



RTEMS Software Engineering

Release 6.fa31da1 (13th November 2020)

© 1988, 2020 RTEMS Project and contributors

CONTENTS

1	Preface	3
2	RTEMS Project Mission Statement	5
2.1	Free Software Project	6
2.2	Design and Development Goals	7
2.3	Open Development Environment	8
3	RTEMS Stakeholders	9
4	Introduction to Pre-Qualification	11
4.1	Stakeholder Involvement	13
5	Software Requirements Engineering	15
5.1	Requirements for Requirements	17
5.1.1	Identification	17
5.1.2	Level of Requirements	18
5.1.2.1	Absolute Requirements	18
5.1.2.2	Absolute Prohibitions	19
5.1.2.3	Recommendations	19
5.1.2.4	Permissions	19
5.1.2.5	Possibilities and Capabilities	19
5.1.3	Syntax	20
5.1.4	Wording Restrictions	20
5.1.5	Separate Requirements	22
5.1.6	Conflict Free Requirements	23
5.1.7	Use of Project-Specific Terms and Abbreviations	23
5.1.8	Justification of Requirements	23
5.1.9	Requirement Validation	23
5.1.10	Resources and Performance	24
5.2	Specification Items	25
5.2.1	Specification Item Hierarchy	25
5.2.2	Specification Item Types	26
5.2.2.1	Root Item Type	26
5.2.2.2	Build Item Type	27
5.2.2.3	Build Ada Test Program Item Type	28
5.2.2.4	Build BSP Item Type	29
5.2.2.5	Build Configuration File Item Type	31
5.2.2.6	Build Configuration Header Item Type	31

5.2.2.7	Build Group Item Type	32
5.2.2.8	Build Library Item Type	33
5.2.2.9	Build Objects Item Type	34
5.2.2.10	Build Option Item Type	35
5.2.2.11	Build Script Item Type	36
5.2.2.12	Build Start File Item Type	38
5.2.2.13	Build Test Program Item Type	38
5.2.2.14	Constraint Item Type	40
5.2.2.15	Glossary Item Type	40
5.2.2.16	Glossary Group Item Type	40
5.2.2.17	Glossary Term Item Type	40
5.2.2.18	Interface Item Type	41
5.2.2.19	Application Configuration Group Item Type	41
5.2.2.20	Application Configuration Option Item Type	42
5.2.2.21	Application Configuration Feature Enable Option Item Type	42
5.2.2.22	Application Configuration Feature Option Item Type	43
5.2.2.23	Application Configuration Value Option Item Type	43
5.2.2.24	Interface Compound Item Type	43
5.2.2.25	Interface Container Item Type	44
5.2.2.26	Interface Define Item Type	44
5.2.2.27	Interface Domain Item Type	44
5.2.2.28	Interface Enum Item Type	45
5.2.2.29	Interface Enumerator Item Type	45
5.2.2.30	Interface Forward Declaration Item Type	45
5.2.2.31	Interface Function Item Type	46
5.2.2.32	Interface Group Item Type	46
5.2.2.33	Interface Header File Item Type	47
5.2.2.34	Interface Macro Item Type	47
5.2.2.35	Interface Typedef Item Type	48
5.2.2.36	Interface Unspecified Item Type	48
5.2.2.37	Interface Variable Item Type	49
5.2.2.38	Requirement Item Type	49
5.2.2.39	Functional Requirement Item Type	50
5.2.2.40	Action Requirement Item Type	50
5.2.2.41	Generic Functional Requirement Item Type	54
5.2.2.42	Non-Functional Requirement Item Type	54
5.2.2.43	Generic Non-Functional Requirement Item Type	55
5.2.2.44	Runtime Performance Requirement Item Type	55
5.2.2.45	Requirement Validation Item Type	58
5.2.2.46	Runtime Measurement Test Item Type	58
5.2.2.47	Specification Item Type	59
5.2.2.48	Test Case Item Type	60
5.2.2.49	Test Platform Item Type	61
5.2.2.50	Test Procedure Item Type	62
5.2.2.51	Test Suite Item Type	62
5.2.3	Specification Attribute Sets and Value Types	63
5.2.3.1	Action Requirement Condition	63
5.2.3.2	Action Requirement Name	63
5.2.3.3	Action Requirement Skip Reasons	64
5.2.3.4	Action Requirement State	64
5.2.3.5	Action Requirement Transition	64

5.2.3.6	Action Requirement Transition Post-Conditions	65
5.2.3.7	Action Requirement Transition Pre-Condition State Set	65
5.2.3.8	Action Requirement Transition Pre-Conditions	66
5.2.3.9	Application Configuration Group Member Link Role	66
5.2.3.10	Application Configuration Option Constraint Set	66
5.2.3.11	Application Configuration Option Name	67
5.2.3.12	Boolean or Integer or String	67
5.2.3.13	Build Assembler Option	67
5.2.3.14	Build C Compiler Option	67
5.2.3.15	Build C Preprocessor Option	68
5.2.3.16	Build C++ Compiler Option	68
5.2.3.17	Build Dependency Link Role	68
5.2.3.18	Build Include Path	68
5.2.3.19	Build Install Directive	69
5.2.3.20	Build Install Path	69
5.2.3.21	Build Link Static Library Directive	70
5.2.3.22	Build Linker Option	70
5.2.3.23	Build Option Action	70
5.2.3.24	Build Option C Compiler Check Action	73
5.2.3.25	Build Option C++ Compiler Check Action	74
5.2.3.26	Build Option Default by Variant	74
5.2.3.27	Build Option Name	74
5.2.3.28	Build Option Set Test State Action	75
5.2.3.29	Build Option Value	75
5.2.3.30	Build Source	75
5.2.3.31	Build Target	76
5.2.3.32	Build Test State	76
5.2.3.33	Build Use After Directive	76
5.2.3.34	Build Use Before Directive	77
5.2.3.35	Constraint Link Role	77
5.2.3.36	Copyright	77
5.2.3.37	Enabled-By Expression	77
5.2.3.38	Glossary Membership Link Role	78
5.2.3.39	Integer or String	78
5.2.3.40	Interface Brief Description	79
5.2.3.41	Interface Compound Definition Kind	79
5.2.3.42	Interface Compound Member Compound	80
5.2.3.43	Interface Compound Member Declaration	80
5.2.3.44	Interface Compound Member Definition	80
5.2.3.45	Interface Compound Member Definition Directive	81
5.2.3.46	Interface Compound Member Definition Variant	81
5.2.3.47	Interface Definition	81
5.2.3.48	Interface Definition Directive	82
5.2.3.49	Interface Definition Variant	82
5.2.3.50	Interface Description	82
5.2.3.51	Interface Enabled-By Expression	83
5.2.3.52	Interface Enum Definition Kind	84
5.2.3.53	Interface Enumerator Link Role	84
5.2.3.54	Interface Function Definition	84
5.2.3.55	Interface Function Definition Directive	85
5.2.3.56	Interface Function Definition Variant	85

5.2.3.57	Interface Function Link Role	85
5.2.3.58	Interface Group Identifier	85
5.2.3.59	Interface Group Membership Link Role	86
5.2.3.60	Interface Include Link Role	86
5.2.3.61	Interface Notes	86
5.2.3.62	Interface Parameter	86
5.2.3.63	Interface Parameter Direction	87
5.2.3.64	Interface Placement Link Role	87
5.2.3.65	Interface Return Directive	87
5.2.3.66	Interface Return Value	88
5.2.3.67	Interface Target Link Role	88
5.2.3.68	Link	88
5.2.3.69	Name	89
5.2.3.70	Optional String	90
5.2.3.71	Placement Order Link Role	90
5.2.3.72	Requirement Reference	90
5.2.3.73	Requirement Reference Type	90
5.2.3.74	Requirement Refinement Link Role	91
5.2.3.75	Requirement Text	91
5.2.3.76	Requirement Validation Link Role	93
5.2.3.77	Requirement Validation Method	93
5.2.3.78	Runtime Measurement Environment	93
5.2.3.79	Runtime Measurement Environment Table	94
5.2.3.80	Runtime Measurement Parameter Set	94
5.2.3.81	Runtime Measurement Request Link Role	94
5.2.3.82	Runtime Measurement Value Kind	94
5.2.3.83	Runtime Measurement Value Table	95
5.2.3.84	Runtime Performance Limit Table	95
5.2.3.85	Runtime Performance Parameter Set	95
5.2.3.86	SPDX License Identifier	95
5.2.3.87	Specification Attribute Set	96
5.2.3.88	Specification Attribute Value	96
5.2.3.89	Specification Boolean Value	97
5.2.3.90	Specification Explicit Attributes	97
5.2.3.91	Specification Floating-Point Assert	97
5.2.3.92	Specification Floating-Point Value	98
5.2.3.93	Specification Generic Attributes	99
5.2.3.94	Specification Information	99
5.2.3.95	Specification Integer Assert	100
5.2.3.96	Specification Integer Value	101
5.2.3.97	Specification List	101
5.2.3.98	Specification Mandatory Attributes	101
5.2.3.99	Specification Member Link Role	102
5.2.3.100	Specification Refinement Link Role	102
5.2.3.101	Specification String Assert	102
5.2.3.102	Specification String Value	103
5.2.3.103	Test Case Action	104
5.2.3.104	Test Case Check	104
5.2.3.105	Test Context Member	104
5.2.3.106	Test Header	105
5.2.3.107	Test Run Parameter	106

5.2.3.108	Test Support Method	106
5.2.3.109	UID	107
5.3	Traceability of Specification Items	108
5.3.1	History of Specification Items	108
5.3.2	Backward Traceability of Specification Items	108
5.3.3	Forward Traceability of Specification Items	108
5.3.4	Traceability between Software Requirements, Architecture and Design	108
5.4	Requirement Management	109
5.4.1	Change Control Board	109
5.4.2	Add a Requirement	110
5.4.3	Modify a Requirement	111
5.4.4	Mark a Requirement as Obsolete	111
5.5	Tooling	112
5.5.1	Tool Requirements	112
5.5.2	Tool Evaluation	112
5.5.3	Best Available Tool - Doorstop	113
5.5.4	Custom Requirements Management Tool	114
5.6	How-To	115
5.6.1	Getting Started	115
5.6.2	Application Configuration Options	115
5.6.2.1	Modify an Existing Group	116
5.6.2.2	Modify an Existing Option	116
5.6.2.3	Add a New Group	116
5.6.2.4	Add a New Option	116
5.6.2.5	Generate Content after Changes	117
5.6.3	Glossary Specification	117
5.6.4	Interface Specification	118
5.6.4.1	Specify an API Header File	118
5.6.4.2	Specify an API Element	118
6	Software Development Management	121
6.1	Software Development (Git Users)	122
6.1.1	Browse the Git Repository Online	122
6.1.2	Using the Git Repository	122
6.1.3	Making Changes	123
6.1.4	Working with Branches	124
6.1.5	Viewing Changes	125
6.1.6	Reverting Changes	125
6.1.7	git reset	126
6.1.8	git revert	126
6.1.9	Merging Changes	127
6.1.10	Rebasing	127
6.1.11	Accessing a developer's repository	128
6.1.12	Creating a Patch	128
6.1.13	Submitting a Patch	129
6.1.14	Configuring git send-email to use Gmail	129
6.1.15	Sending Email	129
6.1.16	Troubleshooting	130
6.1.17	Manage Your Code	130
6.1.18	Private Servers	130
6.1.19	Learn more about Git	131

6.2	Software Development (Git Writers)	132
6.2.1	SSH Access	132
6.2.2	Personal Repository	132
6.2.3	Create a personal repository	132
6.2.3.1	Check your setup	133
6.2.3.2	Push commits to personal repo master from local master	133
6.2.3.3	Push a branch onto personal repo	134
6.2.3.4	Update from upstream master (RTEMS head)	134
6.2.4	GIT Push Configuration	134
6.2.5	Pull a Developer's Repo	135
6.2.6	Committing	135
6.2.6.1	Ticket Updates	135
6.2.6.2	Commands	136
6.2.7	Pushing Multiple Commits	137
6.2.8	Oops!	137
6.3	Coding Standards	138
6.3.1	Coding Conventions	138
6.3.1.1	Source Documentation	138
6.3.1.2	Licenses	138
6.3.1.3	Language and Compiler	138
6.3.1.4	Readability	140
6.3.1.5	Robustness	140
6.3.1.6	Portability	141
6.3.1.7	Maintainability	141
6.3.1.8	Performance	141
6.3.1.9	Miscellaneous	142
6.3.1.10	Header Files	142
6.3.1.11	Layering	142
6.3.1.12	Exceptions to the Rules	143
6.3.1.13	Tools	143
6.3.2	Formatting	143
6.3.2.1	Rules	143
6.3.2.2	Eighty Character Line Limit	144
6.3.2.3	Breaking Long Lines	144
6.3.3	Deprectating Interfaces	146
6.3.4	Doxygen Guidelines	146
6.3.4.1	Group Names	146
6.3.4.2	Use Groups	146
6.3.4.3	Files	147
6.3.4.4	Type Definitions	148
6.3.4.5	Function Declarations	148
6.3.4.6	Header File Examples	150
6.3.5	File Templates	150
6.3.5.1	Copyright and License Block	151
6.3.5.2	C/C++ Header File Template	152
6.3.5.3	C/C++/Assembler Source File Template	153
6.3.5.4	Python File Template	154
6.3.5.5	Shell Scripts	155
6.3.5.6	reStructuredText File Template	155
6.3.6	Naming Rules	156
6.3.6.1	General Rules	156

6.4	Documentation Guidelines	158
6.4.1	Application Configuration Options	158
6.5	Python Development Guidelines	160
6.5.1	Python Language Versions	160
6.5.2	Python Code Formatting	160
6.5.3	Static Analysis Tools	161
6.5.4	Type Annotations	161
6.5.5	Testing	161
6.5.5.1	Test Organization	161
6.5.6	Documentation	161
6.5.7	Existing Code	162
6.5.8	Third-Party Code	162
6.6	Change Management	163
6.7	Issue Tracking	164
7	Software Test Plan Assurance and Procedures	165
7.1	Testing and Coverage	166
7.1.1	Test Suites	166
7.1.1.1	Legacy Test Suites	167
7.1.2	RTEMS Tester	167
8	Software Test Framework	169
8.1	The RTEMS Test Framework	170
8.1.1	Nomenclature	170
8.1.2	Test Cases	171
8.1.3	Test Fixture	172
8.1.4	Test Case Planning	174
8.1.5	Test Case Resource Accounting	175
8.1.6	Test Case Scoped Dynamic Memory	176
8.1.7	Test Case Destructors	178
8.1.8	Test Checks	179
8.1.8.1	Test Check Variant Conventions	179
8.1.8.2	Test Check Parameter Conventions	179
8.1.8.3	Test Check Condition Conventions	180
8.1.8.4	Test Check Type Conventions	180
8.1.8.5	Integers	181
8.1.8.6	Boolean Expressions	182
8.1.8.7	Generic Types	183
8.1.8.8	Pointers	184
8.1.8.9	Memory Areas	184
8.1.8.10	Strings	185
8.1.8.11	Characters	185
8.1.8.12	RTEMS Status Codes	186
8.1.8.13	POSIX Error Numbers	186
8.1.8.14	POSIX Status Codes	186
8.1.9	Log Messages and Formatted Output	187
8.1.10	Utility	188
8.1.11	Time Services	188
8.1.12	Code Runtime Measurements	190
8.1.13	Interrupt Tests	194
8.1.14	Test Runner	196
8.1.15	Test Verbosity	198

8.1.16	Test Reporting	199
8.1.17	Test Report Validation	203
8.1.18	Supported Platforms	204
8.2	Test Framework Requirements for RTEMS	205
8.2.1	License Requirements	205
8.2.2	Portability Requirements	205
8.2.3	Reporting Requirements	206
8.2.4	Environment Requirements	207
8.2.5	Usability Requirements	207
8.2.6	Performance Requirements	210
8.3	Off-the-shelf Test Frameworks	211
8.3.1	bdd-for-c	211
8.3.2	CBDD	211
8.3.3	Google Test	211
8.3.4	Unity	211
8.4	Standard Test Report Formats	212
8.4.1	JUnit XML	212
8.4.2	Test Anything Protocol	212
9	BSP Build System	213
9.1	Goals	214
9.2	Overview	215
9.3	Commands	216
9.3.1	BSP List	216
9.3.2	BSP Defaults	216
9.3.3	Configure	216
9.3.4	Build, Clean, and Install	216
9.4	UID Naming Conventions	217
9.5	Build Specification Items	219
9.6	How-To	221
9.6.1	Find the Right Item	221
9.6.2	Create a BSP Architecture	221
9.6.3	Create a BSP Family	221
9.6.4	Add a Base BSP to a BSP Family	222
9.6.5	Add a BSP Option	223
9.6.6	Extend a BSP Family with a Group	224
9.6.7	Add a Test Program	224
9.6.8	Add a Library	224
10	Software Release Management	227
10.1	Release Process	228
10.1.1	Releases	228
10.1.1.1	Release Layout	228
10.1.1.2	Release Version Numbering	229
10.1.1.3	Release Scripts	229
10.1.1.4	Release Snapshots	230
10.1.2	Release Repositories	230
10.1.3	Pre-Release Procedure	230
10.1.4	Release Branching	231
10.1.4.1	LibBSD Release Branch	231
10.1.4.2	Pre-Branch Procedure	231
10.1.4.3	Branch Procedure	231

10.1.4.4 Post-Branch Procedure	232
10.1.5 Release Procedure	233
10.1.6 Post-Release Procedure	234
10.2 Software Change Report Generation	235
10.3 Version Description Document (VDD) Generation	236
11 User's Manuals	237
11.1 Documentation Style Guidelines	238
12 Licensing Requirements	239
12.1 Rationale	240
13 Appendix: Core Qualification Artifacts/Documents	243
14 Glossary	247
Bibliography	251
Index	253

Copyrights and License

© 2018, 2020 embedded brains GmbH

© 2018, 2020 Sebastian Huber

© 1988, 2015 On-Line Applications Research Corporation (OAR)

This document is available under the [Creative Commons Attribution-ShareAlike 4.0 International Public License](#).

The authors have used their best efforts in preparing this material. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. No warranty of any kind, expressed or implied, with regard to the software or the material contained in this document is provided. No liability arising out of the application or use of any product described in this document is assumed. The authors reserve the right to revise this material and to make changes from time to time in the content hereof without obligation to notify anyone of such revision or changes.

The RTEMS Project is hosted at <https://www.rtems.org>. Any inquiries concerning RTEMS, its related support components, or its documentation should be directed to the RTEMS Project community.

RTEMS Online Resources

Home	https://www.rtems.org
Documentation	https://docs.rtems.org
Mailing Lists	https://lists.rtems.org
Bug Reporting	https://devel.rtems.org/wiki/Developer/Bug_Reporting
Git Repositories	https://git.rtems.org
Developers	https://devel.rtems.org

PREFACE

This manual aims to guide the development of RTEMS itself. You should read this document if you want to participate in the development of RTEMS. Users of RTEMS may find background information in this manual. Please refer to the RTEMS User Manual and RTEMS Classic API Guide if you want to know how the RTEMS development environment is set up and how you can develop applications using RTEMS.

RTEMS PROJECT MISSION STATEMENT

RTEMS development done under the umbrella of the RTEMS Project aims to provide a free and open real-time operating system targeted towards deeply embedded systems which is competitive with proprietary products. The RTEMS Project encourages the support and use of standard APIs in order to promote application portability and ease porting other packages to the RTEMS environment.

The RTEMS development effort uses an open development environment in which all users collaborate to improve RTEMS. The RTEMS cross development tool suite is based upon the free GNU tools and the open source standard C library newlib. RTEMS supports many host platforms and target architectures.

2.1 Free Software Project

The free software goals of the project are:

- RTEMS and supporting components are available under various free licenses with copyrights being held by individual authors.
- All software which executes on the target will not place undue restrictions on embedded applications. See also *Licensing Requirements* (page 239).
- Patches must be legally acceptable for inclusion into the RTEMS Project or the specific project being used.

2.2 Design and Development Goals

- Source based development with all users building from source
- Any suitable host should be supported
- Open testing, tests and test results
- Ports to new architectures and CPU models
- Addition of Board Support Packages for available hardware
- Improved runtime libraries
- Faster debug cycle
- Various other infrastructure improvements

2.3 Open Development Environment

- Encourage cooperation and communication between developers
- Work more closely with “consumers”
- Code available to everyone at any time, and everyone is welcome to participate in development
- Patches will be considered equally based on their technical merits
- All individuals and companies are welcome to contribute as long as they accept the ground rules
- Open mailing lists
- Developer friendly tools and procedures with a focus on keeping them current
- Conflicts of interest exist for many RTEMS developers. The developers contributing to the RTEMS Project must put the interests of the RTEMS Project first.

RTEMS STAKEHOLDERS

You are a potential RTEMS stakeholder. RTEMS is a community based free and open source project. All users are treated as stakeholders. It is hoped that as stakeholders, users will contribute to the project, sponsor core developers, and help fund the infrastructure required to host and manage the project. Please have a look at the *Support and Contributing* chapter of the **ERROR: ::url: `user`**.

INTRODUCTION TO PRE-QUALIFICATION

RTEMS has a long history of being used to support critical applications. In some of these application domains, there are standards (e.g., DO-178C, NPR 7150.2) which define the expectations for the processes used to develop software and the associated artifacts. These standards typically do not specify software functionality but address topics like requirements definition, traceability, having a documented change process, coding style, testing requirements, and a user's manual. During system test, these standards call for a review - usually by an independent entity - that the standard has been adhered to. These reviews cover a broad variety of topics and activities, but the process is generally referred to as qualification, verification, or auditing against the specific standard in use. The RTEMS Project will use the term "qualification" independent of the standard.

The goal of the RTEMS Qualification Project is to make RTEMS easier to review regardless of the standard chosen. Quite specifically, the RTEMS Qualification effort will NOT produce a directly qualified product or artifacts in the format dictated by a specific organization or standard. The goal is to make RTEMS itself, documentation, testing infrastructure, etc. more closely align with the information requirements of these high integrity qualification standards. In addition to improving the items that a mature, high quality open source project will have, there are additional artifacts needed for a qualification effort that no known open source project possesses. Specifically, requirements and the associated traceability to source code, tests, and documentation are needed.

The RTEMS Qualification Project is technically "pre-qualification." True qualification must be performed on the project's target hardware in a system context. The FAA has provided guidance for Reusable Software Components (FAA-AC20-148) and this effort should follow that guidance. The open RTEMS Project, with the assistance of domain experts, will possess and maintain the master technical information needed in a qualification effort. Consultants will provide the services required to tailor the master information, perform testing on specific system hardware, and to guide end users in using the master technical data in the context of a particular standard.

The RTEMS Qualification Project will broadly address two areas. The first area is suggesting areas of improvement for automated project infrastructure and the master technical data that has traditionally been provided by the RTEMS Project. For example, the RTEMS Qualification could suggest specific improvements to code coverage reports. The teams focused on qualification should be able to provide resources for improving the automated project infrastructure and master technical data for RTEMS. The term "resources" is often used by open source projects to refer to volunteer code contributions or funding. Although code contributions in this area are important and always welcome, funding is also important. At a minimum, ongoing funding is

needed for maintenance and upgrades of the RTEMS Project server infrastructure, addition of services to those servers, and core contributors to review submissions

The second area is the creation and maintenance of master technical data that has traditionally not been owned or maintained by the RTEMS Project. The most obvious example of this is a requirements set with proper infrastructure for tracing requirements through code to test and documentation. It is expected that these will be maintained by the RTEMS Qualification Project. They will be evaluated for adoption by the main RTEMS Project but the additional maintenance burden imposed will be a strong factor in this consideration. It behooves the RTEMS Qualification Project to limit dependence on manual checks and ensure that automation and ongoing support for that automation is contributed to the RTEMS Project.

It is expected that the RTEMS Qualification Project will create and maintain maps from the RTEMS master technical data to the various qualification standards. It will maintain “score-cards” which identify how the RTEMS Project is currently doing when reviewed per each standard. These will be maintained in the open as community resources which will guide the community in improving its infrastructure.

4.1 Stakeholder Involvement

Qualification of RTEMS is a specialized activity and only specific users of RTEMS will complete a formal qualification activity. The RTEMS Project cannot self-fund this entire activity and requires stakeholder to invest in an ongoing basis to ensure that any investment they make is maintained and viable in an ongoing basis. The RTEMS core developers view steady support of the qualification effort as necessary to continue to lower the overall costs of qualifying RTEMS.

SOFTWARE REQUIREMENTS ENGINEERING

Software engineering standards for critical software such as ECSS-E-ST-40C demand that software requirements for a software product are collected in a software requirements specification (technical specification in ECSS-E-ST-40C terms). They are usually derived from system requirements (requirements baseline in ECSS-E-ST-40C terms). RTEMS is designed as a reusable software product which can be utilized by application designers to ease the development of their applications. The requirements of the end system (system requirements) using RTEMS are only known to the application designer. RTEMS itself is developed by the RTEMS maintainers and they do not know the requirements of a particular end system in general. RTEMS is designed as a real-time operating system to meet typical system requirements for a wide range of applications. Its suitability for a particular application must be determined by the application designer based on the technical specification provided by RTEMS accompanied with performance data for a particular target platform.

Currently, no technical specification of RTEMS exists in the form of a dedicated document. Since the beginning of the RTEMS evolution in the late 1980s it was developed iteratively. It was never developed in a waterfall model. During initial development the RTEID [Mot88] and later the ORKID [VIT90] draft specifications were used as requirements. These were evolving during the development and an iterative approach was followed often using simple algorithms and coming back to optimise. In 1993 and 1994 a subset of pthreads sufficient to support *GNAT* was added as requirements. At this time the Ada tasking was defined, however, not implemented in *GNAT*, so this involved guessing during the development. Later some adjustments were made when Ada tasking was actually implemented. So, it was consciously iterative with the specifications evolving and feedback from performance analysis. Benchmarks published from other real time operating systems were used for comparison. Optimizations were carried out until the results were comparable. Development was done with distinct contractual phases and tasks for development, optimization, and the addition of priority inheritance and rate monotonic scheduling. The pthreads requirement has grown to be as much POSIX as possible.

Portability from FreeBSD to use its network stack, USB stack, SD/MMC card stack and device drivers resulted in another set of requirements. The addition of support for symmetric multiprocessing (SMP) was a huge driver for change. It was developed step by step and sponsored by several independent users with completely different applications and target platforms in mind. The high performance OpenMP support introduced the Futex as a new synchronization primitive.

Guidance

A key success element of RTEMS is the ability to accept changes driven by user needs and still keep the operating system stable enough for production systems. Procedures that place a high burden on changes are doomed to be discarded by the RTEMS Project. We have to keep this in mind when we introduce a requirements management work flow which should be followed by RTEMS community members and new contributors.

We have to put in some effort first into the reconstruction of software requirements through reverse engineering using the RTEMS documentation, test cases, sources, standard references, mailing list archives, etc. as input. Writing a technical specification for the complete RTEMS code base is probably a job of several person-years. We have to get started with a moderate feature set (e.g. subset of the Classic API) and extend it based on user demands step by step.

The development of the technical specification will take place in two phases. The first phase tries to establish an initial technical specification for an initial feature set. This technical specification will be integrated into RTEMS as a big chunk. In the second phase the technical specification is modified through arranged procedures. There will be procedures

- to modify existing requirements,
- add new requirements, and
- mark requirements as obsolete.

All procedures should be based on a peer review principles.

5.1 Requirements for Requirements

5.1.1 Identification

Each requirement shall have a unique identifier (UID). The question is in which scope should it be unique? Ideally, it should be universally unique. Therefore all UIDs used to link one specification item to another should use relative UIDs. This ensures that the RTEMS requirements can be referenced easily in larger systems though a system-specific prefix. The standard ECSS-E-ST-10-06C recommends in section 8.2.6 that the identifier should reflect the type of the requirement and the life profile situation. Other standards may have other recommendations. To avoid a bias of RTEMS in the direction of ECSS, this recommendation will not be followed.

The *absolute UID* of a specification item (for example a requirement) is defined by a leading / and the path of directories from the specification base directory to the file of the item separated by / characters and the file name without the .yml extension. For example, a specification item contained in the file build/cpukit/librtemscpu.yml inside a spec directory has the absolute UID of /build/cpukit/librtemscpu.

The *relative UID* to a specification item is defined by the path of directories from the file containing the source specification item to the file of the destination item separated by / characters and the file name of the destination item without the .yml extension. For example the relative UID from /build/bsps/sparc/leon3/grp to /build/bsps/bspopts is ../../bspopts.

Basically, the valid characters of an UID are determined by the file system storing the item files. By convention, UID characters shall be restricted to the following set defined by the regular expression [a-zA-Z0-9_-]+. Use - as a separator inside an UID part.

In documents the URL-like prefix spec: shall be used to indicated specification item UIDs.

The UID scheme for RTEMS requirements shall be component based. For example, the UID spec:/classic/task/create-err-invaddr may specify that the rtems_task_create() directive shall return a status of RTEMS_INVALID_ADDRESS if the id parameter is NULL.

A initial requirement item hierarchy could be this:

- build (building RTEMS BSPs and libraries)
- acfg (application configuration groups)
 - opt (application configuration options)
- classic
 - task
 - * create-* (requirements for rtems_task_create())
 - * delete-* (requirements for rtems_task_delete())
 - * exit-* (requirements for rtems_task_exit())
 - * getaff-* (requirements for rtems_task_get_affinity())
 - * getpri-* (requirements for rtems_task_get_priority())
 - * getsched-* (requirements for rtems_task_get_scheduler())
 - * ident-* (requirements for rtems_task_ident())
 - * issusp-* (requirements for rtems_task_is_suspended())

- * iter-* (requirements for `rtems_task_iterate()`)
 - * mode-* (requirements for `rtems_task_mode()`)
 - * restart-* (requirements for `rtems_task_restart()`)
 - * resume* (requirements for `rtems_task_resume()`)
 - * self* (requirements for `rtems_task_self()`)
 - * setaff-* (requirements for `rtems_task_set_affinity()`)
 - * setpri-* (requirements for `rtems_task_set_priority()`)
 - * setsched* (requirements for `rtems_task_set_scheduler()`)
 - * start-* (requirements for `rtems_task_start()`)
 - * susp-* (requirements for `rtems_task_suspend()`)
 - * wkafter-* (requirements for `rtems_task_wake_after()`)
 - * wkwhen-* (requirements for `rtems_task_wake_when()`)
- sema
- * ...
- posix
 - ...

A more detailed naming scheme and guidelines should be established. We have to find the right balance between the length of UIDs and self-descriptive UIDs. A clear scheme for all Classic API managers may help to keep the UIDs short and descriptive.

The specification of the validation of requirements should be maintained also by specification items. For each requirement directory there should be a validation subdirectory named *test*, e.g. `spec/classic/task/test`. A test specification directory may contain also validations by analysis, by inspection, and by design, see *Requirement Validation* (page 23).

5.1.2 Level of Requirements

The level of a requirement shall be expressed with one of the verbal forms listed below and nothing else. The level of requirements are derived from RFC 2119 [Bra97] and ECSS-E-ST-10-06C [ECS09].

5.1.2.1 Absolute Requirements

Absolute requirements shall be expressed with the verbal form *shall* and no other terms.

5.1.2.2 Absolute Prohibitions

Absolute prohibitions shall be expressed with the verbal form *shall not* and no other terms.

Warning: Absolute prohibitions may be difficult to validate. They should not be used.
--

5.1.2.3 Recommendations

Recommendations shall be expressed with the verbal forms *should* and *should not* and no other terms with guidance from RFC 2119:

SHOULD This word, or the adjective “RECOMMENDED”, mean that there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course.

SHOULD NOT This phrase, or the phrase “NOT RECOMMENDED” mean that there may exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

5.1.2.4 Permissions

Permissions shall be expressed with the verbal form *may* and no other terms with guidance from RFC 2119:

MAY This word, or the adjective “OPTIONAL”, mean that an item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because the vendor feels that it enhances the product while another vendor may omit the same item. An implementation which does not include a particular option **MUST** be prepared to interoperate with another implementation which does include the option, though perhaps with reduced functionality. In the same vein an implementation which does include a particular option **MUST** be prepared to interoperate with another implementation which does not include the option (except, of course, for the feature the option provides.)

5.1.2.5 Possibilities and Capabilities

Possibilities and capabilities shall be expressed with the verbal form *can* and no other terms.

5.1.3 Syntax

Use the Easy Approach to Requirements Syntax (*EARS*) to formulate requirements. A recommended reading list to get familiar with this approach is [MWHN09], [MW10], and [MWGU16]. Please also have a look at the EARS quick reference sheet [Uus12]. The sentence types are:

- Ubiquitous

The <system name> shall <system response>.

- Event-driven

When <optional preconditions> <trigger>, the <system name> shall <system response>.

- State-driven

While <in state>, the <system name> shall <system response>.

- Unwanted behaviour

If <optional preconditions> <trigger>, *then* the <system name> shall <system response>.

- Optional

Where <feature>, the <system name> shall <system response>.

The optional sentence type should be only used for application configuration options. The goal is to use the *enabled-by* attribute to enable or disable requirements based on configuration parameters that define the RTEMS artefacts used to build an application executable (header files, libraries, linker command files). Such configuration parameters are for example the architecture, the platform, CPU port options, and build configuration options (e.g. uniprocessor vs. SMP).

5.1.4 Wording Restrictions

To prevent the expression of imprecise requirements, the following terms shall not be used in requirement formulations:

- “acceptable”
- “adequate”
- “almost always”
- “and/or”
- “appropriate”
- “approximately”
- “as far as possible”
- “as much as practicable”
- “best”
- “best possible”
- “easy”

- “efficient”
- “e.g.”
- “enable”
- “enough”
- “etc.”
- “few”
- “first rate”
- “flexible”
- “generally”
- “goal”
- “graceful”
- “great”
- “greatest”
- “ideally”
- “i.e.”
- “if possible”
- “in most cases”
- “large”
- “many”
- “maximize”
- “minimize”
- “most”
- “multiple”
- “necessary”
- “numerous”
- “optimize”
- “ought to”
- “probably”
- “quick”
- “rapid”
- “reasonably”
- “relevant”
- “robust”
- “satisfactory”

- “several”
- “shall be included but not limited to”
- “simple”
- “small”
- “some”
- “state-of-the-art”.
- “sufficient”
- “suitable”
- “support”
- “systematically”
- “transparent”
- “typical”
- “user-friendly”
- “usually”
- “versatile”
- “when necessary”

For guidelines to avoid these terms see Table 11-2, “Some ambiguous terms to avoid in requirements” in [WB13]. There should be some means to enforce that these terms are not used, e.g. through a client-side pre-commit Git hook, a server-side pre-receive Git hook, or some scripts run by special build commands.

5.1.5 Separate Requirements

Requirements shall be stated separately. A bad example is:

spec:/classic/task/create

The task create directive shall evaluate the parameters, allocate a task object and initialize it.

To make this a better example, it should be split into separate requirements:

spec:/classic/task/create

When the task create directive is called with valid parameters and a free task object exists, the task create directive shall assign the identifier of an initialized task object to the `id` parameter and return the `RTEMS_SUCCESSFUL` status.

spec:/classic/task/create-err-toomany

If no free task objects exists, the task create directive shall return the `RTEMS_TOO_MANY` status.

spec:/classic/task/create-err-invaddr

If the `id` parameter is `NULL`, the task create directive shall return the `RTEMS_INVALID_ADDRESS` status.

spec:/classic/task/create-err-invname

If the `name` parameter is invalid, the task create directive shall return the `RTEMS_INVALID_NAME` status.

...

5.1.6 Conflict Free Requirements

Requirements shall not be in conflict with each other inside a specification. A bad example is:

spec:/classic/sema/mtx-obtain-wait

When a mutex is not available, the mutex obtain directive shall enqueue the calling thread on the wait queue of the mutex.

spec:/classic/sema/mtx-obtain-err-unsat

If a mutex is not available, the mutex obtain directive shall return the RTEMS_UNSATISFIED status.

To resolve this conflict, a condition may be added:

spec:/classic/sema/mtx-obtain-wait

When a mutex is not available and the RTEMS_WAIT option is set, the mutex obtain directive shall enqueue the calling thread on the wait queue of the mutex.

spec:/classic/sema/mtx-obtain-err-unsat

If a mutex is not available, when the RTEMS_WAIT option is not set, the mutex obtain directive shall return the RTEMS_UNSATISFIED status.

5.1.7 Use of Project-Specific Terms and Abbreviations

All project-specific terms and abbreviations used to formulate requirements shall be defined in the project glossary.

5.1.8 Justification of Requirements

Each requirement shall have a rationale or justification recorded in a dedicated section of the requirement file. See *rationale* attribute for *Specification Items* (page 25).

5.1.9 Requirement Validation

The validation of each *Requirement Item Type* (page 49) item shall be accomplished by one or more specification items of the types *Test Case Item Type* (page 60) or *Requirement Validation Item Type* (page 58) through a link from the validation item to the requirement item with the *Requirement Validation Link Role* (page 93).

Validation by test is strongly recommended. The choice of any other validation method shall be strongly justified. The requirements author is obligated to provide the means to validate the requirement with detailed instructions.

5.1.10 Resources and Performance

Normally, resource and performance requirements are formulated like this:

- The resource U shall need less than V storage units.
- The operation Y shall complete within X time units.

Such statements are difficult to make for a software product like RTEMS which runs on many different target platforms in various configurations. So, the performance requirements of RTEMS shall be stated in terms of benchmarks. The benchmarks are run on the project-specific target platform and configuration. The results obtained by the benchmark runs are reported in a human readable presentation. The application designer can then use the benchmark results to determine if its system performance requirements are met. The benchmarks shall be executed under different environment conditions, e.g. varying cache states (dirty, empty, valid) and system bus load generated by other processors. The application designer shall have the ability to add additional environment conditions, e.g. system bus load by DMA engines or different system bus arbitration schemes.

To catch resource and performance regressions via test suite runs there shall be a means to specify threshold values for the measured quantities. The threshold values should be provided for each validation platform. How this can be done and if the threshold values are maintained by the RTEMS Project is subject to discussion.

5.2 Specification Items

5.2.1 Specification Item Hierarchy

The specification item types have the following hierarchy:

- *Root Item Type* (page 26)
 - *Build Item Type* (page 27)
 - * *Build Ada Test Program Item Type* (page 28)
 - * *Build BSP Item Type* (page 29)
 - * *Build Configuration File Item Type* (page 31)
 - * *Build Configuration Header Item Type* (page 31)
 - * *Build Group Item Type* (page 32)
 - * *Build Library Item Type* (page 33)
 - * *Build Objects Item Type* (page 34)
 - * *Build Option Item Type* (page 35)
 - * *Build Script Item Type* (page 36)
 - * *Build Start File Item Type* (page 38)
 - * *Build Test Program Item Type* (page 38)
 - *Constraint Item Type* (page 40)
 - *Glossary Item Type* (page 40)
 - * *Glossary Group Item Type* (page 40)
 - * *Glossary Term Item Type* (page 40)
 - *Interface Item Type* (page 41)
 - * *Application Configuration Group Item Type* (page 41)
 - * *Application Configuration Option Item Type* (page 42)
 - *Application Configuration Feature Enable Option Item Type* (page 42)
 - *Application Configuration Feature Option Item Type* (page 43)
 - *Application Configuration Value Option Item Type* (page 43)
 - * *Interface Compound Item Type* (page 43)
 - * *Interface Container Item Type* (page 44)
 - * *Interface Define Item Type* (page 44)
 - * *Interface Domain Item Type* (page 44)
 - * *Interface Enum Item Type* (page 45)
 - * *Interface Enumerator Item Type* (page 45)
 - * *Interface Forward Declaration Item Type* (page 45)

- * *Interface Function Item Type* (page 46)
- * *Interface Group Item Type* (page 46)
- * *Interface Header File Item Type* (page 47)
- * *Interface Macro Item Type* (page 47)
- * *Interface Typedef Item Type* (page 48)
- * *Interface Unspecified Item Type* (page 48)
- * *Interface Variable Item Type* (page 49)
- *Requirement Item Type* (page 49)
 - * *Functional Requirement Item Type* (page 50)
 - *Action Requirement Item Type* (page 50)
 - *Generic Functional Requirement Item Type* (page 54)
 - * *Non-Functional Requirement Item Type* (page 54)
 - *Generic Non-Functional Requirement Item Type* (page 55)
 - *Runtime Performance Requirement Item Type* (page 55)
- *Requirement Validation Item Type* (page 58)
- *Runtime Measurement Test Item Type* (page 58)
- *Specification Item Type* (page 59)
- *Test Case Item Type* (page 60)
- *Test Platform Item Type* (page 61)
- *Test Procedure Item Type* (page 62)
- *Test Suite Item Type* (page 62)

5.2.2 Specification Item Types

5.2.2.1 Root Item Type

The technical specification of RTEMS will contain for example requirements, specializations of requirements, interface specifications, test suites, test cases, and requirement validations. These things will be called *specification items* or just *items* if it is clear from the context.

The specification items are stored in files in *YAML* format with a defined set of key-value pairs called attributes. Each attribute key name shall be a *Name* (page 89). In particular, key names which begin with an underscore (`_`) are reserved for internal use in tools.

This is the root specification item type. All explicit attributes shall be specified. The explicit attributes for this type are:

SPDX-License-Identifier

The attribute value shall be a *SPDX License Identifier* (page 95). It shall be the license of the item.

copyrights

The attribute value shall be a list. Each list element shall be a *Copyright* (page 77). It shall be the list of copyright statements of the item.

enabled-by

The attribute value shall be an *Enabled-By Expression* (page 77). It shall define the conditions under which the item is enabled.

links

The attribute value shall be a list. Each list element shall be a *Link* (page 88).

type

The attribute value shall be a *Name* (page 89). It shall be the item type. The selection of types and the level of detail depends on a particular standard and product model. We need enough flexibility to be in line with ECSS-E-ST-10-06 and possible future applications of other standards. The item type may be refined further with additional type-specific subtypes.

This type is refined by the following types:

- *Build Item Type* (page 27)
- *Constraint Item Type* (page 40)
- *Glossary Item Type* (page 40)
- *Interface Item Type* (page 41)
- *Requirement Item Type* (page 49)
- *Requirement Validation Item Type* (page 58)
- *Runtime Measurement Test Item Type* (page 58)
- *Specification Item Type* (page 59)
- *Test Case Item Type* (page 60)
- *Test Platform Item Type* (page 61)
- *Test Procedure Item Type* (page 62)
- *Test Suite Item Type* (page 62)

5.2.2.2 Build Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is `build`. This set of attributes specifies a build item. All explicit attributes shall be specified. The explicit attributes for this type are:

build-type

The attribute value shall be a *Name* (page 89). It shall be the build item type.

