

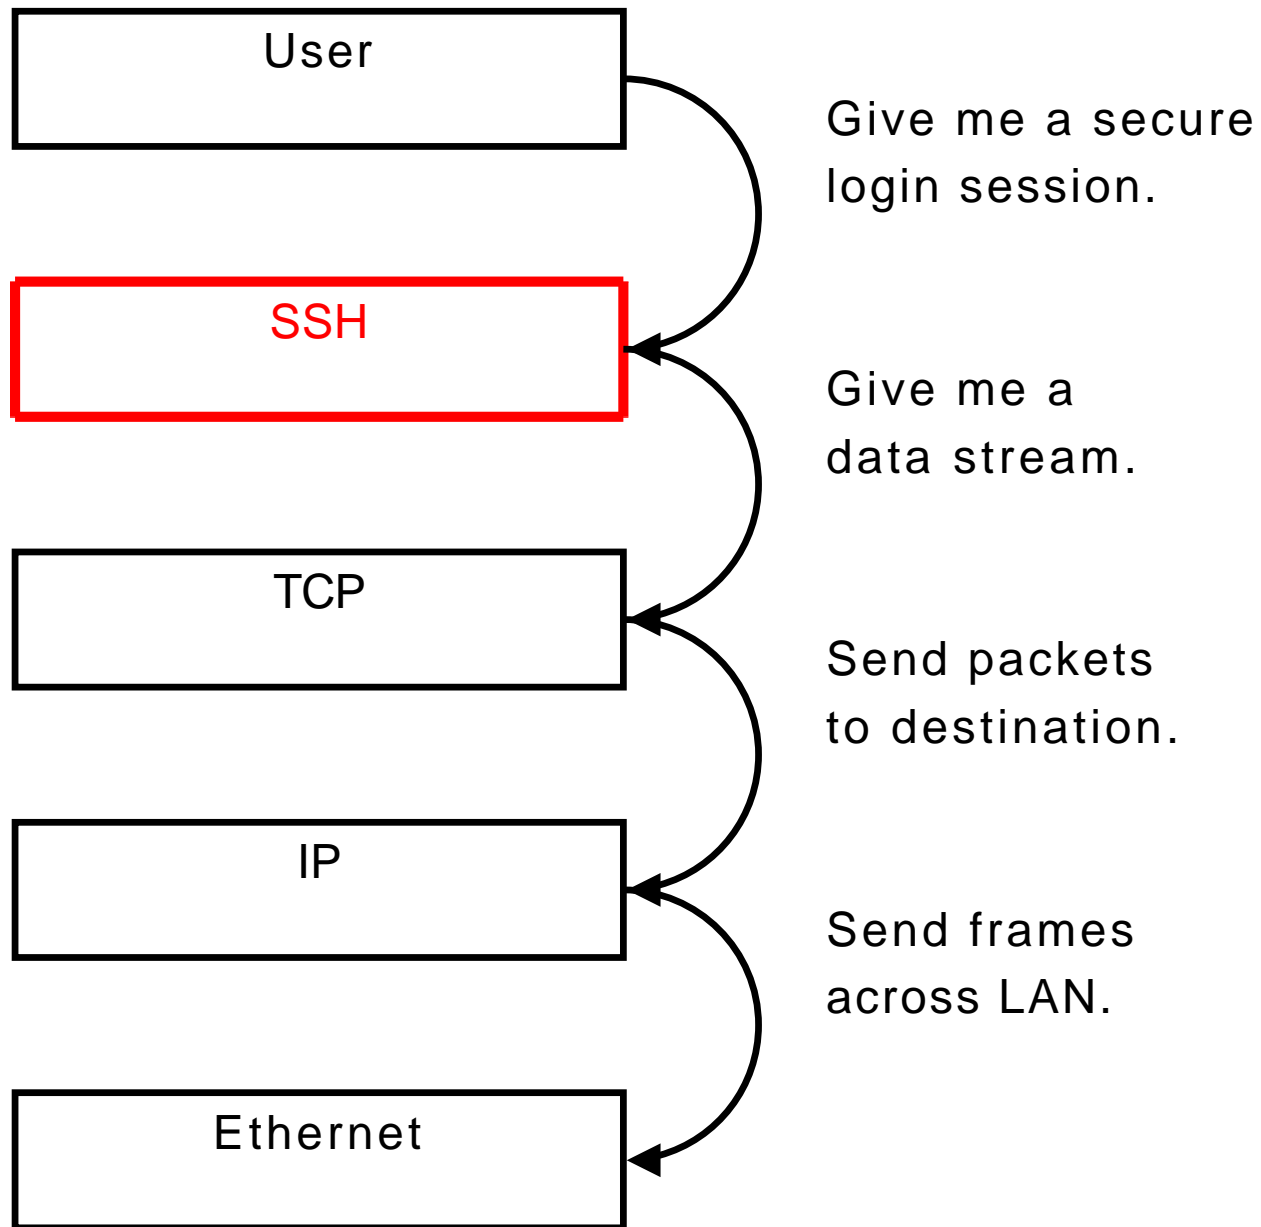
# **Network Security, Part II**

**Jeff King**

**July 3, 2007**

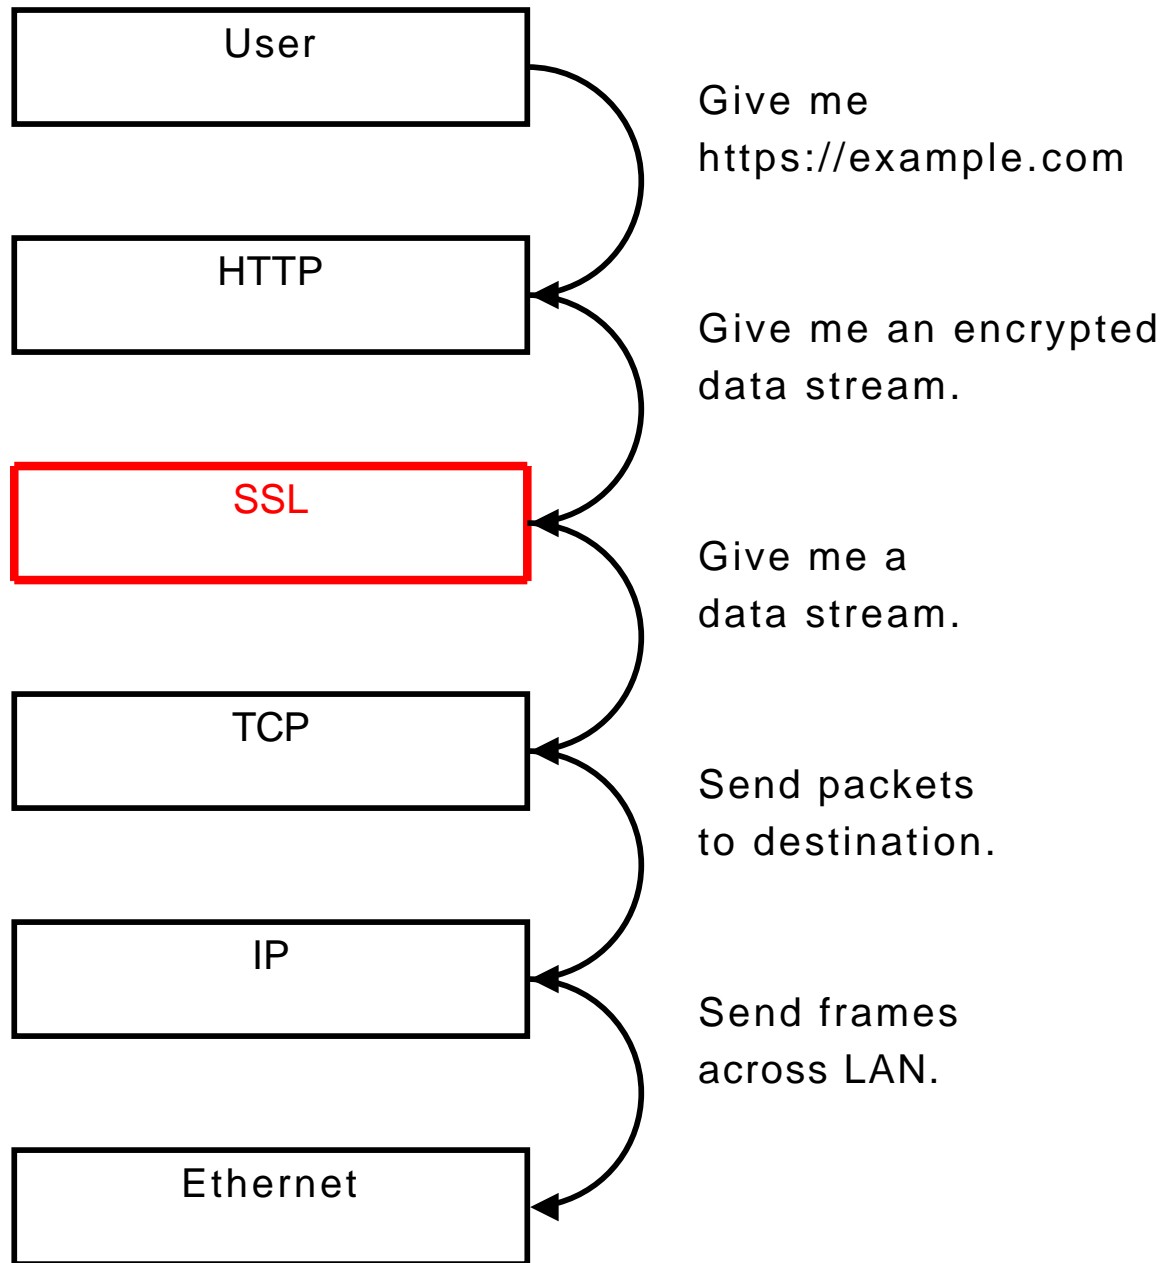
# Cryptography

- Encryption of data as it passes over the network
  - Can provide confidentiality (encryption) and integrity (keyed hash)
  - Can provide authenticity (PKI)
- Encryption can happen at different network layers
- Some application protocols are tied to encryption (SSH)



