# Foundations of Computer Security

## Lecture 18: Non-Interference II

Dr. Bill Young
Department of Computer Sciences
University of Texas at Austin

# Non-Interference Policies

Under BLP, the metapolicy for the system on the right is: *information may flow from L to H, but not vice versa.*

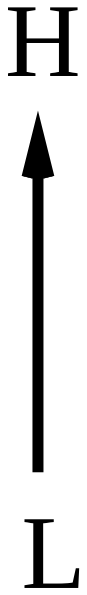

The Non-Interference version is just:

$$L \mapsto H$$

Notice how closely the NI policy mimics the confidentiality metapolicy.

There're no rules about which subjects can read/write which objects. In fact, nothing about objects or actions at all.

# Verifying NI

An NI policy is nicely abstract. *But how could one show that a system satisfies it?*

Suppose L and H were the only users in your system and you need to show that system satisfies the NI policy: $L \mapsto H$.

In a system satisfying that policy, no actions by H should have *any* effect visible to L.

# Verifying NI

Imagine an arbitrary interleaving of actions by the two subjects:

$$l_1, l_2, h_1, l_3, h_2, h_3, \ldots, l_k, h_j, \ldots$$

where $l_i$ and $h_i$ are the $i^{th}$ actions by L and H, respectively.

What L sees after this system runs should be *exactly* what L sees after the system runs the following instruction sequence:

$$l_1, l_2, l_3, \ldots, l_k, \ldots$$

This observation gives a way, at least conceptually, of verifying whether the NI policy is satisfied. If you could prove that L's "view" of the two runs will always be identical, the policy holds.

# Verifying NI

Anything L might “view” are things that H’s actions may not affect.

So, the policy can be made stronger by enlarging L’s “view.”

- Include within L’s view only the contents of files L could see under BLP, then you have exactly BLP.
- Include within L’s view the values of all system flags, then those can’t be used in any covert channel to L.
- Include the system clock, then that can’t be used in any timing covert channel to L.
- If you include *everything* L could ever observe, then there’s *nothing* H can use to send information to L.

# Verifying NI

*So why not include everything L could ever observe within his view?*

- Interferences are very common in real systems.
- Most involve low-level system attributes.
- Many “interferences” are benign, e.g., encrypted files.

Proving NI for realistic systems is extremely difficult.

# Lessons

- Non-Interference is an expressive, intuitive policy that mimics the confidentiality metapolicy.
- There are methods of establishing that a system satisfies NI.
- However, realistic systems have many potential interferences.

**Next lecture:** What is Integrity?