

CS329E: Elements of Security

Policies and Channels

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

Last updated: February 11, 2015 at 10:42

How Do We Define Security?

We said that in the most general terms, *security* seems to mean something like “protection of assets against attack.”

But this question is very specific to the context. Security for a wireless phone system may be very different from security for a military database system or an on-line banking system.

Often, security for a given system is defined in terms of a *security policy*, also sometimes called a *security model*.

The policy is the system *specification* wrt security. Another way to think of it is as a *contract* between the designer/implementor and the customer. That's why it needs to be both achievable and adequate for the intended uses.



The policy defines what “security” means for a given system or family of systems. A policy may be characterized informally, semi-formally, or formally. It may be very abstract or very concrete.

Thought Experiment #1

Your academic records are stored on computers at the university.
Design a security policy to protect them.

Start by asking: What does it mean “to protect them”? What are you protecting and what are the potential threats? Who are the stakeholders, i.e., whose interests are at risk? Do they conflict? Which of the following should you care about: confidentiality, integrity, availability?



<http://catalog.utexas.edu/general-information/appendices/appendix-c/educational-records/> outlines these rules for the university

Metapolicy vs. Policy

I like to make the distinction between the *metapolicy* and the *policy*. This is not a distinction which is often drawn in the security literature, but I think it's a very useful one.

metapolicy: The security goals in the most abstract sense.

policy: A system-specific refinement of the metapolicy adequate to provide guidance to developers and users of the system.



Metapolicy vs. Policy Example

The following is part of the UT Austin policy for protecting educational records:

Policy: faculty/staff may not use student SSNs in documents/files/postings; all older docs containing SSNs must be destroyed unless deemed necessary; documents deemed necessary must be kept in secure storage; etc.

But what are these rules trying to accomplish? I.e., what is the metapolicy?

Metapolicy vs. Policy Example

The following is part of the UT Austin policy for protecting educational records:

Policy: faculty/staff may not use student SSNs in documents/files/postings; all older docs containing SSNs must be destroyed unless deemed necessary; documents deemed necessary must be kept in secure storage; etc.

But what are these rules trying to accomplish? I.e., what is the metapolicy?

Metapolicy: social security numbers of students should be protected from disclosure.

Often, the policy rules may seem obscure and arbitrary without knowing the metapolicy. I.e., what's the underlying goal?

Why Bother with the Policy?

If the metapolicy is what we really care about, why do we still need a policy at all?

Why Bother with the Policy?

If the metapolicy is what we really care about, why do we still need a policy at all?

- The metapolicy is often too general to provide adequate guidance.
- The metapolicy may be subject to multiple interpretations.
- The policy provides specific, enforceable guidelines to the user/developer.
- Multiple acceptable policies may accomplish the security goals.

Aside: Mechanisms for Building a Solution

Often a security solution/policy (access control) is phrased in terms of the following three categories:

- Objects:** the items being protected by the system (documents, files, directories, databases, transactions, etc.)
- Subjects:** entities (users, processes, etc.) that execute activities and request access to objects.
- Actions:** operations, primitive or complex, that can operate on objects and must be controlled.

Aside: Mechanisms for Building a Solution

We often specify the security policy in terms of which subjects can perform which actions on which objects.



For example, in the Unix operating system, processes (subjects) may have permission to perform read, write or execute (actions) on files (objects). In addition, processes can create other processes, create and delete files, etc.

Certain processes (running with root permission) can do almost anything. That is one approach to the security problem.

Security Design Process

Anyone developing a secure system might address (iteratively) the following questions:

- 1 What are you protecting and what are the potential threats? (risk assessment)
- 2 What is the intuitive notion of security for such a system? (metapolicy)
- 3 What are appropriate security rules that attempt to capture this notion for this system? (policy)
- 4 What is a system architecture that supports our security goals? (system design)
- 5 By what specific mechanisms might the security goals be accomplished? (detailed design)

Of course, there are lots of other questions that need to be addressed as well in any development.

Thought Experiment #2

Suppose you have several secure LAN's that are geographically distributed, and must communicate securely over an insecure backbone network (the Internet).

