WIND RIVER

Wind River Hypervisor

Wind River Hypervisor is an embedded hypervisor that provides a virtualization layer that partitions a single or multi-core chip into multiple partitions with varying levels of protection and capabilities. The effective adoption and optimization of virtualization and multi-core will be key differentiating factors in the competitive marketplace for next-generation devices. Wind River provides a comprehensive solution with the breadth and flexibility to provide a future-proof path on the way to multi-core and virtualization adoption.

Virtualization is used in the information technology industry to deliver higher CPU utilization rates by combining many individual systems on a single piece of hardware. This delivers significant benefit from a management and cost perspective, especially with current multi-core processors.

Virtualization is making its way into embedded systems as well, on the back of ever more powerful single and multi-core processors. Virtualization in embedded opens up new opportunities for companies building next-generation products using single and multi-core devices. Device makers can utilize virtualization to consolidate their systems by replacing multiple boards or CPUs with a single board or a single CPU. They can use multiple operating systems (such as a real-time operating system and a general purpose operating system) cooperatively to provide innovative device functionality and adopt multi-core processors with improved scalability and reliability.

Adopting and optimizing virtualization and multi-core in the embedded industry requires a wide array of technologies and skills. It requires the virtualization technology itself (the hypervisor), operating systems, support for varying processor architectures and boards, debugging and analysis tools, and test capabilities. These technologies need to be integrated, easy to use, and supported by a group of embedded experts ready to assist customers. All of these elements together are critical for success on the path to a new product that is using multi-core and virtualization.

Wind River Hypervisor is deterministic, event-driven, small, and scalable and provides direct access to devices and is processor-architecture- and OS-agnostic. It is a type 1 hypervisor that runs directly on the hardware, with a small memory footprint. It is custom developed with the demands of real-time systems in mind.

Table of Contents

Multi-core Software Configurations

Configuring an embedded system in the "old days" was simple. A processor had a single core and for that core you would choose an operating system. Depending on your needs, you would choose a general purpose or a real-time operating system. If both were needed, the design would have to accommodate the two processors.

Today's powerful single and multi-core processors can be used in many different configurations.

