



BASIC INTRO: COMPUTER SECURITY MODELS

HERMANN HÄRTIG, MARCUS VÖLP, 2017

CLOSELY FOLLOWING PRESENTATION IN

MB: MATT BISHOP "COMPUTER SECURITY ART AND SCIENCE" (2003)

Q3: Is there an algorithm to determine for a system with a given setting of access control permissions, whether or not a Subject A can obtain a right on Object B?

Given a System of Entities ("Objects")
acting as Subjects and/or Objects

- with clearly-defined limited access rights among themselves
- can we achieve clearly-defined Security Objectives ?

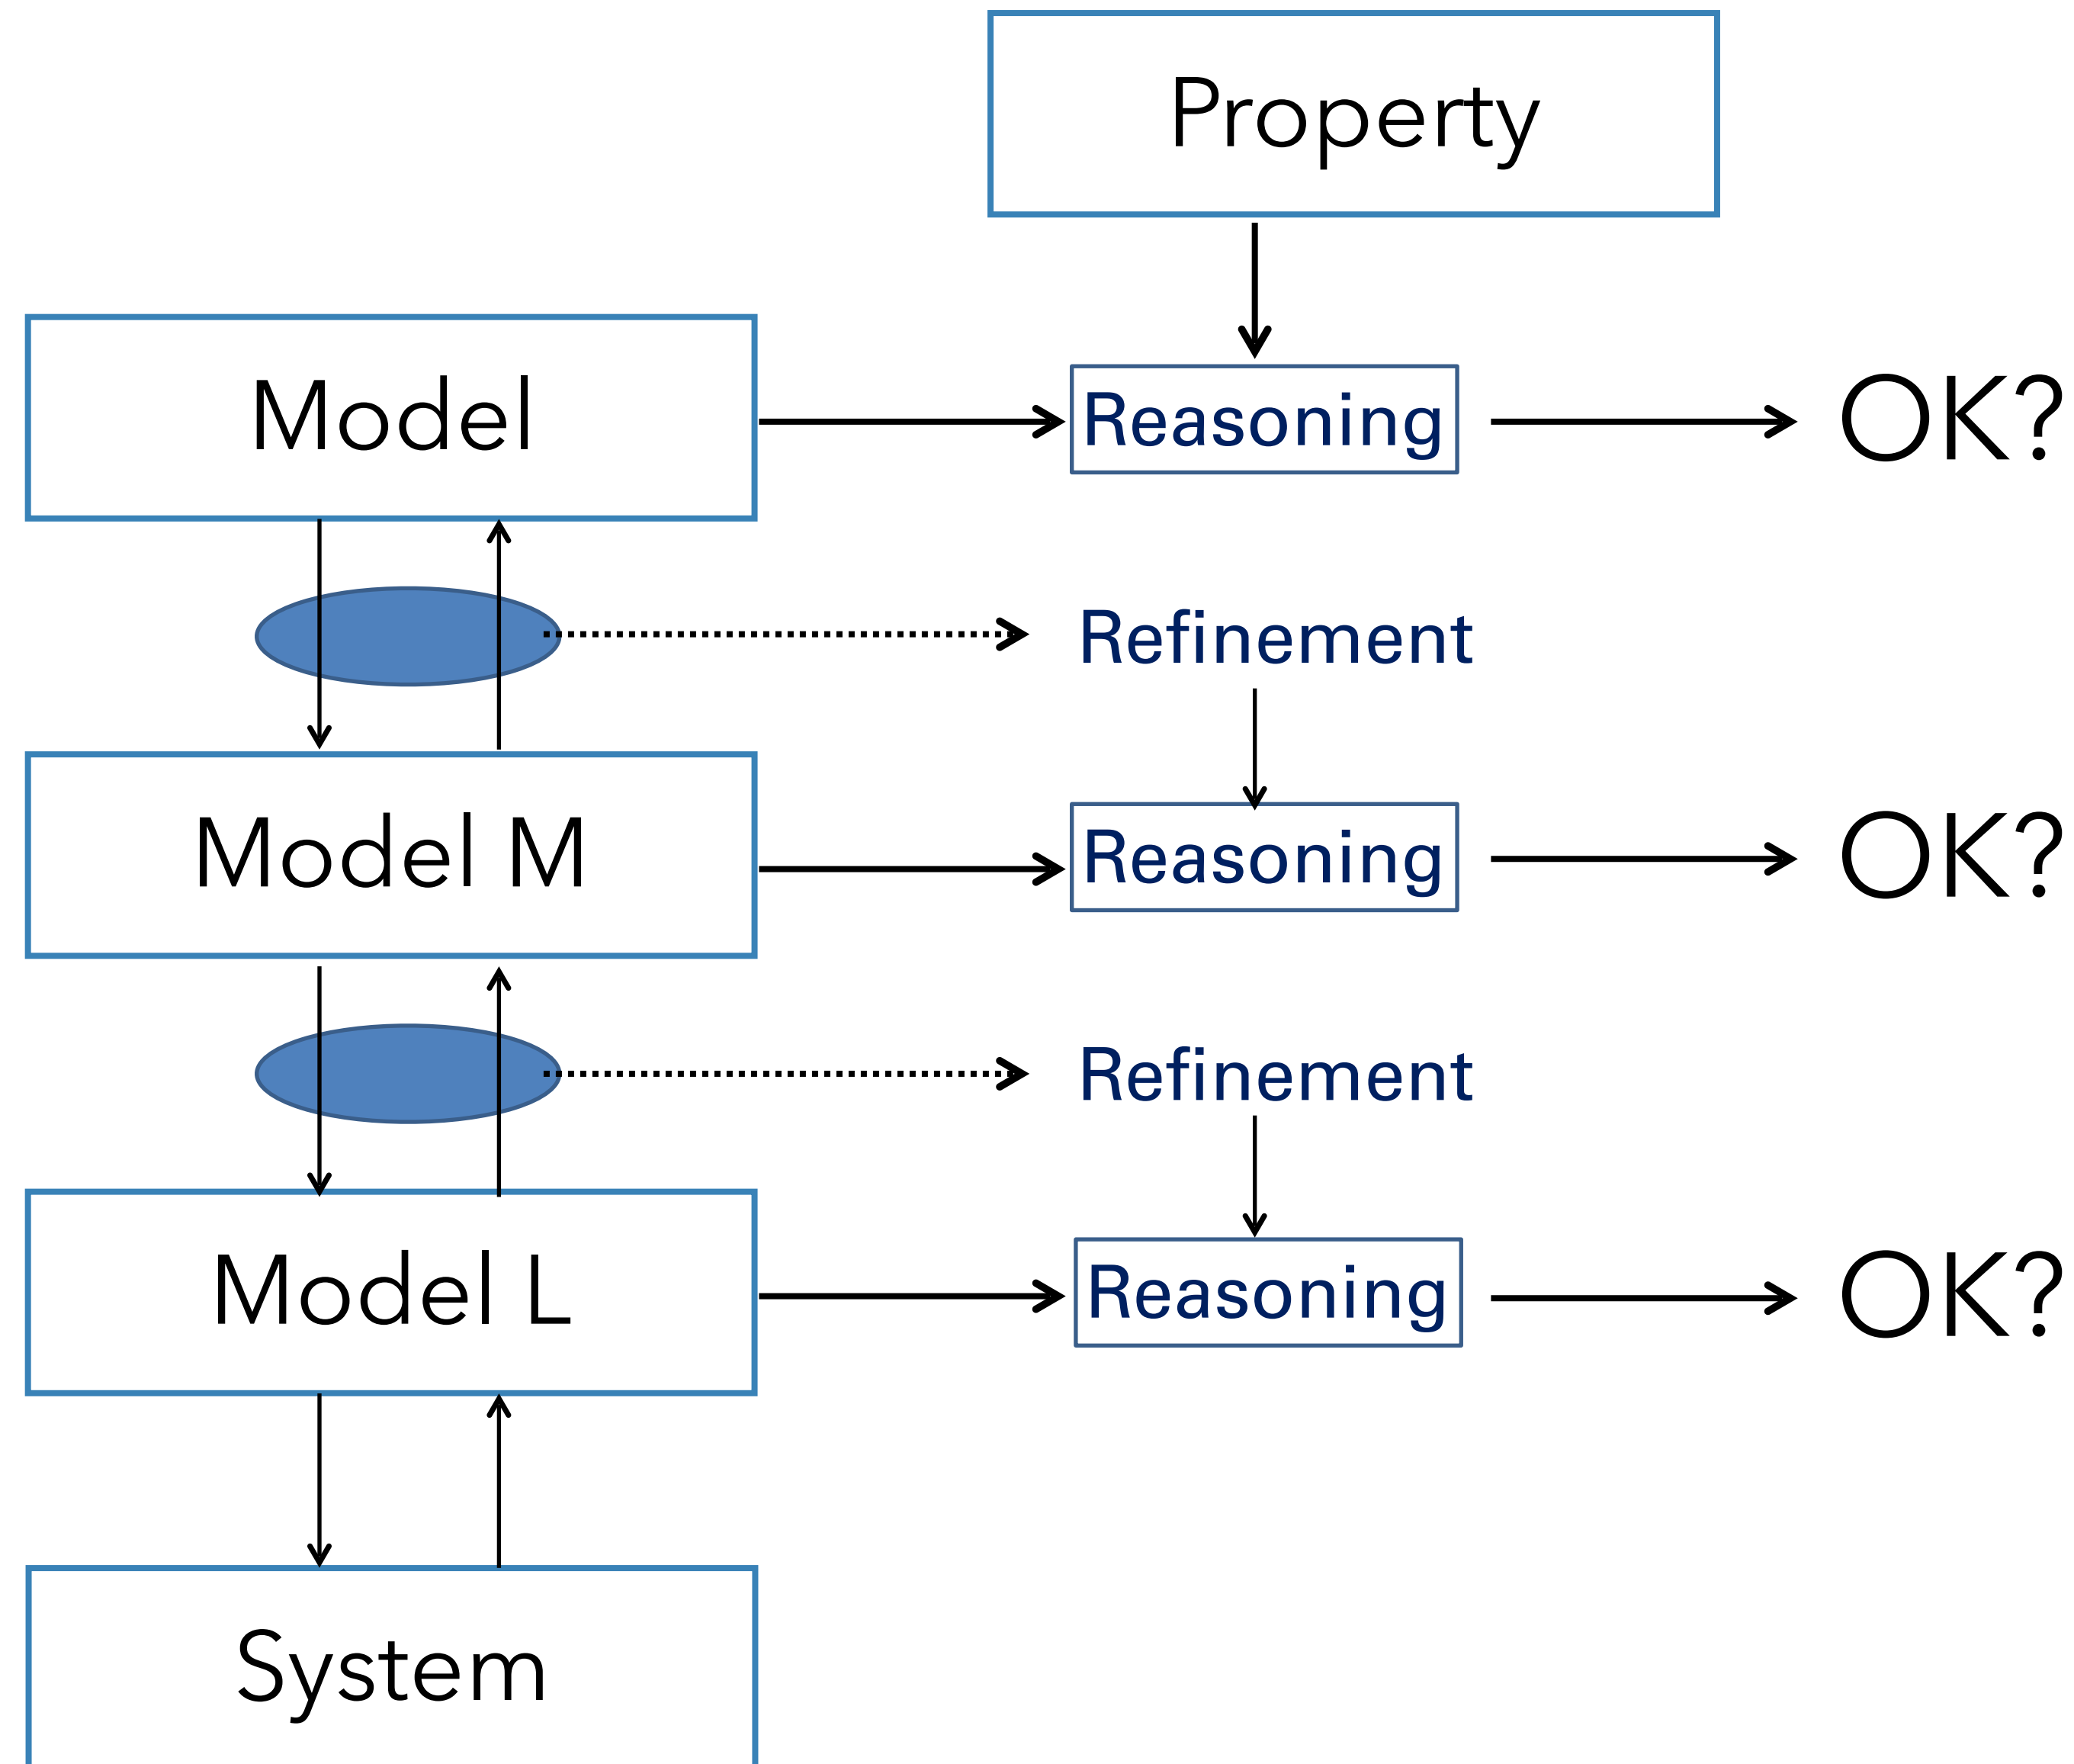
- Definition and Example of “higher-Level” Security Policies (Security Policy Models) (Bell La Padula, Chinese Wall)
- Mechanisms to express/set clearly-defined access rights: Access Control Matrix, ACL, and Capabilities
- Q3 “formalized” in 2 Models: “ACM-based” & “Take Grant”
- Decidable ?
- No proofs (in 2017)

“Reasoning”:

- Common sense
- Formal Verification
- Careful Inspection
- Mathematics ...

