The eXpressive Internet Architecture: From Architecture to Network

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

Winlab FIA, May 14, 2012

Carnegie Mellon BOSTON UNIVERSITY THE UNIVERSITY of WISCONSIN MADISON

How do you Improve on the Internet?

- The Internet has been tremendously successful
  - Has sustained tremendous growth
  - Supports very diverse set of applications and services
  - Integral part of our society and economy
- Lots of exciting research on how to improve Internet
  - Security, routing, wireless/mobile, management, …
  - But Internet architecture constrains what can be modified
- Future Internet Architecture frees researchers to go beyond today's IP architecture and infrastructure
  - Multi-phase, NSF-funded research program
  - Five teams building full scale networks

Predicting the Future is Hard!

- A lot of really smart people don't agree:
  - Named Data Networking: content centric networking - data is a first class entity
  - Mobility First: mobility as the norm rather than the exception – generalizes delay tolerant networking
  - Nebula: Internet centered around cloud computing, data centers that are well connected

We love all of them!

Outline

- Background
- XIA principles
- XIA architecture
- Building XIA
- Conclusion

XIA Vision

We envision a future Internet that:

- Is trustworthy
  - Security broadly defined is the biggest challenge
- Supports long-term evolution of usage models
  - Including host-host, content retrieval, services, …
- Supports long term technology evolution
  - Not just for link technologies, but also for storage and computing capabilities in the network and end-points
- Allows all actors to operate effectively
  - Despite differences in roles, goals and incentives

5

Today's Internet

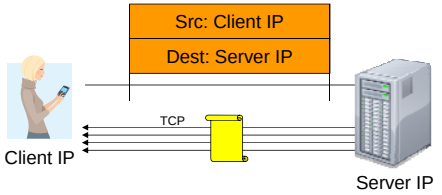

- Client retrieves document from a specific web server
  - But client mostly cares about correctness of content, timeliness
  - Specific server, file name, etc. are not of interest
- Transfer is between wrong principals
  - What if the server fails?
  - Optimizing transfer using local caches is hard
    - Need to use application-specific overlay or transparent proxy – bad!

6

eXpressive Internet Architecture

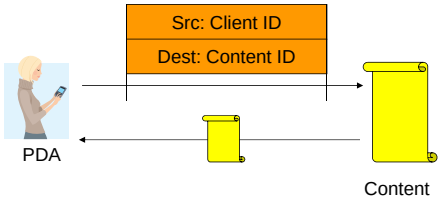

- Client expresses communication intent for content explicitly
  - Network uses content identifier to retrieve content from appropriate location
- How does client know the content is correct?
  - Intrinsic security! Verify content using self-certifying id:
    hash(content) = content id
- How does source know it is talking to the right client?
  - Intrinsic security! Self-certifying host identifiers

7

A Bit More Detail …

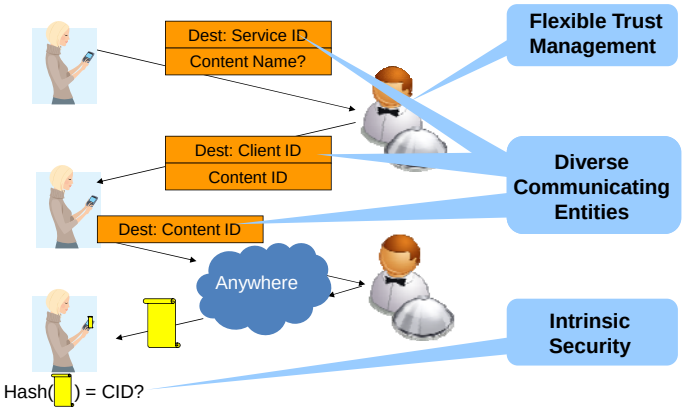

8

Evolvable Set of Principals

- Identifying the intended communicating entities reduces complexity and overhead
  - No need to force all communication at a lower level (hosts), as in today's Internet
- Allows the network to *evolve*