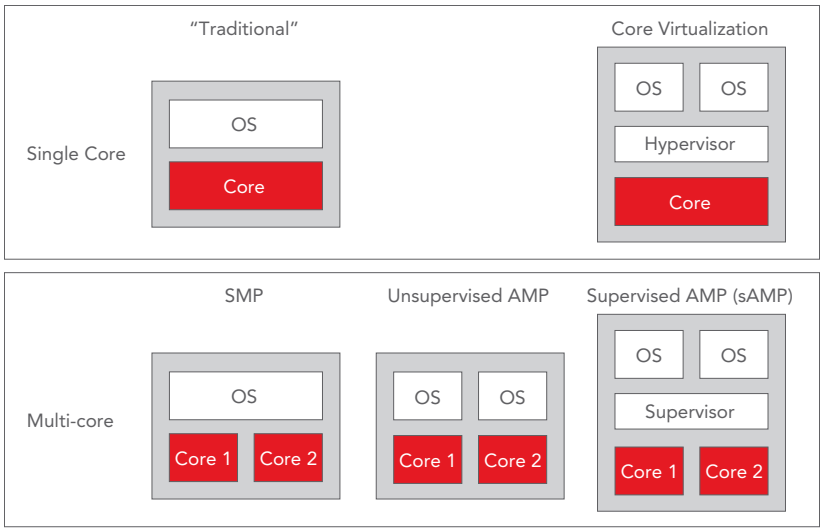

Figure 1: Primary multi-core configurations

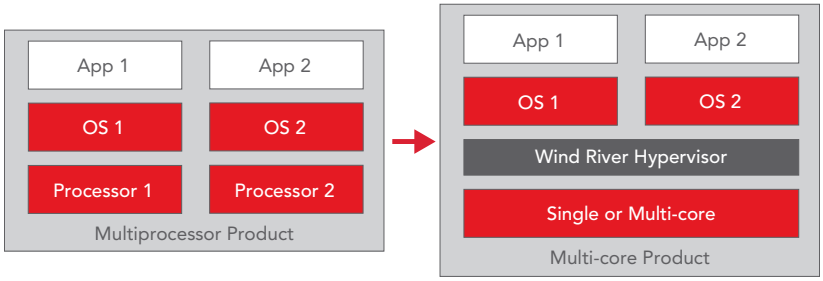

Figure 2: Consolidation of multiple single cores to a single multi-core

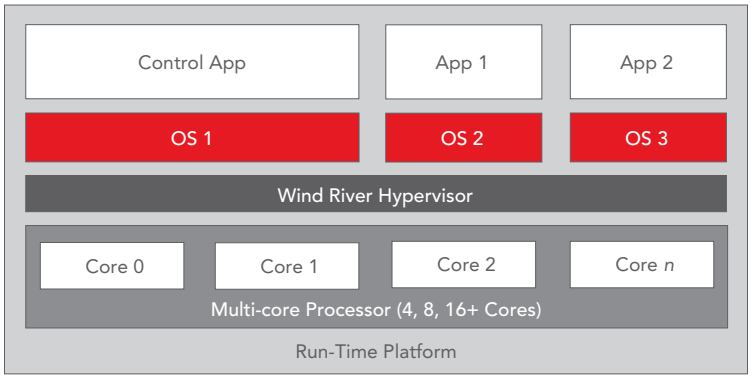

Figure 3: Scalability to high core counts

A multi-core processor can be managed by a single symmetric multiprocessing (SMP) operating system that manages all the cores. Alternatively each core can be given to a separate operating system in an asymmetric multiprocessing configuration (AMP). Both SMP and AMP have their challenges and advantages. SMP does not always scale well, depending on workload; and AMP can be difficult to configure in regard to which operating system gets access to which device.

The addition of virtualization to the system in the form of a hypervisor enables a wide variety of configurations including mixes of AMP, SMP, and core virtualization. The hypervisor manages the hardware and creates partitions to execute operating systems within. Each partition is given access to resources (processing cores, memory, devices) as specified by the development team. Each partition can hold an operating system (also known as the guest OS) and is protected from the other partitions. The hypervisor can execute a single partition on a single core, a single partition across multiple cores, or multiple partitions on a single core.

The combination of SMP, AMP, and core virtualization provides unprecedented opportunities for device developers to innovate.

Use Cases

Embedded hypervisors provide new and exciting capabilities to developers. The following are some use cases that a hypervisor enables.

Consolidation

Many current systems use multiple processors, either on the same board or on multiple boards in a rack. Reasons to use multiple processors include performance and the need for separation between different types of functionality (real-time and non-real-time, safety, security, or GPL/non-GPL).

The use of a hypervisor provides separation on a single or multi-core processor, thereby reducing the bill of materials (cost of goods sold) and reducing the amount of power used (operating costs). Each of the operating systems with a separate processor in the old system could be migrated to a virtual board. The virtual board provides separation and fault containment. Figure 2 shows an example of a migration from a dual processor system to a dual core processor.

Effective and Scalable Adoption of Multi-core

In addition to greater raw processing power, the multi-core chips of today provide hardware acceleration features for specific tasks such as packet processing and graphics. There are many ways that these chips can be configured as offloading engines with a single operating system, a mix of different operating systems, or a combination of operating systems and lightweight executables. Configurations can provide memory separation, the capability to restart, sharing of devices, and more. Cores can be shared between multiple virtual boards, or a single virtual board can span multiple cores and run an SMP operating system.

An embedded hypervisor provides the capability to partition the multi-core device effortlessly. The partitioning can be done to provide highest performance, to leverage existing code, and to provide evolution and scalability to adapt to newer chips with more cores in the future.

