

LINUX 2.6.33.7.2-RT30 ON X86

© Copyright Dedicated Systems Experts NV. All rights reserved, no part of the contents of this document may be reproduced or transmitted in any form or by any means without the written permission of Dedicated Systems Experts NV, Diepenbeemd 5, B-1650 Beersel, Belgium.

Authors: Luc Perneel (1, 2), Hasan Fayyad-Kazan (2) and Martin Timmerman (1, 2, 3)

1: Dedicated Systems Experts, 2: VUB-Brussels, 3: RMA-Brussels

Disclaimer

Although all care has been taken to obtain correct information and accurate test results, Dedicated Systems Experts, VUB-Brussels, RMA-Brussels and the authors cannot be liable for any incidental or consequential damages (including damages for loss of business, profits or the like) arising out of the use of the information provided in this report, even if these organisations and authors have been advised of the possibility of such damages.

<http://www.dedicated-systems.com>

E-mail: info@dedicated-systems.com

EVALUATION REPORT LICENSE

This is a legal agreement between you (the downloader of this document) and/or your company and the company DEDICATED SYSTEMS EXPERTS NV, Diepenbeemd 5, B-1650 Beersel, Belgium.
It is not possible to download this document without registering and accepting this agreement on-line.

1. **GRANT.** Subject to the provisions contained herein, Dedicated Systems Experts hereby grants you a non-exclusive license to use its accompanying proprietary evaluation report for projects where you or your company are involved as major contractor or subcontractor. You are not entitled to support or telephone assistance in connection with this license.
2. **PRODUCT.** Dedicated Systems Experts shall furnish the evaluation report to you electronically via Internet. This license does not grant you any right to any enhancement or update to the document.
3. **TITLE.** Title, ownership rights, and intellectual property rights in and to the document shall remain in Dedicated Systems Experts and/or its suppliers or evaluated product manufacturers. The copyright laws of Belgium and all international copyright treaties protect the documents.
4. **CONTENT.** Title, ownership rights, and an intellectual property right in and to the content accessed through the document is the property of the applicable content owner and may be protected by applicable copyright or other law. This License gives you no rights to such content.
5. **YOU CANNOT:**
 - You cannot, make (or allow anyone else make) copies, whether digital, printed, photographic or others, except for backup reasons. The number of copies should be limited to 2. The copies should be exact replicates of the original (in paper or electronic format) with all copyright notices and logos.
 - You cannot, place (or allow anyone else place) the evaluation report on an electronic board or other form of on line service without authorisation.
6. **INDEMNIFICATION.** You agree to indemnify and hold harmless Dedicated Systems Experts against any damages or liability of any kind arising from any use of this product other than the permitted uses specified in this agreement.
7. **DISCLAIMER OF WARRANTY.** All documents published by Dedicated Systems Experts on the World Wide Web Server or by any other means are provided "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. This disclaimer of warranty constitutes an essential part of the agreement.
8. **LIMITATION OF LIABILITY.** Neither Dedicated Systems Experts nor any of its directors, employees, partners or agents shall, under any circumstances, be liable to any person for any special, incidental, indirect or consequential damages, including, without limitation, damages resulting from use of OR RELIANCE ON the INFORMATION presented, loss of profits or revenues or costs of replacement goods, even if informed in advance of the possibility of such damages.
9. **ACCURACY OF INFORMATION.** Every effort has been made to ensure the accuracy of the information presented herein. However Dedicated Systems Experts assumes no responsibility for the accuracy of the information. Product information is subject to change without notice. Changes, if any, will be incorporated in new editions of these publications. Dedicated Systems Experts may make improvements and/or changes in the products and/or the programs described in these publications at any time without notice. Mention of non-Dedicated Systems Experts products or services is for information purposes only and constitutes neither an endorsement nor a recommendation.
10. **JURISDICTION.** In case of any problems, the court of BRUSSELS-BELGIUM will have exclusive jurisdiction.

Agreed by downloading the document via the internet.