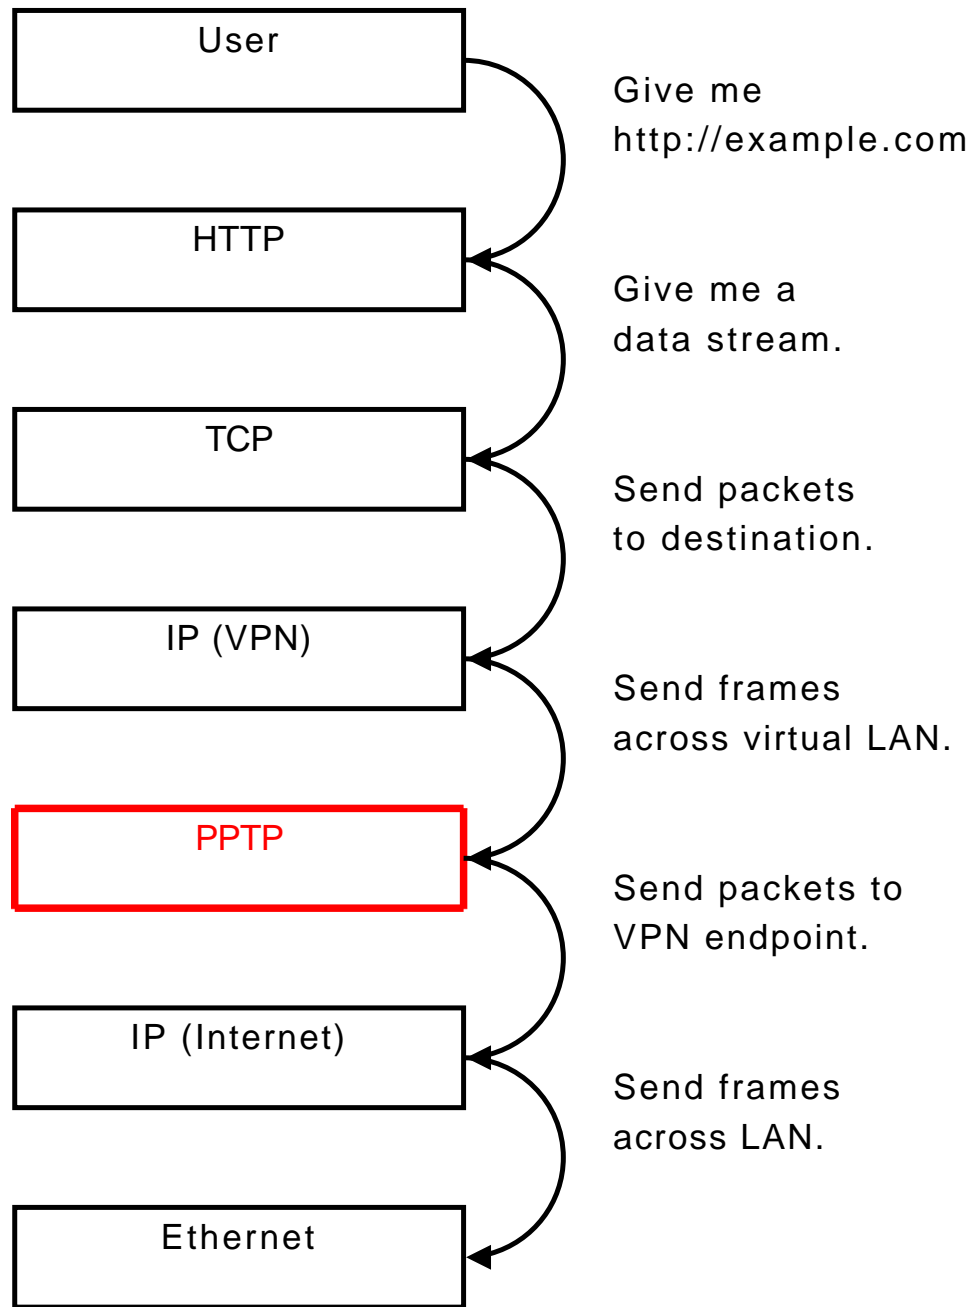
# Cryptographic Tunnels

- Some protocols provide an encrypted “tunnel” at the transport layer
- Secure Socket Layer (SSL), Transport Layer Security (TLS)
- HTTPS, SMTP/POP/IMAP STARTTLS, etc
- Pro: bolts onto existing protocols
- Pro: authentication information is closely associated with session
- Con: two versions of each protocol
- Con: assumes TCP-based protocol



## Cryptographic Tunnels, Take 2

- Q: What happens when we move the tunnel down a layer?
- A: Virtual Private Networks (VPNs)



## IPsec

- Build encryption into the network layer
- Extension to IPv4, mandatory in IPv6
- Transport Mode
  - Payload is encrypted, headers are not
- Tunnel Mode
  - Point-to-point communications
  - IP encapsulated inside IP
  - Entire packet is encrypted

## Public Key Infrastructure (PKI)

- Now that we have all this encryption...
- ...how do we know we're talking to the right endpoint?

## What is PKI?

- Fundamentally: a binding of identity to some bits
- What is an identity?
  - Person? Business? Role? It depends on policy and goals.
- What is a binding?
  - Connects two (or more) things.
  - Implies some authority that is hard to forge.
- Why do we want it?
  - Make sure we encrypt **to** the right entity
  - Correctly verify signatures **from** other entities
  - Concepts used everywhere: tunneling protocols, secure e-mail, web e-commerce, etc

## What do we want out of a good PKI?

- Secure (hard to forge binding)
- Flexible (can specify many aspects of identity)
- Efficient (online vs offline verification)
- Adaptable (changes over time)
- Scalable (handles many users)
- Easy to use (lock icon in browser)

## How do we do PKI?

- Local caching (ssh)
- Out-of-band methods (ssh, PGP, DomainKeys)
- Hierarchical (X.509, DNSSEC)
- Web of trust (PGP)

