
Subject: [RFC][PATCH 5/7] UBC: kernel memory accounting (core)

Posted by [dev](#) on Wed, 16 Aug 2006 15:39:03 GMT

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

Introduce UB_KMEMSIZE resource which accounts kernel objects allocated by task's request.

Reference to UB is kept on struct page or slab object.

For slabs each struct slab contains a set of pointers corresponding objects are charged to.

Allocation charge rules:

1. Pages - if allocation is performed with `__GFP_UBC` flag - page is charged to current's `exec_ub`.
2. Slabs - `kmem_cache` may be created with `SLAB_UBC` flag - in this case each allocation is charged. Caches used by `kmalloc` are created with `SLAB_UBC` | `SLAB_UBC_NOCHARGE` flags. In this case only `__GFP_UBC` allocations are charged.

Signed-Off-By: Pavel Emelianov <xemul@sw.ru>

Signed-Off-By: Kirill Korotaev <dev@sw.ru>

```
include/linux/gfp.h      |  8 +-
include/linux/mm.h       |  6 ++
include/linux/slab.h     |  4 +
include/linux/vmalloc.h  |  1
include/ub/beancounter.h |  4 +
include/ub/kmem.h        | 33 ++++++++
kernel/ub/Makefile       |  1
kernel/ub/beancounter.c  |  3 +
kernel/ub/kmem.c         | 89 +++++
mm/mempool.c             |  2
mm/page_alloc.c          | 11 ++++
mm/slab.c                | 121 +++++
mm/vmalloc.c             |  6 ++
13 files changed, 264 insertions(+), 25 deletions(-)
```

--- ./include/linux/gfp.h.kmemcore 2006-08-16 19:10:38.000000000 +0400

+++ ./include/linux/gfp.h 2006-08-16 19:12:56.000000000 +0400

@@ -46,15 +46,18 @@ struct vm_area_struct;

```
#define __GFP_NOMEMALLOC ((__force gfp_t)0x10000u) /* Don't use emergency reserves */
#define __GFP_HARDWALL   ((__force gfp_t)0x20000u) /* Enforce hardwall cpuset memory
allocs */
```

```
#define __GFP_THISNODE ((__force gfp_t)0x40000u) /* No fallback, no policies */
```

```
+#define __GFP_UBC ((__force gfp_t)0x80000u) /* Charge allocation with UB */
```

```
+#define __GFP_UBC_LIMIT ((__force gfp_t)0x100000u) /* Charge against UB limit */
```

```

#define __GFP_BITS_SHIFT 20 /* Room for 20 __GFP_FOO bits */
#define __GFP_BITS_SHIFT 21 /* Room for 20 __GFP_FOO bits */
#define __GFP_BITS_MASK ((__force gfp_t)((1 << __GFP_BITS_SHIFT) - 1))

/* if you forget to add the bitmask here kernel will crash, period */
#define GFP_LEVEL_MASK (__GFP_WAIT|__GFP_HIGH|__GFP_IO|__GFP_FS| \
    __GFP_COLD|__GFP_NOWARN|__GFP_REPEAT| \
    __GFP_NOFAIL|__GFP_NORETRY|__GFP_NO_GROW|__GFP_COMP| \
    - __GFP_NOMEMALLOC|__GFP_HARDWALL|__GFP_THISNODE)
+ __GFP_NOMEMALLOC|__GFP_HARDWALL|__GFP_THISNODE| \
+ __GFP_UBC|__GFP_UBC_LIMIT)

/* This equals 0, but use constants in case they ever change */
#define GFP_NOWAIT (GFP_ATOMIC & ~__GFP_HIGH)
@@ -63,6 +66,7 @@ struct vm_area_struct;
#define GFP_NOIO (__GFP_WAIT)
#define GFP_NOFS (__GFP_WAIT | __GFP_IO)
#define GFP_KERNEL (__GFP_WAIT | __GFP_IO | __GFP_FS)
#define GFP_KERNEL_UBC (__GFP_WAIT | __GFP_IO | __GFP_FS | __GFP_UBC)
#define GFP_USER (__GFP_WAIT | __GFP_IO | __GFP_FS | __GFP_HARDWALL)
#define GFP_HIGHUSER (__GFP_WAIT | __GFP_IO | __GFP_FS | __GFP_HARDWALL | \
    __GFP_HIGHMEM)
--- ./include/linux/mm.h.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./include/linux/mm.h 2006-08-16 19:10:51.000000000 +0400
@@ -274,8 +274,14 @@ struct page {
    unsigned int gfp_mask;
    unsigned long trace[8];
#endif
#ifdef CONFIG_USER_RESOURCE
+ union {
+ struct user_beancounter *page_ub;
+ } bc;
#endif
};

#define page_ub(page) ((page)->bc.page_ub)
#define page_private(page) ((page)->private)
#define set_page_private(page, v) ((page)->private = (v))

--- ./include/linux/slab.h.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./include/linux/slab.h 2006-08-16 19:10:51.000000000 +0400
@@ -46,6 +46,8 @@ typedef struct kmem_cache kmem_cache_t;
#define SLAB_PANIC 0x00040000UL /* panic if kmem_cache_create() fails */
#define SLAB_DESTROY_BY_RCU 0x00080000UL /* defer freeing pages to RCU */
#define SLAB_MEM_SPREAD 0x00100000UL /* Spread some memory over cpuset */
#define SLAB_UBC 0x00200000UL /* Account with UB */
#define SLAB_UBC_NOCHARGE 0x00400000UL /* Explicit accounting */

