Moving Application Security into the Network

Steven M. Bellovin

smb@cs.columbia.edu

http://www.cs.columbia.edu/˜smb

Dept. of Computer Science, Columbia University

Issues with Security Mechanisms in the Net

- Firewalls
- Logging in
- Proxies
- Communicating policy
- Privacy

Firewalls

- Firewalls are the most obvious form of security device in today's networks
- They're a broad-brush solution, and assume that bad guys are only on the outside
- But — we use them because they provide *scalable* protection
- Equally important, these policies can't be subverted by a random compromised host
- My view in 1994: "Firewalls are not a solution to network problems. They are a network response to a host security problem."
- Most of that is still true

The Trouble with Today’s Firewalls

- They rely on an accident of topology — and of ancient topology; today’s corporate networks are far more interconnected
- “A sort of crunchy shell around a soft, chewy center”
- Too inflexible in the face of new protocols (though sometimes that’s a benefit!)

A Co-operative Firewall Architecture

- Many network elements enforce policy
- Hosts communicate their identity to the network; this identity includes a policy (or a pointer to a policy)
- The policy is cryptographically signed offline, and hence not subvertible by next-generation worms
- “No login or no policy, no service”
- Applications with special needs (i.e., FTP and SIP) communicate explicitly with firewall elements

Logging In

- Login protocols can be computationally expensive
- Can routers be overloaded by malicious hosts?
- Talk to a login server, which talks to the routers by protected channels?
- What about multiple identities?

Multiple Identities

- Users don’t have a single identity
- Example: at the moment, I’m retrieving email from three different servers, all of which use different credentials; I’m also logged on to two remote hosts and three different IM servers, one via a proxy. Who am I?
- Different instantiations of “me” have different privileges. Some of those privileges are dependent on physical connectivity and device being used. (The *only* way to log into my office desktop machine is via physical access or via cryptographic negotiation from exactly two other machines. Remote passwords simply don’t work.)
- Use different network addresses for different identities?

Proxies

- Proxies – web, email, and more — act on behalf of many different users
- They must therefore have different identities to the network when representing different users
- Query: whom should an email gateway claim to be when forwarding inbound email from an unknown outside user?
- How can the network trust a machine whose identity keeps changing?
- Compromised machines act as proxies for the bad guys — but they announce a good guy's identity

Communicating Policy

- Security policy must cover all paths between a host and any possible bad guy
- You may know the path from you to some server — but the bad guy can be anywhere
- Policy enforcement may need to take place at layer 2 as well
- How do applications request a policy variance? When should this be permitted?
- Policy requests can come from many places: the network owner, the machine, the user's organization, and individual applications. How are these merged, reconciled, distributed, etc.?

Privacy

- Does the network know you’re a dog?
- If the network knows who you are, to whom can it announce this?
- Today, proxies can buy a fair amount of privacy. Will tomorrow’s proxies announce your real identity?
- On the other hand, if there is a societal consensus against too much privacy and anonymity, how can network elements retain enough state?

Conclusions

- Moving security into the network has advantages, but it's not easy
- Many of the design options require tradeoffs between equally desirable goals
- Often, we'll want to be able to switch among these goals at different times or in different places; our network architecture shouldn't constrain our choices