TIUPAM: A Framework for Trustworthiness-centric Information Sharing

Shouhuai Xu
Univ. Texas at San Antonio

Joint work with
Qun Ni and Elisa Bertino (Purdue Univ.)
Ravi Sandhu (Univ. Texas at San Antonio)

Roadmap

- Motivation and Goal
- The TIUPAM Framework
  - High-level structure
  - Components
- Related Work

Motivation

- Information is essential to decision making
  - Datamining-like techniques can deal with the information volume issue
- But what if the underlying data is inaccurate, incorrect, inappropriate, **misleading, or maliciously introduced**?
- The problem is further complicated by the shifting from "need to know" to "need to share"
  - E.g., one can obtain data from sources not within previously defined boundary (e.g., internal sourced information) with (reasonable) accountability

Goal

- A systematic framework for information sharing
  - Trustworthiness-centric Identity, Usage, Provenance, and Attack Management (TIUPAM)
  - Four supporting components:
    - Identity management
    - Usage management
    - Provenance management
    - Attack management
  - The framework is centered at the need of trustworthiness and risk management for decision makers

Roadmap

- Motivation and Goal
- The TIUPAM Framework
  - High-level structure
  - Components
- Related Work

Bird's Eye View of TIUPAM

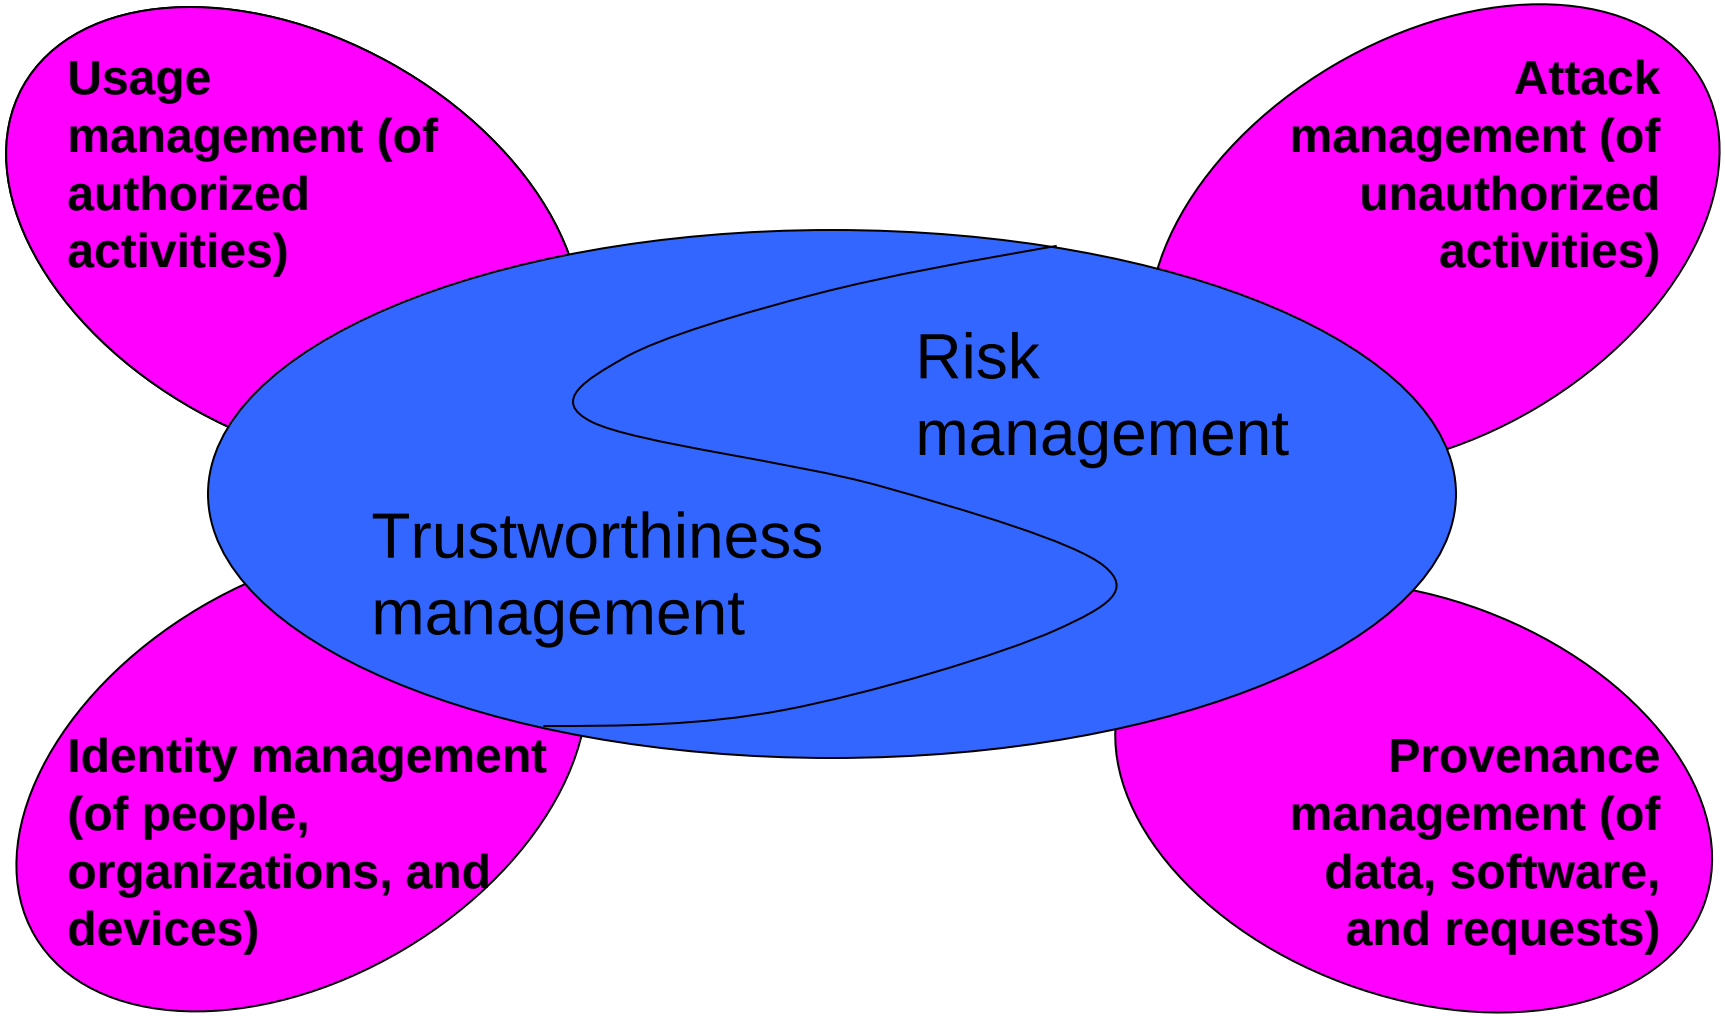

Note: "1 – trustworthiness ≠ risk" in general

