eXpressive Internet Architecture
Fault Management

Dave Andersen, Adrian Perrig, Peter Steenkiste
David Eckhardt, Sara Kiesler, Jon Peha, Srini Seshan,
Marvin Sirbu, Hui Zhang
Carnegie Mellon University
Aditya Akella, University of Wisconsin
John Byers, Boston University

Carnegie Mellon

BOSTON UNIVERSITY

THE UNIVERSITY of WISCONSIN MADISON

XIA 101:

- Client-server access using simple DAG
- Naming, routing, and forwarding interact
  - All impact execution of communication operation
- This AM: fault mgmt
  Impact on routing, performance, … later

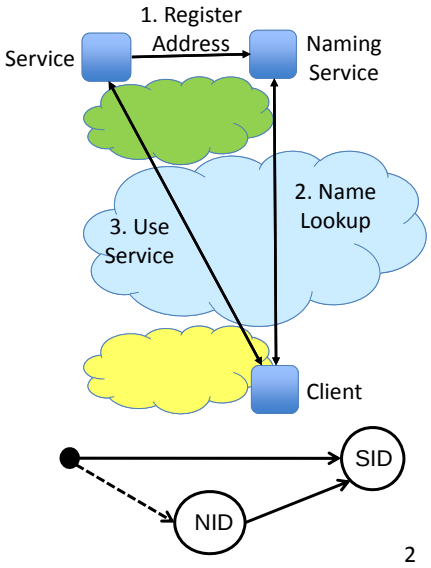

2

Fault Management Question

- "3) Management of faults, which includes detection, localization,

Prevent? Human error, malice

Can't prevent: Backhoes, power failures, s/w bugs

Manage: Contain, detect, localize, mitigate, repair

selection (like NDN) must be able to detect and deal with problems internally. (In ICNs like NDN, flawed data is presumably a relevant class of fault.)"

3

Source of Failures

- Hardware and software faults
  - Router / link failure; incorrect operation or crash
- Operator error
- Adversarial actions
  - Malicious ISP, compromised router or link;
  - Malicious end hosts that perform DDoS attack

4

Availability very generally

- Availability: MTTF / (MTTF + MTTR)
- High availability:
  - Boost MTTF: avoid/prevent failures;
  - Shrink MTTR: enable rapid repair *or masking*
- Masking failures: Redundancy!
  Ensure that higher layers don't see failure

5

XIA: Multi-faceted Approach for Availability

- **Prevent**:
  - Intrinsic security
  - SCION routing
  - STRIDE DDoS defense
- **Enable use of redundancy**:
  - Principal types
  - DAGs

6

Intermezzo: "Levels" of Architecture

- **Principles and Invariants** ("IRTF level"):
  - e.g., the *idea* of using multiple principal types; using intrinsic security;
  - the *requirement* that XIA addresses support fallback for unsupported principal types.
- **Concrete spec**: ("IETF level")
  - What is a host ID principal type? What crypto is used?
- **Implementation** ("Cisco level")

7

Intrinsic Security & Trust Mgmt

- When possible

Trustworthiness should be intrinsic
i.e., not depend on
external databases or info

- Familiar examples: Content hash CIDs and public key hash host/domain IDs (HIDs and ADs)

8

Whither intrinsic security?

- **Prevention:** Ease security-related configuration:
  - e.g., BGP peering: If know peer AD, then crypto follows.
  - Must know peer AD to configure session anyway. *S-BGP can be automatic.*
- *Not a panacea:* Must get intrinsically secure identifier somehow.
- Principle: *Be explicit about the boundary between "human" and self-certifying IDs; cross only once if at all possible!*

9

Principal Types

- **Enable Redundancy:** Different principal types provide independent routing / forwarding planes
  - Failure in one principal type may not affect other types
  - We're still thinking about this one. Most of our focus has been on…

10

Better redundancy for fault masking

- **Example: CIDs vs HIDs:**
  - Content can be served from alternate location (inherent advantage of CCN)
  - Observe: more robust masking than multipath alone; can also mask source failures.