Innovation and Differentiation

Virtualization enables device developers to refresh an existing product in an innovative way by introducing a second operating system type to the device, for example, introducing a Wind River Linux–based partition to an existing product based on VxWorks. This allows the product to benefit from improved capabilities on Linux such as network connectivity, graphics, or other applications in the Linux ecosystem. The real-time nature of Wind River Hypervisor allows the new product to maintain the real-time determinism while providing exciting new capabilities to deliver innovation and differentiation.

Figure 4 depicts a situation where an existing OS and application can be taken from an existing product and be augmented with a new application and operating system for a next-generation product.

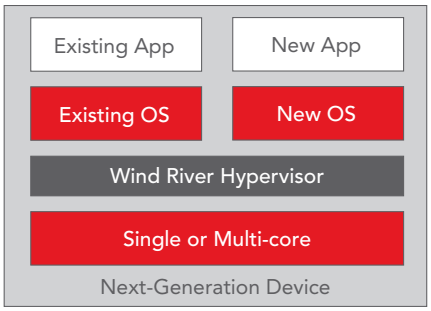

Figure 4: Innovation and differentiation

Fast Boot

Bringing the system up into a state where it can respond to external input is important, especially in many consumer and automotive devices. Wind River Hypervisor boots extremely quickly before booting the system's partitions in the order defined by the designer.

One of the partitions can run an application without an operating system, which is active immediately and can start communicating with external busses or devices, while other partitions with operating systems such as Wind River Linux can be booted in the background and provide graphical user interfaces later. This situation is partially depicted in Figure 4, but instead of a new OS and new app it would contain a partition with no operating system.

Intellectual Property Isolation

The ability to protect intellectual property is very important and becomes more difficult as systems become comprised of open source software, third-party software, and internally developed software. In some cases it is a concern related to the GNU Public License, and in other cases third-party code comes with restrictions on what it can be integrated with. Wind River Hypervisor can be used to partition the systems so that intellectual property is isolated while allowing customers to innovate with multiple operating systems and software from multiple sources.

Features

Virtual Board

Wind River Hypervisor separates a system into multiple virtual boards (partitions). The virtual board can contain a guest operating system with applications or an application without an operating system (virtual board application). The virtual board is managed by the hypervisor. The hypervisor controls which cores the virtual board executes on and which memory and devices it can access. A virtual board can share a single core with another virtual board, run dedicated on a single core, or span multiple cores. The hypervisor can also share a single core between multiple virtual boards.

The virtual board is a fault container. The hypervisor prevents faults that occur within a single virtual board from affecting other virtual boards and the rest of the system. The hypervisor can be designated as the fault handler, or the fault can be propagated to the virtual board for handling within the guest operating system.

Virtual boards can also be individually paused, resumed, and restarted, providing an easy way to recover from critical failures inside a virtual board. The hypervisor's automated life-cycle control is a convenient feature and provides for improved reliability and robustness, particularly on multi-core processors.

Guest Operating System

Wind River Hypervisor executes guest operating systems, which host applications. A guest operating system executes with near native performance. Wind River Hypervisor uses a mix of hardware assist and paravirtualization of the guest operating system to deliver optimized performance and determinism.

Wind River Hypervisor supports real-time operating systems such as Wind River Linux and VxWorks, and its open interface allows other operating systems and executives to run, including those that are open-source-based, proprietary, or internally developed.

Wind River Hypervisor allows the use of Microsoft Windows inside a virtual board when running on top of Intel processors. This enables the use of Windows alongside real-time content hosted by VxWorks or Wind River Linux.

Wind River Hypervisor also allows the execution of a virtual board without an operating system. This is useful for certain tasks that don't need the capabilities of an operating system, such as a fast polling loop. This no-operating system capability (known as a virtual board application) is one of the enablers of the fast-boot scenario.

Hardware Assist and Paravirtualization

Some processor architectures provide features in hardware that facilitate virtualization (e.g., Intel VT-x, PowerPC e500mc). The hypervisor can use these features to provide virtualization services.

