



▶ Polycom RMX™ 2000/4000 Open Source Third Party Licenses

Trademark Information

Polycom®, the Polycom “Triangles” logo, and the names and marks associated with Polycom’s products are trademarks and/or service marks of Polycom, Inc., and are registered and/or common-law marks in the United States and various other countries.

All other trademarks are the property of their respective owners.

Patent Information

The accompanying product is protected by one or more U.S. and foreign patents and/or pending patent applications held by Polycom, Inc.

Portions, aspects and/or features of this product are protected under United States Patent Law in accordance with the claims of United States Patent No: US 6,300,973; US 6,492,216; US 6,496,216; US 6,757,005; US 6,760,750; US 7,054,620; US 7,085,243; US 7,113,200; US 7,269,252; US 7,310,320.

PATENT PENDING

© 2009 Polycom, Inc. All rights reserved.

Polycom, Inc.
4750 Willow Road
Pleasanton, CA 94588-2708
USA

No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of Polycom, Inc. Under the law, reproducing includes translating into another language or format.

As between the parties, Polycom, Inc., retains title to and ownership of all proprietary rights with respect to the software contained within its products. The software is protected by United States copyright laws and international treaty provision. Therefore, you must treat the software like any other copyrighted material (e.g., a book or sound recording).

Every effort has been made to ensure that the information in this manual is accurate. Polycom, Inc., is not responsible for printing or clerical errors. Information in this document is subject to change without notice.

Regulatory Notices

United States Federal Communication Commission (FCC)

Part 15: Class A Statement. This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. Test limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manuals, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his or her own expense.

Part 68: Network Registration Number. This equipment is registered with the FCC in accordance with Part 68 of the FCC Rules. This equipment is identified by the FCC registration number.

If requested, the FCC registration Number and REN must be provided to the telephone company.

Any repairs to this equipment must be carried out by Polycom Inc. or our designated agent. This stipulation is required by the FCC and applies during and after the warranty period.

United States Safety Construction Details:

- All connections are indoor only.
- Unit is intended for RESTRICTED ACCESS LOCATION.
- Unit is to be installed in accordance with the National Electrical Code.
- The branch circuit overcurrent protection shall be rated 20 A for the AC system.
- This equipment has a maximum operating ambient of 40°C, the ambient temperature in the rack shall not exceed this temperature.

To eliminate the risk of battery explosion, the battery should not be replaced by an incorrect type. Dispose of used batteries according to their instructions.

CE Mark R&TTE Directive

Polycom Inc., declares that the Polycom RMX™ 2000 is in conformity with the following relevant harmonized standards:

EN 60950-1:2001

EN 55022: 1998+A1:2000+A2:2003 class A

EN 300 386 V1.3.3: 2005

Following the provisions of the Council Directive 1999/CE on radio and telecommunication terminal equipment and the recognition of its conformity.

Canadian Department of Communications

This Class [A] digital apparatus complies with Canadian ICES-003.

Notice: The Industry Canada label identifies certified equipment. This certification means that the equipment meets telecommunication network protective, operational and safety requirements as prescribed in the appropriate Terminal Equipment Technical Requirements document(s). The Department does not guarantee the equipment will operate to the user's satisfaction.

Before installing this equipment, users should ensure that it is permissible to be connected to the facilities of the local telecommunications company. The equipment must also be installed using an acceptable method of connection. The customer should be aware that compliance with the above conditions may not prevent degradation of service in some situations. Repairs to certified equipment malfunctions, may give the telecommunications company causes to request the user to disconnect the equipment.

Users should ensure for their own protection that the electrical ground connections of the power utility, telephone lines and internal metallic water pipe system, if present, are connected together. This precaution may be particularly important in rural areas.

Caution: Users should not attempt to make such connections themselves, but should contact the appropriate electric inspection authority, or electrician, as appropriate.

Regulatory Notices

RMX 2000: Chinese Communication Certificate

声 明

此为 A 级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

Singapore Certificate

RMX 2000 complies with IDA standards G0916-07

Table of Contents

Third Party Software Licenses 1-1

 Note 1 - Apache 1-18

 Note 2 - BSD and ISC 1-23

 Note 3 - MIT 1-24

 Note 4 - zlib 1-25

 Note 5 - BusyBox 1-25

 Note 6 - glibc 1-26

 Note 7 - ntp 1-31

 Note 8 - openSSH 1-32

 Note 9 - openSSL 1-39

 Note 10 - IPERF 1-42

 Note 11 - DMalloc 1-43

Modified Code Samples 2-1

 apache: httpd-2.2.11.FIPS_mode.patch 2-1

 apache: httpd-2.2.11.pragma.patch 2-3

 bind-9.3.1-cross-compile.diff 2-5

 busybox-1.00-new-passwd-dir.diff 2-5

 busybox-1.00-switch_root.patch 2-6

 busybox-1.00-vi-backspace.diff 2-12

 cfgsh.patch 2-14

 dhcpcd-1.3.22-pl4.patch 2-15

 dhcpcd-2.0.0_rpl_malloc.patch 2-22

 gdb-6.3-cppflags-fix.diff 2-22

 gdb-6.3-ldflags-fix.diff 2-23

 gdb-6.4-termcap-configure.patch 2-25

 genext2fs_1.3-7.1.diff 2-25

 grub-0.95.max8MB_kernel.patch 2-58

 grub-0.95-no-A20-keyboard-controller.patch 2-59

 httpd-2.0.54-return_Pragma_for_not_modified.patch 2-60

 httpd-2.0.54-skip-storage-on-post.patch 2-60

 httpd.2.0.61.pragma.300.patch 2-61

 lm_sensors-2.10.0-01-EXLDFLAGS.patch 2-62

 module-init-tools-3.2-pre1-uclibc.diff 2-62

module-init-tools-oldtoolsname.diff	2-64
nano-1.2.4_cross_compile.diff	2-65
generic-tic-xcompile.diff	2-65
ntp-4.2.0-linux-2.6.14.patch	2-66
generic-configure-ac.patch	2-163
openssl-0.9.7g-configure.diff	2-176
portmap-5beta-compilation_fixes-3.patch	2-180
portmap-5beta-disable-tcp_wrapper.patch	2-203
portmap-5beta-glibc_errno_fix-1.patch	2-203
proftpd-1.3.0rc1-ptx.diff	2-204
smartmontools-5.36-01-snprintf-cross.patch	2-205
squashfs2.2-r2-Makefile.patch	2-205
generic-include-fix.diff	2-206

Third Party Software Licenses

This document provides the copyright statements for third party software products that are part of the Polycom RMX™ 2000/4000.

Table 1-1 *Third Party Licenses*

Software Source Name	License Type	License Conditions
apache (2.2.11)	Apache License	<p>Version 2.0, January 2004 http://www.apache.org/licenses/</p> <p>For additional copyright information, see "Note 1 - Apache" on page 1-18. Changes to the code can be viewed in:</p> <ul style="list-style-type: none">• "apache: httpd-2.2.11.FIPS_mode.patch" on page 2-1.• "apache: httpd-2.2.11.pragma.patch" on page 2-3

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>bind</i>	ISC	Copyright (C) 2004, 2005 Internet Systems Consortium, Inc. ("ISC") Copyright (C) 1996-2003 Internet Software Consortium. For additional copyright information, see "Note 2 - BSD and ISC" on page 1-23.
<i>busybox-1.00</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . For additional copyright information, see "Note 5 - BusyBox" on page 1-25. Changes to the code can be viewed in "busybox-1.00-new-passwd-dir.diff" on page 2-5, "busybox-1.00-switch_root.patch" on page 2-6 and "busybox-1.00-vi-backspace.diff" on page 2-12.
<i>ccmalloc</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>cfgsh</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "cfgsh.patch" on page 2-14.
<i>dhcpcd 1.3.22</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "dhcpcd-1.3.22-pl4.patch" on page 2-15 and "bind-9.3.1-cross-compile.diff" on page 2-5.
<i>Dmalloc</i>	Not applicable	Copyright 1992 to 2007 by Gray Watson.
<i>Electric fence</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>e2fsprogs</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>ethtool-3</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>expat</i>	MIT License	Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd and Clark Cooper. Copyright (c) 2001, 2002, 2003 Expat maintainers. For license terms, see "Note 3 - MIT" on page 1-24.
<i>expat</i>	MIT License	Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd and Clark Cooper. Copyright (c) 2001, 2002, 2003 Expat maintainers. For additional copyright information, see "Note 3 - MIT" on page 1-24.
<i>gcc-3.3.3</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>gdb-6.3</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "gdb-6.3-ldflags-fix.diff" on page 2-23 and "gdb-6.4-termcap-configure.patch" on page 2-25.
<i>genext2fs-1.3</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "genext2fs_1.3-7.1.diff" on page 2-25.
<i>glibc-2.3.5</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . For additional copyrights and Licenses, see "Note 6 - glibc" on page 1-26.

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>grub</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "grub-0.95.max8MB_kernel.patch" on page 2-58 and "grub-0.95-no-A20-keyboard-controller.patch" on page 2-59.
<i>hotplug-2004_03_29</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>IPERF</i>	Not applicable	Copyright (c) 1999-2007, The Board of Trustees of the University of Illinois All Rights Reserved. http://www.ncsa.uiuc.edu For additional copyright information, see "Note 10 - IPERF" on page 1-42
<i>ipkg-build</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>kexec-1.101</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>libpcap-0.8.3</i>	BSD	Copyright (c) 1994, 1995, 1996, 1997, 1998 The Regents of the University of California. For additional copyright information, see "Note 2 - BSD and ISC" on page 1-23.
<i>Libxml2 2.6.12</i>	MIT License	Copyright (C) 1998-2003 Daniel Veillard. For additional copyright information, see "Note 3 - MIT" on page 1-24.
<i>Linux-2.4.30</i>	GPL v2 only	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright (C) Linus Torvalds. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "httpd-2.0.54-return_Pragma_for_not_modified.patch" on page 2-60, "httpd-2.0.54-skip-storage-on-post.patch" on page 2-60 and "httpd.2.0.61.pragma.300.patch" on page 2-61.

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>lm_sensors-2.10.0</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "lm_sensors-2.10.0-01-EXLDFLAGS.patch" on page 2-62.
<i>module-init-tools-3.2-pre1</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "module-init-tools-3.2-pre1-uclibc.diff" on page 2-62 and "module-init-tools-oldtoolsname.diff" on page 2-64.
<i>mtd-20040621-1</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>nano-1.2.4</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "nano-1.2.4_cross_compile.diff" on page 2-65.
<i>ncurses 5.3</i>	MIT License	Copyright (c) 1998-2004,2005 Free Software Foundation, Inc. For additional copyright information, see "Note 3 - MIT" on page 1-24. Changes to the code can be viewed in "generic-tic-xcompile.diff" on page 2-65.
<i>netkit-ftp-0.17</i>	BSD	Copyright (c) 1985, 1989 Regents of the University of California. For additional copyright information, see "Note 2 - BSD and ISC" on page 1-23.
<i>nfs-utils</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>ntp-4.2.0</i>	GPL v2.0 and above	<p>Network Time Protocol Version 4 Distribution. Copyright (c) David L. Mills 1992-2003. See "Note 7 - ntp" on page 1-31. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html. Copyright (C) 1989, 1991 Free Software Foundation, Inc. Changes to the code can be viewed in "ntp-4.2.0-linux-2.6.14.patch" on page 2-66.</p>
<i>OpenIPMI-1.4.19</i>	GPL v2.0 and above	<p>Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html.</p>
<i>openssh</i>	BSD	<p>Copyrights of the various components are detailed in "Note 8 - openSSH" on page 1-32. Distributed under the BSD License. For additional copyright information, see "Note 2 - BSD and ISC" on page 1-23. Changes to the code can be viewed in "generic-configure-ac.patch" on page 2-163.</p>

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>openssl</i>	BSD	<p>Copyright (c) 1998-2007 The OpenSSL Project. All rights reserved.</p> <p>Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.</p> <p>OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. Both licenses are BSD-style Open Source licenses.</p> <p>For additional copyright information, see "Note 9 - openssl" on page 1-39.</p> <p>Changes to the code can be viewed in "openssl-0.9.7g-configure.diff" on page 2-176.</p>
<i>popt-1.7</i>	GPL v2.0 and above	<p>Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html.</p>

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>portmap</i>	BSD	Copyright (c) 1990 The Regents of the University of California. For additional copyright information, see "Note 2 - BSD and ISC" on page 1-23. Changes to the code can be viewed in " portmap-5beta-compilation_fixes-3.patch " on page 2-180, " portmap-5beta-disable-tcp_wrapper.patch " on page 2-203, " portmap-5beta-glibc_errno_fix-1.patch " on page 2-203.
<i>proftpd</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in " proftpd-1.3.0rc1-ptx.diff " on page 2-204.
<i>readline-5.0</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>shorewall-3.0.7</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>smartmontools-5.36</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "smartmontools-5.36-01-sprintf-cross.patch" on page 2-205.
<i>squashfs</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "squashfs2.2-r2-Makefile.patch" on page 2-205.

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>strace-4.5.7</i>	GPL v2.0 and above	<p>Copyright (c) 1991, 1992 Paul Kranenburg <pk@cs.few.eur.nl> Copyright (c) 1993 Branko Lankester <branko@hacktic.nl> Copyright (c) 1993 Ulrich Pegelow <pegelow@moorea.uni-muenster.de> Copyright (c) 1995, 1996 Michael Elizabeth Chastain <mec@duracef.shout.net> Copyright (c) 1993, 1994, 1995, 1996 Rick Sladkey <jrs@world.std.com> Copyright (C) 1998-2001 Wichert Akkerman <wakkerma@deephackmode.org> All rights reserved. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html.</p>
<i>sysfsutils-1.3.0</i>	GPL v2.0 and above and LGPL v2.1	<p>Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html.</p>

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>sysutils-0.1.0</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>tar-1.17</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>tcpdump-3.8</i>	BSD	Copyright (c) 1998-2006 The TCPDUMP project. Original code by Hannes Gredler hannes@juniper.net See "Note 2 - BSD and ISC" on page 1-23.
<i>termcap-1.3.1</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html . Changes to the code can be viewed in "generic-include-fix.diff" on page 2-206.

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>U-Boot</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>utelnetsd</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>util-linux-2.12</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>valgrind</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>vconfig 1.8</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .

Table 1-1 Third Party Licenses (Continued)

Software Source Name	License Type	License Conditions
<i>vlan</i>	GPL v2.0 and above	The MAC vlan primarily the work of, and is copy-righted (GPL) by: Copyright (C) 2001 Alex Zeffertt, Cambridge Broadband Ltd, ajz@cambridgebroadband.com Copyright (C) 1989, 1991 Free Software Foundation, Inc. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>wget-1.9.1</i>	GPL v2.0 and above	Copyright (C) 1989, 1991 Free Software Foundation, Inc. Copyright 2007 as modified by Polycom. Distributed under the GPL License V. 2, which may be found at WWW.gnu.org/copyleft/gpl.html .
<i>zlib-1.2.1</i>	Zlib	Copyright (C) 1995-2003 Jean-loup Gailly and Mark Adler. For additional copyright information, see "Note 4 - zlib" on page 1-25.

The license agreements can be found in the following sites:

<http://httpd.apache.org/>

<http://www.opensource.org/licenses/bsd-license.php>

<http://www.gnu.org/copyleft/gpl.html>

<http://www.isc.org/index.pl>

<http://www.gnu.org/licenses/lgpl.txt>

<http://www.opensource.org/licenses/mit-license.html>

<http://www.opensource.org/licenses/zlib-license.php>



Modified open source code that is under GPL licenses is available upon request from Polycom Support.

Note 1 - Apache

Copyright 2004. Apache License

Version 2.0, January 2004

Licensed under the Apache License, Version 2.0 (the "License"); the code used in the RMX 2000 distribution has been modified by Polycom;

You may not use the Apache software except in compliance with the License. You may obtain a copy of the License at:

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Apache HTTP Server Subcomponents:

The Apache HTTP Server includes a number of subcomponents with separate copyright notices and license terms. Your use of the source code for these subcomponents is subject to the terms and conditions of the following licenses:

- `mod_mime_magic`: MIME type lookup via file magic numbers
Copyright (c) 1996-1997 Cisco Systems, Inc.
 - Some of this code is derived from the free version of the file command originally posted to comp.sources.unix.
Copyright (c) Ian F. Darwin, 1987. Written by Ian F. Darwin.
Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it freely, subject to the following restrictions:

- The author is not responsible for the consequences of use of this software, no matter how awful, even if they arise from flaws in it.
 - The origin of this software must not be misrepresented, either by explicit claim or by omission. Since few users ever read sources, credits must appear in the documentation.
 - Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. Since few users ever read sources, credits must appear in the documentation.
- For the `src/modules/standard/mod_imap.c` component:
"macmartinized" polygon code copyright 1992 by Eric Haines, erich@eye.com
 - For the `src/main/util_md5.c` component:
 - NCSA HTTPd Server
Copyright (C) 1995, Board of Trustees of the University of Illinois
 - md5.c: NCSA HTTPd code which uses the md5c.c RSA Code Original Code Copyright (C) 1994, Jeff Hostetler, Spyglass, Inc.
 - Portions of Content-MD5 code
Copyright (C) 1993, 1994 by Carnegie Mellon University (see below)
 - Portions of Content-MD5 code
Copyright (C) 1991 Bell Communications Research, Inc. (Bellcore) .
BELLCORE MAKES NO REPRESENTATIONS ABOUT THE ACCURACY OR SUITABILITY OF THIS MATERIAL FOR ANY PURPOSE. IT IS PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES.
 - Portions extracted from mpack, John G. Myers - jgm+@cmu.edu
 - Content-MD5 Code contributed by Martin Hamilton (martin@net.lut.ac.uk)these portions extracted from mpack, John G. Myers - jgm+@cmu.edu
Copyright (C) 1993,1994 by Carnegie Mellon University. All Rights Reserved.
CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF

MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

- For the `src/include/ap_md5.h` component:
Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.
- For the `src/ap/ap_md5c.c` component:
This work is derived from material Copyright RSA Data Security, Inc.
Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

The RSA copyright statement and Licence for that original material is included below. This is followed by the Apache copyright statement and licence for the modifications made to that material.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

The `p_MD5Encode()` routine uses much code obtained from the FreeBSD 3.0

— `MD5 crypt()` function, which is licenced as follows:
"THE BEER-WARE LICENSE" (Revision 42):
<phk@login.dknet.dk> wrote this file.

- For the `src/lib/sdbm` component:

sdbm - ndbm work-alike hashed database library

based on Per-Aake Larson's Dynamic Hashing algorithms. BIT 18 (1978).

Author: oz@nexus.yorku.ca

Status: public domain.

- For the src/lib/expat-lite component:

Copyright (c) 1998, 1999 James Clark. Expat is subject to the Mozilla Public License Version 1.1. Alternatively you may use expat under the GNU General Public License instead.

- For the src/regex component:

Copyright 1992, 1993, 1994 Henry Spencer. All rights reserved.

This software is not subject to any license of the American Telephone and Telegraph Company or of the Regents of the University of California.

Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it, subject to the following restrictions:

- The author is not responsible for the consequences of use of this software, no matter how awful, even if they arise from flaws in it.
 - The origin of this software must not be misrepresented, either by explicit claim or by omission. Since few users ever read sources, credits must appear in the documentation.
 - Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. Since few users ever read sources, credits must appear in the documentation.
- For the src/test/zb.c component:
ZeusBench V1.01. Copyright (C) Zeus Technology Limited 1996.
This program may be used and copied freely providing this copyright notice is not removed.
This software is provided "as is" and any express or implied warranties, including but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall Zeus Technology Ltd. be liable for any direct, indirect, incidental, special,

exemplary, or consequential damaged (including, but not limited to, procurement of substitute good or services; loss of use, data, or profits; or business interruption) however caused and on theory of liability. Whether in contract, strict liability or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

Written by Adam Twiss (adam@zeus.co.uk). March 1996.

Thanks to the following people for their input:

Mike Belshe (mbelshe@netscape.com)

Michael Campanella (campanella@stevms.enet.dec.com)

- For the expat xml parser component:

Copyright(c)1998,1999,2000ThaiOpenSourceSoftwareCenterLtd
and Clark Cooper

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Note 2 - BSD and ISC

Copyright (c) 2007. All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Polycom nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

bind-9.3.1

Copyright (C) 2004, 2005 Internet Systems Consortium, Inc. ("ISC")

Copyright (C) 1996-2003 Internet Software Consortium.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

COPYRIGHT, v 1.6.2.2.8.3 2005/01/10 23:51:37 marka Exp
Portions Copyright (C) 1996-2001 Nominum, Inc.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NOMINUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL NOMINUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Note 3 - MIT

Copyright (c) 2007.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Note 4 - zlib

Copyright (C) 1995-2003 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

- 1** The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
- 2** Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
- 3** This notice may not be removed or altered from any source distribution.

Note 5 - BusyBox

BusyBox is distributed under the GNU General Public License. A copy of the GNU General Public License may be found at <http://www.gnu.org/copyleft/gpl.html>.

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

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

Note 6 - glibc

glibc is distributed with the copying permission notices for various files in the GNU C Library distribution that have copyright owners other than the Free Software Foundation. These notices all require that a copy of the notice be included in the accompanying documentation and be distributed with binary distributions of the code, so be sure to include this file along with any binary distributions derived from the GNU C Library.

- All code incorporated from 4.4 BSD is distributed under the following license:

Copyright (C) 1991 Regents of the University of California. All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS

FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- The DNS resolver code, taken from BIND 4.9.5, is copyrighted both by UC Berkeley and by Digital Equipment Corporation. The DEC portions are under the following license:

Portions Copyright (C) 1993 by Digital Equipment Corporation.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Digital Equipment Corporation not be used in advertising or publicity pertaining to distribution of the document or software without specific, written prior permission.

THE SOFTWARE IS PROVIDED ``AS IS" AND DIGITAL EQUIPMENT CORP. DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL DIGITAL EQUIPMENT CORPORATION BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

- The Sun RPC support (from rpcsrc-4.0) is covered by the following license:

Copyright (C) 1984, Sun Microsystems, Inc.

Sun RPC is a product of Sun Microsystems, Inc. and is provided for unrestricted use provided that this legend is included on all tape media and as a part of the software program in whole or part. Users

may copy or modify Sun RPC without charge, but are not authorized to license or distribute it to anyone else except as part of a product or program developed by the user.

SUN RPC IS PROVIDED AS IS WITH NO WARRANTIES OF ANY KIND INCLUDING THE WARRANTIES OF DESIGN, MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE.

Sun RPC is provided with no support and without any obligation on the part of Sun Microsystems, Inc. to assist in its use, correction, modification or enhancement.

SUN MICROSYSTEMS, INC. SHALL HAVE NO LIABILITY WITH RESPECT TO THE INFRINGEMENT OF COPYRIGHTS, TRADE SECRETS OR ANY PATENTS BY SUN RPC OR ANY PART THEREOF.

In no event will Sun Microsystems, Inc. be liable for any lost revenue or profits or other special, indirect and consequential damages, even if Sun has been advised of the possibility of such damages.

- The following CMU license covers some of the support code for Mach, derived from Mach 3.0:

Mach Operating System

Copyright (C) 1991,1990,1989 Carnegie Mellon University All Rights Reserved.

Permission to use, copy, modify and distribute this software and its documentation is hereby granted, provided that both the copyright notice and this permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS ``AS IS" CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Carnegie Mellon requests users of this software to return to
Software Distribution Coordinator,
School of Computer Science
Carnegie Mellon University
Pittsburgh PA 15213-3890
or

Software.Distribution@CS.CMU.EDU any improvements or extensions that they make and grant Carnegie Mellon the rights to redistribute these changes.

- The file if_ppp.h is under the following CMU license:
Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
 - Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
 - Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
 - Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY CARNEGIE MELLON UNIVERSITY AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE UNIVERSITY OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- The following license covers the files from Intel's "Highly Optimized Mathematical Functions for Itanium" collection:
Intel License Agreement
Copyright (c) 2000, Intel Corporation. All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Intel Corporation may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL INTEL OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- The files `inet/getnameinfo.c` and `sysdeps/posix/getaddrinfo.c` are copyright (C) by Craig Metz and are distributed under the following license:

The Inner Net License, Version 2.00

The author(s) grant permission for redistribution and use in source and

binary forms, with or without modification, of the software and documentation provided that the following conditions are met:

- All terms of the all other applicable copyrights and licenses must be followed.
- Redistributions of source code must retain the authors' copyright notice(s), this list of conditions, and the following disclaimer.

- Redistributions in binary form must reproduce the authors' copyright notice(s), this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name(s) of the author(s) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY ITS AUTHORS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Note 7 - ntp

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both the copyright notice and this permission notice appear in supporting documentation, and that the name University of Delaware not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. The University of Delaware makes no representations about the suitability this software for any purpose. It is provided "as is" without express or implied warranty.

Note 8 - openSSH

OpenSSH contains no GPL code.

- Copyright (c) 1995 Tatu Ylonen <ylo@cs.hut.fi>, Espoo, Finland. All rights reserved.

Any derived versions of this software must be clearly marked as such, and if the derived work is incompatible with the protocol description in the RFC file, it must be called by a name other than "ssh" or "Secure Shell".

However, I am not implying to give any licenses to any patents or copyrights held by third parties, and the software includes parts that are not under my direct control. As far as I know, all included source code is used in accordance with the relevant license agreements and can be used freely for any purpose (the GNU license being the most restrictive); see below for details.

[However, none of that term is relevant at this point in time. All of these restrictively licenced software components which he talks about have been removed from OpenSSH, i.e.,

- RSA is no longer included, found in the OpenSSL library
- IDEA is no longer included, its use is deprecated
- DES is now external, in the OpenSSL library
- GMP is no longer used, and instead we call BN code from OpenSSL
- Zlib is now external, in a library
- The make-ssh-known-hosts script is no longer included
- TSS has been removed
- MD5 is now external, in the OpenSSL library
- RC4 support has been replaced with ARC4 support from OpenSSL
- Blowfish is now external, in the OpenSSL library

Note that any information and cryptographic algorithms used in this software are publicly available on the Internet and at any major bookstore, scientific library, and patent office worldwide. More information can be found e.g. at "<http://www.cs.hut.fi/crypto>".

The legal status of this program is some combination of all these permissions and restrictions. Use only at your own responsibility. You will be responsible for any legal consequences yourself; I am not making any claims whether possessing or using this is legal or not in your country, and I am not taking any responsibility on your behalf.

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

- The 32-bit CRC compensation attack detector in deattack.c was contributed by CORE SDI S.A. under a BSD-style license. Cryptographic attack detector for ssh - source code Copyright (c) 1998 CORE SDI S.A., Buenos Aires, Argentina. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that this copyright notice is retained.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES ARE DISCLAIMED. IN NO EVENT SHALL CORE SDI S.A. BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES RESULTING FROM THE USE OR MISUSE OF THIS SOFTWARE.

Ariel Futoransky <futo@core-sdi.com> <<http://www.core-sdi.com>>

- ssh-keygen was contributed by David Mazieres under a BSD-style license.
Copyright 1995, 1996 by David Mazieres <dm@lcs.mit.edu>.
Modification and redistribution in source and binary forms is permitted provided that due credit is given to the author and the OpenBSD project by leaving this copyright notice intact.
- The Rijndael implementation by Vincent Rijmen, Antoon Bosselaers and Paulo Barreto is in the public domain and distributed with the following license:
Version 3.0 (December 2000)
Optimised ANSI C code for the Rijndael cipher (now AES)
Author Vincent Rijmen <vincent.rijmen@esat.kuleuven.ac.be>
Author Antoon Bosselaers <antoon.bosselaers@esat.kuleuven.ac.be>
Author Paulo Barreto <paulo.barreto@terra.com.br>
THIS SOFTWARE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
- One component of the ssh source code is under a 3-clause BSD license, held by the University of California, since we pulled these parts from original Berkeley code.
Copyright (c) 1983, 1990, 1992, 1993, 1995. The Regents of the University of California. All rights reserved.
Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- Remaining components of the software are provided under a standard 2-term BSD licence with the following names as copyright holders:

Markus Friedl
Theo de Raadt
Niels Provos
Dug Song
Aaron Campbell
Damien Miller
Kevin Steves
Daniel Kouril
Wesley Griffin
Per Allansson

Nils Nordman

Simon Wilkinson

- Portable OpenSSH additionally includes code from the following copyright holders, also under the 2-term BSD license:

Ben Lindstrom

Tim Rice

Andre Lucas

Chris Adams

Corinna Vinschen

Cray Inc.

Denis Parker

Gert Doering

Jakob Schlyter

Jason Downs

Juha Yrjölä

Michael Stone

Networks Associates Technology, Inc.

Solar Designer

Todd C. Miller

Wayne Schroeder

William Jones

Darren Tucker

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

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

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR

PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- Portable OpenSSH contains the following additional licenses:
 - a** md5crypt.c, md5crypt.h
"THE BEER-WARE LICENSE" (Revision 42):
Poul-Henning Kamp<phk@login.dknet.dk> wrote this file.
 - b** snprintf replacement
Copyright Patrick Powell 1995
This code is based on code written by Patrick Powell
(papowell@astart.com)
 - c** Compatibility code (openbsd-compat)
Apart from the previously mentioned licenses, various pieces of code in the openbsd-compat/ subdirectory are licensed as follows:
Some code is licensed under a 3-term BSD license, to the following copyright holders:
 - Todd C. Miller
 - Theo de Raadt
 - Damien Miller
 - Eric P. AllmanThe Regents of the University of California
Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
 - Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- Some code is licensed under an ISC-style license, to the following copyright holders:

Internet Software Consortium.

Todd C. Miller

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND TODD C. MILLER DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL TODD C. MILLER BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR

PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

- Some code is licensed under a MIT-style license to the following copyright holders:

Free Software Foundation, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, distribute with modifications, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE ABOVE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name(s) of the above copyright holders shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization.

Note 9 - openssl

- OpenSSL License

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

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

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- All advertising materials mentioning features or use of this software must display the following acknowledgment:
"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
- The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
- Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
- Redistributions of any form whatsoever must retain the following acknowledgment:
"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- Original SSLeay License
This product includes cryptographic software written by Eric Young (eyay@cryptsoft.com).

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com).

The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed.

If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

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

- Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- All advertising materials mentioning features or use of this software must display the following acknowledgement:
"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)"
The word 'cryptographic' can be left out if the rouines from the library being used are not cryptographic related :-).
- If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson
(tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Note 10 - IPERF

Copyright (c) 1999-2007, The Board of Trustees of the University of Illinois
All Rights Reserved.

Iperf performance test

Mark Gates

Ajay Tirumala

Jim Ferguson

Jon Dugan

Feng Qin

Kevin Gibbs

John Estabrook

National Laboratory for Applied Network Research

National Center for Supercomputing Applications

University of Illinois at Urbana-Champaign

<http://www.ncsa.uiuc.edu>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software (Iperf) and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

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

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

Neither the names of the University of Illinois, NCSA, nor the names of its contributors may be used to endorse or promote products derived from this Software without specific prior written permission. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE CONTRIBUTORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE

Note 11 - DMalloc

Copyright 1992 to 2007 by Gray Watson.

Permission to use, copy, modify, and distribute this software for any purpose and without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Gray Watson not be used in advertising or publicity pertaining to distribution of the document or software without specific, written prior permission.

Gray Watson makes no representations about the suitability of the software described herein for any purpose. It is provided "as is" without express or implied warranty.

