Distributed Data Integrity Assurance and Repair Using the LOCKSS Content Audit Protocol (LCAP)

Thib Guicherd-Callin
LOCKSS Program, Stanford University

Outline

1. Context and Use Cases
2. Threat Models
3. LCAP in Action
4. Unlocking LOCKSS for Developers
5. Q&A

Outline

1. Context and Use Cases
2. Threat Models
3. LCAP in Action
4. Unlocking LOCKSS for Developers
5. Q&A

Traditional Research Libraries

- Ownership model
- Many independent replicas
- Features
  - Disaster resistance
  - Disaster recovery
  - Tamper evident
  - Permanent access

Research Libraries in the Web Era

- Leasing model
- One master copy
- Misfeatures
  - Disaster resistance?
  - Disaster recovery?
  - Tamper evident?
  - Permanent access?

LOCKSS Technology Use Cases

- "Lots Of Copies Keep Stuff Safe"
- Global LOCKSS Network (GLN)
- CLOCKSS Archive
- Government documents networks
- Regional and national networks

Key Publications

- "Founding paper"
  - David S.H. Rosenthal, Vicky Reich. "Permanent Web Publishing." Proceedings of the 2000 USENIX Annual Technical Conference FREENIX Track, pg. 129-140, 2000. URL: `https://www.usenix.org/legacy/publications/library/proceedings/usenix2000/freenix/rosenthal.html`

Key Publications

- "Protocol paper"
  - Petros Maniatis, Mema Roussopoulos, TJ Giuli, David S.H. Rosenthal, Mary Baker, and Yanto Muliadi. "Preserving Peer Replicas By Rate-Limited Sampled Voting." Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles (SOSP '03), pg. 44-59, 2003. DOI: 10.1145/945445.945451
  - Petros Maniatis, Mema Roussopoulos, TJ Giuli, David S.H. Rosenthal, Mary Baker, and Yanto Muliadi. "LOCKSS: A Peer-To-Peer Digital Preservation System." Technical report cs.CR/0303026, Stanford University, 2003. URL: http://www.eecs.harvard.edu/~mema/publications/SOSP2003-long.pdf

Key Publications

- "Threat models paper"
  - David S.H. Rosenthal, Thomas S. Robertson, Tom Lipkis, Vicky Reich, Seth Morabito. "Requirements for Digital Preservation Systems: A Bottom-Up Approach." D-Lib Magazine, vol. 11, iss. 11, November 2005. DOI: 10.1045/november2005-rosenthal

Outline

1. Context and Use Cases
2. Threat Models
3. LCAP in Action
4. Unlocking LOCKSS for Developers
5. Q&A

Goal of Digital Preservation

The goal of a digital preservation system is that the information it contains remains accessible to users over a period of time much longer than the lifetime of individual storage media, hardware and software components

Key Properties

- No single point of failure
- Media, hardware and software flow through as they fail or are replaced
- Regular audits frequent enough to keep probability of irrecoverable failure acceptable

Threat Taxonomy (1)

- Media failure
- Hardware failure
- Software failure
- Communication errors
- Failure of network services
- Natural disaster

Threat Taxonomy (2)

- Media and hardware obsolescence
- Software obsolescence

Threat Taxonomy (3)

- Operator error
- Economic failure
- Organizational failure

Threat Taxonomy (4)