```

```

/* flags passed to a constructor func */
#define SLAB_CTOR_CONSTRUCTOR 0x001UL /* if not set, then deconstructor */
@@ -293,6 +295,8 @@ extern kmem_cache_t *bio_cachep;

extern atomic_t slab_reclaim_pages;

+struct user_beancounter;
+struct user_beancounter **kmem_cache_ubp(kmem_cache_t *cachep, void *obj);
#endif /* __KERNEL__ */

#endif /* _LINUX_SLAB_H */
--- ./include/linux/vmalloc.h.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./include/linux/vmalloc.h 2006-08-16 19:10:51.000000000 +0400
@@ -36,6 +36,7 @@ struct vm_struct {
 * Highlevel APIs for driver use
 */
extern void *vmalloc(unsigned long size);
+extern void *vmalloc_ub(unsigned long size);
extern void *vmalloc_user(unsigned long size);
extern void *vmalloc_node(unsigned long size, int node);
extern void *vmalloc_exec(unsigned long size);
--- ./include/ub/beancounter.h.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./include/ub/beancounter.h 2006-08-16 19:10:51.000000000 +0400
@@ -12,7 +12,9 @@
 * Resource list.
 */

-#define UB_RESOURCES 0
+#define UB_KMEMSIZE 0
+
+#define UB_RESOURCES 1

struct ubparm {
/*
--- ./include/ub/kmem.h.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./include/ub/kmem.h 2006-08-16 19:10:51.000000000 +0400
@@ -0,0 +1,33 @@
+/*
+ * include/ub/kmem.h
+ *
+ * Copyright (C) 2006 OpenVZ. SWsoft Inc
+ *
+ */
+
+#ifndef __UB_KMEM_H_
+#define __UB_KMEM_H_
+
+#include <linux/config.h>

```

```

+
+/*
+ * UB_KMEMSIZE accounting
+ */
+
+struct mm_struct;
+struct page;
+struct user_beancounter;
+
+#ifdef CONFIG_USER_RESOURCE
+int ub_page_charge(struct page *page, int order, gfp_t flags);
+void ub_page_uncharge(struct page *page, int order);
+
+int ub_slab_charge(kmem_cache_t *cachep, void *obj, gfp_t flags);
+void ub_slab_uncharge(kmem_cache_t *cachep, void *obj);
+#else
+#define ub_page_charge(pg, o, mask) (0)
+#define ub_page_uncharge(pg, o) do { } while (0)
+#define ub_slab_charge(cachep, o) (0)
+#define ub_slab_uncharge(cachep, o) do { } while (0)
+#endif
+#endif /* __UB_SLAB_H_ */
--- ./kernel/ub/Makefile.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./kernel/ub/Makefile 2006-08-16 19:10:51.000000000 +0400
@@ -7,3 +7,4 @@
obj-$(CONFIG_USER_RESOURCE) += beancounter.o
obj-$(CONFIG_USER_RESOURCE) += misc.o
obj-y += sys.o
+obj-$(CONFIG_USER_RESOURCE) += kmem.o
--- ./kernel/ub/beancounter.c.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./kernel/ub/beancounter.c 2006-08-16 19:10:51.000000000 +0400
@@ -20,6 +20,7 @@ static void init_beancounter_struct(stru
struct user_beancounter ub0;

const char *ub_rnames[] = {
+ "kmemsize", /* 0 */
};

#define ub_hash_fun(x) (((x) >> 8) ^ (x)) & (UB_HASH_SIZE - 1)
@@ -356,6 +357,8 @@ static void init_beancounter_syslimits(s
{
int k;

+ ub->ub_parms[UB_KMEMSIZE].limit = 32 * 1024 * 1024;
+
for (k = 0; k < UB_RESOURCES; k++)
ub->ub_parms[k].barrier = ub->ub_parms[k].limit;
}

