

# 5.2-RELEASE TODO

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/current/2003-09/1666.html>

---

**From:** Robert Watson ([rwatson\\_at\\_FreeBSD.org](mailto:rwatson_at_FreeBSD.org))

**Date:** 09/29/03

Date: Mon, 29 Sep 2003 10:00:18 -0400 (EDT)

To: [current@FreeBSD.org](mailto:current@FreeBSD.org)

This is an automated bi-weekly mailing of the FreeBSD 5.2 open issues list.  
The live version of this list is available at:

<http://www.FreeBSD.org/releases/5.2R/todo.html>

Automated mailing of this list will continue through the release of  
FreeBSD 5.2.

## FreeBSD 5.2 Open Issues

### Open Issues

This is a list of open issues that need to be resolved for FreeBSD 5.2. If  
you have any updates for this list, please e-mail [re@FreeBSD.org](mailto:re@FreeBSD.org).

### Must Resolve Issues for 5.2-RELEASE

Issue	Status	Responsible	Description
KSE M:N threading			
support is			
reaching			
experimental yet			
Julian			usable status on
Production-quality	In	Elischer, David	i386 for
M:N threading	progress	Xu, Daniel	5.1-RELEASE. M:N
Eischen			threading should
be productionable			
and usable on all			
platforms by			
5.2-RELEASE.			
Kernel bits are			
implemented but			
KSE support for	In		untested. Userland
sparc64	progress	Jake Burkholder	bits are not

freebsd-current: 5.2-RELEASE TODO

			implemented.
			Required for
			5.2-RELEASE.
-----+-----+-----+-----+			
			Kernel and
KSE support for		Marcel	userland bits
ia64	Complete.	Moolenaar	implemented but
			unstable. Required
			for 5.2-RELEASE.
-----+-----+-----+-----+			
			Userland bits
			implemented,
KSE support for	In	Marcel	kernel bits not
alpha	progress.	Moolenaar	implemented.
			Required for
			5.2-RELEASE.
-----+-----+-----+-----+			
			Kris Kennaway
			reports high
			instability of
			5-CURRENT on ia64
			machines, such as
ia64 stability	In	Marcel	the pluto\*
Progress	Moolenaar	machines. These	
			problems need to
			be fixed in order
			to get a
			successful package
			build.
-----+-----+-----+-----+			
			A reworking of the
			sio driver is
			needed to support
			serial terminal
	Marcel	devices on sparc64	
New serial UART	In	Moolenaar,	and ia64
framework	progress	Warner Losh	platforms, among
			others. This is
			also an enabler
			for syscons
			support on
			sparc64.
-----+-----+-----+-----+			
			FAST\_IPSEC
			currently cannot
			be used directly
			with the KAME IPv6
			implementation,
			requiring an
			additional level
			of IP tunnel

			indirection to
			protect IPv6
			packets when using
			hardware crypto
FAST\_IPSEC and KAME			acceleration. This
compatibility	--	--	issue must be
			resolved so that
			the two services
			may more easily be
			used together.
			Among other
			things, this will
			require a careful
			review of the
			handling of mbuf
			header copying and
			m\_tag support in
			the KAME IPv6
			code.

---

			The FreeBSD KAME
			IPv6 code is now
			substantially
			dated with respect
			to the KAME vendor
KAME	In	Hajimu UMEMOTO	source. The
Synchronization	progress		FreeBSD Project
			needs to take
			initiative in
			driving the merge
			of new bug fixes,
			features, et al.

---

			Almost all process
			debugging tools
			have been updated
			to use non-procfs
			kernel primitives,
			with the exception
			of truss(1). As
			procfs is
			considered
			deprecated due to
			its inherent
truss support for	In	Robert Drehmel	security risks, it
ptrace	progress		is highly
			desirable to
			update truss to
			operate in a
			post-procfs world.
			Dag-Erling

		Smorgrav had
		prototype patches;
		Robert Drehmel is
		developing and
		testing patches
		now.

---

		Apple's Darwin	
		operating system	
		has fairly	
		extensive	
Merge of Darwin			improvements to
msdosfs, other	--	--	msdosfs and other
fixes			kernel services;
		these fixes must	
		be reviewed and	
		merged to the	
		FreeBSD tree.	

