

XIA: An Architecture for a Trustworthy and Evolvable Internet

Peter Steenkiste

Dave Andersen, David Eckhardt, Sara Kiesler, Jon Peha,
Adrian Perrig, Srinu Seshan, Marvin Sirbu, Hui Zhang
Carnegie Mellon University

Aditya Akella, University of Wisconsin

John Byers, Boston University

KRNet 2012 Conference
June 2012, Seoul, Korea

Carnegie Mellon

BOSTON
UNIVERSITY

THE UNIVERSITY
of
WISCONSIN
MADISON

How do you Improve on the Internet?

- The Internet has been tremendously successful
 - Supports very diverse set of applications and services
 - Integral part of our society and economy
 - But there are also many challenges ...
- Lots of exciting research on how to improve Internet
 - Security, routing, wireless/mobile, management, ...
 - But Internet architecture limits what can be modified
- Future Internet Architecture frees researchers to go beyond today's IP architecture and infrastructure
 - Long-term multi-phase effort funded by US NSF

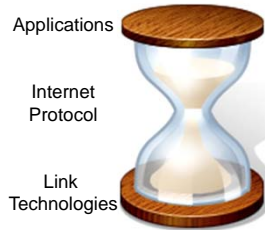
2

NSF Future Internet Architecture

- Can we look beyond the current Internet architecture?
 - Tackle challenges in more fundamental ways?
 - Want to keep the good features of today's network
- Four teams were selected in the second phase:
 - Named Internet Architecture: **content centric** networking - data is a (the) 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
 - eXpressive Internet Architecture: focus on trustworthiness, **evolvability**

3

“Narrow Waist” of the Internet Key to its Success

- Has allowed Internet to evolve dramatically
 - But now an obstacle to addressing challenges:
 - No built-in security
 - Hard to evolve
 - Limited interactions edge-core
- 
- XIA exploring three concepts to address issues:
 - Diverse types of end-points
 - Intrinsic security
 - Flexible addressing

Outline

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

5


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

6


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!