On CPUs where these features do not exist, or if these features impact performance and determinism, the hypervisor implements them in software, and paravirtualization of the operating system is required.

Paravirtualization is the modification of an operating system to collaborate with a hypervisor. This collaboration works through "hypercalls," calls from the guest operating system into the hypervisor. The amount of modification depends highly on the processor architecture. It typically concerns the following:

- **Privilege levels:** A virtualized system typically requires three privilege levels: the hypervisor; the guest operating system; and the application. Many architectures without hardware assist only have two privilege levels. The task of paravirtualization is to emulate the missing privilege level. Privileged instructions in the guest operating system need to be replaced by hypercalls.
- **Device access:** A driver in a native operating system is able to access any device. In a virtualized environment, the hypervisor arbitrates whether a virtual board has access to a device. This includes device interrupts, device I/O, device register access, and direct memory access (DMA). Device access includes access to timers, network cards, graphic cards, and so forth. Environments with hardware assist often allow device drivers to operate without modification.

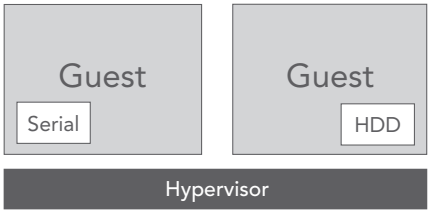

Figure 5: Device drivers in the guest, not in the hypervisor

- **MMU:** The hypervisor controls the memory management unit (MMU) in a virtualized environment. Many processors with hardware assist allow a guest to modify the virtual board MMU. If hardware assist is unavailable, the guest may have to collaborate with the hypervisor to modify the MMU.

Configuration

Wind River Hypervisor provides flexible tooling to configure virtual boards on single core and multi-core. The configuration defines how the virtual boards are spread out over the available processing cores; how they are scheduled; and how the hardware is partitioned between the virtual boards. There are a number of basic building blocks that the developer can use to assemble a system:

- **One virtual board allocated to a single dedicated core:** This provides the same performance as if the guest OS is running natively on that core.
- **Multiple virtual boards allocated to a single core:** The hypervisor provides scheduling services. The hypervisor's scheduler decides how processor cycles are provided to each virtual board. These decisions are either done on a priority basis or a partition basis (*x* time slices per virtual board). The scheduler in Wind River Hypervisor is pluggable, so, if needed, teams can design their own schedulers.
- **Multiple cores allocated to a single virtual board:** The guest operating system is executing in SMP mode across a number of cores.

System designers can decide how to use these building blocks in their designs. They can, for example, carve up an eight-core processor in a virtual board spanning multiple cores and running an SMP guest operating system. They can run several virtual boards, each using a single core, and they can run multiple virtual boards on a single core.

Architectural Design

Wind River has decades of experience in developing real-time embedded operating environments and partitioned systems, and this knowledge has been leveraged in the architecture of Wind River Hypervisor.

Deterministic

Embedded systems need to be fast and deterministic. They need to respond within known time bounds without fail. Embedded developers rely on a deterministic operating environment that needs to include the hypervisor as well as guest operating systems.

Besides the requirement for deterministic response times, Wind River Hypervisor is optimized to provide low overhead, even on multi-core processors with a large amount of cores.

Event Driven

Wind River Hypervisor works as passively as possible. It provides separation, configures resources, and allows the guest operating systems to run at full speed. Wind River Hypervisor has no active threads. It is completely event driven, meaning there is no overhead or activity unless requested by the guest operating systems. Examples of events are system calls from the guest OS or a driver, interrupt delivery, and so forth.

