSpan(contextStrm);

{
    std::unique_ptr<opentracing::Span> childSpan = opentracing::Tracer::Global()->
↳StartSpan(
        "tracerPublishSubscribeTest",
        {opentracing::ChildOf(spanContext.get())});
}
```

Listing 8-6 Tracing mal pub/sub example (Java)

```
/**
 * @brief Initialize tracing (implementation specific to the user
 * environment, mal is using only opentracing API to obtain the tracer)
```

(continues on next page)



(continued from previous page)

```
*
* In this example we are using jaeger implementation of opentracing
* standard.
*
* @param serviceName service name
*/
public static void initializeTracing(String serviceName) {

    /* Configuration specific to jaeger (opentracing API implementation) */
    Configuration.SamplerConfiguration samplerConfig =
    Configuration.SamplerConfiguration.fromEnv().withType("const").withParam(1);

    Configuration.ReporterConfiguration reporterConfig =
    Configuration.ReporterConfiguration.fromEnv().withLogSpans(true);

    Configuration config = new
    Configuration(serviceName).withSampler(samplerConfig).withReporter(
        reporterConfig);

    /* Register jaeger tracer as global opentracing tracer,
    MAL will call opentracing API to obtain this tracer */
    GlobalTracer.registerIfAbsent(config.getTracer());
}
initializeTracing("some service");

/* Create a subscriber and a publisher in usual way and start publishing */
URI uri = new URI("zpb.ps://127.0.0.1:12517/Sample");

try (Subscriber<SimpleSample> subscriber =
    CiiFactory.getInstance().getSubscriber(
        uri, QoS.DEFAULT, this.malProperties, SimpleSample.class);
    Publisher<SimpleSample> publisher =
    CiiFactory.getInstance().getPublisher(
        uri, QoS.DEFAULT, malProperties, SimpleSample.class);
) {

    SimpleSample dataEntity = publisher.createDataEntity();

    dataEntity.setValue(42.0);
    dataEntity.setDaqId(0);

    final int sampleCount = 3;
    for (int i = 0; i < sampleCount; i++) {
        dataEntity.setValue(i);
        dataEntity.setDaqId(i);
    }

    /* Since the global tracer is set (initializeTracing) and
    SimpleSample has struct attribute trace="true",
    MAL will start a new span and finish it immediately and
```

(continues on next page)



(continued from previous page)

```
propagate the span context using sample.setTrace() */

publisher.publish(dataEntity, 100, TimeUnit.MILLISECONDS);
}
/* Get published samples */
ArrayList<SimpleSample> samples = new ArrayList<>();
LocalDateTime now = LocalDateTime.now();
LocalDateTime timeLimit = now.plusSeconds(10);
while (timeLimit.isAfter(now)) {
    List<SimpleSample> readSamples = subscriber.read(0);
    samples.addAll(readSamples);

    if (samples.size() == sampleCount) {
        break;
    }

    now = LocalDateTime.now();
}

SimpleSample simpleSample = samples.get(0);

/* Obtain the span context from sample and create a child span */

Span span = TracingUtil.createSpan("SubscriberDemo",
    simpleSample.getTrace(), Tags.SPAN_KIND_CONSUMER);
```

Context Propagation

Tracing context is automatically propagated through MAL communication. Every request-reply action is instrumented. Upon completion or cancelation of the request, the `Span.finish()` method is called on the created span.

OLDB messaging automatically propagates the tracing context as well. If there is no active span at the time of call of OLDB API function, then no tracing context is transferred to the OLDB. `Span.finish()` method is called just before OLDB call completes but only if there was an active span. The OLDB client can obtain the span context from a data point value by calling a `getTrace` method which will return the context only if an active span exited at the time when write method was called. Otherwise, the context will be empty. It is then the responsibility of the client to create a child span from the context read from the data point value. For the OLDB subscriptions the user is also responsible for obtaining the span context from data point value received through a subscription listener (`NewValue` call) and then creating new child spans.

In both cases, if there is an active span, the new span is created with the active span as a parent.

Any other type of tracing context propagation is the responsibility of the developer (see [22]).



9.8.3 Tracing Services

This section describes the services used by tracing. The services described below are used for collection, transmission, persistence, and querying of distributed CII traces data. They must be running and properly configured for the tracing system to work.

These services should be managed and maintained by system administrators. See the Log TRD document on how to set up the tracing services.

Documentation on configuration and deployment of these services can be found in [23].

Jaeger Agent

Network daemon that listens for spans sent over UDP, which it batches and sends to the collector.

Jaeger Collector

Receives traces from Jaeger Agents and runs them through a processing pipeline. It is responsible for persisting traces to the storage backend. It supports Elasticsearch as a storage backend.

Jaeger Query

This service retrieves traces from storage and hosts a UI to display them

9.8.4 Tracing GUI (Jaeger UI)

The Jaeger UI is a web application provided by the Jaeger Query service. It allows browsing the tracing data produced by the different CII services and applications.

By default, it's available on port 16686.

To access it, open to the <http://ciijaegerqueryhost:16686> on your browser, where `ciijaegerqueryhost` is the hostname of the Jaeger Query service. The Jaeger UI search page will be presented to the user, as shown in Figure 8-6.




Figure 8-6 Jaeger UI search view

On the left side of this view, the different filtering parameters for the traces search can be defined by the user. To perform the search, click the Find Traces button. When this button is clicked, a list of traces matching the search filtering options will be displayed on the right side of the screen, as shown in Figure 8-7.

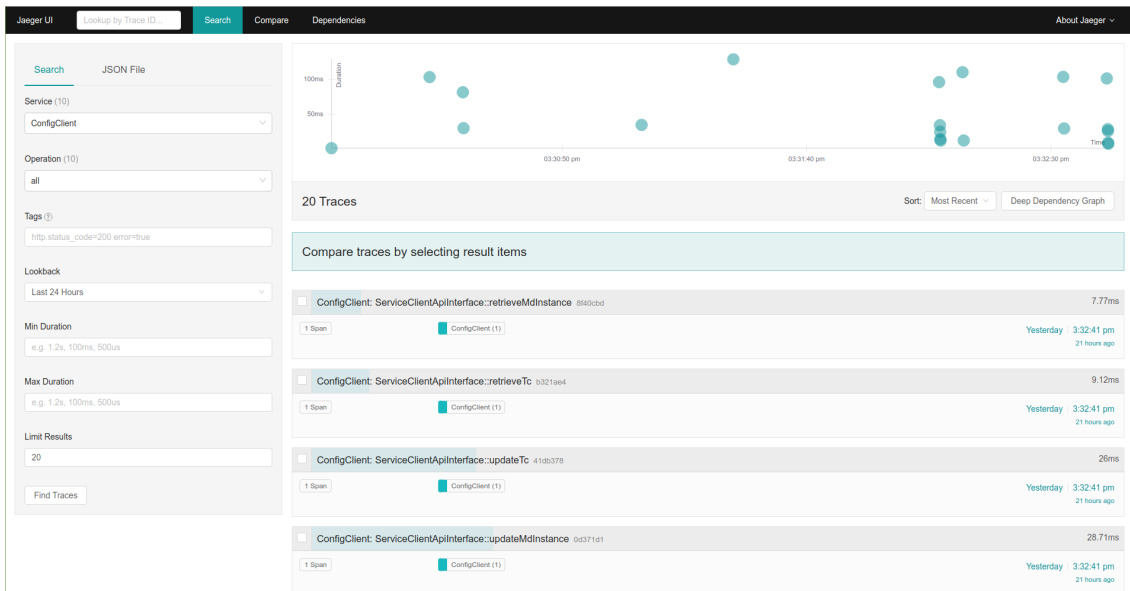


Figure 8-7 Jaeger UI traces search result

It is possible to look for a specific trace using the Lookup by Trace ID input field on the top-left side of the screen. It is also possible to browse a trace directly with the URL



https://ciijaegerqueryhost:16686/trace/<trace_id>¹.

When clicking on a trace on the search results list, or accessing it directly, the trace view will be presented to the user in a trace timeline view, as shown in Figure 8-8. Clicking on each of the spans will display the information about that span.

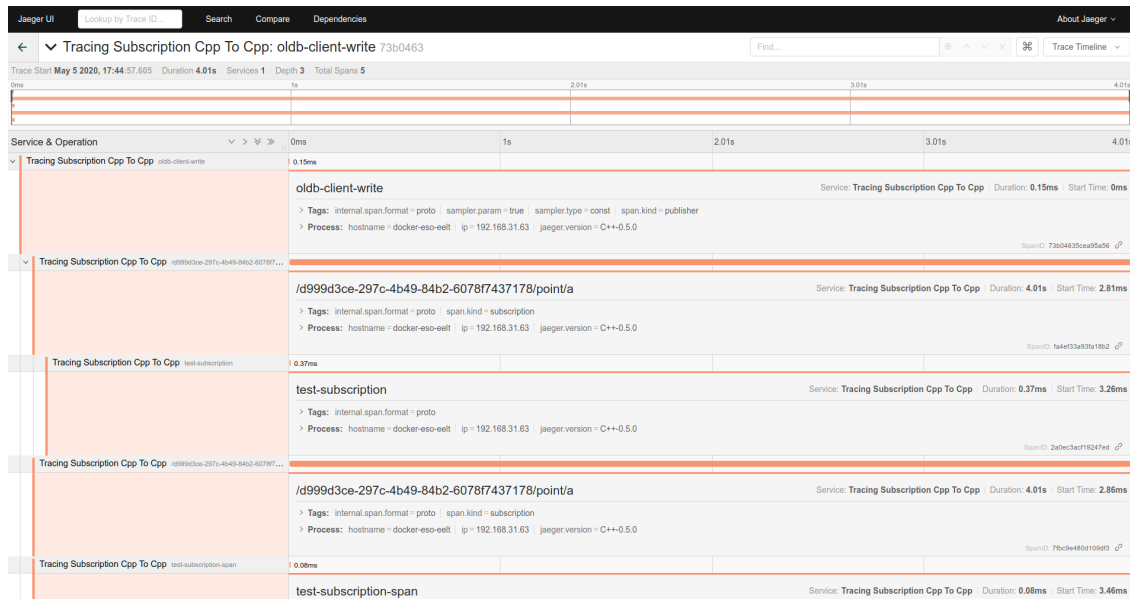


Figure 8-8 Jaeger UI trace view

The trace view offers the following trace visualization modes:

- **Trace Timeline:** Displays the trace and its spans in a timeline visualization (Figure 8-8).
- **Trace Graph:** Displays the trace and its spans as a directed acyclic graph (Figure 8-9).
- **Trace JSON:** Displays the trace and its spans as a JSON object.
- **Trace JSON (Unadjusted):** Displays the trace and its spans as a JSON object with raw data without any adjustments [25].

¹ https://ciijaegerqueryhost:16686/trace/%3ctrace_id%3e



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 376 of 589

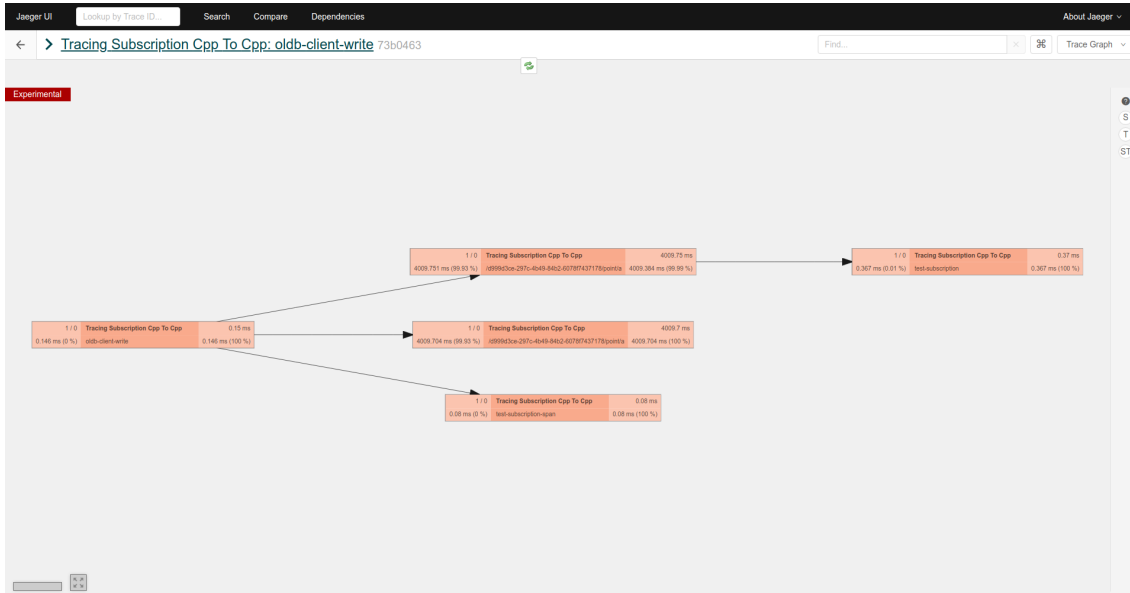


Figure 8-9 Jaeger UI Trace Graph visualization



9.9 CII Log Client API

This are the listings for the CII Log API methods with short descriptions. APIs for Java, C++, and Python are slightly different, therefore a listing for each of the three languages is provided.

9.9.1 Java

Java CII Log API is provided by the CiiLogManager, CiiLogMessageBuilder , CiiLayout, and CiiSimpleLayout.

CiiLogManager

<code>public static Logger getLogger()</code>
CII logger factory method.
Returns:
<ul style="list-style-type: none">• A formatter logger object with the name of the calling class. It automatically initializes and configures the logging context if it has not been initialized previously.

<code>public static Logger getLogger(String loggerName)</code>
CII logger factory method.
Returns:
<ul style="list-style-type: none">• A formatter logger object with the provided name. It automatically initializes and configures the logging context if it has not been initialized previously.

<code>public static void addAppender(Appender appender, Level level, Filter filter)</code>
This method adds an appender with the provided level and filter to the root logger.
Parameters:
<ul style="list-style-type: none">• appender The appender to be added.• level The level threshold for the appender.• filter The filter to be applied to log messages by the appender.

<code>public static void isPostmortemBufferEnabled()</code>
Checks if the postmortem buffer is enabled.
Returns:
<ul style="list-style-type: none">• true if the post moretem buffer is enabled, false otherwise.



`public static boolean flushPostmortemBuffer()`

Flushes buffered logs to appenders that have previously rejected those log messages.

Returns:

- A boolean value indicating if the flushing of the buffer was actually performed. For example, if the method is called when the buffer is disabled, it will return false.

`public static void shutdown()`

Shuts down current logging context.

`public static long elt.log.CiiLogManager.getPostmortemBufferTtlMillis()`

Gets the value of the postmortem Time-To-Live (TTL) for buffered log events.

Returns:

- TTL of the postmortem buffer in ms.

CiiLogMessageBuilder

`public static CiiLogMessageBuilder create(String message)`

Factory method to create a CiiLogMessageBuilder instance with the provided Message field value.

Parameters:

- message The message value.

Returns:

- The message builder.

`public static CiiLogMessageBuilder create()`

Factory method to create a CiiLogMessageBuilder instance.

Returns:

- The message builder.

`public CiiLogMessageBuilder withMessage(String message)`

Sets the Message field value.

Parameters:

- message The message value.

Returns:

- The message builder.



<code>public CiiLogMessageBuilder withUserLogId(String userLogId)</code>
Sets the UserLogID field value.
Parameters:
<ul style="list-style-type: none">• userLogId The user ID value.
Returns:
<ul style="list-style-type: none">• The message builder.

<code>public CiiLogMessageBuilder withAudience(CiiLogAudience audience)</code>
Sets the Audience field value.
Parameters:
<ul style="list-style-type: none">• audience The audience value.
Returns:
<ul style="list-style-type: none">• The message builder.

<code>public CiiLogMessageBuilder withTracingInfo()</code>
Indicates the CiiLogMessageBuilder to include TraceID and SpanID fields from the current tracing context.
Returns:
<ul style="list-style-type: none">• The message builder.

<code>public CiiLogMessageBuilder withTracingInfo(Boolean withTracingInfo)</code>
Indicates the CiiLogMessageBuilder if it should include TraceID and SpanID fields from the current tracing context.
Parameters:
<ul style="list-style-type: none">• withTracingInfo Boolean indicating if tracing info should be added to the built message.
Returns:
<ul style="list-style-type: none">• The message builder.

<code>public CiiLogMessageBuilder withCiiException()</code>
Sets a CiiException to be formatted to the ErrorTypeName, ErrorDateTime, ErrorMessage, and ErrorTrace and added to the log message.
Returns:
The message builder.



public String build()

Constructs the log message string according to the defined parameters.

Returns:

- The formatted string.

public static String createAndBuildWithException(String message, CiiException exception)

Convenience method that creates a formatted log message with the the provided Message field value and the fields containing the CiiException info.

Parameters:

- message The message value.
- Exception The exception to be formatted.

Returns:

- The formatted string.

public static String createAndBuildWithAudience(String message, CiiLogAudience audience)

Convenience method that creates a formatted log message with the the provided Message field value and the provided audience.

Parameters:

- message The message value.
- audience The audience value.

Returns:

- The formatted string.

CiiLayout

public static CiiLayout create()

Factory method to create a CiiLayout instance.

Returns:

- A new CiiLayout instance.



CiiSimpleLayout

<code>public static CiiSimpleLayout create()</code>
Factory method to create a CiiSimpleLayout instance.
Returns: <ul style="list-style-type: none">• A new CiiLayout instance.

9.9.2 C++

C++ CII Log Client API is provided by CiiLogManager, CiiLogMessageBuilder, CiiLogConfigurator, CiiLayout, and CiiSimpleLayout classes.

CiiLogManager

<code>void elt::log::CiiLogManager::configure (const std::map< std::string, std::string > & properties)</code>
Configure logging from properties.
Parameters: <ul style="list-style-type: none">• properties Logging configuration written as properties (name and value pairs)

<code>void elt::log::CiiLogManager::configure (const std::string & logConfigFile)</code>
Configure logging from a file name.
Parameters: <ul style="list-style-type: none">• logConfigFile Logging configuration file

<code>std::string elt::log::CiiLogManager::detailMessage(const std::string & message, const CiiLogAudience audience, const std::string userLogId = "")</code>
Utility method producing a detailed message with all method's parameters included.
Parameters: <ul style="list-style-type: none">• message message part of detailed message• audience audience part of detailed message• userLogId userLogId part of detailed message



<code>log4cplus::Logger elt::log::CiiLogManager::getLogger (const std::string loggerName)</code>
Gets a logger instance.
Parameters:
<ul style="list-style-type: none">• loggerName logger name

CiiLogMessageBuilder

Provides a way to construct a log message according to the CII log format

<code>elt::log::CiiLogMessageBuilder::CiiLogMessageBuilder ()</code>
Default constructor.

<code>elt::log::CiiLogMessageBuilder::CiiLogMessageBuilder (const std::string & message)</code>
Constructor, taking message as argument
Parameters:
<ul style="list-style-type: none">• message Message value

<code>std::string elt::log::CiiLogMessageBuilder::Build () const</code>
Build formatted log message string from information within builder and return it.
Returns:
<ul style="list-style-type: none">• The formatted log message string

<code>static CiiLogMessageBuilder elt::log::CiiLogMessageBuilder::Create ()</code>
Static method to create CiiLogMessageBuilder (Java API compat)
Returns:
<ul style="list-style-type: none">• A new instance of the CiiLogMessageBuilder

<code>static CiiLogMessageBuilder elt::log::CiiLogMessageBuilder::Create (const std::string & message)</code>
Static method to create CiiLogMessageBuilder with message (Java API compat)
Parameters:
<ul style="list-style-type: none">• message Message value
Returns:
<ul style="list-style-type: none">• A new instance of the CiiLogMessageBuilder



static std::string elt::log::CiiLogMessageBuilder::CreateAndBuildWithAudience (const std::string & message, CiiLogAudience audience)

Static method to create a formatted log message string with provided message and audience

Parameters:

- message Message value
- audience Target log audience

Returns:

- The formatted message string

static std::string elt::log::CiiLogMessageBuilder::CreateAndBuildWithException (const std::string & message, const elt::error::CiiException & exception)

Static method to create a formatted log message string with provided message and CiiException

Parameters:

- message Message value
- exception Reference to the instance of the elt::error::CiiException

Returns:

- The formatted message string

static std::string elt::log::CiiLogMessageBuilder::CreateAndBuildWithMessage (const std::string & message)

Static method to create a formatted log message string with provided message

Parameters:

- message Message value

Returns:

- The formatted message string

CiiLogMessageBuilder& elt::log::CiiLogMessageBuilder::WithAudience (CiiLogAudience audience)

Add audience to the builder and return reference to self (chaining support)

Parameters

- message audience CiiLogAudience value

Returns

- A reference to self



<code>CiiLogMessageBuilder& elt::log::CiiLogMessageBuilder::WithException (const elt::error::CiiException & exception)</code>
Add information about an <code>elt::error::CiiException</code> to the builder and return reference to self (chaining support).
Parameters: <ul style="list-style-type: none">• exception Reference to the exception
Returns: <ul style="list-style-type: none">• A reference to self

<code>CiiLogMessageBuilder& elt::log::CiiLogMessageBuilder::WithException (const std::exception & exception)</code>
Add information about <code>std::exception</code> to the builder and return reference to self (chaining support)
Parameters: <ul style="list-style-type: none">• exception Reference to the exception
Returns: <ul style="list-style-type: none">• A reference to self

<code>CiiLogMessageBuilder& elt::log::CiiLogMessageBuilder::WithMessage (const std::string & message)</code>
Add message to the builder and return reference to self (chaining support)
Parameters: <ul style="list-style-type: none">• message message value
Returns: <ul style="list-style-type: none">• A reference to self

<code>CiiLogMessageBuilder& elt::log::CiiLogMessageBuilder::WithTracingInfo (bool tracing_info)</code>
Enable/disable tracing info flag and return reference to self (chaining support)
Parameters: <ul style="list-style-type: none">• tracing_info boolean flag to enable/disable addition of tracing information into the message
Returns: <ul style="list-style-type: none">• A reference to self



```
CiiLogMessageBuilder& elt::log::CiiLogMessageBuilder::WithUserLogId ( const std::string & user_log_id )
```

Add user log ID string to the builder and return reference to self (chaining support)

Parameters:

- user_log_id user defined log id string

Returns:

- A reference to self

CiiLogConfigurator

Configures Log4cplus loggers and appenders from configuration file and injects each configured logger with one PostmortemAppender instance to keep all filtered events for some specified time for later usage.

Configure loggers with configuration from log4cplus.properties file and life time of 1 second with example:

```
//configure log4cplus from log4cplus.properties file and then set the  
//life time of the stored events for 1 second  
CiiLogConfigurator::doConfigure ( 1 , "log4cplus.properties" );
```

If configuration of Log4cplus is initialized without the postmortem buffer enabled, it is possible to enable it once the Log4cplus configuration has been initialized with the following method:

```
//set the life time of the stored events for 1 second  
CiiLogConfigurator::inject( 1 );
```

To dump stored events which still have valid lifetimes call this method:

```
//dump all valid events to parent logger  
CiiLogConfigurator::postmortemDump();
```

CiiLogConfigurator has an option to reset log4cplus configuration be fully reset when calling doConfigure:

```
CiiLogConfigurator::setResetOnConfigure( shouldReset );
```



```
static void elt::log::CiiLogConfigurator::doConfigure( uint32_t life_time, const log4cplus::tstring & config_filename, log4cplus::Hierarchy & h = log4cplus::Logger::getDefaultHierarchy(), unsigned flags = 0 )
```

This method eliminates the need to create a temporary PropertyConfigurator to configure log4cplus. It is equivalent to the following:

```
PropertyConfigurator config("filename");  
config.configure();
```

Parameters:

- life_time The life time in seconds for an event to be stored internally in the Postmortem appender
- config_filename The configuration file name
- h root hierarchy
- flags flags

```
static void elt::log::CiiLogConfigurator::inject ( uint32_t life_time )
```

Inject PostMortemAppender into each configured logger instance without loading property file.

Parameters:

- life_time The life time in seconds for an event to be stored internally in the Postmortem appender

```
static void elt::log::CiiLogConfigurator::postmortemDump ( )
```

Dump all valid events (events which have life time less than globally specified life time) to all the loggers that previously rejected them.

```
static void elt::log::CiiLogConfigurator::setResetOnConfigure ( bool reset )
```

Decides whether the configuration should reset all already configured loggers when a new configuration is applied.

Parameters:

- reset true = enable, false = disable

```
static bool elt::log::CiiLogConfigurator::getResetOnConfigure ( )
```

Returns:

- Boolean indicating if the existing loggers will be reset when a configuration is applied.



CiiLayout

Log layout for CII Log format.

<code>elt::log::layout::CiiLayout::CiiLayout(const log4cplus::helpers::Properties & properties)</code>
Construct CII layout with properties. Parameters: <ul style="list-style-type: none">• properties Layout properties.

CiiSimpleLayout

Simple log layout for CII Log format.

<code>elt::log::layout::CiiSimpleLayout::CiiSimpleLayout(const log4cplus::helpers::Properties & properties)</code>
Construct CII simple layout with properties. Parameters: <ul style="list-style-type: none">• properties Layout properties.

9.9.3 Python

Python CII Log Client API is provided by CiiLogManager, CiiLogConfigurator, and CiiLogMessageBuilder classes.

CiiLogManager

Provides methods to manage and configure logging, as well as methods to obtain logger objects.



log.CiiLogManager.configure(config = None)

Perform logging configuration from provided dictionary/file just like logging.config.dictConfig() function with additional functionality:

- Initialize missing keys in the provided dictionary from CII_DEFAULT_CONFIG, so that the root logger gets initialized with logsink and console handlers.
- Configuration dictionary supports additional top level key `cii_log_buffer`, which must contain a dictionary with keys `enabled` and `ttl_ms`. The value of `enabled` key must be boolean. When `True`, a log buffer will be enabled. The value of `ttl_ms` key must be integer defining captured record Time-To-Live in ms (within the buffer). Note that these settings will be overridden with the values of the environment variables below.
- Environment variables `CII_LOG_BUFFER_ON` and `CII_LOG_BUFFER_TTL_MS` are checked for existence. These environment variables override settings for `cii_log_buffer` provided in the configuration dictionary.
 - If `CII_LOG_BUFFER_ON` exists AND has value `1` or `TRUE` (case insensitive), the log buffer will be enabled. The value of `0` or `FALSE` will disable the log buffer.
 - Environment variable `CII_LOG_BUFFER_TTL_MS` if exists, must be convertible to int and represents the number of milliseconds that the log buffer retains the messages. When this variable does not exist, the default of `60000ms` is assumed.

Parameters:

- `config` Configuration dictionary or filename or `None`. When a filename is provided, it must be a valid JSON file that can be parsed into the config dictionary. When `None` is provided, the configuration dictionary is used from the default `CII_DEFAULT_CONFIG` dictionary.

Returns:

- Dictionary with the actual configuration used.

log.CiiLogManager.get_logger(name = None)

CII logger factory method.

Parameters:

- `name` Name of the logger or `None` to obtain root logger.

Returns:

- Instance of logger object.

log.CiiLogManager.flush_postmortem_buffer()

Flush buffered logs to the handlers that have previously rejected those log messages. Cyclic log buffer is reset.

Returns:

- `True` if flushing was actually performed, `False` otherwise.



`log.CiiLogManager.is_postmortem_buffer_enabled()`

Query whether buffered logger is installed.

Returns:

- True when that is the case, False otherwise.

CiiLogConfigurator

`log.CiiLogConfigurator.load_configuration(filename)`

Load log configuration from provided file.

Log configuration must be saved in JSON format. Note that this is not supported with Python logging subsystem. Parsing of log configuration must produce a dictionary that is compatible with `CiiLogManager.configure()` method

Parameters:

- filename Name of JSON file containing configuration.

Returns:

- Dictionary with log configuration.



log.CiiLogConfigurator.configure(config = None)

Perform logging configuration from provided dictionary/file just like logging.config.dictConfig() function with additional functionality:

- Initialize missing keys in the provided dictionary from CII_DEFAULT_CONFIG, so that the root logger gets initialized with logsink and console handlers.
- Configuration dictionary supports additional top level key `cii_log_buffer`, which must contain a dictionary with keys `enabled` and `ttl_ms`. The value of `enabled` key must be boolean. When `True`, a log buffer will be enabled. The value of `ttl_ms` key must be integer defining captured record Time-To-Live in ms (within the buffer). Note that these settings will be overridden with the values of the environment variables below.
- Environment variables `CII_LOG_BUFFER_ON` and `CII_LOG_BUFFER_TTL_MS` are checked for existence. These environment variables override settings for `cii_log_buffer` provided in the configuration dictionary.
 - If `CII_LOG_BUFFER_ON` exists AND has value `1` or `TRUE` (case insensitive), the log buffer will be enabled. The value of `0` or `FALSE` will disable the log buffer.
 - Environment variable `CII_LOG_BUFFER_TTL_MS` if exists, must be convertible to int and represents the number of milliseconds that the log buffer retains the messages. When this variable does not exist, the default of `60000ms` is assumed.

Parameters:

- `config` Configuration dictionary or filename or `None`. When a filename is provided, it must be a valid JSON file that can be parsed into the config dictionary. When `None` is provided, the configuration dictionary is used from the default `CII_DEFAULT_CONFIG` dictionary.

Returns:

- Dictionary with the actual configuration used.

log.CiiLogConfigurator.check_config()

Check whether `configure()` method was already called. If that was not the case, it calls it to configure defaults.

log.CiiLogConfigurator.save_configuration(config, filename)

Save log configuration dictionary as json into provided filename.

Parameters:

- `config` Dictionary with log configuration.
- `filename` Name of the output file.



`log.CiiLogConfigurator.get_formatter(simple = False)`

Create and return CII specific formatter

Parameters:

- `simple` When set to True, return formatter with simple format, otherwise return formatter with CII full pattern.

Returns:

- Instance of formatter object

`log.CiiLogConfigurator.flush_postmortem_buffer()`

Flush buffered logs to the handlers that have previously rejected those log messages. Cyclic log buffer is reset.

Returns:

- True if flushing was actually performed, False otherwise.

`log.CiiLogConfigurator.flush_postmortem_buffer_to_string()`

Flush content of the cyclic log buffer to string. Cyclic log buffer is reset.

Returns:

- String containing log messages captured by CiiBufferedLogger, or a RuntimeError when CiiBufferedLogger is not installed.

`log.CiiLogConfigurator.reset_postmortem_buffer()`

Empty the log buffer if buffered logger is installed.

Returns:

- True when log buffer was reset, False when not installed.

`log.CiiLogConfigurator.get_effective_logsink_filename()`

Return full name of the filename for log sink (if one was created by calling `get_logsink_filename`).

Returns:

- The filename of the logsink log file or None when `get_logsink_filename` was not called.

`log.CiiLogConfigurator.get_logsink_filename()`

Return full name of the filename for log sink.

Returns:

- The filename of the logsink log file.



<code>log.CiiLogConfigurator.enable_postmortem_buffer(ttl_seconds = 60)</code>
Enable postmortem buffer by installing a buffered logger as a root logger. Parameters: <ul style="list-style-type: none">• <code>ttl_seconds</code> Message Time-To-Live in seconds. Returns: <ul style="list-style-type: none">• True when logger installed with this invocation, False when already installed. Note <code>ttl_seconds</code> parameter has no meaning when BufferedLogger is already installed.

<code>log.CiiLogConfigurator.is_postmortem_buffer_enabled()</code>
Query whether buffered logger is installed. Returns: <ul style="list-style-type: none">• True when that is the case, False otherwise.

CiiLogMessageBuilder

Provides a way to construct a log message according to the CII log format.

<code>__init__(self, message=None)</code>
Constructor Parameters: <ul style="list-style-type: none">• <code>message</code> The message field value.

<code>create(cls, message=None)</code>
Creates a formatted log message string with exception info. Parameters: <ul style="list-style-type: none">• <code>message</code> The message field value. Returns: <ul style="list-style-type: none">• A new instance of the CiiLogMessageBuilder.



`create_and_build_with_exception(cls, message, exception)`

Creates a formatted log message string with exception info.

Parameters:

- message The message field value.
- exception The exception instance to be formatted to the appropriate fields the log message. Exception instance should be derived from the `elt.config.CiiException`.

Returns:

- The formatted string.

`create_and_build_with_audience(cls, message, audience)`

Creates a formatted log message string with audience field.

Parameters:

- message The message field value.
- audience The `CiiLogAudience` value to be formatted to the Audience field.

Returns:

- The formatted string.

`create_and_build_with_args(cls, **kwargs)`

Create a formatted log message string from provided arguments. See `with_args`. The following arguments are supported:

- message: value of the message field (string).
- audience: value of the audience field (`CiiLogAudience`).
- user_log_id: value of the user log ID field.
- exception: exception to be formatted.
- with_tracing_info: boolean value for enabling/disabling of tracing info.

Parameters:

- kwargs keyword arguments.

Returns:

- The formatted string.

`with_message(self, message)`

Sets the message field value.

Parameters:

- message The message field value.

Returns:

- The builder instance itself.



`with_user_log_id(self, user_log_id)`

Sets the user ID field value.

Parameters:

- `user_log_id` Value for the user log ID field.

Returns:

- The builder instance itself.

`with_audience(self, audience)`

Sets the audience for the log message.

Parameters:

- `audience` The intended `CiiLogAudience` of the log message.

Returns:

- The builder instance itself.

`with_tracing_info(self, tracing_info=True)`

Sets the log message to include/exclude tracing info fields.

Returns:

- The builder instance itself.

`with_cii_exception(self, exception)`

Sets the exception to be formatted and included in the log message.

Parameters:

- `exception` Exception to be formatted and added to the log message.

Returns:

- The builder instance itself.

`with_exception(self, exception)`

Sets the exception to be formatted and included in the log message. Alias for `with_cii_exception`.

Parameters:

- `exception` Exception to be formatted and added to the log message.

Returns:

- The builder instance itself.



`with_args(self, message=None, audience=None, user_log_id=None, exception=None, with_tracing_info=None)`

Sets the fields from provided keyword arguments. The following arguments are supported:

- `message`: value of the message field (string).
- `audience`: value of the audience field (`CiiLogAudience`).
- `user_log_id`: value of the user log ID field.
- `exception`: exception to be formatted.
- `with_tracing_info`: boolean value for enabling/disabling of tracing info.

Parameters:

- `kwargs` keyword arguments.

Returns:

- The builder instance itself

`build(self)`

Construct the log message string according to the defined parameters.

Returns:

- The constructed log message following the Cii log format.



9.10 CII Log Fields

All fields in a log message, aside from the date/time fields are treated as strings.

Table 8-3 Log fields

Field	ES field name	Description	Optional
Cref	cii.cref	Source filename/line number.	NO
Date	cii.date	Date and time of log creation down to milliseconds in UTC timescale with the 'T' designator between date and time, eg: 2020-03-16T09:32:09.681+0100. It follows ISO 8601 standard.	NO
Host-Name	cii.hostname	Hostname of the computing node where the log was produced.	NO
Log-Type	cii.logtype	The type of log message (level).	NO
SourceID	cii.sourceid	Class name, thread name.	NO
LogID	ci.logid	ID of the log message source location.	NO
Message	cii.message	Text message of the log event.	NO
Audience	cii.audience	The intended recipient of the log.	YES
User-LogID	cii.userlogid	User provided ID.	YES
TraceID	cii.traceid	ID of the trace for the current tracing context.	YES
SpanID	cii.spanid	ID of the span for the current tracing context.	YES
ErrorType-Name	cii.error.type-name	Name of the CII Error	YES
Error-Date-Time	cii.error.datetime	Unix timestamp of the error	YES
Error-Message	cii.error.message	Message of the error	YES
Error-Trace	cii.error.trace	Stack trace of the error	YES

ALARM

Revision:	
Status:	Released
Repository:	https://gitlab.eso.org/cii/info/cii-docs
Project:	ELT CII
Folder:	userManual/ciiman/src/docs
Document ID:	
File:	alarm.rst
Owner:	Matej Šekoranja
Last modification:	September 14, 2022
Created:	July 5, 2019

Prepared by	Reviewed by	Approved by
Jernej Strniša (CSL) Matej Šekoranja (CSL SWE)	Borut Terpinc (CSL SWE) Jan Pribošek (CSL SWE)	Gregor Čuk (CSL)

Document History



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 398 of 589

Re-vi-sion	Date	Changed/ reviewed	Sec-tion(s)	Modification
1.0	6.5.2020	Ojstrnisa/ bterpinc	All	Document creation.
1.1	5.7.2020	Ojstrnisa/ bterpinc	4	Added plugin tab example
1.2	15.10.2020	Ojstrnisa/ msekora nja	All	RIX updates
1.3	17.11.2020	Ojstrnisa/ msekora nja	4.3.2.	Replaced ias-webserver sender with iasWebServer Sender
1.4	13.1.2021	Ojstrnisa/ bterpinc	4.2	Alarm Configuration GUI updates
1.5	26.1.2021	Ojstrnisa/ bterpinc, msekora nja	4.2	Alarm Configuration GUI updates (RIXes PSI 19, PSI 21). Transfer function alarm priority paragraph added (PSI 39).
2.0	14.9.2022	Ojstrnisa/ msekora nja	all	CII v3: new tooling and workflow for defining and monitoring alarms

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is manual for the Alarm system used by ESO employees involved with the ELT Core Integration Infrastructure Software project.

Audience

This document is aimed at those Cosylab and ESO employees involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms



API	Application Programming Interface
CII	Core Integration Infrastructure
CLI	Command Line Interface
CDB	IAS configuration database
DP	Data Point
GUI	Graphical User Interface
ES	ElasticSearch
JSON	Javascript object notation
OLDB	Online Database
URI	Uniform Resource Identifier
SVN	Subversion
IAS	Integrated Alarm System
DASU	Distributed Alarm System Unit
ASCE	Alarm System Computing Element
IASIO	Integrated Alarm System Input/Output

References

1. ESO, Core Integration Infrastructure Requirements Specification, ESO-192922 Version 7
2. ESO, Integrated Alarm System Architecture, ESO-293482 Version 2
3. ESO, Integrated Alarm System Design, ESO-299387 Version 1
4. ESO, Integrated Alarm System General information, <https://github.com/IntegratedAlarmSystem-Group/ias/wiki>
5. <https://www.eso.org/~eltmgr/CII/latest/manuals/html/docs/services.html>
6. <https://www.eso.org/~eltmgr/CII/latest/manuals/html/docs/oldb.html>
8. <https://github.com/IntegratedAlarmSystem-Group/ias/wiki/Transfer-functions-how-to>
9. <https://github.com/IntegratedAlarmSystem-Group/ias/wiki/Transfer-function-in-python>
10. https://github.com/IntegratedAlarmSystem-Group/integration-tools/blob/develop/docs/DEPLOYMENT_GUIDE.md
11. https://github.com/IntegratedAlarmSystem-Group/integration-tools/blob/develop/docs/WEB_APPS_DOCUMENTATION.md
12. <https://github.com/IntegratedAlarmSystem-Group/ias/wiki/ConfigurationDatabase>



10.1 Overview

This document is a user manual for CII Alarm service. It explains how to set-up, start and stop the alarm system (chapter 3), how to configure alarms (chapter 4), and how to inspect and address alarms (chapter 4), including concrete working examples.

The advanced topics (chapter 5) are mostly only relevant for administrators.



10.2 Introduction

The CII Alarm System is based on the Integrated Alarm System IAS.

It operates on rules that observe the OLDB and decide whether an alarm should be raised. The alarm system is therefore fully configuration-based. An application never directly triggers an alarm. An alarm is represented by an OLDB datapoint as well. These datapoints are owned by the Alarm system, and you cannot write to them. You read these datapoints to find out which alarms are active or not.

IAS is a software whose main purpose is to get alarms and monitor points from different sources and generate alarms to present to the users who can range from operators in the control room up to engineers at their desks. The alarm system is a message-passing facility that routes information about abnormal situations detected from hardware or software to the user. For IAS architecture and design, please refer to the IAS architecture [2] (necessary to read to get to know IAS concepts) and design [3] documents.

IAS is based on pluggable architecture to integrate with external systems. To connect the IAS to the CII, the CII-IAS bridge was developed. The bridge consists of the following components:

- OLDB Plugin - The long-running process that reads from OLDB and writes (via Kafka) to IAS.
- OLDB Sink - The long-running process that reads (via Kafka) from IAS and writes to OLDB.

Changes in CII v3:

The following components were retired

- CII-OLDB Adaptor
- CII-IAS Config Exporter
- Alarm GUI webserver
- WebServer Sender
- Alarm Config GUI

The following features are *not available* in this version of the CII Alarm System

- Acknowledgement
- Shelving



10.3 Prerequisites

This section describes the prerequisites for using the CII Alarm service.

10.3.1 Set-up

An administrator (with root privileges) has to post-install the Alarm system on the host. The procedure is described in #Advanced-Topics/Installation below.

10.3.2 Start-Stop

Start the Alarm system service(s)

```
sudo cii-services start alarm
```

Stop the Alarm system service(s)

```
sudo cii-services stop alarm
```



10.4 Usage

The CII Alarm System operates on rules that observe the OLDB and decide whether an alarm should be raised. An application never directly triggers an alarm. An alarm is represented by an OLDB datapoint as well. These datapoints are owned by the Alarm system, and you cannot write to them. You read these datapoints to find out which alarms are active or not.

The general workflow is this:

- I. You define alarm rules
 - Step 1: You write an alarm definition file
 - Step 2: You pass the definitions to the alarm system
 - Step 3: You request the alarm system to read the new definitions
- II. You monitor and trigger alarms
 - Step 4: You monitor the alarm-datapoints
 - Step 5: You write to the input-datapoints

10.4.1 Write Alarm Definitions

Location Alarm rules and CDB are stored in `/etc/cii_ias/`

Create/Edit this file:

```
/etc/cii_ias/alarm-rules.yaml
```

Example

```
---
alarmId: /alarm/harmoni/sensor1_alarm
shortDesc: Demo alarm for numerical value (lowOn, highOn)
rule: [ minmax, /harmoni/fcs1/sensor1/value, 30, 50 ]
---
alarmId: /alarm/harmoni/errorstates/tracking_alarm
shortDesc: Demo alarm with a boolean input value
rule: [ bool, /harmoni/fcs1/adcl/trackingfailure ]
---
alarmId: /alarm/harmoni/hitemp/dcs_alarm
shortDesc: Detector High Temperature (highOff, highOn)
rule: [ max, /harmoni/dcs1/detector/temperature1, 70.0, 100.0 ]
---
alarmId: /alarm/harmoni/device/error_alarm
shortDesc: Demo alarm with string input
rule: [ regexp, /harmoni/dcs1/device/error_string, Error.* ]
```



Format

In the alarm definitions file, each rule consists of 4 lines:

```
---          #1 three minus signs (required)
alarmId:     #2 rule name. must begin with "/alarm/"
shortDesc:   #3 free text describing the alarm
rule:        #4 the alarm function name and its arguments, in list format
```

Alarm functions

These are the available functions for use in the alarm definitions.

Hints

- Arguments called “dp” are addresses of OLDB datapoints. Give the absolute path inside the oldb, but without the “cii.oldb://” prefix, i.e. the address shall start with a single slash.
- The DP’s oldb-type must be compatible with the function-arg type. E.g., where an alarm function expects a double, avoid passing a String-DP (OldbString), but you can pass an Int-DP (e.g. OldbInt64Std).

```
* min (double dp, double lowOn)
  Alarm activates while dp value is below lowOn

* min (double dp, double lowOn, double lowOff)
  Alarm activates when dp value goes below lowOn, deactivates above lowOff.

* max (double dp, double highOn)
  Alarm activates while dp value is above highOn

* max (double dp, double highOff, double highOn)
  Alarm activates when dp value goes above highOn, deactivates below highOff.

* minmax (double dp, double lowOn, double highOn)
  Alarm activates while dp value is below lowOn or above highOn
  In other words, this is minmax without hysteresis.

* minmax (double dp, double lowOn, double lowOff, double highOff, double highOn)
  Alarm activates when dp value goes below lowOn, deactivates above lowOff.
  Analogously for highOn/Off.
  In other words, this is minmax with hysteresis.

* bool (bool dp)
  Alarm activates while dp value is "true"

* regexp (string dp, string regexp)
  Alarm activates while dp value matches the regexp
```



10.4.2 Inject Definitions

The `ciiAlarmLoader` tool reads the alarm definition file, and creates alarm config as well as other config in the format understood by IAS: supervisor config, converter config, etc. The IAS keeps this information in the IAS configuration database (CDB).

```
ciiAlarmLoader alarm-rules.yaml . --check-inputs
```

The alarm loader will output a message if there are alarm rules for which the input-DPs (datapoint-URLs given as arguments to the alarm functions) are currently not existing in the OLDB. These alarm rules will not be able to work, and the affected alarms will never trigger, until the input-datapoints get created by someone. In real scenarios, the input-datapoints get created by the ECS oldb-loader tool, or by control applications.

Alarm rules for which the input-DPs are existing already at the time of reloading the CDB (see next step) will function instantly. For those with missing inputs, there will be a certain delay after the missing input-DPs were created. The alarm system will discover the new inputs within 60 seconds.

Example

To create the input-DPs for the example rules file above, you can use this shell command:

```
/usr/bin/env python - <<EOF
import elt.oldb, elt.config
elt.oldb.CiiOldbGlobal.set_write_enabled(True)
oldb = elt.oldb.CiiOldbFactory.get_instance()

for k,v in {
    "/harmoni/fcs1/sensor1/value" : 34,
    "/harmoni/fcs1/adcl/trackingfailure": False,
    "/harmoni/dcs1/detector/temperature1": 56.0,
    "/harmoni/dcs1/device/error_string": "None"
}.items():
    try: oldb.create_data_point_by_value (elt.config.Uri("cii.oldb://" + k) , v)
    except: pass
EOF
```

10.4.3 Request a Reload

Assuming the alarm system is already running, request it to re-read its config.