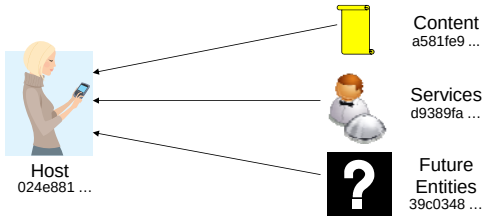

9

Security as Intrinsic as Possible

- Security properties are a direct result of the design of the system
  - Do not rely on correctness of external configurations, actions, data bases
  - Malicious actions can be easily identified

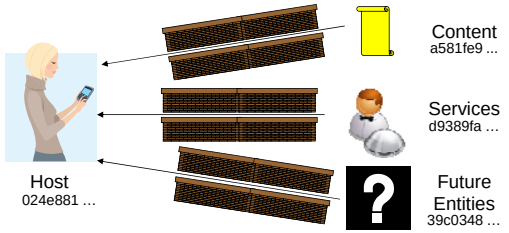

10

Other XIA Principles

- Narrow waist for all principals
  - Defines the API between the principals and the network protocol mechanisms
- Narrow waist for trust management
  - Ensure that the inputs to the intrinsically secure system match the trust assumptions and intensions of the user
  - Narrow waist allows leveraging diverse mechanisms for trust management: CAs, reputation, personal, …
- All other network functions are explicit services
  - Keeps the architecture simple and easy to reason about
  - XIA provides a principal type for services (visible)

Look familiar?

11

XIA: eXpressive Internet Architecture

- Each communication operation expresses the intent of the operation
  - Also: explicit trust management, APIs among actors
- XIA is a single inter-network in which all principals are connected
  - Not a collection of architectures implemented through, e.g., virtualization or overlays
  - Not based on a "preferred" principal (host or content), that has to support all communication

12

What Applications Does XIA Support?

- Since XIA supports host-based communication, today's applications continue to work
  - Will benefit from the intrinsic security properties
- New applications can express the right principal
  - Can also specify other principals (host based) as fallbacks
  - Content-centric applications
  - Explicit reliance on network services
  - Mobile users
  - As yet unknown usage models

13

Outline

- Background
- XIA principles
- XIA architecture
  - Multiple principals
  - DAG-based addressing
  - Intrinsic security
- Building XIA
- Conclusion

14

What Do We Mean by Evolvability?

- Narrow waist of the Internet has allowed the network to evolve significantly
- But need to evolve the waist as well!
  - Can make the waist smarter

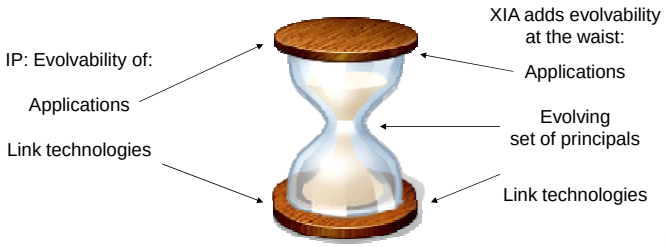

15

Multiple Principal Types

- Hosts XIDs support host-based communication similar to IP – *who?*
- Service XIDs allow the network to route to possibly replicated services – *what does it do?*
  - LAN services access, WAN replication, …
- Content XIDs allow network to retrieve content from "anywhere" – *what is it?*
  - Opportunistic caches, CDNs, …
- Autonomous domains allow scoping, hierarchy
- What are conditions for adding principal types?