Architecture of TIUPAM

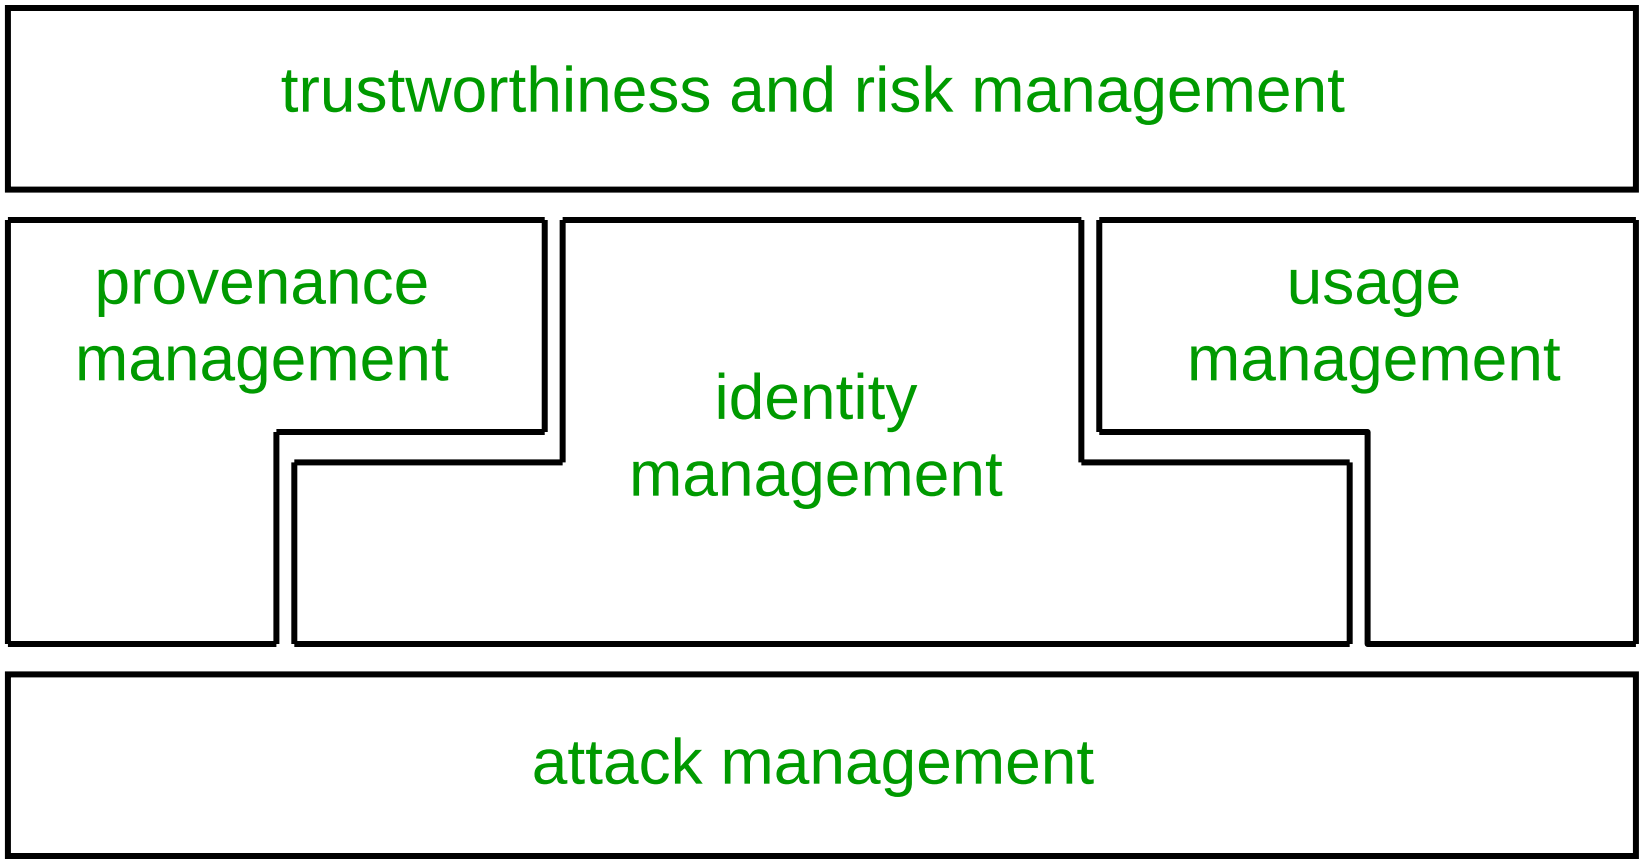

Core questions

- Trustworthiness and risk management
  - Data at hand reflects the current “snapshot” of the world, and may be (in)accurate, (in)correct, misleading, or even maliciously introduced
  - The “snapshots” are dynamical as, for example, incorrect information may later be corrected
  - **Trustworthiness** is a measure against the “snapshot” of one’s up-to-date observation about information in question
  - **Risk** is a measure against the potential consequences caused by the execution of decisions based on not-necessarily-trustworthy information.

Core questions (cont.)

- For Identity Management: How can/should we evaluate the trustworthiness of digital identities and digital credentials of people, organizations and devices?
  - The "snapshots" are derived, in one way or another, from the statements asserted by the relevant people, organizations, and devices
    - A software vendor digitally signs "the output of this algorithm (e.g., datamining) is the desired results"
    - A message digitally signed by an organization would make one tend to accept it as trustworthy
  - However, there are threats like botnets, identity theft

Core questions (cont.)

- For Usage Management: How can/should we deal with "authorized" activities in situations complicated by "need to share" and "not necessarily having prior authorizations"?
  - What is the trustworthiness of a request (in delivering its promise of appropriately using data)?
  - A subject is traditionally deemed as trustworthy (trusted) as long as it passes certain authentication. But what if the authentication credential has been compromised?
  - An object is always trustworthy (trusted) as long as it is in the filesystem or database. But what if the object itself was malicious or incorrectly provided?

Core questions (cont.)

- For Provenance Management: How can/should we evaluate the trustworthiness of data, software, and requests?
  - Provenance of data allows to measure the trustworthiness of information
  - Provenance of software helps evaluate the trustworthiness of programs output
  - Provenance of requests enhances the assurance that they are invoked by the individual or process in question, rather than by malware

Core questions (cont.)

- For Attack Management: How can/should we deal with attacks (e.g., maliciously introduce wrong or misleading information into the system) and “unauthorized” activities in situations complicated by “need to share” and “not necessarily having prior authorizations”?
  - Help manage the trustworthiness of infrastructure-level services provided to the other components as well as their services in the framework (e.g., authentication services)
  - PKI may be trusted, but not necessarily trustworthy