```
ciiAlarmCtl daemon reload
```

The command will write a warning if the alarm system is not running, to inform you that no-one will handle your request. If so, see the #Prerequisites section on how to start it.



With this, the definition of alarms is completed, and you can start using them.

10.4.4 Monitor Alarms

Since the alarms are OLDB datapoints, you can monitor them with the OLDB API, OLDB GUI, or OLDB tools (all described in [6]), for example:

```
oldb-cli read </absolute/path/of/alarm-dp> 2>/dev/null | tail -n1  
oldb-cli subscribe </absolute/path/of/alarm-dp>
```

Moreover, the Alarm System comes with 2 dedicated tools:

1. `ciiAlarmMon` (command line executable)

```
ciiAlarmMon /etc/cii_ias
```

2. Alarm Monitor service (system service)

```
http://<host_where_ias_runs>:5602/
```

The Alarm Monitor service also records all alarm state changes in a (rolling) file

```
/var/log/elt/ci-alarm-mon.log
```

10.4.5 Trigger Alarms

Since the inputs are OLDB datapoints, you can read/write them with the OLDB API, OLDB GUI, or OLDB tools (all described in [6]), for example:

```
oldb-cli read </absolute/path/of/input-dp> 2>/dev/null | tail -n1  
oldb-cli write </absolute/path/of/input-dp> <value> 2>/dev/null | tail -n1
```

Example 1

This examples uses a built-in demo alarm, so will work on every installation

```
oldb-cli read /alarm/sandbox/cii_diag/basic_alarm 2>/dev/null | tail -n1  
Value:  NOMINAL  
  
oldb-cli read /sandbox/cii_diag/alarm_source_bool 2>/dev/null | tail -n1  
Value:  false  
  
oldb-cli write /sandbox/cii_diag/alarm_source_bool true 2>/dev/null | tail -n1  
Value:  true
```

(continues on next page)



(continued from previous page)

```
oldb-cli read /alarm/sandbox/cii_diag/basic_alarm 2>/dev/null | tail -n1  
Value: ALARM_PRIORITY2  
  
oldb-cli write /sandbox/cii_diag/alarm_source_bool false 2>/dev/null | tail -n1  
Value: false
```

Example 2

This example uses the alarms demonstrated earlier, so it can only work if you have followed the above examples.

In terminal A:

```
ciiAlarmMon /etc/cii_ias
```

In terminal B:

```
# Current temperature  
oldb-cli read /harmoni/dcs1/detector/temperature1 2>/dev/null | tail -n1  
Value: 56.0  
  
# Increase to 90 - no alarm  
oldb-cli write /harmoni/dcs1/detector/temperature1 90 2>/dev/null | tail -n1  
  
# Increase temperature above limit - alarm state will change  
oldb-cli write /harmoni/dcs1/detector/temperature1 120.0 2>/dev/null | tail -n1  
  
# Decrease to 90 - alarm still on  
oldb-cli write /harmoni/dcs1/detector/temperature1 90.0 2>/dev/null | tail -n1  
  
# Decrease to normal - alarm state will change  
oldb-cli write /harmoni/dcs1/detector/temperature1 56 2>/dev/null | tail -n1
```



10.5 Advanced Topics

10.5.1 Administration

Installation

The CII-IAS-Bridge and Alarm tools are included in the ELT DevEnv distribution, but need to be post-installed before they can be used. The post-install routine will also install additional third-party libraries that are not included in the ELT DevEnv (these libraries are only needed on hosts where the CII Alarm System is running).

To set-up the alarm system on a host, do (as root):

```
/elt/ciisrv/postinstall/cii-postinstall alarm
```

Auto-start

Auto-start the Alarm system on boot (requires root privs)

```
sudo cii-services --beta enable alarm
```

Health Checks

Check status of Alarm system

```
cii-services --beta info
```

The relevant lines in the output are “ias”, “kafka”, “zookeeper”, and “alarm-mon”

Log Configuration

Redefine log levels (requires root privs)

```
# IAS-Bridge  
/elt/ciisrv/etc/log/cii-ias-bridge-logging.xml  
  
# Alarm Monitor  
/elt/ciisrv/etc/log/cii-alarm-mon-logging.json
```




Logs

Inspect Logs (requires root privs)

```
journalctl -e -u cii_ias -u cii_alarm_mon
```

10.5.2 Troubleshooting

1) Alarm-datapoint has value UNKNOWN (instead of NOMINAL etc.):

Try `ciiAlarmCtl daemon reload`, or (as root) `systemctl restart cii_ias`

2) To see the logs from the ias processes (requires root privs):

Use `journalctl -f -u cii_ias -u cii_alarm_mon -u kafka -u kafka-zookeeper`.

The ias core logs are in `/opt/IasRoot/logs`, `/var/log/kafka`, `/var/log/kafka-zookeeper`.

3) For debugging, it is possible to use the internal tools that come with IAS:

Set up the shell environment as described in #Run-IAS-directly. This enables you to run, e.g.:

- `ciiAlarmCtl check heartbeat`
- `iasDumpKafkaTopic`
- Ex.1: `iasDumpKafkaTopic -t plugin`
- Ex.2: `iasDumpKafkaTopic -t core | grep <AlarmId>`
- `ciiAlarmCtl check cdb`
- `iasCdbChecker -jCdb $ALARM_CDB_ROOT`

4) Visit the CII Knowledge Base: <https://gitlab.eso.org/ecs/eltsw-docs/-/wikis/KnowledgeBase/CII>

10.5.3 Modify CDB directly

The IAS offers a number of additional configuration options, but they require an understanding of the internals of the IAS. Before modifying the IAS CDB directly, read the detailed description of the Alarm system architecture [2], and for all details on alarm configuration attributes see [4] and [12].

Table 1 contains a brief explanation of the configuration elements used in IAS. For the proper explanation of the elements, refer to the IAS documentation [2], [3], [4].

Table 1: Alarm system elements



Entity Name	Description
ASCE	An alarm system computing element that runs inside DASU and contains one transfer function.
DASU	An alarm system element that runs inside Supervisor and contains one or many ASCEs.
IASIO	An alarm system input/output entity.
Supervisor	A standalone process that contains one or many DASUs.
Transfer Function	A transformation function that calculates alarm states. It receives and outputs IASIOs.
Plugin	A plugin for IAS system that interfaces the IAS system with external systems (in CII case OLDB).

General Settings

General settings are configured in `CDB/ias.json`

Set `LogLevel`, refresh rate, etc. according to the IAS CDB documentation [12].

IASIOs

IASIOs are input and output entities of the Alarm system. To define a new IASIO, go to `CDB/IASIO/iasios.json`

For the IASIO, write the URI of the input data point (e.g. `"/alarmtest/device/motor/input_int_dp"`, i.e. the relative path of the input data point `"cii.ldb:/alarmtest/device/motor/input_int_dp"`) and set the type (e.g. INT).

Create another IASIO for the output alarm, set the id (e.g. `"/motor/input_int_dp_alarm"`). The type must be set to "ALARM". You can define if the operator can suppress/hide the alarm by setting `canShelve` to "True".

Other values for both IASIOs can be left empty.

The fields `shortDesc`, `docURL` and `emails` are optional, they provide additional information on the IASIO. Also note that the fields `canShelve`, `sound`, and `emails` are only applicable to IASIOs of type ALARM. Template field allows templating of the IASIO (refer to the IAS documentation on templating).



Transfer Functions

Alarm system output IASIOs are calculated by Transfer Functions. Transfer functions are in the IAS domain and must be implemented and built there. IAS already provides a small set of Transfer functions. Please refer [8] to [9] and for details on Transfer functions. To import a new transfer function, go to `CDB/TF/tfs.json` and add a new entry.

To register a newly implemented transfer function, write the `className` (e.g. `"org.eso.ias.asce.transfer.impls.MinMaxThresholdTF"`) and `implLang` (e.g. `"SCALA"`).

DASUs

DASUs are containers for ASCE entities, that make alarm calculations using the Transfer functions. First DASU needs to be defined followed by its ASCE(s) with Transfer function to be used.

To define a new DASU, go to `CDB/DASU/` and create a new file, or copy an existing one.

Set the name of the DASU file (e.g. `"int_alarm_dasu.json"`). Leave the other properties empty for the time being. The default logging level will be used and there will be no templating of the DASU.

Go to `CDB/ASCE/` and define a new ASCE file (e.g. `"int_alarm_asce.json"`). For the `transferFunctionID` choose one of the transfer function names defined earlier (e.g. `"org.eso.ias.asce.transfer.impls.MinMaxThresholdTF"`). In the `outputId` field enter the output IASIO defined earlier (e.g. `"/motor/input_int_dp_alarm"`). For the `dasuId`, write the name of the DASU defined above, (e.g. `"int_alarm_dasu"`).

For the Inputs, add as an input the IASIO name (e.g. `"/alarmtest/device/motor/input_int_dp"`) to be the (only) input of the ASCE. In the properties, specify the parameters for the transfer function (e.g. `"org.eso.ias.tf.minmaxthreshold.highOn"` with value 5, `"org.eso.ias.tf.minmaxthreshold.highOff"` with value 1)

Note: For the `MinMaxThresholdTF` transfer function, analogous lower limits `minmaxthreshold.lowOff` and `minmaxthreshold.lowOn` can be defined. If both upper and lower limits are defined in the same ASCE, the alarm will not distinguish between the input value going above or below the limits. To distinguish between these two cases, one needs to define two ASCE with the same input, one checking for the high value and the other checking for the low value.

The Alarm priority is determined by the Transfer function. Some implementations allow the priority to be configurable, e.g. `org.eso.ias.asce.transfer.impls.MinMaxThresholdTF` has a property named `org.eso.ias.tf.alarm.priority` for that reason. Valid values are: `SET_CRITICAL`, `SET_HIGH`, `SET_MEDIUM` (default), `SET_LOW`.

To add the ASCE to deploy in the DASU, write the ASCE name into the DASU file (in `CDB/DASU/`) of the respective DASU (i.e. `int_alarm_asce`)

To define the output of the DASU, edit the DASU file, and write the `outputId` as defined in the IASIO file (i.e. `"/motor/input_int_dp_alarm"`).

Finally, save your configuration changes, and ask the Alarm system to re-read its config, as described in section #Usage above.



Supervisor

To deploy DASUs we need to define a Supervisor first. Supervisor is a stand-alone process that hosts DASUs among with its ASCEs including needed Transfer Function processing.

To define a new Supervisor, go to `CDB/Supervisors` and create a new file, or duplicate an existing one.

Set the supervisor's name (e.g. "test_supervisor"), and assign the hostname (e.g. "localhost"). In the field `DASUs to Deploy`, write the name of the DASU defined above, leave `templateId` and `instance` empty, meaning the DASU will not be processed as a template.

The `template` field defines what template to apply and the `instance` field defines the instance number used to apply when processing the template. For more information on templates refer to the IAS documentation.

Finally, save your configuration changes, and ask the Alarm system to re-read its config, as described in section #Usage above.

Plugin

The OLDB Plugin configuration foremost lists the set of IASIO IDs (each of them corresponding to an OLDB data point) to be monitored by the particular Plugin instance.

To configure how frequently the plugin should check for newly created input-DPs, set the property `dpconnRetrySec`, the default value is 60.

To define a new Plugin, go to `CDB/Plugins/` and create a new file, or duplicate an existing one. Name the new plugin "CiiPlugin" and set the monitored system to "OLDB". The `values` defines a set of data points to monitor. Add a new monitored value, and select the desired input. If needed change the "Refresh time" (for how long a value is valid if not updated) or other attributes.

Finally, save your configuration changes, and ask the Alarm system to re-read its config, as described in section #Usage above.

10.5.4 Run IAS directly

This is **not the recommended way**, because it is more involved and complicated. But if you cannot use the alarm system as a system service for some reason, it is possible to run it "in user space".

When you run the alarm system in this way, you have the freedom to store the alarm rules file(s) and CDB outside the standard location, i.e. not in `/etc/cii_ias`.

Run these commands as `eltdev` (or other user with write-permissions on the installation folders).

Set-up

```
mkdir -p $HOME/modulefiles  
UPDATES=https://www.eso.org/~eltnmgr/CII/latest/install
```

(continues on next page)



(continued from previous page)

```
wget -q -O $HOME/modulefiles/cii_ias.lua $UPDATES/ias/cii_ias.lua-4.0.0  
module load cii_ias
```

Start

```
ciiAlarmCtl start all+
```

Check

```
ciiAlarmCtl is-active all+
```

Reload

```
ciiAlarmCtl reload cdb
```

Stop

```
ciiAlarmCtl stop all+
```

TELEMETRY

Revision:	1.7
Status:	Released
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	ELT Telemetry
Folder:	trunk/deliverables/phase11
Document ID:	CSL-DOC-20-181458
File:	MAN_ELT_Telemetry.docx
Owner:	Jan Pribošek
Last modification:	October 15, 2020
Created:	March 8, 2020

Prepared by	Reviewed by	Approved by
Jan Pribošek (CSL SWE) Sara Pia Marinček (CSL) Domen Soklič (CSL)	Borut Terpinc (CSL SWE) Domen Soklič (CSL)	Gregor Čuk (CSL)

Document History



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 415 of 589

Re-vi-sion	Date	Changed/rev iewed	Sec-tion(s)	Modificatio n
0.1	09.03.2020	Marincek, jpribosek, dsoklic	All	Document creation.
1.0	06.05.2020	dsoklic/bte rpinc	All	Review updates. Released document.
1.1	17.6.2020	jpribosek	All	Updating manual after receiving support questions.
1.2	7.7.2020	jpribosek	All	Updating manual according to raised LAT issues.
1.3	30.7.2020	jpribosek	All	Updating manual according to raised LAT issues.
1.4	27.8.2020	jpribosek	All	Updating manual according to raised LAT issues.
1.5	3.9.2020	jpribosek	All	Updating manual according to raised LAT issues. Added additional section for setting up OLDB simulator.
1.6	24.9.2020	jpribosek	All	Updating manual according to raised LAT issues.
1.7	06.10.2020	jpribosek	All	Updating manual according to the review comments.

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is the manual for the Telemetry Service system used by applications in the ELT Core Integration Infrastructure Software.

Audience

This document is aimed at those Cosylab and ESO employees involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.

Glossary of Terms



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 416 of 589

API	Application Programming Interface
CII	Core Integration Infrastructure
CLI	Command Line Interface
DP	Data Point
EA	Engineering Archive
ES	ElasticSearch
JSON	Javascript object notation
OLDB	Online Database
SVN	Subversion
YAML	YAML Ain't Markup Language
ciiconfser-vicehost	Configuration Service hostname that should be resolved either by /etc/hosts file or the local DNS server.
ciarchivehost	Engineering Archive hostname that should be resolved either by /etc/hosts file or the local DNS server

References

1. Cosylab, ELT CII Configuration user manual, CSL-DOC-19-173189, v1.7
2. Cosylab, Middleware Abstraction Layer design document, CSL-DOC-17-147260 v 1.6
3. Elasticsearch website. <https://www.elastic.co/>
4. Cosylab, ELT CII – ICD Generation, User's Manual, CSL-MAN-19-170402, v1.4
1. Cosylab, CII Service Management Interface, CSL-DOC-18-163377, v1.2
2. Cosylab, ELT CII Online Database user manual, CSL-DOC-19-176283, v1.3
3. Cosylab, ELT CII Log user manual, CSL-DOC-20-181435, v1.3
4. Telemetry examples project: <https://svnhq8.hq.eso.org/p8/trunk/CONTROL-SYSTEMS/CentralControlSystem/CII/CODE/telemetry/telemetry-examples/>



11.1 Overview

This document is a user manual for usage of the Telemetry Service system. It explains how to use the Telemetry Archiver API, how to use the CLI tools, and how to configure the Telemetry Archiver.

All examples in this manual will be also presented in the telemetry-examples module.

11.2 Introduction

The purpose of the CII Telemetry Service is archiving the periodic, event-based and ad-hoc data that will be used for calculation of statistics, and the long-term trends and performance of the ELT telescope.

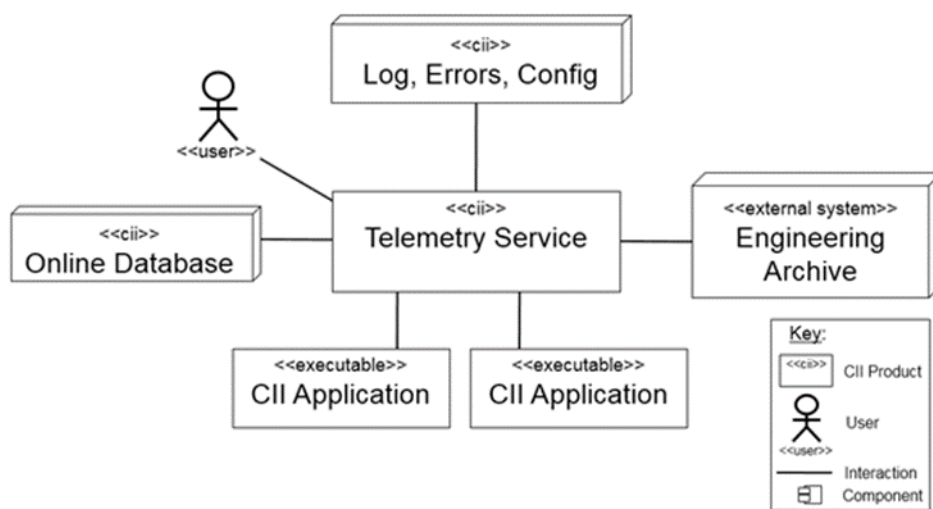


Figure 2-1: Telemetry Service API interactions

CII Telemetry Service is a central archiving service of the CII infrastructure. It is reliant on the other CII services:

- CII Online Database (OLDB) – the main source of data for archiving.
- CII Configuration Service (Config) – storing and retrieving data capture configurations.
- CII Log Service (Log) – used for storing Telemetry Service logs.
- CII Errors Service (Errors) – used for reporting Telemetry Service errors.

The data captured by the CII Telemetry service is stored into the Engineering Archive (EA).

The CII Telemetry Service consists of multiple CII Telemetry Archivers that are responsible for archiving the incoming telemetry data. The Telemetry Archiver is split into two parts:

1. **OLDB Subscription:** this part of the archiver subscribes to data point changes that are coming from the online database and archives them accordingly to the specified values in the data capture configurations. It also handles periodic archiving of the data point values according to the data capture configuration.
2. **Telemetry Client API:** this part of the archiver handles communication between the service and CII Applications. The Client API is split into two endpoints:
 - Management endpoint – management of archiver (stop archiver, restart archiver, refresh archiver’s configuration) and archiver statistics (number of data points that are being archived, data archive rate)

- Archive API endpoint – API used for storing ad-hoc data, querying, and downloading data from the Engineering Archive.

A high-level overview of a single CII Telemetry Archiver can be seen in Figure 2-2.

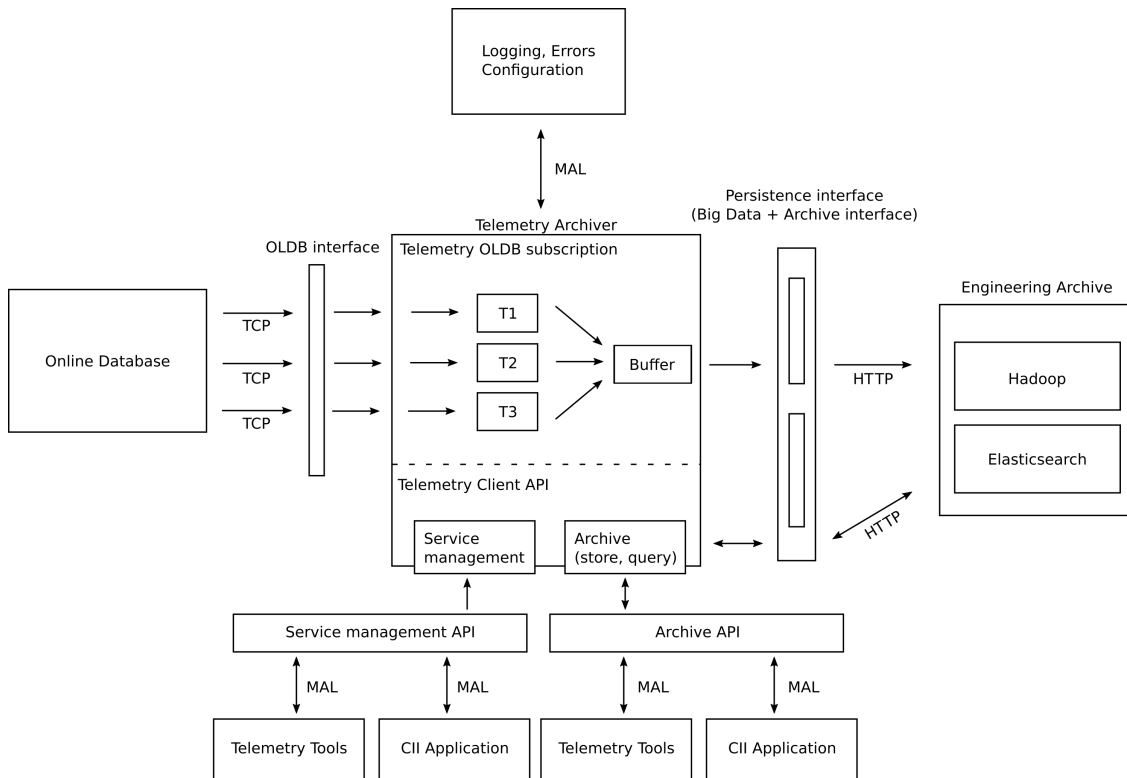


Figure 2-2: High-level overview of one Telemetry Archiver

The CII Telemetry Archiver accesses the CII Configuration service to retrieve its configuration. This configuration is used when the CII Telemetry Archiver receives data points to determine which should be archived and which should be discarded. When a data point that should be archived is received, it is stored in the Engineering Archive.

The Engineering Archive consists of two parts:

- Data Storage, which is used to archive smaller data (up to 2 GB); and
- Big Data Storage, which is used to archive large data.



11.2.1 Telemetry OLDB Subscription

CII Telemetry Archiver subscribes to OLDB data points via OLDB client and archives the data point values according to the data capture configurations that are stored in configuration service. Each data capture configuration is used to provide archiving options for one data point. It is possible to define multiple data capture rules for the same data point URI on the same CII Telemetry Archiver – e.g. data point could be archived on both value change and quality flag change on the same Telemetry Archiver.

Each Telemetry Archiver has defined its unique specified range in the form of a URI (i.e. *serviceRange*) that defines the location in configuration service from where the data capture configurations will be loaded from. For example, if the *serviceRange* field in Telemetry Archiver configuration is set to `cii.config:///telemetry/service1`, then all data capture configurations placed under `/telemetry/service1/datacapture/` URI on configuration service will be used (an example of valid uri would then be `cii.config:///telemetry/service1/datacapture/c722835e-cb22-4819-adb6-7cfe60c1ff39`). Extra note: data capture configurations should be placed under `<SERVICE RANGE>/datacapture/` uri. For each data point that will be archived a data capture configuration must be created. Since the data capture configurations are going to be stored in the CII Configuration service a `CIIBasicDataTypes` can be used for configuration fields. The data capture settings are described in Appendix B.

The changes of the data capture configurations are only propagated when the end-user calls the `refreshConfiguration` method on the Service Management API (see section 5.4). This will reload the specified configuration (or all configurations) without interrupting the ongoing archiving tasks except as required by the modified configuration. The configuration service will not be continuously polled for data capture configuration changes. When the Telemetry Archiver is started (initialized), the data capture configurations are downloaded via CII config client from the location specified in *serviceRange*. The Telemetry OLDB Subscription subscribes to data point changes via OLDB client. Internally the OLDB client maintains a pool of data point subscriptions.

Telemetry Archivers do not know about one another and it's the end user's responsibility to distribute data points across them depending on the expected service workload. This allows easy horizontal scaling of the Telemetry Archivers depending on the amount of data that we would like to archive. If more throughput is needed, additional Telemetry Archivers can be deployed.

Each Telemetry Archiver contains its configuration stored in the CII Configuration Service and connects to the CII Configuration Service using the CII Configuration Client. The client by default uses the `ciiconfservicehost` as the host location of CII Configuration Service (`ciiconfservicehost` is the hostname of the CII Configuration Service that should be resolved either by local DNS or custom `/etc/hosts` file). The default Telemetry Archiver configuration is stored under the `cii.config:///*/telemetry/config/defaultarchiverconfig`.

Multiple service range entries can be defined per one Telemetry Archiver configuration.



To specify the custom configuration for another Telemetry Archiver, a custom Telemetry Archiver configuration YAML must be provided. Appendix C contains a description of all configuration fields and an example of a custom YAML configuration.

Figure 2-3 shows an example of 2 Telemetry Archivers. Telemetry Archiver 1 uses the default configuration and the Telemetry Archiver 2 uses custom configuration. They both communicate with the configuration service to retrieve the service range and data capture configurations. Each Telemetry Archiver configuration specifies service ranges under which the data capture configurations are found. The data capture configuration, which is in subpath of specified service range, specifies the details about the behavior of the data capture (OLDB data point URI, delta value, ...).

When the Telemetry Archiver starts up, the data point values (CiiOldbDpValue) of the subscribed data points (as defined in data capture configurations) are downloaded from OLDB and stored in memory. Every time the data point value of the subscribed data point changes, they are checked against the data capture configurations to determine if the data point is suitable for archiving.

The subscribed data point will be archived depending on the following cases:

- **Value change:** if the data point value changed and if the difference between the old and new data point value has changed for more than the specified delta value and if duration since the last archive is bigger than the minimum interval duration (set in data capture configuration) the data point will be archived.
- **Metadata change:** if the data capture configuration is capturing based on the metadata change, the data point value will be archived whenever the metadata has changed
- **Quality change:** if the data capture configuration is capturing based on the quality change and the data point quality flag has changed, the data point value will be archived.
- **Maximum interval:** the purpose of the maximum interval is to archive the data point value in the absence of a data point value change event. If the maximum interval is set, a periodic task will be submitted to Telemetry Archiver's scheduled executor that will try to archive the last data point value according to the specified period. Timestamp since the last archive event for the given data capture configuration will be stored in the memory of the archiving worker.
- **Minimum interval:** the purpose of the minimum interval is to not archive too many data point value changes in a short timeframe. If the minimum interval is set and the duration between now and time, since the last value was archived, is bigger than the minimum interval duration, the data point value will be archived

It is possible to combine the above rules to make more fine-grained archiving behavior. For example, one can create two data capture configurations:

1. Data capture configuration with value archiving behavior and maximum time interval set to 10 seconds.
2. Data capture configuration with quality archiving behavior.

In such a case, the data point will be archived whenever the quality flag of the data point has changed



or whenever the value has changed or every 10 seconds if the value was not archived in the last 10 seconds.

If the data point is valid for archiving according to the rules described above, the archiving worker thread will transform a data point configuration into a data packet described and insert it into the archiving buffer. This operation preserves the data of the data point (such as quality flags and source timestamps). If no data point source timestamp exists, a current wall clock time is set with microsecond resolution. If no quality flag is associated with the data point data, an Unknown quality flag is assigned.

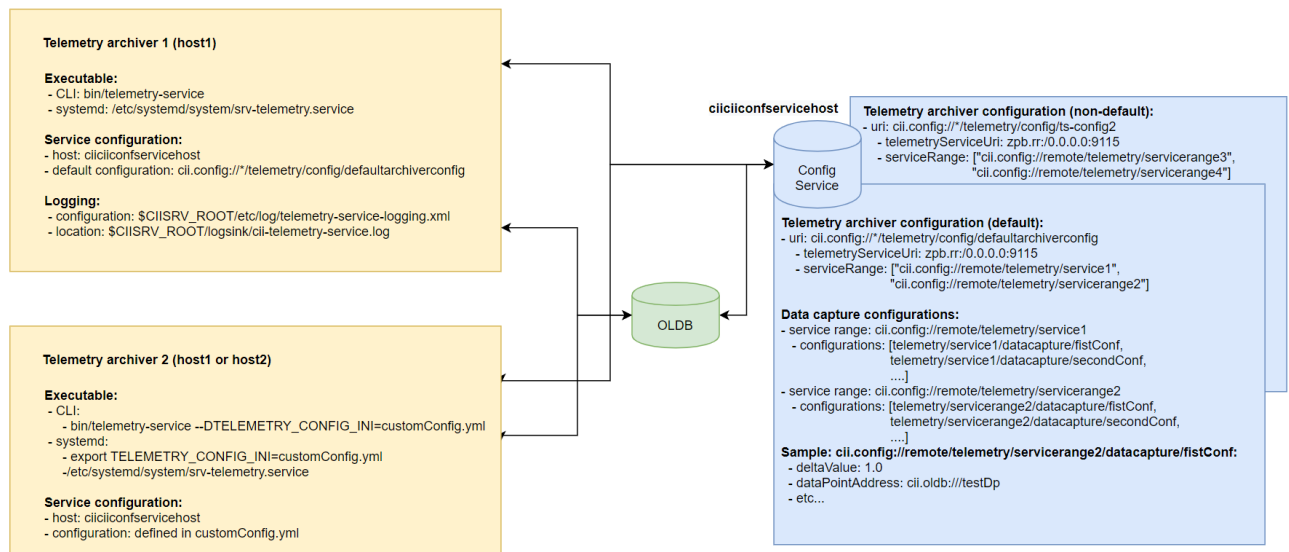


Figure 2-3: Communication and configuration of Telemetry Archivers

Data Packet

The Data Packet is a data transfer object used for transferring data point values and ad-hoc data between Telemetry Archiver and its clients. The Telemetry Archiver clients could be written in one of the three supported languages (Java, C++, Python). For this reason, the data packet is generated with the MAL ICD [4]. The Data Packet contains the following fields:

Table 2-1: Data packet fields description



Field	Type	Description	Example
Timestamp (auto-matically assigned)	Double	Representing the point in time when the data packet was archived (in UNIX time)	1571914537.0
Data (required)	String	Data is stored as a valid JSON string	{"temperature": 30, "comment": "ELT motor temperature"}
Big Data File Reference (optional)	String	A reference / location of the big data file in the large storage service. The client has to store a large file into the large storage service before archiving the data packet	

11.2.2 Telemetry Client API

Telemetry Client API provides an interface for easier interaction with the archiving system. The API could be used by the command line tools and other CII Applications. Telemetry Client API consists of three parts:

- **Archive API** – API used for storing, querying, and retrieving data from the Engineering Archive.
- **Service Management API** – API used for managing Telemetry Archiver and providing important statistics about the Telemetry Archiver internals.
- **Telemetry Tools** – CLI tools that make interaction with the Telemetry Archiver easier.

In order to provide an API to CII Applications that are written in different languages, a MAL ZPB service is used internally for handling the client requests. Client API service exposes two ICD interfaces:

- Archive API interface, the service is exposed on:
zpb.rr://server:9115/telemetry/service/archive
- Service Management interface, the service is exposed on:
zpb.rr://server:9115/telemetry/service/management

Archive API

Archive API shall be used by CLI tools and other CII Applications that need to interact with the Engineering Archive. Internally an Archive API is communicating with the archiving system through the Persistence Interface using the Telemetry Persistence Implementation.

Archive API will receive ad-hoc data from the clients in the form of data packets which will be stored irrespective of the age of the data into the EA.

Archiving API does not support storing data packets that contain large data blobs directly as that would result in a lot of unnecessary network traffic. The clients that will be using the Archiving API will



have to first store the large files into the Large Storage Service and then store the data packet with the location reference of the blob into the Engineering Archive.

To make storing large data files to Large Storage Service easier, a library implementing the Large Storage Service interface is provided. The provided implementation will interact with the Hadoop storage system.

Archiving API handles any valid data package irrespective of the sampling frequency of the data. The actual speed of storing large big data blobs depends on the used hardware.

Service Management API

The CII Telemetry Archiver is a *systemd* service and should be started, stopped or restarted via *systemd* commands:

- **Start** - start the service if it's not already running
- **Restart** - restart the Telemetry Archiver. All currently ongoing tasks will finish before the service will be restarted.
- **Stop** - shutdown the Telemetry Archiver. All currently ongoing tasks will finish before the service is stopped

The instructions for manipulating the CII Telemetry Archiver with *systemd*, please check the CII Telemetry Transfer document.

The Service Management Client API exposes a *RefreshConfiguration* command which refreshes the specified data capture configurations from the configuration service without interrupting the ongoing archiving tasks.

As part of the Service Management API, metrics about the current state of the Telemetry Archiver are exposed via the JMX interface. For more information about the JMX interface see the design document for Service Management Interface [5]. For the Telemetry Archiver, the following metrics are exposed:

General metrics:

- Current time.
- Telemetry Archiver uptime.
- Current Memory consumption.
- Current CPU consumption.
- The current number of archiving worker threads.
- Engineering Archive health (periodic health checks reporting if the archiving system is still on-line).

Archiving worker threads:

- The number of currently subscribed data points.



- **The number of current data points that are being periodically** archived.
- **The number of dropped data point values (archiving buffer was full** and a data point value was dropped).

Archiving buffer statistics:

- Buffer fullness - the percentage of the buffer that is full.
- Archiving throughput - number of archived data points in the last hour.
- The number of errors that occurred while storing packets into the Engineering Archive in the last hour.
- The total number of archived data points for the current day.

Telemetry CLI tools

The Telemetry CLI tools are separated into 2 modules:

- **CLI Manager is a tool that allows a user to send the Telemetry Archiver** a notification to trigger a refresh of its configurations.
- **CLI Archiver is a tool enabling users to store, download, and query** data packets from and to the Engineering Archive. The tool provides the user with three different commands: store, download, and query.



11.3 Prerequisites

This section describes the prerequisites for using the Telemetry Service, API, and applications. Note that the preparation of the environment (i.e. installation of required modules, configuration, running of services) is not in the scope of this document. Consult your administrator and the Telemetry Service TRD document for information on how to establish the environment for usage of the Telemetry Archiver, API, and applications. Before using the CII Telemetry make sure that telemetry configuration is initialized with *telemetry-initEs* command.

11.3.1 Services

Engineering Archive

The Engineering Archive (EA) stores the data that the Telemetry Archiver receives and decides to archive. The Engineering Archive is, in fact, a set of services. The EA has been initially designed to use Persistence interface to access Elasticsearch and Hadoop, but can be modified to use other forms of storage.

Configuration Service

The CII Telemetry Service uses CII Configuration Service to store its configuration. The Configuration service must be accessible for the Telemetry Archiver to successfully initialize.

OLDB Service

The Online Database service is the main source of data that provides data point updates which are checked against the capturing rules defined by the Telemetry Archiver (TA) and later archived in the archiving system. For testing purposes, the OLDB service could be replaced with the OLDB simulator that generates the data point value updates based on the capturing rules.



11.4 Data Capture Configuration

11.4.1 Create Data Capture Configuration

Before the Telemetry Archiver can archive data packets, it must be configured. The data capture configuration can be done using the Configuration GUI.

The following steps are necessary for configuration:

1. In the GUI click on Options and tick the Write enabled checkbox.
2. In the Target Config tab click on the Add button.
3. The dialog shown in Figure 4-1 will open. In this dialog the new configuration can be specified.

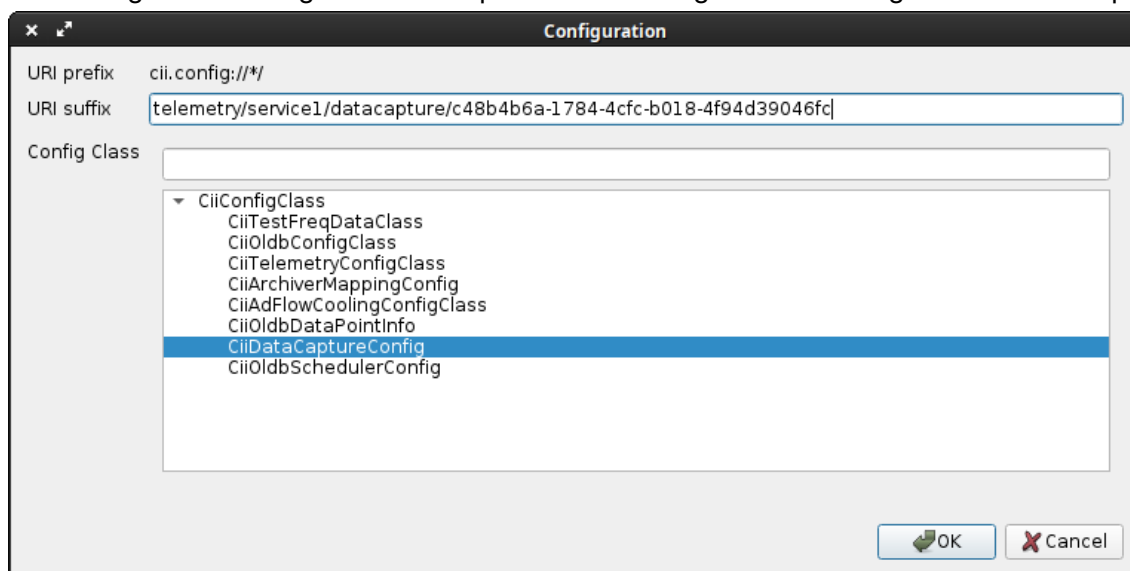


Figure 4-1: Creation of new configuration

4. Specify the following:
 - a. URI suffix: The suffix needs to follow the following format: <SERVICE_RANGE*>*/datacapture/<UUID>. The first part needs to be the range under which the Telemetry Archiver is running (e.g. telemetry/service1 in the example). The second part needs to be a unique ID. A tool such as uuidgen can be used to generate this UID.

NOTE: when generating a unique ID with the uuidgen tool, use only the generated ID that starts with a letter, regenerate an ID if it starts with a number.
 - b. Config Class: Select CiiDataCaptureConfig.
5. Click Ok. The data capture configuration should now be stored on the configuration service.
6. Edit the created data capture configuration fields by following the steps in the next section 4.2.



11.4.2 Edit Data Capture Configuration

1. Select the data capture configuration in the tree menu in the Target Config tab.
2. Fill in the configuration parameters. The example in Figure 4-2 shows a capture configuration that is storing value changes. The individual fields are listed and documented in Table 9.4.

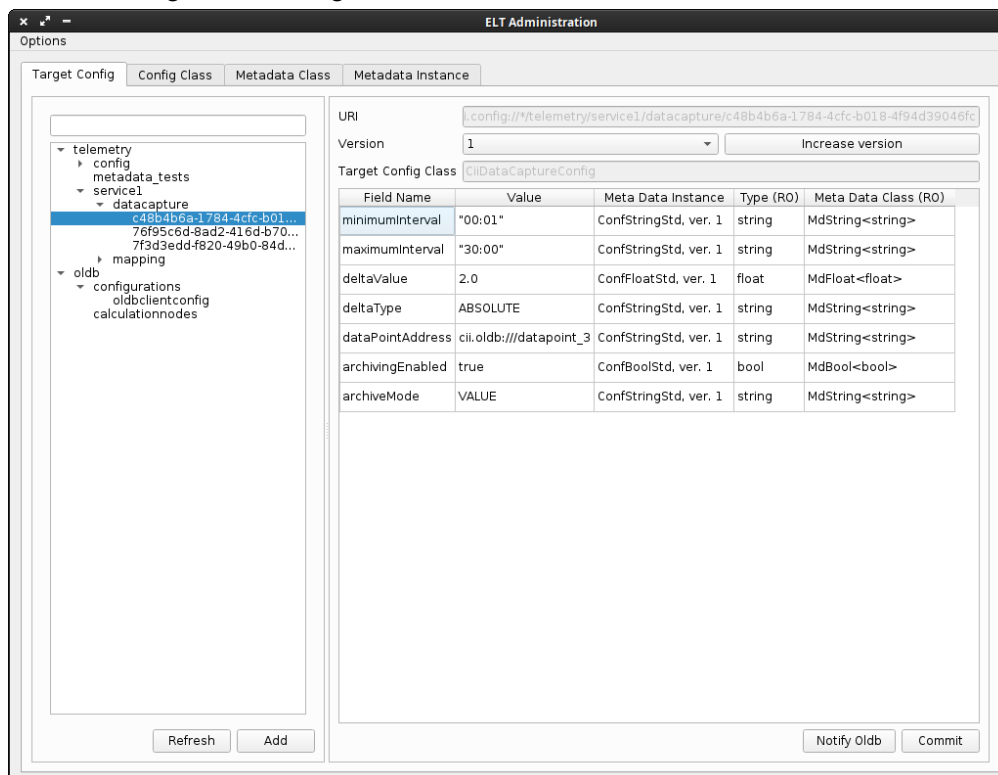


Figure 4-2: Configuring a data capture configuration

3. Click on the Commit button.
4. Before the Telemetry Archiver can use the new configuration, Telemetry Archiver must be notified/refreshed about the updated configuration. For that, one can use the CLI tool described in section 6.1. The ID that the CLI tool expects for refreshing configuration is the URI of the configuration that is selected in the tree view on the left side of the GUI.



11.5 Telemetry Archiver API Usage

This section provides examples on how to use the Archive API and Service management API in an application. The Archive API provides users with the functionality to store, query, and download data as well as write, read, and delete big data that is saved in the Engineering Archive. The Service management API provides users with the functionality to refresh configurations.

All examples assume that CII Telemetry Archiver is running on localhost. This can be changed to the *ciiarchivehost* accordingly.

11.5.1 Basic Example

The following section demonstrates a basic step-by-step example workflow of using the Telemetry API.

The basic example starts by first configuring the data capture. Once the configuration service has processes and saved the new configuration, we refresh the configuration on the Telemetry Archiver. After a short while, the Telemetry Archiver will subscribe to the data point.

At this point, the Telemetry Archiver is archiving data points that match the provided configuration. The basic example continues by connecting to the Archive API and queries it to retrieve the archived data point IDs. The retrieved data point IDs are used to download the data packets. Upon downloading the packets, the data points data is displayed.

The APIs are documented in sections 5.2, 5.3, and 5.4.

Listing 5-1 shows this basic example in Java (*TelemetrySubscriptionExample.java*). This can be done in CPP and Python as well.

Listing 5-1: The basic example in Java

```
import java.net.URI;
import elt.mal.CiiFactory;
import elt.mal.rr.qos.QoS;
import elt.mal.rr.qos.ReplyTime;
import elt.mal.zpb.ZpbMal;
import elt.telemetry.archive.DataPacket;
import elt.telemetry.archive.TelemetryArchiveApiSync;

public class TelemetrySubscriptionExample {
    public static void main(String[] args) throws Exception {
        final URI telemetryServiceRange = URI.create("cii.config://*/telemetry/
↵service1");
        final URI uniqueCaptureLocation = URI
            .create(String.format("%s/%s", telemetryServiceRange, UUID.randomUUID().
↵toString()));
        final URI oldbDatapointUri = URI.create("cii.oldb:///datapoint_1");
        final URI telemetryManagementEndpoint = URI
            .create("zpb.rr://localhost:9115/telemetry/service/management");
```

(continues on next page)



(continued from previous page)

```
final URI telemetryArchiveEndpoint = URI
    .create("zpb.rr://localhost:9115/telemetry/service/archive");
// start of the actual program
CiiConfigClient configClient = CiiConfigClient.getInstance();
configClient.setWriteEnabled(true);
DataCaptureConfiguration captureConfig =
↪createCaptureConfig(olddbDatapointUri);
configClient.saveTargetConfig(uniqueCaptureLocation, captureConfig);
// wait for the configuration service to process and save request
Thread.sleep(2_000);
// refresh specific data capture configuration
refreshConfiguration(Arrays.asList(uniqueCaptureLocation.toString()),
    telemetryManagementEndpoint);
// wait for the telemetry archiver to subscribe to data point (if it's not
↪already)
Thread.sleep(2_000);
// query the stored data point from the archive
TelemetryArchiveApiSync archiveApi =
↪connectToArchiveApi(telemetryArchiveEndpoint);
List<String> ids = archiveApi.queryData("data.uri:*datapoint_1*");
assert ids.size() > 0 : "At least one data point should be present";
// download queried ids
List<DataPacket> packets = archiveApi.downloadData(ids);
for (DataPacket packet : packets) {
    // process packets
    System.out.println(packet.getData());
}
// close resources
configClient.close();
archiveApi.close();
System.exit(0);
}
}

public static void refreshConfiguration(List<String> refreshIds, URI
↪managementEndpoint) {
    CiiFactory factory = null;
    TelemetryServiceManagementSync client = null;
    try {
        factory = CiiFactory.getInstance();
        factory.registerMal("zpb", new ZpbMal());
        client = factory.getClient(managementEndpoint,
            new QoS[]{new ReplyTime(10, TimeUnit.SECONDS)}, new Properties(),
            TelemetryServiceManagementSync.class);
        if (refreshIds.isEmpty()) {
            client.refreshAllConfigurations();
        } else {
            client.refreshConfigurations(refreshIds);
        }
    }
```

(continues on next page)



(continued from previous page)

```
    } catch (CiiSerializableException e) {
        String msg = String.format("Failed to refresh configurations %s", e.
↪getCiiMessage());
        System.out.println(msg);
    } catch (Exception e) {
        System.out.println("Failed to refresh configuration: " + e);
    } finally {
        if (client != null) {
            client.close();
        }
        if (factory != null) {
            factory.close();
        }
    }
}

public static TelemetryArchiveApiSync connectToArchiveApi(URI_
↪telemetryClientEndpoint) throws Exception {
    CiiFactory factory = CiiFactory.getInstance();
    factory.registerMal("zpb", new ZpbMal());
    TelemetryArchiveApiSync client = factory.getClient(telemetryClientEndpoint,
        new QoS[]{new ReplyTime(10, TimeUnit.SECONDS)}, new Properties(),
        TelemetryArchiveApiSync.class);
    CompletableFuture<Void> connectionFuture = client.asyncConnect();
    connectionFuture.get(5, TimeUnit.SECONDS);
    return client;
}

public static DataCaptureConfiguration createCaptureConfig(URI oldbDatapointUri)

    throws CiiConfigDataLimitsException, CiiInvalidTypeException {
    DataCaptureConfiguration captureConfig = new DataCaptureConfiguration(
        "dataCaptureConfigTest", // name
        "", // id is automatically set on store
        oldbDatapointUri.toString(), // to which data point we want to subscribe_
↪to
        1.0f, // delta value change
        "ABSOLUTE", // delta type set to absolute changes
        "VALUE", // archive mode set to value changes
        "00:01", // minimum interval (at most once per second)
        "10:00", // maximum interval (at least once every 10 minutes)
        true // archiving enabled
    );
    return captureConfig;
}
```



11.5.2 Archive API: storing, querying and downloading data

A short description of the part of the Archive API used for storing, querying and downloading data can be found in A.1.

Imports/includes

Usage of the Telemetry Archive API requires the following imports to establish a MAL connection: CiiFactory, QoS, ReplyTime, and ZpbMal. The TelemetryArchiveApiSync interface defines the Telemetry Archive API methods.

An example of an application that uses the Telemetry Archive API can be found in the main method of the *TelemetryArchiveAPIExample.java* file in the telemetry-examples project (see [8]). Similar examples can be found in `cpp/archive_api` and `python/archive_api` folders.

For a list of necessary libraries in order to connect to the API make sure to check the wscript files in `cpp`, `java` and `python` folders of the telemetry-examples project.

Java

