

# Quick Guide to the Kernel Source#

The kernel and the [process manager](#) make up the operating system. They share the same address space. The kernel is non-threaded, and forces mutual-exclusion so that there is only ever one instance of it running, (even on SMP (small lie, see the SMP section)). The kernel "is entered" when anyone invokes a kernel call (or system call). The kernel runs in a privileged supervisor (or "ring0" in some architectures) state. The kernel only very rarely inhibits interrupts (and has an extensive and clever way of dancing out of the way of them). The most notable invoker of the kernel is the process manager or "[proc](#)".

## Where is the Kernel Source?#

[/services/system/ker](#)

## How do I checkout and build the source?#

Actually, one always builds the kernel and the process manager together:

- [checkout](#)
- [build](#)

## Where are the good bits?#

Functional Area	Related Files
kernel public headers	<a href="#">services/system/public/sys</a>
kernel internal headers	<a href="#">services/system/public/kernel</a> <a href="#">kmacros.h</a> <a href="#">kerproto.h</a> <a href="#">externs.h</a>
struct defs for thread, process, channel , interrupt ...	<a href="#">services/system/public/kernel/objects.h</a>
low-level kernel data structures	<a href="#">nano_object.c</a> <a href="#">nano_lookup.c</a> <a href="#">kerext_query.c</a> <a href="#">nano_query.c</a> <a href="#">nano_vector.c</a> <a href="#">nano_alloc.c</a>
debug support, reading kernel internal data	<a href="#">kprintf.c</a> <a href="#">cpu_debug.c</a> <a href="#">kerext_debug.c</a> <a href="#">nano_debug.c</a> <a href="#">nano_kerdebug.c</a> <a href="#">deb_rec.c</a>
floating point	<a href="#">nano_fp_emu.c</a>
kernel api, linkage kernel call entry, locking and exit(4)	<a href="#">kerext_*</a> <a href="#">ker_*</a> <a href="#">kerargs.h</a> <a href="#">ker_call_table.c</a> <a href="#">ker_ring0.c</a> <a href="#">nano_specret.c</a> <a href="#">kerext_misc.c</a> <a href="#">nano_misc.c</a> <a href="#">asmoff.c</a>

	<a href="#">module_list.S</a>
kernel entrypoint, init, config	<a href="#">_main.c</a> <a href="#">init_nto.c</a> <a href="#">kgetopt.c</a> <a href="#">kmain.c</a> <a href="#">idle.c</a> <a href="#">init_objects.c</a> <a href="#">nano_conf.c</a>
instrumentation, tracing	<a href="#">kertrace.h</a> <a href="#">ker_trace.c</a> <a href="#">kerext_trace.c</a> <a href="#">nano_trace.c</a>
interrupts, events, signals	<a href="#">ker_interrupt.c</a> <a href="#">nano_interrupt.c</a> <a href="#">nano_clock.c</a> <a href="#">nano_smp_interrupt.c</a> <a href="#">nano_event.c</a> <a href="#">ker_signal.c</a> <a href="#">nano_signal.c</a>
magic	<a href="#">arm/kernel.S</a> <a href="#">mips/kernel.S</a> <a href="#">ppc/kernel.s</a> <a href="#">sh/kernel.s</a> <a href="#">x86/kernel.S</a>
messaging, send, receive, reply	<a href="#">nano_message.c</a> <a href="#">nano_xfer.c</a> <a href="#">nano_xfer_msg.c</a> <a href="#">nano_xfer_check.c</a> <a href="#">nano_xfer_cpy.c</a> <a href="#">ker_channel.c</a> <a href="#">ker_connect.c</a> <a href="#">ker_fastmsg.c</a> <a href="#">ker_message.c</a>
messaging, asynchronous	<a href="#">nano_asyncmsg.c</a>
messaging, networked	<a href="#">ker_net.c</a>
messaging, pulses	<a href="#">nano_pulse.c</a> <a href="#">nano_xfer_pulse.c</a>
MMU stuff cache control, address spaces, page faulting	<a href="#">kerext_cache.c</a> <a href="#">walk_asinfo.c</a> <a href="#">kerext_bind.c</a> <a href="#">kerext_page.c</a> <a href="#">kerext_stack.c</a>
memory manager stubs(2)	<a href="#">smm_*</a> <a href="#">nano_memphys.c</a>
embedded memory manager(1)	<a href="#">emm_init_mem.c</a> <a href="#">emm_mmap.c</a> <a href="#">emm_munmap.c</a> <a href="#">emm_pmem_add.c</a> <a href="#">emm_vaddr_to_memobj.c</a> <a href="#">nano_memphys.c</a>
Minidriver support(6)	<a href="#">mdriver.c</a>
mutex support	<a href="#">ker_sync.c</a>
thread synchronization	<a href="#">nano_sync.c</a>

power management stub	<a href="#">kerext_cpumode.c</a>
process creation/destruction	<a href="#">kerext_process.c</a>
security (uname,pw), rlimits	<a href="#">kerext_cred.c</a> <a href="#">nano_cred.c</a> <a href="#">kerext_limits.c</a>
scheduler	<a href="#">ker_sched.c</a> <a href="#">nano_sched.c</a> <a href="#">nano_clock.c</a> <a href="#">kexterns.c</a>
Adaptive Partitioning Scheduler	see the <a href="#">source guide</a> for aps
shutdown, reboot	<a href="#">kerext_reboot.c</a> <a href="#">shutdown_nto.c</a>
Symmetric Multiprocessing support(3)	<a href="#">nano_smp_interrupt.c</a> <a href="#">smp_flush_tlb.c</a> <a href="#">smp_get_cpunum.c</a> <a href="#">smp_send_ipi.c</a> <a href="#">services/system/smpswitch.h</a> <a href="#">ker_nop.c</a>
system page firmware/bios interface, key kernel globals	<a href="#">ker_sys.c</a> <a href="#">nano_syspage.c</a>
timers, clock	<a href="#">nano_clock.c</a> <a href="#">ker_timer.c</a> <a href="#">nano_timer.c</a> <a href="#">ker_clock.c</a>
thread creation/deletion	<a href="#">nano_thread.c</a> <a href="#">ker_thread.c</a>

## Footnotes

(1) The embedded memory manager is only used when the OS is built as a stand-alone executive, i.e. without the process manager. QNX never ships it that way. So the `emm_*` files are generally not used. The nominal [memory manager](#) is in the process manager.

(2) The [memory manager](#) implementation is located in `proc`. It overrides these stubs in the kernel.

(3) Actually there is SMP support scattered throughout the kernel. Look for `#if defined(VARIANT_smp)`. Kernel modules, which aren't allowed to have compile variants, will test **NUM\_PROCESSORS** or **runmask** instead.

(4) How we enter and exit the kernel is described [here](#).

(5) There is no footnote 5.

(6) Fixme: wth are minidrivers

## Where the heck is the implementation of `KernelCall()`? #

You won't find it by grepping. To find the implementation of a kernel call spelled like `CamelCaseKernelCall`, look for a function `ker_camel_case_kernel_call`. For example the `SchedCtl()` kernel call is implemented by `ker_sched_ctl()`. How the invocation sequence works is described [here](#).

## What are the file naming conventions?#

filename pattern	Means
ker/ker_*	kernel entry file. Checks parameters and decides when to lock the kernel
ker/nano_*	low-level implementation, in-kernel state and in ring-0 or supervisor mode
ker/arm/* ker/mips/* ker/ppc* ker/sh/* ker/x86/*	processor specific code for ARM, PowerPC, Hitach SH-4, and Intel x86 processor architectures