Try to address the Security Design Process questions for this problem.

Thought Experiment #2

Suppose you have several secure LAN's that are geographically distributed, and must communicate securely over an insecure backbone network (the Internet).

Try to address the Security Design Process questions for this problem.

Caution: Don't jump too quickly to an implementation without considering what you're trying to accomplish.

If you think cryptography will solve your problem, then you don't understand cryptography ... and you don't understand your problem. –Bruce Schneier

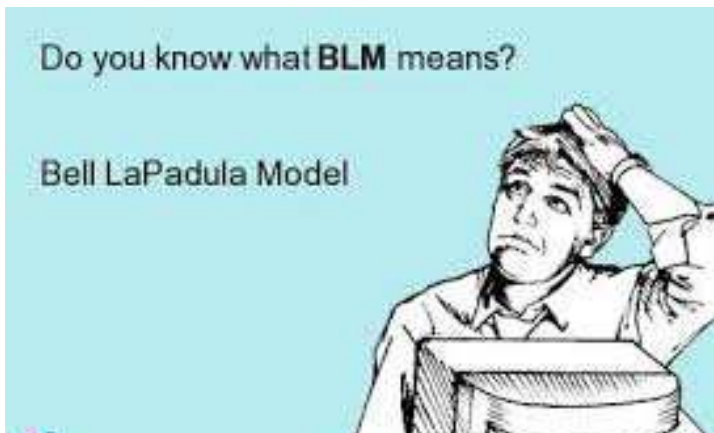
After you've gone through the design process for your system, you have to assess how well you've succeeded. There are some additional questions to be asked:

- 1 Do the system design and implementation accomplish the security goals expressed by the policy?
- 2 How do you know? How certain are you of the assessment? What is the evidence?
- 3 Are there intuitively insecure behaviors that fall outside the range of the policy?
- 4 If so, does that mean that the policy doesn't adequately capture the *metapolicy* or is the metapolicy incomplete?

Thought Experiment #3: MLS

Multi-level Security (MLS) is the following problem: given information at various sensitivity levels and individuals having various degrees of trustworthiness, how do you control access to information within the system.

It's also called *military security* because it models the confidentiality problem with military documents (even before computers).



The initial formalization was developed in 1973 by David Bell and Len LaPadula and is called the *Bell and LaPadula model (BLP)*. It is one of the most influential efforts in the history of computer security.

Aside: Models of Military Systems

In military systems, four models of operation are often defined for computers handling classified information:

- dedicated:** all users cleared for all information on machine; no need for access control (MILS);
- system-high:** all users cleared, but must obey need-to-know compartments (discretionary access control).
- compartmented:** all users cleared, but must be need-to-know compartments (mandatory access control). System must handle requests across classifications.
- multi-level:** not all users cleared for all information; system enforces access control (MLS).

MLS is the most difficult so not widely deployed.

(RAND Report R-609-1, "Security Controls for Computer Systems," (1970) summarizes best practices.)

Multi-Level Security (MLS)

Let's make our MLS thought experiment concrete.

Setting: General Eisenhower's office in 1943 Europe.

The problem: Assume an environment in which there are various pieces of *information* at different sensitivity levels: the war plan, the defense budget, the base softball schedule, the general's laundry list, etc.

Also, there are a variety of *individuals* with access to selected pieces of information: Eisenhower, Patton, privates, colonels, secretaries, janitors, spies, etc.



Multi-Level Security (MLS)

The goal: Understand what “security” might mean in this context and define some rules to implement it.



What are we protecting? Against what threats?

Notice: it's very important that we're only considering confidentiality in this thought experiment, not integrity or availability. Someone bombing the office and destroying the war plan might be a significant threat, but *it's not a threat to confidentiality*.

Confidentiality Questions

Recall the questions we asked about ensuring confidentiality:

- ① How do you group and categorize information?
- ② How do you characterize who is authorized to see what?
- ③ How are the permissions administered and checked?
According to what rules?
- ④ How can authorizations change over time?
- ⑤ How do you control the flow of “permissions” in the system?
Can I authorize others to view data that I am authorized to view?

