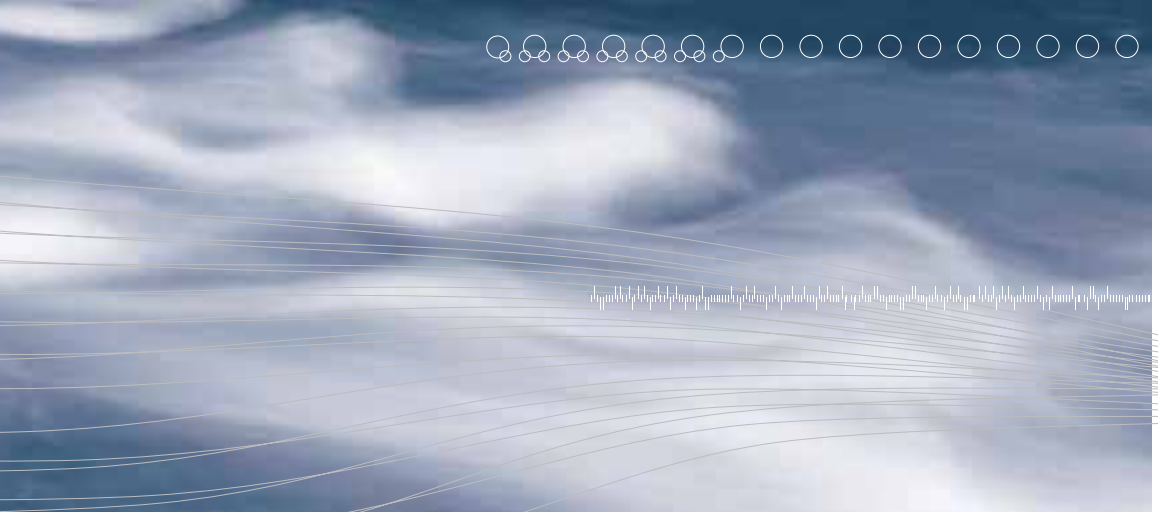

Wind River | *White paper*

An Introduction to Network Security in Embedded Devices

As the sophistication of the Internet has increased, issues of transaction security, user authentication, and authorization have taken on greater and greater importance. Companies are turning to a *network security model* in order to control access to internal networks comprised of many hosts, as well as the services they offer. Embedded devices are becoming increasingly important network components, and users need embedded networking protocols composed of comprehensive, industry-standard security solutions. This paper describes common security-related technologies and networking protocols and their place in lowering potential security risks.

WindRiver®
HOW SMART THINGS THINK™

Executive summary

The Internet began as a rather simple medium for transmitting data between widely dispersed nondesktop computers and their associated "dumb" terminals. Now this information medium has been transformed into a sophisticated platform for business data exchange, and issues of transaction security, user authentication, and authorization have taken on greater and greater importance.

As use of the Internet has increased, business networking environments have become larger and more diverse, and securing them on a host-by-host basis has grown more difficult. As a result, companies are turning to a *network security model.* This model enables companies to control access to internal networks comprised of many hosts, as well as the services they offer. The network security model involves building firewalls to protect internal systems and networks, using strong authentication approaches (such as one-time passwords and digital certificates), and employing strong encryption to protect particularly sensitive data as it transits a network.

Meanwhile, embedded devices are becoming increasingly important network components. State-of-the-art embedded devices are not only network aware, but can also be integrated and made interoperable with network equipment such as routers, switches, hubs, and servers. As a result, users need embedded networking protocols composed of comprehensive, industry-standard security solutions.

This paper introduces common security-related technologies and networking protocols and shows how they are being used and deployed in embedded devices.

Isssues around networking security

When Bob Khan and Vinton Cerf created TCP/IP in 1982, they did not design it with security in mind, but as a medium for exchanging information. Some ten years later, however, hackers were exploiting security weaknesses exposed by the public's exploding use of the worldwide Internet for personal and business applications. Security problems revolved around three issues that the Internet protocol (IP), in itself, does not address:

- *Authentication:* A network user or message must have an identifiable origin.
- *Integrity:* A message or data packet's contents must be verifiably genuine.
- *Confidentiality:* A message or data packet's contents must be guaranteed to remain private.