“Refinement”:

- Abstraction
- Implementation
- Formal Refinement

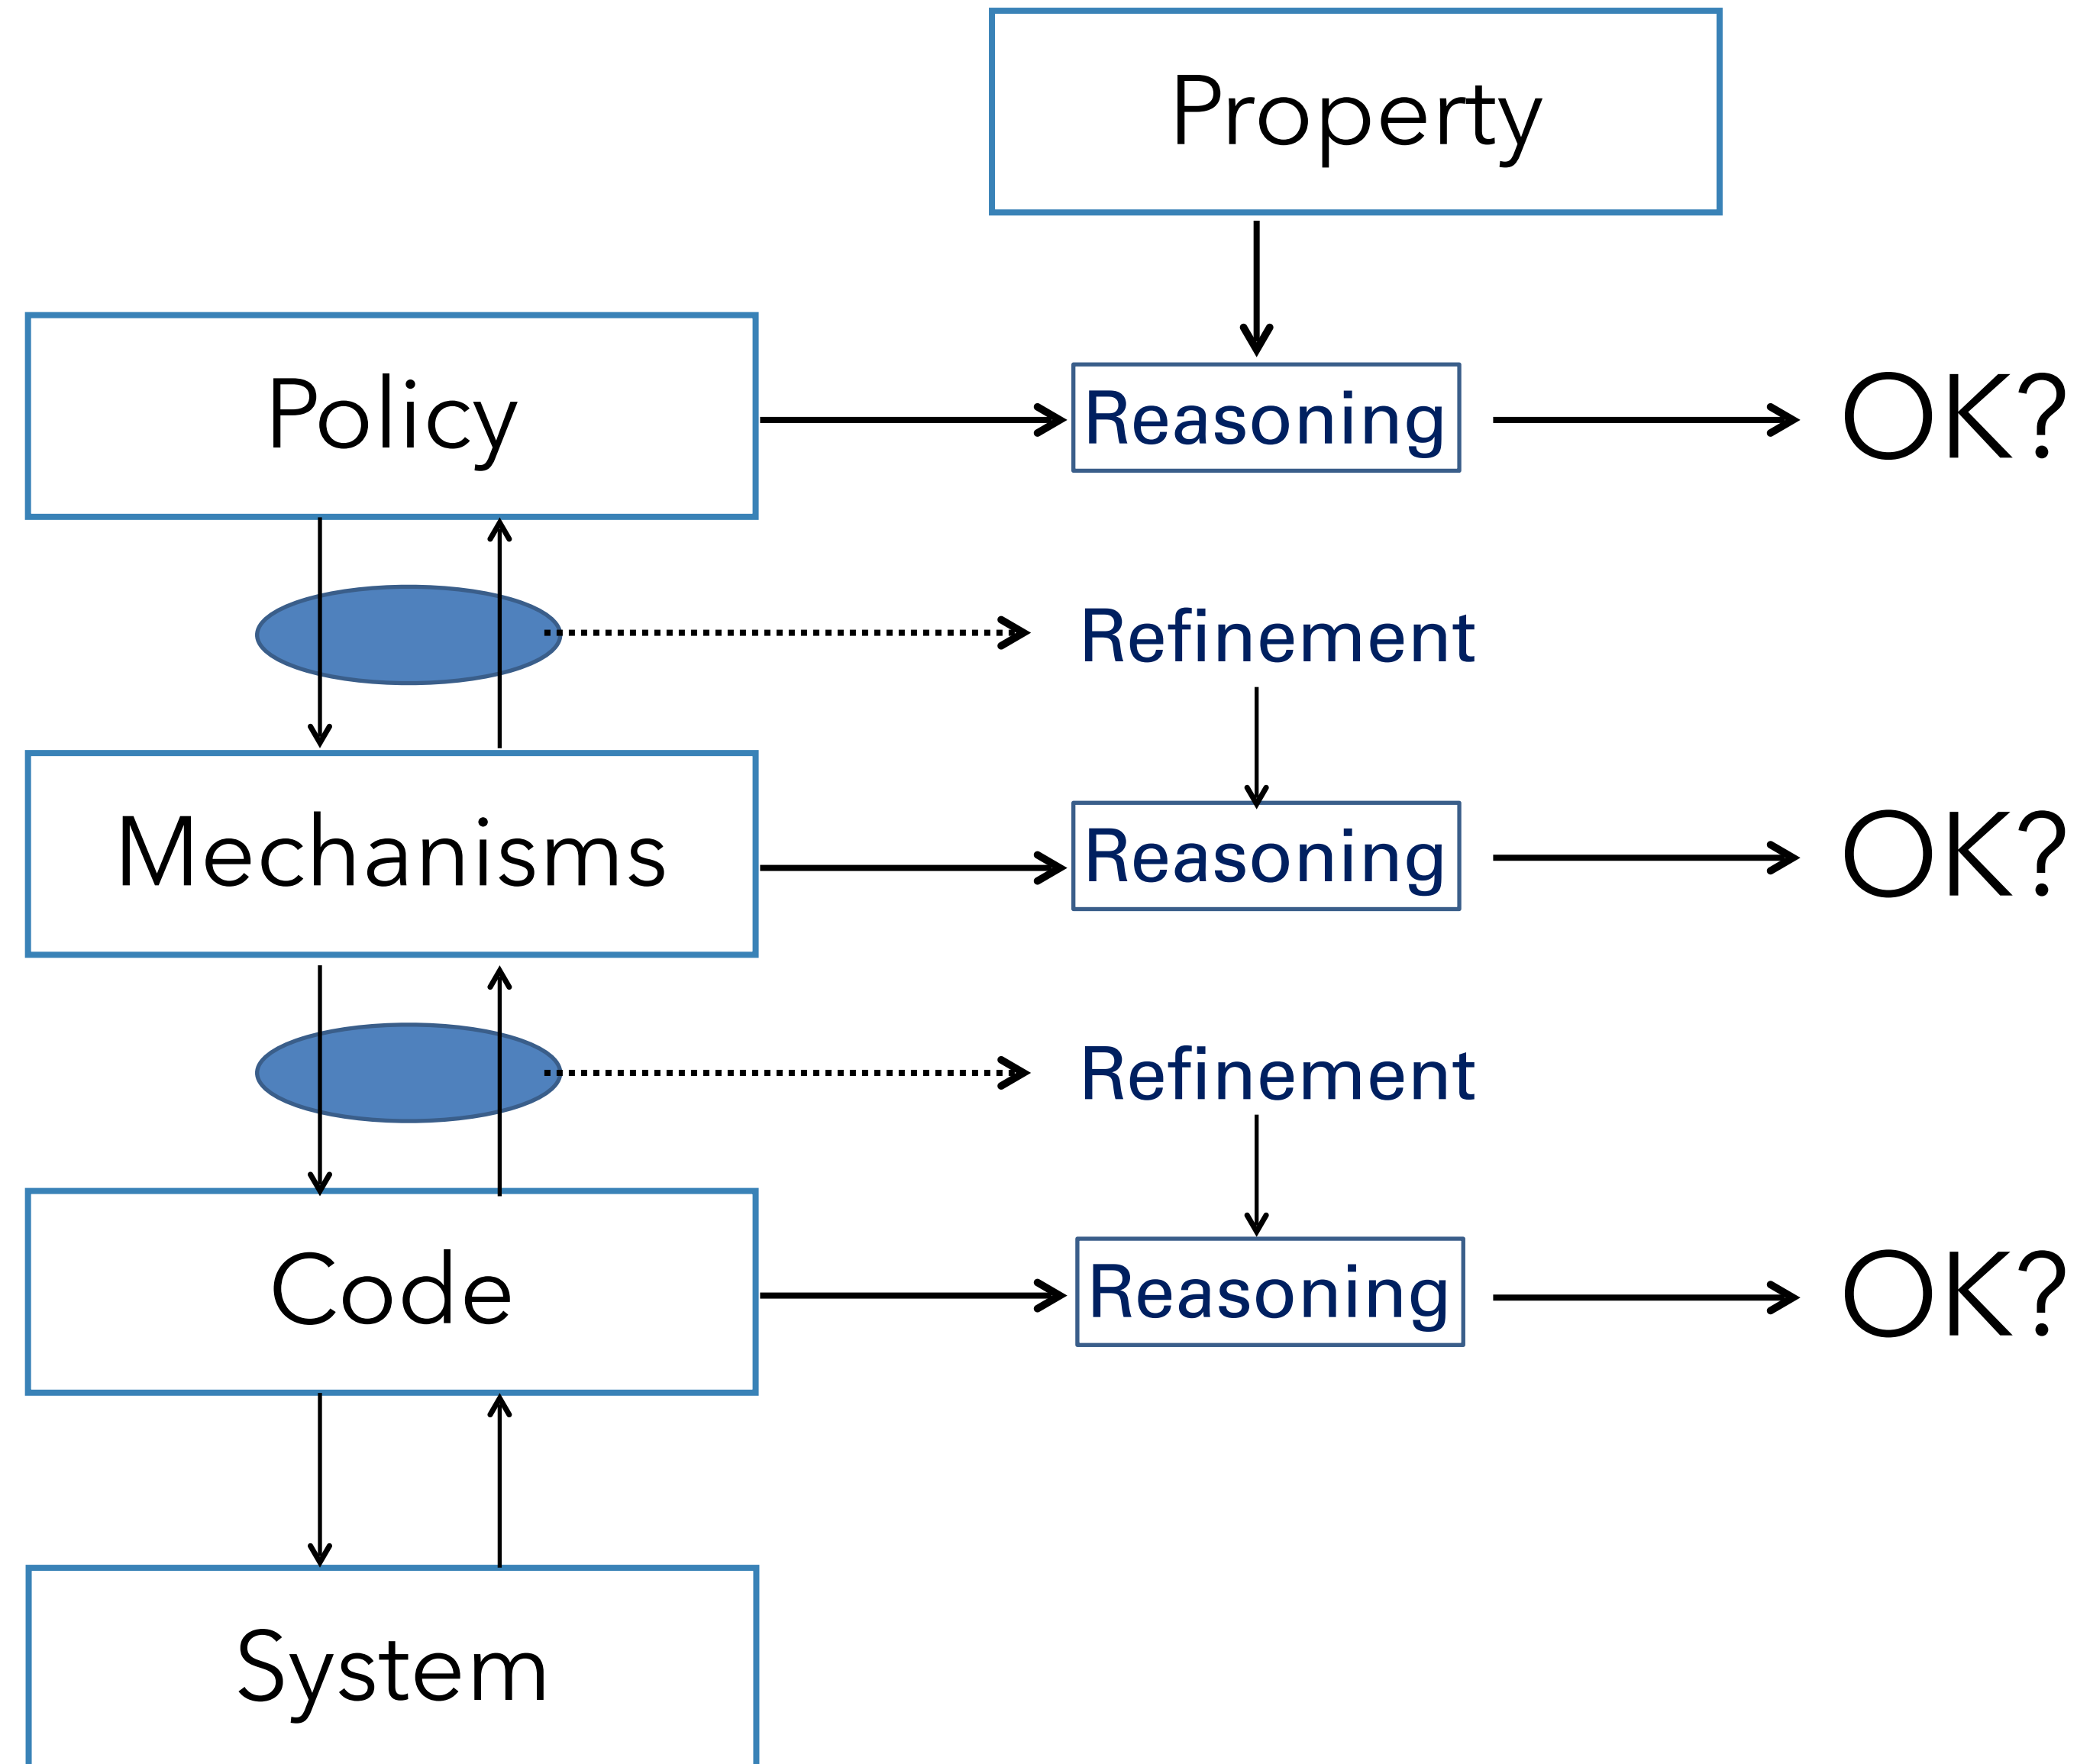


“Reasoning”:

- Common sense
- Formal Verification
- Careful Inspection
- Mathematics ...
- **“Common Criteria Assurance”**

“Refinement”:

- Abstraction
- Implementation
- Formal Refinement

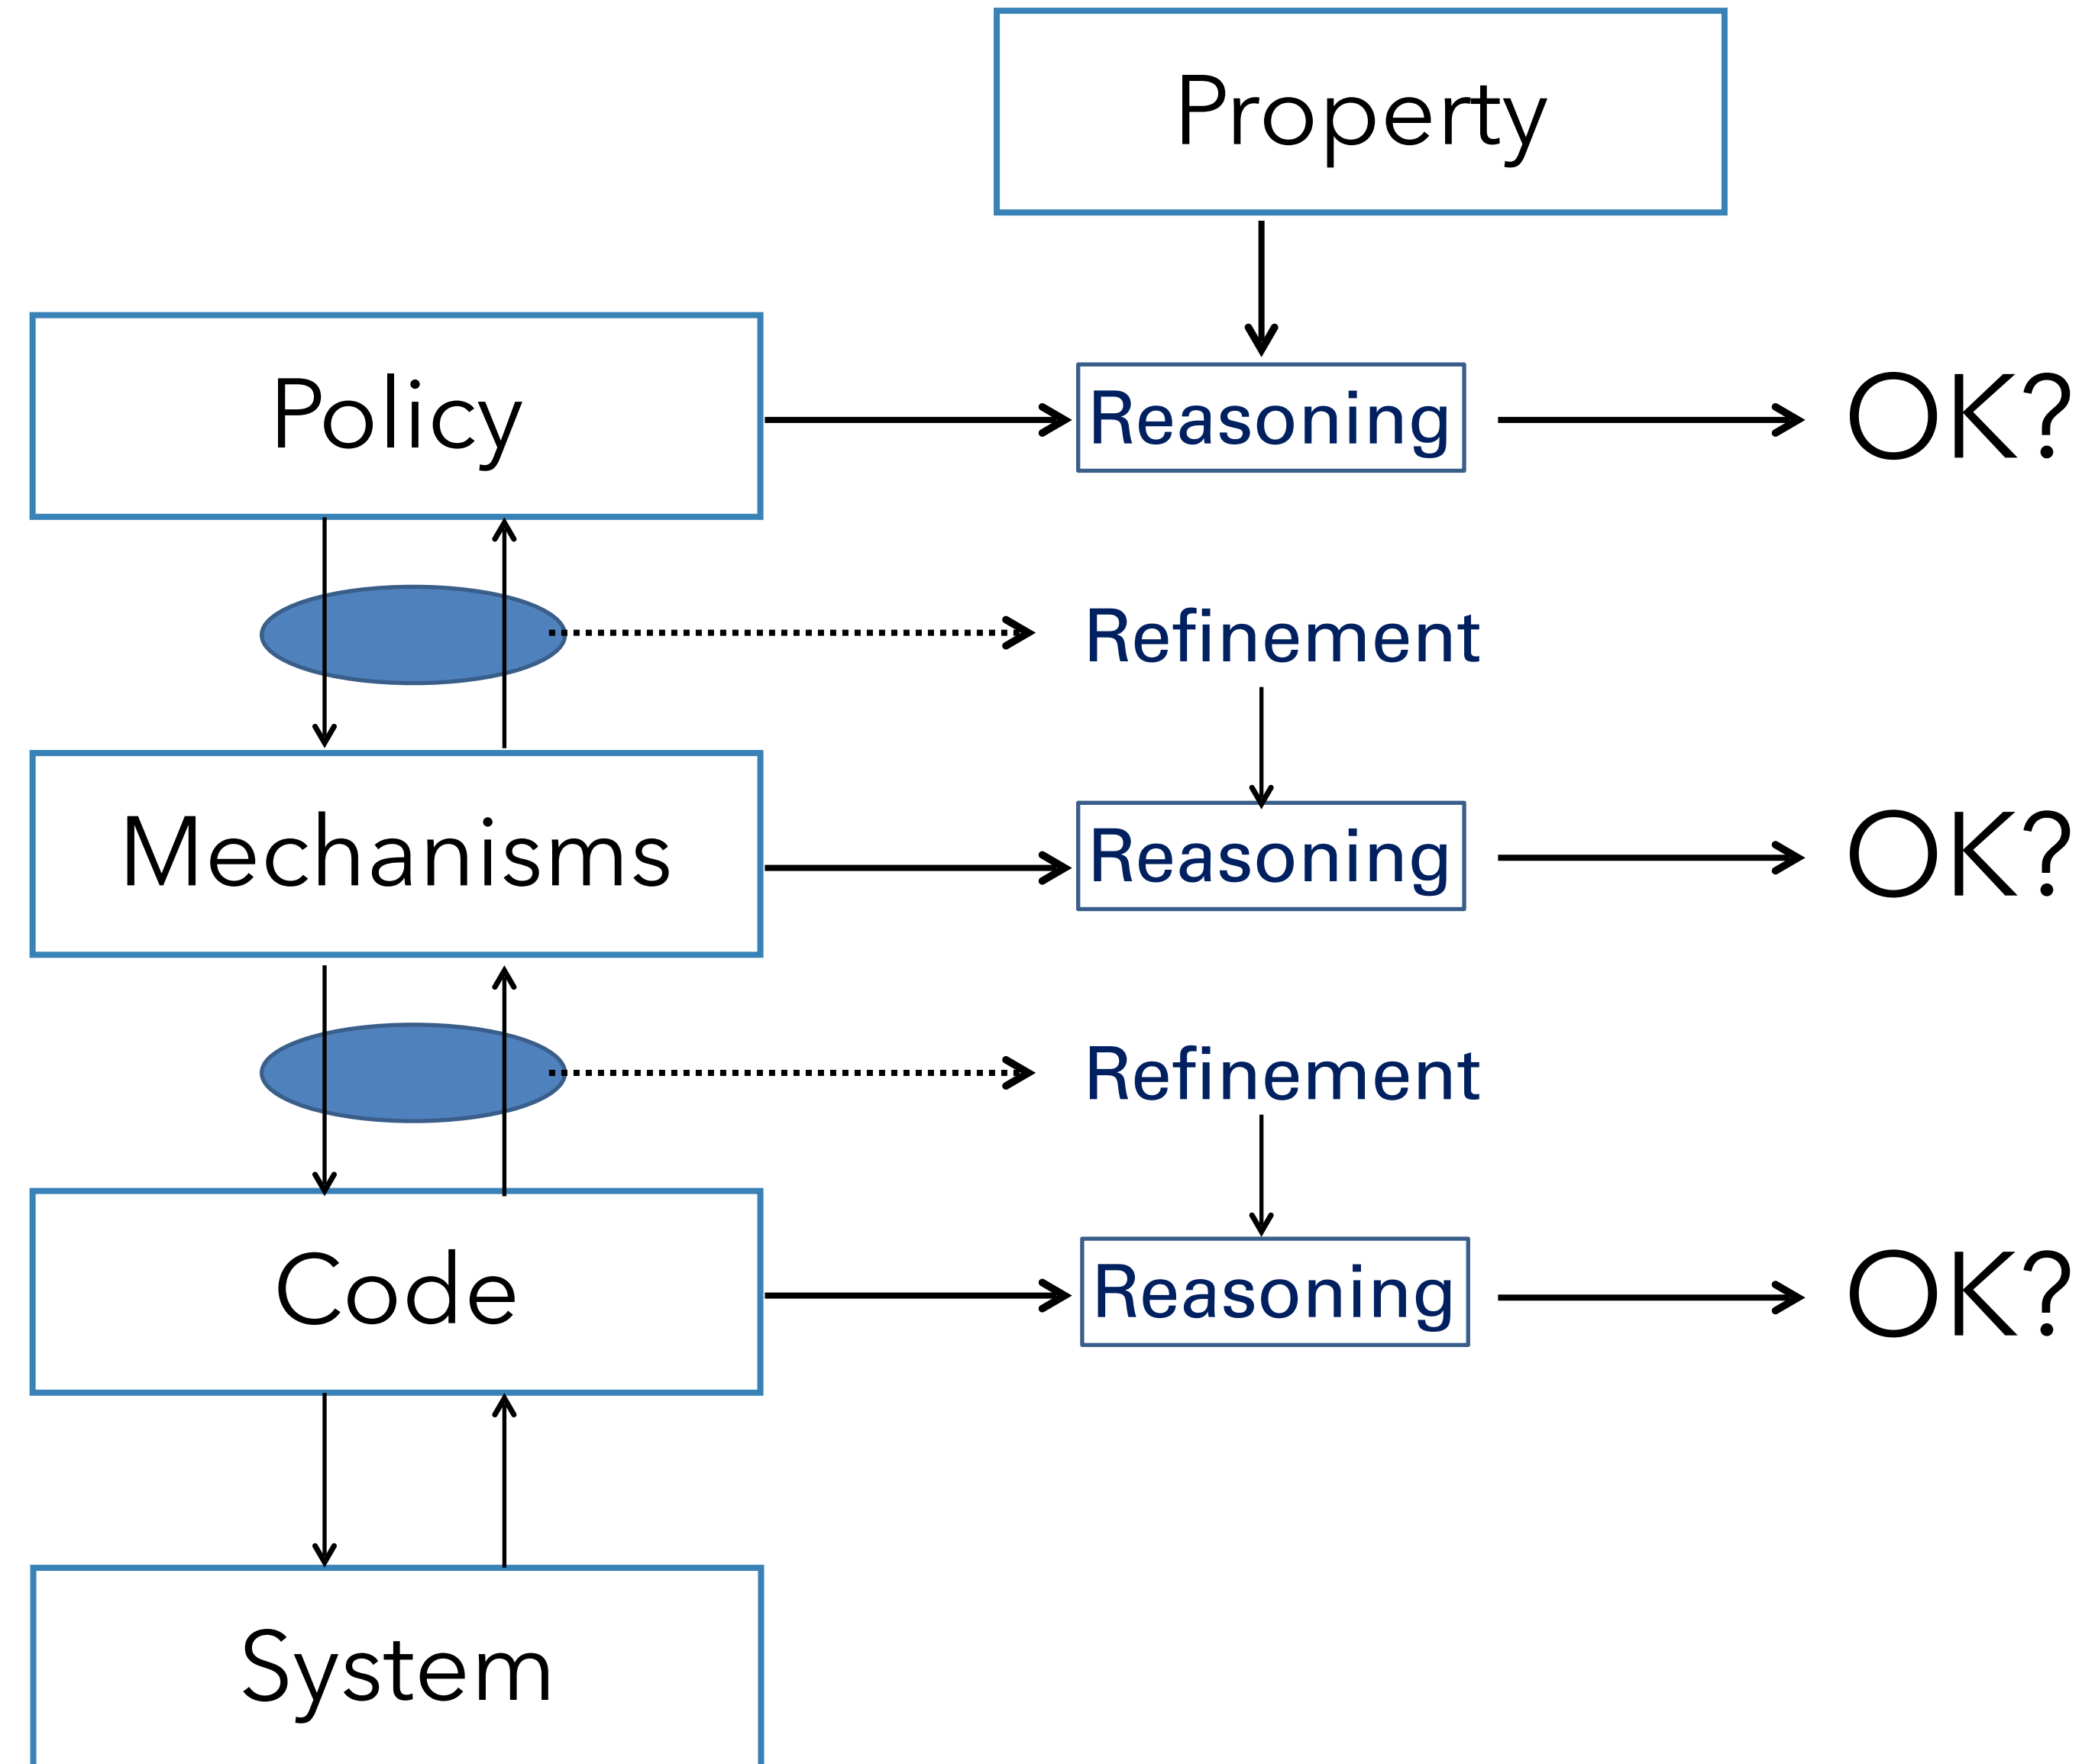


Definiton: Policy

Examples:

Higher-Level Policies
(very short):