For simplicity, let's assume an environment of *static* permissions. That means we'll ignore questions 4 and 5. Let's see if we can figure out some possible answers *for this specific setting* to the other questions.

Dividing Information

Back to our thought experiment: Gen. Eisenhower's office in 1943. The relevant “space” of information contains lots of individual atoms or factoids:

- ① The base softball team has a game tomorrow at 3pm.
- ② The Normandy invasion is scheduled for June 6.
- ③ The cafeteria is serving chopped beef on toast today.
- ④ Col. Jones just got a raise.
- ⑤ Col. Smith didn't get a raise.
- ⑥ The British have broken the German Enigma codes.
- ⑦ ...



Dividing Information (2)

Note that not all information is created equal. How do we group and categorize information rationally (in order to protect what needs protecting)?

Dividing Information (2)

Note that not all information is created equal. How do we group and categorize information rationally (in order to protect what needs protecting)?

Idea: Information is separated into documents/folders/objects/files. How should that be done?

Documents (objects) are *labeled* according to an assessment of their sensitivity level. *We'll assume a certain form for labels; they might be done differently.*



Object Sensitivity Levels

One part of the label is taken from a linearly ordered set. One common scheme has levels: **Unclassified**, **Confidential**, **Secret**, **Top Secret**.

There are also “need-to-know” *categories*, from an unordered set, expressing membership within one or more interest groups, e.g., **Crypto**, **Nuclear**, **Janitorial**, **Softball**, etc.



Some labels are special, but can be treated as need-to-know categories, e.g., **FOUO**, **No Foreign**, **Eyes Only**.

Object Sensitivity Levels

Ideally, the label on any document reflects the sensitivity of the information contained in that document. The label contains both a hierarchical component *and* a set of categories.

For example, two documents might have levels:

(Secret: {Nuclear}),
(Top Secret: {Crypto}).

One can expect that the first contains rather sensitive information related to the category Nuclear. This second contains highly sensitive information in category Crypto.

Some entity/agency/officer makes these labeling decisions. How they are made is outside the scope of our concern.

Object Sensitivity Levels

How should you label a document that contains “mixed information”? Remember the point of the labels!

- Suppose the document contains both sensitive and non-sensitive information?
- Suppose it contains information relating to both the Crypto and Nuclear domains?



Sometimes a decision is made that a document classification should be changed. This is called *downgrading* (or *upgrading*). We'll ignore that possibility for now.

Authorization Levels



Individuals (subjects) have *clearances* or *authorization levels* that are typically of the same form as document sensitivity levels.

That is, each individual has:

- a hierarchical security level indicating the degree of trustworthiness to which he or she has been vetted;
- a *set* of “need-to-know categories” indicating groups to which he or she belongs or areas of interest in which he or she is authorized to operate.

Example: an individual at level (**Top Secret, {Crypto, Nuclear}**) is highly trusted *and* authorized to view information in domains Crypto and Nuclear.

Authorization Levels

The “lowest” security level in the system is called *system low*. For our MLS-type system it is (**Unclassified: { }**).

Higher clearances are assigned by some organization or government entity according to their assessment of the individual's trustworthiness and need for the information.

The “highest” (most permissive) level in the system, if it exists, is called *system high*. What would be system high for our MLS system?

Some levels may be unpopulated, i.e., no individual is cleared at that level.

Least Privilege: An Aside

The need-to-know categories are a reflection that even within a given security level (such as **Top Secret**) there is plenty of information to which not everyone cleared to that level should have access. This is an instance of:

The Principle of Least Privilege:

Any subject should have access to the *minimum* amount of information needed to do its job.

This is as close to an axiom as anything in security. *Why does it make sense?*



Now What?

Given that we have labels for objects and clearances for subjects, how do we decide which subjects are permitted access to which objects?

Surely it's some relationship between the subject level and the object level. But what?

For example, should a subject with clearance **(Secret: {Crypto})** be able to read a document labeled **(Confidential: {Crypto})**?

Should a subject with clearance **(Top Secret: {Crypto, Nuclear})** be able to modify a document labeled **(Confidential: {Crypto})**?