Additional examples of security issues include scenarios of attackers hijacking already authorized connections, assuming the identity of a previously authenticated user, and continuing communications without other users noticing the change. Another well-known weakness is the so-called *man-in-the-middle attack,* where two rightful communications peers don't realize they are actually communicating through a third party who is impersonating both sides of the connection. Similar to this is a *replay attack,* which maliciously or fraudulently repeats a valid data transmission. This can be done either by the data originator or by an adversary who intercepts the data and retransmits it by posing illegitimately as (that is, assuming the identity of) the originating entity.

Outside of actual data traffic, security is required to control access to network resources. Even though a user might be authenticated by a network, he still might not be authorized to access all of the network's resources. It is here that additional security measures are needed in order to control access to particular parts of a network on a per-user or time-of-day/day-of-week basis.

The case for networking security: B2B and B2C

Expanded channels

Security is not just for insurance purposes anymore. To increase revenues and profitability, companies need security solutions to expand trusted relationships with their customers, partners, suppliers, and sales channels. The ability to use security technologies to enable greater access to corporate content deepens and stabilizes these relationships.

Remote access

The connections mobile employees and telecommuters use to access corporate network resources over the Internet are highly vulnerable to network attacks. These connections and associated access rights usually reach far deeper into a corporation's intranet than the ones business partners or customers have.

During the current year, it is expected that almost 60 percent of the workforce will be mobile, and most of them will require secure remote network access (Forrester Research, July 2000). As a result, companies continue to increase their IT budgets for infrastructure security.

Internal threats to the enterprise

Perhaps half of all the damage caused to enterprise information systems comes from authorized personnel who are either untrained or incompetent. Another quarter of the damage seems to come from physical factors such as fire, water, and bad power. Maybe a fifth of the damage comes from dishonest and disgruntled employees. Computer viruses cause another few percent, and perhaps 5 or 10 percent of the damage is caused by external attacks (Web-site defacements, Web-site vandalism, network intrusion, information theft, information destruction, and so forth).

Private Internet user

The average Internet user's awareness of common security threats is growing at a fast pace. Press coverage about internationally spreading viruses, worms, or distributed denial of services attacks is a regular occurrence. People know that an always-on connection such as a DSL or cable modem, combined with accompanying static IP addresses, exposes their PCs directly to the Internet world. What's more, they expect their ISP's equipment and service to protect them.

Other threats to private networking are related to online personal financial transactions such as banking, investing, or shopping. Cases of credit card number thefts were probably among the first incidents confronting private network users with serious security threats from the public Internet.

Interestingly, these days it is not the financial transaction (in which an individual sends a credit card number over the Internet) that is commonly attacked, but rather the servers storing this sensitive data. That's because average network traffic volume has grown so large and encryption techniques so strong that a standard hacker with standard hardware can hardly "sniff" on all potential traffic. Even with just 40-bit encryption, a hacker will have difficulties coping with all the packets. It is therefore more efficient to attack the servers. This is why technology advances today focus on firewalls, intrusion detection systems, and access control.

Modern society

Jim Wolf, in the article "U.S. Draws Attention to Information Warfare Threat," notes that"U.S. officials repeatedly warn about perceived dangers to a United States increasingly stitched together by bits and bytes of computer code. A key stated fear is information warfare, or sneak electronic assaults that could crash power grids, financial networks, transportation systems, and telecommunications, among other vital services. National security experts trace the threat to hostile or potentially hostile governments, as well as drug lords, criminal cartels, and increasingly computer-savvy guerrilla groups. Some of these organizations "are doing reconnaissance today on our networks, mapping them, looking for vulnerabilities," according to Richard Clarke, President Clinton's top aide for infrastructure protection and counterterrorism."

He continues, "Cyberblitzes like those that briefly knocked out major Web sites in February 2000 – including Yahoo! Inc.'s Internet gateway, eBay Inc.'s auction service, and Amazon.com, Inc.'s retail site – could easily be copied on a larger scale, said Clarke, at the time a staff member of the White House National Security Council. Such warnings from Clarke are not new. He has frequently conjured up a "digital Pearl Harbor," a reference to the Japanese surprise attack that threw the United States into the Second World War."[1]

Overview of available networking security technologies

Authentication, authorization, and accounting

Probably the best-known and most widely used mechanism for user authentication is the one that combines username and password. Other authentication mechanisms that build upon this concept include one-time passwords, tokens for one-time password generation on the fly, smartcards, and digital certificates.

