Securing Sensitive Healthcare Information in the U.S. Health Care System: Past, Present and Future

Doug Blough, Georgia Tech

My Recent Health IT Activities

- ***MedVault*** (with Ahamad, Liu: Georgia Tech; Chopra: Children's Healthcare of Atlanta) – security and privacy of electronic health records; sponsor: NSF (from 2007)

- ***ACTSI Biomedical Informatics Security*** (with Ahamad: Georgia Tech; Saltz (Emory Univ. Center for Clinical Informatics) – attribute-based authorization in federated biomedical research applications; sponsor: NIH

- ***GHIE Privacy*** – advisor to the Georgia Health Information Exchange program on privacy issues; GHIE sponsor: HHS

CS Research in Health IT

- ***Personal health*** or ***mobile health:*** home health devices and monitoring, applications to assist with personal health maintenance/improvement, social networking for health, etc.

- ***Health IT in the health system:*** software-controlled medical devices in hospitals and doctor’s offices, *electronic health record (EHR) repositories and exchanges*

US EHR Adoption Rates: Lies, Damn Lies, and Statistics

- ***Harvard School of Public Health 2009 Study***
  - 90% of hospitals in US do not have a functional *comprehensive* EHR system
    - does *not* mean that 90% of hospitals still rely on paper records
    - does mean that, in many cases, the ER uses an IT system from one vendor, the ICU uses a system from a different vendor, radiology uses a system from a 3rd vendor, …, and different systems do not interoperate
  - 83% of doctor’s offices and clinics do not have functional EHR
    - *does* mean that 83% still use paper records, but:
      - 17% had purchased but not yet fully implemented
      - 26% planned to purchase in near future
      - main laggards – very small offices in poor rural areas (health IT divide)

US EHR Adoption Rates, Continued

- ***SK&A 2010 Study of Physician Practices*** (stimulus effect)
  - Jan. 2010 adoption rates: 28% solo, 54% mid-size, 71% large
  - Nov. 2010 adoption rates: 29% solo, 61% mid-size, 73% large

- ***CDC Study of Physician Practices***
  - 2007: 17%, 2008: 21%, 2009: 27%

Past: Closed Health Systems Architecture

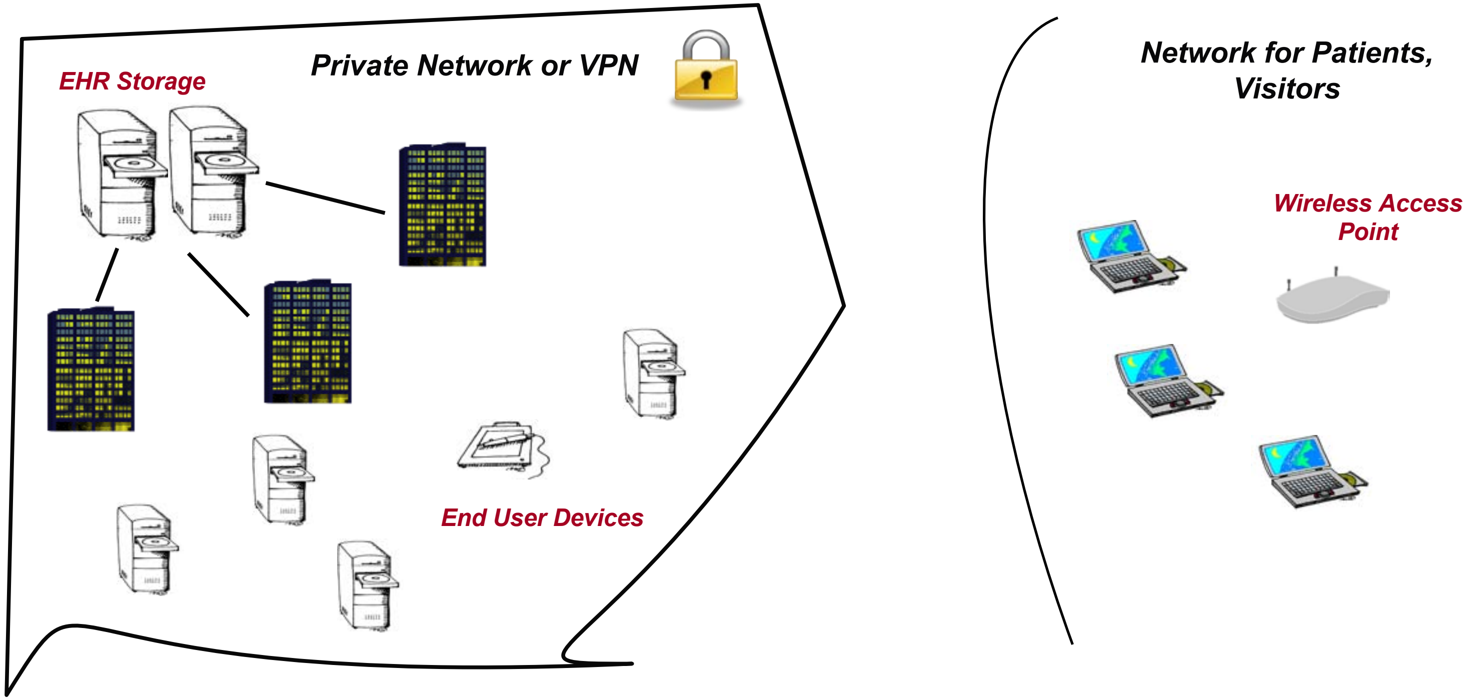

- On-site access only – no external connections allowed (only exception is pushing of billing info)
- Primary threat – insider attack

Typical Security Model for Closed Systems

- ***No access control*** – any health professional (doctor, nurse, technician) can see all records about all patients
- ***Logging*** of all EHR accesses
- ***"Random" auditing*** is performed (in practice, auditing is typically performed on high profile patients or when a complaint has been filed)
- ***Training*** of all staff on policy – EHR info to be accessed only for legitimate medical purposes
- ***Accountability*** – serious consequences, up to and including dismissal, for policy breaches

Representative Incidents in Closed Systems

- ***Lost or stolen components:*** laptops, tapes, servers
  - In 2006, an Indiana man stole a server containing records of 900,000 patients and attempted to extort money from AIG by threatening to release the information on the Internet

- ***Human error***
  - From 1999-2005, Kaiser Permanente posted records of 150 patients on a misconfigured Web site accessible to the public
  - In 2008, a contractor for Grady Hospital in Atlanta posted records of 45 patients on a public Web site for several weeks
  - In 2007, records of 900,000 soldiers, govt. employees, and family members were posted on an SAIC server on the Internet

Representative Incidents, continued

- ***Deliberate internal breaches*** – usually small-scale: curiosity, revenge, blackmail, divorce/custody disputes, etc.
  - In 2009, two Kaiser Permanente employees were fired, 16 resigned, and 9 were otherwise disciplined for accessing the records of Nadya Suleman and her children
  - Two days ago, 3 clinical support staff and 1 nurse at Tucson Univ. Medical Center were fired for unauthorized accesses to records of Gabrielle Giffords and other shooting victims
  - No reliable statistics on this category, many cases likely go undetected