Functions as the Glue

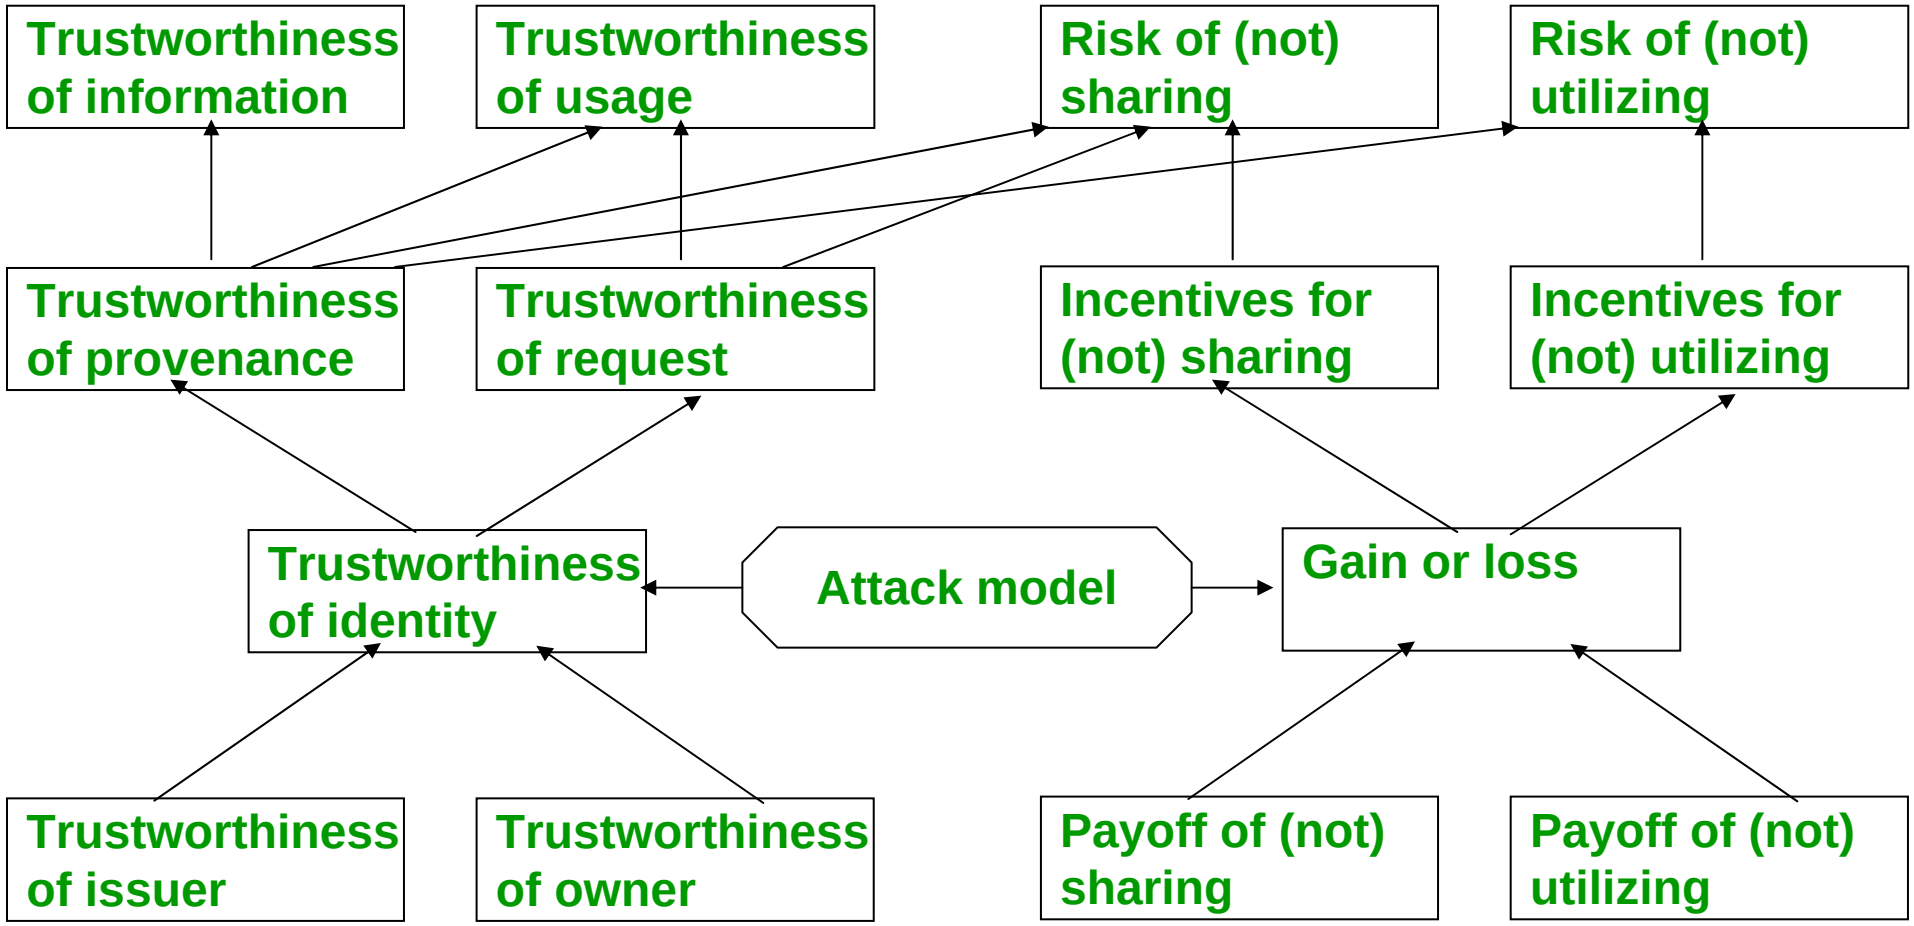

Q: How should we construct/approximate these functions?

Roadmap

- Motivation and Goal
- The TIUPAM Framework
  - High-level structure
  - Components
- Related Work

Component I

- In order to fulfill the envisioned trustworthiness and risk management, what kinds of Identity Management systems we want?

Desired Properties

- Extensibility: Can accommodate or integrate emerging new identity systems.
  - Heterogeneous or not, large-scale or small-scale
- Automated trustworthiness: for higher-layer applications
  - Compromise containment: User-end is relatively easy to deal with when compared with server-end compromise
  - Accountability: Who should be responsible for a transaction (or how good forensics we can hope for)?

Component II

- How can we achieve the Usage Management we envisioned?
- Possible solution:
  - Observation: The dynamic characteristics of usage control are well suited to the problem of trustworthiness-centric information sharing
  - Extending Usage Control (UCON) to Usage Management?

From ABAC to UCON

- Attribute-based access control or ABAC: access is determined by subject and object attributes
- Attributes can be roles, groups, clearances, sensitivity, title, rank, cost, status etcetera
- UCON differentiators:
  - Mutable attributes:
  - Obligations:
  - Conditions:
  - 3 phase enforcement: pre, ongoing, post.