Modified Code Samples

This section includes the modified code for the various open source components included in the RMX software.

apache: httpd-2.2.11.FIPS_mode.patch

```
diff -Nuar httpd-2.2.11.orig/modules/ssl/ssl_engine_init.c httpd-2.2.11/modules/ssl/
ssl_engine_init.c
--- httpd-2.2.11.orig/modules/ssl/ssl_engine_init.c 2008-09-18 17:34:51.000000000 +0300
+++ httpd-2.2.11/modules/ssl/ssl_engine_init.c 2009-07-23 16:07:13.000000000 +0300
@@ -79,6 +79,19 @@
{
    SSLModConfigRec *mc = myModConfig(s);

+   // sagi - mod_ssl patch
+   if (FIPS_mode() && bits < 1024)
+   {
+       mc->pTmpKeys[idx] = NULL;
+       ap_log_error(APLOG_MARK, APLOG_ERR, 0, s,
+                   "Init: Skipping generating temporary "
+                   "%d bit RSA private key in FIPS mode", bits);
+       return OK;
+   }
+   //
+
+   if (!(mc->pTmpKeys[idx] =
+       RSA_generate_key(bits, RSA_F4, NULL, NULL)))
+   {
@@ -96,6 +109,19 @@
{
```

```
SSLModConfigRec *mc = myModConfig(s);

+ // sagi - mod_ssl patch
+ if (FIPS_mode() && bits < 1024)
+ {
+     mc->pTmpKeys[idx] = NULL;
+     ap_log_error(APLOG_MARK, APLOG_ERR, 0, s,
+         "Init: Skipping generating temporary "
+         "%d bit DH parameters in FIPS mode", bits);
+     return OK;
+ }
+ //
+
+
+
+ if (!(mc->pTmpKeys[idx] =
+     ssl_dh_GetTmpParam(bits)))
+ {
+@@ -231,6 +257,26 @@
+     */
+     ssl_rand_seed(base_server, ptemp, SSL_RSCTX_STARTUP, "Init: ");

+ // sagi = mod_ssl fips patch
+ if (!FIPS_mode())
+ {
+
+     if(FIPS_mode_set(1))
+     {
+         ap_log_error(APLOG_MARK, APLOG_NOTICE, 0, s,
+             "FIPS mode enabled");
+     }
+     else {
+         ap_log_error(APLOG_MARK, APLOG_EMERG, 0, s, "FIPS mode failed");
+         ssl_log_ssl_error(APLOG_MARK, APLOG_EMERG, s);
+         ssl_die();
+     }
+     OpenSSL_add_all_ciphers();
+     OpenSSL_add_all_digests();
```

```

+ }
+ //
+
+
+ /*
+  * read server private keys/public certs into memory.
+  * decrypting any encrypted keys via configured SSLPassPhraseDialogs
diff -Nuar httpd-2.2.11.orig/modules/ssl/ssl_private.h httpd-2.2.11/modules/ssl/ssl_private.h
--- httpd-2.2.11.orig/modules/ssl/ssl_private.h2007-11-20 16:16:11.000000000 +0200
+++ httpd-2.2.11/modules/ssl/ssl_private.h2009-07-23 16:02:01.000000000 +0300
@@ -526,6 +526,14 @@
const char *ssl_cmd_SSLProxyMachineCertificatePath(cmd_parms *, void *, const char *);
const char *ssl_cmd_SSLProxyMachineCertificateFile(cmd_parms *, void *, const char *);

+
+
+/* sagi - mod_ssl patch for fips */
+const char *ssl_cmd_SSLFIPS(cmd_parms *cmd, void *dcfg, int flag);
+/* end */
+
+
+
+/** module initialization */
int      ssl_init_Module(apr_pool_t *, apr_pool_t *, apr_pool_t *, server_rec *);
void      ssl_init_Engine(server_rec *, apr_pool_t *);

```

apache: httpd-2.2.11.pragma.patch

```

diff -Nuar httpd-2.2.11.orig/modules/http/http_filters.c httpd-2.2.11/modules/http/http_filters.c
--- httpd-2.2.11.orig/modules/http/http_filters.c2008-12-03 00:28:21.000000000 +0200
+++ httpd-2.2.11/modules/http/http_filters.c2009-07-23 13:54:57.000000000 +0300
@@ -1270,6 +1270,7 @@
    "Warning",
    "WWW-Authenticate",
    "Proxy-Authenticate",
+   "Pragma",
    "Set-Cookie",
    "Set-Cookie2",

```

```
NULL);
```

bind-9.3.1-cross-compile.diff

```
--- bind-9.3.1/configure.orig2005-12-04 13:06:15.000000000 +0200
+++ bind-9.3.1/configure2005-12-04 12:35:14.000000000 +0200
@@ -23677,7 +23677,7 @@

echo "$as_me:$LINENO: checking for inet_ntop with IPv6 support" >&5
echo $ECHO_N "checking for inet_ntop with IPv6 support... $ECHO_C" >&6
-if test "$cross_compiling" = yes; then
+if test "$cross_compiling" = hack_yes; then
    { { echo "$as_me:$LINENO: error: cannot run test program while cross compiling
See `config.log' for more details." >&5
echo "$as_me: error: cannot run test program while cross compiling
```

busybox-1.00-new-passwd-dir.diff

```
--- busybox-1.00/libbb/messages.c.orig2004-03-15 10:28:42.000000000 +0200
+++ busybox-1.00/libbb/messages.c2005-11-08 12:55:54.000000000 +0200
@@ -56,27 +56,27 @@
#endif

#ifdef L_passwd_file
-#define PASSWD_FILE    "/etc/passwd"
+#define PASSWD_FILE    "/etc/volatile/passwd"
const char * const bb_path_passwd_file = PASSWD_FILE;
#endif

#ifdef L_shadow_file
-#define SHADOW_FILE    "/etc/shadow"
+#define SHADOW_FILE    "/etc/volatile/shadow"
const char * const bb_path_shadow_file = SHADOW_FILE;
#endif
```

```
#ifdef L_group_file
#define GROUP_FILE    "/etc/group"
#define GROUP_FILE    "/etc/volatile/group"
const char * const bb_path_group_file = GROUP_FILE;
#endif

#ifdef L_gshadow_file
#define GSHADOW_FILE    "/etc/gshadow"
#define GSHADOW_FILE    "/etc/volatile/gshadow"
const char * const bb_path_gshadow_file = GSHADOW_FILE;
#endif

#ifdef L_nologin_file
#define NOLOGIN_FILE    "/etc/nologin"
#define NOLOGIN_FILE    "/etc/volatile/nologin"
const char * const bb_path_nologin_file = NOLOGIN_FILE;
#endif
```

busybox-1.00-switch_root.patch

```
diff -Nuar busybox-1.00.orig/include/applets.h busybox-1.00/include/applets.h
--- busybox-1.00.orig/include/applets.h2004-08-27 02:01:34.000000000 +0300
+++ busybox-1.00/include/applets.h2006-01-17 12:32:03.000000000 +0200
@@ -544,6 +544,9 @@
#ifdef CONFIG_SWAPONOFF
APPLET(swapon, swap_on_off_main, _BB_DIR_SBIN, _BB_SUID_NEVER)
#endif
+#ifdef CONFIG_SWITCH_ROOT
+APPLET(switch_root, switch_root_main, _BB_DIR_SBIN, _BB_SUID_NEVER)
+#endif
#ifdef CONFIG_SYNC
APPLET(sync, sync_main, _BB_DIR_BIN, _BB_SUID_NEVER)
#endif
```



```

diff -Nuar busybox-1.00.orig/include/usage.h busybox-1.00/include/usage.h
--- busybox-1.00.orig/include/usage.h2004-09-14 19:23:56.000000000 +0300
+++ busybox-1.00/include/usage.h2006-01-17 12:32:03.000000000 +0200
@@ -2280,6 +2280,12 @@
"Options:\n" \
"\t-a\tStart swapping on all swap devices"

+#define switch_root_trivial_usage \
+"NEW_ROOT NEW_INIT [ARGUMENTS_TO_INIT]"
+#define switch_root_full_usage \
+"Use from PID 1 under initramfs to free initramfs, chroot to NEW_ROOT,\n" \
+"and exec NEW_INIT.\n"
+
#define sync_trivial_usage \
""

#define sync_full_usage \
diff -Nuar busybox-1.00.orig/util-linux/Config.in busybox-1.00/util-linux/Config.in
--- busybox-1.00.orig/util-linux/Config.in2004-05-19 14:06:20.000000000 +0300
+++ busybox-1.00/util-linux/Config.in2006-01-17 12:32:03.000000000 +0200
@@ -260,7 +260,28 @@
with some other mounted file system. This allows you to do all sorts
of wild and crazy things with your Linux system and is far more
powerful than 'chroot'.

-
+ Note: This is for initrd in linux 2.4. Under initramfs (introduced
+ in linux 2.6) use switch_root instead.
+
+config CONFIG_SWITCH_ROOT
+bool "switch_root"
+default n
+help
+ The switch_root utility is used from initramfs to select a new
+ root device. Under initramfs, you have to use this instead of

```

```
+ pivot_root. (Stop reading here if you don't care why.)
+
+ Booting with initramfs extracts a gzipped cpio archive into rootfs
+ (which is a variant of ramfs/tmpfs). Because rootfs can't be moved
+ or unmounted*, pivot_root will not work from initramfs. Instead,
+ switch_root deletes everything out of rootfs (including itself),
+ does a mount --move that overmounts rootfs with the new root, and
+ then execs the specified init program.
+
+ * Because the Linux kernel uses rootfs internally as the starting
+ and ending point for searching through the kernel's doubly linked
+ list of active mount points. That's why.
+
config CONFIG_RDATE
bool "rdate"
default n
diff -Nuar busybox-1.00.orig/util-linux/Makefile.in busybox-1.00/util-linux/Makefile.in
--- busybox-1.00.orig/util-linux/Makefile.in2004-10-08 09:46:08.000000000 +0200
+++ busybox-1.00/util-linux/Makefile.in2006-01-17 12:32:03.000000000 +0200
@@ -41,6 +41,7 @@
UTILLINUX-$(CONFIG_MOUNT)+=mount.o
UTILLINUX-$(CONFIG_NFSMOUNT)+=nfsmount.o
UTILLINUX-$(CONFIG_PIVOT_ROOT)+=pivot_root.o
+UTILLINUX-$(CONFIG_SWITCH_ROOT)+=switch_root.o
UTILLINUX-$(CONFIG_RDATE)+=rdate.o
UTILLINUX-$(CONFIG_SWAPONOFF)+=swaponoff.o
UTILLINUX-$(CONFIG_UMOUNT)+=umount.o
diff -Nuar busybox-1.00.orig/util-linux/switch_root.c busybox-1.00/util-linux/switch_root.c
--- busybox-1.00.orig/util-linux/switch_root.c1970-01-01 02:00:00.000000000 +0200
+++ busybox-1.00/util-linux/switch_root.c2006-01-22 10:58:16.000000000 +0200
@@ -0,0 +1,128 @@
+/* vi:set ts=4:*/
+
```

```
+#include <dirent.h>
+#include <fcntl.h>
+#include <stdio.h>
+#include <sys/mount.h>
+#include <sys/stat.h>
+#include <sys/types.h>
+#include <sys/vfs.h>
+#include <unistd.h>
+
+#include "busybox.h"
+
+// Make up for header deficiencies.
+
+#ifndef RAMFS_MAGIC
+#define RAMFS_MAGIC0x858458f6
+#endif
+
+#ifndef TMPFS_MAGIC
+#define TMPFS_MAGIC0x01021994
+#endif
+
+#ifndef MS_MOVE
+#define MS_MOVE8192
+#endif
+
+dev_t rootdev;
+
+// Recursively delete contents of rootfs.
+
+static void delete_contents(char *directory)
+{
+DIR *dir;
+struct dirent *d;
```

```
+struct stat st;
+
+// Don't descend into other filesystems
+if (stat(directory,&st) || st.st_dev != rootdev) return;
+
+// Recursively delete the contents of directories.
+if (S_ISDIR(st.st_mode)) {
+if((dir = opendir(directory))) {
+while ((d = readdir(dir))) {
+char *newdir=d->d_name;
+
+// Skip . and ..
+if(*newdir=='.' && (!newdir[1] || (newdir[1]=='.' && !newdir[2])))
+continue;
+
+// Recurse to delete contents
+newdir = alloca(strlen(directory) + strlen(d->d_name) + 2);
+sprintf(newdir, "%s/%s", directory, d->d_name);
+delete_contents(newdir);
+}
+closedir(dir);
+
+// Directory should now be empty. Zap it.
+rmdir(directory);
+}
+
+// It wasn't a directory. Zap it.
+
+} else unlink(directory);
+}
+
+int switch_root_main(int argc, char *argv[])
+{
```

```
+char *newroot, *console=NULL;
+struct stat st1, st2;
+struct statfs stfs;
+char *args[8];
+
+// Parse args (-c console)
+
+bb_opt_complementaly="-2";
+bb_getopt_ulflags(argc,argv,"c:",&console);
+
+// Change to new root directory and verify it's a different fs.
+
+newroot="/mnt/root";
+
+if (chdir(newroot) || stat(".", &st1) || stat("/", &st2) ||
+st1.st_dev == st2.st_dev)
+{
+bb_error_msg_and_die("bad newroot %s",newroot);
+}
+rootdev=st2.st_dev;
+
+// Additional sanity checks: we're about to rm -rf /, so be REALLY SURE
+// we mean it. (I could make this a CONFIG option, but I would get email
+// from all the people who WILL eat their filesystems.)
+
+//if (stat("/init", &st1) || !S_ISREG(st1.st_mode) || statfs("/", &stfs) ||
+//(stfs.f_type != RAMFS_MAGIC && stfs.f_type != TMPFS_MAGIC) ||
+//getpid() != 1)
+//{
+//bb_error_msg_and_die("not rootfs");
+//}
+
+// Zap everything out of rootdev
```

```
+
+delete_contents("/");
+
+// Overmount / with newdir and chroot into it. The chdir is needed to
+// recalculate "." and ".." links.
+
+if (mount(".", "/", NULL, MS_MOVE, NULL) || chroot(".") || chdir("/"))
+bb_error_msg_and_die("moving root");
+
+// If a new console specified, redirect stdin/stdout/stderr to that.
+
+if (console) {
+close(0);
+if(open(console, O_RDWR) < 0)
+bb_error_msg_and_die("Bad console '%s'",console);
+dup2(0, 1);
+dup2(0, 2);
+}
+
+
+// Exec real init. (This is why we must be pid 1.)
+//execv(argv[optind],argv+optind+1);
+
+args[0]="init";
+execv("/sbin/init", args);
+bb_error_msg_and_die("Bad init '/sbin/init'");
+}
```

busybox-1.00-vi-backspace.diff

```
# posted by paul fox on the busybox list on Thu Nov 4 08:30:58 MST 2004
#
# the following patch fixes vi so that ^H and DEL will both
```

```
# properly backspace when you're in insert mode, whether or not
# they're your tty erasechar. this matches the behavior of vile and
# vim. (the patch also moves some whitespace and fixes a comment
# typo.)
#
# paul
# =-----
# paul fox, pgf at brightstareng.comdiff -ru busybox-1.00.orig/editors/vi.c busybox-1.00/
editors/vi.c
#
--- busybox-1.00.orig/editors/vi.c2004-08-19 15:15:06.000000000 -0400
+++ busybox-1.00/editors/vi.c2004-11-04 10:10:37.000000000 -0500
@@ -1596,7 +1596,7 @@
if ((p[-1] != '\n') && (dot>text)) {
p--;
}
-} else if (c == erase_char) { // Is this a BS
+} else if (c == erase_char || c == 8 || c == 127) { // Is this a BS
// 123456789
if ((p[-1] != '\n') && (dot>text)) {
p--;
@@ -3061,8 +3061,8 @@
break;
case 'h': // h- move left
case VI_K_LEFT: // cursor key Left
-case 8: // ctrl-H- move left (This may be ERASE char)
-case 127: // DEL- move left (This may be ERASE char)
+case 8: // ctrl-H- move left (This may be ERASE char)
+case 127: // DEL- move left (This may be ERASE char)
if (cmdcnt-- > 1) {
do_cmd(c);
} // repeat cnt
@@ -3225,7 +3225,7 @@
```

```
//
// dont separate these two commands. 'f' depends on ';'
//
-//**** fall thru to ... 'i'
+//**** fall thru to ... ';'
case ';/ /' :- look at rest of line for last forward char
if (cmdcnt-- > 1) {
do_cmd(';');
```

cfgsh.patch

```
diff -Naur cfgsh-1.1.1/cfgsh.h cfgsh-1.1.1-fixed/cfgsh.h
--- cfgsh-1.1.1/cfgsh.h2004-06-06 22:22:36.000000000 +0300
+++ cfgsh-1.1.1-fixed/cfgsh.h2005-10-06 09:29:31.800990646 +0200
@@ -35,15 +35,15 @@
#define PATH_MAX (1024)
#endif

-#define TZ_PATH "/etc/localtime"
+#define TZ_PATH "/config/common/etc/localtime"
#define ZONES_PATH "/usr/share/zoneinfo"

#define PING_PATH "/bin/ping"
#define PING_COUNT_PARAM "-c"
#define PING_COUNT "4"

-#define RESOLV_PATH "/etc/resolv.test"
-#define CONFIG_PATH "/etc/system.conf"
+#define RESOLV_PATH "/etc/resolv.conf"
+#define CONFIG_PATH "/config/common/etc/system.conf"
#define STATE_PATH "/cfgsh.state"

#define DHCPD_START_PATH "/usr/sbin/run-dhcpd"
```


dhcpcd-1.3.22-pl4.patch

```

diff -Nuar dhcpcd-1.3.22-pl4.orig/buildmsg.c dhcpcd-1.3.22-pl4/buildmsg.c
--- dhcpcd-1.3.22-pl4.orig/buildmsg.c 2002-12-31 07:15:07.000000000 +0200
+++ dhcpcd-1.3.22-pl4/buildmsg.c 2006-04-25 07:51:40.000000000 +0300
@@ -92,7 +92,7 @@
     memcpy(p,&nleaseTime,4);
     p += 4;
     *p++ = dhcpParamRequest;
- *p++ = 14;
+ *p++ = 16; // 14;
     *p++ = subnetMask;
     *p++ = routersOnSubnet;
     *p++ = dns;
@@ -107,6 +107,8 @@
     *p++ = nisDomainName;
     *p++ = nisServers;
     *p++ = ntpServers;
+ *p++ = dhcpSipServers; // AN
+ *p++ = dhcpAvayaSsopOpt; // AN
     if ( HostName )
     {
         *p++ = hostName;
@@ -179,7 +179,7 @@
     p += 4;
     }
     *p++ = dhcpParamRequest;
- *p++ = 14;
+ *p++ = 16; // 14;
     *p++ = subnetMask;
     *p++ = routersOnSubnet;
     *p++ = dns;
@@ -194,6 +194,6 @@
     *p++ = nisDomainName;

```

```
*p++ = nisServers;
*p++ = ntpServers;
+ *p++ = dhcpSipServers;// AN
+ *p++ = dhcpAvayaSsopOpt;// AN
  if ( HostName )
  {
    *p++ = hostName;
@@ -254,7 +258,7 @@
  }
#endif
  *p++ = dhcpParamRequest;
- *p++ = 14;
+ *p++ = 16;// 14;
  *p++ = subnetMask;
  *p++ = routersOnSubnet;
  *p++ = dns;
@@ -269,6 +273,8 @@
  *p++ = nisDomainName;
  *p++ = nisServers;
  *p++ = ntpServers;
+ *p++ = dhcpSipServers;// AN
+ *p++ = dhcpAvayaSsopOpt;// AN
  if ( HostName )
  {
    *p++ = hostName;
@@ -327,7 +333,7 @@
    p += 4;
  }
  *p++ = dhcpParamRequest;
- *p++ = 14;
+ *p++ = 16;// 14;
  *p++ = subnetMask;
  *p++ = routersOnSubnet;
```

```
*p++ = dns;
@@ -342,6 +348,8 @@
    *p++ = nisDomainName;
    *p++ = nisServers;
    *p++ = ntpServers;
+ *p++ = dhcpSipServers;// AN
+ *p++ = dhcpAvayaSsopOpt;// AN
    if ( HostName )
    {
        *p++ = hostName;
@@ -408,7 +416,7 @@
    memcpy(p,&nleaseTime,4);
    p += 4;
    *p++ = dhcpParamRequest;
- *p++ = 14;
+ *p++ = 16;// 14;
    *p++ = subnetMask;
    *p++ = routersOnSubnet;
    *p++ = dns;
@@ -423,6 +431,8 @@
    *p++ = nisDomainName;
    *p++ = nisServers;
    *p++ = ntpServers;
+ *p++ = dhcpSipServers;// AN
+ *p++ = dhcpAvayaSsopOpt;// AN
    if ( HostName )
    {
        *p++ = hostName;
@@ -556,7 +566,7 @@
    memcpy(p,&dhcpMsgSize,2);
    p += 2;
    *p++ = dhcpParamRequest;
- *p++ = 14;
```

```
+ *p++ = 16; // 14;
  *p++ = subnetMask;
  *p++ = routersOnSubnet;
  *p++ = dns;
@@ -571,6 +581,8 @@
  *p++ = nisDomainName;
  *p++ = nisServers;
  *p++ = ntpServers;
+ *p++ = dhcpSipServers; // AN
+ *p++ = dhcpAvayaSsopOpt; // AN
  if ( HostName )
  {
    *p++ = hostName;
diff -Nuar dhcpcd-1.3.22-pl4.orig/client.h dhcpcd-1.3.22-pl4/client.h
--- dhcpcd-1.3.22-pl4.orig/client.h2002-12-31 07:15:33.000000000 +0200
+++ dhcpcd-1.3.22-pl4/client.h2006-04-25 07:52:10.000000000 +0300
@@ -131,6 +131,8 @@
  dhcpT2value=59,
  dhcpClassIdentifier=60,
  dhcpClientIdentifier=61,
+ dhcpSipServers= 120,
+ dhcpAvayaSsopOpt= 176, // Avaya specific Site Specific Option for GK parameters
  endOption=255
};

diff -Nuar dhcpcd-1.3.22-pl4.orig/dhcpconfig.c dhcpcd-1.3.22-pl4/dhcpconfig.c
--- dhcpcd-1.3.22-pl4.orig/dhcpconfig.c2002-12-30 00:05:44.000000000 +0200
+++ dhcpcd-1.3.22-pl4/dhcpconfig.c2006-04-25 07:51:35.000000000 +0300
@@ -725,6 +725,89 @@
fprintf(f,"CLIENTID=%02X:%02X:%02X:%02X:%02X:%02X\n",
DhcpIface.client_id[3],DhcpIface.client_id[4],DhcpIface.client_id[5],
DhcpIface.client_id[6],DhcpIface.client_id[7],DhcpIface.client_id[8]);
+
```

```
+// AN - dhcpSipServers
+if ( DhcpOptions.len[dhcpSipServers])
+ {
+ int enc;
+
+ enc = (int)((((unsigned char *)DhcpOptions.val[dhcpSipServers])[0]);
+
+ // enc == 1: IPv4 Address
+ if (enc == 1) {
+   fprintf(f,"SIPSERVERS_ADDR=%u.%u.%u.%u",
+   ((unsigned char *)DhcpOptions.val[dhcpSipServers])[1],
+   ((unsigned char *)DhcpOptions.val[dhcpSipServers])[2],
+   ((unsigned char *)DhcpOptions.val[dhcpSipServers])[3],
+   ((unsigned char *)DhcpOptions.val[dhcpSipServers])[4]);
+   for (i=5;i<DhcpOptions.len[dhcpSipServers];i+=4)
+     fprintf(f,"%u.%u.%u.%u",
+     ((unsigned char *)DhcpOptions.val[dhcpSipServers])[i],
+     ((unsigned char *)DhcpOptions.val[dhcpSipServers])[1+i],
+     ((unsigned char *)DhcpOptions.val[dhcpSipServers])[2+i],
+     ((unsigned char *)DhcpOptions.val[dhcpSipServers])[3+i]);
+
+ } else { // enc == 0: domain names
+
+ int j;
+ int strLen;
+ int domainIndex = 0;
+ int bDot = 0;
+ int totalLen;
+
+ unsigned char val;
+ unsigned char sipDomainName[255];
+
+ totalLen = DhcpOptions.len[dhcpSipServers];
```

```
+
+ if (totalLen > 255)
+ totalLen = 255;
+
+ fprintf(f,"SIPSERVERS_NAME=");
+
+for (i = 1, j = 0; i < totalLen;) {
+
+strLen = (int)((((unsigned char *)DhcpOptions.val[dhcpSipServers])[i]);
+
+if (strLen) {
+
+if (bDot)
+sipDomainName[domainIndex++] = '.';
+else {
+sipDomainName[domainIndex++] = '\\';
+bDot = 1;
+}
+
+} else {
+strcpy(&sipDomainName[domainIndex], "\\,");
+domainIndex += (strlen("\\,") + 1);
+i++;
+bDot = 0;
+continue;
+}
+
+for (j = i + 1; j < strLen; j++) {
+val = (((unsigned char *)DhcpOptions.val[dhcpSipServers])[j]);
+sipDomainName[domainIndex++] = val;
+}
+
+i += (strLen + 1);
```

```
+}
+
+sipDomainName[--domainIndex] = '\0';
+ fprintf(f, "%s\n", sipDomainName);
+ }
+ }
+ else
+  printf("\nNO DHCP SIP servers\n");
+
+// gk
+if ( DhcpOptions.len[dhcpAvayaSsopOpt])
+ {
+  fprintf(f, "\nAVAYA_SSOP_OPT=%s\n", (char *) DhcpOptions.val[dhcpAvayaSsopOpt]);
+ } else
+  fprintf(f, "\nNO DHCP AVAYA Options\n");
+
+    fclose(f);
+  }
+  else
```

dhcpcd-2.0.0_rpl_malloc.patch

```
--- dhcpcd-2.0.0/configure.orig2005-12-13 21:27:02.000000000 +0200
+++ dhcpcd-2.0.0/configure2005-12-13 21:25:46.000000000 +0200
@@ -4613,9 +4613,9 @@
esac

-cat >>confdefs.h <<\_ACEOF
-#define malloc rpl_malloc
-_ACEOF
+#cat >>confdefs.h <<\_ACEOF
+##define malloc rpl_malloc
+#_ACEOF

fi
```

gdb-6.3-cppflags-fix.diff

```
diff -ru gdb-6.3/Makefile.in build/gdb-6.3/Makefile.in
--- gdb-6.3/Makefile.in2005-01-06 13:19:03.000000000 +0100
+++ build/gdb-6.3/Makefile.in2005-01-06 13:20:48.000000000 +0100
@@ -156,6 +156,7 @@
CC="$(CC)"; export CC; \
CFLAGS="$(CFLAGS)"; export CFLAGS; \
CONFIG_SHELL="$(SHELL)"; export CONFIG_SHELL; \
+CPPFLAGS="$(CPPFLAGS)"; export CPPFLAGS; \
CXX="$(CXX)"; export CXX; \
CXXFLAGS="$(CXXFLAGS)"; export CXXFLAGS; \
AR="$(AR)"; export AR; \
@@ -198,7 +199,7 @@
CC="$(CC_FOR_TARGET)"; export CC; \
CFLAGS="$(CFLAGS_FOR_TARGET)"; export CFLAGS; \
CONFIG_SHELL="$(SHELL)"; export CONFIG_SHELL; \
```



```
-CPPFLAGS="$(CFLAGS_FOR_TARGET)"; export CPPFLAGS; \  
+CPPFLAGS="$(CPPFLAGS_FOR_TARGET)"; export CPPFLAGS; \  
CXXFLAGS="$(CXXFLAGS_FOR_TARGET)"; export CXXFLAGS; \  
GCJ="$(GCJ_FOR_TARGET)"; export GCJ; \  
GFORTRAN="$(GFORTRAN_FOR_TARGET)"; export GFORTRAN; \  
@@ -326,6 +327,7 @@  
CC = @CC@  
CFLAGS = @CFLAGS@  
LIBCFLAGS = $(CFLAGS)  
+CPPFLAGS = @CPPFLAGS@  
  
CXX = @CXX@  
CXXFLAGS = @CXXFLAGS@  
@@ -388,6 +390,7 @@  
# (which we know are built with gcc) are built with optimizations so  
# prepend -O2 when setting CFLAGS_FOR_TARGET.  
CFLAGS_FOR_TARGET = -O2 $(CFLAGS)  
+CPPFLAGS_FOR_TARGET = @CPPFLAGS_FOR_TARGET@  
# If GCC_FOR_TARGET is not overridden on the command line, then this  
# variable is passed down to the gcc Makefile, where it is used to  
# build libgcc2.a. We define it here so that it can itself be  
Only in build/gdb-6.3: Makefile.in~
```

gdb-6.3-ldflags-fix.diff

```
#  
# Submitted:  
#  
# Marc Kleine-Budde <kleine-budde@gmx.de>, 2004-12-01  
#  
# Error:  
#  
# checking for library containing waddstr... no
```

```
# checking for library containing tgetent... no
# configure: error: no termcap library found
# make[1]: *** [configure-gdb] Error 1
# make[1]: Leaving directory `/ptx/work/mkl/build/cameron/ptxdist/build/gdb-6.3-build'
# make: *** [/ptx/work/mkl/build/cameron/ptxdist/state/gdb.compile] Error 2
#
# Description:
#
# LDFLAGS and LDFLAGS_FOR_TARGET are overwritten by toplevel Makefile
# instead preserving the ones from the configure-run
#
# State:
#
# going-to-be-transmitted-to-gdb-mailinglist-real-soon
#
diff -ruN gdb-6.3-orig/Makefile.in gdb-6.3/Makefile.in
--- gdb-6.3-orig/Makefile.in2004-11-08 17:21:22.000000000 +0100
+++ gdb-6.3/Makefile.in2004-12-01 20:52:05.000000000 +0100
@@ -336,7 +336,7 @@
NM = @NM@

LD = @LD@
-LDFLAGS =
+LDFLAGS = @LDFLAGS@

RANLIB = @RANLIB@

@@ -434,7 +434,7 @@
    fi; \
    fi`

-LDFLAGS_FOR_TARGET =
+LDFLAGS_FOR_TARGET = @LDFLAGS_FOR_TARGET@
```

```
NM_FOR_TARGET=@NM_FOR_TARGET@
CONFIGURED_NM_FOR_TARGET=@CONFIGURED_NM_FOR_TARGET@
```

gdb-6.4-termcap-configure.patch

```
--- gdb-6.4/gdb/configure.orig 2005-12-13 19:39:35.000000000 +0200
+++ gdb-6.4/gdb/configure 2005-12-13 19:39:00.000000000 +0200
@@ -8198,6 +8198,9 @@
     ac_cv_search_tgetent="none required"
     CONFIG_OBS="$CONFIG_OBS win32-termcap.o"
     ;;
+ *)
+ LIBS="$with_termcap/libtermcap.a $LIBS"
+ ac_cv_search_tgetent="$with_termcap/libtermcap.a"
esac
```

These are the libraries checked by Readline.

genext2fs_1.3-7.1.diff

```
#
# debian patch 1.3-7.1
# e.g. http://debian.tu-bs.de/debian/pool/main/g/genext2fs/genext2fs_1.3-7.1.diff.gz
# fixes problem with getline under cygwin
#
--- genext2fs-1.3.orig/Makefile
+++ genext2fs-1.3/Makefile
@@ -0,0 +1,32 @@
+all: genext2fs
+INSTALL=install
+CFLAGS=-Wall -O2
```

```
+
+install:
+$(INSTALL) -d $(DESTDIR)/usr/bin/
+$(INSTALL) -m 755 genext2fs $(DESTDIR)/usr/bin/
+$(INSTALL) -d $(DESTDIR)/usr/share/man/man8/
+$(INSTALL) -m 644 genext2fs.8 $(DESTDIR)/usr/share/man/man8/
+
+clean:
+-rm genext2fs
+rm -rf test ext2.img
+
+check: all
+mkdir -p test
+dd if=/dev/zero of=test/zero count=1
+./genext2fs -b 4096 -d test ext2.img
+
+md5_correct=d048d86c381a4ad361c9652f86b1eba9; \
+md5_tested=`md5sum ext2.img 2>/dev/null | cut -f 1 -d " "; \
+if [ "$md5_tested" = "$md5_correct" ]; then \
+echo "test succeeded."; \
+else \
+md5_tested=`md5 ext2.img 2>/dev/null | cut -f 4 -d " "; \
+if [ "$md5_tested" = "$md5_correct" ]; then \
+echo "test succeeded."; \
+else \
+echo "test failed."; \
+fi \
+fi
+
--- genext2fs-1.3.orig/debian/control
+++ genext2fs-1.3/debian/control
@@ -0,0 +1,19 @@
+Source: genext2fs
```

```
+Section: admin
+Priority: optional
+Maintainer: David Kimdon <dwhedon@debian.org>
+Build-Depends: debhelper (> 3.0.0)
+Standards-Version: 3.5.2
+
+Package: genext2fs
+Architecture: any
+Depends: ${shlibs:Depends}
+Description: ext2 filesystem generator for embedded systems
+ `genext2fs' is meant to generate an ext2 filesystem
+ as a normal (non-root) user. It doesn't require you to mount
+ the image file to copy files on it. It doesn't even require
+ you to be the superuser to make device nodes.
+ .
+ Warning ! `genext2fs' has been designed for embedded
+ systems. As such, it will generate a filesystem for single-user
+ usage: all files/directories/etc... will belong to UID/GID 0
--- genext2fs-1.3.orig/debian/rules
+++ genext2fs-1.3/debian/rules
@@ -0,0 +1,70 @@
+#!/usr/bin/make -f
+# Sample debian/rules that uses debhelper.
+# GNU copyright 1997 to 1999 by Joey Hess.
+
+# Uncomment this to turn on verbose mode.
+#export DH_VERBOSE=1
+
+# This is the debhelper compatability version to use.
+export DH_COMPAT=2
+
+configure: configure-stamp
+configure-stamp:
```

```
+dh_testdir
+# Add here commands to configure the package.
+# ./configure --prefix=/usr --mandir=/usr/share/man/
+
+touch configure-stamp
+
+build: configure-stamp build-stamp
+build-stamp:
+dh_testdir
+
+# Add here commands to compile the package.
+$(MAKE)
+$(MAKE) check
+
+touch build-stamp
+
+clean:
+dh_testdir
+dh_testroot
+rm -f build-stamp configure-stamp
+
+# Add here commands to clean up after the build process.
+$(MAKE) clean
+
+dh_clean
+
+install: build
+dh_testdir
+dh_testroot
+dh_clean -k
+dh_installdirs
+
+# Add here commands to install the package into debian/genext2fs.
```

```
+$ (MAKE) install DESTDIR=`pwd`/debian/genext2fs
+
+
+## Build architecture-independent files here.
+binary-indep: build install
+## We have nothing to do by default.
+
+## Build architecture-dependent files here.
+binary-arch: build install
+dh_testdir
+dh_testroot
+dh_installdocs
+dh_installchangelogs
+dh_link
+dh_strip
+dh_compress
+dh_fixperms
+dh_installdeb
+dh_shlibdeps
+dh_gencontrol
+dh_md5sums
+dh_builddeb
+
+binary: binary-indep binary-arch
+.PHONY: build clean binary-indep binary-arch binary install configure
--- genext2fs-1.3.orig/debian/changelog
+++ genext2fs-1.3/debian/changelog
@@ -0,0 +1,70 @@
+genext2fs (1.3-7.1) unstable; urgency=high
+
+ * Non-maintainer upload.
+ * Apply patch from Finn Thain to fix byteswapping code issues that broke
+   d-i on several arches. Thanks. (Closes: #266039)
```

```
+
+ -- Joshua Kwan <joshk@triplehelix.org> Tue, 24 Aug 2004 19:09:01 -0700
+
+genext2fs (1.3-7) unstable; urgency=low
+
+ * Fix block count rounding. The -b switch again specifies the
+ size of the resulting filesystem image.
+ Patch thanks to Finn Thain <ft01@webmastery.com.au>
+ (closes: #262352)
+
+ -- David Kimdon <dwhedon@debian.org> Sat, 7 Aug 2004 09:51:37 +0200
+
+genext2fs (1.3-6) unstable; urgency=low
+
+ * Fix block counts, now created image fsck clean. (closes: #232023)
+ Thanks to Finn Thain <ft01@webmastery.com.au> for the fix.
+
+ -- David Kimdon <dwhedon@debian.org> Sat, 10 Jul 2004 17:37:17 +0200
+
+genext2fs (1.3-5) unstable; urgency=low
+
+ * - Compile with -O2.
+ - get_workblk() should return zero'd memory. This way we get
+ nul terminations in the produced image.
+ - get_workblk() now allocates the block dynamically on the heap.
+ The previous method of allocating the block as 'static' apparently
+ results in undefined behavior since get_workblk() is 'inline'.
+ Thanks to Vincent Sanders <vince@debian.org> for the patch.
+ (closes: #248987)
+
+ -- David Kimdon <dwhedon@debian.org> Sat, 15 May 2004 19:37:54 +0200
+
+genext2fs (1.3-4) unstable; urgency=low
```