## Local caching

```
$ ssh example.com
```

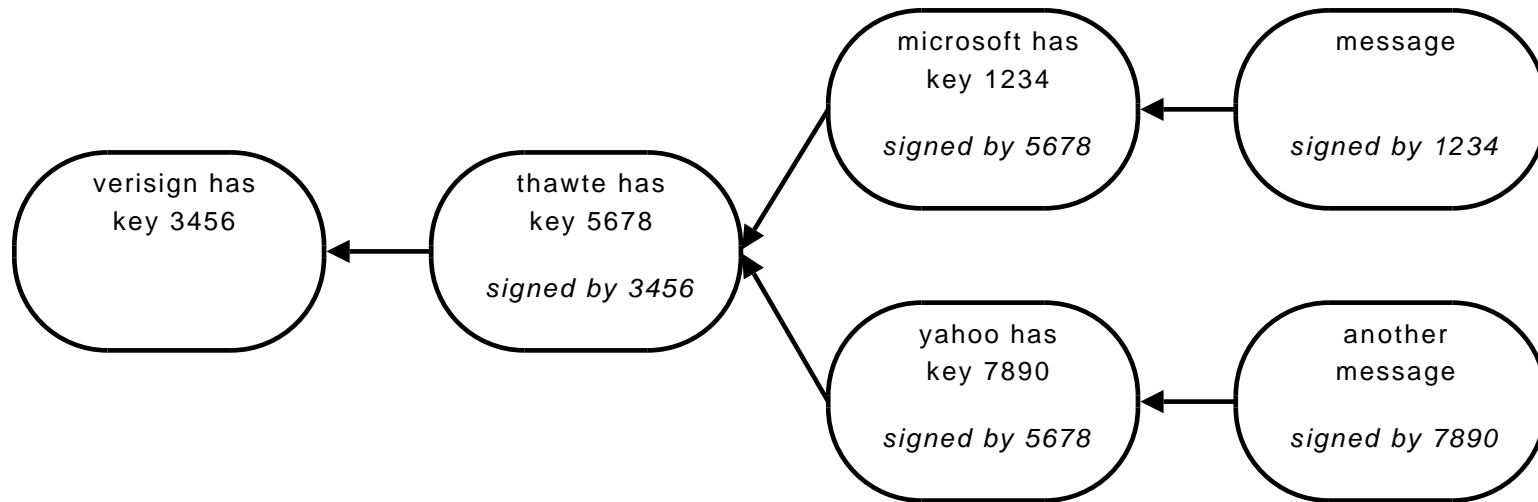
```
The authenticity of host 'example.com (10.0.0.1)' can't be established.  
DSA key fingerprint is 5f:ff:23:0d:be:dc:c2:33:e2:13:a6:62:9f:65:9d:cc.  
Are you sure you want to continue connecting (yes/no)?
```

- Warn the user during the first session.
- Aggressively cache (read: never delete) public keys.
- Check cached keys and abort if mismatch.
- Secure? Flexible? Efficient? Scalable? Adaptable? Easy to use?

## Out-of-band Methods

- ssh – verify key fingerprint by email/phone
- ssh – seed public keys from disk
- PGP – key signing parties (check government id!)
- DomainKeys – DNS hierarchy of keys
  
- Secure? Flexible? Efficient? Scalable? Adaptable? Easy to use?

## Hierarchical Methods



- Vocabulary word: **certificate** – a signature that binds an identity (and some authorization) to a public key
- Chain of trust going to root
- Standards of practice for certificate authorities (CAs)
- Why do we trust the root?
- Secure? Flexible? Efficient? Scalable? Adaptable? Easy to use?

## Web of Trust

- Imagine a hierarchy tree without a root
- Alice knows Bob, and Bob knows Carol, so Bob introduces Alice to Carol (cryptographically)
- Bob tells Alice how well he knows Carol; Alice decides how much to trust the introduction
- Introductions are transitive (but weaken with distance)
- Secure? Flexible? Efficient? Scalable? Adaptable? Easy to use?

## Why is PKI hard?

- Trust enforcement is suspect
  - Verification procedures are not standardized
  - “Hello, this is Microsoft; we need some new certificates...”
- Trust isn't infinitely transitive
  - Complex policy rules for who can sign whom
- Weak links in trust chains (single point of failure)
  - Break/steal root key, break them all
  - Trust Verisign's handling

## Why is PKI hard?

- Identities == names?
  - Names aren't always unique ("Jeff King" the baseball player, TV producer, and Iditarod champion)
  - Entities use many names; which are legitimate?  
(bankofamerica.com, bankofamerica.com.hk, bankofamerica.cx)
  - Unambiguous names don't always last forever (yearly domain registrations)

## Why is PKI hard?


- Revocation
  - Sign revocation with cert key (or pre-sign with private key)
  - Revocation lists must be checked online
  - Require negative response from list (replay issues)
  - Scalability issues

## Why is PKI hard?

- User issues
  - Trust decisions are hard to make
  - Key storage / recovery issues
- Political issues
  - Export regulations
  - Crypto usage regulations in some countries
  - Carnivore, wiretapping
  - Has the US government given up on crypto restrictions?