Data integrity

In order to verify data origin and integrity, hashing algorithms like SHA-1, MD5, and RSA are used to generate a keyed checksum of a message. The receiving party runs the same algorithm with the same (secret) key and can compare the two checksums.

Data confidentiality

Encryption algorithms provide data confidentiality. Algorithms like DES, 3DES, RSA, RC-4 and the advanced encryption standard (AES – the RIJNDAEL algorithm) can scramble messages to make them unreadable by third parties.

Encryption and keys

There are two types of encryption keys: *symmetric* and *asymmetric* key pairs. With symmetric encryption, one and the same key is used for encryption and decryption. However, the same key must exist on both sides of the connection and must be kept secret. With asymmetric encryption, a key pair composed of a private and a public key is used – one key for encryption, the other for decryption. The key pair is associated with a particular user who keeps the private key secret and makes the public key available to everyone who wants to communicate with that user.

Asymmetric encryption can be used in several ways, including authentication, data integrity checks, or confidentiality. A message encrypted with the user's public key can only be decrypted with the user's secret private key. The message is therefore unreadable by other parties that only have the user's public key. A received message that is encrypted with the user's private key can only be decrypted with his public key. In this way, the message's origin can be verified. If the public key cannot successfully decrypt a message, then this message was not encrypted by the user's private key, but by some other key.

In some cases symmetric encryption performs up to 1000 times faster than asymmetric encryption. This is why asymmetric encryption is often used only for authentication, integrity checking, and symmetric key exchange. Actual bulk data encryption is preferably performed by the much faster symmetric encryption.

[1] *Reuters,* Dec. 26, 2000

Key generation and exchange

In order to use hashing or encryption algorithms, the respective secret keys need to be generated and safely transmitted over what are often unsecured network connections. Algorithms providing this kind of capability include Diffie-Hellman, RSA, and El Gamal. Digital certificates allow a verifiable distribution of users' public keys.

Hardware encryption and security processors

Cryptography is a highly CPU-intensive process. Standard CPUs often cannot handle the required number of simultaneous encrypted connections. The popular DES algorithm, for example, was actually designed to be more efficiently performed in hardware. As a result, cryptographic functions are more and more often offloaded to dedicated hardware. This makes possible better network throughput and lower utilization of the main CPU. Today's security processors do not just perform encryption algorithms, but also hashing and key generation. Latest generations even perform parts of IP header processing.

Digital certificates

Digital certificates bind a user ID to its respective public key. Only the owner of the certificate knows the corresponding private key. The digital signature component of a certificate is an electronic identity card. The digital signature tells the recipient that the information actually came from that particular sender and has not been forged or tampered with.

Digital certificates are used within a so-called public key infrastructure (PKI). Central points in a PKI are trusted third parties – so-called certificate authorities (CAs) that issue, sign, and distribute users' digital certificates. Large corporations that offer CA services include Verisign, Entrust, Equifax, Thawte, GTE, Microsoft, Netscape, and Deutsche Telekom. CAs sign certificates with their secret keys and end users verify certificate and its content with the CA's public key. CAs publish their public keys on a large scale (for example, in every copy of Web-browsers like Microsoft Internet Explorer or Netscape Navigator, all kinds of media, and various other methods of distribution), which enables every end user to obtain the genuine public key.

Digital certificates are not only issued to individuals but also, for example, to corporations and Web sites. When a Web-browser visits a secure Web site (one whose address starts with https), the site automatically sends its certificate. The current ITU standard format for digital certificates is X.509v3.

Tunneling

Tunnels are used to create a private communications path over a public network infrastructure like the Internet. Data packets (not necessarily IP based) or frames like PPP are wrapped in IP packets that disguise the data among other traffic and hide original source and destination information. Today's most popular tunneling protocol is the layer 2 tunneling protocol (L2TP). Even IPsec can operate in tunnel mode, in addition to its encryption capabilities.

With IPsec operating in tunnel mode and encrypting the entire original packet, including its header, tunneling provides yet another advantage: protection against traffic analysis attacks. All that can be seen of the traffic between two routers is encrypted traffic being sent back and forth. The attacker might possibly be able to discover the router in front of a server, but since the actual IP addresses are encrypted inside the tunnel, the attacker cannot determine the server's IP address. Identifying a server's IP address goes a long way towards determining which system is interesting enough to invest the effort in attacking.