UCON Differentiators

- Mutable attributes: attributes can change due to access, allows consumable rights.
  - Only three ATM withdrawals in a day
- Obligations: in addition to authorization.
  - Click accept button on license agreement
- Conditions: system wide conditions.
  - Threat level: red, orange, yellow, blue, green
- 3 phase enforcement: pre, ongoing, post.
  - Change in threat level can kill ongoing access (ongoing condition)
  - When phone card runs out of money the call is killed (ongoing authorization)
  - After access is completed the minutes used are billed to the account (post obligation)

UCON Model

- unified model integrating
  - authorization
  - obligation
  - conditions
- and incorporating
  - continuity of decisions
  - mutability of attributes

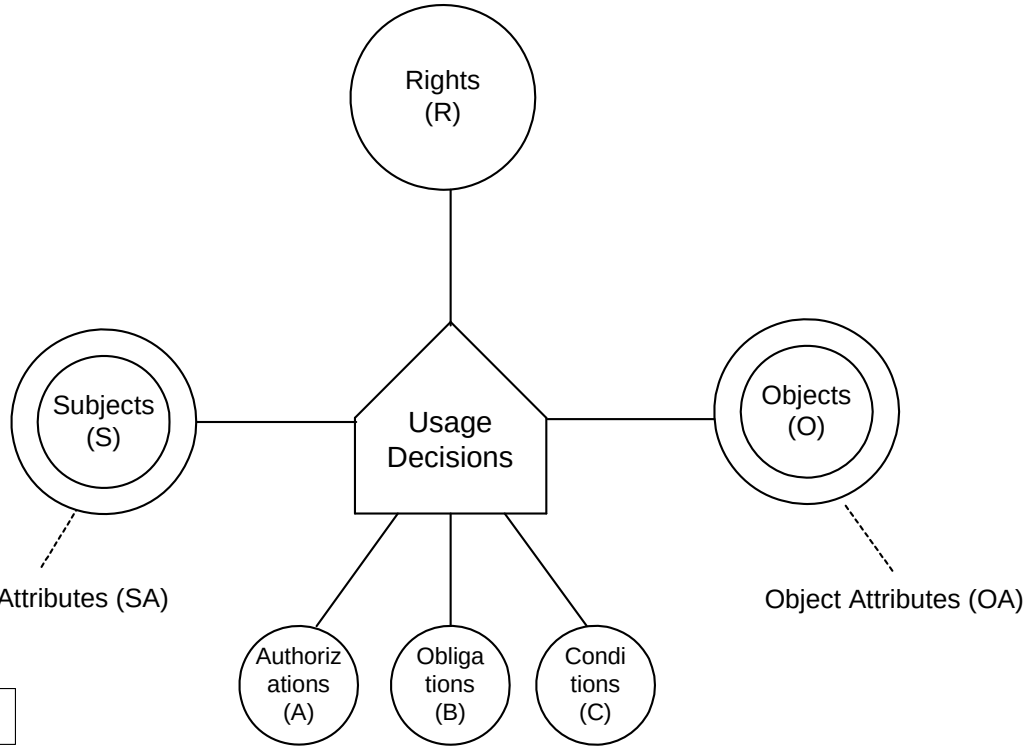

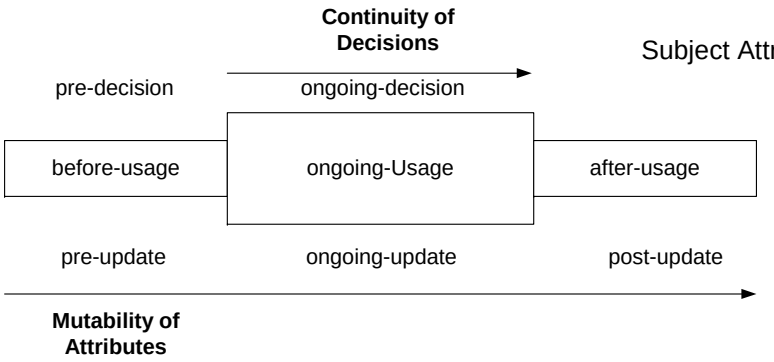

UCON is Attribute-Based Access Control on Steroids

UCON Limitations

- Future obligations
  - Access exception is made but should be explained
  - Example: Physician is allowed access in emergency situation but has future obligation to write an explanatory note within 3 weeks
- System obligations
  - Obligation is on the system not on the user
  - Example: data must be deleted after use
- Handle trustworthiness and risk

Component III

- Need Secure Provenance Management systems
- Security aspect is not well understood
  - [Braun-Shinnar-Seltzer HotSec'08] highlighted the difficulties in managing access control to provenance information
  - We look at a broader picture:
    - How to use provenance for trustworthiness and risk management?
    - How to secure provenance information itself?

Provenance

- Why-provenance: why is a certain piece of information here?
- Source-provenance: what is the source of a certain piece of information?
- How-provenance: how does a certain piece of information get here?

Secure Provenance Management

- What would secure provenance management systems --- say, as an analogy to secure DBMS --- look like?
  - Functional requirements
  - Security requirements
- How should we design and implement them?
  - Facets of secure provenance management

Functional Requirements

- Secure provenance management system should cover the entire lifecycle of information (in this context) as well as their associated provenance