This type is refined by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build BSP Item Type* (page 29)
- *Build Configuration File Item Type* (page 31)
- *Build Configuration Header Item Type* (page 31)

- *Build Group Item Type* (page 32)
- *Build Library Item Type* (page 33)
- *Build Objects Item Type* (page 34)
- *Build Option Item Type* (page 35)
- *Build Script Item Type* (page 36)
- *Build Start File Item Type* (page 38)
- *Build Test Program Item Type* (page 38)

5.2.2.3 Build Ada Test Program Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is `ada-test-program`. This set of attributes specifies an Ada test program executable to build. Test programs may use additional objects provided by *Build Objects Item Type* (page 34) items. Test programs have an implicit `enabled-by` attribute value which is controlled by the option action *set-test-state* (page 35). If the test state is set to `exclude`, then the test program is not built. All explicit attributes shall be specified. The explicit attributes for this type are:

ada-main

The attribute value shall be a string. It shall be the path to the Ada main body file.

ada-object-directory

The attribute value shall be a string. It shall be the path to the Ada object directory (`-D` option value for `gnatmake`).

adaflags

The attribute value shall be a list of strings. It shall be a list of options for the Ada compiler.

adaincludes

The attribute value shall be a list of strings. It shall be a list of Ada include paths.

cflags

The attribute value shall be a list. Each list element shall be a *Build C Compiler Option* (page 67).

cppflags

The attribute value shall be a list. Each list element shall be a *Build C Preprocessor Option* (page 68).

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

ldflags

The attribute value shall be a list. Each list element shall be a *Build Linker Option* (page 70).

source

The attribute value shall be a list. Each list element shall be a *Build Source* (page 75).

stlib

The attribute value shall be a list. Each list element shall be a *Build Link Static Library Directive* (page 70).

target

The attribute value shall be a *Build Target* (page 76).

use-after

The attribute value shall be a list. Each list element shall be a *Build Use After Directive* (page 76).

use-before

The attribute value shall be a list. Each list element shall be a *Build Use Before Directive* (page 77).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 ada-main: testsuites/ada/samples/hello/hello.adb
3 ada-object-directory: testsuites/ada/samples/hello
4 adaflags: []
5 adaincludes:
6 - cpukit/include/adainclude
7 - testsuites/ada/support
8 build-type: ada-test-program
9 cflags: []
10 copyrights:
11 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
12 cppflags: []
13 enabled-by: true
14 includes: []
15 ldflags: []
16 links: []
17 source:
18 - testsuites/ada/samples/hello/init.c
19 stlib: []
20 target: testsuites/ada/ada_hello.exe
21 type: build
22 use-after: []
23 use-before: []

```

5.2.2.4 Build BSP Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is `bsp`. This set of attributes specifies a base BSP variant to build. All explicit attributes shall be specified. The explicit attributes for this type are:

arch

The attribute value shall be a string. It shall be the target architecture of the BSP.

bsp

The attribute value shall be a string. It shall be the base BSP variant name.

cflags

The attribute value shall be a list. Each list element shall be a *Build C Compiler Option* (page 67).

cppflags

The attribute value shall be a list. Each list element shall be a *Build C Preprocessor Option* (page 68).

family

The attribute value shall be a string. It shall be the BSP family name. The name shall be the last directory of the path to the BSP sources.

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

install

The attribute value shall be a list. Each list element shall be a *Build Install Directive* (page 69).

source

The attribute value shall be a list. Each list element shall be a *Build Source* (page 75).

Please have a look at the following example:

```
1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 arch: myarch
3 bsp: mybsp
4 build-type: bsp
5 cflags: []
6 copyrights:
7 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
8 cppflags: []
9 enabled-by: true
10 family: mybsp
11 includes: []
12 install:
13 - destination: ${BSP_INCLUDEDIR}
14   source:
15 - bsp/myarch/mybsp/include/bsp.h
16 - bsp/myarch/mybsp/include/tm27.h
17 - destination: ${BSP_INCLUDEDIR}/bsp
18   source:
19 - bsp/myarch/mybsp/include/bsp/irq.h
20 - destination: ${BSP_LIBDIR}
21   source:
22 - bsp/myarch/mybsp/start/linkcmds
23 links:
24 - role: build-dependency
25   uid: ../../obj
26 - role: build-dependency
27   uid: ../../opto2
28 - role: build-dependency
29   uid: abi
30 - role: build-dependency
31   uid: obj
32 - role: build-dependency
33   uid: ../start
34 - role: build-dependency
35   uid: ../../bspopts
36 source:
37 - bsp/myarch/mybsp/start/bspstart.c
38 type: build
```

5.2.2.5 Build Configuration File Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is `config-file`. This set of attributes specifies a configuration file placed in the build tree. The configuration file is generated during the `configure` command execution and is placed in the build tree. All explicit attributes shall be specified. The explicit attributes for this type are:

content

The attribute value shall be a string. It shall be the content of the configuration file. A `${VARIABLE}` substitution is performed during the `configure` command execution using the variables of the configuration set. Use `$$` for a plain `$` character. To have all variables from sibling items available for substitution it is recommended to link them in the proper order.

install-path

The attribute value shall be a *Build Install Path* (page 69).

target

The attribute value shall be a *Build Target* (page 76).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 build-type: config-file
3 content: |
4   # ...
5   Name: ${ARCH}-rtems${__RTEMS_MAJOR__}-${BSP_NAME}
6   # ...
7 copyrights:
8 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
9 enabled-by: true
10 install-path: ${PREFIX}/lib/pkgconfig
11 links: []
12 target: ${ARCH}-rtems${__RTEMS_MAJOR__}-${BSP_NAME}.pc
13 type: build

```

5.2.2.6 Build Configuration Header Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is `config-header`. This set of attributes specifies configuration header file. The configuration header file is generated during `configure` command execution and is placed in the build tree. All collected configuration defines are written to the configuration header file during the `configure` command execution. To have all configuration defines from sibling items available it is recommended to link them in the proper order. All explicit attributes shall be specified. The explicit attributes for this type are:

guard

The attribute value shall be a string. It shall be the header guard define.

include-headers

The attribute value shall be a list of strings. It shall be a list of header files to include via `#include <...>`.

install-path

The attribute value shall be a *Build Install Path* (page 69).

target

The attribute value shall be a *Build Target* (page 76).

5.2.2.7 Build Group Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is group. This set of attributes provides a means to aggregate other build items and modify the build item context which is used by referenced build items. The `includes`, `ldflags`, `objects`, and `use variables` of the build item context are updated by the corresponding attributes of the build group. All explicit attributes shall be specified. The explicit attributes for this type are:

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

install

The attribute value shall be a list. Each list element shall be a *Build Install Directive* (page 69).

ldflags

The attribute value shall be a list of strings. It shall be a list of options for the linker. They are used to link executables referenced by this item.

use-after

The attribute value shall be a list. Each list element shall be a *Build Use After Directive* (page 76).

use-before

The attribute value shall be a list. Each list element shall be a *Build Use Before Directive* (page 77).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 build-type: group
3 copyrights:
4 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
5 enabled-by:
6 - BUILD_TESTS
7 - BUILD_SAMPLES
8 includes:
9 - testsuites/support/include
10 install: []
11 ldflags:
12 - -Wl,--wrap=printf
13 - -Wl,--wrap=puts
14 links:
15 - role: build-dependency
16   uid: ticker
17 type: build
18 use-after: []
19 use-before:
20 - rtemstest

```

5.2.2.8 Build Library Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is `library`. This set of attributes specifies a static library. Library items may use additional objects provided by *Build Objects Item Type* (page 34) items through the build dependency links of the item. All explicit attributes shall be specified. The explicit attributes for this type are:

cflags

The attribute value shall be a list. Each list element shall be a *Build C Compiler Option* (page 67).

cppflags

The attribute value shall be a list. Each list element shall be a *Build C Preprocessor Option* (page 68).

cxxflags

The attribute value shall be a list. Each list element shall be a *Build C++ Compiler Option* (page 68).

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

install

The attribute value shall be a list. Each list element shall be a *Build Install Directive* (page 69).

install-path

The attribute value shall be a *Build Install Path* (page 69).

source

The attribute value shall be a list. Each list element shall be a *Build Source* (page 75).

target

The attribute value shall be a *Build Target* (page 76). It shall be the name of the static library, e.g. `z` for `libz.a`.

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 build-type: library
3 cflags:
4 - -Wno-pointer-sign
5 copyrights:
6 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
7 cppflags: []
8 cxxflags: []
9 enabled-by: true
10 includes:
11 - cpukit/libfs/src/jffs2/include
12 install:
13 - destination: ${BSP_INCLUDEDIR}/rtems
14   source:
15 - cpukit/include/rtems/jffs2.h
16 install-path: ${BSP_LIBDIR}
17 links: []
18 source:
19 - cpukit/libfs/src/jffs2/src/build.c
20 target: jffs2
21 type: build

```

5.2.2.9 Build Objects Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is objects. This set of attributes specifies a set of object files used to build static libraries or test programs. All explicit attributes shall be specified. The explicit attributes for this type are:

cflags

The attribute value shall be a list. Each list element shall be a *Build C Compiler Option* (page 67).

cppflags

The attribute value shall be a list. Each list element shall be a *Build C Preprocessor Option* (page 68).

cxxflags

The attribute value shall be a list. Each list element shall be a *Build C++ Compiler Option* (page 68).

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

install

The attribute value shall be a list. Each list element shall be a *Build Install Directive* (page 69).

source

The attribute value shall be a list. Each list element shall be a *Build Source* (page 75).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 build-type: objects
3 cflags: []
4 copyrights:
5 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
6 cppflags: []
7 cxxflags: []
8 enabled-by: true
9 includes: []
10 install:
11 - destination: ${BSP_INCLUDEDIR}/bsp
12   source:
13     - bsp/include/bsp/bootcard.h
14     - bsp/include/bsp/default-initial-extension.h
15     - bsp/include/bsp/fatal.h
16 links: []
17 source:
18 - bsp/shared/start/bootcard.c
19 - bsp/shared/rtems-version.c
20 type: build

```

5.2.2.10 Build Option Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is option. This set of attributes specifies a build option. The following explicit attributes are mandatory:

- `actions`
- `default`
- `default-by-variant`
- `description`

The explicit attributes for this type are:

actions

The attribute value shall be a list. Each list element shall be a *Build Option Action* (page 70). Each action operates on the *action value* handed over by a previous action and action-specific attribute values. The actions pass the processed action value to the next action in the list. The first action starts with an action value of None. The actions are carried out during the configure command execution.

default

The attribute value shall be a *Build Option Value* (page 75). It shall be the default value of the option if no variant-specific default value is specified. Use `null` to specify that no default value exists. The variant-specific default values may be specified by the `default-by-variant` attribute.

default-by-variant

The attribute value shall be a list. Each list element shall be a *Build Option Default by Variant* (page 74). The list is processed from top to bottom. If a matching variant is found, then the processing stops.

description

The attribute value shall be an optional string. It shall be the description of the option.

format

The attribute value shall be an optional string. It shall be a [Python format string](#), for example `'{ }'` or `'{: #010x}'`.

name

The attribute value shall be a *Build Option Name* (page 74).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 actions:
3 - get-integer: null
4 - define: null
5 build-type: option
6 copyrights:
7 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
8 default: 115200
9 default-by-variant:
10 - value: 9600
11 variants:
12 - m68k/m5484FireEngine

```

(continues on next page)

(continued from previous page)

```

13 - powerpc/hsc_cm01
14 - value: 19200
15   variants:
16     - m68k/COBRA5475
17   description: |
18     Default baud for console and other serial devices.
19   enabled-by: true
20   format: '{} '
21   links: []
22   name: BSP_CONSOLE_BAUD
23   type: build

```

5.2.2.11 Build Script Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is script. This set of attributes specifies a build script. The optional attributes may be required by commands executed through the scripts. The following explicit attributes are mandatory:

- `do-build`
- `do-configure`
- `prepare-build`
- `prepare-configure`

The explicit attributes for this type are:

asflags

The attribute value shall be a list. Each list element shall be a *Build Assembler Option* (page 67).

cflags

The attribute value shall be a list. Each list element shall be a *Build C Compiler Option* (page 67).

cppflags

The attribute value shall be a list. Each list element shall be a *Build C Preprocessor Option* (page 68).

cxxflags

The attribute value shall be a list. Each list element shall be a *Build C++ Compiler Option* (page 68).

do-build

The attribute value shall be an optional string. If this script shall execute, then it shall be Python code which is executed via `exec()` in the context of the `do_build()` method of the `wscript`. A local variable `bld` is available with the `waf` build context. A local variable `bic` is available with the build item context.

do-configure

The attribute value shall be an optional string. If this script shall execute, then it shall be Python code which is executed via `exec()` in the context of the `do_configure()` method of the `wscript`. A local variable `conf` is available with the `waf` configuration context. A local variable `cic` is available with the configuration item context.

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

ldflags

The attribute value shall be a list. Each list element shall be a *Build Linker Option* (page 70).

prepare-build

The attribute value shall be an optional string. If this script shall execute, then it shall be Python code which is executed via `exec()` in the context of the `prepare_build()` method of the `wscript`. A local variable `bld` is available with the `waf` build context. A local variable `bic` is available with the build item context.

prepare-configure

The attribute value shall be an optional string. If this script shall execute, then it shall be Python code which is executed via `exec()` in the context of the `prepare_configure()` method of the `wscript`. A local variable `conf` is available with the `waf` configuration context. A local variable `cic` is available with the configuration item context.

stlib

The attribute value shall be a list. Each list element shall be a *Build Link Static Library Directive* (page 70).

use-after

The attribute value shall be a list. Each list element shall be a *Build Use After Directive* (page 76).

use-before

The attribute value shall be a list. Each list element shall be a *Build Use Before Directive* (page 77).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 build-type: script
3 copyrights:
4 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
5 default: null
6 default-by-variant: []
7 do-build: |
8     bld.install_as(
9         "${BSP_LIBDIR}/linkcmds",
10        "bsps/" + bld.env.ARCH + "/" + bld.env.BSP_FAMILY +
11        "/start/linkcmds." + bld.env.BSP_BASE
12    )
13 do-configure: |
14     conf.env.append_value(
15         "LINKFLAGS",
16         ["-qolinkcmds", "-T", "linkcmds." + conf.env.BSP_BASE]
17    )
18 enabled-by: true
19 links: []
20 prepare-build: null
21 prepare-configure: null
22 type: build

```

5.2.2.12 Build Start File Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is `start-file`. This set of attributes specifies a start file to build. A start file is used to link an executable. All explicit attributes shall be specified. The explicit attributes for this type are:

asflags

The attribute value shall be a list. Each list element shall be a *Build Assembler Option* (page 67).

cppflags

The attribute value shall be a list. Each list element shall be a *Build C Preprocessor Option* (page 68).

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

install-path

The attribute value shall be a *Build Install Path* (page 69).

source

The attribute value shall be a list. Each list element shall be a *Build Source* (page 75).

target

The attribute value shall be a *Build Target* (page 76).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 asflags: []
3 build-type: start-file
4 copyrights:
5 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
6 cppflags: []
7 enabled-by: true
8 includes: []
9 install-path: ${BSP_LIBDIR}
10 links: []
11 source:
12 - bsp/sparc/shared/start/start.S
13 target: start.o
14 type: build

```

5.2.2.13 Build Test Program Item Type

This type refines the *Build Item Type* (page 27) though the `build-type` attribute if the value is `test-program`. This set of attributes specifies a test program executable to build. Test programs may use additional objects provided by *Build Objects Item Type* (page 34) items. Test programs have an implicit `enabled-by` attribute value which is controlled by the option action *set-test-state* (page 35). If the test state is set to `exclude`, then the test program is not built. All explicit attributes shall be specified. The explicit attributes for this type are:

cflags

The attribute value shall be a list. Each list element shall be a *Build C Compiler Option* (page 67).

cppflags

The attribute value shall be a list. Each list element shall be a *Build C Preprocessor Option* (page 68).

cxxflags

The attribute value shall be a list. Each list element shall be a *Build C++ Compiler Option* (page 68).

features

The attribute value shall be a string. It shall be the waf build features for this test program.

includes

The attribute value shall be a list. Each list element shall be a *Build Include Path* (page 68).

ldflags

The attribute value shall be a list. Each list element shall be a *Build Linker Option* (page 70).

source

The attribute value shall be a list. Each list element shall be a *Build Source* (page 75).

stlib

The attribute value shall be a list. Each list element shall be a *Build Link Static Library Directive* (page 70).

target

The attribute value shall be a *Build Target* (page 76).

use-after

The attribute value shall be a list. Each list element shall be a *Build Use After Directive* (page 76).

use-before

The attribute value shall be a list. Each list element shall be a *Build Use Before Directive* (page 77).

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 build-type: test-program
3 cflags: []
4 copyrights:
5 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
6 cppflags: []
7 cxxflags: []
8 enabled-by: true
9 features: c cprogram
10 includes: []
11 ldflags: []
12 links: []
13 source:
14 - testsuites/samples/ticker/init.c
15 - testsuites/samples/ticker/tasks.c
16 stlib: []
17 target: testsuites/samples/ticker.exe
18 type: build
19 use-after: []
20 use-before: []

```

5.2.2.14 Constraint Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is constraint. This set of attributes specifies a constraint. All explicit attributes shall be specified. The explicit attributes for this type are:

rationale

The attribute value shall be an optional string. If the value is present, then it shall state the rationale or justification of the constraint.

scope

The attribute value shall be a string. It shall be the scope of the constraint.

text

The attribute value shall be a *Requirement Text* (page 91). It shall state the constraint.

5.2.2.15 Glossary Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is glossary. This set of attributes specifies a glossary item. All explicit attributes shall be specified. The explicit attributes for this type are:

glossary-type

The attribute value shall be a *Name* (page 89). It shall be the glossary item type.

This type is refined by the following types:

- *Glossary Group Item Type* (page 40)
- *Glossary Term Item Type* (page 40)

5.2.2.16 Glossary Group Item Type

This type refines the *Glossary Item Type* (page 40) though the *glossary-type* attribute if the value is group. This set of attributes specifies a glossary group. All explicit attributes shall be specified. The explicit attributes for this type are:

name

The attribute value shall be a string. It shall be the human readable name of the glossary group.

text

The attribute value shall be a string. It shall state the requirement for the glossary group.

5.2.2.17 Glossary Term Item Type

This type refines the *Glossary Item Type* (page 40) though the *glossary-type* attribute if the value is term. This set of attributes specifies a glossary term. All explicit attributes shall be specified. The explicit attributes for this type are:

term

The attribute value shall be a string. It shall be the glossary term.

text

The attribute value shall be a string. It shall be the definition of the glossary term.

5.2.2.18 Interface Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is interface. This set of attributes specifies an interface specification item. Interface items shall specify the interface of the software product to other software products and the hardware. Use *Interface Domain Item Type* (page 44) items to specify interface domains, for example the *API*, C language, compiler, interfaces to the implementation, and the hardware. All explicit attributes shall be specified. The explicit attributes for this type are:

index-entries

The attribute value shall be a list of strings. It shall be a list of additional document index entries. A document index entry derived from the interface name is added automatically.

interface-type

The attribute value shall be a *Name* (page 89). It shall be the interface item type.

This type is refined by the following types:

- *Application Configuration Group Item Type* (page 41)
- *Application Configuration Option Item Type* (page 42)
- *Interface Compound Item Type* (page 43)
- *Interface Container Item Type* (page 44)
- *Interface Define Item Type* (page 44)
- *Interface Domain Item Type* (page 44)
- *Interface Enum Item Type* (page 45)
- *Interface Enumerator Item Type* (page 45)
- *Interface Forward Declaration Item Type* (page 45)
- *Interface Function Item Type* (page 46)
- *Interface Group Item Type* (page 46)
- *Interface Header File Item Type* (page 47)
- *Interface Macro Item Type* (page 47)
- *Interface Typedef Item Type* (page 48)
- *Interface Unspecified Item Type* (page 48)
- *Interface Variable Item Type* (page 49)

5.2.2.19 Application Configuration Group Item Type

This type refines the *Interface Item Type* (page 41) though the interface-type attribute if the value is `appl-config-group`. This set of attributes specifies an application configuration group. All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be a string. It shall be the description of the application configuration group.

name

The attribute value shall be a string. It shall be human readable name of the application configuration group.

text

The attribute value shall be a *Requirement Text* (page 91). It shall state the requirement for the application configuration group.

5.2.2.20 Application Configuration Option Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `appl-config-option`. This set of attributes specifies an application configuration option. All explicit attributes shall be specified. The explicit attributes for this type are:

appl-config-option-type

The attribute value shall be a *Name* (page 89). It shall be the application configuration option type.

description

The attribute value shall be an *Interface Description* (page 82). The *Application Configuration Value Option Item Type* (page 43) items have an attribute for constraints.

name

The attribute value shall be an *Application Configuration Option Name* (page 67).

notes

The attribute value shall be an *Interface Notes* (page 86).

text

The attribute value shall be a *Requirement Text* (page 91). It shall state the requirement for the application configuration option.

This type is refined by the following types:

- *Application Configuration Feature Enable Option Item Type* (page 42)
- *Application Configuration Feature Option Item Type* (page 43)
- *Application Configuration Value Option Item Type* (page 43)

5.2.2.21 Application Configuration Feature Enable Option Item Type

This type refines the *Application Configuration Option Item Type* (page 42) though the `appl-config-option-type` attribute if the value is `feature-enable`. This set of attributes specifies an application configuration feature enable option.

5.2.2.22 Application Configuration Feature Option Item Type

This type refines the *Application Configuration Option Item Type* (page 42) though the `appl-config-option-type` attribute if the value is `feature`. This set of attributes specifies an application configuration feature option. All explicit attributes shall be specified. The explicit attributes for this type are:

default

The attribute value shall be a string. It shall describe what happens if the configuration option is undefined.

5.2.2.23 Application Configuration Value Option Item Type

This type refines the following types:

- *Application Configuration Option Item Type* (page 42) though the `appl-config-option-type` attribute if the value is `initializer`
- *Application Configuration Option Item Type* (page 42) though the `appl-config-option-type` attribute if the value is `integer`

This set of attributes specifies application configuration initializer or integer option. All explicit attributes shall be specified. The explicit attributes for this type are:

constraints

The attribute value shall be an *Application Configuration Option Constraint Set* (page 66).

default-value

The attribute value shall be an *Integer or String* (page 78). It shall describe the default value of the application configuration option.

5.2.2.24 Interface Compound Item Type

This type refines the following types:

- *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `struct`
- *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `union`

This set of attributes specifies a compound (struct or union). All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition

The attribute value shall be a list. Each list element shall be an *Interface Compound Member Definition Directive* (page 81).

definition-kind

The attribute value shall be an *Interface Compound Definition Kind* (page 79).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the compound (struct or union).

notes

The attribute value shall be an *Interface Notes* (page 86).

5.2.2.25 Interface Container Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `container`. Items of this type specify an interface container. The item shall have exactly one link with the *Interface Placement Link Role* (page 87) to an *Interface Domain Item Type* (page 44) item. This link defines the interface domain of the container.

5.2.2.26 Interface Define Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `define`. This set of attributes specifies a `define`. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition

The attribute value shall be an *Interface Definition Directive* (page 82).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the `define`.

notes

The attribute value shall be an *Interface Notes* (page 86).

5.2.2.27 Interface Domain Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `domain`. This set of attributes specifies an interface domain. Items of the types *Interface Container Item Type* (page 44) and *Interface Header File Item Type* (page 47) are placed into domains through links with the *Interface Placement Link Role* (page 87). All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be a string. It shall be the description of the domain

name

The attribute value shall be a string. It shall be the human readable name of the domain.

5.2.2.28 Interface Enum Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `enum`. This set of attributes specifies an enum. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition-kind

The attribute value shall be an *Interface Enum Definition Kind* (page 84).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the enum.

notes

The attribute value shall be an *Interface Description* (page 82).

5.2.2.29 Interface Enumerator Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `enumerator`. This set of attributes specifies an enumerator. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition

The attribute value shall be an *Interface Definition Directive* (page 82).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the enumerator.

notes

The attribute value shall be an *Interface Notes* (page 86).

5.2.2.30 Interface Forward Declaration Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `forward-declaration`. Items of this type specify a forward declaration. The item shall have exactly one link with the *Interface Target Link Role* (page 88) to an *Interface Compound Item Type* (page 43) item. This link defines the type declared by the forward declaration.

5.2.2.31 Interface Function Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `function`. This set of attributes specifies a function. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition

The attribute value shall be an *Interface Function Definition Directive* (page 85).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the function.

notes

The attribute value shall be an *Interface Notes* (page 86).

params

The attribute value shall be a list. Each list element shall be an *Interface Parameter* (page 86).

return

The attribute value shall be an *Interface Return Directive* (page 87).

5.2.2.32 Interface Group Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `group`. This set of attributes specifies an interface group. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

description

The attribute value shall be an *Interface Description* (page 82).

identifier

The attribute value shall be an *Interface Group Identifier* (page 85).

name

The attribute value shall be a string. It shall be the human readable name of the interface group.

text

The attribute value shall be a *Requirement Text* (page 91). It shall state the requirement for the interface group.

5.2.2.33 Interface Header File Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `header-file`. This set of attributes specifies a header file. The item shall have exactly one link with the *Interface Placement Link Role* (page 87) to an *Interface Domain Item Type* (page 44) item. This link defines the interface domain of the header file. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

path

The attribute value shall be a string. It shall be the path used to include the header file. For example `rtems/confdefs.h`.

prefix

The attribute value shall be a string. It shall be the prefix directory path to the header file in the interface domain. For example `cpukit/include`.

5.2.2.34 Interface Macro Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `macro`. This set of attributes specifies a macro. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition

The attribute value shall be an *Interface Definition Directive* (page 82).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the macro.

notes

The attribute value shall be an *Interface Notes* (page 86).

params

The attribute value shall be a list. Each list element shall be an *Interface Parameter* (page 86).

return

The attribute value shall be an *Interface Return Directive* (page 87).

5.2.2.35 Interface Typedef Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `typedef`. This set of attributes specifies a typedef. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition

The attribute value shall be an *Interface Definition Directive* (page 82).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the typedef.

notes

The attribute value shall be an *Interface Notes* (page 86).

5.2.2.36 Interface Unspecified Item Type

This type refines the following types:

- *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `unspecified`
- *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `unspecified-define`
- *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `unspecified-function`
- *Interface Item Type* (page 41) though the `interface-type` attribute if the value is `unspecified-type`

This set of attributes specifies an unspecified interface. All explicit attributes shall be specified. The explicit attributes for this type are:

name

The attribute value shall be a string. It shall be the name of the unspecified interface.

reference

The attribute value shall be an optional string. If the value is present, then it shall be an URL to the standard or specification of the interface.

5.2.2.37 Interface Variable Item Type

This type refines the *Interface Item Type* (page 41) though the `interface-type` attribute if the value is variable. This set of attributes specifies a variable. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

definition

The attribute value shall be an *Interface Definition Directive* (page 82).

description

The attribute value shall be an *Interface Description* (page 82).

name

The attribute value shall be a string. It shall be the name of the variable.

notes

The attribute value shall be an *Interface Notes* (page 86).

5.2.2.38 Requirement Item Type

This type refines the *Root Item Type* (page 26) though the `type` attribute if the value is requirement. This set of attributes specifies a requirement. All explicit attributes shall be specified. The explicit attributes for this type are:

rationale

The attribute value shall be an optional string. If the value is present, then it shall state the rationale or justification of the requirement.

references

The attribute value shall be a list. Each list element shall be a *Requirement Reference* (page 90).

requirement-type

The attribute value shall be a *Name* (page 89). It shall be the requirement item type.

text

The attribute value shall be a *Requirement Text* (page 91). It shall state the requirement.

This type is refined by the following types:

- *Functional Requirement Item Type* (page 50)
- *Non-Functional Requirement Item Type* (page 54)

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 copyrights:
3 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de
4 enabled-by: true
5 functional-type: capability
6 links: []
7 rationale: |
8   It keeps you busy.
9 requirement-type: functional
10 text: |

```

(continues on next page)

(continued from previous page)

```

11 The system shall do crazy things.
12 type: requirement

```

5.2.2.39 Functional Requirement Item Type

This type refines the *Requirement Item Type* (page 49) though the requirement-type attribute if the value is functional. This set of attributes specifies a functional requirement. All explicit attributes shall be specified. The explicit attributes for this type are:

functional-type

The attribute value shall be a *Name* (page 89). It shall be the functional type of the requirement.

This type is refined by the following types:

- *Action Requirement Item Type* (page 50)
- *Generic Functional Requirement Item Type* (page 54)

5.2.2.40 Action Requirement Item Type

This type refines the *Functional Requirement Item Type* (page 50) though the functional-type attribute if the value is action. This set of attributes specifies functional requirements and corresponding validation test code. The functional requirements of an action are specified. An action performs a step in a finite state machine. An action is implemented through a function or a macro. The action is performed through a call of the function or an execution of the code of a macro expansion by an actor. The actor is for example a task or an interrupt service routine.

For action requirements which specify the function of an interface, there shall be exactly one link with the *Interface Function Link Role* (page 85) to the interface of the action.

The action requirements are specified by

- a list of pre-conditions, each with a set of states,
- a list of post-conditions, each with a set of states,
- the transition of pre-condition states to post-condition states through the action.

Along with the requirements, the test code to generate a validation test is specified. For an action requirement it is verified that all variations of pre-condition states have a set of post-condition states specified in the transition map. All transitions are covered by the generated test code. All explicit attributes shall be specified. The explicit attributes for this type are:

post-conditions

The attribute value shall be a list. Each list element shall be an *Action Requirement Condition* (page 63).

pre-conditions

The attribute value shall be a list. Each list element shall be an *Action Requirement Condition* (page 63).

skip-reasons

The attribute value shall be an *Action Requirement Skip Reasons* (page 64).

test-action

The attribute value shall be a string. It shall be the test action code.

test-brief

The attribute value shall be an optional string. If the value is present, then it shall be the test case brief description.

test-cleanup

The attribute value shall be an optional string. If the value is present, then it shall be the test cleanup code. The code is placed in the test action loop body after the test post-condition checks.

test-context

The attribute value shall be a list. Each list element shall be a *Test Context Member* (page 104).

test-context-support

The attribute value shall be an optional string. If the value is present, then it shall be the test context support code. The context support code is placed at file scope before the test context definition.

test-description

The attribute value shall be an optional string. If the value is present, then it shall be the test case description.

test-header

The attribute value shall be a *Test Header* (page 105).

test-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include <...>`.

test-local-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include "..."`.

test-prepare

The attribute value shall be an optional string. If the value is present, then it shall be the early test preparation code. The code is placed in the test action loop body before the test pre-condition preparations.

test-setup

The attribute value shall be a *Test Support Method* (page 106).

test-stop

The attribute value shall be a *Test Support Method* (page 106).

test-support

The attribute value shall be an optional string. If the value is present, then it shall be the test case support code. The support code is placed at file scope before the test case code.

test-target

The attribute value shall be a string. It shall be the path to the generated test case source file.

test-teardown

The attribute value shall be a *Test Support Method* (page 106).

transition-map

The attribute value shall be a list. Each list element shall be an *Action Requirement Transition* (page 64).

Please have a look at the following example:

```
1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 copyrights:
3 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
4 enabled-by: true
5 functional-type: action
6 links: []
7 post-conditions:
8 - name: Status
9   states:
10 - name: Success
11   test-code: |
12     /* Check that the status is SUCCESS */
13   text: |
14     The status shall be SUCCESS.
15 - name: Error
16   test-code: |
17     /* Check that the status is ERROR */
18   text: |
19     The status shall be ERROR.
20 test-epilogue: null
21 test-prologue: null
22 - name: Data
23   states:
24 - name: Unchanged
25   test-code: |
26     /* Check that the data is unchanged */
27   text: |
28     The data shall be unchanged by the action.
29 - name: Red
30   test-code: |
31     /* Check that the data is red */
32   text: |
33     The data shall be red.
34 - name: Green
35   test-code: |
36     /* Check that the data is green */
37   text: |
38     The data shall be green.
39 test-epilogue: null
40 test-prologue: null
41 pre-conditions:
42 - name: Data
43   states:
44 - name: NullPtr
45   test-code: |
46     /* Set data pointer to NULL */
47   text: |
48     The data pointer shall be NULL.
49 - name: Valid
50   test-code: |
51     /* Set data pointer to reference a valid data buffer */
52   text: |
53     The data pointer shall reference a valid data buffer.
54 test-epilogue: null
55 test-prologue: null
```

(continues on next page)

(continued from previous page)

```

56 - name: Option
57 states:
58 - name: Red
59   test-code: |
60     /* Set option to RED */
61   text: |
62     The option shall be RED.
63 - name: Green
64   test-code: |
65     /* Set option to GREEN */
66   text: |
67     The option shall be GREEN.
68 test-epilogue: null
69 test-prologue: null
70 requirement-type: functional
71 skip-reasons: {}
72 test-action: |
73   /* Call the function of the action */
74 test-brief: null
75 test-cleanup: null
76 test-context:
77 - brief: null
78   description: null
79   member: void *data
80 - brief: null
81   description: null
82   member: option_type option
83 test-context-support: null
84 test-description: null
85 test-header: null
86 test-includes: []
87 test-local-includes: []
88 test-prepare: null
89 test-setup: null
90 test-stop: null
91 test-support: null
92 test-target: tc-red-green-data.c
93 test-teardown: null
94 transition-map:
95 - enabled-by: true
96   post-conditions:
97     Status: Error
98     Data: Unchanged
99   pre-conditions:
100     Data: NullPtr
101     Option: all
102 - enabled-by: true
103   post-conditions:
104     Status: Success
105     Data: Red
106   pre-conditions:
107     Data: Valid
108     Option: Red
109 - enabled-by: true
110   post-conditions:
111     Status: Success

```

(continues on next page)

(continued from previous page)

```

112   Data: Green
113   pre-conditions:
114     Data: Valid
115     Option: Green
116   rationale: null
117   references: []
118   text: |
119     ${./text-template}
120   type: requirement

```

5.2.2.41 Generic Functional Requirement Item Type

This type refines the following types:

- *Functional Requirement Item Type* (page 50) though the functional-type attribute if the value is capability
- *Functional Requirement Item Type* (page 50) though the functional-type attribute if the value is dependability-function
- *Functional Requirement Item Type* (page 50) though the functional-type attribute if the value is function
- *Functional Requirement Item Type* (page 50) though the functional-type attribute if the value is operational
- *Functional Requirement Item Type* (page 50) though the functional-type attribute if the value is safety-function

Items of this type state a functional requirement with the functional type defined by the specification type refinement.

5.2.2.42 Non-Functional Requirement Item Type

This type refines the *Requirement Item Type* (page 49) though the requirement-type attribute if the value is non-functional. This set of attributes specifies a non-functional requirement. All explicit attributes shall be specified. The explicit attributes for this type are:

non-functional-type

The attribute value shall be a *Name* (page 89). It shall be the non-functional type of the requirement.

This type is refined by the following types:

- *Generic Non-Functional Requirement Item Type* (page 55)
- *Runtime Performance Requirement Item Type* (page 55)

5.2.2.43 Generic Non-Functional Requirement Item Type

This type refines the following types:

- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is build-configuration
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is constraint
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is design
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is documentation
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is interface
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is interface-requirement
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is maintainability
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is performance
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is portability
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is quality
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is reliability
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is resource
- *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is safety

Items of this type state a non-functional requirement with the non-functional type defined by the specification type refinement.

5.2.2.44 Runtime Performance Requirement Item Type

This type refines the *Non-Functional Requirement Item Type* (page 54) though the non-functional-type attribute if the value is performance-runtime. The item shall have exactly one link with the *Runtime Measurement Request Link Role* (page 94). A requirement text processor shall support a substitution of `${./limit-kind}`:

- For a *Runtime Measurement Value Kind* (page 94) of min-lower-bound or min-upper-bound, the substitution of `${./limit-kind}` shall be "minimum".
- For a *Runtime Measurement Value Kind* (page 94) of mean-lower-bound or mean-upper-bound, the substitution of `${./limit-kind}` shall be "mean".

- For a *Runtime Measurement Value Kind* (page 94) of `max-lower-bound` or `max-upper-bound`, the substitution of `${./limit-kind}` shall be "maximum".

A requirement text processor shall support a substitution of `${./limit-condition}`:

- For a *Runtime Measurement Value Kind* (page 94) of `min-lower-bound`, `mean-lower-bound`, or `max-lower-bound`, the substitution of `${./limit-condition}` shall be "greater than or equal to <value>" with <value> being the value of the corresponding entry in the *Runtime Measurement Value Table* (page 95).
- For a *Runtime Measurement Value Kind* (page 94) of `min-upper-bound`, `mean-upper-bound`, or `max-upper-bound`, the substitution of `${./limit-condition}` shall be "less than or equal to <value>" with <value> being the value of the corresponding entry in the *Runtime Measurement Value Table* (page 95).

A requirement text processor shall support a substitution of `${./environment}`. The value of the substitution shall be "<environment> environment" with <environment> being the environment of the corresponding entry in the *Runtime Measurement Environment Table* (page 94).

This set of attributes specifies non-functional runtime performance requirements and the corresponding validation test code. Along with the requirement, the test code to execute a measure runtime request is specified. All explicit attributes shall be specified. The explicit attributes for this type are:

limits

The attribute value shall be a *Runtime Performance Limit Table* (page 95).

params

The attribute value shall be a *Runtime Performance Parameter Set* (page 95).

test-body

The attribute value shall be a *Test Support Method* (page 106). It shall provide the code of the measure runtime body handler. In contrast to other methods, this method is mandatory.

test-cleanup

The attribute value shall be a *Test Support Method* (page 106). It may provide the code to clean up the measure runtime request. This method is called before the cleanup method of the corresponding *Runtime Measurement Test Item Type* (page 58) item and after the request.

test-prepare

The attribute value shall be a *Test Support Method* (page 106). It may provide the code to prepare the measure runtime request. This method is called after the prepare method of the corresponding *Runtime Measurement Test Item Type* (page 58) item and before the request.

test-setup

The attribute value shall be a *Test Support Method* (page 106). It may provide the code of the measure runtime setup handler.

test-teardown

The attribute value shall be a *Test Support Method* (page 106). It may provide the code of the measure runtime teardown handler.

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 copyrights:
3 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
4 enabled-by: true

```

(continues on next page)

(continued from previous page)

```

5 links:
6 - role: runtime-measurement-request
7   uid: ../val/performance
8 limits:
9   sparc/leon3:
10     DirtyCache:
11       max-upper-bound: 0.000005
12       mean-upper-bound: 0.000005
13     FullCache:
14       max-upper-bound: 0.000005
15       mean-upper-bound: 0.000005
16     HotCache:
17       max-upper-bound: 0.000005
18       mean-upper-bound: 0.000005
19     Load/1:
20       max-upper-bound: 0.00001
21       mean-upper-bound: 0.00001
22     Load/2:
23       max-upper-bound: 0.00001
24       mean-upper-bound: 0.00001
25     Load/3:
26       max-upper-bound: 0.00001
27       mean-upper-bound: 0.00001
28     Load/4:
29       max-upper-bound: 0.00001
30       mean-upper-bound: 0.00001
31 params: {}
32 rationale: null
33 references: []
34 test-body:
35   brief: |
36     Get a buffer.
37   code: |
38     ctx->status = rtems_partition_get_buffer( ctx->part_many, &ctx->buffer );
39   description: null
40 test-cleanup: null
41 test-prepare: null
42 test-setup: null
43 test-teardown:
44   brief: |
45     Return the buffer.
46   code: |
47     rtems_status_code sc;
48
49     T_quiet_rsc_success( ctx->status );
50
51     sc = rtems_partition_return_buffer( ctx->part_many, ctx->buffer );
52     T_quiet_rsc_success( sc );
53
54     return tic == toc;
55   description: null
56 text: |
57   When a partition has exactly ${../val/performance:/params/buffer-count} free
58   buffers, the ${.:limit-kind} runtime of exactly
59   ${../val/performance:/params/sample-count} successful calls to
60   ${../if/get-buffer:/name} in the ${../environment} shall be

```

(continues on next page)

(continued from previous page)

```

61  ${.:limit-condition}.
62  non-functional-type: performance-runtime
63  requirement-type: non-functional
64  type: requirement

```

5.2.2.45 Requirement Validation Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is validation. This set of attributes provides a requirement validation evidence. The item shall have exactly one link to the validated requirement with the *Requirement Validation Link Role* (page 93). All explicit attributes shall be specified. The explicit attributes for this type are:

method

The attribute value shall be a *Requirement Validation Method* (page 93). Validation by test is done through *Test Case Item Type* (page 60) items.

text

The attribute value shall be a string. It shall provide the validation evidence depending on the validation method:

- *By analysis*: A statement shall be provided how the requirement is met, by analysing static properties of the *software product*.
- *By inspection*: A statement shall be provided how the requirement is met, by inspection of the *source code*.
- *By review of design*: A rationale shall be provided to demonstrate how the requirement is satisfied implicitly by the software design.

5.2.2.46 Runtime Measurement Test Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is runtime-measurement-test. This set of attributes specifies a runtime measurement test case. All explicit attributes shall be specified. The explicit attributes for this type are:

params

The attribute value shall be a *Runtime Measurement Parameter Set* (page 94).

test-brief

The attribute value shall be an optional string. If the value is present, then it shall be the test case brief description.

test-cleanup

The attribute value shall be a *Test Support Method* (page 106). If the value is present, then it shall be the measure runtime request cleanup method. The method is called after each measure runtime request.

test-context

The attribute value shall be a list. Each list element shall be a *Test Context Member* (page 104).

test-context-support

The attribute value shall be an optional string. If the value is present, then it shall be the test context support code. The context support code is placed at file scope before the test context definition.

test-description

The attribute value shall be an optional string. If the value is present, then it shall be the test case description.

test-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include <...>`.

test-local-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include "..."`.

test-prepare

The attribute value shall be a *Test Support Method* (page 106). If the value is present, then it shall be the measure runtime request prepare method. The method is called before each measure runtime request.

test-setup

The attribute value shall be a *Test Support Method* (page 106). If the value is present, then it shall be the test case setup fixture method.

test-stop

The attribute value shall be a *Test Support Method* (page 106). If the value is present, then it shall be the test case stop fixture method.

test-support

The attribute value shall be an optional string. If the value is present, then it shall be the test case support code. The support code is placed at file scope before the test case code.

test-target

The attribute value shall be a string. It shall be the path to the generated test case source file.

test-teardown

The attribute value shall be a *Test Support Method* (page 106). If the value is present, then it shall be the test case teardown fixture method.

5.2.2.47 Specification Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is spec. This set of attributes specifies specification types. All explicit attributes shall be specified. The explicit attributes for this type are:

spec-description

The attribute value shall be an optional string. It shall be the description of the specification type.

spec-example

The attribute value shall be an optional string. If the value is present, then it shall be an example of the specification type.

spec-info

The attribute value shall be a *Specification Information* (page 99).

spec-name

The attribute value shall be an optional string. It shall be the human readable name of the specification type.

spec-type

The attribute value shall be a *Name* (page 89). It shall the specification type.

Please have a look at the following example:

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 copyrights:
3 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
4 enabled-by: true
5 links:
6 - role: spec-member
7   uid: root
8 - role: spec-refinement
9   spec-key: type
10  spec-value: example
11  uid: root
12 spec-description: null
13 spec-example: null
14 spec-info:
15   dict:
16     attributes:
17       an-example-attribute:
18         description: |
19           It shall be an example.
20         spec-type: optional-str
21       example-number:
22         description: |
23           It shall be the example number.
24         spec-type: int
25       description: |
26         This set of attributes specifies an example.
27       mandatory-attributes: all
28 spec-name: Example Item Type
29 spec-type: spec
30 type: spec

```

5.2.2.48 Test Case Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is test-case. This set of attributes specifies a test case. All explicit attributes shall be specified. The explicit attributes for this type are:

test-actions

The attribute value shall be a list. Each list element shall be a *Test Case Action* (page 104).

test-brief

The attribute value shall be a string. It shall be the test case brief description.

test-description

The attribute value shall be an optional string. It shall be the test case description.

test-epilogue

The attribute value shall be an optional string. If the value is present, then it shall be the test case epilogue code. The epilogue code is placed in the test case body after the test case actions.

test-fixture

The attribute value shall be an optional string. If the value is present, then it shall be a pointer to the test case fixture. The test case fixture pointer declaration may be provided by the test case support code or via an included header file.

test-header

The attribute value shall be a *Test Header* (page 105).

test-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include <...>`.

test-local-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include "..."`.

test-prologue

The attribute value shall be an optional string. If the value is present, then it shall be the test case prologue code. The prologue code is placed in the test case body before the test case actions. A use case is the declaration of local variables used by the test case actions.

test-support

The attribute value shall be an optional string. If the value is present, then it shall be the test case support code. The support code is placed at file scope before the test case code.

test-target

The attribute value shall be a string. It shall be the path to the generated target test case source file.

