

COALY

Context aware logging system for RUST

Requirement Specification

Copyright (c) 2022, Frank Sommer.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE.

FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Table of contents

AIM OF THE PROJECT.....2

REQUIREMENTS.....3

 GENERAL.....3

 CONFIGURATION.....5

 OPERATION.....12

GLOSSARY.....14

Aim of the project

The project's goal is a flexible library for logging and tracing functionality in Rust applications.

There is definitely no lack of such libraries at all, and for the vast majority of applications one of the common logging crates is a preferable choice.

However, especially for complex applications where direct debugging isn't possible and there's a high probability that spurious errors must be tracked down, this library may be helpful due to some unique features:

- **dynamic operational mode**
the operational mode (filtering of output records according to their level, e. g. error or warning), is usually defined once upon application start and on a per-module basis. In Coaly, the operational mode is set upon application start and may change whenever configurable events like certain function calls or structure creations occur.
- **custom output record formatting**
the format of output records may be configured with a set of placeholder variables, on a per record level basis
- **several output resource kinds**
support for output resource kinds file, memory mapped file, console, syslog and network
- **output resource splitting**
file based output resources may be process-, thread- or record level-specific
- **built-in rollover**
file based output resources can be rolled over to a new version depending on either file size or application run time
- **fault tolerance**
automatic caching of output records in case of temporary failures

Requirements

General

RQ-1 Supported platforms

The system shall support the operating systems Linux, BSD and Windows; on all architectures where a Rust ecosystem is available.

RQ-2 Use cases

The system shall be usable for standalone applications.

The system shall be usable for dedicated logging servers.

The system shall use the same configuration file structure for both standalone and logging server applications.

RQ-3 Security

The system may use unsafe code only for absolute performance critical functions, or in cases where a needed functionality can't be accessed otherwise.

If the system must use unsafe code, extra comprehensive tests must be implemented to ensure stability and minimize security risks.

The system shall not encrypt messages sent to a logging server, it is not planned to expose a logging server to the internet.

The system shall provide simple access control to a logging server to prevent accidental interference with other applications in the network.

RQ-4 Localization

The system shall support localization of all messages issued by the system itself.

The system shall supply and use English locale as default.

The system shall use the first supported locale found in the following order:

1. Locale referenced by environment variable `COALY_LANG`
2. Platform locale
3. English locale

RQ-5 Networking

The system shall support the protocols UDP and TCP (both on IPv4 and IPv6) and Unix sockets.

The system will require any URL to be specified in the form `protocol://address:port`, where port and its leading colon are optional.

The system will require IPv6 addresses (excluding port) to be specified within square brackets.

RQ-6 API

The system shall provide a function to initialize the logging system.

The system shall provide a function to shutdown the logging system.

The system shall provide macros to issue output records.

The system shall provide a trait to recognize logging relevant application objects.

RQ-7 Performance

The system shall process all tasks in a background thread as far as possible.

The system shall minimize memory copy operations.

RQ-8 Stability

The system must not panic by itself.

RQ-9 Auxiliary functionality

The system shall provide a script to insert function traces into a Rust source file.

The system shall provide a script to convert the contents of a memory mapped file to a regular file.

The system shall provide a script to sort all output records in a regular file by time stamp.

Configuration

RQ-101 Configuration options

The system shall be configurable by a configuration file read upon system initialization.

The system shall report problems in the configuration file with line number and comprehensive, localized error messages.

The system shall not provide functionality for configuration changes after system initialization.

The system shall provide default settings for all configuration parameters, allowing usage without a configuration file.

RQ-102 Configuration file format

The system will accept configuration files in TOML format only.

RQ-103 Configuration file precedence

The system shall apply the settings defined in the first configuration file found using the following order:

1. File name passed to library initialization function
2. File referenced by environment variable `COALY_CFG_FILE`
3. File `coaly.toml` in the application's working directory

The system shall use system defaults if any of the files above is not usable.

RQ-104 Unspecified configuration parameters

The system shall use system default settings for parameters not specified by the relevant configuration file, i. e. all settings defined in a configuration file with lower precedence than the relevant one are ignored.

RQ-105 Configuration parameter “File version”

The system shall provide a configuration parameter to specify the configuration file version.

