

# Cyber Security Spring '06 Final Project Scenarios

## *Second Draft*

### **1. Common Requirements**

There are three final project scenarios. Each scenario has a customer assigned. You can ask that customer or Prof. Hinrichs for clarification of the requirements.

In all scenarios, your group will be responsible for creating

- A security policy and a threat model. What are the goals of the architecture? What are the threats that design is concerned with?
- A security architecture design. This design should identify what technologies are used and where. It should discuss the implementation and maintenance issues (e.g. key management and access changes in the in face of a changing population). Where appropriate, the design should discuss the tradeoffs and the motivations for choosing one technology or technique over another. The design should include an overview diagram which can be hand drawn.
- A laboratory implementation for a subset of the design. Depending on what is implemented, you should submit an implementation design, configuration files, and supporting log data.

#### **1.1 Important Dates**

April 11: group members and scenario assignments finalized

April 11 – April 18: groups meet with Prof. Hinrichs for initial design review and identification of lab implementation subset. An initial security policy is due at this time.

April 27 and May 2: In class presentation of design.

May 5: Final design and lab due

### **2. Collaborative Information Sharing Scenario**

Customer is Chad Hanson (chanson@TrustedCS.com).

In this scenario, a number of different organizations are collaborating to address an urgent problem. Each organization has strong information labeling and information flow constraints. Each organization has a separate user authentication space.

The primary goals for this architecture are:

- Flexible but high assurance entity authentication
- Flexible but high assurance information sharing.

#### **2.1 Collaborative Environment**

In response to an emergency, we need a scheme to quickly map how the labeling schemes relate and have an automated means to share information between the different

organizations. The emergency may be by a natural disaster like Katrina, a terrorist act like 9-11, or a regional war like in Iraq or Sudan. In all cases, people from a variety of organizations will need to share information starting very quickly for the period of weeks to years. This can be very sensitive information, so the design must also be careful to not drop security so much that the malicious entity can take advantage of the chaos of the event to gain access to restricted information.

Several approaches have been taken to share data between organizations. One approach is to have each member of the coalition maintain their own portion of the data and use access control or a guard approach to automatically enable a process of upgrading/downgrading data between different labels.

Alternatively, the coalition could create a joint data repository or community of interest that is accessed by all organizations. The joint authority can either be hosted by a "lead" organization (this is reasonable in a military setup), by a trusted third party (not easy to find), or maintained with a consensus based policy approach. Some recent work on the joint repository approach is described in the following papers:

- Laura Pearlman, Von Welch, Ian Foster, Carl Kesselman, and Steven Tuecke. [A Community Authorization Service for group collaboration](#). In *Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks*, 2002.
- Rakesh Bobba, Serban Gavrilă, Virgil Gligor, Himanshu Khurana, and Radostina Koleva. [Administering Access Control in Dynamic Coalitions](#). In *Proceedings of the 19th USENIX Large Installation System Administration Conference (LISA)*, Tucson, AZ, December 2005.

In addition to enabling information sharing, your design will also need to address how people are authenticated into the system. Since these collaborations are dynamic and not pre-planned a basic password scheme is not going to be sufficient. Most technologies that attempt scalable authentication use some form of certificates plus strong multi-factor authentication. Safely deploying and maintaining long-lived certificates is a major concern.

In practice only limited forms of multi-factor authentication may be viable for coalition environments. This is because each organization is likely to retain its own identity certification process that is trusted by other domains in the coalition. Therefore, trusting multiple factors for the authentication gets complicated. At the same time, the coalition as a whole may decide that the local domains enforce strong authentication requirements for users that are granted access to coalition data.

## **2.2 Collaborative Infrastructure Requirements**

- High-assurance environment.
- Strong, flexible cross-organization authentication

- Certificate-based
- Strong, flexible cross-organization data sharing.
  - Automated, safe data-sharing

### **3. Managed Blog Service Scenario**

Customer is Alan M. Carroll ([amc@thought-mesh.net](mailto:amc@thought-mesh.net)).

The customer is setting up a managed service provider company, *Pundits R Us*. This company will provide web log services for its clients. The company is targeting high-profile web loggers with large bandwidth needs and a strong potentially polarized following.