The Dominates Relation

Given a set of security labels (L, S) , comprising hierarchical levels and categories, we can define a *partial order* among them.

Definition: (L_1, S_1) *dominates* (L_2, S_2) iff

- ① $L_1 \geq L_2$ in the ordering on levels, and
- ② $S_2 \subseteq S_1$.

We usually write $(L_1, S_1) \geq (L_2, S_2)$.

Note that this is *not* a total order. There are security labels A and B , such that neither $A \geq B$ nor $B \geq A$.



Partial Order

A partial order is a relation that is reflexive, transitive, and antisymmetric.

Reflexive: $x \geq x$

Transitive: $[x \geq y \wedge y \geq z] \rightarrow x \geq z$

Anti-symmetric: $[x \geq y \wedge y \geq x] \rightarrow x = y$

It's easy to prove that dominates is a partial order.

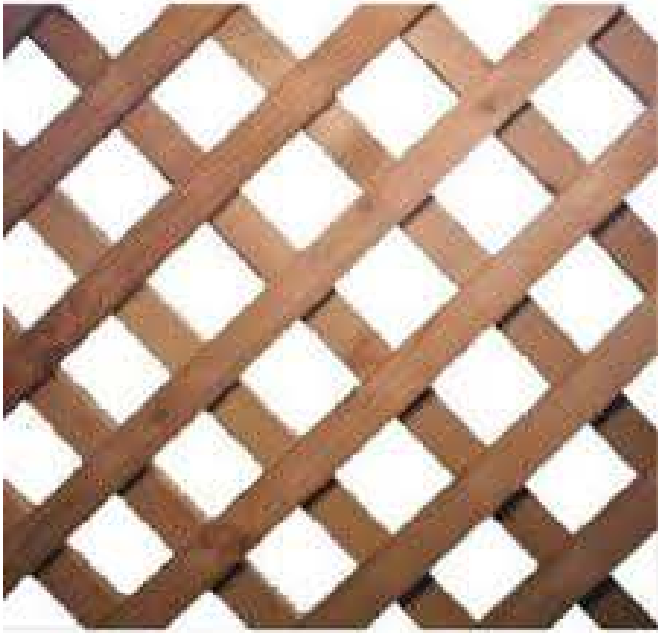
Aside: Lattices

Algebraically, the (full) set of labels with their ordering would form a *lattice*. This is sometimes called “lattice-based security.”

Below is the lattice for hierarchical levels {**Secret**, **Top Secret**} and need to know categories {**Crypto**, **Nuclear**}.



Aside: Lattices



In mathematics, a lattice is a partially ordered set (or poset), in which all nonempty finite subsets have both a supremum (join or lub) and an infimum (meet or glb). Lattices can also be characterized as algebraic structures that satisfy certain identities.

Exercise: Suppose you have two hierarchical levels H and L such that $L < H$, and two categories A and B . Using the dominates relation as the partial order, draw the lattice of levels in this system. *How many levels are possible?*

Some Answers?

Given our mechanisms for classifying objects (data / files) according to security labels, and personnel according to clearances, what are the answers to these questions?

How do you group and categorize information? The grouping is done as documents (files) and categorized according to labels.

What does that mean? Who assigns the labels? What about documents that contain “mixed” information?

Some Answers: Continued

Suppose we've given labels to our files (objects) and clearances to our individuals (subjects). **How do you characterize who is authorized to see what?**

The answer seems to be a relationship between the sensitivity level of a document (file) and the authorization level of the individual.

- What is the appropriate relationship?
- How do we codify it as a system of rules for access within this system?
- Does permission depend on the type of access requested? For example, are read and write access interchangeable?

Secure Reading

Suppose subject S with authorization (L_S, C_S) asks to *read* an object O with classification (L_O, C_O) . Under what conditions should the request be granted by the system?



For example, suppose Lisa has clearance **(Secret: {Crypto})**. Which of the following should she be able to read?

- document labeled **(Confidential: {Crypto})**
- document labeled **(Top Secret: {Crypto})**
- document labeled **(Secret: {Nuclear})**
- document labeled **(Secret: {Crypto, Nuclear})**

So what is the formal rule?