Minimal Footprint

The footprint of Wind River Hypervisor is minimal. This ensures that unnecessary code is not built or executed as part of the hypervisor-based system. The hypervisor is the bottom layer of the system. The more functionality added to the hypervisor, the greater the chance that a small unforeseen problem can bring down the entire system. This is the reason there are no Ethernet drivers or graphics drivers in Wind River Hypervisor. The virtualization layer is scaled down to the level where it only contains the minimal functionality required to provide its services. Any additional code, such as device drivers, is not contained inside the hypervisor but inside the guests that use it.

Scalable

Multi-core processors are making their entry into embedded systems. Top-of-the-line systems (typically in the networking space) may use up to 32 cores. This is certain to increase in the coming years. Many hypervisors and SMP operating systems will not scale up to this number of cores. But Wind River Hypervisor was designed from the ground up to be scalable, ensuring that performance is not compromised as the core count increases.

Direct Access to Devices

Embedded devices are all about device access. In some devices I/O throughput is critical; in others it is latency or timing

constraints. IT virtualization layers separate the guest from the devices they need to access. This is not so with Wind River Hypervisor. Memory, PCI attributes, and interrupts can be directly mapped into the guest. The hypervisor is not involved in the data path to or from the device. This delivers performance that is equal to native performance.

Architecture and OS Agnostic

Fragmentation of processor types and operating systems is common in the embedded industry, particularly across a product line. Many device developers use low-cost processors for small devices and higher-performing processors for their high-end products. Wind River Hypervisor, together with Wind River's operating systems, provide a single API and environment that can be used across many different devices. The hypervisor itself is completely OS-agnostic, meaning that other operating systems, either third-party or in-house, can be added as a guest.

The Larger Picture

Virtualization is one part of a comprehensive solution consisting of virtualization capabilities, operating systems, development tools, and testing solutions. Wind River has recognized the importance of multi-core and virtualization and has defined a broad and flexible product portfolio to support it.

The flexibility in the portfolio stems from Wind River providing virtualized and nonvirtualized alternatives, to effectively take advantage of the multi-core processors of today and tomorrow. Wind River's portfolio allows a device developer to start, for example, with a VxWorks SMP configuration on a dual core processor and in a later device evolution migrate to a VxWorks and Linux combined solution in an AMP configuration. The same application can also be used in a virtualized configuration, either on a single core or spread out over multiple cores.

The term "same application" is key. Wind River offers the same technologies over the breadth of its product portfolio. That is what gives designers the flexibility to choose an appropriate multi-core software configuration and know it is future-proof, as processor technology and device requirements evolve over time. The following are the important technologies in the portfolio.

Multi-OS Inter-process Communication

Communication between virtual boards is an important capability when building embedded systems. Multi-OS inter-process communication (MIPC) provides a socketlike API to send messages between virtual boards inside of a multi-core environment. Sending messages uses hardware acceleration where possible and sends using zero copy. MIPC can also be used as an inter-process communication feature to debug other virtual boards, either as an extended serial port or as a sophisticated debug connection for Wind River Workbench.

MIPC Serial Device

The MIPC serial device (MSD) allows many different guests to use a single serial port. The port can be used to share the consoles of multiple different guests, or can be used as a multiplexed printf channel.

MIPC Network Device

MIPC Network Device (MND) provides for an extremely fast TCP/IP layer between the guests executing on top of the hypervisor. All of the guests have access to the TCP/IP bus and can use standard interfaces to communicate. MND can be used, for example, to display graphics using the X11 protocol or for Network File System (NFS) or other data- or performance-intensive tasks.

Development Tools

An integrated development environment is typically the cockpit that the development team leverages to build embedded systems. Wind River Workbench provides an Eclipse-based development environment that improves developer productivity for tasks such as configuring, diagnosing, and analyzing hypervisor-based systems.

Configuration

Systems based on Wind River Hypervisor require configuration to partition system resources allocated to its virtual boards. Wind River Workbench provides editors for modifying the XML files required to configure the system, including direct links to the documentation for additional guidance.

Build and Debug

Workbench provides a single-click system build function including the virtual boards and the complete system file. Once the build completes, the system image can be automatically downloaded for debugging and analysis. Debugging in a multi-core and virtualized environment often requires the developer to keep track of many things at the same time. Workbench enables this through the multi-context debugger.