Present: Limited External Access

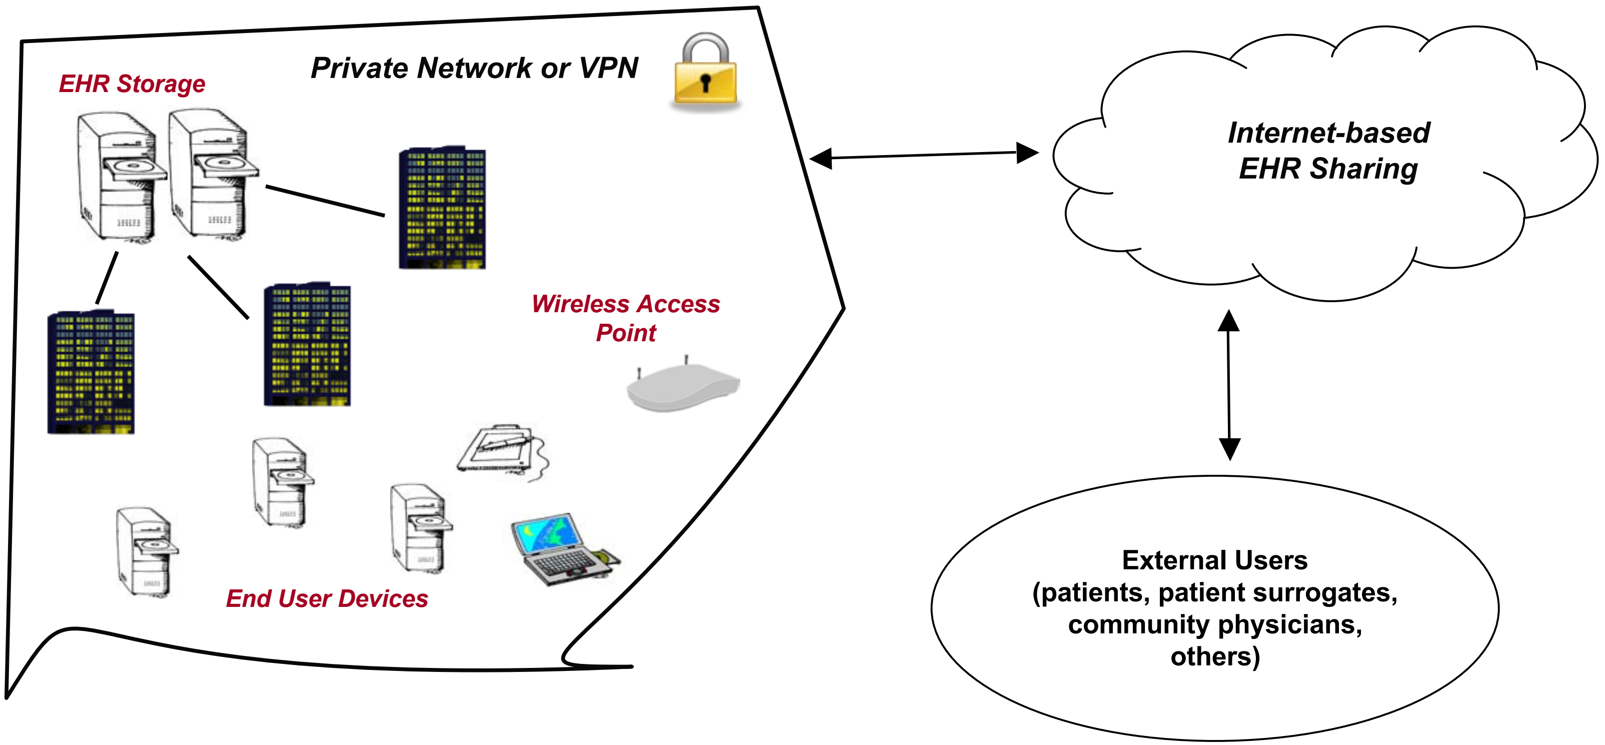
Private Network or VPN
EHR Storage
Wireless Access Point
End User Devices
Internet-based EHR Sharing
External Users (patients, patient surrogates, community physicians, others)

- Common patient uses: access to test results; communication with doctors, nurses; limited views of EHR info; prescription refills
- Currently, limited health information exchange, usually within a highly localized region and between systems from same vendor, is occurring

Future: Large-Scale Health Information Exchange (HIE)

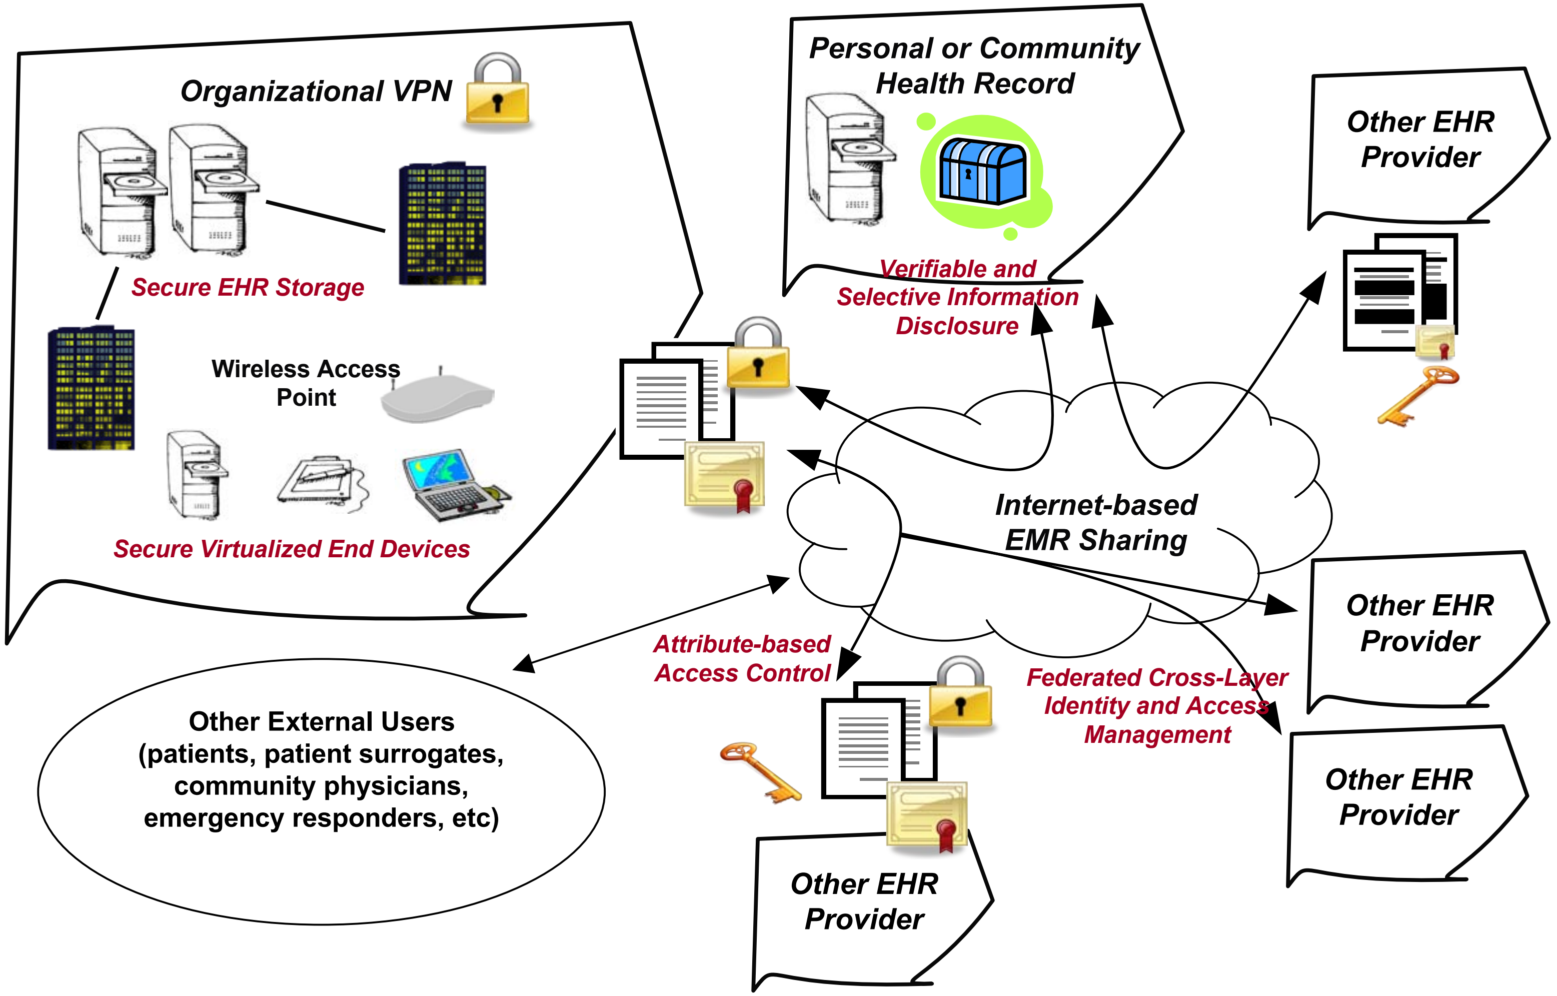
Organizational VPN
Secure EHR Storage
Wireless Access Point
Secure Virtualized End Devices
Personal or Community Health Record
Verifiable and Selective Information Disclosure
Other EHR Provider
Internet-based EMR Sharing
Other EHR Provider
Attribute-based Access Control
Federated Cross-Layer Identity and Access Management
Other EHR Provider
Other External Users (patients, patient surrogates, community physicians, emergency responders, etc)
Other EHR Provider