```

```

--- ./kernel/ub/kmem.c.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./kernel/ub/kmem.c 2006-08-16 19:10:51.000000000 +0400
@@ -0,0 +1,89 @@
+/*
+ * kernel/ub/kmem.c
+ *
+ * Copyright (C) 2006 OpenVZ. SWsoft Inc
+ *
+ */
+
+#include <linux/sched.h>
+#include <linux/gfp.h>
+#include <linux/slab.h>
+#include <linux/mm.h>
+
+#include <ub/beancounter.h>
+#include <ub/kmem.h>
+#include <ub/task.h>
+
+/*
+ * Slab accounting
+ */
+
+int ub_slab_charge(kmem_cache_t *cachep, void *objp, gfp_t flags)
+{
+ unsigned int size;
+ struct user_beancounter *ub, **slab_ulp;
+
+ ub = get_exec_ub();
+ if (ub == NULL)
+ return 0;
+
+ size = kmem_cache_size(cachep);
+ if (charge_beancounter(ub, UB_KMEMSIZE, size,
+ (flags & __GFP_UBC_LIMIT ? UB_LIMIT : UB_BARRIER)))
+ return -ENOMEM;
+
+ slab_ulp = kmem_cache_ulp(cachep, objp);
+ *slab_ulp = get_beancounter(ub);
+ return 0;
+}
+
+void ub_slab_uncharge(kmem_cache_t *cachep, void *objp)
+{
+ unsigned int size;
+ struct user_beancounter *ub, **slab_ulp;
+
+ slab_ulp = kmem_cache_ulp(cachep, objp);

```

```

+ if (*slab_ubp == NULL)
+ return;
+
+ ub = *slab_ubp;
+ size = kmem_cache_size(cache);
+ uncharge_beancounter(ub, UB_KMEMSIZE, size);
+ put_beancounter(ub);
+ *slab_ubp = NULL;
+}
+
+/*
+ * Pages accounting
+ */
+
+int ub_page_charge(struct page *page, int order, gfp_t flags)
+{
+ struct user_beancounter *ub;
+
+ BUG_ON(page_ub(page) != NULL);
+
+ ub = get_exec_ub();
+ if (ub == NULL)
+ return 0;
+
+ if (charge_beancounter(ub, UB_KMEMSIZE, PAGE_SIZE << order,
+ (flags & __GFP_UBC_LIMIT ? UB_LIMIT : UB_BARRIER)))
+ return -ENOMEM;
+
+ page_ub(page) = get_beancounter(ub);
+ return 0;
+}
+
+void ub_page_uncharge(struct page *page, int order)
+{
+ struct user_beancounter *ub;
+
+ ub = page_ub(page);
+ if (ub == NULL)
+ return;
+
+ uncharge_beancounter(ub, UB_KMEMSIZE, PAGE_SIZE << order);
+ put_beancounter(ub);
+ page_ub(page) = NULL;
+}
--- ./mm/mempool.c.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./mm/mempool.c 2006-08-16 19:10:51.000000000 +0400
@@ -119,6 +119,7 @@ int mempool_resize(mempool_t *pool, int
    unsigned long flags;