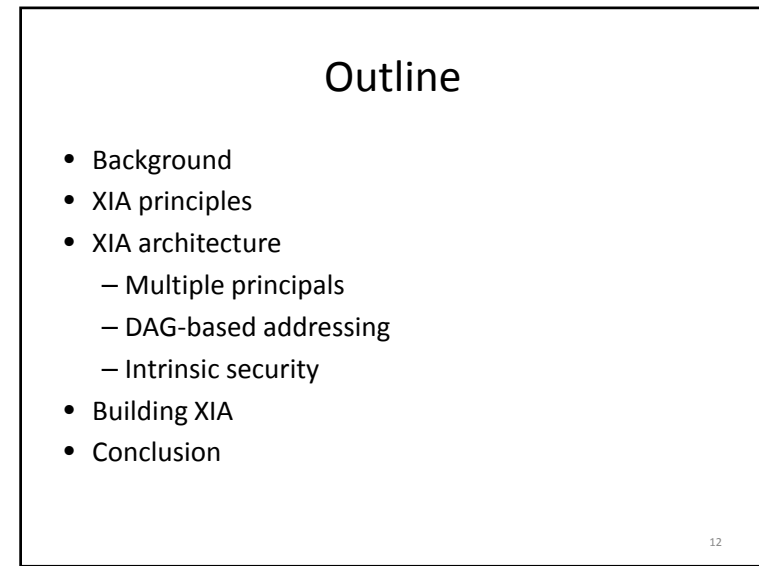
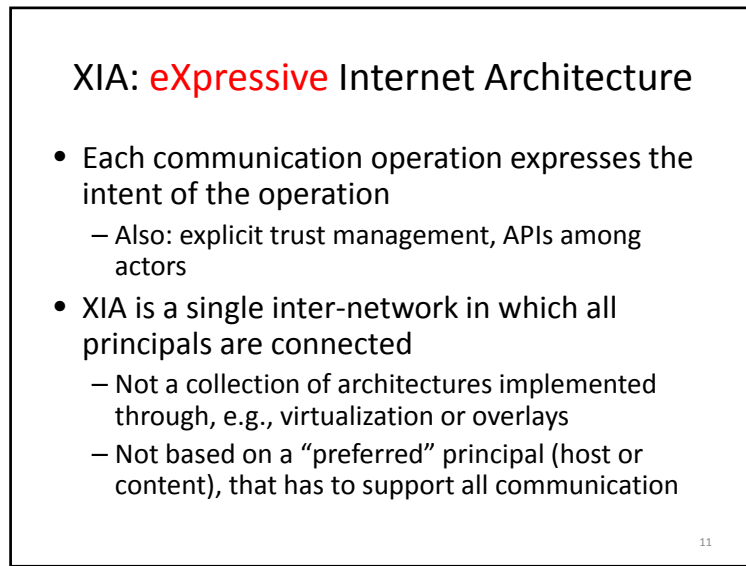
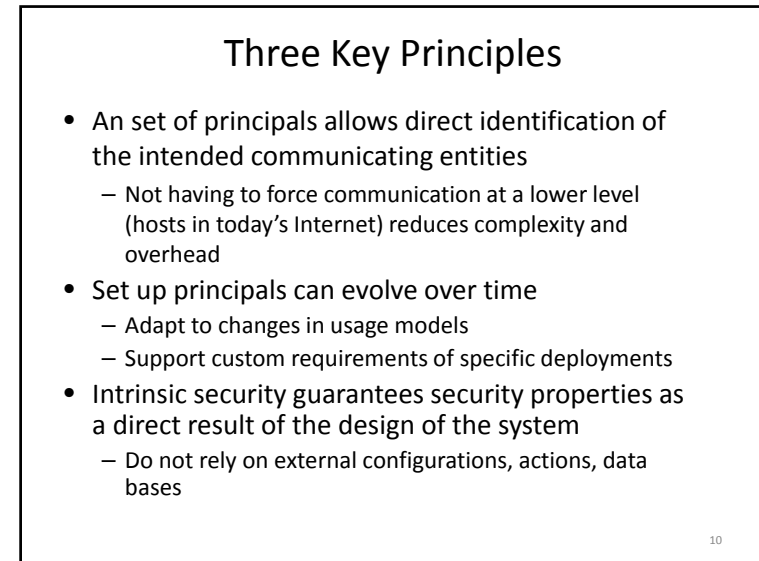
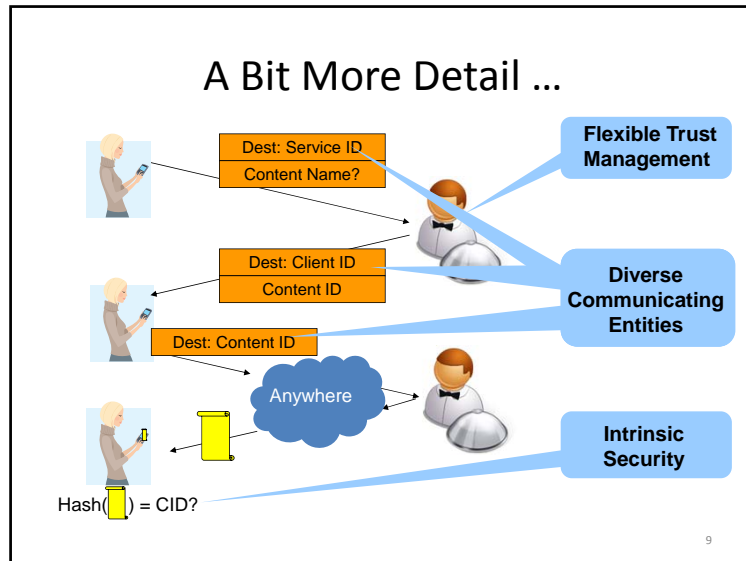
7

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

8



Developing XIA v0.1

- Principles do not make an architecture!
- Meet the core XIA team:



Five happy professors cheering:
John Byers, Aditya Akella, Dave Anderson,
Srinu Seshan, Peter Steenkiste




- Next: quick look at multiple principals, fallbacks and DAGs, intrinsic security

13

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



Multiple Principal Types


- Associated with different forwarding semantics
 - Support heterogeneity in usage and deployment models
- 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

15

Multiple Principal Types

Choice involves tradeoffs:

- Control • Trust
- Efficiency • Privacy



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

| |
|---------|
| |
| CID |
| AD:HID |
| AD:HID |
| |
| Payload |

Dest

Src

17

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

18

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

```

graph LR
    S(( )) --> HID_S((HID_S))
            
```

Dummy source:
special node indicating
packet sender

Intent:
final destination of packet
with no outgoing edges

- DAGs help routers deal with routing problems

19

Support for Fallbacks with DAG

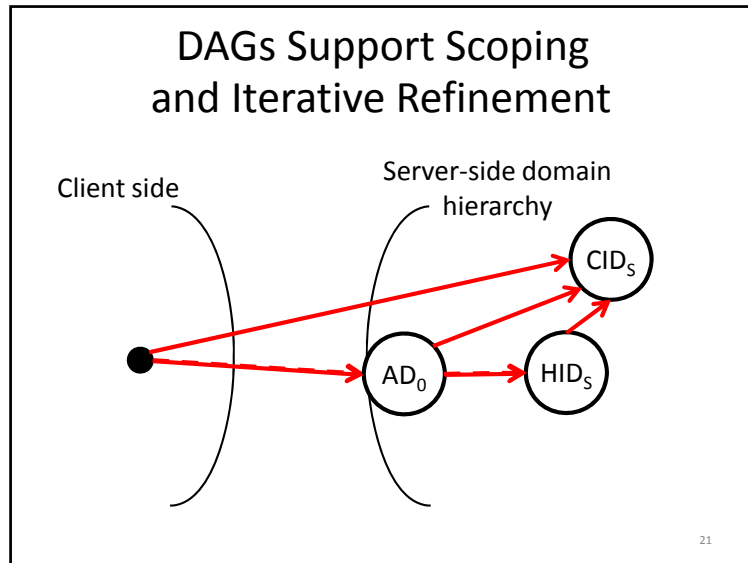
- A node can have **multiple outgoing edges**

```

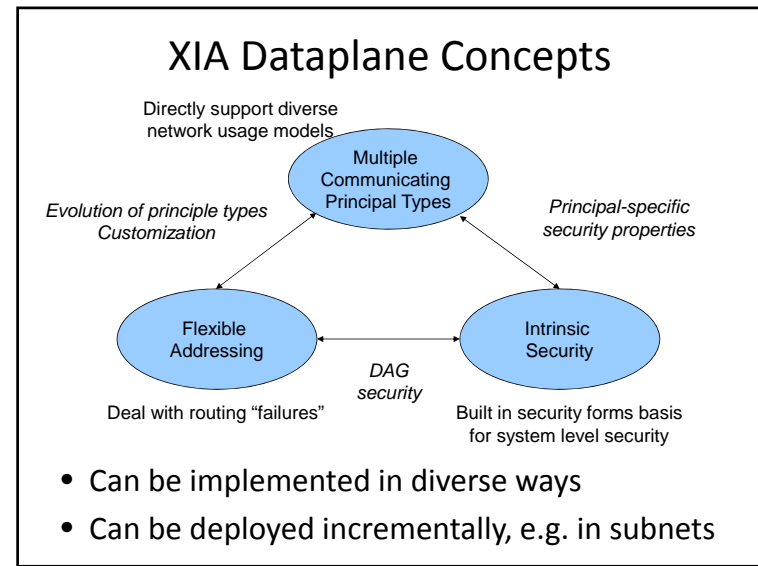
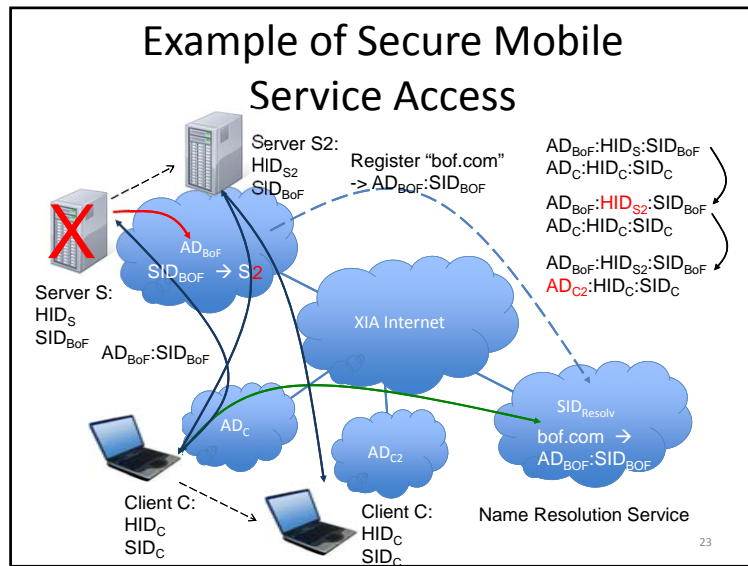
graph LR
    S(( )) -- Primary edges --> CID_A((CID_A))
    S -. Fallback edge (low priority edge) .-> HID_S((HID_S))
    HID_S -- Intermediate node --> CID_A
            
```