Firewalls

In building construction, a firewall is designed to keep a fire from spreading from one part of a building to another. In theory, an Internet firewall serves a similar purpose. It prevents the dangers of the Internet from spreading to an internal network. Firewalls restrict unauthorized users from entering a carefully controlled network. An Internet firewall is most often installed at the point where a protected internal network connects to the Internet. Firewalls are built into devices that are located at the edge of a network, such as a network access server and – a new device that's proliferating rapidly – a SOHO (small office/home office) router gateway. Furthermore, there is a clear industry trend that is moving away from dedicated, standalone firewall solutions towards firewall capabilities built into gateway IP routers.

Firewalls on the edge of a network prevent attackers from getting close to other network defenses. By controlling inbound and outbound traffic they can also restrict authorized users from leaving a carefully controlled network. There are various types of firewalls differentiated from each other by their level of sophistication. Basic firewalls use packet filtering based on filtering criteria similar to that used by IPsec. Other firewalls are aware of user-created connections and allow associated inbound and outbound traffic. This kind of firewall is commonly called *stateful* packet filtering. More intelligent solutions are also application aware and filter traffic based on specifically set rules. The most complex solutions today include virus scanners, network address translation (NAT), and IPsec awareness.

Networking security protocols

The following is an overview of industry-standard networking protocols using the technologies discussed in the previous section.

Remote authentication dial-in user service (RADIUS)

RADIUS is a popular networking protocol for user authentication, authorization, and accounting. Users who want to log on to a network provide their user credentials to a RADIUS client, which then forwards them to a remote RADIUS server. The server verifies the data and notifies the client of the success of the authentication. RADIUS can also provide authorization information for access control and keep accounting data of network usage. The client/server concept removes the server and all its security-sensitive user information out of the first line of network defense, leaving only the client directly exposed to potential attacks.

Layer 2 tunneling protocol (L2TP)

The IETF derived an open standard for layer 2 tunneling of PPP from the older Microsoft point-to-point tunneling protocol (PPTP) and Cisco Systems' layer 2 forwarding (L2F). The new IETF standard, L2TP, runs over multiple media, such as ATM, frame relay, and X.25, and tunnels various protocols, including PPP, IPX, NBF, and AppleTalk. L2TP is primarily carried over UDP/IP, but the standards are such that it can be carried over any transport protocol. L2TP provides multiple tunnels between the same source and destination. What's more, the tunnels can receive different priorities for Quality of Service purposes. While L2TP authenticates tunnel endpoints, it does not ensure the integrity of packets in the tunnel. For this reason and to add privacy through encryption, it can be used in conjunction with IPsec.

IP security (IPsec)

IPsec provides full security services as an optional part of the Internet protocol. IPsec is independent of the application protocol. The application does not need to change to use IPsec, nor does it even need to be aware that IPsec is involved. IPsec is a protocol suite that makes possible tunneling, data hashing, and encryption to all IP communications on the network layer. It can provide authenticity, integrity, and confidentiality. Features include the following:

- *Authentication header (AH) and encapsulating security payload (ESP):* AH is one of the protocol components within IPsec. AH provides data integrity and authenticity, but not privacy. It allows virtual private network (VPN) entities to authenticate themselves to each other by using keyed-hash functions (HMAC-MD5 or HMAC-SHA-1). ESP, the other component of IPsec, provides modular confidentiality, authentication, and integrity by using encryption to scramble data in each packet. ESP and AH can be used independently and in combination (Figure 1).
- *Transport and tunneling modes:* In addition to AH and ESP, IPsec provides a choice of two operational modes (Figure 2). In transport mode, only the IP payload is protected. This is useful in conjunction with L2TP for end-to-end communication between two hosts, and can also be used by routers. In tunnel mode, an outer packet encapsulates and protects the entire IP packet. This mode is used between gateways and/or hosts. Both modes are supported in AH and ESP packets.
- *Security databases:* IPsec filters IP packets based on criteria like source and destination IP address, source and destination port numbers, and transport protocol (if not encrypted and therefore inaccessible). The network administrator specifies these filtering selectors in the security policy database and associates

