Get the inside track on 27 years of microkernel innovation#

Sebastien and Colin talked about the history of the QNX Microkernel, the new Hybrid Development Model and then got into some details of how the kernel and the process manager actually work.

Archived Web Broadcast#

The On-Demand version of the broadcast is available here - http://seminar2.techonline.com/s/qnx_oct1707

Slides#

Here are the slides from the webinar - sorry the are in PowerPoint format. Oct27_Microkernel_Innovation/Webinar_kernel_oct07_final.ppt

Questions From The Webinar#

There were loads of questions during the webinar!

Are there are QNX Kernel development books in process of writing or available now?#

Not that I'm aware of - CB

Does the Momentics version contain the develoment system and also the OS for the BSP.#

What tools one needs to try this out?#

is the uK student version FUll or minimal ?#

There is only one version - the student and non-commercial aspect is simply the license agreement you select when you download - CB

Has or will QNX publish a set of development standards or procedures (e.g., coding conventions)?#

Yes - see the developers info page (OSDeveloperInformation) page for our coding guidelines - CB

how would the "QNX Comunity" work with the hybrid SW model? how do people out-side of your company contribute?#

The OSDeveloperInformation page covers how to contribute - CB

could hybrid source model adversely effect stability of customer product ? i.e. enforced product releases for bug fixes....#

Our regular releases are still going to be as thoroughly tested as before, there should be no stability problems introduced by sharing our code. - CB

Now that the kernel source is available to the world, does this make it more vulnerable to malicious attacks?#

Is there a overview or roadmap for the Momentics system profiler capabilities?#

See the IDE Project for early releases of the next IDE - CB

I am having trouble converting a Makefile for a Linux/POSIX application to a QNX IDE project. I have figured out how to get the files into a QNX C/C++ project but build doesn't do anything. Where can I find help for this?#

is release of io-pkt aimed at replacing use of io-net ?#

is it targeted to provide managed code platforms on uK like JAVA ?#

"uK doesn't run on its own" more explanation plz#

In one sense, you can consider the microkernel to be like a library of code. Just as a subroutine library doesn't do anything on its own, neither will the microkernel. There are only three things that will cause the microkernel to begin executing: a hardware interrupt, a CPU exception, or a user thread makes a system call. - BJS

Does the microkernel do any dynamic memory allocation?#

Yes it does. For the most part, this done by object_alloc() calls in the source. - BJS

kernel call operations squence more details plz#

Since only 1 process can be in kernel mode at a given time, what happens if a process executing in kernel mode is preempted by a higher priority process, and then the latter makes a kernel call?#

The kernel has a 'restart' model for kernel call preemption. At the point that the preemption occured, the context state of the thread making the call was rearranged so that the next time it runs, it will re-execute the kernel call from the begining (this is done by essentially backing up the saved PC to point at the system call instruction in libc). The new, higher priority thread now exits from the kernel code and begins running. If it makes a kernel call, no problem - there's nobody in there. At a later point in time, the original thread will become the highest priority READY thread again and begin running again. The first thing it will do is re-execute the system call instruction, which will restart the kernel call that was being worked on when it was so rudely interrupted in the first place. - BJS

How does the kernel manage the lot many number of context switches which is an overhead in a determenistic manner?#

what kernel primitives are used to ensure locking and sync with the system?#

Outside of the kernel itself, the same posix mutexes/semaphores/condition variables/pipes/message queues or QNX message passing can be used to provide synchronization between multiple execution contexts. For operations inside the kernel, because we only allow one thread into it at at time on an SMP system, not too much interior locking is required. In places where it is, InterruptLock/Unlock is used. - BJS

How does the kernel manage the overhead of copying data on message passing in a seamlessly buredenless manner ?#

Does QNX support the concept of a "hypervisor" in a multi-core SoC?#

Is there an advanced memory protection mechanism? How about performenca metrics, because it is a real time system?#

I have a new enhanced binary search algorithm#

does the instrumented kernal, effect latency ?#

what is the solution of the priortiy inversion problem?#

In the book "Getting Started with QNX Neutrino 2", it states on page 171 that priority inheritance is done only one level deep (i.e. no propagation of priorities). Is this still the case? If so, are there any plans to improve that?#

I believe this is no longer true - the priority inherited is the current priority, not the non boosted priority. - CB

do you have any kind of partitioning to virtually split one or multiple physical CPU into virtual one.?#

how about context switch time? is it fast enough? can you compare with integrity, vxworks, etc?#

How difficult would it be to greatly expand the number of priority levels that QNX can handle (like 2^31 or 2^63 or thereabouts)? Would it be totally infeasible?#

The current priority value is defined by an unsigned char (uint8_t), thus restricting the priority range from 0-255. The current max priority is defined by NUM_PRI (in public/kernel/macros.h). In order to go higher than this, you would have to make changes to the DISPATCH_* macros in ker/kermacros.h and you would also have to change the priority field(s) type in struct thread_entry, not to mention struct sched_param, procfs_info etc etc (which would break binary compat). So not infeasible, just very inconvenient) - CB

Can you please brief about the new QNX multimedia architecture? Is it possible to get some documentation on this?#

which HW platforms targeted by Trinity 2?#

How are procmgr thread priorities defined?#

Process Manager threads listen on a channel for incoming messages. When they receive a message they 'float' to the same priority as the client thread. As such they have no fixed priority.

In slide 40 it has been mentioned the user address space might be adopted. What is the use of this and does'nt this affect the memory protection?#

The process manager is a 'trusted' process, and runs in the same privileged address space as the kernel. All users applications run in a non-privileged virtual address range. Since they don't overlap, the process manager can take on the client's address space (see ProcessBind in ker/kerext_bind.c) - CB

Does the Process Manager as described here correspond to Proc32 in QNX 4.x? (Sorry, vintage user :)#

Yes, it does - CB

Are there any plans to port QNX to the ADI Blackfin? (The ADI Blackfin has memory protection hardware, but does not implement virtual addressing.)#

Not at this time. Right now we don't support any processors that don't have virtual addressing. - CB

You've mentioned plans to support the IBM 970 processor. Any thoughts on supporting IBM's cell or Power6 processors.#

Not at this time. - BJS

is BSP planned for MPC8313-RDB evaluation kit ?#

What about 64 bits OS support?#

We've been thinking about it for a while, and tried to design the API's to be prepared for it, but we just haven't had the time/resources to actually implement it. - BJS

From an OS security architectural perspective, the "Micro-Kernel" and its ancillary components support the concept of a "reference-monitor?"#

What does QNX stand for?#

Originally, QNX was named QUNIX - for Quick Unix I think. At some point it was shorted to QNX, possible because of a trademark issue. - CB