5.2.2.49 Test Platform Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is `test-platform`. Please note:

Warning: This item type is work in progress.

This set of attributes specifies a test platform. All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be a string. It shall be the description of the test platform.

name

The attribute value shall be a string. It shall be the human readable name of the test platform.

5.2.2.50 Test Procedure Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is `test-procedure`. Please note:

Warning: This item type is work in progress.

This set of attributes specifies a test procedure. All explicit attributes shall be specified. The explicit attributes for this type are:

name

The attribute value shall be a string. It shall be the human readable name of the test procedure.

purpose

The attribute value shall be a string. It shall state the purpose of the test procedure.

steps

The attribute value shall be a string. It shall describe the steps of the test procedure execution.

5.2.2.51 Test Suite Item Type

This type refines the *Root Item Type* (page 26) though the type attribute if the value is `test-suite`. This set of attributes specifies a test suite. All explicit attributes shall be specified. The explicit attributes for this type are:

test-brief

The attribute value shall be a string. It shall be the test suite brief description.

test-code

The attribute value shall be a string. It shall be the test suite code. The test suite code is placed at file scope in the target source file.

test-description

The attribute value shall be an optional string. It shall be the test suite description.

test-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include <...>`.

test-local-includes

The attribute value shall be a list of strings. It shall be a list of header files included via `#include "..."`.

test-target

The attribute value shall be a string. It shall be the path to the generated target test suite source file.

5.2.3 Specification Attribute Sets and Value Types

5.2.3.1 Action Requirement Condition

This set of attributes defines an action pre-condition or post-condition. All explicit attributes shall be specified. The explicit attributes for this type are:

name

The attribute value shall be an *Action Requirement Name* (page 63).

states

The attribute value shall be a list. Each list element shall be an *Action Requirement State* (page 64).

test-epilogue

The attribute value shall be an optional string. If the value is present, then it shall be the test epilogue code. The epilogue code is placed in the test condition preparation or check before the state-specific code. The code may use a local variable `ctx` which points to the test context, see *Test Context Member* (page 104).

test-prologue

The attribute value shall be an optional string. If the value is present, then it shall be the test prologue code. The prologue code is placed in the test condition preparation or check after the state-specific code. The code may use a local variable `ctx` which points to the test context, see *Test Context Member* (page 104).

This type is used by the following types:

- *Action Requirement Item Type* (page 50)

5.2.3.2 Action Requirement Name

The value shall be a string. It shall be the name of a condition or a state of a condition used to define pre-conditions and post-conditions of an action requirement. It shall be formatted in CamelCase. It should be brief and abbreviated. The rationale for this is that the names are used in tables and the horizontal space is limited by the page width. The more conditions you have in an action requirement, the shorter the names should be. The name NA is reserved and indicates that a condition is not applicable. The value

- shall match with the regular expression “`^[A-Z][a-zA-Z0-9]+$`”,
- and, shall be not equal to “NA”.

This type is used by the following types:

- *Action Requirement Condition* (page 63)
- *Action Requirement Skip Reasons* (page 64)
- *Action Requirement State* (page 64)
- *Action Requirement Transition Post-Conditions* (page 65)
- *Action Requirement Transition Pre-Conditions* (page 66)

5.2.3.3 Action Requirement Skip Reasons

This set of attributes specifies skip reasons used to justify why transitions in the transition map are skipped. Generic attributes may be specified. Each generic attribute key shall be an *Action Requirement Name* (page 63). Each generic attribute value shall be a string. The key defines the name of a skip reason. The name can be used in *Action Requirement Transition Post-Conditions* (page 65) to skip the corresponding transitions. The value shall give a reason why the transitions are skipped.

This type is used by the following types:

- *Action Requirement Item Type* (page 50)

5.2.3.4 Action Requirement State

This set of attributes defines an action pre-condition or post-condition state. All explicit attributes shall be specified. The explicit attributes for this type are:

name

The attribute value shall be an *Action Requirement Name* (page 63).

test-code

The attribute value shall be a string. It shall be the test code to prepare or check the state of the condition. The code may use a local variable `ctx` which points to the test context, see *Test Context Member* (page 104).

text

The attribute value shall be a *Requirement Text* (page 91). It shall define the state of the condition.

This type is used by the following types:

- *Action Requirement Condition* (page 63)

5.2.3.5 Action Requirement Transition

This set of attributes defines the transition from multiple sets of states of pre-conditions to a set of states of post-conditions through an action in an action requirement. The ability to specify multiple sets of states of pre-conditions which result in a common set of post-conditions may allow a more compact specification of the transition map. For example, let us suppose you want to specify the action of a function with a pointer parameter. The function performs an early check that the pointer is NULL and in this case returns an error code. The pointer condition dominates the action outcome if the pointer is NULL. Other pre-condition states can be simply set to `all` for this transition. All explicit attributes shall be specified. The explicit attributes for this type are:

enabled-by

The attribute value shall be an *Enabled-By Expression* (page 77). The transition map may be customized to support configuration variants through this attribute. The default transitions (`enabled-by: true`) shall be specified before the customized variants in the list.

post-conditions

The attribute value shall be an *Action Requirement Transition Post-Conditions* (page 65).

pre-conditions

The attribute value shall be an *Action Requirement Transition Pre-Conditions* (page 66).

This type is used by the following types:

- *Action Requirement Item Type* (page 50)

5.2.3.6 Action Requirement Transition Post-Conditions

A value of this type shall be of one of the following variants:

- The value may be a set of attributes. This set of attributes defines for each post-condition the state after the action for a transition in an action requirement. Generic attributes may be specified. Each generic attribute key shall be an *Action Requirement Name* (page 63). Each generic attribute value shall be an *Action Requirement Name* (page 63). There shall be exactly one generic attribute key for each post-condition. The key name shall be the post-condition name. The value of each generic attribute shall be the state of the post-condition.
- The value may be a string. It shall be the name of a skip reason. If a skip reason is given instead of a listing of post-condition states, then this transition is skipped and no test code runs for this transition. The value
 - shall match with the regular expression “`^[A-Z][a-zA-Z0-9]+$`”,
 - and, shall be not equal to “NA”.

This type is used by the following types:

- *Action Requirement Transition* (page 64)

5.2.3.7 Action Requirement Transition Pre-Condition State Set

A value of this type shall be of one of the following variants:

- The value may be a list. Each list element shall be an *Action Requirement Name* (page 63). The list defines the set of states of the pre-condition in the transition.
- The value may be a string. The value `all` represents all states of the pre-condition in this transition. The value `N/A` marks the pre-condition as not applicable in this transition. The value shall be an element of
 - “all”, and
 - “N/A”.

This type is used by the following types:

- *Action Requirement Transition Pre-Conditions* (page 66)

5.2.3.8 Action Requirement Transition Pre-Conditions

This set of attributes defines for each pre-condition the set of states before the action for a transition in an actin requirement. Generic attributes may be specified. Each generic attribute key shall be an *Action Requirement Name* (page 63). Each generic attribute value shall be an *Action Requirement Transition Pre-Condition State Set* (page 65). There shall be exactly one generic attribute key for each pre-condition. The key name shall be the pre-condition name. The value of each generic attribute shall be a set of states of the pre-condition.

This type is used by the following types:

- *Action Requirement Transition* (page 64)

5.2.3.9 Application Configuration Group Member Link Role

This type refines the *Link* (page 88) though the role attribute if the value is `appl-config-group-member`. It defines the application configuration group membership role of links.

5.2.3.10 Application Configuration Option Constraint Set

This set of attributes defines application configuration option constraints. Additional constraints can be added through the links of the item using the *Constraint Link Role* (page 77). None of the explicit attributes is mandatory, they are all optional. The explicit attributes for this type are:

max

The attribute value shall be an *Integer or String* (page 78). It shall be the maximum value of the application configuration option.

min

The attribute value shall be an *Integer or String* (page 78). It shall be the minimum value of the application configuration option.

set

The attribute value shall be a list. Each list element shall be an *Integer or String* (page 78). It shall be the set of valid values for the application configuration option.

texts

The attribute value shall be a list. Each list element shall be a *Requirement Text* (page 91). It shall be a list of constraints specific to this application configuration option. For general constraints, use a link with the *Constraint Link Role* (page 77) to a constraint item.

This type is used by the following types:

- *Application Configuration Value Option Item Type* (page 43)

5.2.3.11 Application Configuration Option Name

The value shall be a string. It shall be the name of an application configuration option. The value shall match with the regular expression “`^(CONFIGURE_|BSP_)[A-Z0-9_]+$`”.

This type is used by the following types:

- *Application Configuration Option Item Type* (page 42)

5.2.3.12 Boolean or Integer or String

A value of this type shall be of one of the following variants:

- The value may be a boolean.
- The value may be an integer number.
- The value may be a string.

This type is used by the following types:

- *Build Option Action* (page 70)
- *Interface Return Value* (page 88)

5.2.3.13 Build Assembler Option

The value shall be a string. It shall be an option for the assembler. The options are used to assemble the sources of this item. The options defined by this attribute succeed the options presented to the item by the build item context.

This type is used by the following types:

- *Build Script Item Type* (page 36)
- *Build Start File Item Type* (page 38)

5.2.3.14 Build C Compiler Option

The value shall be a string. It shall be an option for the C compiler. The options are used to compile the sources of this item. The options defined by this attribute succeed the options presented to the item by the build item context.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build BSP Item Type* (page 29)
- *Build Library Item Type* (page 33)
- *Build Objects Item Type* (page 34)
- *Build Option C Compiler Check Action* (page 73)
- *Build Script Item Type* (page 36)
- *Build Test Program Item Type* (page 38)

5.2.3.15 Build C Preprocessor Option

The value shall be a string. It shall be an option for the C preprocessor. The options are used to preprocess the sources of this item. The options defined by this attribute succeed the options presented to the item by the build item context.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build BSP Item Type* (page 29)
- *Build Library Item Type* (page 33)
- *Build Objects Item Type* (page 34)
- *Build Script Item Type* (page 36)
- *Build Start File Item Type* (page 38)
- *Build Test Program Item Type* (page 38)

5.2.3.16 Build C++ Compiler Option

The value shall be a string. It shall be an option for the C++ compiler. The options are used to compile the sources of this item. The options defined by this attribute succeed the options presented to the item by the build item context.

This type is used by the following types:

- *Build Library Item Type* (page 33)
- *Build Objects Item Type* (page 34)
- *Build Option C++ Compiler Check Action* (page 74)
- *Build Script Item Type* (page 36)
- *Build Test Program Item Type* (page 38)

5.2.3.17 Build Dependency Link Role

This type refines the *Link* (page 88) though the role attribute if the value is `build-dependency`. It defines the build dependency role of links.

5.2.3.18 Build Include Path

The value shall be a string. It shall be a path to header files. The path is used by the C preprocessor to search for header files. It succeeds the includes presented to the item by the build item context. For an *Build Group Item Type* (page 32) item the includes are visible to all items referenced by the group item. For *Build BSP Item Type* (page 29), *Build Objects Item Type* (page 34), *Build Library Item Type* (page 33), *Build Start File Item Type* (page 38), and *Build Test Program Item Type* (page 38) items the includes are only visible to the sources specified by the item itself and they do not propagate to referenced items.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build BSP Item Type* (page 29)
- *Build Group Item Type* (page 32)
- *Build Library Item Type* (page 33)
- *Build Objects Item Type* (page 34)
- *Build Script Item Type* (page 36)
- *Build Start File Item Type* (page 38)
- *Build Test Program Item Type* (page 38)

5.2.3.19 Build Install Directive

This set of attributes specifies files installed by a build item. All explicit attributes shall be specified. The explicit attributes for this type are:

destination

The attribute value shall be a string. It shall be the install destination directory.

source

The attribute value shall be a list of strings. It shall be the list of source files to be installed in the destination directory. The path to a source file shall be relative to the directory of the wscript.

This type is used by the following types:

- *Build BSP Item Type* (page 29)
- *Build Group Item Type* (page 32)
- *Build Library Item Type* (page 33)
- *Build Objects Item Type* (page 34)

5.2.3.20 Build Install Path

A value of this type shall be of one of the following variants:

- There may be no value (null).
- The value may be a string. It shall be the installation path of a *Build Target* (page 76).

This type is used by the following types:

- *Build Configuration File Item Type* (page 31)
- *Build Configuration Header Item Type* (page 31)
- *Build Library Item Type* (page 33)
- *Build Start File Item Type* (page 38)

5.2.3.21 Build Link Static Library Directive

The value shall be a string. It shall be an external static library identifier. The library is used to link programs referenced by this item, e.g. `m` for `libm.a`. The library is added to the build command through the `stlib` attribute. It shall not be used for internal static libraries. Internal static libraries shall be specified through the `use-after` and `use-before` attributes to enable a proper build dependency tracking.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build Script Item Type* (page 36)
- *Build Test Program Item Type* (page 38)

5.2.3.22 Build Linker Option

The value shall be a string. It shall be an option for the linker. The options are used to link executables. The options defined by this attribute succeed the options presented to the item by the build item context.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build Script Item Type* (page 36)
- *Build Test Program Item Type* (page 38)

5.2.3.23 Build Option Action

This set of attributes specifies a build option action. Exactly one of the explicit attributes shall be specified. The explicit attributes for this type are:

append-test-cppflags

The attribute value shall be a string. It shall be the name of a test program. The action appends the action value to the `CPPFLAGS` of the test program. The name shall correspond to the name of a *Build Test Program Item Type* (page 38) item. Due to the processing order of items, there is no way to check if the name specified by the attribute value is valid.

assert-aligned

The attribute value shall be an integer number. The action asserts that the action value is aligned according to the attribute value.

assert-eq

The attribute value shall be a *Boolean or Integer or String* (page 67). The action asserts that the action value is equal to the attribute value.

assert-ge

The attribute value shall be an *Integer or String* (page 78). The action asserts that the action value is greater than or equal to the attribute value.

assert-gt

The attribute value shall be an *Integer or String* (page 78). The action asserts that the action value is greater than the attribute value.

assert-int16

The attribute shall have no value. The action asserts that the action value is a valid signed 16-bit integer.

assert-int32

The attribute shall have no value. The action asserts that the action value is a valid signed 32-bit integer.

assert-int64

The attribute shall have no value. The action asserts that the action value is a valid signed 64-bit integer.

assert-int8

The attribute shall have no value. The action asserts that the action value is a valid signed 8-bit integer.

assert-le

The attribute value shall be an *Integer or String* (page 78). The action asserts that the action value is less than or equal to the attribute value.

assert-lt

The attribute value shall be an *Integer or String* (page 78). The action asserts that the action value is less than the attribute value.

assert-ne

The attribute value shall be a *Boolean or Integer or String* (page 67). The action asserts that the action value is not equal to the attribute value.

assert-power-of-two

The attribute shall have no value. The action asserts that the action value is a power of two.

assert-uint16

The attribute shall have no value. The action asserts that the action value is a valid unsigned 16-bit integer.

assert-uint32

The attribute shall have no value. The action asserts that the action value is a valid unsigned 32-bit integer.

assert-uint64

The attribute shall have no value. The action asserts that the action value is a valid unsigned 64-bit integer.

assert-uint8

The attribute shall have no value. The action asserts that the action value is a valid unsigned 8-bit integer.

check-cc

The attribute value shall be a *Build Option C Compiler Check Action* (page 73).

check-cxx

The attribute value shall be a *Build Option C++ Compiler Check Action* (page 74).

define

The attribute value shall be an optional string. The action adds a define to the configuration set. If the attribute value is present, then it is used as the name of the define, otherwise the name of the item is used. The value of the define is the action value. If the action value is a string, then it is quoted.

define-condition

The attribute value shall be an optional string. The action adds a conditional define to the configuration set. If the attribute value is present, then it is used as the name of the define, otherwise the name of the item is used. The value of the define is the action value.

define-unquoted

The attribute value shall be an optional string. The action adds a define to the configuration set. If the attribute value is present, then it is used as the name of the define, otherwise the name of the item is used. The value of the define is the action value. If the action value is a string, then it is not quoted.

env-append

The attribute value shall be an optional string. The action appends the action value to an environment of the configuration set. If the attribute value is present, then it is used as the name of the environment variable, otherwise the name of the item is used.

env-assign

The attribute value shall be an optional string. The action assigns the action value to an environment of the configuration set. If the attribute value is present, then it is used as the name of the environment variable, otherwise the name of the item is used.

env-enable

The attribute value shall be an optional string. If the action value is true, then a name is appended to the ENABLE environment variable of the configuration set. If the attribute value is present, then it is used as the name, otherwise the name of the item is used.

find-program

The attribute shall have no value. The action tries to find the program specified by the action value. Uses the `${PATH}` to find the program. Returns the result of the find operation, e.g. a path to the program.

find-tool

The attribute shall have no value. The action tries to find the tool specified by the action value. Uses the tool paths specified by the `--rtems-tools` command line option. Returns the result of the find operation, e.g. a path to the program.

format-and-define

The attribute value shall be an optional string. The action adds a define to the configuration set. If the attribute value is present, then it is used as the name of the define, otherwise the name of the item is used. The value of the define is the action value. The value is formatted according to the format attribute value.

get-boolean

The attribute shall have no value. The action gets the action value for subsequent actions from a configuration file variable named by the items name attribute. If no such variable exists in the configuration file, then the default value is used. The value is converted to a boolean.

get-env

The attribute value shall be a string. The action gets the action value for subsequent actions from the environment variable of the configuration set named by the attribute value.

get-integer

The attribute shall have no value. The action gets the action value for subsequent actions from a configuration file variable named by the items name attribute. If no such variable exists in the configuration file, then the default value is used. The value is converted to an integer.

get-string

The attribute shall have no value. The action gets the action value for subsequent actions from a configuration file variable named by the items name attribute. If no such variable exists in the configuration file, then the default value is used. The value is converted to a string.

script

The attribute value shall be a string. The action executes the attribute value with the Python `eval()` function in the context of the script action handler.

set-test-state

The attribute value shall be a *Build Option Set Test State Action* (page 75).

set-value

The attribute value shall be a *Build Option Value* (page 75). The action sets the action value for subsequent actions to the attribute value.

split

The attribute shall have no value. The action splits the action value.

substitute

The attribute shall have no value. The action performs a `#{VARIABLE}` substitution on the action value. Use `$$` for a plain `$` character.

This type is used by the following types:

- *Build Option Item Type* (page 35)

5.2.3.24 Build Option C Compiler Check Action

This set of attributes specifies a check done using the C compiler. All explicit attributes shall be specified. The explicit attributes for this type are:

cflags

The attribute value shall be a list. Each list element shall be a *Build C Compiler Option* (page 67).

fragment

The attribute value shall be a string. It shall be a code fragment used to check the availability of a certain feature through compilation with the C compiler. The resulting object is not linked to an executable.

message

The attribute value shall be a string. It shall be a description of the feature to check.

This type is used by the following types:

- *Build Option Action* (page 70)

5.2.3.25 Build Option C++ Compiler Check Action

This set of attributes specifies a check done using the C++ compiler. All explicit attributes shall be specified. The explicit attributes for this type are:

cxxflags

The attribute value shall be a list. Each list element shall be a *Build C++ Compiler Option* (page 68).

fragment

The attribute value shall be a string. It shall be a code fragment used to check the availability of a certain feature through compilation with the C++ compiler. The resulting object is not linked to an executable.

message

The attribute value shall be a string. It shall be a description of the feature to check.

This type is used by the following types:

- *Build Option Action* (page 70)

5.2.3.26 Build Option Default by Variant

This set of attributes specifies build option default values by variant. All explicit attributes shall be specified. The explicit attributes for this type are:

value

The attribute value shall be a *Build Option Value* (page 75). Its value shall be the default value for the matching variants.

variants

The attribute value shall be a list of strings. It shall be a list of Python regular expression matching with the desired variants.

This type is used by the following types:

- *Build Option Item Type* (page 35)

5.2.3.27 Build Option Name

The value shall be a string. It shall be the name of the build option. The value shall match with the regular expression “`^[a-zA-Z_][a-zA-Z0-9_]*$`”.

This type is used by the following types:

- *Build Option Item Type* (page 35)

5.2.3.28 Build Option Set Test State Action

This set of attributes specifies test states for a set of test programs. Generic attributes may be specified. Each generic attribute key shall be a *Name* (page 89). Each generic attribute value shall be a *Build Test State* (page 76). The keys shall be test program names. The names shall correspond to the name of a *Build Test Program Item Type* (page 38) or *Build Ada Test Program Item Type* (page 28) item. Due to the processing order of items, there is no way to check if the name specified by the attribute key is valid.

This type is used by the following types:

- *Build Option Action* (page 70)

5.2.3.29 Build Option Value

A value of this type shall be of one of the following variants:

- The value may be a boolean.
- The value may be an integer number.
- The value may be a list. Each list element shall be a string.
- There may be no value (null).
- The value may be a string.

This type is used by the following types:

- *Build Option Action* (page 70)
- *Build Option Default by Variant* (page 74)
- *Build Option Item Type* (page 35)

5.2.3.30 Build Source

The value shall be a string. It shall be a source file. The path to a source file shall be relative to the directory of the `wscript`.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build BSP Item Type* (page 29)
- *Build Library Item Type* (page 33)
- *Build Objects Item Type* (page 34)
- *Build Start File Item Type* (page 38)
- *Build Test Program Item Type* (page 38)

5.2.3.31 Build Target

The value shall be a string. It shall be the target file path. The path to the target file shall be relative to the directory of the wscript. The target file is located in the build tree.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build Configuration File Item Type* (page 31)
- *Build Configuration Header Item Type* (page 31)
- *Build Library Item Type* (page 33)
- *Build Start File Item Type* (page 38)
- *Build Test Program Item Type* (page 38)

5.2.3.32 Build Test State

The value shall be a string. This string defines a test state. The value shall be an element of

- “benchmark”,
- “exclude”,
- “expected-fail”,
- “indeterminate”, and
- “user-input”.

This type is used by the following types:

- *Build Option Set Test State Action* (page 75)

5.2.3.33 Build Use After Directive

The value shall be a string. It shall be an internal static library identifier. The library is used to link programs referenced by this item, e.g. z for libz.a. The library is placed after the use items of the build item context.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build Group Item Type* (page 32)
- *Build Script Item Type* (page 36)
- *Build Test Program Item Type* (page 38)

5.2.3.34 Build Use Before Directive

The value shall be a string. It shall be an internal static library identifier. The library is used to link programs referenced by this item, e.g. z for libz.a. The library is placed before the use items of the build item context.

This type is used by the following types:

- *Build Ada Test Program Item Type* (page 28)
- *Build Group Item Type* (page 32)
- *Build Script Item Type* (page 36)
- *Build Test Program Item Type* (page 38)

5.2.3.35 Constraint Link Role

This type refines the *Link* (page 88) though the role attribute if the value is constraint. It defines the constraint role of links. The link target shall be a constraint.

5.2.3.36 Copyright

The value shall be a string. It shall be a copyright statement of a copyright holder of the specification item. The value

- shall match with the regular expression “`^\s*Copyright\s+(C)\s+[0-9]+\s*[0-9]+\s+.\s*$`”,
- or, shall match with the regular expression “`^\s*Copyright\s+(C)\s+[0-9]+\s+.\s*\s*$`”,
- or, shall match with the regular expression “`^\s*Copyright\s+(C)\s+.\s*$`”.

This type is used by the following types:

- *Root Item Type* (page 26)

5.2.3.37 Enabled-By Expression

A value of this type shall be an expression which defines under which conditions the specification item or parts of it are enabled. The expression is evaluated with the use of an *enabled set*. This is a set of strings which indicate enabled features.

A value of this type shall be of one of the following variants:

- The value may be a boolean. This expression evaluates directly to the boolean value.
- The value may be a set of attributes. Each attribute defines an operator. Exactly one of the explicit attributes shall be specified. The explicit attributes for this type are:

and

The attribute value shall be a list. Each list element shall be an *Enabled-By Expression* (page 77). The *and* operator evaluates to the *logical and* of the evaluation results of the expressions in the list.

not

The attribute value shall be an *Enabled-By Expression* (page 77). The *not* operator evaluates to the *logical not* of the evaluation results of the expression.

or

The attribute value shall be a list. Each list element shall be an *Enabled-By Expression* (page 77). The *or* operator evaluates to the *logical or* of the evaluation results of the expressions in the list.

- The value may be a list. Each list element shall be an *Enabled-By Expression* (page 77). This list of expressions evaluates to the *logical or* of the evaluation results of the expressions in the list.
- The value may be a string. If the value is in the *enabled set*, this expression evaluates to true, otherwise to false.

This type is used by the following types:

- *Action Requirement Transition* (page 64)
- *Enabled-By Expression* (page 77)
- *Interface Include Link Role* (page 86)
- *Root Item Type* (page 26)

Please have a look at the following example:

```

1 enabled-by:
2   and:
3   - RTEMS_NETWORKING
4   - not: RTEMS_SMP

```

5.2.3.38 Glossary Membership Link Role

This type refines the *Link* (page 88) though the *role* attribute if the value is *glossary-member*. It defines the glossary membership role of links.

5.2.3.39 Integer or String

A value of this type shall be of one of the following variants:

- The value may be an integer number.
- The value may be a string.

This type is used by the following types:

- *Application Configuration Option Constraint Set* (page 66)
- *Application Configuration Value Option Item Type* (page 43)
- *Build Option Action* (page 70)

5.2.3.40 Interface Brief Description

A value of this type shall be of one of the following variants:

- There may be no value (null).
- The value may be a string. It shall be the brief description of the interface. It should be a single sentence. The value shall not match with the regular expression “\n\n”.

This type is used by the following types:

- *Interface Compound Item Type* (page 43)
- *Interface Compound Member Definition* (page 80)
- *Interface Define Item Type* (page 44)
- *Interface Enum Item Type* (page 45)
- *Interface Enumerator Item Type* (page 45)
- *Interface Function Item Type* (page 46)
- *Interface Group Item Type* (page 46)
- *Interface Header File Item Type* (page 47)
- *Interface Macro Item Type* (page 47)
- *Interface Typedef Item Type* (page 48)
- *Interface Variable Item Type* (page 49)

5.2.3.41 Interface Compound Definition Kind

The value shall be a string. It specifies how the interface compound is defined. It may be a typedef only, the struct or union only, or a typedef with a struct or union definition. The value shall be an element of

- “struct-only”,
- “typedef-and-struct”,
- “typedef-and-union”,
- “typedef-only”, and
- “union-only”.

This type is used by the following types:

- *Interface Compound Item Type* (page 43)

5.2.3.42 Interface Compound Member Compound

This type refines the following types:

- *Interface Compound Member Definition* (page 80) though the kind attribute if the value is struct
- *Interface Compound Member Definition* (page 80) though the kind attribute if the value is union

This set of attributes specifies an interface compound member compound. All explicit attributes shall be specified. The explicit attributes for this type are:

definition

The attribute value shall be a list. Each list element shall be an *Interface Compound Member Definition Directive* (page 81).

5.2.3.43 Interface Compound Member Declaration

This type refines the *Interface Compound Member Definition* (page 80) though the kind attribute if the value is member. This set of attributes specifies an interface compound member declaration. All explicit attributes shall be specified. The explicit attributes for this type are:

definition

The attribute value shall be a string. It shall be the interface compound member declaration. On the declaration a context-sensitive substitution of item variables is performed.

5.2.3.44 Interface Compound Member Definition

This set of attributes specifies an interface compound member definition. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an *Interface Brief Description* (page 79).

description

The attribute value shall be an *Interface Description* (page 82).

kind

The attribute value shall be a string. It shall be the interface compound member kind.

name

The attribute value shall be a string. It shall be the interface compound member name.

This type is refined by the following types:

- *Interface Compound Member Compound* (page 80)
- *Interface Compound Member Declaration* (page 80)

This type is used by the following types:

- *Interface Compound Member Definition Directive* (page 81)
- *Interface Compound Member Definition Variant* (page 81)

5.2.3.45 Interface Compound Member Definition Directive

This set of attributes specifies an interface compound member definition directive. All explicit attributes shall be specified. The explicit attributes for this type are:

default

The attribute value shall be an *Interface Compound Member Definition* (page 80). The default definition will be used if no variant-specific definition is enabled.

variants

The attribute value shall be a list. Each list element shall be an *Interface Compound Member Definition Variant* (page 81).

This type is used by the following types:

- *Interface Compound Item Type* (page 43)
- *Interface Compound Member Compound* (page 80)

5.2.3.46 Interface Compound Member Definition Variant

This set of attributes specifies an interface compound member definition variant. All explicit attributes shall be specified. The explicit attributes for this type are:

definition

The attribute value shall be an *Interface Compound Member Definition* (page 80). The definition will be used if the expression defined by the *enabled-by* attribute evaluates to true. In generated header files, the expression is evaluated by the C preprocessor.

enabled-by

The attribute value shall be an *Interface Enabled-By Expression* (page 83).

This type is used by the following types:

- *Interface Compound Member Definition Directive* (page 81)

5.2.3.47 Interface Definition

A value of this type shall be of one of the following variants:

- There may be no value (null).
- The value may be a string. It shall be the definition. On the definition a context-sensitive substitution of item variables is performed.

This type is used by the following types:

- *Interface Definition Directive* (page 82)
- *Interface Definition Variant* (page 82)

5.2.3.48 Interface Definition Directive

This set of attributes specifies an interface definition directive. All explicit attributes shall be specified. The explicit attributes for this type are:

default

The attribute value shall be an *Interface Definition* (page 81). The default definition will be used if no variant-specific definition is enabled.

variants

The attribute value shall be a list. Each list element shall be an *Interface Definition Variant* (page 82).

This type is used by the following types:

- *Interface Define Item Type* (page 44)
- *Interface Enumerator Item Type* (page 45)
- *Interface Macro Item Type* (page 47)
- *Interface Typedef Item Type* (page 48)
- *Interface Variable Item Type* (page 49)

5.2.3.49 Interface Definition Variant

This set of attributes specifies an interface definition variant. All explicit attributes shall be specified. The explicit attributes for this type are:

definition

The attribute value shall be an *Interface Definition* (page 81). The definition will be used if the expression defined by the *enabled-by* attribute evaluates to true. In generated header files, the expression is evaluated by the C preprocessor.

enabled-by

The attribute value shall be an *Interface Enabled-By Expression* (page 83).

This type is used by the following types:

- *Interface Definition Directive* (page 82)

5.2.3.50 Interface Description

A value of this type shall be of one of the following variants:

- There may be no value (null).
- The value may be a string. It shall be the description of the interface. The description should be short and concentrate on the average case. All special cases, usage notes, constraints, error conditions, configuration dependencies, references, etc. should be described in the *Interface Notes* (page 86).

This type is used by the following types:

- *Application Configuration Option Item Type* (page 42)
- *Interface Compound Item Type* (page 43)

- *Interface Compound Member Definition* (page 80)
- *Interface Define Item Type* (page 44)
- *Interface Enum Item Type* (page 45)
- *Interface Enumerator Item Type* (page 45)
- *Interface Function Item Type* (page 46)
- *Interface Group Item Type* (page 46)
- *Interface Macro Item Type* (page 47)
- *Interface Parameter* (page 86)
- *Interface Return Value* (page 88)
- *Interface Typedef Item Type* (page 48)
- *Interface Variable Item Type* (page 49)

5.2.3.51 Interface Enabled-By Expression

A value of this type shall be an expression which defines under which conditions an interface definition is enabled. In generated header files, the expression is evaluated by the C preprocessor.

A value of this type shall be of one of the following variants:

- The value may be a boolean. It is converted to 0 or 1. It defines a symbol in the expression.
- The value may be a set of attributes. Each attribute defines an operator. Exactly one of the explicit attributes shall be specified. The explicit attributes for this type are:

and

The attribute value shall be a list. Each list element shall be an *Interface Enabled-By Expression* (page 83). The *and* operator defines a *logical and* of the expressions in the list.

not

The attribute value shall be an *Interface Enabled-By Expression* (page 83). The *not* operator defines a *logical not* of the expression.

or

The attribute value shall be a list. Each list element shall be an *Interface Enabled-By Expression* (page 83). The *or* operator defines a *logical or* of the expressions in the list.

- The value may be a list. Each list element shall be an *Interface Enabled-By Expression* (page 83). It defines a *logical or* of the expressions in the list.
- The value may be a string. It defines a symbol in the expression.

This type is used by the following types:

- *Interface Compound Member Definition Variant* (page 81)
- *Interface Definition Variant* (page 82)
- *Interface Enabled-By Expression* (page 83)
- *Interface Function Definition Variant* (page 85)

5.2.3.52 Interface Enum Definition Kind

The value shall be a string. It specifies how the enum is defined. It may be a typedef only, the enum only, or a typedef with an enum definition. The value shall be an element of

- “enum-only”,
- “typedef-and-enum”, and
- “typedef-only”.

This type is used by the following types:

- *Interface Enum Item Type* (page 45)

5.2.3.53 Interface Enumerator Link Role

This type refines the *Link* (page 88) through the role attribute if the value is interface-enumerator. It defines the interface enumerator role of links.

5.2.3.54 Interface Function Definition

This set of attributes specifies a function definition. All explicit attributes shall be specified. The explicit attributes for this type are:

attributes

The attribute value shall be an optional string. If the value is present, then it shall be the function attributes. On the attributes a context-sensitive substitution of item variables is performed. A function attribute is for example the indication that the function does not return to the caller.

body

The attribute value shall be an optional string. If the value is present, then it shall be the definition of a static inline function. On the function definition a context-sensitive substitution of item variables is performed. If no value is present, then the function is declared as an external function.

params

The attribute value shall be a list of strings. It shall be the list of parameter declarations of the function. On the function parameter declarations a context-sensitive substitution of item variables is performed.

return

The attribute value shall be a string. It shall be the function return type. On the return type a context-sensitive substitution of item variables is performed.

This type is used by the following types:

- *Interface Function Definition Directive* (page 85)
- *Interface Function Definition Variant* (page 85)

5.2.3.55 Interface Function Definition Directive

This set of attributes specifies a function definition directive. All explicit attributes shall be specified. The explicit attributes for this type are:

default

The attribute value shall be an *Interface Function Definition* (page 84). The default definition will be used if no variant-specific definition is enabled.

variants

The attribute value shall be a list. Each list element shall be an *Interface Function Definition Variant* (page 85).

This type is used by the following types:

- *Interface Function Item Type* (page 46)

5.2.3.56 Interface Function Definition Variant

This set of attributes specifies a function definition variant. All explicit attributes shall be specified. The explicit attributes for this type are:

definition

The attribute value shall be an *Interface Function Definition* (page 84). The definition will be used if the expression defined by the *enabled-by* attribute evaluates to true. In generated header files, the expression is evaluated by the C preprocessor.

enabled-by

The attribute value shall be an *Interface Enabled-By Expression* (page 83).

This type is used by the following types:

- *Interface Function Definition Directive* (page 85)

5.2.3.57 Interface Function Link Role

This type refines the *Link* (page 88) though the *role* attribute if the value is *interface-function*. It defines the interface function role of links. It is used to indicate that a *Action Requirement Item Type* (page 50) item specifies functional requirements of an *Interface Function Item Type* (page 46) or a *Interface Macro Item Type* (page 47) item.

5.2.3.58 Interface Group Identifier

The value shall be a string. It shall be the identifier of the interface group. The value shall match with the regular expression “`^[A-Z][a-zA-Z0-9]*$`”.

This type is used by the following types:

- *Interface Group Item Type* (page 46)

5.2.3.59 Interface Group Membership Link Role

This type refines the *Link* (page 88) though the `role` attribute if the value is `interface-ingroup`. It defines the interface group membership role of links.

5.2.3.60 Interface Include Link Role

This type refines the *Link* (page 88) though the `role` attribute if the value is `interface-include`. It defines the interface include role of links and is used to indicate that an interface container includes another interface container. For example, one header file includes another header file. All explicit attributes shall be specified. The explicit attributes for this type are:

enabled-by

The attribute value shall be an *Enabled-By Expression* (page 77). It shall define under which conditions the interface container is included.

5.2.3.61 Interface Notes

A value of this type shall be of one of the following variants:

- There may be no value (null).
- The value may be a string. It shall be the notes for the interface.

This type is used by the following types:

- *Application Configuration Option Item Type* (page 42)
- *Interface Compound Item Type* (page 43)
- *Interface Define Item Type* (page 44)
- *Interface Enumerator Item Type* (page 45)
- *Interface Function Item Type* (page 46)
- *Interface Macro Item Type* (page 47)
- *Interface Typedef Item Type* (page 48)
- *Interface Variable Item Type* (page 49)

5.2.3.62 Interface Parameter

This set of attributes specifies an interface parameter. All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be an *Interface Description* (page 82).

dir

The attribute value shall be an *Interface Parameter Direction* (page 87).

name

The attribute value shall be a string. It shall be the interface parameter name.

This type is used by the following types:

- *Interface Function Item Type* (page 46)
- *Interface Macro Item Type* (page 47)

5.2.3.63 Interface Parameter Direction

A value of this type shall be of one of the following variants:

- There may be no value (null).
- The value may be a string. It specifies the interface parameter direction. The value shall be an element of
 - “in”,
 - “out”, and
 - “inout”.

This type is used by the following types:

- *Interface Parameter* (page 86)
- *Test Run Parameter* (page 106)

5.2.3.64 Interface Placement Link Role

This type refines the *Link* (page 88) through the role attribute if the value is interface-placement. It defines the interface placement role of links. It is used to indicate that an interface definition is placed into an interface container, for example a header file.

5.2.3.65 Interface Return Directive

This set of attributes specifies an interface return. All explicit attributes shall be specified. The explicit attributes for this type are:

return

The attribute value shall be an optional string. It shall describe the interface return for unspecified return values.

return-values

The attribute value shall be a list. Each list element shall be an *Interface Return Value* (page 88).

This type is used by the following types:

- *Interface Function Item Type* (page 46)
- *Interface Macro Item Type* (page 47)

5.2.3.66 Interface Return Value

This set of attributes specifies an interface return value. All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be an *Interface Description* (page 82).

value

The attribute value shall be a *Boolean or Integer or String* (page 67). It shall be the described interface return value.

This type is used by the following types:

- *Interface Return Directive* (page 87)

5.2.3.67 Interface Target Link Role

This type refines the *Link* (page 88) though the role attribute if the value is interface-target. It defines the interface target role of links. It is used for interface forward declarations.

5.2.3.68 Link

This set of attributes specifies a link from one specification item to another specification item. The links in a list are ordered. The first link in the list is processed first. All explicit attributes shall be specified. The explicit attributes for this type are:

role

The attribute value shall be a *Name* (page 89). It shall be the role of the link.

uid

The attribute value shall be an *UID* (page 107). It shall be the absolute or relative UID of the link target item.

This type is refined by the following types:

- *Application Configuration Group Member Link Role* (page 66)
- *Build Dependency Link Role* (page 68)
- *Constraint Link Role* (page 77)
- *Glossary Membership Link Role* (page 78)
- *Interface Enumerator Link Role* (page 84)
- *Interface Function Link Role* (page 85)
- *Interface Group Membership Link Role* (page 86)
- *Interface Include Link Role* (page 86)
- *Interface Placement Link Role* (page 87)
- *Interface Target Link Role* (page 88)
- *Placement Order Link Role* (page 90)
- *Requirement Refinement Link Role* (page 91)

- *Requirement Validation Link Role* (page 93)
- *Runtime Measurement Request Link Role* (page 94)
- *Specification Member Link Role* (page 102)
- *Specification Refinement Link Role* (page 102)

This type is used by the following types:

- *Root Item Type* (page 26)
- *Test Case Action* (page 104)
- *Test Case Check* (page 104)

5.2.3.69 Name

The value shall be a string. A string is a valid name if it matches with the `^[a-z][a-z0-9-]*|SPDX-License-Identifier)$` regular expression.

This type is used by the following types:

- *Application Configuration Option Item Type* (page 42)
- *Build Item Type* (page 27)
- *Build Option Set Test State Action* (page 75)
- *Functional Requirement Item Type* (page 50)
- *Glossary Item Type* (page 40)
- *Interface Item Type* (page 41)
- *Link* (page 88)
- *Non-Functional Requirement Item Type* (page 54)
- *Requirement Item Type* (page 49)
- *Root Item Type* (page 26)
- *Runtime Measurement Parameter Set* (page 94)
- *Runtime Performance Parameter Set* (page 95)
- *Specification Attribute Value* (page 96)
- *Specification Explicit Attributes* (page 97)
- *Specification Generic Attributes* (page 99)
- *Specification Item Type* (page 59)
- *Specification List* (page 101)
- *Specification Refinement Link Role* (page 102)

5.2.3.70 Optional String

A value of this type shall be of one of the following variants:

- There may be no value (null).
- The value may be a string.

5.2.3.71 Placement Order Link Role

This type refines the *Link* (page 88) through the `role` attribute if the value is `placement-order`. This link role defines the placement order of items in a container item (for example an interface function in a header file or a documentation section).

5.2.3.72 Requirement Reference

This set of attributes specifies a requirement reference. All explicit attributes shall be specified. The explicit attributes for this type are:

identifier

The attribute value shall be a string. It shall be the type-specific identifier of the reference target. For *group* references use the Doxygen group identifier.

type

The attribute value shall be a *Requirement Reference Type* (page 90).

This type is used by the following types:

- *Requirement Item Type* (page 49)

5.2.3.73 Requirement Reference Type

The value shall be a string. It specifies the type of a requirement reference. The value shall be an element of

- “define”,
- “file”,
- “function”,
- “group”,
- “macro”, and
- “variable”.

This type is used by the following types:

- *Requirement Reference* (page 90)

5.2.3.74 Requirement Refinement Link Role

This type refines the *Link* (page 88) though the role attribute if the value is requirement-refinement. It defines the requirement refinement role of links.