The major security goals for the company are:

- Authentication scheme and access control.
- Protection from weblog-spamming.
- Protection from denial of service attacks from disgruntled readers.

#### **3.1 Managed Service Provider Environment**

High uptime and good quality are critical for the managed service provider to keep their customer base. Anything that denies or degrades service such as lowered bandwidth or garbage entries must be avoided. This implies that avoiding Denial of Service attacks is important. It also means that the service provider avoids gaining a reputation of being a source of spam or other attacks. Otherwise, other service providers will block traffic and links from *Pundits R Us* resulting in bad service for paying customers.

The customer is assuming that he will be starting from an Apache and Moveable Type base, but he is open to other suggestions.

Another unique aspect of the service provider environment is limited trust between customers and between the customer and the service provider, which requires segmenting customers from each other and from the service provider. This segmentation takes several forms:

- Protecting customers from other infected customers (virus, spam).
- Preserving integrity of information between customers and between customers and the service provider itself.

The customer wants to provide free accounts with limited services and paid for by advertising. He also wants to provide a “premier” paid account that provides more services.

The customer wants to differentiate his weblog service by having a very flexible authentication and access control scheme. Authors should be able to delegate privileges to their blogs to new users or authors of other blogs. The author should also be able to attach access controls to individual entries and comments.

### **3.2 Service Provider Infrastructure Requirements**

- Authentication scheme to register authors and a delegation scheme to enable authors to create subauthors and delegate authorship to their blog to existing author's of other blogs.
- Free blog and paid blog with different levels of service.
- Protection from weblog-spamming.
- Weblog and post access controls for readers.
- Protection from denial of service attacks from disgruntled readers.

## **4. Research Organization Scenario**

Customer is Dave Musselman (mussulma@cs.uiuc.edu).

In this scenario, you are responsible for designing the next generation security architecture for the *Information Trust Institute*, a large research organization that collaborates with many other organizations.

The major goals for this design are:

- Protect core organization assets
- Prevent organization assets from being subverted and used as launch points for broader attacks
- Enable flexibility setting up trusted members of the environment. Visitors should be quickly and easily given access to the network.
- Enable security research with direct access to the internet or with known malicious software, but constrain the spread of such experiments

### **4.1 The Research Environment**

The academic environment is inherently very dynamic both in terms of people and technology. People range from very technically savvy to rather technically naïve.

Unlike the commercial environments you cannot dictate the hardware platforms and OS images or versions deployed. A somewhat standard windows environment is used by the administrative staff. Research labs will introduce a wide variety of somewhat esoteric hardware and software. Some of these labs can be isolated from the outside world, but in the age of the internet, some labs must be connected to the greater world.

Labs studying network and computer security place special constraints on the research organizations infrastructure. Some experiments must be performed on a “dirty” network, e.g., honey pots will not capture new viruses if they are on well protected networks. Other labs will involve experimenting with hazardous pieces of malware that must not be allowed to escape a constrained environment.

Visitors and students will bring in a wide variety of laptops running a wide variety of OS images and programs. Students tend to intentionally or accidentally try a wide variety of programs that can have unexpected consequences.

Recently students also raise legal concerns from the music and entertainment industry. Department officials hope to ignore the whole issue, but they must protect themselves if a large entertainment company puts the university into its sights. Recent changes to the CALEA interpretation within a university environment also places additional constraints in the security environment. The infrastructure must be secure, but we must be able to tap into the network with required by law enforcement.

The research organization must interface with the broader university community. There is a university-wide infrastructure for authentication, storage, and computation. It may or may not make sense to leverage this infrastructure, but your design must at least co-exist with the university infrastructure. Your organization's infrastructure must also be able to support visitors from other departments. Many researchers are collaborating across disciplines and such cross department people must be able to work in any of their home departments.

## ***4.2 Research Infrastructure Requirements***

Specifically, the research organization must provide the following cyber infrastructures

- Wireless connectivity for all members of the community plus easy access for guests.
- Wired connectivity to offices and appropriate research labs. Enforced isolation for other labs.
- E-mail service for members of the community.
- Relatively small number of authentication mechanisms.