```
import java.net.URI;
import elt.mal.CiiFactory;
import elt.mal.rr.qos.QoS;
import elt.mal.rr.qos.ReplyTime;
import elt.mal.zpb.ZpbMal;
import elt.telemetry.archive.DataPacket;
import elt.telemetry.archive.TelemetryArchiveApiSync;
import com.fasterxml.jackson.databind.ObjectMapper;
```

CPP

```
#include <mal/Cii.hpp>
#include <mal/utility/LoadMal.hpp>
#include <mal/rr/qos/ReplyTime.hpp>
#include <CiiTelemetryArchiveApi.hpp>
```




Python

```
import elt.pymal as mal
from ModCiiTelemetryArchiveApi.Elt.Telemetry.Archive import DataPacket
from ModCiiTelemetryArchiveApi.Elt.Telemetry.Archive.
↳TelemetryArchiveApi import TelemetryArchiveApiSync
```

Connection to Telemetry Archiver using MAL

To establish a MAL connection between our application and the Archive API endpoint, we first get an instance of CiiFactory, register a MAL connection on it and get a TelemetryArchiveApiSync client that connects to the URI **zpb.rr://localhost:9115/telemetry/service/archive**. We then wait until the connection is established.

Listing 5-2, Listing 5-3 and Listing 5-4 demonstrate how to establish a MAL connection between an application and the Archive API endpoint.

Connection in Java

Listing 5-2: Example of connection to the Telemetry Archiver using Java

```
CiiFactory factory = CiiFactory.getInstance();
Mal mal = new ZpbMal();
factory.registerMal("zpb", mal);
final URI endpointUri = URI.create("zpb.rr:// localhost:9115/telemetry/service/
↳archive");
TelemetryArchiveApiSync client = factory.getClient(endpointUri,
    new QoS[]{new ReplyTime(10, TimeUnit.SECONDS)}, new Properties(),
    TelemetryArchiveApiSync.class);
CompletableFuture<Void> connectionFuture = client.asyncConnect();
connectionFuture.get(10, TimeUnit.SECONDS);
```

Connection in CPP

Listing 5-3: Example of connection to the Telemetry Archiver using CPP

```
int status = 0;
const elt::mal::Uri archiveApiEndpoint("zpb.rr://localhost:9115/telemetry/
↳service/archive");

// Register MAL and obtain client for archive API
std::shared_ptr<::elt::mal::Mal> malInstance = elt::mal::loadMal("zpb", {});
::elt::mal::CiiFactory &factory = elt::mal::CiiFactory::getInstance();
factory.registerMal("zpb", malInstance);
```

(continues on next page)



(continued from previous page)

```
std::unique_ptr<elt::telemetry::archive::TelemetryArchiveApiSync> client =
    factory.getClient<elt::telemetry::archive::TelemetryArchiveApiSync>(
        archiveApiEndpoint,
        {std::make_shared<elt::mal::rr::qos::ReplyTime>(std::chrono::seconds(10))},
        {});

if (client == nullptr) {
    std::cerr << "Could not obtain client for: " << archiveApiEndpoint.string() <<
    ↪std::endl;
    status = 1;
} else {
    ::elt::mal::future<void> connectionFuture = client->asyncConnect();
    std::future_status futureStatus = connectionFuture.wait_
    ↪for(boost::chrono::seconds(10));
    if (futureStatus != boost::future_status::ready) {
        std::cerr << "Connection timeout to " << archiveApiEndpoint.string() <<
    ↪std::endl;
        status = 2;
    }
}
```

Connection in Python

Listing 5-4: Example of connection to the Telemetry Archiver using Python

```
archive_api_endpoint = mal.Uri('zpb.rr://localhost:9115/telemetry/service/archive
↪')

# Initialize Mal
factory = mal.CiiFactory.getInstance()
mal_instance = mal.loadMal('zpb', {})
if mal_instance is None:
    raise RuntimeError('Could not load zpb mal')
factory.registerMal('zpb', mal_instance)
qos = [mal.rr.qos.ReplyTime(datetime.timedelta(seconds=10))]

# Obtain client and connect to Telemetry service
client = factory.getClient(archive_api_endpoint, TelemetryArchiveApiSync,
                           qos, {})
connection_future = client.asyncConnect()
if connection_future.wait_for(datetime.timedelta(seconds=10)) == mal.
↪FutureStatus.TIMEOUT:
    raise RuntimeError('Could not to connect to %s, timeout' % (archive_api_
↪endpoint.string(),))
_ = connection_future.get()
```



Storing a data packet

The data packet should contain the following fields:

- archived: timestamp when the data packet was archived
- bigDataLocation: reference to the big data file stored in Hadoop. This is an optional field.
- data field that contains:
 - timestamp: data point timestamp in microsecond resolution
 - uri: data point URI
 - value: data point value
 - quality: data point quality flag
 - metadataReference: data point metadata reference

When storing a data packet via the Archive API, the data point URI could be an arbitrary string and the URI protocol does not have to start with `cii.oldb`.

Here is an example of how the subscribed data point value update will be serialized in ElasticSearch (fields prefixed with `_` are autogenerated ElasticSearch internal fields and are not fetched when retrieving data):

Listing 5-5: Archived metadata

```
{
  "_id" : "5HEb9m8Ba7LLWT1elmN_",
  "_source" : {
    "archived" : 1580382019114000,
    "data" : {
      "@type" : "MdOldbNumber",
      "value" : 2.0,
      "comment" : "",
      "checked" : false,
      "metadataInstanceVersion" : 1,
```



```
"@name" : "telemetry_metal",  
"@genType" : "SINGLE",  
"metadataInstance" : "telemetry_metal"  
}  
}  
}
```

Listing 5-6: Archived data point value

```
{  
  "_id" : "5XEb9m8Ba7LLWT1emGP8",  
  "_source" : {  
    "archived" : 1580382019737000,  
    "data" : {  
      "uri" : "cii.oldb:///datapoint_1",  
      "metadataReference" : "5HEb9m8Ba7LLWT1elmN_",  
      "quality" : "OK",  
      "value" : 0.5,  
      "timestamp" : 1580382018598000  
    }  
  }  
  "bigDataLocation" : "4a9a3810-1efb-4d39-b8fb-751f09c91ea8"  
}
```



Java client

To store data, we first create a `DataPacket` object containing metadata. In the provided example, the static method `TelemetryArchiveAPIExample.createTestMetadata` serves this purpose. It first creates the data packet class and then sets its fields: a timestamp is applied, the location of the data in the big data database of the Engineering Archive is specified and the data is added in JSON format.

The following sample code is an excerpt from the code file named `TelemetryArchiveAPIExample.java`, stored in the `java/sample-app` folder.

Listing 5-7: Example of data packet metadata creation in Java

```
public static DataPacket createTestMetadata(Mal mal) {
    DataPacket dataPacket = mal.createDataEntity(DataPacket.class);
    dataPacket.setArchived(Instant.now().toEpochMilli());
    dataPacket.setBigDataLocation("");

    Map<String, Object> map = new HashMap<>();
    map.put("@type", "MyOldbNumber");
    map.put("value", 2.0);
    map.put("comment", "");
    map.put("checked", false);
    map.put("metadataInstanceVersion", 1);
    map.put("@name", "telemetry_metal");
    map.put("@genType", "SINGLE");
    map.put("metadataInstance", "telemetry_metal");
    JSONObject json = new JSONObject(map);
    dataPacket.setData(json.toString());

    return dataPacket;
}
```

To store the created data packet, we call the `storeData` method of the `TelemetryArchiveApi` interface. It returns a string specifying the assigned ID of the data packet containing metadata.

Listing 5-8: Example of storing of metadata data packet in Java

```
DataPacket metadata = TelemetryArchiveAPIExample.createTestMetadata(mal);
String metadataReference = client.storeData(metadata);
```

We proceed by creating another `DataPacket` object representing the data point value by passing the ID of the data packet containing the corresponding metadata (`metadataReference`) to the static method `TelemetryArchiveAPIExample.createTestDataPointValue`.

Listing 5-9: Example of creating a data packet in Java

```
public static DataPacket createTestDataPointValue(Mal mal, String
↪metadataReference) {

    DataPacket dataPacket = mal.createDataEntity(DataPacket.class);
```

(continues on next page)



(continued from previous page)

```
dataPacket.setArchived(Instant.now().toEpochMilli());
dataPacket.setBigDataLocation("bigDataLocation");

Map<String, Object> map = new HashMap<>();
map.put("uri", "cii.client:///datapoint_1");
map.put("metadataReference", metadataReference);
map.put("quality", "OK");
map.put("value", 0.5);
map.put("timestamp", Instant.now().toEpochMilli() - 10_000);
JSONObject json = new JSONObject(map);
dataPacket.setData(json.toString());

return dataPacket;
}
```

To store the created data packet, we call the `storeData` method of the `TelemetryArchiveApi` interface. It returns a string specifying the assigned ID of the data packet in the Engineering Archive.

Listing 5-10: Example of storing of data packet in Java

```
DataPacket dataPointValue =
    TelemetryArchiveAPIExample.createTestDataPointValue(mal, metadataReference);
String id = client.storeData(dataPointValue);
```

CPP client

The following listings contain a sample CPP client that stores a datapacket. The sample code is an excerpt from the code file named `app.cpp` stored in the `cpp/archive-api` folder. The code follows the same flow as the Java code.

First, a metadata `DataPacket` is created.

Listing 5-11: Example of data packet metadata creation in CPP

```
std::shared_ptr<elt::telemetry::archive::DataPacket> createTestMetadata(
    std::unique_ptr<elt::telemetry::archive::TelemetryArchiveApiSync> &client) {
    std::shared_ptr<elt::mal::Mal> malInstance = client->getMal();
    std::shared_ptr<elt::telemetry::archive::DataPacket> dataPacket =
        malInstance->createDataEntity<elt::telemetry::archive::DataPacket>();
    long timestamp = timestampMillis();
    dataPacket->setArchived(timestamp);
    dataPacket->setBigDataLocation("");
    boost::property_tree::ptree tree;
    tree.put("@type", "MyOldbNumber");
    tree.put("value", 2.0);
    tree.put("comment", "");
    tree.put("checked", false);
}
```

(continues on next page)



(continued from previous page)

```
tree.put("metadataInstanceVersion", 1);
tree.put("@name", "telemetry_metal");
tree.put("@genType", "SINGLE");
tree.put("metadataInstance", "telemetry_metal");
std::stringstream data;
boost::property_tree::json_parser::write_json(data, tree);
dataPacket->setData(data.str());
return dataPacket;
}
```

To store the data packet, we call the `storeData` method of the client.

Listing 5-12: Example of storing of metadata data packet in CPP

```
std::shared_ptr<elt::telemetry::archive::DataPacket> metadata =
↳createTestMetadata(client);
std::string metadataReference = client->storeData(metadata);
```

We proceed by creating another `DataPacket` object representing the data point value.

Listing 5-13: Example of creating a data packet in CPP

```
std::shared_ptr<elt::telemetry::archive::DataPacket> createTestDatapointValue(
    std::unique_ptr<elt::telemetry::archive::TelemetryArchiveApiSync> &client,
    const std::string &metadataReference) {
    std::shared_ptr<elt::mal::Mal> malInstance = client->getMal();
    std::shared_ptr<elt::telemetry::archive::DataPacket> dataPacket = malInstance->
↳createDataEntity<elt::telemetry::archive::DataPacket>();
    long timestamp = timestampMillis();
    dataPacket->setArchived(timestamp);
    dataPacket->setBigDataLocation("bigDataLocation");
    boost::property_tree::ptree tree;
    tree.put("uri", "cii.client:///datapoint_1");
    tree.put("metadataReference", metadataReference);
    tree.put("quality", "OK");
    tree.put("value", 9.5);
    tree.put("timestamp", timestamp - 10000.0);
    std::stringstream data;
    boost::property_tree::json_parser::write_json(data, tree);
    dataPacket->setData(data.str());
    return dataPacket;
}
```

To store the data packet, we call the `storeData` method of the client.

Listing 5-14: Example of storing of the data packet in CPP

```
std::shared_ptr<elt::telemetry::archive::DataPacket> datapoint =
↳createTestDatapointValue(client, metadataRef);
std::string dataReference = client->storeData(datapoint);
```



Python client

The following listings contain a sample Python client that stores a data packet. The sample code is an excerpt from the code file named `cii-telemetry-examples-archive-py.py` stored in the `python/archive_api` folder. The code follows the same flow as the Java code.

First, a metadata `DataPacket` is created.

Listing 5-15: Example of data packet metadata creation in Python

```
def create_test_metadata(client):
    mal_instance = client.getMal()
    data_packet = mal_instance.createDataEntity(DataPacket)
    data_packet.setArchived(datetime.datetime.now().timestamp() * 1000.0)
    data_packet.setBigDataLocation('')
    data = {'@type': 'MyOldbNumber',
           'value': 2.0,
           'comment': '',
           'checked': False,
           'metadataInstanceVersion': 1,
           '@name': 'telemetry_metal',
           '@genType': 'SINGLE',
           'metadataInstance': 'telemetry_metal'}
    data_packet.setData(json.dumps(data))
    return data_packet
```

To store the data packet, we call the `storeData` method of the client.

Listing 5-16: Example of storing of metadata data packet in Python

```
metadata = create_test_metadata(client)
metadata_reference = client.storeData(metadata)
```

We proceed by creating another `DataPacket` object representing the data point value.

Listing 5-17: Example of creating a data packet in Python

```
def create_test_datapoint_value(client, metadata_reference):
    mal_instance = client.getMal()
    # timestamp must be in millis
    timestamp = datetime.datetime.now().timestamp() * 1000.0
    data_packet = mal_instance.createDataEntity(DataPacket)
    data_packet.setArchived(timestamp)
    data_packet.setBigDataLocation('bigDataLocation')
    data = {'uri': 'cii.client:///datapoint_1',
           'metadataReference': metadata_reference, # WATCH OUT
           'quality': 'OK',
           'value': 0.5,
           'timestamp': timestamp - 10000.0}
    data_packet.setData(json.dumps(data))
    return data_packet
```




To store the data packet, we call the `storeData` method of the client.

Listing 5-18: Example of storing of data packet in Python

```
datapoint_value = create_test_datapoint_value(client, metadata_reference)
packet_id = client.storeData(datapoint_value)
```

Storing a list of data packets

To store a list of data packets, we first create a list of `DataPacket` objects representing data point values. Then we call the `storeDataList` method of the `TelemetryArchiveApi` interface. It returns a list of strings specifying the assigned IDs of the data packets in the Engineering Archive.

Listing 5-19, Listing 5-20 and Listing 5-21 contain a simple client that stores a list of data packets in the Java, CPP and Python programming languages respectively.

Java client

Listing 5-19: Example of storing a list of data packets in Java

```
List<DataPacket> dataPackets = new ArrayList<>();
metadata = TelemetryArchiveAPIExample.createTestMetadata(mal);
metadataReference = client.storeData(metadata);
for (int i = 0; i < 2; i++) {
    dataPointValue = TelemetryArchiveAPIExample.
    ↪createTestDataPointValue(metadataReference);
    dataPackets.add(dataPointValue);
}
List<String> ids = client.storeDataList(dataPackets);
```

CPP client

Listing 5-20: Example of storing a list of data packets in CPP

```
std::vector<std::shared_ptr<elt::telemetry::archive::DataPacket>> dataPackets;
std::shared_ptr<elt::telemetry::archive::DataPacket> metadataInstance =
    ↪createTestMetadata(client);
std::string metadataRef = client->storeData(metadataInstance);
for (std::size_t i = 0; i < 2; i++) {
    std::shared_ptr<elt::telemetry::archive::DataPacket> datapoint =
    ↪createTestDataPointValue(client, metadataRef);
    dataPackets.push_back(datapoint);
}
std::vector<std::string> ids = client->storeDataList(dataPackets);
```



Python Client

Listing 5-21: Example of storing a list of data packets in Python

```
data_packets = []
metadata = create_test_metadata(client)
metadata_reference = client.storeData(metadata)
for i in range(2):
    datapoint_value = create_test_datapoint_value(client, metadata_reference)
    data_packets.append(datapoint_value)
ids = client.storeDataList(data_packets)
```

Querying data packets

To query data packets in the Engineering Archive we call the `queryData` method of the `TelemetryArchiveApi` interface. It takes a string representing an ES Query DSL [3] query and returns a list of strings specifying the assigned IDs of the data packets in the Engineering Archive. In the provided example, we query all data packets that have the `data.metadataReference` field set to the ID of the data packet containing metadata that was stored last. When querying the EA using the ES Query DSL, the current limitation is that it currently does not support sorting the results.

Note: When the Elasticsearch is used as the Engineering Archive, some time needs to elapse between storing the data and querying it. This time can be up to 1 second depending on the settings of the Elasticsearch.

Listing 5-22, Listing 5-23 and Listing 5-24 contain a sample Java client that queries the Archive API for a data packet.

Java client

Listing 5-22: Example of querying data packets in Java

```
List<String> queriedIDs = client.queryData("data.metadataReference:" +
↳ metadataReference);
```

CPP client

Listing 5-23: Example of querying data packets in CPP

```
std::stringstream query;
query << "data.metadataReference:" << metadataReference;
std::vector<std::string> queriedIDs = client->queryData(query.str());
```



Python client

Listing 5-24: Example of querying data packets in Python

```
queried_ids = client.queryData("data.metadataReference:%s" % (metadata_reference,  
↪))
```

If the search query contains Elasticsearch reserved characters such as + - = && || > < ! () { } [] ^ " ~ * ? : \ / . , they have to be escaped e.g: \. Additional resources on Elasticsearch queries: <https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-query-string-query.html>

Additional Query Examples

Querrying data for a specific data point in a specified time interval (please note the quotes around the data point URI). When querying for subscribed URIs the URI protocol should be set to cii.olddb. When querying for data stored via Archiver API, the client might have used a different protocol:

```
data.uri:"cii.olddb:///trk/ctrl/my_data_point" AND data.timestamp:(>=↪  
↪1593786692807000 AND <= 1593786892807000)
```

Searching for data point with an exact data point value

```
data.uri:"cii.olddb:///trk/ctrl/my_data_point" AND data.value:0.0
```

Searching for all data points with a quality flag set to BAD in specified time interval

```
data.quality:"BAD" AND data.timestamp:(>= 1593786692807000 AND <=↪  
↪1593786892807000)
```

Searching for all data points with a quality flag set to BAD or OK and data point value >= 1.0

```
(data.quality:"BAD" OR data.quality:"OK") AND data.value:>=1.0
```

Searching for all data points uris that start with a specific path:

```
data.uri:"cii.olddb:///trk/ctrl/*"
```

The limitation in this case is that one cannot use wildcard in the middle of the path (such as: cii.olddb:///trk/c* as Elasticsearch splits the term into subpaths by splitting on the '/' character).



Downloading data packets

To download data packets from the Engineering Archive we call the `downloadData` method of the `TelemetryArchiveApi` interface. It takes a list of strings representing the IDs of the data packets and returns a list of `DataPacket` objects. In the provided example, we download the data packets specified by the IDs that we acquired above.

There exists a separate API intended for downloading big data. The examples for big data can be found in section 5.3.

Java client

The following sample code demonstrates how to download data packets. The code is an excerpt from the code file named `TelemetryArchiveAPIExample.java` stored in the `java/sample-app` folder.

Listing 5-25: Example of downloading data packets in Java

```
List<DataPacket> downloadedDataPackets = client.downloadData(queriedIDs);  
downloadedDataPackets.forEach(TelemetryArchiveAPIExample::displayDataPacket);
```

The static method `TelemetryArchiveAPIExample.displayDataPackets` takes care of printing the downloaded data packets to the standard output.

Listing 5-26: Example of displaying data packet data in Java

```
public static void displayDataPacket(DataPacket dataPacket) {  
  
    Map<String, Object> map = new HashMap<>();  
  
    map.put("archived", dataPacket.getArchived());  
    map.put("bigDataLocation", dataPacket.getBigDataLocation());  
    map.put("data", dataPacket.getData());  
  
    JSONObject json = new JSONObject(map);  
    System.out.println("\t" + json);  
}
```

CPP client

In a similar manner to the example using Java, the packets can be downloaded in CPP as follows:

Listing 5-27: Example of downloading data packets in CPP

```
std::vector<std::shared_ptr<::elt::telemetry::archive::DataPacket>>  
↳downloadedDataPackets = client->downloadData(queriedIDs);  
for (auto &packet : downloadedDataPackets) {  
    displayDataPacket(packet);  
}
```



The data packets can be displayed using the following function:

Listing 5-28: Example of displaying data packet data in CPP

```
void displayDataPacket (std::shared_ptr<elt::telemetry::archive::DataPacket> &
↳dataPacket) {
    std::cout << "... downloaded data packet: " << std::endl;
    std::cout << "    archived: " << dataPacket->getArchived() << std::endl;
    std::cout << "    bigDataLocation: " << dataPacket->getBigDataLocation() <<
↳std::endl;
    std::cout << "    data: " << dataPacket->getData() << std::endl;
}
```

Python client

In a similar manner to the example using Java, the packets can be downloaded in Python as follows:

Listing 5-29: Example of downloading data packets in Python

```
downloaded_data_packets = client.downloadData(queried_ids)
_ = [display_data_packet(data_packet) for data_packet in downloaded_data_packets]
```

The data packets can be displayed using the following function:

Listing 5-30: Example of displaying data packet data in Python

```
def display_data_packet (data_packet):
    print(' '*10, 'Data packet:')
    print(' '*15, 'archived: ', data_packet.getArchived())
    print(' '*15, 'bigDataLocation: ', data_packet.getBigDataLocation())
    print(' '*15, 'data: ', data_packet.getData())
    print()
```

Data packet deserialization

The following examples are presenting one way of deserializing data packet data field that contains the serialized OLDB data point.

Deserializing in Java

First, we have to define the deserialization structure (getters and setters are omitted for clarity):

```
static class ExampleDataPoint {
    private URI uri;
    private String metadataReference;
    private String quality;
```

(continues on next page)



(continued from previous page)

```
private double value;  
private double timestamp;  
}
```

Once we have the structure defined, we can deserialize the data packet:

Listing 5-31: Example of deserializing data packet containing data point in Java

```
ObjectMapper mapper = new ObjectMapper();
```

```
double archived = packet.getArchived();
```

```
String bigDataLocation = packet.getBigDataLocation();
```

```
ExampleDataPoint dataPoint = mapper.readValue(packet.getData(), ExampleDataPoint.  
↪class)
```

Deserializing in CPP

First, we have to define the deserialization structure:

```
// structure that was stored in the data field  
struct ExampleDataPoint {  
    std::string uri;  
    std::string metadataReference;  
    std::string quality;  
    double value;  
    double timestamp;  
  
    // Serialize itself to json string  
    std::string serialize() const {  
        boost::property_tree::ptree tree;  
        tree.put("uri", uri);  
        tree.put("metadataReference", metadataReference);  
        tree.put("quality", quality);  
        tree.put("value", value);  
        tree.put("timestamp", timestamp);  
        std::stringstream data;  
        boost::property_tree::json_parser::write_json(data, tree);  
        return data.str();  
    }  
  
    // Method to construct ExampleDataPoint from json string  
    static ExampleDataPoint deserialize(const std::string &json) {  
        boost::property_tree::ptree tree;  
        // Parse json and store its contents in the property tree
```

(continues on next page)



(continued from previous page)

```
std::istream stream(json);
boost::property_tree::json_parser::read_json(stream, tree);
// Move the values from the tree into new ExampleDataPoint instance
ExampleDataPoint newInstance;
newInstance.uri = tree.get<std::string>("uri");
newInstance.metadataReference = tree.get<std::string>("metadataReference");
newInstance.quality = tree.get<std::string>("quality");
newInstance.value = tree.get("value", 0.0);
newInstance.timestamp = tree.get("timestamp", 0.0);
return newInstance;
}
}
```

Once we have the deserialization structure defined, we can deserialize the data field:

Listing 5-32: Example of deserializing data packet containing data point in CPP

```
// deserialize the downloaded data packet
double archived = packet->getArchived();
std::string bigDataLocation = packet->getBigDataLocation();
ExampleDataPoint dataPoint = ExampleDataPoint::deserialize(packet->getData());
```

Deserializing in Python

Listing 5-33: Example of deserializing data packet containing data point in Python

```
archived = data_packet.getArchived()
bigDataLocation = data_packet.getBigDataLocation()
data = json.loads(data_packet.getData())
for key, value in data.items():
    print(key, '=', value)
```

Closing the MAL connection

Once we are finished storing, querying, and downloading data packets, we should close the MAL and CII Config client connection by calling the close method.



Closing the connection in Java

The close method needs to be called on both on the TelemetryArchiveAPISync object client as well as the CiiFactory object factory.

```
client.close();  
factory.close();
```

Additionally, both objects are also auto-closeable, so they can be used in try-with-resources statements if that is practical for your application.

Closing the connection in CPP

In CPP the connections can be closed as follows:

```
client->close();
```

Closing the connection in Python

In Python the connections can be closed as follows:

```
client.close()
```

11.5.3 Archive API: Big Data Storage

Normally the Engineering Archive is agnostic of the underlying databases used for storage. However, to have higher performance, the Hadoop storage must be accessed directly when working with big data. The provided example uses an agnostic Big Data interface, but a specific implementation of that interface for Hadoop.

An example of an application that uses Big Data Storage library can be found in the telemetry-examples project. The examples project shows the usage of the Big Data Storage in the:

- *HadoopStorageExample.java* file in the java/sample-app directory for Java,
- *app.cpp* file in the cpp/storage directory for CPP, and the
- *cii-telemetry-examples-storage-py.py* file in the python/storage directory for Python.

The examples can also be found in Appendix D.

A short description of the part of the Archive API used for writing, reading and deleting big data can be found in Appendix A.2.



Imports/Includes

Usage of the HadoopStorage library requires importing the Hadoop Storage class that implements the methods of the BigDataStorage interface.

Java

```
import java.net.URI;  
import elt.storage.HadoopStorage;  
import elt.storage.BigDataStorage;
```

CPP

```
#include <ciiHadoopStorage.hpp>  
#include <ciiException.hpp>
```

Python

```
import elt.storage
```

Creating the client instance

First, we need to create an instance of the Hadoop Storage class by providing the URI of the Hadoop database of the Engineering Archive.

Java example

The following listing shows how to create an instance of the Big Data Storage in Java. The sample code is an excerpt from the code file named TelemetryArchiveAPIExample.java stored in the java/sample-app folder.

Listing 5-34: Example of Big Data Storage client creation in Java

```
final BigDataStorage hadoopStorage =  
    new HadoopStorage(new URI("http://ciihdfshost:9870"), "/esoLs");  
// contains path to the archive  
final URI uri = new URI("cii.config://*/archive");
```



CPP example

The following listing shows how to create an instance of the Big Data Storage in CPP. The sample code is an excerpt from the code file named `app.cpp` stored in the `cpp/storage` folder. The code follows the same flow as the Java code.

Listing 5-35: Example of Big Data Storage client creation in CPP

```
const elt::mal::Uri hadoopLocation("http://ciihdfshost:9870");  
const std::string hadoopEndpoint("/esoSample");  
elt::storage::CiiHadoopStorage storage(hadoopLocation, hadoopEndpoint);
```

Python example

The following listing shows how to create an instance of the Big Data Storage in Python. The sample code is an excerpt from the code file named `cii-telemetry-examples-storage-py.py` stored in the `python/storage` folder. The code follows the same flow as the Java code.

Listing 5-36: Example of Big Data Storage client creation in Python

```
hadoop_location = 'http://ciihdfshost:9870'  
hadoop_endpoint = 'esoSample'  
storage = elt.storage.CiiHadoopStorage(hadoop_location, hadoop_endpoint)
```

Writing data

To write data to the big data database of the Engineering Archive we call the `write` method of the `HadoopStorage` class. It takes the big data archive URI and the data as parameters. Once the big data file is stored, the `write` method returns the automatically generated UUID file name.

If one would like to store the big data file reference into the Engineering Archive, one has to create the data packet (see section 5.2.3), set created big data file reference into the `bigDataLocation` data packet field and store the data packet into the Engineering Archive as described in section 5.2.3.

Java example

The following listing shows how to write data into the Big Data Storage in Java. The sample code is an excerpt from the code file named `TelemetryArchiveAPIExample.java` stored in the `java/sample-app` folder.

Listing 5-37: Example of writing data to the Big Data Storage in Java

```
byte[] data = "data".getBytes();  
String filename = hadoopStorage.write(uri, data);
```



CPP example

The following listing shows how to write data into the Big Data Storage in CPP. The sample code is an excerpt from the code file named `app.cpp` stored in the `cpp/storage` folder. The code follows the same flow as the Java code.

Listing 5-38: Example of writing data to the Big Data Storage in CPP

```
const std::string remoteDirectory = "archive";  
const std::string remoteFilename = "file.bin";  
const std::string contentToStore(generateContent());  
std::stringstream stream(contentToStore);  
storage.write(remoteDirectory, remoteFilename, stream);
```

Python example

The following listing shows how to write data into the Big Data Storage in Python. The sample code is an excerpt from the code file named `cii-telemetry-examples-storage-py.py` stored in the `python/storage` folder. The code follows the same flow as the Java code.

Listing 5-39: Example of writing data to the Big Data Storage in Python

```
remote_directory = 'archive'  
remote_filename = 'file.bin'  
content_to_store = generate_content()  
stream = io.BytesIO(content_to_store.encode('utf-8'))  
storage.write_from(stream, remote_directory, remote_filename)
```

Finding UUIDs of Archived Big Data Files

The big data file could be stored in the big data database either through OLDB-TA subscription mechanism or through the client code (see section 5.3.3). At some point after storing the data in the big data database, one might want to read that file for further analysis. In order to do that a UUID of the big data file must be known.

There are two ways to get the UUID:

- When you store the big data file via Big Data Storage library, an UUID of the file is returned which you can use for further actions.
- When you are trying to retrieve a historic big data file (data that was archived some time ago), you have to find it by querying the Engineering Archive. Big data files are usually associated with other data points values/metadata, so you have to form your query according to the data point you are interested in; see section 5.2.5 on how to query the data. From the queried data you have to extract the big data file reference (UUID that you are looking for) and use that UUID for downloading or deleting big data file.



Note on extracting: if you are using the Archiving API for querying, the API will return you all the data packets relevant to your query. To get the big data file UUIDs you have to iterate over received data packets and call `getBigDataLocation()` method.

Reading data

To read the data from the big data database of the Engineering Archive we call the `read` method of the `HadoopStorage` class. It takes the big data archive URI and the UUID file name as parameters, and returns data encoded into a sequence of bytes. In the provided example we read the data that we have previously written. If one would like to find the big data file UUID that was archived in the past, see section 5.3.4 on how to do that.

Listing 5-40, Listing 5-41, and Listing 5-42 show examples of reading data from the Big Data Storage in Java, CPP, and Python respectively.

Java example

Listing 5-40: Example of reading data from the Big Data Storage in Java

```
byte[] readDataBytes = hadoopStorage.read(uri, filename);
```

CPP example

Listing 5-41: Example of reading data from the Big Data Storage in CPP

```
std::ostream ostream;  
storage.read(remoteDirectory, remoteFilename, ostream);
```

Python example

Listing 5-42: Example of reading data from the Big Data Storage in Python

```
ostream = io.BytesIO()  
storage.read_into(ostream, remote_directory, remote_filename)
```



Deleting files

To delete a file from the big data database of the Engineering Archive we call the delete method of the HadoopStorage class. It takes the big data archive URI and the UUID file name as parameters. In the provided example we delete the file containing the data that we have previously written and read. If one would like to find the big data file UUID that was archived in the past, see section 5.3.4 on how to do that.

Listing 5-43, Listing 5-44 and Listing 5-45 show examples of deleting data from the Big Data Storage in Java, CPP and Python respectively.

Java example

Listing 5-43: Example of deleting files in the Big Data Storage in Java

```
hadoopStorage.delete(uri, filename);
```

CPP example

Listing 5-44: Example of deleting files in the Big Data Storage in CPP

```
storage.remove(remoteDirectory, remoteFilename);
```

Python example

Listing 5-45: Example of deleting files in the Big Data Storage in Python

```
storage.remove(remote_directory, remote_filename)
```

11.5.4 Telemetry Archiver Management API

After changing configurations on the Configuration service, the Service Management API must be used to refresh the configuration on the Telemetry Archiver.

A short description of the part of the Telemetry Archiver Management API can be found in A.3.

An example of an application that uses the Telemetry Archiver Management API to refresh a configuration can be found in the *TelemetryServiceManagementAPIExample.java* file in the telemetry-examples project.



Imports/includes

Usage of the Telemetry Archiver Management API requires the following imports to establish a MAL connection: CiiFactory, QoS, ReplyTime, and ZpbMal. The TelemetryServiceManagementSync interface defines Telemetry Archiver Management methods.

Java

```
import java.net.URI;  
import elt.mal.CiiFactory;  
import elt.mal.rr.qos.QoS;  
import elt.mal.rr.qos.ReplyTime;  
import elt.mal.zpb.ZpbMal;  
import elt.telemetry.management.TelemetryServiceManagementSync;
```

CPP

```
#include <mal/Cii.hpp>  
#include <mal/utility/LoadMal.hpp>  
#include <mal/rr/qos/ReplyTime.hpp>  
#include <CiiTelemetryManagement.hpp>
```

Python

```
import elt.pymal as mal  
from ModCiiTelemetryManagement.Elt.Telemetry.Management.  
↳TelemetryServiceManagement import TelemetryServiceManagementSync
```



Connection to Telemetry Archiver using MAL and closing the connection

The connection can be created in the same way as described in section 5.2.2. The MAL connection needs to be closed as shown in section 5.2.8.

Refreshing configurations

To refresh selected configurations, we call the `refreshConfigurations` method of the `TelemetryServiceManagement` interface. It takes a list of strings specifying the URIs of the configurations that we want to refresh as parameter. The URIs of the relevant data capture configurations could be obtained:

- Via configuration GUI, for more information see section 4.2.
- When the data capture configuration is saved in the config-service: in order to save the data capture configuration in the config-service, the location/URI where the data capture configuration will be stored must be specified. In order to inform the Telemetry Archiver of the new/updated configuration one has to call `refreshConfigurations` method with the same URI.
- Programmatically by querying the config-service for the relevant data capture configuration. An example of programmatic querying in Java for all data captures stored under Telemetry Archiver's service range could be seen below. For more info on how to query the config-service in other languages, see [1].

```
CiiConfigClient configClient = CiiConfigClient.getInstance();
String archiverServiceRange = ElasticsearchUtilities.createElasticSearchId(
    URI.create("cii.config://*/telemetry/service1/datacapture/"));
List<String> dataCaptureConfigIds =
    configClient.remoteIndexSearch("configuration_instance",
        String.format("id:%s*", archiverServiceRange));

// dataCaptureConfigIds could now be as an input for refreshConfigurations
// method
```

In order to refresh all the configurations, we call the `refreshAllConfigurations` method of the `TelemetryServiceManagement` interface.

Listing 5-46, Listing 5-47 and Listing 5-48 show examples refreshing configurations in Java, CPP and Python respectively. Note that the example code assumes that a configuration with the URL `"cii.config://telemetry/service1/48fe831c-436a-401b-ab27-18ee401735af"` already exists.



Java example

Listing 5-46: Example of refreshing configurations in Java

```
// refresh selected configurations
List<String> ids = new ArrayList<>();
ids.add("cii.config://*/telemetry/service1/48fe831c-436a-401b-ab27-18ee401735af
↵");
client.refreshConfigurations(ids);

// refresh all configurations
client.refreshAllConfigurations();
```

CPP example

Listing 5-47: Example of refreshing configurations in CPP

```
// Refresh selected configurations
std::vector<std::string> dataCaptureIds
{"cii.config://*/telemetry/service1/48fe831c-436a-401b-ab27-18ee401735af"};
client->refreshConfigurations(dataCaptureIds);
// Refresh all configurations
client->refreshAllConfigurations();
```

Python example

Listing 5-48: Example of refreshing configurations in Python

```
# Refresh selected configurations
ids = ['cii.config://*/telemetry/service1/48fe831c-436a-401b-ab27-18ee401735af']
client.refreshConfigurations(ids)

# Refresh all configurations
client.refreshAllConfigurations()
```




11.6 Telemetry CLI tools

This section describes how to use Telemetry Archiver CLI tools. These command line tools simplify basic operations without needing to write custom applications for these tasks.

11.6.1 CLI Manager

CLI Manager allows a user to send the Telemetry Archiver a notification to trigger a refresh of its configurations.

The default syntax for the CLI Manager is:

```
telemetry-manager refresh [-c <config> <config>] [-h] [-s <service endpoint>]
```

The -h argument prints a help description to the standard output.

The refresh command options are:

<code>-c, --configurations <config> <config></code>	specifies list of configuration ids, based on which telemetry service configurations are refreshed. If this option is omitted, all configurations are refreshed
<code>-h, --help</code>	displays this help message
<code>-s, --service <service endpoint></code>	specifies custom service endpoint of telemetry archiver for manager to connect to.

Usage example

```
$ telemetry-manager refresh -c config_1 config_2
```

Here only specific configurations are refreshed (configurations with ids `config_1` and `config_2`).

```
$ telemetry-manager refresh
```

Here all existing configurations are refreshed.

11.6.2 CLI Archiver

CLI Archiver is a tool enabling users to store, download, and query data packets from and to the Engineering Archive. The CLI Archiver is not optimized for speed and is meant to be used for one-time slow archive related tasks. A single instance of the telemetry-tool doesn't support full required rate of telemetry service.

The telemetry-tool provides the user with three different commands: store, download, and query. All commands have specific options available, further specifying the behavior of the command. The help



and service options can be used with all commands. The help option is used for display the help message. The service option enables the user to specify a custom address of the Telemetry Archiver.

The default syntax for CLI archiver is:

```
telemetry-tool [command] [-h | specific options] | [-h] [command]
```

The -h argument prints a help description to the standard output.

Each command has a list of options that can be used with it. All options can be viewed with the -h option.

Storing data packets

Storing of data packages is possible by using the store command. Additionally, by using different provided options the user can store different files (data packets containing metadata and/or big binary files), as well as specifying different service endpoints for the archiver to connect to (service endpoint option).

The Store command syntax is as follows:

```
telemetry-tool store [-b <file path> <file path>] -d <file path> <file path> | -  
↪h [-s <service endpoint>]
```

The Store command options are:

```
options:  
-b,--bigdata <file path> <file path> specifies file path to  
bigData binary file to be stored  
-d,--data <file path> <file path> specifies file path that  
contains data for creation of new  
dataPacket  
-h,--help displays help message  
-s,--service <service endpoint> specifies custom service endpoint  
of telemetry archiver to connect to.
```

Multiple data packets can be created and stored with a single store command. If multiple file paths are provided as parameters to the option. The bigdata option is optional, but can also specify multiple file paths. Additionally, specified big data file paths correspond to their respective data packets provided with the data option (this means, that not all data packets necessarily have to contain corresponding big data files attached to them and that both types of packets are paired based on the order they were passed as parameters: 1 data packet goes with 1 big data file, etc.).

The data file is meant to be any kind of JSON based file up to 2 GB in size. While the data file does not have to follow the data packet structure (any valid json will do), it is highly recommended that it follows it (the data packet structure), otherwise analyzing and operating on such highly dynamic data will be extremely problematic.

The bigdata file could be anything (text file, image, video, compiled program) and doesn't have a file size limit.



Storing metadata

The example below shows how to store the data point metadata into the Engineering Archive:

```
$ telemetry-tool store -d metaFile1.json  
mcx7jHQBRBKWEuv9m9Gc
```

The output from the telemetry tool represents the stored metadata data packet ID (mcx7jHQBRBKWEuv9m9Gc), which the user could use later on to download the archived metadata or reference the metadata when archiving data point values (see section 6.2.1.2). The metaFile1.json is following the structure of the CII Config metadata and it is displayed below:

Listing 6-1: an example of metaFile1.json file

```
{  
  "@type": "MdOldbNumber",  
  "value": 2.0,  
  "comment": "",  
  "checked": false,  
  "metadataInstanceVersion": 1,  
  "@name": "telemetry_meta1",  
  "@genType": "SINGLE",  
  "metadataInstance": "telemetry_meta1"  
}
```

Storing data point values

The data point values are stored via the following command:

```
$telemetry-tool store -d data/dpValue1.json data/dpValue2.json -b BigPicture.jpg  
j_HWn3ABUTxPCroFrXwo n8yTjHQBRBKWEuv9ltGT
```

The dataPacket1.json and dataPacket2.json are files that should follow the structure below (the structure is not enforced and the end user can archive any valid JSON):

Listing 6-2: structure of the dpValue1.json file

```
{  
  "uri": "cii.client:///my_datapoint_id_1",  
  "metadataReference": "mcx7jHQBRBKWEuv9m9Gc",  
  "quality": "OK",  
  "value": 0.5,  
  "timestamp": 1600084745793000  
}
```

Before storing the data point value (e.g., dpValue1.json) into the Engineering Archive, make sure that the metadata in the metadataReference field is already present in the Engineering Archive. The metadata reference can be obtained by either searching in the Engineering Archive (see section 6.2.3)



or storing the metadata first (see section 6.2.1.1) and then storing the data packet with the reference of the archived metadata.

If the `-b` (`--bigdata`) flag is present, the referenced big data file is stored to the Big Data Storage and its ID is stored to corresponding data packet's `bigDataLocation` field. Big data files are paired with data packets in order they are provided. In the provided example above, the `dpValue1.json` data packet has a reference to big data file attached (reference to archived `BigPicture.jpg`), while the `dpValue2.json` data packet does not have a reference to big data file.

The telemetry tool returns the IDs of the archived data point data packets, which are used for downloading the archived data packets from the Engineering Archive (see section 6.2.2).

Downloading data packets

Data packets and their corresponding big data files can be retrieved from the Engineering Archive using the download command. Data packets can be either downloaded to local disk or displayed based on options provided.

The Download command syntax is as follows:

```
telemetry-tool download -h | -p <packet id> <packet id> [-o <file path>] [-s <service endpoint>]
```

The Download command options are:

```
options:
-h, --help                displays help message
-o, --output <folder path> specifies the path to folder where
                           data packets will be stored. If
                           this option is ommited, downloaded
                           packets data is displayed instead.
-p, --packets <packet id> <packet id> specifies list of packet IDs that
                           will be downloaded
-s, --service <service endpoint> specifies custom service endpoint
                           of telemetry archiver to connect
                           to.
```

Multiple packets can be displayed/downloaded using a single download command. The output option (`-o,--output`) controls if the data packets will be displayed or downloaded. This option is not required, so by default, the specified data packets will only be displayed to the user. If this option is present, packets will be stored inside the folder denoted by this option's parameter, specifying the location on the local disk where the download folder will be created. If the data packet that should be downloaded contains the location (URI denoting the location in the EA) of the big data file, that file is also retrieved and stored in the same folder.

Below two examples show the usage of both modes (display and download) for the download command.

Usage example



Displaying data packet downloaded from the Engineering Archive:

```
$ telemetry-tool download -p j_HWn3ABUTxPCroFrXwo  
{ "archived":1.583229611294E15,  
  "data":{"key":"testValue1"},  
  "bigDataLocation":"cii.config:\\\\*\\archive\\0a383f2f-8b37-4e11-b6c5-  
↪1d798e278da0" }
```

Here the specified data packet is retrieved from the EA and its content is displayed. It is not stored on the local disk.

Usage example

Downloading a data packet from Engineering Archive to local disk:

```
$ telemetry-tool download -p j_HWn3ABUTxPCroFrXwo -o ~/downloadFolder
```

Here a new folder named downloadFolder is created inside the user's home directory. A new file is created inside containing the information stored in the data field of the downloaded data packet. If the packet has a big data file location defined, another binary file will be created. The data packet ID also specifies the name of the newly created file, where the big data binary file additionally has .bin appended to the end.

Querying data packets

The user can search for specific data packets using the query command. The ES Query DSL [3] is used for querying, allowing the user to find packets matching certain criteria specified by the query request (e.g. all packets having certain field set to specified value). The result of the query command is a list of IDs representing data packets that match the sent query.

The Query command syntax is as follows:

```
telemetry-tool query -h | -r <request definition> [-s <service endpoint>]
```

The Query command options are:

```
options:  
-h, --help                displays help message  
-r, --request <request definition> Elasticsearch query based on which  
                             the returned packet IDs are selected  
-s, --service <service endpoint> specifies a custom service endpoint of  
                             the telemetry archiver to connect to.
```

When specifying options, either request option (-r,--request) or help options (-h,--help) must be provided. The example below displays the usage of the query command.

Usage example



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 462 of 589