---

		Many systems		
		supporting		
		POSIX.1e ACLs		
		permit a minor		
		violation to that		
		specification, in		
		which the ACL\_MASK		
		entry overrides		
ACL\_MASK override	In			the umask, rather
of umask support in	progress	Robert Watson	than being	
UFS			intersected with	
		it. The resulting		
		semantics can be		
		useful in		
		group-oriented		
		environments, and		
		as such would be		
		very helpful on		
		FreeBSD.		

---

		Significant parts
		of the network
		stack (especially
		IPv4 and IPv6) now
		have fine-grained
		locking of their
		data structures.
		However, it is not
		yet possible for
		the netisr threads
		to run without
		Giant, due to

Fine-grained			dependencies on	
network stack	In	Jeffrey Hsu,	sockets, routing,	
locking without	progress	Seigo Tanimura,	etc. A 5.2-RELEASE	
Giant			Sam Leffler	goal is to have
		the network stack		
		running largely		
		without Giant,		
		which should		
		substantially		
		improve		
		performance of the		
		stack, as well as		
		other system		
		components by		
		reducing		
		contention on		
		Giant.		

---

		Productionable	
		support for the	
		AMD64 platform.	
		Currently, AMD64	
		runs fully in	
		32-bit emulation	
Tier-1 Support for	In	Peter Wemm,	mode, and boots to
AMD64 Hammer	progress	David O'Brien	single-user in
		64-bit mode. We	
		expect full	
		production support	
		for the AMD64	
		architecture in	
		5.2-RELEASE.	

---

		Kernel modules are
		currently built
		independently from
		a kernel
		configuration, and
		independently from
		one another,
		resulting in
		substantially
		redundant
		compilation of
		objects, as well
		as the inability
		to easily manage
		compile-time
		options for kernel
		objects (such as
		MAC, PAE, etc)

Revised kld build	--	--	that may require
infrastructure		conditional	
		compilation in the	
		kernel modules. In	
		order to improve	
		build performance	
		and better support	
		options of this	
		sort, the KLD	
		build	
		infrastructure	
		needs to be	
		revamped. Peter	
		Wemm has done some	
		initial	
		prototyping, and	
		should be	
		contacted before	
		starting on this	
		work.	

---

		Currently, there	
		are two classes of	
		interrupt handlers	
		in 5.x: fast	
		interrupt handlers	
		which run entirely	
		in interrupt	
		context, and	
		heavy-weight	
		handlers which	
		execute in a	
		full-weight kernel	
		interrupt thread.	
		It is possible to	
		optimize interrupt	
		thread context	
		management such	
		that a	
		light-weight	
		context switch is	
		performed to begin	
		execution of the	
		interrupt thread	
		in the handler	
		context, and only	
Light-weight		when a full-weight	
interrupt threads,	--	--	context is
context switches		required (such as	
		sleeping on a	
		lock) is that cost	

		required. This
		optimization
		should
		substantially
		improve interrupt
		latency. There are
		also additional
		kernel thread
		context switch
		optimizations that
		can be made to
		improve the
		performance of
		thread workers in
		the kernel, such
		as found in the
		network stack,
		crypto worker
		threads, and GEOM.
		Bosko Milekic has
		done substantial
		prototyping work,
		and should be
		coordinated with.

---

		The existing APIC	
		interrupt code	
		does not support	
Complete the APIC			PCI interrupt
PCI interrupt	In	John Baldwin	routing properly.
routing support	progress		As a result, PCI
		interrupts cannot	
		be routed either	
		via ACPI or across	
		PCI-PCI bridges.	

---

		Currently, gbde	
		must be manually	
		configured at	
		run-time each time	
		an encrypted disk	
		device is mounted.	
		This prevents easy	
Run-time			integration into
autoconfiguration			/etc/fstab and
of GBDE and related	--	--	easy automated
transforms			deployment.
		Improved	
		integration with	
		the configuration,	
		mounting, and boot	

## freebsd-current: 5.2-RELEASE TODO

		process is
		required to make
		this feature more
		easily accessible.