+
+ * Fix assertion failure when there is a 268K file on the
+ filesystem. Patch thanks to Matt Kraai <kraai@alumni.cmu.edu>
+ (closes: #201277)
+
+ -- David Kimdon <dwhedon@debian.org> Mon, 17 Nov 2003 22:53:15 +0100
+
+genext2fs (1.3-3) unstable; urgency=low
+
+ * Update to current upstream cvs
+ - support triple indirect blocks (removes 8MB limit)
+ - support for groups.
+
+ -- David Kimdon <dwhedon@debian.org> Sat, 11 Jan 2003 18:45:58 -0800
+
+genext2fs (1.3-2) unstable; urgency=low
+
+ * apply fix from upstream cvs that appears to fix endian bug
+ (closes: #122411)
+ * mention filesystem size limit in manpage (closes: #122729)
+ * mention that hard links are not supported in manpage
+ (closes: #155464)
+ * add sanity check at the end of the build
+
+ -- David Kimdon <dwhedon@debian.org> Fri, 8 Mar 2002 23:17:36 -0800
+
+genext2fs (1.3-1) unstable; urgency=low
+
+ * Initial Release. (closes: #105263)
+
+ -- David Kimdon <dwhedon@debian.org> Sat, 14 Jul 2001 13:24:49 -0700
+
--- genext2fs-1.3.orig/debian/copyright

```
+++ genext2fs-1.3/debian/copyright
@@ -0,0 +1,15 @@
+This package was debianized by David Kimdon <dwhedon@debian.org> on
+Sat, 14 Jul 2001 13:24:49 -0700.
+
+It was downloaded from http://freshmeat.net/projects/genext2fs/
+Upstream Author(s): Xavier Bestel <xbestel@aplio.fr>
+
+Copyright (C) 2000 Xavier Bestel <xavier.bestel@free.fr>
+
+This program is free software; you can redistribute it and/or
+modify it under the terms of the GNU General Public License
+as published by the Free Software Foundation; version
+2 of the License.
+
+On Debian systems, the complete text of the GNU General Public
+License can be found in /usr/share/common-licenses/GPL file.
--- genext2fs-1.3.orig/genext2fs.8
+++ genext2fs-1.3/genext2fs.8
@@ -0,0 +1,112 @@
+.\"                Hey, EMACS: -*- nroff -*-
+.\" First parameter, NAME, should be all caps
+.\" Second parameter, SECTION, should be 1-8, maybe w/ subsection
+.\" other parameters are allowed: see man(7), man(1)
+.TH GENEXT2FS 8 "July 14, 2001"
+.\" Please adjust this date whenever revising the manpage.
+.\"
+.\" Some roff macros, for reference:
+.\" .nh      disable hyphenation
+.\" .hy      enable hyphenation
+.\" .ad l    left justify
+.\" .ad b    justify to both left and right margins
+.\" .nf      disable filling
```

+.\ " .fi enable filling
+.\ " .br insert line break
+.\ " .sp <n> insert n+1 empty lines
+.\ " for manpage-specific macros, see man(7)
+.SH NAME
+genext2fs \- ext2 filesystem generator for embedded systems
+.SH SYNOPSIS
+.B genext2fs
+.RI [options] " image"
+.SH DESCRIPTION
+.\fbgenext2fs\fp generates an ext2 filesystem
+as a normal (non-root) user. It doesn't require you to mount
+the image file to copy files on it. It doesn't even require
+you to be the superuser to make device nodes.
+.SH OPTIONS
+.TP
+.BI -x \ image
+Use this image as a starting point
+.TP
+.BI -d \ directory
+Add this directory as source
+.TP
+.BI -f \ file
+Add nodes (e.g. devices) from this spec file
+.TP
+.BI -b \ blocks
+Size in blocks
+.TP
+.BI -i \ inodes
+Number of inodes
+.TP
+.BI -r \ reserved
+Number of reserved blocks

```
+ .TP
+ .BI -g \ path
+ Generate a block map file for this path
+ .TP
+ .BI -e \ value
+ Fill unallocated blocks with value
+ .TP
+ .BI -z
+ Make files with holes
+ .TP
+ .BI -v
+ Print resulting filesystem structure
+ .TP
+ .BI -h
+ Display help
+ .TP
+ .SH EXAMPLES
+
+ .EX
+ .B
+ genext2fs -b 1440 -d src /dev/fd0
+ .EE
+
+ All files in the
+ .I src
+ directory will be written to
+ .B /dev/fd0
+ as a new ext2 filesystem image. You can then mount the floppy as
+ usual.
+
+ .EX
+ .B
+ genext2fs -b 1024 -d src -f dev.txt flashdisk.img
```

```
+EE
+
+This example builds a filesystem from all the files in
+.I src
+, then device nodes are created based on the contents of the device file
+.I dev.txt.
+An example device file follows:
+
+.EX
+ drwx      /dev
+ crw- 10,190 /dev/lcd
+ brw- 1,0   /dev/ram0
+.EE
+
+This device list builds the /dev directory, a character device
+node /dev/lcd (major 10, minor 190) and a block device node
+/dev/ram0 (major 1, minor 0)
+.SH BUGS
+\fbgenext2fs\fp has been designed for embedded
+systems. As such, it will generate a filesystem for single-user
+usage: all files/directories/etc... will belong to UID/GID 0
+
+\fbgenext2fs\fp does not support hard links. Hard links present in the input
+tree will be represented as separate files in the ext2 image.
+
+.SH SEE ALSO
+.BR mkfs (8),
+.BR genromfs (8),
+.BR mkisofs (8).
+.br
+.SH AUTHOR
+This manual page was written by David Kimdon <dwshedon@debian.org>,
+for the Debian GNU/Linux system (but may be used by others).
```

```
--- genext2fs-1.3.orig/genext2fs.c
+++ genext2fs-1.3/genext2fs.c
@@ -26,6 +26,15 @@
// Bugfix: getcwd values for Solarisxavier.gueguen@col.bsf.alcatel.fr
// Bugfix: ANSI scanf for non-GNU Cxavier.gueguen@col.bsf.alcatel.fr
// 28 Jun 2001Bugfix: getcwd differs for Solaris/GNUMike@sowbug.com
+// 8 Mar 2002Bugfix: endianness swap of x-indirects
+// 23 Mar 2002Bugfix: test for IFCHR or IFBLK was flawed
+// 10 Oct 2002Added comments,makefile targets,vsundar@ixiacom.com
+// endianness swap assert check.
+// Copyright (C) 2002 Ixia communications
+// 12 Oct 2002Added support for triple indirectionvsundar@ixiacom.com
+// Copyright (C) 2002 Ixia communications
+// 14 Oct 2002Added support for groupsvsundar@ixiacom.com
+// Copyright (C) 2002 Ixia communications

// `genext2fs' is a mean to generate an ext2 filesystem
@@ -67,6 +76,7 @@
#include <stdarg.h>
#include <unistd.h>
#include <sys/stat.h>
+#include <assert.h>

@@ -76,10 +86,14 @@
#define BLOCKSIZE      1024
#define BLOCKS_PER_GROUP 8192
#define BYTES_PER_INODE (8*BLOCKSIZE)
+/* Percentage of blocks that are reserved.*/
#define RESERVED_INODES 5/100
```

```
// inode block size (why is it != BLOCKSIZE !?)
+/* The field i_blocks in the ext2 inode stores the number of data blocks
+ but in terms of 512 bytes. That is what INODE_BLOCKSIZE represents.
+ INOBLK is the number of such blocks in an actual disk block */

#define INODE_BLOCKSIZE 512
#define INOBLK (BLOCKSIZE / INODE_BLOCKSIZE)
@@ -147,6 +161,39 @@

#define OP_HOLES 0x01 // make files with holes

+/* Defines for accessing group details */
+
+// Number of groups in the filesystem
+#define GRP_NBGROUPS(fs) ( (fs)->sb.s_blocks_count+(fs)->sb.s_blocks_per_group-1)/
+(fs)->sb.s_blocks_per_group )
+
+// Get group block bitmap (bbm) given the group number
+#define GRP_GET_GROUP_BBM(fs,grp) ( get_blk((fs),(fs)->gd[(grp)].bg_block_bitmap) )
+
+// Get group inode bitmap (ibm) given the group number
+#define GRP_GET_GROUP_IBM(fs,grp) ( get_blk((fs),(fs)->gd[(grp)].bg_inode_bitmap) )
+
+// Given an inode number find the group it belongs to
+#define GRP_GROUP_OF_INODE(fs,nod) ( ((nod)-1) / (fs)->sb.s_inodes_per_group)
+
+// Given an inode number get the inode bitmap that covers it
+#define GRP_GET_INODE_BITMAP(fs,nod) \
+( GRP_GET_GROUP_IBM((fs),GRP_GROUP_OF_INODE((fs),(nod)))) )
+
+// Given an inode number find its offset within the inode bitmap that covers it
+#define GRP_IBM_OFFSET(fs,nod) \
```

```
+( (nod) - GRP_GROUP_OF_INODE((fs),(nod))*(fs)->sb.s_inodes_per_group )
+
+// Given a block number find the group it belongs to
+#define GRP_GROUP_OF_BLOCK(fs,blk) ( ((blk)-1) / (fs)->sb.s_blocks_per_group)
+
+// Given a block number get the block bitmap that covers it
+#define GRP_GET_BLOCK_BITMAP(fs,blk) \
+( GRP_GET_GROUP_BBM((fs),GRP_GROUP_OF_BLOCK((fs),(blk)))) )
+
+// Given a block number find its offset within the block bitmap that covers it
+#define GRP_BBM_OFFSET(fs,blk) \
+( (blk) - GRP_GROUP_OF_BLOCK((fs),(blk))*(fs)->sb.s_blocks_per_group )
+

// used types

@@ -287,7 +334,6 @@
{
    groupdescriptor_decl
    uint32 bg_reserved[3];
    -uint32 bg_pad_to_bk[(BLOCKSIZE-32)/sizeof(uint32)];
    } groupdescriptor;

typedef struct
@@ -304,6 +350,32 @@

typedef uint8 block[BLOCKSIZE];

+/* blockwalker fields:
+ The blockwalker is used to access all the blocks of a file (including
+ the indirection blocks) through repeated calls to walk_bw.
+
+ bpdirent -> index into the inode->i_block[]. Indicates level of indirection.
```


- + bnum -> total number of blocks so far accessed. including indirection
- + blocks.
- + bbind,bpdind,bptind -> index into indirection blocks.
- +
- + bbind, bpdind, bptind do *NOT* index into single, double and triple
- + indirect blocks resp. as you might expect from their names. Instead
- + they are in order the 1st, 2nd & 3rd index to be used
- +
- + As an example..
- + To access data block number 70000:
- + bbind: 15 (we are doing triple indirection)
- + bbind: 0 (index into the triple indirection block)
- + bpdind: 16 (index into the double indirection block)
- + bptind: 99 (index into the single indirection block)
- + $70000 = 12 + 256 + 256*256 + 16*256 + 100$ (indexing starts from zero)
- +
- + So,for double indirection bbind will index into the double indirection
- + block and bpdind into the single indirection block. For single indirection
- + only bbind will be used.
- + */
- +
- typedef struct
- {
- uint32 bnum;
- @@ -313,15 +385,14 @@
- uint32 bptind;
- } blockwalker;
- +
- + /* Filesystem structure that support groups */
- #if BLOCKSIZE == 1024
- typedef struct
- {

```
block zero;    // The famous block 0
superblock sb; // The superblock
-groupdescriptor gd; // The group descriptor
-block bbm;    // The block bitmap
-block ibm;    // The inode bitmap
-inode itab[0]; // The inode table
+groupdescriptor gd[0]; // The group descriptors
} filesystem;
#else
#error UNHANDLED BLOCKSIZE
@@ -416,14 +487,15 @@
// temporary working block
inline uint8 * get_workblk(void)
{
-static block b;
+unsigned char* b=calloc(1,BLOCKSIZE);
return b;
}
inline void free_workblk(block b)
{
+free(b);
}

-// rounds a quantity up to a blocksize
+/* Rounds qty upto a multiple of siz. siz should be a power of 2 */
uint32 rndup(uint32 qty, uint32 siz)
{
return (qty + (siz - 1)) & ~(siz - 1);
@@ -444,7 +516,13 @@
// return a given inode from a filesystem
inline inode * get_nod(filesystem *fs, uint32 nod)
{
-return &fs->itab[nod-1];
```

```

+int grp,offset;
+inode *itab;
+
+offset = GRP_IBM_OFFSET(fs,nod);
+grp = GRP_GROUP_OF_INODE(fs,nod);
+itab = (inode *)get_blk(fs, fs->gd[grp].bg_inode_table);
+return itab+offset-1;
}

// allocate a given block/inode in the bitmap
@@ -479,29 +557,57 @@
}

// allocate a block
-uint32 alloc_blk(filesystem *fs)
+uint32 alloc_blk(filesystem *fs, uint32 nod)
{
-uint32 bk;
-if(!(bk = allocate(fs->bbm, 0)))
+uint32 bk=0;
+uint32 grp,nbgroups;
+
+grp = nod/fs->sb.s_inodes_per_group;
+nbgroups = ( fs->sb.s_blocks_count - fs->sb.s_first_data_block + fs->sb.s_blocks_per_group -
1 ) /
+fs->sb.s_blocks_per_group;
+if(!(bk = allocate(get_blk(fs,fs->gd[grp].bg_block_bitmap), 0))) {
+for(grp=0;grp<nbgroups && !bk;grp++)
+bk=allocate(get_blk(fs,fs->gd[grp].bg_block_bitmap),0);
+grp--;
+}
+if (!bk)
+errx(1,"couldn't allocate a block (no free space)");

```

```
-if(!(fs->gd.bg_free_blocks_count--))
-errexit("group descr. free blocks count == 0 (corrupted fs?)");
+if(!(fs->gd[grp].bg_free_blocks_count--))
+errexit("group descr %d. free blocks count == 0 (corrupted fs?)",grp);
if(!(fs->sb.s_free_blocks_count--))
errexit("superblock free blocks count == 0 (corrupted fs?)");
-return bk;
+return fs->sb.s_blocks_per_group*grp + bk;
}

// allocate an inode
uint32 alloc_nod(filesystem *fs)
{
-uint32 nod;
-if(!(nod = allocate(fs->ibm, 0)))
+uint32 nod=0,best_group=0;
+uint32 grp,nbgroups,avefreei;
+
+nbgroups = ( fs->sb.s_blocks_count - fs->sb.s_first_data_block + fs->sb.s_blocks_per_group -
1 ) /
+fs->sb.s_blocks_per_group;
+
+/* Distribute inodes amongst all the blocks */
+/* For every block group with more than average number of free inodes */
+/* find the one with the most free blocks and allocate node there */
+/* Idea from find_group_dir in fs/ext2/ialloc.c in 2.4.19 kernel */
+/* We do it for all inodes. */
+avefreei = fs->sb.s_free_inodes_count / nbgroups;
+for(grp=0;grp<nbgroups && !nod;grp++) {
+if (fs->gd[grp].bg_free_inodes_count < avefreei)
+continue;
+if (!best_group ||
+fs->gd[grp].bg_free_blocks_count > fs->gd[best_group].bg_free_blocks_count)
```

```

+best_group = grp;
+}
+if (!(nod = allocate(get_blk(fs,fs->gd[best_group].bg_inode_bitmap),0)))
    errexit("couldn't allocate an inode (no free inode)");
-if(!(fs->gd.bg_free_inodes_count--))
+if(!(fs->gd[best_group].bg_free_inodes_count--))
    errexit("group descr. free blocks count == 0 (corrupted fs?)");
if(!(fs->sb.s_free_inodes_count--))
    errexit("superblock free blocks count == 0 (corrupted fs?)");
-return nod;
+return fs->sb.s_inodes_per_group*best_group+nod;
}

// print a bitmap allocation
@@ -546,14 +652,14 @@
{
    bkref = &get_nod(fs, nod)->i_block[bw->bpdire = 0];
    if(extend) // allocate first block
    -*bkref = hole ? 0 : alloc_blk(fs);
    +*bkref = hole ? 0 : alloc_blk(fs,nod);
}
// direct block
else if(bw->bpdire < EXT2_NDIR_BLOCKS)
{
    bkref = &get_nod(fs, nod)->i_block[++bw->bpdire];
    if(extend) // allocate block
    -*bkref = hole ? 0 : alloc_blk(fs);
    +*bkref = hole ? 0 : alloc_blk(fs,nod);
}
// first block in indirect block
else if(bw->bpdire == EXT2_NDIR_BLOCKS)
@@ -562,11 +668,11 @@
    bw->bpdire = EXT2_IND_BLOCK;

```

```
bw->bpind = 0;
if(extend) // allocate indirect block
-get_nod(fs, nod)->i_block[bw->bpdir] = alloc_blk(fs);
+get_nod(fs, nod)->i_block[bw->bpdir] = alloc_blk(fs,nod);
b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdir]);
bkref = &b[bw->bpind];
if(extend) // allocate first block
-*bkref = hole ? 0 : alloc_blk(fs);
+*bkref = hole ? 0 : alloc_blk(fs,nod);
}
// block in indirect block
else if((bw->bpdir == EXT2_IND_BLOCK) && (bw->bpind < BLOCKSIZE/4 - 1))
@@ -575,7 +681,7 @@
b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdir]);
bkref = &b[bw->bpind];
if(extend) // allocate block
-*bkref = hole ? 0 : alloc_blk(fs);
+*bkref = hole ? 0 : alloc_blk(fs,nod);
}
// first block in first indirect block in first double indirect block
else if(bw->bpdir == EXT2_IND_BLOCK)
@@ -585,14 +691,14 @@
bw->bpind = 0;
bw->bpdind = 0;
if(extend) // allocate double indirect block
-get_nod(fs, nod)->i_block[bw->bpdir] = alloc_blk(fs);
+get_nod(fs, nod)->i_block[bw->bpdir] = alloc_blk(fs,nod);
b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdir]);
if(extend) // allocate first indirect block
-b[bw->bpind] = alloc_blk(fs);
+b[bw->bpind] = alloc_blk(fs,nod);
b = (uint32*)get_blk(fs, b[bw->bpind]);
bkref = &b[bw->bpdind];
```

```

if(extend) // allocate first block
-*bkref = hole ? 0 : alloc_blk(fs);
+*bkref = hole ? 0 : alloc_blk(fs,nod);
}
// block in indirect block in double indirect block
else if((bw->bpdire == EXT2_DIND_BLOCK) && (bw->bpdind < BLOCKSIZE/4 - 1))
@@ -602,7 +708,7 @@
b = (uint32*)get_blk(fs, b[bw->bpdind]);
bkref = &b[bw->bpdind];
if(extend) // allocate block
-*bkref = hole ? 0 : alloc_blk(fs);
+*bkref = hole ? 0 : alloc_blk(fs,nod);
}
// first block in indirect block in double indirect block
else if((bw->bpdire == EXT2_DIND_BLOCK) && (bw->bpdind < BLOCKSIZE/4 - 1))
@@ -612,19 +718,99 @@
bw->bpdind++;
b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdire]);
if(extend) // allocate indirect block
-b[bw->bpdind] = alloc_blk(fs);
+b[bw->bpdind] = alloc_blk(fs,nod);
b = (uint32*)get_blk(fs, b[bw->bpdind]);
bkref = &b[bw->bpdind];
if(extend) // allocate first block
-*bkref = hole ? 0 : alloc_blk(fs);
+*bkref = hole ? 0 : alloc_blk(fs,nod);
+}
+
+/* Adding support for triple indirection */
+/* Just starting triple indirection. Allocate the indirection
+ blocks and the first data block
+ */
+else if (bw->bpdire == EXT2_DIND_BLOCK)

```

```
+{
+ bw->bnum += 3;
+bw->bpdirext = EXT2_TIND_BLOCK;
+bw->bpind = 0;
+bw->bpdind = 0;
+bw->bptind = 0;
+if(extend) // allocate triple indirect block
+get_nod(fs, nod)->i_block[bw->bpdirext] = alloc_blk(fs,nod);
+b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdirext]);
+if(extend) // allocate first double indirect block
+b[bw->bpind] = alloc_blk(fs,nod);
+b = (uint32*)get_blk(fs, b[bw->bpind]);
+if(extend) // allocate first indirect block
+b[bw->bpdind] = alloc_blk(fs,nod);
+b = (uint32*)get_blk(fs, b[bw->bpdind]);
+bkref = &b[bw->bptind];
+if(extend) // allocate first data block
+*bkref = hole ? 0 : alloc_blk(fs,nod);
+}
+/* Still processing a single indirect block down the indirection
+ chain. Allocate a data block for it
+ */
+else if ( (bw->bpdirext == EXT2_TIND_BLOCK) &&
+ (bw->bptind < BLOCKSIZE/4 -1) )
+{
+bw->bptind++;
+b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdirext]);
+b = (uint32*)get_blk(fs, b[bw->bpind]);
+b = (uint32*)get_blk(fs, b[bw->bpdind]);
+bkref = &b[bw->bptind];
+if(extend) // allocate data block
+*bkref = hole ? 0 : alloc_blk(fs,nod);
+}
```



```
+/* Finished processing a single indirect block. But still in the
+ same double indirect block. Allocate new single indirect block
+ for it and a data block
+ */
+else if ( (bw->bpdire == EXT2_TIND_BLOCK) &&
+ (bw->bpdind < BLOCKSIZE/4 -1) )
+{
+bw->bnum++;
+bw->bptind = 0;
+bw->bpdind++;
+b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdire]);
+b = (uint32*)get_blk(fs, b[bw->bpdind]);
+if (extend) // allocate single indirect block
+b[bw->bpdind] = alloc_blk(fs,nod);
+b = (uint32*)get_blk(fs, b[bw->bpdind]);
+bkref = &b[bw->bptind];
+if(extend) // allocate first data block
+*bkref = hole ? 0 : alloc_blk(fs,nod);
+}
+/* Finished processing a double indirect block. Allocate the next
+ double indirect block and the single,data blocks for it
+ */
+else if ( (bw->bpdire == EXT2_TIND_BLOCK) &&
+ (bw->bpind < BLOCKSIZE/4 - 1) )
+{
+bw->bnum += 2;
+bw->bpdind = 0;
+bw->bptind = 0;
+bw->bpind++;
+b = (uint32*)get_blk(fs, get_nod(fs, nod)->i_block[bw->bpdire]);
+if(extend) // allocate double indirect block
+b[bw->bpind] = alloc_blk(fs,nod);
+b = (uint32*)get_blk(fs, b[bw->bpind]);
```

```
+if(extend) // allocate single indirect block
+b[bw->bpdind] = alloc_blk(fs,nod);
+b = (uint32*)get_blk(fs, b[bw->bpdind]);
+bkref = &b[bw->bptind];
+if(extend) // allocate first block
+*bkref = hole ? 0 : alloc_blk(fs,nod);
+
-// I don't do triple indirect - it's such a small filesystem ...
else
-errexit("file too big ! blocks list for inode %d extends past double indirect blocks!", nod);
+errexit("file too big ! blocks list for inode %d extends past triple indirect blocks!", nod);
+/* End change for walking triple indirection */
+
+if(*bkref)
+{
+  bw->bnum++;
-if(!allocated(fs->bbm, *bkref))
+if(!allocated(GRP_GET_BLOCK_BITMAP(fs,*bkref), GRP_BBM_OFFSET(fs,*bkref)))
+  errexit("[block %d of inode %d is unallocated !]", *bkref, nod);
+}
+if(extend)
@@ -780,7 +966,7 @@
+add2dir(fs, parent_nod, nod, name);
+add2dir(fs, nod, nod, ".");
+add2dir(fs, nod, parent_nod, "..");
-fs->gd.bg_used_dirs_count++;
+fs->gd[GRP_GROUP_OF_INODE(fs,nod)].bg_used_dirs_count++;
+return nod;
+}

@@ -981,44 +1167,98 @@
// endianness swap of x-indirect blocks
void swap_goodblocks(filesystem *fs, inode *nod)
```

```

{
-int i;
+int i,j,done=0;
+uint32 *b,*b2;
+
  int nblk = nod->i_blocks / INOBLK;
-if((nod->i_size && !nblk) || (nod->i_mode & (FM_IFBLK | FM_IFCHR)))
+if((nod->i_size && !nblk) || ((nod->i_mode & FM_IFBLK) == FM_IFBLK) || ((nod->i_mode
& FM_IFCHR) == FM_IFCHR))
  for(i = 0; i <= EXT2_TIND_BLOCK; i++)
    nod->i_block[i] = swab32(nod->i_block[i]);
  if(nblk <= EXT2_IND_BLOCK)
    return;
  swap_block(get_blk(fs, nod->i_block[EXT2_IND_BLOCK]));
-if(nblk <= EXT2_IND_BLOCK + BLOCKSIZE/4)
+if(nblk <= EXT2_DIND_BLOCK + BLOCKSIZE/4)
  return;
+/* Currently this will fail b'cos the number of blocks as stored
+ in i_blocks also includes the indirection blocks (see
+ walk_bw). But this function assumes that i_blocks only
+ stores the count of data blocks ( Actually according to
+ "Understanding the Linux Kernel" (Table 17-3 p502 1st Ed)
+ i_blocks IS supposed to store the count of data blocks). so
+ with a file of size 268K nblk would be 269.The above check
+ will be false even though double indirection hasn't been
+ started.This is benign as 0 means block 0 which has been
+ zeroed out and therefore points back to itself from any offset
+ */
+assert(nod->i_block[EXT2_DIND_BLOCK] != 0);
  for(i = 0; i < BLOCKSIZE/4; i++)
-if(nblk > EXT2_IND_BLOCK + BLOCKSIZE/4 + i)
+if(nblk > EXT2_IND_BLOCK + BLOCKSIZE/4 + (BLOCKSIZE/4)*i )
    swap_block(get_blk(fs, ((uint32*)get_blk(fs, nod->i_block[EXT2_DIND_BLOCK]))[i]));

```

```
swap_block(get_blk(fs, nod->i_block[EXT2_DIND_BLOCK]));
if(nblk <= EXT2_IND_BLOCK + BLOCKSIZE/4 + BLOCKSIZE/4 * BLOCKSIZE/4)
return;
-errexit("too big file on the filesystem");
+/* Adding support for triple indirection */
+b = (uint32*)get_blk(fs,nod->i_block[EXT2_TIND_BLOCK]);
+for(i=0;i < BLOCKSIZE/4 && !done ; i++) {
+b2 = (uint32*)get_blk(fs,b[i]);
+for(j=0; j<BLOCKSIZE/4;j++) {
+if (nblk > ( EXT2_IND_BLOCK + BLOCKSIZE/4 +
+ (BLOCKSIZE/4)*(BLOCKSIZE/4) +
+ i*(BLOCKSIZE/4)*(BLOCKSIZE/4) +
+ j*(BLOCKSIZE/4)) )
+ swap_block(get_blk(fs,b2[j]));
+else {
+ done = 1;
+ break;
+}
+}
+swap_block((uint8 *)b2);
+}
+swap_block((uint8 *)b);
+return;
}

void swap_badblocks(filesystem *fs, inode *nod)
{
-int i;
+int i,j,done=0;
+uint32 *b,*b2;
+
int nblk = nod->i_blocks / INOBLK;
-if((nod->i_size && !nblk) || (nod->i_mode & (FM_IFBLK | FM_IFCHR)))
```

```

+if((nod->i_size && !nblk) || ((nod->i_mode & FM_IFBLK) == FM_IFBLK) || ((nod->i_mode
& FM_IFCHR) == FM_IFCHR))
for(i = 0; i <= EXT2_TIND_BLOCK; i++)
nod->i_block[i] = swab32(nod->i_block[i]);
if(nblk <= EXT2_IND_BLOCK)
return;
swap_block(get_blk(fs, nod->i_block[EXT2_IND_BLOCK]));
-if(nblk <= EXT2_IND_BLOCK + BLOCKSIZE/4)
+if(nblk <= EXT2_DIND_BLOCK + BLOCKSIZE/4)
return;
+/* See comment in swap_goodblocks */
+assert(nod->i_block[EXT2_DIND_BLOCK] != 0);
swap_block(get_blk(fs, nod->i_block[EXT2_DIND_BLOCK]));
for(i = 0; i < BLOCKSIZE/4; i++)
-if(nblk > EXT2_IND_BLOCK + BLOCKSIZE/4 + i)
+if(nblk > EXT2_IND_BLOCK + BLOCKSIZE/4 + (BLOCKSIZE/4)*i )
swap_block(get_blk(fs, ((uint32*)get_blk(fs, nod->i_block[EXT2_DIND_BLOCK]))[i]));
if(nblk <= EXT2_IND_BLOCK + BLOCKSIZE/4 + BLOCKSIZE/4 * BLOCKSIZE/4)
return;
-errexit("too big file on the filesystem");
+/* Adding support for triple indirection */
+b = (uint32*)get_blk(fs,nod->i_block[EXT2_TIND_BLOCK]);
+swap_block((uint8 *)b);
+for(i=0;i < BLOCKSIZE/4 && !done ; i++) {
+b2 = (uint32*)get_blk(fs,b[i]);
+swap_block((uint8 *)b2);
+for(j=0;j<BLOCKSIZE/4;j++) {
+if (nblk > ( EXT2_IND_BLOCK + BLOCKSIZE/4 +
+ (BLOCKSIZE/4)*(BLOCKSIZE/4) +
+ i*(BLOCKSIZE/4)*(BLOCKSIZE/4) +
+ j*(BLOCKSIZE/4)) )
+ swap_block(get_blk(fs,b2[j]));
+else {

```

```
+ done = 1;
+ break;
+}
+}
+}
+return;
}

// endianness swap of the whole filesystem
@@ -1045,7 +1285,8 @@
swap_goodblocks(fs, nod);
swap_nod(nod);
}
-swap_gd(&fs->gd);
+for(i=0;i<GRP_NBGROUPS(fs);i++)
+swap_gd(&(fs->gd[i]));
swap_sb(&fs->sb);
}

@@ -1053,7 +1294,8 @@
{
int i;
swap_sb(&fs->sb);
-swap_gd(&fs->gd);
+for(i=0;i<GRP_NBGROUPS(fs);i++)
+swap_gd(&(fs->gd[i]));
for(i = 1; i < fs->sb.s_inodes_count; i++)
{
inode *nod = get_nod(fs, i);
@@ -1084,53 +1326,117 @@
directory *d;
uint8 * b;
uint32 nod;
```

```

+uint32 nbgroups,nbinodes_per_group,overhead_per_group,free_blocks,
+free_blocks_per_group,nbblocks_per_group;
+uint32 gd,itbl,ibmpos,bbmpos,itblpos;
+int j;
+uint8 *bbm,*ibm;
+inode *itab0;

if(nbblocks < 16) // totally arbitrary
errexit("too small filesystem");
-if(nbblocks > BLOCKS_PER_GROUP) // I build only one group
-errexit("too big filesystem");
+
+/* nbblocks is the total number of blocks in the filesystem. First
+ * calculate the size of each group assuming each group has
+ * BLOCKS_PER_GROUP blocks (which is the maximum). Then recalculate
+ * blocks per group so that each group (except possibly the last one)
+ * has the same number of blocks. nbinodes is the total number of
+ * inodes in the system. These are divided between all groups.
+ * Then calculate the overhead blocks - inode table blocks, bitmap
+ * blocks, group descriptor blocks etc.
+ */
+
+nbgroups = (nbblocks + BLOCKS_PER_GROUP - 1) / BLOCKS_PER_GROUP;
+nbblocks_per_group = rndup((nbblocks + nbgroups - 1)/nbgroups, 8);
+nbinodes_per_group = rndup((nbinodes + nbgroups - 1)/nbgroups,
+(BLOCKSIZE/sizeof(inode)));
+if (nbinodes_per_group < 16)
+nbinodes_per_group = 16; // minimum number b'cos the first 10 are reserved
+
+gd = rndup(nbgroups*sizeof(groupdescriptor),BLOCKSIZE)/BLOCKSIZE;
+itbl = nbinodes_per_group * sizeof(inode)/BLOCKSIZE;
+overhead_per_group = 3 /*sb,ibm,bbm*/ + itbl + gd;
+free_blocks = nbblocks - overhead_per_group*nbgroups - 1 /*boot block*/;

