

---

Subject: [PATCH 2/13] BC: Kconfig and Makefile  
Posted by [dev](#) on Thu, 09 Nov 2006 16:47:51 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Add kernel/bc/Kconfig file with BC options and  
include it into arch Kconfigs

Signed-off-by: Pavel Emelianov <xemul@sw.ru>

Signed-off-by: Kirill Korotaev <dev@sw.ru>

---

```
init/Kconfig | 4 +++
kernel/Makefile | 1 +
kernel/bc/Kconfig | 17 ++++++=====
kernel/bc/Makefile | 11 ++++++=====
4 files changed, 33 insertions(+)
```

```
--- ./init/Kconfig.bckconfig 2006-11-09 11:29:12.000000000 +0300
```

```
+++ ./init/Kconfig 2006-11-09 11:30:21.000000000 +0300
```

```
@@ @ -585,6 +585,10 @@ config STOP_MACHINE
```

    Need stop\_machine() primitive.

```
endmenu
```

```
+menu "Beancounters"
```

```
+source "kernel/bc/Kconfig"
```

```
+endmenu
```

```
+
```

```
menu "Block layer"
```

```
source "block/Kconfig"
```

```
endmenu
```

```
--- ./kernel/Makefile.bckconfig 2006-11-09 11:29:12.000000000 +0300
```

```
+++ ./kernel/Makefile 2006-11-09 11:30:21.000000000 +0300
```

```
@@ @ -12,6 +12,7 @@ obj-y = sched.o fork.o exec_domain.o
```

```
obj-$(CONFIG_STACKTRACE) += stacktrace.o
```

```
obj-y += time/
```

```
+obj-$(CONFIG_BEANCOUNTERS) += bc/
```

```
obj-$(CONFIG_DEBUG_MUTEXES) += mutex-debug.o
```

```
obj-$(CONFIG_LOCKDEP) += lockdep.o
```

```
ifeq ($(CONFIG_PROC_FS),y)
```

```
--- /dev/null 2006-07-18 14:52:43.075228448 +0400
```

```
+++ ./kernel/bc/Kconfig 2006-11-09 11:30:21.000000000 +0300
```

```
@@ @ -0,0 +1,17 @@
```

```
+config BEANCOUNTERS
```

```
+ bool "Enable resource accounting/control"
```

```
+ default n
```

```
+ depends on CONFIGFS_FS
```

```
+ help
+ When Y this option provides accounting and allows configuring
+ limits for user's consumption of exhaustible system resources.
+ The most important resource controlled by this patch is unswappable
+ memory (either mlock'ed or used by internal kernel structures and
+ buffers). The main goal of this patch is to protect processes
+ from running short of important resources because of accidental
+ misbehavior of processes or malicious activity aiming to ``kill''
+ the system. It's worth mentioning that resource limits configured
+ by setrlimit(2) do not give an acceptable level of protection
+ because they cover only a small fraction of resources and work on a
+ per-process basis. Per-process accounting doesn't prevent malicious
+ users from spawning a lot of resource-consuming processes.
--- /dev/null 2006-07-18 14:52:43.075228448 +0400
+++ ./kernel/bc/Makefile 2006-11-09 11:31:24.000000000 +0300
@@ -0,0 +1,11 @@
+#
+## kernel/bc/Makefile
+##
+## Copyright (C) 2006 OpenVZ SWsoft Inc.
+#
+#
+obj-y = beancounter.o vmpages.o rsspages.o kmem.o misc.o
+
+obj-$(CONFIG_CONFIGFS_FS) += bc_if.o
+
+bc_if-objs := configfs.o
```

---