---

		Brian Feldman has	
		submitted patches	
		to improve the	
		consistency of the	
		pathnames passed	
		into the MAC	
MAC Framework devfs	In	Robert Watson	Framework devfs
path fixes	progress		labeling entry
		points. These	
		patches need to be	
		thoroughly	
		reviewed and	
		tested, then	
		merged.	

---

		A process cannot	
		be interrupted	
		while waiting on a	
		lock. Fixing this	
rpc.lockd(8)	In	Robert Watson	requires that the
stability	progress		rpc code be taught
		how to deal with	
		lock cancellation	
		and interruption	
		events.	

### Desired Features for 5.2-RELEASE

---

Issue	Status	Responsible	Description
Truss appears to			
contain a race			
condition during the			
start-up of debugging,			
which can result in			
truss failing to attach			
to the process before			
it exits. The symptom			
is that truss reports			
that it cannot open the			
procfs node supporting			
the process being			
Race     debugged. A bug also			
conditions in   Errata   Robert Drehmel   appears to exist where			

truss	candidate		in truss will hang if
		execve() returns	
		ENOENT. A further race	
		appears to exist in	
		which truss will return	
		"PIOCWAIT: Input/output	
		error" occasionally on	
		startup. The fix for	
		this sufficiently	
		changes process	
		execution handling that	
		we will defer the fix	
		to post-5.0 and	
		consider this errata.	

---

gdb -k			gdb -k doesn't work on
support for	--	Mark Peek	alpha
alpha			

---

		Currently, MAC	
		protections are	
		enforced only on	
		locally originated file	
		system operations	
		(VOPs), and not on RPCs	
		generated via the NFS	
MAC support			server. Improvements in
for NFS	--	Robert Watson	NFS server credential
Server			handling are required
		to correct this	
		problem, as well as the	
		introduction of new	
		entry points to	
		properly label NFS	
		credentials and perform	
		enforcement properly.	

---

		All PCI drivers must	
		use busdma for DMA; no	
busdma in all			use of vtophys() will
PCI drivers	--	--	be permitted for any
		recent device driver.	
		ISA drivers may be	
		exempt.	

---

		With improved support	
		for threading	
		primitives, support is	
		now required to ease	
GDB thread	--	--	debugging of threaded
support			applications. Ideally,

		this support will work
		for both libthr and
		libkse threading
		models.

-----+-----+-----+-----+  
		Prebinding reduces	
		executable startup time	
		by lowering the expense	
		of symbol lookup,	
		binding and relocation.	
		This is accomplished by	
		a prebinding data file	
		or ELF segment that	
		contains intermediate	
		lookup results allowing	
		fast symbol binding and	
		relocation, provided	
Per object		that dependent objects	
ELF	In progress	Matthew Dodd	remain unchanged since
Prebinding		the prebinding	
support		information was	
		generated.	
		The benefits of	
		prebinding are realized	
		when running	
		executables that use a	
		large (>10) number of	
		shared libraries. C++	
		applications also	
		benefit as they contain	
		a large number of	
		relocations.	

-----+-----+-----+-----+  
Documentation items that must be resolved for 5.2

-----+-----+-----+-----+  
| Issue | Status | Responsible | Description |  
-----+-----+-----+-----+  
Bluetooth		It'd be nice to have some	
documentation	--	Pav Lucistnik	Bluetooth documentation for
		the Handbook.	

-----+-----+-----+-----+  
Testing focuses for 5.2-RELEASE

-----+-----+-----+-----+  
| Issue | Status | Responsible | Description |  
-----+-----+-----+-----+  
||| New ATA model has arrived. |

## freebsd-current: 5.2-RELEASE TODO

ATA driver			Much testing is needed to
structural	Complete	So/ren Schmidt	ensure no regressions.
improvements,			ATAPI-CAM seems especially
MPsafty			fragile, more testing is
		encouraged.	

---

home | contact | legal | (c) 1995-2003 The FreeBSD Project.  
All rights reserved.  
Last modified: 2003/09/01 06:09:17

---

freebsd-current@freebsd.org mailing list  
<http://lists.freebsd.org/mailman/listinfo/freebsd-current>  
To unsubscribe, send any mail to "freebsd-current-unsubscribe@freebsd.org"