16

Multiple Principal Types

Choice involves tradeoffs:
- Control
- Efficiency
- Trust
- Privacy

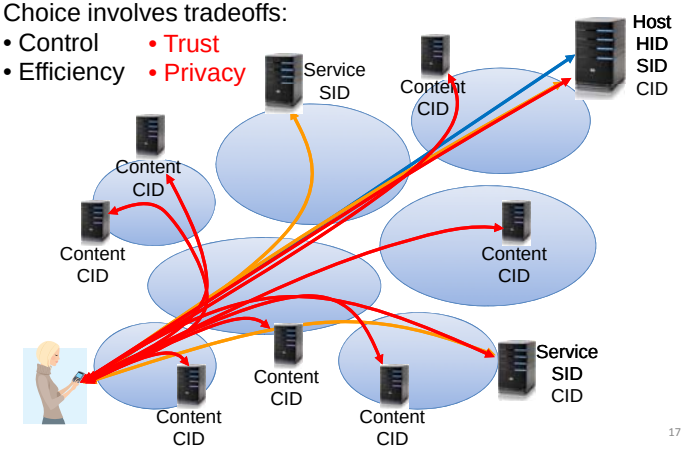

17

Supporting Evolvability

- Introduction of a new principal type will be incremental – no "flag day"!
  - Not all routers and ISPs will provide support from day one
- Creates chicken and egg problem - what comes first: network support or use in applications
- Solution is to provide an *intent* and *fallback* address
  - Intent address allows in-network optimizations based on user intent
  - Fallback address is guaranteed to be reachable

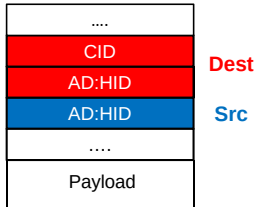

18

Addressing Requirements

- Fallback: intent that may not be globally understood must include a backwards compatible address
  - Incremental introduction of new XID types
- Scoping: support reachability for non-globally routable XID types or XIDs
  - Needed for scalability
  - Generalize scoping based on network identifiers
  - But we do not want to give up leveraging intent
- Iterative refinement: give each XID in the hierarchy option of using intent

19

Our Solution: DAG-Based Addressing

- Uses direct acyclic graph (DAG)
  - Nodes: typed IDs (XID; expressive identifier)
  - Outgoing edges: possible routing choices
- Simple example: Sending a packet to HID_S

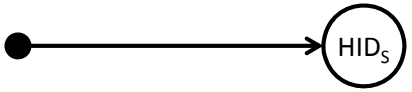

Dummy source: special node indicating packet sender

Intent: final destination of packet with no outgoing edges

20

Support for Fallbacks with DAG

- A node can have multiple outgoing edges

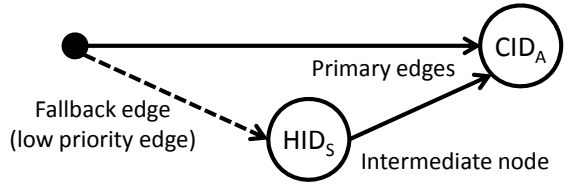

- Outgoing edges have priority among them
  - Forwarding to HID$_S$ is attempted if forwarding to CID$_A$ is not possible – Realization of fallbacks