- External attack
- Internal attack

Outline

1. Context and Use Cases
2. Threat Models
3. LCAP in Action
4. Unlocking LOCKSS for Developers
5. Q&A

Basic Principle

X

The peers hold identical replicas of X
Peer P1 calls a poll on content X

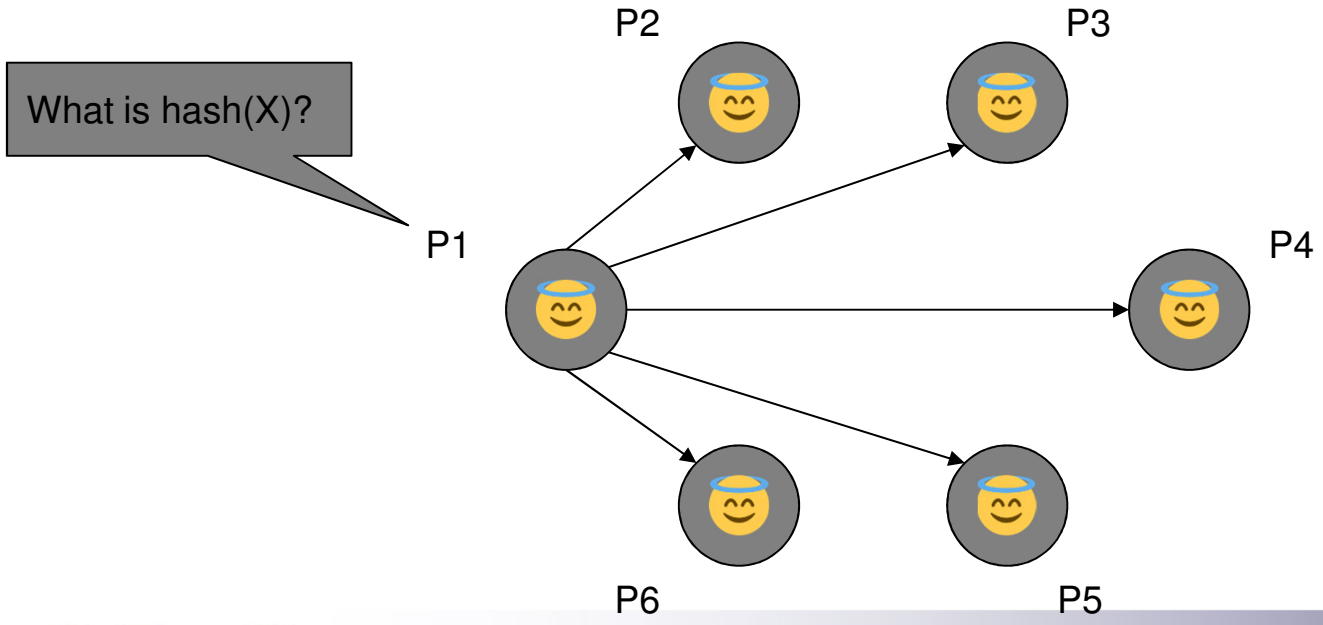

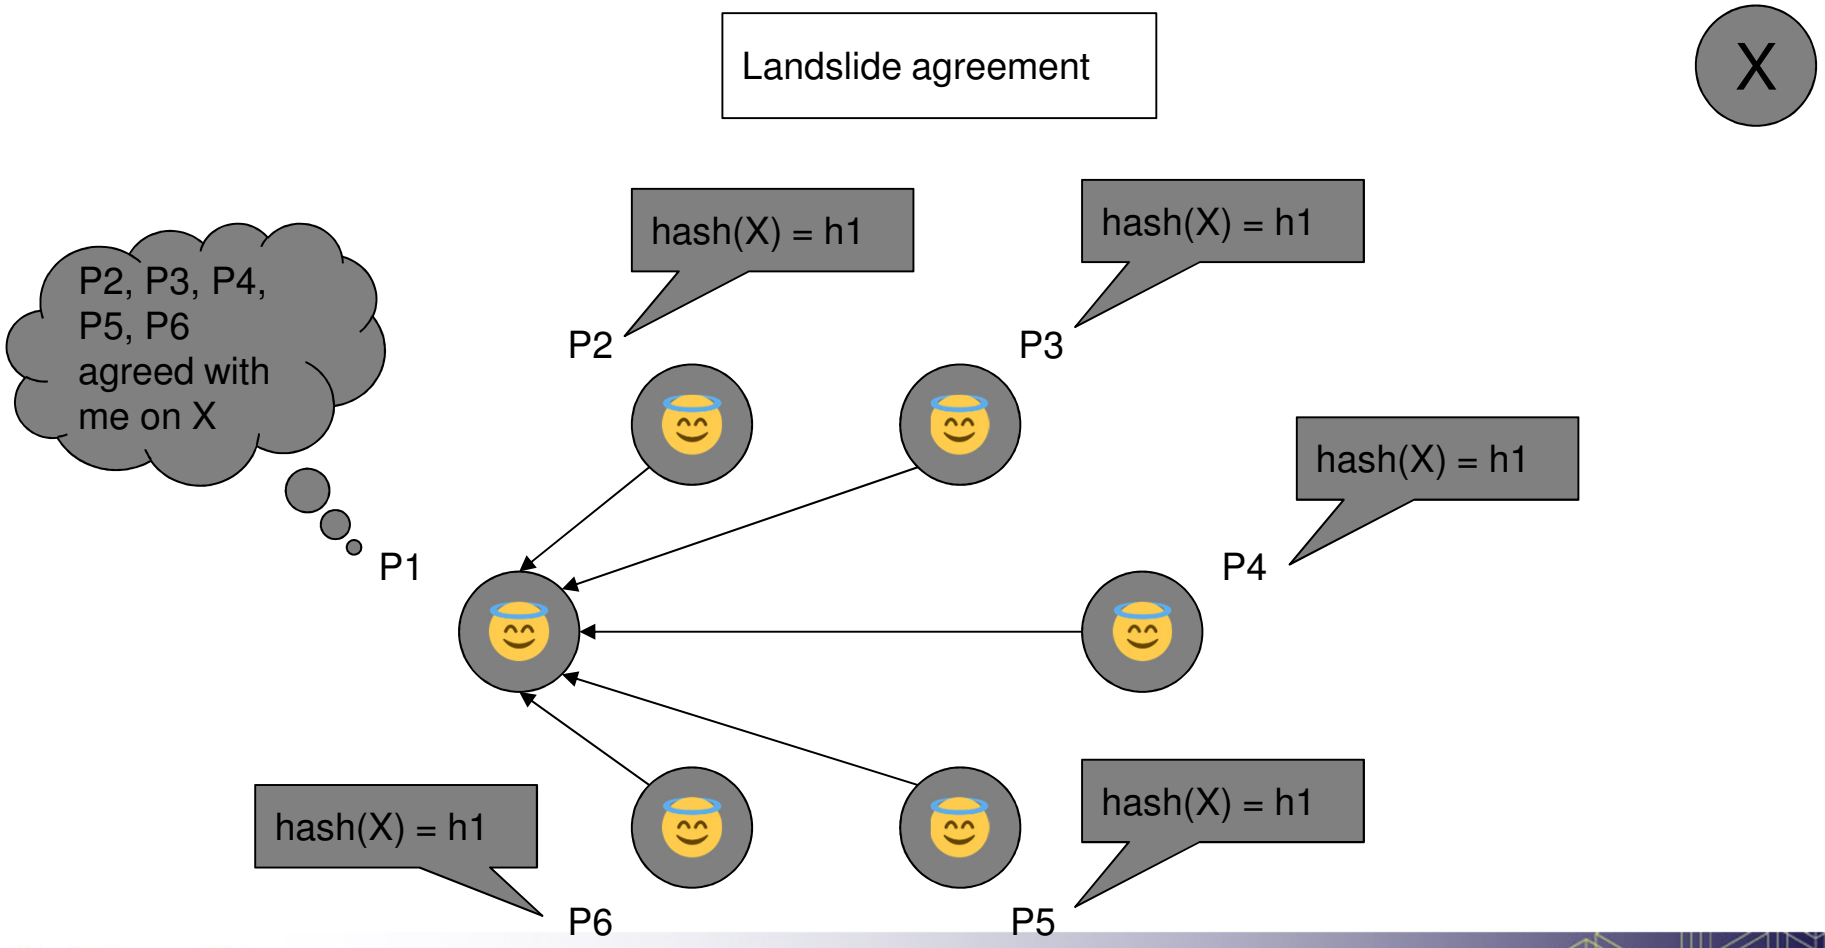

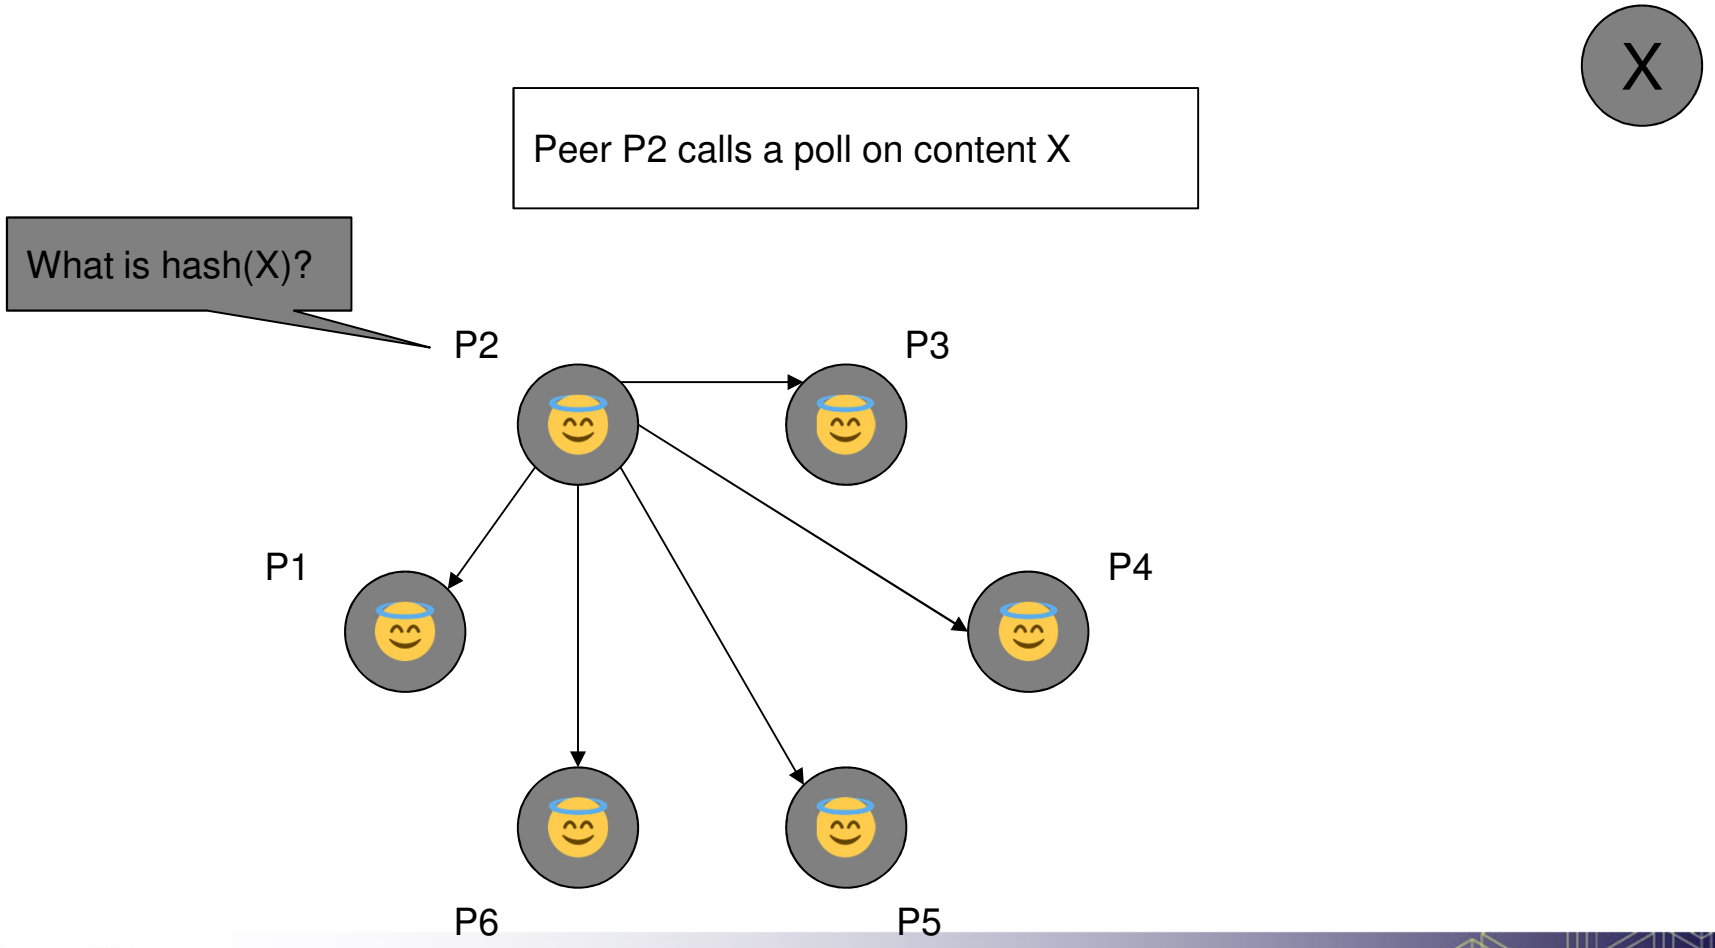