- Bell La Padula
- Chinese Wall



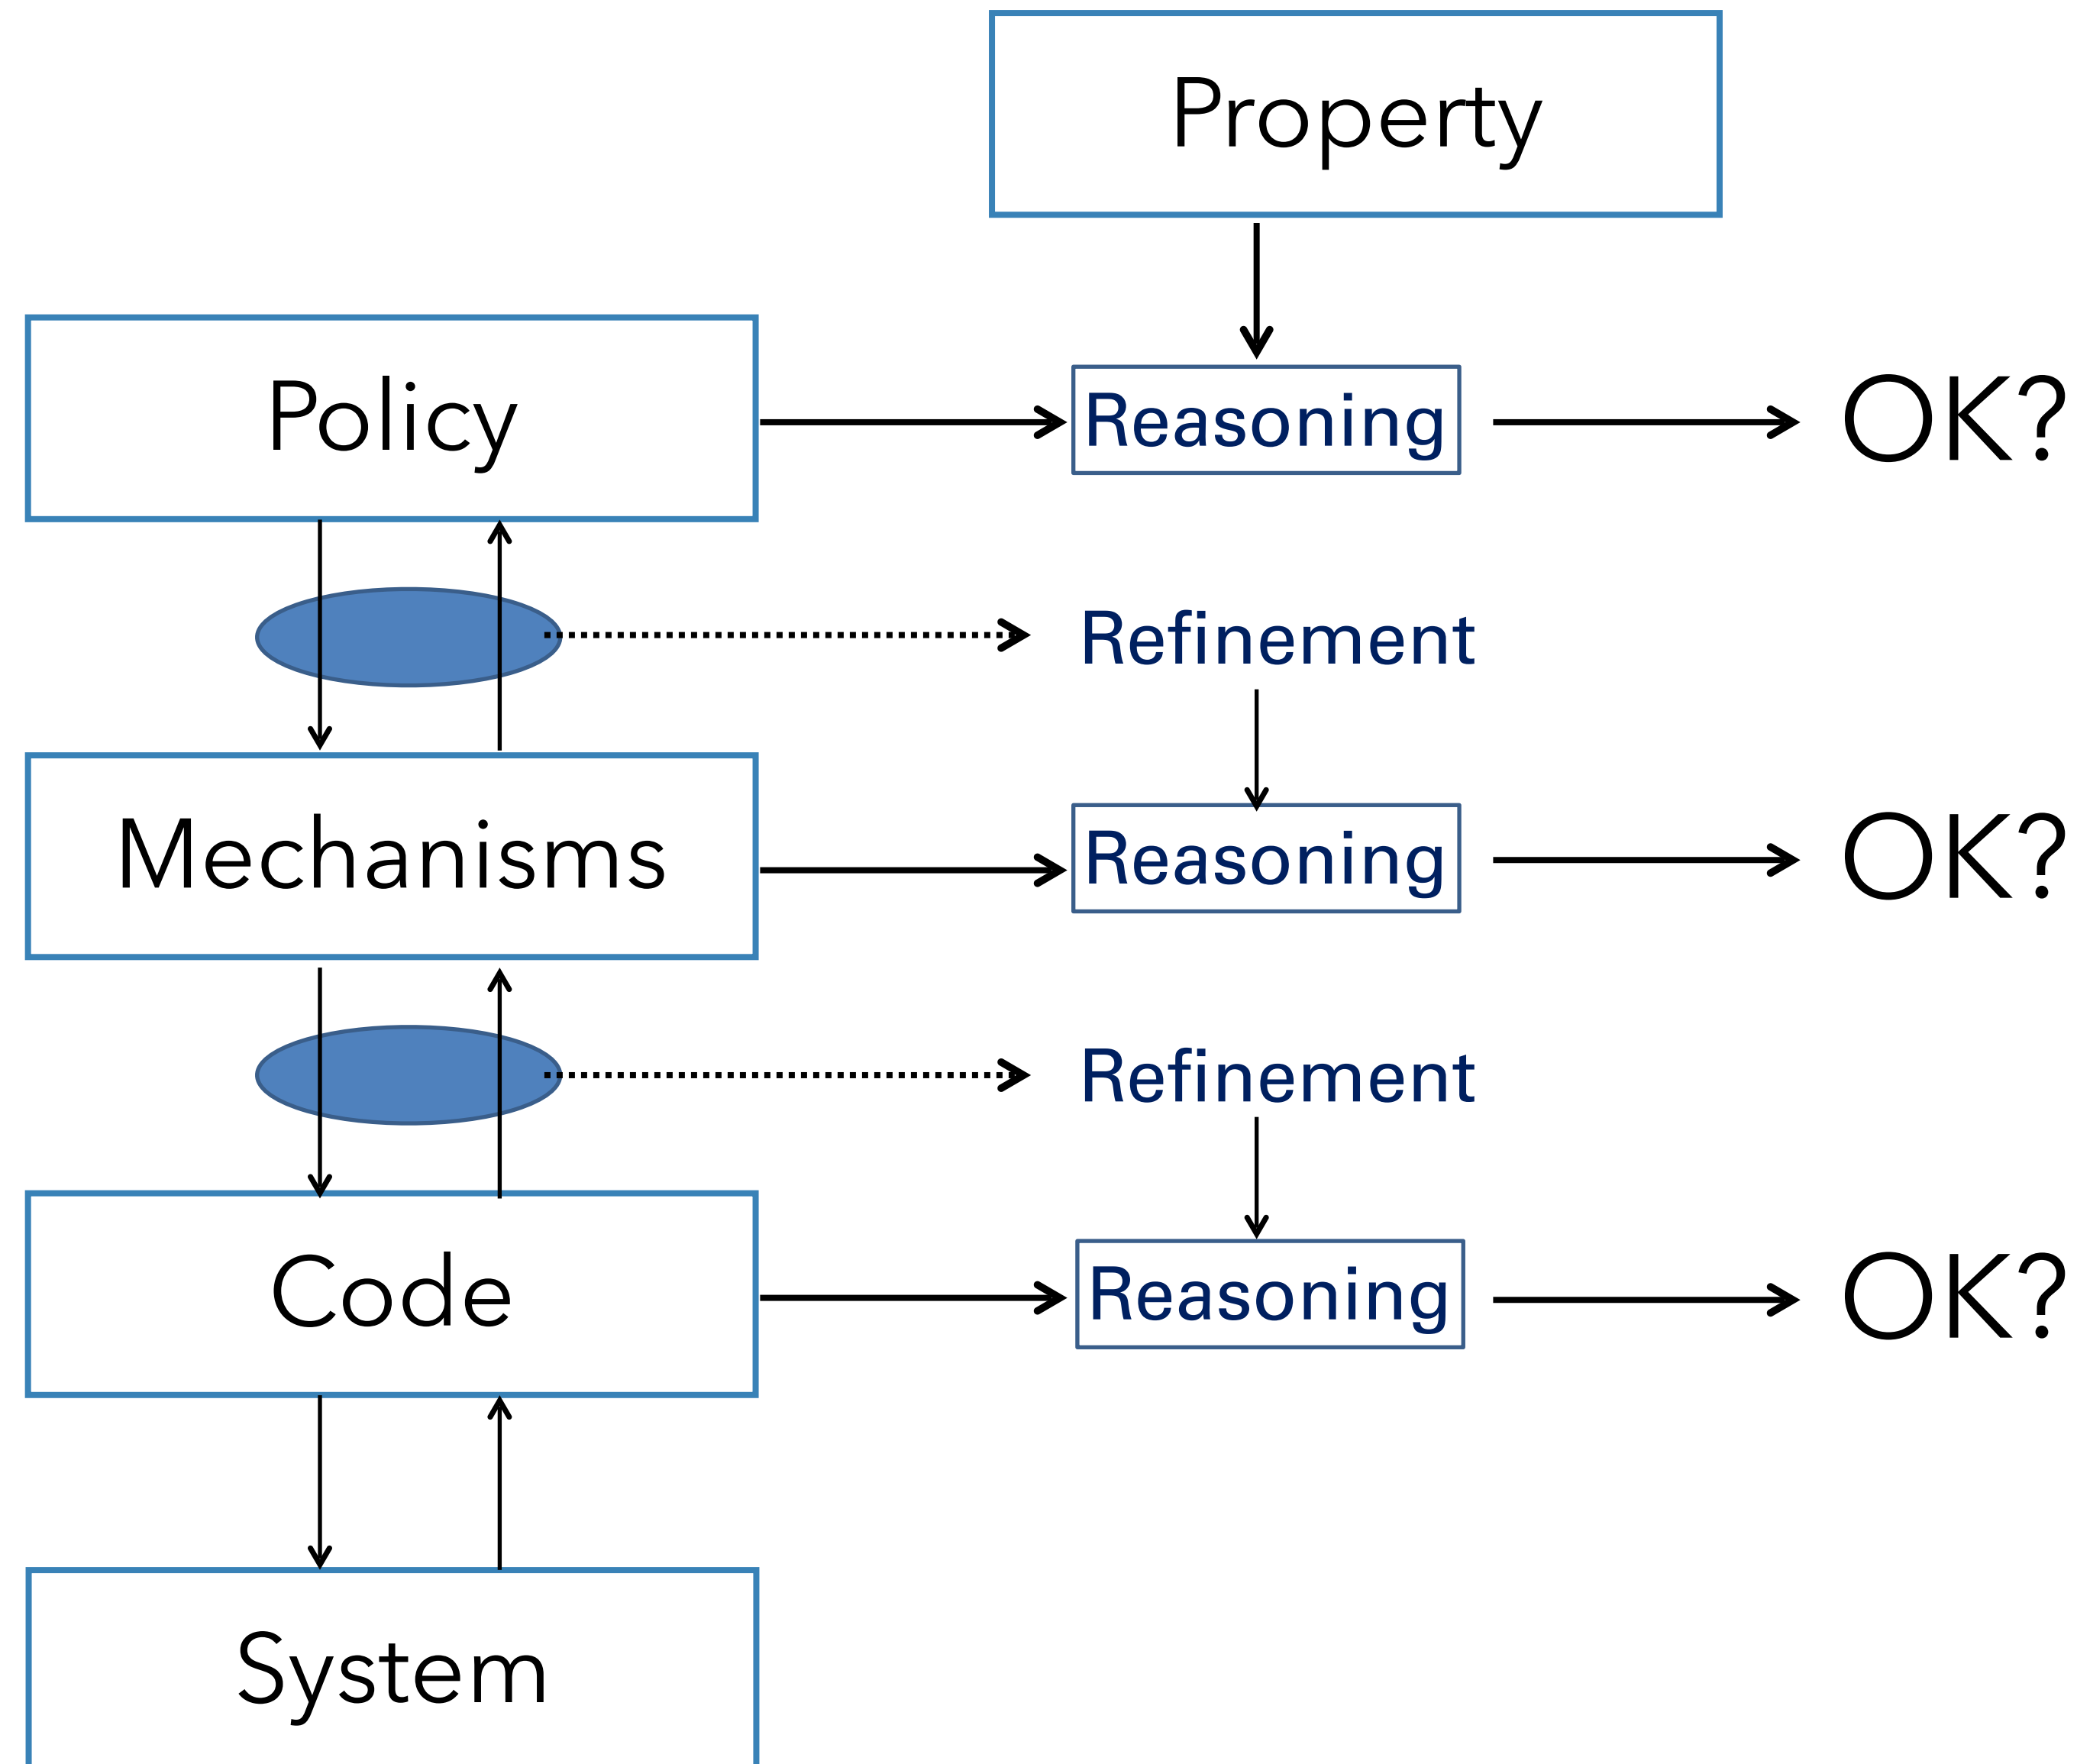
Operating Sys. Mechanisms:

- Access Control List
- Capabilities

Explain Q3 and ...
formalize per model!

Models:

- *based on Access Control Matrix*
- *"take grant" model*



Security Policy

A security policy P is a statement that partitions the states S of a system into a set of authorized (or secure) states (e.g., $\Sigma_{\text{sec}} := \{ \sigma \in \Sigma \mid P(\sigma) \}$) and a set of unauthorized (or non-secure) states.

Secure System

A secure system is a system that starts in an authorized state and that cannot enter an unauthorized state (i.e., $\Sigma_{\text{reachable}} \subseteq \Sigma_{\text{sec}}$)

Reference: Matt Bishop: Computer Security Art and Science

Definitions:

Information or data I is **confidential**

- with respect to a set of entities X if no member of X can obtain information about I .

Information I or data is **integer** if (2 definitions in text books)

- (1) it is current, correct and complete
- (2) it is either is current, correct, and complete or it is possible to detect that these properties do not hold.

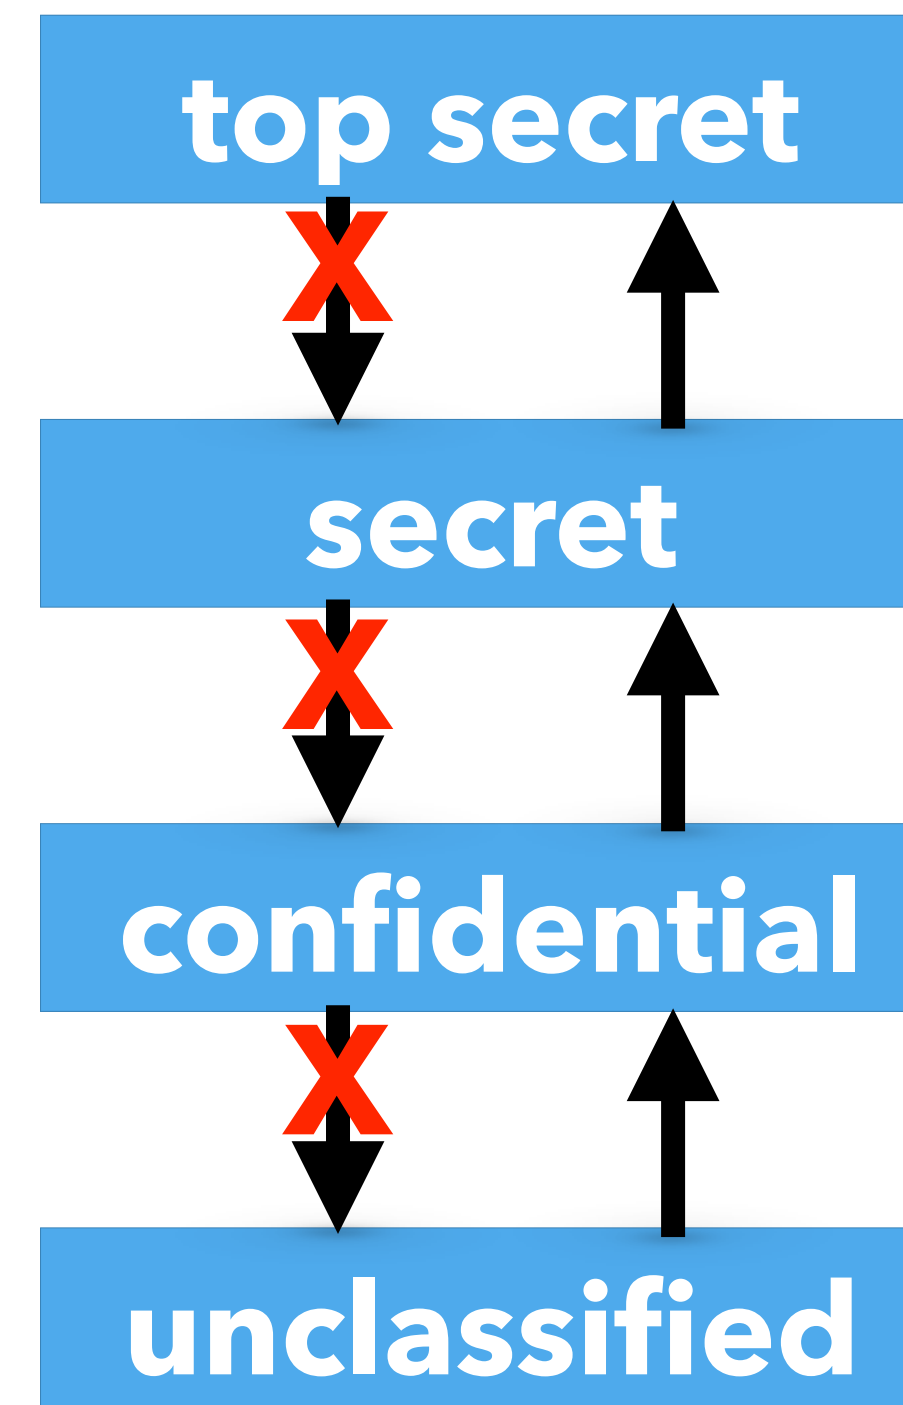
Model for Confidentiality

Secrecy Levels:


- Classification (documents)
- Clearance (persons)
- The higher the level the more sensitive the data
- totally ordered

Categories

information



operations

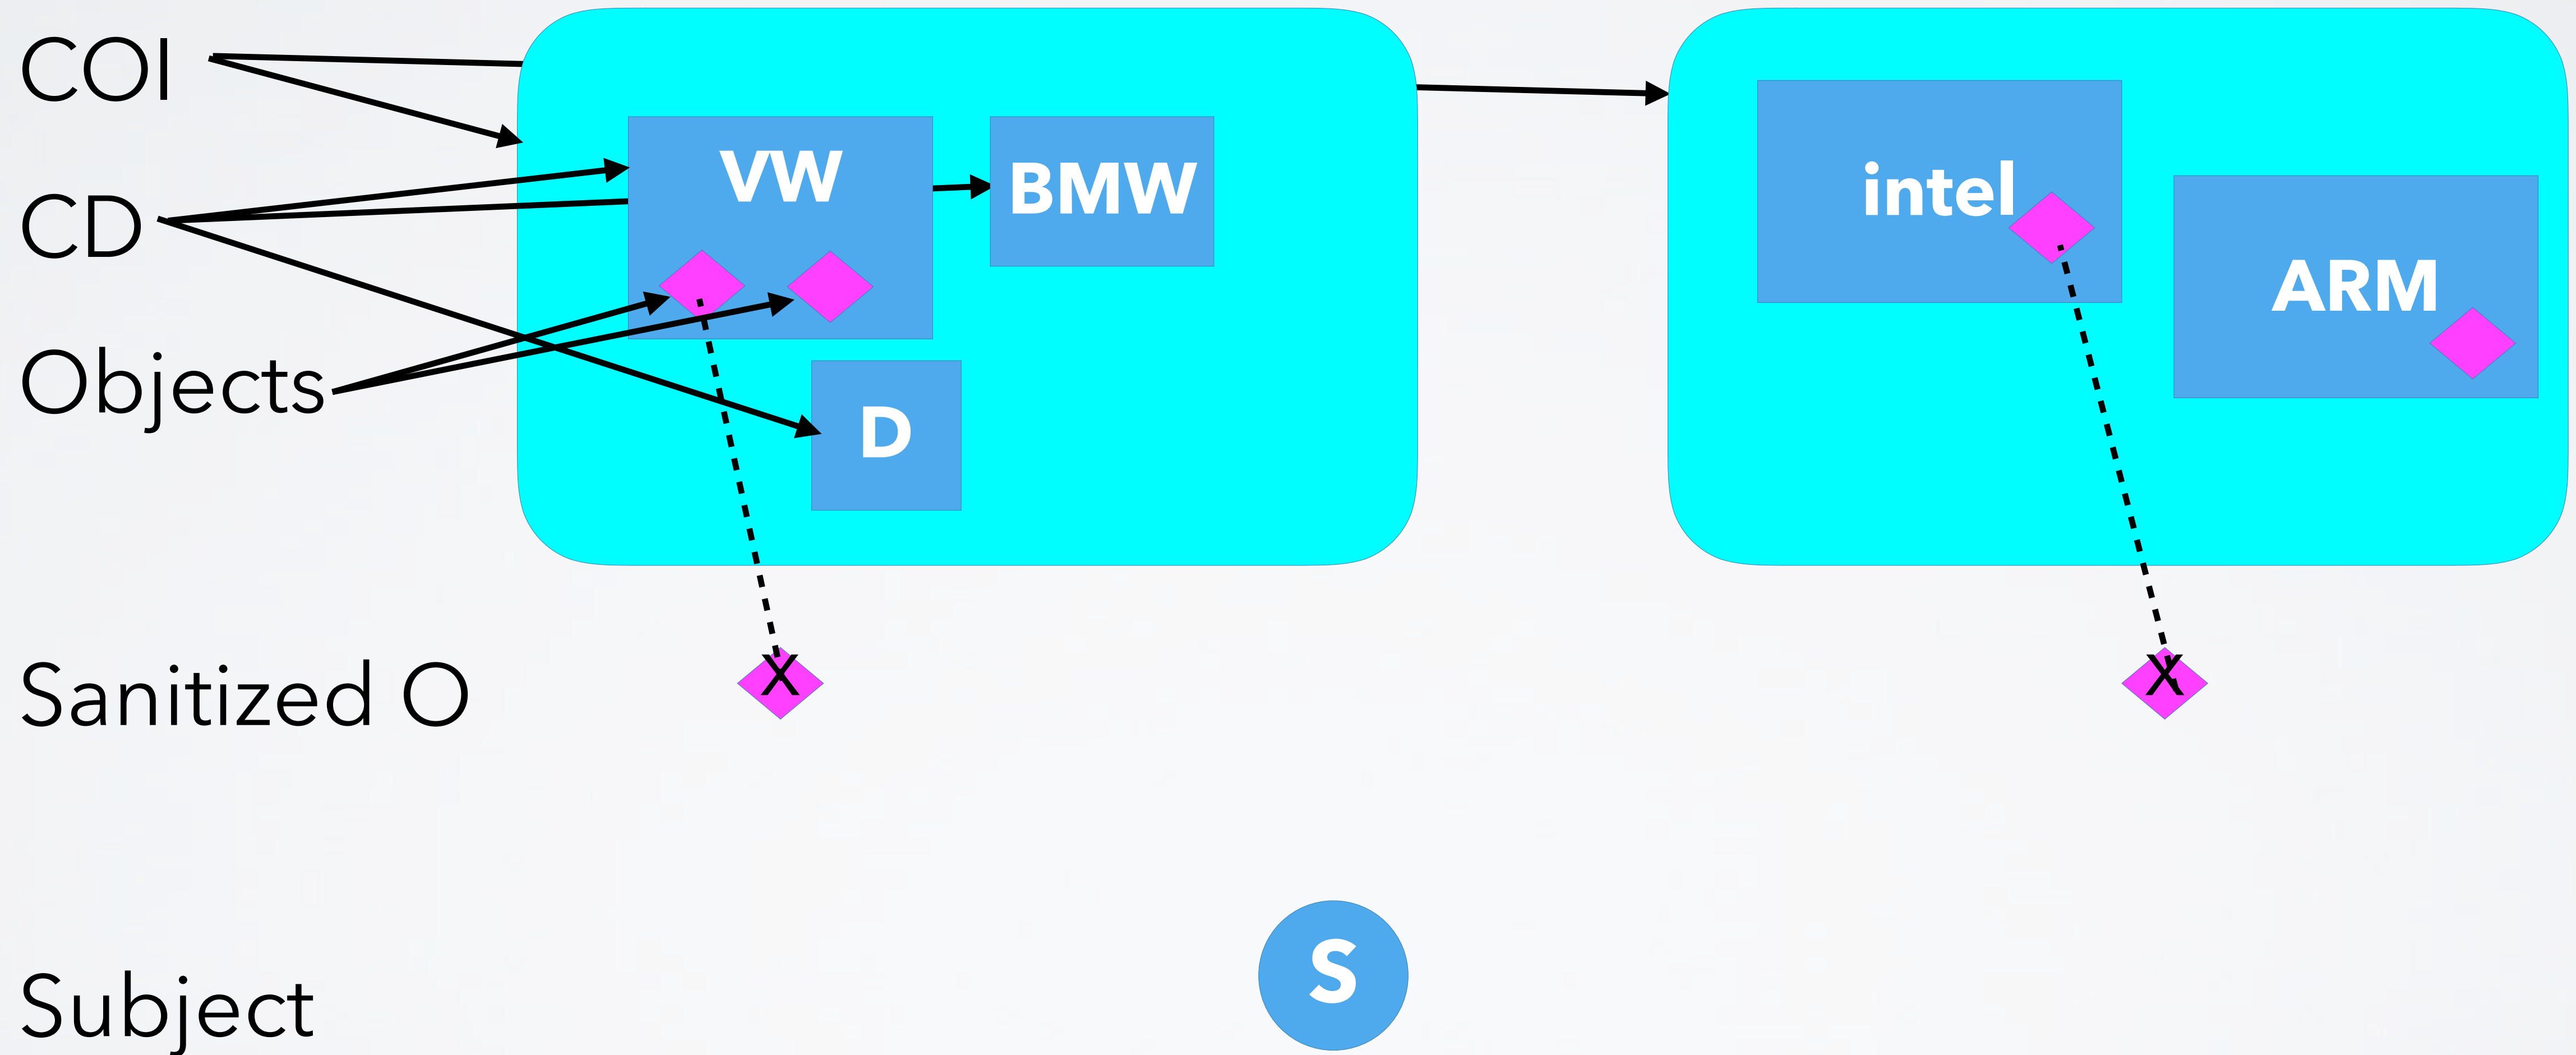


- categories: NATO, Nuclear
levels/clearance: top secret, secret, confidential, unclassified
- document: Nato, secret
- person clearance: read
secret, Nato -> allowed
secret, Nuclear -> not allowed
confidential, Nato -> not allowed

Confidentiality & Integrity

- Subjects
- Objects: pieces of information of a company
- CD: Company Data Sets
objects related to single company
- COI: Conflict of Interest class
data sets of competing companies
- Sanitized Objects
version of object that does contain critical information