Simple-Security Property

According to the Bell-LaPadula Model (BLP) of security, one of the earliest formal security policies, the first formal rule governing access is:

The Simple-Security Property: *Subject S with clearance (L_S, C_S) may be granted **read** access to object O with classification (L_O, C_O) only if (L_S, C_S) dominates (L_O, C_O) .*

We will often write “ (L_S, C_S) dominates (L_O, C_O) ” as “ $(L_S, C_S) \geq (L_O, C_O)$,” but recall that it involves both hierarchical levels and need-to-know categories.

Simple-Security Property

The Simple Security Property models read access in the world of military documents *and* attempts to codify it for the world of electronic information storage.

The Simple-Security Property: Subject S with clearance (L_S, C_S) may be granted *read* access to object O with classification (L_O, C_O) only if $(L_S, C_S) \geq (L_O, C_O)$.

- Why is it “only if” and not “if and only if”?
- Does this work in an electronic context?
- Is it all that is needed? Why or why not?

The Simple-Security property codifies restrictions on *read* access to documents. What about *write* access?

Is the problem different with respect to writing in the electronic context than it is in the world of military documents? Why or why not?

More generally, what assumptions can be made about *persons* in the world of military paper documents that cannot be made about *subjects* (processes) in the context of computers?

Secure Writing

Subjects in the world of military documents are assumed to be *persons* trusted not to disclose (write) to unauthorized parties information to which they have legitimate access.

Subjects in the world of computing are often *programs* operating on behalf of a trusted user (and with his or her clearance).



The program may have embedded malicious logic (a “trojan horse”) that causes it to collude with other users or programs to “leak” information without the knowledge or consent of the authorized user.

For that reason, it is necessary to place mandatory controls on the write accesses of subjects that might not be necessary for persons. This is sometimes called the *confinement problem*.

What is the appropriate restriction on writing? What do we want to prevent?

The *-Property

In the Bell-LaPadula Model of security, the following rule is enforced to restrict *write* access:

The *-Property: *Subject S with clearance (L_S, C_S) may be granted *write* access to object O with classification (L_O, C_O) only if $(L_S, C_S) \leq (L_O, C_O)$.*

This is pronounced “the star property.”

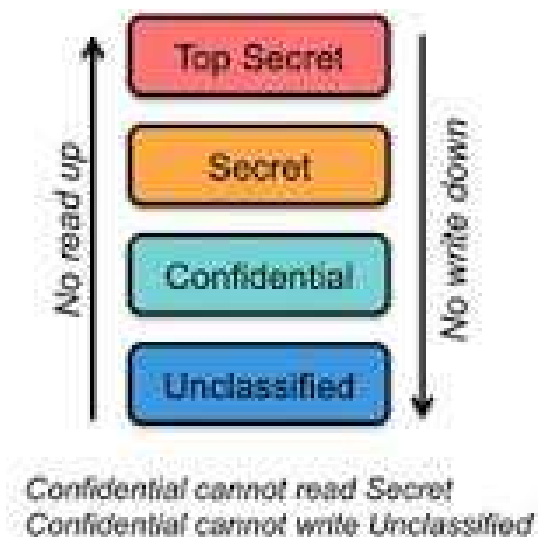
The *-Property

Does this rule make sense? Is it too restrictive? Is it too lax?

According to the *-property, can a commanding general with a top secret clearance email marching orders to a foot soldier? No!

According to the *-property, can a corporal with no clearance overwrite the war plan? Yes, *but that's an integrity problem!*

The simple-security and *-property are sometimes characterized as “read down” and “write up,” respectively. Alternatively, they’re characterized as “no read up” and “no write down.”



Aside: MAC vs. DAC

Security systems should distinguish between:

Mandatory Access Controls (MAC): security rules that are enforced on every attempted access and not at the discretion of any system user;

Discretionary Access Controls (DAC): security rules that are enforced by the system at the discretion and behest of some users.

Example: the Unix file protection system implements DAC since the protections can be modified by the file owner.

