

Collaborative security — the site's perspective

James A. Rome

jar@ornl.gov

Presented at Globus World

January 2004

The submitted manuscript has been authored by a contractor of the U.S. Government under Contract No. DE-AC05-00OR22725. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

The issue

- **In this era of heightened security awareness, collaboration must fit into the security requirements of the sites involved in the collaboration.**
- **It is not enough to create and implement a security infrastructure because it must also meet the needs of the site security officials.**
- **The policies of the collaboration must be reconciled with those of the participating sites.**

External constraints

The Federal government, and DOE in particular, are raising the ante on cyber security. Among the new requirements are:

- **All users must have an official computer account and a strict requirement for some sort of “vetting” before a computer account is created. This processing can take months for foreign nationals.**
- **An audit trail must be maintained of who has access to what devices. Increasingly, it is not enough to say this is done, there are now metrics of compliance and you must also prove it.**

Some rules we live by

- **National policy on the transfer of technical information** <http://www.fas.org/irp/offdocs/nsdd/nsdd-189.htm>.
- **Sensitive but unclassified data.** Things like proprietary information, CRADAs, work for NASA (which has release restrictions) must be protected from access by foreigners.
- **DOE security directives focus on protecting confidentiality with minimal regard for ensuring appropriate release of information.**
- **Department of Commerce export control list.**
- **Audit findings.**

Cross-site data flows

In collaboration, data flows back and forth between sites.

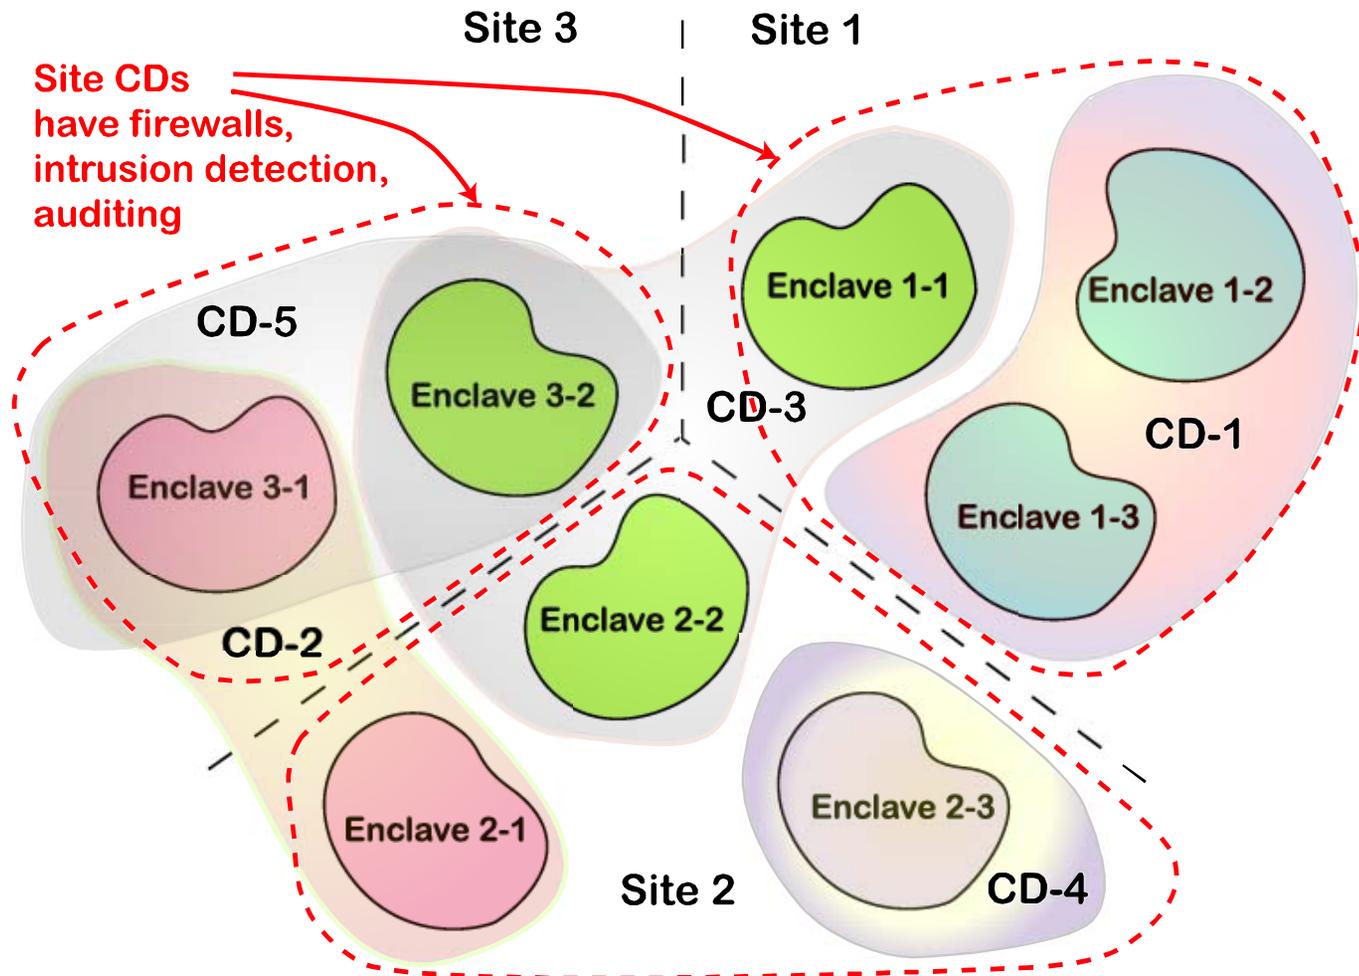
- **There must be a clear line (accepted by DOE) as to which site is liable for what, and when that liability ceases.**
- **If different sites have different security levels, how are these reconciled? Indeed, the collaboration may have security requirements that are different than the member sites (possibly higher).**
- **How to access remote resources at a different security level must be defined. This is especially difficult with shared resources such as supercomputers. Other users on the resource may pose a threat.**