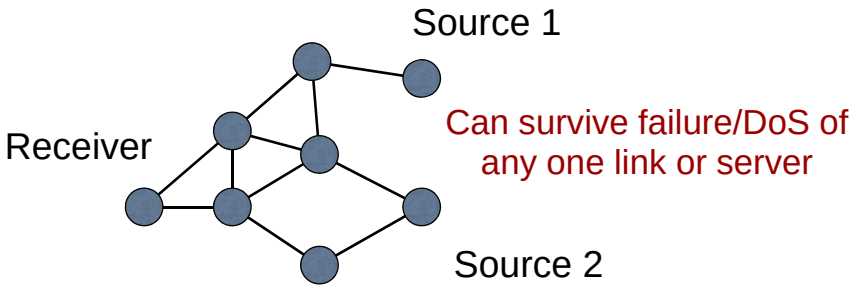

Can survive failure/DoS of any one link or server

Service redundancy, too:

- Facilitates service replication: hash(service pub key)
  - Today's DNS-based service replication limited: binds to specific IP address, no way to support mid-stream failover
  - Expressing service allows any service replica to handle request (like anycast)
- Future: Combine with trusted computing, integrity to create a TrustedAkamai for wide-area replication?
  - Akamai hosts server-side scripts for large providers, runs Java/etc. High degree of implicit trust.
  - XIA's service IDs could enable, e.g., verifiable delegation to a trusted computing-based VM running on TrustedAkamai (tm)
- Example: Replicated DNS servers without the trust/complexity issues of BGP anycast and DNSSEC config!

DAGs

- DAG in packet header offers opportunity for fallback to alternate principal type(s)
- "IRTF"-level observation: DAGs could be used to enable multipath
  - "IETF"-level question: Is this the *right* place for multipath? Or should it be at a lower level?
  - If not, use SCION-level multipath…

13

SCION (1)

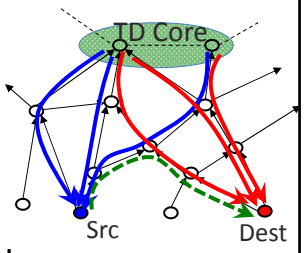

- Host-to-host communication
- **Prevention:** Signed routing beacons prevent announcement of invalid paths
- Cryptographically protected forwarding information prevents alteration of information
  - Alas, information can be replaced by malicious ISP, hence, no forwarding path validation as of now
- **Masking/Redundancy:** Inherent multi-path operation
  - End hosts can simultaneously use multiple paths

14

SCION (2)

- **Prevention: Misconfigurations.** Operator error would likely result in invalid path which will be ignored
  - But intentful "end host" can route around faulty ISP
- Fault localization protocol (currently under development)
  - Enables end host to detect location of faulty link / ISP
- "IETF/cisco" question: Is detection really end-host, or do we mean border router/etc.? TBD!

15

SCION (3)

- **Reduce MTTR:** Active link failure recovery
  - ICMP-like message to inform of link failure: host can immediately pick a new path
  - Sending of new routing beacon
- Passive link failure recovery
  - Periodic beaconing (every 15 seconds) of new working paths

16

STRIDE

- **Prevention:** DDoS defense protocol for SCION
- Resource allocation based on propagation of routing beacons
  - Enabled by tree-based topology of routing beacon propagation emanating from a core
- Per-flow stateless enforcement

17

XIA Fault Mgmt Wrapup

- Prevent faults:
  - Reduce chance of errors: Intrinsic security
  - Prevent some malice: Intrinsic security, SCION
  - Prevent resource attacks: STRIDE, ongoing
- Reduce MTTR: Faster routing and beaconing in SCION
- Enable redundancy for masking: DAGs, expressive principal types to satisfy intent, SCION multipath

18