

SafetyNet: Secure Communications for Embedded High-Performance Computing

Craig S. Steele, Jeff Draper, and Jeff Koller

{steele, draper, koller}@isi.edu

USC Information Sciences Institute

4676 Admiralty Way, Marina del Rey, CA 90292

Abstract

ISI's Advanced Scalable Network Technology (ASNT) project employs multiple specialized network components in an integrated high-performance multicomputer system. Large multicomputer systems are vulnerable to programming error, hardware faults, and potentially, malicious attacks. Allowing direct user access to network interfaces is desirable for performance, but can reduce protection and reliability. Even privileged codes, such as general-purpose operating systems, are untrustworthy and buggy. ASNT's security and reliability subsystems, collectively known as SafetyNet, embed vital distributed-system communications and authentication functions to prevent even OS code from breaching security or bringing the system down. Such mechanisms are necessary to enable deployment of embedded HPC systems in mission-critical applications. In this paper we discuss the motivation and features of the SafetyNet system for embedded high-performance computing.

Extended Abstract

1.0 Motivation

Current packaging and network technology provides excellent performance for compact multicomputers, but the usability of such systems is diminished by reliability and responsiveness considerations. Shared multicomputers have generally proven to be increasingly unstable and unpredictable in performance as the size of the system and number of concurrent users rose. This fact has tended to limit their application to well-behaved scientific applications, often executed in a dedicated single-user batch environment. This mode of operation is inapplicable to large-scale command-and-control system requirements, which are inherently dynamic in resource demands and may not be fully testable against all possible situational challenges.

Traditional engineering solutions for embedded HPC typically produce complex single-point solutions requiring costly reliability engineering for many custom components. Lengthy design and qualification cycles reduce effective performance and increase lifecycle costs enormously. In contrast, a general-purpose machine, if it can be made sufficiently reliable, secure, and flexible, is much more cost-effective and versatile. Flexible use of modular components of a large sharable computational system can reduce development-cycle time and cost. Performance can be improved, in the long term by configuration upgrade to use newer-vintage processors, and in the short term by the ability to dynamically reallocate available resources for the requirements of the moment.

On the other hand, reliability of a large space-shared multicomputer system is compromised by the fact that failures from any source, node hardware, user or system software, tend to propagate more readily in an environment with fairly close hardware coupling and software which performs little error checking. Error checking and defensive protocols are not yet particularly common or effective in multicomputer system software because of the great engineering effort required to try to plug all possible holes and because such checking can materially reduce performance.

The SafetyNet security and reliability subsystems of ASNT make large multicomputers suitable and safe for sharing by a multiplicity of tasks. Under the control of the physically isolated SafetyNet security and configuration network, the multicomputer can be partitioned into distinct subsets of nodes, effectively separated by virtual firewalls on all data-moving networks. These virtual firewalls are implemented by programmable filtering hardware located at each network interface. All network packets are checked for permitted destinations and functions at the source-node network interface. Likewise, all incoming messages are validated by source and function at each destination node network interface. Unless specifically authorized, no message can be sent to a node outside the immediate virtual-firewalled set. In ASNT, computational node software does not even know the physical topology of the system; all such configuration is under the exclusive control of the SafetyNet subsystem.

The ability to flexibly and dynamically partition a large multicomputer is advantageous for an application such as a shipboard central computing system, where many different operational modes need to be supported reliably and securely. The same sensor inputs that record the number of air-defense missiles in a turret should feed both inventory- and fire-control software systems, but the former must not compromise the effectiveness of the latter in a combat mode of operation. With a fast and trustworthy dynamic reconfiguration, the same set of standardized computational resources can be employed for logistical tasks and real-time command-and-control applications.

In addition to the firewall mechanism, the demonstration implementation of ASNT supports selective real-time encryption of the I/O network data, which may be from physically distant sources such as sensor arrays or another multicomputer. The other ASNT networks are assumed to be physically compact and secure. In contrast the I/O network is assumed to be potentially vulnerable to observation, and its data can be encrypted.

The I/O subsystem consists of the I/O network and a number of I/O nodes. The I/O nodes are usually paired, in a closely-coupled "Siamese" fashion, with computational nodes (Figure 1.) The I/O node supports a PCI local peripheral expansion bus which may contain storage, network, display, or sensor interfaces. The I/O node memory subsystem supports PCI transfers concurrently with memory-to-memory DMA transfers to the computational node.

Encrypted data may be transferred between I/O nodes or moved to storage media in encrypted form. Even the I/O node operating system is not privy to the encryption and authentication keys, so the data is resistant to tampering or disclosure. The data is