Enclaves and collaborative domains

To define and resolve security policies across organizations, it is useful to break the problem into enclaves and collaborative domains (CDs).

- **An enclave is specific to one site and is a collection of resources that are governed by a common site security policy.**
- **A collaborative domain is the fabric that instantiates the collaboration and connects the enclaves. Membership in a CD gives the user some special privileges and obligations.**
 - ⇒ **Examples: TeraGrid, a collaboratory**

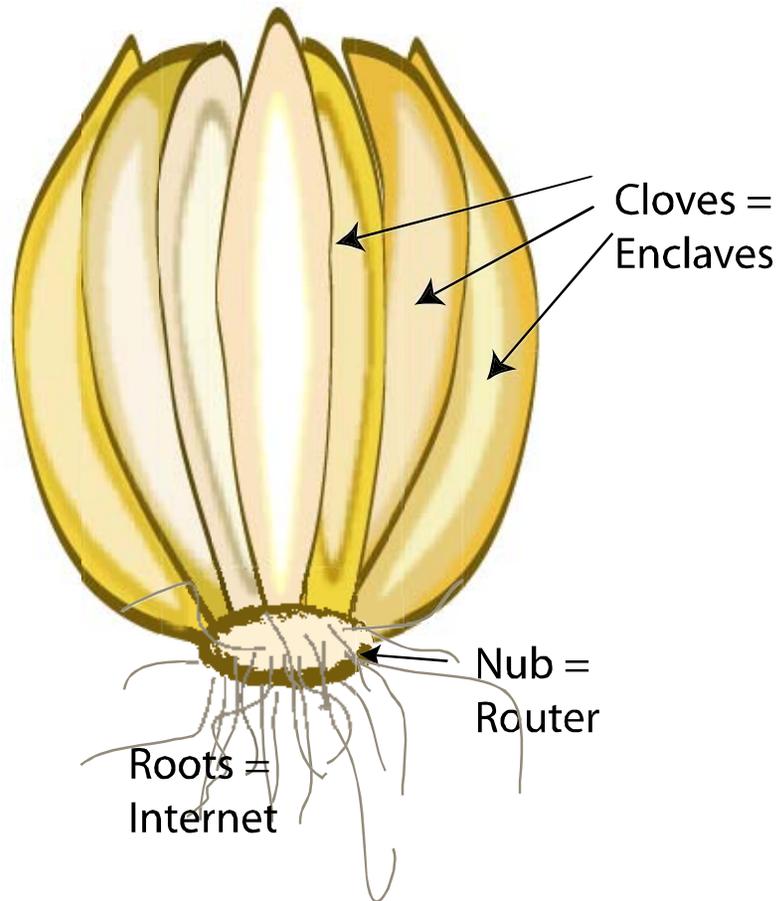
Enclaves and CDs diagrammed



Enclave principles

- **Every computing resource must be in one and only one enclave.**
- **A user (or a process initiated by a user) enters an enclave when a resource in the enclave is used.**
- **“Entering” a different enclave from a CD or another enclave must entail some sort of access control.**
- **Data can only be moved between enclaves by a user (or a user process) who is a member of both.**
- **For all its enclaves, a CD must satisfy the enclave security requirements imposed on the CD by the enclaves plus those unique to the CD.**

The garlic model of site security



- This is in contrast to the usual onion model.
- The goal is to prevent data protected at one level from having to pass through domains with a different protection level.
- It is difficult to fully implement this on shared resources (supercomputers).

The two-edged sword of encryption

Encryption provides confidentiality and integrity of the traffic, but it prevents the site intrusion detection systems from seeing what is actually happening in these encrypted streams.

- **Sites have been attacked through encrypted tunnels (e.g., ssh) using real user IDs and passwords that were stolen via keyboard or tty sniffers.**
- **It is difficult to tell whether a calculation is legitimate or being done by a terrorist (with stolen credentials).**
- **Central ssh servers can decrypt the traffic and re-encrypt it on the link to its destination.**
 - ⇒ **Only one ssh firewall exception.**

One-time password tokens

Given the problem with sniffed passwords, many sites are implementing one-time-password (OTP) tokens.

- **How does a machine use a OTP token?**
- **If sites have different tokens, how many does a person need to access all the sites? How can they be managed? How do they keep track of the different PINs?**

Cross-realm token trust

Let the user authenticate with his site's token at his site and somehow transfer this trust.

- **FermiLab uses tokens to get a Kerberos TGT that contains a HW_AUTH flag. This ticket can be used to get Grid credentials, but currently, the certificates do not contain the information that they were obtained using a token.**
- **RADIUS proxies can allow one site to authenticate against a token from another site.**

Present solutions do not scale

The general security infrastructure of the Labs is approaching many limits because the present methods do not scale.

- **ORNL captures nearly 100 GB/day of traffic that must be analyzed for nefarious activity. Currently, we are using only small fraction of our available bandwidth, and our bandwidth is about to go up by a factor of 10.**
- **Computers talking to each other doing distributed computing can generate much more traffic than we mere mortals.**
- **Preparing for audits is increasingly onerous.**

Summary

Site Security Officers have a broader and different set of security issues that they must reconcile with their user's need to “get their work done.”

- **There are technological solutions to most of these issues, but solving the political issues involved in getting general concurrence among all the sites is a daunting task.**
- **Fielding a technical solution without reconciling policy conflicts is not acceptable.**
- **There needs to be a reusable framework developed that will allow new collaborations to get started without redoing all of the red tape.**