The system shall provide a default setting of 1 for the configuration file version.

The system shall ignore the parameter in the initial release.

The system will use the version to keep backward compatibility in future releases.

RQ-106 Configuration parameter “Application ID”

The system shall provide a configuration parameter to specify an application ID.

The system shall provide a default setting of 0 (not relevant/any application) for the application ID.

The system may use the application ID for access control and grouping functionality within a dedicated logging server.

RQ-107 Configuration parameter “Application name”

The system shall provide a configuration parameter to specify an application name.

The system shall provide a default setting of “” (empty string) for the application name.

RQ-108 Configuration parameter “Output path”

The system shall provide a configuration parameter to specify the directory for file based output resources.

The system shall determine the default directory in the following order:

1. directory referenced by environment variable `COALY_OUTPUT_PATH`
2. directory where the application executable is located

RQ-109 Configuration parameter “Fallback path”

The system shall provide a configuration parameter to specify an alternative path for cases, where the output path is not writable or temporary network problems occur.

The system shall determine the default directory in the following order:

1. directory referenced by environment variable `COALY_FALLBACK_PATH`
2. platform’s temp directory

RQ-110 Configuration parameter “Record levels”

The system shall provide a configuration parameter to specify name and ID character for all supported record levels.

The system shall provide the following default values:

- emergency – ID: ‘Y’, name: “EMGCY”
- alert – ID: ‘A’, name: “ALERT”
- critical – ID: ‘C’, name: “CRIT”
- error – ID: ‘E’, name: “ERROR”
- warning – ID: ‘W’, name: “WARNING”
- notice – ID: ‘N’, name: “NOTICE”
- info – ID: ‘I’, name: “INFO”
- debug – ID: ‘D’, name: “DEBUG”
- function – ID: ‘F’, name: “FUNC”
- module – ID: ‘M’, name: “MOD”
- object – ID: ‘O’, name: “OBJ”

RQ-111 Configuration parameter “Mode change stack size”

The system shall provide a configuration parameter to specify the size of the internal stack for operational mode changes.

The system shall provide a default value of 32768 for the size of the mode change stack.

The system will use the value to cope with overflows during run time.

RQ-112 Configuration parameter “Operational mode”

The system shall provide a configuration parameter to specify the initial operational mode.

The system shall enable all record levels from emergency through info as default.

The system shall not use buffering for any record level as default.

RQ-113 Configuration parameter “Date/Time format”

The system shall provide a configuration parameter to specify one or more formats defining how dates, times or timestamps are written to an output resource.

The system shall allow specification of independent date, time or timestamp formats by supplying appropriate sub-parameters.

The system shall require each sub-parameter to be a string value containing a subset of the variables specified for Rust’s `chrono::format::strftime` module, and treat any other contents as literal characters.

The system shall accept format specifications where any part (date, time, timestamp) is missing and use the default setting for the part instead.

The system shall identify a date/time format by a user supplied name.

The system shall ignore a date/time format using the same name as an already defined format.

The system shall not allow custom date/time formats to be used for the naming of file based output resources.

The system shall allow the variables `%d` (2-digit day), `%m` (2-digit month), `%y` (2-digit year) and `%Y` (4-digit year) within date and timestamp format specifications.

The system shall allow the variables %H (hour 00-23), %I (hour 00-12), %p (am or pm), %P (AM or PM), %M (minute 00-59), %S (second 00-59) and %f (fractional seconds, optionally with precision 1-9, e. g. %3f for 3 digits) within time and timestamp format specifications.

The system shall allow the variables %z (timezone without colon) and %Z (timezone with colon) within timestamp format specifications.

RQ-114 Configuration parameter “Output format”

The system shall provide a configuration parameter to specify one or more formats defining how output records are written to an output resource.

The system shall provide a sub-parameter to specify one or more record levels for which the format should be applied.

The system shall provide a sub-parameter to specify one or more record triggers for which the format should be applied.

The system shall provide a sub-parameter to specify the name of the date/time format to be used.

The system shall provide a sub-parameter to specify the items of the output format.

The system shall require the sub-parameter for the items specification to be string containing any of the case sensitive variables below, and treat any other characters as literals.