1	Document Intention.....	6
1.1	Purpose and scope	6
1.2	Document issue: the 2.9 framework	6
1.3	Conventions	6
	Related documents	7
2	Introduction	8
2.1	Overview	8
2.2	Evaluated (RTOS) product.....	9
2.2.1	Software	9
2.2.2	Hardware	9
3	Evaluation results summary.....	10
3.1	Positive points	10
3.2	Negative points	10
3.3	Ratings	11
4	Test Results	12
4.1	Calibration system test (CAL)	12
4.1.1	Tracing overhead (CAL-P-TRC).....	12
4.1.2	CPU power (CAL-P-CPU)	13
4.2	Clock tests (CLK)	15
4.2.1	Operating system clock setting (CLK-B-CFG)	15
4.2.2	Clock tick processing duration (CLK-P-DUR).....	15
4.3	Thread tests (THR)	17
4.3.1	Thread creation behaviour (THR-B-NEW)	17
4.3.2	Round robin behaviour (THR-B-RR)	18
4.3.3	Thread switch latency between same priority threads (THR-P-SLS).....	19
4.3.4	Thread creation and deletion time (THR-P-NEW).....	22
4.4	Semaphore tests (SEM).....	26
4.4.1	Semaphore locking test mechanism (SEM-B-LCK)	26
4.4.2	Semaphore releasing mechanism (SEM-B-REL).....	26
4.4.3	Time needed to create and delete a semaphore (SEM-P-NEW).....	26
4.4.4	Test acquire-release timings: contention case (SEM-P-ARN)	29
4.4.5	Test acquire-release timings: contention case (SEM-P-ARC)	31
4.5	Mutex tests (MUT).....	34
4.5.1	Priority inversion avoidance mechanism (MUT-B-ARC)	34
4.5.2	Mutex acquire-release timings: contention case (MUT-P-ARC)	34
4.5.3	Mutex acquire-release timings: no-contention case (MUT-P-ARN)	36
4.6	Interrupt tests (IRQ)	38
4.6.1	Interrupt latency (IRQ_P_LAT).....	38
4.6.2	Interrupt dispatch latency (IRQ_P_DLT)	39
4.6.3	Interrupt to thread latency (IRQ_P_TLT).....	40
4.6.4	Maximum sustained interrupt frequency (IRQ_S_SUS).....	41

4.7	Memory tests	41
4.7.1	Memory leak test (MEM_B_LEK)	42
5	Support	43
6	Appendix B: Acronyms	44

SAMPLE

1 Document Intention

1.1 Purpose and scope

This document presents the quantitative evaluation results of the real-time **Linux** operating system (**Linux** with its real-time patches). The testing results of this operating system employed on an x86 processors can be found on our website. (www.dedicated-systems.com)

The layout of this report follows the one depicted in “The OS evaluation template” [Doc. 4]. The test specifications can be found in “The evaluation test report definition” [Doc. 3]. See section 0 of this document for more detailed references. These documents have to be seen as an integral part of this report!

Due to the tightly coupling between these documents, the framework version of “The evaluation test report definition” has to match the framework version of this evaluation report (which is 2.9). More information about the documents and tests versions together with their corresponding relation between both can be found in “The evaluation framework”, see [Doc. 1] in section 0 of this document.

The generic test code used to perform these tests can be downloaded on our website by using the link in the related documents section.

1.2 Document issue: the 2.9 framework

This document shows the test results in the scope of the evaluation framework 2.9.

1.3 Conventions

Throughout this document, we use certain typographical conventions to distinguish technical terms. Our used conventions are the following:

- ❖ ***Bold Italic*** for OS Objects
- ❖ **Bold** for Libraries, packets, directories, software, OSs...
- ❖ `Courier New` for system calls (APIs...)

2 Introduction

This chapter talks about: 1) the OS that we are going to test and evaluate, 2) the real time patch integrated in this OS to achieve some real time performance and behaviour tests, 3) the library used for interaction between the testing applications and the kernel, 4) the hardware on which the under testing OS will be employed.

2.1 Overview

The evaluation project started in 1995 and as such accumulates a long experience with different (RT) OS. Today more and more embedded systems are equipped with **Linux** solutions using more or less real-time patches. Different vendors like MontaVista, Windriver, and Lynuxworks have now **Linux** variants in their product portfolio.

Since the kernel version 2.4, a lot of improvements regarding real-time behaviour found their way into the standard "**Vanilla**" kernel. There is a well maintained real-time patch available (both have their origins from Ingo Molnar) called **RT_PREEMPT** patch. Remark that some real-time features (like priority-inheritance **mutexes**, introduced in version 2.6.18) are already in the **Vanilla** kernel.