Nationwide Health Information Network (NHIN) Standard for HIE

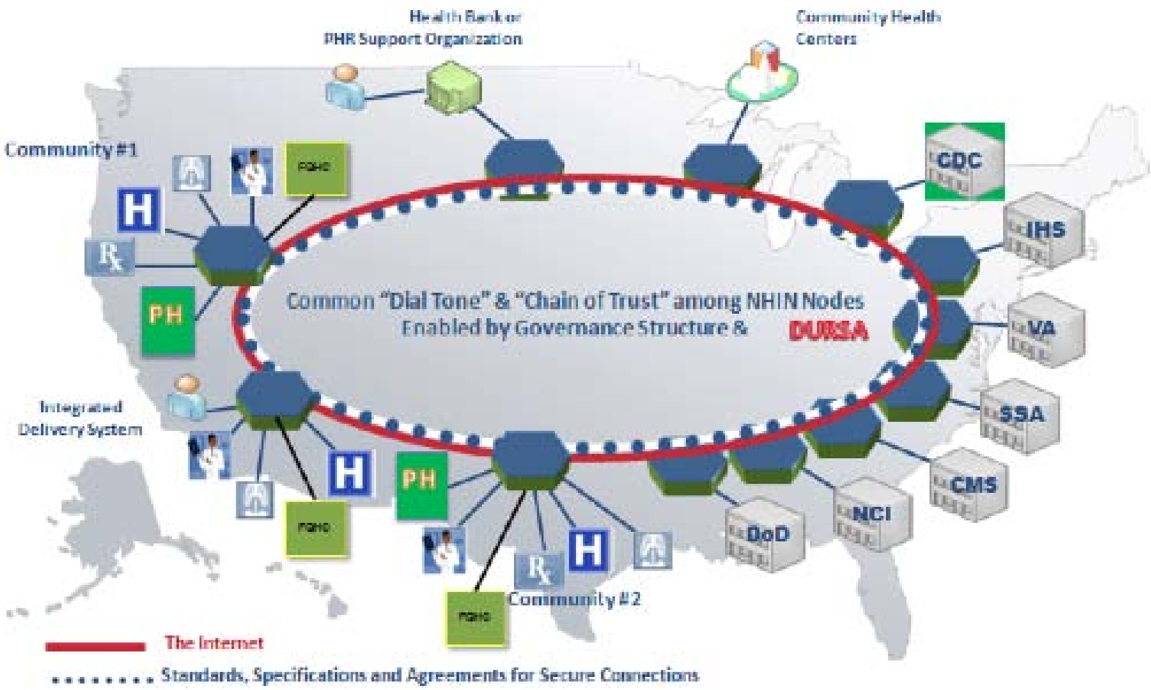
Health Bank or PHR Support Organization
Community Health Centers
Community #1
FQHC
H
Rx
PH
CDC
IHS
VA
SSA
CMS
NCI
DoD
Common "Dial Tone" & "Chain of Trust" among NHIN Nodes Enabled by Governance Structure & DURSA
Integrated Delivery System
Community #2
The Internet
Standards, Specifications and Agreements for Secure Connections

⬢ = NHIN Gateway: an instantiation of the NHIN technical specifications that supports secure, interoperable health information exchange across the NHIN

NHIN Standard, continued

- Based on Web services and using the Internet
- Includes specifications for authorization framework, messaging platform, patient discovery, document retrieval, document submission, access consent policies, etc.
- Use governed by Data Use and Reciprocal Support Agreement (DURSA), a multi-party-trust legal agreement based upon a set of policy assumptions that bridge varying state and federal laws and regulations
- Authorization framework uses SAML assertions to pass security tokens, XACML policies to describe access policies, and a XACML engine for authorization decisions
- Open source implementation – NHIN Connect