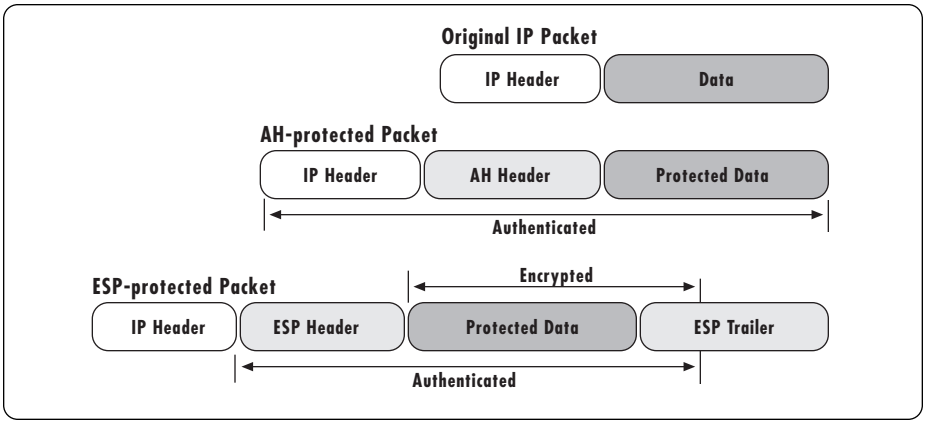

Figure 1: AH and ESP components of IPsec

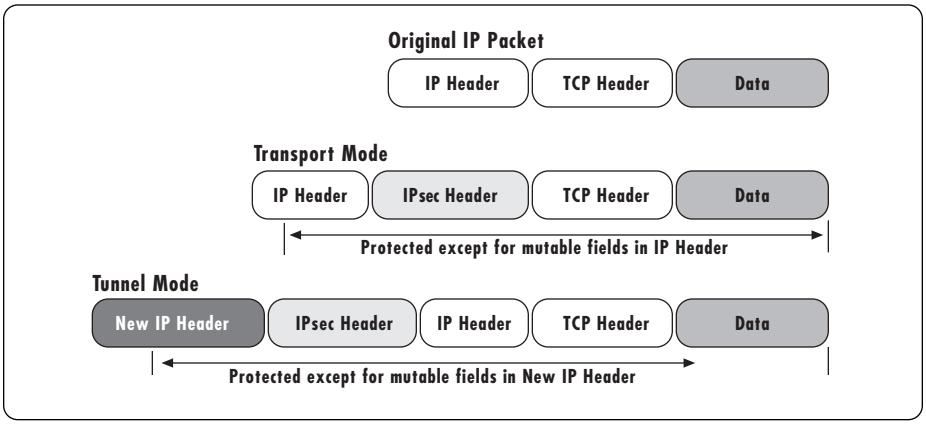

Figure 2: IPsec transport and tunnel modes

them with actions that specify whether a packet is to be dropped, may bypass, or is to be IPsec processed. Another database stores security associations (SAs) that define a relationship from sender to receiver. SAs specify and provide mutually agreed security services (IPsec mode, encryption algorithm, encryption key, and so forth) and are negotiated by IKE .

- *Internet key exchange (IKE):* IKE is the default key management protocol for IPsec. Among the many tasks performed by IKE is authenticating IKE peers using the algorithms SHA-1 and MD5 or digital certificates. The algorithm used to negotiate the session keys for IPsec is Diffie-Hellman. IKE also establishes the security association for IPsec.

Network address translation (NAT)

NAT is a mechanism for conserving registered IP addresses in large networks and simplifying IP-addressing management tasks. As its name implies, NAT translates IP addresses within private internal networks to legal IP addresses for transport over public external networks such as the Internet. Incoming traffic is translated back for delivery within the internal network. Thus, NAT allows an organization with unregistered private addresses to connect to the Internet by translating those addresses into globally registered IP addresses. NAT also increases network privacy by hiding internal IP addresses from external networks.

NAT/IPsec conflict

There is a fundamental conflict in the design of these two IP protocols that results in a variety of incompatibilities. Yet at the same time, both protocols are popular for remote network access applications such as virtual private networks, home gateways, and network servers.

The limited number of available IPv4 addresses makes NAT an absolute requirement in many networking applications. As a result, NAT became a major barrier to wide deployment of one of IPsec's principal uses. The few compatible uses of NAT and IPsec limit IPsec's flexibility and restrict the potential number of operational modes and configurations. Next are some examples of configurations that break each other and examples that work together.