# Firewalls

- Fundamentally: Keep bad packets out, let good ones in.
- How can we tell them apart?
  - Source / Destination address, ports
  - Header fields
  - Application layer reassembly
- How do we create rules?
  - By hand
  - Rudimentary learning



## Why are firewalls good?


- Remember that networks are automated; they usually try hard to move data around
- Firewalls allow us to change the default-pass behavior
- Why not just turn off listening services?
  - Centralize policy in one place for ease of management
  - Protect poorly-maintained systems
  - Firewalled systems are not necessarily under the same administrative control

## What is Intrusion Detection?

- Intrusion – a set of actions aimed to compromise security goals
- Intrusion detection – the process of identifying and responding to intrusion activities

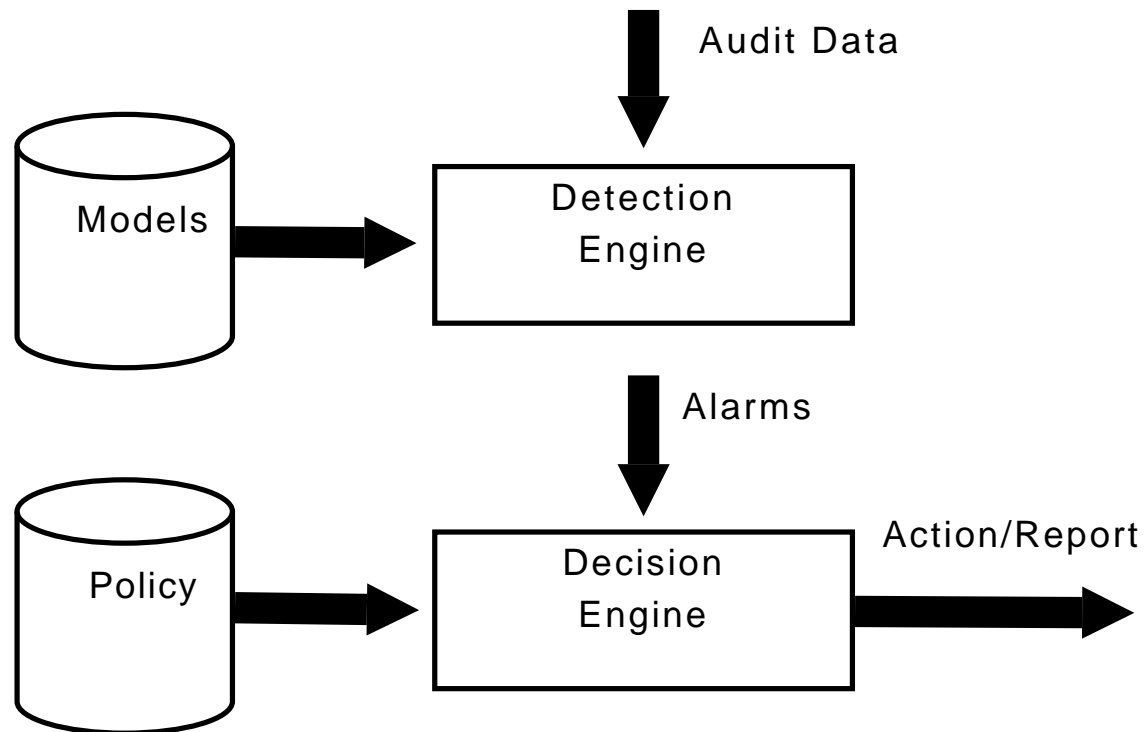
## Why do we need intrusion detection?

- Layered security
- In practice, firewalls will not block all malicious traffic
- There is value in knowing about malicious traffic on network



## How does intrusion detection work?

- Assume: system activities are observable
- Assume: normal and intrusive events have distinct evidence



## How does intrusion detection work?

- Misuse detection (signature based)
  - Hand-code model of bad behavior
  - Can't detect new attacks
- Anomaly detection (statistical based)
  - Compare behavior to normal (trained) profile
  - Higher false positive rate

## How do we rate intrusion detection systems?

- Alarm: A; Intrusion: I
- Detection rate:  $P(A|I)$
- False negative: an attack you fail to notice,  $P(\neg A|I)$
- False positive: an alarm when there is no attack,  $P(A|\neg I)$
- Bayesian detection rate:  $P(I|A)$
- Base-rate fallacy: Even if false alarm rate is very low, Bayesian detection rate is still low if the base-rate  $P(I)$  is low
- E.g., if  $P(A|I) = 1$ ,  $P(A|\neg I) = 10^{-5}$ ,  $P(I) = 2 \times 10^{-5}$ , then  $P(I|A) = 66\%$