21

Support for Scoping with DAG

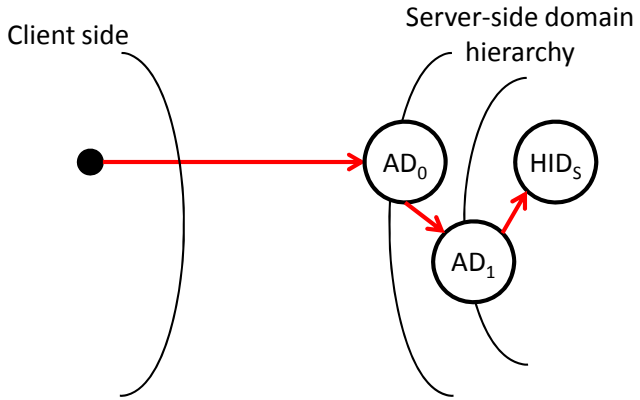

Support scalable routing, binding, migration, mobility, …

22

Iterative Refinement: Scoping while Maintaining Intent

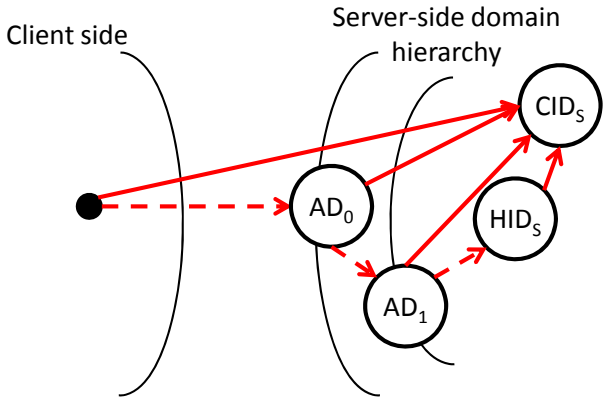

23

DAG Addressing Research Questions

- DAG addressing supports is flexible …
  - Fallback, binding, source routing, mobility, ..
- … but many questions remain:
  - Is it expensive to process?
  - How big will the addresses be?
  - How do ISPs verify policy compliance?
  - Can they be used to attack network?
  - Can it be deployed incrementally?

24

Intrinsic Security in XIA

- XIA uses self-certifying identifiers that guarantee security properties for communication operation
  - Host ID is a hash of its public key – accountability (AIP)
  - Content ID is a hash of the content – correctness
  - Does not rely on external configurations
- Intrinsic security is specific to the principal type
- Example: retrieve content using …
  - Content XID: content is correct
  - Service XID: the right service provided content
  - Host XID: content was delivered from right host

25

Example of Secure Mobile Service Access

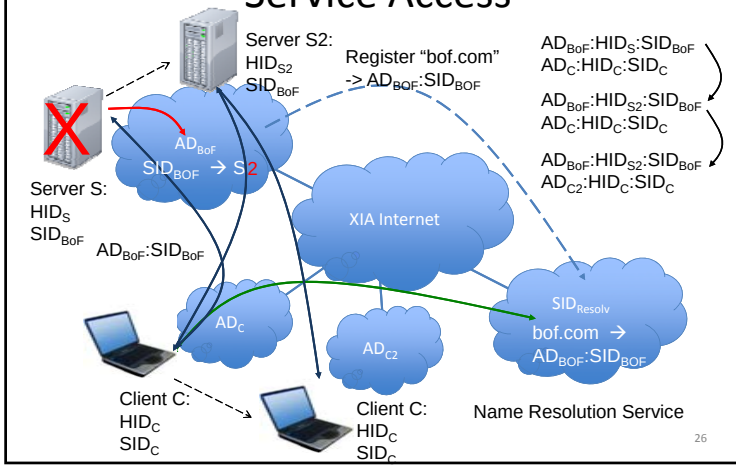

26

Path Selection in SCION Architecture Overview

- Source/destination can choose among up/down hill paths
  - Path control shared between ISPs, receivers, senders
- Desirable security properties:
  - High availability, even in presence of malicious parties
  - Explicit trust for operations
  - Minimal TCB: limit number of entities that must be trusted
  - No single root of trust
  - Simplicity, efficiency, flexibility, and scalability

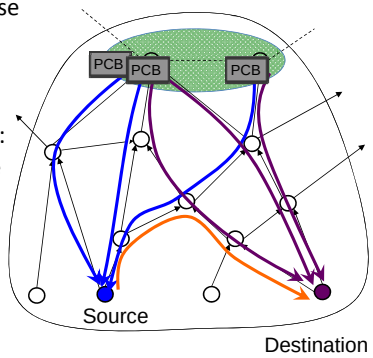

27

Distributed Control in XIA

- Customers have more choices:
  - Choice of XID type, i.e. how is communication operation performed; involves different tradeoffs
  - DAGs add flexibility: fallback, services, …
  - Scion offers some control over path selection
- Service providers have choices as well
  - Use of XID types to optimize new services
  - Scion allows new path optimization options
  - Use DAGs for binding, scoping, mobility, …
- Provides opportunities for customizing interactions to context