```

```
+free_blocks_per_group = nbblocks_per_group - overhead_per_group;
+
if(!(fs = (filesystem*)calloc(nbblocks, BLOCKSIZE)))
errexit("not enough memory for filesystem");

// create the superblock for an empty filesystem
-fs->sb.s_inodes_count = rndup(nbinodes, BLOCKSIZE/ sizeof(inode));
+fs->sb.s_inodes_count = nbinodes_per_group * nbgroups;
fs->sb.s_blocks_count = nbblocks;
fs->sb.s_r_blocks_count = nbresrvd;
-fs->sb.s_free_blocks_count = nbblocks;
+fs->sb.s_free_blocks_count = free_blocks;
fs->sb.s_free_inodes_count = fs->sb.s_inodes_count - EXT2_FIRST_INO + 1;
fs->sb.s_first_data_block = (BLOCKSIZE == 1024);
fs->sb.s_log_block_size = BLOCKSIZE >> 11;
fs->sb.s_log_frag_size = BLOCKSIZE >> 11;
-fs->sb.s_blocks_per_group = BLOCKS_PER_GROUP;
-fs->sb.s_frags_per_group = BLOCKS_PER_GROUP;
-fs->sb.s_inodes_per_group = fs->sb.s_inodes_count;
+fs->sb.s_blocks_per_group = nbblocks_per_group;
+fs->sb.s_frags_per_group = nbblocks_per_group;
+fs->sb.s_inodes_per_group = nbinodes_per_group;
fs->sb.s_magic = EXT2_MAGIC_NUMBER;

// set up groupdescriptors
-fs->sb.s_free_blocks_count -= 5 + fs->sb.s_inodes_count * sizeof(inode) / BLOCKSIZE;
-fs->gd.bg_free_blocks_count = fs->sb.s_free_blocks_count;
-fs->gd.bg_free_inodes_count = fs->sb.s_free_inodes_count;
-fs->gd.bg_used_dirs_count = 1;
-fs->gd.bg_block_bitmap = 3;
-fs->gd.bg_inode_bitmap = 4;
-fs->gd.bg_inode_table = 5;
-
```



```
-// mark non-filesystem blocks and inodes as allocated
-for(i = fs->sb.s_blocks_count; i <= BLOCKSIZE * 8; i++)
-allocate(fs->bbm, i);
-for(i = fs->sb.s_inodes_count + 1; i <= BLOCKSIZE * 8; i++)
-allocate(fs->ibm, i);
-
-// mark system blocks and inodes as allocated
-for(i = 1; i <= 4 + fs->sb.s_inodes_count * sizeof(inode) / BLOCKSIZE; i++)
-allocate(fs->bbm, i);
-for(i = 1; i < EXT2_FIRST_INO; i++)
-allocate(fs->ibm, i);
+for(i = 0,bbmpos=2+gd,ibmpos=3+gd,itblpos =4+gd;
+i<nbgroups;
+i++, bbmpos += nbblocks_per_group,ibmpos += nbblocks_per_group,
+itblpos += nbblocks_per_group) {
+
+if( free_blocks > free_blocks_per_group ) {
+fs->gd[i].bg_free_blocks_count = free_blocks_per_group;
+free_blocks -= free_blocks_per_group;
+} else {
+fs->gd[i].bg_free_blocks_count = free_blocks;
+free_blocks = 0; // this is the last block group
+}
+if( i )
+fs->gd[i].bg_free_inodes_count = nbinoes_per_group;
+else
+fs->gd[i].bg_free_inodes_count = nbinoes_per_group -
+EXT2_FIRST_INO + 2;
+fs->gd[i].bg_used_dirs_count = 0;
+fs->gd[i].bg_block_bitmap = bbmpos;
+fs->gd[i].bg_inode_bitmap = ibmpos;
+fs->gd[i].bg_inode_table = itblpos;
+}
```

```
+
+/* Mark non-filesystem blocks and inodes as allocated */
+/* Mark system blocks and inodes as allocated      */
+for(i = 0; i<nbgroups;i++) {
+
+/* Block bitmap */
+//j=fs->sb.s_inodes_per_group;
+//printf("j: %d\n",j);
+bbm = get_blk(fs,fs->gd[i].bg_block_bitmap);
+//non-filesystem blocks
+for(j = fs->gd[i].bg_free_blocks_count
++ overhead_per_group + 1; j <= BLOCKSIZE * 8; j++)
+allocate(bbm, j);
+//system blocks
+for(j = 1; j <= overhead_per_group; j++)
+allocate(bbm, j);
+
+/* Inode bitmap */
+ibm = get_blk(fs,fs->gd[i].bg_inode_bitmap);
+//non-filesystem inodes
+for(j = fs->sb.s_inodes_per_group+1; j <= BLOCKSIZE * 8; j++)
+allocate(ibm, j);
+//system inodes
+if( !i )
+for(j = 1; j < EXT2_FIRST_INO; j++)
+allocate(ibm, j);
+}

// make root inode and directory
-fs->itab[EXT2_ROOT_INO-1].i_mode = FM_IFDIR | FM_IRWXU | FM_IRWXG |
FM_IRWXO;
-fs->itab[EXT2_ROOT_INO-1].i_size = BLOCKSIZE;
-fs->itab[EXT2_ROOT_INO-1].i_links_count = 2;
```

```

+/* We have groups now. Add the root filesystem in group 0 */
+/* Also increment the directory count for group 0 */
+fs->gd[0].bg_free_inodes_count--;
+fs->gd[0].bg_used_dirs_count = 1;
+itab0 = (inode *)get_blk(fs,fs->gd[0].bg_inode_table);
+itab0[EXT2_ROOT_INO-1].i_mode = FM_IFDIR | FM_IRWXU | FM_IRWXG | FM_IRWXO;
+itab0[EXT2_ROOT_INO-1].i_size = BLOCKSIZE;
+itab0[EXT2_ROOT_INO-1].i_links_count = 2;
+
b = get_workblk();
d = (directory*)b;
d->d_inode = EXT2_ROOT_INO;
@@ -1406,7 +1712,7 @@
s = (nod >= EXT2_FIRST_INO) ? "normal" : "unknown reserved";
}
printf("inode %d (%s, %d links): ", nod, s, get_nod(fs, nod)->i_links_count);
-if(!allocated(fs->ibm, nod))
+if(!allocated(GRP_GET_INODE_BITMAP(fs,nod), GRP_IBM_OFFSET(fs,nod)))
{
printf("unallocated\n");
return;
@@ -1445,19 +1751,21 @@
// describes various fields in a filesystem
void print_fs(filesystem *fs)
{
-int i;
+// int i;

printf("%d blocks (%d free, %d reserved), first data block: %d\n", fs->sb.s_blocks_count, fs->sb.s_free_blocks_count, fs->sb.s_r_blocks_count, fs->sb.s_first_data_block);
printf("%d inodes (%d free)\n", fs->sb.s_inodes_count, fs->sb.s_free_inodes_count);
printf("block size = %d, frag size = %d\n", fs->sb.s_log_block_size ? (fs->sb.s_log_block_size << 11) : 1024, fs->sb.s_log_frag_size ? (fs->sb.s_log_frag_size << 11) : 1024);
printf("%d blocks per group, %d frags per group, %d inodes per group\n", fs->sb.s_blocks_per_group, fs->sb.s_frags_per_group, fs->sb.s_inodes_per_group);

```

```
-printf("block bitmap: block %d, inode bitmap: block %d, inode table: block %d\n", fs-
>gd.bg_block_bitmap, fs->gd.bg_inode_bitmap, fs->gd.bg_inode_table);
+// *TBD* printf("block bitmap: block %d, inode bitmap: block %d, inode table: block %d\n",
fs->gd.bg_block_bitmap, fs->gd.bg_inode_bitmap, fs->gd.bg_inode_table);
printf("block bitmap allocation:\n");
-print_bm(fs->bbm, fs->sb.s_blocks_count);
+// *TBD* print_bm(fs->bbm, fs->sb.s_blocks_count);
printf("inode bitmap allocation:\n");
+/* *TBD*
print_bm(fs->ibm, fs->sb.s_inodes_count);
for(i=1; i<=fs->sb.s_inodes_count; i++)
if(allocated(fs->ibm, i))
print_inode(fs, i);
+*/
}
```

```
void dump_fs(filesystem *fs, FILE * fh, int swapit)
@@ -1622,7 +1930,7 @@
}
if(emptyval)
for(i = 1; i < fs->sb.s_blocks_count; i++)
-if(!allocated(fs->bbm, i))
+if(!allocated(GRP_GET_BLOCK_BITMAP(fs,i),GRP_BBM_OFFSET(fs,i)))
memset(get_blk(fs, i), emptyval, BLOCKSIZE);
if(verbose)
print_fs(fs);
```

grub-0.95.max8MB_kernel.patch

```
diff -Nuar grub-0.95.orig/stage2/boot.c grub-0.95/stage2/boot.c
--- grub-0.95.orig/stage2/boot.c2004-03-30 13:44:08.000000000 +0200
+++ grub-0.95/stage2/boot.c2008-08-20 09:12:23.091296013 +0300
@@ -267,6 +267,9 @@
    data_len = setup_sects << 9;
```

```

text_len = filemax - data_len - SECTOR_SIZE;

+   if (text_len > 1024*1024*8)
+       text_len = 1024*1024*8; // sagi , patch - quick reading only kernel
+
linux_data_tmp_addr = (char *) LINUX_BZIMAGE_ADDR + text_len;

    if (! big_linux
@@ -280,7 +283,7 @@
errnum = ERR_WONT_FIT;
    else
    {
-   grub_printf (" [Linux-%s, setup=0x%x, size=0x%x]\n",
+   grub_printf (" [Linux-%s, setup=0x%x, size=0x%x] - RMX patched for 8MB max size\n",
        (big_linux ? "bzImage" : "zImage"), data_len, text_len);

/* Video mode selection support. What a mess! */

```

grub-0.95-no-A20-keyboard-controller.patch

```

diff -Naur grub-0.95/stage2.orig/builtins.c grub-0.95/stage2/builtins.c
--- grub-0.95/stage2.orig/builtins.c 2004-05-14 22:30:52.000000000 +0300
+++ grub-0.95/stage2/builtins.c 2005-11-24 15:34:46.000000000 +0200
@@ -287,7 +287,7 @@
    set_int13_handler (bios_drive_map);
}

-   gateA20 (0);
+   //gateA20 (0);
    boot_drive = saved_drive;
    chain_stage1 (0, BOOTSEC_LOCATION, boot_part_addr);
    break;
diff -Naur grub-0.95/stage2.orig/common.c grub-0.95/stage2/common.c

```

```
--- grub-0.95/stage2.orig/common.c2004-03-27 18:25:44.000000000 +0200
+++ grub-0.95/stage2/common.c2005-11-24 15:34:52.000000000 +0200
@@ -160,7 +160,7 @@
     * to 0. Not too desirable.
     */

- gateA20 (1);
+ //gateA20 (1);

/* Store the size of extended memory in EXTENDED_MEMORY, in order to
tell it to non-Multiboot OSes. */
```

httpd-2.0.54-return_Pragma_for_not_modified.patch

```
diff -Nuar httpd-2.0.54.orig/modules/http/http_protocol.c httpd-2.0.54/modules/http/
http_protocol.c
--- httpd-2.0.54.orig/modules/http/http_protocol.c 2005-02-04 22:21:18.000000000 +0200
+++ httpd-2.0.54/modules/http/http_protocol.c 2007-01-04 14:42:57.616581920 +0200
@@ -1657,6 +1657,7 @@
     "Warning",
     "WWW-Authenticate",
     "Proxy-Authenticate",
+    "Pragma",
     NULL);
+
+    }
+    else {
```

httpd-2.0.54-skip-storage-on-post.patch

```
diff -Nuar httpd-2.0.54.orig/server/request.c httpd-2.0.54/server/request.c
--- httpd-2.0.54.orig/server/request.c2005-02-04 22:21:18.000000000 +0200
+++ httpd-2.0.54/server/request.c2006-04-06 12:42:11.431183072 +0300
@@ -144,12 +144,17 @@
```

```

/* Reset to the server default config prior to running map_to_storage
*/
+
r->per_dir_config = r->server->lookup_defaults;

- if ((access_status = ap_run_map_to_storage(r))) {
-     /* This request wasn't in storage (e.g. TRACE) */
-     return access_status;
+ if (strcmp(r->method,"POST") != 0) /* avoid disk access in POST - sagi */
+ {
+     if ((access_status = ap_run_map_to_storage(r))) {
+         /* This request wasn't in storage (e.g. TRACE) */
+         return access_status;
+     }
+ }
+

/* Excluding file-specific requests with no 'true' URI...
*/

```

httpd.2.0.61.pragma.300.patch

```

diff -Nuar httpd-2.0.61.orig/modules/http/http_protocol.c httpd-2.0.61/modules/http/
http_protocol.c
--- httpd-2.0.61.orig/modules/http/http_protocol.c 2006-07-12 10:40:55.000000000 +0300
+++ httpd-2.0.61/modules/http/http_protocol.c 2008-06-23 17:24:55.733514521 +0300
@@ -1752,6 +1752,7 @@
     "Warning",
     "WWW-Authenticate",
     "Proxy-Authenticate",
+    "Pragma",
+    NULL);
}
else {

```

lm_sensors-2.10.0-01-EXLDFLAGS.patch

```
diff -Nuar lm_sensors-2.10.0.orig/Makefile lm_sensors-2.10.0/Makefile
--- lm_sensors-2.10.0.orig/Makefile2006-02-15 03:46:46.000000000 +0200
+++ lm_sensors-2.10.0/Makefile2006-08-16 12:52:03.806332418 +0300
@@ -117,7 +117,7 @@
# library files (both static and shared) will be installed.
LIBDIR := $(PREFIX)/lib

-EXLDFLAGS := -Wl,-rpath,$(LIBDIR)
+EXLDFLAGS := -Wl,-rpath,$(LIBDIR),$(POLYCOM_LDFLAGS)

# You should not need to change this. It is the directory into which the
# executable program files will be installed. BINDIR for programs that are
diff -Nuar lm_sensors-2.10.0.orig/lib/Module.mk lm_sensors-2.10.0/lib/Module.mk
--- lm_sensors-2.10.0.orig/lib/Module.mk2006-02-15 03:46:48.000000000 +0200
+++ lm_sensors-2.10.0/lib/Module.mk2006-08-16 12:56:55.383425297 +0300
@@ -66,7 +66,7 @@
# How to create the shared library
ifdef SYSFS_SUPPORT
$(MODULE_DIR)/$(LIBSHLIBNAME): $(LIBSHOBJECTS)
-$(CC) -shared -Wl,-soname,$(LIBSHSONAME) -o $$ $^ -lc -lm -lsysfs
+$(CC) -shared -Wl,-soname,$(LIBSHSONAME) $(EXLDFLAGS) -o $$ $^ -lc -lm -lsysfs
else
$(MODULE_DIR)/$(LIBSHLIBNAME): $(LIBSHOBJECTS)
$(CC) -shared -Wl,-soname,$(LIBSHSONAME) -o $$ $^ -lc -lm
```

module-init-tools-3.2-pre1-uclibc.diff

```
#
# Submitted-By: Robert Schwebel, 2005-05-08
# Committed-By: Robert Schwebel
#
# Error:
```



```
#
# config.sub claims that this is no valid configuration.
#
# Description:
#
# config.sub claims that this is no valid configuration.
#
# State:
#
# unknown
#

diff -urN module-init-tools-3.2-pre1/config.sub module-init-tools-3.2-pre1-ptx/config.sub
--- module-init-tools-3.2-pre1/config.sub2002-11-26 02:38:55.000000000 +0100
+++ module-init-tools-3.2-pre1-ptx/config.sub2005-05-08 13:34:24.000000000 +0200
@@ -281,10 +281,10 @@
     basic_machine=$basic_machine-pc
 ;;
 # Object if more than one company name word.
 -*_*)
 -echo Invalid configuration \`$1\': machine \`$basic_machine\' not recognized 1>&2
 -exit 1
 -;;
 +#*_*_*)
 +#echo Invalid configuration \`$1\': machine \`$basic_machine\' not recognized 1>&2
 +#exit 1
 +#;;
 # Recognize the basic CPU types with company name.
 580-* \
 | a29k-* \
@@ -1099,7 +1099,7 @@
 # The portable systems comes first.
 # Each alternative MUST END IN A *, to match a version number.
```

```
# -sysv* is not here because it comes later, after sysvr4.
--gnu* | -bsd* | -mach* | -minix* | -genix* | -ultrix* | -irix* \
+-gnu* | -bsd* | -mach* | -minix* | -genix* | -ultrix* | -irix* | -uclibc* \
  | -*vms* | -sco* | -esix* | -isc* | -aix* | -sunos | -sunos[34]* \
  | -hpux* | -unos* | -osf* | -luna* | -dgux* | -solaris* | -sym* \
  | -amigaos* | -amigados* | -msdos* | -newsos* | -unicos* | -aof* \
```

module-init-tools-oldtoolsname.diff

```
#
# Submitted-By: Marc Kleine-Budde, 2005-04-21
#
# Error:
#
# when using module-init-tools's depmod against 2.4 kernels, it will
# try to call 'depmod.old', regardless it's own name.
#
# Description:
#
# e.g.: arm-softfloat-linux-gnu-depmod calls depmod.old
# the target-tuple gets lost. This patch fixes that.
# now arm-softfloat-linux-gnu-depmod will call
# arm-softfloat-linux-gnu-depmod.old
#
diff -ru module-init-tools-3.2-pre1-orig/depmod.c module-init-tools-3.2-pre1/depmod.c
--- module-init-tools-3.2-pre1-orig/depmod.c2005-01-18 06:54:59.000000000 +0100
+++ module-init-tools-3.2-pre1/depmod.c2005-04-21 12:51:29.000000000 +0200
@@ -212,15 +212,10 @@

static void exec_old_depmod(char *argv[])
{
-char *sep;
-char pathname[strlen(argv[0])+1];
```

```
-char oldname[strlen("depmod") + strlen(argv[0]) + sizeof(".old")];  
-  
-memset(pathname, 0, strlen(argv[0])+1);  
-sep = strrchr(argv[0], '/');  
-if (sep)  
-memcpy(pathname, argv[0], sep - argv[0]+1);  
-sprintf(oldname, "%s%s.old", pathname, "depmod");  
+char *old = ".old";  
+char oldname[strlen(argv[0]) + strlen(old) + 1];  
+  
+sprintf(oldname, "%s%s", argv[0], old);  
  
/* Recursion detection: we need an env var since we can't  
   change argv[0] (as older modutils uses it to determine
```

nano-1.2.4_cross_compile.diff

```
--- nano-1.2.4/configure.orig2005-12-04 11:19:38.000000000 +0200  
+++ nano-1.2.4/configure2005-12-04 11:16:32.000000000 +0200  
@@ -5683,7 +5683,7 @@  
if test $ac_cv_header_regex_h = yes; then  
    echo "$as_me:$LINENO: checking for broken regexec" >&5  
    echo $ECHO_N "checking for broken regexec... $ECHO_C" >&6  
-   if test "$cross_compiling" = yes; then  
+   if test "$cross_compiling" = hack_yes; then  
        { { echo "$as_me:$LINENO: error: cannot run test program while cross compiling  
See `config.log' for more details." >&5  
        echo "$as_me: error: cannot run test program while cross compiling
```

generic-tic-xcompile.diff

```
#  
# this patch tweaks ncurses not to use the compiled version of tic,  
# but a version found on the host system
```

```
#
# very helpfull if we are crosscompiling
#
# source: a perl one-liner in eric andersen's buildroot
#      ncurses.make by Ken Restivo
#
# http://www.uclibc.org/cgi-bin/cvsweb/*checkout*/buildroot/make/
# ncurses.mk?rev=1.27
#
# converted to a patch
#
diff -ruN ncurses-5.3-orig/misc/run_tic.in ncurses-5.3/misc/run_tic.in
--- ncurses-5.3-orig/misc/run_tic.inFri Sep 26 00:27:10 2003
+++ ncurses-5.3/misc/run_tic.inFri Sep 26 00:28:57 2003
@@ -111,7 +111,7 @@
problems for older ncurses applications.

EOF
-if ( $srcdir/shlib tic$suffix -s -o $TERMINFO $source )
+if ( /usr/bin/tic -s -o $TERMINFO $source )
then
echo '** built new '$TERMINFO
else

*/
```

ntp-4.2.0-linux-2.6.14.patch

```
diff -Nuar ntp-4.2.0.orig/include/ntp_config.h ntp-4.2.0/include/ntp_config.h
--- ntp-4.2.0.orig/include/ntp_config.h2003-09-09 18:42:31.000000000 +0300
+++ ntp-4.2.0/include/ntp_config.h2005-12-20 10:15:40.000000000 +0200
@@ -14,51 +14,51 @@
/*
 * Types of entries we understand.
```

```
*/
-#define CONFIG_UNKNOWN0
+#define NTP_CONFIG_UNKNOWN0

/*
 * Command keywords
 */
-#define CONFIG_PEER1
-#define CONFIG_SERVER2
-#define CONFIG_AUTOMAX3
-#define CONFIG_DRIFTFILE4
-#define CONFIG_BROADCAST5
-#define CONFIG_BROADCASTCLIENT6
-#define CONFIG_AUTHENTICATE7
-#define CONFIG_KEYS8
-#define CONFIG_REVOKE9
-#define CONFIG_PPS10
-#define CONFIG_RESTRICT11
-#define CONFIG_BDELAY12
-#define CONFIG_TRUSTEDKEY13
-#define CONFIG_REQUESTKEY14
-#define CONFIG_CONTROLKEY15
-#define CONFIG_TRAP16
-#define CONFIG_FUDGE17
-#define CONFIG_TINKER18
-#define CONFIG_STATSDIR19
-#define CONFIG_FILEGEN20
-#define CONFIG_STATISTICS21
-#define CONFIG_PIDFILE22
-#define CONFIG_SETVAR23
-#define CONFIG_DISCARD24
-#define CONFIG_ADJ25
-#define CONFIG_MULTICASTCLIENT26
```

```
-#define CONFIG_ENABLE27
-#define CONFIG_DISABLE28
-#define CONFIG_PHONE29
-#define CONFIG_LOGFILE30
-#define CONFIG_LOGCONFIG31
-#define CONFIG_MANYCASTCLIENT32
-#define CONFIG_MANYCASTSERVER33
-#define CONFIG_TOS34
-#define CONFIG_TTL35
-#define CONFIG_INCLUDEFILE 36
-#define CONFIG_KEYSDIR37
-#define CONFIG_CDELAY38
+#define NTP_CONFIG_PEER1
+#define NTP_CONFIG_SERVER2
+#define NTP_CONFIG_AUTOMAX3
+#define NTP_CONFIG_DRIFTFILE4
+#define NTP_CONFIG_BROADCAST5
+#define NTP_CONFIG_BROADCASTCLIENT6
+#define NTP_CONFIG_AUTHENTICATE7
+#define NTP_CONFIG_KEYS8
+#define NTP_CONFIG_REVOKE9
+#define NTP_CONFIG_PPS10
+#define NTP_CONFIG_RESTRICT11
+#define NTP_CONFIG_BDELAY12
+#define NTP_CONFIG_TRUSTEDKEY13
+#define NTP_CONFIG_REQUESTKEY14
+#define NTP_CONFIG_CONTROLKEY15
+#define NTP_CONFIG_TRAP16
+#define NTP_CONFIG_FUDGE17
+#define NTP_CONFIG_TINKER18
+#define NTP_CONFIG_STATS DIR19
+#define NTP_CONFIG_FILEGEN20
+#define NTP_CONFIG_STATISTICS21
```

```
+#define NTP_CONFIG_PIDFILE22
+#define NTP_CONFIG_SETVAR23
+#define NTP_CONFIG_DISCARD24
+#define NTP_CONFIG_ADJ25
+#define NTP_CONFIG_MULTICASTCLIENT26
+#define NTP_CONFIG_ENABLE27
+#define NTP_CONFIG_DISABLE28
+#define NTP_CONFIG_PHONE29
+#define NTP_CONFIG_LOGFILE30
+#define NTP_CONFIG_LOGCONFIG31
+#define NTP_CONFIG_MANYCASTCLIENT32
+#define NTP_CONFIG_MANYCASTSERVER33
+#define NTP_CONFIG_TOS34
+#define NTP_CONFIG_TTL35
+#define NTP_CONFIG_INCLUDEFILE 36
+#define NTP_CONFIG_KEYSDIR37
+#define NTP_CONFIG_CDELAY38
#ifdef OPENSSEL
-#define CONFIG_CRYPT039
+#define NTP_CONFIG_CRYPT039
#endif /* OPENSSEL */

/*
diff -Nuar ntp-4.2.0.orig/ntpd/ntp_config.c ntp-4.2.0/ntpd/ntp_config.c
--- ntp-4.2.0.orig/ntpd/ntp_config.c2003-10-07 10:35:14.000000000 +0200
+++ ntp-4.2.0/ntpd/ntp_config.c2005-12-20 10:16:29.000000000 +0200
@@ -61,46 +61,46 @@
 * Command keywords
 */
static struct keyword keywords[] = {
-{"automax", CONFIG_AUTOMAX },
-{"broadcast", CONFIG_BROADCAST },
-{"broadcastclient", CONFIG_BROADCASTCLIENT },
```

```
-{ "broadcastdelay",CONFIG_BDELAY },
-{ "calldelay",CONFIG_CDELAY},
+{ "automax",NTP_CONFIG_AUTOMAX },
+{ "broadcast",NTP_CONFIG_BROADCAST },
+{ "broadcastclient",NTP_CONFIG_BROADCASTCLIENT },
+{ "broadcastdelay",NTP_CONFIG_BDELAY },
+{ "calldelay",NTP_CONFIG_CDELAY},
#ifdef OPENSLL
-{ "crypto",CONFIG_CRYPT0 },
+{ "crypto",NTP_CONFIG_CRYPT0 },
#endif /* OPENSLL */
-{ "controlkey",CONFIG_CONTROLKEY },
-{ "disable",CONFIG_DISABLE },
-{ "driftfile",CONFIG_DRIFTFILE },
-{ "enable",CONFIG_ENABLE },
-{ "filegen",CONFIG_FILEGEN },
-{ "fudge",CONFIG_FUDGE },
-{ "includefile",CONFIG_INCLUDEFILE },
-{ "keys",CONFIG_KEYS },
-{ "keysdir",CONFIG_KEYSDIR },
-{ "logconfig",CONFIG_LOGCONFIG },
-{ "logfile",CONFIG_LOGFILE },
-{ "manycastclient",CONFIG_MANYCASTCLIENT },
-{ "manycastserver",CONFIG_MANYCASTSERVER },
-{ "multicastclient",CONFIG_MULTICASTCLIENT },
-{ "peer",CONFIG_PEER },
-{ "phone",CONFIG_PHONE },
-{ "pidfile",CONFIG_PIDFILE },
-{ "discard",CONFIG_DISCARD },
-{ "requestkey",CONFIG_REQUESTKEY },
-{ "restrict",CONFIG_RESTRICT },
-{ "revoke",CONFIG_REVOKE },
-{ "server",CONFIG_SERVER },
```



```
-{ "setvar",CONFIG_SETVAR },
-{ "statistics",CONFIG_STATISTICS },
-{ "statsdir",CONFIG_STATSDIR },
-{ "tick",CONFIG_ADJ },
-{ "tinker",CONFIG_TINKER },
-{ "tos",CONFIG_TOS },
-{ "trap",CONFIG_TRAP },
-{ "trustedkey",CONFIG_TRUSTEDKEY },
-{ "ttl",CONFIG_TTL },
-{ "",CONFIG_UNKNOWN }
+{ "controlkey",NTP_CONFIG_CONTROLKEY },
+{ "disable",NTP_CONFIG_DISABLE },
+{ "driftfile",NTP_CONFIG_DRIFTFILE },
+{ "enable",NTP_CONFIG_ENABLE },
+{ "filegen",NTP_CONFIG_FILEGEN },
+{ "fudge",NTP_CONFIG_FUDGE },
+{ "includefile",NTP_CONFIG_INCLUDEFILE },
+{ "keys",NTP_CONFIG_KEYS },
+{ "keysdir",NTP_CONFIG_KEYS DIR },
+{ "logconfig",NTP_CONFIG_LOGCONFIG },
+{ "logfile",NTP_CONFIG_LOGFILE },
+{ "manycastclient",NTP_CONFIG_MANYCASTCLIENT },
+{ "manycastserver",NTP_CONFIG_MANYCASTSERVER },
+{ "multicastclient",NTP_CONFIG_MULTICASTCLIENT },
+{ "peer",NTP_CONFIG_PEER },
+{ "phone",NTP_CONFIG_PHONE },
+{ "pidfile",NTP_CONFIG_PIDFILE },
+{ "discard",NTP_CONFIG_DISCARD },
+{ "requestkey",NTP_CONFIG_REQUESTKEY },
+{ "restrict",NTP_CONFIG_RESTRICT },
+{ "revoke",NTP_CONFIG_REVOKE },
+{ "server",NTP_CONFIG_SERVER },
+{ "setvar",NTP_CONFIG_SETVAR },
```

```
+{ "statistics",NTP_CONFIG_STATISTICS },
+{ "statsdir",NTP_CONFIG_STATS DIR },
+{ "tick",NTP_CONFIG_ADJ },
+{ "tinker",NTP_CONFIG_TINKER },
+{ "tos",NTP_CONFIG_TOS },
+{ "trap",NTP_CONFIG_TRAP },
+{ "trustedkey",NTP_CONFIG_TRUSTEDKEY },
+{ "ttl",NTP_CONFIG_TTL },
+{ "",NTP_CONFIG_UNKNOWN }
};
```

```
/*
@@ -118,7 +118,7 @@
{ "prefer",CONF_MOD_PREFER },
{ "ttl",CONF_MOD_TTL }, /* NTP peers */
{ "version",CONF_MOD_VERSION },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};
```

```
/*
@@ -138,7 +138,7 @@
{ "notrust",CONF_RES_NOTRUST },
{ "ntpport",CONF_RES_NTPPORT },
{ "version",CONF_RES_VERSION },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};
```

```
/*
@@ -147,7 +147,7 @@
static struct keyword trap_keywords[] = {
{ "port",CONF_TRAP_PORT },
```

```
{ "interface",CONF_TRAP_INTERFACE },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};
```

```
/*
@@ -162,7 +162,7 @@
{ "stratum",CONF_FDG_STRATUM },
{ "time1",CONF_FDG_TIME1 },
{ "time2",CONF_FDG_TIME2 },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};
```

```
/*
@@ -175,7 +175,7 @@
{ "link",CONF_FGEN_FLAG_LINK },
{ "nolink",CONF_FGEN_FLAG_NOLINK },
{ "type",CONF_FGEN_TYPE },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};
```

```
/*
@@ -189,7 +189,7 @@
{ "pid",FILEGEN_PID },
{ "week",FILEGEN_WEEK },
{ "year",FILEGEN_YEAR },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};
```

```
/*
```

```
@@ -204,7 +204,7 @@
{ "ntp",PROTO_NTP },
{ "pps",PROTO_PPS },
{ "stats",PROTO_FILEGEN },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};

/*
@@ -214,7 +214,7 @@
{ "average",CONF_DISCARD_AVERAGE },
{ "minimum",CONF_DISCARD_MINIMUM },
{ "monitor",CONF_DISCARD_MONITOR },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};

/*
@@ -228,7 +228,7 @@
{ "allan",CONF_CLOCK_ALLAN },
{ "huffpuff",CONF_CLOCK_HUFFPUFF },
{ "freq",CONF_CLOCK_FREQ },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};

/*
@@ -240,7 +240,7 @@
{ "floor",CONF_TOS_FLOOR },
{ "ceiling",CONF_TOS_CEILING },
{ "cohort",CONF_TOS_COHORT },
-{ "",CONFIG_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
```

```
};

#ifdef OPENSSEL
@@ -257,7 +257,7 @@
{ "pw",CONF_CRYPTOP_W },
{ "randfile",CONF_CRYPTOP_RAND },
{ "sign",CONF_CRYPTOP_SIGN },
-{ "",CONFIGN_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};
#endif /* OPENSSEL */

@@ -268,7 +268,7 @@
static struct keyword addr_type[] = {
{ "-4",CONF_ADDR_IPV4 },
{ "-6",CONF_ADDR_IPV6 },
-{ "",CONFIGN_UNKNOWN }
+{ "",NTP_CONFIG_UNKNOWN }
};

/*
@@ -492,7 +492,7 @@
char line[MAXLINE];
char *(tokens[MAXTOKENS]);
int ntokens = 0;
-int tok = CONFIGN_UNKNOWN;
+int tok = NTP_CONFIG_UNKNOWN;
struct interface *localaddr;
struct refclockstat clock_stat;
FILEGEN *filegen;
@@ -580,7 +580,7 @@
tok = gettokens_netinfo(config_netinfo, tokens, &ntokens);
#endif /* HAVE_NETINFO */
```

```
-if (tok == CONFIG_UNKNOWN) {
+if (tok == NTP_CONFIG_UNKNOWN) {
    if (includelevel > 0) {
    fclose(fp[includelevel--]);
    continue;
@@ -590,15 +590,15 @@
    }

    switch(tok) {
-   case CONFIG_PEER:
-   case CONFIG_SERVER:
-   case CONFIG_MANYCASTCLIENT:
-   case CONFIG_BROADCAST:
+   case NTP_CONFIG_PEER:
+   case NTP_CONFIG_SERVER:
+   case NTP_CONFIG_MANYCASTCLIENT:
+   case NTP_CONFIG_BROADCAST:
-if (tok == CONFIG_PEER)
+if (tok == NTP_CONFIG_PEER)
    hmode = MODE_ACTIVE;
-else if (tok == CONFIG_SERVER)
+else if (tok == NTP_CONFIG_SERVER)
    hmode = MODE_CLIENT;
-else if (tok == CONFIG_MANYCASTCLIENT)
+else if (tok == NTP_CONFIG_MANYCASTCLIENT)
    hmode = MODE_CLIENT;
else
    hmode = MODE_BROADCAST;
@@ -644,8 +644,8 @@
    * and unicast address for manycastclient!
    */
    if (peeraddr.ss_family == AF_INET) {
```

```

-if (((tok == CONFIG_SERVER) | |
-   (tok == CONFIG_PEER)) &&
+if (((tok == NTP_CONFIG_SERVER) | |
+   (tok == NTP_CONFIG_PEER)) &&
#ifdef REFCLOCK
    !ISREFCLOCKADR(&peeraddr) &&
#endif
@@ -655,7 +655,7 @@
    stoa(&peeraddr));
    break;
}
-if ((tok == CONFIG_MANYCASTCLIENT) &&
+if ((tok == NTP_CONFIG_MANYCASTCLIENT) &&
    !IN_CLASSD(ntohl(((struct sockaddr_in*)&peeraddr)->sin_addr.s_addr))) {
    msyslog(LOG_ERR,
    "attempt to configure invalid address %s",
@@ -664,8 +664,8 @@
    }
}
else if(peeraddr.ss_family == AF_INET6) {
-    if (((tok == CONFIG_SERVER) | |
-        (tok == CONFIG_PEER)) &&
+    if (((tok == NTP_CONFIG_SERVER) | |
+        (tok == NTP_CONFIG_PEER)) &&
#ifdef REFCLOCK
        !ISREFCLOCKADR(&peeraddr) &&
#endif
    #endif
@@ -675,7 +675,7 @@
        stoa(&peeraddr));
        break;
    }
-    if ((tok == CONFIG_MANYCASTCLIENT) &&
+    if ((tok == NTP_CONFIG_MANYCASTCLIENT) &&