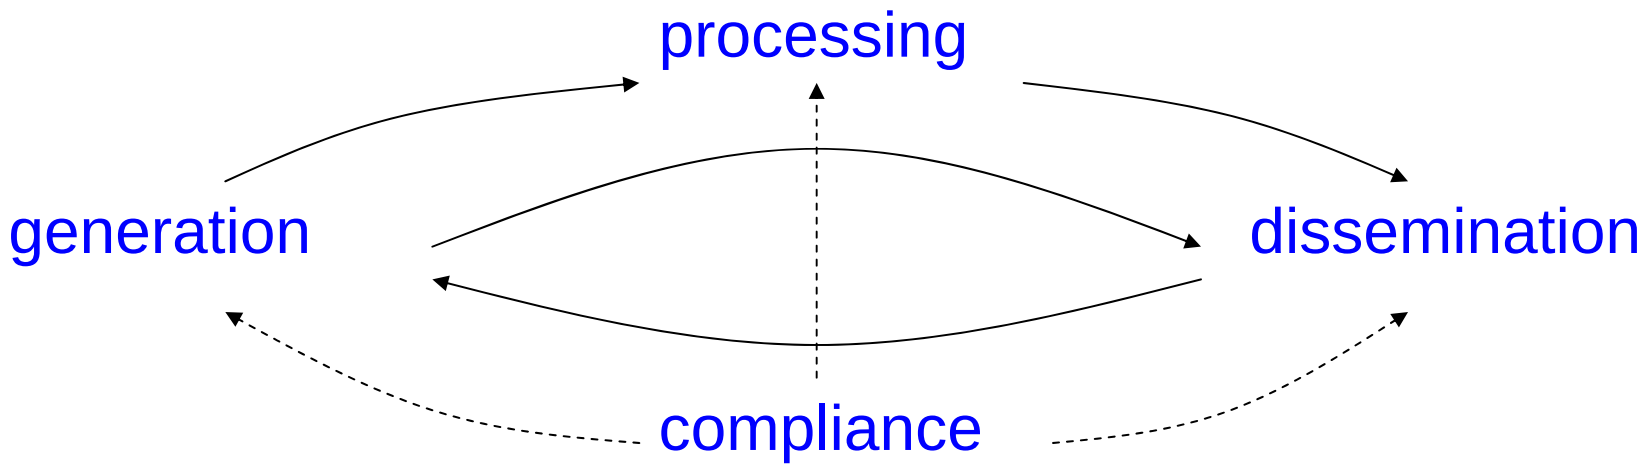

Functional Requirements

- Information generation: first time enter the system
- Information processing: new information items may be derived (e.g., datamining output)
- Dissemination: data and information can be disseminated in the presence of malicious attacks
- Compliance: who could read/write/modify and who have read/written/modified which data items