The system shall identify an output format by a user supplied name.

The system shall ignore an output format using the same name as an already defined format.

The system shall provide the following variables that can be used for the items of an output format:

- “\$AppId” – application ID
- “\$AppName” – application name
- “\$Date” – current date
- “\$Env[...]” - value of specified environment variable
- “\$HostName” – host name
- “\$IpAddress” – host’s IP address, to be used in a logging server only
- “\$Level” – name of the record level
- “\$LevelId” – ID character of the record level
- “\$Message” – log or trace text as issued by the application
- “\$ObserverName” – name of the record trigger (function, module or application object)
- “\$ObserverValue” – custom value of the record trigger (application object only)
- “\$ProcessId” – the application’s process ID
- “\$ProcessName” – name of the application’s executable
- “\$PureSourceFileName” – name of the source file that issued an output record, without path
- “\$SourceFileName” – name of the source file that issued an output record, including path
- “\$SourceLineNr” – line number in the source file, where an output record was issued
- “\$ThreadId” – the ID of the thread that issued the output record
- “\$ThreadName” – the name of the thread that issued the output record
- “\$Time” – the current time
- “\$TimeStamp” – the current date and time

RQ-115 Configuration parameter “Rollover policy”

The system shall provide a configuration parameter to specify one or more policies to handle rollover of file based resources.

The system shall identify a rollover policy by a user supplied name.

The system shall provide a sub-parameter to specify the condition when a rollover shall occur.

The system shall accept a string value with format “size > value” for the sub-parameter to specify that a rollover shall occur, if the file size exceeds the given value. The value must be

specified as an integer number, optionally followed by a unit indicator K, M or G for Kilobytes, Megabytes or Gigabytes.

The system shall accept a string value with format “every [n] [second(s) | minute(s) | hour(s) | day(s)]” for the sub-parameter to specify that a rollover shall occur, whenever a multiple of the specified time span since system initialization has elapsed. The number of time units must be specified as a positive integer number.

The system shall accept a string value with format “every [n] [hour(s) | day(s) | week(s) | month(s)] at *timestamp*” for the sub-parameter to specify that a rollover shall occur, whenever a multiple of the specified time span at the given timestamp has elapsed. The number of time units must be specified as a positive integer number. The system shall require the timestamp to be specified as a unique value within the interval unit as follows:

- for unit hour the timestamp must contain minute and second
- for unit day the timestamp must contain hour, minute and second
- for unit week the timestamp must contain day of week, hour, minute and second. Day of week must be specified in English.
- for unit month the timestamp must contain day, hour, minute and second. Ultimo may be used for the last day of a month.

The system shall provide a sub-parameter to specify the number of old rollover files to keep before deletion; the system shall provide 9 as a default value.

The system shall provide a sub-parameter to specify the compression algorithm to use for rollover files; the system shall support the compression algorithms zip, gzip, bzip2 and lzma and provide gzip as a default value.

The system shall not provide a default rollover policy.

RQ-116 Configuration parameter “Buffer policy”

The system shall provide a configuration parameter to specify one or more policies to handle buffering of loutput records in main memory.

The system shall identify a buffer policy by a user supplied name.

The system shall provide a sub-parameter to configure the size of a buffer in bytes; the system shall provide a default value of 32 Megabytes.

The system shall provide a sub-parameter to configure the maximum number of records the buffer may hold; the system shall provide a default value of 1 M.

The system shall provide a sub-parameter to configure the maximum size of output records without being truncated; the system shall provide a default value of 4096 Bytes.

The system shall provide a sub-parameter to configure the list of events causing the buffer to be flushed to its associated output resource; the system shall assume a default value of [exit, rollover, emergency, critical, alert, warning].

The system shall provide a default buffer policy holding the default values for all sub-parameters.

RQ-117 Configuration parameter “Output resource”

The system shall provide a configuration parameter to specify one or more output resources.

The system shall provide a mandatory sub-parameter for the resource kind. The system shall support resource kinds file, memory mapped file, standard output, standard error, network and syslog.

The system shall provide a mandatory sub-parameter for the record levels associated with the resource.

The system shall provide a sub-parameter for the application IDs associated with the resource. The system shall provide 0 (all applications) as a default value.