5.2.3.75 Requirement Text

The value shall be a string. It shall state a requirement or constraint. The value shall not contain an element of

- “acceptable”,
- “adequate”,
- “almost always”,
- “and/or”,
- “appropriate”,
- “approximately”,
- “as far as possible”,
- “as much as practicable”,
- “best”,
- “best possible”,
- “easy”,
- “efficient”,
- “e.g.”,
- “enable”,
- “enough”,
- “etc.”,
- “few”,
- “first rate”,
- “flexible”,
- “generally”,
- “goal”,
- “graceful”,
- “great”,
- “greatest”,
- “ideally”,
- “i.e.”,
- “if possible”,
- “in most cases”,

- “large”,
- “many”,
- “maximize”,
- “minimize”,
- “most”,
- “multiple”,
- “necessary”,
- “numerous”,
- “optimize”,
- “ought to”,
- “probably”,
- “quick”,
- “rapid”,
- “reasonably”,
- “relevant”,
- “robust”,
- “satisfactory”,
- “several”,
- “shall be included but not limited to”,
- “simple”,
- “small”,
- “some”,
- “state of the art”,
- “sufficient”,
- “suitable”,
- “support”,
- “systematically”,
- “transparent”,
- “typical”,
- “user friendly”,
- “usually”,
- “versatile”, and
- “when necessary”.

This type is used by the following types:

- *Action Requirement State* (page 64)
- *Application Configuration Group Item Type* (page 41)
- *Application Configuration Option Constraint Set* (page 66)
- *Application Configuration Option Item Type* (page 42)
- *Constraint Item Type* (page 40)
- *Interface Group Item Type* (page 46)
- *Requirement Item Type* (page 49)

5.2.3.76 Requirement Validation Link Role

This type refines the *Link* (page 88) though the role attribute if the value is validation. It defines the requirement validation role of links.

5.2.3.77 Requirement Validation Method

The value shall be a string. This value type characterizes a requirement validation method (except validation by test). The value shall be an element of

- “by-analysis”,
- “by-inspection”, and
- “by-review-of-design”.

This type is used by the following types:

- *Requirement Validation Item Type* (page 58)

5.2.3.78 Runtime Measurement Environment

The value shall be a string. It specifies the runtime measurement environment. The value

- shall be an element of
 - “FullCache”,
 - “HotCache”, and
 - “DirtyCache”,
- or, shall match with the regular expression “^Load/[1-9][0-9]*\$”.

This type is used by the following types:

- *Runtime Measurement Environment Table* (page 94)

5.2.3.79 Runtime Measurement Environment Table

This set of attributes provides runtime performance limits for a set of runtime measurement environments. Generic attributes may be specified. Each generic attribute key shall be a *Runtime Measurement Environment* (page 93). Each generic attribute value shall be a *Runtime Measurement Value Table* (page 95).

This type is used by the following types:

- *Runtime Performance Limit Table* (page 95)

5.2.3.80 Runtime Measurement Parameter Set

This set of attributes defines parameters of the runtime measurement test case. All explicit attributes shall be specified. The explicit attributes for this type are:

sample-count

The attribute value shall be an integer number. It shall be the sample count of the runtime measurement context.

In addition to the explicit attributes, generic attributes may be specified. Each generic attribute key shall be a *Name* (page 89). The attribute value may have any type.

This type is used by the following types:

- *Runtime Measurement Test Item Type* (page 58)

5.2.3.81 Runtime Measurement Request Link Role

This type refines the *Link* (page 88) though the role attribute if the value is *runtime-measurement-request*. It defines the runtime measurement request role of links. The link target shall be a *Runtime Measurement Test Item Type* (page 58) item.

5.2.3.82 Runtime Measurement Value Kind

The value shall be a string. It specifies the kind of a runtime measurement value. The value shall be an element of

- “max-lower-bound”,
- “max-upper-bound”,
- “mean-lower-bound”,
- “mean-upper-bound”,
- “min-lower-bound”, and
- “min-upper-bound”.

This type is used by the following types:

- *Runtime Measurement Value Table* (page 95)

5.2.3.83 Runtime Measurement Value Table

This set of attributes provides a set of runtime measurement values each of a specified kind. The unit of the values shall be one second. Generic attributes may be specified. Each generic attribute key shall be a *Runtime Measurement Value Kind* (page 94). Each generic attribute value shall be a floating-point number.

This type is used by the following types:

- *Runtime Measurement Environment Table* (page 94)

5.2.3.84 Runtime Performance Limit Table

This set of attributes provides runtime performance limits for BSP variants specified by "<arch>/<bsp>" with <arch> being the architecture of the BSP and <bsp> being the base name of the BSP. Generic attributes may be specified. Each generic attribute key shall be a string. Each generic attribute value shall be a *Runtime Measurement Environment Table* (page 94).

This type is used by the following types:

- *Runtime Performance Requirement Item Type* (page 55)

5.2.3.85 Runtime Performance Parameter Set

This set of attributes defines parameters of the runtime performance requirement. Generic attributes may be specified. Each generic attribute key shall be a *Name* (page 89). The attribute value may have any type.

This type is used by the following types:

- *Runtime Performance Requirement Item Type* (page 55)

5.2.3.86 SPDX License Identifier

The value shall be a string. It defines the license of the item expressed though an SPDX License Identifier. The value

- shall be equal to "CC-BY-SA-4.0 OR BSD-2-Clause",
- or, shall be equal to "BSD-2-Clause",
- or, shall be equal to "CC-BY-SA-4.0".

This type is used by the following types:

- *Root Item Type* (page 26)

5.2.3.87 Specification Attribute Set

This set of attributes specifies a set of attributes. The following explicit attributes are mandatory:

- attributes
- description
- mandatory-attributes

The explicit attributes for this type are:

attributes

The attribute value shall be a *Specification Explicit Attributes* (page 97). It shall specify the explicit attributes of the attribute set.

description

The attribute value shall be an optional string. It shall be the description of the attribute set.

generic-attributes

The attribute value shall be a *Specification Generic Attributes* (page 99). It shall specify the generic attributes of the attribute set.

mandatory-attributes

The attribute value shall be a *Specification Mandatory Attributes* (page 101). It shall specify the mandatory attributes of the attribute set.

This type is used by the following types:

- *Specification Information* (page 99)

5.2.3.88 Specification Attribute Value

This set of attributes specifies an attribute value. All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be an optional string. It shall be the description of the attribute value.

spec-type

The attribute value shall be a *Name* (page 89). It shall be the specification type of the attribute value.

This type is used by the following types:

- *Specification Explicit Attributes* (page 97)

5.2.3.89 Specification Boolean Value

This attribute set specifies a boolean value. Only the description attribute is mandatory. The explicit attributes for this type are:

assert

The attribute value shall be a boolean. This optional attribute defines the value constraint of the specified boolean value. If the value of the assert attribute is true, then the value of the specified boolean value shall be true. If the value of the assert attribute is false, then the value of the specified boolean value shall be false. In case the assert attribute is not present, then the value of the specified boolean value may be true or false.

description

The attribute value shall be an optional string. It shall be the description of the specified boolean value.

This type is used by the following types:

- *Specification Information* (page 99)

5.2.3.90 Specification Explicit Attributes

Generic attributes may be specified. Each generic attribute key shall be a *Name* (page 89). Each generic attribute value shall be a *Specification Attribute Value* (page 96). Each generic attribute specifies an explicit attribute of the attribute set. The key of the each generic attribute defines the attribute key of the explicit attribute.

This type is used by the following types:

- *Specification Attribute Set* (page 96)

5.2.3.91 Specification Floating-Point Assert

A value of this type shall be an expression which asserts that the floating-point value of the specified attribute satisfies the required constraints.

A value of this type shall be of one of the following variants:

- The value may be a set of attributes. Each attribute defines an operator. Exactly one of the explicit attributes shall be specified. The explicit attributes for this type are:

and

The attribute value shall be a list. Each list element shall be a *Specification Floating-Point Assert* (page 97). The *and* operator evaluates to the *logical and* of the evaluation results of the expressions in the list.

eq

The attribute value shall be a floating-point number. The *eq* operator evaluates to true, if the value to check is equal to the value of this attribute, otherwise to false.

ge

The attribute value shall be a floating-point number. The *ge* operator evaluates to true, if the value to check is greater than or equal to the value of this attribute, otherwise to false.

gt

The attribute value shall be a floating-point number. The *gt* operator evaluates to true, if the value to check is greater than the value of this attribute, otherwise to false.

le

The attribute value shall be a floating-point number. The *le* operator evaluates to true, if the value to check is less than or equal to the value of this attribute, otherwise to false.

lt

The attribute value shall be a floating-point number. The *lt* operator evaluates to true, if the value to check is less than the value of this attribute, otherwise to false.

ne

The attribute value shall be a floating-point number. The *ne* operator evaluates to true, if the value to check is not equal to the value of this attribute, otherwise to false.

not

The attribute value shall be a *Specification Floating-Point Assert* (page 97). The *not* operator evaluates to the *logical not* of the evaluation results of the expression.

or

The attribute value shall be a list. Each list element shall be a *Specification Floating-Point Assert* (page 97). The *or* operator evaluates to the *logical or* of the evaluation results of the expressions in the list.

- The value may be a list. Each list element shall be a *Specification Floating-Point Assert* (page 97). This list of expressions evaluates to the *logical or* of the evaluation results of the expressions in the list.

This type is used by the following types:

- *Specification Floating-Point Assert* (page 97)
- *Specification Floating-Point Value* (page 98)

5.2.3.92 Specification Floating-Point Value

This set of attributes specifies a floating-point value. Only the description attribute is mandatory. The explicit attributes for this type are:

assert

The attribute value shall be a *Specification Floating-Point Assert* (page 97). This optional attribute defines the value constraints of the specified floating-point value. In case the *assert* attribute is not present, then the value of the specified floating-point value may be every valid floating-point number.

description

The attribute value shall be an optional string. It shall be the description of the specified floating-point value.

This type is used by the following types:

- *Specification Information* (page 99)

5.2.3.93 Specification Generic Attributes

This set of attributes specifies generic attributes. Generic attributes are attributes which are not explicitly specified by *Specification Explicit Attributes* (page 97). They are restricted to uniform attribute key and value types. All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be an optional string. It shall be the description of the generic attributes.

key-spec-type

The attribute value shall be a *Name* (page 89). It shall be the specification type of the generic attribute keys.

value-spec-type

The attribute value shall be a *Name* (page 89). It shall be the specification type of the generic attribute values.

This type is used by the following types:

- *Specification Attribute Set* (page 96)

5.2.3.94 Specification Information

This set of attributes specifies attribute values. At least one of the explicit attributes shall be specified. The explicit attributes for this type are:

bool

The attribute value shall be a *Specification Boolean Value* (page 97). It shall specify a boolean value.

dict

The attribute value shall be a *Specification Attribute Set* (page 96). It shall specify a set of attributes.

float

The attribute value shall be a *Specification Floating-Point Value* (page 98). It shall specify a floating-point value.

int

The attribute value shall be a *Specification Integer Value* (page 101). It shall specify an integer value.

list

The attribute value shall be a *Specification List* (page 101). It shall specify a list of attributes or values.

none

The attribute shall have no value. It specifies that no value is required.

str

The attribute value shall be a *Specification String Value* (page 103). It shall specify a string.

This type is used by the following types:

- *Specification Item Type* (page 59)

5.2.3.95 Specification Integer Assert

A value of this type shall be an expression which asserts that the integer value of the specified attribute satisfies the required constraints.

A value of this type shall be of one of the following variants:

- The value may be a set of attributes. Each attribute defines an operator. Exactly one of the explicit attributes shall be specified. The explicit attributes for this type are:

and

The attribute value shall be a list. Each list element shall be a *Specification Integer Assert* (page 100). The *and* operator evaluates to the *logical and* of the evaluation results of the expressions in the list.

eq

The attribute value shall be an integer number. The *eq* operator evaluates to true, if the value to check is equal to the value of this attribute, otherwise to false.

ge

The attribute value shall be an integer number. The *ge* operator evaluates to true, if the value to check is greater than or equal to the value of this attribute, otherwise to false.

gt

The attribute value shall be an integer number. The *gt* operator evaluates to true, if the value to check is greater than the value of this attribute, otherwise to false.

le

The attribute value shall be an integer number. The *le* operator evaluates to true, if the value to check is less than or equal to the value of this attribute, otherwise to false.

lt

The attribute value shall be an integer number. The *lt* operator evaluates to true, if the value to check is less than the value of this attribute, otherwise to false.

ne

The attribute value shall be an integer number. The *ne* operator evaluates to true, if the value to check is not equal to the value of this attribute, otherwise to false.

not

The attribute value shall be a *Specification Integer Assert* (page 100). The *not* operator evaluates to the *logical not* of the evaluation results of the expression.

or

The attribute value shall be a list. Each list element shall be a *Specification Integer Assert* (page 100). The *or* operator evaluates to the *logical or* of the evaluation results of the expressions in the list.

- The value may be a list. Each list element shall be a *Specification Integer Assert* (page 100). This list of expressions evaluates to the *logical or* of the evaluation results of the expressions in the list.

This type is used by the following types:

- *Specification Integer Assert* (page 100)
- *Specification Integer Value* (page 101)

5.2.3.96 Specification Integer Value

This set of attributes specifies an integer value. Only the description attribute is mandatory. The explicit attributes for this type are:

assert

The attribute value shall be a *Specification Integer Assert* (page 100). This optional attribute defines the value constraints of the specified integer value. In case the assert attribute is not present, then the value of the specified integer value may be every valid integer number.

description

The attribute value shall be an optional string. It shall be the description of the specified integer value.

This type is used by the following types:

- *Specification Information* (page 99)

5.2.3.97 Specification List

This set of attributes specifies a list of attributes or values. All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be an optional string. It shall be the description of the list.

spec-type

The attribute value shall be a *Name* (page 89). It shall be the specification type of elements of the list.

This type is used by the following types:

- *Specification Information* (page 99)

5.2.3.98 Specification Mandatory Attributes

It defines which explicit attributes are mandatory.

A value of this type shall be of one of the following variants:

- The value may be a list. Each list element shall be a *Name* (page 89). The list defines the mandatory attributes through their key names.
- The value may be a string. It defines how many explicit attributes are mandatory. If *none* is used, then none of the explicit attributes is mandatory, they are all optional. The value shall be an element of
 - “all”,
 - “at-least-one”,
 - “at-most-one”,
 - “exactly-one”, and
 - “none”.

This type is used by the following types:

- *Specification Attribute Set* (page 96)

5.2.3.99 Specification Member Link Role

This type refines the *Link* (page 88) though the *role* attribute if the value is *spec-member*. It defines the specification membership role of links.

5.2.3.100 Specification Refinement Link Role

This type refines the *Link* (page 88) though the *role* attribute if the value is *spec-refinement*. It defines the specification refinement role of links. All explicit attributes shall be specified. The explicit attributes for this type are:

spec-key

The attribute value shall be a *Name* (page 89). It shall be the specification type refinement attribute key of the specification refinement.

spec-value

The attribute value shall be a *Name* (page 89). It shall be the specification type refinement attribute value of the specification refinement.

5.2.3.101 Specification String Assert

A value of this type shall be an expression which asserts that the string of the specified attribute satisfies the required constraints.

A value of this type shall be of one of the following variants:

- The value may be a set of attributes. Each attribute defines an operator. Exactly one of the explicit attributes shall be specified. The explicit attributes for this type are:

and

The attribute value shall be a list. Each list element shall be a *Specification String Assert* (page 102). The *and* operator evaluates to the *logical and* of the evaluation results of the expressions in the list.

contains

The attribute value shall be a list of strings. The *contains* operator evaluates to true, if the string to check converted to lower case with all white space characters converted to a single space character contains a string of the list of strings of this attribute, otherwise to false.

eq

The attribute value shall be a string. The *eq* operator evaluates to true, if the string to check is equal to the value of this attribute, otherwise to false.

ge

The attribute value shall be a string. The *ge* operator evaluates to true, if the string to check is greater than or equal to the value of this attribute, otherwise to false.

gt

The attribute value shall be a string. The *gt* operator evaluates to true, if the string to check is greater than the value of this attribute, otherwise to false.

in

The attribute value shall be a list of strings. The *in* operator evaluates to true, if the string to check is contained in the list of strings of this attribute, otherwise to false.

le

The attribute value shall be a string. The *le* operator evaluates to true, if the string to check is less than or equal to the value of this attribute, otherwise to false.

lt

The attribute value shall be a string. The *lt* operator evaluates to true, if the string to check is less than the value of this attribute, otherwise to false.

ne

The attribute value shall be a string. The *ne* operator evaluates to true, if the string to check is not equal to the value of this attribute, otherwise to false.

not

The attribute value shall be a *Specification String Assert* (page 102). The *not* operator evaluates to the *logical not* of the evaluation results of the expression.

or

The attribute value shall be a list. Each list element shall be a *Specification String Assert* (page 102). The *or* operator evaluates to the *logical or* of the evaluation results of the expressions in the list.

re

The attribute value shall be a string. The *re* operator evaluates to true, if the string to check matches with the regular expression of this attribute, otherwise to false.

uid

The attribute shall have no value. The *uid* operator evaluates to true, if the string is a valid UID, otherwise to false.

- The value may be a list. Each list element shall be a *Specification String Assert* (page 102). This list of expressions evaluates to the *logical or* of the evaluation results of the expressions in the list.

This type is used by the following types:

- *Specification String Assert* (page 102)
- *Specification String Value* (page 103)

5.2.3.102 Specification String Value

This set of attributes specifies a string. Only the *description* attribute is mandatory. The explicit attributes for this type are:

assert

The attribute value shall be a *Specification String Assert* (page 102). This optional attribute defines the constraints of the specified string. In case the *assert* attribute is not present, then the specified string may be every valid string.

description

The attribute value shall be an optional string. It shall be the description of the specified string attribute.

This type is used by the following types:

- *Specification Information* (page 99)

5.2.3.103 Test Case Action

This set of attributes specifies a test case action. All explicit attributes shall be specified. The explicit attributes for this type are:

action

The attribute value shall be a string. It shall be the test case action code.

checks

The attribute value shall be a list. Each list element shall be a *Test Case Check* (page 104).

description

The attribute value shall be an optional string. It shall be the test case action description.

links

The attribute value shall be a list. Each list element shall be a *Link* (page 88).

This type is used by the following types:

- *Test Case Item Type* (page 60)

5.2.3.104 Test Case Check

This set of attributes specifies a test case check. All explicit attributes shall be specified. The explicit attributes for this type are:

check

The attribute value shall be a string. It shall be the test case check code.

description

The attribute value shall be an optional string. It shall be the test case check description.

links

The attribute value shall be a list. Each list element shall be a *Link* (page 88).

This type is used by the following types:

- *Test Case Action* (page 104)

5.2.3.105 Test Context Member

A value of this type shall be of one of the following variants:

- The value may be a set of attributes. This set of attributes defines an action requirement test context member. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an optional string. It shall be the test context member brief description.

description

The attribute value shall be an optional string. It shall be the test context member description.

member

The attribute value shall be a string. It shall be the test context member definition. It shall be a valid C structure member definition without a trailing ; .

- There may be no value (null).

This type is used by the following types:

- *Action Requirement Item Type* (page 50)
- *Runtime Measurement Test Item Type* (page 58)

5.2.3.106 Test Header

A value of this type shall be of one of the following variants:

- The value may be a set of attributes. This set of attributes specifies a test header. In case a test header is specified, then instead of a test case a test run function will be generated. The test run function will be declared in the test header target file and defined in the test source target file. The test run function can be used to compose test cases. The test header file is not automatically included in the test source file. It should be added to the includes or local includes of the test. All explicit attributes shall be specified. The explicit attributes for this type are:

code

The attribute value shall be an optional string. If the value is present, then it shall be the test header code. The header code is placed at file scope after the general test declarations and before the test run function declaration.

includes

The attribute value shall be a list of strings. It shall be a list of header files included by the header file via `#include <...>`.

local-includes

The attribute value shall be a list of strings. It shall be a list of header files included by the header file via `#include "..."`.

run-params

The attribute value shall be a list. Each list element shall be a *Test Run Parameter* (page 106).

target

The attribute value shall be a string. It shall be the path to the generated test header file.

- There may be no value (null).

This type is used by the following types:

- *Action Requirement Item Type* (page 50)
- *Test Case Item Type* (page 60)

5.2.3.107 Test Run Parameter

This set of attributes specifies a parameter for the test run function. In case this parameter is used in an *Action Requirement Item Type* (page 50) item, then the parameter is also added as a member to the test context, see *Test Context Member* (page 104). All explicit attributes shall be specified. The explicit attributes for this type are:

description

The attribute value shall be a string. It shall be the description of the parameter.

dir

The attribute value shall be an *Interface Parameter Direction* (page 87).

name

The attribute value shall be a string. It shall be the parameter name.

specifier

The attribute value shall be a string. It shall be the complete function parameter specifier. Use `${.:name}` for the parameter name, for example "int `${.:name}`".

This type is used by the following types:

- *Test Header* (page 105)

5.2.3.108 Test Support Method

A value of this type shall be of one of the following variants:

- The value may be a set of attributes. This set of attributes defines an action requirement test support method. All explicit attributes shall be specified. The explicit attributes for this type are:

brief

The attribute value shall be an optional string. It shall be the test support method brief description.

code

The attribute value shall be a string. It shall be the test support method code. The code may use a local variable `ctx` which points to the test context, see *Test Context Member* (page 104).

description

The attribute value shall be an optional string. It shall be the test support method description.

- There may be no value (null).

This type is used by the following types:

- *Action Requirement Item Type* (page 50)
- *Runtime Measurement Test Item Type* (page 58)
- *Runtime Performance Requirement Item Type* (page 55)

5.2.3.109 UID

The value shall be a string. The string shall be a valid absolute or relative item UID.

This type is used by the following types:

- *Link* (page 88)

5.3 Traceability of Specification Items

The standard ECSS-E-ST-10-06C demands that requirements shall be under configuration management, backwards-traceable and forward-traceable [ECS09]. Requirements are a specialization of specification items in RTEMS.

5.3.1 History of Specification Items

The RTEMS specification items should be placed in the RTEMS sources using Git for version control. The history of specification items can be traced with Git. Special commit procedures for changes in specification item files should be established. For example, it should be allowed to change only one specification item per commit. A dedicated Git commit message format may be used as well, e.g. use of `Approved-by:` or `Reviewed-by:` lines which indicate an agreed statement (similar to the [Linux kernel patch submission guidelines](#)). Git commit procedures may be ensured through a server-side pre-receive hook. The history of requirements may be also added to the specification items directly in a *revision* attribute. This would make it possible to generate the history information for documents without having the Git repository available, e.g. from an RTEMS source release archive.

5.3.2 Backward Traceability of Specification Items

Providing backward traceability of specification items means that we must be able to find the corresponding higher level specification item for each refined specification item. A custom tool needs to verify this.

5.3.3 Forward Traceability of Specification Items

Providing forward traceability of specification items means that we must be able to find all the refined specification items for each higher level specification item. A custom tool needs to verify this. The links from parent to child specification items are implicitly defined by links from a child item to a parent item.

5.3.4 Traceability between Software Requirements, Architecture and Design

The software requirements are implemented in custom YAML files, see *Specification Items* (page 25). The software architecture and design is written in Doxygen markup. Doxygen markup is used throughout all header and source files. A Doxygen filter program may be provided to place Doxygen markup in assembler files. The software architecture is documented via Doxygen groups. Each Doxygen group name should have a project-specific name and the name should be unique within the project, e.g. `RTEMS_TopLevel_MidLevel_LowLevel`. The link from a Doxygen group to its parent group is realized through the `@ingroup` special command. The link from a Doxygen group or *software component* to the corresponding requirement is realized through a `@satisfy{req}` **custom command** which needs the identifier of the requirement as its one and only parameter. Only links to parents are explicitly given in the Doxygen markup. The links from a parent to its children are only implicitly specified via the link from a child to its parent. So, a tool must process all files to get the complete hierarchy of software requirements, architecture and design. Links from a software component to another software component are realized through automatic Doxygen references or the `@ref` and `@see` special commands.

5.4 Requirement Management

5.4.1 Change Control Board

Working with requirements usually involves a Change Control Board (*CCB*). The *CCB* of the RTEMS Project is the [RTEMS developer mailing list](#).


There are the following actors involved:

- *RTEMS users*: Everyone using the RTEMS real-time operating system to design, develop and build an application on top of it.
- *RTEMS developers*: The persons developing and maintaining RTEMS. They write patches to add or modify code, requirements, tests and documentation.
- *RTEMS maintainers*: They are listed in the [MAINTAINERS](#) file and have write access to the project repositories.


Adding and changing requirements follows the normal patch review process. The normal patch review process is described in the [RTEMS User Manual](#). Reviews and comments may be submitted by anyone, but a maintainer review is required to approve *significant* changes. In addition for significant changes, there should be at least one reviewer with a sufficient independence from the author which proposes a new requirement or a change of an existing requirement. Working in another company on different projects is sufficiently independent. RTEMS maintainers do not know all the details, so they trust in general people with experience on a certain platform. Sometimes no review comments may appear in a reasonable time frame, then an implicit agreement to the proposed changes is assumed. Patches can be sent at anytime, so controlling changes in RTEMS requires a permanent involvement on the RTEMS developer mailing list.

For a qualification of RTEMS according to certain standards, the requirements may be approved by an RTEMS user. The approval by RTEMS users is not the concern of the RTEMS Project, however, the RTEMS Project should enable RTEMS users to manage the approval of requirements easily. This information may be also used by a independent authority which comes into play with an Independent Software Verification and Validation (*ISVV*). It could be used to select a subset of requirements, e.g. look only at the ones approved by a certain user. RTEMS users should be able to reference the determinative content of requirements, test procedures, test cases and justification reports in their own documentation. Changes in the determinative content should invalidate all references to previous versions.

5.4.2 Add a Requirement



5.4.3 Modify a Requirement



5.4.4 Mark a Requirement as Obsolete

Requirements shall be never removed. They shall be marked as obsolete. This ensures that requirement identifiers are not reused. The procedure to obsolete a requirement is the same as the one to *modify a requirement* (page 111).

5.5 Tooling

5.5.1 Tool Requirements

To manage requirements some tool support is helpful. Here is a list of requirements for the tool:

- The tool shall be open source.
- The tool should be actively maintained during the initial phase of the RTEMS requirements specification.
- The tool shall use plain text storage (no binary formats, no database).
- The tool shall support version control via Git.
- The tool should export the requirements in a human readable form using the Sphinx documentation framework.
- The tool shall support traceability of requirements to items external to the tool.
- The tool shall support traceability between requirements.
- The tool shall support custom requirement attributes.
- The tool should ensure that there are no cyclic dependencies between requirements.
- The tool should provide an export to *ReqIF*.

5.5.2 Tool Evaluation

During an evaluation phase the following tools were considered:

- [aNimble](#)
- *Doorstop*
- [OSRMT](#)
- [Papyrus](#)
- [ProR](#)
- [ReqIF Studio](#)
- [Requirement Heap](#)
- [rmToo](#)

The tools aNimble, OSRMT and Requirement Heap were not selected since they use a database. The tools Papyrus, ProR and ReqIF are Eclipse based and use complex XML files for data storage. They were difficult to use and lack good documentation/tutorials. The tools rmToo and Doorstop turned out to be the best candidates to manage requirements in the RTEMS Project. The Doorstop tool was selected as the first candidate mainly due a recommendation by an RTEMS user.

5.5.3 Best Available Tool - Doorstop

Doorstop is a requirements management tool. It has a modern, object-oriented and well-structured implementation in Python 3.6 under the LGPLv3 license. It uses a continuous integration build with style checkers, static analysis, documentation checks, code coverage, unit test and integration tests. In 2019, the project was actively maintained. Pull requests for minor improvements and new features were reviewed and integrated within days. Each requirement is contained in a single file in *YAML* format. Requirements are organized in documents and can be linked to each other [BA14].

Doorstop consists of three main parts

- a stateless command line tool *doorstop*,
- a file format with a pre-defined set of attributes (*YAML*), and
- a primitive GUI tool (not intended to be used).

For RTEMS, its scope could be extended to manage specifications in general. The primary reason for a close consideration of Doorstop as the requirements management tool for the RTEMS Project was its data format which allows a high degree of customization. Doorstop uses a directed, acyclic graph (DAG) of items. The items are files in *YAML* format. Each item has a set of **standard attributes** (key-value pairs).

The use case for the standard attributes is requirements management. However, Doorstop is capable to manage custom attributes as well. We will heavily use custom attributes for the specification items. Enabling Doorstop to effectively use custom attributes was done specifically for the RTEMS Project in several patch sets which in the end turned out to be not enough to use Doorstop for the RTEMS Project.

A key feature of Doorstop is the **fingerprint of items**. For the RTEMS Project, the fingerprint hash algorithm was changed from MD5 to SHA256. In 2019, it can be considered cryptographically secure. The fingerprint should cover the normative values of an item, e.g. comments etc. are not included. The fingerprint would help RTEMS users to track the significant changes in the requirements (in contrast to all the changes visible in Git). As an example use case, a user may want to assign a project-specific status to specification items. This can be done with a table which contains columns for

1. the UID of the item,
2. the fingerprint, and
3. the project-specific status.

Given the source code of RTEMS (which includes the specification items) and this table, it can be determined which items are unchanged and which have another status (e.g. unknown, changed, etc.).

After some initial work with Doorstop some issues surfaced ([#471](#)). It turned out that Doorstop is not designed as a library and contains too much policy. This results in a lack of flexibility required for the RTEMS Project.

1. Its primary use case is requirements management. So, it has some standard attributes useful in this domain, like derived, header, level, normative, ref, reviewed, and text. However, we want to use it more generally for specification items and these attributes make not always sense. Having them in every item is just overhead and may cause confusion.

2. The links cannot have custom attributes, e.g. role, enabled-by. With link-specific attributes you could have multiple DAGs formed up by the same set of items.
3. Inside a document (directory) items are supposed to have a common type (set of attributes). We would like to store at a hierarchy level also distinct specializations.
4. The verification of the items is quite limited. We need verification with type-based rules.
5. The UIDs in combination with the document hierarchy lead to duplication, e.g. a/b/c/a-b-c-d.yml. You have the path (a/b/c) also in the file name (a-b-c). You cannot have relative UIDs in links (e.g. ../parent-req) . The specification items may contain multiple requirements, e.g. min/max attributes. There is no way to identify them.
6. The links are ordered by Doorstop alphabetically by UID. For some applications, it would be better to use the order specified by the user. For example, we want to use specification items for a new build system. Here it is handy if you can express things like this: A is composed of B and C. Build B before C.

5.5.4 Custom Requirements Management Tool

No requirements management tool was available that fits the need of the RTEMS Qualification Project. The decision was to develop a custom requirements management tool written in Python 3.6 or later. The design for it is heavily inspired by Doorstop.

5.6 How-To

5.6.1 Getting Started

The RTEMS specification items and qualification tools are work in progress. The first step to work with the RTEMS specification and the corresponding tools is a clone of the following repository:

```
1 git clone git://git.rtems.org/rtems-central.git
2 git submodule init
3 git submodule update
```

The tools need a virtual Python 3 environment. To set it up use:

```
1 cd rtems-central
2 make env
```

Each time you want to use one of the tools, you have to activate the environment in your shell:

```
1 cd rtems-central
2 . env/bin/activate
```

5.6.2 Application Configuration Options

The application configuration options and groups are maintained by specification items in the directory `spec/if/acfg`. Application configuration options are grouped by *Application Configuration Group Item Type* (page 41) items which should be stored in files using the `spec/if/acfg/group-*.yml` pattern. Each application configuration option shall link to exactly one group item with the *Application Configuration Group Member Link Role* (page 66). There are four application option item types available which cover all existing options:

- The *feature enable options* let the application enable a feature option. If the option is not defined, then the feature is simply not available or active. There should be no feature-specific code linked to the application if the option is not defined. Examples are options which enable a device driver like `CONFIGURE_APPLICATION_NEEDS_CLOCK_DRIVER`. These options are specified by *Application Configuration Feature Enable Option Item Type* (page 42) items.
- The *feature options* let the application enable a specific feature option. If the option is not defined, then a default feature option is used. Regardless whether the option is defined or not defined, feature-specific code may be linked to the application. Examples are options which disable a feature if the option is defined such as `CONFIGURE_APPLICATION_DISABLE_FILESYSTEM` and options which provide a stub implementation of a feature by default and a full implementation if the option is defined such as `CONFIGURE_IMFS_ENABLE_MKFIFO`. These options are specified by *Application Configuration Feature Option Item Type* (page 43) items.
- The *integer value options* let the application define a specific value for a system parameter. If the option is not defined, then a default value is used for the system parameter. Examples are options which define the maximum count of objects available for application use such as `CONFIGURE_MAXIMUM_TASKS`. These options are specified by *Application Configuration Value Option Item Type* (page 43) items.

- The *initializer options* let the application define a specific initializer for a system parameter. If the option is not defined, then a default setting is used for the system parameter. An example option of this type is `CONFIGURE_INITIAL_EXTENSIONS`. These options are specified by *Application Configuration Value Option Item Type* (page 43) items.

Sphinx documentation sources and header files with Doxygen markup are generated from the specification items. The descriptions in the items shall use a restricted Sphinx formatting. Emphasis via one asterisk (“*”), strong emphasis via two asterisk (“**”), code samples via blockquotes (“`”), code blocks (“.. code-block:: c”) and lists are allowed. References to interface items are also allowed, for example “`{appl-needs-clock-driver:/name}`” and “`{./rtems/tasks/create:/name}`”. References to other parts of the documentation are possible, however, they are currently provided by hard-coded tables in `rtemsspec/applconfig.py`.

5.6.2.1 Modify an Existing Group

Search for the group by its section header and edit the specification item file. For example:

```
1 $ grep -rl "name: General System Configuration" spec/if/acfg
2 spec/if/acfg/group-general.yml
3 $ vi spec/if/acfg/group-general.yml
```

5.6.2.2 Modify an Existing Option

Search for the option by its C preprocessor define name and edit the specification item file. For example:

```
1 $ grep -rl CONFIGURE_APPLICATION_NEEDS_CLOCK_DRIVER spec/if/acfg
2 spec/if/acfg/appl-needs-clock-driver.yml
3 $ vi spec/if/acfg/appl-needs-clock-driver.yml
```

5.6.2.3 Add a New Group

Let `new` be the UID name part of the new group. Create the file `spec/if/acfg/group-new.yml` and provide all attributes for an *Application Configuration Group Item Type* (page 41) item. For example:

```
1 $ vi spec/if/acfg/group-new.yml
```

5.6.2.4 Add a New Option

Let `my-new-option` be the UID name of the option. Create the file `if/acfg/my-new-option.yml` and provide all attributes for an appropriate refinement of *Application Configuration Option Item Type* (page 42). For example:

```
1 $ vi spec/if/acfg/my-new-option.yml
```


5.6.2.5 Generate Content after Changes

Once you are done with the modifications of an existing item or the creation of a new item, the changes need to be propagated to generated source files. This is done by the `spec2modules.py` script. Before you call this script, make sure the Git submodules are up-to-date.

```
1 $ ./spec2modules.py
```

The script modifies or creates source files in `modules/rtems` and `modules/rtems-docs`. Create patch sets for these changes just as if these were work done by a human and follow the normal patch review process described in the *RTEMS User Manual*. When the changes are integrated, update the Git submodules and check in the changed items.

5.6.3 Glossary Specification

The glossary of terms for the RTEMS Project is defined by *Glossary Term Item Type* (page 40) items in the `spec/glossary` directory. For a new glossary term add a glossary item to this directory. As the file name use the term in lower case with all white space and special characters removed or replaced by alphanumeric characters, for example `spec/glossary/magicpower.yml` for the term *magic power*.

Use `${uid:/attribute}` substitutions to reference other parts of the specification.

```
1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 copyrights:
3 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
4 enabled-by: true
5 glossary-type: term
6 links:
7 - role: glossary-member
8   uid: ../glossary-general
9 term: magic power
10 text: |
11   Magic power enables a caller to create magic objects using a
12   ${magicwand:/term}.
13 type: glossary
```

Define acronyms with the phrase *This term is an acronym for **. in the text attribute:

```
1 ...
2 term: MP
3 ...
4 text: |
5   This term is an acronym for Magic Power.
6 ...
```

Once you are done with the glossary items, run the script `spec2modules.py` to generate the derived documentation content. Send patches for the generated documentation and the specification to the [Developers Mailing List](#) and follow the normal patch review process.

5.6.4 Interface Specification

5.6.4.1 Specify an API Header File

The RTEMS API header files are specified under `spec:/if/rtems/*`. Create a subdirectory with a corresponding name for the API, for example in `spec/if/rtems/foo` for the `foo` API. In this new subdirectory place an *Interface Header File Item Type* (page 47) item named `header.yml` (`spec/if/rtems/foo/header.yml`) and populate it with the required attributes.

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 copyrights:
3 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
4 enabled-by: true
5 interface-type: header-file
6 links:
7 - role: interface-placement
8   uid: /if/domains/api
9 path: rtems/rtems/foo.h
10 prefix: cpukit/include
11 type: interface

```

5.6.4.2 Specify an API Element

Figure out the corresponding header file item. If it does not exist, see *Specify an API Header File* (page 118). Place a specialization of an *Interface Item Type* (page 41) item into the directory of the header file item, for example `spec/if/rtems/foo/bar.yml` for the `bar()` function. Add the required attributes for the new interface item. Do not hard code interface names which are used to define the new interface. Use `#{uid-of-interface-item:/name}` instead. If the referenced interface is specified in the same directory, then use a relative UID. Using interface references creates implicit dependencies and helps the header file generator to resolve the interface dependencies and header file includes for you. Use *Interface Unspecified Item Type* (page 48) items for interface dependencies to other domains such as the C language, the compiler, the implementation, or user-provided defines. To avoid cyclic dependencies between types you may use an *Interface Forward Declaration Item Type* (page 45) item.

```

1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 brief: Tries to create a magic object and returns it.
3 copyrights:
4 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
5 definition:
6   default:
7     body: null
8     params:
9     - #{magic-wand:/name} #{./params[0]/name}
10    return: #{magic-type:/name} *
11    variants: []
12 description: |
13   The magic object is created out of nothing with the help of a magic wand.
14 enabled-by: true
15 interface-type: function
16 links:
17 - role: interface-placement
18   uid: header

```

(continues on next page)

(continued from previous page)

```
19 - role: interface-ingroup
20   uid: /groups/api/classic/foo
21 name: bar
22 notes: null
23 params:
24 - description: is the magic wand.
25   dir: null
26   name: magic_wand
27 return:
28   return: Otherwise, the magic object is returned.
29   return-values:
30   - description: The caller did not have enough magic power.
31     value: ${if/c/null}
32 type: interface
```


**SOFTWARE DEVELOPMENT
MANAGEMENT**

6.1 Software Development (Git Users)

6.1.1 Browse the Git Repository Online

You can browse all available repositories online by accessing <https://git.rtems.org/>.

6.1.2 Using the Git Repository

The following examples demonstrate how to use the RTEMS' Git repos. These examples are provided for the main rtems module, but they are also valid for the other modules.

First, we need to obtain our own local copy of the RTEMS Git repository:

```
1 git clone git://git.rtems.org/rtems.git rtems
```

This command will create a folder named rtems in the current directory. This folder will contain a full-featured RTEMS' Git repository and the current HEAD revision checked out. Since all the history is available we can check out any release of RTEMS. Major RTEMS releases are available as separate branches in the repo.

To see all available remote branches issue the following command:

```
1 git branch -r
```

We can check out one of those remote branches (e.g. rtems-4.10 branch) using the command:

```
1 git checkout -b rtems410 origin/4.10
```

This will create a local branch named "rtems410", containing the rtems-4.10 release, that will track the remote branch "rtems-4-10-branch" in origin (git://git.rtems.org/rtems.git). The git branch command prints a list of the current local branches, indicating the one currently checked out.

If you want to switch between local branches:

```
1 git checkout <branch-name>
```

With time your local repository will diverge from the main RTEMS repository. To keep your local copy up to date you need to issue:

```
1 git pull origin
```

This command will update all your local branches with any new code revisions available on the central repository.

6.1.3 Making Changes

Git allows you to make changes in the RTEMS source tree and track those changes locally. We recommend you make all your changes in local branches. If you are working on a few different changes or a progression of changes it is best to use a local branch for each change.

A branch for each change lets your repo's master branch track the upstream RTEMS' master branch without interacting with any of the changes you are working on. A completed change is emailed to the developer's list for review and this can take time. While this is happening the upstream's master branch may be updated and you may need to rebase your work and test again if you are required to change or update your patch. A local branch isolates a specific change from others and helps you manage the process.

First, you need to clone the repository:

```
1 git clone git://git.rtems.org/rtems.git rtems
```

Or if you already cloned it before, then you might want to update to the latest version before making your changes:

```
1 cd rtems
2 git pull
```

Create a local branch to make your changes in, in this example, the change is faster-context-switch:

```
1 git checkout -b faster-context-switch
```

Next, make your changes to files. If you add, delete or move/rename files you need to inform Git

```
1 git add /some/new/file
2 git rm /some/old/file
3 git mv /some/old/file /some/new/file
```

When you're satisfied with the changes you made, commit them (locally)

```
1 git commit -a
```

The `-a` flag commits all the changes that were made, but you can also control which changes to commit by individually adding files as you modify them by using `git add`. You can also specify other options to commit, such as a message with the `-m` flag.

```
1 git add /some/changed/files
2 git commit
```

Create a patch from your branch, in this case, we have two commits we want to send for review:

```
1 git format-patch -2
2
3 There are new changes pushed to the RTEMS' master branch and our local branch
4 needs to be updated:
```

```
1 git checkout master
2 git pull
```

(continues on next page)

(continued from previous page)

```
3 git checkout faster-context-switch
4 git rebase master
```

6.1.4 Working with Branches

Branches facilitate trying out new code and creating patches.

The previous releases of RTEMS are available through remote branches. To check out a remote branch, first query the Git repository for the list of branches:

```
1 git branch -r
```

Then check out the desired remote branch, for example:

```
1 git checkout -b rtems410 origin/4.10
```

Or if you have previously checked out the remote branch then you should see it in your local branches:

```
1 git branch
```

You can change to an existing local branch easily:

```
1 git checkout rtems410
```

You can also create a new branch and switch to it:

```
1 git branch temporary
2 git checkout temporary
```

Or more concisely:

```
1 git checkout -b temporary
```

If you forget which branch you are on

```
1 git branch
```

shows you by placing a * next to the current one.

When a branch is no longer useful you can delete it.

```
1 git checkout master
2 git branch -d temporary
```

If you have unmerged changes in the old branch Git complains and you need to use `-D` instead of `-d`.

6.1.5 Viewing Changes

To view all changes since the last commit:

```
1 git diff HEAD
```

To view all changes between the current branch and another branch, say master:

```
1 git diff master..HEAD
```

To view descriptions of committed changes:

```
1 git log
```

Or view the changeset for some file (or directory):

```
1 git log /some/file
```

To view the changesets made between two branches:

```
1 git log master..HEAD
```

Or for a more brief description use shortlog:

```
1 git shortlog master..HEAD
```

6.1.6 Reverting Changes

To remove all (uncommitted) changes on a branch

```
1 git checkout -f
```

Or to selectively revert (uncommitted) files, for example if you accidentally deleted ./some/file

```
1 git checkout -- ./some/file
```

or

```
1 git checkout HEAD ./some/file
```

To remove commits there are two useful options, reset and revert. `git reset` should only be used on local branches that no one else is accessing remotely. `git revert` is cleaner and is the right way to revert changes that have already been pushed/pulled remotely.

6.1.7 git reset

`git reset` is a powerful and tricky command that should only be used on local (un-pushed) branches): A good description of what it enables to do can be found [here](#). The following are a few useful examples. Note that adding a `~` after `HEAD` refers to the most recent commit, and you can add a number after the `~` to refer to commits even further back; `HEAD` by itself refers to the current working directory (changes since the last commit).