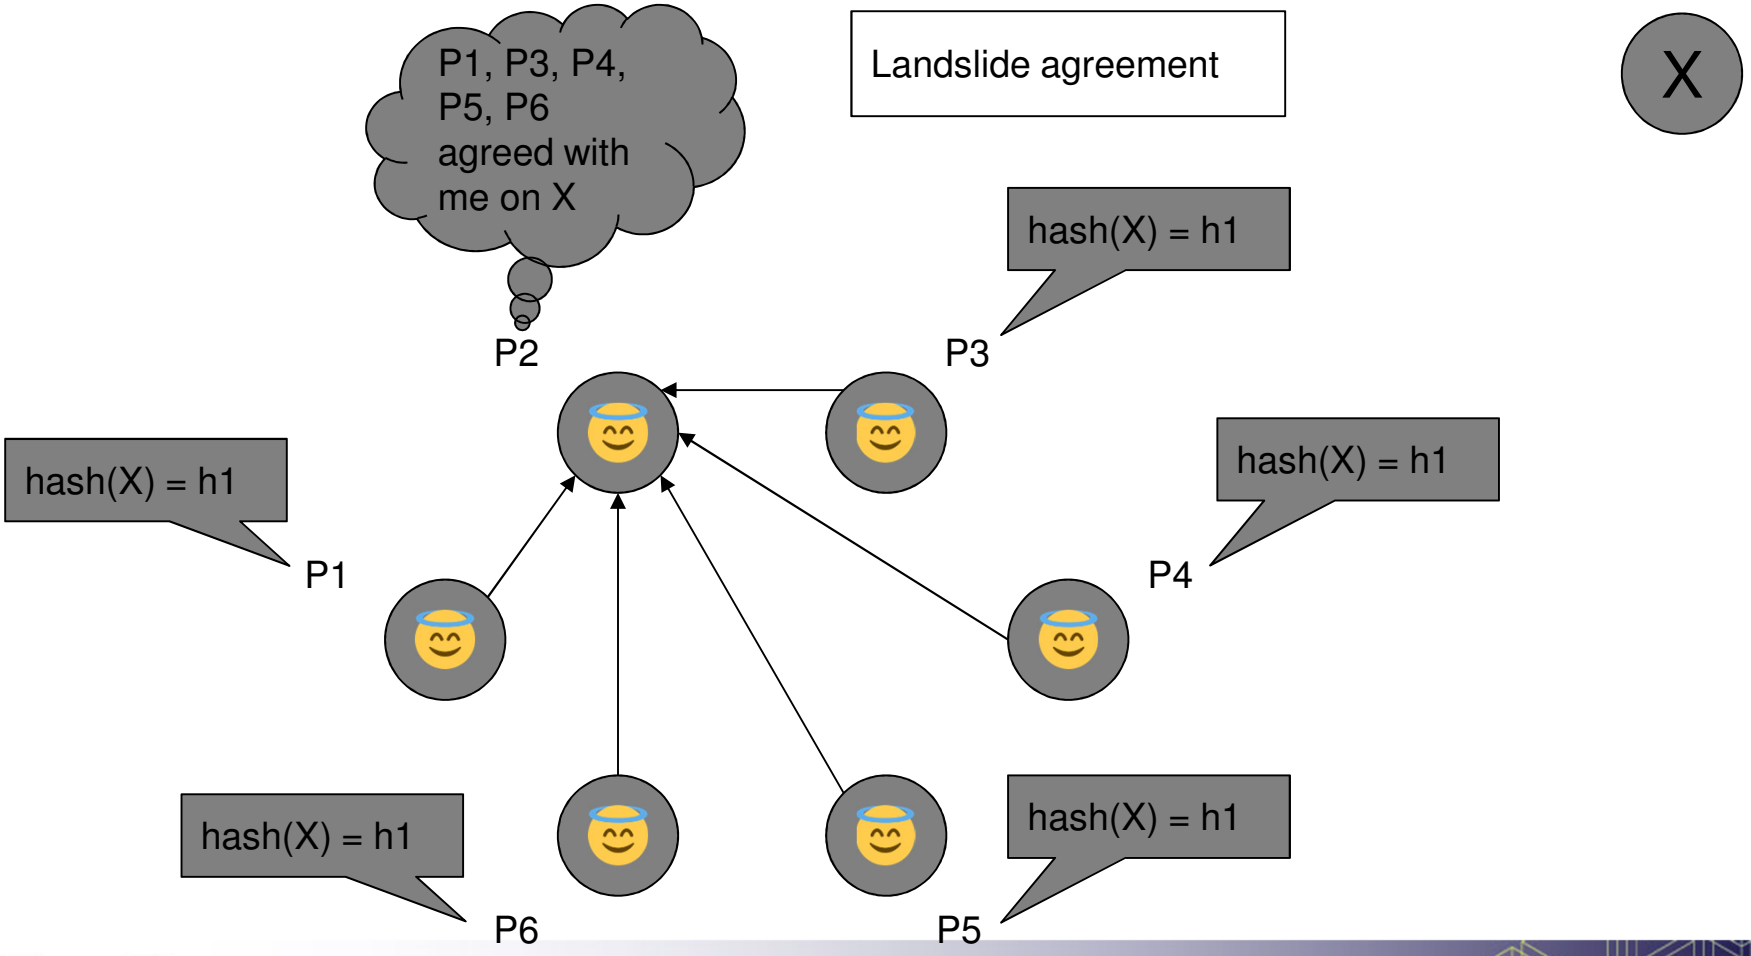

Peer P1 incurs damage on content X
Peer P1 later calls a poll on content X

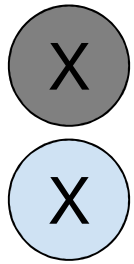