NHIN Standard, continued

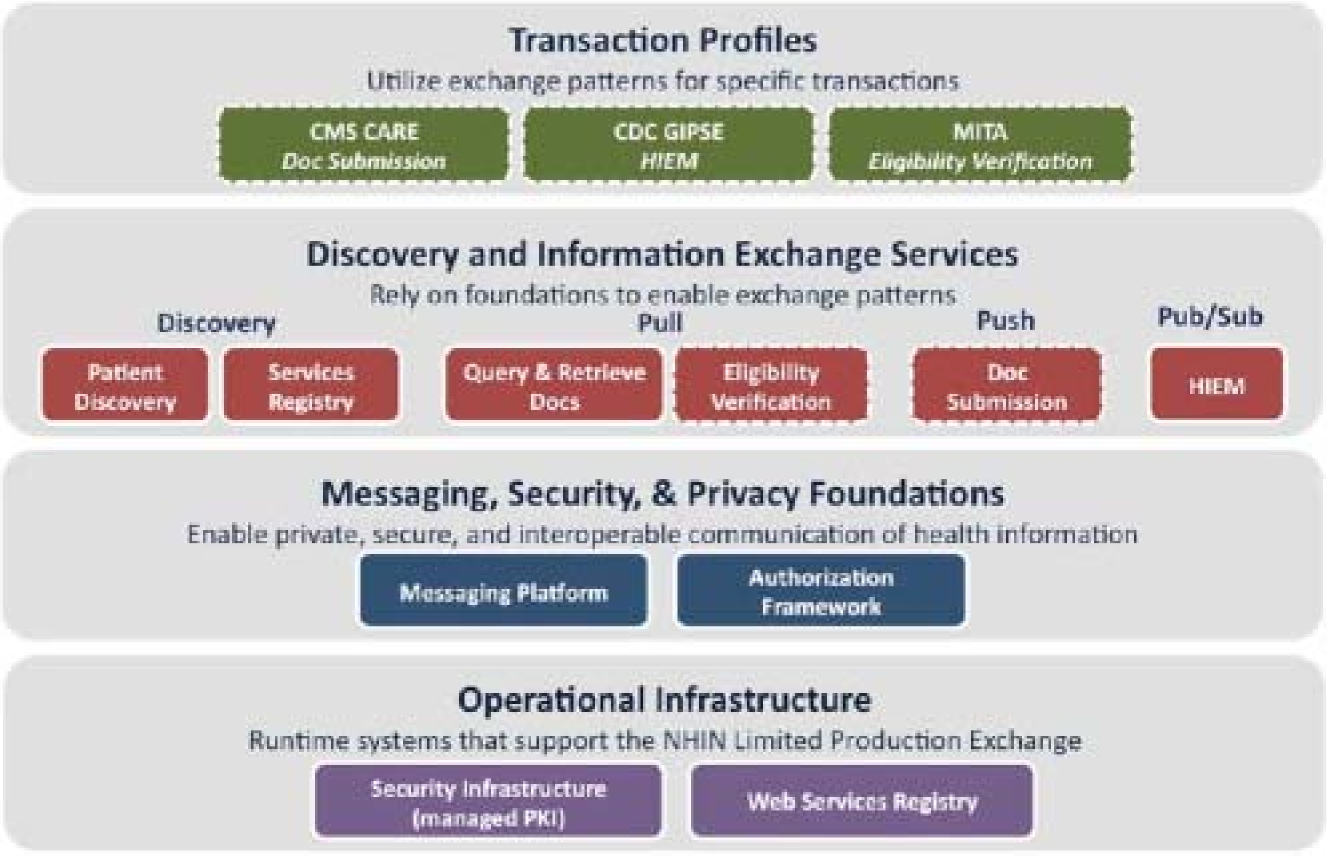

Transaction Profiles

Utilize exchange patterns for specific transactions

- CMS CARE *Doc Submission*
- CDC GIPSE *HIEM*
- MITA *Eligibility Verification*

Discovery and Information Exchange Services

Rely on foundations to enable exchange patterns

- Discovery: Patient Discovery; Services Registry
- Pull: Query & Retrieve Docs; Eligibility Verification
- Push: Doc Submission
- Pub/Sub: HIEM

Messaging, Security, & Privacy Foundations

Enable private, secure, and interoperable communication of health information

- Messaging Platform
- Authorization Framework

Operational Infrastructure

Runtime systems that support the NHIN Limited Production Exchange

- Security Infrastructure (managed PKI)
- Web Services Registry

NHIN Gateway

NHIN Infra-structure

NHIN Layers

Emerging Threat Examples

- ***Insider threat explosion:*** if closed system access model is extended to HIE, insiders will include most health care professionals in the US, public health officials, and many others
- ***Malware:*** in 2009, 300 Windows boxes controlling MRI machines in the US were found to be infected with the Conficker worm
- ***Monetization of stolen health information:*** in Oct. 2010, the largest Medicare fraud in history, carried out by a large organized crime syndicate, was announced
  - identities of thousands of Medicare beneficiaries and licensed physicians were stolen
  - perpetrators used stolen info to bill Medicare for services never provided
  - $163M in fraudulent billings through 118 phony clinics in 25 states

Examples of Security and Privacy Research Topics

- Transparent situationally-aware access control
- Anomaly detection
- Provenance – both for auditing and compliance
- Dynamic detection and resolution of policy conflicts
- Metadata privacy – inference attacks
- Identity management: both users (health care professionals) and non-users (patients)
- Digital consent
- Intrusion detection/tolerance, malware analysis and detection, etc.

Research Highlight

Trust and Control in Health Information Exchange

- HIE participants exchange patient info using HL7 standard Continuity of Care Documents (CCDs)

- CCDs are XML docs with standard format including sections for purpose, payers, advance directives, problems, family history, social history, alerts, medications, medical equipment, immunizations, vital signs, test results, procedures, etc.

- Example CCD from John Halamka available on his blog: http://geekdoctor.blogspot.com/

- CCDs are passed between providers with info filled in by different parties along the way

Trust and Control in HIE, continued

- ***Problems***
  1. How to know the source of particular info in the CCD and verify that it has not been modified by any other party?
  2. How to allow selective disclosure of CCD info to respect patient's privacy wishes?

- ***Motivating example***
  1. Parents send their daughter to religious summer camp, where they need to show proof of immunizations
  2. Parents get child's CCD including immunization record from their health care provider
  3. Parents forward CCD to camp but do not want to reveal any record of the child's HPV vaccination
  4. Camp receives partial CCD and needs to verify that it came from a recognized health care provider and was not modified

Redactable Signatures for Selective Disclosure

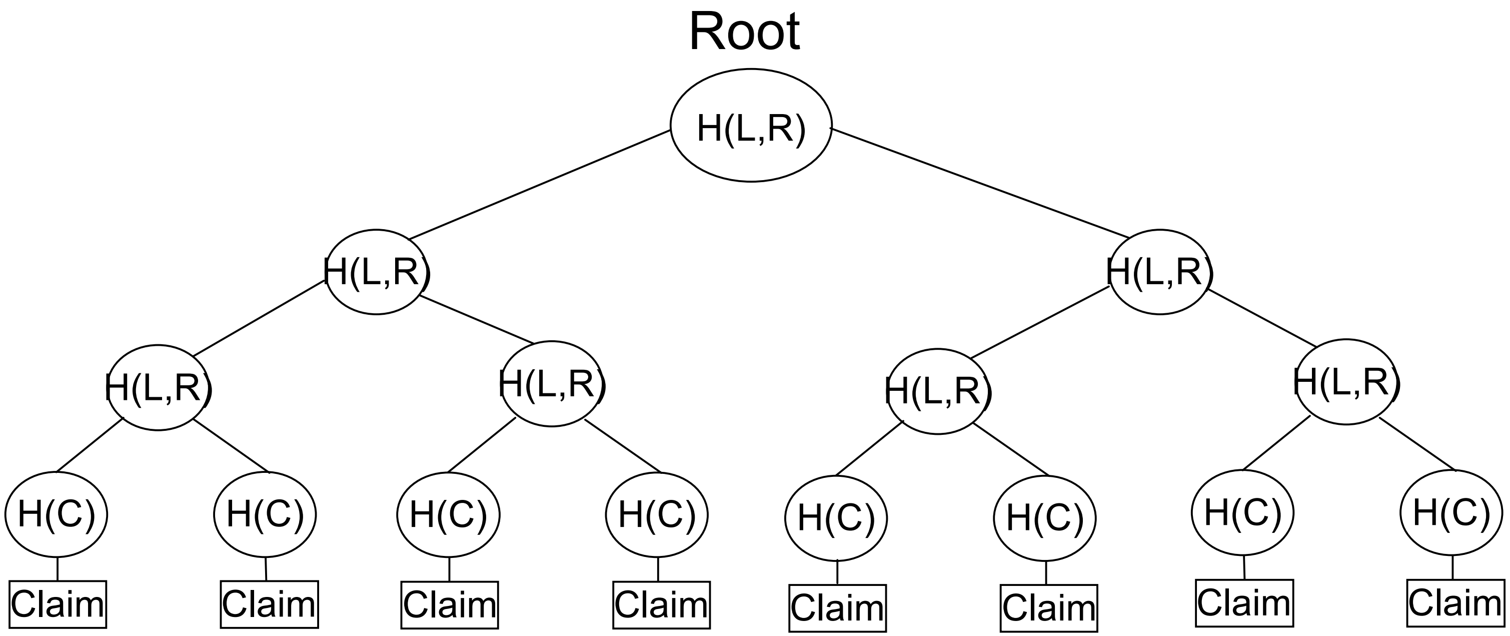

- sign root of hash tree
- release arbitrary subset of claims plus hash values necessary to reconstruct root hash
- Johnson, et. al., "Homomorphic Signature Schemes“, 2002.