For MLS, we'll focus on mandatory controls. (Note that the acronym "MAC" is used for several different notions in computing and security, so don't get confused.)

Trusted Subjects

Often, to get around the more onerous restrictions of a mandatory policy, an implementation may add *trusted subjects*, specialized subjects permitted to operate “outside the rules of the policy” in very constrained ways.

Example: The *-property implies that the general can never send an email to the private. We add a special *downgrader* subject to the system and extend the *-property with the proviso that an object’s level can be reduced in specific ways only if the object’s contents are reviewed by the downgrader subject *including visual inspection by a trained human being*.

Notice: this technically violates the naive *-property, but prevents any malicious *program* from leaking information (unless it is clever enough to fool the downgrader). See Steganography.

Notice that Simple Security and the *-Property control two ways in which information can flow from A to B.

- ① B can “pull” information from A by reading objects in A’s space. Simple Security is designed to constrain that type of information flow.
- ② A can “push” information to B by writing objects in B’s space. The *-Property is designed to constrain that type of information flow.

Are there additional ways that information can flow from A to B that don’t involve either of those mechanisms.

Our discussion of the Bell and LaPadula model explicitly included **Read** and **Write** access, but not **Create**, **Destroy**, **Execute**, **Append**, others. How might we add these operations to our BLP framework?

In particular, is **Execute** effectively a modify (write) operation? A reference (read) operation? Neither? Both?

Our discussion of the Bell and LaPadula model explicitly included **Read** and **Write** access, but not **Create**, **Destroy**, **Execute**, **Append**, others. How might we add these operations to our BLP framework?

In particular, is **Execute** effectively a modify (write) operation? A reference (read) operation? Neither? Both?

Maybe that's the wrong way to think about execute. Maybe it *creates a subject* with the creator's permission levels. Then, aren't Simple Security and the *-Property adequate?

Basic Security Theorem

According to BLP, security is essentially defined as follows:

Definition: A system is *secure* if it always satisfies the simple security condition and the *-property.

Bell and LaPadula proved a theorem about a formalization of their model that they considered to be very important.

The Basic Security Theorem: Let Σ be a system with a secure initial state σ_0 , and let T be a set of state transitions. If every element of T preserves the simple security condition and the *-property, then every $\sigma_i, i \geq 0$, is secure.

The proof is a simple induction over i

John McLean (NRL) pointed out that the Basic Security Theorem isn't very useful, because it says what is true in the *states* of the system, but doesn't constrain *transitions* that may occur in the system.

Consider a system (System Z) in which any attempt to read a file causes all objects and subjects in the system to be downgraded to security level system-low.

The Basic Security Theorem can still be proved for this system but it is obviously insecure.

What's the Point

McLean argued that reasoning merely about *states* isn't adequate. It is also necessary to reason about *transitions*.

Bell responded that McLean had misunderstood the nature of the model. The model is only a formalism that provides a framework for reasoning about secure systems. It doesn't provide a definition of security.

The Lesson: Formal definitions and theorems don't guarantee anything unless they are validated against reality. Any *interpretation* of the formalism is as valid as any other. This controversy raised questions about the “foundations” of computer security research.

An obvious hole in the BLP model would be the ability of a subject to change its own security level or that of an object under its control. One could add either:

The Strong Tranquility Property: Subjects and objects do not change levels during the lifetime of the system.

The Weak Tranquility Property: Subjects and objects do not change levels *in a way that violates the “spirit” of the security policy.*

Is this useful? Is it overly restrictive? What if a user needs to operate at different levels during the course of the day?

The Weak Tranquility Property: Subjects and objects do not change levels *in a way that violates the “spirit” of the security policy.*

What does this mean?

- Suppose your system includes a command to *lower* the level of a subject/object. Does that violate the goals of simple security or the *-property?
- Suppose your system includes a command to *raise* the level of a subject/object. Does that violate the goals of simple security or the *-property?

The Bell and LaPadula Model in its original incarnation was somewhat more complex, but a thumbnail version is simply:

- Simple Security Property
- the *-Property
- some version of the Tranquility Property.