Security Requirements

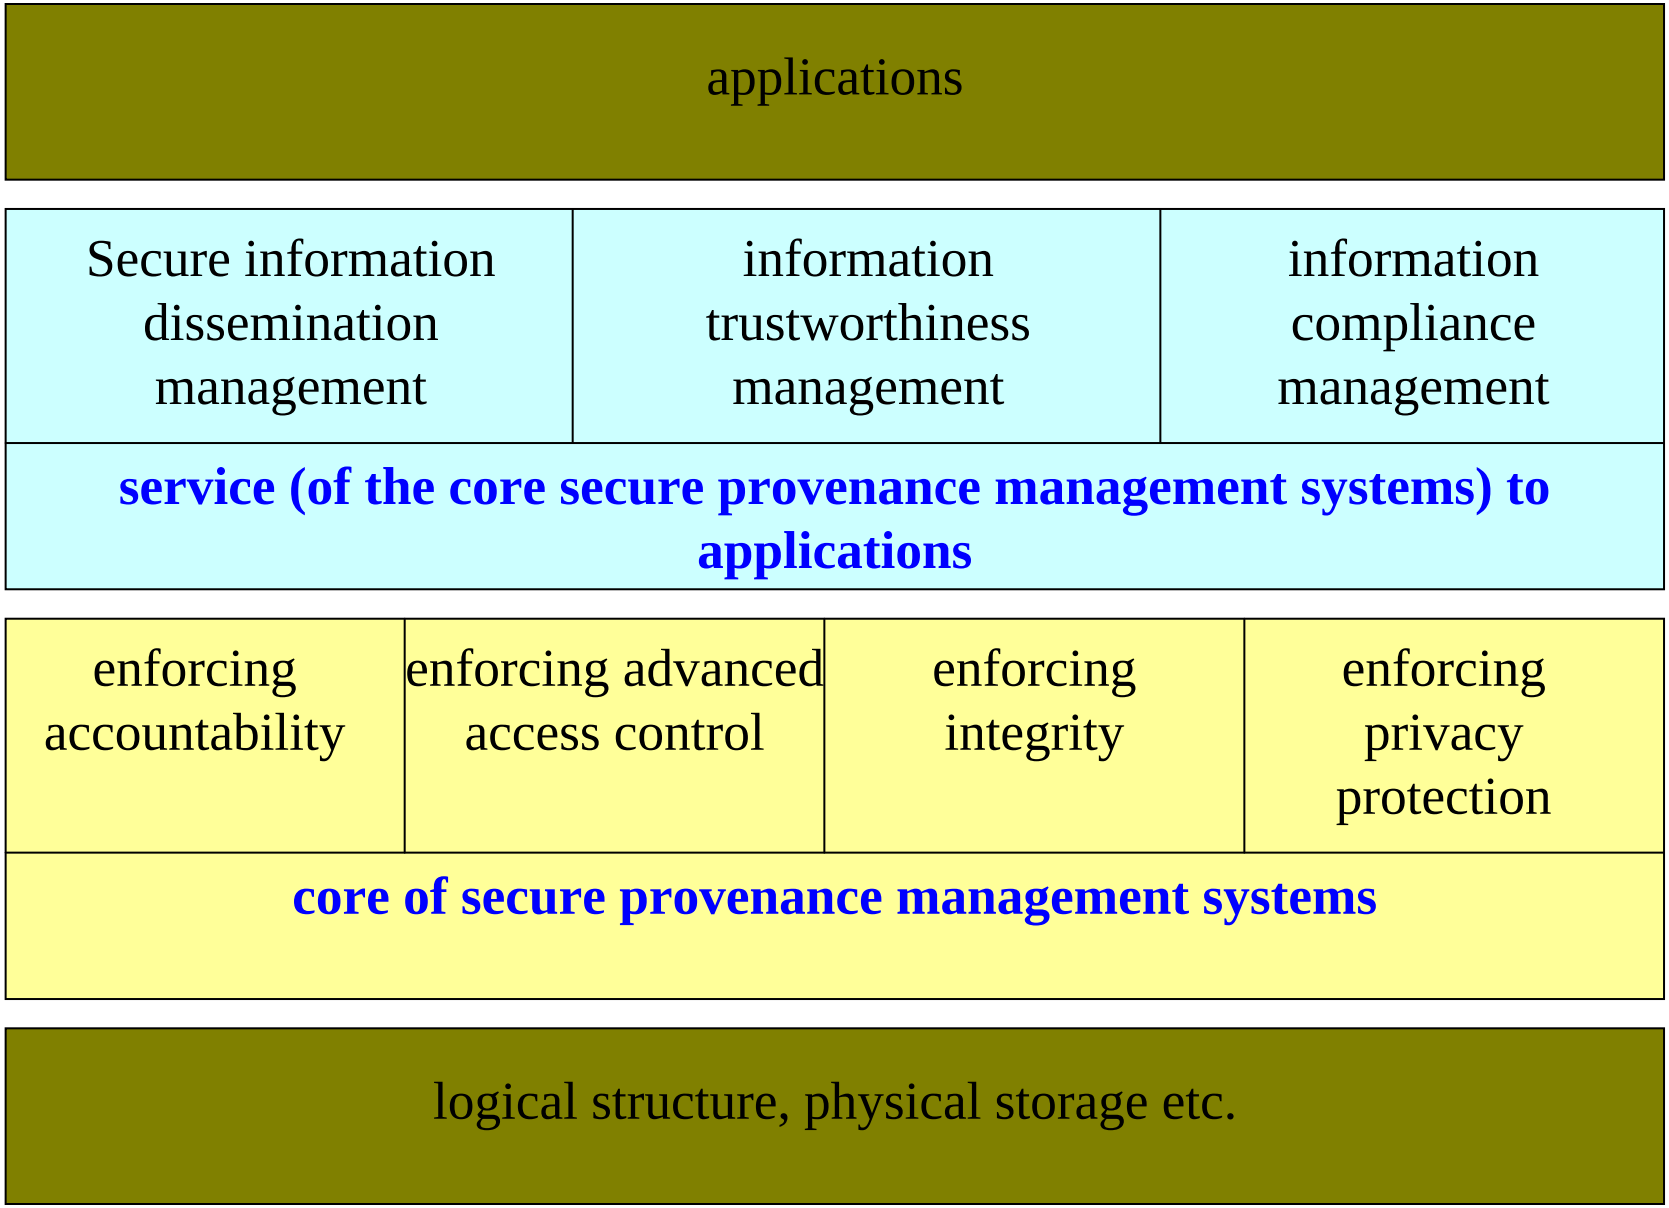

Security Requirements

- Information trustworthiness management
  - For a source, what is the trustworthiness of an information item that has to be entered into the system?
  - For an intermediate node, what is the trustworthiness of both the source and the prior intermediate nodes so that, e.g., a decision may be made whether to re-disseminate the processed information?
  - For an information consumer, how to evaluate the trustworthiness of an incoming information item?
  - For an administrator, who has greater influence in the information network?

Security Requirements

- Secure information dissemination management
  - What if there are malicious insiders/attackers in the dissemination system?
  - How can we enforce re-dissemination control?
  - Is the dissemination process privacy-preserving?

Security Requirements

- Information compliance management
  - Who has read/written/modified and could read/write/modify a certain data item?
    - This might help detect malicious insider who leaked a certain confidential information.
  - Who has read/written/modified and could read/write/modify a certain provenance item?
    - This might help detect malicious insider who leaked a certain confidential information.
    - This could help detect the party who leaked, for example, “who has participated in which operation/process?”

Security Requirements

- Securing data provenance
  - Enforcing advanced access control: provenance of data item if often a Directed Acyclic Graph (DAG). Traditional access control models and their straightforward adaptations are not sufficient [Braun-Shinnar-Seltzer HotSec'08].
  - Enforcing integrity.
  - Enforcing accountability.
  - Enforcing privacy protection.