```

```

BUG_ON(new_min_nr <= 0);
+ gfp_mask &= ~__GFP_UBC;

spin_lock_irqsave(&pool->lock, flags);
if (new_min_nr <= pool->min_nr) {
@@ -212,6 +213,7 @@ void * mempool_alloc(mempool_t *pool, gf
gfp_mask |= __GFP_NOMEMALLOC; /* don't allocate emergency reserves */
gfp_mask |= __GFP_NORETRY; /* don't loop in __alloc_pages */
gfp_mask |= __GFP_NOWARN; /* failures are OK */
+ gfp_mask &= ~__GFP_UBC; /* do not charge */

gfp_temp = gfp_mask & ~(__GFP_WAIT|__GFP_IO);

--- ./mm/page_alloc.c.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./mm/page_alloc.c 2006-08-16 19:10:51.000000000 +0400
@@ -38,6 +38,8 @@
#include <linux/mempolicy.h>
#include <linux/stop_machine.h>

+#include <ub/kmem.h>
+
#include <asm/tlbflush.h>
#include <asm/div64.h>
#include "internal.h"
@@ -484,6 +486,8 @@ static void __free_pages_ok(struct page
if (reserved)
return;

+ ub_page_uncharge(page, order);
+
kernel_map_pages(page, 1 << order, 0);
local_irq_save(flags);
__count_vm_events(PGFREE, 1 << order);
@@ -764,6 +768,8 @@ static void fastcall free_hot_cold_page(
if (free_pages_check(page))
return;

+ ub_page_uncharge(page, 0);
+
kernel_map_pages(page, 1, 0);

pcp = &zone_pcp(zone, get_cpu())->pcp[cold];
@@ -1153,6 +1159,11 @@ nopage:
show_mem();
}
got_pg:
+ if ((gfp_mask & __GFP_UBC) &&

```

```

+ ub_page_charge(page, order, gfp_mask)) {
+ __free_pages(page, order);
+ page = NULL;
+ }
#ifdef CONFIG_PAGE_OWNER
    if (page)
        set_page_owner(page, order, gfp_mask);
--- ./mm/slab.c.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./mm/slab.c 2006-08-16 19:10:51.000000000 +0400
@@ -108,6 +108,8 @@
#include <linux/mutex.h>
#include <linux/rtmutex.h>

+#include <ub/kmem.h>
+
#include <asm/uaccess.h>
#include <asm/cacheflush.h>
#include <asm/tlbflush.h>
@@ -175,11 +177,13 @@
    SLAB_CACHE_DMA | \
    SLAB_MUST_HWCACHE_ALIGN | SLAB_STORE_USER | \
    SLAB_RECLAIM_ACCOUNT | SLAB_PANIC | \
+   SLAB_UBC | SLAB_UBC_NOCHARGE | \
    SLAB_DESTROY_BY_RCU | SLAB_MEM_SPREAD)
#else
# define CREATE_MASK (SLAB_HWCACHE_ALIGN | \
    SLAB_CACHE_DMA | SLAB_MUST_HWCACHE_ALIGN | \
    SLAB_RECLAIM_ACCOUNT | SLAB_PANIC | \
+   SLAB_UBC | SLAB_UBC_NOCHARGE | \
    SLAB_DESTROY_BY_RCU | SLAB_MEM_SPREAD)
#endif

@@ -801,9 +805,33 @@ static struct kmem_cache *kmem_find_gene
return __find_general_cachep(size, gfpflags);
}

-static size_t slab_mgmt_size(size_t nr_objs, size_t align)
+static size_t slab_mgmt_size_raw(size_t nr_objs)
+{
+ return sizeof(struct slab) + nr_objs * sizeof(kmem_bufctl_t);
+}
+
+#ifdef CONFIG_USER_RESOURCE
+#define UB_EXTRASIZE sizeof(struct user_beancounter *)
+static inline size_t slab_mgmt_size_noalign(int flags, size_t nr_objs)
+{
+ size_t size;
+

```

```

+ size = slab_mgmt_size_raw(nr_objs);
+ if (flags & SLAB_UBC)
+ size = ALIGN(size, UB_EXTRASIZE) + nr_objs * UB_EXTRASIZE;
+ return size;
+}
+
+##else
+##define UB_EXTRASIZE 0
+static inline size_t slab_mgmt_size_noalign(int flags, size_t nr_objs)
+{
+ return slab_mgmt_size_raw(nr_objs);
+}
+##endif
+
+static inline size_t slab_mgmt_size(int flags, size_t nr_objs, size_t align)
+{
+ - return ALIGN(sizeof(struct slab)+nr_objs*sizeof(kmem_bufctl_t), align);
+ return ALIGN(slab_mgmt_size_noalign(flags, nr_objs), align);
+}

/*
@@ -848,20 +876,21 @@ static void cache_estimate(unsigned long
 * into account.
 */
nr_objs = (slab_size - sizeof(struct slab)) /
- (buffer_size + sizeof(kmem_bufctl_t));
+ (buffer_size + sizeof(kmem_bufctl_t) +
+ (flags & SLAB_UBC ? UB_EXTRASIZE : 0));

/*
 * This calculated number will be either the right
 * amount, or one greater than what we want.
 */
- if (slab_mgmt_size(nr_objs, align) + nr_objs*buffer_size
+ if (slab_mgmt_size(flags, nr_objs, align) + nr_objs*buffer_size
    > slab_size)
    nr_objs--;

if (nr_objs > SLAB_LIMIT)
    nr_objs = SLAB_LIMIT;

- mgmt_size = slab_mgmt_size(nr_objs, align);
+ mgmt_size = slab_mgmt_size(flags, nr_objs, align);
}
*num = nr_objs;
*left_over = slab_size - nr_objs*buffer_size - mgmt_size;
@@ -1420,7 +1449,8 @@ void __init kmem_cache_init(void)
    sizes[INDEX_AC].cs_cachep = kmem_cache_create(names[INDEX_AC].name,
        sizes[INDEX_AC].cs_size,