```

```
                                !IN6_IS_ADDR_MULTICAST(&((struct sockaddr_in6*)&peeraddr)-
>sin6_addr)) {
                                msyslog(LOG_ERR,
                                "attempt to configure in valid address %s",

@@ -803,7 +803,7 @@
    ttl = atoi(tokens[++i]);
    break;

-case CONFIG_UNKNOWN:
+case NTP_CONFIG_UNKNOWN:
    errflg = 1;
    break;
}
@@ -821,7 +821,7 @@
"configuration of %s failed",
stoa(&peeraddr));
}
- if (tok == CONFIG_MANYCASTCLIENT)
+ if (tok == NTP_CONFIG_MANYCASTCLIENT)
proto_config(PROTO_MULTICAST_ADD,
    0, 0., &peeraddr);

@@ -832,21 +832,21 @@
}
break;

- case CONFIG_DRIFTFILE:
+ case NTP_CONFIG_DRIFTFILE:
if (ntokens >= 2)
    stats_config(STATS_FREQ_FILE, tokens[1]);
else
    stats_config(STATS_FREQ_FILE, (char *)0);
break;
```



```
- case CONFIG_PIDFILE:
+ case NTP_CONFIG_PIDFILE:
if (ntokens >= 2)
    stats_config(STATS_PID_FILE, tokens[1]);
else
    stats_config(STATS_PID_FILE, (char *)0);
break;

- case CONFIG_INCLUDEFILE:
+ case NTP_CONFIG_INCLUDEFILE:
if (ntokens < 2) {
    msyslog(LOG_ERR, "includefile needs one argument");
    break;
@@ -865,7 +865,7 @@
fp[++includelevel] = includefile;
break;

- case CONFIG_LOGFILE:
+ case NTP_CONFIG_LOGFILE:
if (ntokens >= 2) {
    FILE *new_file;

@@ -889,7 +889,7 @@
    msyslog(LOG_ERR, "logfile needs one argument");
    break;

- case CONFIG_LOGCONFIG:
+ case NTP_CONFIG_LOGCONFIG:
for (i = 1; i < ntokens; i++)
{
    int add = 1;
@@ -924,12 +924,12 @@
```

```
}
break;

- case CONFIG_BROADCASTCLIENT:
+ case NTP_CONFIG_BROADCASTCLIENT:
proto_config(PROTO_BROADCASTCLIENT, 1, 0., NULL);
break;

- case CONFIG_MULTICASTCLIENT:
- case CONFIG_MANYCASTSERVER:
+ case NTP_CONFIG_MULTICASTCLIENT:
+ case NTP_CONFIG_MANYCASTSERVER:
if (ntokens > 1) {
  istory = 1;
  memset((char *)&peeraddr, 0, sizeof(peeraddr));
  @@ -962,19 +962,19 @@
} else
  proto_config(PROTO_MULTICAST_ADD,
    0, 0., NULL);
-if (tok == CONFIG_MULTICASTCLIENT)
+if (tok == NTP_CONFIG_MULTICASTCLIENT)
  sys_bclient = 1;
-else if (tok == CONFIG_MANYCASTSERVER)
+else if (tok == NTP_CONFIG_MANYCASTSERVER)
  sys_mancastserver = 1;
break;

- case CONFIG_KEYS:
+ case NTP_CONFIG_KEYS:
if (ntokens >= 2) {
  getauthkeys(tokens[1]);
}
break;
```

```
- case CONFIG_KEYSDIR:
+ case NTP_CONFIG_KEYSDIR:
if (ntokens < 2) {
    msyslog(LOG_ERR,
    "Keys directory name required");
@@ -984,7 +984,7 @@
    strcpy(keysdir, tokens[1]);
    break;

- case CONFIG_TINKER:
+ case NTP_CONFIG_TINKER:
for (i = 1; i < ntokens; i++) {
    int temp;
    double ftemp;
@@ -1030,7 +1030,7 @@
}
break;

- case CONFIG_TOS:
+ case NTP_CONFIG_TOS:
for (i = 1; i < ntokens; i++) {
    int temp;
    double ftemp;
@@ -1068,14 +1068,14 @@
}
break;

- case CONFIG_TTL:
+ case NTP_CONFIG_TTL:
for (i = 1; i < ntokens && i < MAX_TTL; i++) {
    sys_ttl[i - 1] = (u_char) atoi(tokens[i]);
    sys_ttlmax = i - 1;
```

```
}
break;

- case CONFIG_DISCARD:
+ case NTP_CONFIG_DISCARD:
for (i = 1; i < ntokens; i++) {
    int temp;

@@ -1109,17 +1109,17 @@
break;

#ifdef OPENSSSL
- case CONFIG_REVOKE:
+ case NTP_CONFIG_REVOKE:
if (ntokens >= 2)
    sys_revoke = (u_char) max(atoi(tokens[1]), KEY_REVOKE);
break;

- case CONFIG_AUTOMAX:
+ case NTP_CONFIG_AUTOMAX:
if (ntokens >= 2)
    sys_automax = 1 << max(atoi(tokens[1]), 10);
break;

- case CONFIG_CRYPT0:
+ case NTP_CONFIG_CRYPT0:
if (ntokens == 1) {
    crypto_config(CRYPTO_CONF_NONE, NULL);
    break;
@@ -1191,7 +1191,7 @@
break;
#endif /* OPENSSSL */
```

```
- case CONFIG_RESTRICT:
+ case NTP_CONFIG_RESTRICT:
if (ntokens < 2) {
msyslog(LOG_ERR, "restrict requires an address");
break;
@@ -1289,7 +1289,7 @@
peerversion |= RES_LIMITED;
break;

- case CONFIG_UNKNOWN:
+ case NTP_CONFIG_UNKNOWN:
errflg++;
break;
}
@@ -1301,7 +1301,7 @@
(int)peerkey, peerversion);
break;

- case CONFIG_BDELAY:
+ case NTP_CONFIG_BDELAY:
if (ntokens >= 2) {
double tmp;

@@ -1315,7 +1315,7 @@
}
break;

- case CONFIG_CDELAY:
+ case NTP_CONFIG_CDELAY:
if (ntokens >= 2) {
u_long ui;

@@ -1327,7 +1327,7 @@
```

```
}
break;

- case CONFIG_TRUSTEDKEY:
+ case NTP_CONFIG_TRUSTEDKEY:
for (i = 1; i < ntokens; i++) {
    keyid_t tkey;

@@ -1342,7 +1342,7 @@
}
break;

- case CONFIG_REQUESTKEY:
+ case NTP_CONFIG_REQUESTKEY:
if (ntokens >= 2) {
    if (!atouint(tokens[1], &ul)) {
        msyslog(LOG_ERR,
@@ -1363,7 +1363,7 @@
    }
    break;

- case CONFIG_CONTROLKEY:
+ case NTP_CONFIG_CONTROLKEY:
if (ntokens >= 2) {
    keyid_t ckey;

@@ -1378,7 +1378,7 @@
}
break;

- case CONFIG_TRAP:
+ case NTP_CONFIG_TRAP:
if (ntokens < 2) {
```

```
msyslog(LOG_ERR,  
"no address for trap command, line ignored");  
@@ -1452,7 +1452,7 @@  
    }  
    break;  
  
-case CONFIG_UNKNOWN:  
+case NTP_CONFIG_UNKNOWN:  
    errflg++;  
    break;  
    }  
@@ -1472,7 +1472,7 @@  
    }  
    break;  
  
- case CONFIG_FUDGE:  
+ case NTP_CONFIG_FUDGE:  
    if (ntokens < 2) {  
        msyslog(LOG_ERR,  
"no address for fudge command, line ignored");  
@@ -1577,7 +1577,7 @@  
        clock_stat.flags |= c;  
        break;  
  
- case CONFIG_UNKNOWN:  
+ case NTP_CONFIG_UNKNOWN:  
    errflg = -1;  
    break;  
    }  
@@ -1596,12 +1596,12 @@  
#endif  
break;
```

```
- case CONFIG_STATSDIR:
+ case NTP_CONFIG_STATSDIR:
if (ntokens >= 2)
stats_config(STATS_STATSDIR,tokens[1]);
break;

- case CONFIG_STATISTICS:
+ case NTP_CONFIG_STATISTICS:
for (i = 1; i < ntokens; i++) {
filegen = filegen_get(tokens[i]);

@@ -1620,7 +1620,7 @@
}
break;

- case CONFIG_FILEGEN:
+ case NTP_CONFIG_FILEGEN:
if (ntokens < 2) {
msyslog(LOG_ERR,
"no id for filegen command, line ignored");
@@ -1667,7 +1667,7 @@
}
peerkey = matchkey(tokens[++i],
fgen_types, 1);
-if (peerkey == CONFIG_UNKNOWN) {
+if (peerkey == NTP_CONFIG_UNKNOWN) {
msyslog(LOG_ERR,
"filegen %s unknown type \"%s\"",
tokens[1], tokens[i]);
@@ -1698,7 +1698,7 @@
(u_char)peerkey, (u_char)peerflags);
break;
```



```
- case CONFIG_SETVAR:
+ case NTP_CONFIG_SETVAR:
  if (ntokens < 2) {
    msyslog(LOG_ERR,
      "no value for setvar command - line ignored");
    @@ -1713,12 +1713,12 @@
  }
  break;

- case CONFIG_ENABLE:
+ case NTP_CONFIG_ENABLE:
  for (i = 1; i < ntokens; i++) {
    int flag;

    flag = matchkey(tokens[i], flags_keywords, 1);
    -if (flag == CONFIG_UNKNOWN) {
    +if (flag == NTP_CONFIG_UNKNOWN) {
      msyslog(LOG_ERR,
        "enable unknown flag %s",
        tokens[i]);
      @@ -1729,12 +1729,12 @@
    }
    break;

- case CONFIG_DISABLE:
+ case NTP_CONFIG_DISABLE:
  for (i = 1; i < ntokens; i++) {
    int flag;

    flag = matchkey(tokens[i], flags_keywords, 1);
    -if (flag == CONFIG_UNKNOWN) {
    +if (flag == NTP_CONFIG_UNKNOWN) {
      msyslog(LOG_ERR,
```

```
"disable unknown flag %s",
tokens[i]);
@@ -1745,7 +1745,7 @@
}
break;

- case CONFIG_PHONE:
+ case NTP_CONFIG_PHONE:
for (i = 1; i < ntokens && i < MAXPHONE; i++) {
(void)strncpy(sys_phone[i - 1],
tokens[i], MAXDIAL);
@@ -1753,7 +1753,7 @@
sys_phone[i - 1][0] = '\0';
break;

- case CONFIG_ADJ: {
+ case NTP_CONFIG_ADJ: {
double ftemp;

sscanf(tokens[1], "%lf", &ftemp);
@@ -1923,7 +1923,7 @@
}

/* No list; we're done here. */
- if (!val_list) return CONFIG_UNKNOWN;
+ if (!val_list) return NTP_CONFIG_UNKNOWN;

/*
* We have a list of values for the current property.
@@ -2001,7 +2001,7 @@
}
if (cp == NULL) {
*ntokens = 0;
```

```

-return CONFIG_UNKNOWN; /* hack. Is recognized as EOF */
+return NTP_CONFIG_UNKNOWN; /* hack. Is recognized as EOF */
}

/*
@@ -2029,7 +2029,7 @@
 */
*ntokens = ntok + 1;
ntok = matchkey(tokenlist[0], keywords, 1);
-if (ntok == CONFIG_UNKNOWN)
+if (ntok == NTP_CONFIG_UNKNOWN)
goto again;
return ntok;
}
@@ -2047,12 +2047,12 @@
)
{
for (;;) {
-if (keys->keytype == CONFIG_UNKNOWN) {
+if (keys->keytype == NTP_CONFIG_UNKNOWN) {
if (complain)
msyslog(LOG_ERR,
"configure: keyword \"%s\" unknown, line ignored",
word);
-return CONFIG_UNKNOWN;
+return NTP_CONFIG_UNKNOWN;
}
if (STRSAME(word, keys->text))
return keys->keytype;
diff -Nuar ntp-4.2.0.orig/ntpd/ntp_config.c.orig ntp-4.2.0/ntpd/ntp_config.c.orig
--- ntp-4.2.0.orig/ntpd/ntp_config.c.orig 1970-01-01 02:00:00.000000000 +0200
+++ ntp-4.2.0/ntpd/ntp_config.c.orig 2003-10-07 10:35:14.000000000 +0200
@@ -0,0 +1,2371 @@

```

```
+/*
+ * ntp_config.c - read and apply configuration information
+ */
+#ifdef HAVE_CONFIG_H
+# include <config.h>
+#endif
+
+#ifdef HAVE_NETINFO
+# include <netinfo/ni.h>
+#endif
+
+#include "ntpd.h"
+#include "ntp_io.h"
+#include "ntp_unixtime.h"
+#include "ntp_refclock.h"
+#include "ntp_filegen.h"
+#include "ntp_stdlib.h"
+#include "ntp_config.h"
+#include "ntp_cmdargs.h"
+
+#include <stdio.h>
+#include <ctype.h>
+#ifdef HAVE_SYS_PARAM_H
+#include <sys/param.h>
+#endif
+#include <signal.h>
+#ifndef SIGCHLD
+# define SIGCHLD SIGCLD
+#endif
+#if !defined(VMS)
+# ifdef HAVE_SYS_WAIT_H
+# include <sys/wait.h>
+# endif
```

```
+ #endif /* VMS */
+
+ #ifdef SYS_WINNT
+ #include <io.h>
+ extern HANDLE ResolverThreadHandle;
+ #endif /* SYS_WINNT */
+
+ #include <netdb.h>
+
+ extern int priority_done;
+
+ /*
+  * These routines are used to read the configuration file at
+  * startup time. An entry in the file must fit on a single line.
+  * Entries are processed as multiple tokens separated by white space
+  * Lines are considered terminated when a '#' is encountered. Blank
+  * lines are ignored.
+  */
+ /*
+  * Translation table - keywords to function index
+  */
+ struct keyword {
+   const char *text;
+   int keytype;
+ };
+
+ /*
+  * Command keywords
+  */
+ static struct keyword keywords[] = {
+   { "automax", CONFIG_AUTOMAX },
+   { "broadcast", CONFIG_BROADCAST },
+   { "broadcastclient", CONFIG_BROADCASTCLIENT },
```

```
+{ "broadcastdelay",CONFIG_BDELAY },
+{ "calldelay",CONFIG_CDELAY},
+#ifdef OPENSSSL
+{ "crypto",CONFIG_CRYPTO },
+#endif /* OPENSSSL */
+{ "controlkey",CONFIG_CONTROLKEY },
+{ "disable",CONFIG_DISABLE },
+{ "driftfile",CONFIG_DRIFTFILE },
+{ "enable",CONFIG_ENABLE },
+{ "filegen",CONFIG_FILEGEN },
+{ "fudge",CONFIG_FUDGE },
+{ "includefile",CONFIG_INCLUDEFILE },
+{ "keys",CONFIG_KEYS },
+{ "keysdir",CONFIG_KEYSDIR },
+{ "logconfig",CONFIG_LOGCONFIG },
+{ "logfile",CONFIG_LOGFILE },
+{ "manycastclient",CONFIG_MANYCASTCLIENT },
+{ "manycastserver",CONFIG_MANYCASTSERVER },
+{ "multicastclient",CONFIG_MULTICASTCLIENT },
+{ "peer",CONFIG_PEER },
+{ "phone",CONFIG_PHONE },
+{ "pidfile",CONFIG_PIDFILE },
+{ "discard",CONFIG_DISCARD },
+{ "requestkey",CONFIG_REQUESTKEY },
+{ "restrict",CONFIG_RESTRICT },
+{ "revoke",CONFIG_REVOKE },
+{ "server",CONFIG_SERVER },
+{ "setvar",CONFIG_SETVAR },
+{ "statistics",CONFIG_STATISTICS },
+{ "statsdir",CONFIG_STATS DIR },
+{ "tick",CONFIG_ADJ },
+{ "tinker",CONFIG_TINKER },
+{ "tos",CONFIG_TOS },
```

```
+{ "trap",CONFIG_TRAP },
+{ "trustedkey",CONFIG_TRUSTEDKEY },
+{ "ttl",CONFIG_TTL },
+{ "",CONFIG_UNKNOWN }
+};
+
+/*
+ * "peer", "server", "broadcast" modifier keywords
+ */
+staticstruct keyword mod_keywords[] = {
+{ "autokey",CONF_MOD_SKEY },
+{ "burst",CONF_MOD_BURST },
+{ "iburst",CONF_MOD_IBURST },
+{ "key",CONF_MOD_KEY },
+{ "maxpoll",CONF_MOD_MAXPOLL },
+{ "minpoll",CONF_MOD_MINPOLL },
+{ "mode",CONF_MOD_MODE }, /* refclocks */
+{ "noselect",CONF_MOD_NOSELECT },
+{ "prefer",CONF_MOD_PREFER },
+{ "ttl",CONF_MOD_TTL }, /* NTP peers */
+{ "version",CONF_MOD_VERSION },
+{ "",CONFIG_UNKNOWN }
+};
+
+/*
+ * "restrict" modifier keywords
+ */
+staticstruct keyword res_keywords[] = {
+{ "ignore",CONF_RES_IGNORE },
+{ "limited",CONF_RES_LIMITED },
+{ "kod",CONF_RES_DEMOBILIZE },
+{ "lowpriotrap",CONF_RES_LPTRAP },
+{ "mask",CONF_RES_MASK },
```

```
+{ "nomodify",CONF_RES_NOMODIFY },
+{ "nopeer",CONF_RES_NOPEER },
+{ "noquery",CONF_RES_NOQUERY },
+{ "noserve",CONF_RES_NOSERVE },
+{ "notrap",CONF_RES_NOTRAP },
+{ "notrust",CONF_RES_NOTRUST },
+{ "ntpport",CONF_RES_NTPPORT },
+{ "version",CONF_RES_VERSION },
+{ "",CONFIG_UNKNOWN }
+};
+
+/*
+ * "trap" modifier keywords
+ */
+staticstruct keyword trap_keywords[] = {
+{ "port",CONF_TRAP_PORT },
+{ "interface",CONF_TRAP_INTERFACE },
+{ "",CONFIG_UNKNOWN }
+};
+
+/*
+ * "fudge" modifier keywords
+ */
+staticstruct keyword fudge_keywords[] = {
+{ "flag1",CONF_FDG_FLAG1 },
+{ "flag2",CONF_FDG_FLAG2 },
+{ "flag3",CONF_FDG_FLAG3 },
+{ "flag4",CONF_FDG_FLAG4 },
+{ "refid",CONF_FDG_REFID },
+{ "stratum",CONF_FDG_STRATUM },
+{ "time1",CONF_FDG_TIME1 },
+{ "time2",CONF_FDG_TIME2 },
+{ "",CONFIG_UNKNOWN }
```



```
+};  
+  
+/*  
+ * "filegen" modifier keywords  
+ */  
+staticstruct keyword filegen_keywords[] = {  
+{ "disable",CONF_FGEN_FLAG_DISABLE },  
+{ "enable",CONF_FGEN_FLAG_ENABLE },  
+{ "file",CONF_FGEN_FILE },  
+{ "link",CONF_FGEN_FLAG_LINK },  
+{ "nolink",CONF_FGEN_FLAG_NOLINK },  
+{ "type",CONF_FGEN_TYPE },  
+{ "",CONFIG_UNKNOWN }  
+};  
+  
+/*  
+ * "type" modifier keywords  
+ */  
+staticstruct keyword fgen_types[] = {  
+{ "age",FILEGEN_AGE },  
+{ "day",FILEGEN_DAY },  
+{ "month",FILEGEN_MONTH },  
+{ "none",FILEGEN_NONE },  
+{ "pid",FILEGEN_PID },  
+{ "week",FILEGEN_WEEK },  
+{ "year",FILEGEN_YEAR },  
+{ "",CONFIG_UNKNOWN}  
+};  
+  
+/*  
+ * "enable", "disable" modifier keywords  
+ */  
+static struct keyword flags_keywords[] = {
```

```
+{ "auth",PROTO_AUTHENTICATE },
+{ "bclient",PROTO_BROADCLIENT },
+{ "calibrate",PROTO_CAL },
+{ "kernel",PROTO_KERNEL },
+{ "monitor",PROTO_MONITOR },
+{ "ntp",PROTO_NTP },
+{ "pps",PROTO_PPS },
+{ "stats",PROTO_FILEGEN },
+{ "",CONFIG_UNKNOWN }
+};
+
+/*
+ * "discard" modifier keywords
+ */
+static struct keyword discard_keywords[] = {
+{ "average",CONF_DISCARD_AVERAGE },
+{ "minimum",CONF_DISCARD_MINIMUM },
+{ "monitor",CONF_DISCARD_MONITOR },
+{ "",CONFIG_UNKNOWN }
+};
+
+/*
+ * "tinker" modifier keywords
+ */
+static struct keyword tinker_keywords[] = {
+{ "step",CONF_CLOCK_MAX },
+{ "panic",CONF_CLOCK_PANIC },
+{ "dispersion",CONF_CLOCK_PHI },
+{ "stepout",CONF_CLOCK_MINSTEP },
+{ "allan",CONF_CLOCK_ALLAN },
+{ "huffpuff",CONF_CLOCK_HUFFPUFF },
+{ "freq",CONF_CLOCK_FREQ },
+{ "",CONFIG_UNKNOWN }
```

```
+};  
+  
+/*  
+ * "tos" modifier keywords  
+ */  
+static struct keyword tos_keywords[] = {  
+{ "minclock",CONF_TOS_MINCLOCK },  
+{ "minsane",CONF_TOS_MINSANE },  
+{ "floor",CONF_TOS_FLOOR },  
+{ "ceiling",CONF_TOS_CEILING },  
+{ "cohort",CONF_TOS_COHORT },  
+{ "",CONFIG_UNKNOWN }  
+};  
+  
+#ifdef OPENSSEL  
+/*  
+ * "crypto" modifier keywords  
+ */  
+static struct keyword crypto_keywords[] = {  
+{ "cert",CONF_CRYPTOCERT },  
+{ "gqpar",CONF_CRYPTOGQPAR },  
+{ "host",CONF_CRYPTORSA },  
+{ "iffpar",CONF_CRYPTOIFFPAR },  
+{ "leap",CONF_CRYPTOLEAP },  
+{ "mvpar",CONF_CRYPTOMVPAR },  
+{ "pw",CONF_CRYPTOPW },  
+{ "randfile",CONF_CRYPTORAND },  
+{ "sign",CONF_CRYPTOSIGN },  
+{ "",CONFIG_UNKNOWN }  
+};  
+#endif /* OPENSSEL */  
+  
+/*
```

```
+ * Address type selection, IPv4 or IPv4.
+ * Used on various lines.
+ */
+static struct keyword addr_type[] = {
+{ "-4",CONF_ADDR_IPV4 },
+{ "-6",CONF_ADDR_IPV6 },
+{ "",CONFIG_UNKNOWN }
+};
+
+/*
+ * "logconfig" building blocks
+ */
+struct masks {
+const char *name;
+unsigned long mask;
+};
+
+static struct masks logcfg_class[] = {
+{ "clock",NLOG_OCLOCK },
+{ "peer",NLOG_OPEER },
+{ "sync",NLOG_OSYNC },
+{ "sys",NLOG_OSYS },
+{ (char *)0,0 }
+};
+
+static struct masks logcfg_item[] = {
+{ "info",NLOG_INFO },
+{ "allinfo",NLOG_SYINFO | NLOG_PEERINFO | NLOG_CLOCKINFO | NLOG_SYNCINFO },
+{ "events",NLOG_EVENT },
+{ "allevents",
+NLOG_SYSEVENT | NLOG_PEEREVENT | NLOG_CLOCKEVENT | NLOG_SYNCEVENT },
+{ "status",NLOG_STATUS },
```

```

+{ "allstatus",
NLOG_SYSSTATUS | NLOG_PEERSTATUS | NLOG_CLOCKSTATUS | NLOG_SYNCSTATUS
},
+{ "statistics", NLOG_STATIST },
+{ "allstatistics",
NLOG_SYSSTATIST | NLOG_PEERSTATIST | NLOG_CLOCKSTATIST | NLOG_SYNCSTATIS
T },
+{ "allclock",
(NLOG_INFO | NLOG_STATIST | NLOG_EVENT | NLOG_STATUS) << NLOG_OCLOCK },
+{ "allpeer",
(NLOG_INFO | NLOG_STATIST | NLOG_EVENT | NLOG_STATUS) << NLOG_OPEER },
+{ "allsys",
(NLOG_INFO | NLOG_STATIST | NLOG_EVENT | NLOG_STATUS) << NLOG_OSYS },
+{ "allsync",
(NLOG_INFO | NLOG_STATIST | NLOG_EVENT | NLOG_STATUS) << NLOG_OSYNC },
+{ "all", NLOG_SYSMASK | NLOG_PEERMASK | NLOG_CLOCKMASK | NLOG_SYNCMASK
},
+{ (char *)0, 0 }
+};
+
+/*
+ * Limits on things
+ */
+#define MAXTOKENS20 /* 20 tokens on line */
+#define MAXLINE1024 /* maximum length of line */
+#define MAXPHONE5 /* maximum number of phone strings */
+#define MAXPPS20 /* maximum length of PPS device string */
+#define MAXINCLUDELEVEL5 /* maximum include file levels */
+
+/*
+ * Miscellaneous macros
+ */
+#define STRSAME(s1, s2) (*s1 == *s2) && strcmp((s1), (s2)) == 0
+#define ISEOL(c) ((c) == '#' || (c) == '\n' || (c) == '\0')
+#define ISSPACE(c) ((c) == ' ' || (c) == '\t')

```

```
+ #define STREQ(a, b)(*(a) == *(b) && strcmp((a), (b)) == 0)
+
+ #define KEY_TYPE_MD54
+
+ /*
+  * File descriptor used by the resolver save routines, and temporary file
+  * name.
+  */
+ int call_resolver = 1; /* ntp-genkeys sets this to 0, for example */
+ static FILE *res_fp;
+ #ifndef SYS_WINNT
+ static char res_file[20]; /* enough for /tmp/ntpXXXXXX\0 */
+ #define RES_TEMPFILE "/tmp/ntpXXXXXX"
+ #else
+ static char res_file[MAX_PATH];
+ #endif /* SYS_WINNT */
+
+ /*
+  * Definitions of things either imported from or exported to outside
+  */
+ char const *progname;
+ charsys_phone[MAXPHONE][MAXDIAL]; /* ACTS phone numbers */
+ char *keysdir = NTP_KEYSDIR; /* crypto keys directory */
+ charpps_device[MAXPPS + 1]; /* PPS device name */
+ #if defined(HAVE_SCHED_SETSCHEDULER)
+ intconfig_priority_override = 0;
+ intconfig_priority;
+ #endif
+
+ const char *config_file;
+ #ifdef HAVE_NETINFO
+ struct netinfo_config_state *config_netinfo = NULL;
+ int check_netinfo = 1;
```

```
+ #endif /* HAVE_NETINFO */
+ #ifdef SYS_WINNT
+ char *alt_config_file;
+ LPTSTR temp;
+ char config_file_storage[MAX_PATH];
+ char alt_config_file_storage[MAX_PATH];
+ #endif /* SYS_WINNT */
+
+ #ifdef HAVE_NETINFO
+ /*
+  * NetInfo configuration state
+  */
+ struct netinfo_config_state {
+ void *domain; /* domain with config */
+ ni_id config_dir; /* ID config dir */
+ int prop_index; /* current property */
+ int val_index; /* current value */
+ char **val_list; /* value list */
+ };
+ #endif
+
+ /*
+  * Function prototypes
+  */
+ static unsigned long get_pfxmatch P((char **, struct masks *));
+ static unsigned long get_match P((char *, struct masks *));
+ static unsigned long get_logmask P((char *));
+ #ifdef HAVE_NETINFO
+ static struct netinfo_config_state *get_netinfo_config P((void));
+ static void free_netinfo_config P((struct netinfo_config_state *));
+ static int gettokens_netinfo P((struct netinfo_config_state *, char **, int *));
+ #endif
+ static int gettokens P((FILE *, char *, char **, int *));
```

```
+static int matchkey P((char *, struct keyword *, int));
+static int getnetnum P((const char *, struct sockaddr_storage *, int));
+static void save_resolve P((char *, int, int, int, int, u_int, int,
+    keyid_t, u_char *));
+static void do_resolve_internal P((void));
+static void abort_resolve P((void));
+#if !defined(VMS) && !defined(SYS_WINNT)
+static RETSIGTYPE catchchild P((int));
+#endif /* VMS */
+
+/*
+ * get_pfxmatch - find value for prefixmatch
+ * and update char * accordingly
+ */
+static unsigned long
+get_pfxmatch(
+char ** s,
+struct masks *m
+)
+{
+while (m->name) {
+if (strcmp(*s, m->name, strlen(m->name)) == 0) {
+*s += strlen(m->name);
+return m->mask;
+} else {
+m++;
+}
+}
+return 0;
+}
+
+/*
+ * get_match - find logmask value
```



```
+ */
+static unsigned long
+get_match(
+char *s,
+struct masks *m
+)
+{
+while (m->name) {
+if (strcmp(s, m->name) == 0) {
+return m->mask;
+} else {
+m++;
+}
+}
+return 0;
+}
+
+/*
+ * get_logmask - build bitmask for ntp_syslogmask
+ */
+static unsigned long
+get_logmask(
+char *s
+)
+{
+char *t;
+unsigned long offset;
+unsigned long mask;
+
+t = s;
+offset = get_pfxmatch(&t, logcfg_class);
+mask = get_match(t, logcfg_item);
+
```

```
+if (mask)
+return mask << offset;
+else
+msyslog(LOG_ERR, "logconfig: illegal argument %s - ignored", s);
+
+return 0;
+}
+
+
+/*
+ * getconfig - get command line options and read the configuration file
+ */
+void
+getconfig(
+int argc,
+char *argv[])
+{
+register int i;
+int c;
+int errflg;
+int istory;
+int peerversion;
+int minpoll;
+int maxpoll;
+int ttl;
+long stratum;
+unsigned long ul;
+keyid_t peerkey;
+_char *peerkeystr;
+_long fudgeflag;
+_int peerflags;
+int hmode;
```

```
+struct sockaddr_storage peeraddr;
+struct sockaddr_storage maskaddr;
+FILE *fp[MAXINCLUDELEVEL+1];
+FILE *includefile;
+int includelevel = 0;
+char line[MAXLINE];
+char *(tokens[MAXTOKENS]);
+int ntokens = 0;
+int tok = CONFIG_UNKNOWN;
+struct interface *localaddr;
+struct refclockstat clock_stat;
+FILEGEN *filegen;
+
+/*
+ * Initialize, initialize
+ */
+errflg = 0;
+/* HMS: don't initialize debug to 0 here! */
+#ifndef SYS_WINNT
+config_file = CONFIG_FILE;
+#else
+temp = CONFIG_FILE;
+if (!ExpandEnvironmentStrings((LPCTSTR)temp, (LPTSTR)config_file_storage,
(DWORD)sizeof(config_file_storage))) {
+msyslog(LOG_ERR, "ExpandEnvironmentStrings CONFIG_FILE failed: %m\n");
+exit(1);
+}
+config_file = config_file_storage;
+
+temp = ALT_CONFIG_FILE;
+if (!ExpandEnvironmentStrings((LPCTSTR)temp, (LPTSTR)alt_config_file_storage,
(DWORD)sizeof(alt_config_file_storage))) {
+msyslog(LOG_ERR, "ExpandEnvironmentStrings ALT_CONFIG_FILE failed: %m\n");
```

```
+exit(1);
+}
+alt_config_file = alt_config_file_storage;
+
+#endif /* SYS_WINNT */
+progname = argv[0];
+res_fp = NULL;
+memset((char *)sys_phone, 0, sizeof(sys_phone));
+ntp_syslogmask = NLOG_SYNCMASK; /* set more via logconfig */
+
+/*
+ * install a non default variable with this daemon version
+ */
+(void) sprintf(line, "daemon_version=\ "%s\ """, Version);
+set_sys_var(line, strlen(line)+1, RO);
+
+/*
+ * Say how we're setting the time of day
+ */
+(void) sprintf(line, "settimeofday=\ "%s\ """, set_tod_using);
+set_sys_var(line, strlen(line)+1, RO);
+
+/*
+ * Initialize the loop.
+ */
+loop_config(LOOP_DRIFTINIT, 0);
+
+getCmdOpts(argc, argv);
+
+if (
+  (fp[0] = fopen(FindConfig(config_file), "r")) == NULL
+#ifdef HAVE_NETINFO
+  /* If there is no config_file, try NetInfo. */
```

```

+  && check_netinfo && !(config_netinfo = get_netinfo_config())
+#endif /* HAVE_NETINFO */
+  ) {
+fprintf(stderr, "getconfig: Couldn't open <%s>\n", FindConfig(config_file));
+msyslog(LOG_INFO, "getconfig: Couldn't open <%s>", FindConfig(config_file));
+#ifdef SYS_WINNT
+/* Under WinNT try alternate_config_file name, first NTP.CONF, then NTP.INI */
+
+if ((fp[0] = fopen(FindConfig(alt_config_file), "r")) == NULL) {
+
+/*
+ * Broadcast clients can sometimes run without
+ * a configuration file.
+ */
+
+fprintf(stderr, "getconfig: Couldn't open <%s>\n", FindConfig(alt_config_file));
+msyslog(LOG_INFO, "getconfig: Couldn't open <%s>", FindConfig(alt_config_file));
+return;
+}
+#else /* not SYS_WINNT */
+return;
+#endif /* not SYS_WINNT */
+}
+
+for (;;) {
+if (fp[includelevel])
+tok = gettokens(fp[includelevel], line, tokens, &ntokens);
+#ifdef HAVE_NETINFO
+else
+tok = gettokens_netinfo(config_netinfo, tokens, &ntokens);
+#endif /* HAVE_NETINFO */
+
+if (tok == CONFIG_UNKNOWN) {