We believed that it is the time to test this kernel by our standard real-time behaviour evaluation framework and find out how well it behaves.

For this evaluation, we used the standard **glibc** library as the **uClibc** package does not include yet the **Native POSIX Thread Library (NPTL)**. Moreover, **uClibc** also does not use **futexes** which means that it also does not have any support for priority inheritance (which must be available when considering real-time behaviour). Further **uClibc** uses internal protection systems (**mutex**, **semaphores**) and signals for the **pthread POSIX** layer, which behaves differently while compared to the usage of direct **NPTL** calls. From a real-time point of view, using **uClibc** in its current form makes the kernel real-time support unavailable in user space. It has to be said that there is an active **NPTL** branch in the **uClibc** code base. Therefore we suspect that it is only a matter of time before it will become available in the official releases.

It is a pity that the **uClibc** wagon is not yet on the **NPTL** rail. Using the **buildroot**, **uClibc**, and **busybox** combo makes it easier to have an embedded **Linux** platform with a small storage footprint. But without the **NPTL** support, real-time applications cannot be used in user space, unless you use direct system calls to the kernel.

Remark that the **RT_PREEMPT** patch degrades throughput performance which means that it should be used only when your project has low latency requirements. This is normal and a fundamental rule in real-time software: latency improvements have a negative impact on throughput and vice versa. Some quick measurements using an NFS mount stressing network and disk showed a negative throughput impact between 5 and 10% by enabling the **RT_PREEMPT** patch!

2.2 Evaluated (RTOS) product

2.2.1 Software

The operating system OS that will be evaluated is **Vanilla Linux** 2.6.33.7 with real-time patch v30. This RT patch was the latest version officially released by OSADL (the Open Source Automation Development Lab) on December 21, 2010. Being as OSADL's latest stable release was our main reason for testing this version. The RT patches can be found at <http://www.kernel.org/pub/linux/kernel/projects/rt/>.

The evaluation of this kernel version (2.6.33.7.2-rt30) was performed using several performance and behaviour tests. The testing results are applicable only to this version as other versions may have other significant performance figures and behaviour.

The library used between the testing applications and the kernel is the **glibc** version 2.11.1 as mentioned before. This interfacing library is important because user applications (when using POSIX calls) can access the real-time features of the kernel only if this library supports them. Otherwise, direct system calls in user space applications are needed.

2.2.2 Hardware

The operating system was tested on our standard x86 Pentium MMX evaluation platform running at 200MHz with 128MB RAM. So we can compare the results with other RTOS tested on the same platform.

This platform is already pretty old (from 1997 to be correct), but its advantage is that it makes tests comparable as we have tested all RTOS for more than a decade on this platform. Also there is no impact from BIOS interrupts (motherboard/BIOS is too old).

3 Evaluation results summary

Remember that the tested and evaluated product is **Vanilla Linux 26.33.7** with **RT_PREEMPT** patch v30. If correctly used and configured, the **RT_PREEMPT Linux** system has the internals to provide some real-time characteristics.

Compared with the traditional RTOS that supports also memory protection between processes, the worst case latencies in **Linux RT_PREEMPT** are still around 5 to 10 times slower (depending on the RTOS you compare with). Our study and measurements show the latencies are bound and therefore this **Linux** version may be labelled Real-Time.

TAKE CARE: Using a wrong driver or wrong configuration can destroy real-time behaviour. You need to follow the detailed rules described in the relevant document (Doc 5).

3.1 Positive points








- No license fees
- Source code available
- Extensible

3.2 Negative points

- The real-time characteristics of the OS are present only when everything is configured and built correctly (and not for all drivers)
- GPL license is not completely free...
- Setting up a complete embedded target from scratch is a daunting task.
- **uClibc**, which is used a lot in embedded systems, does not have currently NTP support and as such cannot provide real-time characteristics to the user level. Thus, **glibc** should be used.

3.3 Ratings

For a description of the ratings, see [Doc. 3].

RTOS Architecture	0		10
OS Documentation	0		10
OS Configuration	0		10
Internet Components	0		10
Development Tools	0		10
Installation and BSP	0		10
Test Results	0		10
Support	0	N.A.	10

Although [Doc. 3] gives a description of the ratings, comparison with other reports on other OS should help you understand the scoring.

SAMPLE