28

Outline

- Background
- XIA principles
- XIA architecture
- Building XIA
  - Forwarding packets
  - Building a network
  - Prototype
- Conclusion

29

Putting Address into Packet Headers

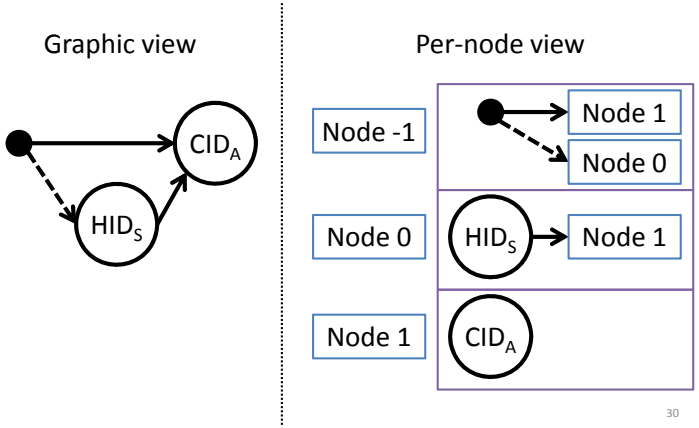

30

XIP Packet Header

- DAGs represent source and destination addresses
- Array of nodes with pointers
- Maintains a *LastNode* field in the header
  - Routers to know where to begin forwarding lookups

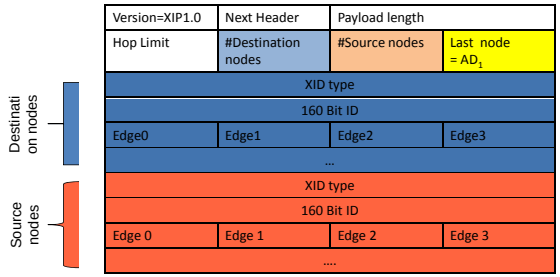

Router's View on Packet Forwarding

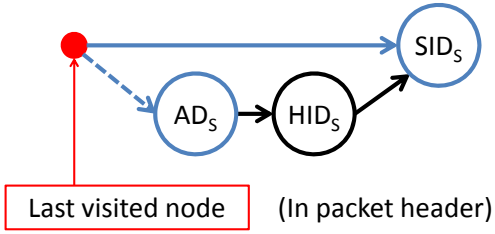

1. Forward to SID_S if possible
2. Otherwise, forward to AD_S
   - If router is AD_S itself, update last visited node to AD_S

32

Packet Processing Pipeline

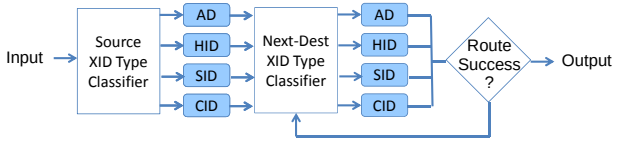

- Principle-independent processing defines how to interpret the DAG
  - The core XIA architecture
- Principle-dependent processing realizes forwarding semantics for each XID type
- Optimizations possible: fast path processing, packet level and intra-packet parallelism

33

Evaluation Setup

- Router
- Packet generator

Software:

PacketShader I/O Engine

Click modular router – multithreaded(12 threads)

Hardware:

10Gbit NIC : 4 ports (multi-queue support)