Fundamentally, NAT operates by modifying end node addresses en route (within the IP header). The IPsec AH standard, on the other hand, is explicitly designed to detect alterations to IP packet headers. So when NAT alters the address information in the IP header, the destination host receiving the altered packet will invalidate the packet. As a result, the IPsec AH-secured packet traversing NAT will not reach the target application.

In the case of TCP/UDP packets, NAT needs to update the checksum in TCP/UDP headers when an address in the IP header is changed. However, since IPsec's ESP component encrypts the TCP/UDP header, NAT is not able to update the checksum. As a result, TCP/UDP packets encrypted in transport mode ESP, traversing a NAT device, will fail the TCP/UDP checksum validation on the receiving end and will not reach the target application. ESP-encrypted IPsec packets may be altered by a NAT device only in a limited number of cases.

IPsec ESP tunnels do not cover the outer IP header within the authentication hash, and so will not suffer hash invalidation due to address translation. IPsec tunnels also need not be concerned about checksum invalidation (unlike L2TP).

However, IPv6 solves the problem of limited static IP address spaces. As a result, NAT will become obsolete and IPsec a built-in feature of IPv6.

Secure socket layer (SSL) and transport layer security (TLS)

Netscape originally developed SSL, then handed over version 3.0 to the IETF, which renamed it TLS. There are only minor differences between SSLv3 and TLSv1, but enough to make the two incompatible. TLS represents an additional layer in the Internet protocol. TLS is inserted between the transport (TCP) and application layers, and requires very few changes in the protocols above and below. HTTP applications interface with TLS in nearly the same way they would with TCP. To TCP, TLS is just another application using its services. TLS cannot operate using a connectionless transport protocol like UDP. Like IPsec, TLS provides isolation between the application and security, but TLS allows some interaction between the two. For example, an application such as HTTP need not change when security is added, but the application typically has to make the decision to use TLS or not. Supported applications other than HTTP include NNTP, FTP, and LDAP. TLS uses X.509 certificates for authentication and RSA for key exchange. The most popular encryption algorithm with TLS in HTTP applications is RC-4. SSL 2.0 is widely considered to have many flaws, which makes TLS 1.0 (formerly SSL 3.0) mandatory for most deployments.

SSL/TLS is popular with Web-based device management, such as Wind River's RapidControl™ for Web.

Virtual private networks – an example of applied networking security

Virtual private networks (VPNs) are prime examples of applying the networking security technologies and protocols discussed in the previous sections. Whenever networking infrastructure is shared with other users who cannot necessarily be trusted (as in the public Inter-

net) one is exposed to many networking security threats. This is why VPNs became very popular with corporations.

A private network is built on the assumption that it is shared only by a controlled number of trusted users and therefore is safe and secure. However, many corporations' networks are not located in just one controlled location, but are physically separated and distributed over multiple sites worldwide. In order to remain one network, these sites need to be connected. A similar situation can be found with remote users like business travelers or telecommuters who want to connect to their home network. Most often, the connections are not running over costly private infrastructure but over the public Internet. As a result, these networks are only virtually private. In order to still have real privacy on these connections, the data needs to be tunneled, for example by L2TP. RADIUS can be used to authenticate the tunnel endpoints. In addition to being tunneled, the data traffic can be encrypted by IPsec.

The advantages of VPNs include reducing the need for building private networks and, by sharing the infrastructure, lowering the costs of networking equipment. Even though VPNs run partly over the public Internet, they still restrict network access to trusted parties. VPNs make it possible to co-locate virtualized equipment that enables remote access users to dial into their local ISP rather than make costly long-distance connections.

On the other hand, VPNs open the network to more security issues. After all, network data and internal services are again exposed to the public. Other disadvantages include, potentially, hiding Quality of Service markings through tunneling and the requirement to use NAT to accommodate duplicate address spaces in the various networks.

Summary

In general, every measure taken to increase security can be only so effective. None of the security technologies known and used today provides a guarantee of absolute security. All these technologies can do is lower the potential security risk. As a rule of thumb, costs of improving security go up exponentially with the achieved gain in protection.

For information about the availability of these networking security technologies on the VxWorks® real-time operating system, please contact your Wind River representative.

Wind River Worldwide Headquarters
500 Wind River Way
Alameda, CA 94501 USA
Toll free 1-800-545-WIND
Phone 1-510-748-4100
Fax 1-510-749-2010
Inquiries@windriver.com
Nasdaq: WIND

For additional contact information,
please see our Web site at www.windriver.com.