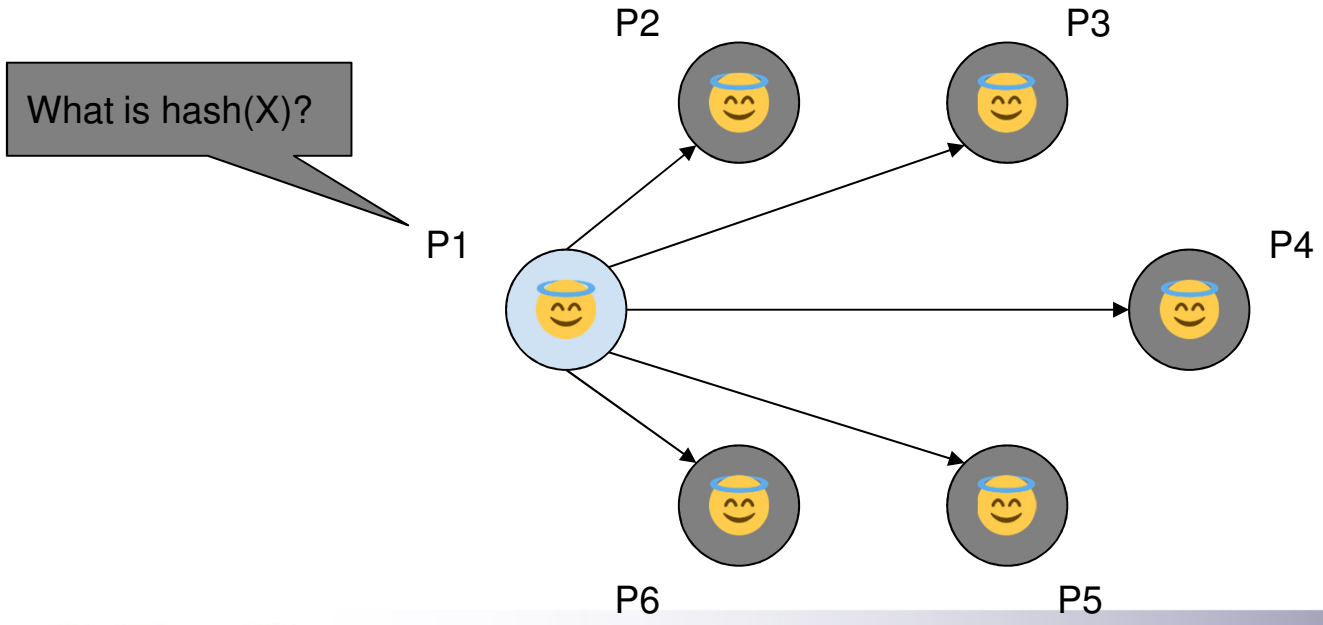

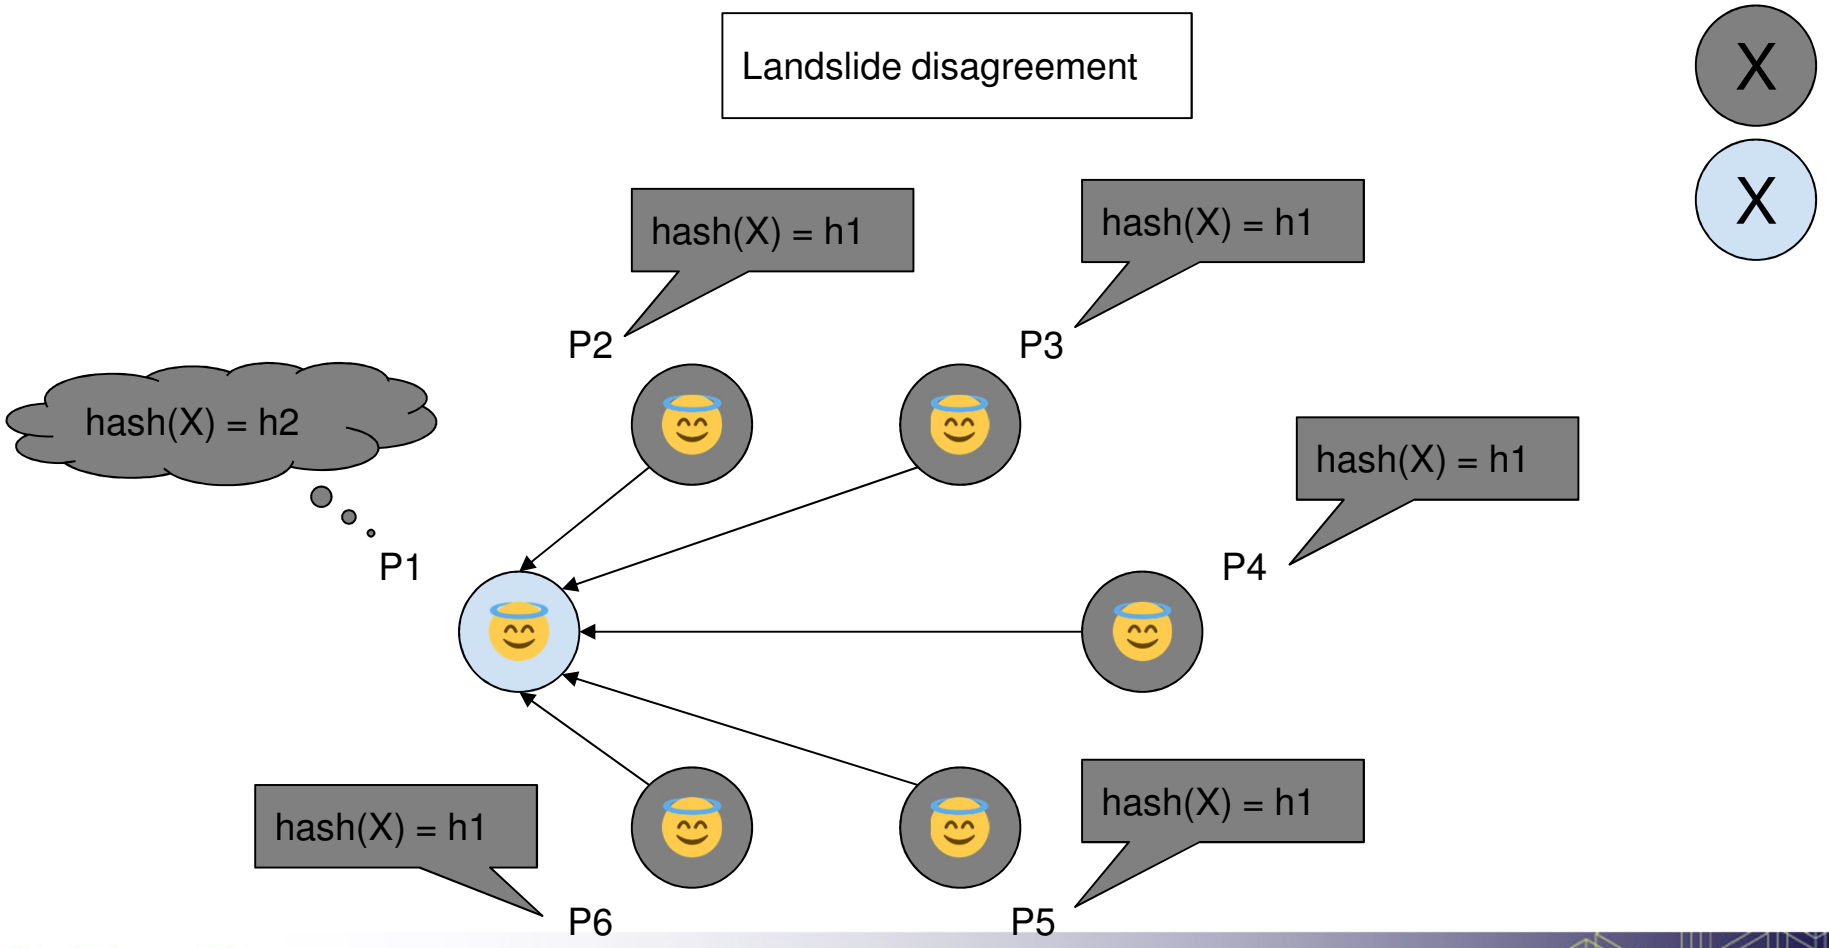

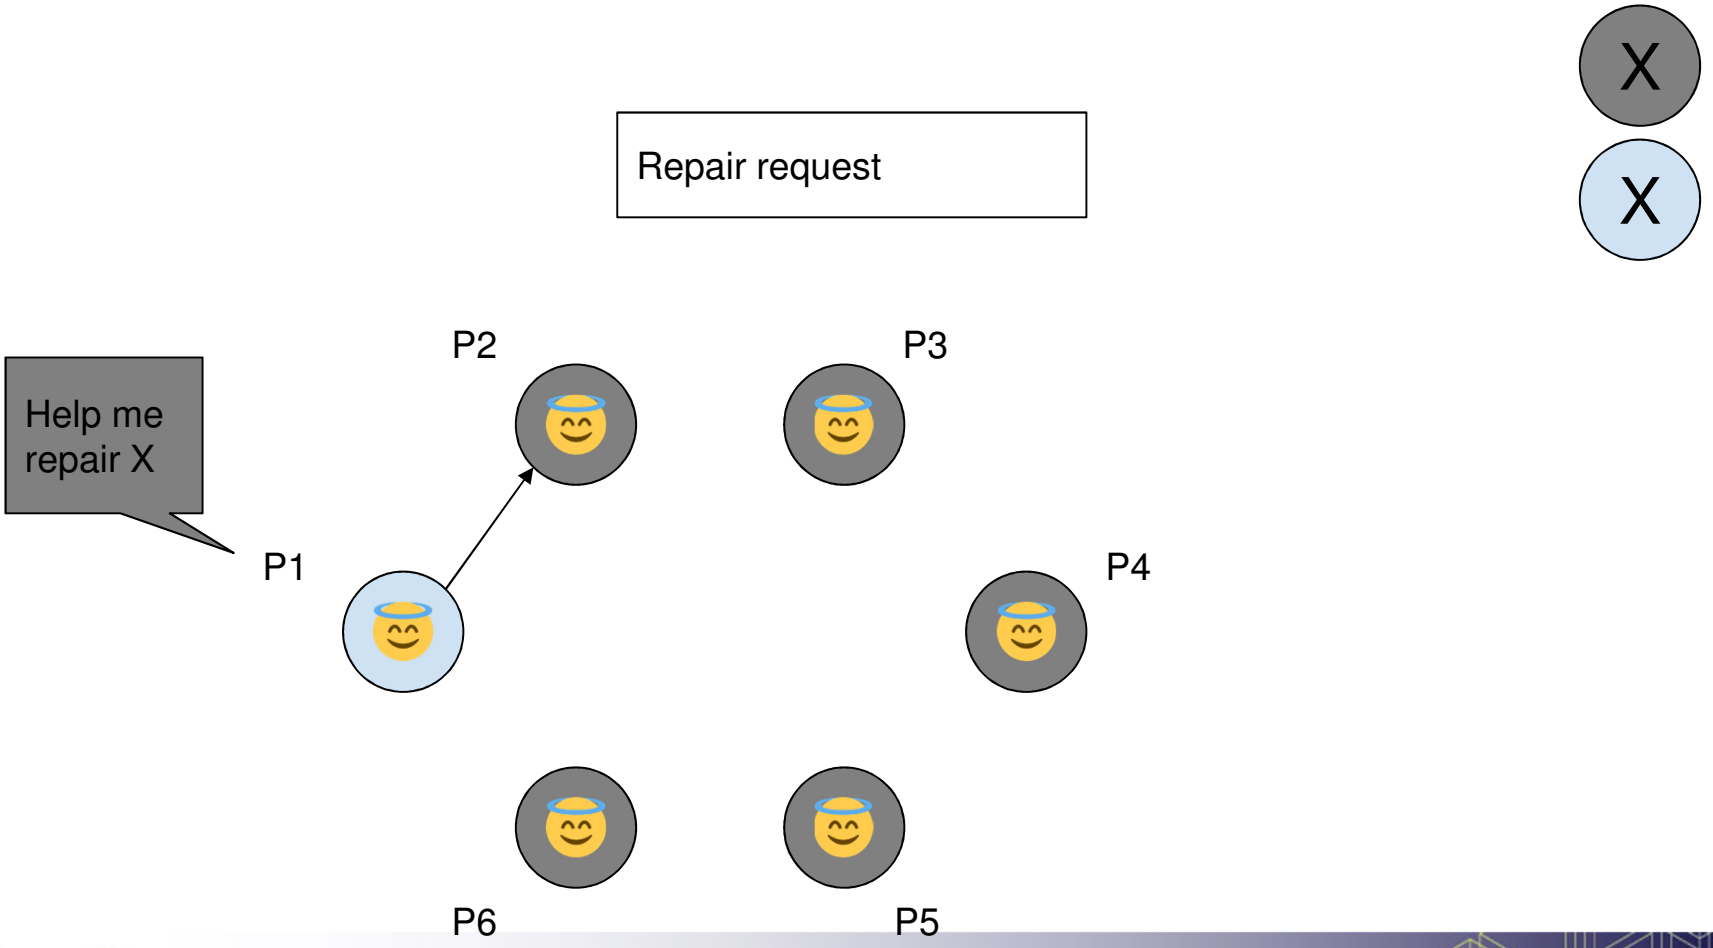

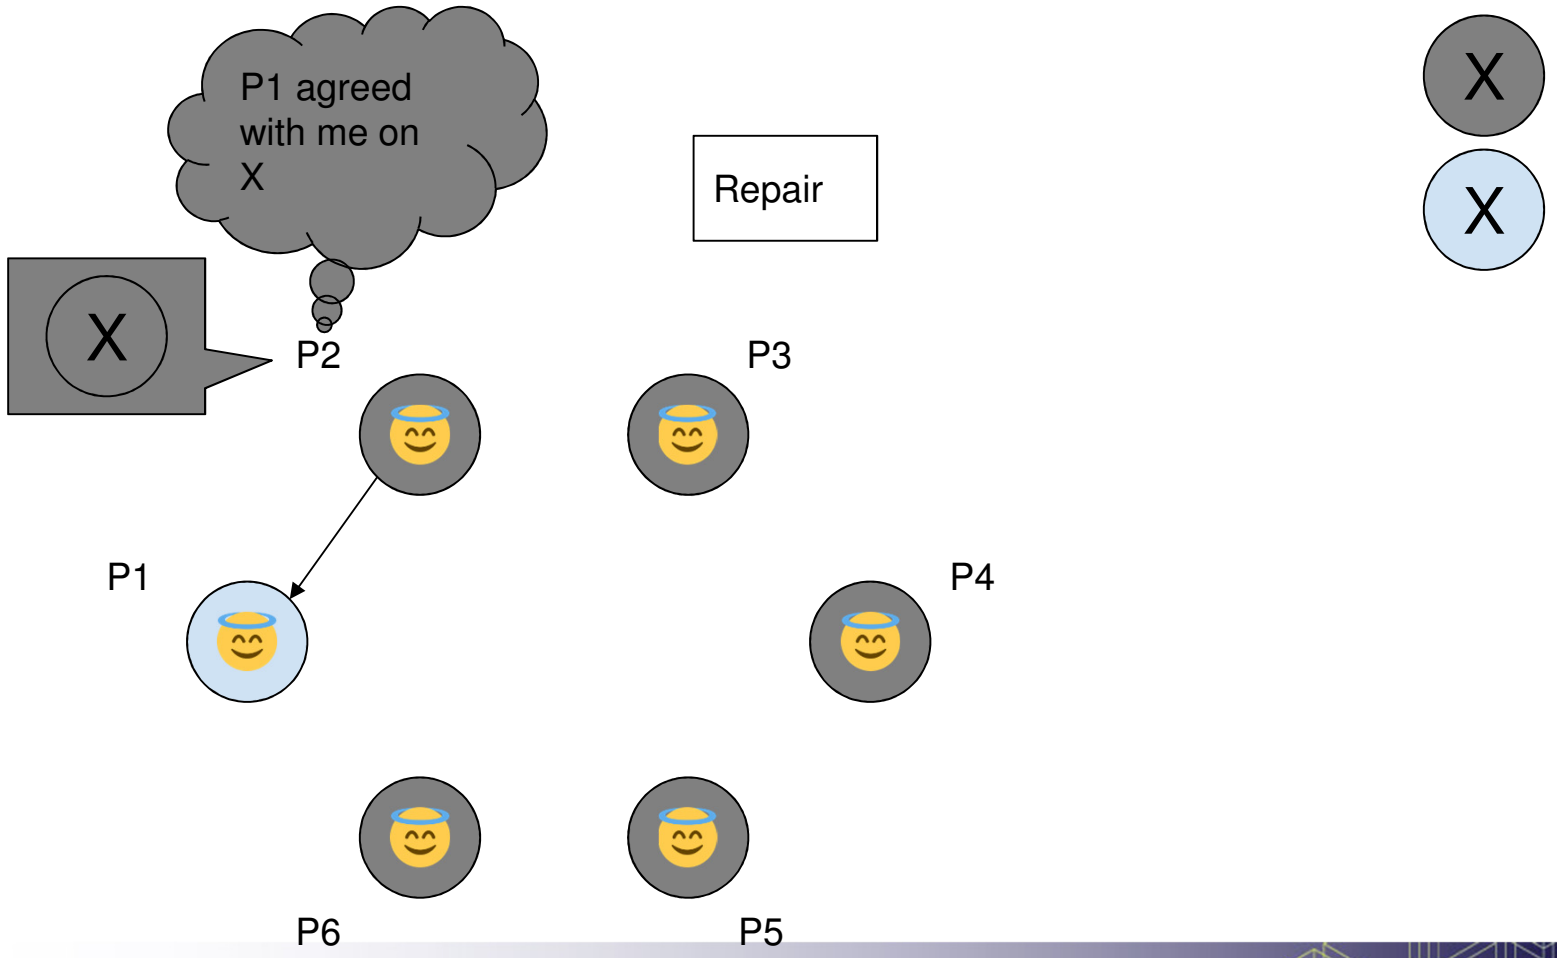

The peers hold identical replicas of X