```

```

    ARCH_KMALLOC_MINALIGN,
-   ARCH_KMALLOC_FLAGS|SLAB_PANIC,
+   ARCH_KMALLOC_FLAGS | SLAB_UBC |
+   SLAB_UBC_NOCHARGE | SLAB_PANIC,
    NULL, NULL);

    if (INDEX_AC != INDEX_L3) {
@@ -1428,7 +1458,8 @@ void __init kmem_cache_init(void)
    kmem_cache_create(names[INDEX_L3].name,
        sizes[INDEX_L3].cs_size,
        ARCH_KMALLOC_MINALIGN,
-   ARCH_KMALLOC_FLAGS|SLAB_PANIC,
+   ARCH_KMALLOC_FLAGS | SLAB_UBC |
+   SLAB_UBC_NOCHARGE | SLAB_PANIC,
        NULL, NULL);
    }

@@ -1446,7 +1477,8 @@ void __init kmem_cache_init(void)
    sizes->cs_cachep = kmem_cache_create(names->name,
        sizes->cs_size,
        ARCH_KMALLOC_MINALIGN,
-   ARCH_KMALLOC_FLAGS|SLAB_PANIC,
+   ARCH_KMALLOC_FLAGS | SLAB_UBC |
+   SLAB_UBC_NOCHARGE | SLAB_PANIC,
        NULL, NULL);
    }

@@ -1943,7 +1975,8 @@ static size_t calculate_slab_order(struc
    * looping condition in cache_grow().
    */
    offslab_limit = size - sizeof(struct slab);
-   offslab_limit /= sizeof(kmem_bufctl_t);
+   offslab_limit /= (sizeof(kmem_bufctl_t) +
+   (flags & SLAB_UBC ? UB_EXTRASIZE : 0));

    if (num > offslab_limit)
        break;
@@ -2251,8 +2284,8 @@ kmem_cache_create (const char *name, siz
    cachep = NULL;
    goto oops;
}
- slab_size = ALIGN(cachep->num * sizeof(kmem_bufctl_t)
-   + sizeof(struct slab), align);
+
+ slab_size = slab_mgmt_size(flags, cachep->num, align);

/*
    * If the slab has been placed off-slab, and we have enough space then