```
1 git reset HEAD~
```

Will undo the last commit and unstage those changes. Your working directory will remain the same, therefore a `git status` will yield any changes you made plus the changes made in your last commit. This can be used to fix the last commit. You will need to add the files again.

```
1 git reset --soft HEAD~
```

Will just undo the last commit. The changes from the last commit will still be staged (just as if you finished `git add`ing them). This can be used to amend the last commit (e.g. You forgot to add a file to the last commit).

```
1 git reset --hard HEAD~
```

Will revert everything, including the working directory, to the previous commit. This is dangerous and can lead to you losing all your changes; the `--hard` flag ignores errors.

```
1 git reset HEAD
```

Will unstage any change. This is used to revert a wrong `git add`. (e.g. You added a file that shouldn't be there, but you haven't 'committed')

Will revert your working directory to a `HEAD` state. You will lose any change you made to files after the last commit. This is used when you just want to destroy all changes you made since the last commit.

6.1.8 git revert

`git revert` does the same as `reset` but creates a new commit with the reverted changes instead of modifying the local repository directly.

```
1 git revert HEAD
```

This will create a new commit which undoes the change in `HEAD`. You will be given a chance to edit the commit message for the new commit.

6.1.9 Merging Changes

Suppose you commit changes in two different branches, `branch1` and `branch2`, and want to create a new branch containing both sets of changes:

```
1 git checkout -b merged
2 git merge branch1
3 git merge branch2
```

Or you might want to bring the changes in one branch into the other:

```
1 git checkout branch1
2 git merge branch2
```

And now that `branch2` is merged you might get rid of it:

```
1 git branch -d branch2
```

If you have done work on a branch, say `branch1`, and have gone out-of-sync with the remote repository, you can pull the changes from the remote repo and then merge them into your branch:

```
1 git checkout master
2 git pull
3 git checkout branch1
4 git merge master
```

If all goes well the new commits you pulled into your master branch will be merged into your `branch1`, which will now be up-to-date. However, if `branch1` has not been pushed remotely then rebasing might be a good alternative to merging because the merge generates a commit.

6.1.10 Rebasing

An alternative to the merge command is rebase, which replays the changes (commits) on one branch onto another. `git rebase` finds the common ancestor of the two branches, stores each commit of the branch you are on to temporary files and applies each commit in order.

For example

```
1 git checkout branch1
2 git rebase master
```

or more concisely

```
1 git rebase master branch1
```

will bring the changes of master into `branch1`, and then you can fast-forward master to include `branch1` quite easily

```
1 git checkout master
2 git merge branch1
```

Rebasing makes a cleaner history than merging; the log of a rebased branch looks like a linear history as if the work was done serially rather than in parallel. A primary reason to rebase is to ensure commits apply cleanly on a remote branch, e.g. when submitting patches to RTEMS that

you create by working on a branch in a personal repository. Using rebase to merge your work with the remote branch eliminates most integration work for the committer/maintainer.

There is one caveat to using rebase: Do not rebase commits that you have pushed to a public repository. Rebase abandons existing commits and creates new ones that are similar but different. If you push commits that others pull down, and then you rewrite those commits with `git rebase` and push them up again, the others will have to re-merge their work and trying to integrate their work into yours can become messy.

6.1.11 Accessing a developer's repository

RTEMS developers with Git commit access have personal repositories on <https://git.rtems.org/> that can be cloned to view cutting-edge development work shared there.

6.1.12 Creating a Patch

Before submitting a patch read about [Contributing](#) to RTEMS and the [Commit Message](#) formatting we require.

The recommended way to create a patch is to branch the Git repository master and use one commit for each logical change. Then you can use `git format-patch` to turn your commits into patches and easily submit them.

```
1 git format-patch master
```

Creates a separate patch for each commit that has been made between the master branch and the current branch and writes them in the current directory. Use the `-o` flag to redirect the files to a different directory.

If you are re-submitting a patch that has previously been reviewed, you should specify a version number for your patch, for example, use

```
1 git format-patch -v2 ...
```

to indicate the second version of a patch, `-v3` for a third, and so forth.

Patches created using `git format-patch` are formatted so they can be emailed and rely on having Git configured with your name and email address, for example

```
1 git config --global user.name "Your Name"  
2 git config --global user.email name@domain.com
```

Please use a real name, we do not allow pseudonyms or anonymous contributions.

6.1.13 Submitting a Patch

Using `git send-email` you can easily contribute your patches. You will need to install `git send-email` first:

```
1 sudo yum install git-email
```

or

```
1 sudo dnf install git-email
```

or

```
1 sudo apt install git-email
```

Then you will need to configure an SMTP server. You could install one on your localhost, or you can connect to a mail server such as Gmail.

6.1.14 Configuring git send-email to use Gmail

Configure Git to use Gmail:

```
1 git config --global sendemail.smtpserver smtp.gmail.com
2 git config --global sendemail.smtpserverport 587
3 git config --global sendemail.smtpencryption tls
4 git config --global sendemail.smtpuser your_email@gmail.com
```

It will ask for your password each time you use `git send-email`. Optionally you can also put it in your `git config`:

```
1 git config --global sendemail.smtppass your_password
```

6.1.15 Sending Email

To send your patches just

```
1 git send-email /path/to/patch --to devel@rtems.org
```

To send multiple related patches (if you have more than one commit in your branch) specify a path to a directory containing all of the patches created by `git format-patch`. `git send-email` has some useful options such as:

- `--annotate` to show/edit your patch
- `--cover-letter` to prepend a summary
- `--cc=<address>` to cc someone

You can configure the to address:

```
1 git config --global sendemail.to devel@rtems.org
```

So all you need is:

```
1 git send-email /path/to/patch
```

6.1.16 Troubleshooting

Some restrictive corporate firewalls block access through the Git protocol (`git://`). If you are unable to reach the server `git://git.rtems.org/` you can try accessing through `http`. To clone the `rtems` repository using the `http` protocol use the following command:

```
1 git clone http://git.rtems.org/rtems/ rtems
```

This access through `http` is slower (way slower!) than through the `git` protocol, therefore, the `Git` protocol is preferred.

6.1.17 Manage Your Code

You may prefer to keep your application and development work in a `Git` repository for all the good reasons that come with version control. For public repositories, you may like to try [GitHub](#) or [BitBucket](#). RTEMS maintains [mirrors on GitHub](#) which can make synchronizing with upstream changes relatively simple. If you need to keep your work private, you can use one of those services with private repositories or manage your own server. The details of setting up a server are outside the scope of this document, but if you have a server with `SSH` access you should be able to [find instructions](#) on how to set up `Git` access. Once you have `git` configured on the server, adding repositories is a snap.

6.1.18 Private Servers

In the following, replace `@USER@` with your username on your server, `@REPO@` with the name of your repository, and `@SERVER@` with your server's name or address.

To push a mirror to your private server, first create a bare repository on your server.

```
1 cd /home/@USER@
2 mkdir git
3 mkdir git/@REPO@.git
4 cd git/@REPO@.git
5 git --bare init
```

Now from your client machine (e.g. your work laptop/desktop), push a `git`, perhaps one you cloned from elsewhere, or one that you made locally with `git init`, by adding a remote and pushing:

```
1 git remote add @SERVER@ ssh://@SERVER@/home/@USER@/git/@REPO@.git
2 git push @SERVER@ master
```

You can replace the `@SERVER@` with another name for your remote if you like. And now you can push other branches that you might have created. Now you can push and pull between your client and your server. Use `SSH` keys to authenticate with your server if you want to save on password typing; remember to put a passphrase on your `SSH` key if there is a risk the private key file might get compromised.

The following is an example scenario that might be useful for RTEMS users that uses a slightly different approach than the one just outlined:

```
1 ssh @SERVER@
2 mkdir git
3 git clone --mirror git://git.rtems.org/rtems.git
4 ## Add your ssh key to ~/.ssh/authorized_keys
5 exit
6 git clone ssh://@SERVER@/home/@USER@/git/rtems.git
7 cd rtems
8 git remote add upstream git://git.rtems.org/rtems.git
9 git fetch upstream
10 git pull upstream master
11 git push
12 ## If you want to track RTEMS on your personal master branch,
13 ## you should only push changes to origin/master that you pull
14 ## from upstream. The basic workflow should look something like:
15 git checkout master
16 git pull upstream master
17 git push
18 git checkout -b anewbranch
19 ## Repeat: do work, git commit -a
20 git push origin anewbranch
21
22 ## delete a remote branch
23 git push origin :anewbranch
24 ## delete a local branch
25 git branch -d anewbranch
```

6.1.19 Learn more about Git

Links to the sites with good Git information:

- <http://gitready.com/> - An excellent resource from beginner to very advanced.
- <http://progit.org/book/> - Covers Git basics and some advanced features. Includes some useful workflow examples.
- <https://lab.github.com/> - Learn to use Git and GitHub while doing a series of projects.
- <https://git-scm.com/docs> - The official Git reference.

6.2 Software Development (Git Writers)

6.2.1 SSH Access

Currently all committer's should have an ssh account on the main git server, `dispatch.rtems.org`. If you have been granted commit access and do have an account on `dispatch.rtems.org` one should be requested on the `devel@` list. SSH access for git uses key logins instead of passwords. The key should be at least 1024 bits in length.

The public repositories can be cloned with

```
1 git clone ssh://user@dispatch.rtems.org/data/git/rtems.git
```

Or replace `rtems.git` with another repo to clone another one.

6.2.2 Personal Repository

Personal repositories keep the clutter away from the master repository. A user with a personal repository can make commits, create and delete branches, plus more without interfering with the master repository. Commits to the master repository generate email to the `vc@` list and development type commits by a developer would only add noise and lessen the effectiveness of the commit list

A committer should maintain a personal clone of the RTEMS repository through which all changes merged into the RTEMS head are sent. The personal repository is also a good place for committers to push branches that contain works in progress. The following instructions show how to setup a personal repository that by default causes commits to go to your private local repository and pushes to go to your publicly visible personal repository. The RTEMS head is configured as a remote repository named 'upstream' to which you can push changes that have been approved for merging into RTEMS.

Branches aren't automatically pushed until you tell git to do the initial push after which the branch is pushed automatically. In order to keep code private just put it on a branch in your local clone and do not push the branch.

6.2.3 Create a personal repository

Set up the server side repository. In the following substitute user with your username.

```
1 # ssh git.rtems.org
2 [user@git ~]$ ln -s /data/git/user git
3 [user@git ~]$ ls -l
4 lrwxrwxrwx 1 user rtems 16 Feb  1 11:52 git -> /data/git/user
5 [user@git ~]$ cd git
6 [user@git git]$ git clone --mirror /data/git/rtems.git
```

Provide a description for the repository, for example "Clone of master repository."

```
1 [user@git git]$ echo "Clone of master repository." > rtems.git/description
2 [user@git git]$ logout
```


Clone the repository on your local machine


```
1 # git clone ssh://user@dispatch.rtems.org/home/user/git/rtems.git
2 # cd rtems
```

Add the RTEMS repository as a remote repository and get the remote tags and branches

```
1 # git remote add upstream ssh://user@dispatch.rtems.org/data/git/rtems.git
2 # git fetch upstream
```

After a little while you should be able to see your personal repo at <https://git.rtems.org/@USER@/rtems.git/> and you can create other repositories in your git directory that will propagate to <https://git.rtems.org/@USER@/> if you need. For example, *joel's* personal repos appear at <https://git.rtems.org/joel/>.



6.2.3.1 Check your setup

```
1 git remote show origin
```

Should print something similar to

```
1 * remote origin
2 Fetch URL: ssh://user@dispatch.rtems.org/home/user/git/rtems.git
3 Push URL: ssh://user@dispatch.rtems.org/home/user/git/rtems.git
4 HEAD branch: master
5 Remote branches:
6   4.10 tracked
7   4.8 tracked
8   4.9 tracked
9   master tracked
10 Local branch configured for 'git pull':
11   master merges with remote master
12 Local ref configured for 'git push':
13   master pushes to master (up to date)
```

6.2.3.2 Push commits to personal repo master from local master

```
1 # git push
```

6.2.3.3 Push a branch onto personal repo

```
1 # git push origin branchname
```

6.2.3.4 Update from upstream master (RTEMS head)

When you have committed changes on a branch that is private (hasn't been pushed to your personal repo) then you can use rebase to obtain a linear history and avoid merge commit messages.

```
1 # git checkout new_features
2 # git pull --rebase upstream master
```

If you cannot do a fast-forward merge then you could use the `--no-commit` flag to prevent merge from issuing an automatic merge commit message.

When you have committed changes on a branch that is public/shared with another developer you should not rebase that branch.

6.2.4 GIT Push Configuration

People with write access to the main repository should make sure that they push the right branch with the `git push` command. The above setup ensures that `git push` will not touch the main repository, which is identified as upstream, unless you specify the upstream (by `git push upstream master`).

Lets suppose we have a test branch intended for integration into the master branch of the main repository.

```
1 # git branch
2   master
3 * test
```

There are two options for pushing with the branch. First,

```
1 # git push origin test
```

Will push the test branch to the personal repository. To delete the remote branch

```
1 # git push origin :test
```

You'll still need to delete your local branch if you are done with it.

If you are going to work exclusively with one branch for a while, you might want to configure git to automatically push that branch when you use `git push`. By default `git push` will use the local master branch, but you can use the `test` branch as the source of your changes:

```
1 # git config remote.origin.push test:master
```

Now `git push` will merge into your master branch on your personal repository. You can also setup a remote branch:

```
1 # git config remote.origin.push test:test
```

You can see what branch is configured for pushing with

```
1 # git remote show origin
```

And reset to the default

```
1 # git config remote.origin.push master
```

6.2.5 Pull a Developer's Repo

The procedures for creating personal repositories ensure that every developer can post branches that anyone else can review. To pull a developer's personal repository into your local RTEMS git clone, just add a new remote repo:

```
1 # git remote add devname git://dispatch.rtems.org/devname/rtems.git
2 # git fetch devname
3 # git remote show devname
4 # git branch -a
```

Replace devname with the developer's user name on git, which you can see by accessing <https://git.rtems.org>. Now you can switch to the branches for this developer.

Use a tracking branch if the developer's branch is changing:

```
1 # git branch --track new_feature devname/new_feature
```

6.2.6 Committing

6.2.6.1 Ticket Updates

Our trac instance supports updating a related ticket with the commit message.

Any references to a ticket for example #1234 will insert the message into the ticket as an 'update'. No command is required.

Closing a ticket can be done by prefixing the ticket number with any of the following commands:

close, closed, closes, fix, fixed, or fixes

For example:

```
closes #1234
```

This is a random update it closes #1234 and updates #5678

6.2.6.2 Commands

When merging someone's work, whether your own or otherwise, we have some suggested procedures to follow.

- Never work in the master branch. Checkout a new branch and apply patches/commits to it.
- Before pushing upstream: - Update master by fetching from the server - Rebase the working branch against the updated master - Push the working branch to the server master

The basic workflow looks like

```

1 # git checkout -b somebranch upstream/master
2 # patch .. git add/rm/etc
3 # git commit ...
4 # git pull --rebase upstream master
5 # git push upstream somebranch:master

```

If someone pushed since you updated the server rejects your push until you are up to date.

For example a workflow where you will commit a series of patches from `../patches/am/` directory:

```

1 # git checkout -b am
2 # git am ../patches/am*
3 # git pull --rebase upstream master
4 # git push upstream am:master
5 # git checkout master
6 # git pull upstream master
7 # git log
8 # git branch -d am
9 # git push

```

The git log stage will show your newly pushed patches if everything worked properly, and you can delete the am branch created. The git push at the end will push the changes up to your personal repository.

Another way to do this which pushes directly to the upstream is shown here in an example which simply (and quickly) applies a patch to the branch:

```

1 git checkout -b rtems4.10 --track remotes/upstream/4.10
2 cat /tmp/sp.diff | patch
3 vi sparc.t
4 git add sparc.t
5 git commit -m "sparc.t: Correct for V8/V9"
6 git push upstream rtems4.10:4.10
7 git checkout master
8 git log
9 git branch -d rtems4.10

```

6.2.7 Pushing Multiple Commits

A push with more than one commit results in Trac missing them. Please use the following script to push a single commit at a time:

```
1 #!/bin/sh
2 commits=$(git log --format='%h' origin/master..HEAD | tail -r)
3 for c in $commits
4 do
5     cmd=$(echo $c | sed 's%\(.*\)%git push origin \1:master%')
6     echo $cmd
7 $cmd
8 done
```

6.2.8 Oops!

So you pushed something upstream and broke the repository. First things first: stop what you're doing and notify devel@... so that (1) you can get help and (2) no one pulls from the broken repo. For an extended outage also notify users@... Now, breathe easy and let's figure out what happened. One thing that might work is to just **undo the push**. To get an idea of what you did, run `git reflog`, which might be useful for getting assistance in undoing whatever badness was done.

6.3 Coding Standards

TBD - Write introduction, re-order, identify missing content

6.3.1 Coding Conventions

The style of RTEMS is generally consistent in the core areas. This section attempts to capture generally accepted practices. When in doubt, consult the code around you, look in the RTEMS sources in the directories `cpukit/include/rtems/score` and `cpukit/score`, or ask on the [Developers Mailing List](#).

6.3.1.1 Source Documentation

- Use Doxygen according to our *Doxygen Guidelines* (page 146).
- Use the file templates, see *File Templates* (page 150).
- Use `/* */` comments.
- Do not use `//` comments.
- Use comments wisely within function bodies, to explain or draw attention without being verbose.
- Use English prose and strive for good grammar, spelling, and punctuation.
- Use TODO with a comment to indicate code that needs improvement. Make it clear what there is to do. Add a ticket and add a link to it.
- Use XXX or FIXME to indicate an error/bug/broken code. Add a ticket and add a link to it.

6.3.1.2 Licenses

The RTEMS Project has strict requirements on the types of software licenses that apply to software it includes and distributes. Submissions will be summarily rejected that do not follow the correct license or file header requirements.

- Refer to *Licensing Requirements* (page 239) for a discussion of the acceptable licenses and the rationale.
- Refer to *Copyright and License Block* (page 151) for example copyright/license comment blocks for various languages.

6.3.1.3 Language and Compiler

- Use C99.
- Treat warnings as errors: eliminate them.
- Favor C, but when assembly language is required use inline assembly if possible.
- Do not use compiler extensions.

- Use the RTEMS macros defined in `<rtems/score/basedefs.h>` for abstracting compiler-specific features. For using attributes see the [GCC attribute syntax](#). Prefer to place attributes in front of the declarator. Try to be in line with [C++11 attributes](#) and C11 keywords such as `_Noreturn`.
- Use NULL for the null pointer, and prefer to use explicit checks against NULL, e.g.,

```
1 if ( ptr != NULL )
```

instead of

```
1 if ( !ptr )
```

- Use explicit checks for bits in variables.
 - Example 1: Use

```
1 if ( XBITS == (var & XBITS) )
```

to check for a set of defined bits.

- Example 2: Use

```
1 if ( (var & X_FLAGS) != 0 )
```

instead of

```
1 if ( !(var & X_FLAGS) )
```

to check for at least 1 defined bit in a set.

- Use `(void) unused;` to mark unused parameters and set-but-unused variables immediately after being set.
- Do not put function prototypes in C source files, any global functions should have a prototype in a header file and any private function should be declared static.
- Declare global variables in exactly one header file. Define global variables in at most one source file. Include the header file declaring the global variable as the first include file if possible to make sure that the compiler checks the declaration and definition and that the header file is self-contained.
- Do not cast arguments to any `printf()` or `printk()` variant. Use `<inttypes.h>` PRI constants for the types supported there. Use `<rtems/inttypes.h>` for the other POSIX and RTEMS types that have PRI constants defined there. This increases the portability of the `printf()` format.
- Do not use the register keyword. It is deprecated since C++14.

6.3.1.4 Readability

- Understand and follow the [naming rules](#)..
- Use typedef to remove 'struct', but do not use typedef to hide pointers or arrays. * Exception: typedef can be used to simplify function pointer types.
- Do not mix variable declarations and code.
- Declare variables at the start of a block.
- Only use primitive initialization of variables at their declarations. Avoid complex initializations or function calls in variable declarations.
- Do not put unrelated functions or data in a single file.
- Do not declare functions inside functions.
- Avoid deep nesting by using early exits e.g. return, break, continue. * Parameter checking should be done first with early error returns. * Avoid allocation and critical sections until error checking is done. * For error checks that require locking, do the checks early after acquiring locks. * Use of 'goto' requires good reason and justification.
- Test and action should stay close together.
- Avoid complex logic in conditional and loop statements.
- Put conditional and loop statements on the line after the expression.
- Favor inline functions to hide [compile-time feature-conditioned compilation](#)..
- Define non-inline functions in a .c source file.
- Declare all global (non-static) functions in a .h header file.
- Declare and define inline functions in one place. Usually, this is a `*impl.h` header file.
- Declare and define static functions in one place. Usually, this is toward the start of a .c file. Minimize forward declarations of static functions.
- Function declarations should include variable names.
- Avoid excess parentheses. Learn the [operator precedence](#). rules.
- Always use parentheses with sizeof. This is an exception to the rule about excess parentheses.

6.3.1.5 Robustness

- Check all return statuses.
- Validate input parameters.
- Use debug assertions (assert).
- Use const when appropriate for read-only function parameters and compile-time constant values.
- Do not hard code limits such as maximum instances into your code.
- Prefer to use `sizeof(variable)` instead of `sizeof(type)`.
- Favor C automatic variables over global or static variables.

- Use global variables only when necessary and ensure atomicity of operations.
- Do not shadow variables.
- Avoid declaring large buffers or structures on the stack.
- Avoid using zero (0) as a valid value. Memory often defaults to being zero.
- Favor mutual exclusion primitives over disabling preemption.
- Avoid unnecessary dependencies, such as by not calling “printf()” on error paths.
- Avoid inline functions and macros with complicated logic and decision points.
- Prefer inline functions, enum, and const variables instead of CPP macros.
- CPP macros should use a leading underscore for parameter names and **avoid macro pitfalls**..

6.3.1.6 Portability

- Think portable! RTEMS supports a lot of target hardware.
- For integer primitives, prefer to use precise-width integer types from C99 `stdint.h`.
- Write code that is 16-bit, 32-bit, and 64-bit friendly.

6.3.1.7 Maintainability

- Minimize modifications to **third-party code**..
- Keep it simple! Simple code is easier to debug and easier to read than clever code.
- Share code with other architectures, CPUs, and BSPs where possible.
- Do not duplicate standard OS or C Library routines.

6.3.1.8 Performance

- Prefer algorithms with the **lowest order of time and space**. for fast, deterministic execution times with small memory footprints.
- Understand the constraints of **real-time programming**.. Limit execution times in interrupt contexts and critical sections, such as Interrupt and Timer Service Routines (TSRs).
- Functions used only through function pointers should be declared ‘static inline’ (RTEMS_INLINE_ROUTINE)
- Prefer to ++preincrement instead of postincrement++.
- Avoid using floating point except where absolutely necessary.

6.3.1.9 Miscellaneous

- If you need to temporarily change the execution mode of a task/thread, restore it.
- If adding code to “cpukit” be sure the filename is unique since all files under that directory get merged into a single library.

6.3.1.10 Header Files

- Do not add top-level header files. Place the header files in a directory, for example `#include <rtems/*>`, `#include <bsp/*>`, `#include <dev/*>`, etc.
- Use the extension `.h` for C header files.
- Use the extension `.hpp` for C++ header files.
- Use the file template for header files, see *C/C++ Header File Template* (page 152).
- Use separate header files for the API and the implementation.
- Use `foobar.h` for the header file of the foobar module which defines API components.
- Use `foobardata.h` for the header file of the foobar module which defines interfaces used by the application configuration.
- Use `foobarimpl.h` for the header file of the foobar module which defines interfaces, macros, and inline functions used by the implementation.
- Do not place inline functions which are only used in one implementation source file into the implementation header file. Add these inline functions directly to the corresponding source file.
- Document all elements in header files with comments in Doxygen markup, see *Doxygen Guidelines* (page 146).
- Only place header files which should be directly included by the user with an `@file` Doxygen directive into the API documentation group. Place internal API header files with an `@file` Doxygen command into the implementation documentation group even if they define API elements. The API documentation group should only list public header files and no internal header files.

6.3.1.11 Layering

- TBD: add something about the dependencies and header file layering.
- Understand the ‘RTEMS Software Architecture’ <https://devel.rtems.org/wiki/TBR/UserManual/RTEMS_Software_Architecture>’.

6.3.1.12 Exceptions to the Rules

- Minimize reformatting existing code in RTEMS unless the file undergoes substantial non-style changes.
- **Third-party code**. should not be reformatted to fit RTEMS style. Exception: unmaintained third-party code adopted and maintained by RTEMS may be reformatted, subject to the above rules.

6.3.1.13 Tools

Some of the above can be assisted by tool support. Feel free to add more tools, configurations, etc here.

- **Uncrustify**. Configuration for RTEMS: [rtems.uncrustify](#).

6.3.2 Formatting

6.3.2.1 Rules

- Adhere to the *Eighty Character Line Limit* (page 144).
- Use spaces instead of tabs.
- Use two spaces for one indentation level.
- Put function return types and names on one line if they fit.
- Put function calls on one line if they fit.
- No space between a function name or function-like macro and the opening parenthesis.
- Put braces on the same line as and one space after the conditional expression ends.
- Put the opening brace of a function definition one line after the closing parenthesis of its prototype.
- Put a single space inside and outside of each parenthesis of a conditional expression. Exception: never put a space before a closing semi-colon.
- Put a single space before and after ternary operators.
- Put a single space before and after binary operators.
- Put no space between unary operators (e.g. *, &, !, ~, ++, --) and their operands.
- No spaces around dereferencing operators (-> and .).
- Do not use more than one blank line in a row.
- Do not use trailing white space at the end of a line.

6.3.2.2 Eighty Character Line Limit

Code should look good for everyone under some standard width assumptions. Where a line wraps should be the same for anyone reading the code. For historical reasons, RTEMS uses 80 characters as the maximum width for each line of code. The newline (`\n`) character terminating the line does not count for the 80 character limit.

If you find yourself with code longer than 80 characters, first ask yourself whether the nesting level is too deep, names too long, compound expressions too complicated, or if some other guideline for improving readability can help to shrink the line length. Refactoring nested blocks into functions can help to alleviate code width problems while improving code readability. Making names descriptive yet terse can also improve readability. If absolutely necessary to have a long line, follow the rules on this page to break the line up to adhere to the 80 characters per line rule.

6.3.2.3 Breaking Long Lines

The `if`, `while`, and `for` control statements have their condition expressions aligned and broken on separate lines. When the conditions have to be broken, none go on the first line with the `if`, `while`, or `for` statement, and none go on the last line with the closing parenthesis and the curly brace. Long statements are broken up and indented at operators, with an operator always being the last token on a line. No blank spaces should be left at the end of any line. The continuation of a broken line is indented by one level. Here is an example with a `for` loop.

```

1 for ( initialization = statement; a + really + longish + statement + that + evaluates +
  ↪to < a + boolean; another + statement ) {
2     some_variable = a + really + longish + statement + that + needs + two + lines + gets +
  ↪indented + four + more + spaces + on + the + second + and + subsequent + lines + and +
  ↪broken + up + at + operators;
3 }

```

Should be replaced with

```

1 for (
2     initialization = statement;
3     a + really + longish + statement + that + evaluates + to <
4     a + boolean;
5     another + statement
6 ) {
7     some_variable = a + really + longish + statement + that + needs +
8     two + lines + gets + indented + four + more +
9     spaces + on + the + second + and + subsequent +
10    lines + and + broken + up + at + operators;
11 }

```

Similarly,

```

1 if ( this + that < those && this + these < that && this + those < these && this < those &&
  ↪those < that ) {

```

should be broken up like

```

1 if (
2     this + that < those &&

```

(continues on next page)

(continued from previous page)

```

3  this + these < that &&
4  this + those < these &&
5  this < those &&
6  those < that
7 ) {

```

Note that each expression that resolves to a boolean goes on its own line. Where you place the boolean operator is a matter of choice.

When a line is long because of a comment at the end, move the comment to just before the line, for example

```

1 #define A_LONG_MACRO_NAME (AND + EXPANSION) /* Plus + a + really + long + comment */

```

can be replaced with

```

1 /* Plus + a + really + long + comment */
2 #define A_LONG_MACRO_NAME (AND + EXPANSION)

```

C Preprocessor macros need to be broken up with some care, because the preprocessor does not understand that it should eat newline characters. So

```

1 #define A_LONG_MACRO_NAME (AND + EXCESSIVELY + LONG + EXPANSION + WITH + LOTS + OF + \
↪EXTRA + STUFF + DEFINED)

```

would become

```

1 #define A_LONG_MACRO_NAME ( \
2   AND + EXCESSIVELY + LONG + EXPANSION + WITH + LOTS + OF + EXTRA + STUFF + \
3   DEFINED \
4 )

```

Notice that each line is terminated by a backslash. The backslash tells the preprocessor to eat the newline. Of course, if you have such a long macro, you should consider not using a macro.

Function declarations can be broken up at each argument, for example

```

1 int this_is_a_function( int arg1, int arg2, int arg3, int arg4, int arg5, int arg6, int_
↪arg7, int arg8, int arg9 );

```

would be broken up as

```

1 int this_is_a_function(
2   int arg1,
3   int arg2,
4   int arg3,
5   int arg4,
6   int arg5,
7   int arg6,
8   int arg7,
9   int arg8,
10  int arg9
11 );

```

Excessively long comments should be broken up at a word boundary or somewhere that makes sense, for example

```

1 /* Excessively long comments should be broken up at a word boundary or somewhere that
   ↪makes sense, for example */

```

would be

```

1 /*
2  * Excessively long comments should be broken up at a word boundary or
3  * somewhere that makes sense, for example.
4  */

```

Note that multiline comments have a single asterisk aligned with the asterisk in the opening `/*`. The closing `*/` should appear on a line by itself at the end.

6.3.3 Deprectating Interfaces

TBD - Convert the following to Rest and insert into this file TBD - <https://devel.rtems.org/wiki/Developer/Coding/Deprecating>

6.3.4 Doxygen Guidelines

6.3.4.1 Group Names

Doxygen group names shall use **CamelCase**. In the RTEMS source code, CamelCase is rarely used, so this makes it easier to search and replace Doxygen groups. It avoids ambiguous references to functions, types, defines, macros, and groups. All groups shall have an RTEMS prefix. This makes it possible to include the RTEMS files with Doxygen comments in a larger project without name conflicts.

```

1 /**
2  * @defgroup RTEMSScoreThread
3  *
4  * @ingroup RTEMSScore
5  *
6  * ...
7  */

```

6.3.4.2 Use Groups

Every file, function declaration, type definition, typedef, define, macro and global variable declaration shall belong to at least one Doxygen group. Use `@defgroup` and `@addtogroup` with `@{` and `@}` brackets to add members to a group. A group shall be defined at most once. Each group shall be documented with an `@brief` description and an optional detailed description. The `@brief` description shall use **Title Case**. Use grammatically correct sentences for the detailed descriptions.

```

1 /**
2  * @defgroup RTEMSScoreThread
3  *
4  * @ingroup RTEMSScore
5  *

```

(continues on next page)

(continued from previous page)

```

6  * @brief Thread Handler
7  *
8  * ...
9  *
10 * @{
11 */
12
13 ... declarations, defines ...
14
15 /** @} */

```

```

1 /**
2  * @addtogroup RTEMScoreThread
3  *
4  * @{
5  */
6
7  ... declarations, defines ...
8
9  /** @} */

```

6.3.4.3 Files

Each source or header file shall have an @file block at the top of the file. The @file block should precede the license header separated by one blank line. This placement reduces the chance of merge conflicts in imported third-party code. The @file block shall be put into a group with @ingroup GroupName. The @file block should have an @brief description and a detailed description if it is considered helpful. Use @brief @copybrief GroupName as a default to copy the @brief description from the corresponding group and omit the detailed description.

```

1 /**
2  * @file
3  *
4  * @ingroup RTEMScoreThread
5  *
6  * @brief @copybrief RTEMScoreThread
7  */

```

```

1 /**
2  * @file
3  *
4  * @ingroup RTEMScoreThread
5  *
6  * @brief Some helpful brief description.
7  *
8  * Some helpful detailed description.
9  */

```

6.3.4.4 Type Definitions

Each type defined in a header file shall be documented with an @brief description and an optional detailed description. Each type member shall be documented with an @brief description and an optional detailed description. Use grammatically correct sentences for the detailed descriptions.

```

1 /**
2  * @brief The information structure used to manage each API class of objects.
3  *
4  * If objects for the API class are configured, an instance of this structure
5  * is statically allocated and pre-initialized by OBJECTS_INFORMATION_DEFINE()
6  * through <rtems/confdefs.h>. The RTEMS library contains a statically
7  * allocated and pre-initialized instance for each API class providing zero
8  * objects, see OBJECTS_INFORMATION_DEFINE_ZERO().
9  */
10 typedef struct {
11     /**
12      * @brief This is the maximum valid ID of this object API class.
13      *
14      * This member is statically initialized and provides also the object API,
15      * class and multiprocessing node information.
16      *
17      * It is used by _Objects_Get() to validate an object ID.
18      */
19     Objects_Id maximum_id;
20
21     ... more members ...
22 } Objects_Information;

```

6.3.4.5 Function Declarations

Each function declaration or function-like macros in a header file shall be documented with an @brief description and an optional detailed description. Use grammatically correct sentences for the brief and detailed descriptions. Each parameter shall be documented with an @param entry. List the @param entries in the order of the function parameters. For *non-const pointer* parameters

- use @param[out], if the referenced object is modified by the function, or
- use @param[in, out], if the referenced object is read and modified by the function.

For other parameters (e.g. *const pointer* and *scalar* parameters) do not use the [in], [out] or [in, out] parameter specifiers. Each return value or return value range shall be documented with an @retval entry. Document the most common return value first. Use a placeholder name for value ranges, e.g. pointer in the _Workspace_Allocate() example below. In case the function returns only one value, then use @return, e.g. use @retval only if there are at least two return values or return value ranges. Use grammatically correct sentences for the parameter and return value descriptions.

```

1 /**
2  * @brief Sends a message to the message queue.
3  *
4  * This directive sends the message buffer to the message queue indicated by

```

(continues on next page)

(continued from previous page)

```

5  * ID. If one or more tasks is blocked waiting to receive a message from this
6  * message queue, then one will receive the message. The task selected to
7  * receive the message is based on the task queue discipline algorithm in use
8  * by this particular message queue. If no tasks are waiting, then the message
9  * buffer will be placed at the rear of the chain of pending messages for this
10 * message queue.
11 *
12 * @param id The message queue ID.
13 * @param buffer The message content buffer.
14 * @param size The size of the message.
15 *
16 * @retval RTEMS_SUCCESSFUL Successful operation.
17 * @retval RTEMS_INVALID_ID Invalid message queue ID.
18 * @retval RTEMS_INVALID_ADDRESS The message buffer pointer is @c NULL.
19 * @retval RTEMS_INVALID_SIZE The message size is larger than the maximum
20 * message size of the message queue.
21 * @retval RTEMS_TOO_MANY The new message would exceed the message queue limit
22 * for pending messages.
23 */
24 rtems_status_code rtems_message_queue_send(
25     rtems_id    id,
26     const void *buffer,
27     size_t      size
28 );

```

```

1  /**
2  * @brief Receives a message from the message queue
3  *
4  * This directive is invoked when the calling task wishes to receive a message
5  * from the message queue indicated by ID. The received message is to be placed
6  * in the buffer. If no messages are outstanding and the option set indicates
7  * that the task is willing to block, then the task will be blocked until a
8  * message arrives or until, optionally, timeout clock ticks have passed.
9  *
10 * @param id The message queue ID.
11 * @param[out] buffer The buffer for the message content. The buffer must be
12 * large enough to store maximum size messages of this message queue.
13 * @param[out] size The size of the message.
14 * @param option_set The option set, e.g. RTEMS_NO_WAIT or RTEMS_WAIT.
15 * @param timeout The number of ticks to wait if the RTEMS_WAIT is set. Use
16 * RTEMS_NO_TIMEOUT to wait indefinitely.
17 *
18 * @retval RTEMS_SUCCESSFUL Successful operation.
19 * @retval RTEMS_INVALID_ID Invalid message queue ID.
20 * @retval RTEMS_INVALID_ADDRESS The message buffer pointer or the message size
21 * pointer is @c NULL.
22 * @retval RTEMS_TIMEOUT A timeout occurred and no message was received.
23 */
24 rtems_status_code rtems_message_queue_receive(
25     rtems_id    id,
26     void        *buffer,
27     size_t      *size,
28     rtems_option option_set,
29     rtems_interval timeout
30 );

```

```

1 /**
2  * @brief Allocates a memory block of the specified size from the workspace.
3  *
4  * @param size The size of the memory block.
5  *
6  * @retval pointer The pointer to the memory block. The pointer is at least
7  *   aligned by CPU_HEAP_ALIGNMENT.
8  * @retval NULL No memory block with the requested size is available in the
9  *   workspace.
10 */
11 void *_Workspace_Allocate( size_t size );

```

```

1 /**
2  * @brief Rebalances the red-black tree after insertion of the node.
3  *
4  * @param[in, out] the_rbtrees The red-black tree control.
5  * @param[in, out] the_node The most recently inserted node.
6  */
7 void _RBTrees_Insert_color(
8   RBTrees_Control *the_rbtrees,
9   RBTrees_Node    *the_node
10 );

```

```

1 /**
2  * @brief Builds an object ID from its components.
3  *
4  * @param the_api The object API.
5  * @param the_class The object API class.
6  * @param node The object node.
7  * @param index The object index.
8  *
9  * @return Returns the object ID constructed from the arguments.
10 */
11 #define _Objects_Build_id( the_api, the_class, node, index )

```

6.3.4.6 Header File Examples

The `<rtems/score/thread.h>` and `<rtems/score/threadimpl.h>` header files are a good example of how header files should be documented.

6.3.5 File Templates

Every source code file shall have a copyright and license block. Corresponding to the license, every file shall have an **SPDX License Identifier** in the first possible line of the file. C/C++ files should have a Doxygen file comment block.

The preferred license for source code is **BSD-2-Clause**. The preferred license for documentation is **CC-BY-SA-4.0**.

6.3.5.1 Copyright and License Block

You are the copyright holder. Use the following copyright and license block for your source code contributions to the RTEMS Project. Place it after the SPDX License Identifier line and the optional file documentation block. Replace the <FIRST YEAR> placeholder with the year of your first substantial contribution to this file. Update the <LAST YEAR> with the year of your last substantial contribution to this file. If the first and last years are the same, then remove the <LAST YEAR> placeholder with the comma. Replace the <COPYRIGHT HOLDER> placeholder with your name.

In case you are a real person, then use the following format for <COPYRIGHT HOLDER>: <FIRST NAME> <MIDDLE NAMES> <LAST NAME>. The <FIRST NAME> is your first name (also known as given name), the <MIDDLE NAMES> are your optional middle names, the <LAST NAME> is your last name (also known as family name).

If more than one copyright holder exists for a file, then sort the copyright lines by the first year (earlier years are below later years) followed by the copyright holder in alphabetical order (A is above B in the file).

Use the following template for a copyright and license block. Do not change the license text.

```

1 Copyright (C) <FIRST YEAR>, <LAST YEAR> <COPYRIGHT HOLDER>
2
3 Redistribution and use in source and binary forms, with or without
4 modification, are permitted provided that the following conditions
5 are met:
6 1. Redistributions of source code must retain the above copyright
7   notice, this list of conditions and the following disclaimer.
8 2. Redistributions in binary form must reproduce the above copyright
9   notice, this list of conditions and the following disclaimer in the
10  documentation and/or other materials provided with the distribution.
11
12 THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
13 AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
14 IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
15 ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
16 LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
17 CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
18 SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
19 INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
20 CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
21 ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
22 POSSIBILITY OF SUCH DAMAGE.
```

Check the top-level COPYING file of the repository. If you are a new copyright holder, then add yourself to the top of the list. If your last year of a substantial contribution changed, then please update your copyright line.

6.3.5.2 C/C++ Header File Template

Use the following guidelines and template for C and C++ header files (here <foo/bar/baz.h>):

- Place the SPDX License Identifier in the first line of the file.
- Add a Doxygen file documentation block.
- Place the copyright and license comment block after the documentation block.
- For the <FIRST YEAR>, <LAST YEAR>, and <COPYRIGHT HOLDER> placeholders see *Copyright and License Block* (page 151).
- Separate comment blocks by exactly one blank line.
- Separate the Doxygen comment block from the copyright and license, so that the copyright and license information is not included in the Doxygen output.
- For C++ header files discard the extern "C" start and end sections.

```

1 /* SPDX-License-Identifier: BSD-2-Clause */
2
3 /**
4  * @file
5  *
6  * @ingroup TheGroupForThisFile
7  *
8  * @brief Short "Table of Contents" Description of File Contents
9  *
10 * A short description of the purpose of this file.
11 */
12
13 /*
14 * Copyright (C) <FIRST YEAR>, <LAST YEAR> <COPYRIGHT HOLDER>
15 *
16 * Redistribution and use in source and binary forms, with or without
17 * modification, are permitted provided that the following conditions
18 * are met:
19 * 1. Redistributions of source code must retain the above copyright
20 *    notice, this list of conditions and the following disclaimer.
21 * 2. Redistributions in binary form must reproduce the above copyright
22 *    notice, this list of conditions and the following disclaimer in the
23 *    documentation and/or other materials provided with the distribution.
24 *
25 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
26 * AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
27 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
28 * ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
29 * LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
30 * CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
31 * SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
32 * INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
33 * CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
34 * ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
35 * POSSIBILITY OF SUCH DAMAGE.
36 */
37
38 #ifndef _FOO_BAR_BAZ_H
39 #define _FOO_BAR_BAZ_H

```

(continues on next page)

(continued from previous page)

```

40
41 #include <foo/bar/xyz.h>
42
43 /* Remove for C++ code */
44 #ifdef __cplusplus
45 extern "C" {
46 #endif
47
48 /* Declarations, defines, macros, inline functions, etc. */
49
50 /* Remove for C++ code */
51 #ifdef __cplusplus
52 }
53 #endif
54
55 #endif /* _FOO_BAR_BAZ_H */

```

6.3.5.3 C/C++/Assembler Source File Template

Use the following template for C, C++, and assembler source files (here implementation of <foo/bar/baz.h>). For the <FIRST YEAR>, <LAST YEAR>, and <COPYRIGHT HOLDER> placeholders see *Copyright and License Block* (page 151).