Workbench provides a single integrated environment to configure, build, debug, and analyze systems based on Wind River Linux, VxWorks, and Wind River Hypervisor, which simplifies a developer's workflow and increases developer productivity in comparison to switching back and forth between tools from different vendors.

On-Chip Debugging

Embedded development is difficult, embedded debugging even more so. Being able to have fine-grained control over an embedded target is paramount for success. Wind River Workbench On-Chip Debugging allows debugging of drivers, interrupt-routines, board bring-up, and other low-level scenarios. However, it also allows a developer to stop an entire multi-core chip. Often when debugging a single thread, it is necessary to stop other threads at the same time to avoid buffer overruns or underruns that would change the debug

scenario. Workbench On-Chip Debugging allows the developer to stop all cores on multi-core with short lag time, providing an overview of the entire context.

Integration

Wind River Hypervisor is part of the Wind River multi-core software solution, integrating with the other components: VxWorks, Wind River Linux, Wind River Workbench, and Wind River Test Management.

Wind River Hypervisor has an open and well-defined interface to the virtual boards that run on it. Any operating system can be executed inside a virtual board, with the appropriate paravirtualization modifications. Wind River Workbench provides the interface to configure, build, and debug small and large hypervisor-based systems through graphical editors, elaborate build configurations, and multi-context debuggers.

MIPC can be used to communicate between the virtual boards during run-time. MIPC provides a fast, zero-copy communication channel with a socketlike API. MIPC can also be used as an extended serial connection to virtual boards that do not have access to a serial port or can be used to connect to Wind River Workbench for more sophisticated debugging scenarios.

Professional Services

The introduction of multi-core and virtualization provides many more choices to design teams delivering next-generation devices. But there is often a lack of experience and expertise in these new technologies. Wind River Professional Services is on the leading-edge of the multi-core and virtualization revolution and can help accelerate the introduction of these technologies while reducing risk.

Wind River Professional Services, a CMMI Level 3–rated organization, enables you to focus on development activities that add value and differentiate your design. Wind River offers industry-specific services practices, with focused offerings that help you meet strict market deadlines while keeping development costs down. Our experienced team delivers device software expertise globally to solve key development challenges and directly contributes to our clients' success.

Backed by our commercial-grade project methodology, Wind River Professional Services includes the following:

- Multi-core/virtualization architectural review
- Requirements discovery and definition
- Multi-core and virtualized board support package (BSP) and driver optimization
- Software system and middle integration
- Application and infrastructure development

Typical projects range from two to four man-weeks for driver and BSP implementation; one man-month to one man-year for hardware design or extensions to an existing software solution; and multi-man-year programs that bring customer concepts to reality through design, creation, and system test and verification.

Professional Services has extensive experience with platform design, including safety-critical systems and navigation/infotainment systems. Professional Services has implemented both hardware and software solutions for the embedded device market and continues to work with standards organizations to establish the next-generation platforms.

Installation and Orientation Services

The Wind River multi-core software solution provides a host of technologies that may seem overwhelming at first. Wind River offers Installation and Orientation Services to ensure your project starts on time and without hassle by delivering the following:

- **Onsite installation:** Guided install on your hardware and host platform, along with a sample multi-core software configuration, demonstrations, and examples of customizations
- **Hands-on orientation:** Architecture, development file system, open source packages, porting drivers, design issues

Advice: Introduction to Wind River support channels and processes, additional services, project reviews, and consultation

Wind River Installation and Orientation Services will expedite your path to productivity, allow you to rest assured that we have eliminated a common source of user error, and help you maximize the benefit from multi-core and virtualization.