Are simple-security, the *-property, and the tranquility property adequate to ensure confidentiality within the system?

What about the following issues:

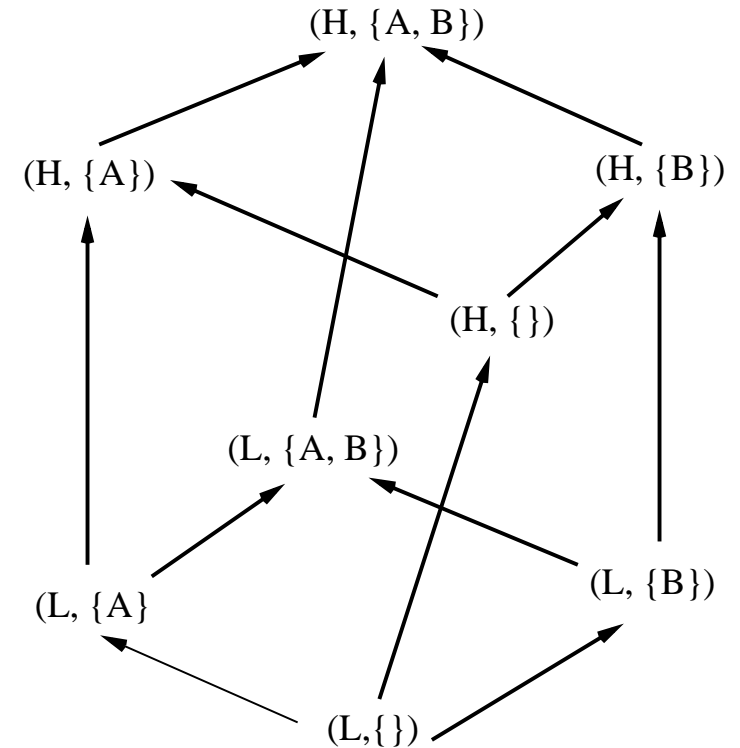
- What about other types of operations? Can every operation be thought of as a *read* or *write*? Can some be both?
- What useful operations can you imagine that might subvert the protections offered?
- Our restrictions control access by subjects to objects. Are there ways in which information might be compromised without explicit read or write operations?

A Lattice

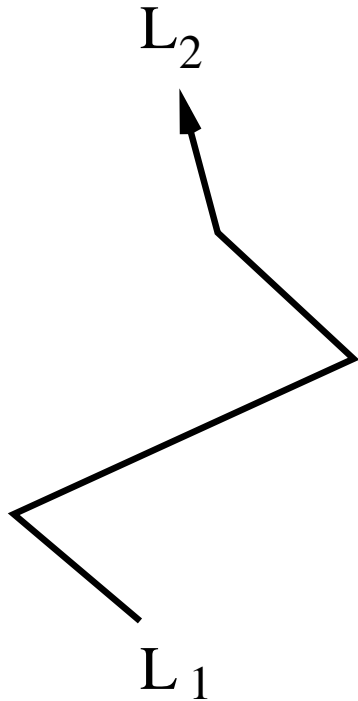
Assume a BLP system with hierarchical levels $\{H, L\}$ (with $H > L$) and categories $\{A, B\}$. On the right is a directed graph representation of the resulting lattice of labels.

The arrows represent (some of) the dominates relationships among the labels. If there is an path from L_1 to L_2 in the graph, then $L_1 \leq L_2$.

To simplify the picture, it does not include the reflexive or transitive arrows.



The BLP Metapolicy



A path in the graph from L_1 to L_2 means that “information is allowed to flow” from level L_1 to level L_2 . That can happen in either of two ways:

- 1 a subject at level L_2 can read a level L_1 object, or
- 2 a subject at level L_1 can write a level L_2 object.