CHINESE WALL, EXAMPLE

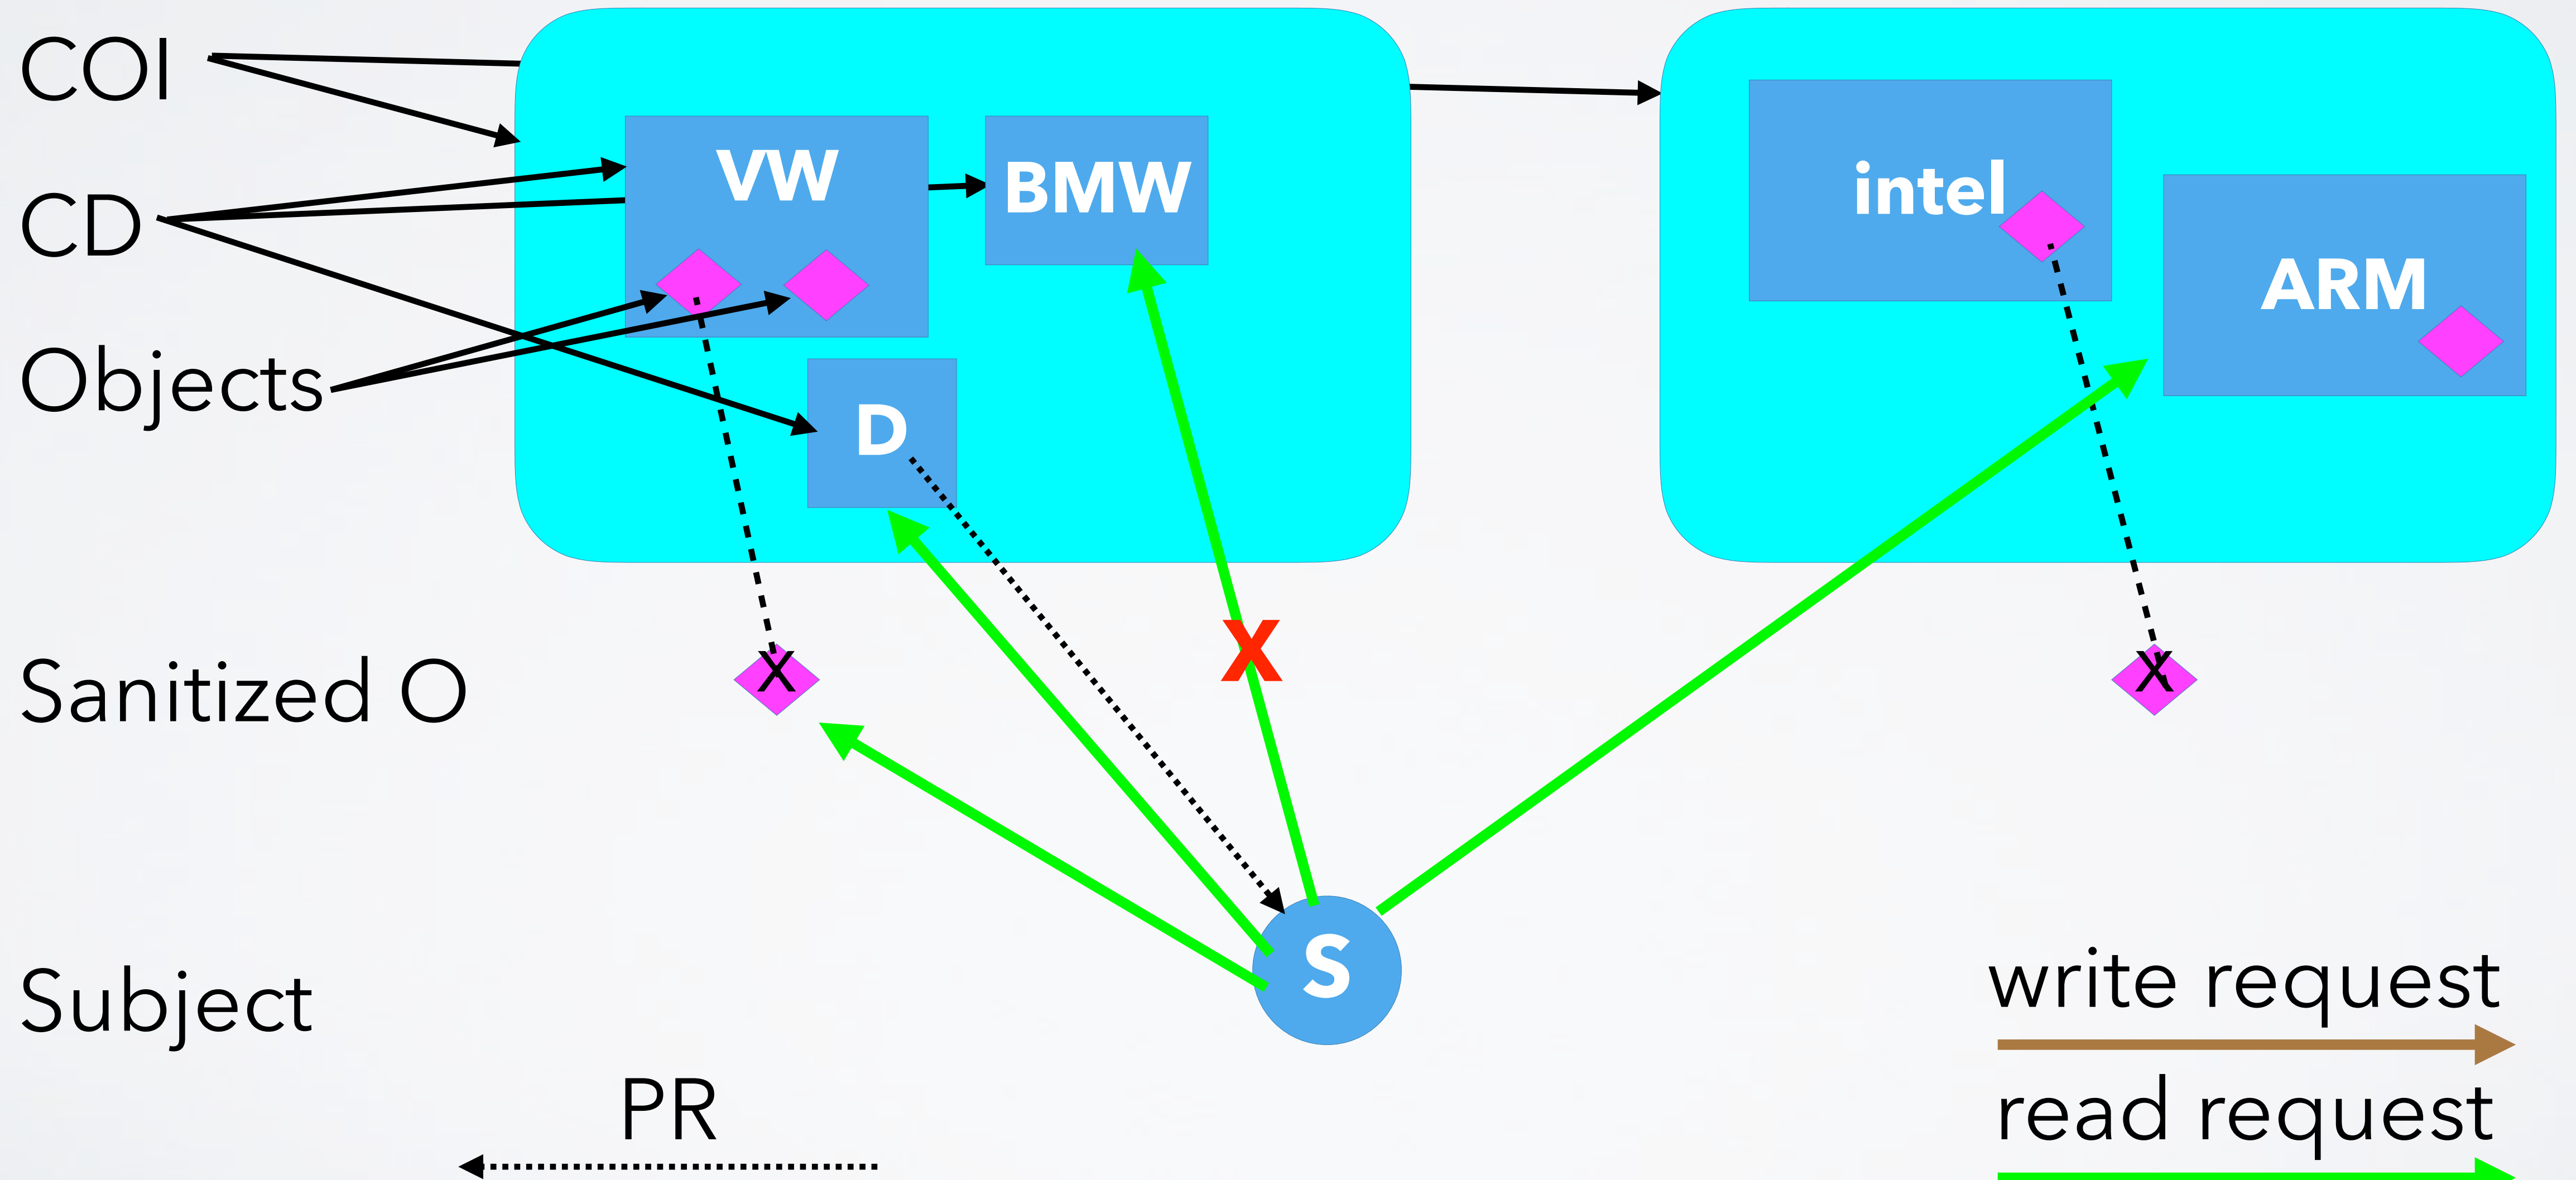


PR(S): set of Objects previously read by S

S can read O, if any of the following holds

- first-time read
- $\forall O', O' \in PR(S) \Rightarrow COI(O) \neq COI(O')$
- O is a sanitized Object