Education Services

Education is fundamentally connected not only to individual performance but also to the success of a project or an entire company. Lack of product knowledge can translate into longer development schedules, poor quality, and higher costs. The ability to learn—and to convert that learning into improved performance—creates extraordinary value for individuals, teams, and organizations. To help your team achieve that result, Wind River offers flexible approaches to delivering product education that best fits your time, budget, and skills development requirements.

Public Courses

Wind River's public courses are scheduled for your geographical convenience. They are conducted over one to five days, using a mixed lecture and interactive lab classroom format that leverages the experience of Wind River instructors and other course participants. Courses provide a fast, cost-effective way for students to become more productive in Wind River technology.

Benefits of public courses include the following:

- A conceptual introduction that orients students to the subject matter
- A selective examination of the details, focusing on the most commonly used areas, or on areas with which users tend to be least familiar
- Personal guidance and hands-on application of individual tools and course concepts
- The chance to grasp device software concepts, as well as the fundamental issues involved in real-time design
- The knowledge needed to develop device drivers, perform hardware porting, or develop applications

- Answers to specific questions about topics addressed in the course

Consult your local Wind River sales representative for course schedules and fees.

Onsite Education

If you have a large project team or a number of new users, you may benefit from custom onsite education. Instructors will consult with you and, based on the workshop series curriculum, determine which topics should be included and emphasized. This type of education offers an opportunity for one-on-one discussions with our instructors about your specific project needs, technical requirements, and challenges, all in the comfort of your own office.

Advantages of onsite education include the following:

- Your entire team gains a common knowledge base.
- Knowledge and skills are transferred directly from the classroom to your workplace.
- Use of your location saves employees travel expenses and time away from the office.

Consult your local Wind River sales representative for further information about onsite education.

Support Services

Wind River Customer Support, a Service Capability and Performance (SCP)–certified organization, provides support for Wind River Hypervisor.

Your subscription includes full maintenance and support delivered through Wind River's Online Support (OLS) website and our worldwide technical support team. While under subscription or support agreement, customers receive both maintenance updates and major updates.

Visit Wind River Online Support at www.windriver.com/support for fast access to product manuals, downloadable software, and other problem-solving resources. OLS offers a comprehensive knowledge base with robust search features for locating information quickly.

Additional support features, including proactive email alerts covering particular technologies, platforms, or product patches and technical tips for common problems, are available for all customers by subscription. OLS visitors can also access a community of developers to discuss their issues and experiences.

If you cannot find the information you need through Online Support, contact our global support team for access to the industry's most knowledgeable and experienced support staff.

For more details on our support processes including escalations and defect resolution, consult Wind River's Customer Support User's Guide at www.windriver.com/support/resources/csug.pdf.

Wind River Customer Support contact details are at www.windriver.com/support/.

North America, South America, Asia/Pacific

support@windriver.com
Toll-free tel.: 800-872-4977 (800-USA-4WRS)
Tel.: 510-748-4100
Fax: 510-749-2164
Hours: 6:00 a.m.–5:00 p.m. (Pacific time)

Japan

support-jp@windriver.com
Tel.: +81 3 5778 6001
Fax: +81 3 5778 6003
Hours: 9:00 a.m.–5:30 p.m. (local time)

Europe, Middle East, Africa

support-ec@windriver.com
Toll-free tel.: +800 4977 4977
France tel.: +33 1 64 86 66 66
France fax: +33 1 64 86 66 10
Germany tel.: +49 899 624 45 444
Germany fax: +49 899 624 45 999
Italy tel.: +39 011 2448 411
Italy fax: +39 011 2448 499
Middle East region tel.: +972 9741 9561
Middle East region fax: +972 9746 0867
Nordic tel.: +46 9 594 611 20
Nordic fax: +46 8 594 611 49
UK tel.: +44 1793 831 393
UK fax: +44 1793 831 808
Hours: 9:00 a.m.–6:00 p.m. (local time)

WIND RIVER

Wind River is a world leader in embedded and mobile software. We enable companies to develop, run, and manage device software faster, better, at lower cost, and more reliably. www.windriver.com