2x 6 Core Intel Xeon @ 2.26GHz

Forwarding Performance Comparison

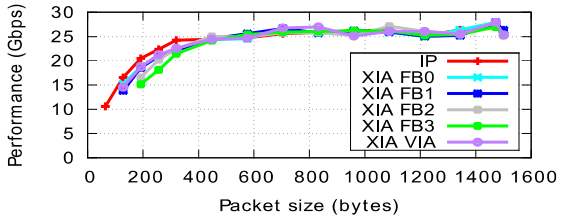

351K FIB entries
Workload: Identifiers generated using Pareto distribution

XIP forwarding is fast!
@128 byte FB0 is 8% slower than IP
@192 byte FB3 is 26% slower than IP

Fast Path Performance

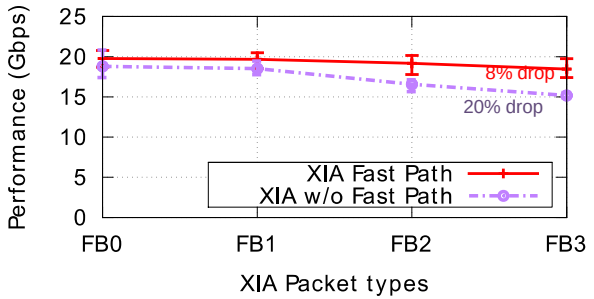

Look-aside cache of 1024 entries

Using fast-path processing, the gap between FB0 and FB3 is reduced significantly !

Path Selection in SCION Architecture Overview

- Source/destination can choose among up/down hill paths
  - Path control shared between ISPs, receivers, senders
- Desirable security properties:
  - High availability, even in presence of malicious parties
  - Explicit trust for operations
  - Minimal TCB: limit number of entities that must be trusted
  - No single root of trust
  - Simplicity, efficiency, flexibility, and scalability

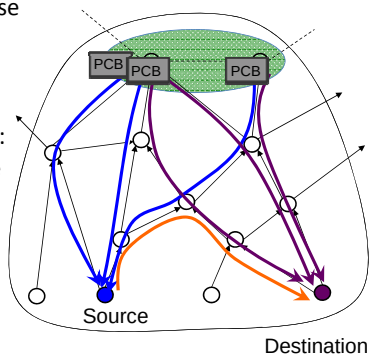

37

XIP Protocol Stack

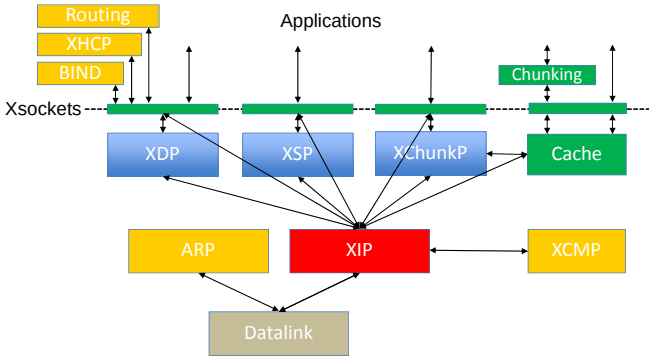

- Open source release of complete prototype this month
- Support for GENI and VM-based experiments

38

XIA Components and Interactions

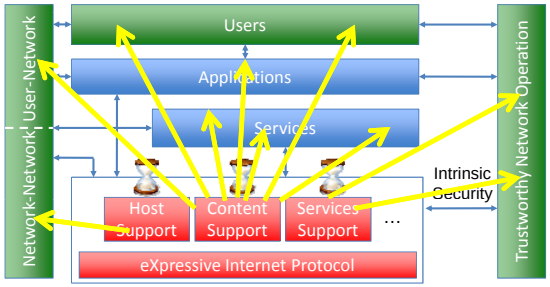

39

Conclusion

- XIA supports evolution, expressiveness, and trustworthy operation.
  - Multiple principal types, flexible addressing, and intrinsic security
- But research has just started!
  - Transport protocols, applications, services, …
  - Trustworthy protocols that fully utilizes intrinsic security of XIA
- More information on http://www.cs.cmu.edu/~xia