The system shall provide a sub-parameter to specify the name of the output format to use for the resource. The system shall assume the default output format, if the sub-parameter is not present.

The system shall provide a sub-parameter to specify the name of the buffer policy to use for the resource. The system shall assume the default buffer policy, if the sub-parameter is not present. The system shall ignore the sub-parameter for memory mapped files.

The system shall provide a sub-parameter for the local URL to use for a network or syslog resource. The system shall provide “udp://127.0.0.1:0” as default. The system shall ignore the sub-parameter for resources other than network and syslog.

The system shall provide a sub-parameter for the logging server URL to use for a network or syslog resource. The system shall provide “udp://127.0.0.1:3690” (network) resp. “udp://127.0.0.1:514” (syslog) as default. The system shall ignore the sub-parameter for resources other than network and syslog.

The system shall provide a sub-parameter to specify the syslog facility. The system shall provide 1 (user level) as default. The system shall ignore the sub-parameter for resources other than syslog.

The system shall provide a sub-parameter to specify the name of the rollover policy to use for the resource. If the sub-parameter is not present, the system shall not apply rollover to the resource. The system shall ignore the sub-parameter for resources other than file and memory mapped file.

The system shall provide a sub-parameter for the size of a memory mapped file. The system shall provide 32 Mbytes as default. The system shall ignore the sub-parameter for resources other than memory mapped files.

The system shall provide a sub-parameter for the name of a file or memory mapped file. The system shall require the sub-parameter for files and memory mapped files. The system shall ignore the sub-parameter for resources other than files and memory mapped files. The system shall provide the following variables that can be used for the file name:

- “\$AppId” – application ID
- “\$AppName” – application name
- “\$Date” – date at file creation
- “\$Env[...]” - value of specified environment variable
- “\$HostName” – host name
- “\$ProcessId” – the application’s process ID
- “\$ProcessName” – name of the application’s executable
- “\$ThreadId” – the ID of the thread that issued the output record
- “\$ThreadName” – the name of the thread that issued the output record
- “\$Time” – time at file creation
- “\$TimeStamp” – date and time at file creation

RQ-118 Configuration parameter “Mode change”

The system shall provide a configuration parameter to specify one or more mode change triggers.

The system shall provide a mandatory sub-parameter for the trigger kind. The system shall support the trigger kinds function, module and logging relevant application object.

The system shall provide a sub-parameter for the trigger name. The system will match the specified value against the qualified function or module name resp. the outcome of application object’s coaly_name function.

The system shall provide a sub-parameter for the trigger value, if the trigger kind is logging relevant application object. The system will match the specified value against the outcome of application object’s coaly_value function.

The system shall require sub-parameter name to be specified for trigger kinds function and module.

The system shall require at least one of the sub-parameters name and value to be specified for trigger kind logging relevant application object. The system will match name and value, if both are specified.

The system will allow regular expression patterns for name and value specification.

The system shall provide a sub-parameter to configure the enabled levels after the mode change; the system shall provide no change of current setting as default.

The system shall provide a sub-parameter to configure the buffered levels after the mode change; the system shall provide no change of current setting as default.

The system shall provide a sub-parameter to configure the scope of a mode change, either thread specific or process wide; the system shall provide thread specific as default.

The system shall enable levels emergency through info, disable levels debug through object and disable buffering as default.

The system shall not define any mode changes during run time as default.

RQ-119 Configuration parameter “Server data URL”

The system shall provide a configuration parameter to specify the URL where a logging server listens for output records.

The system shall provide a default URL “udp://127.0.0.1:3690”.

RQ-120 Configuration parameter “Server admin URL”

The system shall provide a configuration parameter to specify the URL where a logging server listens for administrative commands.

The system shall provide no default URL, disabling administration over network.

The system shall provide 3691 as default port, if not specified in the URL.

The system shall require a secret key to be sent by any client, if administration over network is enabled.

RQ-121 Configuration parameter “Server maximum connections”

The system shall provide a configuration parameter to specify the maximum number of simultaneous connections for a logging server.

The system shall provide 10 as default.

RQ-122 Configuration parameter “Server connection expiry”

The system shall provide a configuration parameter to specify the time span elapsed after reception of the last message, after which a logging server considers an UDP connection as expired.