Redactable Signatures for Health Info

- Multi-authority redactable signatures (*Bauer, Blough, and Cash, ACM CCS Workshop on Digital Identity Management, 2008*) allow a CCD to be signed by all providers that contribute to it

- Redactable signatures with disclosure dependencies (*Bauer, Blough, and Mohan, ACM CCS Workshop on Privacy in the Electronic Society, 2009*) allow providers to put limited constraints on how information can be hidden after they release a CCD

Multi-Authority Redactable Signatures

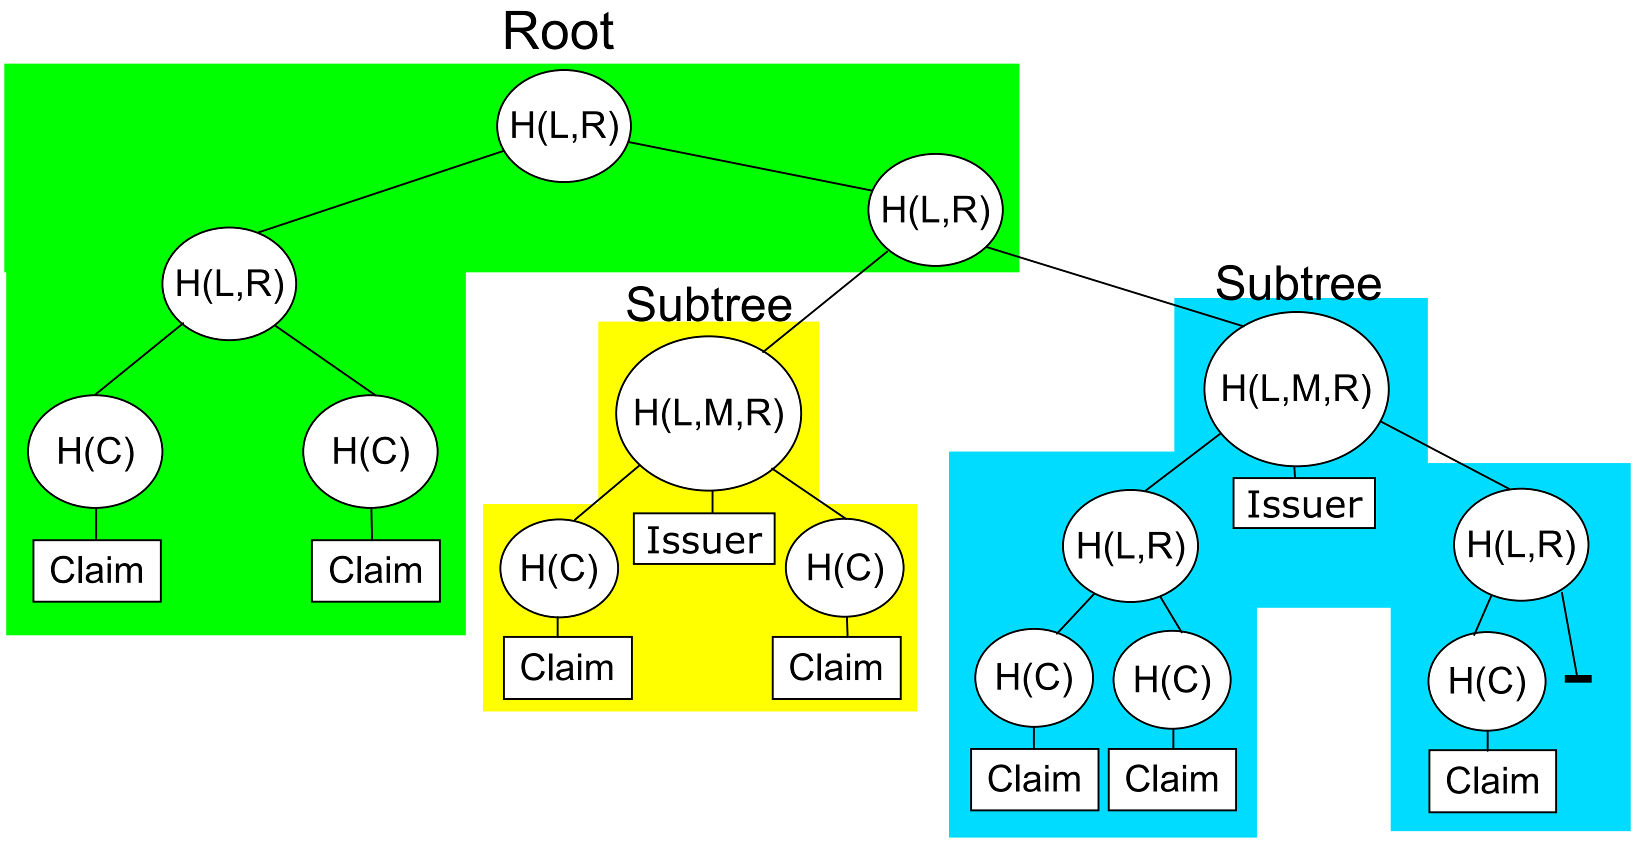

Root signatory certifies its own pieces of record and that subtree root hash values are correct (correctness of pieces in a subtree certified by its signatory)

Redactable Signatures with Disclosure Dependencies

- Motivation: provider of signed record, e.g. health care provider, wants to enforce some release constraints
- Example: do not release a diagnosis without also releasing the test result on which it was based
- Example: do not release medical image without its metadata (however, might want to release metadata without image for searching/browsing)

Allowable Subsets

- Release policies may allow many options
  - “Release A only if also releasing B or C” gives 4 options: “AB”, “AC”, “B”, and “C”
  - Adding the rule “Release B or C only if releasing D or E” gives 10 options

- Number of allowable subsets can be exponential (enumerating possibilities in hash tree is extremely inefficient)

Dependencies as Logical Expressions

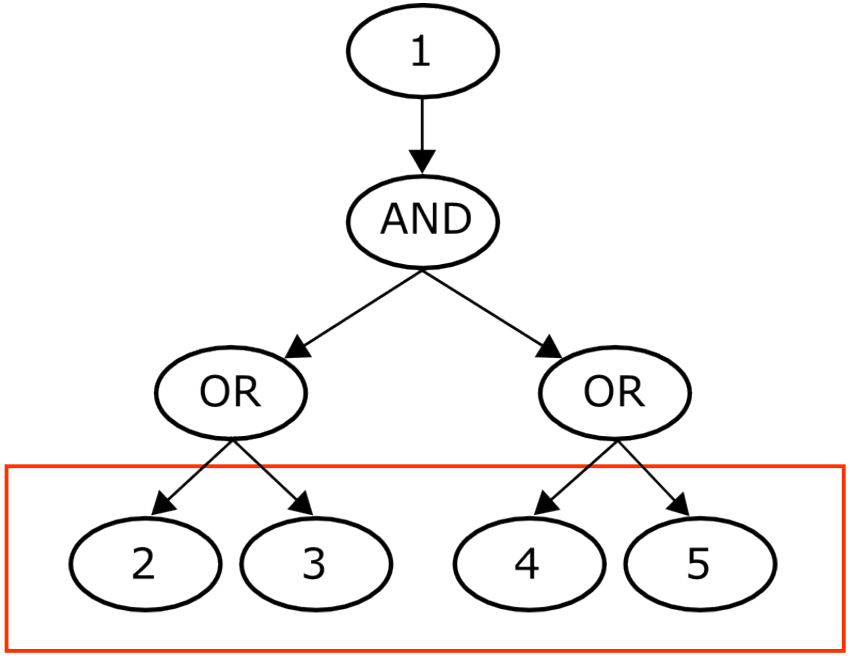

"1" cannot be released without also releasing either "2" or "3" along with either "4" or "5"