```

1 /* SPDX-License-Identifier: BSD-2-Clause */
2
3 /**
4  * @file
5  *
6  * @ingroup TheGroupForThisFile
7  *
8  * @brief Short "Table of Contents" Description of File Contents
9  *
10 * A short description of the purpose of this file.
11 */
12
13 /*
14 * Copyright (C) <FIRST YEAR>, <LAST YEAR> <COPYRIGHT HOLDER>
15 *
16 * Redistribution and use in source and binary forms, with or without
17 * modification, are permitted provided that the following conditions
18 * are met:
19 * 1. Redistributions of source code must retain the above copyright
20 *    notice, this list of conditions and the following disclaimer.
21 * 2. Redistributions in binary form must reproduce the above copyright
22 *    notice, this list of conditions and the following disclaimer in the
23 *    documentation and/or other materials provided with the distribution.
24 *
25 * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
26 * AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
27 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
28 * ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
29 * LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
30 * CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
31 * SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS

```

(continues on next page)

(continued from previous page)

```

32 * INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
33 * CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
34 * ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
35 * POSSIBILITY OF SUCH DAMAGE.
36 */
37
38 #ifdef HAVE_CONFIG_H
39 #include "config.h"
40 #endif
41
42 #include <foo/bar/baz.h>
43
44 /* Definitions, etc. */

```

6.3.5.4 Python File Template

Use the following template for Python source files. For the <FIRST YEAR>, <LAST YEAR>, and <COPYRIGHT HOLDER> placeholders see *Copyright and License Block* (page 151).

The File documentation block is a [Python docstring \(PEP 257\)](#) module documentation block. RTEMS uses `"""` for Python docstrings.

```

1 # SPDX-License-Identifier: BSD-2-Clause
2 """File documentation block"""
3
4 # Copyright (C) <FIRST YEAR>, <LAST YEAR> <COPYRIGHT HOLDER>
5 #
6 # Redistribution and use in source and binary forms, with or without
7 # modification, are permitted provided that the following conditions
8 # are met:
9 # 1. Redistributions of source code must retain the above copyright
10 #    notice, this list of conditions and the following disclaimer.
11 # 2. Redistributions in binary form must reproduce the above copyright
12 #    notice, this list of conditions and the following disclaimer in the
13 #    documentation and/or other materials provided with the distribution.
14 #
15 # THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
16 # AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
17 # IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
18 # ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
19 # LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
20 # CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
21 # SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
22 # INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
23 # CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
24 # ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
25 # POSSIBILITY OF SUCH DAMAGE.

```

If the Python source file is a command line command add the following as the first line of the file:

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

A command line Python module does not need to have the `.py` file extension.

Only specify python as the command to env. A system that does not provide the python command can install a virtual python environment or the user can prepend the specific Python versioned command to the Python script on the command line when invoking the command.

6.3.5.5 Shell Scripts

Use the following template for shell script source files and other files which accept a #-style comment block. For the <FIRST YEAR>, <LAST YEAR>, and <COPYRIGHT HOLDER> placeholders see *Copyright and License Block* (page 151).

```

1 #!/bin/sh
2 # SPDX-License-Identifier: BSD-2-Clause
3
4 # File documentation block
5
6 # Copyright (C) <FIRST YEAR>, <LAST YEAR> <COPYRIGHT HOLDER>
7 #
8 # Redistribution and use in source and binary forms, with or without
9 # modification, are permitted provided that the following conditions
10 # are met:
11 # 1. Redistributions of source code must retain the above copyright
12 #    notice, this list of conditions and the following disclaimer.
13 # 2. Redistributions in binary form must reproduce the above copyright
14 #    notice, this list of conditions and the following disclaimer in the
15 #    documentation and/or other materials provided with the distribution.
16 #
17 # THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS"
18 # AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
19 # IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
20 # ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
21 # LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
22 # CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
23 # SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
24 # INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
25 # CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
26 # ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
27 # POSSIBILITY OF SUCH DAMAGE.

```

6.3.5.6 reStructuredText File Template

Use the following template for reStructuredText (reST) source files. For the <FIRST YEAR>, <LAST YEAR>, and <COPYRIGHT HOLDER> placeholders see *Copyright and License Block* (page 151).

```

1 .. SPDX-License-Identifier: CC-BY-SA-4.0
2
3 .. Copyright (C) <FIRST YEAR>, <LAST YEAR> <COPYRIGHT HOLDER>

```

6.3.6 Naming Rules

6.3.6.1 General Rules

- Avoid abbreviations.
 - Exception: when the abbreviation is more common than the full word.
 - Exception: For well-known acronyms.
- Use descriptive language.
- File names should be lower-case alphabet letters only, plus the extension. Avoid symbols in file names.
- Prefer to use underscores to separate words, rather than **CamelCase** or **!TitleCase**.
- Local-scope variable names are all lower case with underscores between words.
- CPP macros are all capital letters with underscores between words.
- Enumerated (enum) values are all capital letters with underscores between words, but the type name follows the regular rules of other type names.
- Constant (const) variables follow the same rules as other variables. An exception is that a const that replaces a CPP macro might be all capital letters for backward compatibility.
- Type names, function names, and global scope names have different rules depending on whether they are part of the public API or are internal to RTEMS, see below.

User-Facing APIs

The public API routines follow a standard API like POSIX or BSD or start with *rtems_*. If a name starts with *rtems_*, then it should be assumed to be available for use by the application and be documented in the User's Guide.

If the method is intended to be private, then make it static to a file or start the name with a leading *_*.

Classic API

- Public facing APIs start with *rtems_* followed by a word or phrase to indicate the Manager or functional category the method or data type belongs to.
- Non-public APIs should be static or begin with a leading *_*. The required form is the use of a leading underscore, functional area with leading capital letter, an underscore, and the method with a leading capital letter.

POSIX API

- Follow the rules of POSIX.

RTEMS Internal Interfaces

Super Core

The **Super Core** is organized in an Object-Oriented fashion. Each score Handler is a Package, or Module, and each Module contains type definitions, functions, etc. The following summarizes our conventions for using names within **SuperCore** Modules.

- Use “Module_name_Particular_type_name” for type names.
- Use “_Module_name_Particular_function_name” for functions names.

- Use “_Module_name_Global_or_file_scope_variable_name” for global or file scope variable names.

Within a structure:

- Use “Name” for struct aggregate members.
- Use “name” for reference members.
- Use “name” for primitive type members.

As shown in the following example:

```
1 typedef struct {
2     Other_module_Struct_type    Aggregate_member_name;
3     Other_module_Struct_type    *reference_member_name;
4     Other_module_Primitive_type primitive_member_name;
5 } The_module_Type_name;
```

BSP

- TODO.

6.4 Documentation Guidelines

6.4.1 Application Configuration Options

Add at least an index entry and a label for the configuration option. Use a pattern of `CONFIGURE_[A-Z0-9_]+` for the option name. Use the following template for application configuration feature options:

```

1 .. index:: CONFIGURE_FEATURE
2
3 .. _CONFIGURE_FEATURE:
4
5 CONFIGURE_FEATURE
6 -----
7
8 CONSTANT:
9     ``CONFIGURE_FEATURE``
10
11 OPTION TYPE:
12     This configuration option is a boolean feature define.
13
14 DEFAULT CONFIGURATION:
15     If this configuration option is undefined, then the described feature is not
16     enabled.
17
18 DESCRIPTION:
19     In case this configuration option is defined, then feature happens.
20
21 NOTES:
22     Keep the description short. Add all special cases, usage notes, error
23     conditions, configuration dependencies, references, etc. here to the notes.

```

Use the following template for application configuration integer and initializer options:

```

1 .. index:: CONFIGURE_VALUE
2
3 .. _CONFIGURE_VALUE:
4
5 CONFIGURE_VALUE
6 -----
7
8 CONSTANT:
9     ``CONFIGURE_VALUE``
10
11 OPTION TYPE:
12     This configuration option is an integer (or initializer) define.
13
14 DEFAULT VALUE:
15     The default value is X.
16
17 VALUE CONSTRAINTS:
18     The value of this configuration option shall satisfy all of the following
19     constraints:
20
21     * It shall be greater than or equal to A.
22

```

(continues on next page)

(continued from previous page)

```
23 * It shall be less than or equal to B.  
24  
25 * It shall meet C.  
26  
27 DESCRIPTION:  
28     The value of this configuration option defines the Y of Z in W.  
29  
30 NOTES:  
31     Keep the description short. Add all special cases, usage notes, error  
32     conditions, configuration dependencies, references, etc. here to the notes.
```

6.5 Python Development Guidelines

Python is the preferred programming language for the RTEMS Tools. The RTEMS Tools run on the host computer of an RTEMS user or maintainer. These guidelines cover the Python language version, the source code formatting, use of static analysis tools, type annotations, testing, code coverage, and documentation. There are exceptions for existing code and third-party code. It is recommended to read the [PEP 8 - Style Guide for Python Code](#) and the [Google Python Style Guide](#).

6.5.1 Python Language Versions

Although the official end-of-life of Python 2.7 was on January 1, 2020, the RTEMS Project still cares about Python 2.7 compatibility for some tools. Every tool provided by the RTEMS Project which an RTEMS user may use to develop applications with RTEMS should be Python 2.7 compatible. Examples are the build system, the RTEMS Source Builder, and the RTEMS Tester. The rationale is that there are still some maintained Linux distributions in the wild which ship only Python 2.7 by default. An example is CentOS 7 which gets maintenance updates until June 2024. Everything an RTEMS maintainer uses should be written in Python 3.6.

6.5.2 Python Code Formatting

Good looking code is important. Unfortunately, what looks good is a bit subjective and varies from developer to developer. Arguing about the code format is not productive. Code reviews should focus on more important topics, for example functionality, testability, and performance. Fortunately, for Python there are some good automatic code formatters available. All new code specifically developed for the RTEMS Tools should be piped through the [yapf](#) Python code formatter before it is committed or sent for review. Use the default settings of the tool ([PEP 8](#) coding style).

You can disable the automatic formatting by the tool in a region starting with the `# yapf: disable` comment until the next `# yapf: enable` comment, for example

```
1 # yapf: disable
2 FOO = {
3     # ... some very large, complex data literal.
4 }
5
6 BAR = [
7     # ... another large data literal.
8 ]
9 # yapf: enable
```

For a single literal, you can disable the formatting like this:

```
1 BAZ = {
2     (1, 2, 3, 4),
3     (5, 6, 7, 8),
4     (9, 10, 11, 12),
5 } # yapf: disable
```

6.5.3 Static Analysis Tools

Use the `flake8` and `pylint` static analysis tools for Python. Do not commit your code or send it for review if the tools find some rule violations. Run the tools with the default configuration. If you have problems to silence the tools, then please ask for help on the [Developers Mailing List](#). Consult the tool documentation to silence false positives.

6.5.4 Type Annotations

For Python 3.6 or later code use type annotations. All public functions of your modules should have [PEP 484](#) type annotations. Check for type issues with the `mypy` static type checker.

6.5.5 Testing

Write tests for your code with the `pytest` framework. Use the `monkeypatch` mocking module. Do not use the standard Python `unittest` and `unittest.mock` modules. Use `coverage run -m pytest` to run the tests with code coverage support. If you modify existing code or contribute new code to a subproject which uses tests and the code coverage metric, then do not make the code coverage worse.

6.5.5.1 Test Organization

Do not use test classes to group tests. Use separate files instead. Avoid deep test directory hierarchies. For example, place tests for `mymodule.py` in `tests/test_mymodule.py`. For class-specific tests use:

- `mymodule.py:class First` → `tests/test_mymodule_first.py`
- `mymodule.py:class Second` → `tests/test_mymodule_second.py`
- `mymodule.py:class Third` → `tests/test_mymodule_third.py`

You can also group tests in other ways, for example:

- `mymodule.py` → `tests/test_mymodule_input.py`
- `mymodule.py` → `tests/test_mymodule_output.py`

6.5.6 Documentation

Document your code using the [PEP 257 - Docstring Conventions](#). Contrary to PEP 257, use the descriptive-style (`"""Fetches rows from a Bigtable."""`) instead of imperative-style (`"""Fetch rows from a Bigtable."""`) as recommended by [Comments and Docstrings - Functions and Methods](#). Use the `Sphinx` format. The `sphinx-autodoc-typehints` helps to reuse the type annotations for the documentation. Test code does not need docstrings in general.

6.5.7 Existing Code

Existing code in the RTEMS Tools may not follow the preceding guidelines. The RTEMS Project welcomes contributions which bring existing code in line with the guidelines. Firstly, run the yapf code formatter through the existing code of interest. Add `# yapf: disable` comments to avoid reformatting in some areas if it makes sense. If the existing code has no unit tests, then add unit tests before you modify existing code by hand. With the new unit tests aim at a good code coverage especially in the areas you intend to modify. While you review the code add docstrings. Run the static analysers and fix the rule violations. Please keep in mind that also trivial modifications can break working code. Make sure you have some unit tests. Add type annotations unless the code should be Python 2.7 compatible. Concentrate on the public interfaces.

6.5.8 Third-Party Code

Try to not modify imported third-party code. In case there are issues with third-party code, then at least write a bug report or otherwise contact the upstream project. Reimport the third-party code after the issue is fixed in the upstream project. Only temporarily modify imported third-party code until a solution integrated in the upstream is available.

6.6 Change Management

Major decisions about RTEMS are made by the core developers in concert with the user community, guided by the Mission Statement. We provide access to our development sources via a Git Repository (see these Instructions for details).

TBD - ??? what in the Wiki could go here

6.7 Issue Tracking

The RTEMS Project uses Trac to manage all change requests and problem reports and refers to either as a ticket.

The bug reporting procedure is documented in the [RTEMS User Manual](#).

TBD Review process, workflows, etc.

SOFTWARE TEST PLAN ASSURANCE AND PROCEDURES

7.1 Testing and Coverage

Testing to verify that requirements are implemented is a critical part of the high integrity processes. Similarly, measuring and reporting source and decision path coverage of source code is critical.

Needed improvements to the RTEMS testing infrastructure should be done as part of the open project. Similarly, improvements in RTEMS coverage reporting should be done as part of the open project. Both of these capabilities are part of the RTEMS Tester toolset.

Assuming that a requirements focused test suite is added to the open RTEMS, tools will be needed to assist in verifying that requirements are “fully tested.” A fully tested requirement is one which is implemented and tested with associated logical tracing. Tools automating this analysis and generating reporting and alerts will be a critical part of ensuring the master technical data does not bit rot.

Must use tools from:

TBD - Change URL to [git.rtems.org](https://github.com/RTEMS/rtems-tools) and list support tools RTEMS Tools Project: <https://github.com/RTEMS/rtems-tools>

Scope, Procedures, Methodologies, Tools TBD - Write content

7.1.1 Test Suites

All RTEMS source distributions include the complete RTEMS test suites. These tests must be compiled and linked for a specific BSP. Some BSPs are for freely available simulators and thus anyone may test RTEMS on a simulator. Most of the BSPs which can execute on a simulator include scripts to help automate running them.

The RTEMS Project welcomes additions to the various test suites and sample application collections. This helps improve coverage of functionality as well as ensure user use cases are regularly tested.

The following functional test suites are included with RTEMS.

- Classic API Single Processor Test Suite
- POSIX API Test Suite
- File System Test Suite
- Support Library Test Suite (libtests)
- Symmetric Multiprocessing Test Suite
- Distributed Multiprocessing Test Suite
- Classic API Ada95 Binding Test Suite

The following timing test suites are included with RTEMS.

- Classic API Timing Test Suite
- POSIX API Timing Test Suite
- Rhealstone Collection
- Benchmarks Collection

The RTEMS source distribution includes two collections of sample applications.

- Sample Applications (built as RTEMS tests)
- Example Applications (built as RTEMS user applications)

The RTEMS libbsd package includes its own test suite.

7.1.1.1 Legacy Test Suites

The following are available for the legacy IPV4 Network Stack:

- Network Demonstration Applications

Post RTEMS 4.10, ITRON API support was removed. The following test suites are only available if the ITRON API support is present in RTEMS.

- ITRON API Test Suite
- ITRON API Timing Test Suite

7.1.2 RTEMS Tester

TBD - Convert the following to Rest and insert into this file TBD <https://devel.rtems.org/wiki/Testing/Tester> TBD - Above file is horribly out of date. Find new docs and refer to them

SOFTWARE TEST FRAMEWORK

8.1 The RTEMS Test Framework

The *RTEMS Test Framework* helps you to write test suites. It has the following features:

- Implemented in standard C11
- Tests can be written in C or C++
- Runs on at least FreeBSD, MSYS2, Linux and RTEMS
- Test runner and test case code can be in separate translation units
- Test cases are automatically registered at link-time
- Test cases may have a test fixture
- Test checks for various standard types
- Supports test case planning
- Test case scoped dynamic memory
- Test case destructors
- Test case resource accounting to show that no resources are leaked during the test case execution
- Supports early test case exit, e.g. in case a `malloc()` fails
- Individual test case and overall test suite duration is reported
- Procedures for code runtime measurements in RTEMS
- Easy to parse test report to generate for example human readable test reports
- Low overhead time measurement of short time sequences (using cycle counter hardware if available)
- Configurable time service provider for a monotonic clock
- Low global memory overhead for test cases and test checks
- Supports multi-threaded execution and interrupts in test cases
- A simple (polled) put character function is sufficient to produce the test report
- Only text, global data and a stack pointer must be set up to run a test suite
- No dynamic memory is used by the framework itself
- No memory is aggregated throughout the test case execution

8.1.1 Nomenclature

A *test suite* is a collection of test cases. A *test case* consists of individual test actions and checks. A *test check* determines if the outcome of a test action meets its expectation. A *test action* is a program sequence with an observable outcome, for example a function invocation with a return status. If a test action produces the expected outcome as determined by the corresponding test check, then this test check passes, otherwise this test check fails. The test check failures of a test case are summed up. A test case passes, if the failure count of this test case is zero, otherwise the test case fails. The test suite passes if all test cases pass, otherwise it fails.

8.1.2 Test Cases

You can write a test case with the `T_TEST_CASE()` macro followed by a function body:

```

1 T_TEST_CASE(name)
2 {
3     /* Your test case code */
4 }
```

The test case *name* must be a valid C designator. The test case names must be unique within the test suite. Just link modules with test cases to the test runner to form a test suite. The test cases are automatically registered via static C constructors.

Listing 8.1: Test Case Example

```

1 #include <t.h>
2
3 static int add(int a, int b)
4 {
5     return a + b;
6 }
7
8 T_TEST_CASE(a_test_case)
9 {
10    int actual_value;
11
12    actual_value = add(1, 1);
13    T_eq_int(actual_value, 2);
14    T_true(false, "a test failure message");
15 }
```

Listing 8.2: Test Case Report

```

1 B:a_test_case
2 P:0:8:UI1:test-simple.c:13
3 F:1:8:UI1:test-simple.c:14:a test failure message
4 E:a_test_case:N:2:F:1:D:0.001657
```

The *B* line indicates the begin of test case *a_test_case*. The *P* line shows that the test check in file *test-simple.c* at line 13 executed by task *UI1* on processor 0 as the test step 0 passed. The invocation of *add()* in line 12 is the test action of test step 0. The *F* lines shows that the test check in file *test-simple.c* at line 14 executed by task *UI1* on processor 0 as the test step 1 failed with a message of “a test failure message”. The *E* line indicates the end of test case *a_test_case* resulting in a total of two test steps (*N*) and one test failure (*F*). The test case execution duration (*D*) was 0.001657 seconds. For test report details see: *Test Reporting* (page 199).

8.1.3 Test Fixture

You can write a test case with a test fixture with the `T_TEST_CASE_FIXTURE()` macro followed by a function body:

```
1 T_TEST_CASE_FIXTURE(name, fixture)
2 {
3     /* Your test case code */
4 }
```

The test case *name* must be a valid C designator. The test case names must be unique within the test suite. The *fixture* must point to a statically initialized read-only object of type `T_fixture`.

```
1 typedef struct T_fixture {
2     void (*setup)(void *context);
3     void (*stop)(void *context);
4     void (*teardown)(void *context);
5     void (*scope)(void *context, char *buffer, size_t size);
6     void *initial_context;
7 } T_fixture;
```

The test fixture provides methods to setup, stop, and teardown a test case as well as the scope for log messages. A context is passed to each of the methods. The initial context is defined by the read-only fixture object. The context can be obtained by the `T_fixture_context()` function. The context can be changed within the scope of one test case by the `T_set_fixture_context()` function. The next test case execution using the same fixture will start again with the initial context defined by the read-only fixture object. Setting the context can be used for example to dynamically allocate a test environment in the setup method.

The test case fixtures of a test case are organized as a stack. Fixtures can be dynamically added to a test case and removed from a test case via the `T_push_fixture()` and `T_pop_fixture()` functions.

```
1 void *T_push_fixture(T_fixture_node *node, const T_fixture *fixture);
2
3 void T_pop_fixture(void);
```

The `T_push_fixture()` function needs an uninitialized fixture node which must exist until `T_pop_fixture()` is called. It returns the initial context of the fixture. At the end of a test case all pushed fixtures are popped automatically. A call of `T_pop_fixture()` invokes the teardown method of the fixture and must correspond to a previous call to `T_push_fixture()`.

Listing 8.3: Test Fixture Example

```
1 #include <t.h>
2
3 static int initial_value = 3;
4
5 static int counter;
6
7 static void
8 setup(void *ctx)
9 {
10     int *c;
11
```

(continues on next page)

(continued from previous page)

```

12 T_log(T_QUIET, "setup begin");
13 T_eq_ptr(ctx, &initial_value);
14 T_eq_ptr(ctx, T_fixture_context());
15 c = ctx;
16 counter = *c;
17 T_set_fixture_context(&counter);
18 T_eq_ptr(&counter, T_fixture_context());
19 T_log(T_QUIET, "setup end");
20 }
21
22 static void
23 stop(void *ctx)
24 {
25     int *c;
26
27     T_log(T_QUIET, "stop begin");
28     T_eq_ptr(ctx, &counter);
29     c = ctx;
30     ++(*c);
31     T_log(T_QUIET, "stop end");
32 }
33
34 static void
35 teardown(void *ctx)
36 {
37     int *c;
38
39     T_log(T_QUIET, "teardown begin");
40     T_eq_ptr(ctx, &counter);
41     c = ctx;
42     T_eq_int(*c, 4);
43     T_log(T_QUIET, "teardown end");
44 }
45
46 static const T_fixture fixture = {
47     .setup = setup,
48     .stop = stop,
49     .teardown = teardown,
50     .initial_context = &initial_value
51 };
52
53 T_TEST_CASE_FIXTURE(fixture, &fixture)
54 {
55     T_assert_true(true, "all right");
56     T_assert_true(false, "test fails and we stop the test case");
57     T_log(T_QUIET, "not reached");
58 }

```

Listing 8.4: Test Fixture Report

```

1 B:fixture
2 L:setup begin
3 P:0:0:UI1:test-fixture.c:13
4 P:1:0:UI1:test-fixture.c:14
5 P:2:0:UI1:test-fixture.c:18

```

(continues on next page)

(continued from previous page)

```

6 L:setup end
7 P:3:0:UI1:test-fixture.c:55
8 F:4:0:UI1:test-fixture.c:56:test fails and we stop the test case
9 L:stop begin
10 P:5:0:UI1:test-fixture.c:28
11 L:stop end
12 L:teardown begin
13 P:6:0:UI1:test-fixture.c:40
14 P:7:0:UI1:test-fixture.c:42
15 L:teardown end
16 E:fixture:N:8:F:1

```

8.1.4 Test Case Planning

A non-quiet test check fetches and increments the test step counter atomically. For each test case execution the planned steps can be specified with the *T_plan()* function.

```
1 void T_plan(unsigned int planned_steps);
```

This function must be invoked at most once in each test case execution. If the planned test steps are set with this function, then the final test steps after the test case execution must be equal to the planned steps, otherwise the test case fails.

Use the *T_step_*(step, ...)* test check variants to ensure that the test case execution follows exactly the planned steps.

Listing 8.5: Test Planning Example

```

1 #include <t.h>
2
3 T_TEST_CASE(wrong_step)
4 {
5     T_plan(2);
6     T_step_true(0, true, "all right");
7     T_step_true(2, true, "wrong step");
8 }
9
10 T_TEST_CASE(plan_ok)
11 {
12     T_plan(1);
13     T_step_true(0, true, "all right");
14 }
15
16 T_TEST_CASE(plan_failed)
17 {
18     T_plan(2);
19     T_step_true(0, true, "not enough steps");
20     T_quiet_true(true, "quiet test do not count");
21 }
22
23 T_TEST_CASE(double_plan)
24 {
25     T_plan(99);
26     T_plan(2);

```

(continues on next page)

(continued from previous page)

```

27 }
28
29 T_TEST_CASE(steps)
30 {
31     T_step(0, "a");
32     T_plan(3);
33     T_step(1, "b");
34     T_step(2, "c");
35 }

```

Listing 8.6: Test Planning Report

```

1 B:wrong_step
2 P:0:0:UI1:test-plan.c:6
3 F:1:0:UI1:test-plan.c:7:planned step (2)
4 E:wrong_step:N:2:F:1
5 B:plan_ok
6 P:0:0:UI1:test-plan.c:13
7 E:plan_ok:N:1:F:0
8 B:plan_failed
9 P:0:0:UI1:test-plan.c:19
10 F:*:0:UI1:*:*:actual steps (1), planned steps (2)
11 E:plan_failed:N:1:F:1
12 B:double_plan
13 F:*:0:UI1:*:*:planned steps (99) already set
14 E:double_plan:N:0:F:1
15 B:steps
16 P:0:0:UI1:test-plan.c:31
17 P:1:0:UI1:test-plan.c:33
18 P:2:0:UI1:test-plan.c:34
19 E:steps:N:3:F:0

```

8.1.5 Test Case Resource Accounting

The framework can check if various resources have leaked during a test case execution. The resource checkers are specified by the test run configuration. On RTEMS, checks for the following resources are available

- workspace and heap memory,
- file descriptors,
- POSIX keys and key value pairs,
- RTEMS barriers,
- RTEMS user extensions,
- RTEMS message queues,
- RTEMS partitions,
- RTEMS periods,
- RTEMS regions,
- RTEMS semaphores,

- RTEMS tasks, and
- RTEMS timers.

Listing 8.7: Resource Accounting Example

```

1 #include <t.h>
2
3 #include <stdlib.h>
4
5 #include <rtems.h>
6
7 T_TEST_CASE(missing_sema_delete)
8 {
9     rtems_status_code sc;
10    rtems_id id;
11
12    sc = rtems_semaphore_create(rtems_build_name('S', 'E', 'M', 'A'), 0,
13        RTEMS_COUNTING_SEMAPHORE, 0, &id);
14    T_rsc_success(sc);
15 }
16
17 T_TEST_CASE(missing_free)
18 {
19     void *p;
20
21    p = malloc(1);
22    T_not_null(p);
23 }

```

Listing 8.8: Resource Accounting Report

```

1 B:missing_sema_delete
2 P:0:0:UI1:test-leak.c:14
3 F:*:0:UI1:*:*:RTEMS semaphore leak (1)
4 E:missing_sema_delete:N:1:F:1:D:0.004013
5 B:missing_free
6 P:0:0:UI1:test-leak.c:22
7 F:*:0:UI1:*:*:memory leak in workspace or heap
8 E:missing_free:N:1:F:1:D:0.003944

```

8.1.6 Test Case Scoped Dynamic Memory

You can allocate dynamic memory which is automatically freed after the current test case execution. You can provide an optional destroy function to `T_zalloc()` which is called right before the memory is freed. The `T_zalloc()` function initializes the memory to zero.

```

1 void *T_malloc(size_t size);
2
3 void *T_calloc(size_t nelem, size_t elsize);
4
5 void *T_zalloc(size_t size, void (*destroy)(void *));
6
7 void T_free(void *ptr);

```

Listing 8.9: Test Case Scoped Dynamic Memory Example

```
1 #include <t.h>
2
3 T_TEST_CASE(malloc_free)
4 {
5     void *p;
6
7     p = T_malloc(1);
8     T_assert_not_null(p);
9     T_free(p);
10 }
11
12 T_TEST_CASE(malloc_auto)
13 {
14     void *p;
15
16     p = T_malloc(1);
17     T_assert_not_null(p);
18 }
19
20 static void
21 destroy(void *p)
22 {
23     int *i;
24
25     i = p;
26     T_step_eq_int(2, *i, 1);
27 }
28
29 T_TEST_CASE(zalloc_auto)
30 {
31     int *i;
32
33     T_plan(3);
34     i = T_zalloc(sizeof(*i), destroy);
35     T_step_assert_not_null(0, i);
36     T_step_eq_int(1, *i, 0);
37     *i = 1;
38 }
```

Listing 8.10: Test Case Scoped Dynamic Memory Report

```

1 B:malloc_free
2 P:0:0:UI1:test-malloc.c:8
3 E:malloc_free:N:1:F:0:D:0.005200
4 B:malloc_auto
5 P:0:0:UI1:test-malloc.c:17
6 E:malloc_auto:N:1:F:0:D:0.004790
7 B:zalloc_auto
8 P:0:0:UI1:test-malloc.c:35
9 P:1:0:UI1:test-malloc.c:36
10 P:2:0:UI1:test-malloc.c:26
11 E:zalloc_auto:N:3:F:0:D:0.006583

```

8.1.7 Test Case Destructors

You can add test case destructors with `T_add_destructor()`. The destructors are called automatically at the test case end before the resource accounting takes place. Optionally, a registered destructor can be removed before the test case end with `T_remove_destructor()`. The `T_destructor` structure of a destructor must exist after the return from the test case body. It is recommended to use statically allocated memory. Do not use stack memory or dynamic memory obtained via `T_malloc()`, `T_calloc()` or `T_zalloc()` for the `T_destructor` structure.

```

1 void T_add_destructor(T_destructor *destructor,
2     void (*destroy)(T_destructor *));
3
4 void T_remove_destructor(T_destructor *destructor);

```

Listing 8.11: Test Case Destructor Example

```

1 #include <t.h>
2
3 static void
4 destroy(T_destructor *dtor)
5 {
6     (void)dtor;
7     T_step(0, "destroy");
8 }
9
10 T_TEST_CASE(destructor)
11 {
12     static T_destructor dtor;
13
14     T_plan(1);
15     T_add_destructor(&dtor, destroy);
16 }

```

Listing 8.12: Test Case Destructor Report

```

1 B:destructor
2 P:0:0:UI1:test-destructor.c:7
3 E:destructor:N:1:F:0:D:0.003714

```

8.1.8 Test Checks

A *test check* determines if the actual value presented to the test check has the expected properties. The actual value should represent the outcome of a test action. If a test action produces the expected outcome as determined by the corresponding test check, then this test check passes, otherwise this test check fails. A failed test check does not stop the test case execution immediately unless the *T_assert_**() test variant is used. Each test check increments the test step counter unless the *T_quiet_**() test variant is used. The test step counter is initialized to zero before the test case begins to execute. The *T_step_*(step, ...)* test check variants verify that the test step counter is equal to the planned test step value, otherwise the test check fails.

8.1.8.1 Test Check Variant Conventions

The *T_quiet_**() test check variants do not increment the test step counter and only print a message if the test check fails. This is helpful in case a test check appears in a tight loop.

The *T_step_*(step, ...)* test check variants check in addition that the test step counter is equal to the specified test step value, otherwise the test check fails.

The *T_assert_**() and *T_step_assert_*(step, ...)* test check variants stop the current test case execution if the test check fails.

8.1.8.2 Test Check Parameter Conventions

The following names for test check parameters are used throughout the test checks:

step

The planned test step for this test check.

a

The actual value to check against an expected value. It is usually the first parameter in all test checks, except in the *T_step_*(step, ...)* test check variants, here it is the second parameter.

e

The expected value of a test check. This parameter is optional. Some test checks have an implicit expected value. If present, then this parameter is directly after the actual value parameter of the test check.

fmt

A printf()-like format string. Floating-point and exotic formats may be not supported.

8.1.8.3 Test Check Condition Conventions

The following names for test check conditions are used:

eq

The actual value must equal the expected value.

ne

The actual value must not equal the value of the second parameter.

ge

The actual value must be greater than or equal to the expected value.

gt

The actual value must be greater than the expected value.

le

The actual value must be less than or equal to the expected value.

lt

The actual value must be less than the expected value.

If the actual value satisfies the test check condition, then the test check passes, otherwise it fails.

8.1.8.4 Test Check Type Conventions

The following names for test check types are used:

ptr

The test value must be a pointer (*void **).

mem

The test value must be a memory area with a specified length.

str

The test value must be a null byte terminated string.

nstr

The length of the test value string is limited to a specified maximum.

char

The test value must be a character (*char*).

schar

The test value must be a signed character (*signed char*).

uchar

The test value must be an unsigned character (*unsigned char*).

short

The test value must be a short integer (*short*).

ushort

The test value must be an unsigned short integer (*unsigned short*).

int

The test value must be an integer (*int*).

uint

The test value must be an unsigned integer (*unsigned int*).

long

The test value must be a long integer (*long*).

ulong

The test value must be an unsigned long integer (*unsigned long*).

ll

The test value must be a long long integer (*long long*).

ull

The test value must be an unsigned long long integer (*unsigned long long*).

i8

The test value must be a signed 8-bit integer (*int8_t*).

u8

The test value must be an unsigned 8-bit integer (*uint8_t*).

i16

The test value must be a signed 16-bit integer (*int16_t*).

u16

The test value must be an unsigned 16-bit integer (*uint16_t*).

i32

The test value must be a signed 32-bit integer (*int32_t*).

u32

The test value must be an unsigned 32-bit integer (*uint32_t*).

i64

The test value must be a signed 64-bit integer (*int64_t*).

u64

The test value must be an unsigned 64-bit integer (*uint64_t*).

iptr

The test value must be of type *intptr_t*.

uptr

The test value must be of type *uintptr_t*.

ssz

The test value must be of type *ssize_t*.

sz

The test value must be of type *size_t*.

8.1.8.5 Integers

Let *xyz* be the type variant which shall be one of *schar*, *uchar*, *short*, *ushort*, *int*, *uint*, *long*, *ulong*, *ll*, *ull*, *i8*, *u8*, *i16*, *u16*, *i32*, *u32*, *i64*, *u64*, *iptr*, *uptr*, *ssz*, and *sz*.

Let *I* be the type name which shall be compatible to the type variant.

The following test checks for integers are available:

```

1 void T_eq_xyz(I a, I e);
2 void T_assert_eq_xyz(I a, I e);
3 void T_quiet_eq_xyz(I a, I e);
4 void T_step_eq_xyz(unsigned int step, I a, I e);
5 void T_step_assert_eq_xyz(unsigned int step, I a, I e);
6
7 void T_ne_xyz(I a, I e);
8 void T_assert_ne_xyz(I a, I e);
9 void T_quiet_ne_xyz(I a, I e);
10 void T_step_ne_xyz(unsigned int step, I a, I e);
11 void T_step_assert_ne_xyz(unsigned int step, I a, I e);
12
13 void T_ge_xyz(I a, I e);
14 void T_assert_ge_xyz(I a, I e);
15 void T_quiet_ge_xyz(I a, I e);
16 void T_step_ge_xyz(unsigned int step, I a, I e);
17 void T_step_assert_ge_xyz(unsigned int step, I a, I e);
18
19 void T_gt_xyz(I a, I e);
20 void T_assert_gt_xyz(I a, I e);
21 void T_quiet_gt_xyz(I a, I e);
22 void T_step_gt_xyz(unsigned int step, I a, I e);
23 void T_step_assert_gt_xyz(unsigned int step, I a, I e);
24
25 void T_le_xyz(I a, I e);
26 void T_assert_le_xyz(I a, I e);
27 void T_quiet_le_xyz(I a, I e);
28 void T_step_le_xyz(unsigned int step, I a, I e);
29 void T_step_assert_le_xyz(unsigned int step, I a, I e);
30
31 void T_lt_xyz(I a, I e);
32 void T_assert_lt_xyz(I a, I e);
33 void T_quiet_lt_xyz(I a, I e);
34 void T_step_lt_xyz(unsigned int step, I a, I e);
35 void T_step_assert_lt_xyz(unsigned int step, I a, I e);

```

An automatically generated message is printed in case the test check fails.

8.1.8.6 Boolean Expressions

The following test checks for boolean expressions are available:

```

1 void T_true(bool a, const char *fmt, ...);
2 void T_assert_true(bool a, const char *fmt, ...);
3 void T_quiet_true(bool a, const char *fmt, ...);
4 void T_step_true(unsigned int step, bool a, const char *fmt, ...);
5 void T_step_assert_true(unsigned int step, bool a, const char *fmt, ...);
6
7 void T_false(bool a, const char *fmt, ...);
8 void T_assert_false(bool a, const char *fmt, ...);
9 void T_quiet_false(bool a, const char *fmt, ...);
10 void T_step_false(unsigned int step, bool a, const char *fmt, ...);
11 void T_step_assert_false(unsigned int step, bool a, const char *fmt, ...);

```

The message is only printed in case the test check fails. The format parameter is mandatory.

Listing 8.13: Boolean Test Checks Example

```

1 #include <t.h>
2
3 T_TEST_CASE(example)
4 {
5     T_true(true, "test passes, no message output");
6     T_true(false, "test fails");
7     T_quiet_true(true, "quiet test passes, no output at all");
8     T_quiet_true(false, "quiet test fails");
9     T_step_true(2, true, "step test passes, no message output");
10    T_step_true(3, false, "step test fails");
11    T_assert_false(true, "this is a format %s", "string");
12 }

```

Listing 8.14: Boolean Test Checks Report

```

1 B:example
2 P:0:0:UI1:test-example.c:5
3 F:1:0:UI1:test-example.c:6:test fails
4 F:*:0:UI1:test-example.c:8:quiet test fails
5 P:2:0:UI1:test-example.c:9
6 F:3:0:UI1:test-example.c:10:step test fails
7 F:4:0:UI1:test-example.c:11:this is a format string
8 E:example:N:5:F:4

```

8.1.8.7 Generic Types

The following test checks for data types with an equality (==) or inequality (!=) operator are available:

```

1 void T_eq(T a, T e, const char *fmt, ...);
2 void T_assert_eq(T a, T e, const char *fmt, ...);
3 void T_quiet_eq(T a, T e, const char *fmt, ...);
4 void T_step_eq(unsigned int step, T a, T e, const char *fmt, ...);
5 void T_step_assert_eq(unsigned int step, T a, T e, const char *fmt, ...);
6
7 void T_ne(T a, T e, const char *fmt, ...);
8 void T_assert_ne(T a, T e, const char *fmt, ...);
9 void T_quiet_ne(T a, T e, const char *fmt, ...);
10 void T_step_ne(unsigned int step, T a, T e, const char *fmt, ...);
11 void T_step_assert_ne(unsigned int step, T a, T e, const char *fmt, ...);

```

The type name *T* specifies an arbitrary type which must support the corresponding operator. The message is only printed in case the test check fails. The format parameter is mandatory.

8.1.8.8 Pointers

The following test checks for pointers are available:

```

1 void T_eq_ptr(const void *a, const void *e);
2 void T_assert_eq_ptr(const void *a, const void *e);
3 void T_quiet_eq_ptr(const void *a, const void *e);
4 void T_step_eq_ptr(unsigned int step, const void *a, const void *e);
5 void T_step_assert_eq_ptr(unsigned int step, const void *a, const void *e);
6
7 void T_ne_ptr(const void *a, const void *e);
8 void T_assert_ne_ptr(const void *a, const void *e);
9 void T_quiet_ne_ptr(const void *a, const void *e);
10 void T_step_ne_ptr(unsigned int step, const void *a, const void *e);
11 void T_step_assert_ne_ptr(unsigned int step, const void *a, const void *e);
12
13 void T_null(const void *a);
14 void T_assert_null(const void *a);
15 void T_quiet_null(const void *a);
16 void T_step_null(unsigned int step, const void *a);
17 void T_step_assert_null(unsigned int step, const void *a);
18
19 void T_not_null(const void *a);
20 void T_assert_not_null(const void *a);
21 void T_quiet_not_null(const void *a);
22 void T_step_not_null(unsigned int step, const void *a);
23 void T_step_assert_not_null(unsigned int step, const void *a);

```

An automatically generated message is printed in case the test check fails.

8.1.8.9 Memory Areas

The following test checks for memory areas are available:

```

1 void T_eq_mem(const void *a, const void *e, size_t n);
2 void T_assert_eq_mem(const void *a, const void *e, size_t n);
3 void T_quiet_eq_mem(const void *a, const void *e, size_t n);
4 void T_step_eq_mem(unsigned int step, const void *a, const void *e, size_t n);
5 void T_step_assert_eq_mem(unsigned int step, const void *a, const void *e, size_t n);
6
7 void T_ne_mem(const void *a, const void *e, size_t n);
8 void T_assert_ne_mem(const void *a, const void *e, size_t n);
9 void T_quiet_ne_mem(const void *a, const void *e, size_t n);
10 void T_step_ne_mem(unsigned int step, const void *a, const void *e, size_t n);
11 void T_step_assert_ne_mem(unsigned int step, const void *a, const void *e, size_t n);

```

The `memcmp()` function is used to compare the memory areas. An automatically generated message is printed in case the test check fails.

8.1.8.10 Strings

The following test checks for strings are available:

```

1 void T_eq_str(const char *a, const char *e);
2 void T_assert_eq_str(const char *a, const char *e);
3 void T_quiet_eq_str(const char *a, const char *e);
4 void T_step_eq_str(unsigned int step, const char *a, const char *e);
5 void T_step_assert_eq_str(unsigned int step, const char *a, const char *e);
6
7 void T_ne_str(const char *a, const char *e);
8 void T_assert_ne_str(const char *a, const char *e);
9 void T_quiet_ne_str(const char *a, const char *e);
10 void T_step_ne_str(unsigned int step, const char *a, const char *e);
11 void T_step_assert_ne_str(unsigned int step, const char *a, const char *e);
12
13 void T_eq_nstr(const char *a, const char *e, size_t n);
14 void T_assert_eq_nstr(const char *a, const char *e, size_t n);
15 void T_quiet_eq_nstr(const char *a, const char *e, size_t n);
16 void T_step_eq_nstr(unsigned int step, const char *a, const char *e, size_t n);
17 void T_step_assert_eq_nstr(unsigned int step, const char *a, const char *e, size_t n);
18
19 void T_ne_nstr(const char *a, const char *e, size_t n);
20 void T_assert_ne_nstr(const char *a, const char *e, size_t n);
21 void T_quiet_ne_nstr(const char *a, const char *e, size_t n);
22 void T_step_ne_nstr(unsigned int step, const char *a, const char *e, size_t n);
23 void T_step_assert_ne_nstr(unsigned int step, const char *a, const char *e, size_t n);

```

The *strcmp()* and *strncmp()* functions are used to compare the strings. An automatically generated message is printed in case the test check fails.

8.1.8.11 Characters

The following test checks for characters (*char*) are available:

```

1 void T_eq_char(char a, char e);
2 void T_assert_eq_char(char a, char e);
3 void T_quiet_eq_char(char a, char e);
4 void T_step_eq_char(unsigned int step, char a, char e);
5 void T_step_assert_eq_char(unsigned int step, char a, char e);
6
7 void T_ne_char(char a, char e);
8 void T_assert_ne_char(char a, char e);
9 void T_quiet_ne_char(char a, char e);
10 void T_step_ne_char(unsigned int step, char a, char e);
11 void T_step_assert_ne_char(unsigned int step, char a, char e);

```

An automatically generated message is printed in case the test check fails.

8.1.8.12 RTEMS Status Codes

The following test checks for RTEMS status codes are available:

```

1 void T_rsc(rtems_status_code a, rtems_status_code e);
2 void T_assert_rsc(rtems_status_code a, rtems_status_code e);
3 void T_quiet_rsc(rtems_status_code a, rtems_status_code e);
4 void T_step_rsc(unsigned int step, rtems_status_code a, rtems_status_code e);
5 void T_step_assert_rsc(unsigned int step, rtems_status_code a, rtems_status_code e);
6
7 void T_rsc_success(rtems_status_code a);
8 void T_assert_rsc_success(rtems_status_code a);
9 void T_quiet_rsc_success(rtems_status_code a);
10 void T_step_rsc_success(unsigned int step, rtems_status_code a);
11 void T_step_assert_rsc_success(unsigned int step, rtems_status_code a);

```

An automatically generated message is printed in case the test check fails.

8.1.8.13 POSIX Error Numbers

The following test checks for POSIX error numbers are available:

```

1 void T_eno(int a, int e);
2 void T_assert_eno(int a, int e);
3 void T_quiet_eno(int a, int e);
4 void T_step_eno(unsigned int step, int a, int e);
5 void T_step_assert_eno(unsigned int step, int a, int e);
6
7 void T_eno_success(int a);
8 void T_assert_eno_success(int a);
9 void T_quiet_eno_success(int a);
10 void T_step_eno_success(unsigned int step, int a);
11 void T_step_assert_eno_success(unsigned int step, int a);

```

The actual and expected value must be a POSIX error number, e.g. EINVAL, ENOMEM, etc. An automatically generated message is printed in case the test check fails.

8.1.8.14 POSIX Status Codes

The following test checks for POSIX status codes are available:

```

1 void T_psx_error(int a, int eno);
2 void T_assert_psx_error(int a, int eno);
3 void T_quiet_psx_error(int a, int eno);
4 void T_step_psx_error(unsigned int step, int a, int eno);
5 void T_step_assert_psx_error(unsigned int step, int a, int eno);
6
7 void T_psx_success(int a);
8 void T_assert_psx_success(int a);
9 void T_quiet_psx_success(int a);
10 void T_step_psx_success(unsigned int step, int a);
11 void T_step_assert_psx_success(unsigned int step, int a);

```

The *eno* value must be a POSIX error number, e.g. EINVAL, ENOMEM, etc. An actual value of zero indicates success. An actual value of minus one indicates an error. An automatically generated message is printed in case the test check fails.

Listing 8.15: POSIX Status Code Example

```

1 #include <t.h>
2
3 #include <sys/stat.h>
4 #include <errno.h>
5
6 T_TEST_CASE(stat)
7 {
8     struct stat st;
9     int status;
10
11     errno = 0;
12     status = stat("foobar", &st);
13     T_psx_error(status, ENOENT);
14 }
```

Listing 8.16: POSIX Status Code Report

```

1 B:stat
2 P:0:0:UI1:test-psx.c:13
3 E:stat:N:1:F:0
```

8.1.9 Log Messages and Formatted Output

You can print log messages with the *T_log()* function:

```

1 void T_log(T_verbosity verbosity, char const *fmt, ...);
```

A newline is automatically added to terminate the log message line.

Listing 8.17: Log Message Example

```

1 #include <t.h>
2
3 T_TEST_CASE(log)
4 {
5     T_log(T_NORMAL, "a log message %i, %i, %i", 1, 2, 3);
6     T_set_verbosity(T_QUIET);
7     T_log(T_NORMAL, "not verbose enough");
8 }
```

Listing 8.18: Log Message Report

```

1 B:log
2 L:a log message 1, 2, 3
3 E:log:N:0:F:0

```

You can use the following functions to print formatted output:

```

1 int T_printf(char const *, ...);
2
3 int T_vprintf(char const *, va_list);
4
5 int T_snprintf(char *, size_t, const char *, ...);

```

In contrast to the corresponding standard C library functions, floating-point and exotic formats may not be supported. On some architectures supported by RTEMS, floating-point operations are only supported in special tasks and may be forbidden in interrupt context. The formatted output functions provided by the test framework work in every context.

8.1.10 Utility

You can stop a test case via the `T_stop()` function. This function does not return. You can indicate unreachable code paths with the `T_unreachable()` function. If this function is called, then the test case stops.

You can busy wait with the `T_busy()` function:

```

1 void T_busy(uint_fast32_t count);

```

It performs a busy loop with the specified iteration count. This function is optimized to not perform memory accesses and should have a small jitter. The loop iterations have a processor-specific duration.

You can get an iteration count for the `T_busy()` function which corresponds roughly to one clock tick interval with the `T_get_one_clock_tick_busy()` function:

```

1 uint_fast32_t T_get_one_clock_tick_busy(void);

```

This function requires a clock driver. It must be called from thread context with interrupts enabled. It may return a different value each time it is called.

8.1.11 Time Services

The test framework provides two unsigned integer types for time values. The `T_ticks` unsigned integer type is used by the `T_tick()` function which measures time using the highest frequency counter available on the platform. It should only be used to measure small time intervals. The `T_time` unsigned integer type is used by the `T_now()` function which returns the current monotonic clock value of the platform, e.g. `CLOCK_MONOTONIC`.