The system shall provide one day as default.

RQ-123 Configuration parameter “Server maximum message size”

The system shall provide a configuration parameter to specify the maximum size of a network message, before being truncated by a logging server.

The system shall provide 64 Kbytes as default.

RQ-124 Configuration parameter “Server admin key”

The system shall provide a configuration parameter to specify the secret key that must be presented by clients in order to send administrative messages to a logging server.

The system will not provide a default.

RQ-125 Configuration parameter “Server data client restriction”

The system shall provide a configuration parameter to specify a list of clients allowed to send output records to a logging server.

The system shall provide a sub-parameter to specify the source URL of a client.

The system shall accept an asterisk to match a whole subnet within the sub-parameter specifying the source URL of a client.

The system shall assume any port to be allowed, if no port was given in the sub-parameter specifying the source URL of a client.

The system shall provide a sub-parameter to specify a list of allowed application IDs.

The system shall allow restrict connections to local host, if a restriction specifies a list of application IDs, but omits the source URL.

The system shall accept any connections from local host with any application ID as default.

RQ-126 Configuration parameter “Server admin client restriction”

The system shall provide a configuration parameter to specify a list of client URLs allowed to send administrative commands to a logging server.

The system shall accept an asterisk to match a whole subnet.

The system shall assume any port to be allowed, if no port was given in the URL.

The system shall accept any connections from local host as default.

Operation

RQ-201 Emergency error handling

The system shall provide one output resource for emergency cases and use the following sequence to determine that resource:

1. file referenced by environment variable `COALY_EMERGENCY_LOGFILE`
2. file `coaly.log` in the application's working directory
3. file `coaly.log` in the system's temp directory
4. standard output

If no emergency output resource is writable, the system shall silently discard any output, and not panic.

If an output record can't be written to standard output or standard error device, the system shall silently discard the record.

RQ-202 Fallback error handling

If an output record can't be written to a file or memory mapped file, the system shall try to write the record to a file with the same name in the fallback directory instead.

If an output record can't be written to network, the system shall try to write the record to a file with a name built from the network address in the fallback directory instead.

The system shall transfer all output written to fallback resource to its original resource and delete the fallback resource as soon as the problem ends.

The system shall log begin and end of a fallback period in the system's emergency resource.

If an error occurs during a write operation to a fallback resource, the system shall silently discard the output record, and log the error in the system's emergency resource, but for the first error within a fallback period only.

RQ-203 File names

The system shall not allow path separator characters in file names, neither direct nor through evaluation of variables of the file name specification. The system will replace all path separator characters with underscores.

The system shall recognize thread specific output files by the presence of at least one of the variables `$ThreadId` and `$ThreadName` in the file name specification. The system shall give precedence to `$ThreadId` over `$ThreadName` in the mapping of threads to output resources, since two threads may have the same name, but never the same ID.

The system shall recognize application specific output files by the presence of variable `$AppId` in the file name specification. The system can therefore map output records of different applications to one output file on a logging server.

RQ-204 Buffering

The system shall allow output records to be buffered in main memory for all output resource kinds except for memory mapped files, where the operating system provides buffering.

The system shall organize main memory buffers circular, so that the oldest record(s) are overwritten if a buffer can't hold a new output record.

The system shall truncate output records larger than the buffer size to exactly buffer size.

RQ-205 Buffer flush

The system shall allow a buffer to be flushed upon application termination.

The system shall allow a buffer to be flushed when it is full.

The system shall allow a buffer to be flushed when an output record with a certain level is written, these levels can be configured in the list of flush events.

If a buffer's associated output resource is a regular file, the system shall allow a buffer to be flushed upon file rollover. If the rollover occurred due to file size, the buffer contents is written to the new output file, after rollover takes place. If the rollover occurred due to time, the buffer contents is written to the old output resource, before rollover takes place.

If no flush event occurs, the system shall discard the buffer contents upon application termination. The system shall not provide functionality to keep output records in chronological order (a buffer may hold records being older than those written to the associated output resource).

RQ-206 Rollover

The system shall allow file based output resources to be archived automatically by closing and renaming the current file, then open a new empty file.