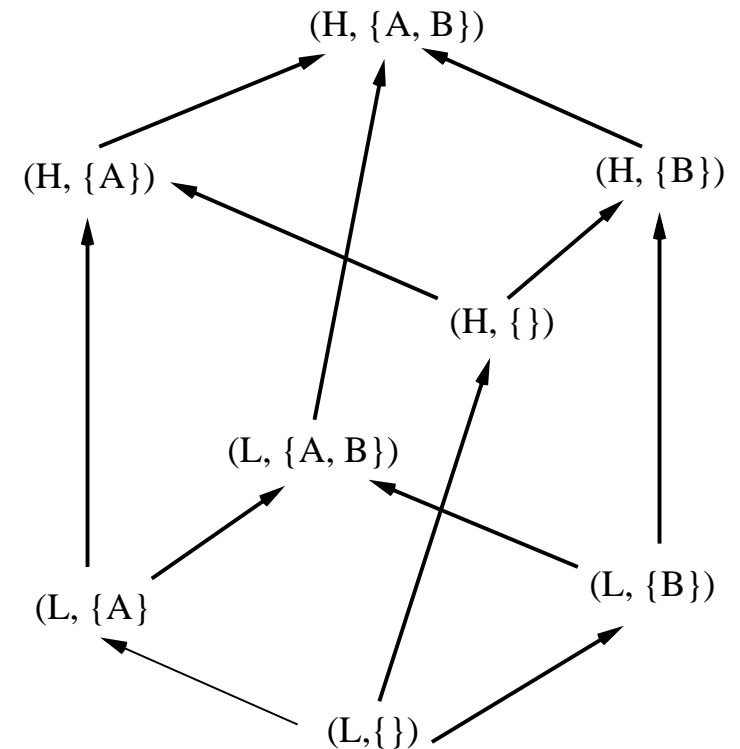
If no such path exists from L_1 to L_2 , then Simple Security should prevent 1 and the *-Property should prevent 2.

So What is the Metapolicy?

Recall that a metapolicy is the collection of *overall security goals of the system*.

So for any Bell and LaPadula system, we only want information to flow “upward” in the lattice of security levels. Equivalently, information may flow from L_1 to L_2 only if $L_2 \geq L_1$.

Any other flow indicates a violation of the security goals.



The Bottom Line

The metapolicy of any BLP system is to constrain the flow of information among the different security levels.

Recall that the metapolicy is *what we really care about* from the security standpoint.

So, if we can build a system that satisfies the BLP rules yet still violates the metapolicy, the BLP rules must not be enough!

The Bell and LaPadula Model is an example of an *Access Control Policy*. This is a popular way of conceptualizing and implementing security.

The basic idea is to introduce rules that control what accesses system *subjects* have to system *objects*.

This is an important aspect of security. The problem is that there may be information channels in the system that don't involve the access of subjects to objects.

Levels of Concern

For any secure system, we have to consider the following different areas of concern:

Policy: What is the notion of security that is being enforced by the system?

Mechanism: How is that policy enforced in the system?

Assurance: How certain can we be that the policy is enforced by the mechanisms we have put in place?

The Bell and LaPadula rules are somewhat ambiguous about whether they constitute a policy or a mechanism. They must be enforced by various different mechanisms in real systems. The level of assurance is a measure of the care and rigor with which the system is evaluated with respect to the policy. Sometimes it is difficult to judge.

An Aside: Firewalls

What is a firewall? Essentially, it's just an access control mechanism applied by structuring the system in a particular way.

Your first programming assignment involves integrating the access control checks tightly into the semantics of the operations READ and WRITE. Assume instead that your system is modeled as a server receiving commands from outside. Now suppose you applied your access control checks *at the system boundary* before accepting any command. Note that this may involve a separate processor.

Then the semantics of the individual operations could be simpler because any request would be guaranteed to be legal. *That is the basic idea of a firewall.*

Some Questions

- 1 Can our system satisfy BLP's security properties and still be intuitively non-secure?
- 2 Are there ways in which a high level subject could pass *information* to a low level subject without violating our security property?
- 3 Are there other instructions that we might naturally want to include in our system? Eg., suppose we add a CREATE instruction that allows a subject to create a new object. What constraints should be placed on their operation?
- 4 How should we handle exceptions? Eg., what should happen if ill-formed instructions are included in the instruction stream?
- 5 Is there a “stronger” security policy that we might apply?