CHINESE WALL, EXAMPLE



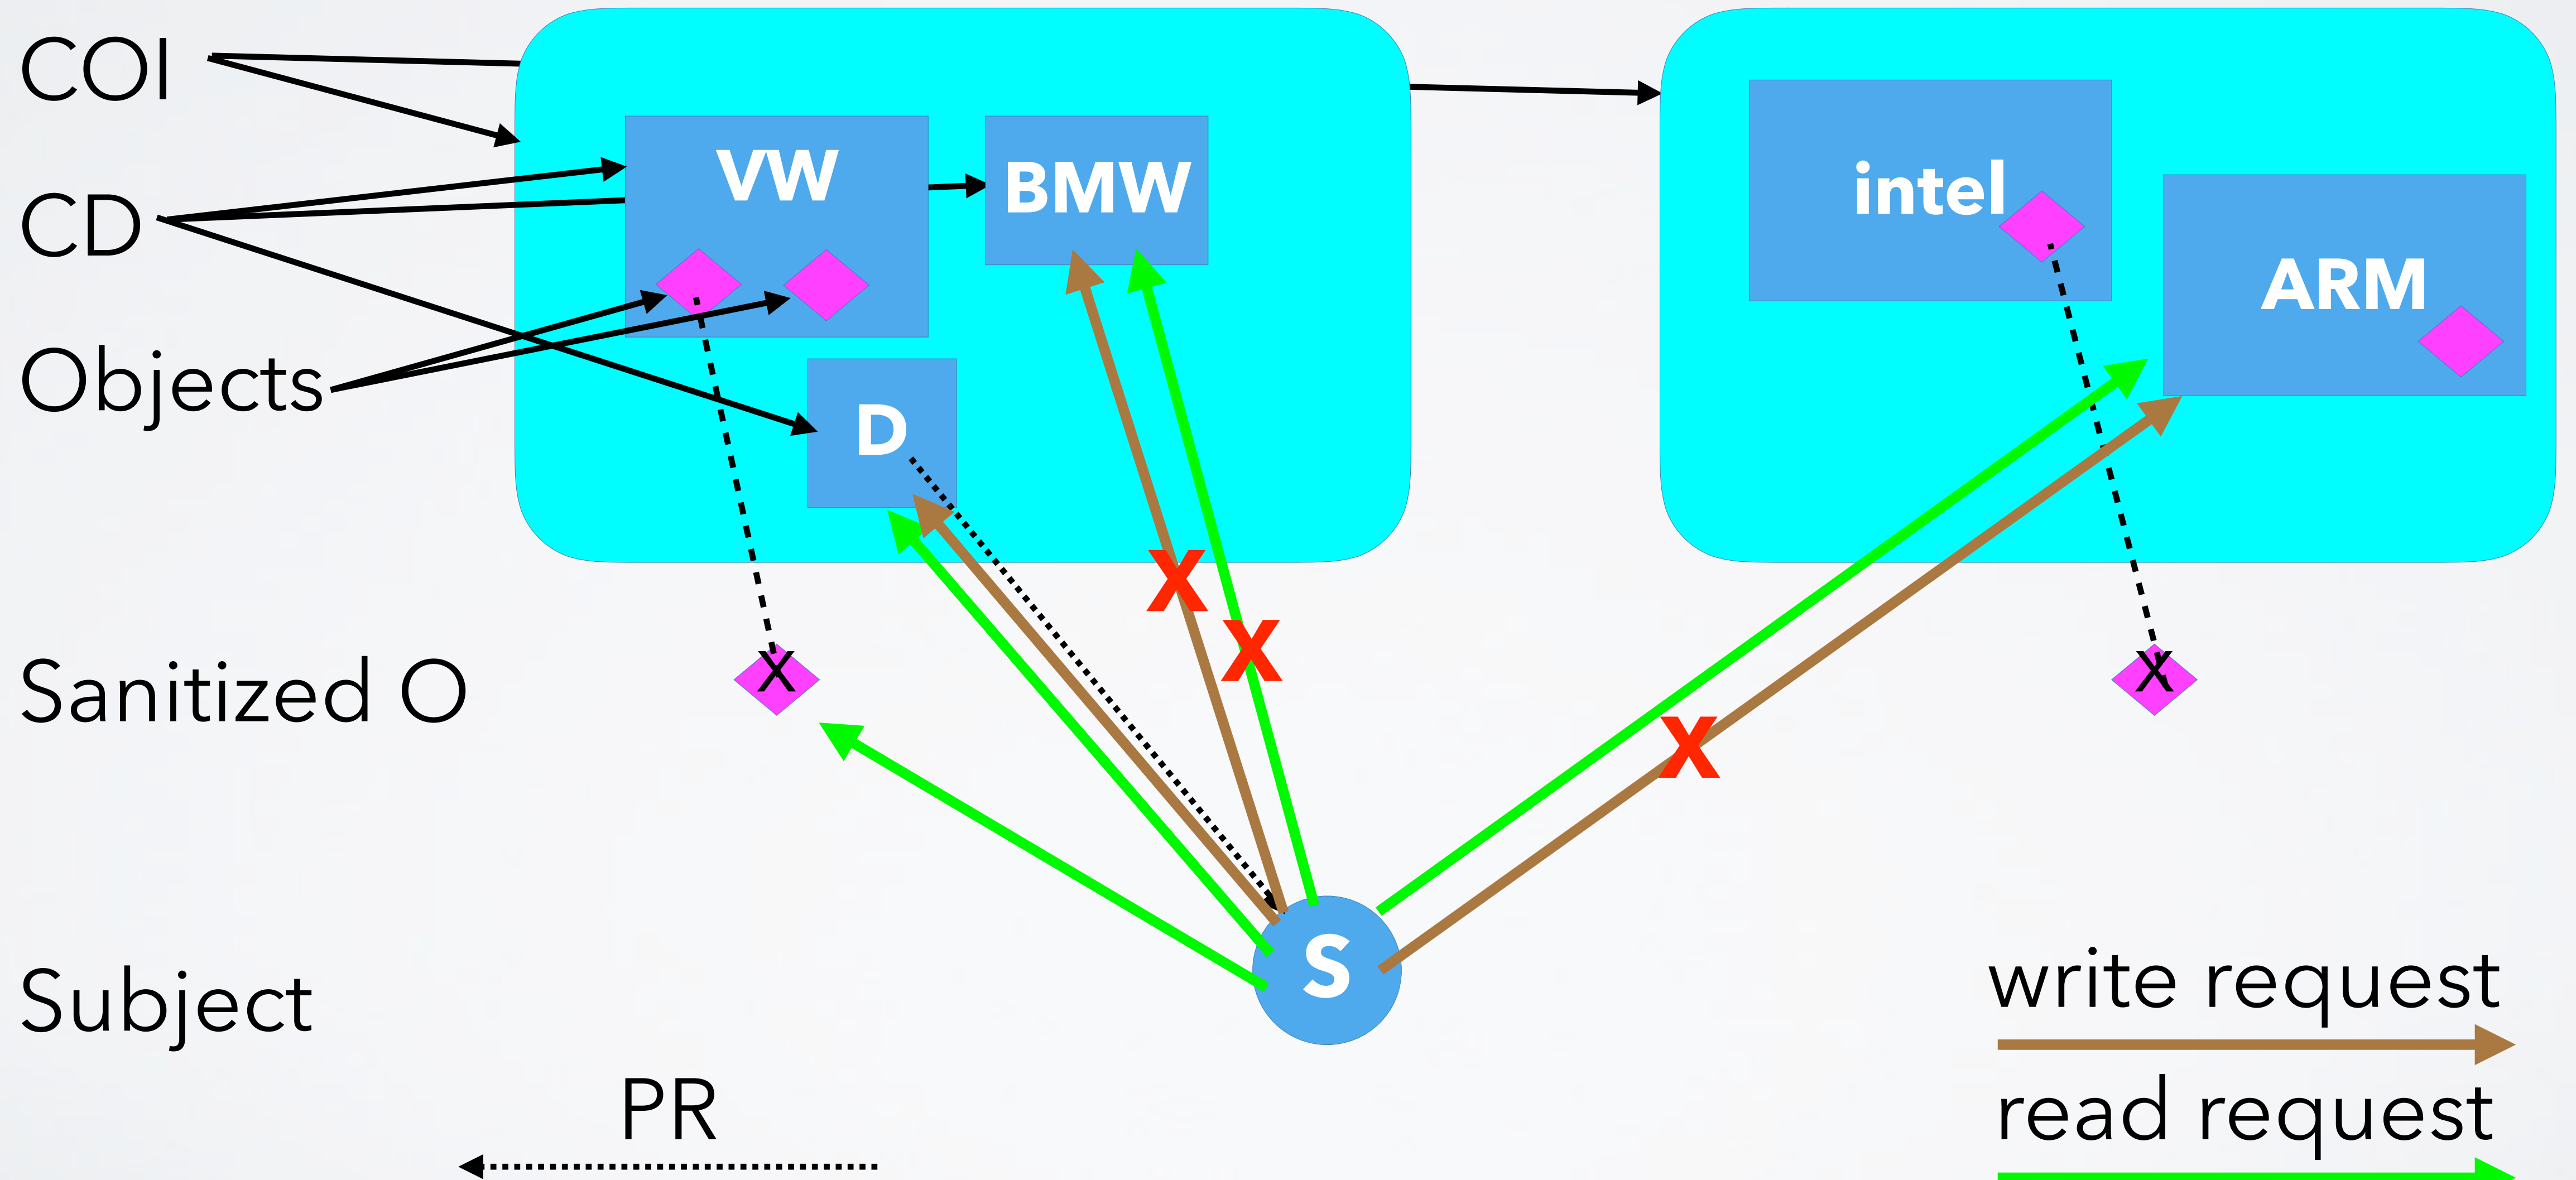
$PR(S)$: set of Objects read by S

S can write O , if

- "S can read O "

- \forall unsanitized O' , "S can read O " $\Rightarrow CD(O) = CD(O')$

CHINESE WALL, EXAMPLE



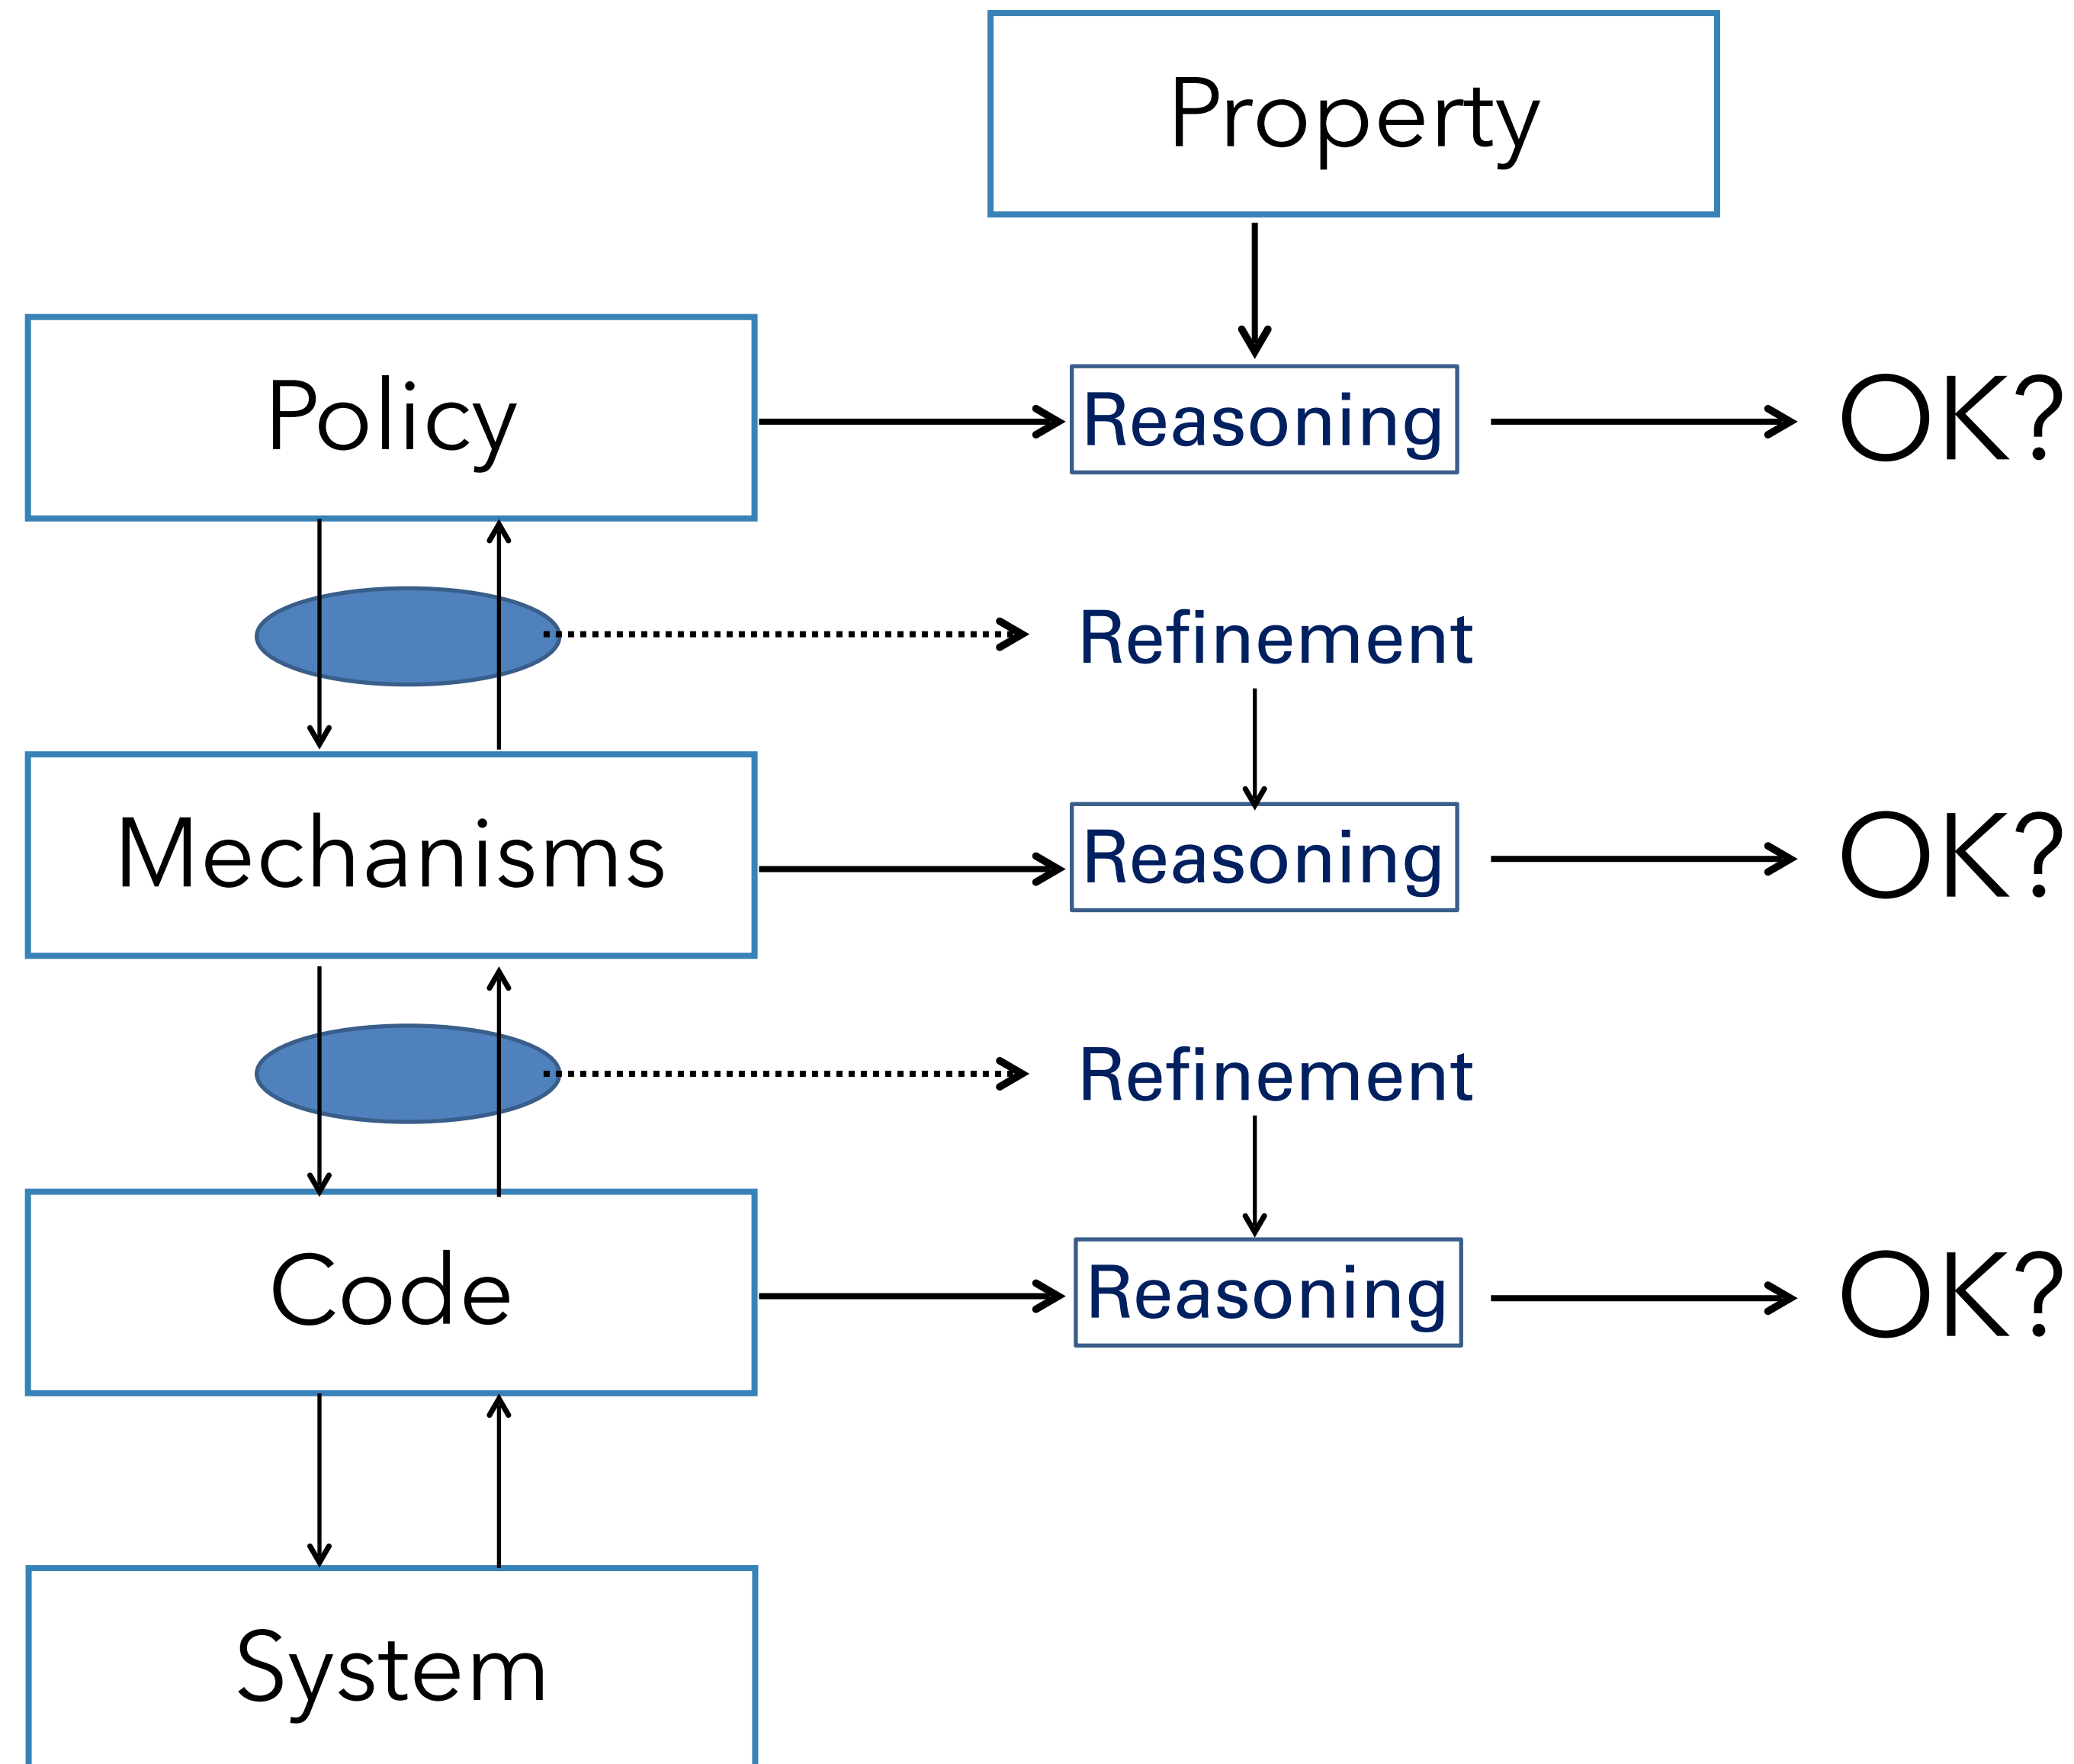
Operating Sys. Mechanisms:

- Access Control List
- Capabilities

Explain Q3 and ...
formalize per model!

Models:

- *based on Access Control Matrix*
- *"take grant" model*



Subjects: S

Objects: O

Entities: $E = S \cup O$

Rights: {read, write, own, ...}

Matrix: $S \times E \times R$

Simple ACM Operations:

create subject / object

destroy subject / object

enter / delete R into cell (s,o)

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own

ACM

- Access Control List (ACL)

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own

- Capabilities

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own

- Define Protection Mechanisms of an Operating System in terms of primitive ACM operations
- only the defined mechanism provided by the OS can be used

- "Leakage":
 - an access right is placed into S/O that has not been there before
 - it does not matter whether or not that is allowed
- Is leakage decidable ?

Examples for OS-Mechanisms defined by ACM-Operations:

UNIX create file (S1,F)

create object

enter own into A(S1,F)

enter read into A(S1,F)

enter write into A(S1,F)

	O1	O2	S1	S2	F
S1	r,w,own	r,w	r,w,own	--	r,w,own
S2	r,w	r,w,own	-	r,w,own	-
S3	r,w	r	w	--	-

Examples for OS-Mechanisms defined by ACM-Operations:

UNIX `chmod -w (S2,F)`

if `own ∈ A(caller,F)`

then delete `w` in `A(S2,F)`

	O1	O2	S1	S2	F
S1	r,w,own	r,w	r,w,own	--	r,w,own
S2	r,w	r,w,own	-	r,w,own	r,-
S3	r,w	r	w	--	-

Q3:

Given an OS with a ACM-based description of protection mechanisms is "Leakage" decidable for any R in $A(x,y)$?

Decidable

- no subjects/objects can be created

or ■ only **one** primitive ACM operation per OS-Mechanism
by exhaustive search !

Q3 in general:

- undecidable (proof: reduction to Turing machine)


Directed Graph:

Subjects: ●

Objects: ○

Either S or O: ⊗

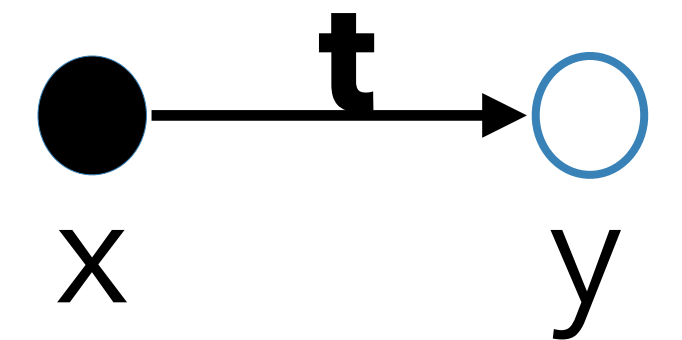
x has capability
with set of rights α on y:



t take right

x has cap with set of rights


τ that includes t



g grant right

x has cap with set of rights


γ that includes **g**



Rules:


take rule ($\alpha \subseteq \beta$)

a takes (α to y) from z



grant rule ($\alpha \subseteq \beta$)


z grants (α to y) to x



Rules:


create rule

x create (α to new vertex) y



remove rule

x removes (α to) y



Application of rules \vdash^* creates sequences of Graphs G_i

CanShare(α, x, y, G_0):


there exists a sequence of $G_0 \dots G_n$ with $G_0 \vdash^* G_n$

and there is an edge in G_n : 

Q3/ 2: CAREFUL: LEMMA


take rule ($\alpha \subseteq \beta$)

a takes (α to y) from z




grant rule ($\alpha \subseteq \beta$)

z grants (α to y) to x




Question:



Q3/ 2: CAREFUL: LEMMA


$(\alpha \subseteq \beta)$

create rule

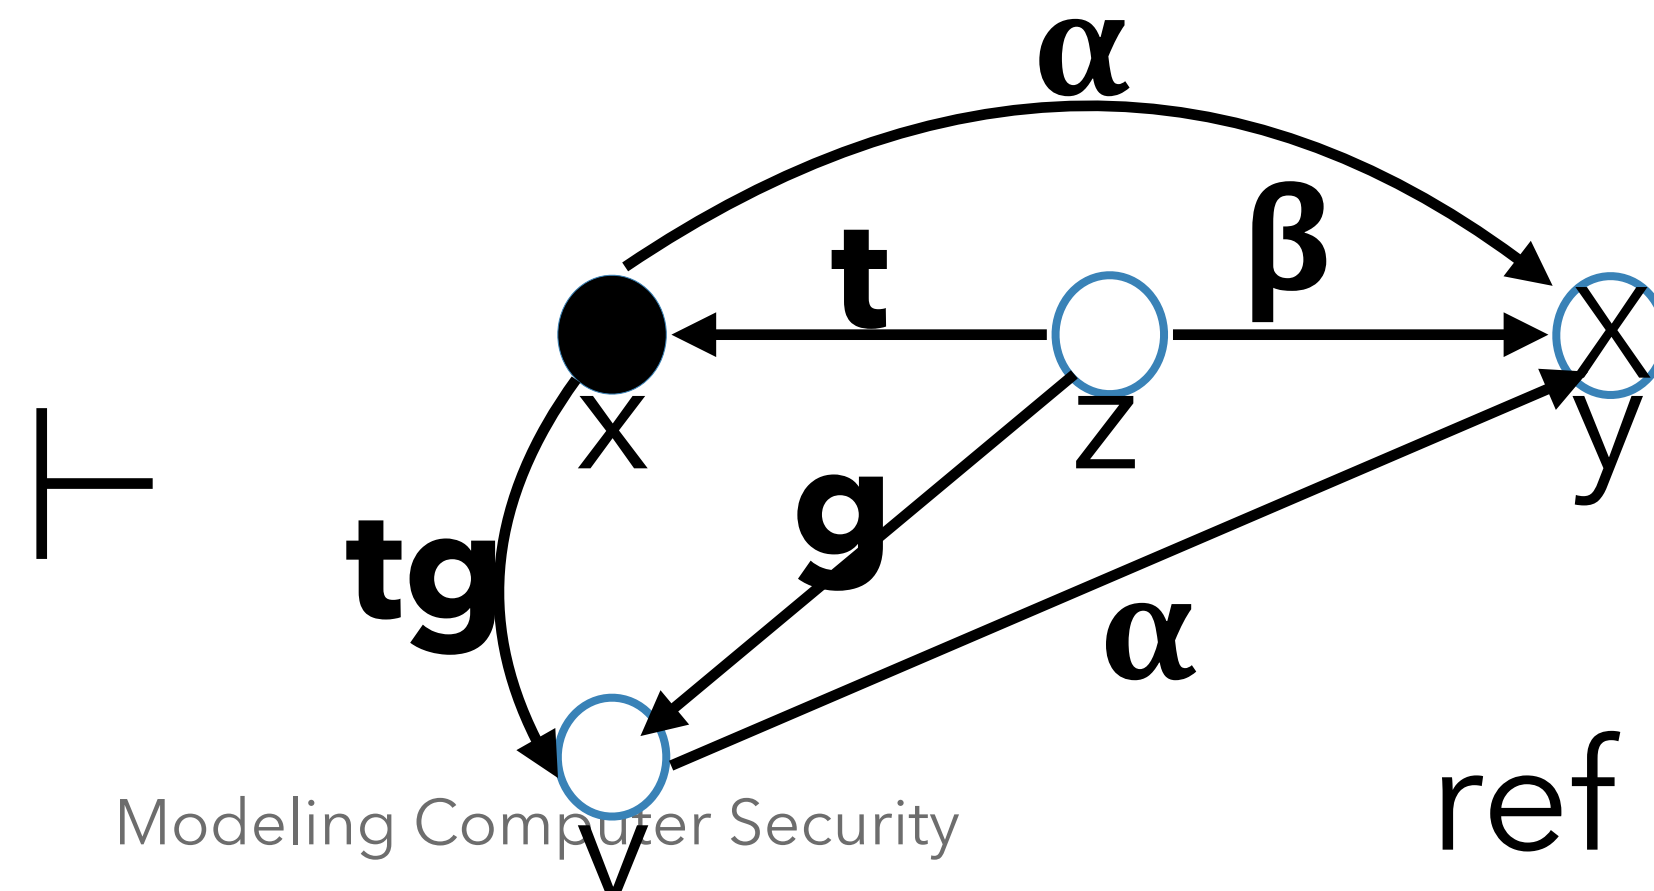


z takes (g to v) from x

z grants (α to y) to v




x takes (α to y) from v



CanShare(α, x, y, G_0):

there exists a sequence of $G_0 \dots G_n$ with $G_0 \vdash^* G_n$

and there is an edge:



CanShare decidable in linear time !

- three questions, 2 models per question, different answers !!!
- modeling is powerful
- need to look extremely carefully into understanding models !!!

- Q1/M1:

Pfitzmann A., Härtig H. (1982) Grenzwerte der Zuverlässigkeit von Parallel-Serien-Systemen. In: Nett E., Schwärtzel H. (eds) Fehlertolerierende Rechnersysteme. Informatik-Fachberichte, vol 54. Springer, Berlin, Heidelberg (in German only)

- Q1/M2:

John v. Neuman, PROBABILISTIC LOGICS AND THE SYNTHESIS OF RELIABLE. ORGANISMS FROM UNRELIABLE COMPONENTS.

- Q2: most textbooks on distributed systems

- Q3: textbook: Matt Bishop, Computer Security, Art and Science, Addison Wesley 2002