Facets of Secure Provenance Management

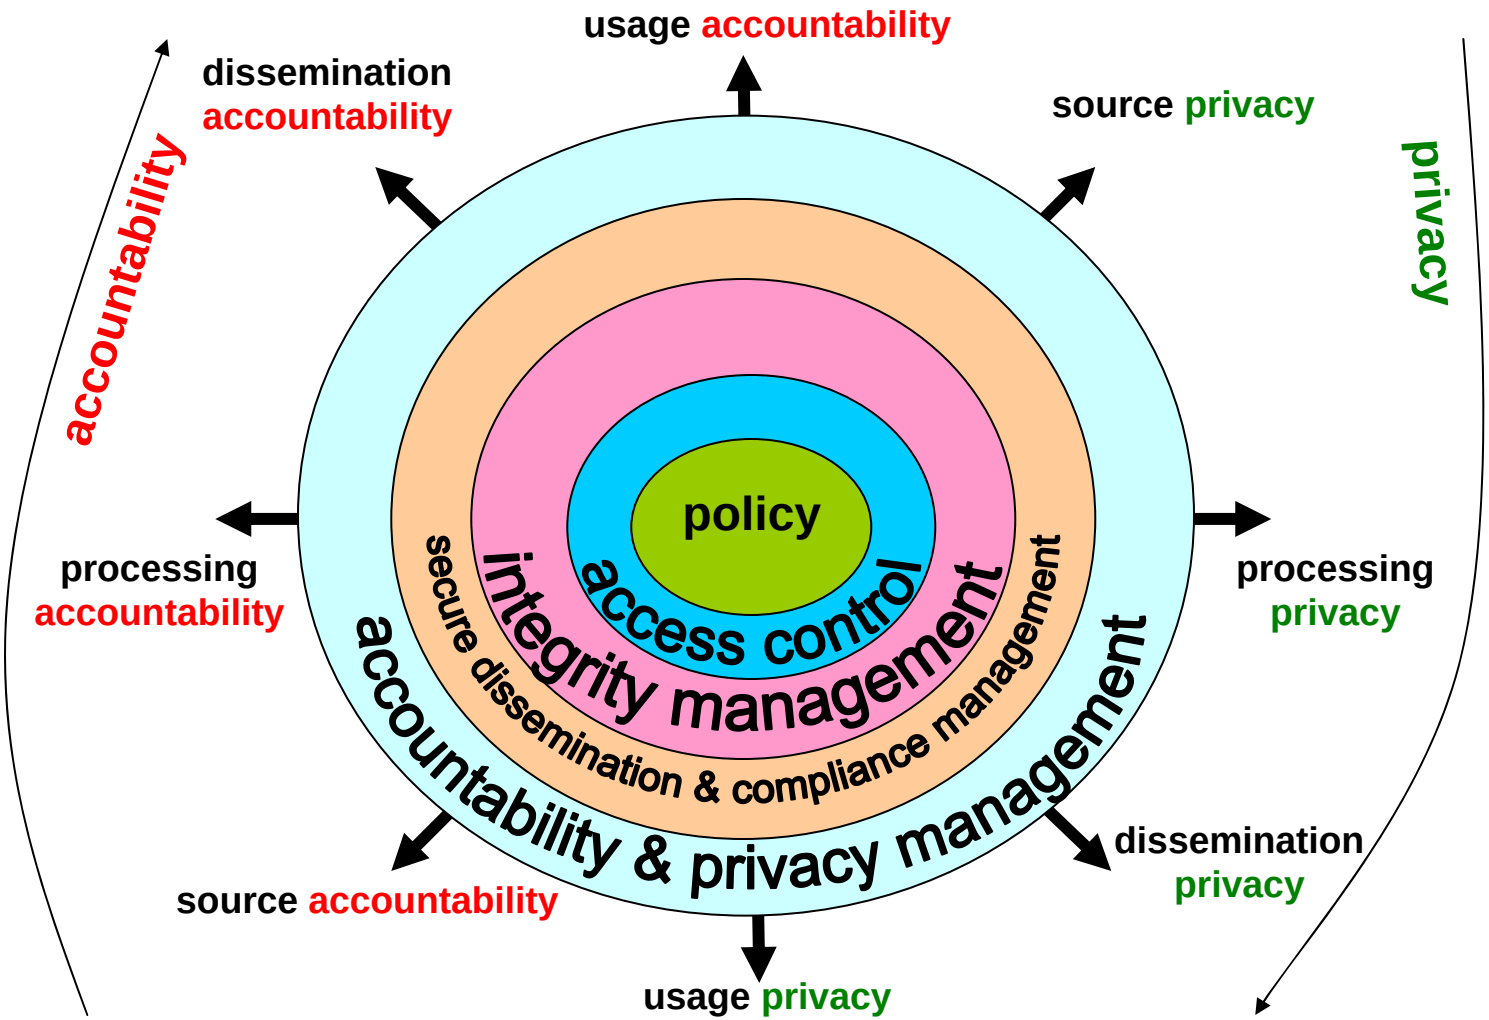

Attack Management

- Attacks attempting to manipulate the trustworthiness of information, including compromise of provenance information
  - Can be done by attack
- Attacks attempting to manipulate the trustworthiness of usage
  - By attacking trustworthiness of information provenance (e.g., malicious information spreads to many users) or trustworthiness of request (e.g., compromising credential)

Component IV

- If attacks are not controllable, can we manage them?
- Attacks attempting to manipulate the trustworthiness of information, including compromise of provenance information
- Attacks attempting to manipulate the trustworthiness of usage
  - By attacking trustworthiness of information provenance (e.g., malicious information spreading to many users) or trustworthiness of request (e.g., compromising credential)

Attack Management

- Any security statement is with respect to a specific adversarial model
- If we want to measure security, trustworthiness, and assurance, we must state the (holistic) model explicitly
- Cryptography has made very progress in this aspect, but we need holistic model

Related Work

- This is an ambitious project
- The framework itself likely will be refined
- We are investigating mechanisms for fitting into the framework as well
- [Braun-Shinnar-Seltzer HotSec'08] presented an excellent discussion on the challenge of access control of provenance data

Thanks, and Questions?