

5.2-RELEASE TODO

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/current/2003-06/1746.html>

From: Robert Watson (rwatson_at_FreeBSD.org)

Date: 06/29/03

Date: Sun, 29 Jun 2003 10:00:15 -0400 (EDT)

To: 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.

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.			
Currently,			the MD
elements			of KSE are
present			only for
the i386			platform,
limiting			use of KSE

freebsd-current: 5.2-RELEASE TODO

		to the i386	
		platform. It is	
		highly desirable to	
KSE support for		Jake	make KSE available
sparc64, alpha,	--	Burkholder, --,	on non-i386
ia64		--	platforms for
		5.2-RELEASE so that	
		KSE can see more	
		broad exposure, and	
		the performance	
		benefits of KSE can	
		be visible to users	
		of the 64-bit	
		FreeBSD	
		architectures.	

		Kris Kennaway	
		reports high	
		instability of	
		5-CURRENT on ia64	
	In	Marcel	machines, such as
ia64 stability	Progress	Moolenaar	the pluto*
		machines. These	
		problems need to be	
		fixed in order to	
		get a successful	
		package build.	

		ia64 serial console	
		support is reported	
		to not be	
		functional on HP	
	In	Marcel	Itanium2 platforms.
ia64 sio support	progress	Moolenaar,	A reworking of the
		Warner Losh	sio driver to
		improve platform	
		independence and	
		bus handling is	
		likely needed.	

		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			KAME vendor source.
Synchronization	--	--	The 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		
truss support for	In			its inherent
ptrace	progress	Robert Drehmel	security risks, it	
		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.	

		Port syscons to	
		sparc64. Add device	
		drivers for sun	
		mice and keyboards.	
		Allow for more than	
sparc64 adaptation	In		3 bits of
of syscons	progress	Jake Burkholder	background colour
		in syscons. Creator	
		frame buffer device	
		driver. In the	
		process, generally	
		improve the MI-ness	
		of syscons.	

		Many systems	
		supporting POSIX.1e	
		ACLs permit a minor	
		violation to that	
		specification, in	
		which the ACL_MASK	
		entry overrides the	
ACL_MASK override	In		umask, rather than
of umask support in	progress	Robert Watson	being intersected
UFS			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			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.	

		Move ATA commands		
		into requests, so		
		that they can be		
		linked together,		
		centralize request		
		queue management		
		permitting direct		
ATA driver			DMA of ATA	
structural	In	So/ren Schmidt	commands. These	
improvements,	progress			architectural
MPSafety			changes will also	
		facilitate complete		
		MP-safety of the		
		ATA driver suite,		
		and dramatically		
		improve support for		
		Promise RAID		
		controllers.		

		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.

		To properly support	
		AMD64, a compiler	
		upgrade is	
		required; however,	
		the new GCC 3.3 has	
		David O'Brien	a substantial
GCC 3.3 upgrade	In	Alexander	number of new
progress	Kabaev	warnings that will	
		disrupt application	
		support, preventing	
		inclusion in	
		5.1-RELEASE. This	
		is a must-have for	
		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	
Revised kld build	--	--	MAC, PAE, etc) that
infrastructure			may require
		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 required
		context switches			(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.

		With the	
		introduction of	
		extensive PAM and	
		NSS support in	
		5.0-RELEASE and	
		5.1-RELEASE,	
Support for a fully	In	Gordon Tetlow	support for a fully
dynamic system	progress		dynamically linked
		system is desired.	
		Gordon Tetlow has	
		in-progress patches	
		to post for review	
		following the 5.1	
		release cycle.	

		Existing interrupt	
		routing code is not	
		able to correctly	
		determine the	
		interrupt routing	
Complete the APIC			of PCI devices that
PCI interrupt	--	--	are behind
routing support			PCI-to-PCI bridges
		when using APIC	
		Full Table routing	
		(with non-legacy	
		interrupts above	
		15).	

		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 easy
of GBDE and related	--	--	automated
transforms			deployment.
		Improved	
		integration with	
		the configuration,	

freebsd-current: 5.2-RELEASE TODO

		mounting, and boot
		process is required
		to make this
		feature more easily
		accessible.

		Brian Feldman has	
		submitted patches	
		to improve the	
		consistency of the	
		pathnames passed	
MAC Framework devfs	In		into the MAC
path fixes	progress	Robert Watson	Framework devfs
		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		
		debugged. A		
		bug also		
Race			appears to	
conditions	Errata	Robert Drehmel	exist where	
in 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. ||
 |-----|
 ||| Kris Kennaway ||
 ||| reports ||
 ||| deadlocks ||
 ||| involving the ||
 ||| use of nullfs ||
 ||| in the bento ||
 ||| environment: ||
 ||| buildworld ||
 ||| -j4 with src ||
 ||| and obj ||

```

| | | mounted via | | |
| | | nullfs; the | |
| nullfs | -- | -- | gcc processes | |
| deadlocks | | | eventually | |
| | | deadlocked in | |
| | | the ufs | |
| | | state. DDB | |
| | | traceback | |
| | | showed two | |
| | | different | |
| | | codepaths. | |
| | | I've just | |
| | | repeated | |
| | | this, so the | |
| | | bug still | |
| | | exists. | | |
|---|---|---|---|---|
| gdb -k | | | gdb -k | |
| support | -- | Mark Peek | doesn't work | |
| for alpha | | | on alpha | |
|-----|
| | | Currently, | |
| | | MAC | |
| | | protections | |
| | | are enforced | |
| | | only on | |
| | | locally | |
| | | originated | |
| | | file system | |
| | | operations | |
| | | (VOPs), and | |
| | | not on RPCs | |
| | | generated via | |
| | | the NFS | |
| MAC | | | server. | |
| support | | | Improvements | |
| for NFS | -- | Robert Watson | in NFS server | |
| Server | | | credential | |
| | | 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 ||
||| use of ||
| busdma in || vtophys() ||
| all PCI | -- | -- | will be ||
| drivers || 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 ||
| support | -- | -- | threaded ||
||| 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, | The benefits |
||| binding and | of |
||| relocation. | prebinding |
||| This is | are realized |
||| accomplished | when running |
||| by a | executables |
||| prebinding | that use a |
||| data file or | large (>10) |
| Per object || ELF segment | number of |
| ELF | In | Matthew Dodd | that contains | shared |

```

freebsd-current: 5.2-RELEASE TODO

Prebinding	progress		intermediate	libraries.
support			lookup	C++
		results	applications	
		allowing fast	also benefit	
		symbol	as they	
		binding and	contain a	
		relocation,	large number	
		provided that	of	
		dependent	relocations.	
		objects		
		remain		
		unchanged		
		since the		
		prebinding		
		information		
		was		
		generated.		

-----+
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 |
-----+

freebsd-qa@FreeBSD.ORG

Copyright (c) 1995-2003 The FreeBSD Project. All rights reserved.

Last modified: 2003/05/31 07:19:32

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"