```
$ telemetry-tool query -r "data.metadataInstance: telemetry* AND data.value: 2"  
j_HWn3ABUTxPCroFrXwo jvHWn3ABUTxPCroFrXwo
```

Here query requests data packets where the *data.value* field equals 2 and *data.metadataInstance* matches all *telemetry** entries. As a result, IDs of two data packets (j_HWn3ABUTxPCroFrXwo and jvHWn3ABUTxPCroFrXwo in the above example), that match the set query are returned.



11.7 Telemetry Subscription

This section describes how to correctly prepare a Telemetry Archiver (TA) for archiving data point changes that are published via OLDB subscription. Telemetry Archiver subscribes to the data point changes depending on the data capture configurations which are stored in the config-service.

Each data capture configuration represents a set of rules that determine whether or not a specific data point change should be captured. It is possible to define multiple data capture configurations for one data point in order to make a more fine-grained archiving behaviour.

Each Telemetry Archiver has its own specified range that defines the location in the Configuration service from which the data capture configurations are being downloaded from. An example on how this works is shown below:

Say we have a Telemetry Archiver (TA1) and we would like it to archive data point (DP1) value changes. In order to do that we have to:

1. Create a data capture configuration (CiiDataCaptureConfig) with desired archiving rules set (mentioned below in the examples section).
2. Store data capture configuration into config-service under Telemetry Archiver's service range.
3. Reload the data capture configurations of TA1 by calling one of the refreshConfiguration methods on the TA1's Service Management API that will seamlessly reload the in-memory capture configurations with the new ones coming from the config-service.

The data capture configuration fields are documented in Appendix B.

The data capture configurations could be saved/updated with both the config-client and config-gui. See provided manuals [1] on how to effectively use config-client or config-gui. An example written in the Java programming language is shown below (similarly one can do the same via C++ or Python):

Listing 7-1: Example of creating and saving a data capture configuration

```
final URI serviceRange = URI.create("cii.config;/**/telemetry/service1");
final URI oldbDatapointUri = URI.create("cii.oldb:///datapoint_1");
DataCaptureConfiguration captureConfig = new DataCaptureConfiguration(
    oldbDatapointUri, // oldb datapoint uri
    1.0f,             // delta value change
    "ABSOLUTE",      // delta type set to absolute changes
    "VALUE",          // archive mode set to value changes
    "00:01",          // minimum interval (at most once per second)
    "10:00",          // maximum interval (at least once every 10 minutes)
    true              // archiving enabled
);
// data capture configuration id should be unique
captureConfig.setId(UUID.randomUUID().toString());

// save data capture configuration to config-service
try {
```

(continues on next page)



(continued from previous page)

```
CiiConfigClient client = CiiConfigClient.getInstance();
client.saveTargetConfig(serviceRange, captureConfig);
} catch (CiiSerializableException e) {
    System.out.println("Error while storing data capture configuration: "
        + e.getCiiMessage());
}
```

11.7.1 Data Capture Configuration Examples

We assume the Telemetry Archiver (TA1) which will subscribe to data point (DP1) changes with a service range set to `cii.config://*/telemetry/archiver1`:

Archive data point on absolute value change or quality change

In this example we would like to archive data point (DP1) in the following situation:

- Data point should be archived on absolute value change (+/- 2.0)
- Data point should be archived at least once every 5 minutes
- Data point should be archived at most every second
- Data point should be archived if the quality flag has changed

We have to create 2 capture configurations:

1. Data capture (DC1) that will define the rules for archiving absolute value change. The data capture configuration should have the following fields set:

Field	Value
dataPointAddress	cii.olddb:///datapoints/dp1
archiveMode	VALUE
deltaType	ABSOLUTE
deltaValue	2.0
minimumInterval	1
maximumInterval	05:00
archivingEnabled	true

2. Data capture (DC2) that will define the rules for archiving quality change:



Field	Value
data-PointAddress	cii.olddb:///datapoints/dp1
archive-Mode	QUALITY
deltaType	""
deltaValue	0
minimumInterval	0 (minimum interval when capturing quality changes is not used)
maximumInterval	0 (maximum interval for quality change can be set as well, but in this example this is already handled in the DC1 config example)
archivingEnabled	true

Archive data point on relative value change or metadata change

We would like to archive the data point based on the following rules:

- Data point should be archived if the value changes by 4% of the current value
- Data point should be archived if the metadata has changed

Again, we create two data capture configurations:

1. Data capture configuration (DC1) that will define the rules for relative value change. The process is the same as in section 7.1.1, except the delta type field should be set to "RELATIVE".

Field	Value
dataPointAddress	cii.olddb:///datapoints/dp1
archiveMode	VALUE
deltaType	RELATIVE
deltaValue	0.04
minimumInterval	0 (all updates will be stored)
maximumInterval	0 (maximum interval is not set)
archivingEnabled	true

2. Data capture configuration (DC2) that will define the rule for metadata change. The data capture configuration that handles metadata changes does not respect the set minimum and maximum interval and will archive the data point on every metadata change. It is expected that metadata will be rarely changed.



Field	Value
dataPointAddress	cii.olddb:///datapoints/dp1
archiveMode	METADATA
deltaType	"" (irrelevant for metadata change)
deltaValue	0 (irrelevant for metadata change)
minimumInterval	0 (minimum interval is not respected; metadata change is always archived)
maximumInterval	0 (maximum interval is not set)
archivingEnabled	true

Stop archiving the data point

There are two options for TA1 to stop archiving changes of the specified data point (DP1):

1. Set the archivingEnabled field to false for all data capture configurations that are referencing a DP1 (via dataPointAddress field) and call one of the refreshConfiguration methods on the TA1's Service Management API.
2. Delete the data capture configuration referencing a DP1 (via dataPointAddress field) from the configuration service and call one of the refreshConfiguration methods on the TA1's Service Management API.

11.7.2 OLDB Simulation

For testing purposes of the Telemetry Archiver, it might be useful to simulate the data point value updates. This could be done by:

- Using OLDB mock.
- Using OLDB simulator.



OLDB Mock

For fine grained testing purposes, it might be useful to use your own OLDB mock service instead of the production OLDB service. Internally the Telemetry Archiver is using OLDB client, which subscribes to the OLDB service for data point value updates. Consult the 3.3.1 OLDB Client Configuration section of CII Online Database user manual [6], on how to configure the OLDB client to connect to your own OLDB mock service which you are going to use as part of the testing purposes.

OLDB Simulator

For performance testing purposes it might be useful to use the OLDB simulator which allows you to continuously simulate data point value updates depending on the data capture configurations that are stored in the CII Configuration Service. The OLDB data point value simulation works in the following fashion:

- The user starts the Telemetry Archiver with the TA's simulateCapture configuration field set to true.
- The OLDB simulator that is started within the Telemetry Archiver downloads the data capture configurations from the CII Config Service from the serviceRange specified in the TA's configuration.
- The data point info is parsed from the downloaded data capture configurations. The URI from the data capture configuration defines the URI of the data point for which the updates will be simulated.
- The OLDB simulator starts generating data point updates based on parsed data point info.

The OLDB simulator is currently supporting 2 archive modes from the data capture configuration:

- Value change
- Quality change

If the data capture configuration archiveMode is set to VALUE_CHANGE, the ABSOLUTE value of the data point will be continuously generated with the minimum interval based the data capture configuration deltaChange field (delta change represents the value change necessary for data point to be archived). This ensures that every generated data point value is suitable for archiving. The data point RELATIVE value updates are currently not supported by the simulator.

If the data capture configuration archiveMode is set to QUALITY_CHANGE, the quality flag of the data point will be flipped (from true to false and vice versa) with the minimum interval period as defined in the data capture configuration.

The URIs in the data capture configurations could be anything as long as they follow the OLDB URI convention (cii.olddb:///uri), however the OLDB simulator supports a special behavior in case the data point URI is one of the following:



ELT CII - User Manual

Datapoint URI	Behavior
cii.olddb:///simulated_datapoint_1	generate data point updates based on the data capture configuration, without any side effects. This is the regular behaviour of the OLDB simulator.
cii.olddb:///simulated_datapoint_2	generate data point updates based on the data capture configuration, but omit the data point timestamp field (timestamp field of the generated data point is null).
cii.olddb:///simulated_datapoint_3	generate data point updates based on the data capture configuration, but omit the data point quality flag field (quality flag of the generated data point is null).



11.8 Telemetry Archiver Deployment

Telemetry Archivers (TA) are designed to be horizontally scalable. In order to achieve this, they do not know about one another and the end user has to manually specify which TA will be archiving which data points under which data capture rules.

This data point archiving specification is done through the Telemetry Archiver's configuration (see Appendix C). Every Telemetry Archiver has to have its own configuration file, because the Telemetry Archiver configurations have to differ in at least two configuration fields (both fields have to be unique when running two services on the same machine):

- **telemetryServiceURI**: location that defines the network interface and port on which the TA will be bound to (e.g. `zpb.rr://0.0.0.0:9115/`).
- **serviceRange**: location in configuration service where the specific Archiver's data capture configurations are located (e.g. `cii.config/*telemetry1/archiver1/`)

Each Telemetry Archiver contains its configuration stored in the CII Configuration Service. The configuration file itself could be deployed in either local or remote configuration service database. The location of the configuration file is determined by either command line flag or environment variable (`TELEMETRY_CONFIG`). The variable should be following the CII Configuration Service URI rules (`cii.config/location/path/to/config`).

The order of configuration file location is the following:

1. Command line flag (`-DTELEMETRY_CONFIG`)
2. Environment variable (`TELEMETRY_CONFIG`)
3. If none of the above is set, default configuration will be used (`cii.config/*telemetry/config/defaultArchiverConfig`)

If the configuration is stored on remote Configuration service, a `CONFIG_SERVICE_LOCATION` variable has to be set as well (e.g. `zpb.rr://hostname:port/`).

The order of the Configuration service location flag is the following:

1. Command line flag (`-DCONFIG_SERVICE_LOCATION`)
2. Environment variable (`CONFIG_SERVICE_LOCATION`)
3. If none of the above is set, default Configuration service location will be used (`zpb.rr://ciiconfservicehost:9116/`)

If one would like to run multiple TAs on the same server (or virtual machine), one has to use command line flag (`-DTELEMETRY_CONFIG`) as otherwise two services will try to bind on the same port and only one will succeed.

Make sure that the `serviceRange` field in the TAs configuration is unique, otherwise different archivers will subscribe to the same data points with the same capturing rules and you will end up with duplicated data in the archive.



11.8.1 Example

Say we have our custom TA configuration located in the CII configuration service under the `cii.config://*/telemetry/config/custom` location.

- Start TA with `-DTELEMETRY_CONFIG="cii.config://*/telemetry/config/custom/"` command line flag.
- Store data capture configurations in configuration service under location that is specified in the `serviceRange` (see also section 4). If `serviceRange` is specified as `cii.config//telemetry/archiver1`, the data capture configuration should be stored under `cii.config//telemetry/archiver1/datacapture/<data capture UUID>`

11.8.2 Systemd Deployment

If you are running multiple services via `systemd`, you have to deploy multiple `systemd` unit files that will start the Telemetry Archiver executables with different `-DTELEMETRY_CONFIG` flag set.

11.8.3 Scaling Telemetry Archiver

The actual load that one Telemetry Archiver can handle entirely depends on the hardware on which the TA is running and the desired performance needs (number of subscribed data points and their update frequency). The TA was designed in a way to support horizontal scaling, meaning if more throughput is necessary you can easily deploy more archivers.

For maximum performance of the TA it is recommended to deploy it on a different machine than the one where CII OLDB Service or Engineering Archive are deployed on. Nevertheless, if your desired performance needs are lower than one server can handle then you are free to run an entire infrastructure on one server.

The only way to get the actual performance numbers of one TA node is to run performance tests on the production hardware and measurements. This is usually done by slowly increasing the load of the TA and observing the metrics (CPU, RAM, JVM metrics and TA metrics through JMX). The following are some of the general problems and solutions that are likely to occur during performance measurements:

- If archiving buffer starts filling up in the TA (check JMX metrics), that means you have to increase the performance of the Engineering Archive. Before you start tinkering with the Engineering Archive, you might want to try and increase or decrease the archiving buffer batch size through the TA's configuration. It's possible that decreasing or increasing the batch size might improve the archiving throughput. During the performance tests that we performed, the Engineering Archive was always a bottleneck.
- If CPU is the bottleneck and fully utilized add more cores or share the load by moving some of the subscribed data points to another TA node. For maximum performance make sure the TA does not share CPU resources with another virtual machine that is running on the same server.



- If RAM is the bottleneck, add more RAM or share the load with another TA. A rule of thumb is to size the heap in a way that after a garbage collection, cca. 30% of the heap is still occupied. You can try to modify the heap size via JVM flags (-Xmx, -Xms). For maximum performance ensure that the system is not swapping (on high performance servers swap is usually disabled).
- If the network interface can't handle the incoming data, then you have to act accordingly to your infrastructure capabilities.
- Running the TA on a more recent JVM (11) might improve performance by 10-15%, due to the improvements made in G1GC garbage collector. If you notice long stop the world GC pauses, you can try tuning the GC via JVM flags. If long GC pauses are proven as problematic, you might want to try switching to one of the newer concurrent garbage collectors (Shenandoah, ZGC) that promise garbage collector pauses lower than 10ms.

To get a better feeling of what is really going on with the system, one might want to use the CLI tools that are usually installed on a Linux machine:

- vmstat - reports information about processes, memory, paging, block IO, traps, disks and CPU activity.
- iostat - reports Central Processing Unit (CPU) statistics and input/output statistics for devices, partitions and network filesystems (NFS).
- nicstat - prints out network statistics for all network cards (NICs), including packets, kilobytes per second, average packet sizes and more.

11.8.4 Logging

For logging purposes, the Telemetry Archivers are using CiiLogManager which is configured through logging configuration file. For more information about logging and possible configuration options, see the CII Log user manual [7].



11.9 Advanced Topics

11.9.1 Telemetry Archiver statistics

Telemetry Archiver exposes basic statistic using the Java MBeans JMX interface. Statistics can be obtained with the usage of jconsole. The user should start the jconsole with the same user account that started Telemetry Archiver and running the following command:

```
jconsole&
```

Telemetry Archiver can be chosen from the list of processes and connected to by clicking on the Connect button (Figure 9-1).

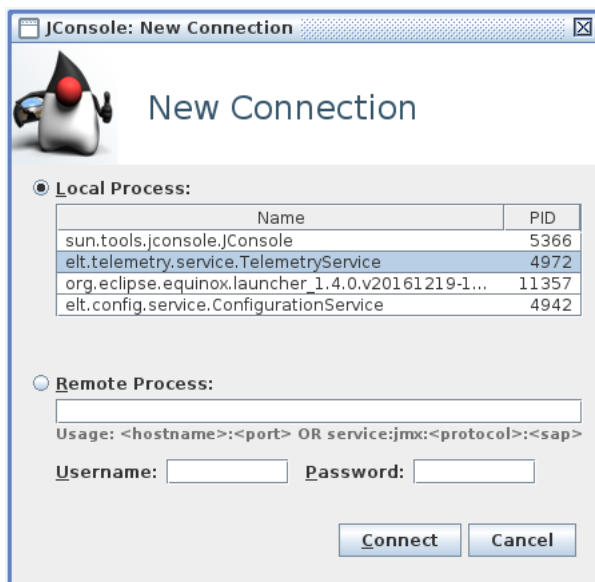


Figure 9-1: Connect to Telemetry Archiver with Jconsole

After the connection is made, the console with statistics is displayed. The console displays heap usage, CPU usage, number of threads, and similar java statistics. The specific Calculation service statistics are exposed under the elt.telemetry.service.stat namespace.

The statistics show multiple metrics:

- current time,
- service up-time,
- number of working threads,
- EA status,



- number of subscribed data points,
- number of dropped data points,
- buffer fullness,
- number of archived data points in last hour,
- number of archiving errors,
- number of archived data points in a day,
- number of transferred data packets,
- number of data packets not inserted,
- number of failed data point subscriptions,
- number of invalid data packets.

To obtain statistics, the user should click on the MBeans tab and browse the `elt.telemetry.service.stat` namespace (Figure 9-2).

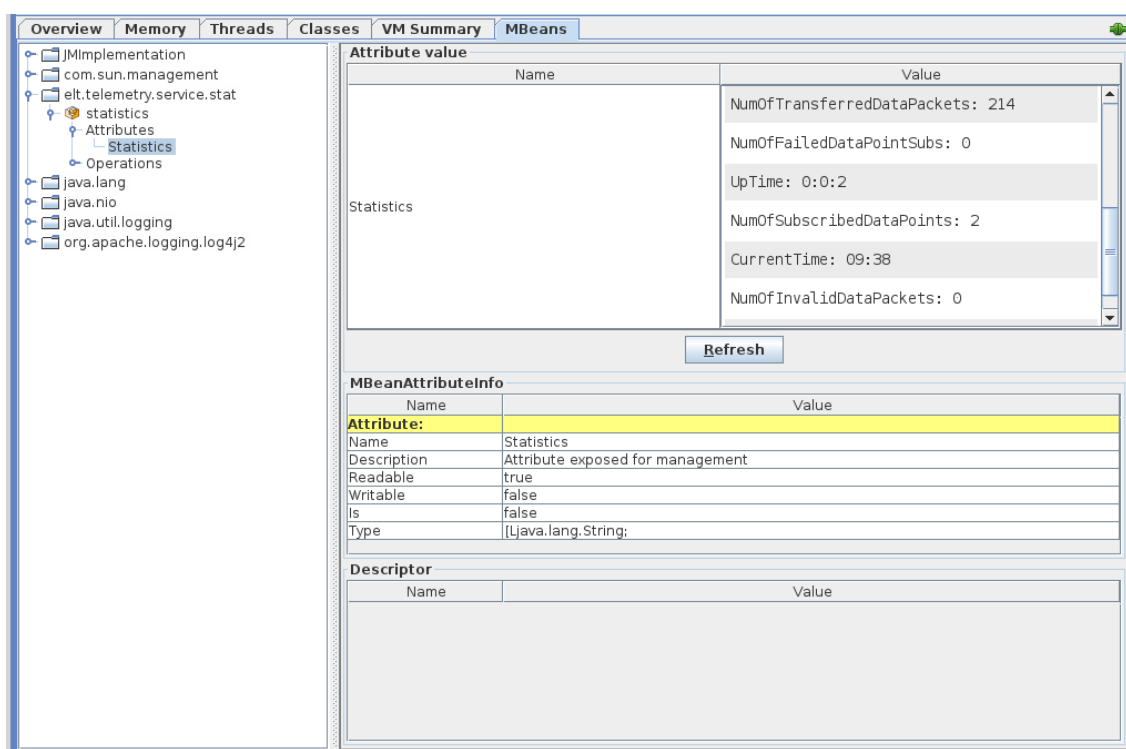


Figure 9-2: Telemetry Archiver Statistics



11.10 Archive API

11.10.1 Archive API: storing, querying and downloading data

Telemetry Archive API provides users with the functionality to store, query and download data that is saved in the Engineering Archive. The methods, specified in the TelemetryArchiveApi interface are implemented in the TelemetryArchiveApiSyncImpl class.

Table 9.1: Archive API

Return Type	Method and Description
String	<code>storeData(DataPacket dataPacket)</code> Stores a data packet (containing a timestamp, location of the data in the big data database and data in JSON format) to the Engineering Archive. Returns a string specifying the assigned ID of the data packet.
List<String>	<code>storeDataList(List<DataPacket> dataPackets)</code> Stores a list of data packets to the Engineering Archive. Returns a list of strings specifying the assigned IDs of the data packets.
List<String>	<code>queryData(String queryString)</code> Queries data packets in the Engineering Archive based on the provided query string (ES Query DSL should be used). Returns a list of strings specifying the IDs of the corresponding data packets.
List<DataPacket>	<code>downloadData(List<String> dataPacketIDs)</code> Downloads data packets with provided IDs from the Engineering Archive.

11.10.2 Archive API: Big Data Storage

The Hadoop Storage library provides users with the functionality of writing data to and reading data from the big data database of the Engineering Archive. The HadoopStorage class implements the methods of the BigDataStorage interface.

Table 9.2: Big Data Storage API



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 475 of 589

Re- turn Type	Method and Description
String	<code>write(Uri uri, byte[] data)</code> Writes the data to the file on the location, specified by the big data archive URI and the automatically generated UUID file name. Returns the file name.
String	<code>write(Uri uri, InputStream stream)</code> Writes the data from the stream to the file on the specified location and automatically generates UUID file name. Returns the file name.
void	<code>write(Uri uri, int version, String filename, byte[] data)</code> Writes the data to the file on the location, specified by the big data archive URI, version and file name.
void	<code>write(Uri uri, String filename, InputStream stream)</code> Writes the data to the file on the location, specified by the big data archive URI and file name.
void	<code>write(Uri uri, int version, String filename, InputStream stream)</code> Writes the data to the file on the location, specified by the big data archive URI, version and file name.
byte[]	<code>read(Uri uri)</code> Reads data from the file on the location specified by the URI.
byte[]	<code>read(Uri uri, String filename)</code> Reads data from the file on the location, specified by the big data archive URI and file name.
byte[]	<code>read(Uri uri, int version, String filename)</code> Reads data from the file on the location, specified by the big data archive URI, version and file name.
byte[]	<code>read(Uri uri, OutputStream stream)</code> Reads data from the file on the location specified by the URI. The data is written to the output stream.
void	<code>read(Uri uri, String filename, OutputStream stream)</code> Reads data from the file on the location specified by the URI and file name. The data is written to the output stream.
void	<code>read(Uri uri, int version, String filename, OutputStream stream)</code> Read data from the file on the location specified by the URI, the version and file name. The data is written to the output stream.
void	<code>delete(Uri uri, String filename)</code> Deletes file on the location, specified by the big data archive URI and file name.
void	<code>delete(Uri uri, int version, String filename)</code> Deletes file on the location, specified by the big data archive URI, version and file name.



11.10.3 Telemetry Archiver Management API

The Telemetry Archiver Management API provides users with the functionality to refresh configurations. The methods, specified in the TelemetryServiceManagement interface are implemented in the TelemetryServiceManagementSynclImpl class.

Table 9.3: Service Management API

Return Type	Method and Description
void	refreshConfigurations(List<String > ids) Refreshes configurations, specified by their IDs. The IDs are the UUIDs of the data capture configuration that can be obtained either through querying the config-service or GUI, for more info see section 5.4.3
void	refreshAllConfigurations() Refreshes all configurations.



11.11 Data Capture Configuration Options

The data capture configurations contain the following fields, which could be tweaked in order to get the desired archiving behaviour:

Table 9.4: Data capture configuration options

Field	Description	Example
data-Point-Address	Address of data point in OLDB	cii.olddb:///datapoints/dp1/
archiveMode	Defines on which on-change event the data point is archived (data point value change, quality flag change, time based – e.g. once per day)	VALUE_CHANGE or QUALITY_CHANGE or META-DATA_CHANGE
deltaValue	Delta value change required to archive (in the case of data value change event)	+/- 1.0
deltaType	Type that the delta value field represent (absolute change or relative change) If RELATIVE delta type is chosen, the delta value represents the percentage of the data point value for which the data point value should change in order to be archived. The value 1 represents 100% change and value 0.01 represents 1% change. If ABSOLUTE delta type is chosen, the delta value represents the actual value for which the data point value should have changed in order to be archived.	ABSOLUTE or RELATIVE
minimumInterval	Minimum time between specific data point archiving in the <i>hh:mm:ss</i> format. If the data point value has changed multiple times in a short interval, it ensures we don't store too many values. If minimum interval is set to 0, every data point value change will be archived according to the delta field.	00:00:10
maximumInterval	Maximum interval between two data point archiving events in the <i>hh:mm:ss</i> format. If the data point value has not changed in the specified maximum interval, the data point value will still be archived. Setting it to 0 will disable it.	01:15:00 (data point should be archived at least on every 1 hour and 15 minutes)
archiveEnabled	Enable archiving (True) or disable archiving (False) for this specific data capture configuration	True/False



11.12 Telemetry Archiver Configuration

Listing 9-1 Telemetry Archiver configuration YAML

```
engineeringArchiveBackupURI: "http://ciiarchivehost:9200/"  
engineeringArchiveURI: "http://ciiarchivehost:9200/"  
largeStorageServiceURI: "http://ciihdfshost:9870/"  
serviceRange:  
  - "cii.config://remote/telemetry/serviceRange2"
```

Each Telemetry Archiver contains its configuration stored in the CII Configuration Service.

Table 9.5 contains a description of all configuration fields. If Telemetry Archiver needs to use custom configuration, a java parameter `DTELEMETRY_CONFIG_INI` or environment variable `TELEMETRY_CONFIG_INI` must be specified with the location to the custom configuration YAML. An example of configuration YAML is shown in Listing 9-1.

Table 9.5: Telemetry Archiver Configuration settings



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
 Doc. Version: 1
 Released on: -
 Page: 479 of 589

Field	Description	Example
telemetry-ServiceURI	Defines an IP address and port on which the Telemetry Archiver is bound to	zpb.rr://0.0.0.0:9115 / Default is zpb.rr://0.0.0.0:9115 /
simulateArchive	If set to true, all data archive operations will be simulated. If set to false, data will be archived to Engineering archive	True/False Default is False
simulate-Capture	If set to true, all data capture is simulated (using simulator instead of real OLDB)	True/False Default is False
archive-Quality-Mask	Permits selective archiving of a data points based on data points quality flag Possible values: UNKOWN, INVALID, VALID, SIMULATED, ALL, NONE	UNKOWN, VALID Default is ALL
engineeringArchiveURI	Engineering Archive location	Default is https://ciiarchive:9200/
engineeringArchive-BackupURI	Engineering Archive backup URI	Default is https://ciiarchive2:9200/
maxSearch-Size	The maximum number of data packets stored in the EA that can be retrieved at once	5000 Default is 10000
authenticationToken	Authentication token that allows service to interact with the archiving system	Default is empty string
serviceRange (required)	Defines the configuration service URI location from where the data capture configuration will be loaded. Multiple service ranges are allowed here	[cii.config://*/telemetry/service1]
archiving-Buffer-Size	Maximum number of data packets that the archiving buffer can contain	50000 Default is 100000
bufferBatch-Size	Maximum number of data packets that one archiving batch can contain	150 Default is 400
maximum-BatchWait	Maximum number of seconds that the archiving buffer will wait for data before sending a batch to the Engineering Archive (even if the batch is not full).	5 Default is 1



11.13 Big Data Storage Examples

The following listings show examples of the usage of the Big Data Storage for the Java, CPP and Python programming languages.

Listing 9-2: Example usage of the Big Data Storage in Java

```
package sampleapp;

import java.io.IOException;
import java.io.ByteArrayInputStream;
import java.io.ByteArrayOutputStream;
import java.net.URI;
import java.net.URISyntaxException;
import elt.storage.HadoopStorage;
import elt.storage.BigDataStorage;

/**
 * Hadoop Storage example.
 */
public class HadoopStorageExample {

    /**
     * Private method to generate content to be written to the remote file
     */
    private static String generateContent() {
        String s = "AAAA";
        for (int i = 0; i < 100; i++) {
            switch(i % 5) {
                case 0:
                    s += "CCCC";
                    break;
                case 1:
                    s = s + s;
                    break;
                case 2:
                    s += "XXXXXXXXXX";
                    break;
                case 3:
                    s += "!!!!!!";
                    break;
                case 4:
                    s += "<<<>>>";
                    break;
            }
        }
        return s;
    }
}
```

(continues on next page)



(continued from previous page)

```
* HadoopStorage main method.
*
* @param args Command line arguments (none are required).
*/
public static void main(String[] args) throws URISyntaxException,
↳IOException, InterruptedException {

    int status = 0;
    final BigDataStorage hadoopStorage =
        new HadoopStorage(new URI("http://ciihdfshost:9870"), "/esoSample
↳");

    // contains path to the archive (remote directory)
    final URI uri = new URI("cii.config://*/archive");
    // remote filename
    final String remoteFilename = "file1.bin";
    // content to store
    final String contentToStore = generateContent();

    try {
        // save data to Hadoop, use binary stream
        ByteArrayInputStream stream = new ByteArrayInputStream(contentToStore.
↳getBytes());
        System.out.println("Writing blob to hadoop storage...");
        hadoopStorage.write(uri, remoteFilename, stream);

        // read data from Hadoop
        ByteArrayOutputStream ostream = new ByteArrayOutputStream();
        System.out.println("Reading blob from hadoop storage...");
        hadoopStorage.read(uri, remoteFilename, ostream);
        ostream.close();

        if (contentToStore.equals(ostream.toString())) {
            System.out.println("OK, saved and read content are the same");
        } else {
            System.err.println("FAIL: data from remote not same as source");
            status = 1;
        }

        // delete a file from Hadoop
        hadoopStorage.delete(uri, remoteFilename);
    } catch (Throwable th) {
        System.err.println(String.format("Exception during remote operation: %s",
↳ th));
        status = 5;
    }
    try {
        // try to read from non existent file
        System.out.println("Access to non existing remote file (should fail)...
↳");
    }
```

(continues on next page)



(continued from previous page)

```
        byte[] data = hadoopStorage.read(uri, remoteFilename);
    } catch (Throwable th) {
        System.err.println(String.format("REMOTE READ FAILURE (expected): %s",
↪th));
    }
    System.exit(status);
}
}
```

Listing 9-3: Example usage of the Big Data Storage in CPP

```
#include <iostream>
#include <sstream>
#include <string>
#include <cstdlib>
#include <ciiHadoopStorage.hpp>
#include <ciiException.hpp>

/**
 * Generates sample string
 * @return string, generated content
 */
static std::string generateContent() {
    std::string s = "AAAA";
    for (std::size_t i = 0; i < 100; i++) {
        switch (i % 5) {
            case 0:
                s += "BBBB";
                break;
            case 1:
                s = s + s;
                break;
            case 2:
                s += "XXXXXXXXXX";
                break;
            case 3:
                s += ".....";
                break;
            case 4:
                s += "<<<>>>";
                break;
        }
    }
    return s;
}

/**
 * main method
```

(continues on next page)



(continued from previous page)

```
*/  
int main(int argc, char **argv) {  
    int status = 0;  
    const elt::mal::Uri hadoopLocation("http://ciihdfshost:9870");  
    const std::string hadoopEndpoint("/esoSample");  
    const std::string remoteDirectory = "archive";  
    const std::string remoteFilename = "file.bin";  
    const std::string contentToStore(generateContent());  
  
    // Initialize hadoop storage,  
    // First parameter is Uri of the webhdfs server  
    // Second parameter is name of the top directory under which all files for this  
    // instance of hadoop storage will be stored.  
    elt::storage::CiiHadoopStorage storage(hadoopLocation, hadoopEndpoint);  
  
    // In this example, stream is used  
    std::stringstream stream(contentToStore);  
    // Create receiving stream  
    std::ostreamstream ostream;  
  
    try {  
        // Write content of the stream to the remote file  
        std::cout << "Writing blob to hadoop storage..." << std::endl;  
        storage.write(remoteDirectory, remoteFilename, stream);  
  
        // Read content of the stream from the remote file  
        std::cout << "Reading blob from hadoop storage..." << std::endl;  
        storage.read(remoteDirectory, remoteFilename, ostream);  
  
        // Remove remote file  
        storage.remove(remoteDirectory, remoteFilename);  
  
        if (ostream.str() == contentToStore) {  
            std::cout << "OK, saved and read content are the same" << std::endl;  
        } else {  
            std::cerr << "FAIL: data from remote not same as source" << std::endl;  
            status = 1;  
        }  
    } catch (const elt::error::CiiException &e) {  
        std::cerr << "Cii Exception during remote operation: " << e.what() << "  
↪std::endl;  
        status = 2;  
    } catch (const std::exception &e) {  
        std::cerr << "Std Exception during remote operation: " << e.what() << "  
↪std::endl;  
        status = 5;  
    }  
  
    // Read from non existing remote file
```

(continues on next page)



(continued from previous page)

```
try {
    std::cout << "Access to non existing remote file (should fail)..." <<
↪std::endl;
    storage.read(remoteDirectory, remoteFilename, ostream);
} catch (const std::exception &e) {
    std::cerr << "REMOTE READ FAILURE (expected): " << e.what() << std::endl;
}
return status;
}
```

Listing 9-4: Example usage of the Big Data Storage in Python

```
#!/usr/bin/env python

#pylint: disable=E1101,C0103,R0914,C0330,W0612
#This script demonstrates usage of HadoopStorage

import sys
import io
import traceback
import elt.storage

def generate_content():
    """Helper function that generates content string"""
    s = 'AAAA'
    for i in range(100):
        modulo = i % 5
        if modulo == 0:
            s += 'bbbb'
        elif modulo == 1:
            s += s
        elif modulo == 2:
            s += 'XXXXXXXXXX'
        elif modulo == 3:
            s += ',,,,,'
        else:
            s += '<<<>>>'
    return s

# Define URI for remote webhdfs server to be used
hadoop_location = 'http://ciihdfshost:9870'
# Define top level remote directory
hadoop_endpoint = 'esoSample'
# Define remote directory
remote_directory = 'archive'
# Define remote filename
remote_filename = 'file.bin'
```

(continues on next page)



(continued from previous page)

```
def _main():
    """main method implementation"""
    status = 0
    # Initialize hadoop storage.
    # First parameter is string containing Uri of the webhdfs server
    # Second parameter is name of the top directory under which all files
    # for this instance of hadoop storage will be stored.
    storage = elt.storage.CiiHadoopStorage(hadoop_location, hadoop_endpoint)

    # In this example, stream is used
    content_to_store = generate_content()
    # Content to write must be raw binary values. Use io.BytesIO!
    stream = io.BytesIO(content_to_store.encode('utf-8'))
    try:
        print('Writing blob to hadoop storage...')
        # Write content of the stream to the remote file
        storage.write_from(stream, remote_directory, remote_filename)

        # Read content of the remote file into the stream
        ostream = io.BytesIO()
        print('Reading blob from hadoop storage...')
        storage.read_into(ostream, remote_directory, remote_filename)

        # Delete remote file
        storage.remove(remote_directory, remote_filename)

        if content_to_store == ostream.getvalue().decode('utf-8'):
            print("OK, saved and read content are the same")
        else:
            print("FAIL: data from remote not same as source", file=sys.stderr)
            status = 1
    # pylint: disable=W0703
    except Exception as e:
        traceback.print_tb(sys.exc_info()[2])
        print("Exception during remote operation: ", e, file=sys.stderr)
        status = 5

    # try to read from non existing remote file
    try:
        print('Access to non existing remote file (should fail)...')
        ostream = io.BytesIO()
        storage.read_into(ostream, remote_directory, remote_filename)
    # pylint: disable=W0703
    except Exception as e:
        print("REMOTE READ FAILURE (expected): ", e, file=sys.stderr)
    return status

def main():
    """main method wrapper"""
```

(continues on next page)



(continued from previous page)

```
result = 0
try:
    result = _main()
    #pylint: disable=W0703
except Exception as e:
    print(e)
    traceback.print_tb(sys.exc_info()[2])
    result = 5
return result

if __name__ == '__main__':
    sys.exit(main())
```

SERVICE MANAGEMENT

Revision:	1.0
Status:	New
Repository:	https://gitlab.eso.org/cii/info/cii-docs
Project:	ELT CII
Document ID:	
Owner:	Marcus Schilling
Last modification:	October 4, 2021

Document History

Revision	Date	Changed/reviewed	Section(s)	Modification
1.0	04.10.2021	mschilli	All	Created.

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is a user manual for the ELT Core Integration Infrastructure Software project.

Audience

This document is aimed at those ESO employees and ESO partners involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 488 of 589

12.1 Overview

This document is a user manual for CII Services management tooling.



12.2 Introduction

This document describes the tools available for the management of CII Services

- `cii-postinstall` - for setting up CII services on a host
 - `cii-services` - for starting / stopping / monitoring of the CII Services
-



12.3 cii-postinstall

You run the CII-PostInstall on a freshly installed DevEnv machine, or after an upgrade of the DevEnv. You choose a role for the host, and the utility will apply the corresponding configuration steps. For most configuration steps, root privileges are needed.

Usage:

```
$ /elt/ciisrv/postinstall/cii-postinstall <role>
```

Run the command without arguments to see its help.

12.3.1 Roles

The role “ownserver” configures the services for local development on a single machine.:

```
$ /elt/ciisrv/postinstall/cii-postinstall role_ownserver
```

The role “groupserver” does a few more things than the one above. It runs the services with more resources and makes the services/databases accessible from other hosts.:

```
$ /elt/ciisrv/postinstall/cii-postinstall role_groupserver
```

The role “groupclient” configures the host to use the services provided by a groupserver. For example, assuming host “elthlccd20” has been post-installed with the role “groupserver”, then, to let applications on your host use the CII Services on that groupserver, you do:

```
$ /elt/ciisrv/postinstall/cii-postinstall role_groupclient elthlccd20
```

Switching Roles?

It is possible to run the utility again later, but your mileage will vary: some role-switches will work, others not.

12.3.2 Subcommands

A “role” is a collection of configuration steps, and those configuration steps can also be applied one-by-one, by executing them as a “subcommand”. To see the available subcommands, just pass “help”:

```
$ cii-postinstall help
CII PostInstall (20210730)
Available subcommands: role_ownserver  role_groupserver  role_groupclient  ↵
↵ ECII397 logsink  filebeat  logstash  kibana  logtrail  serverlocation  elastic ↵
↵ schemas  unrestrict  oldb  schemas_local  help
```



12.3.3 Other Functions

In addition, there are some configuration steps that are outside of any role. You can run them as normal user (no root privileges) to configure your personal environment.

Example: Set up a Local-DB so you can use the “Config from Local Files” feature of the CII Config Service.:

```
$ cii-postinstall schemas_local /home/eltdev/INTROOT/localdb  
[...]
```

Tip If you create the Local-DB at "\$INTROOT/localdb" or "\$HOME/localdb", the ↪
↪Config Service will find it there.
If you choose another location, also define this variable: export CONFIG_LOCAL_
↪DB=/tmp/delphi/localdb



12.4 cii-services

The cii-services utility lets you start/stop/monitor the CII services. For some operations, it requires root-privileges -> unless you are running it with root-privileges, it will show a password prompt when needed.

Usage:

```
$ cii-services
Available subcommands: info status logs start stop
```

12.4.1 Monitoring/Status Check

info This is a deployment-centric view that tells you whether the services are running and where they are.:

```
$ cii-services info
CII Services Tool (20210706)
Collecting information.....
Collecting information.....

Installations on this host:
ElasticSearch |active:yes |install:/usr/share/elasticsearch/bin/
->elasticsearch
Redis |active:yes |install:/opt/anaconda3/bin/redis-server
MinIO |active:no |install:/usr/local/bin/minio
ConfigService |active:yes |install:/elt/ciisrv/bin/config-service
->|ini:
TelemService |active:no |install:
Filebeat |active:yes |install:/usr/share/filebeat/bin/filebeat
Logstash |active:yes |install:/usr/share/logstash/bin/logstash
Kibana |active:yes |install:/usr/share/kibana/bin/kibana
Jaeger |active:no |install:
AlarmServices |active:no |install:
AlarmPlugin |active:no
AlarmConsumer |active:no
AlarmConverter |active:no
AlarmSupervisor |active:no
ConfigClient |ini:

Addresses of Services:
Config / Local-DB |access:yes |file:/home/eltdev/localdb
Config / ElasticSearch |access:yes |host:ciielastichost (IP:127.0.0.1)
Config / MinIO |access:no |host:ciihdfshost (IP:)
Config / Service |access:yes |host:ciiconfservicehost (IP:127.0.0.1)
->|host2:ciiconfservicehost2 (IP:127.0.0.1)
OLDB / Redis |access:yes |host:localhost (IP:::1)
OLDB / MinIO |access:no |host:ciihdfshost (IP:)
```

(continues on next page)



(continued from previous page)

```
Telem / Service          |access:no          |host:ciiarchivehost (IP:127.0.0.1)

Statistics of Services:
Config / Service         |NrOfRequests: 2    |UpTime: 4 minutes
OLDB / Redis             |total_connections_received:3 |rejected_connections:0
Telemetry / Service
```

status This is a feature-centric view, that basically tells you which features you have available. For example, “Blob Values” means that the Distributed File System MinIO (in previous versions: Hadoop) is available for storage of large values and binaries.:

```
$ cii-services status
CII Services Tool (20210706)
Collecting information.....

Status Summary:
OK Config from Central Server
OK Config from Local Files
-- Config with Blob Values
OK OLDB Normal Mode
-- OLDB with Blob Values
-- Telemetry
```

logs Privileged Operation. Run the command as root.:

```
$ su -c "cii-services logs"
Password:
CII Services Tool (20210706)

=== ConfigService Journal ===
-- Logs begin at Tue 2021-08-31 18:06:15 UTC, end at Tue 2021-09-07 15:32:55 UTC.
↪ --
Sep 07 15:05:28 elthlccd20 config-service[29961]: 2021-09-07 15:05:28,979 main_
↪DEBUG Registering MBean org.apache.logging.log4j2:type=764c12b6,
↪component=StatusLogger
Sep 07 15:05:28 elthlccd20 config-service[29961]: 2021-09-07 15:05:28,980 main_
↪DEBUG Registering MBean org.apache.logging.log4j2:type=764c12b6,
↪component=ContextSelector
Sep 07 15:05:28 elthlccd20 config-service[29961]: 2021-09-07 15:05:28,981 main_
↪DEBUG Registering MBean org.apache.logging.log4j2:type=764c12b6,
↪component=Loggers,name=
Sep 07 15:05:28 elthlccd20 config-service[29961]: 2021-09-07 15:05:28,985 main_
↪DEBUG Registering MBean org.apache.logging.log4j2:type=764c12b6,
↪component=Appenders,name=console
[...]

=== Logstash Journal ===
-- Logs begin at Tue 2021-08-31 18:06:15 UTC, end at Tue 2021-09-07 15:32:55 UTC.
↪ --
```

(continues on next page)



(continued from previous page)

```
Sep 07 15:32:11 elthlccd20 systemd[1]: Started logstash.  
Sep 07 15:32:24 elthlccd20 logstash[33887]: Sending Logstash logs to /var/log/  
↳logstash which is now configured via log4j2.properties  
Sep 07 15:32:25 elthlccd20 logstash[33887]: [2021-09-07T15:32:25,370][INFO_]  
↳][logstash.runner          ] Starting Logstash {"logstash.version"=>"6.8.3"}  
Sep 07 15:32:26 elthlccd20 logstash[33887]: [2021-09-07T15:32:26,183][INFO_]  
↳][logstash.config.source.local.configpathloader] No config files found in path  
↳{:path=>"/etc/logstash/conf.d/*.conf"}  
[...]  
  
=== TelemService Journal =====  
-- Logs begin at Tue 2021-08-31 18:06:15 UTC, end at Tue 2021-09-07 15:32:55 UTC.  
↳ --  
-- No entries --
```

12.4.2 Start / Stop

Privileged Operations. You will be asked for the root password.:

```
$ cii-services start  
CII Services Tool (20210706)  
Possible names: logKibana redis elasticsearch logLogstash configService_]  
↳logFilebeat, or a prefix thereof, or 'all'  
  
$ cii-services start all  
CII Services Tool (20210706)  
About to start logKibana redis elasticsearch logLogstash configService_]  
↳logFilebeat ...  
Password:  
  
$ cii-services stop all  
CII Services Tool (20210706)  
About to stop logKibana redis elasticsearch logLogstash configService_]  
↳logFilebeat ...  
Password:
```



12.5 Example

(base) elthlccd20 eltddev:~ 8 > cii-services info:

```
CII Services Tool (20210706)
Collecting information.....
Collecting information.....

Installations on this host:
ElasticSearch |active:no |install:/usr/share/elasticsearch/bin/
->elasticsearch
Redis |active:no |install:/opt/anaconda3/bin/redis-server
MinIO |active:no |install:/usr/local/bin/minio
ConfigService |active:no |install:/elt/ciisrv/bin/config-service
->|ini:
TelemService |active:no |install:
Filebeat |active:no |install:/usr/share/filebeat/bin/filebeat
Logstash |active:no |install:/usr/share/logstash/bin/logstash
Kibana |active:no |install:/usr/share/kibana/bin/kibana
Jaeger |active:no |install:
AlarmServices |active:no |install:
AlarmPlugin |active:no
AlarmConsumer |active:no
AlarmConverter |active:no
AlarmSupervisor |active:no
ConfigClient |ini:

Addresses of Services:
Config / Local-DB |access:no |file:/home/eltddev/localdb
Config / ElasticSearch |access:no |host:ciielastichost (IP:)
Config / MinIO |access:no |host:ciihdfshost (IP:)
Config / Service |access:no |host:ciiconfservicehost (IP:)
->|host2:ciiconfservicehost2 (IP:)
OLDB / Redis |access:no |host:localhost (IP:::1)
OLDB / MinIO |access:no |host:ciihdfshost (IP:)
Telem / Service |access:no |host:ciiarhivehost (IP:)

Statistics of Services:
Config / Service
OLDB / Redis
Telemetry / Service
```

How to read this output:

This section reports some info about the tool itself:

```
CII Services Tool (20210706)
Collecting information.....
Collecting information.....
```

=> Tool version and progress dots. => If an error occurs while collecting information, we count the



progress dots to spot the problem.

This section reports where applications on your host will find the CII services, and whether the `cii-services` tool was able to access them (= talk to them). The services may be running on your host (see next paragraph) or on another host.:

```
Addresses of Services:
Config / Local-DB      |access:no      |file:/localdb
Config / ElasticSearch |access:no      |host:ciielastichost (IP:)
Config / MinIO         |access:no      |host:ciihdfshost (IP:)
Config / Service       |access:no      |host:ciiconfservicehost (IP:)
OLDB / Redis           |access:no      |host:localhost (IP:::1)
OLDB / MinIO           |access:no      |host:ciihdfshost (IP:)
Telem / Service        |access:no      |host:ciiarhivehost (IP:)
```

=> The empty parentheses “(IP:)” tell us that hostname resolution is not configured. => This means applications running on your host are not able to locate the CII services. => You need to configure your host, see `cii-postinstall`

This section reports which CII services are installed on your host, and whether they are running.:

```
Installations on this host:
ElasticSearch |active:no      |install:/usr/share/elasticsearch/bin/
↳elasticsearch
Redis          |active:no      |install:/opt/anaconda3/bin/redis-server
MinIO          |active:no      |install:/usr/local/bin/minio
ConfigService  |active:no      |install:/elt/ciisrv/bin/config-service
↳|ini:
TelemService  |active:no      |install:
Filebeat       |active:no      |install:/usr/share/filebeat/bin/filebeat
Logstash       |active:no      |install:/usr/share/logstash/bin/logstash
Kibana         |active:no      |install:/usr/share/kibana/bin/kibana
Jaeger         |active:no      |install:
AlarmServices  |active:no      |install:
AlarmPlugin    |active:no      |
AlarmConsumer  |active:no      |
AlarmConverter |active:no      |
AlarmSupervisor|active:no      |
ConfigClient   |ini:           |
```

=> Several services (not: Telemetry and Alarms) are installed on your host. None of them is running on your host. => If you plan to use CII services running on another host, this is fine. => Otherwise, you need to start the services, see above.

This section reports some statistics from the services.:

```
Statistics of Services:
Config / Service
OLDB / Redis
Telemetry / Service
```




=> An empty output like this again indicates that the the services are not accessible, but you knew that already.

Preview: After configuring and starting the services, the output will change to something like this::

```
CII Services Tool (20210706)
Collecting information.....
Collecting information.....