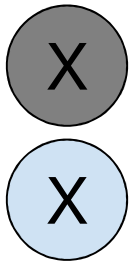

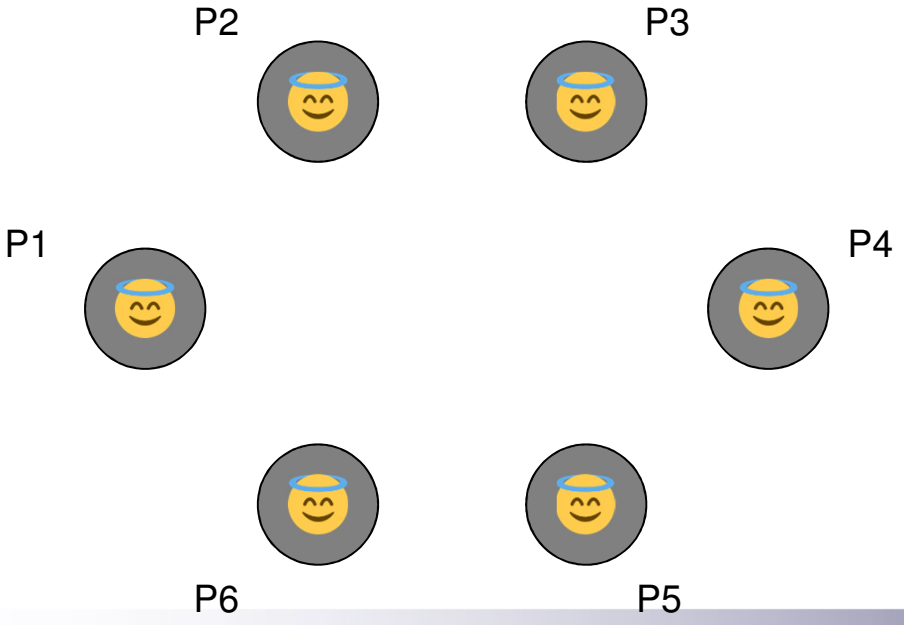

Stealth Modification Gap

- Landslide agreement: take no action (high confidence in outcome)
- Inconclusive agreement: take no action and raise alarm (low confidence in outcome)
- Landslide disagreement: seek repair and notify (high confidence in outcome)