```

```
+ if (includelevel > 0) {
+fclose(fp[includelevel--]);
+continue;
+ } else {
+break;
+ }
+}
+
+switch(tok) {
+ case CONFIG_PEER:
+ case CONFIG_SERVER:
+ case CONFIG_MANYCASTCLIENT:
+ case CONFIG_BROADCAST:
+if (tok == CONFIG_PEER)
+ hmode = MODE_ACTIVE;
+else if (tok == CONFIG_SERVER)
+ hmode = MODE_CLIENT;
+else if (tok == CONFIG_MANYCASTCLIENT)
+ hmode = MODE_CLIENT;
+else
+ hmode = MODE_BROADCAST;
+
+if (ntokens < 2) {
+msyslog(LOG_ERR,
+"No address for %s, line ignored",
+tokens[0]);
+break;
+}
+
+istart = 1;
+memset((char *)&peeraddr, 0, sizeof(peeraddr));
+switch (matchkey(tokens[istart], addr_type, 0)) {
+case CONF_ADDR_IPV4:
```

```
+peeraddr.ss_family = AF_INET;
+istart++;
+break;
+case CONF_ADDR_IPV6:
+peeraddr.ss_family = AF_INET6;
+istart++;
+break;
+}
+
+if (!getnetnum(tokens[istart], &peeraddr, 0)) {
+errflg = -1;
+} else {
+errflg = 0;
+
+if (
+#ifdef REFCLOCK
+!ISREFCLOCKADR(&peeraddr) &&
+#endif
+ISBADADR(&peeraddr)) {
+msyslog(LOG_ERR,
+"attempt to configure invalid address %s",
+stoa(&peeraddr));
+break;
+}
+/*
+ * Shouldn't be able to specify multicast
+ * address for server/peer!
+ * and unicast address for manycastclient!
+ */
+if (peeraddr.ss_family == AF_INET) {
+if (((tok == CONFIG_SERVER) ||
+ (tok == CONFIG_PEER)) &&
+#ifdef REFCLOCK
```

```
+ !ISREFCLOCKADR(&peeraddr) &&
+ #endif
+ IN_CLASSD(ntohl(((struct sockaddr_in*)&peeraddr)->sin_addr.s_addr))) {
+ msyslog(LOG_ERR,
+ "attempt to configure invalid address %s",
+ stoa(&peeraddr));
+ break;
+ }
+ if ((tok == CONFIG_MANYCASTCLIENT) &&
+ !IN_CLASSD(ntohl(((struct sockaddr_in*)&peeraddr)->sin_addr.s_addr))) {
+ msyslog(LOG_ERR,
+ "attempt to configure invalid address %s",
+ stoa(&peeraddr));
+ break;
+ }
+ }
+ else if(peeraddr.ss_family == AF_INET6) {
+         if (((tok == CONFIG_SERVER) ||
+         (tok == CONFIG_PEER)) &&
+ #ifdef REFCLOCK
+         !ISREFCLOCKADR(&peeraddr) &&
+ #endif
+         IN6_IS_ADDR_MULTICAST(&((struct sockaddr_in6*)&peeraddr)-
+ >sin6_addr)) {
+         msyslog(LOG_ERR,
+         "attempt to configure in valid address %s",
+         stoa(&peeraddr));
+         break;
+     }
+     if ((tok == CONFIG_MANYCASTCLIENT) &&
+         !IN6_IS_ADDR_MULTICAST(&((struct sockaddr_in6*)&peeraddr)-
+ >sin6_addr)) {
+         msyslog(LOG_ERR,
```



```
+             "attempt to configure in valid address %s",
+             stoa(&peeraddr));
+             break;
+}
+}
+}
+
+peerversion = NTP_VERSION;
+minpoll = NTP_MINDPOLL;
+maxpoll = NTP_MAXDPOLL;
+peerkey = 0;
+peerkeystr = (u_char *)"";
+peerflags = 0;
+ttl = 0;
+istart++;
+for (i = istart; i < ntokens; i++)
+  switch (matchkey(tokens[i], mod_keywords, 1)) {
+case CONF_MOD_VERSION:
+  if (i >= ntokens-1) {
+    msyslog(LOG_ERR,
+    "peer/server version requires an argument");
+    errflg = 1;
+    break;
+  }
+  peerversion = atoi(tokens[++i]);
+  if ((u_char)peerversion > NTP_VERSION
+  || (u_char)peerversion < NTP_OLDVERSION) {
+    msyslog(LOG_ERR,
+    "inappropriate version number %s, line ignored",
+    tokens[i]);
+    errflg = 1;
+  }
+  break;
```

```
+
+case CONF_MOD_KEY:
+  if (i >= ntokens-1) {
+    msyslog(LOG_ERR,
+    "key: argument required");
+    errflg = 1;
+    break;
+  }
+  peerkey = (int)atol(tokens[++i]);
+  peerflags |= FLAG_AUTHENABLE;
+  break;
+
+case CONF_MOD_MINPOLL:
+  if (i >= ntokens-1) {
+    msyslog(LOG_ERR,
+    "minpoll: argument required");
+    errflg = 1;
+    break;
+  }
+  minpoll = atoi(tokens[++i]);
+  if (minpoll < NTP_MINPOLL) {
+    msyslog(LOG_INFO,
+    "minpoll: provided value (%d) is below minimum (%d)",
+    minpoll, NTP_MINPOLL);
+minpoll = NTP_MINPOLL;
+  }
+  break;
+
+case CONF_MOD_MAXPOLL:
+  if (i >= ntokens-1) {
+    msyslog(LOG_ERR,
+    "maxpoll: argument required"
+    );
```

```
+   errflg = 1;
+   break;
+ }
+   maxpoll = atoi(tokens[++i]);
+   if (maxpoll > NTP_MAXPOLL) {
+       msyslog(LOG_INFO,
+       "maxpoll: provided value (%d) is above maximum (%d)",
+       maxpoll, NTP_MAXPOLL);
+       maxpoll = NTP_MAXPOLL;
+   }
+   break;
+
+case CONF_MOD_PREFER:
+   peerflags |= FLAG_PREFER;
+   break;
+
+case CONF_MOD_NOSELECT:
+   peerflags |= FLAG_NOSELECT;
+   break;
+
+case CONF_MOD_BURST:
+   peerflags |= FLAG_BURST;
+   break;
+
+case CONF_MOD_IBURST:
+   peerflags |= FLAG_IBURST;
+   break;
+#ifdef OPENSSEL
+case CONF_MOD_SKEY:
+   peerflags |= FLAG_SKEY |
+   FLAG_AUTHENABLE;
+   break;
+#endif /* OPENSSEL */
```

```
+
+case CONF_MOD_TTL:
+  if (i >= ntokens-1) {
+msyslog(LOG_ERR,
+  "ttl: argument required");
+    errflg = 1;
+    break;
+  }
+  ttl = atoi(tokens[++i]);
+  if (ttl >= MAX_TTL) {
+msyslog(LOG_ERR,
+  "ttl: invalid argument");
+errflg = 1;
+  }
+  break;
+
+case CONF_MOD_MODE:
+  if (i >= ntokens-1) {
+msyslog(LOG_ERR,
+  "mode: argument required");
+errflg = 1;
+break;
+  }
+  ttl = atoi(tokens[++i]);
+  break;
+
+case CONFIG_UNKNOWN:
+  errflg = 1;
+  break;
+  }
+if (minpoll > maxpoll) {
+msyslog(LOG_ERR,
+  "config error: minpoll > maxpoll");
```

```
+errflg = 1;
+}
+if (errflg == 0) {
+  if (peer_config(&peeraddr,
+ANY_INTERFACE_CHOOSE(&peeraddr), hmode,
+peerversion, minpoll, maxpoll, peerflags,
+ttl, peerkey, peerkeystr) == 0) {
+msyslog(LOG_ERR,
+"configuration of %s failed",
+stoa(&peeraddr));
+  }
+  if (tok == CONFIG_MANYCASTCLIENT)
+proto_config(PROTO_MULTICAST_ADD,
+  0, 0., &peeraddr);
+
+} else if (errflg == -1) {
+save_resolve(tokens[1], hmode, peerversion,
+  minpoll, maxpoll, peerflags, ttl,
+  peerkey, peerkeystr);
+}
+break;
+
+  case CONFIG_DRIFTFILE:
+if (ntokens >= 2)
+  stats_config(STATS_FREQ_FILE, tokens[1]);
+else
+  stats_config(STATS_FREQ_FILE, (char *)0);
+break;
+
+  case CONFIG_PIDFILE:
+if (ntokens >= 2)
+  stats_config(STATS_PID_FILE, tokens[1]);
+else
```

```
+ stats_config(STATS_PID_FILE, (char *)0);
+break;
+
+ case CONFIG_INCLUDEFILE:
+if (ntokens < 2) {
+  syslog(LOG_ERR, "includefile needs one argument");
+  break;
+}
+if (includelevel >= MAXINCLUDELEVEL) {
+  fprintf(stderr, "getconfig: Maximum include file level exceeded.\n");
+  syslog(LOG_INFO, "getconfig: Maximum include file level exceeded.");
+  break;
+}
+includefile = fopen(FindConfig(tokens[1]), "r");
+if (includefile == NULL) {
+  fprintf(stderr, "getconfig: Couldn't open <%s>\n", FindConfig(tokens[1]));
+  syslog(LOG_INFO, "getconfig: Couldn't open <%s>", FindConfig(tokens[1]));
+  break;
+}
+fp[++includelevel] = includefile;
+break;
+
+ case CONFIG_LOGFILE:
+if (ntokens >= 2) {
+FILE *new_file;
+
+new_file = fopen(tokens[1], "a");
+if (new_file != NULL) {
+NLOG(NLOG_SYSINFO) /* conditional if clause for conditional syslog */
+  syslog(LOG_NOTICE, "logging to file %s", tokens[1]);
+if (syslog_file != NULL &&
+  fileno(syslog_file) != fileno(new_file))
+  (void)fclose(syslog_file);
```

```
+
+syslog_file = new_file;
+syslogit = 0;
+}
+else
+  msyslog(LOG_ERR,
+  "Cannot open log file %s",
+  tokens[1]);
+}
+else
+  msyslog(LOG_ERR, "logfile needs one argument");
+break;
+
+ case CONFIG_LOGCONFIG:
+for (i = 1; i < ntokens; i++)
+{
+int add = 1;
+int equals = 0;
+char * s = &tokens[i][0];
+
+switch (*s) {
+  case '+':
+  case '-':
+  case '=':
+add = *s == '+';
+equals = *s == '=';
+s++;
+break;
+
+  default:
+break;
+}
+if (equals) {
```

```
+ntp_syslogmask = get_logmask(s);
+} else {
+if (add) {
+ntp_syslogmask |= get_logmask(s);
+} else {
+ntp_syslogmask &= ~get_logmask(s);
+}
+}
+#ifdef DEBUG
+if (debug)
+  printf("ntp_syslogmask = 0x%08lx (%s)\n", ntp_syslogmask, tokens[i]);
+#endif
+}
+break;
+
+ case CONFIG_BROADCASTCLIENT:
+proto_config(PROTO_BROADCASTCLIENT, 1, 0., NULL);
+break;
+
+ case CONFIG_MULTICASTCLIENT:
+ case CONFIG_MANYCASTSERVER:
+if (ntokens > 1) {
+istart = 1;
+memset((char *)&peeraddr, 0, sizeof(peeraddr));
+switch (matchkey(tokens[istart],
+  addr_type, 0)) {
+case CONF_ADDR_IPV4:
+peeraddr.ss_family = AF_INET;
+istart++;
+break;
+case CONF_ADDR_IPV6:
+peeraddr.ss_family = AF_INET6;
+istart++;
```



```
+break;
+}
+/*
+ * Abuse maskaddr to store the preferred ip
+ * version.
+ */
+memset((char *)&maskaddr, 0, sizeof(maskaddr));
+maskaddr.ss_family = peeraddr.ss_family;
+
+for (i = istart; i < ntokens; i++) {
+memset((char *)&peeraddr, 0,
+ sizeof(peeraddr));
+peeraddr.ss_family = maskaddr.ss_family;
+if (getnetnum(tokens[i], &peeraddr, 1))
+ proto_config(PROTO_MULTICAST_ADD,
+ 0, 0., &peeraddr);
+}
+} else
+ proto_config(PROTO_MULTICAST_ADD,
+ 0, 0., NULL);
+if (tok == CONFIG_MULTICASTCLIENT)
+sys_bclient = 1;
+else if (tok == CONFIG_MANYCASTSERVER)
+sys_maycastserver = 1;
+break;
+
+ case CONFIG_KEYS:
+if (ntokens >= 2) {
+getauthkeys(tokens[1]);
+}
+break;
+
+ case CONFIG_KEYSDIR:
```

```
+if (ntokens < 2) {
+  msyslog(LOG_ERR,
+"Keys directory name required");
+  break;
+}
+keysdir = emalloc(strlen(tokens[1]) + 1);
+strcpy(keysdir, tokens[1]);
+break;
+
+ case CONFIG_TINKER:
+for (i = 1; i < ntokens; i++) {
+  int temp;
+  double ftemp;
+
+  temp = matchkey(tokens[i++], tinker_keywords, 1);
+  if (i > ntokens - 1) {
+msyslog(LOG_ERR,
+ "tinker: missing argument");
+errflg++;
+break;
+  }
+  sscanf(tokens[i], "%lf", &ftemp);
+  switch(temp) {
+
+ case CONF_CLOCK_MAX:
+
+         loop_config(LOOP_MAX, ftemp);
+break;
+
+ case CONF_CLOCK_PANIC:
+loop_config(LOOP_PANIC, ftemp);
+break;
+
+ case CONF_CLOCK_PHI:
```

```
+loop_config(LOOP_PHI, ftemp);
+break;
+
+ case CONF_CLOCK_MINSTEP:
+loop_config(LOOP_MINSTEP, ftemp);
+break;
+
+ case CONF_CLOCK_ALLAN:
+loop_config(LOOP_ALLAN, ftemp);
+break;
+
+ case CONF_CLOCK_HUFFPUFF:
+loop_config(LOOP_HUFFPUFF, ftemp);
+break;
+
+ case CONF_CLOCK_FREQ:
+loop_config(LOOP_FREQ, ftemp);
+break;
+ }
+}
+break;
+
+ case CONFIG_TOS:
+for (i = 1; i < ntokens; i++) {
+ int temp;
+ double ftemp;
+
+ temp = matchkey(tokens[i++], tos_keywords, 1);
+ if (i > ntokens - 1) {
+msyslog(LOG_ERR,
+ "tinker: missing argument");
+errflg++;
+break;
```

```
+ }
+ sscanf(tokens[i], "%lf", &ftemp);
+ switch(temp) {
+
+ case CONF_TOS_MINCLOCK:
+proto_config(PROTO_MINCLOCK, 0, ftemp, NULL);
+break;
+
+ case CONF_TOS_MINSANE:
+proto_config(PROTO_MINSANE, 0, ftemp, NULL);
+break;
+
+ case CONF_TOS_FLOOR:
+proto_config(PROTO_FLOOR, 0, ftemp, NULL);
+break;
+
+ case CONF_TOS_CEILING:
+proto_config(PROTO_CEILING, 0, ftemp, NULL);
+break;
+
+ case CONF_TOS_COHORT:
+proto_config(PROTO_COHORT, 0, ftemp, NULL);
+break;
+ }
+}
+break;
+
+ case CONFIG_TTL:
+for (i = 1; i < ntokens && i < MAX_TTL; i++) {
+ sys_ttl[i - 1] = (u_char) atoi(tokens[i]);
+ sys_ttlmax = i - 1;
+}
+break;
```

```
+
+ case CONFIG_DISCARD:
+for (i = 1; i < ntokens; i++) {
+ int temp;
+
+ temp = matchkey(tokens[i++],
+discard_keywords, 1);
+ if (i > ntokens - 1) {
+msyslog(LOG_ERR,
+ "discard: missing argument");
+errflg++;
+break;
+ }
+ switch(temp) {
+ case CONF_DISCARD_AVERAGE:
+res_avg_interval = atoi(tokens[i++]);
+break;
+
+ case CONF_DISCARD_MINIMUM:
+res_min_interval = atoi(tokens[i++]);
+break;
+
+ case CONF_DISCARD_MONITOR:
+mon_age = atoi(tokens[i++]);
+break;
+
+ default:
+msyslog(LOG_ERR,
+ "discard: unknown keyword");
+break;
+ }
+}
+break;
```

```
+
+ #ifdef OPENSLL
+  case CONFIG_REVOKE:
+    if (ntokens >= 2)
+      sys_revoke = (u_char) max(atoi(tokens[1]), KEY_REVOKE);
+    break;
+
+  case CONFIG_AUTOMAX:
+    if (ntokens >= 2)
+      sys_automax = 1 << max(atoi(tokens[1]), 10);
+    break;
+
+  case CONFIG_CRYPT0:
+    if (ntokens == 1) {
+      crypto_config(CRYPTO_CONF_NONE, NULL);
+      break;
+    }
+    for (i = 1; i < ntokens; i++) {
+      int temp;
+
+      temp = matchkey(tokens[i++],
+        crypto_keywords, 1);
+      if (i > ntokens - 1) {
+        msyslog(LOG_ERR,
+          "crypto: missing argument");
+        errflg++;
+        break;
+      }
+      switch(temp) {
+
+      case CONF_CRYPT0_CERT:
+        crypto_config(CRYPTO_CONF_CERT,
+          tokens[i]);
```

```
+break;
+
+ case CONF_CRYPTTO_RSA:
+crypto_config(CRYPTO_CONF_PRIV,
+ tokens[i]);
+break;
+
+ case CONF_CRYPTTO_IFFPAR:
+crypto_config(CRYPTO_CONF_IFFPAR,
+ tokens[i]);
+break;
+
+ case CONF_CRYPTTO_GQPAR:
+crypto_config(CRYPTO_CONF_GQPAR,
+ tokens[i]);
+break;
+
+ case CONF_CRYPTTO_MVPAR:
+crypto_config(CRYPTO_CONF_MVPAR,
+ tokens[i]);
+break;
+
+ case CONF_CRYPTTO_LEAP:
+crypto_config(CRYPTO_CONF_LEAP,
+ tokens[i]);
+break;
+
+ case CONF_CRYPTTO_PW:
+crypto_config(CRYPTO_CONF_PW,
+ tokens[i]);
+break;
+
+ case CONF_CRYPTTO_RAND:
```

2-126


```
+break;
+}
+
+/*
+ * Assume default means an IPv4 address, except
+ * if forced by a -4 or -6.
+ */
+if (STREQ(tokens[istart], "default")) {
+if (peeraddr.ss_family == 0)
+peeraddr.ss_family = AF_INET;
+} else if (!getnetnum(tokens[istart], &peeraddr, 1))
+break;
+
+
+/*
+ * Use peerversion as flags, peerkey as mflags. Ick.
+ */
+peerversion = 0;
+peerkey = 0;
+errflg = 0;
+SET_HOSTMASK(&maskaddr, peeraddr.ss_family);
+istart++;
+for (i = istart; i < ntokens; i++) {
+switch (matchkey(tokens[i], res_keywords, 1)) {
+ case CONF_RES_MASK:
+if (i >= ntokens-1) {
+msyslog(LOG_ERR,
+"mask keyword needs argument");
+errflg++;
+break;
+}
+i++;
+if (!getnetnum(tokens[i], &maskaddr, 1))
+ errflg++;
```

```
+break;
+
+ case CONF_RES_IGNORE:
+peerversion |= RES_IGNORE;
+break;
+
+ case CONF_RES_NOSERVE:
+peerversion |= RES_DONTSERVE;
+break;
+
+ case CONF_RES_NOTRUST:
+peerversion |= RES_DONTTRUST;
+break;
+
+ case CONF_RES_NOQUERY:
+peerversion |= RES_NOQUERY;
+break;
+
+ case CONF_RES_NOMODIFY:
+peerversion |= RES_NOMODIFY;
+break;
+
+ case CONF_RES_NOPEER:
+peerversion |= RES_NOPEER;
+break;
+
+ case CONF_RES_NOTRAP:
+peerversion |= RES_NOTRAP;
+break;
+
+ case CONF_RES_LPTRAP:
+peerversion |= RES_LPTRAP;
+break;
```

```
+
+ case CONF_RES_NTPPORT:
+peerkey |= RESM_NTPONLY;
+break;
+
+ case CONF_RES_VERSION:
+peerversion |= RES_VERSION;
+break;
+
+ case CONF_RES_DEMOBILIZE:
+peerversion |= RES_DEMOBILIZE;
+break;
+
+ case CONF_RES_LIMITED:
+peerversion |= RES_LIMITED;
+break;
+
+ case CONFIG_UNKNOWN:
+errflg++;
+break;
+}
+}
+if (SOCKNUL(&peeraddr))
+ ANYSOCK(&maskaddr);
+if (!errflg)
+ hack_restrict(RESTRICT_FLAGS, &peeraddr, &maskaddr,
+ (int)peerkey, peerversion);
+break;
+
+ case CONFIG_BDELAY:
+if (ntokens >= 2) {
+double tmp;
+
+
```

```
+if (sscanf(tokens[1], "%lf", &tmp) != 1) {
+msyslog(LOG_ERR,
+"broadcastdelay value %s undecodable",
+tokens[1]);
+} else {
+proto_config(PROTO_BROADDELAY, 0, tmp, NULL);
+}
+}
+break;
+
+ case CONFIG_CDELAY:
+     if (ntokens >= 2) {
+         u_long ui;
+
+if (sscanf(tokens[1], "%ld", &ui) != 1)
+msyslog(LOG_ERR,
+ "illegal value - line ignored");
+else
+proto_config(PROTO_CALLDELAY, ui, 0, NULL);
+}
+break;
+
+ case CONFIG_TRUSTEDKEY:
+for (i = 1; i < ntokens; i++) {
+keyid_t tkey;
+
+
+tkey = atol(tokens[i]);
+if (tkey == 0) {
+msyslog(LOG_ERR,
+"trusted key %s unlikely",
+tokens[i]);
+} else {
+authtrust(tkey, 1);
```

```
+}
+}
+break;
+
+ case CONFIG_REQUESTKEY:
+if (ntokens >= 2) {
+if (!atouint(tokens[1], &ul)) {
+msyslog(LOG_ERR,
+ "%s is undecodable as request key",
+ tokens[1]);
+} else if (ul == 0) {
+msyslog(LOG_ERR,
+ "%s makes a poor request keyid",
+ tokens[1]);
+} else {
+ #ifdef DEBUG
+if (debug > 3)
+ printf(
+ "set info_auth_key to %08lx\n", ul);
+ #endif
+info_auth_keyid = (keyid_t)ul;
+}
+}
+break;
+
+ case CONFIG_CONTROLKEY:
+if (ntokens >= 2) {
+keyid_t ckey;
+
+ ckey = atol(tokens[1]);
+if (ckey == 0) {
+msyslog(LOG_ERR,
+ "%s makes a poor control keyid",
```

```
+tokens[1]);
+} else {
+ctl_auth_keyid = ckey;
+}
+}
+break;
+
+ case CONFIG_TRAP:
+if (ntokens < 2) {
+msyslog(LOG_ERR,
+"no address for trap command, line ignored");
+break;
+}
+istart = 1;
+memset((char *)&peeraddr, 0, sizeof(peeraddr));
+switch (matchkey(tokens[istart], addr_type, 0)) {
+case CONF_ADDR_IPV4:
+peeraddr.ss_family = AF_INET;
+istart++;
+break;
+case CONF_ADDR_IPV6:
+peeraddr.ss_family = AF_INET6;
+istart++;
+break;
+}
+
+if (!getnetnum(tokens[istart], &peeraddr, 1))
+ break;
+
+/*
+ * Use peerversion for port number. Barf.
+ */
+errflg = 0;
```

```
+peerversion = 0;
+localaddr = 0;
+istart++;
+for (i = istart; i < ntokens-1; i++)
+  switch (matchkey(tokens[i], trap_keywords, 1)) {
+case CONF_TRAP_PORT:
+  if (i >= ntokens-1) {
+  msyslog(LOG_ERR,
+  "trap port requires an argument");
+  errflg = 1;
+  break;
+  }
+  peerversion = atoi(tokens[++i]);
+  if (peerversion <= 0
+  || peerversion > 32767) {
+  msyslog(LOG_ERR,
+  "invalid port number %s, trap ignored",
+  tokens[i]);
+  errflg = 1;
+  }
+  break;
+
+case CONF_TRAP_INTERFACE:
+  if (i >= ntokens-1) {
+  msyslog(LOG_ERR,
+  "trap interface requires an argument");
+  errflg = 1;
+  break;
+  }
+
+  memset((char *)&maskaddr, 0,
+  sizeof(maskaddr));
+  maskaddr.ss_family = peeraddr.ss_family;
```

```
+ if (!getnetnum(tokens[++i],
+ &maskaddr, 1)) {
+   errflg = 1;
+   break;
+ }
+
+ localaddr = findinterface(&maskaddr);
+ if (localaddr == NULL) {
+   msyslog(LOG_ERR,
+ "can't find interface with address %s",
+ stoa(&maskaddr));
+   errflg = 1;
+ }
+   break;
+
+case CONFIG_UNKNOWN:
+   errflg++;
+   break;
+ }
+
+if (!errflg) {
+if (peerversion != 0)
+  ((struct sockaddr_in6*)&peeraddr)->sin6_port = htons( (u_short) peerversion);
+else
+  ((struct sockaddr_in6*)&peeraddr)->sin6_port = htons(TRAPPORT);
+if (localaddr == NULL)
+  localaddr = ANY_INTERFACE_CHOOSE(&peeraddr);
+if (!ctlsettrap(&peeraddr, localaddr, 0,
+NTP_VERSION))
+  msyslog(LOG_ERR,
+ "can't set trap for %s, no resources",
+ stoa(&peeraddr));
+}
```



```
+break;
+
+ case CONFIG_FUDGE:
+if (ntokens < 2) {
+msyslog(LOG_ERR,
+"no address for fudge command, line ignored");
+break;
+}
+memset((char *)&peeraddr, 0, sizeof(peeraddr));
+if (!getnetnum(tokens[1], &peeraddr, 1))
+ break;
+
+if (!ISREFCLOCKADR(&peeraddr)) {
+msyslog(LOG_ERR,
+"%s is inappropriate address for the fudge command, line ignored",
+stoa(&peeraddr));
+break;
+}
+
+memset((void *)&clock_stat, 0, sizeof clock_stat);
+fudgeflag = 0;
+errflg = 0;
+for (i = 2; i < ntokens-1; i++) {
+switch (c = matchkey(tokens[i],
+ fudge_keywords, 1)) {
+ case CONF_FDG_TIME1:
+if (sscanf(tokens[++i], "%lf",
+ &clock_stat.fudgetime1) != 1) {
+msyslog(LOG_ERR,
+"fudge %s time1 value in error",
+stoa(&peeraddr));
+errflg = i;
+break;
```

```
+}
+clock_stat.haveflags |= CLK_HAVETIME1;
+break;
+
+ case CONF_FDG_TIME2:
+if (sscanf(tokens[++i], "%lf",
+ &clock_stat.fudgetime2) != 1) {
+msyslog(LOG_ERR,
+"fudge %s time2 value in error",
+stoa(&peeraddr));
+errflg = i;
+break;
+}
+clock_stat.haveflags |= CLK_HAVETIME2;
+break;
+
+
+ case CONF_FDG_STRATUM:
+ if (!atoint(tokens[++i], &stratum))
+{
+msyslog(LOG_ERR,
+"fudge %s stratum value in error",
+stoa(&peeraddr));
+errflg = i;
+break;
+}
+clock_stat.fudgeval1 = stratum;
+clock_stat.haveflags |= CLK_HAVEVAL1;
+break;
+
+
+ case CONF_FDG_REFID:
+/* HMS: Endianness and 0 bytes? */
+/* XXX */
```

```
+strncpy((char *)&clock_stat.fudgeval2,  
+tokens[++i], 4);  
+clock_stat.haveflags |= CLK_HAVEVAL2;  
+break;  
+  
+ case CONF_FDG_FLAG1:  
+ case CONF_FDG_FLAG2:  
+ case CONF_FDG_FLAG3:  
+ case CONF_FDG_FLAG4:  
+if (!atouint(tokens[++i], &fudgeflag)  
+ || fudgeflag > 1) {  
+msyslog(LOG_ERR,  
+"fudge %s flag value in error",  
+stoa(&peeraddr));  
+errflg = i;  
+break;  
+}  
+switch(c) {  
+ case CONF_FDG_FLAG1:  
+c = CLK_FLAG1;  
+clock_stat.haveflags |=CLK_HAVEFLAG1;  
+break;  
+ case CONF_FDG_FLAG2:  
+c = CLK_FLAG2;  
+clock_stat.haveflags |=CLK_HAVEFLAG2;  
+break;  
+ case CONF_FDG_FLAG3:  
+c = CLK_FLAG3;  
+clock_stat.haveflags |=CLK_HAVEFLAG3;  
+break;  
+ case CONF_FDG_FLAG4:  
+c = CLK_FLAG4;  
+clock_stat.haveflags |=CLK_HAVEFLAG4;
```

```
+break;
+}
+if (fudgeflag == 0)
+  clock_stat.flags &= ~c;
+else
+  clock_stat.flags |= c;
+break;
+
+  case CONFIG_UNKNOWN:
+errflg = -1;
+break;
+}
+}
+
+#ifdef REFCLOCK
+/*
+ * If reference clock support isn't defined the
+ * fudge line will still be accepted and syntax
+ * checked, but will essentially do nothing.
+ */
+if (!errflg) {
+refclock_control(&peeraddr, &clock_stat,
+  (struct refclockstat *)0);
+}
+#endif
+break;
+
+  case CONFIG_STATSDIR:
+if (ntokens >= 2)
+stats_config(STATS_STATSDIR,tokens[1]);
+break;
+
+  case CONFIG_STATISTICS:
```

```
+for (i = 1; i < ntokens; i++) {
+filegen = filegen_get(tokens[i]);
+
+if (filegen == NULL) {
+msyslog(LOG_ERR,
+"no statistics named %s available",
+tokens[i]);
+continue;
+}
+#ifdef DEBUG
+if (debug > 3)
+ printf("enabling filegen for %s statistics \"%s%s\"\\n",
+ tokens[i], filegen->prefix, filegen->basename);
+#endif
+filegen->flag |= FGGEN_FLAG_ENABLED;
+}
+break;
+
+ case CONFIG_FILEGEN:
+if (ntokens < 2) {
+msyslog(LOG_ERR,
+"no id for filegen command, line ignored");
+break;
+}
+
+filegen = filegen_get(tokens[1]);
+if (filegen == NULL) {
+msyslog(LOG_ERR,
+"unknown filegen \"%s\" ignored",
+tokens[1]);
+break;
+}
+/*
```

```
+ * peerversion is (ab)used for filegen file (index)
+ * peerkey   is (ab)used for filegen type
+ * peerflags is (ab)used for filegen flags
+ */
+peerversion = 0;
+peerkey = filegen->type;
+peerflags = filegen->flag;
+errflg = 0;
+
+for (i = 2; i < ntokens; i++) {
+switch (matchkey(tokens[i],
+ filegen_keywords, 1)) {
+ case CONF_FGEN_FILE:
+if (i >= ntokens - 1) {
+msyslog(LOG_ERR,
+"filegen %s file requires argument",
+tokens[1]);
+errflg = i;
+break;
+}
+peerversion = ++i;
+break;
+ case CONF_FGEN_TYPE:
+if (i >= ntokens - 1) {
+msyslog(LOG_ERR,
+"filegen %s type requires argument",
+tokens[1]);
+errflg = i;
+break;
+}
+peerkey = matchkey(tokens[++i],
+ fgen_types, 1);
+if (peerkey == CONFIG_UNKNOWN) {
```

```
+msyslog(LOG_ERR,
+"filegen %s unknown type \"%s\"",
+tokens[1], tokens[i]);
+errflg = i;
+break;
+}
+break;
+
+ case CONF_FGEN_FLAG_LINK:
+peerflags |= FGEN_FLAG_LINK;
+break;
+
+ case CONF_FGEN_FLAG_NOLINK:
+peerflags &= ~FGEN_FLAG_LINK;
+break;
+
+ case CONF_FGEN_FLAG_ENABLE:
+peerflags |= FGEN_FLAG_ENABLED;
+break;
+
+ case CONF_FGEN_FLAG_DISABLE:
+peerflags &= ~FGEN_FLAG_ENABLED;
+break;
+}
+}
+if (!errflg)
+filegen_config(filegen, tokens[peerversion],
+ (u_char)peerkey, (u_char)peerflags);
+break;
+
+ case CONFIG_SETVAR:
+if (ntokens < 2) {
+msyslog(LOG_ERR,
```

```
+ "no value for setvar command - line ignored");
+ } else {
+   set_sys_var(tokens[1], strlen(tokens[1])+1,
+   (u_short) (RW |
+   (((ntokens > 2)
+   && !strcmp(tokens[2],
+   "default"))))
+   ? DEF
+   : 0));
+ }
+ break;
+
+ case CONFIG_ENABLE:
+ for (i = 1; i < ntokens; i++) {
+ int flag;
+
+ flag = matchkey(tokens[i], flags_keywords, 1);
+ if (flag == CONFIG_UNKNOWN) {
+ msyslog(LOG_ERR,
+ "enable unknown flag %s",
+ tokens[i]);
+ errflg = 1;
+ break;
+ }
+ proto_config(flag, 1, 0., NULL);
+ }
+ break;
+
+ case CONFIG_DISABLE:
+ for (i = 1; i < ntokens; i++) {
+ int flag;
+
+ flag = matchkey(tokens[i], flags_keywords, 1);
```



```
+if (flag == CONFIG_UNKNOWN) {
+msyslog(LOG_ERR,
+"disable unknown flag %s",
+tokens[i]);
+errflg = 1;
+break;
+}
+proto_config(flag, 0, 0., NULL);
+}
+break;
+
+ case CONFIG_PHONE:
+for (i = 1; i < ntokens && i < MAXPHONE; i++) {
+(void)strncpy(sys_phone[i - 1],
+ tokens[i], MAXDIAL);
+}
+sys_phone[i - 1][0] = '\0';
+break;
+
+ case CONFIG_ADJ: {
+ double ftemp;
+
+ sscanf(tokens[1], "%lf", &ftemp);
+ proto_config(PROTO_ADJ, 0, ftemp, NULL);
+}
+break;
+
+}
+}
+if (fp[0])
+(void)fclose(fp[0]);
+
+#ifdef HAVE_NETINFO
```

```
+if (config_netinfo)
+free_netinfo_config(config_netinfo);
+#endif /* HAVE_NETINFO */
+
+#if !defined(VMS) && !defined(SYS_VXWORKS)
+/* find a keyid */
+if (info_auth_keyid == 0)
+req_keyid = 65535;
+else
+req_keyid = info_auth_keyid;
+
+/* if doesn't exist, make up one at random */
+if (!authhavekey(req_keyid)) {
+char rankey[9];
+int j;
+
+for (i = 0; i < 8; i++)
+for (j = 1; j < 100; ++j) {
+rankey[i] = (char) (RANDOM & 0xff);
+if (rankey[i] != 0) break;
+}
+rankey[8] = 0;
+authusekey(req_keyid, KEY_TYPE_MD5, (u_char *)rankey);
+authtrust(req_keyid, 1);
+if (!authhavekey(req_keyid)) {
+msyslog(LOG_ERR, "getconfig: Couldn't generate a valid random key!");
+/* HMS: Should this be fatal? */
+}
+}
+
+/* save keyid so we will accept config requests with it */
+info_auth_keyid = req_keyid;
+#endif /* !defined(VMS) && !defined(SYS_VXWORKS) */
```

```
+
+if (res_fp != NULL) {
+if (call_resolver) {
+/*
+ * Need name resolution
+ */
+do_resolve_internal();
+}
+}
+}
+
+
+
+#ifdef HAVE_NETINFO
+
+/*
+ * get_netinfo_config - find the nearest NetInfo domain with an ntp
+ * configuration and initialize the configuration state.
+ */
+static struct netinfo_config_state *
+get_netinfo_config()
+{
+ni_status status;
+void *domain;
+ni_id config_dir;
+    struct netinfo_config_state *config;
+
+if (ni_open(NULL, ".", &domain) != NI_OK) return NULL;
+
+while ((status = ni_pathsearch(domain, &config_dir, NETINFO_CONFIG_DIR)) ==
NI_NODIR) {
+void *next_domain;
+if (ni_open(domain, "..", &next_domain) != NI_OK) {
+ni_free(next_domain);
```

```
+break;
+}
+ni_free(domain);
+domain = next_domain;
+}
+if (status != NI_OK) {
+ni_free(domain);
+return NULL;
+}
+
+    config = (struct netinfo_config_state *)malloc(sizeof(struct netinfo_config_state));
+    config->domain = domain;
+    config->config_dir = config_dir;
+    config->prop_index = 0;
+    config->val_index = 0;
+    config->val_list = NULL;
+
+return config;
+}
+
+
+
+/*
+ * free_netinfo_config - release NetInfo configuration state
+ */
+static void
+free_netinfo_config(struct netinfo_config_state *config)
+{
+ni_free(config->domain);
+free(config);
+}
+
+
```

```
+
+/*
+ * gettokens_netinfo - return tokens from NetInfo
+ */
+static int
+gettokens_netinfo (
+struct netinfo_config_state *config,
+char **tokenlist,
+int *ntokens
+)
+{
+int prop_index = config->prop_index;
+int val_index = config->val_index;
+char **val_list = config->val_list;
+
+/*
+ * Iterate through each keyword and look for a property that matches it.
+ */
+again:
+if (!val_list) {
+    for (; prop_index < (sizeof(keywords)/sizeof(keywords[0])); prop_index++)
+    {
+        ni_namelist namelist;
+struct keyword current_prop = keywords[prop_index];
+
+/*
+ * For each value associated in the property, we're going to return
+ * a separate line. We squirrel away the values in the config state
+ * so the next time through, we don't need to do this lookup.
+ */
+        NI_INIT(&namelist);
+        if (ni_lookupprop(config->domain, &config->config_dir, current_prop.text, &namelist)
+== NI_OK) {
```

```
+ni_index index;
+
+/* Found the property, but it has no values */
+if (namelist.ni_namelist_len == 0) continue;
+
+if (! (val_list = config->val_list = (char**)malloc(sizeof(char*) * (namelist.ni_namelist_len +
1))))
+{ msyslog(LOG_ERR, "out of memory while configuring"); break; }
+
+for (index = 0; index < namelist.ni_namelist_len; index++) {
+char *value = namelist.ni_namelist_val[index];
+
+if (! (val_list[index] = (char*)malloc(strlen(value)+1)))
+{ msyslog(LOG_ERR, "out of memory while configuring"); break; }
+
+strcpy(val_list[index], value);
+}
+val_list[index] = NULL;
+
+break;
+}
+ni_namelist_free(&namelist);
+}
+config->prop_index = prop_index;
+}
+
+/* No list; we're done here. */
+    if (!val_list) return CONFIG_UNKNOWN;
+
+/*
+ * We have a list of values for the current property.
+ * Iterate through them and return each in order.
+ */
```

```
+if (val_list[val_index])
+{
+int ntok = 1;
+int quoted = 0;
+char *tokens = val_list[val_index];
+
+msyslog(LOG_INFO, "%s %s", keywords[prop_index].text, val_list[val_index]);
+
+(const char*)tokenlist[0] = keywords[prop_index].text;
+for (ntok = 1; ntok < MAXTOKENS; ntok++) {
+tokenlist[ntok] = tokens;
+while (!ISEOL(*tokens) && (!ISPACE(*tokens) || quoted))
+quoted ^= (*tokens++ == '"');
+
+if (ISEOL(*tokens)) {
+*tokens = '\\0';
+break;
+} else { /* must be space */
+*tokens++ = '\\0';
+while (ISPACE(*tokens)) tokens++;
+if (ISEOL(*tokens)) break;
+}
+}
+*ntokens = ntok + 1;
+
+config->val_index++;
+
+return keywords[prop_index].keytype;
+}
+
+/* We're done with the current property. */
+prop_index = ++config->prop_index;
+
```

```
+/* Free val_list and reset counters. */
+for (val_index = 0; val_list[val_index]; val_index++)
+free(val_list[val_index]);
+    free(val_list);val_list = config->val_list = NULL; val_index = config->val_index = 0;
+
+goto again;
+}
+
+#endif /* HAVE_NETINFO */
+
+
+/*
+ * gettokens - read a line and return tokens
+ */
+static int
+gettokens (
+FILE *fp,
+char *line,
+char **tokenlist,
+int *ntokens
+)
+{
+register char *cp;
+register int ntok;
+register int quoted = 0;
+
+/*
+ * Find start of first token
+ */
+again:
+while ((cp = fgets(line, MAXLINE, fp)) != NULL) {
+cp = line;
+while (ISSPACE(*cp))
```



```
+cp++;
+if (!ISEOL(*cp))
+break;
+}
+if (cp == NULL) {
+*ntokens = 0;
+return CONFIG_UNKNOWN; /* hack. Is recognized as EOF */
+}
+
+/*
+ * Now separate out the tokens
+ */
+for (ntok = 0; ntok < MAXTOKENS; ntok++) {
+tokenlist[ntok] = cp;
+while (!ISEOL(*cp) && (!ISSPACE(*cp) || quoted))
+quoted ^= (*cp++ == '"');
+
+if (ISEOL(*cp)) {
+*cp = '\0';
+break;
+} else { /* must be space */
+*cp++ = '\0';
+while (ISSPACE(*cp))
+cp++;
+if (ISEOL(*cp))
+break;
+}
+}
+
+/*
+ * Return the match
+ */
+*ntokens = ntok + 1;
```

```
+ntok = matchkey(tokenlist[0], keywords, 1);
+if (ntok == CONFIG_UNKNOWN)
+goto again;
+return ntok;
+}
+
+
+
+/*
+ * matchkey - match a keyword to a list
+ */
+static int
+matchkey(
+register char *word,
+register struct keyword *keys,
+int complain
+)
+{
+for (;;) {
+if (keys->keytype == CONFIG_UNKNOWN) {
+if (complain)
+msyslog(LOG_ERR,
+ "configure: keyword \"%s\" unknown, line ignored",
+ word);
+return CONFIG_UNKNOWN;
+}
+if (STRSAME(word, keys->text))
+return keys->keytype;
+keys++;
+}
+}
+
+
```

```
+/*
+ * getnetnum - return a net number (this is crude, but careful)
+ */
+static int
+getnetnum(
+const char *num,
+struct sockaddr_storage *addr,
+int complain
+)
+{
+struct addrinfo hints;
+struct addrinfo *ptr;
+
+/* Get host address. Looking for UDP datagram connection */
+memset(&hints, 0, sizeof (hints));
+if (addr->ss_family == AF_INET || addr->ss_family == AF_INET6)
+    hints.ai_family = addr->ss_family;
+else
+    hints.ai_family = AF_UNSPEC;
+
+hints.ai_socktype = SOCK_DGRAM;
+#ifdef DEBUG
+if (debug > 3)
+printf("getaddrinfo %s\n", num);
+#endif
+if (getaddrinfo(num, "ntp", &hints, &ptr)!=0) {
+if (complain)
+msyslog(LOG_ERR,
+"getaddrinfo: \"%s\" invalid host address, line ignored",
+num);
+#ifdef DEBUG
+if (debug > 3)
+printf(
```

```
+ "getaddrinfo: \"%s\" invalid host address%s.\n",
+ num, (complain)
+ ? ", line ignored"
+ : "");
+ #endif
+ return 0;
+ }
+
+ memcpy(addr, ptr->ai_addr, ptr->ai_addrlen);
+ #ifdef DEBUG
+ if (debug > 1)
+ printf("getnetnum given %s, got %s\n",
+   num, stoa(addr));
+ #endif
+   freeaddrinfo(ptr);
+ return 1;
+ }
+
+
+ #if !defined(VMS) && !defined(SYS_WINNT)
+ /*
+  * catchchild - receive the resolver's exit status
+  */
+ static RETSIGTYPE
+ catchchild(
+ int sig
+ )
+ {
+ /*
+  * We only start up one child, and if we're here
+  * it should have already exited. Hence the following
+  * shouldn't hang. If it does, please tell me.
+  */
```

```
+ #if !defined (SYS_WINNT) && !defined(SYS_VXWORKS)
+ (void) wait(0);
+ #endif /* SYS_WINNT && VXWORKS */
+ }
+ #endif /* VMS */
+
+
+ /*
+  * save_resolve - save configuration info into a file for later name resolution
+  */
+ static void
+ save_resolve(
+ char *name,
+ int mode,
+ int version,
+ int minpoll,
+ int maxpoll,
+ u_int flags,
+ int ttl,
+ keyid_t keyid,
+ u_char *keystr
+ )
+ {
+ #ifndef SYS_VXWORKS
+ if (res_fp == NULL) {
+ #ifndef SYS_WINNT
+ (void) strcpy(res_file, RES_TEMPFILE);
+ #else
+ /* no /tmp directory under NT */
+ {
+ if(!(GetTempPath((DWORD)MAX_PATH, (LPTSTR)res_file))) {
+ msyslog(LOG_ERR, "cannot get pathname for temporary directory: %m");
+ return;
+ }
+ }
+ }
+ }
+ }
```

```
+}
+(void) strcat(res_file, "ntpdXXXXXX");
+}
+#endif /* SYS_WINNT */
+#ifdef HAVE_MKSTEMP
+{
+int fd;
+
+res_fp = NULL;
+if ((fd = mkstemp(res_file)) != -1)
+res_fp = fdopen(fd, "r+");
+}
+#else
+(void) mktemp(res_file);
+res_fp = fopen(res_file, "w");
+#endif
+if (res_fp == NULL) {
+msyslog(LOG_ERR, "open failed for %s: %m", res_file);
+return;
+}
+}
+#ifdef DEBUG
+if (debug) {
+printf("resolving %s\n", name);
+}
+#endif
+
+(void)fprintf(res_fp, "%s %d %d %d %d %d %d %u %s\n", name,
+ mode, version, minpoll, maxpoll, flags, ttl, keyid, keystr);
+#ifdef DEBUG
+if (debug > 1)
+printf("config: %s %d %d %d %d %x %d %u %s\n", name, mode,
+ version, minpoll, maxpoll, flags, ttl, keyid, keystr);
```

```
+ #endif
+
+ #else /* SYS_VXWORKS */
+ /* save resolve info to a struct */
+ #endif /* SYS_VXWORKS */
+ }
+
+
+ /*
+  * abort_resolve - terminate the resolver stuff and delete the file
+  */
+ static void
+ abort_resolve(void)
+ {
+ /*
+  * In an ideal world we would might reread the file and
+  * log the hosts which aren't getting configured. Since
+  * this is too much work, however, just close and delete
+  * the temp file.
+  */
+ if (res_fp != NULL)
+ (void) fclose(res_fp);
+ res_fp = NULL;
+
+ #ifndef SYS_VXWORKS /* we don't open the file to begin with */
+ #if !defined(VMS)
+ (void) unlink(res_file);
+ #else
+ (void) delete(res_file);
+ #endif /* VMS */
+ #endif /* SYS_VXWORKS */
+ }
+
```

```
+
+/*
+ * do_resolve_internal - start up the resolver function (not program)
+ */
+/*
+ * On VMS, this routine will simply refuse to resolve anything.
+ *
+ * Possible implementation: keep `res_file' in memory, do async
+ * name resolution via QIO, update from within completion AST.
+ * I'm unlikely to find the time for doing this, though. -wjm
+ */
+static void
+do_resolve_internal(void)
+{
+int i;
+
+if (res_fp == NULL) {
+/* belch */
+msyslog(LOG_ERR,
+"do_resolve_internal: Fatal: res_fp == NULL");
+exit(1);
+}
+
+/* we are done with this now */
+(void) fclose(res_fp);
+res_fp = NULL;
+
+#if !defined(VMS) && !defined(SYS_VXWORKS)
+req_file = res_file; /* set up pointer to res file */
+#ifndef SYS_WINNT
+(void) signal_no_reset(SIGCHLD, catchchild);
+
+#ifndef SYS_VXWORKS
```



```
+i = fork();
+if (i == 0) {
+/*
+ * this used to close everything
+ * I don't think this is necessary
+ */
+/*
+ * To the unknown commenter above:
+ * Well, I think it's better to clean up
+ * after oneself. I have had problems with
+ * refclock-io when intres was running - things
+ * where fine again when ntpintres was gone.
+ * So some systems react erratic at least.
+ *
+ *Frank Kardel
+ *
+ * 94-11-16:
+ * Further debugging has proven that the above is
+ * absolutely harmful. The internal resolver
+ * is still in the SIGIO process group and the lingering
+ * async io information causes it to process requests from
+ * all file descriptor causing a race between the NTP daemon
+ * and the resolver. which then eats data when it wins 8-(.
+ * It is absolutly necessary to kill any IO associations
+ * shared with the NTP daemon.
+ *
+ * We also block SIGIO (currently no ports means to
+ * disable the signal handle for IO).
+ *
+ * Thanks to wgstuken@informatik.uni-erlangen.de to notice
+ * that it is the ntp-resolver child running into trouble.
+ *
+ * THUS:
```

```
+ */
+
+closelog();
+kill_asyncio(0);
+
+(void) signal_no_reset(SIGCHLD, SIG_DFL);
+
+#ifdef DEBUG
+if (0)
+    debug = 2;
+#endif
+
+#ifndef LOG_DAEMON
+openlog("ntpd_initres", LOG_PID);
+# else /* LOG_DAEMON */
+
+# ifndef LOG_NTP
+#    define LOG_NTP LOG_DAEMON
+# endif
+openlog("ntpd_initres", LOG_PID | LOG_NDELAY, LOG_NTP);
+#ifndef SYS_CYGWIN32
+# ifdef DEBUG
+if (debug)
+    setlogmask(LOG_UPTO(LOG_DEBUG));
+else
+# endif /* DEBUG */
+    setlogmask(LOG_UPTO(LOG_DEBUG)); /* @@@ was INFO */
+# endif /* LOG_DAEMON */
+#endif
+
+ntp_initres();
+
+/*
```

```
+ * If we got here, the intres code screwed up.
+ * Print something so we don't die without complaint
+ */
+msyslog(LOG_ERR, "call to ntp_intres lost");
+abort_resolve();
+exit(1);
+}
+#else
+ /* vxWorks spawns a thread... -casey */
+ i = sp (ntp_intres);
+ /*i = taskSpawn("ntp_intres",100,VX_FP_TASK,20000,ntp_intres);*/
+#endif
+if (i == -1) {
+msyslog(LOG_ERR, "fork() failed, can't start ntp_intres: %m");
+(void) signal_no_reset(SIGCHLD, SIG_DFL);
+abort_resolve();
+}
+#else /* SYS_WINNT */
+{
+/* NT's equivalent of fork() is _spawn(), but the start point
+ * of the new process is an executable filename rather than
+ * a function name as desired here.
+ */
+DWORD dwThreadId;
+fflush(stdout);
+ResolverThreadHandle = CreateThread(
+NULL, /* no security attributes*/
+0, /* use default stack size*/
+(LPTHREAD_START_ROUTINE) ntp_intres, /* thread function*/
+NULL, /* argument to thread function */
+0, /* use default creation flags */
+&dwThreadId); /* returns the thread identifier */
+if (ResolverThreadHandle == NULL) {
```

```
+msyslog(LOG_ERR, "CreateThread() failed, can't start ntp_intres");
+abort_resolve();
+}
+}
+#endif /* SYS_WINNT */
+#else /* VMS VX_WORKS */
+msyslog(LOG_ERR,
+"Name resolution not implemented for VMS - use numeric addresses");
+abort_resolve();
+#endif /* VMS VX_WORKS */
+}
diff -Nuar ntp-4.2.0.orig/ntpd/ntp_config.c.rej ntp-4.2.0/ntpd/ntp_config.c.rej
--- ntp-4.2.0.orig/ntpd/ntp_config.c.rej1970-01-01 02:00:00.000000000 +0200
+++ ntp-4.2.0/ntpd/ntp_config.c.rej2005-12-20 10:15:40.000000000 +0200
@@ -0,0 +1,17 @@
+*****
+*** 924,930 ****
+
+ }
+ break;
+
+- case CONFIG_BROADCASTCLIENT:
+ if (ntokens == 1) {
+ proto_config(PROTO_BROADCASTCLIENT, 1, 0., NULL);
+ } else {
+--- 924,930 ----
+ }
+ break;
+
++ case NTP_CONFIG_BROADCASTCLIENT:
+ if (ntokens == 1) {
+ proto_config(PROTO_BROADCASTCLIENT, 1, 0., NULL);
+ } else {
```