Installations on this host:
ElasticSearch |active:yes |install:/usr/share/elasticsearch/bin/
->elasticsearch
Redis |active:yes |install:/opt/anaconda3/bin/redis-server
MinIO |active:no |install:/usr/local/bin/minio
ConfigService |active:yes |install:/elt/ciisrv/bin/config-service
->|ini:
TelemService |active:no |install:
Filebeat |active:yes |install:/usr/share/filebeat/bin/filebeat
Logstash |active:yes |install:/usr/share/logstash/bin/logstash
Kibana |active:yes |install:/usr/share/kibana/bin/kibana
Jaeger |active:no |install:
AlarmServices |active:no |install:
AlarmPlugin |active:no
AlarmConsumer |active:no
AlarmConverter |active:no
AlarmSupervisor |active:no
ConfigClient |ini:

Addresses of Services:
Config / Local-DB |access:yes |file:/home/eltdev/localdb
Config / ElasticSearch |access:yes |host:ciielastichost (IP:127.0.0.1)
Config / MinIO |access:no |host:ciihdfshost (IP:)
Config / Service |access:yes |host:ciiconfservicehost (IP:127.0.0.1)
->|host2:ciiconfservicehost2 (IP:127.0.0.1)
OLDB / Redis |access:yes |host:localhost (IP:::1)
OLDB / MinIO |access:no |host:ciihdfshost (IP:)
Telem / Service |access:no |host:ciiarhivehost (IP:127.0.0.1)

Statistics of Services:
Config / Service |NrOfRequests: 2 |UpTime: 4 minutes
OLDB / Redis |total_connections_received:3 |rejected_connections:0
Telemetry / Service
```

INTERNAL CONFIG SYSTEM

Revision:	1.8
Status:	Released
Repository:	https://gitlab.eso.org/cii/srv/cii-srv
Project:	ELT CII
Folder:	trunk/deliverables/phase7
Document ID:	
File:	MAN-ELT_Configuration.docx
Owner:	Borut Terpinc
Last modification:	March 17, 2022
Created:	March 16, 2019

Prepared by	Reviewed by	Approved by
Borut Terpinc (CSL SWE) Marko Novak (CSL)	Miha Vitorovič (CSL)	Gregor Čuk (CSL)

Document History



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 499 of 589

Re- vi- sion	Date	Changed/ iewed	Section(s)	Modificatio n
0.1	10.04. 2019	bter- pinc	All	Document creation
0.2	15.04. 2019	manovak	GUI	GUI part added
0.9	26.4. 2019	bter- pinc	All	Update after internal review, release
1.0	3.6. 2019	bter- pinc	6	Package support update
1.1	19.7. 2019	bter- pinc	3	Added sample application section.
1.2	11.9. 2019	bter- pinc/jp ribosek	All	Updated after internal review. Added sample applica- tion workflow and search examples.
1.3	25.9. 2019	bter- pinc	Appendix C	Added sample code produced by the generators.
1.4	09.10. 2019	bter- pinc	Introductio n 4	Added introductio n section Added simple custom metadata sample application
1.5	12.02. 2020	bter- pinc	All	General rewrite of the manual. Sections added and restructure d.
1.6	20.02. 2020	kstam- par	All	Update of cpp file names and added flags when gen- erating python bindings.
1.7	13.07. 2020	bter- pinc	Appendix C Appendix G Appendix H 5.4 2.1	Removed the MdInt64 duplicate. Added client settings Added service settings Added cache example Additional note on class versioning.
1.8	17.03. 2022	mschilli	Doc Title Overview	Module is now only for CII-internal usage

Confidentiality

This document is classified as a confidential document. As such, it or parts thereof must not be made accessible to anyone not listed in the Audience section, neither in electronic nor in any other form.

Scope

This document is manual for the configuration system used by applications in the ELT Core Integration Infrastructure Software.

Audience

This document is aimed at those Cosylab and ESO employees involved with the ELT Core Integration Infrastructure Software project, as well as other Users and Maintainers of the ELT Core Integration Infrastructure Software.



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 500 of 589

References

1. Cosylab, ELT CII Online Database, User Manual, CSL-DOC-19-176283, v 1.1
1. Cosylab, ELT Config, Error Handling transfer document, CSL-DOC-19-172644, Version 1.9
2. JsonPath, <https://github.com/json-path/JsonPath>
3. Elastic search query, <https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-query-string-query.html#query-string-syntax>



13.1 Overview

This document is a manual for the CII-internal configuration system. It is used by CII to manage the metadata of the Online database, the capture configurations of the Telemetry service, and others.

13.2 Introduction

The Core Integration Infrastructure (CII) Internal Configuration system provides the mechanisms for distributing configuration data to the CII services software and to allow the collection of configuration data from these systems.

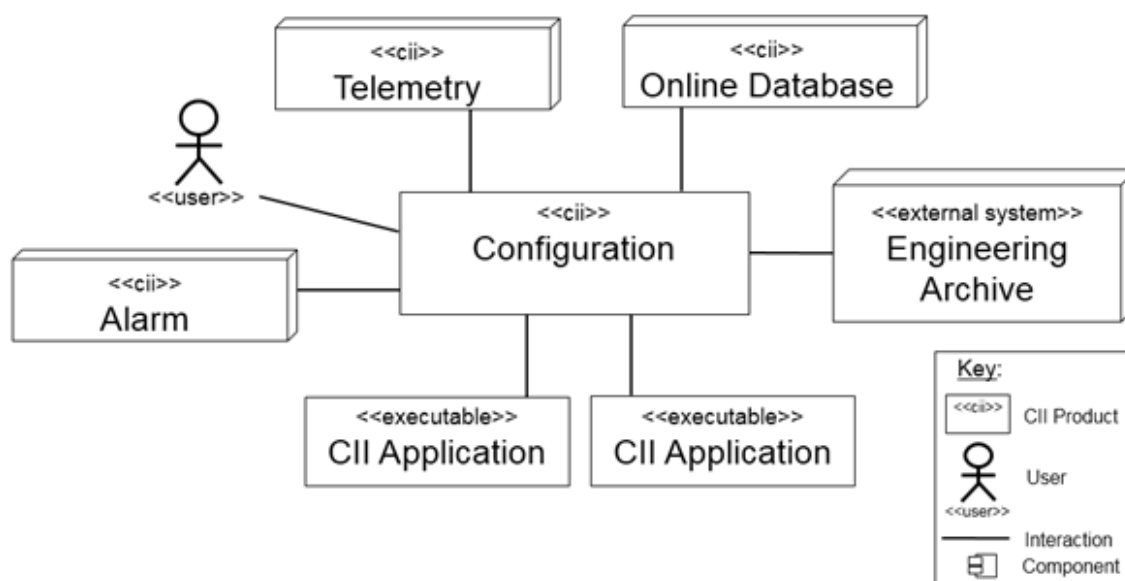


Figure 2-1 Configuration interactions

Configuration data is stored as JSON documents, and supports three different storage modes:

- **Cache:** Data is stored into client API local memory. This mode is not persistent.
- **Local DB:** Data is stored into local data storage on client computers. This mode is persistent.
- **Remote DB:** Data is stored in the Engineering Archive using the config service. As config service together with Engineering Archive storage provides a complete solution to access and store data to remote storage, it is also referred to as “remote DB”. This mode is persistent.

The configuration is a set of **data** and **metadata**. The types and structure of configuration is defined in classes. Metadata classes define the structure of config points and metadata, while configuration classes define the structure of a target configuration (configuration instance). Both configuration and metadata classes support inheritance.

The configuration and metadata classes are based on YAML structures used to generate language-specific classes. The structure of the YAML defines the class names, class member fields, and member types.

Configuration data is composed of metadata instances containing the metadata values and default config values, and target configuration containing config values of the target config points. The configuration instances and metadata instances are JSON data structures. Instances define how values of the actual data-bind to corresponding config and metadata classes.

Configuration instances are instances of the configuration classes which present the actual target

configurations. Therefore “target configuration” and “configuration instance” have the same meaning.

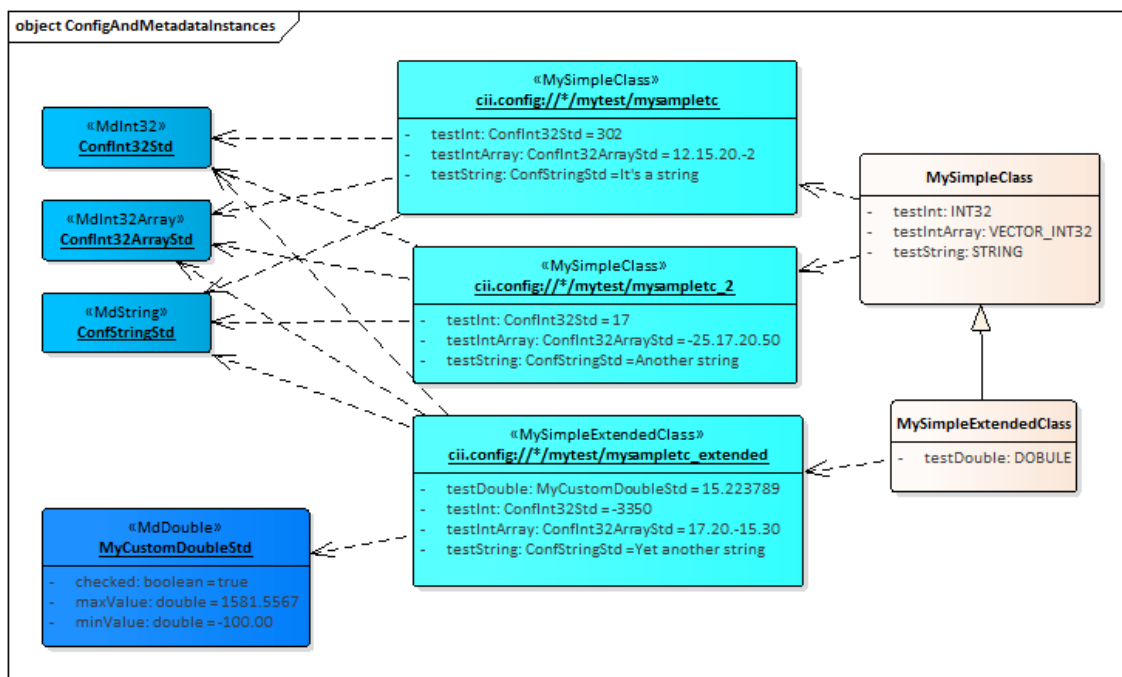


Figure 2-2 Configuration/metadata classes and instances

Figure 2-2 shows the example of relationship between configuration system (configuration and metadata) classes and instances. The figure shows configuration classes in orange, target configurations in light blue and metadata instances in dark blue colour. The configuration classes **MySimpleClass** and **MySimpleExtendedClass** define the member fields (`testInt`, `testIntArray`, `testString`, `testDouble`) of chosen type.

Based on these classes, target configurations (configuration instances) at specified endpoints are created. The following target configurations are created: *cii.config://*/mytest/mysampletc**, *cii.config://*/mytest/mysampletc_2**, *cii.config://*/mytest/mysampletc_extended**. Target configurations hold the configuration values for the fields defined in the configuration classes.

In this example, every target configuration except **MyCustomDoubleStd** uses default metadata instances to set the metadata of the specific field. The default metadata instances don't hold any particular metadata information and exist for the ease of usage. The **MyCustomDoubleStd** is a custom metadata instance that defines the minimum and maximum value limits. The instance also sets the checked flag to true. This means that the set limits will be used when the value is set to `testDouble` member field. In case that the value is of the limits, **CiiConfigLimitOutOfBoundsException** will be thrown.

More detailed description of the system can be found in the sections below.



13.2.1 Classes

Classes only present the structure of the configuration and thus cannot contain any values. Classes, in general, have the following properties:

- Classes present the structure of the configuration.
- Classes don't contain any values.
- Classes definitions are generated into selected programming language class bindings.
- Changing the class structure forces recompilation of the client code.
- Classes cannot be versioned, but they can be updated and deleted.
- Classes are generated from remote storage or a corresponding local YAML file.

Classes cannot be versioned, but they support inheritance. In the standard workflow, classes cannot be deleted or updated. For defining a new class structure, a new derived class must be created. The new derived class must inherit from the root (configuration or metadata) class which is already a part of CII configuration. This root class has no members, so it doesn't define any class structure. Members must be added to the derived class.

Note on class versioning: As classes present the structure and are a base for the class bindings, the versioning of the classes would present additional complexity to the configuration. Every change of the class version would mean the recompilation of the code. A part of a good design principle is to have fixed structures defined. Versioning of classes would break this concept, as every version could completely change the structure. As the instances are bound to the structure, this would also mean completely changing the instances under the existing TC URI. As the versioning of the classes is not supported, the structure can be extended by using the inheritance. Inheritance maintains the base structure, but adds the ability to add additional attributes to the configuration.

Deleting and updating the classes can only be used in the process of development and testing.

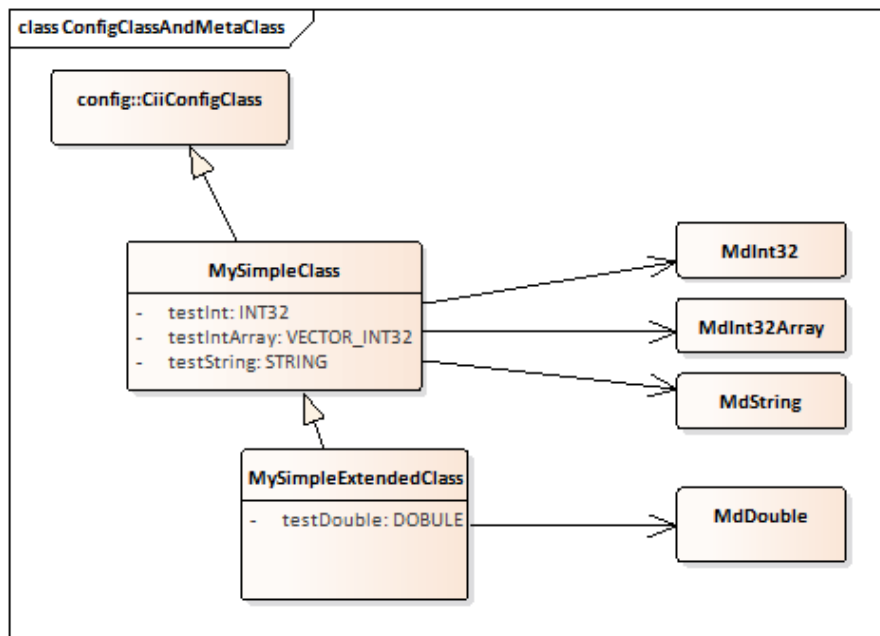


Figure 2-3 Configuration and metadata classes relationship

Figure 2-3 shows relationship diagram of config and metadata classes on basic example. The class MySimpleClass derives from the config base class (CiiConfigClass), which is the root class for all the Cii configuration classes. The MySimpleClass contains 3 member fields, testInt, testIntArray and testString. Field types are defined by metadata classes (MdInt32, MdInt32Array, MdString). For the list of the supported types refer to 2.1.2 and Appendix C.

Additionally, MySimpleExtendedClass derives from the MySimpleClass and adds additional field testDouble which is of type MdDouble. All other fields are inherited from the base class (MySimpleClass).

Configuration classes

Configuration classes are structures which define the type and member fields for generated specific language classes. They have the following properties:

- Configuration classes present nodes in the config structure.
- Configuration class member fields present the configuration data points – config points.
- Config points can be metadata classes or references to another configuration class.
- Configuration classes cannot contain complex types or list of complex types, the only exception is a reference to another class.
- All configuration classes derive from basic class (CiiConfigClass).
- All configuration classes must be stored in the elt.config.classes package.



Metadata classes

Metadata classes are structures which are used to generate language specific classes that contain metadata values. Metadata classes contain custom metadata fields and a special value field, which holds the actual value of the config point. The value field is added to the metadata classes for simplification and more straightforward usage of Java API. The actual value field is not serialized in metadata instances, but in configuration instances, as the value is not directly part of metadata. Usage of the value field in the metadata classes simplifies typing of metadata and config points, as the types for value, default value and metadata are all contained in one class.

Config points (configuration class member fields) are in most cases metadata classes. The naming convention for this classes follows the rule: Md + chosen name (e.g.: MdNumber). Metadata classes have the following properties:

- Metadata classes can only contain a limited set of CiiBaseTypes (CII Basic Data Type Set), which are synced with corresponding language types.
- Metadata classes cannot contain complex types.
- Metadata member fields cannot contain big arrays. These are limited to the special file field, which only exist in MdArray and all its derivatives.
- Metadata classes support generic member types, so the type of desired members can be set according to the defined generic type of the class (as Java generics).
- All metadata classes derive from basic class (MdBase).
- All metadata classes must be stored in the `elt.config.classes.meta` package.

Set of basic metadata classes is pre-defined. These classes derive directly from MdBase and cannot be deleted or changed as they are part of the configuration system. Basic metadata classes also contain check functions. The check functions are executed when calling the `setValue` method. The exception is thrown when the value is not properly evaluated by the check function.

MdBase is a root metadata class for all metadata. All other metadata classes must derive from this class. Basic types are defined by metadata classes. The metadata types are shown in Figure 2-4 and in the table in Appendix C.

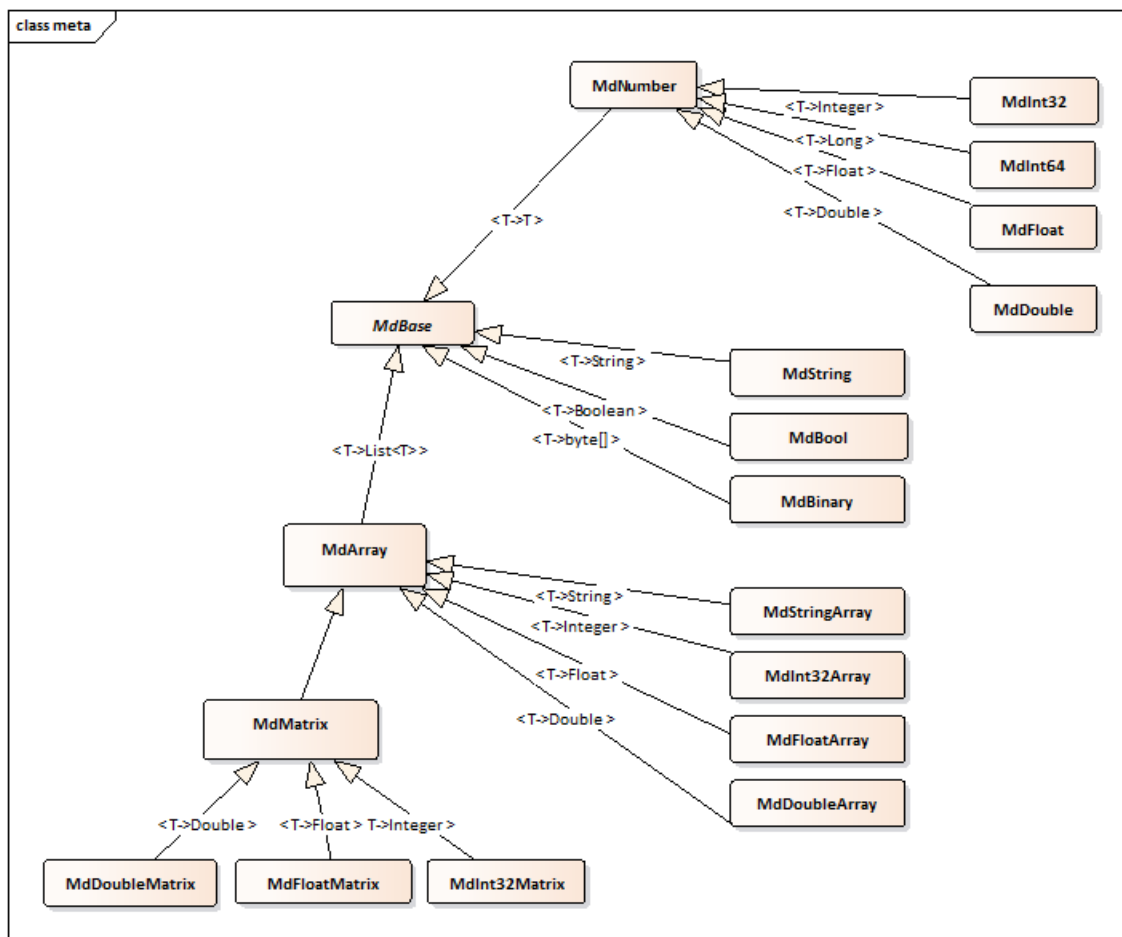


Figure 2-4 Metadata type system

To store array and matrix values, MdArray and MdMatrix derivatives must be used.

The data for matrix is represented as a single array, with specified width (member field of the MdMatrix class), so matrix is represented as a continuous array of elements. The width attribute determines at which array element number the array should be cut into row to form a matrix. All matrices are row based; this means that the data is split into rows. When using default values, the width is set to 50.

Binary values are always stored as file references. Like for all the other types of data, the value for binary is set by using the setValue and retrieved by using the getValue method. On setValue method call, the data is written to a file and a reference is stored in the target configuration JSON file. On getValue method call, the file reference is resolved and data from the reference is retrieved. There is no limit specified for the binary data values.

Vector (array) and matrix values are stored as a string up to the limit of 100KB. All values exceeding this limit are stored as file references.

For more complex configurations, custom metadata classes can be defined. The classes can extend existing metadata classes.

The data of the config and metadata classes cannot be manipulated using the config client API.



Classes and metadata classes data shall be centrally stored in the remote DB. The config GUI is used to manipulate the classes on the remote DB. For the development and testing purposes the classes can also be stored on local developer or test machines as YAML definitions.

13.2.2 Instances

Instances are actual value holders for the corresponding configuration and metadata classes. Instances are entities of JSON structure stored in local or remote DB and bind data to corresponding class defined fields. Class and instances member fields must be linked by name, otherwise exception `CiiConfigTcBindException` is thrown.

Metadata instances

Metadata instances define metadata for the config point. Metadata instances have the following properties:

- Metadata instances are identified by name. The names must be unique.
- Metadata instances can be created, updated and deleted.
- They can be stored to cache, local DB or remote DB. The asterisk "*" authority in metadata instances is not supported. The proper storage location for reading the reference directly from the TC is determined by the TC URI authority.
- Config point metadata is configured by metadata instance. Configuration instances only hold references to metadata instances. As metadata instances are not directly part of configuration instances, they are stored as separate entities.
- Metadata instances can be versioned.
- Set of default (predefined) metadata instances is provided (Appendix C). These default metadata instances cannot be saved or updated.

When `saveTargetConfigWithMetadata` method is used, metadata instances are automatically saved with target configuration. For other types of metadata operations, API contains methods that support operations just with metadata instances. These methods enable saving, updating and deleting metadata instances.




Figure 2-5 Metadata instances

Figure 2-5 shows the relationship between the metadata classes and metadata instances. The metadata class MdInt32 which derives from the MdNumber and root metadata class MdBase defines the structure for the metadata instances. Three metadata instances are created based on the MdInt32 class. The ConfInt32Std is a basic predefined metadata instance and has no special values defined. The custom metadata instances ConfCustomInstance1 and ConfCustomInstance2 have custom values set. The values of the class member fields (defaultValue, minValue, maxValue, checked, etc.) can be set for these instances.

Configuration instance

Target configuration (configuration instance) contains actual configuration data for the specific configuration class. The configuration instance is identified with URI location. Each URI defines the location from which the configuration instance is obtained. The authority part of URI defines the storage location. The following authorities are supported:

- Cache: 'cii.config://cache/exampleconfig/c1'
- Local DB: 'cii.config://local/exampleconfig/c1'
- Remote DB: 'cii.config://remote/exampleconfig/c1'
- All locations: "cii.config://*/exampleconfig/c1"



Note: URI locations that start with number are not allowed. Examples of not valid URIs: `cii.config:///exampleconfig/1abc` or `cii.config:///1exampleconfig/abc`

The asterisk "*" authority can be used for retrieving configuration instance. When saving the instances, it is better to provide a specific location. When "*" authority is used, target configuration is obtained in the following order: cache, local DB, remote DB. In case there is no data stored under specified URI in the specific storage location, the location is skipped.

If there is no explicit version specified for retrieving the target configuration, the last existing version stored in the specific storage location will be retrieved. These versions can differ between the storage locations. As the config service can be used also without the presence of the remote DB, an additional setting was added to the client, to control the exception case.

The setting `exceptionWhenRemoteNotAccessible` can be set. Appendix G contains details on how to apply this setting to the client.

The following rules apply when retrieving the Target configuration with the "*" authority is used:

- `exceptionWhenRemoteNotAccessible=true`: If no remote DB is present – throw an exception.
- `exceptionWhenRemoteNotAccessible=false`: If no remote is present, check local and cache.

When saving or updating the Target configuration using the "*" authority, the versions between the different storage locations can become out of sync. In general, the following rules apply:

- `saveTC(*)`: save always creates a new version.
- `updateTC(*)`: always update the last version.
- `updateTC(specificVersion)` – updates specific version. It is always expected for a specific version on location (cache, local, remote) to exist. If it doesn't exist an exception is thrown.

As the configuration can be used without the remote DB, the following specifics apply:

- `exceptionWhenRemoteNotAccessible=true`:
 - `saveTC(*)`: remote DB is version maintainer. Max version off all the locations is retrieved, using the rule (cache|local|remote) + 1. This version is used to make new cache|remote|local Target configuration. This way versions between the locations are always in sync, but version holes can exist.
 - `updateTC(*)`: remote DB is version maintainer (the last version is retrieved from the remote DB). It is expected that the same version exists in local and cache and that this is the last version (important). If this is not the case, an exception is thrown.
 - `updateTC(specificVersion)`: All versions are expected to exist. Updates for the specific version to all locations are made (cache|local|remote). If there is no specific version present in any of the DB, an exception is thrown.
- `exceptionWhenRemoteNotAccessible=false`:

If remote DB is present, the behavior is the same as above. If not, the following actions are applied after the timeout:
- `saveTC(*)`: local DB is version maintainer. We get max version from (cache|local) + 1. We use



this version to make a new cache|local TC.

- `updateTC("**")`: local DB is version maintainer (the last version is retrieved from the local DB). It is expected that the same version exists in cache and local and that this is the last version. If this is not the case, an exception is thrown.
 - `updateTC(specificVersion)`: All versions are expected to exist. Updates for the specific version to these locations are made (cache|local). If there is no specific version present in any of the DB, an exception is thrown.

To use the configuration in Local DB mode, YAML definition files for classes and values must be created. The *config-tool* shall be used to generate language-specific classes and to deploy configuration data into the local DB.

When working with Remote DB, *config-gui* shall be used to manipulate the data. The remote DB also contains the config and metadata class definitions.

A cache can be used to store the data into memory. From the operations point of view, the cache acts the same as the local DB or remote DB and it is transparent to the client. The major difference to the other storage location is that cache is isolated per process and will not provide any permanent storage. The major benefit of using the cache is the fast storage and retrieval of configuration data. The cache is cleared when the process terminates.

For the configuration instances that use the config point fields of `MdArray` and `MdBinary` class, the data is not stored as a part of the JSON, but it is stored to a large storage location. The data to the large storage location is written as raw data in the big-endian order.



13.3 Prerequisites

In order to use the configuration system, the following packages must be built and installed:

- MAL, MAL ZPB (elt-mal)
- Client APIs (client-APIs)
- Config service (srv-config)
- Config tools (config-tools)
- Config GUI (config-gui)

Config service depends on the ElasticSearch. ElasticSearch server must be accessible to the config service. The default configuration assumes that ElasticSearch is available on <http://ciielastichost:9200>. Appendix H contains details on how to configure the config service.

For remote large data storage, the Hadoop HDFS system is used. HDFS system must be accessible in order to store large data files. The default configuration assumes that HDFS is accessible on <http://ciihdfshost:9870/>.

To use the remote database, the remote config service has to be installed and started.

See Error Handling and Configuration Transfer Document [2] for the instructions on how to install and configure the configuration system and its components. The document also contains a description of how to set the ElasticSearch and Hadoop hosts.

13.3.1 Local DB initialization

Local DB stores it's data into the **\$INTROOT/localdb** folder. To use the default metadata instances used in the examples, the initial predefined metadata instances must be deployed:

```
config-tool init
```




13.3.2 Includes/Imports

For basic usage of the Config API the user needs the CiiConfigClient class and all generated binding configuration classes that will be used. If any custom metadata classes are used, the imports of these classes are also needed. Next, the language specific URI class is needed. Java also needs import of all the checked exceptions used. Following sections present the code imports and includes:

Java imports

```
import elt.config.client.CiiConfigClient;  
import elt.config.exceptions.CiiConfigInitializationError;  
import elt.config.exceptions.CiiConfigNoTcException;  
import elt.config.exceptions.CiiConfigWriteDisabledException;  
import elt.config.exceptions.CiiConfigSaveException;  
import elt.config.exceptions.CiiConfigWriteDisabledException;  
import elt.error.CiiInvalidTypeException;  
import elt.error.CiiInvalidURIException;  
import java.net.URI;  
import java.net.URISyntaxException;
```

CPP Includes

```
#include <ciiDataProviderFactory.hpp>  
#include <ciiLocalDataProvider.hpp>  
#include <ciiCacheDataProvider.hpp>  
#include <ciiRemoteDataProvider.hpp>  
#include <ciiConfigApi.hpp>
```

Python imports

```
import os  
import elt.config
```



13.4 Basic usage

This section presents basic usage examples. The examples in this section can be found in the *config-examples* project (M*ySimpleClient.java, MySimpleClient.py, mySimpleClient.cpp*). Refer to the *config-examples* project for details about files, folder structure and the WAF wscript files configuration. The examples project is build using the standard WAF build procedure:

```
waf distclean configure build install
```

To run the examples, use the provided executables (*run-java*, *cpp-config-sample-app*, *python MySimpleClient.py*).

The examples show the usage of Config tool (6.1) and Config GUI (7). The Config tool only deploys/undeploys data from YAML files to local DB. Config tool is intended to be used in the development process, where developer can quickly change definitions in YAML files on local machine, while Config GUI will work with remote with all the defined TCs. Config tool cannot be used to deploy data to remote DB. This can however be done with usage of Client API (the data can be read from local DB and stored to remote DB). For this reason, the examples show usage of both Config tools and Config GUI in for the same example.

Note: In case examples are run directly from *config-examples* project, make sure that all the needed local DB configuration instance files are deployed. The deploy commands are listed in *README.md* file of the project.

13.4.1 Example description

This section presents simple CII configuration example. All the YAML definitions in this example are also stored in the *yaml* folder of the *config-examples* project.

In the example we create a simple configuration with one string, one double and one integer type field. The simple configuration uses default metadata instances. The following YAML class definition is used (mySimpleClass.cdt.yaml):

```
---
configuration:
  configClassesConfiguration:
    classes:
      - name: MySimpleClass
        parent: CiiConfigClass
        __comment__: Personal config test class
        members:
          - name: testIntValue
            type: MdInt32
          - name: testDoubleValue
            type: MdDouble
```

(continues on next page)



(continued from previous page)

```
- name: testStringValue  
  type: MdString
```

The types names and language mappings for all types can be found in Appendix A. The syntax of YAML schema is explained in Appendix E. Not all words are allowed for the field names of config or metadata classes. The table in Appendix D contains all the reserved words.

In the example we create a sample application. The application is placed in the root of config-examples project. To prepare the directory structure we create the following subdirectories: java, cpp, cpp-app, cpp-pybindings, python and YAML:

```
mkdir java  
mkdir cpp  
mkdir cpp-app  
mkdir cpp-pybindings  
mkdir python  
mkdir yaml
```

13.4.2 Creating a simple configuration using YAML definitions

In the simple class definition, we defined the structure of the configuration class. We store this definition in YAML file named *mySimpleClass.cdt.yaml* in the *yaml* directory. Based on this YAML definition we must first generate the classes used by the application:

```
config-tool generate -i yaml/mySimpleClass.cdt.yaml -o java/  
config-tool generate -i yaml/mySimpleClass.cdt.yaml -o cpp/ -l cpp  
config-tool generate -i yaml/mySimpleClass.cdt.yaml -o cpp-pybindings/ -l python_  
↩-ws
```

The sample code generated by the config-tool can be found in Appendix I.

Once the class is prepared, we must also define target configuration YAML. The target configuration will use the generated class *MySimpleClass* and default metadata instances. List of default metadata instances with internal types is listed in Appendix C. The following TC YAML is prepared:

```
config:  
  instance:  
    __comment__: Configuration for cryo cooler  
    data:  
      "@type": MySimpleClass  
      fields:  
        testStringValue:  
          "@type": MdString  
          metadataInstance: ConfStringStd  
          metadataInstanceVersion: 1  
          value: 'My string defined value'  
        testIntValue:
```

(continues on next page)



(continued from previous page)

```
"@type": MdInt32
metadataInstance: ConfInt32Std
metadataInstanceVersion: 1
value: 155
testDoubleValue:
  "@type": MdDouble
  metadataInstance: ConfDoubleStd
  metadataInstanceVersion: 1
  value: 155.33
```

We store the target configuration under the *testTC.cdi.yaml* file in *yaml* directory.

Note: Types defined in target configuration must be the same as types defined in the classes in order for the configuration to work properly. The depending on the language the JSON (de)serializers will throw exception with indication that (de)serialization could not be executed. For the types with embedded arrays, the types used in target configuration must have the vector types defined. Please refer to Appendix C for details.

To have the default metadata instances ready, we must first initialize the local DB:

```
config-tool init
```

Then we deploy the target configuration to local DB under **cii.config://mytest/mysampletc**:

```
config-tool deploy -i yaml/testTC.cdi.yaml -a cii.config://mytest/mysampletc
```

After the target config is deployed to local DB, we can start using the clients. For successful code builds, proper build dependencies must be set for each of the clients. Please check the *config-examples* project for details about the WAF dependencies.

Results of deploying *testTC.cdi.yaml* to local DB (JSON file) can be found in Listing 7-5.

Java client

For the Java sample, we will put our code in the *java/src* directory. We name our client sample code *mySimpleClient.java*. We use "*" authority for retrieving the stored target config. Data is stored under **cii.config://local/mytest/mysampletc** location.

Listing 4-1 contains the sample code. In the example we first read the data that was previously deployed in local DB. The client API will not allow save, update and delete operations. To enable these operations, we set the attribute `setWriteEnabled` to true. The data is then changed and stored under a different location. At the end the data is read in again (Listing 4-1).

Listing 4-1 Simple client Java sample code



```
public static void main(String[] args) throws Exception {

    URI myUri = new URI("cii.config://*/mytest/mysampletc");

    //Retrieve config from local or remote DB
    try(CiiConfigClient client=CiiConfigClient.getInstance()){
        MySimpleClass data = client.retrieveConfig(myUri).getData(MySimpleClass.
↪class);
        System.out.println("My double value:" + data.getTestDoubleValue().
↪getValue());
        System.out.println("My int value:" + data.getTestIntValue().getValue());
        System.out.println("My string value:" + data.getTestStringValue().
↪getValue());

        //Store changed data into new TC location -- local DB
        client.setWriteEnabled(true);
        URI localUri = new URI("cii.config://local/mytest/mysampletc2");
        data.getTestDoubleValue().setValue(121442242.224242424);
        CiiConfigClient.getInstance().saveTargetConfig(localUri, data);

        //Read changed data
        MySimpleClass changedData = client.retrieveConfig(localUri).
↪getData(MySimpleClass.class);
        System.out.println("My double value:" + changedData.getTestDoubleValue().
↪getValue());
        System.out.println("My int value:" + changedData.getTestIntValue().
↪getValue());
        System.out.println("My string value:" + changedData.getTestStringValue().
↪getValue());
    }
}
}
```

CPP client

Listing 4-2 contains a sample CPP client code. The code file is named *mySimpleClient.cpp* and stored in the *cpp-app* folder. The code follows the same flow as the Java code.

Listing 4-2 Simple client CPP sample code

```
int main(int ac, char* av[]) {

::elt::config::classes::MySimpleClass mySimpleClassDataCh;

std::shared_ptr<::elt::config::classes::MySimpleClass> retrievedData =
    elt::config::CiiConfigClient::getConfigData<
        ::elt::config::classes::MySimpleClass>(
```

(continues on next page)



(continued from previous page)

```
        elt::mal::Uri("cii.config://*/mytest/mysampletc"));  
  
std::cout << "My int value:" << retrievedData->get_testIntValue() << std::endl;  
std::cout << "My double value:" << retrievedData->get_testDoubleValue() <<  
↳std::endl;  
std::cout << "My string value:" << retrievedData->get_testStringValue() <<  
↳std::endl;  
  
//Store changed data into new TC location -- local DB  
::elt::config::CiiConfigClient::setWriteEnabled(true);  
mySimpleClassDataCh.set_testStringValue("Changed string value");  
mySimpleClassDataCh.set_testIntValue(retrievedData->get_testIntValue());  
mySimpleClassDataCh.set_testDoubleValue(retrievedData->get_testDoubleValue());  
  
int version = elt::config::CiiConfigClient::saveTargetConfig(  
    elt::mal::Uri("cii.config://local/mytest/mysampletc2"),  
↳mySimpleClassDataCh);  
  
//Read changed data  
std::shared_ptr<::elt::config::classes::MySimpleClass> retrievedData2 =  
    elt::config::CiiConfigClient::getConfigData<  
        ::elt::config::classes::MySimpleClass>(  
        elt::mal::Uri("cii.config://local/mytest/mysampletc2"));  
  
std::cout << "My int value:" << retrievedData2->get_testIntValue() << std::endl;  
std::cout << "My double value:" << retrievedData2->get_testDoubleValue() <<  
↳std::endl;  
std::cout << "My string value:" << retrievedData2->get_testStringValue() <<  
↳std::endl;  
  
}
```

Python client

Listing 4-3 contains a sample Python code. The code file is named *mySimpleClient.py* and stored in the *python* folder. The code follows the same flow as the CPP code.

Listing 4-3 Simple client Python sample code

```
import os  
import elt.config  
from MySimpleClassPyB import MySimpleClass  
  
uri = elt.config.Uri("cii.config://*/mytest/mysampletc")  
uri2 = elt.config.Uri("cii.config://local/mytest/mysampletc2")
```

(continues on next page)



(continued from previous page)

```
client = elt.config.CiiConfigClient

#Retrieve config configuration
retrieved = client.retrieve_config(uri, -1)

print("My string value:" + retrieved.get_testStringValue())
print("My integer value: {}".format(retrieved.get_testIntValue()))
print("My double value: {}".format(retrieved.get_testDoubleValue()))

#Store changed data into new TC location -- local DB
client.set_write_enabled(True)

retrievedNew = MySimpleClass.get_new_node_instance()
retrievedNew.set_testStringValue("New test value")
retrievedNew.set_testIntValue(retrieved.get_testIntValue())
retrievedNew.set_testDoubleValue(retrieved.get_testDoubleValue())

version = client.save_target_config(uri2, retrievedNew)
retrieved2 = client.retrieve_config(uri2, -1)

print("My string value:" + retrieved2.get_testStringValue())
print("My integer value: {}".format(retrieved2.get_testIntValue()))
print("My double value: {}".format(retrieved2.get_testDoubleValue()))
```

13.4.3 Creating a simple configuration using GUI

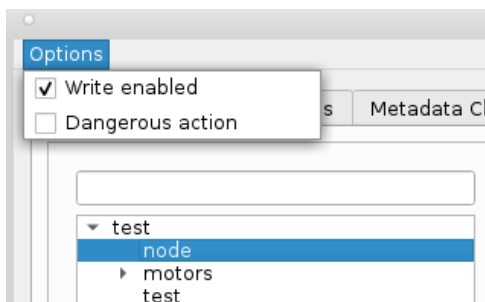
In this section GUI is used to define classes and target configuration on remote DB.

We first start the configuration GUI:

```
config-gui
```

Make sure that GUI has write tick-box enabled (Figure 4-1):

Figure 4-1 Write enabled tick box

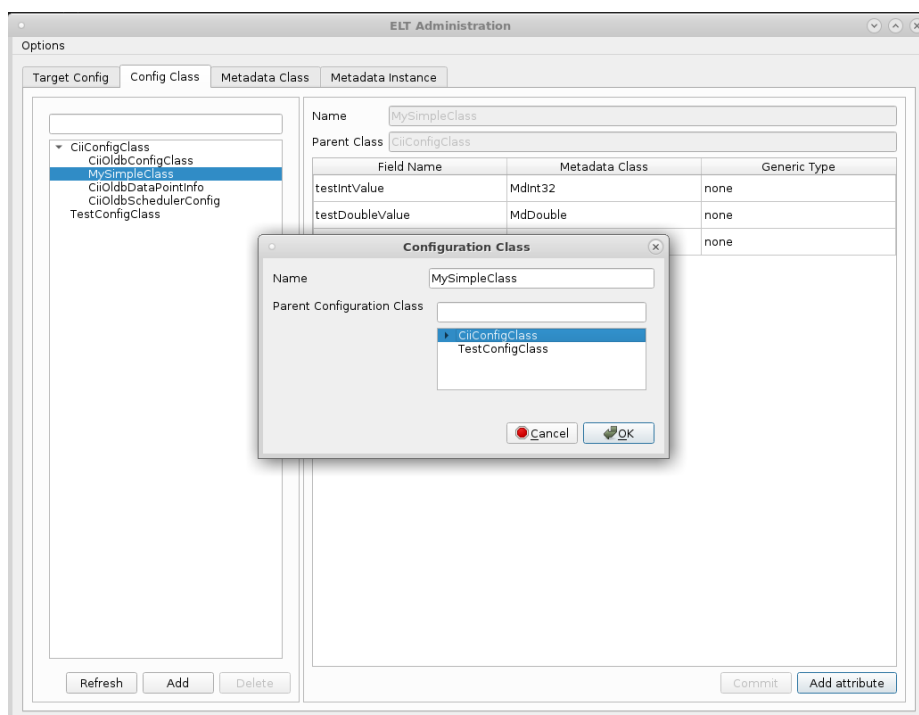




Creating a simple class in GUI

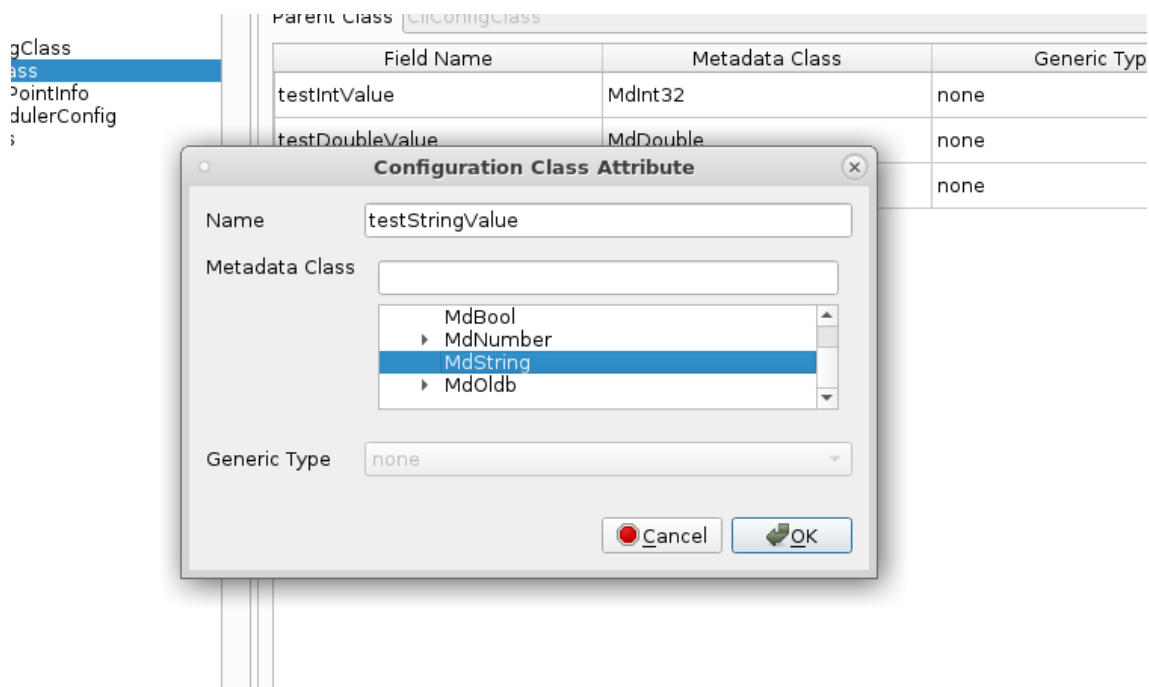
Click on the Configuration class tab, and click *Add* button to add a configuration class (Figure 4-2). Select the *CiiConfigClass* to be the parent class, and name the class *MySimpleClass*.

Figure 4-2 Adding configuration class in GUI



After the class is created, we add attributes to the class. To add attributes (members), we click on the *Add attribute* button and fill in the details of every class attribute (Figure 4-3). We use the same names and types as in the YAML class definition in section 4.1.

Figure 4-3 Adding class members in GUI

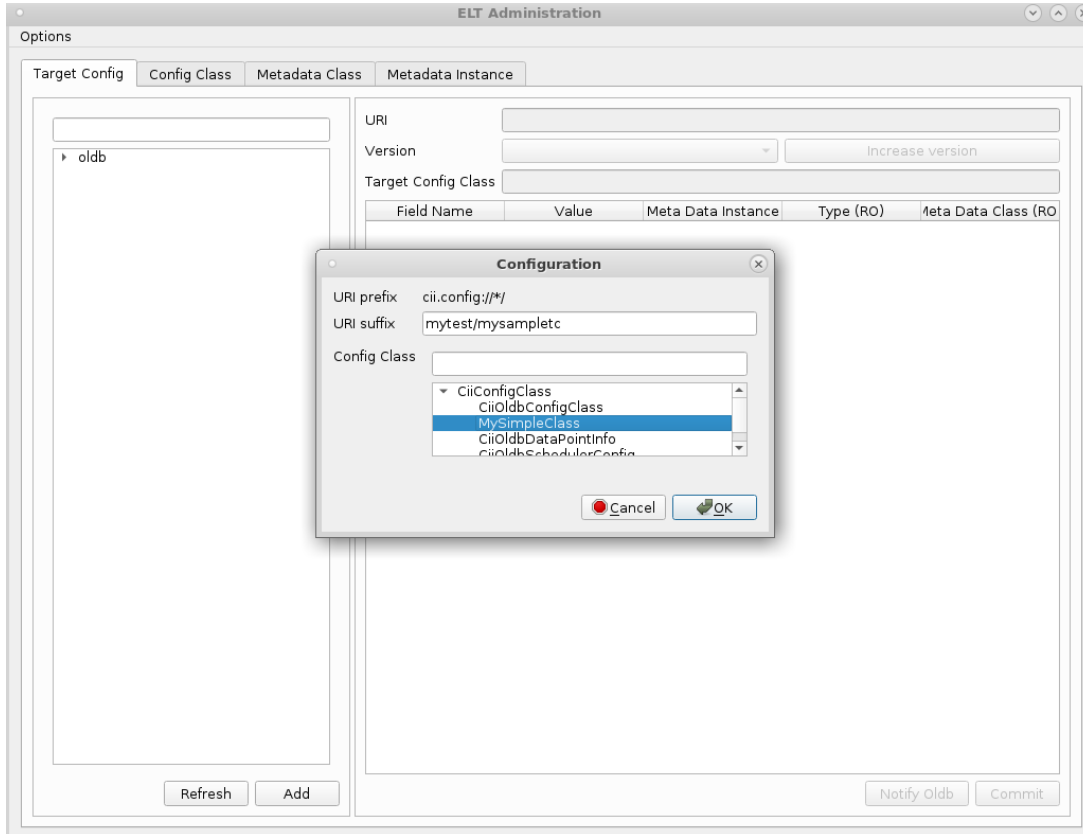


After we added all the attributes (*testIntValue*, *testDoubleValue*, *testStringValue*), definition changes must be confirmed and propagated to remote DB by clicking on the *Commit* button.

Creating a target configuration for simple class in GUI

To create a TC we click on the Target Config tab and then the Add button. We choose the proper Configuration class and define a URI suffix that corresponds to the URI used when deploying to the local DB, for example **mytest/mysampletc**. (Figure 4-4)

Figure 4-4 Creating target configuration



After the target configuration is created, we must refresh the tree by clicking the “Refresh button”. Next, we bind the metadata instances used by the TC to the member fields. To select the metadata instance, we click on the cells (Figure 4-6) under Meta Data Instance column. For every class member field name, we choose metadata instance (Figure 4-5). Predefined default metadata instances are used (ConfInt32Std, ConfDoubleStd, ConfStringStd).

Figure 4-5 Target configuration definition in GUI





Figure 4-6 Target configuration definition in GUI



After metadata instances are set, we set the values for every member (cells under Value column), and confirm the changes by clicking on the *Commit* button (Figure 4-6).

Generating classes from the classes defined in GUI

To generate classes defined in GUI (stored in remote DB), the config tool is called without an input argument:

```
config-tool generate -o java/  
config-tool generate -o cpp/ -l cpp  
config-tool generate -o cpp-pybindings/ -l python -ws
```

To read the data with the client, we use the same code as shown in the previous example (Listing 4-1, Listing 4-2, Listing 4-3). To read the data from the remote DB, the “remote” authority must be used. In case “*” authority is used, all data will first be read from the local DB. Only if the data is not present in local DB, the remote DB will be used.



13.4.4 CRUD operations

The following section presents the code for creating, reading, updating and deleting metadata and configuration instances. The examples for all languages follow the same flow, and can be easily understood from the code samples. The examples first show the operations on the metadata instance, followed by operation on target configuration. The TC is saved under the `cii.config://remote/mytest/my_test_instance` URI. Config API also supports the listing of config points stored under certain path. To get all the children, the `getChildren` method can be used. Refer to Appendix A for details about the API.

Note: When setting only one attribute of a configuration programmatically, the metadata for the other attributes are not set automatically. It is expected that all the fields of the configuration class are set by user.

The same `MySimpleClass` as it was used in section 4.2 is also used in the following examples.

Java CRUD

Listing 4-4 Java metadata CRUD example

```
//Metadata CRUD example
final String mdiName = "MyCustomMetaInstanceName";
MyCustomMetaClass<Integer> customMetaClass = new MyCustomMetaClass<>(mdiName,
↪ "This is my custom metadata instance name",
    123, false, "m/s", 9);
URI metadataLocal = new URI("cii.config://local");
//Save metadata under local DB with name MyCustomMetaInstanceName
int version2 = client.saveMetadata(metadataLocal, customMetaClass);

//Retrieve and update the same version
MyCustomMetaClass<Integer> customMetaClassChange = client.
↪ retrieveMetadata(metadataLocal, mdiName, -1, MyCustomMetaClass.class);
System.out.println("Metadata int:" + customMetaClassChange.getMetaOfInt());

customMetaClassChange.setMetaOfInt(15);
client.updateMetadata(metadataLocal, mdiName, version2, customMetaClassChange);

//delete the metadata
System.out.println("Deleting metadata: " + mdiName + ", version: " + version2);
client.deleteMetadata(metadataLocal, mdiName, version2);
```

Listing 4-5 Java target config CRUD example

```
//TC CRUD example
//Create another instance of MySimple class and store it to remote
MySimpleClass newSimpleInstance = new MySimpleClass();
newSimpleInstance.setTestDoubleValueValue(12.0d);
newSimpleInstance.setTestIntValueValue(100);
```

(continues on next page)



(continued from previous page)

```
newSimpleInstance.setTestStringValueValue("My test string");

URI newInstanceUri = new URI("cii.config://remote/mytest/my_test_instance");
int version1 = client.saveTargetConfig(newInstanceUri, newSimpleInstance);

//Retrieve and update the same version
MySimpleClass newSimpleInstanceChange = client.retrieveConfig(newInstanceUri).
↳getData(MySimpleClass.class);
newSimpleInstanceChange.setTestStringValueValue("My test string changed");
client.updateConfig(newInstanceUri, version1, newSimpleInstanceChange);

System.out.println("My string value:" + newSimpleInstanceChange.
↳getTestStringValue().getValue());

//Delete TC
client.deleteConfig(newInstanceUri, -1);
```

CPP CRUD

Listing 4-6 CPP metadata CRUD example

```
//Metadata CRUD example
std::string mdiName("MyCustomMetaInstanceName");
std::shared_ptr<elt::config::classes::meta::MyCustomMetaClass<std::int32_t>>
↳customMetaClass(
    elt::config::CiiDataPointMetadataFactory::getNewMetadataInstance
↳<elt::config::classes::meta::MyCustomMetaClass<std::int32_t>>(mdiName));
customMetaClass->set_metaOfBool(false);
customMetaClass->set_metaOfString("m/s");
customMetaClass->set_metaOfInt(9);

//Save metadata under local DB with name MyCustomMetaInstanceName
int version2 = elt::config::CiiConfigClient::saveMetadata("local",
↳*customMetaClass);

//Retrieve and update the same version
std::shared_ptr<elt::config::classes::meta::MyCustomMetaClass<std::int32_t>>
↳customMetaClassChange =
    std::dynamic_pointer_cast<elt::config::classes::meta::MyCustomMetaClass
↳<std::int32_t>>(elt::config::CiiConfigClient::retrieveMetadata("local",
↳mdiName));

std::cout << "Metadata int:" << customMetaClassChange->get_metaOfInt() <<
↳std::endl;
customMetaClassChange->set_metaOfInt(15);

elt::config::CiiConfigClient::updateMetadata("local", version2,
↳*customMetaClassChange);
```

(continues on next page)



(continued from previous page)

```
//delete the metadata
std::cout << "Deleting metadata: " << mdiName << ", version: " << version2 <<
↳std::endl;
elt::config::CiiConfigClient::deleteMetadata("local", mdiName, version2);
```

Listing 4-7 CPP target config CRUD example

```
//TC CRUD example
//Create another instance of MySimple class and store it to remote
std::shared_ptr<::elt::config::classes::MySimpleClass> newSimpleInstance =
    CiiNodeFactory::getNewNodeInstance<::elt::config::classes::MySimpleClass>
↳();
newSimpleInstance->set_testStringValue("My test string");
newSimpleInstance->set_testIntValue(100);
newSimpleInstance->set_testDoubleValue(12.0);

elt::mal::Uri newInstanceUri("cii.config://remote/mytest/my_test_instance");
int version1 = elt::config::CiiConfigClient::saveTargetConfig(newInstanceUri,
↳*newSimpleInstance);

//Retrieve and update the same version
std::shared_ptr<::elt::config::classes::MySimpleClass> newSimpleInstanceChange =
elt::config::CiiConfigClient::getConfigData
↳<::elt::config::classes::MySimpleClass>(newInstanceUri);

newSimpleInstanceChange->set_testStringValue("My test string changed");
elt::config::CiiConfigClient::updateConfig(newInstanceUri, version1,
↳*newSimpleInstanceChange);

std::cout << "My string value:" << newSimpleInstanceChange->get_
↳testStringValue() << std::endl;

//Delete TC
elt::config::CiiConfigClient::deleteConfig(newInstanceUri, -1);
```

Python CRUD

Listing 4-8 Python metadata CRUD example

```
#Metadata CRUD example
mdiName = "MyCustomMetaInstanceName"
customMetaClass = MyCustomMetaClassINT32.get_new_metadata_instance(mdiName)
customMetaClass.set_metaOfBool(False);
customMetaClass.set_metaOfString("m/s");
customMetaClass.set_metaOfInt(9);
```

(continues on next page)



(continued from previous page)

```
#Save metadata under local DB with name MyCustomMetaInstanceName
version2 = client.save_metadata("local", customMetaClass);

#Retrieve and update the same version
customMetaClassChange = client.retrieve_metadata("local", mdiName)
customMetaClassChange = MyCustomMetaClassINT32.cast(customMetaClassChange)

#Cannot print as wrong type is returned.
print("Metadata int: %i" % customMetaClassChange.get_metaOfInt())
customMetaClassChange.set_metaOfInt(15)
print("Metadata int changed: %i" % customMetaClassChange.get_metaOfInt())
client.update_metadata("local", version2, customMetaClassChange);

#delete the metadata
print("Deleting metadata: %s, version: %i" % (mdiName, version2))
client.delete_metadata("local", mdiName, version2);
```

Listing 4-9 Python target config CRUD example

```
#TC CRUD example
#Create another instance of MySimple class and store it to remote
newSimpleInstance = MySimpleClass.get_new_node_instance()
newSimpleInstance.set_testStringValue("My test string")
newSimpleInstance.set_testIntValue(100)
newSimpleInstance.set_testDoubleValue(12.0)

newInstanceUri = elt.config.Uri("cii.config://remote/mytest/my_test_instance")
version1 = client.save_target_config(newInstanceUri, newSimpleInstance)

#Retrieve and update the same version
newSimpleInstanceChange = client.retrieve_config(newInstanceUri)

newSimpleInstanceChange.set_testStringValue("My test string changed");
client.update_config(newInstanceUri, version1, newSimpleInstanceChange);

print("My integer value: {}".format(newSimpleInstanceChange.get_
↪testStringValue()))

#Delete TC
client.delete_config(newInstanceUri, -1);
```



13.4.5 Using the cache operations

The following examples present the usage of cache operations. From the operations point of view, the cache acts the same as the local DB or remote DB and it is transparent to the client. The cache is not permanent storage, as it lives with the process. The cache can be used for inter-process fast storing/retrieval of Target Configurations.

The example first copies Target Configuration from the remote DB (cii.config://remote/mytest/my_test_instance), then it stores it to the cache. Then the values are changed and stored back to same location in cache (cii.config://cache/mytest/my_test_instance). Finally, these values are stored back to the original remote DB location.

In the example, the same path is used for remote DB and cache. If the same path is used, the Target Configuration can be retrieved using the "*" authority. It will first check the cache, then the local and remote DB. The same Configuration class data can be stored to any URI, as there is no limitation for cache to be on the same path as remote and local DB.

Java cache example

Listing 4-10 Java cache example

```
//The cache workflow -- data in cache only lives as long as the process is alive
//Retrieve the TC from the remote, and store it to cache.

URI newInstanceUri = new
    URI("cii.config://remote/mytest/my_test_instance");
URI cacheInstance = new URI("cii.config://cache/mytest/my_test_instance");
client.saveTargetConfig(cacheInstance,
    client.retrieveConfig(newInstanceUri).getData(MySimpleClass.class));

MySimpleClass dataFromCache =
    client.retrieveConfig(cacheInstance).getData(MySimpleClass.class);

System.out.println("Original cache data: " +
    dataFromCache.getTestIntValueValue() + ", " +
    dataFromCache.getTestStringValueValue());
//manipulate cache data
dataFromCache.setTestIntValueValue(15);
dataFromCache.setTestStringValueValue("My string value");
System.out.println("Changed cache data: " +
    dataFromCache.getTestIntValueValue() + ", " +
    dataFromCache.getTestStringValueValue());

//store data into cache
client.saveTargetConfig(cacheInstance, dataFromCache);

//retrieve data form cache and store it to remote.
MySimpleClass dataFromCacheChanged =
```

(continues on next page)



(continued from previous page)

```
client.retrieveConfig(cacheInstance).getData(MySimpleClass.class);
System.out.println("Retrieved changed cache data: " +
    dataFromCacheChanged.getTestIntValueValue() + ", " +
    dataFromCacheChanged.getTestStringValueValue());

//save the changed data to remote
client.saveTargetConfig(newInstanceUri, dataFromCacheChanged);

//Delete TC
client.deleteConfig(newInstanceUri, -1);
```

CPP Cache

Listing 4-11 CPP cache example

```
//The cache workflow -- data in cache only lives as long as the process is alive
//Retrieve the TC from the remote, and store it to cache.
elt::mal::Uri newInstanceUri
    ("cii.config://remote/mytest/my_test_instance");
elt::mal::Uri cacheInstance("cii.config://cache/mytest/my_test_instance");

//as there are no metadata instances in cache, they must be saved
elt::config::CiiConfigClient::SaveTargetConfig(cacheInstance,
    *elt::config::CiiConfigClient::GetConfigData<::elt::config::classes::
    MySimpleClass>(newInstanceUri));

std::shared_ptr<::elt::config::classes::MySimpleClass> dataFromCache =
    elt::config::CiiConfigClient::GetConfigData<::elt::config::classes::
    MySimpleClass>(cacheInstance);

std::cout << "Original cache data: " << dataFromCache->get_testIntValue()
    << ", " << dataFromCache->get_testStringValue() << std::endl;
//manipulate cache data
dataFromCache->set_testIntValue(15);
dataFromCache->set_testStringValue("My string value");

//store data into cache
elt::config::CiiConfigClient::SaveTargetConfig(cacheInstance, *dataFromCache);

//retrieve data form cache and store it to remote.
std::shared_ptr<::elt::config::classes::MySimpleClass> dataFromCacheChanged =
    elt::config::CiiConfigClient::GetConfigData<::elt::config::classes::
    MySimpleClass>(cacheInstance);
std::cout << "Changed cache data: " << dataFromCacheChanged->get_testIntValue() <
    << ", " << dataFromCacheChanged->get_testStringValue() << std::endl;
```

(continues on next page)



(continued from previous page)

```
//save the changed data to remote
elt::config::CiiConfigClient::SaveTargetConfig(newInstanceUri,
↳*dataFromCacheChanged);

//Delete TC
elt::config::CiiConfigClient::DeleteConfig(newInstanceUri, -1);
```

Python cache example

Listing 4-12 Python cache example

```
#The cache workflow -- data in cache only lives as long as the process is alive
#Retrieve the TC from the remote, and store it to cache.
newInstanceUri = elt.config.Uri("cii.config://remote/mytest/my_test_instance")
cacheInstance = elt.config.Uri("cii.config://cache/mytest/my_test_instance")
client.save_target_config(cacheInstance, client.retrieve_config(newInstanceUri))

dataFromCache = client.retrieve_config(cacheInstance)

print("Original cache data: %i, %s" % (dataFromCache.get_testIntValue(),
↳dataFromCache.get_testStringValue()))

#manipulate cache data
dataFromCache.set_testStringValue("My string value")
dataFromCache.set_testIntValue(15)

#store data into cache
client.save_target_config(cacheInstance, dataFromCache)

#retrieve data form cache and store it to remote.
dataFromCacheChanged = client.retrieve_config(cacheInstance)
print("Changed cache data: %i, %s" % (dataFromCacheChanged.get_testIntValue(),
↳dataFromCacheChanged.get_testStringValue()))

#save the changed data to remote
client.save_target_config(newInstanceUri, dataFromCacheChanged)

#Delete TC
client.delete_config(newInstanceUri, version1)
```



13.5 Advanced usage

This section presents advanced usage examples. The examples in this section are suffixed with advanced in the config-examples project (M*yAdvancedClient.java, MyAdvancedClient.py, myAdvancedClient.cpp*). The examples project is build using the standard WAF build procedure:

```
waf distclean configure build install
```

To run the examples, use the provided executables (*run-java-advanced*, *cpp-config-sample-app-advanced*, *python MyAdvancedClient.py*).

Note: In a case when examples are run directly from *config-examples* project, make sure that all the needed local DB configuration instance files are deployed. The deploy commands are listed in *Readme.md* file of the project.

13.5.1 Custom metadata class

In this example we create a custom metadata class, which will be used as a member field in the configuration class. Listing 5-1 shows the YAML code of the metadata class definition. We name the custom metadata class *MyCustomMetaClass* and save the metadata class definition as *myCustom-Metadata.cdt.yaml* in the folder named *yaml*. The class derives from the *MdNumber*. This class has generic type set to *T*. This means that the actual type of the configuration class member field can be set at the class configuration class definition.

Listing 5-1 MyCustomMetaClass definition

```
---
configuration:
  metadataClassesConfiguration:
    classes:
      - name: MyCustomMetaClass
        parent: MdNumber
        __comment__: Custom metadata class
        genericType: T
        members:
          - name: metaOfBool
            type: BOOLEAN
          - name: metaOfString
            type: STRING
          - name: metaOfInt
            type: INT32
```

The class has 3 metadata fields defined: *metaOfBool* (type Boolean), *metaOfString* (type string), *metaOfInt* (type int32). All other metadata fields that can be used by the metadata instances (*minValue*, *maxValue*, *checked*, *defaultValue*, ...) are already defined by the *MdNumber* and its parent *MdBase*.

Next, we create a configuration class, which will use this custom metadata class. We name the



class `MyClassWithCustomMeta` and save the configuration class definition as `myClassWithCustomMeta.cdt.yaml` in the folder named `yaml`. The class derives from the parent configuration class `CiiConfigClass` and has only one Config point defined: `intPointWithCustomMeta`. As the `MyCustomMetaClass` definition has the type defined as `T`, we need to set the generic type of the member field to proper type. We set the `intPointWithCustomMeta` to be `INT32`. Listing 5-2 shows the YAML definition of the class.

Listing 5-2 `MyClassWithCustomMeta` definition

```
---
configuration:
  configClassesConfiguration:
    classes:
      - name: MyClassWithCustomMeta
        parent: CiiConfigClass
        __comment__: Class with custom metadata
        members:
          - name: intPointWithCustomMeta
            type: MyCustomMetaClass
            genericType: INT32
```

After the configuration classes are defined, we must create YAML files which will define the metadata instance and TC values. Listing 5-3 shows the YAML definition which defines the values of the custom metadata instance member fields for the `MyCustomMetaClass`. We name the definition `myCustomMDI.mdi.yaml` and save it in the `yaml` folder. The definition sets the values and sets the type. As metadata default value can be of any number type, we must also define the type with `@genType` directive in the metadata instances definition file.

Listing 5-3 Definition of custom metadata instance (`myCustomMDI`)

```
metadata:
  instances:
    - "@type": MyCustomMetaClass
      "@name": myCustomMDI
      "@genType": INT32
      comment: "Integer MDI with values"
      metaOfBool: "true"
      metaOfString: "The string metadata example"
      metaOfInt: "12"
```

Finally, we must define the TC values file. We name this file `simpleCustomMetaTC.cdi.yaml` and save it in the `yaml` folder. The file contains the values for linking the metadata instance `myCustomMDI` to the config point `intPointWithCustomMeta`. The file also sets the value of the point to 13. Listing 5-4 shows the file definition.

Listing 5-4 Definition of target config with custom meta data (`simpleCustomMetaTC`)

```
config:
  instance:
    __comment__: Configuration for cryo cooler
```

(continues on next page)



(continued from previous page)

```
data:
  "@type": MyClassWithCustomMeta
  fields:
    intPointWithCustomMeta:
      "@type": MyCustomMetaClass
      "@genType": INT32
      metadataInstance: myCustomMDI
      metadataInstanceVersion: 1
      value: 13
```

Once all the classes are defined, we must generate the metadata binding classes. We generate the classes for all three programming languages with the config tool:

```
config-tool generate -i yaml/myCustomMetadata.mdt.yaml -t metadata -o java/
config-tool generate -i yaml/myCustomMetadata.mdt.yaml -t metadata -o cpp/ -l cpp
config-tool generate -i yaml/myCustomMetadata.mdt.yaml -t metadata -o cpp-
↳ pybindings/ -l python
```

Next, we generate the configuration classes with our custom metadata:

```
config-tool generate -i yaml/myClassWithCustomMeta.cdt.yaml -o java/
config-tool generate -i yaml/myClassWithCustomMeta.cdt.yaml -o cpp/ -l cpp
config-tool generate -i yaml/myClassWithCustomMeta.cdt.yaml -o cpp-pybindings/ -
↳ l python -ws
```

After the language bindings for the configuration and metadata classes are generated, the produced code files can be built. Please note that *config-tool* generator will not generate all the WAF wscript files. For the proper build, the wscript files with dependencies have to be added to the *java*, *cpp*, and *cpp-pybindings* directory. The python bindings will assume, that CPP code and python bindings are located in *cpp* and *cpp-pybindings* directory respectively. The config tool provides option “-ws” to generate wscript file for generated python binding modules.

To retrieve the data from the local DB, the YAML value files have to be deployed. The following example shows the CLI commands to deploy the metadata instance values and the configuration instance values. The target configuration is deployed under the **cii.config://mytest/myCustomMetaTC**:

```
config-tool deploy -i yaml/myCustomMDI.mdi.yaml -t metadata
config-tool deploy -i yaml/simpleCustomMetaTC.cdi.yaml -a cii.config://mytest/my_
↳ custom_meta_tc
```

Note: Config-tool can be used without the authority (*cii.config://mytest/*) or with the “*” authority (*cii.config://mytest/*).* In case of “*” is used, the tool will discard authority and in all cases deploy data to local DB.

Finally, we create the Java (Listing 5-5), CPP (Listing 5-6), and Python (Listing 5-7) sample applications to read the metadata values.



Listing 5-5 Java custom metadata class code example

```
public class mySimpleClient{  
  
public static void main(String[] args) throws Exception {  
  
    URI myUri = new URI("cii.config://local/mytest/my_custom_meta_tc");  
  
    try(CiiConfigClient client=CiiConfigClient.getInstance()){  
  
        MyClassWithCustomMeta data = client.retrieveConfig(myUri).  
↪getData(MyClassWithCustomMeta.class);  
        System.out.println("The value:" + data.getIntPointWithCustomMeta().  
↪getValue());  
        System.out.println("Metadata one:" + data.getIntPointWithCustomMeta().  
↪getMetaOfBool());  
        System.out.println("Metadata two:" + data.getIntPointWithCustomMeta().  
↪getMetaOfString());  
        System.out.println("Metadata three:" + data.getIntPointWithCustomMeta().  
↪getMetaOfInt());  
  
    }  
}  
}
```

Listing 5-6 CPP custom metadata class code example

```
int main(int ac, char* av[]) {  
  
::elt::config::classes::MyClassWithCustomMeta MyClassWithCustomMetaCh;  
  
std::shared_ptr<::elt::config::classes::MyClassWithCustomMeta> retrievedData =  
    elt::config::CiiConfigClient::getConfigData<  
        ::elt::config::classes::MyClassWithCustomMeta>(  
        elt::mal::Uri("cii.config://local/mytest/my_custom_meta_tc"));  
  
std::cout << "The value:" << retrievedData->get_intPointWithCustomMeta() <<  
↪std::endl;  
std::cout << "Metadata one:" << retrievedData->get_metadata_  
↪intPointWithCustomMeta()->get_metaOfBool() << std::endl;  
std::cout << "Metadata two:" << retrievedData->get_metadata_  
↪intPointWithCustomMeta()->get_metaOfString() << std::endl;  
std::cout << "Metadata three:" << retrievedData->get_metadata_  
↪intPointWithCustomMeta()->get_metaOfInt() << std::endl;  
  
}
```

Listing 5-7 Python custom metadata class code example



```
uriCm = elt.config.Uri("cii.config://*/mytest/my_custom_meta_tc")

#Retrieve config configuration
retrievedCm = client.retrieve_config(uriCm, -1)

print("The value: %i" % retrievedCm.get_stringPointWithCustomMeta())
print("The first metadata value : %s" % retrievedCm.get_metadata_
↪stringPointWithCustomMeta().get_metaOfBool())
print("The second metadata value: %s" % retrievedCm.get_metadata_
↪stringPointWithCustomMeta().get_metaOfString())
print("The third metadata value: %s" % retrievedCm.get_metadata_
↪stringPointWithCustomMeta().get_metaOfInt())
```

13.5.2 Using referenced classes

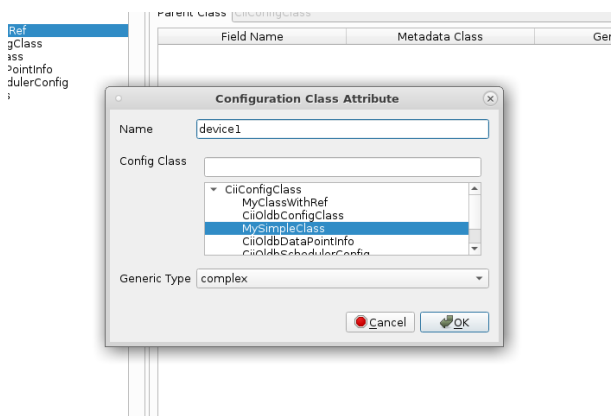
In this example we create a class that has a field which links a referenced class. This means that the field will act as a placeholder for all class members of the class that is defined as referenced. The members can be accessed using the "." dot keyword in the instance YAML definition. The class which we are referencing to must already exist or be defined in the same YAML definition. The class uses **isRef** directive to identify the reference type in the configuration class YAML definition.

In the example we name the class `MyClassWithRef`. The following class definition YAML is prepared (*configClassWithRef.cdt.yaml*):

```
configuration:
  configClassesConfiguration:
    classes:
      - name: MyClassWithRef
        parent: CiiConfigClass
        __comment__: Configuration class with referenced class
        members:
          - name: device1
            type: MySimpleClass
            isRef: true
```

The configuration expects that referenced class, `MySimpleClass` is already present in the source code. To create the same reference field in GUI, the Generic Type "complex" must be used (Figure 5-1).

Figure 5-1 Adding referenced class in GUI



Once the class is prepared, we must define target configuration YAML. The target configuration will use the generated class `MyClassWithRef`. The following TC YAML is prepared (`testRefTc.cdi.yaml`):

```
config:
  instance:
    __comment__: Configuration for cryo cooler
    data:
      "@type": MyClassWithRef
      fields:
        device1.testStringValue:
          "@type": MySimpleClass.MdString
          metadataInstance: ConfStringStd
          metadataInstanceVersion: 1
          value: 'Ref string defined value'
        device1.testIntValue:
          "@type": MySimpleClass.MdInt32
          metadataInstance: ConfInt32Std
          metadataInstanceVersion: 1
          value: 100
        device1.testDoubleValue:
          "@type": MySimpleClass.MdDouble
          metadataInstance: ConfDoubleStd
          metadataInstanceVersion: 1
          value: 200.33
```

When using the referenced classes, target configuration YAML fields and class types must be separated with “.” dot. For example: To set the value of `testDobuleValue` from the referenced field `device1`, the `device1.testDoubleValue` is used for the field value definition. Also, both types must be set, the class type and the metadata type (`MySimpleClass.MdDouble`).

After the class is generated (chapter 4.2 shows the example), we deploy the target configuration under the `cii.config://mytest/myreftc`:

```
config-tool deploy -i yaml/testRefTc.cdi.yaml -a cii.config://mytest/myreftc
```

The following sections show usage examples of referenced class in all languages.



Java referenced type

Listing 5-8 Java referenced type example

```
/** Usage of the referenced class */
CiiConfigClient client=CiiConfigClient.getInstance();

URI myUri = new URI("cii.config://*/mytest/myreftc");

//Retrieve config from local or remote DB
MyClassWithRef data=
    client.retrieveConfig(myUri).getData(MyClassWithRef.class);

//Retrieve referenced class
MySimpleClass dataFromRef = data.getDevice1();

System.out.println("My double value:" +
    dataFromRef.getTestDoubleValueValue());
System.out.println("My int value:" +
    dataFromRef.getTestIntValue().getValue());
System.out.println("My string value:" +
    dataFromRef.getTestStringValue().getValue());
```

CPP referenced type

Listing 5-9 CPP referenced type example

```
/** Usage of the referenced class */
std::shared_ptr<elt::config::classes::MyClassWithRef> retrievedData =
    elt::config::CiiConfigClient::getConfigData<
        :elt::config::classes::MyClassWithRef>(elt::mal::Uri("cii.config://*/mytest/
↵myreftc"));

std::shared_ptr<elt::config::classes::MySimpleClass> retrievedDataRef =
↵retrievedData->get_device1();

std::cout << "My int value:" << retrievedDataRef->get_testIntValue() <<
↵std::endl;
std::cout << "My double value:" << retrievedDataRef->get_testDoubleValue() <<
↵std::endl;
std::cout << "My string value:" << retrievedDataRef->get_testStringValue() <<
↵std::endl;
```



Python referenced type

Listing 5-10 Python referenced type example

```
# Usage of the referenced class
uri = elt.config.Uri("cii.config://*/mytest/myreftc")

client = elt.config.CiiConfigClient

#Retrieve config configuration
retrieved = client.retrieve_config(uri, -1)
retFromRef = retrieved.get_device1()

print("My string value:" + retFromRef.get_testStringValue())
print("My integer value: {}".format(retFromRef.get_testIntValue()))
print("My double value: {}".format(retFromRef.get_testDoubleValue()))
```

13.5.3 Using check values

The metadata classes support a check for setting minimum and maximum values on *number* types. The limits can be set by setting the values of the `minLimit` (`setMinLimit`), `maxLimit` (`setMaxLimit`) metadata instance member fields. The check is disabled by default. To enable the checks, the flag `checked` (`setChecked`) has to be set to true.

The `minLimit` uses “value \geq minLimit” check logic, while `maxLimit` uses “value $<$ maxLimit” check logic. In case the value exceeds the limit, the `CiiConfigLimitOutOfBoundsException` exception is thrown. The check is executed when setting the value using the API.

The `MdString` class also supports allowed values check. This check verifies that the values for config points are part of a set of allowed values. The array of allowed values has to be set. In case the value is not part of the allowed values, the `CiiConfigAllowedValuesException` is thrown.

The following examples show the usage of min and max limits. The usage of allowed values follows the same principle.

Java checked functions

Listing 5-11 Java checked functions example

```
/**
 * Using the check functions
 */
try {
MdInt32 checkedIntClass = MdInt32.getDefaultMDI();
checkedIntClass.setMinLimit(1);
checkedIntClass.setMaxLimit(200);
checkedIntClass.setChecked(true);
```

(continues on next page)



(continued from previous page)

```
//value in limits
checkedIntClass.SetValue(50);

//value out of limits
checkedIntClass.SetValue(200);

MySimpleClass testClass = new MySimpleClass();
testClass.SetTestIntValue(checkedIntClass);
testClass.SetTestDoubleValue(12.0d);
testClass.SetTestStringValue("My test string");

} catch (CiiConfigLimitOutOfBoundsException e) {
    System.out.println(e.getMessage());
}
```

CPP checked functions

Listing 5-12 CPP checked functions example

```
/**
 * Using the check functions
 */

std::shared_ptr<MdNumber<std::int32_t>> checkedIntClass =
    std::dynamic_pointer_cast<elt::config::MdNumber<std::int32_t>>(
    ↪ CiiDataPointMetadataFactory::getMetadataInstance(elt::common::CiiBasicDataType::INT32)
    ↪

checkedIntClass->set_checked(true);
checkedIntClass->set_min_limit(1);
checkedIntClass->set_max_limit(200);

try{
    std::shared_ptr<::elt::config::classes::MySimpleClass> newSimpleInstance =
        CiiNodeFactory::getNewNodeInstance<::elt::config::classes::MySimpleClass>
    ↪ ();

    //value in limits
    newSimpleInstance->set_testIntValue(110, checkedIntClass);

    //value out of limits
    newSimpleInstance->set_testIntValue(201, checkedIntClass);

    newSimpleInstance->set_testStringValue("My test string");
    newSimpleInstance->set_testDoubleValue(12.0);
```

(continues on next page)



(continued from previous page)

```
}  
catch(const elt::config::CiiConfigLimitOutOfBoundsException& e) {  
    std::cout << e.what() << std::endl;  
}  
}
```

Python checked functions

Listing 5-13 Python checked functions example

```
#Using the check functions  
  
# this has to be called  
checkedIntClass = elt.config.MdInt32.get_new_metadata_instance('')  
checkedIntClass.set_checked(True)  
checkedIntClass.set_min_limit(1)  
checkedIntClass.set_max_limit(200)  
  
try:  
    newSimpleInstance = MySimpleClass.get_new_node_instance()  
    newSimpleInstance.set_testStringValue("My test string")  
    #value in limits  
    newSimpleInstance.set_testIntValue(10, checkedIntClass)  
    #value out of limits  
    newSimpleInstance.set_testIntValue(202, checkedIntClass)  
    newSimpleInstance.set_testDoubleValue(12.0)  
except elt.config.CiiConfigLimitOutOfBoundsException as cix:  
    print(cix)
```



13.6 Additional information

13.6.1 Config tool

The configuration tool (config-tool) is used to generate classes and to deploy and undeploy the YAML defined configurations to/from the local DB JSON structure. The tool also supports class generation from remote DB.

The **config-tool** has the following usage pattern: **config-tool <operation> <options>**.

To use the predefined simple metadata classes (ConfInt32Std, ConfStringStd, etc.), the classes must be deployed in the local DB. The **config-tool init** operation supports this. To initialize the local DB use:

```
config-tool init
```

For remote DB the config-initES.sh must be run. Please refer to [2] for instructions.

Class generation

The tool supports programming language configuration and metadata class files generation. The files are generated from YAML files (optionally JSON). To generate classes the **generate** option must be used. The following options are supported:

- lang: cpp, java, python. Usage: -l <arg>. Default: java
- type: config, metadata. Usage: -t <arg>. Default: config
- input file. Usage: -i <arg>
- output directory. Usage: -o <arg>
- input file is in JSON. Usage: -j
- generate wscript files for python bindings. Usage: -ws

The argument “o” is mandatory, other arguments are optional. When generating metadata classes, the type “t” must be properly set. To generate the classes from the remote DB, run the tool without “i” argument.

It is not allowed to mix the metadata and configuration classes in one YAML file. As configuration classes change frequently and metadata classes will not change much, the metadata and configuration class generation are separated. The “t” argument forces the user to explicitly know that one is actually generating metadata classes.

The class generation tool has no support for checking the names of existing classes and field names of parent classes. It is up to the user, to properly check the existing fields of the parent classes. However, the tool will check for the reserved words for field names which also contain the field names of the predefined classes. Appendix D contains the list of reserved words.



Deployment and undeployment

The tool supports deployment and undeployment of instances files to/from local DB. To deploy the instances the **deploy** option must be used. The following options are supported:

- address: target configuration URI. Option: -a <arg>
- input file. Option: -i <arg>
- type: config, metadata. Option: -t <arg> Default: config
- input file is in JSON. Option: -j
- version. Option: -v <arg>

The arguments “a” and “i” are mandatory. When deploying metadata instances the type “t” must be properly set. If no version “v” is specified the instances are always deployed under version 1.

To undeploy instances, the **undeploy** option must be used. The following options are supported:

- address: target configuration URI. Option: -a <arg>
- instancenames: Metadata instance names to be deleted,

separated with “,”. Option: -n <args>

- type: config, metadata. Option: -t <arg> Default: config
- input file is in JSON. Option: -j
- version to be deleted. Option: -v <arg>

The version “v” must be explicitly specified (for safety reasons). To delete all versions, use -1.

The argument “n” must be used with type *metadata*. In this case the argument “a” is ignored.

To get a list of all available operations, run the:

```
config-tool -h
```

To get help on the chosen operation, run **config-tool <operation> -h**. Example:

```
config-tool undeploy -h
```

13.6.2 Generic types

Metadata classes partly support generic type inference. The generic type can be used in all the metadata classes, but the actual type used must be defined in the configuration class definition. There is no limit for the actual type to be defined before in the hierarchy of metadata. When the generic type is fixed, there is no need to define the type in the configuration class definition. Example of these classes are all classes with defined type (MdString, MdInt32). These classes have the actual type defined in the metadata class definition, so the type cannot be defined in the configuration class definition. To set the type of the metadata class to be generic, the **T** keyword must be used (no other



keywords are allowed). The keyword tells the metadata class definition that the actual type will be set by the configuration class. Using the defined generic types in the metadata class definition (e.g.: INT32) will determine the generic type of the class. If the type is fixed in the metadata class hierarchy, the *genericType* attribute of configuration class must be left out or set to *genericType*: NONE.

Listing 6-1 shows example of custom generic metadata class (MyCustomMetaClass) definition in YAML. Listing 6-2 shows use of the custom generic metadata class in configuration class YAML definition (MyClassWithCustomMeta). The type INT32 is set in the configuration class definition.

Listing 6-1 Generic type metadata class

```
configuration:
  metadataClassesConfiguration:
    classes:
      - name: MyCustomMetaClass
        parent: MdNumber
        __comment__: Custom metadata class
        genericType: T
        members:
          - name: metaOfBool
            type: BOOLEAN
          - name: metaOfString
            type: STRING
          - name: metaOfInt
            type: INT32
```

Listing 6-2 Usage of generic type metadata class

```
configuration:
  configClassesConfiguration:
    classes:
      - name: MyClassWithCustomMeta
        parent: CiiConfigClass
        __comment__: Class with custom metadata
        members:
          - name: stringPointWithCustomMeta
            type: MyCustomMetaClass
            genericType: INT32
```

CII Configuration provides an option to have generic type metadata classes and to set the type of the config point metadata class in the configuration class definition. This option can be used for classes, however the instances definition YAML files must have generic type specified for serializers to properly determine the type of data provided.

Types defined in instances must be the same as types defined in the classes in order for configuration to work properly. Listing 6-3 and Listing 6-4 show the examples of instances YAML definition files.

Listing 6-3 Config instance when generic type is used

```
metadata:
```

(continues on next page)



(continued from previous page)

```
instances:
- "@type": MyCustomMetaClass
  "@name": myCustomMDI
  "@genType": INT32
  comment: "String MDI with values"
  metaOfBool: "true"
  metaOfString: "The string metadata example"
  metaOfInt: "12"
```

Listing 6-4 Configuration instances with generic type YAML

```
config:
  instance:
    __comment__: Configuration for cryo cooler
    data:
      "@type": MyClassWithCustomMeta
      fields:
        stringPointWithCustomMeta:
          "@type": MyCustomMetaClass
          "@genType": INT32
          metadataInstance: myCustomMDI
          metadataInstanceVersion: 1
          value: 13
```

13.6.3 Java search example

Listing 6-5 shows the Java example for configuration search. The same API can also be used for CPP and Python. A different search language is used for local and remote DB:

- When searching in local database, a json path expression should be used: <https://github.com/json-path/JsonPath>
- When searching in remote database, an Elasticsearch query DSL should be used: <https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-query-string-query.html#query-string-syntax>

Searching with both languages is shown in the example.

Listing 6-5 Java search example

```
// remote search example
CiiConfigClient client = CiiConfigClient.getInstance();

// search or configuration with name 'temperatureSensor'
String configurationName = "temperatureSensor";
final String searchExpression = String.format("data.name:%s",
↪configurationName);
try {
```

(continues on next page)



(continued from previous page)

```
List<CiiTargetConfig> foundConfigurations = client
    .searchConfigRepo(ApiSearch.REMOTE, searchExpression);
} catch (CiiConfigNoTcException e) {
    // handle exception in case there was no such configuration
} catch (CiiConfigSearchException e) {
    // handle exception in case an error occurred while searching for
↪configuration
}

// local search example
CiiConfigClient client = CiiConfigClient.getInstance();

// search locally for configuration with name 'temperatureSensor'
String configurationName = "temperatureSensor";
final String search = String.format("$.?[?(@.name == '%s')]",
↪configurationName);
try {
    List<CiiTargetConfig> foundConfigurations = client.
↪searchConfigRepo(ApiSearch.LOCAL, search);
} catch (CiiConfigNoTcException e) {
    // handle exception in case there was no such configuration
} catch (CiiConfigSearchException e) {
    // handle exception in case an error occurred while searching for
↪configuration
}
```

13.6.4 Large data – binary files

The MdArray and MdBinary types support storing of large data. The CII Configuration will store all binary data (MdBinary) as file references. All the arrays (MdArray) exceeding 100kb (this limit is hard coded) will be stored as binary files and only references to files will be passed to predefined metadata field “file”. Data is stored as a raw data in big-endian order.

Configuration instance YAML definition supports deployment of the raw data files. To properly deploy the file, it must be placed in the subfolder named “files”. The name of the file must be set in the file field of the config point. In Listing 6-6 the *picture1* config point will use the file *MyBigData.raw* as file reference. When the instance is deployed the file will be copied to local DB or remote large data storage.

Listing 6-6 Configuration instance field with file

```
picture1:
  "@type": MdBinary
  metadataInstance: ConfBinaryStd
  metadataInstanceVersion: 1
  value:
```

(continues on next page)



(continued from previous page)

```
file: MyBigData.raw
```

13.6.5 Extending metadata classes with additional logic

Cii Configuration provides built-in check functions (5.3), but has no direct support for extending the metadata classes with additional logic. The only way to add additional check functions is to modify the source code of the CII Configuration. For such an action, deep knowledge of internals is needed. As it is not intended for the end user to modify the source code, here only some short guidelines are provided.

To add additional logic in Java, the logic has to be put into one of the existing metadata classes. The logic has to be programmed in a separate private method, and then added to existing setValue method. The recommended classes for additional logic are MdNumber, MdArray, MdString. The logic method can also be put in MdBase (as protected), but it has to be called in setValue method of the MdNumber, MdArray or MdString classes.

For CPP additional logic has to be added as a private method in the MdBase.hpp file. The custom logic private method then has to be added to the existing validate method. As the method must support different types of an input value, it has to be a templated method.

As the Python bindings to CPP are used, no additional work is needed for the Python.

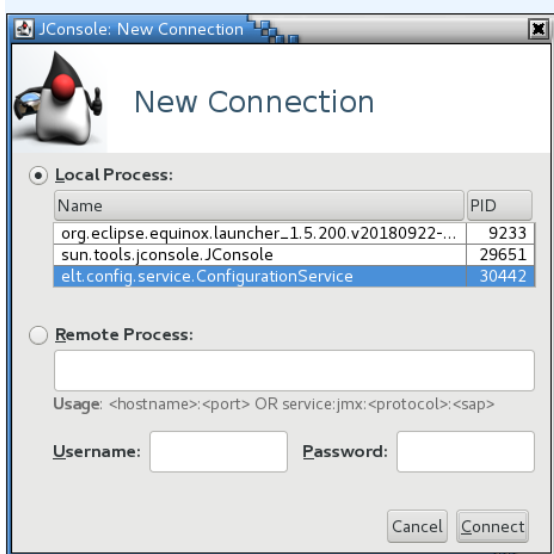
13.6.6 Accessing statistics using the JMX

Config service exposes basic statistic using the Java MBeans JMX interface. Statistics can be obtained using various JMX tools. The following example shows usage of jconsole. First start the jconsole with:

```
jconsole&
```

Connect to the Config service process:

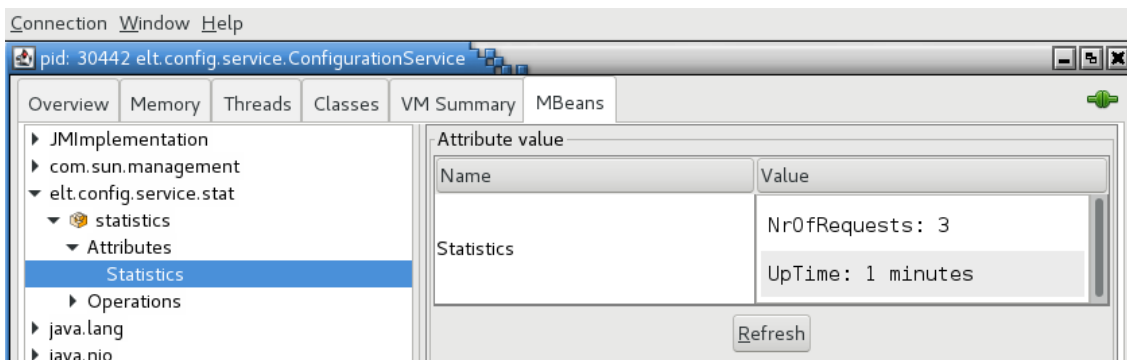
Figure 6-1 Jconsole JMX connection



After connection is established, the console with statistics is displayed. The console displays heap usage, CPU usage, number of threads and similar java statistics. The specific Config service statistics is exposed under the `elt.config.service.stat` namespace.

To obtain statistics click on the MBeans tab and browse the `elt.config.service.stat` namespace:

Figure 6-2 Jconsole configuration service statistics





13.7 GUI

13.7.1 General

Configuration GUI is a graphical editor application that can be used for browsing, querying and editing the configuration and metadata classes, target configurations and metadata instances. None of the configuration data is stored locally, everything is being sent to the remote configuration service. Config GUI is thus intended to work with the remote DB only.

Configuration GUI can be executed by issuing the following command from the shell:

```
config-gui
```

13.7.2 Configuration GUI settings

By default, the configuration GUI will look at **ciiconfservicehost** for config service and at **http://ciihdfshost:9870/webhdfs/v1/esoLs** for HDFS. For logging, default settings of CII logging will be used. Optionally, the executable takes 2 arguments:

- **path_to_application_config_file**: a path to the file that contains the settings for the Configuration GUI application. Usually, the configuration file for the Configuration GUI application is named **app-config.ini**. A detailed description of the application configuration file can be found in Appendix F.
- **path_to_log_config_file**: a path to the file that contains the settings for the log4cplus logging library. If this argument is not provided, the basic log4cplus configuration is used. Usually, the configuration file for the log4cplus library is named **log4cplus.properties**. Example value for the attribute: `../log4cplus.properties`.

To display the help, “-h” parameter shall be used.

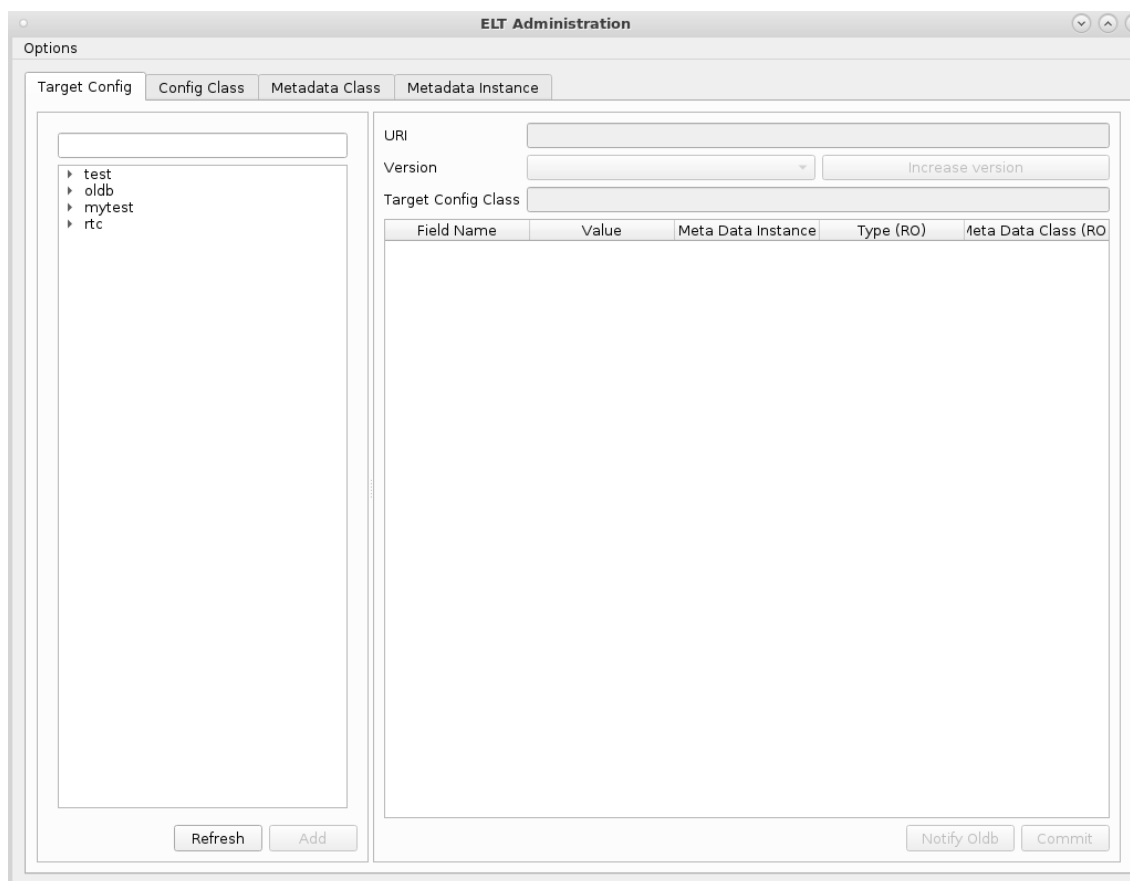
13.7.3 Main Window

The main window is displayed when we execute the Configuration GUI. The application screenshot is shown in Figure 7-1. As we can see from the figure, the functionality of the Configuration GUI is organized into 4 groups, each group having its own tab:

- **Target Config**: this tab is used for viewing and managing configurations. It is described in more detail in section 7.5.
- **Configuration class**: this tab is used for viewing and managing configuration classes. It is described in more detail in section 7.6.
- **Metadata Class**: this tab is used for viewing and managing metadata classes. It is described in more detail in section 7.7.
- **Metadata Instance**: this tab is used for viewing and managing metadata instances. It is described in more detail in section 7.8.



Figure 7-1 Main window.



Every tab is divided into 2 parts. The left part displays the list of entities (i.e. configurations, configuration classes, metadata instances, metadata classes) retrieved from the remote service and the button for adding new entities. In some cases, also the delete button is present. The user can select entities by clicking them. A list of entries can be refreshed by clicking on the **Refresh** button.

The right part is used for displaying the data about the selected entity: name, parent, list of attributes, etc. For some tabs, the right part of the GUI also contains the controls for selecting and managing entity versions.

Filtering the values

For easier searching of the entries, a search filter is implemented. The filter can be used for searching the classes and instances. The filter is presented as input box on each of the tabs (Figure 7-2). Filter can also be used in the dialogs (Figure 7-3). The user can input a part of the entire and all the entries whose name doesn't contain the string input are filtered out (the string comparison is case insensitive). In order to display all the entries again, the user needs to clear the filtering text.

Figure 7-2 Filter input box location on each tab

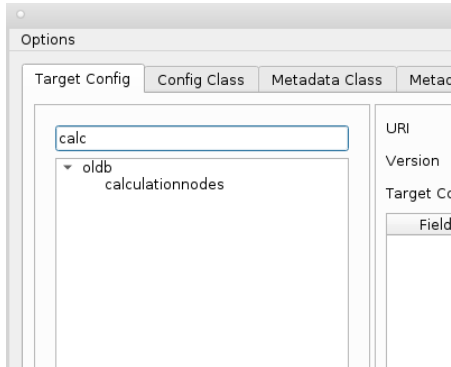
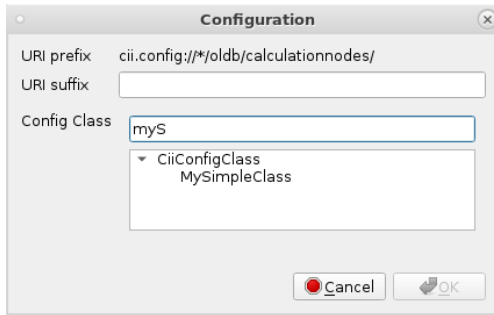


Figure 7-3 Filter input box location on a dialog



13.7.4 Menu

The Configuration GUI application only contains a single menu entry, the **Options**. This menu entry contains the **Write Enabled** and **Dangerous action** items for enabling the modifying of configurations. By default, the modification of configurations is disabled in Configuration GUI, so the user doesn't break them by mistake. Before being able to do any modifications, the user has to check the "Write Enabled" menu item. After **Write Enabled** is ticked, the GUI controls that are related to the modification of the configurations will be enabled. Enabling **Dangerous action** tick box will enable deletion of the metadata classes.

13.7.5 Target Config Tab

The Target Config tab enables the user to do the following actions:

- Navigate and view the target configurations in the remote database.
- Control configuration values – manage configuration control operations.
- Create target configurations on the desired URI endpoints.
- View revision history of target configurations and increase versions.
- Filter the navigation tree (section 7.3.1)



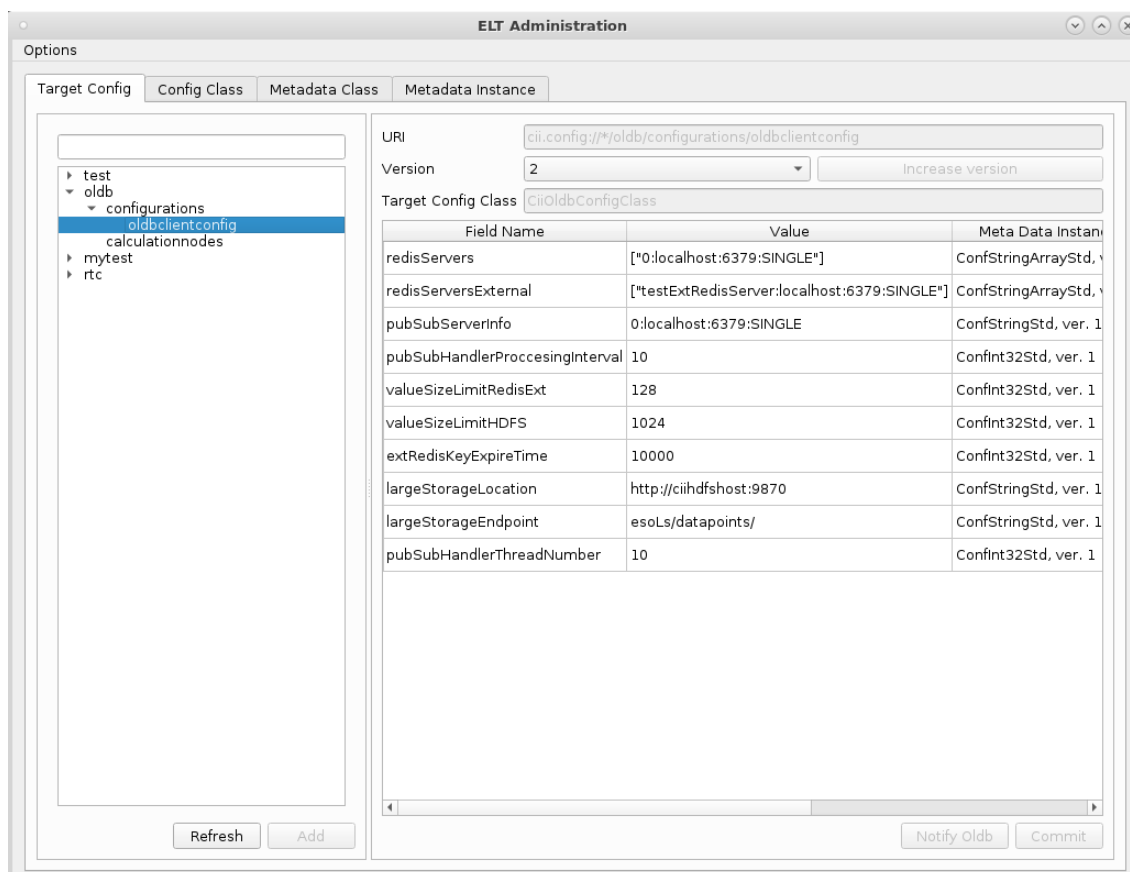
Selecting configuration

The user selects a configuration by clicking the leaf item in the configuration hierarchy tree (Figure 7-4). If the outer tree node is clicked or **Refresh** button is pressed, the configuration is deselected. The configurations are organized in the tree structure based on their URI paths (2.2.2).

When a configuration is selected, the following data is loaded in the right part of the GUI:

- URI: the unique URI that identifies the given configuration.
- Version: the most recent version number of the configuration,
- Target Configuration class: the assigned configuration class. It determines the list of attributes the configuration contains.
- The list of configuration attributes and their values: the list of attributes is comprised of the configuration class's attributes and the attributes of all its parents' classes.

Figure 7-4 Selecting configuration.





Adding new configuration entity

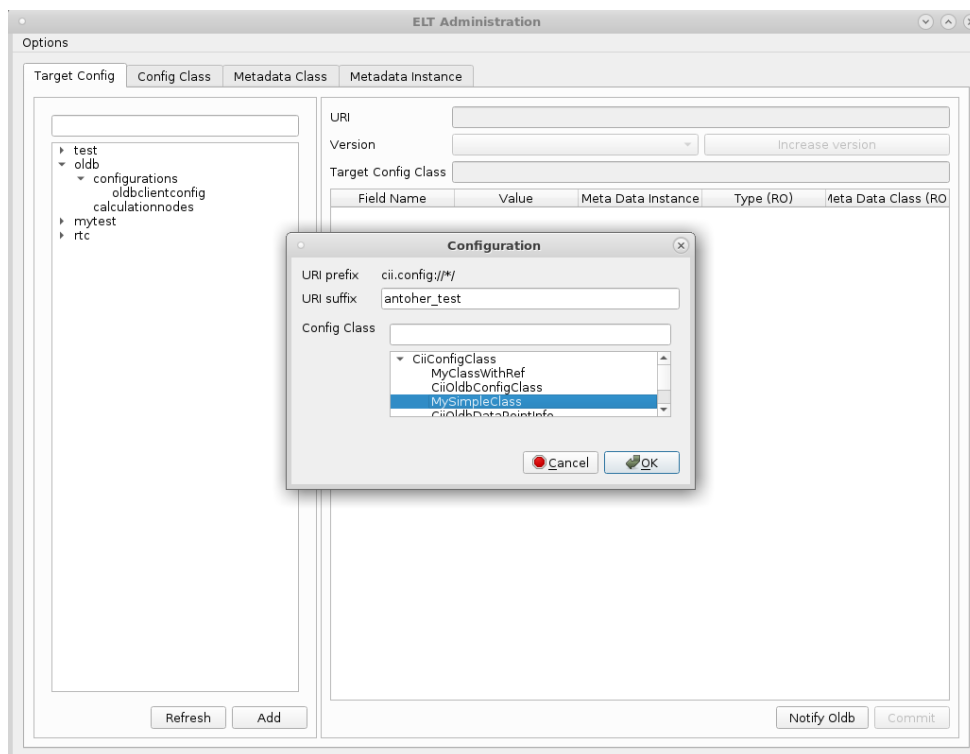
The adding of a new configuration entity is initiated by clicking the **Add** button below the hierarchy tree. The button is enabled only in the “Write Enabled” mode.

When the “Add” button is clicked, a dialog for inserting configuration data is displayed (Figure 7-5). The dialog contains the following fields:

- URI prefix of TC: a read-only field that contains the URI prefix (*cii.config:///**).
- URI suffix of TC: the input field where the user can input the URI suffix. The URI suffix actually defines the TC endpoint. The URI suffix can contain multiple “/” characters: in such case, multiple intermediate nodes will be created in the tree hierarchy. The final URI of the configuration is formed by concatenating the URI prefix and suffix. For example URI suffix */test_node/test_point* will create TC under the *cii.config://test_node/test_point**.
- Configuration class: here, the user selects the configuration class for the newly created configuration.

After the user clicks the **OK** button, the dialog is closed and a new configuration is created on the remote server.

Figure 7-5 Adding new configuration entity.





Setting configuration value

The setting of the value for the configuration attribute is performed using one of the input data dialogs (Figure 7-6) that are displayed when clicking the **Value** cell of the given attribute.

There are 3 types of input data dialogs:

- Primitive data input dialog: this dialog is used for entering the majority of data, except for vector and binary data. It contains a simple input field where the user can input the data value.
- Vector data input dialog: this dialog is used for inputting vector data and matrices (see organisation of matrices in 2.1.2). The dialog contains the input field where the user can input the content for smaller vectors. Compared to primitive data input dialog, this dialog enables input of larger data. The vector input data must be formatted as JSON arrays. Figure 7-7 shows the example of vector input data type.

Besides the input field, the dialog also contains the “Upload File” button which can be used for uploading the file containing the vector elements to the remote HDFS server. After the successful upload of the file, its URL is inserted in the input field. Below the upload button there is the “Download” button that can be used to download the file from the remote HDFS server.

- Binary data input dialog: The dialog contains the “Upload File” button which can be used for uploading the binary file to the large storage server. After the successful upload of the file, file name is inserted in the input field. Below the upload button there is the “Download” button that can be used to download the file from the large storage server.

After the user clicks the “OK” button, the dialog is closed, however the updates aren’t stored on the remote DB yet. They are only displayed in the attribute table. In order to trigger the updating on the remote server, the user needs to click the **Commit** button.

Note: When inputting vector data as a file, it is treated the same way as binary data input. All files must contain raw data in big-endian ordering.

Figure 7-6 Setting configuration attribute value.

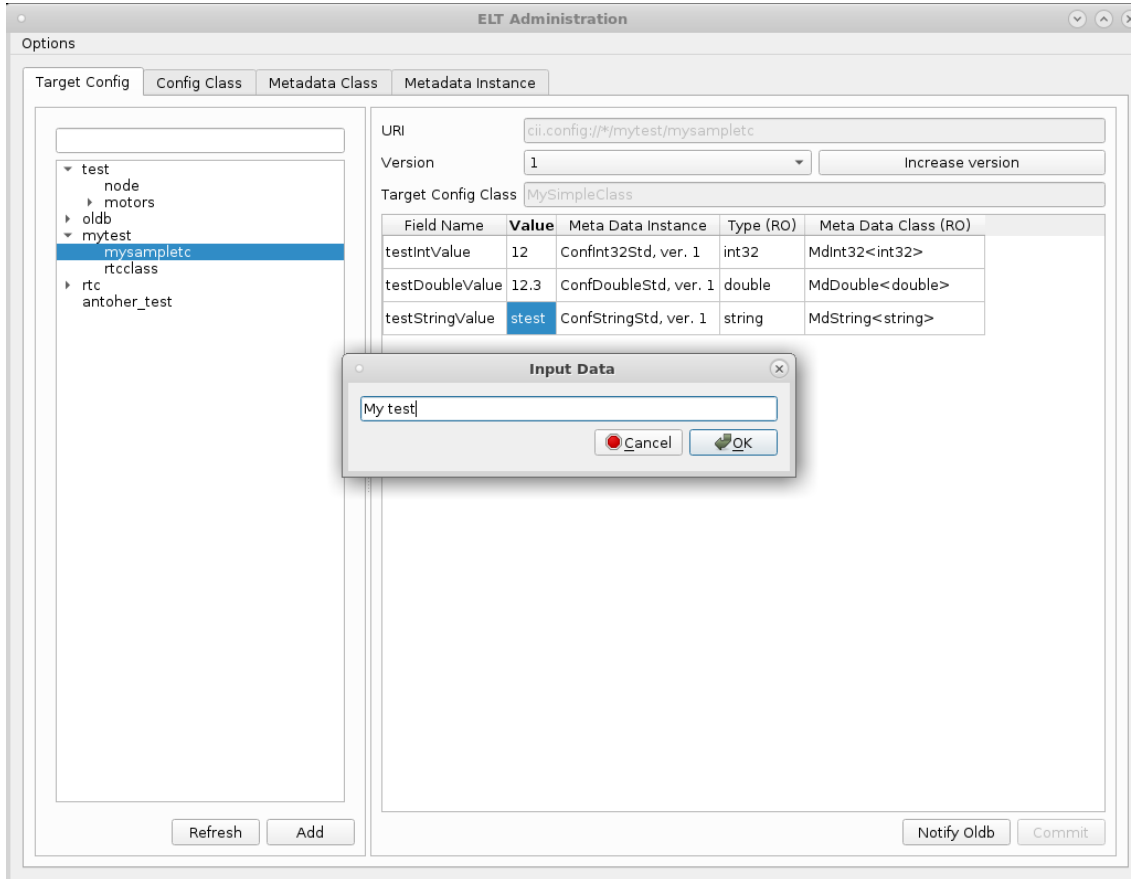
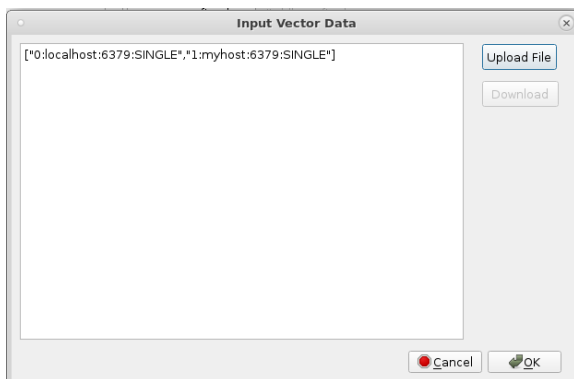


Figure 7-7 List box input for arrays.





Setting metadata instance and version

The setting of the metadata instance and a version for the configuration attribute is performed using the dialog that is displayed when clicking the **Meta Data Instance** cell of the given attribute. In the dialog, the user selects the metadata instance and its version. After the user clicks the “OK” button, the dialog is closed, however, the updates aren’t stored on the remote server yet. They are only displayed in the attribute table. In order to trigger the updating on the remote server, the user needs to click the **Commit** button.

Version management

The switching between different versions of the configuration is performed by using the **Version** combo box in the right section of the tab. When selecting a particular version number, the data of the corresponding configuration version will be displayed.

The user can also increase the version of the configuration by clicking the **Increase Version** button. When doing this, the most recent version gets copied on the remote server, the version number gets increased and the new latest version gets loaded in the GUI.

13.7.6 Configuration class Tab

The Configuration class tab enables the user to do the following actions:

- Create, delete, view and browse configuration classes.
- Define the structure and member fields of configuration classes.
- Adding additional member fields to the configuration class.
- Reference another configuration class (using the complex type).
- Filter the navigation tree (section 7.3.1)

Selecting configuration class

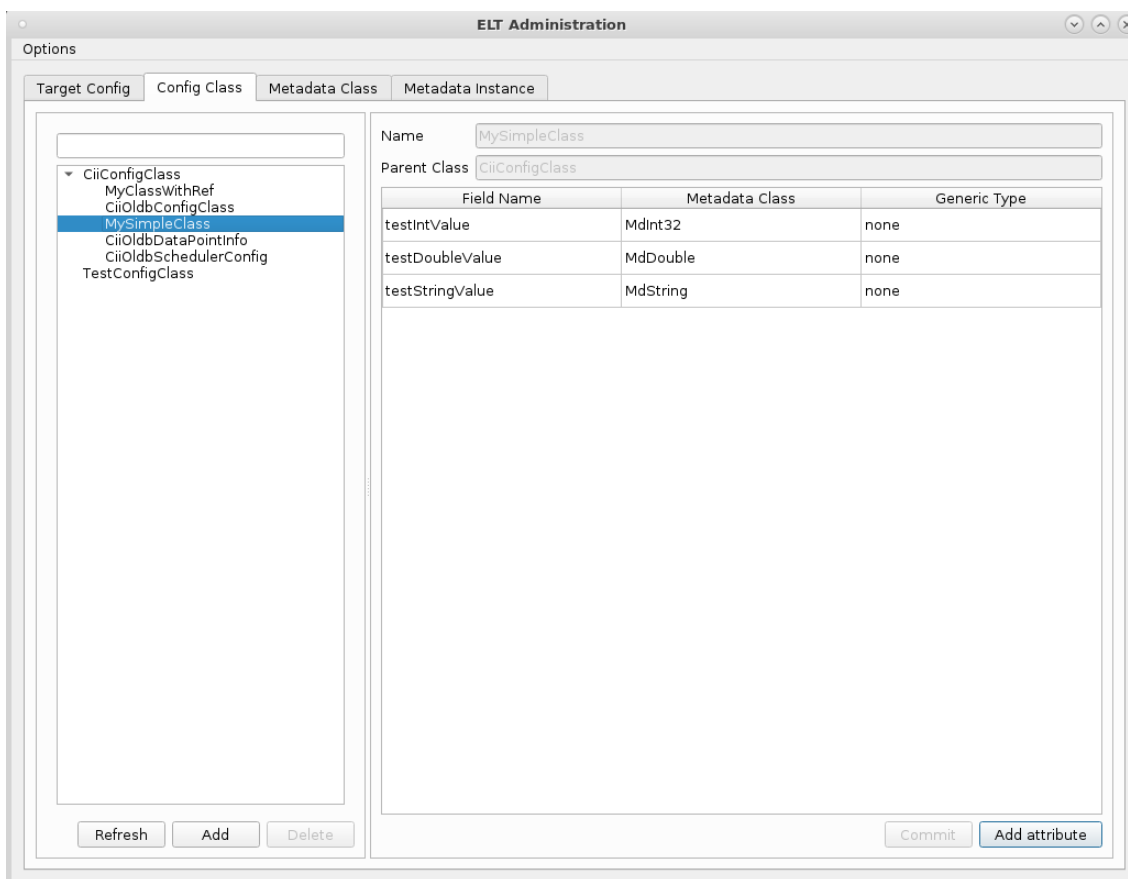
The user selects a configuration class by clicking the item in the configuration class hierarchy tree (Figure 7-8). The configuration classes are organized in the tree structure based on their “parent” attribute.

When a configuration class is selected, the following data is loaded in the right part of the GUI:

- Name: the name of the configuration class.
- Parent Class: the name of the configuration class’s parent class. This attribute is empty for root configuration classes.
- the list of configuration class attributes: the list of attributes is comprised only of the selected configuration class’s attributes. In order to see the configuration class’s parents’ attributes, the parent class has to be selected in the hierarchy tree.



Figure 7-8 Selecting configuration class.



Adding new configuration class entity

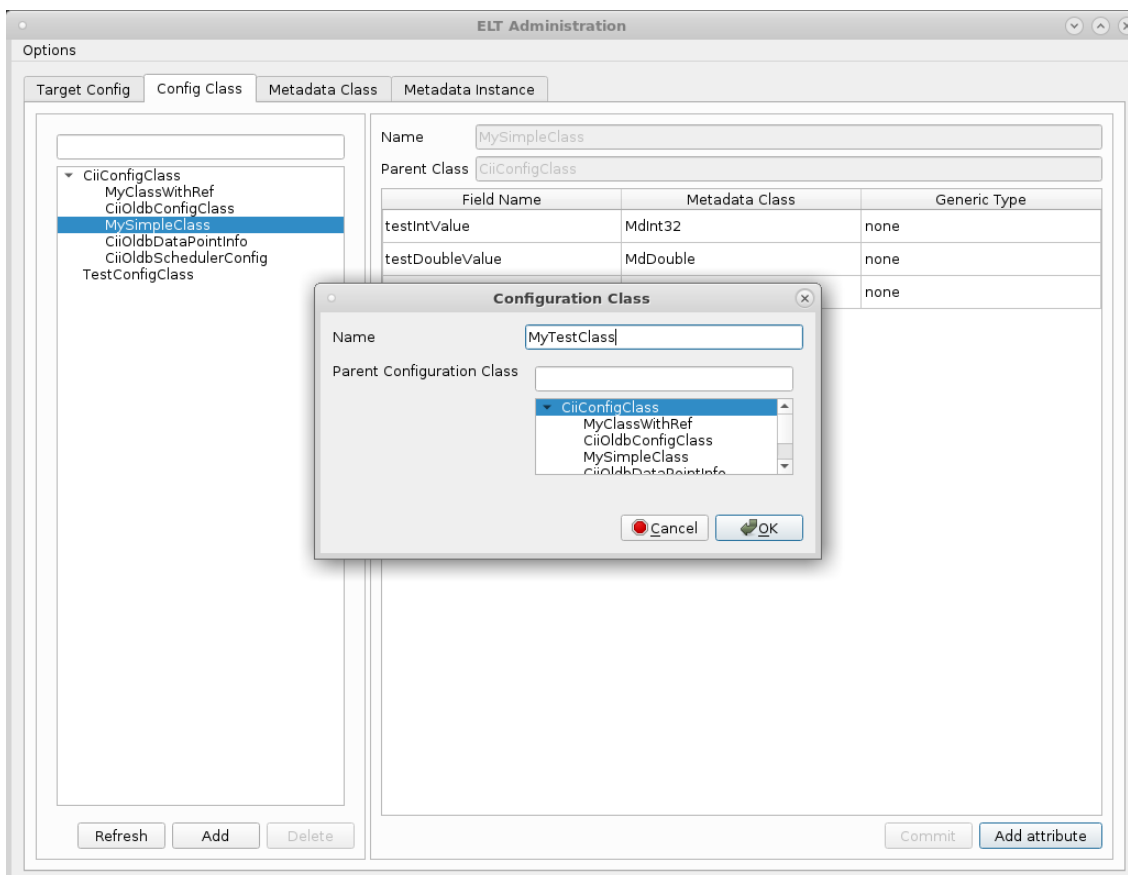
Adding a new configuration class entity is initiated by clicking the **Add** button below the hierarchy tree. The button is only enabled in the “Write Enabled” mode.

When the “Add” button is clicked, a dialog for inserting a configuration class data is displayed (Figure 7-9). The dialog contains the following fields:

- Name: the name of the newly-created configuration class.
- Parent Configuration Class: here, the user selects the parent configuration class for the newly-created configuration class.

After the user clicks the “OK” button, the dialog is closed and a new configuration class is created on the remote server.

Figure 7-9 Adding new configuration class entity.



Adding member fields to the configuration class

Adding of new member field (configuration class attribute) is initiated by clicking the **Add attribute** button below the attributes table. The button is only enabled in the “Write Enabled” mode.

When the “Add attribute” button is clicked, a dialog for inserting a configuration class attribute data is displayed (Figure 7-10). The dialog contains the following fields:

- Name: the name of the newly-created member field.
- Generic Type: here, the user selects generic type for the member field. This attribute is optional: for the non-generic member fields the user should select the “none” option from the combo box. When the type of the member field is a **reference to another** configuration class, this option must be set to **complex** (Figure 7-11).
- Metadata Class: here, the user selects the metadata class for the newly created member field.

After the user clicks the “OK” button, the dialog is closed. However, the updates aren’t stored on the remote server yet. They are only displayed in the attribute table. In order to trigger the updating on the remote server, the user needs to click the **Commit** button.

Figure 7-10 Adding a new configuration class attribute.

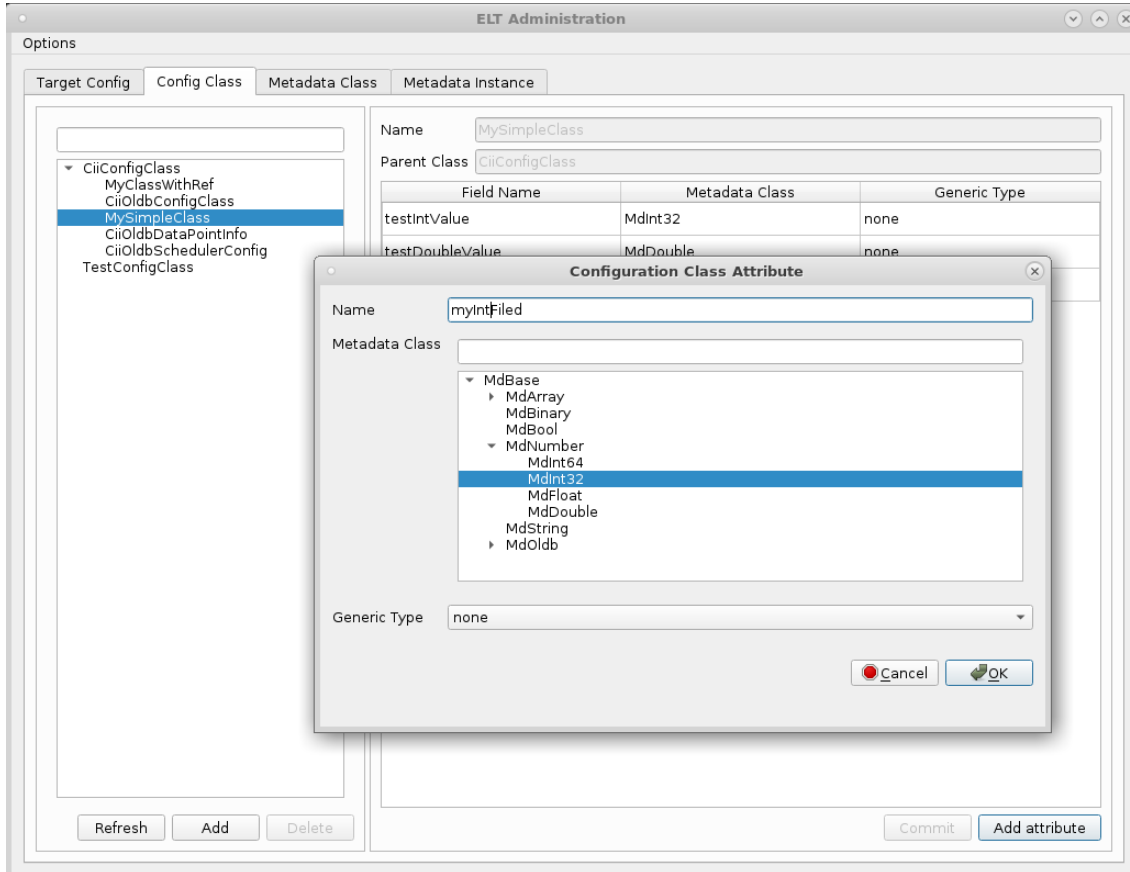
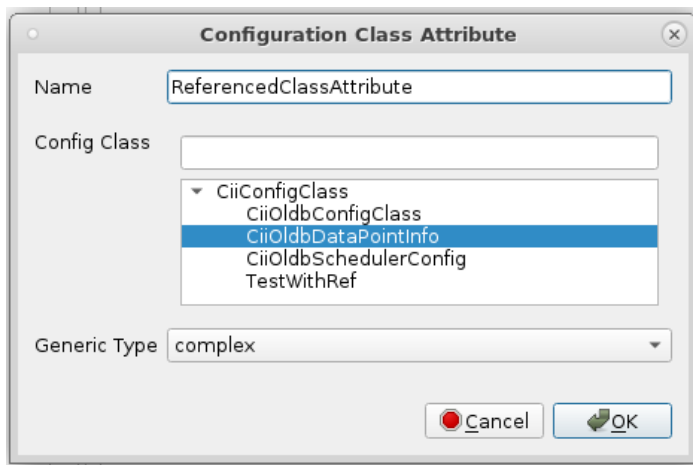


Figure 7-11 Adding referenced configuration class attribute.

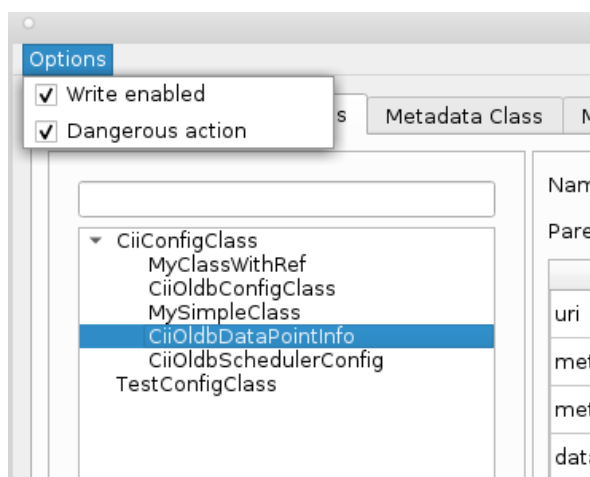




Delete configuration class

To delete the configuration class, first make sure that **Dangerous action** is ticked under the Options menu (Figure 7-12). After this is enabled, click on the desired configuration class and press **Delete** button. A warning is displayed. Click **OK** to finally delete the class.

Figure 7-12 Deleting configuration class attribute.



13.7.7 Metadata Class Tab

The Metadata Class tab enables the user to do the following actions:

- Create, delete, view and browse metadata classes.
- Define the structure and member fields of metadata classes.
- Filter the navigation tree (section 7.3.1)

Selecting metadata class

The user selects a metadata class by clicking the item in the metadata class hierarchy tree (Figure 7-13). The metadata classes are organized in the tree structure based on their “parent” attribute.

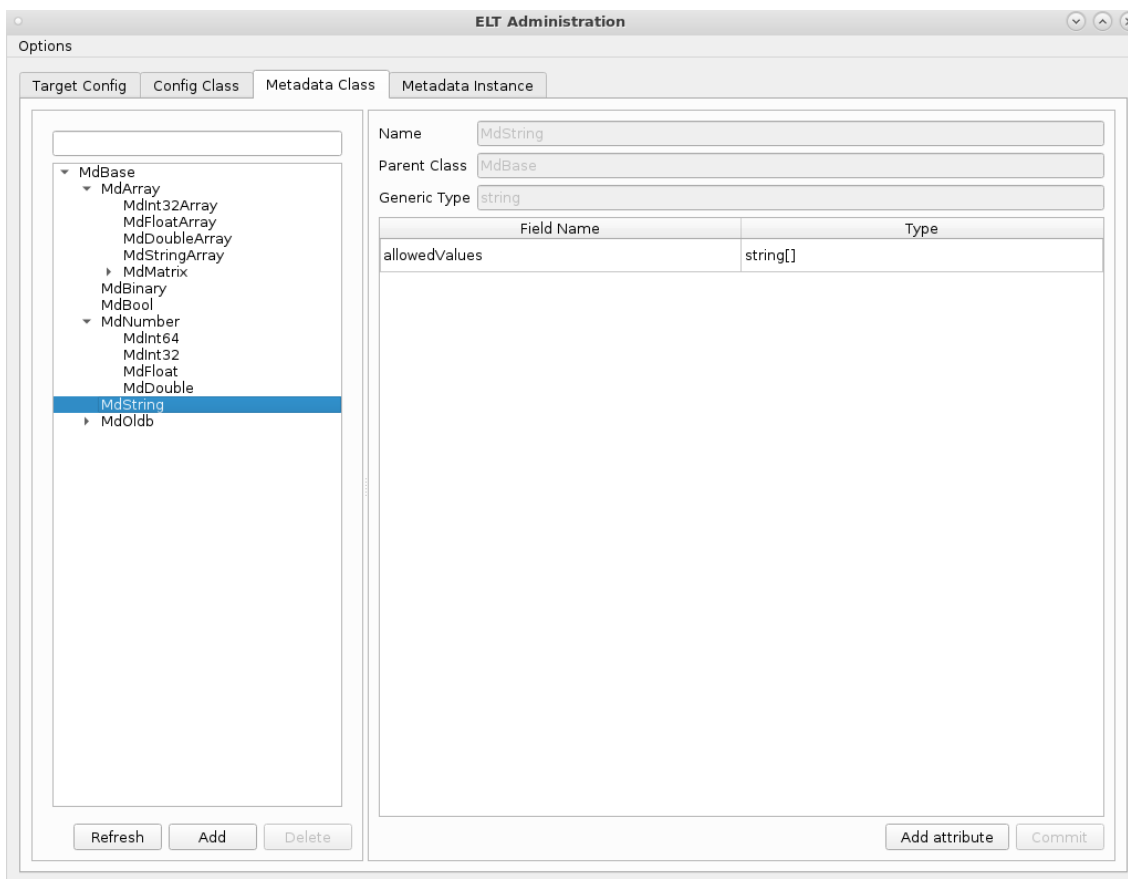
When a metadata class is selected, the following data is loaded in the right part of the GUI:

- Name: the name of the metadata class.
- Parent Class: the name of the metadata class’s parent class. This attribute is empty for root metadata classes.
- Generic Type: the generic type for the metadata class.



- the list of metadata class attributes: the list of attributes is comprised only of the selected metadata class's attributes. In order to see the metadata parents' attributes, the parent class has to be selected in the hierarchy tree.

Figure 7-13 Selecting metadata class.



Adding new metadata class entity

Adding of a new metadata class entity is initiated by clicking the **Add** button below the hierarchy tree. The button is only enabled in the “Write Enabled” mode.

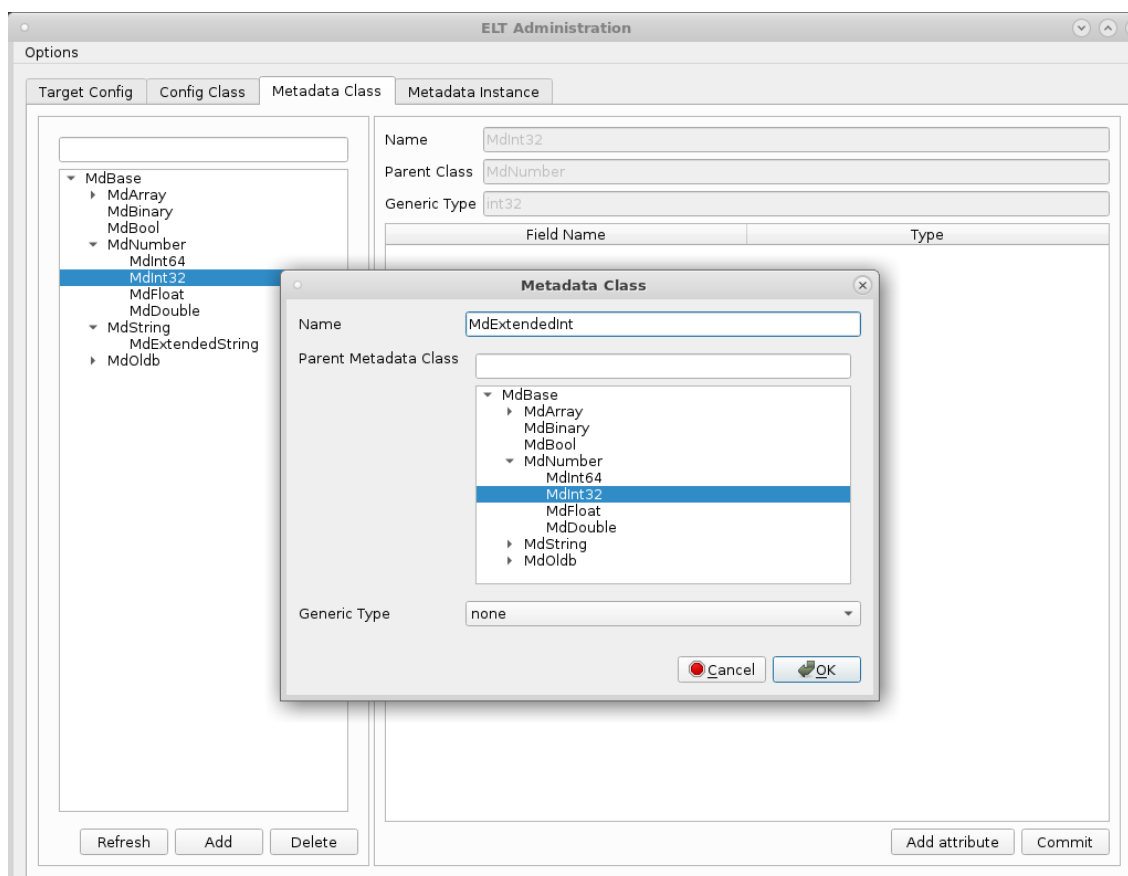
When the “Add” button is clicked, a dialog for inserting a metadata class data is displayed (Figure 7-14). The dialog contains the following fields:

- Name: the name of the newly-created metadata class.
- Parent Metadata Class: here, the user selects the parent metadata class for the newly-created metadata class.
- Generic Type: here, the user selects a generic type for the newly-created metadata class. This attribute is optional: for instances whose metadata class is non-generic the “none” value should be selected.



After the user clicks the “OK” button, the dialog is closed and a new metadata class is created on the remote server.

Figure 7-14 Adding metadata class entity.



Adding metadata class attribute

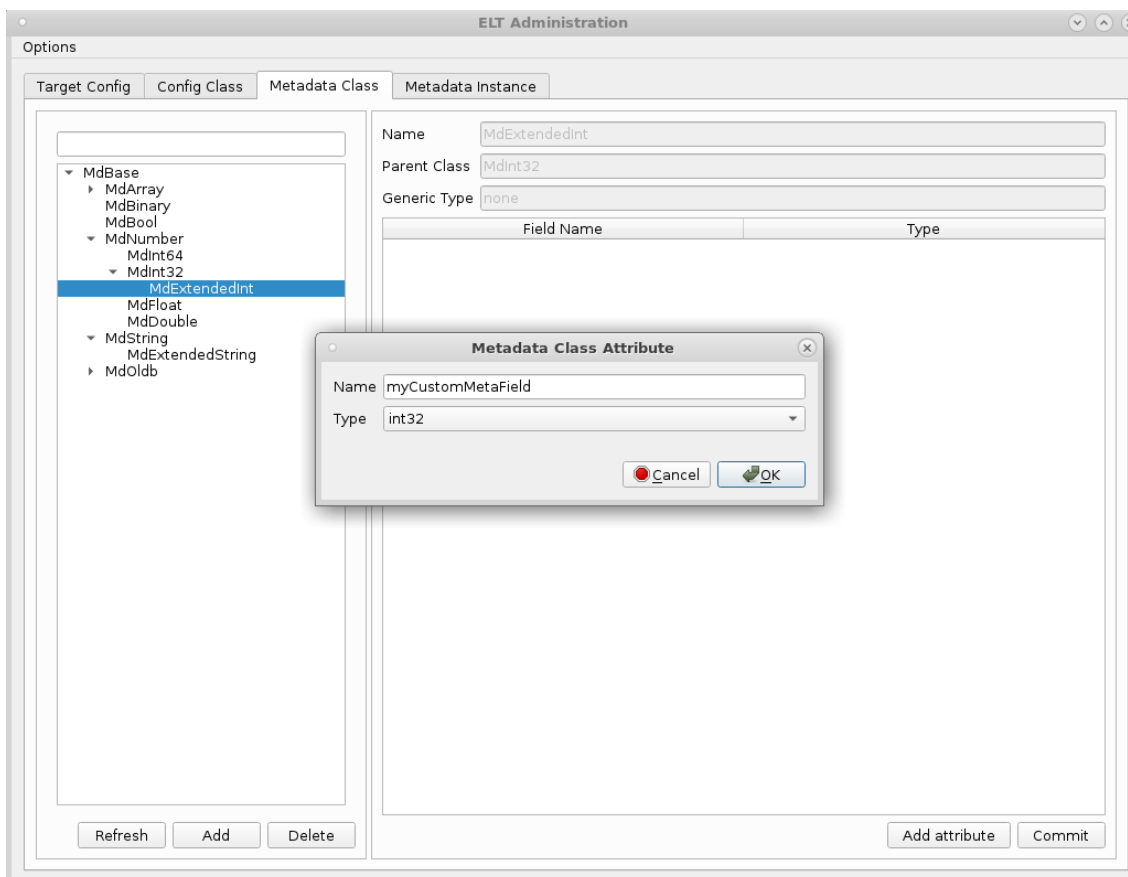
Adding of new metadata class attribute is initiated by clicking the **Add attribute** button below the attributes table. The button is only enabled in the “Write Enabled” mode.

When the “Add attribute” button is clicked, a dialog for inserting a metadata class attribute data is displayed (Figure 7-15). The dialog contains the following fields:

- Name: the name of the newly-created metadata class attribute.
- Type: the type of the newly-created metadata class attribute. The user is able to choose from a pre-defined list of basic types.

After the user clicks the “OK” button, the dialog is closed, however, the updates aren’t stored on the remote server yet. They are only displayed in the attribute table. In order to trigger the updating on the remote server, the user needs to click the **Commit** button.

Figure 7-15 Adding metadata class attribute.



Delete metadata class

To delete the metadata class, first make sure that **Dangerous action** is ticked under Options menu (Figure 7-12). After this is enabled, click on the desired metadata class and press **Delete** button. The warning is displayed. Click **OK** to finally delete the class.

13.7.8 Metadata Instance Tab

The Metadata Instance tab enables the user to do the following actions:

- Define instances of metadata classes.
- Assign metadata values to metadata member fields.
- Update metadata instances values.
- View revision history of metadata instances and increase versions.
- Filter the navigation list.



Selecting metadata instance

User selects a metadata instance by clicking the item in the metadata instance list (Figure 7-16).

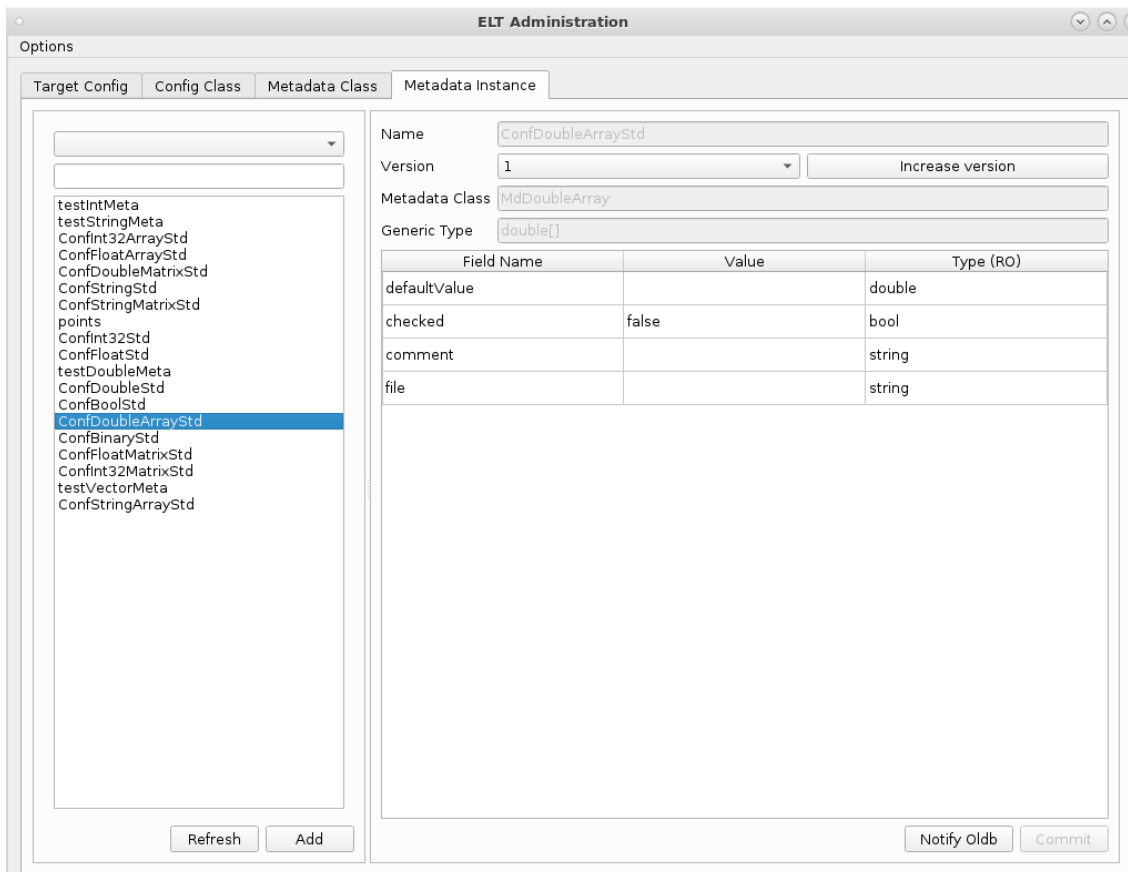
For the easier searching of metadata instances, search and dropdown filters are implemented:

- With the search filter, the user can input a part of the name and all the metadata instances whose name doesn't contain the string input are filtered out (the string comparison is case insensitive). In order to display all the metadata instances again, the user needs to clear the filtering text.
- In the dropdown filter there is a list of all metadata classes. When the user chooses one of the metadata classes from the dropdown filter only the metadata instances with that metadata class stay on the list, the others get filtered out. In order to display all the metadata instances again, the user needs to choose the empty option in the dropdown filter.

When a metadata instance is selected, the following data is loaded in the right part of the GUI:

- Name: the name of the metadata instance.
- Version: the most recent version of the metadata instance.
- Metadata Class: the metadata class assigned to a given metadata instance. It determines the list of attributes the metadata instance contains.
- Generic Type: the generic type for the metadata instance. The generic type for the metadata instance will be automatically retrieved from the metadata class generic type.
- The list of metadata instance attributes, their values and their types: the list of attributes is comprised of the metadata class's attributes and the attributes of all its parents' classes.

Figure 7-16 Selecting metadata instance.



Adding new metadata instance

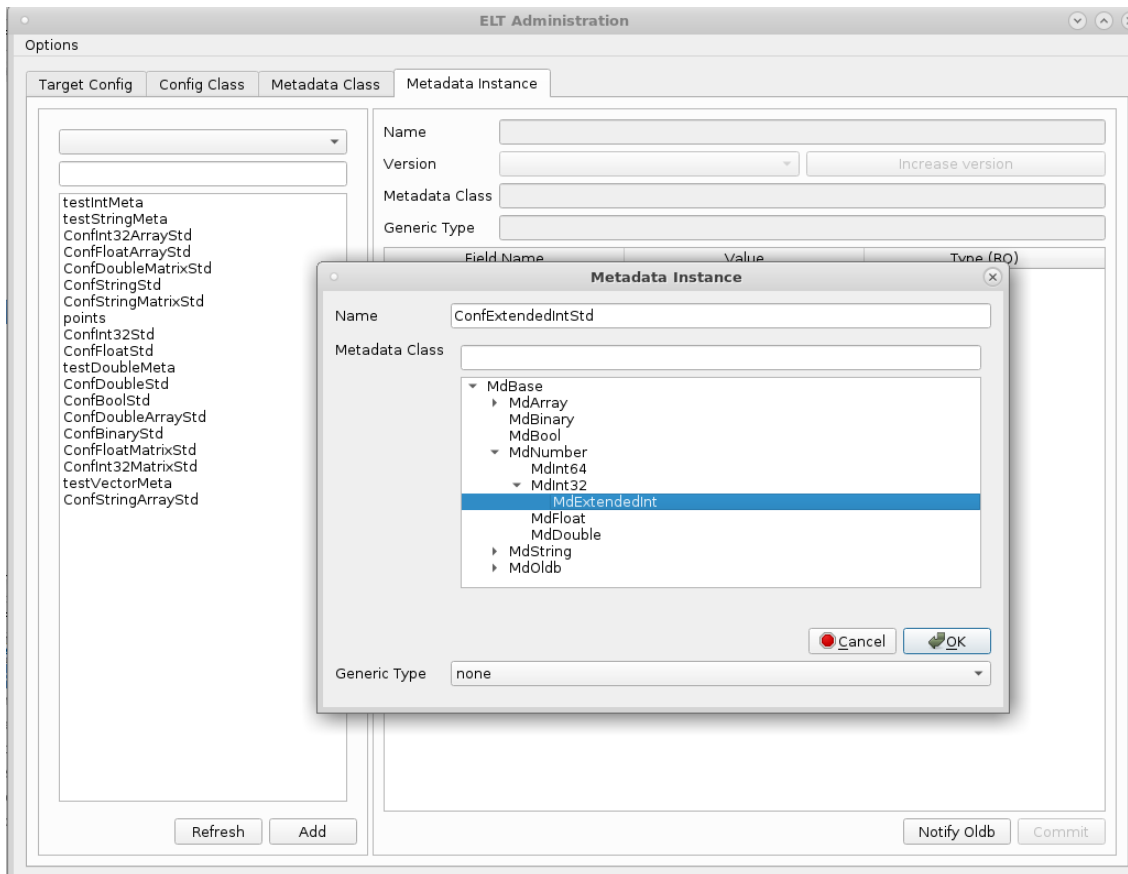
The adding of new metadata instance entity is initiated by clicking the **Add** button below the list of instances. The button is only enabled in the “Write Enabled” mode.

When the “Add” button is clicked, a dialog for inserting a metadata instance data is displayed (Figure 7-17). The dialog contains the following fields:

- Name: the name of the newly-created metadata instance.
- Metadata Class: here, the user selects a metadata class for the newly-created metadata instance.

After the user clicks the “OK” button, the dialog is closed and a new metadata instance is created on the remote server.

Figure 7-17 Adding new metadata instance.

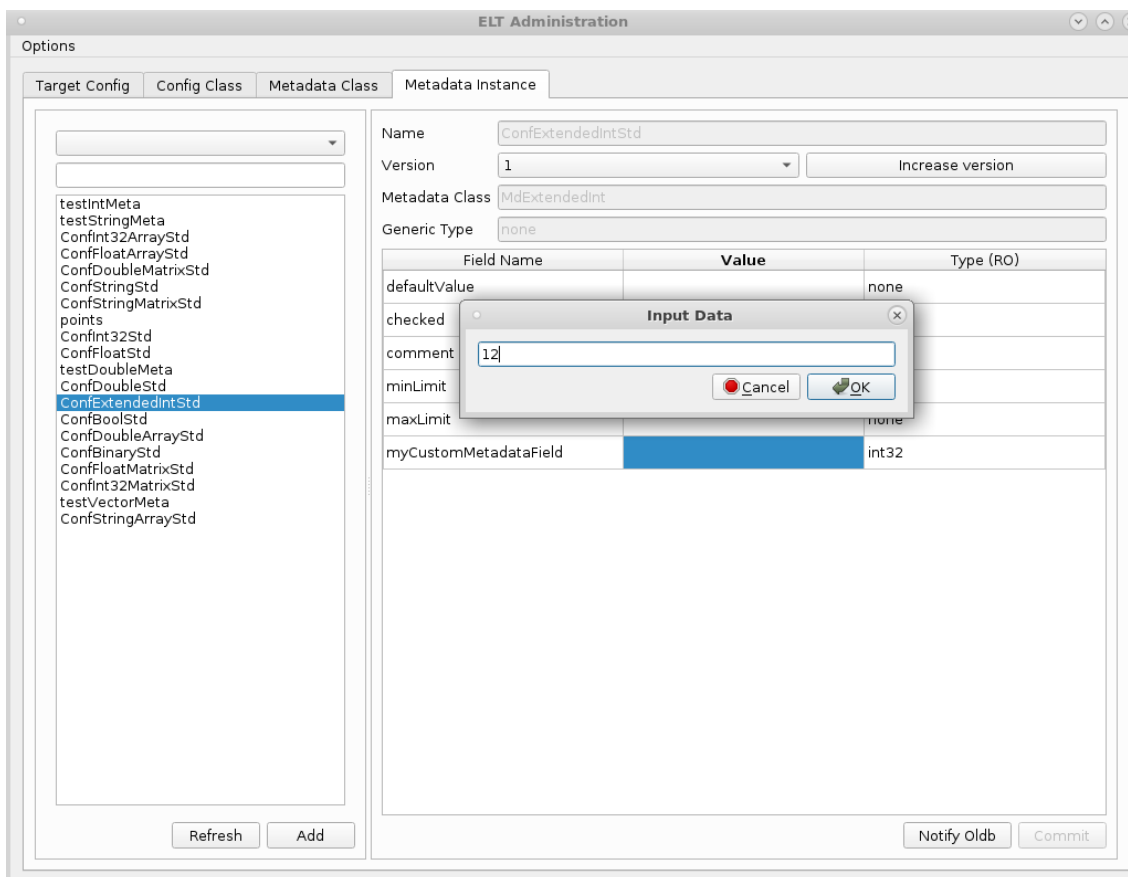


13.7.9 Setting metadata instance value

The setting of the value for the metadata instance attribute is performed using the primitive data input dialog (Figure 7-18) that is displayed when clicking the **Value** cell of the given attribute. This dialog contains a simple input field where the user can input the data value.

After the user clicks the “OK” button, the dialog is closed, however, the updates aren’t stored on the remote server, yet. They are only displayed in the attribute table. In order to trigger the updating on the remote server, the user needs to click the **Commit** button.

Figure 7-18 Setting metadata instance attribute value.



Version management

The switching between different versions of the metadata instances is performed by using the **Version** combo box in the right section of the tab. When selecting a particular version number, the data of the corresponding metadata instance version will be displayed.

The user can also increase the version of the metadata instance by clicking the **Increase Version** button. When doing this, the most recent version gets copied on the remote server, the version number gets increased and the new latest version gets loaded in the GUI.

13.7.10 Notify OLDB button

OLDB [1] internally uses the CII Configuration remote DB (target configuration and metadata instances) for its configuration. When configuration for OLDB changes, the calculation nodes and scheduler service must be notified about the change to use the changed values. When using OLDB client API, the services are automatically notified. But if the configuration settings are changed using the config GUI, the services are not automatically notified, as the configuration has no dependency on the OLDB. To overcome this problem, **Notify OLDB** button was added. This button can be used to notify OLDB services about the changes in the target configuration or metadata. Proper target



configuration or metadata instance must be selected and button **Notify OLDB** clicked to trigger the notification. The usage of this button only makes sense on the config points (URI `cii.config://*/oldb`) and metadata instances (prefixed with `Oldb`) used by OLDB.

Figure 7-19 Notify OLDB button





13.8 Config client API listing

Config client API provides methods for the CRUD operations of configuration and metadata instances. After the listing with methods description, the source code of specific programming language API is provided (A.1,A.2,A.3). Additional methods for searching the configurations are also present. The following methods are provided:

getInstance()
This is Java only method that returns instance of Config client API

saveTargetConfig(uri, rootNode)
Saves the configuration class to the specified URI location (cache, local, remote) and returns a version that was assigned to the saved configuration. Used programming language instance of configuration class must be set to rootNode parameter.

saveTargetConfigWithMetadata(uri, rootNode)
Saves the configuration class to the specified URI location (cache, local, remote) and returns a version that was assigned to the saved configuration. The method will also save all the metadata instances that are programmatically created and used in the programming language configuration class instance. Used programming language instance of configuration class must be set to rootNode parameter.

retrieveConfig(uri) retrieveConfig(uri, version)
Retrieves the chosen version of target configuration from the specified location (determined by the uri). If the version is set to -1, the last target configuration will be retrieved (configuration with the highest version that is stored at the specified location. The overload with no version will return the latest version.

updateConfig(uri, rootNode) updateConfig(uri, version, rootNode)
Updates target configuration at the specified URI with the data from the Configuration class set to the rootNode parameter. Overload with no version updates the last known configuration version.

deleteConfig(uri, version)
Deletes exact version of target config at defined URI location. If the version is set to -1, all configurations of the given URI will be deleted.



`saveMetadata(host, metadataInstance)`

Saves new metadata instance into the specified location. Parameter host determines the location under which metadata will be saved to. In Java the first part of the path that determines the location must be provided (`cii.config://cache/`, `cii.config://local/` or `cii.config://remote/`). The CPP and Python API uses string representation of location (“cache”, “local”, “remote”).

`retrieveMetadata(host, instanceName) | retrieveMetadata(host, instanceName, version)`

Retrieves specific version of the metadata from the chosen location saved under the chosen instanceName. Parameter host determines the location where metadata will be retrieved from. The overload without version will return the last saved version. In Java the first part of the path that determines the location must be provided (`cii.config://cache/`, `cii.config://local/` or `cii.config://remote/`). The CPP and Python API uses string representation of location (“cache”, “local”, “remote”).

`updateMetadata(host, instanceName, version)`

Updates specific version of the metadata with instanceName, under the chosen host location. Parameter host determines the location in which metadata will be updated. In Java the first part of the path that determines the location must be provided (`cii.config://cache/`, `cii.config://local/` or `cii.config://remote/`). The CPP and Python API uses string representation of location (“cache”, “local”, “remote”).

`deleteMetadata (host, instanceName, version)`

Deletes specific version of the metadata with instanceName, under the chosen host location. Parameter host determines the location where metadata will be deleted from In Java the first part of the path that determines the location must be provided (`cii.config://cache/`, `cii.config://local/` or `cii.config://remote/`). The CPP and Python API uses string representation of location (“cache”, “local”, “remote”).

`searchConfigRepo (host, query)`

Searches for target configurations that can be found using the chosen search expression. Depending on the location of the searches different search expressions should be used. For local searches a search expression should be written in JSON path format [3]. When searching in remote database a search expression should be written in elastic search query DSL format [4].
Host: In Java the first part of the path that determines the location must be provided (`cii.config://local/` or `cii.config://remote/`). The CPP and Python API uses string representation of location (“local”, “remote”).

`getChildren (uri)`

The method returns all the children config point URIs under the certain path. The method will return list of URIs with the config points.



setWriteEnabled (enabled) | isWriteEnabled()

Sets or checks the write enabled state flag. If write is enabled we can save, update and delete target configurations. If write is not enabled we can only retrieve the data.

13.8.1 Java API code