- Release policy is a **graph**
  - Each claim is a **node**
  - Each AND/OR is a node
  - No limit on fan-out or fan-in
- May have many top-level and bottom-level nodes
- Bottom (leaf) nodes are signed directly by a redactable signature
  - Other nodes are not

OR Dependencies

- Consider z → x OR y
- S(x) is called the string for node x
- S(x) = H(S(z) + x)
  - H is a hash function
  - "+" is concatenation
  - x is the actual data
  - S(z) is the string for node z
- S(y) = H(S(z) + y)
- S(z) = z

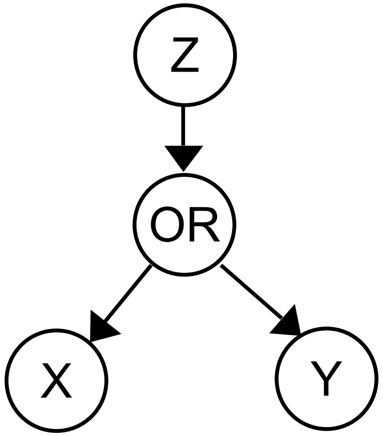

AND Dependencies

- OR Nodes disappear; AND nodes don’t
- AND nodes handled with secret sharing
- Consider z → w AND x AND y
- Generate random strings A1, A2
- S(AND) = H(S(z) + A1 + A2)
- A3 = S(AND) XOR A1 XOR A2
- S(w) = H(A1 + w)
- S(x) = H(A2 + x)
- S(y) = H(A3 + y)

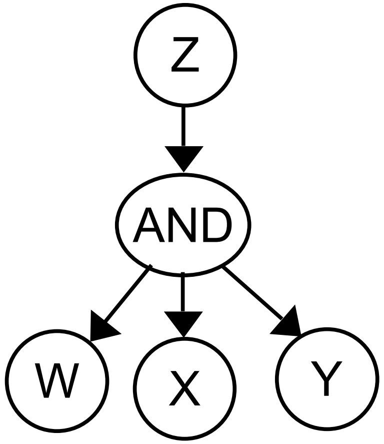

Prior Example

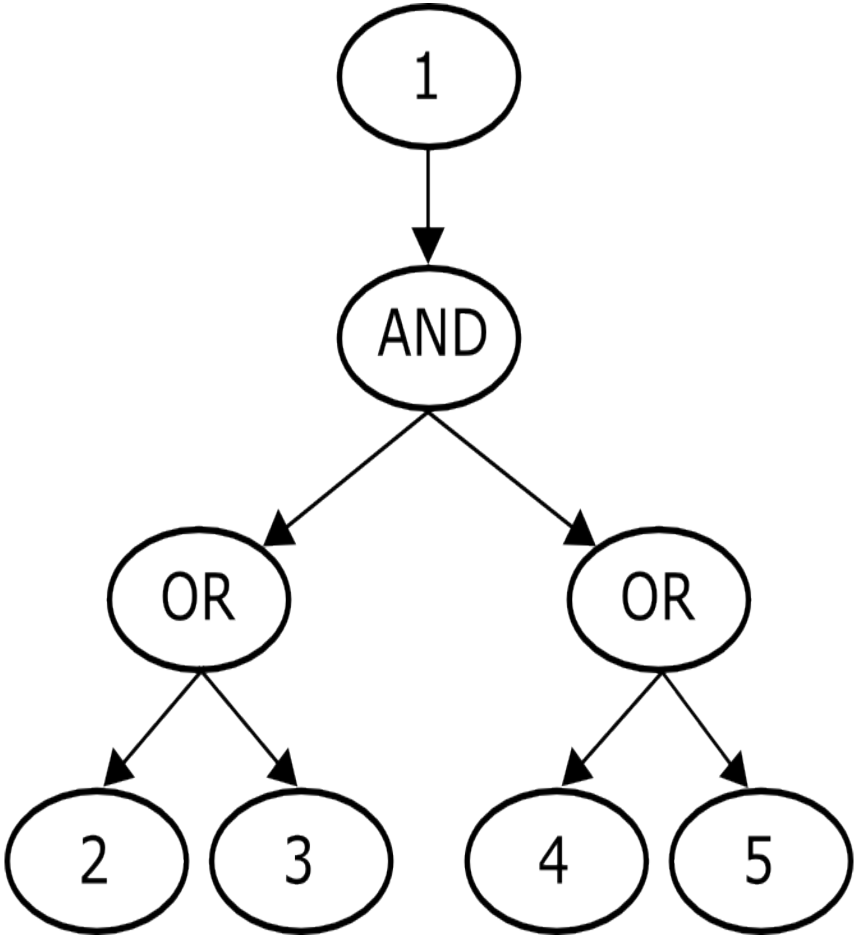

- S(1) = 1
- A1 = random string
- S(AND) = H(S(1) + A1)
- A2 = S(AND) XOR A1
- S(2) = H(A1 + 2)
- S(3) = H(A1 + 3)
- S(4) = H(A2 + 4)
- S(5) = H(A2 + 5)

Other Forms of Constraints

- Threshold functions, e.g. release A only if at least k out of n other claims are also released: can be handled by replacing xor secret sharing with Shamir's threshold secret sharing scheme

- Other functions - open problem (of course, can generate all allowable subsets and treat as a logical sum of products)

Other Application Areas

- Allow quoting from documents while preventing out of context quotes

- Verifying source tree signatures while ensuring that software module dependencies are honored (have done extensive testing on Ubuntu source trees)

- Any area where source signs data and wants to permit limited redaction once data leaves their sphere of control