generic-configure-ac.patch

```
#
# Submitted-By: Ladislav Michl, 2005-02-08
# Committed-By: Robert Schwebel, 2005-02-09
#
# Error:
#
# openssh-3.9p1 fails crosscompile
#
# Description:
#
# Ported patch for 3.7.1p2 to 3.9p1.
#
# State:
#
# Mostly fixed upstream (few oneliners left)
#
--- openssh-3.9p1/configure.ac2004-08-16 15:12:06.000000000 +0200
+++ openssh-3.9p1/configure.ac2005-02-08 16:59:34.000000000 +0100
@@ -14,7 +14,7 @@
# ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING
OUT OF
# OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

-AC_INIT
+AC_INIT(OpenSSH, Portable)
AC_CONFIG_SRCDIR([ssh.c])

AC_CONFIG_HEADER(config.h)
@@ -249,18 +249,23 @@
*-*-linux*)
no_dev_ptmx=1
```

```
check_for_libcrypt_later=1
-check_for_openpty_ctty_bug=1
AC_DEFINE(DONT_TRY_OTHER_AF)
AC_DEFINE(PAM_TTY_KLUDGE)
AC_DEFINE(LOCKED_PASSWD_PREFIX, "!")
AC_DEFINE(SPT_TYPE,SPT_REUSEARGV)
AC_DEFINE(LINK_OPNOTSUPP_ERRNO, EPERM)
inet6_default_4in6=yes
-case `uname -r` in
-1.* | 2.0.*)
-AC_DEFINE(BROKEN_CMSG_TYPE)
-;;
-esac
+if test ! -z "$cross_compiling" && test "x$cross_compiling" = "xyes"; then
+AC_MSG_WARN([cross compiling: Disabling openpty test])
+# Let's hope noone sane is using that...
+case `uname -r` in
+1.* | 2.0.*)
+AC_DEFINE(BROKEN_CMSG_TYPE)
+;;
+esac
+else
+check_for_openpty_ctty_bug=1
+fi
;;
mips-sony-bsd | mips-sony-newsos4)
AC_DEFINE(HAVE_NEWS4)
@@ -347,12 +352,14 @@
# Attention: always take care to bind libsocket and libnsl before libc,
# otherwise you will find lots of "SIOCGPGRP errno 22" on syslog
;;
+# UnixWare 1.x, UnixWare 2.x, and others based on code from Univel.
*-*-sysv4.2*)
```

```
AC_DEFINE(USE_PIPES)
AC_DEFINE(SETUID_BREAKS_SETUID)
AC_DEFINE(BROKEN_SETREUID)
AC_DEFINE(BROKEN_SETREGID)
;;
+# UnixWare 7.x, OpenUNIX 8
*-*-sysv5*)
AC_DEFINE(USE_PIPES)
AC_DEFINE(SETUID_BREAKS_SETUID)
@@ -361,6 +368,7 @@
;;
*-*-sysv*)
;;
+# SCO UNIX and OEM versions of SCO UNIX
*-*-sco3.2v4*)
CPPFLAGS="$CPPFLAGS -Dftruncate=chsize"
LIBS="$LIBS -los -lprot -lcrypt_i -lx -ltinfo -lm"
@@ -380,6 +388,7 @@
do_sco3_extra_lib_check=yes
TEST_SHELL=ksh
;;
+# SCO OpenServer 5.x
*-*-sco3.2v5*)
if test -z "$GCC"; then
CFLAGS="$CFLAGS -belf"
@@ -506,15 +515,17 @@
)

AC_MSG_CHECKING(compiler and flags for sanity)
-AC_TRY_RUN([
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([
#include <stdio.h>
```

```
int main(){exit(0);}
-],
+)],
[AC_MSG_RESULT(yes) ],
[
AC_MSG_RESULT(no)
AC_MSG_ERROR([*** compiler cannot create working executables, check config.log ***])
-]
+],
+[AC_MSG_WARN([cross compiling: not checking compiler sanity]) ]
)
```

Checks for header files.

```
@@ -640,7 +651,7 @@
)
```

AC_MSG_CHECKING(for zlib 1.1.4 or greater)

```
-AC_TRY_RUN([
+AC_RUN_IFELSE([AC_LANG_SOURCE([[
#include <zlib.h>
int main()
{
@@ -652,7 +663,7 @@
exit(0);
exit(2);
}
-],
+]]]),
AC_MSG_RESULT(yes),
[ AC_MSG_RESULT(no)
if test -z "$zlib_check_nonfatal" ; then
@@ -665,7 +676,8 @@
else
```



```
AC_MSG_WARN([zlib version may have security problems])
fi
-]
+],
+[AC_MSG_WARN([cross compiling: not checking zlib version]) ]
)
```

```
dnl UnixWare 2.x
@@ -719,16 +731,20 @@
)
```

```
AC_MSG_CHECKING([whether struct dirent allocates space for d_name])
-AC_TRY_RUN(
-[
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([[
#include <sys/types.h>
#include <dirent.h>
int main(void){struct dirent d;exit(sizeof(d.d_name)<=sizeof(char));}
-],
+]])],
[AC_MSG_RESULT(yes)],
[
AC_MSG_RESULT(no)
AC_DEFINE(BROKEN_ONE_BYTE_DIRENT_D_NAME)
+],
+[
+AC_MSG_WARN([cross compiling: assuming BROKEN_ONE_BYTE_DIRENT_D_NAME])
+AC_DEFINE(BROKEN_ONE_BYTE_DIRENT_D_NAME)
]
)
```

```
@@ -888,28 +904,32 @@
```

```
AC_CHECK_FUNCS(setresuid, [
dnl Some platforms have setresuid that isn't implemented, test for this
AC_MSG_CHECKING(if setresuid seems to work)
-AC_TRY_RUN([
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([[
#include <stdlib.h>
#include <errno.h>
int main(){errno=0; setresuid(0,0,0); if (errno==ENOSYS) exit(1); else exit(0);}
-],
+]]),
[AC_MSG_RESULT(yes)],
[AC_DEFINE(BROKEN_SETRESUID)
- AC_MSG_RESULT(not implemented)]
+ AC_MSG_RESULT(not implemented)],
+[AC_MSG_WARN([cross compiling: not checking setresuid])]
)
])
```

```
AC_CHECK_FUNCS(setresgid, [
dnl Some platforms have setresgid that isn't implemented, test for this
AC_MSG_CHECKING(if setresgid seems to work)
-AC_TRY_RUN([
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([[
#include <stdlib.h>
#include <errno.h>
int main(){errno=0; setresgid(0,0,0); if (errno==ENOSYS) exit(1); else exit(0);}
-],
+]]),
[AC_MSG_RESULT(yes)],
[AC_DEFINE(BROKEN_SETRESGID)
- AC_MSG_RESULT(not implemented)]
```

```
+ AC_MSG_RESULT(not implemented)],  
+[AC_MSG_WARN([cross compiling: not checking setresuid])]  
)  
])
```

```
@@ -935,17 +955,18 @@
```

```
# Check for broken snprintf  
if test "x$ac_cv_func_snprintf" = "xyes" ; then  
  AC_MSG_CHECKING([whether snprintf correctly terminates long strings])  
  -AC_TRY_RUN(  
  -[  
  +AC_RUN_IFELSE(  
  +[AC_LANG_SOURCE([[  
    #include <stdio.h>  
    int main(void){char b[5];snprintf(b,5,"123456789");exit(b[4]!='\0');}  
  -],  
  +]])],  
  [AC_MSG_RESULT(yes)],  
  [  
    AC_MSG_RESULT(no)  
    AC_DEFINE(BROKEN_SNPRINTF)  
    AC_MSG_WARN([***** Your snprintf() function is broken, complain to your vendor])  
  -]  
  +],  
  +[ AC_MSG_WARN([cross compiling: Assuming working snprintf()]) ]  
  )  
  fi
```

```
@@ -990,7 +1011,7 @@
```

```
fi
```

```
dnl make sure that openpty does not reacquire controlling terminal  
-if test ! -z "$check_for_openpty_ctty_bug"; then
```

```
+if test "x$check_for_openpty_ctty_bug" = "x1"; then
AC_MSG_CHECKING(if openpty correctly handles controlling tty)
AC_TRY_RUN(
[
@@ -1206,8 +1227,8 @@

# Determine OpenSSL header version
AC_MSG_CHECKING([OpenSSL header version])
-AC_TRY_RUN(
-[
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([
#include <stdio.h>
#include <string.h>
#include <openssl/opensslv.h>
@@ -1225,7 +1246,7 @@

exit(0);
}
-],
+]]),
[
ssl_header_ver=`cat conftest.sslincver`
AC_MSG_RESULT($ssl_header_ver)
@@ -1233,13 +1254,16 @@
[
AC_MSG_RESULT(not found)
AC_MSG_ERROR(OpenSSL version header not found.)
+],
+[
+AC_MSG_WARN([cross compiling: not checking])
]
)
```

```
# Determine OpenSSL library version
AC_MSG_CHECKING([OpenSSL library version])
-AC_TRY_RUN(
-[
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([[
#include <stdio.h>
#include <string.h>
#include <openssl/opensslv.h>
@@ -1258,7 +1282,7 @@

exit(0);
}
-],
+]])),
[
ssl_library_ver=`cat conftest.sslibver`
AC_MSG_RESULT($ssl_library_ver)
@@ -1266,17 +1290,20 @@
[
AC_MSG_RESULT(not found)
AC_MSG_ERROR(OpenSSL library not found.)
+],
+[
+AC_MSG_WARN([cross compiling: not checking])
]
)

# Sanity check OpenSSL headers
AC_MSG_CHECKING([whether OpenSSL's headers match the library])
-AC_TRY_RUN(
-[
```

```
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([[
#include <string.h>
#include <openssl/opensslv.h>
int main(void) { exit(SSLeay() == OPENSSL_VERSION_NUMBER ? 0 : 1); }
-],
+]]]),
[
AC_MSG_RESULT(yes)
],
@@ -1285,6 +1312,9 @@
AC_MSG_ERROR([Your OpenSSL headers do not match your library.
Check config.log for details.
Also see contrib/findssl.sh for help identifying header/library mismatches.])
+],
+[
+AC_MSG_WARN([cross compiling: not checking])
]
)

@@ -1305,12 +1335,12 @@

# Check wheter OpenSSL seeds itself
AC_MSG_CHECKING([whether OpenSSL's PRNG is internally seeded])
-AC_TRY_RUN(
-[
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([[
#include <string.h>
#include <openssl/rand.h>
int main(void) { exit(RAND_status() == 1 ? 0 : 1); }
-],
+]]]),
```

```
[
OPENSSL_SEEDS_ITSELF=yes
AC_MSG_RESULT(yes)
@@ -1320,6 +1350,12 @@
# Default to use of the rand helper if OpenSSL doesn't
# seed itself
USE_RAND_HELPER=yes
+],
+[
+AC_MSG_WARN([cross compiling: assuming yes])
+# This is safe, since all recent OpenSSL versions will
+# complain at runtime if not seeded correctly.
+OPENSSL_SEEDS_ITSELF=yes
]
)
```

```
@@ -1886,8 +1922,8 @@
exit 1;
else
dnl test snprintf (broken on SCO w/gcc)
-AC_TRY_RUN(
-[
+AC_RUN_IFELSE(
+[AC_LANG_SOURCE([[
#include <stdio.h>
#include <string.h>
#ifdef HAVE_SNPRINTF
@@ -1910,7 +1946,8 @@
#else
main() { exit(0); }
#endif
-], [ true ], [ AC_DEFINE(BROKEN_SNPRINTF) ]
+]])), [ true ], [ AC_DEFINE(BROKEN_SNPRINTF) ],
```

```
+AC_MSG_WARN([cross compiling: Assuming working snprintf()])
)
fi
```

```
@@ -2015,13 +2052,14 @@
```

dnl make sure we're using the real structure members and not defines

```
AC_CACHE_CHECK([for msg_accrights field in struct msghdr],
ac_cv_have_accrights_in_msghdr, [
-AC_TRY_RUN(
+AC_COMPILE_IFELSE(
```

```
[
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/uio.h>
int main() {
#ifdef msg_accrights
+error "msg_accrights is a macro"
exit(1);
#endif
struct msghdr m;
```

```
@@ -2039,13 +2077,14 @@
```

```
AC_CACHE_CHECK([for msg_control field in struct msghdr],
ac_cv_have_control_in_msghdr, [
```

```
-AC_TRY_RUN(
+AC_COMPILE_IFELSE(
```

```
[
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/uio.h>
int main() {
#ifdef msg_control
+error "msg_control is a macro"
```



```
exit(1);
#endif
struct msghdr m;
@@ -2379,6 +2418,10 @@
AC_DEFINE_UNQUOTED(MAIL_DIRECTORY, "$maildir")
fi

+if test ! -z "$cross_compiling" && test "x$cross_compiling" = "xyes"; then
+AC_MSG_WARN([cross compiling: Disabling /dev/ptmx test])
+disable_ptmx_check=yes
+fi
if test -z "$no_dev_ptmx" ; then
if test "x$disable_ptmx_check" != "xyes" ; then
AC_CHECK_FILE("/dev/ptmx",
@@ -2389,12 +2432,17 @@
)
fi
fi
-AC_CHECK_FILE("/dev/ptc",
-[
-AC_DEFINE_UNQUOTED(HAVE_DEV_PTS_AND_PTC)
-have_dev_ptc=1
-]
-)
+
+if test ! -z "$cross_compiling" && test "x$cross_compiling" != "xyes"; then
+AC_CHECK_FILE("/dev/ptc",
+[
+AC_DEFINE_UNQUOTED(HAVE_DEV_PTS_AND_PTC)
+have_dev_ptc=1
+]
+)
+else
```

```
+AC_MSG_WARN([cross compiling: Disabling /dev/ptc test])
+fi

# Options from here on. Some of these are preset by platform above
AC_ARG_WITH(mantype,
@@ -2491,13 +2539,17 @@
# check for /etc/default/login and use it if present.
AC_ARG_ENABLE(etcd-default-login,
[ --disable-etcd-default-login    Disable using PATH from /etc/default/login [no]],
-[
-AC_CHECK_FILE("/etc/default/login", [ external_path_file=/etc/default/login ])
+[ AC_CHECK_FILE("/etc/default/login",
+ [ external_path_file=/etc/default/login ])

-if test "x$external_path_file" = "x/etc/default/login"; then
-AC_DEFINE(HAVE_ETC_DEFAULT_LOGIN)
-fi
-])
+if test ! -z "$cross_compiling" && test "x$cross_compiling" = "xyes";
+then
+AC_MSG_WARN([cross compiling: Disabling /etc/default/login test])
+elif test "x$external_path_file" = "x/etc/default/login"; then
+AC_DEFINE(HAVE_ETC_DEFAULT_LOGIN)
+fi
+]
+)

dnl BSD systems use /etc/login.conf so --with-default-path= has no effect
if test $ac_cv_func_login_getcapbool = "yes" -a \
```

openssl-0.9.7g-configure.diff

diff -urN openssl-0.9.7g/Configure openssl-0.9.7g-ptx/Configure

```

--- openssl-0.9.7g/Configure2005-04-07 18:06:01.000000000 +0200
+++ openssl-0.9.7g-ptx/Configure2005-05-05 17:22:01.000000000 +0200
@@ -390,6 +390,41 @@
# assembler versions -- currently defunct:
##"OpenBSD-alpha","gcc:-DTERMIO -O3 -fomit-frame-
pointer:::(unknown):SIXTY_FOUR_BIT_LONG DES_INT DES_PTR
DES_RISC2:${alpha_asm}",

+# Sane Linux configuration values, stolen from the Debian package....
+"linux-alpha","gcc:-DTERMIO -O3 -fomit-frame-pointer -Wall:-D_REENTRANT:-
ldl:SIXTY_FOUR_BIT_LONG RC4_CHUNK DES_RISC1
DES_UNROLL:${alpha_asm}:dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-alpha-ev4","gcc:-DTERMIO -O3 -mcpu=ev4 -fomit-frame-pointer -Wall:-
D_REENTRANT:-ldl:SIXTY_FOUR_BIT_LONG RC4_CHUNK DES_RISC1
DES_UNROLL:${alpha_asm}:dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-alpha-ev5","gcc:-DTERMIO -O3 -mcpu=ev5 -fomit-frame-pointer -Wall:-
D_REENTRANT:-ldl:SIXTY_FOUR_BIT_LONG RC4_CHUNK DES_RISC1
DES_UNROLL:${alpha_asm}:dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-arm","gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall:-
D_REENTRANT:-ldl:BN_LLONG DES_RISC1:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-armeb","gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall:-
D_REENTRANT:-ldl:BN_LLONG DES_RISC1:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-freebsd-alpha","gcc:-DTERMIO -O -fomit-frame-
pointer:::(unknown)::SIXTY_FOUR_BIT_LONG RC4_CHUNK DES_INT DES_PTR
DES_RISC2:::dlfcn:bsd-gcc-shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-freebsd-i386", "gcc:-DTERMIO -DL_ENDIAN -fomit-frame-pointer -O3 -m486 -
Wall:-pthread -D_REENTRANT -D_THREAD_SAFE -D_THREADSAFE::BN_LLONG
${x86_gcc_des} ${x86_gcc_opts}:${x86_elf_asm}:dlfcn:bsd-gcc-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-hppa","gcc:-DB_ENDIAN -DTERMIO -O2 -Wall:-D_REENTRANT:-ldl:BN_LLONG
MD2_CHAR RC4_INDEX:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",

```

```
+ "linux-hurd-i386", "gcc:-DL_ENDIAN -DTERMIOS -O3 -fomit-frame-pointer -m486 -Wall:-  
D_REENTRANT::-ldl:BN_LLONG ${x86_gcc_des} ${x86_gcc_opts}:${x86_elf_asm}:dlfcn:linux-  
shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-ia64", "gcc:-DTERMIO -O3 -fomit-frame-pointer -Wall::-D_REENTRANT::-  
ldl:SIXTY_FOUR_BIT_LONG RC4_CHUNK RC4_CHAR:asm/ia64.o:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "#linux-i386", "gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -m486 -Wall:-  
D_REENTRANT::-ldl:BN_LLONG ${x86_gcc_des} ${x86_gcc_opts}:dlfcn:linux-shared:-  
fPIC",  
+ "linux-i386", "gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall:-  
D_REENTRANT::-ldl:BN_LLONG ${x86_gcc_des} ${x86_gcc_opts}:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-i386-i486", "gcc:-DL_ENDIAN -DTERMIO -O3 -march=i486 -mcpu=i486 -fomit-frame-  
pointer -Wall::-D_REENTRANT::-ldl:BN_LLONG ${x86_gcc_des}  
${x86_gcc_opts}:${x86_elf_asm}:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-i386-i586", "gcc:-DL_ENDIAN -DTERMIO -O3 -march=i586 -mcpu=i586 -fomit-frame-  
pointer -Wall::-D_REENTRANT::-ldl:BN_LLONG ${x86_gcc_des}  
${x86_gcc_opts}:${x86_elf_asm}:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-i386-i686/cmov", "gcc:-DL_ENDIAN -DTERMIO -O3 -march=i686 -mcpu=i686 -fomit-  
frame-pointer -Wall::-D_REENTRANT::-ldl:BN_LLONG ${x86_gcc_des}  
${x86_gcc_opts}:${x86_elf_asm}:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-m68k", "gcc:-DB_ENDIAN -DTERMIO -O2 -Wall::-D_REENTRANT::-ldl:BN_LLONG  
MD2_CHAR RC4_INDEX:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-mips", "gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall:-  
D_REENTRANT::-ldl:BN_LLONG:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-mipsel", "gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall:-  
D_REENTRANT::-ldl:BN_LLONG:dlfcn:linux-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",  
+ "linux-netbsd-i386", "gcc:-DL_ENDIAN -DTERMIOS -O3 -fomit-frame-pointer -m486 -  
Wall:(unknown)::BN_LLONG ${x86_gcc_des} ${x86_gcc_opts}:dlfcn:bsd-gcc-shared:-  
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
```

```

+"linux-netbsd-m68k", "gcc:-DB_ENDIAN -DTERMIOS -O3 -fomit-frame-pointer -
Wall::(unknown)::BN_LLONG MD2_CHAR RC4_INDEX DES_UNROLL:::dlfcn:bsd-gcc-
shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-netbsd-sparc", "gcc:-DB_ENDIAN -DTERMIOS -O3 -fomit-frame-pointer -mv8 -
Wall::(unknown)::BN_LLONG MD2_CHAR RC4_INDEX DES_UNROLL:::dlfcn:bsd-gcc-
shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-openbsd-alpha", "gcc:-DTERMIOS -O3 -fomit-frame-
pointer::(unknown)::SIXTY_FOUR_BIT_LONG DES_INT DES_PTR
DES_RISC2:::dlfcn:bsd-gcc-shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-openbsd-i386", "gcc:-DL_ENDIAN -DTERMIOS -O3 -fomit-frame-pointer -
m486::(unknown)::BN_LLONG ${x86_gcc_des} ${x86_gcc_opts}:${x86_out_asm};dlfcn:bsd-
gcc-shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-openbsd-mips", "gcc:-O2 -DL_ENDIAN::(unknown)::BN_LLONG MD2_CHAR
RC4_INDEX RC4_CHAR DES_UNROLL DES_RISC2 DES_PTR BF_PTR:::dlfcn:bsd-gcc-
shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-powerpc", "gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG DES_UNROLL DES_RISC2 DES_PTR MD2_CHAR
RC4_INDEX:::dlfcn:linux-shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-s390", "gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-sh3", "gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-sh4", "gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-sh3eb", "gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-sh4eb", "gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+"linux-sparc", "gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG RC4_CHAR RC4_CHUNK DES_UNROLL
BF_PTR:::dlfcn:linux-shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",

```

```
+ "linux-sparc-v8", "gcc:-DB_ENDIAN -DTERMIO -O3 -mcpu=v8 -fomit-frame-pointer -Wall -
DBN_DIV2W::-D_REENTRANT::-ldl:BN_LLONG RC4_CHAR RC4_CHUNK DES_UNROLL
BF_PTR:asm/sparcv8.o:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+ "linux-sparc-v9", "gcc:-DB_ENDIAN -DTERMIO -O3 -mcpu=v9 -Wa,-Av8plus -fomit-frame-
pointer -Wall -DULTRASPARC -DBN_DIV2W::-D_REENTRANT::-ldl:BN_LLONG
RC4_CHAR RC4_CHUNK DES_UNROLL BF_PTR:asm/sparcv8plus.o:::asm/md5-
sparcv8plus.o:::dlfcn:linux-shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+ "linux-cris", "gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-ldl:BN_LLONG:::dlfcn:linux-shared:-
fpic::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
# The intel boxes :-), It would be worth seeing if bsdi-gcc can use the
# bn86-elf.o file since it is hand tweaked assembler.
"linux-elf", "gcc:-DL_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -m486 -Wall::-
D_REENTRANT::-ldl:BN_LLONG ${x86_gcc_des} ${x86_gcc_opts}:${x86_elf_asm}:dlfcn:linux-
shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
@@ -438,6 +473,7 @@
# ARM comes in both little- and big-endian flavors. The following line is
# endian neutral, but ./config is free to throw in -D[BL]_ENDIAN...
"linux-elf-arm", "gcc:-DTERMIO -O3 -fomit-frame-pointer -Wall::-D_REENTRANT::-
ldl:BN_LLONG:::dlfcn:linux-shared:-fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",
+ "linux-elf-armeb", "gcc:-DB_ENDIAN -DTERMIO -O3 -fomit-frame-pointer -Wall::-
D_REENTRANT::-BN_LLONG:::dlfcn:linux-shared:-
fPIC::so.\$(SHLIB_MAJOR).\$(SHLIB_MINOR)",

# SCO/Caldera targets.
#
```

portmap-5beta-compilation_fixes-3.patch

Submitted By: Kevin P. Fleming <kpflaming at linuxfromscratch dot org>

Date: 2004-05-02

Initial Package Version: 5beta

Origin: <http://archives.linuxfromscratch.org/mail-archives/blfs-dev/2003-January/001960.html>

Description: The patch was created from the portmap modified package by Mark Heerdink.