```
public static synchronized CiiConfigClient getInstance() throws_  
↳CiiConfigInitializationError { ... }  
  
public List<CiiConfigStatus> getConfigStatus(Uri uri) throws_  
↳CiiInvalidUriException { ... }  
  
public int saveTargetConfig(Uri uri, CiiConfigClass rootNode)  
    throws CiiConfigSaveException, CiiInvalidUriException,_  
↳CiiConfigWriteDisabledException { ... }  
  
public int saveTargetConfigWithMetadata(Uri uri, CiiConfigClass rootNode)  
    throws CiiConfigWriteDisabledException, CiiInvalidUriException,_  
↳CiiConfigSaveException { ... }  
  
public void updateConfig(Uri uri, CiiConfigClass rootNode)  
    throws CiiConfigNoTcException, CiiConfigUpdateException,_  
↳CiiInvalidUriException, CiiConfigWriteDisabledException { ... }  
  
public void updateConfig(Uri uri, int version, CiiConfigClass rootNode)  
    throws CiiConfigNoTcException, CiiConfigUpdateException,_  
↳CiiInvalidUriException, CiiConfigWriteDisabledException { ... }  
  
public CiiTargetConfig retrieveConfig(Uri uri)  
    throws CiiConfigNoTcException, CiiInvalidUriException { ... }  
  
public CiiTargetConfig retrieveConfig(Uri uri, int version)  
    throws CiiConfigNoTcException, CiiInvalidUriException { ... }  
  
public void deleteConfig(Uri uri, int version)  
    throws CiiConfigDeleteException, CiiInvalidUriException,_  
↳CiiConfigWriteDisabledException { ... }  
  
public List<CiiTargetConfig> searchConfigRepo(ApiSearch searchPath, String_  
↳searchExp)  
    throws CiiConfigNoTcException, CiiConfigSearchException { ... }  
  
public List<String> remoteIndexSearch(String index, String query)  
    throws CiiConfigSearchException { ... }  
  
public <T extends MdBase> T retrieveMetadata(Uri searchLocation,
```

(continues on next page)



(continued from previous page)

```
String metadataInstanceName, int version, Class<T> clazz)
    throws IOException, CiiInvalidURIException { ... }
public <T extends MdBase> int saveMetadata(Uri saveLocation, T metadataInstance)
    throws CiiConfigSaveException, CiiInvalidURIException,
↳CiiConfigWriteDisabledException { ... }

public void deleteMetadata(Uri deleteLocation, String metadataInstanceName, int
↳version)
    throws CiiConfigWriteDisabledException, CiiInvalidURIException,
↳CiiConfigDeleteException { ... }

public List<Uri> getChildren(Uri uri) throws CiiInvalidURIException,
↳CiiConfigSearchException { ... }

public void setWriteEnabled(boolean writeEnabled) { ... }

public boolean isWriteEnabled() { ... }
```

13.8.2 CPP API code

```
static int saveTargetConfig(const elt::mal::Uri& uri, const CiiConfigClass& root_
↳node);

static int saveTargetConfigWithMetadata(const elt::mal::Uri& uri, CiiConfigClass&
↳root_node);

static std::list<std::shared_ptr<CiiDataPointMetadataBase>>
↳getTargetConfigMetadata(CiiConfigurationBase& root_node);

static int saveMetadata(const std::string& host, CiiDataPointMetadataBase&
↳metadata);

static std::shared_ptr<CiiConfigClass> retrieveConfig(
    const elt::mal::Uri& uri, int version = -1);
static std::shared_ptr<CiiDataPointMetadataBase> retrieveMetadata(
    const std::string& host, const std::string& instance_name, int
↳version = -1);

static void updateConfig (const elt::mal::Uri& uri, int version,
    const CiiConfigClass& targetConfig);

static void updateMetadata(const std::string& host, const
↳CiiDataPointMetadataBase& metadata);

static void updateMetadata(const std::string& host, int version, const
↳CiiDataPointMetadataBase& metadata);
```

(continues on next page)



(continued from previous page)

```
static void deleteConfig (const elt::mal::Uri& uri, int version = -1);

static void deleteMetadata (const std::string& host, const std::string& instance_
↳name, int version = -1);

template <typename T>
static std::shared_ptr<T> getConfigData(const elt::mal::Uri& uri, int version = -
↳1) {
    return std::dynamic_pointer_cast<T>(retrieveConfig(uri, version));
}

template <typename T>
static std::shared_ptr<T> getMetadata(const std::string& host, const std::string&
↳ instance_name, int version = -1) {
    return std::dynamic_pointer_cast<T>(retrieveMetadata(host, instance_name,
↳version));
}

static const std::list<elt::mal::Uri> getChildren(const elt::mal::Uri& parent_
↳uri);

static const std::list<std::shared_ptr<CiiConfigClass>> searchConfigRepo(const
↳std::string& host, const std::string& query);

static void setWriteEnabled(bool enabled);

static bool isWriteEnabled();
```

13.8.3 Python bindings code

```
py::class_<elt::config::CiiConfigClient, std::shared_ptr
↳<elt::config::CiiConfigClient>>(m, "CiiConfigClient",
    "Configuration Client API")
    .def_static("save_target_config", &
↳elt::config::CiiConfigClient::saveTargetConfig,
    "Save a new version of configuration class")
    .def_static("save_target_config_with_metadata", &
↳elt::config::CiiConfigClient::saveTargetConfigWithMetadata,
    "Save configuration node and the metadata object it references")
    .def_static("retrieve_config", &elt::config::CiiConfigClient::retrieveConfig,
↳"Retrieve a configuration object previously saved via_
↳saveTargetConfig")
    .def_static("retrieve_config",
↳[] (const ::elt::mal::Uri &uri) { return
↳elt::config::CiiConfigClient::retrieveConfig(uri, -1); },
    "Retrieve a configuration object previously saved via_
↳saveTargetConfig")
```

(continues on next page)



(continued from previous page)

```
.def_static("update_config", &elt::config::CiiConfigClient::updateConfig,
            "Overwrite previously saved configuration or save it using a
↪specific version")
.def_static("delete_config", &elt::config::CiiConfigClient::deleteConfig,
            "Delete a configuration object")
.def_static("delete_config", [](const elt::mal::Uri &uri) {
    elt::config::CiiConfigClient::deleteConfig(uri);
}, "Delete a configuration object")
.def_static("get_config_data", &elt::config::CiiConfigClient::getConfigData
↪<elt::config::CiiConfigClass>,
            "Retrieve a configuration object")
.def_static("save_metadata", &elt::config::CiiConfigClient::saveMetadata,
            "Save a new version of a metadata instance")
.def_static("update_metadata", &elt::config::CiiConfigClient::updateMetadata,
            "Retrieve a metadata object previously saved via saveMetadata()")
.def_static("delete_metadata", &elt::config::CiiConfigClient::deleteMetadata,
            "Delete a metadata object")
.def_static("delete_metadata", [](const std::string &host, const std::string&
↪instance_name) {
    ::elt::config::CiiConfigClient::deleteMetadata(host, instance_name);
}, "Delete a metadata object")
.def_static("retrieve_metadata", &
↪elt::config::CiiConfigClient::retrieveMetadata,
            "Retrieve a metadata object previously saved via saveMetadata()")
.def_static("retrieve_metadata", [](const std::string &host, const std::string&
↪&instance_name) {
    return elt::config::CiiConfigClient::retrieveMetadata(host, instance_name);
}, "Retrieve a metadata object previously saved via saveMetaData")
.def_static("search_config_repo", &
↪elt::config::CiiConfigClient::searchConfigRepo,
            "Search a data provider for specific configurations")
.def_static("set_write_enabled", &
↪elt::config::CiiConfigClient::setWriteEnabled,
            "Enables or disables the possibility of writing into the
↪configuration databases")
.def_static("is_write_enabled", &elt::config::CiiConfigClient::isWriteEnabled,
            "Returns True if the API allows writing into configuration
↪databases, false otherwise")
.def_static("get_children", &elt::config::CiiConfigClient::getChildren,
            "Query children nodes")
```



13.9 Config point value type to language types mapping

Below is the mapping table between config point types and data types of all programming languages that the configuration can use.

Table 7-1 Config point types mapping table

ELT CII BASIC TYPE	Java	CPP	Python
INT32	Integer	std::int32_t	int
INT64	Long	std::int64_t	int
UINT8	Integer	std::uint8_t	int
UINT16	Long	std::uint16_t	int
UINT32	Long	std::uint32_t	int
SINGLE	Float	float	float
DOUBLE	Double	double	float
BOOLEAN	Boolean	bool	bool
STRING	String	std::string	str
BINARY	byte[]	std::vector<std::uint8_t>	bytearray



13.10 Default metadata instance mapping

Table 7-2 shows the default metadata instance and class type mappings. The columns "metadata instance @genType" and "class genericType" are CII Basic types. These columns are used in YAML class definition for metadata classes and metadata instances. The fields are described in Appendix E.

Table 7-2 Default metadata mappings

Table with 6 columns: Metadata class name, Metadata instance @genType, Class generic-Type, Is Array, Default metadata instance. Rows include mappings for MdInt32, MdInt64, MdFloat, MdDouble, MdBool, MdString, MdStringArray, MdInt32Array, MdFloatArray, MdDoubleArray, MdInt32Matrix, MdFloatMatrix, MdDoubleMatrix, and MdBinary.



13.11 Class definition reserved words

Table 7-3 shows list of words that cannot be used as field names in config or metadata class definitions.

Table 7-3 List of words not allowed for config and metadata class names

Field name
name
uri
defaultValue
value
comment
ciiType
genType
metadataChanged
minLimit
maxLimit
allowedValues
size
field
file
genericType
parent
__comment__
isRef



13.12 JSON/YAML Schema

Schemas are defined for classes and instances. All local and remote DB JSON files must be in line with the defined schemas. Target configuration and instances schemas are used for value bindings. Class schemas are used for class generation. For local and remote DB, YAML structures are translated into JSON structures with nested elements.

The following schema descriptions are indented in the same manner as the actual YAML files. The indentation follows the YAML files indentation structure.

All class schemas start with the starting structure, which is presented only here and not listed in the subchapters. This structure contains the following elements: "field: type, required/optional; description". The definition is as follows:

- **configuration**: string, required; defines the root for all configuration schemas
 - **configClassesConfiguration** or **metadataClassesConfiguration** string, required; defines the type of configuration schema.

13.12.1 Configuration class schema

Configuration class schema (field: type, required/optional; description):

- **classes**: array, required; List of classes to be generated.
 - **name**: string, required; Name of the configuration class.
 - **parent**: string, required; Parent configuration class name.
 - **__comment__**: string, optional; Class comment. This comment is changed to programming language specific comment on bindings generation.
 - **members**: array, required; member fields of classes (config points)
 - * **name**: string, required; config point name.
 - * **type**: string, required; config point metadata class type.
 - * **genericType**: string, optional; generic type of config point (CII data type)
 - * **isRef**: boolean, optional; True if this member is a reference to another class.



13.12.2 Metadata class schema

Metadata class schema (field: type, required; description):

- **classes**: array, required; List of metadata classes to be generated.
 - **name**: string, required; Name of the metadata class.
 - **parent**: string, required; Parent metadata class name.
 - **__comment__**: string, optional; Metadata class comment. This comment is changed to programming language specific comment on bindings generation.
 - **members**: array, required; member fields of metadata classes
 - * **name**: string, required; metadata member field name.
 - * **type**: string, required; CII data type of metadata class.

13.12.3 Configuration instance schema

Target configuration (configuration instance) schema has the following starting structure:

- **config**: string, required; Root node.
 - **instance**: string, required; Definition of config instance.

The starting structure is followed by the schema structure for the details. The details schema is child node of **instance** (field: type, required/optional; description):

- **__comment__**: string, optional; YAML/JSON comment for the target configuration.
- **data**: Object, required; Data object.
 - **@type**: string, required; Configuration class type.
 - **fields**: string, required; List of all members that have the same member name and type as defined in the corresponding class definition. Member name and value pairs are added to YAML (“field”: “value”):
 - * **@type**: string, required; Metadata class type.
 - * **metadataInstance**: string, required; Metadata instance name.
 - * **metadataInstanceVersion**: string, optional; Version of metadata instance. Last existing version is used if the version is not specified.
 - * **value**: string | number | boolean | array, required; the actual value of config point.
 - * **@genType**: string, optional; generic type used in the instance.

To address the fields of referenced classes, special name rule is used. The name is a concatenation of reference instance name followed by the reference field name, separated by a dot “.” character is used. The same logic is used when one reference class references another class.



Examples:

1. Configuration class contains referenced adFlowClass. The “adFlow1.sensorsSize” will address the sensorsSize field in the adFlow1 referenced class instance.
2. Configuration class contains referenced adFlowClass, which references currentFlowClass. The “adFlow1.currFlow1.running” will address the running field in the currFlow1 referenced class instance, which is referenced from adFlow1.

13.12.4 Metadata instances schema

Metadata instances schema has the following starting structure:

- **metadata:** string, required; Root node.
 - **instances:** string, required; Definition of metadata instance.

The starting structure is followed by the schema structure for the details. The details are child nodes of **instances** node: (field: type, required; description):

- **@name:** string, required; Name of metadata instance.
- **@type:** string, required; Metadata class of metadata instance.
- **@genType:** string, required; Generic type used by the instance. This type must be the same as the defined metadata class type.
- List of all member names that are in line with metadata class. Member name and value pair are added to YAML (“field”: “value”):
 - **field:** string, required; metadata class field name.
 - **value:** string | number | boolean | array, required; the actual value of metadata field.



13.13 GUI configuration file

The configuration file for the Configuration GUI application contains the information needed to connect to the remote configuration service (i.e. the data source) and file storage (i.e. HDFS service). The configuration file contains 4 sections:

- **DataSource**: this section contains all the data needed to connect to the primary configuration service. The **hostname** attribute provides the address of the service (in a form of an IP address or a host name) and the **port** attribute provides a port on which the service listens for incoming connections.
- **DataSourceBackup**: this section contains the same parameters as **DataSource**, only the values are set to the ones of the backup configuration service. In case you don't have a backup configuration service, insert the values of the primary configuration service into this section.
- **HdfsService**: this section contains only the **urlPrefix** attribute. This prefix is used when uploading files to the HDFS service. This is a REST endpoint to which Configuration GUI sends files using the HTTP POST calls.
- **HdfsServiceBackup**: this section contains the same parameters as **HdfsService**, only the values are set to the ones of the backup file storage. In case you don't have backup file storage, insert the values of the primary file storage into this section.

The following is an example of the GUI configuration file:

```
# configuration for a config gui
[DataSource]
hostname = ciiconfservicehost
port = 9116

[DataSourceBackup]
hostname = ciiconfservicehost2
port = 9116

[HdfsService]
urlPrefix = http://ciihdfshost:9870/webhdfs/v1/esoLs

[HdfsServiceBackup]
urlPrefix = http://ciihdfshost:9871/webhdfs/v1/esoLs
```



13.14 Configuration client settings

Configuration client API needs some internal setting to work properly. These settings are key/value pairs. Default settings are a part of the client but can be overridden. The following settings are defined:

- **configLocalDbRoot** – root folder where all the local DB JSON files are stored.
- **configServiceAddress** – URI address and port. E.g : `zpb.rr://ciiconfservicehost:9116/`
- **configServiceBackupAddress** - (Optional) URI address and port. E.g.: `zpb.rr://ciiconfservicehost2:9116/`
- **largeStorageAddress** - URI address and port. E.g.: `http://ciihdfshost:9087`
- **largeStorageBackupAddress** - (Optional) URI address and port. E.g.: `http://ciihdfshost2:9870`
- **largeStorageEndpointPath** – large storage server root path for the data. Default: `/esOLs`
- **exceptionWhenRemoteNotAccessible** - Exception is thrown by the Config client if the remote DB is not accessible. In case this setting is false, remote DB will be ignored when the "*" is used. Default: false.
- **connectionTimeout** – Timeout for the connection to the remote DB in seconds. Default: 30.**
**

Default value of **configLocalDbRoot** is defined by the following procedure:

1. If **CONFIG_LOCAL_DB** environment variable is set, this variable will define the location.
2. If **CONFIG_LOCAL_DB** environment variable is not set, it will check for existence of **INTROOT** environment variable. If this variable is set, then location fallback **\$INTROOT/localdb**.
3. If none of the variables is set, the default location will be **/localdb**.

To override the default values, path to config.ini file can be specified using the **CONFIG_CLIENT_INI** environment variable. The following listing shows a simple example:

```
configServiceAdress:http://configServer.elt.eso:12345
largeStorageAdress:http://hdfs.elt.eso:9087
```



13.15 Configuration service settings

Configuration service needs some settings to work properly. Settings have predefined default values. The configuration file can be created to override predefined values. The file must contain key/value pair settings needed by the configuration service client. The following settings are defined:

- **metadata_class_index** – default name of the Elasticsearch index, which contains metadata classes. Default value: `metadata_class`.
- **metadata_instance_index** – default name of the Elasticsearch index, which contains metadata instances. Default value: `metadata_instance`.
- **configuration_class_index** – default name of the Elasticsearch index, which contains config classes. Default value: `configuration_class`.
- **configuration_instance_indices** – list of ES indices that contain target configurations. Default value: `configuration_instance`.
- **configServiceAddress** – address to which the service is bound to. Default: `zpb.rr://0.0.0.0:9116/`
- **serviceClientEndpointPath** – endpoint URI used by the Config clients. Default: `configuration/service/clientApi`
- **serviceControlEndpointPath** – endpoint URI used by the GUI. Default: `configuration/service/control`
- **config_cache_eviction_in_seconds** – Config service internal cache eviction time for Target configurations. Default: 100
- **metadata_cache_eviction_in_seconds** – Config service internal cache eviction time for meta-data. Default: 100
- **max_search_size** – the maximum number of documents that can be retrieved from Elasticsearch at once. Default value: 10000.
- **elastic_URI** – address and port to the Elasticsearch service. Default value: <http://ciielastichost:9200>.
- **elastic_backup_URI** – Optional; address and port to the Elasticsearch backup service. Default value: <http://ciielastichost2:9200>.
- **elastic_refresh_policy** – Index refresh policy used by the Elasticsearch. Default: None.
- **largeStorageAddress** – address of the large storage server. Default: <http://ciihdfshost:9870>
- **largeStorageBackupAddress** – backup address of the large storage server. Default: <http://ciihdfshost:9870>
- **largeStorageEndpointPath** – large storage server root path for the data. Default: `/esoLs`

The path to `configservice.ini` file must be set, to override default values. This is done by using the **CONFIG_SERVICE_INI** environment variable. The following listing shows the example of the `configservice.ini` file.



ELT CII - User Manual

Doc. Number: ESO-XXXXXX
Doc. Version: 1
Released on: -
Page: 583 of 589

```
metadata_class_index:metadata_class  
metadata_instance_index:metadata_instance  
configuration_class_index:configuration_class  
configuration_instance_index: configuration_instance  
max_search_size:10000  
elastic_URI: 'http://localhost:9200'  
elastic_backup_URI: 'http://localhost:9200'  
elastic_refresh_policy: NONE
```



13.16 Code produced by the generators

Listing 7-1 Java generated code by the config tool

```
/*  
 * @copyright (c) Copyright ESO 2019 All Rights Reserved ESO (eso.org) is an  
 * Intergovernmental  
 * Organisation, and therefore special legal conditions apply.  
 */  
package elt.config.classes;  
  
import com.fasterxml.jackson.databind.ObjectMapper;  
import com.fasterxml.jackson.databind.ObjectWriter;  
import com.fasterxml.jackson.annotation.JsonIgnore;  
import com.fasterxml.jackson.core.JsonProcessingException;  
import java.util.List;  
import elt.error.CiiInvalidTypeException;  
import elt.config.exceptions.CiiConfigException;  
import elt.config.CiiConfigClass;  
  
import elt.config.classes.meta.MdInt32;  
import elt.config.classes.meta.MdDouble;  
import elt.config.classes.meta.MdString;  
  
//Personal config test class  
public class MySimpleClass extends CiiConfigClass {  
  
    private MdInt32 testIntValue;  
    private MdDouble testDoubleValue;  
    private MdString testStringValue;  
  
    //default constructor is needed by the ser/deser engine.  
    public MySimpleClass() {  
        super();  
        this.testIntValue = MdInt32.getDefaultMDI();  
        this.testDoubleValue = MdDouble.getDefaultMDI();  
        this.testStringValue = MdString.getDefaultMDI();  
    }  
  
    public MySimpleClass(String name,  
        MdInt32 testIntValue,  
        MdDouble testDoubleValue,  
        MdString testStringValue) {  
        super(name);  
        this.testIntValue = testIntValue;  
        this.testDoubleValue = testDoubleValue;  
    }  
}
```

(continues on next page)



(continued from previous page)

```
    this.testStringValue = testStringValue;
}

public MySimpleClass(String name,
    int testIntValue,
    double testDoubleValue,
    String testStringValue) throws CiiInvalidTypeException {

    super(name);
    this.testIntValue = new MdInt32(testIntValue);
    this.testDoubleValue = new MdDouble(testDoubleValue);
    this.testStringValue = new MdString(testStringValue);
}

public MdInt32 getTestIntValue() {
    return testIntValue;
}

public void setTestIntValue(MdInt32 testIntValue) {
    this.testIntValue = testIntValue;
}

@JsonIgnore
public int getTestIntValueValue() {
    return testIntValue.getValue();
}

@JsonIgnore
public void setTestIntValueValue(int testIntValue) throws CiiConfigException {
    this.testIntValue.setValue(testIntValue);
}

public MdDouble getTestDoubleValue() {
    return testDoubleValue;
}

public void setTestDoubleValue(MdDouble testDoubleValue) {
    this.testDoubleValue = testDoubleValue;
}

@JsonIgnore
public double getTestDoubleValueValue() {
    return testDoubleValue.getValue();
}

@JsonIgnore
public void setTestDoubleValueValue(double testDoubleValue) throws
↪CiiConfigException {
```

(continues on next page)



(continued from previous page)

```
        this.testDoubleValue.setValue(testDoubleValue);
    }
    public MdString getTestStringValue() {
        return testStringValue;
    }

    public void setTestStringValue(MdString testStringValue) {
        this.testStringValue = testStringValue;
    }

    @JsonIgnore
    public String getTestStringValueValue() {
        return testStringValue.getValue();
    }

    @JsonIgnore
    public void setTestStringValueValue(String testStringValue) throws CiiConfigException {
        this.testStringValue.setValue(testStringValue);
    }

    public String toString() {
        try {
            ObjectWriter writer = new ObjectMapper().writer().
            withDefaultPrettyPrinter();
            return writer.writeValueAsString(this);
        } catch (JsonProcessingException e) {
            e.printStackTrace();
            return "";
        }
    }
}
```

Listing 7-2 HPP generated code from by config tool

```
#ifndef CII_SVCS_CONFIG_TEST_MYSIMPLECLASS_HPP_
#define CII_SVCS_CONFIG_TEST_MYSIMPLECLASS_HPP_
/*****
 * @copyright (c) Copyright ESO 2019 All Rights Reserved ESO (eso.org) is an
 * Intergovernmental
 * Organisation, and therefore special legal conditions apply.
 *****/

#include <ciiNodeContainer.hpp>
#include <ciiNodeFactory.hpp>

#include <string>
#include <vector>
```

(continues on next page)



(continued from previous page)

```
using namespace elt::config;

namespace elt{
namespace config{
namespace classes{
class MySimpleClass : public ::elt::config::CiiConfigClass {

public:
    DECLARE_NODE_CLASS();

    DECLARE_NODE_FIELD_METADATA(MdInt32, testIntValue);
    DECLARE_NODE_FIELD_METADATA(MdDouble, testDoubleValue);
    DECLARE_NODE_FIELD_METADATA(MdString, testStringValue);
};
// class MySimpleClass
} //namespace elt
} //namespace config
} //namespace classes

#endif // CII_SVCS_CONFIG_TEST_MYSIMPLECLASS_HPP_
```

Listing 7-3 CPP generated code from the config tool

```
/*
 * @copyright (c) Copyright ESO 2019 All Rights Reserved ESO (eso.org) is an
 ↪Intergovernmental
 * Organisation, and therefore special legal conditions apply.
 */
#include <ciiNodeFactory.hpp>
#include "mySimpleClass.hpp"

namespace elt{
namespace config{
namespace classes{

IMPLEMENT_NODE_CLASS(MySimpleClass, MySimpleClass)

} //namespace elt
} //namespace config
} //namespace classes
```

Listing 7-4 Python generated code from the config tool

```
py::class_<MySimpleClass, std::shared_ptr<MySimpleClass>,
::elt::config::CiiConfigClass>(m, "MySimpleClass")
.def_static("getNewNodeInstance", []() {
    return ::elt::config::CiiNodeFactory::getNewNodeInstance<MySimpleClass>();
```

(continues on next page)



(continued from previous page)

```
    })
    .def("get_testIntValue", &MySimpleClass::get_testIntValue)
    .def("set_testIntValue", (void (MySimpleClass::*)(const_
↪CiiCPIntegerDataClass::data_type_t&))
                                                    &MySimpleClass::set_
↪testIntValue)
    .def("set_testIntValue", (void (MySimpleClass::*)(const_
↪CiiCPIntegerDataClass::data_type_t&,
                                                    const std::shared_ptr
↪<CiiCPIntegerDataClass>&))
                                                    &MySimpleClass::set_
↪testIntValue)

    .def("get_testDoubleValue", &MySimpleClass::get_testDoubleValue)
    .def("set_testDoubleValue", (void (MySimpleClass::*)(const_
↪CiiCPDoubleDataClass::data_type_t&))
                                                    &MySimpleClass::set_
↪testDoubleValue)
    .def("set_testDoubleValue", (void (MySimpleClass::*)(const_
↪CiiCPDoubleDataClass::data_type_t&,
                                                    const std::shared_ptr
↪<CiiCPDoubleDataClass>&))
                                                    &MySimpleClass::set_
↪testDoubleValue)

    .def("get_testStringValue", &MySimpleClass::get_testStringValue)
    .def("set_testStringValue", (void (MySimpleClass::*)(const MdString::data_type_
↪t&))
                                                    &MySimpleClass::set_
↪testStringValue)
    .def("set_testStringValue", (void (MySimpleClass::*)(const MdString::data_type_
↪t&,
                                                    const std::shared_ptr
↪<MdString>&))
                                                    &MySimpleClass::set_
↪testStringValue)
;
}
```

Listing 7-5 Deployed JSON file from the config-tool

```
{
  "__comment__" : "Configuration for cryo cooler",
  "data" : {
    "@type" : "MySimpleClass",
    "testStringValue" : {
      "@type" : "MdString",
      "metadataInstance" : "ConfStringStd",
```

(continues on next page)



(continued from previous page)

```
"metadataInstanceVersion" : 1,  
  "value" : "My string defined value",  
  "@genType" : "STRING"  
},  
"testIntValue" : {  
  "@type" : "MdNumber",  
  "metadataInstance" : "ConfInt32Std",  
  "metadataInstanceVersion" : 1,  
  "@genType" : "INT32",  
  "value" : 155  
},  
"testDoubleValue" : {  
  "@type" : "MdNumber",  
  "metadataInstance" : "ConfFloatStd",  
  "metadataInstanceVersion" : 1,  
  "@genType" : "SINGLE",  
  "value" : 155.33  
}  
}  
}
```