```

1 T_ticks T_tick(void);
2
3 T_time T_now(void);

```


The reference time point for these two clocks is unspecified. You can obtain the test case begin time with the `T_case_begin_time()` function.

```
1 T_time T_case_begin_time(void);
```

You can convert time into ticks with the `T_time_to_ticks()` function and vice versa with the `T_ticks_to_time()` function.

```
1 T_time T_ticks_to_time(T_ticks ticks);
2
3 T_ticks T_time_to_ticks(T_time time);
```

You can convert seconds and nanoseconds values into a combined time value with the `T_seconds_and_nanoseconds_to_time()` function. You can convert a time value into separate seconds and nanoseconds values with the `T_time_to_seconds_and_nanoseconds()` function.

```
1 T_time T_seconds_and_nanoseconds_to_time(uint32_t s, uint32_t ns);
2
3 void T_time_to_seconds_and_nanoseconds(T_time time, uint32_t *s, uint32_t *ns);
```

You can convert a time value into a string representation. The time unit of the string representation is seconds. The precision of the string representation may be nanoseconds, microseconds, milliseconds, or seconds. You have to provide a buffer for the string (`T_time_string`).

```
1 const char *T_time_to_string_ns(T_time time, T_time_string buffer);
2
3 const char *T_time_to_string_us(T_time time, T_time_string buffer);
4
5 const char *T_time_to_string_ms(T_time time, T_time_string buffer);
6
7 const char *T_time_to_string_s(T_time time, T_time_string buffer);
```

Listing 8.19: Time String Example

```
1 #include <t.h>
2
3 T_TEST_CASE(time_to_string)
4 {
5     T_time_string ts;
6     T_time t;
7     uint32_t s;
8     uint32_t ns;
9
10    t = T_seconds_and_nanoseconds_to_time(0, 123456789);
11    T_eq_str(T_time_to_string_ns(t, ts), "0.123456789");
12    T_eq_str(T_time_to_string_us(t, ts), "0.123456");
13    T_eq_str(T_time_to_string_ms(t, ts), "0.123");
14    T_eq_str(T_time_to_string_s(t, ts), "0");
15
16    T_time_to_seconds_and_nanoseconds(t, &s, &ns);
17    T_eq_u32(s, 0);
18    T_eq_u32(ns, 123456789);
19 }
```

Listing 8.20: Time String Report

```

1 B:time_to_string
2 P:0:0:UI1:test-time.c:11
3 P:1:0:UI1:test-time.c:12
4 P:2:0:UI1:test-time.c:13
5 P:3:0:UI1:test-time.c:14
6 P:4:0:UI1:test-time.c:17
7 P:5:0:UI1:test-time.c:18
8 E:time_to_string:N:6:F:0:D:0.005250

```

You can convert a tick value into a string representation. The time unit of the string representation is seconds. The precision of the string representation may be nanoseconds, microseconds, milliseconds, or seconds. You have to provide a buffer for the string (*T_time_string*).

```

1 const char *T_ticks_to_string_ns(T_ticks ticks, T_time_string buffer);
2
3 const char *T_ticks_to_string_us(T_ticks ticks, T_time_string buffer);
4
5 const char *T_ticks_to_string_ms(T_ticks ticks, T_time_string buffer);
6
7 const char *T_ticks_to_string_s(T_ticks ticks, T_time_string buffer);

```

8.1.12 Code Runtime Measurements

You can measure the runtime of code fragments in several execution environment variants with the *T_measure_runtime()* function. This function needs a context which must be created with the *T_measure_runtime_create()* function. The context is automatically destroyed after the test case execution.

```

1 typedef struct {
2     size_t sample_count;
3 } T_measure_runtime_config;
4
5 typedef struct {
6     const char *name;
7     int flags;
8     void (*setup)(void *arg);
9     void (*body)(void *arg);
10    bool (*teardown)(void *arg, T_ticks *delta, uint32_t tic, uint32_t toc,
11        unsigned int retry);
12    void *arg;
13 } T_measure_runtime_request;
14
15 T_measure_runtime_context *T_measure_runtime_create(
16     const T_measure_runtime_config *config);
17
18 void T_measure_runtime(T_measure_runtime_context *ctx,
19     const T_measure_runtime_request *request);

```

The runtime measurement is performed for the *body* request handler of the measurement request (*T_measure_runtime_request*). The optional *setup* request handler is called before each invocation of the *body* request handler. The optional *teardown* request handler is called after each invocation of the *body* request handler. It has several parameters and a return status. If

it returns true, then this measurement sample value is recorded, otherwise the measurement is retried. The *delta* parameter is the current measurement sample value. It can be altered by the *teardown* request handler. The *tic* and *toc* parameters are the system tick values before and after the request body invocation. The *retry* parameter is the current retry counter. The runtime of the operational *setup* and *teardown* request handlers is not measured.

You can control some aspects of the measurement through the request flags (use zero for the default):

T_MEASURE_RUNTIME_ALLOW_CLOCK_ISR

Allow clock interrupts during the measurement. By default, measurements during which a clock interrupt happened are discarded unless it happens two times in a row.

T_MEASURE_RUNTIME_REPORT_SAMPLES

Report all measurement samples.

T_MEASURE_RUNTIME_DISABLE_FULL_CACHE

Disable the *FullCache* execution environment variant.

T_MEASURE_RUNTIME_DISABLE_HOT_CACHE

Disable the *HotCache* execution environment variant.

T_MEASURE_RUNTIME_DISABLE_DIRTY_CACHE

Disable the *DirtyCache* execution environment variant.

T_MEASURE_RUNTIME_DISABLE_MINOR_LOAD

Disable the *Load* execution environment variants with a load worker count less than the processor count.

T_MEASURE_RUNTIME_DISABLE_MAX_LOAD

Disable the *Load* execution environment variant with a load worker count equal to the processor count.

The execution environment variants (*M:V*) are:

FullCache

Before the *body* request handler is invoked a memory area with twice the size of the outermost data cache is completely read. This fills the data cache with valid cache lines which are unrelated to the *body* request handler. The cache is full with valid data and loading memory used by the handler needs to evict cache lines.

You can disable this variant with the *T_MEASURE_RUNTIME_DISABLE_FULL_CACHE* request flag.

HotCache

Before the *body* request handler is invoked the *body* request handler is called without measuring the runtime. The aim is to load all data used by the *body* request handler to the cache.

You can disable this variant with the *T_MEASURE_RUNTIME_DISABLE_HOT_CACHE* request flag.

DirtyCache

Before the *body* request handler is invoked a memory area with twice the size of the outermost data cache is completely written with new data. This should produce a data cache with dirty cache lines which are unrelated to the *body* request handler. In addition, the entire instruction cache is invalidated.

You can disable this variant with the *T_MEASURE_RUNTIME_DISABLE_DIRTY_CACHE* request flag.

Load/<WorkerCount>

This variant tries to get close to worst-case conditions. The cache is set up according to the *DirtyCache* variant. In addition, other processors try to fully load the memory system. The load is produced through writes to a memory area with twice the size of the outer-most data cache. The load variant is performed multiple times with a different set of active load worker threads. The <WorkerCount> value is the count of active workers which ranges from one to the processor count.

You can disable these variants with the `T_MEASURE_RUNTIME_DISABLE_MINOR_LOAD` and `T_MEASURE_RUNTIME_DISABLE_MAX_LOAD` request flags.

On SPARC, the *body* request handler is called with a register window setting so that window overflow traps will occur in the next level function call.

Each execution in an environment variant produces a sample set of *body* request handler runtime measurements. The minimum (*M:MI*), first quartile (*M:Q1*), median (*M:Q2*), third quartile (*M:Q3*), maximum (*M:MX*), median absolute deviation (*M:MAD*), and the sum of the sample values (*M:D*) is reported.

Listing 8.21: Code Runtime Measurement Example

```

1 #include <t.h>
2
3 static void
4 empty(void *arg)
5 {
6     (void)arg;
7 }
8
9 T_TEST_CASE(measure_empty)
10 {
11     static const T_measure_runtime_config config = {
12         .sample_count = 1024
13     };
14     T_measure_runtime_context *ctx;
15     T_measure_runtime_request req;
16
17     ctx = T_measure_runtime_create(&config);
18     T_assert_not_null(ctx);
19
20     memset(&req, 0, sizeof(req));
21     req.name = "Empty";
22     req.body = empty;
23     T_measure_runtime(ctx, &req);
24 }

```

Listing 8.22: Code Runtime Measurement Report

```

1 B:measure_empty
2 P:0:0:UI1:test-rtms-measure.c:18
3 M:B:Empty
4 M:V:FullCache
5 M:N:1024
6 M:MI:0.000000000
7 M:Q1:0.000000000
8 M:Q2:0.000000000

```

(continues on next page)

(continued from previous page)

```
9 M:Q3:0.000000000
10 M:MX:0.000000009
11 M:MAD:0.000000000
12 M:D:0.000000485
13 M:E:Empty:D:0.208984183
14 M:B:Empty
15 M:V:HotCache
16 M:N:1024
17 M:MI:0.000000003
18 M:Q1:0.000000003
19 M:Q2:0.000000003
20 M:Q3:0.000000003
21 M:MX:0.000000006
22 M:MAD:0.000000000
23 M:D:0.000002626
24 M:E:Empty:D:0.000017046
25 M:B:Empty
26 M:V:DirtyCache
27 M:N:1024
28 M:MI:0.000000007
29 M:Q1:0.000000007
30 M:Q2:0.000000007
31 M:Q3:0.000000008
32 M:MX:0.000000559
33 M:MAD:0.000000000
34 M:D:0.000033244
35 M:E:Empty:D:1.887834875
36 M:B:Empty
37 M:V:Load/1
38 M:N:1024
39 M:MI:0.000000000
40 M:Q1:0.000000002
41 M:Q2:0.000000002
42 M:Q3:0.000000003
43 M:MX:0.000000288
44 M:MAD:0.000000000
45 M:D:0.000002421
46 M:E:Empty:D:0.001798809
47 [... 22 more load variants ...]
48 M:E:Empty:D:0.021252583
49 M:B:Empty
50 M:V:Load/24
51 M:N:1024
52 M:MI:0.000000001
53 M:Q1:0.000000002
54 M:Q2:0.000000002
55 M:Q3:0.000000003
56 M:MX:0.000001183
57 M:MAD:0.000000000
58 M:D:0.000003406
59 M:E:Empty:D:0.015188063
60 E:measure_empty:N:1:F:0:D:14.284869
```

8.1.13 Interrupt Tests

In the operating system implementation you may have two kinds of critical sections. Firstly, there are low-level critical sections protected by interrupts disabled and maybe also some SMP spin lock. Secondly, there are high-level critical sections which are protected by disabled thread dispatching. The high-level critical sections may contain several low-level critical sections. Between these low-level critical sections interrupts may happen which could alter the code path taken in the high-level critical section.


The test framework provides support to write test cases for high-level critical sections through the `T_interrupt_test()` function:

```

1 typedef enum {
2     T_INTERRUPT_TEST_INITIAL,
3     T_INTERRUPT_TEST_ACTION,
4     T_INTERRUPT_TEST_BLOCKED,
5     T_INTERRUPT_TEST_CONTINUE,
6     T_INTERRUPT_TEST_DONE,
7     T_INTERRUPT_TEST_EARLY,
8     T_INTERRUPT_TEST_INTERRUPT,
9     T_INTERRUPT_TEST_LATE,
10    T_INTERRUPT_TEST_TIMEOUT
11 } T_interrupt_test_state;
12
13 typedef struct {
14     void (*prepare)(void *arg);
15     void (*action)(void *arg);
16     T_interrupt_test_state (*interrupt)(void *arg);
17     void (*blocked)(void *arg);
18     uint32_t max_iteration_count;
19 } T_interrupt_test_config;
20
21 T_interrupt_test_state T_interrupt_test(
22     const T_interrupt_test_config *config,
23     void *arg
24 );

```

This function returns `T_INTERRUPT_TEST_DONE` if the test condition was satisfied within the maximum iteration count, otherwise it returns `T_INTERRUPT_TEST_TIMEOUT`. The interrupt test run uses the specified configuration and passes the specified argument to all configured handlers. The function shall be called from thread context with interrupts enabled.



The interrupt test uses an *adaptive bisection algorithm* to try to hit the code section under test by an interrupt. In each test iteration, it waits for a time point one quarter of the clock tick interval after a clock tick using the monotonic clock. Then it performs a busy wait using `T_busy()` with a busy count controlled by the adaptive bisection algorithm. The test maintains a sample set of upper and lower bound busy wait count values. Initially, the lower bound values are zero and the upper bound values are set to a value returned by `T_get_one_clock_tick_busy()`. The busy wait count for an iteration is set to the middle point between the arithmetic mean of the lower and upper bound sample values. After the action handler returns, the set of lower and upper bound sample values is updated based on the test state. If the test state is `T_INTERRUPT_TEST_EARLY`, then the oldest upper bound sample value is replaced by the busy wait count used to delay the action and the latest lower bound sample value is slightly decreased. Reducing the lower bound helps to avoid a zero length interval between the upper and lower bounds. If the test state is `T_INTERRUPT_TEST_LATE`, then the oldest lower bound sample value is replaced by the busy wait count used to delay the action and the latest upper bound sample value is slightly increased. In all other test states the timing values remain as is. Using the arithmetic mean of a sample set dampens the effect of each test iteration and is an heuristic to mitigate the influence of jitters in the action code execution.

The optional *prepare* handler should prepare the system so that the *action* handler can be called. It is called in a tight loop, so all the time consuming setup should be done before `T_interrupt_test()` is called. During the preparation the test state is `T_INTERRUPT_TEST_INITIAL`. The preparation handler shall not change the test state.

The *action* handler should call the function which executes the code section under test. The execution path up to the code section under test should have a low jitter. Otherwise, the adaptive bisection algorithm may not find the right spot.

The *interrupt* handler should check if the test condition is satisfied or a new iteration is necessary. This handler is called in interrupt context. It shall return `T_INTERRUPT_TEST_DONE` if the test condition is satisfied and the test run is done. It shall return `T_INTERRUPT_TEST_EARLY` if the interrupt happened too early to satisfy the test condition. It shall return `T_INTERRUPT_TEST_LATE` if the interrupt happened too late to satisfy the test condition. It shall return `T_INTERRUPT_TEST_CONTINUE` if the test should continue with the current timing settings. Other states shall not be returned. It is critical to return the early and late states if the test condition was not satisfied, otherwise the adaptive bisection algorithm may not work. The returned state is used to try to change the test state from `T_INTERRUPT_TEST_ACTION` to the returned state.

The optional *blocked* handler is invoked if the executing thread blocks during the action processing. It should remove the blocking condition of the thread so that the next iteration can start. It can use `T_interrupt_change_state()` to change the interrupt test state.

The *max iteration count* configuration member defines the maximum iteration count of the test loop. If the maximum iteration count is reached before the test condition is satisfied, then `T_interrupt_test()` returns `T_INTERRUPT_TEST_TIMEOUT`.

The *interrupt* and *blocked* handlers may be called in arbitrary test states.

The *action*, *interrupt*, and *blocked* handlers can use `T_interrupt_test_get_state()` to get the current test state:

```
1 T_interrupt_test_state T_interrupt_test_get_state(void);
```

The *action*, *interrupt*, and *blocked* handlers can use `T_interrupt_test_change_state()` to try to change the test state from an expected state to a desired state:

```
1 T_interrupt_test_state T_interrupt_test_change_state(  
2     T_interrupt_test_state expected_state,  
3     T_interrupt_test_state desired_state  
4 );
```

The function returns the previous state. If it **differs from the expected state**, then the requested state **change to the desired state did not take place**. In an SMP configuration, do not call this function in a tight loop. It could lock up the test run. To busy wait for a state change, use `T_interrupt_test_get_state()`.

The *action* handler can use `T_interrupt_test_busy_wait_for_interrupt()` to busy wait for the interrupt:

```
1 void T_interrupt_test_busy_wait_for_interrupt(void);
```

This is useful if the action code does not block to wait for the interrupt. If the action handler just returns the test code immediately prepares the next iteration and may miss an interrupt which happens too late.

8.1.14 Test Runner

You can call the `T_main()` function to run all registered test cases.

```
1 int T_main(const T_config *config);
```

The `T_main()` function returns 0 if all test cases passed, otherwise it returns 1. Concurrent execution of the `T_main()` function is undefined behaviour.

You can ask if you execute within the context of the test runner with the `T_is_runner()` function:

```
1 bool T_is_runner(void);
```

It returns *true* if you execute within the context of the test runner (the context which executes for example `T_main()`). Otherwise it returns *false*, for example if you execute in another task, in interrupt context, nobody executes `T_main()`, or during system initialization on another processor.

On RTEMS, you have to register the test cases with the `T_register()` function before you call `T_main()`. This makes it possible to run low level tests, for example without the operating system directly in `boot_card()` or during device driver initialization. On other platforms, the `T_register()` is a no operation.

```
1 void T_register(void);
```

You can run test cases also individually. Use `T_run_initialize()` to initialize the test runner. Call `T_run_all()` to run all or `T_run_by_name()` to run specific registered test cases. Call `T_case_begin()` to begin a freestanding test case and call `T_case_end()` to finish it. Finally, call `T_run_finalize()`.

```
1 void T_run_initialize(const T_config *config);
2
3 void T_run_all(void);
4
5 void T_run_by_name(const char *name);
6
7 void T_case_begin(const char *name, const T_fixture *fixture);
8
9 void T_case_end(void);
10
11 bool T_run_finalize(void);
```

The `T_run_finalize()` function returns *true* if all test cases passed, otherwise it returns *false*. Concurrent execution of the runner functions (including `T_main()`) is undefined behaviour. The test suite configuration must be persistent throughout the test run.

```
1 typedef enum {
2     T_EVENT_RUN_INITIALIZE,
3     T_EVENT_CASE_EARLY,
4     T_EVENT_CASE_BEGIN,
5     T_EVENT_CASE_END,
6     T_EVENT_CASE_LATE,
7     T_EVENT_RUN_FINALIZE
8 } T_event;
9
10 typedef void (*T_action)(T_event, const char *);
11
12 typedef void (*T_putchar)(int, void *);
13
14 typedef struct {
15     const char *name;
16     char *buf;
17     size_t buf_size;
18     T_putchar putchar;
19     void *putchar_arg;
20     T_verbosity verbosity;
21     T_time (*now)(void);
22     size_t action_count;
23     const T_action *actions;
24 } T_config;
```

With the test suite configuration you can specify the test suite name, the put character handler used to output the test report, the initial verbosity, the monotonic time provider and an optional set of test suite actions. Only the test runner calls the put character handler, other tasks

or interrupt handlers write to a buffer which is emptied by the test runner on demand. You have to specify this buffer in the test configuration. The test suite actions are called with the test suite name for test suite run events (*T_EVENT_RUN_INITIALIZE* and *T_EVENT_RUN_FINALIZE*) and the test case name for the test case events (*T_EVENT_CASE_EARLY*, *T_EVENT_CASE_BEGIN*, *T_EVENT_CASE_END* and *T_EVENT_CASE_LATE*).

8.1.15 Test Verbosity

Three test verbosity levels are defined:

T_QUIET

Only the test suite begin, system, test case end, and test suite end lines are printed.

T_NORMAL

Prints everything except passed test lines.

T_VERBOSE

Prints everything.

The test verbosity level can be set within the scope of one test case with the *T_set_verbosity()* function:

```
1 T_verbosity T_set_verbosity(T_verbosity new_verbosity);
```

The function returns the previous verbosity. After the test case, the configured verbosity is automatically restored.

An example with *T_QUIET* verbosity:

```
1 A:xyz
2 S:Platform:RTEMS
3 [...]
4 E:a:N:2:F:1
5 E:b:N:0:F:1
6 E:c:N:1:F:1
7 E:d:N:6:F:0
8 Z:xyz:C:4:N:9:F:3
```

The same example with *T_NORMAL* verbosity:

```
1 A:xyz
2 S:Platform:RTEMS
3 [...]
4 B:a
5 F:1:0:UI1:test-verbosity.c:6:test fails
6 E:a:N:2:F:1
7 B:b
8 F:*:0:UI1:test-verbosity.c:12:quiet test fails
9 E:b:N:0:F:1
10 B:c
11 F:0:0:UI1:test-verbosity.c:17:this is a format string
12 E:c:N:1:F:1
13 B:d
14 E:d:N:6:F:0
15 Z:xyz:C:4:N:9:F:3
```

The same example with *T_VERBOSE* verbosity:

```

1 A:xyz
2 S:Platform:RTEMS
3 [...]
4 B:a
5 P:0:0:UI1:test-verbosity.c:5
6 F:1:0:UI1:test-verbosity.c:6:test fails
7 E:a:N:2:F:1
8 B:b
9 F:*:0:UI1:test-verbosity.c:12:quiet test fails
10 E:b:N:0:F:1
11 B:c
12 F:0:0:UI1:test-verbosity.c:17:this is a format string
13 E:c:N:1:F:1
14 B:d
15 P:0:0:UI1:test-verbosity.c:22
16 P:1:0:UI1:test-verbosity.c:23
17 P:2:0:UI1:test-verbosity.c:24
18 P:3:0:UI1:test-verbosity.c:25
19 P:4:0:UI1:test-verbosity.c:26
20 P:5:0:UI1:test-verbosity.c:27
21 E:d:N:6:F:0
22 Z:xyz:C:4:N:9:F:3

```

8.1.16 Test Reporting

The test reporting is line based which should be easy to parse with a simple state machine. Each line consists of a set of fields separated by colon characters (:). The first character of the line determines the line format:

A

A test suite begin line. It has the format:

A:<TestSuite>

A description of the field follows:

<TestSuite>

The test suite name. Must not contain colon characters (:).

S

A test suite system line. It has the format:

S:<Key>:<Value>

A description of the fields follows:

<Key>

A key string. Must not contain colon characters (:).

<Value>

An arbitrary key value string. May contain colon characters (:).

B

A test case begin line. It has the format:

B:<TestCase>

A description of the field follows:

<TestCase>

A test case name. Must not contain colon characters (:).

P

A test pass line. It has the format:

P:<Step>:<Processor>:<Task>:<File>:<Line>

A description of the fields follows:

<Step>

Each non-quiet test has a unique test step counter value in each test case execution. The test step counter is set to zero before the test case executes. For quiet test checks, there is no associated test step and the character * instead of an integer is used to indicate this.

<Processor>

The processor index of the processor which executed at least one instruction of the corresponding test.

<Task>

The name of the task which executed the corresponding test if the test executed in task context. The name *ISR* indicates that the test executed in interrupt context. The name ? indicates that the test executed in an arbitrary context with no valid executing task.

<File>

The name of the source file which contains the corresponding test. A source file of * indicates that no test source file is associated with the test, e.g. it was produced by the test framework itself.

<Line>

The line of the test statement in the source file which contains the corresponding test. A line number of * indicates that no test source file is associated with the test, e.g. it was produced by the test framework itself.

F

A test failure line. It has the format:

F:<Step>:<Processor>:<Task>:<File>:<Line>:<Message>

A description of the fields follows:

<Step> <Processor> <Task> <File> <Line>

See above **P** line.

<Message>

An arbitrary message string. May contain colon characters (:).

L

A log message line. It has the format:

L:<Message>

A description of the field follows:

<Message>

An arbitrary message string. May contain colon characters (:).

E

A test case end line. It has the format:

E:<TestCase>:N:<Steps>:F:<Failures>:D:<Duration>

A description of the fields follows:

<TestCase>

A test case name. Must not contain colon characters (:).

<Steps>

The final test step counter of a test case. Quiet test checks produce no test steps.

<Failures>

The count of failed test checks of a test case.

<Duration>

The test case duration in seconds.

Z

A test suite end line. It has the format:

Z:<TestSuite>:C:<TestCases>:N:<OverallSteps>:F:<OverallFailures>:D:<Duration>

A description of the fields follows:

<TestSuite>

The test suite name. Must not contain colon characters (:).

<TestCases>

The count of test cases in the test suite.

<OverallSteps>

The overall count of test steps in the test suite.

<OverallFailures>

The overall count of failed test cases in the test suite.

<Duration>

The test suite duration in seconds.

Y

Auxiliary information line. Issued after the test suite end. It has the format:

Y:ReportHash:SHA256:<Hash>

A description of the fields follows:

<Hash>

The SHA256 hash value of the test suite report from the begin to the end of the test suite.

M

A code runtime measurement line. It has the formats:

M:B:<Name>

M:V:<Variant>

M:N:<SampleCount>

M:S:<Count>:<Value>

M:MI:<Minimum>

M:Q1:<FirstQuartile>

M:Q2:<Median>

M:Q3:<ThirdQuartile>

M:MX: <Maximum>

M:MAD: <MedianAbsoluteDeviation>

M:D: <SumOfSampleValues>

M:E: <Name> :D: <Duration>

A description of the fields follows:

<Name>

A code runtime measurement name. Must not contain colon characters (:).

<Variant>

The execution variant which is one of **FullCache**, **HotCache**, **DirtyCache**, or **Load/WorkerCount**. The **<WorkerCount>** is the count of active workers which ranges from one to the processor count.

<SampleCount>

The sample count as defined by the runtime measurement configuration.

<Count>

The count of samples with the same value.

<Value>

A sample value in seconds.

<Minimum>

The minimum of the sample set in seconds.

<FirstQuartile>

The first quartile of the sample set in seconds.

<Median>

The median of the sample set in seconds.

<ThirdQuartile>

The third quartile of the sample set in seconds.

<Maximum>

The maximum of the sample set in seconds.

<MedianAbsoluteDeviation>

The median absolute deviation of the sample set in seconds.

<SumOfSampleValues>

The sum of all sample values of the sample set in seconds.

<Duration>

The runtime measurement duration in seconds. It includes time to set up the execution environment variant.

Listing 8.23: Example Test Report

```

1 A:xyz
2 S:Platform:RTEMS
3 S:Compiler:7.4.0 20181206 (RTEMS 5, RSB e0aec65182449a4e22b820e773087636edaf5b32, Newlib_
  ↪1d35a003f)
4 S:Version:5.0.0.820977c5af17c1ca2f79800d64bd87ce70a24c68
5 S:BSP:erc32

```

(continues on next page)

(continued from previous page)

```

6 S:RTEMS_DEBUG:1
7 S:RTEMS_MULTIPROCESSING:0
8 S:RTEMS_POSIX_API:1
9 S:RTEMS_PROFILING:0
10 S:RTEMS_SMP:1
11 B:timer
12 P:0:0:UI1:test-rtems.c:26
13 P:1:0:UI1:test-rtems.c:29
14 P:2:0:UI1:test-rtems.c:33
15 P:3:0:ISR:test-rtems.c:14
16 P:4:0:ISR:test-rtems.c:15
17 P:5:0:UI1:test-rtems.c:38
18 P:6:0:UI1:test-rtems.c:39
19 P:7:0:UI1:test-rtems.c:42
20 E:timer:N:8:F:0:D:0.019373
21 B:rsc_success
22 P:0:0:UI1:test-rtems.c:59
23 F:1:0:UI1:test-rtems.c:60:RTEMS_INVALID_NUMBER == RTEMS_SUCCESSFUL
24 F:*:0:UI1:test-rtems.c:62:RTEMS_INVALID_NUMBER == RTEMS_SUCCESSFUL
25 P:2:0:UI1:test-rtems.c:63
26 F:3:0:UI1:test-rtems.c:64:RTEMS_INVALID_NUMBER == RTEMS_SUCCESSFUL
27 E:rsc_success:N:4:F:3:D:0.011128
28 B:rsc
29 P:0:0:UI1:test-rtems.c:48
30 F:1:0:UI1:test-rtems.c:49:RTEMS_INVALID_NUMBER == RTEMS_INVALID_ID
31 F:*:0:UI1:test-rtems.c:51:RTEMS_INVALID_NUMBER == RTEMS_INVALID_ID
32 P:2:0:UI1:test-rtems.c:52
33 F:3:0:UI1:test-rtems.c:53:RTEMS_INVALID_NUMBER == RTEMS_INVALID_ID
34 E:rsc:N:4:F:3:D:0.011083
35 Z:xyz:C:3:N:16:F:6:D:0.047201
36 Y:ReportHash:SHA256:e5857c520dd9c9b7c15d4a76d78c21ccc46619c30a869ecd11bbcd1885155e0b

```

8.1.17 Test Report Validation

You can add the `T_report_hash_sha256()` test suite action to the test suite configuration to generate and report the SHA256 hash value of the test suite report. The hash value covers everything reported by the test suite run from the begin to the end. This can be used to check that the report generated on the target is identical to the report received on the report consumer side. The hash value is reported after the end of test suite line (Z) as auxiliary information in a Y line. Consumers may have to reverse a `\n` to `\r\n` conversion before the hash is calculated. Such a conversion could be performed by a particular put character handler provided by the test suite configuration.

8.1.18 Supported Platforms

The framework runs on FreeBSD, MSYS2, Linux and RTEMS.

8.2 Test Framework Requirements for RTEMS

The requirements on a test framework suitable for RTEMS are:

8.2.1 License Requirements

TF.License.Permissive

The test framework shall have a permissive open source license such as BSD-2-Clause.

8.2.2 Portability Requirements

TF.Portability

The test framework shall be portable.

TF.Portability.RTEMS

The test framework shall run on RTEMS.

TF.Portability.POSIX

The test framework shall be portable to POSIX compatible operating systems. This allows to run test cases of standard C/POSIX/etc. APIs on multiple platforms.

TF.Portability.POSIX.Linux

The test framework shall run on Linux.

TF.Portability.POSIX.FreeBSD

The test framework shall run on FreeBSD.

TF.Portability.C11

The test framework shall be written in C11.

TF.Portability.Static

Test framework shall not use dynamic memory for basic services.

TF.Portability.Small

The test framework shall be small enough to support low-end platforms (e.g. 64KiB of RAM/ROM should be sufficient to test the architecture port, e.g. no complex stuff such as file systems, etc.).

TF.Portability.Small.LinkTimeConfiguration

The test framework shall be configured at link-time.

TF.Portability.Small.Modular

The test framework shall be modular so that only necessary parts end up in the final executable.

TF.Portability.Small.Memory

The test framework shall not aggregate data during test case executions.

8.2.3 Reporting Requirements

TF.Reporting

Test results shall be reported.

TF.Reporting.Verbosity

The test report verbosity shall be configurable. This allows different test run scenarios, e.g. regression test runs, full test runs with test report verification against the planned test output.

TF.Reporting.Verification

It shall be possible to use regular expressions to verify test reports line by line.

TF.Reporting.Compact

Test output shall be compact to avoid long test runs on platforms with a slow output device, e.g. 9600 Baud UART.

TF.Reporting.PutChar

A simple output one character function provided by the platform shall be sufficient to report the test results.

TF.Reporting.NonBlocking

The output functions shall be non-blocking.

TF.Reporting.Printf

The test framework shall provide printf()-like output functions.

TF.Reporting.Printf.WithFP

There shall be a printf()-like output function with floating point support.

TF.Reporting.Printf.WithoutFP

There shall be a printf()-like output function without floating point support on RTEMS.

TF.Reporting.Platform

The test platform shall be reported.

TF.Reporting.Platform.RTEMS.Git

The RTEMS source Git commit shall be reported.

TF.Reporting.Platform.RTEMS.Arch

The RTEMS architecture name shall be reported.

TF.Reporting.Platform.RTEMS.BSP

The RTEMS BSP name shall be reported.

TF.Reporting.Platform.RTEMS.Tools

The RTEMS tool chain version shall be reported.

TF.Reporting.Platform.RTEMS.Config.Debug

The shall be reported if RTEMS_DEBUG is defined.

TF.Reporting.Platform.RTEMS.Config.Multiprocessing

The shall be reported if RTEMS_MULTIPROCESSING is defined.

TF.Reporting.Platform.RTEMS.Config.POSIX

The shall be reported if RTEMS_POSIX_API is defined.

TF.Reporting.Platform.RTEMS.Config.Profiling

The shall be reported if RTEMS_PROFILING is defined.

TF.Reporting.Platform.RTEMS.Config.SMP

The shall be reported if RTEMS_SMP is defined.

TF.Reporting.TestCase

The test cases shall be reported.

TF.Reporting.TestCase.Begin

The test case begin shall be reported.

TF.Reporting.TestCase.End

The test case end shall be reported.

TF.Reporting.TestCase.Tests

The count of test checks of the test case shall be reported.

TF.Reporting.TestCase.Failures

The count of failed test checks of the test case shall be reported.

TF.Reporting.TestCase.Timing

Test case timing shall be reported.

TF.Reporting.TestCase.Tracing

Automatic tracing and reporting of thread context switches and interrupt service routines shall be optionally performed.

8.2.4 Environment Requirements

TF.Environment

The test framework shall support all environment conditions of the platform.

TF.Environment.SystemStart

The test framework shall run during early stages of the system start, e.g. valid stack pointer, initialized data and cleared BSS, nothing more.

TF.Environment.BeforeDeviceDrivers

The test framework shall run before device drivers are initialized.

TF.Environment.InterruptContext

The test framework shall support test case code in interrupt context.

8.2.5 Usability Requirements

TF.Usability

The test framework shall be easy to use.

TF.Usability.TestCase

It shall be possible to write test cases.

TF.Usability.TestCase.Independence

It shall be possible to write test cases in modules independent of the test runner.

TF.Usability.TestCase.AutomaticRegistration

Test cases shall be registered automatically, e.g. via constructors or linker sets.

TF.Usability.TestCase.Order

It shall be possible to sort the registered test cases (e.g. random, by name) before they are executed.

TF.Usability.TestCase.Resources

It shall be possible to use resources with a life time restricted to the test case.

TF.Usability.TestCase.Resources.Memory

It shall be possible to dynamically allocate memory which is automatically freed once the test case completed.

TF.Usability.TestCase.Resources.File

It shall be possible to create a file which is automatically unlinked once the test case completed.

TF.Usability.TestCase.Resources.Directory

It shall be possible to create a directory which is automatically removed once the test case completed.

TF.Usability.TestCase.Resources.FileDescriptor

It shall be possible to open a file descriptor which is automatically closed once the test case completed.

TF.Usability.TestCase.Fixture

It shall be possible to use a text fixture for test cases.

TF.Usability.TestCase.Fixture.SetUp

It shall be possible to provide a set up handler for each test case.

TF.Usability.TestCase.Fixture.TearDown

It shall be possible to provide a tear down handler for each test case.

TF.Usability.TestCase.Context

The test case context shall be verified a certain points.

TF.Usability.TestCase.Context.VerifyAtEnd

After a test case execution it shall be verified that the context is equal to the context at the test case begin. This helps to ensure that test cases are independent of each other.

TF.Usability.TestCase.Context.VerifyThread

The test framework shall provide a function to ensure that the test case code executes in normal thread context. This helps to ensure that operating system service calls return to a sane context.

TF.Usability.TestCase.Context.Configurable

The context verified in test case shall be configurable at link-time.

TF.Usability.TestCase.Context.ThreadDispatchDisableLevel

It shall be possible to verify the thread dispatch disable level.

TF.Usability.TestCase.Context.ISRNestLevel

It shall be possible to verify the ISR nest level.

TF.Usability.TestCase.Context.InterruptLevel

It shall be possible to verify the interrupt level (interrupts enabled/disabled).

TF.Usability.TestCase.Context.Workspace

It shall be possible to verify the workspace.

TF.Usability.TestCase.Context.Heap

It shall be possible to verify the heap.

TF.Usability.TestCase.Context.OpenFileDescriptors

It shall be possible to verify the open file descriptors.

TF.Usability.TestCase.Context.Classic

It shall be possible to verify Classic API objects.

TF.Usability.TestCase.Context.Classic.Barrier

It shall be possible to verify Classic API Barrier objects.

TF.Usability.TestCase.Context.Classic.Extensions

It shall be possible to verify Classic API User Extensions objects.

TF.Usability.TestCase.Context.Classic.MessageQueues

It shall be possible to verify Classic API Message Queue objects.

TF.Usability.TestCase.Context.Classic.Partitions

It shall be possible to verify Classic API Partition objects.

TF.Usability.TestCase.Context.Classic.Periods

It shall be possible to verify Classic API Rate Monotonic Period objects.

TF.Usability.TestCase.Context.Classic.Regions

It shall be possible to verify Classic API Region objects.

TF.Usability.TestCase.Context.Classic.Semaphores

It shall be possible to verify Classic API Semaphore objects.

TF.Usability.TestCase.Context.Classic.Tasks

It shall be possible to verify Classic API Task objects.

TF.Usability.TestCase.Context.Classic.Timers

It shall be possible to verify Classic API Timer objects.

TF.Usability.TestCase.Context.POSIX

It shall be possible to verify POSIX API objects.

TF.Usability.TestCase.Context.POSIX.Keys

It shall be possible to verify POSIX API Key objects.

TF.Usability.TestCase.Context.POSIX.KeyValuePairs

It shall be possible to verify POSIX API Key Value Pair objects.

TF.Usability.TestCase.Context.POSIX.MessageQueues

It shall be possible to verify POSIX API Message Queue objects.

TF.Usability.TestCase.Context.POSIX.Semaphores

It shall be possible to verify POSIX API Named Semaphores objects.

TF.Usability.TestCase.Context.POSIX.Shms

It shall be possible to verify POSIX API Shared Memory objects.

TF.Usability.TestCase.Context.POSIX.Threads

It shall be possible to verify POSIX API Thread objects.

TF.Usability.TestCase.Context.POSIX.Timers

It shall be possible to verify POSIX API Timer objects.

TF.Usability.Assert

There shall be functions to assert test objectives.

TF.Usability.Assert.Safe

Test assert functions shall be safe to use, e.g. `assert(a == b)` vs. `assert(a = b)` vs. `assert_eq(a, b)`.

TF.Usability.Assert.Continue

There shall be assert functions which allow the test case to continue in case of an assertion failure.

TF.Usability.Assert.Abort

There shall be assert functions which abort the test case in case of an assertion failure.

TF.Usability.EasyToWrite

It shall be easy to write test code, e.g. avoid long namespace prefix `rtems_test_*`.

TF.Usability.Threads

The test framework shall support multi-threading.

TF.Usability.Pattern

The test framework shall support test patterns.

TF.Usability.Pattern.Interrupts

The test framework shall support test cases which use interrupts, e.g. `spintrcritical*`.

TF.Usability.Pattern.Parallel

The test framework shall support test cases which want to run code in parallel on SMP machines.

TF.Usability.Pattern.Timing

The test framework shall support test cases which want to measure the timing of code sections under various platform conditions, e.g. dirty cache, empty cache, hot cache, with load from other processors, etc..

TF.Usability.Configuration

The test framework shall be configurable.

TF.Usability.Configuration.Time

The timestamp function shall be configurable, e.g. to allow test runs without a clock driver.

8.2.6 Performance Requirements

TF.Performance.RTEMS.No64BitDivision

The test framework shall not use 64-bit divisions on RTEMS.

8.3 Off-the-shelf Test Frameworks

There are several [off-the-shelf test frameworks for C/C++](#). The first obstacle for test frameworks is the license requirement (*TF.License.Permissive*).

8.3.1 bdd-for-c

In the [bdd-for-c](#) framework the complete test suite must be contained in one file and the main function is generated. This violates *TF.Usability.TestCase.Independence*.

8.3.2 CBDD

The [CBDD](#) framework uses the [C blocks](#) extension from clang. This violates *TF.Portability.C11*.

8.3.3 Google Test

[Google Test 1.8.1](#) is supported by RTEMS. Unfortunately, it is written in C++ and is too heavy weight for low-end platforms. Otherwise it is a nice framework.

8.3.4 Unity

The [Unity Test API](#) does not meet our requirements. There was a [discussion on the mailing list in 2013](#).

8.4 Standard Test Report Formats

8.4.1 JUnit XML

A common test report format is [JUnit XML](#).