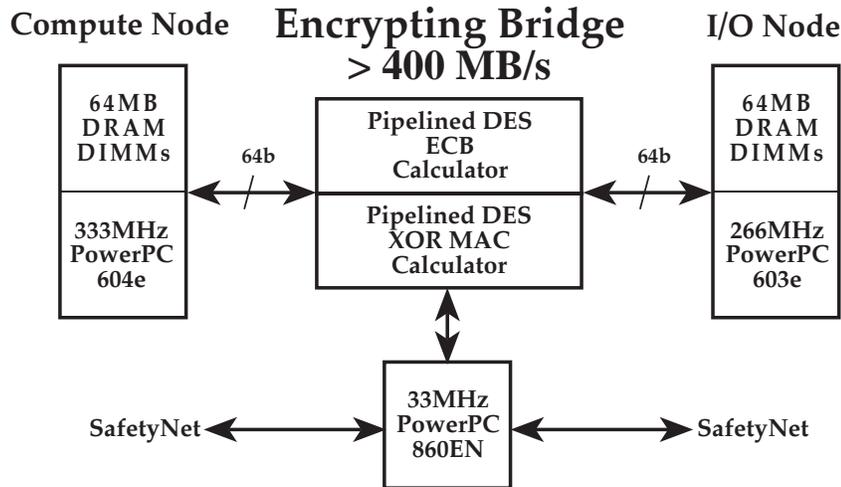


Figure 1. Major components of SafetyNet. The encrypting bridge connects the compute and I/O nodes, which have their own specialized networks (not shown). The embedded processor controls the bridge and network filters, and has a physically separate network connecting it and other bridge nodes to the system security nodes.

decrypted and authenticated before use by the computational node processor, when it crosses the bus bridge between the siamesed nodes.

2.0 SafetyNet Component Overview

The SafetyNet is composed of several hardware and software components (Figure 1.)

1. Bridge nodes – The bridge nodes are interposed between each computational and I/O node pair, and control all initialization and configuration of both nodes and their respective network interfaces. They are in turn controlled by the security nodes, receiving commands via the SafetyNet network. The processor embedded in each bridge node maintains the descriptors for secure objects required by the encrypting bus bridge in each node pair.
2. SafetyNet network – This is a physically distinct network providing communications between the SafetyNet embedded processors located on each computational and I/O node pair and the security nodes
3. SafetyNet security nodes – These are general-purpose computers which manage the initialization, node resource management, and security functions of the ASNT prototype system. For development and demonstration purposes the security nodes are standard PC or workstation computers, but the implementation supports dedication of one or more I/O nodes to this purpose for embedded applications.

4. Network Interface Filters – The network interfaces of the control, data, and I/O networks filter messages sent and received from the corresponding networks to maintain the logical partitioning and segregation of the physical system into logical subsets.
5. Encrypting bus bridge – The 64b-wide busses of the computational and I/O nodes are linked by a bus bridge (part of the bridge node) which is capable of moving, encrypting, and authenticating secure-object data at full bus bandwidth. The address mappings between objects images in the encrypted I/O node memory and the unencrypted computational node memory, along with the associated object DES keys, are maintained in the bridge node memory.
6. SafetyNet system software – SafetyNet has a substantial software component. The bridge nodes have diagnostic and initialization functions in addition to support functions for the network interfaces and encrypting bridge. Reliable network communications with the SafetyNet network, including robust flow control, and intermediation between I/O nodes and security nodes require substantial and reliable code development. Security node software requirements are even more numerous, and includes diagnostic and performance monitoring functions in addition to global resource management and top-level key-distribution functions.
7. SafetyNet application-code interface – A simple programming interface for code in the computational and I/O nodes is provided to support creation, access, and destruction of secure objects

3.0 SafetyNet Functional Overview

The SafetyNet subsystem has multiple functions in the ASNT prototype system.

1. Initialization – The SafetyNet subsystem initializes the computational and I/O node hardware and the other networks and network interfaces. Initialization functions include control of hardware reset functions of the computational and I/O nodes, and initialization of programmable logic components of node and network interface hardware. The SafetyNet serves as a maintenance and diagnostic network in addition to its security functions.
2. Configuration – The SafetyNet subsystem controls the allocation of physical resources of the ASNT prototype system, and configures the node and network-interface hardware to reflect specific assignments of logical identities and logical-to-physical mappings for network operations. Specifically, this includes setting up node ID registers and routing tables for network transmissions.
3. Flexible Firewall Partitioning – This is the strongest security and reliability mechanism in ASNT, and is an elaboration of the configuration functions of SafetyNet. The multicomputer nodes are effectively partitioned into

non-overlapping groups, with restricted or no communication. This is accomplished by programming computational and I/O node network interfaces to filter both outgoing and incoming messages. This ensures that source and destination identities are within the authorized subsets of nodes defining a logical computational partition.

4. Secure & Authentic Object management – SafetyNet provides an object-oriented system for encrypting and authenticating data. The SafetyNet components include a very-high-performance bus bridge which performs DES encryption[1] and DES XOR MAC authentication [2] for transmitted data. SafetyNet initializes and maintains the object descriptors for the bridge hardware, and provides a hardware interface to computational and I/O node processors to create, share, and destroy secure objects.
5. DES key management – Support of secure objects requires creation, distribution, expiration, and evolution of cryptographic keys. All of these functions are performed by the SafetyNet. No key is ever visible to a computational or I/O node processor.
6. Error reporting – Errors detected by SafetyNet hardware components are logged by embedded processors and transmitted via the SafetyNet network component to SafetyNet security nodes

4.0 Acknowledgments

The ASNT project is sponsored by the DARPA Information Technology Office.

5.0 References

- [1] Bruce Schneier, *Applied Cryptography*, second edition, John Wiley & Sons, 1996.
- [2] M. Bellare, R. Gu erin, P. Rogaway, "XOR MACs: New Methods for Message Authentication Using Finite Pseudorandom Functions," *Advances in Cryptology – Crypto 95 Proceedings*, Lecture Notes in Computer Science Vol. 963, D. Coppersmith ed., Springer-Verlag, 1995.