The system shall sort rollover files by file name only, not by file attributes like modification time stamp or similar.

RQ-207 Rollover triggers

The system shall provide functionality to perform a rollover, when the content size of a file or memory mapped file exceeds a limit.

If an output resource is configured for size based rollover, the system shall check whether the rollover size limit was exceeded upon every output operation.

The system shall provide functionality to perform a rollover, when a multiple of a time interval since system initialization has been elapsed. The system shall allow time units second, minute, hour and day to be used in that case.

The system shall provide functionality to perform a rollover, when a multiple of a time interval at a certain timestamp has been elapsed. The system shall allow time units hour, day, week and month to be used in that case.

If an output resource is configured for periodic rollover, the system shall check every second whether a multiple of the specified interval has elapsed.

RQ-208 Mode change applying

Whenever a mode change trigger is detected, the system shall immediately apply the new operational mode for the actual output record.

Whenever a mode change trigger ends, the system shall apply the current operational mode for the actual output record and thereafter restore the mode that was active before the trigger.

RQ-209 Thread specific mode change

The system shall remember active thread specific mode changes in a stack up to a configurable limit.

If the limit for the size of the mode change stack is exceeded, the system shall ignore further thread specific mode changes.

If a process wide mode change is active, the system shall normally update the mode stack for a thread specific mode change.

The system shall ignore consecutive thread specific mode changes of the same trigger, both for trigger activation and end.

The system shall not provide functionality to detect indirect recursion for thread specific mode changes.

RQ-210 Process wide mode change

The system shall remember only one process wide mode change at a time.

The system shall give precedence for a process wide mode change over thread specific changes.

If a process wide mode change occurs while another one being active, the system shall ignore the the new process wide mode change.

Glossary

- Fallback directory** Directory to use for output in case of a temporary failures like network down or file system full. Hence, the directory should be located in a different file system.
- Logging data** Human readable information about essential events in an application, like system startup and shutdown, important user interactions, incoming and outgoing data or error conditions.
- Mode change trigger** The trigger causing a change of the operational mode. May be a function, a module or a logging relevant application object.
- Operational mode** Setting, how output records of a certain record level shall be handled. Records may be written directly to a physical output resource (level enabled), buffered in main memory (level buffered) or discarded (level disabled).
The operational mode can change upon crossing a certain function or module boundary; or creation resp. drop of a certain application specific object.
- Output buffering** For all resource kinds except for memory mapped files, a block in main memory can be reserved to buffer logging or tracing data before writing it to the physical output resource. Buffering improves performance at the cost of data loss in case of an application crash.
Output buffers are managed circular with oldest information overwritten if the buffer is full.
Depending on their description in the Coaly configuration file, output buffers are flushed on events like occurrence of a certain record level, buffer full, file rollover or application termination.
- Output record** Logging or tracing data issued by the application, enhanced with configurable attributes from the Coaly system, like time stamp or name of the source file, where the data originated. An output record may span over more than one line in a physical output resource.

Output resource	<p>An output resource denotes a target, where logging or tracing data is written to. Such a target has a specific resource kind and may be either logical or physical.</p> <p>A physical output resource corresponds to exactly one operating system resource, a logical output resource correlates to its description in the Coaly configuration.</p> <p>For resource kinds standard output/error and network socket, logical and physical resource are identical.</p> <p>For resource kinds plain and memory mapped file, a logical resource may map to more than one physical resource and eventually change over time. If the logical resource is configured to be thread-specific, the logical resource maps to as many physical resources as there are threads in the application. The actual physical output file will change, if a time or size based rollover mechanism has been specified for the logical output resource.</p>
Record level	<p>Severity level of an output record. Levels are based on syslog severities as specified in RFC 5424, extended with debug related levels for entering or leaving functions or modules and tracking of certain application specific objects.</p>
Record trigger	<p>The cause, why an output record was issued. May be an explicit message issued by the application, enter of a function or module, leaving of a function or module, creation of a logging relevant application object or drop of a logging relevant application object.</p>
Resource kind	<p>Coaly supports the resource kinds file, memory mapped file, standard output or error device, syslog and network socket.</p>
Tracing data	<p>Human readable information about program flow and status. To be used to track down software problems only.</p>