Security Approach for a Resource Management System

Cynthia E. Irvine Timothy Levin
Department of Computer Science
Naval Postgraduate School, Monterey, CA 93943
(irvine,levin)@cs.nps.navy.mil

ABSTRACT. We present an overview of our approach to including security in the scheduling mechanism of the Management System for Heterogeneous Networks (MSHN) resource management system.

Resource Management Systems (RMSs) are responsible for efficiently scheduling multiple tasks onto computing and network resources in a distributed heterogeneous computing environment. RMSs support QoS by scheduling to meet user requirements for performance and security, and by providing support for tasks to adapt to changing network resource availability.

Whereas network operating systems typically have exclusive control over the access to and utilization level of resources, the MSHN RMS sacrifices such control in favor of compatibility with existing applications and operating systems. The RMS constructs task schedules based on its network infrastructure model. This model includes the resource and security requirements of current and waiting tasks, and the security requirements and availability of network, computing and storage resources. The resulting schedules are provided to task handlers who run the tasks and provide feedback to the scheduler. If the model is inaccurate (e.g., security or resource availability changes), the RMS adjusts its model and potentially reschedules the tasks.

RMS schedule construction consists of two logical phases: reduction and optimization. In the reduction phase, the scheduler finds the *realizable* resource assignments for the task by discarding the possible assignments that will not work according to the model. In addition to resource capacity, availability and type matching (e.g., task bandwidth requirements vs. resource capacity), security plays a key role. The security requirements of the network resources are compared to the task's security characteristics to determine where the task can run. Additionally, the task's security requirements (e.g., reflecting the user's QoS security specification) are compared to the services available from the resources and the infrastructure. The result is a set of resource-assignment ”solutions,” where each solution identifies various resources sufficient to run the task.

In the optimization phase, an ”optimum” solution is heuristically selected. The criteria for selection is to (attempt to) minimize costs and to maximize the QoS benefit to the users. I.e., using realizable resources from the reduction phase, the scheduler attempts to create a schedule to meet QoS requirements for all of its tasks. In order to support as many tasks as possible, the scheduler must meet the typical task scheduling constraints while minimizing resource usage costs. Different users of a task may request different degrees of support from each requested security service, and the scheduler may adapt security support, within system- and user-defined ranges, in order to schedule the tasks most efficiently. Therefore, near-optimal solution selection depends on the accurate estimation of per-task, per-resource, cost of security.

Our research to support the MSHN RMS involves: how to map user security requirements to network security mechanism abstractions; understanding the range of security services and mechanisms the RMS scheduler must consider; understanding security-variant policies and mechanisms of emerging technologies; how to measure quality of security service, including the effects of redundant security mechanisms; and how to specify the cost of security including resource overhead, economic costs and other factors.