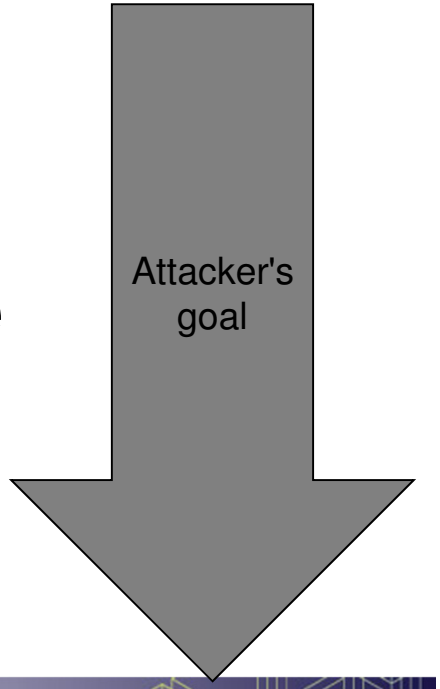

Nonces

- For each voter in a poll over X, the poller supplies a poller nonce P and the voter a voter nonce V
- Rather than hash(X), it is the value of hash(PVX) that is computed
- Nonces must be fresh

Repair Verification

- Byzantine fault
- Bait and switch

Physical Fixity vs. Logical Fixity

- What if the peers hold the same content even though not all of, or even none of, the replicas are byte-identical?

LCAP vs. Threats

- Media failure
- Hardware failure
- Software failure
- Communication errors
- Failure of network services
- Natural disaster
- Media and hardware obsolescence
- Software obsolescence
- Operator error
- Economic failure
- Organizational failure
- External attack
- Internal attack

Outline

1. Context and Use Cases
2. Threat Models
3. LCAP in Action
4. Unlocking LOCKSS for Developers
5. Q&A

Narrow Origins

- Audience: research libraries
- Target: Web content
- Context: appliance model
- "Monolithic stack"

LAAWS Initiative

- "LOCKSS Architected As Web Services"
- Two year Mellon Foundation grant
- Modernization effort

Re-Architecture

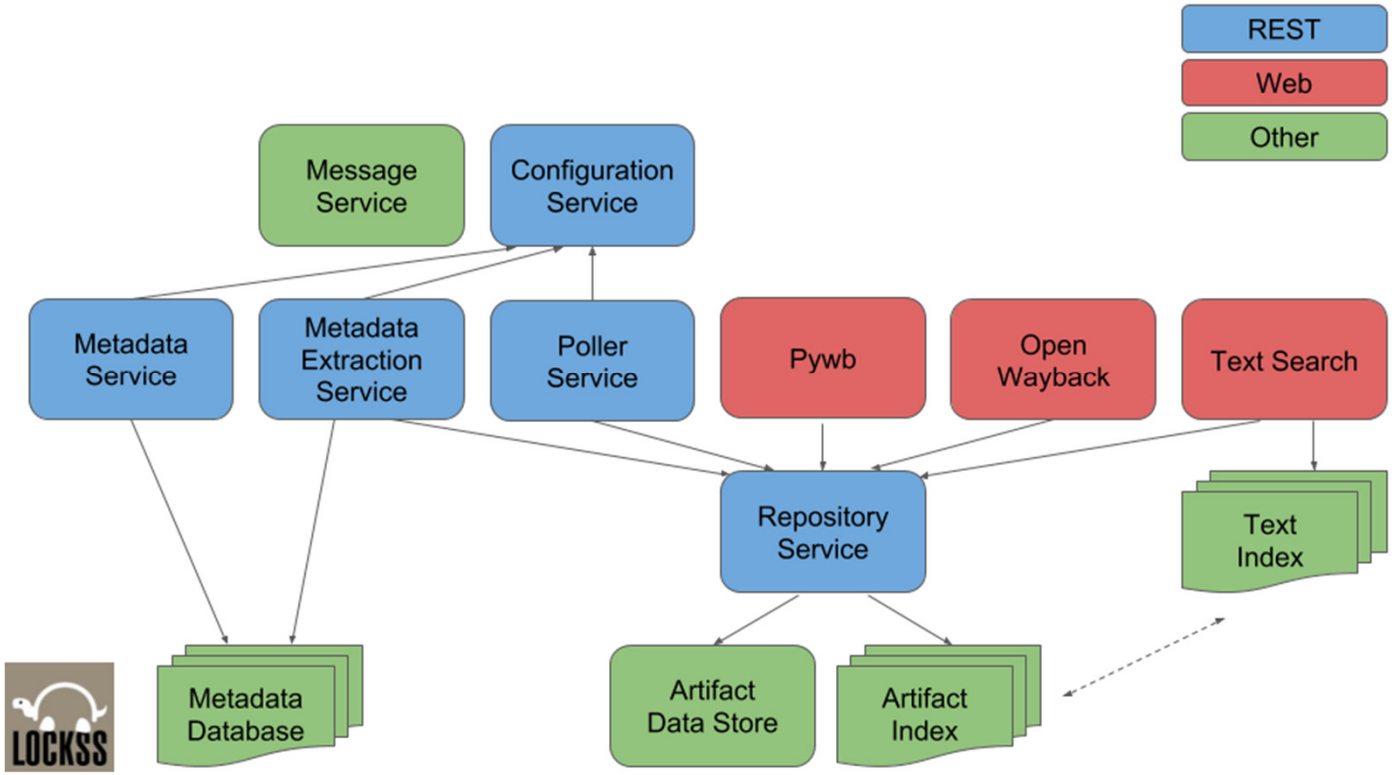

REST APIs and Software Assets

- Repository service
  - https://github.com/lockss/laaws-repository-service
- Configuration service
  - https://github.com/lockss/laaws-configuration-service
- Poller service
  - https://github.com/lockss/laaws-poller
- Metadata extraction service
  - https://github.com/lockss/laaws-metadataextractor
- Metadata service
  - https://github.com/lockss/laaws-metadataservice
- `develop` branch → `docs/swagger.yaml`

Dev/Demo Environment

- `https://github.com/lockss/laaws-demo`
- `feature-mgdemo` branch
- Contains:
  - Docker support infrastructure
  - Docker flavor
  - JAR flavor

Outline

1. Context and Use Cases
2. Threat Models
3. LCAP in Action
4. Unlocking LOCKSS for Developers
5. Q&A


SDC 18
September 24-27, 2018
Santa Clara, CA
www.storagedeveloper.org

Thank you