



Information System Security Operation Cooperative Intrusion Traceback and Response Architecture

Framework to Integrate Diverse Components into a Cooperative, Self-Protecting Network

Motivation

The Cooperative Intrusion Traceback and Response Architecture (CITRA) is an infrastructure for integrating intrusion detection systems, security management systems, firewalls, routers, hosts, and other components to trace attacks back to their true source and block the attacks close to that source. The CITRA components interact to:

- Trace intrusions across network boundaries;
- Prevent or mitigate subsequent damage from intrusions;
- Consolidate and report intrusion activities; and
- Coordinate intrusion responses on a system-wide basis.
- Automating these tasks reduces the costs and damage done by an intrusion, since a response can be taken more quickly than a human administrator could be summoned, perform complex analysis of the problem, and take manual action.

A CITRA “community” is the administrative domain for control of a network. A CITRA community contains the following components:

- *Detectors* - identify intrusions and report them to local Responders and to the Discovery Coordinator.
- *Responders* - take action against intrusions and report their immediate actions to the Discovery Coordinator.
- *Audit processes* - running on each boundary controller and optionally on other

hosts, keep a history of network activity from, to, and through the CITRA node.

- *Discovery Coordinator* - sees reports for all intrusions and their immediate responses, determines the best overall response for the network, and directs Responders in a coordinated effort.

Certain components can be both detectors and responders.

CITRA provides an architecture that makes it possible to integrate existing, diverse COTS components into a cooperative, self-protecting network. The CITRA framework facilitates the incorporation of these components into the CITRA network.

A CITRA community can be subdivided into smaller interconnected entities, known as neighborhoods. Routers, firewalls, and other boundary controllers are the points at which the subdivisions are interconnected. Intrusions are traced back to the first neighborhood in the community to have seen the intrusion, and possibly back to the host and process, if the relevant host is instrumented with CITRA. The intrusion is traced back through hosts, firewalls and other nodes even though IP addresses and ports may change as the intrusion makes its way through the community.

The actions taken when an intrusion is detected are controlled by the policy established by the administrator. The various elements of the policy can be specific to neighborhoods, specific to hosts, or global to the community. The CITRA philosophy frequently has responders take a short-lived, but fairly severe local response to an intrusion as soon as the intrusion is detected. This stops the intrusive action immediately, even at the expense of legitimate operations. The

This work was sponsored by DARPA through Boeing Phantom Works, with McAfee Research, which is now the Security Research Division of SPARTA.



Cooperative Intrusion Traceback and Response Architecture

Framework to Integrate Diverse Components into a Cooperative, Self-Protecting Network

CITRA components trace the intrusion back to all the neighborhoods and hosts that have seen the attack. The Discovery Coordinator consolidates the reports and presents a global view of the attack to the operator. Once the operator has determined a more optimum response to the attack, he can issue new directives that can potentially override the earlier local responses.

- Some of the more recent additions to CITRA permit separately administered communities to cooperate in tracking down an intrusion that spans networks. Each community's policy identifies the other communities with which it is willing to cooperate, how much it is willing to trust the other community, and how much of its information it is willing to share.

The CITRA components communicate with each other via the Intruder Detection and Isolation Protocol (IDIP). IDIP has two layers: an application layer and a message layer.

The IDIP application layer protocol accomplishes intrusion tracking and containment through three major message types:

- *Trace requests* - determine whether an attack has been seen by a node, and to suggest a short-lived immediate response to the attack when it has been seen.
- *Reports (copies of trace requests with the response taken attached)* - inform the Discovery Coordinator of the attack and the immediate response.
- *Discovery Coordinator directives* - are used to direct nodes to undo responses or take another response to the attack, based on a more comprehensive view of the attack.

The IDIP message layer is designed to provide survivable, reliable, secure transport of messages. The message layer uses cryptography, modeled after IP Security (IPSec), to protect against attacker eavesdropping and to provide integrity of delivered messages.

CITRA Nodes Trace Attacks From One Neighborhood To Another Within A Community