A Patient-centric, Source-verifiable Health Records Repository

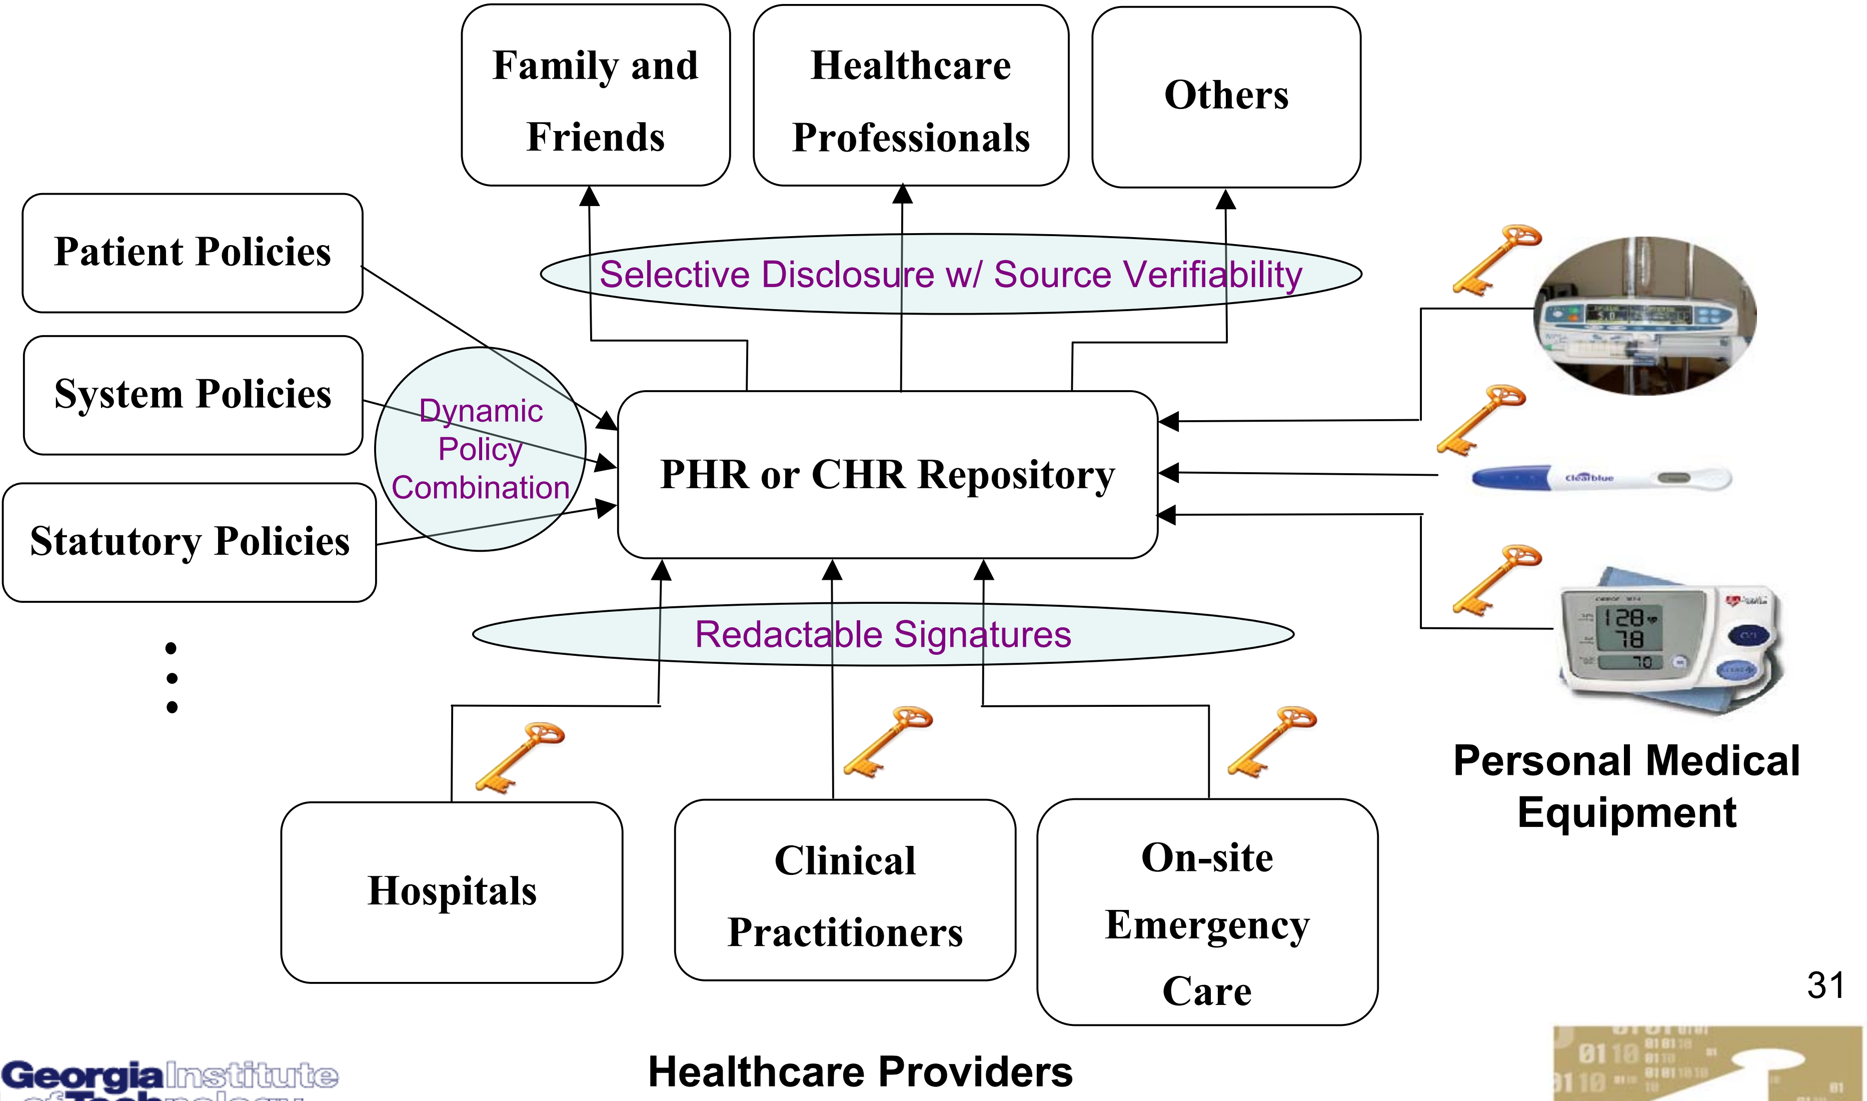

Family and Friends

Healthcare Professionals

Others

Selective Disclosure w/ Source Verifiability

Patient Policies

System Policies

Statutory Policies

⋮

Dynamic Policy Combination

PHR or CHR Repository

Redactable Signatures

Hospitals

Clinical Practitioners

On-site Emergency Care

Healthcare Providers

Personal Medical Equipment

Questions??

Backup Slides

Dynamic Attribute-Based Authorization

- In HIE, authorization is subject to multiple policies and must be highly dynamic to account for different policies that depend on dynamic context, e.g. normal or emergency situation, access for treatment or research, etc.

- We proposed new dynamic policy combination mechanisms for these situations (*Mohan and Blough, IDTrust 2010*) and integrated them into Sun's open-source XACML authorization engine

- Policies are specified according to dynamic attributes, e.g. emergency or non-emergency situation, location of patient and/or accessing entity, purpose of request, etc.