```

1 <?xml version="1.0" encoding="UTF-8" ?>
2 <testsuites id="xyz" name="abc" tests="225" failures="1262" time="0.001">
3   <testsuite id="def" name="ghi" tests="45" failures="17" time="0.001">
4     <testcase id="jkl" name="mno" time="0.001">
5       <failure message="pqr" type="stu"></failure>
6       <system-out>stdout</system-out>
7       <system-err>stderr</system-err>
8     </testcase>
9   </testsuite>
10 </testsuites>

```

The major problem with this format is that you have to output the failure count of all test suites and the individual test suite before the test case output. You know the failure count only after a complete test run. This runs contrary to requirement *TF.Portability.Small.Memory*. It is also a bit verbose (*TF.Reporting.Compact*).

It is easy to convert a full test report generated by *The RTEMS Test Framework* (page 170) to the JUnit XML format.

8.4.2 Test Anything Protocol

The [Test Anything Protocol](#) (TAP) is easy to consume and produce.

```

1 1..4
2 ok 1 - Input file opened
3 not ok 2 - First line of the input valid
4 ok 3 - Read the rest of the file
5 not ok 4 - Summarized correctly # TODO Not written yet

```

You have to know in advance how many test statements you want to execute in a test case. The problem with this format is that there is no standard way to provide auxiliary data such as test timing or a tracing report.

It is easy to convert a full test report generated by *The RTEMS Test Framework* (page 170) to the TAP format.

BSP BUILD SYSTEM

The purpose of the build system is to produce and install artefacts from the RTEMS sources such as static libraries, start files, linker command files, configuration header files, header files, test programs, package description files, and third-party build system support files for a specific BSP in a user controlled configuration.

9.1 Goals

The system should meet the following goals:

- The install location of artefacts should be the same as in previous build systems
- Easy build configuration of BSPs through configuration options
- Enable the BSP build configuration to be placed in a version control system (e.g. no local paths)
- Fast builds (also on Windows)
- Easy to maintain, e.g. add/remove a BSP, add/change/remove configuration options, add/remove a library, add/remove an object to/from a library, add/remove tests
- Reusable build specifications (e.g. generate documentation for BSP options for the user manual)
- Validation of built artefacts (e.g. ensure that the objects are built as specified using the DWARF debug information)
- Support building of BSPs external to the project
- Customization of the build (e.g. build only a subset of the RTEMS functions)
- Support alternative compilers such as clang instead of GCC
- Ability to unit test the build system
- Version control system friendly change sets (e.g. most changes are line based in text files)

Configurable things which depend on the host computer environment such as paths to tools are intended to be passed as command line options to the waf command line tool. Configurable things which define what is built and how the artefacts are configured are intended to be placed in configuration files that can be configuration controlled. The waf build script file called wscript should know nothing about the layout of the sources. What is built and how it is built should be completely defined by the user configuration and the build specification items.

9.2 Overview

For an overview of the build system, see the *BSP Build System* chapter of the [RTEMS User Manual](#).

9.3 Commands

This section explains how the *Build Item Type* (page 27) items determine what the following waf commands do.

9.3.1 BSP List

In the `./waf bsp_list` command, the BSP list is generated from the *Build BSP Item Type* (page 29) items.

9.3.2 BSP Defaults

In the `./waf bsp_defaults` command, the BSP defaults are generated from the *Build BSP Item Type* (page 29) and *Build Option Item Type* (page 35) items. Build specification items contribute to the command through the `do_defaults()` method in the `wscript`.

9.3.3 Configure

In the `./waf configure` command, the build specification items contribute to the command through the `prepare_configure()` and `do_configure()` methods in the `wscript`.

9.3.4 Build, Clean, and Install

In the `./waf`, `./waf clean`, and `./waf install` commands, the build specification items contribute to the commands through the `prepare_build()` and `do_build()` methods in the `wscript`.

9.4 UID Naming Conventions

Use the following patterns for *UID names* (page 17):

abi

Use the name `abi` for the GCC-specific ABI flags item of a BSP family. Each BSP family should have exactly one *Build Option Item Type* (page 35) item which defines the GCC-specific ABI flags for all base BSPs of the family. For an architecture named *arch* and a BSP family named *family*, the file path is `spec/build/bsps/arch/family/abi.yml`.

abclang

Use the name `abclang` for the clang-specific ABI flags item of a BSP family. Each BSP family may have at most one *Build Option Item Type* (page 35) item which defines the clang-specific ABI flags for all base BSPs of the family. For an architecture named *arch* and a BSP family named *family*, the file path is `spec/build/bsps/arch/family/abclang.yml`.

bsp*

Use the prefix `bsp` for base BSPs.

cfg*

Use the prefix `cfg` for `config.h` header options.

grp*

Use the prefix `grp` for groups.

lib*

Use the prefix `lib` for libraries.

linkcmds*

Use the prefix `linkcmds` for linker command files.

obj*

Use the prefix `obj` for objects. Use

- `objmpci` for objects which are enabled by `RTEMS_MULTIPROCESSING`,
- `objnet` for objects which are enabled by `RTEMS_NETWORKING`,
- `objnetnosmp` for objects which are enabled by `RTEMS_NETWORKING` and not `RTEMS_SMP`, and
- `objsmp` for objects which are enabled by `RTEMS_SMP`.

opt*

Use the prefix `opt` for options. Use

- `optclock*` for options which have something to do with the clock driver,
- `optconsole*` for options which have something to do with the console driver,
- `optirq*` for options which have something to do with interrupt processing,
- `optmem*` for options which have something to do with the memory configuration, map, settings, etc., and
- `optosc*` for options which have something to do with oscillators.

start

Use the name `start` for BSP start file items. Each architecture or BSP family should have a *Build Start File Item Type* (page 38) item which builds the start file of a BSP. For an architecture

named *arch* and a BSP family named *family*, the file path is `spec/build/bsps/arch/start.yml` or `spec/build/bsps/arch/family/start.yml`. It is preferable to have a shared start file for the architecture instead of a start file for each BSP family.

tst*

Use the prefix `tst` for test states.

9.5 Build Specification Items

Specification items of refinements of the *Build Item Type* (page 27) are used by the `wscript` to determine what it should do. The `wscript` does not provide default values. All values are defined by specification items. The entry point to what is built are the *Build BSP Item Type* (page 29) items and the top-level *Build Group Item Type* (page 32) item. The user configuration defines which BSPs are built. The top-level group defaults to `/grp` and can be changed by the `--rtems-top-level` command line option given to the `waf configure` command.

The top-level group is a trade-off between the specification complexity and a strict dependency definition. Dependencies are normally explicit though the item links. However, using only explicit dependencies would make the specification quite complex, see Fig. 9.1. The top-level group and explicit *Build BSP Item Type* (page 29) items reduce the specification complexity since they use a priori knowledge of global build dependencies, see Fig. 9.2 for an example. This approach makes the build system easier, but less general.

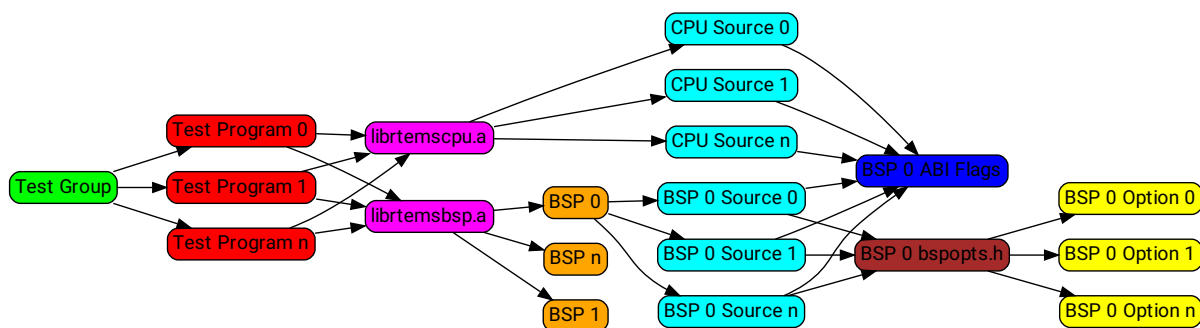


Fig. 9.1: Example with Explicit Item Links

This example shows how build item dependencies are specified explicitly by item links. In this example, a user wants to build a group of tests. Each test program has a dependency on the standard RTEMS libraries. The first issue is that the `librtembsp.a` needs dependencies to all base BSP variants (more than 100). The dependencies are the values of the `links` attribute in the library item files. External BSPs would have to modify the library item files. This is quite undesirable. The second issue is that the source files of the `librtemscpu.a` need a dependency to the ABI compiler flags specified by each BSP.

The third issue is that each BSP has to define its own `bspopts.h` configuration header item.

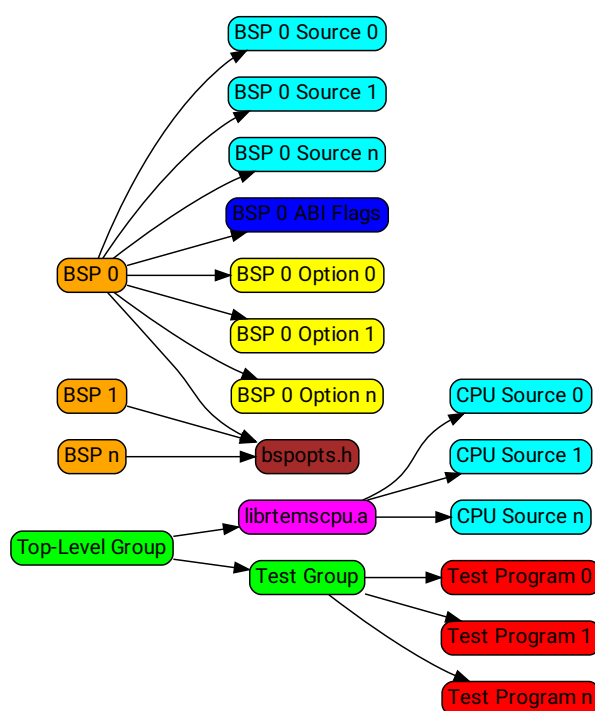


Fig. 9.2: Example with Implicit Ordering Rules

This example shows how build item dependencies are specified by dedicated BSP items, a top-level group, and ordered item links. The BSP is configured after the top-level group item and built before the top-level group item (defined by `wscript` source code). The library group is configured and built before the test group as specified by the item link order in the top-level group. The BSP options are processed before the results are written to the configuration header `bspopts.h` as defined by the BSP item link order.

9.6 How-To

This section presents how to do common maintenance tasks in the build system.

9.6.1 Find the Right Item

You find all build specification items in the RTEMS sources under the `spec/build` directory. You can use the `grep` command line tool to search in the build specification items.

9.6.2 Create a BSP Architecture

Let *arch* be the name of the architecture. You can add the new BSP architecture with:

```
1 $ mkdir spec/build/bsps/arch
```

For a new architecture try to use a shared start file which can be used by all BSPs of the architecture. Add a *Build Start File Item Type* (page 38) item for it:

```
1 $ vi spec/build/bsps/arch/start.yml
```

9.6.3 Create a BSP Family

Let *family* be the BSP family name and *arch* the name of the architecture. You can add the new BSP family with:

```
1 $ mkdir spec/build/bsps/arch/family
```

Add a *Build Option Item Type* (page 35) item for the ABI flags of the BSP family:

```
1 $ vi spec/build/bsps/arch/family/abi.yml
```

Define the ABI flags for each base BSP of the family. The `#{ABI_FLAGS}` are used for the `#{ASFLAGS}`, `#{CFLAGS}`, `#{CXXFLAGS}`, and `#{LDFLAGS}` build environment variables. Please have a look at the following example which shows the GCC-specific ABI flags item of the `sparc/leon3` BSP family:

```
1 SPDX-License-Identifier: CC-BY-SA-4.0 OR BSD-2-Clause
2 actions:
3 - get-string: null
4 - split: null
5 - env-append: null
6 build-type: option
7 copyrights:
8 - Copyright (C) 2020 embedded brains GmbH (http://www.embedded-brains.de)
9 default:
10 - -mcpu=leon3
11 default-by-variant:
12 - value:
13   - -mcpu=leon3
14   - -mfix-ut700
15 variants:
```

(continues on next page)

(continued from previous page)

```

16 - sparc/ut700
17 - value:
18 - -mcpu=leon
19 - -mfix-ut699
20 variants:
21 - sparc/ut699
22 - value:
23 - -mcpu=leon3
24 - -mfix-gr712rc
25 variants:
26 - sparc/gr712rc
27 description: |
28 ABI flags
29 enabled-by: gcc
30 links: []
31 name: ABI_FLAGS
32 type: build

```

If the architecture has no shared start file, then add a *Build Start File Item Type* (page 38) item for the new BSP family:

```

1 $ vi spec/build/bsps/arch/family/start.yml

```

9.6.4 Add a Base BSP to a BSP Family




Fig. 9.3: This example shows a BSP family named *family* in the architecture *arch* which consists of only one base BSP named *xyz*. The BSP sources and installation information is contained in the `spec:/build/bsps/arch/family/bspxyz` BSP item. The items linked by the BSP item are shown using relative UIDs.

Let *family* be the BSP family name, *arch* the name of the architecture, and *new* the name of the new base BSP. You can add the new base BSP with:

```

1 $ vi spec/build/bsps/arch/family/bspnew.yml

```

Define the attributes of your new base BSP according to *Build BSP Item Type* (page 29).

In case the BSP family has no group, then create a group if it is likely that the BSP family will contain more than one base BSP (see *Extend a BSP Family with a Group* (page 224)).




Fig. 9.4: This example shows a BSP family named *family* in the architecture *arch* which consists of three base BSPs named *rst*, *uvw*, and *xyz*. The BSP sources and installation information is contained in the *obj* objects item. The group *grp* defines the main BSP constituents. The base BSP items `spec:/build/bsps/arch/family/bsprst`, `spec:/build/bsps/arch/family/bspuvw`, and `spec:/build/bsps/arch/family/bspxyz` just define the name of the base BSP and set a link to the group item. The base BSP names can be used for example in the default-by-variant attribute of *Build Option Item Type* (page 35) items. The items linked by the BSP items are shown using relative UIDs.

If the BSP family has a group, then link the new base BSP to the group with:

```
1 $ vi spec/build/bsps/arch/family/grp.yml
```

Add a link using a relative UID to the new base BSP item:

```
1 links:
2 - role: build-dependency
3   uid: bspnew
```

9.6.5 Add a BSP Option

Let *family* be the BSP family name, *arch* the name of the architecture, and *new* the name of the new BSP option. You can add the new BSP option with:

```
1 $ vi spec/build/bsps/arch/family/optnew.yml
```

Define the attributes of your new BSP option according to *Build Option Item Type* (page 35). Link the option item to the corresponding group or BSP item using a relative UID:

```
1 links:
2 - role: build-dependency
3   uid: optnew
```

9.6.6 Extend a BSP Family with a Group

Let *family* be the BSP family name and *arch* the name of the architecture. If you want to have more than one base BSP in a BSP family, then you have to use a group item (see *Build Group Item Type* (page 32)). Add the group item named *grp* to the family with:

```
1 $ vi spec/build/bmps/arch/family/grp.yml
```

Define the attributes of your new group according to *Build Group Item Type* (page 32) and move the links of the existing base BSP item to the group item. Add a link to *obj*.

Add an objects item named *obj* to the family with:

```
1 $ vi spec/build/bmps/arch/family/obj.yml
```

Define the attributes of your new objects item according to *Build Objects Item Type* (page 34) and move the *cflags*, *cppflags*, *includes*, *install* and *source* attribute values of the existing base BSP item to the objects item.

9.6.7 Add a Test Program

Let *collection* be the name of a test program collection and *new* the name of the new test program. You can add the new test program with:

```
1 $ vi spec/build/testsuites/collection/new.yml
```

Define the attributes of your new test program according to *Build Test Program Item Type* (page 38).

Edit corresponding group item of the test program collection:

```
1 $ vi spec/build/testsuites/collection/grp.yml
```

Add a link to the new test program item using a relative UID:

```
1 links:
2 - role: build-dependency
3   uid: new
```

9.6.8 Add a Library

Let *new* be the name of the new library. You can add the new library with:

```
1 $ vi spec/build/cpukit/libnew.yml
```

Define the attributes of your new library according to *Build Library Item Type* (page 33).

Edit corresponding group item:

```
1 $ vi spec/build/cpukit/grp.yml
```

Add a link to the new library item using a relative UID:

```
1 links:  
2 - role: build-dependency  
3   uid: libnew
```


SOFTWARE RELEASE MANAGEMENT

10.1 Release Process

The release process creates an RTEMS release. The process has a number of stages that happen before a release can be made, during the creation of the release and after the release has been made.

10.1.1 Releases

RTEMS is released as a collection of ready to use source code and built documentation. Releases are publicly available on the RTEMS servers under <https://ftp.rtems.org/pub/rtems/releases>.

Releases are group under the major version number as a collection of directories consisting of the version number. This is termed a release series. A release may also contain release snapshots.

All releases must have a three digit version number and this can be optionally followed by a dash character (-) and an identifier, e.g. 5.1.0-acme-1.

The RTEMS Project reserves releases with only the three digit version number, e.g. 5.1.0. This identifies an RTEMS Project release.

10.1.1.1 Release Layout

- All released source archives are XZ compressed tar files.
- Top level contains:

README.txt:

A set of brief release instructions.

contrib:

Contributed sources. For example the release scripts used to create the release.

docs:

Compressed documentation build in HTML, Single page HTML and PDF formats. Provide compressed files for each document and a single archive of all the documentation. Provide an SHA512 check sum file.

rtems-<VERSION>-release-notes.pdf:

RTEMS Release notes document the changes in a release. This is a capture of the Trac report for the release's milestone in PDF format.

sha512sum.txt:

SHA512 checksum of all files in this directory.

sources:

All source code referenced by the release.

10.1.1.2 Release Version Numbering

The release numbering scheme changed with RTEMS 5.

The master branch has the version N.0.0 with N being the next major release number. The release branch in a repository will be just the major number.

The first release of this series will have the version number N.1.0. The first bugfix release (minor release) of this series will have the version number N.2.0.

The release branch will have the version number N.M.1 with M being the last minor release of this series. Tools will use N as the version number and must be compatible with all releases and the release branch of the N series.

Examples:

- 5.0.0 is the version number of the development master for the 5 series
- 5.1.0 is the first release of the 5 series
- 5.1.1 is the version number of the 5 series release branch right after the 5.1.0 release until 5.2.0 is released
- 5.2.0 is the first bugfix release of the 5 series
- 5.2.1 is the version number of the 5 series release branch right after the 5.2.0 release until 5.3.0 is released
- 6.0.0 is the version number of the development master for the 6 series

10.1.1.3 Release Scripts

- The release scripts are held in the top level repository <https://git.rtems.org/rtems-release.git>.
- The scripts are written for FreeBSD and can run on FreeBSD 10 through FreeBSD 12. No other host operating system is supported for the release scripts. Updates are welcome if the changes do not affect the operation on FreeBSD.
- A Python virtualenv environment is required for a working Sphinx documentation building environment. Follow the procedure in the `rtems-docs.git` top level README file.
- Building a standard release requires you provide the release major number and the release's remaining version string including any additional identifiers:

```
1 ./rtems-release 5 1.0
```

To create a release snapshot:

```
1 ./rtems-release 5 0.0-m2003
```

- A 3rd option of a release URL can be provided to create a test or deployable release. The URL is a base path the RSB uses to download the release source files from:

```
1 ./rtems-release \  
2 -u https://ftp.rtems.org/pub/rtems/people/chrisj/releases \  
3 5 0.0-m2003-2
```

- Building the release notes requires the Web Toolkit HTML to PDF converter be installed. The FreeBSD package is wkhtmltopdf.

10.1.1.4 Release Snapshots

- Release snapshots are only created for the current development version of RTEMS. For example RTEMS 5 snapshot path is 5/5.0.0/5.0.0-m2003.
- Release snapshots are based on the development sources and may be unstable or not suitable for use in production.
- A release snapshot is created each month and is named as <major>/<version>/<version>-<YYMM> where YY is the last two digits of the current year and MM is the month as a two digit number.
- In the lead up to a release more than one snapshot can be created by appending -<count> to the snapshot version string where <count> is incremented starting from 1. The first snapshot without a count is considered number 0.
- Release snapshots maybe removed from the RTEMS servers at the discretion of the RTEMS project

10.1.2 Release Repositories

The following are the repositories that a release effects. Any repository action is to be performed in the following repositories:

1. rtems.git
2. rtems-docs.git
3. rtems-examples.git
4. rtems-libbsd.git
5. rtems-source-builder.git
6. rtems-tools.git
7. rtems_waf.git
8. rtems-release.git

10.1.3 Pre-Release Procedure

- All tickets must be resolved, closed or moved to a later milestone. Tickets can exist that are specific to the branch and are to be resolved before the first release is made.
- Create release snapshots and post suitable build and test results.

10.1.4 Release Branching

A release has a release branch in each of the release repositories. A release is created from a release branch. The release branch label is the RTEMS major version number.

10.1.4.1 LibBSD Release Branch

The `rtems-libbsd.git` is an exception as it has two active release branches. The repository has a release branch based on the master like all the release repositories and it can have a FreeBSD version specific release branch that is used in the release.

LibBSD runs two branches during its development cycle. The master branch tracks the FreeBSD master branch. This means LibBSD tracks FreeBSD's development. LibBSD also tracks a FreeBSD branch for the RTEMS release. For example RTEMS 5 tracks FreeBSD 12 as its release base. This provides functional stability to the RTEMS 5 release by allowing a control process to track bug fixes in FreeBSD 12.

10.1.4.2 Pre-Branch Procedure

- All tickets assigned to the release's first milestone must be resolved. Tickets can exist that are specific to the branch and are to be resolved before the first release is made.
- The following BSP must build using the RSB:
 - `arm/beagleboneblack`
- Check and make sure the RSB kernel, libbsd and tools configurations reference the master when the branch is made.

The RSB GIT builds reference a specific commit so it is important the relevant configurations are valid.

10.1.4.3 Branch Procedure

- Branch labels are the major number as branch releases increment the minor number. A branch is only created when the first major release is made.

The commands to set a remote branch for a release in a repository are:

```
1 git clone <URL>/<REPO> <REPO>
2 cd <REPO>
3 git checkout -b <VERSION> origin/master
4 git push origin <VERSION>
```

Example:

```
1 git clone ssh://chrisj@dispatch.rtems.org/data/git/rtems.git rtems.git
2 cd rtems.git
3 git checkout -b 5 origin/master
4 git push origin 5
```

- Check and make sure the RSB kernel, libbsd and tools reference the branch commit.

10.1.4.4 Post-Branch Procedure

1. Create a release page for the next RTEMS release in Trac.

2. Update the releases table. The page link is:

<https://devel.rtems.org/wiki/Release>

Update the table adding the new development release to the top moving down the previous releases.

Label the new release branch as “Releasing”. The documentation link is left pointing to master until the release is made and the documentation is installed on the RTEMS Documentation web site.

3. Update the release table in the front page of the Trac Wiki. The page link is:

<https://devel.rtems.org/wiki/>

4. Add the milestones for the new development branch. The Trac page is:

```

1 = 6.1 (open)
2
3 == Statistics
4
5 || '''Total'''|[[TicketQuery(milestone=6.1,count)]]
6 || Fixed|[[TicketQuery(status=closed&milestone=6.1,resolution=fixed,
7 || Invalid|[[TicketQuery(status=closed&milestone=6.1,resolution=invalid,
8 || Works for me|[[TicketQuery(status=closed&milestone=6.1,
9 || Duplicate|[[TicketQuery(status=closed&milestone=6.1,
10 || Won't fix|[[TicketQuery(status=closed&milestone=6.1,resolution=wontfix,
11
12 == Distribution
13 [[TicketQuery(milestone=6.1&group=type,format=progress)]]
14
15 == Summary
16 [[TicketQuery(milestone=6.1)]]
17
18 == Details
19 [[TicketQuery(col=id|time|resolution|component|reporter|owner|changetime,
20 ||status=closed&milestone=6.1,rows=summary|description,table)]]

```

Replace 6.1 with the required milestone.

1. Create the RC1 release candidate with the source as close the branch point as possible.
2. Create a ticket to the clean the RSB for the release. The RSB’s master branch carries a number of older configurations and new release configurations. These can be confusing to a new user and add no value to a released RSB. For example leaving RTEMS 6 tool building configurations in the RTEMS 5 release.

Post-Branch Version Number Updates

After the release repository have been branched the master branch has to have the major version number updated. The follow is a list of the needed changes.

1. RTEMS requires the following files be changed:
 - `aclocal/version.m4`
 - `c/src/aclocal/version.m4`
 - `cpukit/aclocal/version.m4`
 - Doxyfile
 - `testsuites/aclocal/version.m4`
 - `rtems-bsps`
2. RTEMS Documentation the following files be changed:
 - `wscript`: Update `rtems_major_version`.
3. RSB requires the following files be changed:
 - `source-builder/sb/version.py`: Update `_version`.
4. RTEMS Tools requires the following files be changed:
 - `config/rtems-version.ini`: Update revision.
 - `tester/rtems/version.cfg`: Update `rtems_version`.
5. `rtems-libbsd` requires the following files and branches be changed:
 - `README.md`: Update Branches section.
 - `wscript`: Update `rtems_version`.
 - Create a new branch for tracking the FreeBSD stable version, for example `6-freebsd-12`.
6. `rtems-examples` requires the following files be changed:
 - `wscript`: Update `rtems_version`.

10.1.5 Release Procedure

The release procedure can be performed on any FreeBSD machine and uploaded to the RTEMS FTP server. You will need ssh access to the RTEMS server `dispatch.rtems.org` and suitable permissions to write into the FTP release path on the RTEMS server.

1. The release process starts by branching the repositories. To branch run the script:

```
1 ./rtems-release-branch [-p] <USER> <VERSION> <REVISION>
```

Example:

```

1 cd
2 mkdir -p development/rtems/releases
3 cd development/rtems/releases
4 git clone git://git.rtems.org/rtems-release.git rtems-release.git
5 cd rtems-release.git
6 ./rtems-release-branch -p chrisj 5

```

You need to have suitable commit access to the repositories.

2. To create the RTEMS release run the release script:

```
1 ./rtems-release <VERSION> <REVISION>
```

Example:

```
1 ./rtems-release 5 1.0
```

3. Copy the release to the RTEMS FTP server:

```

1 ssh <user>@dispatch.rtems.org mkdir -p /data/ftp/pub/rtems/releases/<VERSION>
2 scp -r <VERSION>.<REVISION> <user>@dispatch.rtems.org:/data/ftp/pub/rtems/releases/
  ↵<VERSION>/

```

Example:

```

1 ssh chrisj@dispatch.rtems.org mkdir -p /data/ftp/pub/rtems/releases/5
2 scp -r 5.1.0 chrisj@dispatch.rtems.org:/data/ftp/pub/rtems/releases/5/

```

4. Verify the release has been uploaded by checking the link:

<https://ftp.rtems.org/pub/rtems/releases/<VERSION>/<VERSION>.<REVISION>>

5. Tag the release repositories with the following command:

```

1 git checkout -b origin/<VERSION>
2 git tag <TAG>
3 git push origin <TAG>

```

Example:

```

1 git clone ssh://chrisj@dispatch.rtems.org/data/git/rtems.git rtems.git
2 cd rtems.git
3 git checkout -b origin/5
4 git tag 5.1.0
5 git push origin 5.1.0

```

10.1.6 Post-Release Procedure

The following procedures are performed after a release has been created.

1. Update the release to the RTEMS servers:

```
1 rsync --rsh=ssh -arv 5.1.0 chrisj@dispatch.rtems.org:/data/ftp/pub/rtems/releases/5/.
```

2. Test a build of the beagleboneblack BSP.

10.2 Software Change Report Generation

TBD - What goes here?

10.3 Version Description Document (VDD) Generation

TBD - discuss how generated. Preferably Dannie's project

This URL may be of use but it probably Trac auto-generated and can only be referenced: <https://devel.rtems.org/wiki/TracChangeLog>

USER'S MANUALS

TBD - write and link to useful documentation, potential URLs:

Reference the RTEMS Classic API Guide

- https://docs.rtems.org/doc-current/share/rtems/pdf/c_user.pdf

Reference any other existing user documentation

- <https://docs.rtems.org/doxygen/cpukit/html/index.html>
- <https://devel.rtems.org/>
- <http://www.rtems.com/>
- <https://www.rtems.org/onlinedocs.html>
- <https://devel.rtems.org/wiki/Developer/Contributing>
- <https://docs.rtems.org/releases/rtemsdocs-4.10.1/share/rtems/html/>

11.1 Documentation Style Guidelines

TBD - write me

LICENSING REQUIREMENTS

All artifacts shall adhere to RTEMS Project licensing requirements. Currently, the preferred licenses are:

- “Two Clause BSD” (BSD-2-Clause) for source code, and
- CC-BY-SA-4.0 license for documentation

Historically, RTEMS has been licensed under the GPL v2 with linking exception (<https://www.rtems.org/license>). It is preferred that new submissions be under one of the two preferred licenses. If you have previously submitted code to RTEMS under a historical license, please grant the project permission to relicense. See <https://devel.rtems.org/ticket/3053> for details.

For example templates for what to include in source code and documentation, see *Copyright and License Block* (page 151).

12.1 Rationale

RTEMS is intended for use in real-time embedded systems in which the application is statically linked with the operating system and all support libraries. Given this use case, the RTEMS development team evaluated a variety of licenses with the goal of promoting use while protecting both users and the developers.

Using the GNU General Public License Version 2 (GPLv2) unmodified was considered but discarded because the GPL can only be linked statically with other GPL code. Put simply, linking your application code statically with GPL code would cause your code to become GPL code. This would force both licensing and redistribution requirements onto RTEMS users. This was completely unacceptable.

The GNU Lesser General Public License Version 2 (LGPLv2) was also considered and deemed to not be a suitable license for RTEMS. This is because it either requires use of a shared library that can be re-linked, or release of the linked (application) code. This would require an RTEMS-based embedded system to provide a “relinking kit.” Again, this license would force an unacceptable requirement on RTEMS users and deemed unacceptable.

Newer versions of the GPL (i.e. version 3) are completely unsuitable for embedded systems due to the additions which add further restrictions on end user applications.

The historical RTEMS [License](#) is a modified version of the GPL version 2 that includes an exception to permit including headers and linking against RTEMS object files statically. This was based on the license used by GCC language runtime libraries at that time. This license allows the static linking of RTEMS with applications without forcing obligations and restrictions on users.

A problem for RTEMS is there are no copyleft licenses that are compatible with the deployment model of RTEMS. Thus, RTEMS Project has to reject any code that uses the GPL or LGPL, even though RTEMS has historically appeared to use the GPL itself – but with the exception for static linking, and also because an upstream GPL version 2 project could at any time switch to GPL version 3 and become totally unusable. In practice, RTEMS can only accept original code contributed under the RTEMS License and code that has a permissive license.

As stated above, the RTEMS Project has defined its preferred licenses. These allow generation of documentation and software from specification as well as allow end users to statically link with RTEMS and not incur obligations.

In some cases, RTEMS includes software from third-party projects. In those cases, the license is carefully evaluated to meet the project licensing goals. The RTEMS Project can only include software under licenses which follow these guidelines:

- 2- and 3-clause BSD, MIT, and other OSI-approved non-copyleft licenses that permit statically linking with the code of different licenses are acceptable.
- The historical RTEMS [License](#) is acceptable for software already in the tree. This software is being relicensed to BSD-2-Clause, rewritten, or removed.
- GPL licensed code is NOT acceptable, neither is LGPL.
- Software which is dual-licensed in a manner which prevents free use in commercial applications is not acceptable.
- Advertising obligations are not acceptable.

- Some license restrictions may be permissible. These will be considered on a case-by-case basis.

In practice, these guidelines are not hard to follow. Critically, they protect the freedom of the RTEMS source code and that of end users to select the license and distribution terms they prefer for their RTEMS-based application.

APPENDIX: CORE QUALIFICATION ARTIFACTS/DOCUMENTS

An effort at NASA has been performed to suggest a core set of artifacts (as defined by **BOTH** NASA NPR 7150.2B and DO-178B) that can be utilized by a mission as a baselined starting point for “pre-qualification” for (open-source) software that is intended to be utilized for flight purposes. This effort analyzed the overlap between NPR 7150.2B and DO-178B and highlighted a core set of artifacts to serve as a starting point for any open-source project. These artifacts were also cross-referenced with similar activities for other NASA flight software qualification efforts, such as the open-source Core Flight System (cFS). Along with the specific artifact, the intent of the artifact was also captured; in some cases open-source projects, such as RTEMS, are already meeting the intent of the artifacts with information simply needing organized and formalized. The table below lists the general category, artifact name, and its intent. Please note that this table does **NOT** represent all the required artifacts for qualification per the standards; instead, this table represents a subset of the most basic/core artifacts that form a strong foundation for a software engineering qualification effort.

Table 13.1: Table 1. Core Qualification Artifacts

Category	Artifact	Intent
Requirements	Software Requirements Specification (SRS) Requirements Management	The project shall document the software requirements. The project shall collect and manage changes to the software requirements. The project shall identify, initiate corrective actions, and track until closure inconsistencies among requirements, project plans, and software products.
	Requirements Test and Traceability Matrix	The project shall perform, document, and maintain bidirectional traceability between the software requirement and the higher-level requirement.
	Validation	The project shall perform validation to ensure that the software will perform as intended in the customer environment.
Design and Implementation	Software Development or Management Plan	A plan for how you will develop the software that you are intent upon developing and delivering. The Software Development Plan includes the objectives, standards and life cycle(s) to be used in the software development process. This plan should include: Standards: Identification of the Software Requirements Standards, Software Design Standards, and Software Code Standards for the project.
	Software Configuration Management Plan	To identify and control major software changes, ensure that change is being properly implemented, and report changes to any other personnel or clients who may have an interest.
	Implementation	The project shall implement the software design into software code. Executable Code to applicable tested software.
	Coding Standards Report	The project shall ensure that software coding methods, standards, and/or criteria are adhered to and verified.
	Version Description Document (VDD)	The project shall provide a Software Version Description document for each software release.
Testing and Software Assurance Activities	Software Test Plan	Document describing the testing scope and activities.
	Software Assurance/Testing Procedures	To define the techniques, procedures, and methodologies that will be used.
	Software Change Report / Problem Report	The project shall regularly hold reviews of software activities, status, and results with the project stakeholders and track issues to resolution.
244	Software Schedule	Milestones have schedule and schedule is updated accordingly. Chapter 13. Appendix: Core Qualification Artifacts/Documents
	Software Test Report / Verification	The project shall record, address, and track to closure the results of software verification activities.

In an effort to remain lightweight and sustainable for open-source projects, Table 1 above was condensed into a single artifact outline that encompasses the artifacts' intents. The idea is that this living qualification document will reside under RTEMS source control and be updated with additional detail accordingly. The artifact outline is as follows:

GLOSSARY

API

This term is an acronym for Application Programming Interface.

assembler language

The assembler language is a programming language which can be translated very easily into machine code and data. For this project assembler languages are restricted to languages accepted by the *GNU* assembler program for the target architectures.

C language

The C language for this project is defined in terms of *C11*.

C11

The standard ISO/IEC 9899:2011.

CCB

This term is an acronym for Change Control Board.

Doorstop

Doorstop is a requirements management tool.

EARS

This term is an acronym for Easy Approach to Requirements Syntax.

ELF

This term is an acronym for **Executable and Linkable Format**.

GCC

This term is an acronym for **GNU Compiler Collection**.

GNAT

GNAT is the *GNU* compiler for Ada, integrated into the *GCC*.

GNU

This term is an acronym for **GNU's Not Unix**.

interrupt service

An *interrupt service* consists of an *Interrupt Service Routine* which is called with a user provided argument upon reception of an interrupt service request. The routine is invoked in interrupt context. Interrupt service requests may have a priority and an affinity to a set of processors. An *interrupt service* is a *software component*.

Interrupt Service Routine

An ISR is invoked by the CPU to process a pending interrupt.

ISVV

This term is an acronym for Independent Software Verification and Validation.

ReqIF

This term is an acronym for [Requirements Interchange Format](#).

RTEMS

This term is an acronym for Real-Time Executive for Multiprocessor Systems.

software component

This term is defined by ECSS-E-ST-40C 3.2.28 as a “part of a software system”. For this project a *software component* shall be any of the following items and nothing else:

- *software unit*
- explicitly defined *ELF* symbol in a *source code* file
- *assembler language* data in a source code file
- *C language* object with static storage duration
- C language object with thread-local storage duration
- *thread*
- *interrupt service*
- collection of *software components* (this is a software architecture element)

Please note that explicitly defined ELF symbols and assembler language data are considered a software component only if they are defined in a *source code* file. For example, this rules out symbols and data generated as side-effects by the toolchain (compiler, assembler, linker) such as jump tables, linker trampolines, exception frame information, etc.

software product

The *software product* is the *RTEMS* real-time operating system.

software unit

This term is defined by ECSS-E-ST-40C 3.2.24 as a “separately compilable piece of source code”. For this project a *software unit* shall be any of the following items and nothing else:

- *assembler language* function in a *source code* file
- *C language* function (external and internal linkage)

A *software unit* is a *software component*.

source code

This project uses the *source code* definition of the [Linux Information Project](#): “Source code (also referred to as source or code) is the version of software as it is originally written (i.e., typed into a computer) by a human in plain text (i.e., human readable alphanumeric characters).”

task

This project uses the [thread definition of Wikipedia](#): “a thread of execution is the smallest sequence of programmed instructions that can be managed independently by a scheduler, which is typically a part of the operating system.”

It consists normally of a set of registers and a stack. The scheduler assigns processors to a subset of the ready tasks. The terms task and *thread* are synonym in RTEMS. The term task is

used throughout the Classic API, however, internally in the operating system implementation and the POSIX API the term *thread* is used.

A task is a software component.

thread

This term has the same meaning as *task*.

YAML

This term is an acronym for **YAML Ain't Markup Language**.

BIBLIOGRAPHY

- [Bra97] Scott Bradner. Key words for use in RFCs to Indicate Requirement Levels. BCP 14, RFC Editor, March 1997. <http://www.rfc-editor.org/rfc/rfc2119.txt>. URL: <http://www.rfc-editor.org/rfc/rfc2119.txt>.
- [BA14] Jace Browning and Robert Adams. Doorstop: Text-Based Requirements Management Using Version Control. *Journal of Software Engineering and Applications*, 7:187–194, 2014. URL: http://www.scirp.org/pdf/JSEA_2014032713545074.pdf.
- [ECS09] ECSS. *ECSS-E-ST-10-06C - Technical requirements specification*. European Cooperation for Space Standardization, 2009. URL: <https://ecss.nl/standard/ecss-e-st-10-06c-technical-requirements-specification/>.
- [MW10] Alistair Mavin and Philip Wilkinson. Big Ears (The Return of Easy Approach to Requirements Engineering). In *18th Requirements Engineering Conference*, 277–282. 11 2010. URL: https://www.researchgate.net/profile/Alistair_Mavin/publication/224195362_Big_Ears_The_Return_of_Easy_Approach_to_Requirements_Engineering/links/568ce39808ae197e426a075e/Big-Ears-The-Return-of-Easy-Approach-to-Requirements-Engineering.pdf, doi:10.1109/RE.2010.39.
- [MWGU16] Alistair Mavin, Philip Wilkinson, Sarah Gregory, and Eero Uusitalo. Listens Learned (8 Lessons Learned Applying EARS). In *24th International Requirements Engineering Conference*. September 2016. URL: https://www.researchgate.net/profile/Alistair_Mavin/publication/308970788_Listens_Learned_8_Lessons_Learned_Applying_EARS/links/5ab0d42caca2721710fe5017/Listens-Learned-8-Lessons-Learned-Applying-EARS.pdf, doi:10.1109/RE.2016.38.
- [MWHN09] Alistair Mavin, Philip Wilkinson, Adrian Harwood, and Mark Novak. Easy approach to requirements syntax (EARS). In *17th Requirements Engineering Conference*, 317–322. 10 2009. URL: https://www.researchgate.net/profile/Alistair_Mavin/publication/224079416_Easy_approach_to_requirements_syntax_EARS/links/568ce3bf08aeb488ea311990/Easy-approach-to-requirements-syntax-EARS.pdf, doi:10.1109/RE.2009.9.
- [Mot88] Motorola. *Real Time Executive Interface Definition*. Motorola Inc., Microcomputer Division and Software Components Group, Inc., January 1988. DRAFT 2.1. URL: https://ftp.rtems.org/pub/rtems/publications/RTEID-ORKID/RTEID-2.1/RTEID-2_1.pdf.

- [Uus12] Eero Uusitalo. *EARS quick reference sheet*. January 2012. URL: https://aaltodoc.aalto.fi/bitstream/handle/123456789/12861/D5_uusitalo_eero_2012.pdf.
- [VIT90] VITA. *Open Real-Time Kernel Interface Definition*. VITA, the VMEbus International Trade Association, August 1990. Draft 2.1. URL: https://ftp.rtems.org/pub/rtems/publications/RTEID-ORKID/ORKID-2.1/ORKID-2_1.pdf.
- [WB13] Karl Wieggers and Joy Beatty. *Software Requirements*. Microsoft Press, 3 edition, 2013. ISBN 0735679665, 9780735679665.

INDEX

A

API, **247**
assembler language, **247**

C

C language, **247**
C11, **247**
CCB, **247**

D

Doorstop, **247**

E

EARS, **247**
ELF, **247**

G

GCC, **247**
GNAT, **247**
GNU, **247**

I

interrupt service, **247**
Interrupt Service Routine, **247**
ISVV, **248**

R

ReqIF, **248**
RTEMS, **248**

S

software component, **248**
software product, **248**
software unit, **248**
source code, **248**

T

task, **248**
thread, **249**

Y

YAML, **249**