```

```
@@ -2263,11 +2296,9 @@ kmem_cache_create(const char *name, siz
    left_over -= slab_size;
}

- if (flags & CFLGS_OFF_SLAB) {
+ if (flags & CFLGS_OFF_SLAB)
    /* really off slab. No need for manual alignment */
-   slab_size =
-       cachep->num * sizeof(kmem_bufctl_t) + sizeof(struct slab);
- }
+   slab_size = slab_mgmt_size_noalign(flags, cachep->num);

    cachep->colour_off = cache_line_size();
    /* Offset must be a multiple of the alignment. */
@@@ -2513,6 +2544,30 @@@ int kmem_cache_destroy(struct kmem_cache
}
EXPORT_SYMBOL(kmem_cache_destroy);

+static inline kmem_bufctl_t *slab_bufctl(struct slab *slabp)
+{
+ return (kmem_bufctl_t *) (slabp + 1);
+}
+
+ #ifdef CONFIG_USER_RESOURCE
+ static inline struct user_beancounter **slab_ub_ptrs(kmem_cache_t *cachep,
+ struct slab *slabp)
+ {
+ return (struct user_beancounter **) ALIGN((unsigned long)
+ (slab_bufctl(slabp) + cachep->num), UB_EXTRASIZE);
+ }
+
+ struct user_beancounter **kmem_cache_ulp(kmem_cache_t *cachep, void *objp)
+ {
+ struct slab *slabp;
+ struct user_beancounter **ubs;
+
+ slabp = virt_to_slab(objp);
+ ubs = slab_ub_ptrs(cachep, slabp);
+ return ubs + obj_to_index(cachep, slabp, objp);
+ }
+ #endif
+
+ /*
+  * Get the memory for a slab management obj.
+  * For a slab cache when the slab descriptor is off-slab, slab descriptors
@@@@ -2533,7 +2588,8 @@ static struct slab *alloc_slabmgmt(struc
if (OFF_SLAB(cachep)) {
    /* Slab management obj is off-slab. */
```

```

    slabp = kmem_cache_alloc_node(cachep->slabp_cache,
-        local_flags, nodeid);
+        local_flags & (~__GFP_UBC),
+        nodeid);
    if (!slabp)
        return NULL;
    } else {
@@ -2544,14 +2600,14 @@ static struct slab *alloc_slabmgmt(struct
    slabp->colour_off = colour_off;
    slabp->s_mem = objp + colour_off;
    slabp->nodeid = nodeid;
+ifdef CONFIG_USER_RESOURCE
+ if (cachep->flags & SLAB_UBC)
+ memset(slab_ub_ptrs(cachep, slabp), 0,
+   cachep->num * UB_EXTRASIZE);
+endif
    return slabp;
    }

-static inline kmem_bufctl_t *slab_bufctl(struct slab *slabp)
-{
- return (kmem_bufctl_t *) (slabp + 1);
-}
-
static void cache_init_objs(struct kmem_cache *cachep,
    struct slab *slabp, unsigned long ctor_flags)
{
@@ -2729,7 +2785,7 @@ static int cache_grow(struct kmem_cache
    * Get mem for the objs. Attempt to allocate a physical page from
    * 'nodeid'.
    */
- objp = kmem_getpages(cachep, flags, nodeid);
+ objp = kmem_getpages(cachep, flags & (~__GFP_UBC), nodeid);
    if (!objp)
        goto failed;

@@ -3077,6 +3133,19 @@ static inline void *____cache_alloc(stru
    return objp;
    }

+static inline int ub_should_charge(kmem_cache_t *cachep, gfp_t flags)
+{
+ifdef CONFIG_USER_RESOURCE
+ if (!(cachep->flags & SLAB_UBC))
+   return 0;
+ if (flags & __GFP_UBC)
+   return 1;
+ if (!(cachep->flags & SLAB_UBC_NOCHARGE))

```

```

+ return 1;
+ #endif
+ return 0;
+ }
+
+ static __always_inline void *__cache_alloc(struct kmem_cache *cachep,
+      gfp_t flags, void *caller)
+ {
@@ -3090,6 +3159,12 @@ static __always_inline void *__cache_all
+ local_irq_restore(save_flags);
+ objp = cache_alloc_debugcheck_after(cachep, flags, objp,
+      caller);
+
+ if (objp && ub_should_charge(cachep, flags))
+ if (ub_slab_charge(cachep, objp, flags)) {
+ kmem_cache_free(cachep, objp);
+ objp = NULL;
+ }
+ prefetchw(objp);
+ return objp;
+ }
@@ -3287,6 +3362,8 @@ static inline void __cache_free(struct k
+ struct array_cache *ac = cpu_cache_get(cachep);
+
+ check_irq_off();
+ if (cachep->flags & SLAB_UBC)
+ ub_slab_uncharge(cachep, objp);
+ objp = cache_free_debugcheck(cachep, objp, __builtin_return_address(0));
+
+ if (cache_free_alien(cachep, objp))
--- ./mm/vmalloc.c.kmemcore 2006-08-16 19:10:38.000000000 +0400
+++ ./mm/vmalloc.c 2006-08-16 19:10:51.000000000 +0400
@@ -520,6 +520,12 @@ void *vmalloc(unsigned long size)
+ }
+ EXPORT_SYMBOL(vmalloc);
+
+ void *vmalloc_ub(unsigned long size)
+ {
+ return __vmalloc(size, GFP_KERNEL_UBC | __GFP_HIGHMEM, PAGE_KERNEL);
+ }
+ EXPORT_SYMBOL(vmalloc_ub);
+
+ /**
+  * vmalloc_user - allocate virtually contiguous memory which has
+  *    been zeroed so it can be mapped to userspace without

```