This patch provides the following improvements:

- * Link against dynamic tcp_wrappers.
- * Create an install target for portmap.
- * Compilation and security fixes.
- * Documentation fixes.

Originally created by Tushar Teredesai, updated by kpflaming to ensure portmap will compile without tcp_wrappers installed.

```
diff -Naur portmap_5beta/BLURB portmap_5beta.gimli/BLURB
--- portmap_5beta/BLURB1996-07-06 16:09:46.000000000 -0500
+++ portmap_5beta.gimli/BLURB2002-01-07 09:13:58.000000000 -0600
@@ -1,3 +1,6 @@
+
+#####
+#####
+
+@(#) BLURB 1.5 96/07/06 23:09:45
```

This is the fifth replacement portmapper release.

```
diff -Naur portmap_5beta/Makefile portmap_5beta.gimli/Makefile
--- portmap_5beta/Makefile1996-07-06 16:06:19.000000000 -0500
+++ portmap_5beta.gimli/Makefile2002-07-15 16:00:07.000000000 -0500
@@ -8,7 +8,7 @@
# if you disagree. See `man 3 syslog' for examples. Some syslog versions
# do not provide this flexibility.
#
-FACILITY=LOG_MAIL
+FACILITY=LOG_DAEMON

# To disable tcp-wrapper style access control, comment out the following
# macro definitions. Access control can also be turned off by providing
@@ -16,7 +16,8 @@
# daemon, is always treated as an authorized host.
```

```
HOSTS_ACCESS= -DHOSTS_ACCESS
-WRAP_LIB = $(WRAP_DIR)/libwrap.a
+#WRAP_LIB = $(WRAP_DIR)/libwrap.a
+WRAP_LIB = -lwrap

# Comment out if your RPC library does not allocate privileged ports for
# requests from processes with root privilege, or the new portmap will
@@ -71,7 +72,7 @@
# With verbose logging on, HP-UX 9.x and AIX 4.1 leave zombies behind when
# SIGCHLD is not ignored. Enable next macro for a fix.
#
-# ZOMBIES = -DIGNORE_SIGCHLD# AIX 4.x, HP-UX 9.x
+ZOMBIES = -DIGNORE_SIGCHLD# AIX 4.x, HP-UX 9.x

# Uncomment the following macro if your system does not have u_long.
#
@@ -81,7 +82,7 @@
# libwrap.a object library. WRAP_DIR should specify the directory with
# that library.

-WRAP_DIR= ../tcp_wrappers
+WRAP_DIR= $(TCPD_DIR)

# Auxiliary object files that may be missing from your C library.
#
@@ -99,22 +100,31 @@

# Comment out if your compiler talks ANSI and understands const
#
-CONST = -Dconst=
+#CONST = -Dconst=
```


End of configurable stuff.

#####

+GLIBC=\$(shell grep -s -c __GLIBC__ /usr/include/features.h)

+

+ifeq (\$(GLIBC),0)

+LIBS += # -lbsd

+else

+LIBS += -lnsl

+endif

+

+

SHELL= /bin/sh

-COPT= \$(CONST) -Dperror=xperror \$(HOSTS_ACCESS) \$(CHECK_PORT) \

+COPT= \$(CONST) \$(HOSTS_ACCESS) \$(CHECK_PORT) \

\$(SYS) -DFACILITY=\$(FACILITY) \$(ULONG) \$(ZOMBIES) \$(SA_LEN) \

\$(LOOPBACK) \$(SETPGRP)

-CFLAGS= \$(COPT) -O \$(NSARCHS)

+CFLAGS= -Wall \$(COPT) -O2 \$(NSARCHS)

OBJECTS= portmap.o pmap_check.o from_local.o \$(AUX)

all:portmap pmap_dump pmap_set

-portmap: \$(OBJECTS) \$(WRAP_DIR)/libwrap.a

+portmap: \$(OBJECTS) # \$(WRAP_DIR)/libwrap.a

\$(CC) \$(CFLAGS) -o \$@ \$(OBJECTS) \$(WRAP_LIB) \$(LIBS)

pmap_dump: pmap_dump.c

@@ -129,6 +139,17 @@

get_myaddress: get_myaddress.c

cc \$(CFLAGS) -DTEST -o \$@ get_myaddress.c \$(LIBS)

```
+install: all
+install -o root -g root -m 0755 -s portmap ${BASEDIR}/sbin
+install -o root -g root -m 0755 -s pmap_dump ${BASEDIR}/usr/sbin
+install -o root -g root -m 0755 -s pmap_set ${BASEDIR}/usr/sbin
+install -o root -g root -m 0644 portmap.8 ${BASEDIR}/usr/share/man/man8
+install -o root -g root -m 0644 pmap_dump.8 ${BASEDIR}/usr/share/man/man8
+install -o root -g root -m 0644 pmap_set.8 ${BASEDIR}/usr/share/man/man8
+#cat README BLURB >${BASEDIR}/usr/share/doc/portmap/portmapper.txt
+#gzip -9f ${BASEDIR}/usr/share/doc/portmap/portmapper.txt
+
+
lint:
lint $(COPT) $(OBJECTS:%.o=%.c)

diff -Naur portmap_5beta/daemon.c portmap_5beta.gimli/daemon.c
--- portmap_5beta/daemon.c1992-06-11 15:53:12.000000000 -0500
+++ portmap_5beta.gimli/daemon.c2002-01-07 09:22:24.000000000 -0600
@@ -36,16 +36,13 @@
 #endif /* LIBC_SCCS and not lint */

#include <fcntl.h>
-
-/* From unistd.h */
-#define STDIN_FILENO0
-#define STDOUT_FILENO1
-#define STDERR_FILENO2
+#include <unistd.h>
+#include <sys/types.h>

/* From paths.h */
#define _PATH_DEVNULL "/dev/null"

-daemon(nochdir, noclose)
```

```
+int daemon(nochdir, noclose)
int nochdir, noclose;
{
int cpid;
diff -Naur portmap_5beta/from_local.c portmap_5beta.gimli/from_local.c
--- portmap_5beta/from_local.c1996-05-31 08:52:58.000000000 -0500
+++ portmap_5beta.gimli/from_local.c2002-01-07 09:25:49.000000000 -0600
@@ -35,7 +35,7 @@
 * Mountain View, California 94043
 */

-#ifndef lint
+#ifdef lint
static char sccsid[] = "@(#) from_local.c 1.3 96/05/31 15:52:57";
#endif

@@ -51,6 +51,9 @@
#include <net/if.h>
#include <sys/ioctl.h>
#include <syslog.h>
#include <stdlib.h>
#include <string.h>
#include <unistd.h>

#ifdef TRUE
#define TRUE 1
@@ -96,7 +99,7 @@

/* find_local - find all IP addresses for this host */

find_local()
int find_local()
{
```

```
    struct ifconf ifc;
    struct ifreq ifreq;
@@ -154,7 +157,7 @@

/* from_local - determine whether request comes from the local system */

-from_local(addr)
+int from_local(addr)
struct sockaddr_in *addr;
{
    int i;
diff -Naur portmap_5beta/pmap_check.c portmap_5beta.gimli/pmap_check.c
--- portmap_5beta/pmap_check.c1996-07-07 03:49:10.000000000 -0500
+++ portmap_5beta.gimli/pmap_check.c2002-01-07 09:37:58.000000000 -0600
@@ -32,7 +32,7 @@
    * Computing Science, Eindhoven University of Technology, The Netherlands.
    */

-#ifndef lint
+#ifdef lint
static char sccsid[] = "@(#) pmap_check.c 1.8 96/07/07 10:49:10";
#endif

@@ -45,6 +45,11 @@
#include <netinet/in.h>
#include <rpc/rpcent.h>
#endif
+#include <sys/types.h>
+#include <unistd.h>
+#ifdef HOSTS_ACCESS
+#include <tcpd.h>
+#endif
```

```
extern char *inet_ntoa();
```

```
@@ -110,7 +113,7 @@
```

```
/* check_default - additional checks for NULL, DUMP, GETPORT and unknown */
```

```
-check_default(addr, proc, prog)
+int check_default(addr, proc, prog)
struct sockaddr_in *addr;
u_long  proc;
u_long  prog;
@@ -128,7 +131,7 @@
```

```
/* check_privileged_port - additional checks for privileged-port updates */
```

```
-check_privileged_port(addr, proc, prog, port)
+int check_privileged_port(addr, proc, prog, port)
struct sockaddr_in *addr;
u_long  proc;
u_long  prog;
@@ -173,7 +176,7 @@
```

```
#else
```

```
-check_setunset(addr, proc, prog, port)
+int check_setunset(addr, proc, prog, port)
struct sockaddr_in *addr;
u_long  proc;
u_long  prog;
@@ -197,7 +200,7 @@
```

```
/* check_callit - additional checks for forwarded requests */
```

```
-check_callit(addr, proc, prog, aproc)
+int check_callit(addr, proc, prog, aproc)
struct sockaddr_in *addr;
u_long proc;
u_long prog;
@@ -249,13 +252,13 @@
    };
    struct proc_map *procp;
    static struct proc_map procmap[] = {
-PMAPPROC_CALLIT, "callit",
-PMAPPROC_DUMP, "dump",
-PMAPPROC_GETPORT, "getport",
-PMAPPROC_NULL, "null",
-PMAPPROC_SET, "set",
-PMAPPROC_UNSET, "unset",
-0, 0,
+{ PMAPPROC_CALLIT, "callit" },
+{ PMAPPROC_DUMP, "dump" },
+{ PMAPPROC_GETPORT, "getport" },
+{ PMAPPROC_NULL, "null" },
+{ PMAPPROC_SET, "set" },
+{ PMAPPROC_UNSET, "unset" },
+{ 0, 0 }
    };

    /*
@@ -269,7 +272,7 @@

    if (prognum == 0) {
        proptime = "";
-} else if (rpc = getrpcbyname((int) prognum)) {
+} else if ((rpc = getrpcbyname((int) prognum)) != NULL) {
        proptime = rpc->r_name;
```

```
} else {
    sprintf(progname = progbuf, "%lu", prognum);
diff -Naur portmap_5beta/pmap_dump.8 portmap_5beta.gimli/pmap_dump.8
--- portmap_5beta/pmap_dump.81969-12-31 18:00:00.000000000 -0600
+++ portmap_5beta.gimli/pmap_dump.82002-01-07 09:13:58.000000000 -0600
@@ -0,0 +1,24 @@
+.TH PMAP_DUMP 8 "21th June 1997" Linux "Linux Programmer's Manual"
+.SH NAME
+pmmap_dump \- print a list of all registered RPC programs
+.SH SYNOPSIS
+.B pmap_dump
+.SH DESCRIPTION
+The
+.B pmap_dump
+command can be used to restart a running portmapper or to print
+a list of all registered RPC programs on the local host. If you
+want to use the program to restart the portmapper you have to
+redirect the output of
+.B pmap_dump
+to a file. After this you can restart the portmapper and restore
+the list of the registered RPC programs by feeding the output
+of
+.B pmap_dump
+to the
+.B pmap_set
+command.
+.SH SEE ALSO
+.BR pmap_set (8),
+.BR rpc.portmap (8)
+
diff -Naur portmap_5beta/pmap_dump.c portmap_5beta.gimli/pmap_dump.c
--- portmap_5beta/pmap_dump.c1992-06-11 15:53:16.000000000 -0500
+++ portmap_5beta.gimli/pmap_dump.c2002-01-07 09:20:19.000000000 -0600
```

@@ -5,7 +5,7 @@

```
* Computing Science, Eindhoven University of Technology, The Netherlands.  
*/
```

-#ifndef lint

+#ifdef lint

```
static char sccsid[] = "@(#) pmap_dump.c 1.1 92/06/11 22:53:15";  
#endif
```

@@ -23,7 +23,20 @@

```
static char *protoname();
```

-main(argc, argv)

+#ifndef INADDR_LOOPBACK

+#define INADDR_LOOPBACK ntohl(inet_addr("127.0.0.1"))

+#endif

+

+static void get_myloopaddress(addrp)

+struct sockaddr_in *addrp;

+{

+ memset((char *) addrp, 0, sizeof(*addrp));

+ addrp->sin_family = AF_INET;

+ addrp->sin_port = htons(PMAPPORT);

+ addrp->sin_addr.s_addr = htonl(INADDR_LOOPBACK);

+}

+

+int main(argc, argv)

int argc;

char **argv;

{

@@ -31,7 +44,7 @@

```
register struct pmaplist *list;
```



```
register struct rpcent *rpc;

- get_myaddress(&addr);
+ get_myloopaddress(&addr);

for (list = pmap_getmaps(&addr); list; list = list->pml_next) {
rpc = getrpcbynumber((int) list->pml_map.pm_prog);
diff -Naur portmap_5beta/pmap_set.8 portmap_5beta.gimli/pmap_set.8
--- portmap_5beta/pmap_set.81969-12-31 18:00:00.000000000 -0600
+++ portmap_5beta.gimli/pmap_set.82002-01-07 09:13:58.000000000 -0600
@@ -0,0 +1,24 @@
+.TH PMAP_SET 8 "21th June 1997" Linux "Linux Programmer's Manual"
+.SH NAME
+pmap_set \- set the list of registered RPC programs
+.SH SYNOPSIS
+.B pmap_set
+.SH DESCRIPTION
+The
+.B pmap_set
+command can be used to restart a running portmapper or to set
+the list of registered RPC programs on the local host. If you
+want to use the program to restart the portmapper you have to
+redirect the output of
+.B pmap_dump
+to a file. After this you can restart the portmapper and restore
+the list of the registered RPC programs by feeding the output
+of
+.B pmap_dump
+to the
+.B pmap_set
+command.
+.SH SEE ALSO
+.BR pmap_dump (8),
```

```
+BR rpc.portmap (8)
+
diff -Naur portmap_5beta/pmap_set.c portmap_5beta.gimli/pmap_set.c
--- portmap_5beta/pmap_set.c1996-07-06 16:06:23.000000000 -0500
+++ portmap_5beta.gimli/pmap_set.c2002-01-07 09:22:10.000000000 -0600
@@ -5,7 +5,7 @@
    * Computing Science, Eindhoven University of Technology, The Netherlands.
    */

-#ifndef lint
+#ifdef lint
static char sccsid[] = "@(#) pmap_set.c 1.2 96/07/06 23:06:23";
#endif

@@ -17,7 +17,9 @@
#include <rpc/rpc.h>
#include <rpc/pmap_clnt.h>

-main(argc, argv)
+int parse_line(char *buf, u_long *prog, u_long *vers, int *prot, unsigned *port);
+
+int main(argc, argv)
int    argc;
char **argv;
{
@@ -40,16 +42,16 @@

/* parse_line - convert line to numbers */

-parse_line(buf, prog, vers, prot, port)
+int parse_line(buf, prog, vers, prot, port)
char  *buf;
u_long *prog;
```

```
u_long *vers;
int *prot;
unsigned *port;
{
- char proto_name[BUFSIZ];
+ char proto_name[256];

- if (sscanf(buf, "%lu %lu %s %u", prog, vers, proto_name, port) != 4) {
+ if (sscanf(buf, "%lu %lu %25s %u", prog, vers, proto_name, port) != 4) {
return (0);
}
if (strcmp(proto_name, "tcp") == 0) {
diff -Naur portmap_5beta/portmap.8 portmap_5beta.gimli/portmap.8
--- portmap_5beta/portmap.81969-12-31 18:00:00.000000000 -0600
+++ portmap_5beta.gimli/portmap.82002-01-07 09:13:58.000000000 -0600
@@ -0,0 +1,146 @@
+.\" Copyright (c) 1987 Sun Microsystems
+.\" Copyright (c) 1990, 1991 The Regents of the University of California.
+.\" All rights reserved.
+.\"
+.\" Redistribution and use in source and binary forms, with or without
+.\" modification, are permitted provided that the following conditions
+.\" are met:
+.\" 1. Redistributions of source code must retain the above copyright
+.\" notice, this list of conditions and the following disclaimer.
+.\" 2. Redistributions in binary form must reproduce the above copyright
+.\" notice, this list of conditions and the following disclaimer in the
+.\" documentation and/or other materials provided with the distribution.
+.\" 3. All advertising materials mentioning features or use of this software
+.\" must display the following acknowledgement:
+.\" This product includes software developed by the University of
+.\" California, Berkeley and its contributors.
+.\" 4. Neither the name of the University nor the names of its contributors
```

```
+.\ " may be used to endorse or promote products derived from this software
+.\ " without specific prior written permission.
+.\ "
+.\ " THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS"
AND
+.\ " ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE
+.\ " IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE
+.\ " ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE
LIABLE
+.\ " FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL
+.\ " DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE
GOODS
+.\ " OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
+.\ " HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN
CONTRACT, STRICT
+.\ " LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN
ANY WAY
+.\ " OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
+.\ " SUCH DAMAGE.
+.\ "
+.\ " from: @(#)portmap.85.3 (Berkeley) 3/16/91
+.\ "$Id: portmap-5beta-compilation_fixes-3.patch,v 1.1 2004/06/08 04:53:09 jim Exp $
+.\ "
+.Dd March 16, 1991
+.Dt PORTMAP 8
+.Os BSD 4.3
+.Sh NAME
+.Nm portmap
+.Nd
+.Tn DARPA
+port to
+.Tn RPC
```

+program number mapper
+.Sh SYNOPSIS
+.Nm portmap
+.Op Fl d
+.Op Fl v
+.Sh DESCRIPTION
+.Nm Portmap
+is a server that converts
+.Tn RPC
+program numbers into
+.Tn DARPA
+protocol port numbers.
+It must be running in order to make
+.Tn RPC
+calls.
+.Pp
+When an
+.Tn RPC
+server is started, it will tell
+.Nm portmap
+what port number it is listening to, and what
+.Tn RPC
+program numbers it is prepared to serve.
+When a client wishes to make an
+.Tn RPC
+call to a given program number,
+it will first contact
+.Nm portmap
+on the server machine to determine
+the port number where
+.Tn RPC
+packets should be sent.
+.Pp

+Nm Portmap
+must be started before any
+.Tn RPC
+servers are invoked.
+.Pp
+Normally
+.Nm portmap
+forks and dissociates itself from the terminal
+like any other daemon.
+.Nm Portmap
+then logs errors using
+.Xr syslog 3 .
+.Pp
+Option available:
+.Bl -tag -width Ds
+.It Fl d
+(debug) prevents
+.Nm portmap
+from running as a daemon,
+and causes errors and debugging information
+to be printed to the standard error output.
+.It Fl v
+(verbose) run
+.Nm portmap
+in verbose mode.
+.El
+
+This
+.Nm portmap
+version is protected by the
+.Nm tcp_wrapper
+library. You have to give the clients access to
+.Nm portmap

+if they should be allowed to use it. To allow connects from clients of
+the .bar.com domain you could use the following line in /etc/hosts.allow:
+
+portmap: .bar.com
+
+You have to use the daemon name
+.Nm portmap
+for the daemon name (even if the binary has a different name). For the
+client names you can only use the keyword ALL or IP addresses (NOT
+host or domain names).
+
+For further information please have a look at the
+.Xr tcpd 8 ,
+.Xr hosts_allow 5
+and
+.Xr hosts_access 5
+manual pages.
+
+.Sh SEE ALSO
+.Xr inetd.conf 5 ,
+.Xr rpcinfo 8 ,
+.Xr pmap_set 8 ,
+.Xr pmap_dump 8 ,
+.Xr inetd 8
+.Xr tcpd 8
+.Xr hosts_access 5
+.Xr hosts_options 5
+.Sh BUGS
+If
+.Nm portmap
+crashes, all servers must be restarted.
+.Sh HISTORY
+The

```
+ .Nm
+ command appeared in
+ .Bx 4.3
diff -Naur portmap_5beta/portmap.c portmap_5beta.gimli/portmap.c
--- portmap_5beta/portmap.c1996-07-06 16:06:24.000000000 -0500
+++ portmap_5beta.gimli/portmap.c2002-01-07 09:26:41.000000000 -0600
@@ -37,7 +37,7 @@
     All rights reserved.\n";
#endif /* not lint */

-#ifndef lint
+#ifdef lint
static char sccsid[] = "@(#) portmap.c 1.6 96/07/06 23:06:23";
#endif /* not lint */

@@ -80,6 +80,9 @@
    * Mountain View, California  94043
    */

+#if defined(__GLIBC__)
+#include <rpc/xdr.h>
+#endif /* __GLIBC__ */
#include <rpc/rpc.h>
#include <rpc/pmap_prot.h>
#include <stdio.h>
@@ -94,6 +97,8 @@
#ifdef SYSV40
#include <netinet/in.h>
#endif
#include <sys/types.h>
#include <unistd.h>

extern char *strerror();
```



```
#include <stdlib.h>
@@ -148,7 +153,7 @@
#endif
#endif

-main(argc, argv)
+int main(argc, argv)
int argc;
char **argv;
{
@@ -350,7 +355,7 @@
    /*
    /* remote host authorization check */
    check_default(svc_getcaller(xprt), rqstp->rq_proc, (u_long) 0);
    -if (!svc_sendreply(xprt, xdr_void, (caddr_t)0) && debugging) {
    +if (!svc_sendreply(xprt, (xdrproc_t) xdr_void, (caddr_t)0) && debugging) {
    abort();
    }
    break;
@@ -359,7 +364,7 @@
    /*
    * Set a program, version to port mapping
    */
    -if (!svc_getargs(xprt, xdr_pmap, &reg))
    +if (!svc_getargs(xprt, (xdrproc_t) xdr_pmap, (caddr_t) &reg))
    svcerr_decode(xprt);
    else {
    /* reject non-local requests, protect priv. ports */
@@ -401,7 +406,7 @@
    ans = 1;
    }
    done:
    -if (!svc_sendreply(xprt, xdr_int, (caddr_t)&ans)) &&
```

```
+if ((!svc_sendreply(xprt, (xdrproc_t) xdr_int, (caddr_t)&ans)) &&
    debugging) {
(void) fprintf(stderr, "svc_sendreply\n");
abort();
@@ -413,7 +418,7 @@
/*
 * Remove a program,version to port mapping.
 */
-if (!svc_getargs(xprt, xdr_pmap, &reg))
+if (!svc_getargs(xprt, (xdrproc_t) xdr_pmap, (caddr_t) &reg))
    svcerr_decode(xprt);
else {
    ans = 0;
@@ -447,7 +452,7 @@
    prevpml->pml_next = pml;
    free(t);
}
-if (!svc_sendreply(xprt, xdr_int, (caddr_t)&ans)) &&
+if ((!svc_sendreply(xprt, (xdrproc_t) xdr_int, (caddr_t)&ans)) &&
    debugging) {
(void) fprintf(stderr, "svc_sendreply\n");
abort();
@@ -459,7 +464,7 @@
/*
 * Lookup the mapping for a program,version and return its port
 */
-if (!svc_getargs(xprt, xdr_pmap, &reg))
+if (!svc_getargs(xprt, (xdrproc_t) xdr_pmap, (caddr_t) &reg))
    svcerr_decode(xprt);
else {
    /* remote host authorization check */
@@ -474,7 +479,7 @@
    port = fnd->pml_map.pm_port;
```

```
else
port = 0;
-if (!svc_sendreply(xprt, xdr_int, (caddr_t)&port)) &&
+if (!svc_sendreply(xprt, (xdrproc_t) xdr_int, (caddr_t)&port)) &&
    debugging) {
(void) fprintf(stderr, "svc_sendreply\n");
abort();
@@ -486,7 +491,7 @@
/*
 * Return the current set of mapped program,version
 */
-if (!svc_getargs(xprt, xdr_void, NULL))
+if (!svc_getargs(xprt, (xdrproc_t) xdr_void, (caddr_t) NULL))
    svcerr_decode(xprt);
else {
/* remote host authorization check */
@@ -497,7 +502,7 @@
} else {
p = pmaplist;
}
-if (!svc_sendreply(xprt, xdr_pmaplist,
+if (!svc_sendreply(xprt, (xdrproc_t) xdr_pmaplist,
    (caddr_t)&p)) && debugging) {
(void) fprintf(stderr, "svc_sendreply\n");
abort();
@@ -645,7 +650,7 @@
timeout.tv_sec = 5;
timeout.tv_usec = 0;
a.rmt_args.args = buf;
-if (!svc_getargs(xprt, xdr_rmtcall_args, &a))
+if (!svc_getargs(xprt, (xdrproc_t) xdr_rmtcall_args, (caddr_t) &a))
return;
/* host and service access control */
```

```
if (!check_callit(svc_getcaller(xprt),
@@ -674,9 +679,9 @@
    au->aup_uid, au->aup_gid, au->aup_len, au->aup_gids);
}
a.rmt_port = (u_long)port;
-if (clnt_call(client, a.rmt_proc, xdr_opaque_parms, &a,
-   xdr_len_opaque_parms, &a, timeout) == RPC_SUCCESS) {
-svc_sendreply(xprt, xdr_rmtcall_result, (caddr_t)&a);
+if (clnt_call(client, a.rmt_proc, (xdrproc_t) xdr_opaque_parms, (char*) &a,
+   (xdrproc_t) xdr_len_opaque_parms, (char*) &a, timeout) == RPC_SUCCESS) {
+svc_sendreply(xprt, (xdrproc_t) xdr_rmtcall_result, (caddr_t)&a);
}
AUTH_DESTROY(client->cl_auth);
clnt_destroy(client);
```

portmap-5beta-disable-tcp_wrapper.patch

```
--- portmap_5beta/Makefile.orig 2005-11-21 14:56:22.000000000 +0200
+++ portmap_5beta/Makefile 2005-11-21 14:57:09.000000000 +0200
@@ -15,9 +15,9 @@
# no access control tables. The local system, since it runs the portmap
# daemon, is always treated as an authorized host.

-HOSTS_ACCESS= -DHOSTS_ACCESS
+#HOSTS_ACCESS= -DHOSTS_ACCESS
#WRAP_LIB = $(WRAP_DIR)/libwrap.a
-WRAP_LIB = -lwrap
+#WRAP_LIB = -lwrap

# Comment out if your RPC library does not allocate privileged ports for
# requests from processes with root privilege, or the new portmap will
```

portmap-5beta-glibc_errno_fix-1.patch

```
Submitted By: Tushar Teredesai <tushar@linuxfromscratch.org>
Date: 2003-10-04
Initial Package Version: 5beta
Origin: None
Description: Fix compilation with recent glibc versions.
--- portmap_5beta/portmap.c 2003-03-10 12:32:26.000000000 -0600
+++ portmap_5beta/portmap.c.new 2003-03-10 12:38:01.000000000 -0600
@@ -129,7 +129,8 @@
static void callit();
struct pmaplist *pmaplist;
int debugging = 0;
-extern int errno;
+#include <errno.h>
+/* extern int errno; */
#include "pmap_check.h"
```

proftpd-1.3.0rc1-ptx.diff

```
#
# Submitted:
#
# Robert Schwebel, 2004-08-12
#
# Error:
#
# Makefile is trying to compile host tool with cross compiler
#
# Description:
#
# Makefile is trying to compile host tool with cross compiler
#
# State:
#
# unknown, not not sent upstream yet
#

diff -urN proftpd-1.3.0rc1-orig/lib/libcap/Makefile proftpd-1.3.0rc1/lib/libcap/Makefile
--- proftpd-1.3.0rc1-orig/lib/libcap/Makefile2003-05-15 02:49:13.000000000 +0200
+++ proftpd-1.3.0rc1/lib/libcap/Makefile2005-05-22 09:51:20.000000000 +0200
@@ -9,6 +9,10 @@
#
topdir=$(shell pwd)/..
include ../../Make.rules
+ifndef CC_FOR_BUILD
+CC_FOR_BUILD=$(CC)
+endif
+
#
# Library version
#
```

```
@@ -26,7 +30,7 @@
```

```
all: $(LIBNAME)
```

```
_makenames: _makenames.c cap_names.sed
```

```
-$ (CC) $(CFLAGS) $(LDFLAGS) $< -o $@
```

```
+$ (CC_FOR_BUILD) $< -o $@
```

```
cap_names.h: _makenames
```

```
./_makenames > cap_names.h
```

smartmontools-5.36-01-sprintf-cross.patch

```
--- smartmontools-5.36/configure.orig2006-06-14 16:05:13.972167068 +0300
```

```
+++ smartmontools-5.36/configure2006-06-14 16:04:14.906349560 +0300
```

```
@@ -6442,7 +6442,7 @@
```

```
echo "$as_me:$LINENO: checking for working sprintf" >&5
```

```
echo $ECHO_N "checking for working sprintf... $ECHO_C" >&6
```

```
-if test "$cross_compiling" = yes; then
```

```
+if test "$cross_compiling" = sagi_yes; then
```

```
  { { echo "$as_me:$LINENO: error: cannot run test program while cross compiling
```

```
See \ `config.log' for more details." >&5
```

```
echo "$as_me: error: cannot run test program while cross compiling
```

squashfs2.2-r2-Makefile.patch

```
diff -Nuar squashfs2.2-r2/squashfs-tools/Makefile squashfs2.2-r2.patched/squashfs-tools/  
Makefile
```

```
--- squashfs2.2-r2/squashfs-tools/Makefile2005-09-01 02:21:14.000000000 +0300
```

```
+++ squashfs2.2-r2.patched/squashfs-tools/Makefile2006-01-09 12:08:34.000000000 +0200
```

```
@@ -3,7 +3,7 @@
```

```
CFLAGS := -I$(INCLUDEDIR) -D_FILE_OFFSET_BITS=64 -D_LARGEFILE_SOURCE -g
```

```
mksquashfs: mksquashfs.o read_fs.o sort.o
```

```
-(CC) mksquashfs.o read_fs.o sort.o -lz -o $@  
+$(CC) $(CFLAGS) $(LDFLAGS) mksquashfs.o read_fs.o sort.o -lz -o $@
```

```
mksquashfs.o: mksquashfs.c mksquashfs.h
```

generic-include-fix.diff

```
#  
# with this patch the includes get installed into <prefix>/include  
# not just /usr/include  
#  
diff -ruN termcap-1.3.1-orig/Makefile.in termcap-1.3.1/Makefile.in  
--- termcap-1.3.1-orig/Makefile.inThu Aug 17 02:54:29 1995  
+++ termcap-1.3.1/Makefile.inFri Nov 14 15:49:14 2003  
@@ -46,7 +46,7 @@  
# so compilers besides gcc can find it by default.  
# If it is empty or not defined, termcap.h will only be installed in  
# includedir.  
-oldincludedir = /usr/include  
+oldincludedir = $(prefix)/include  
  
# Directory in which to install the documentation info files.  
infodir = $(prefix)/info
```