Architecture for Dynamic Authorization

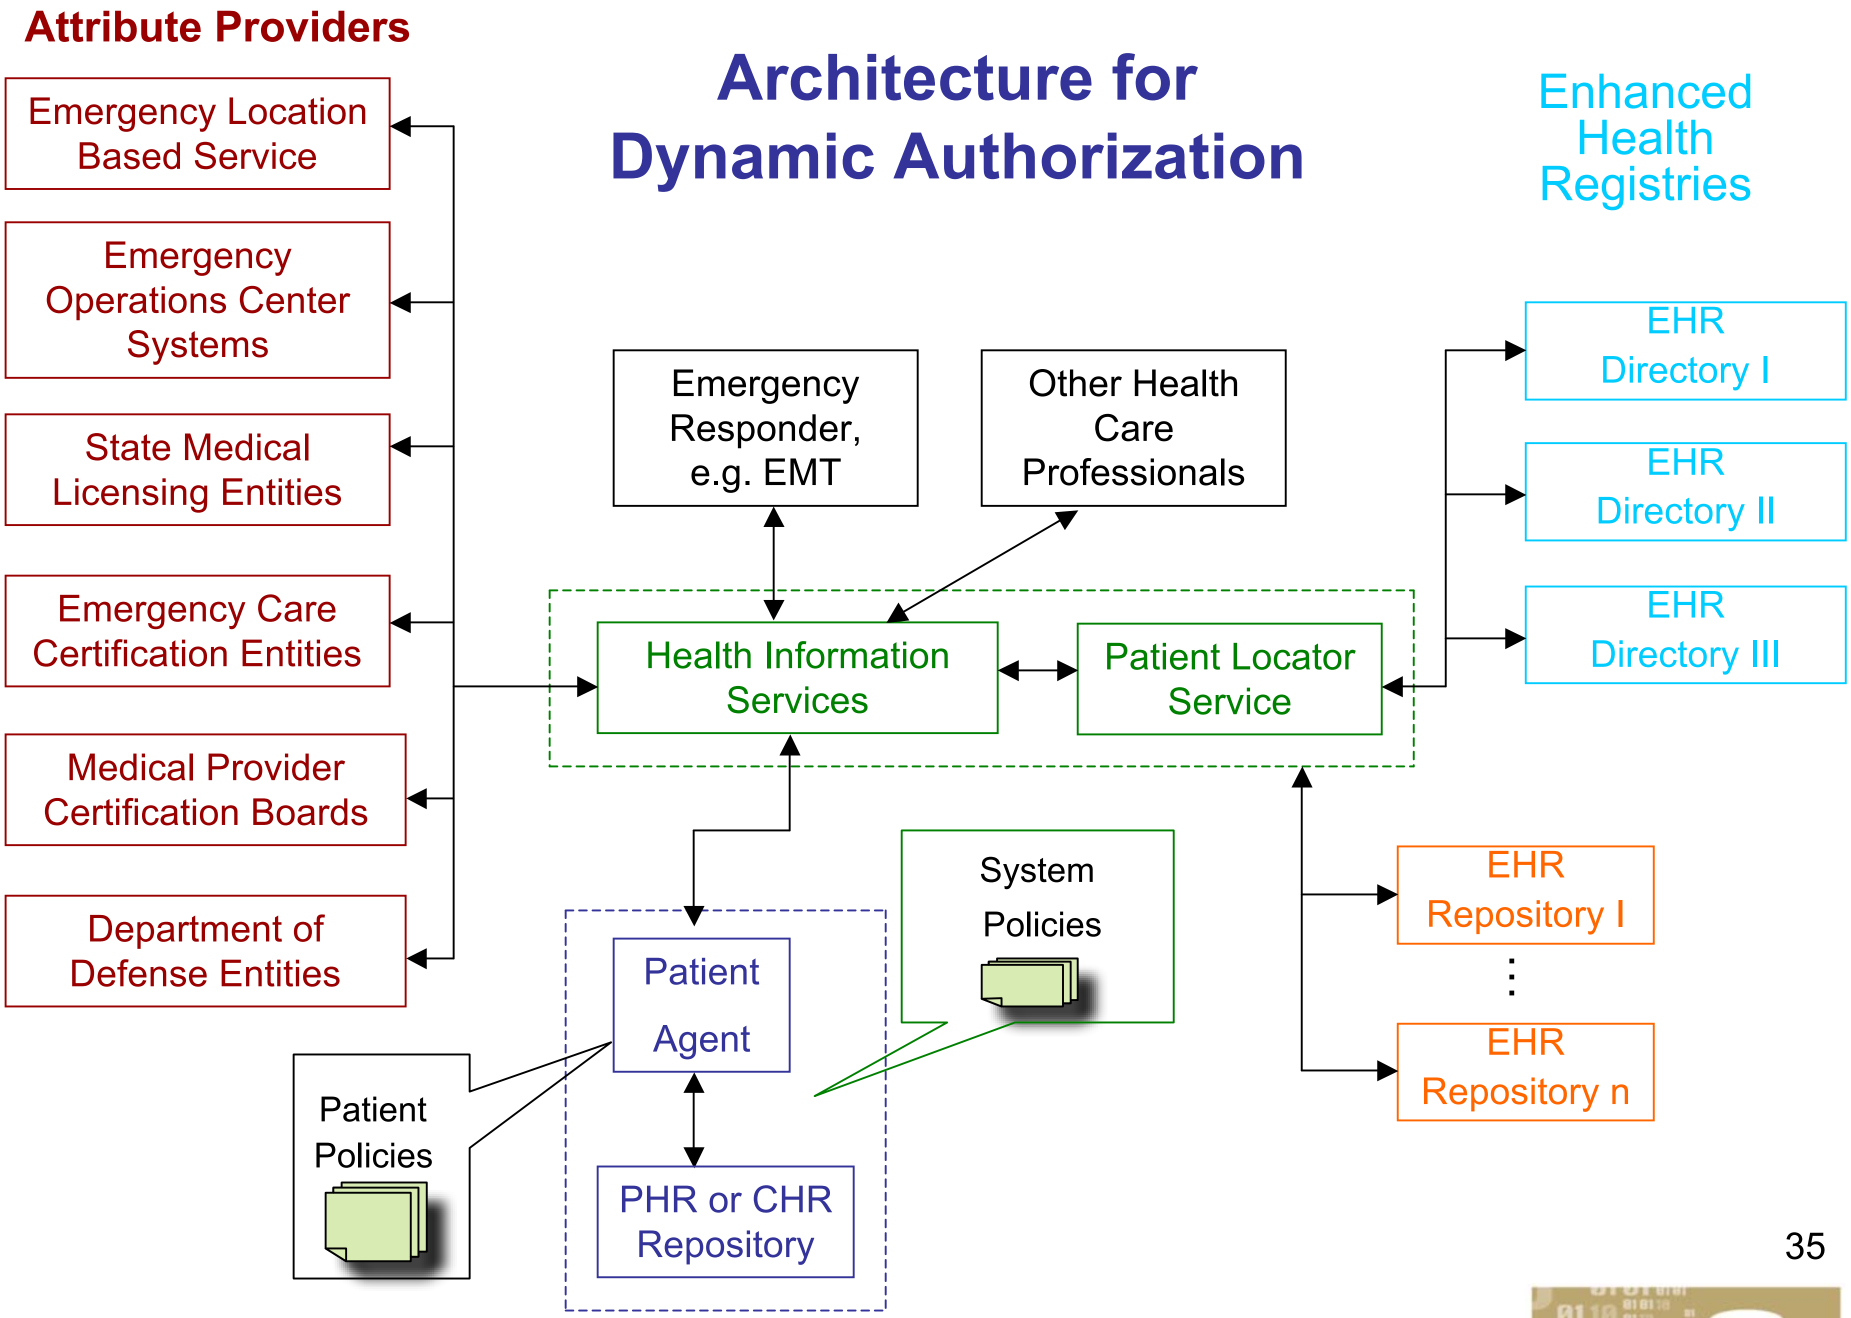

Attribute Providers
Emergency Location Based Service
Emergency Operations Center Systems
State Medical Licensing Entities
Emergency Care Certification Entities
Medical Provider Certification Boards
Department of Defense Entities
Emergency Responder, e.g. EMT
Other Health Care Professionals
Health Information Services
Patient Locator Service
Enhanced Health Registries
EHR Directory I
EHR Directory II
EHR Directory III
EHR Repository I
⋮
EHR Repository n
System Policies
Patient Agent
Patient Policies
PHR or CHR Repository