- Outgoing edges have **priority** among them
 - Forwarding to HID_S is attempted if forwarding to CID_A is not possible – Realization of fallbacks

20



- ### 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
- 22



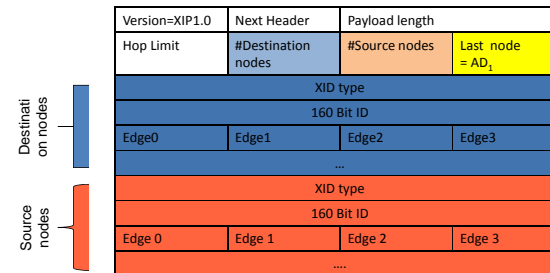
Outline

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


25

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 in AD_s itself, update last visited node to AD_s

27

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

28

Evaluation Setup

- Use packet generator to evaluate throughput
- 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
- Optimizations apply: fast path processing, packet level and intra-packet parallelism


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

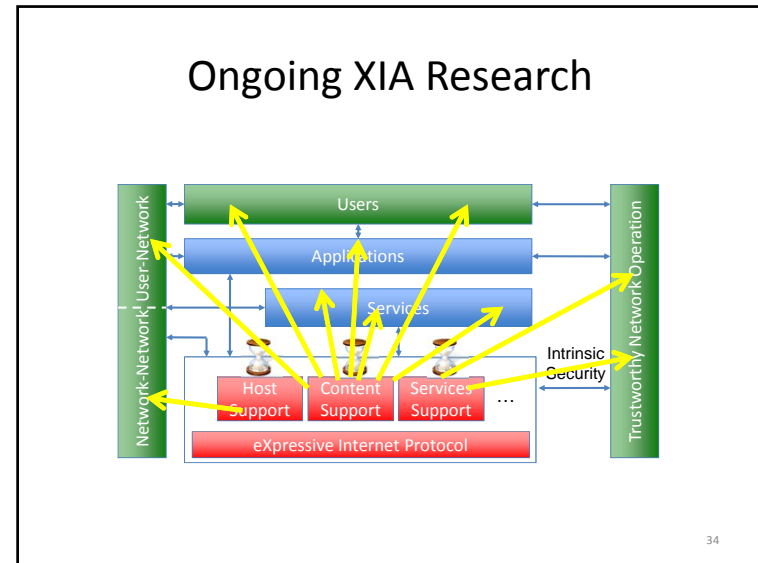
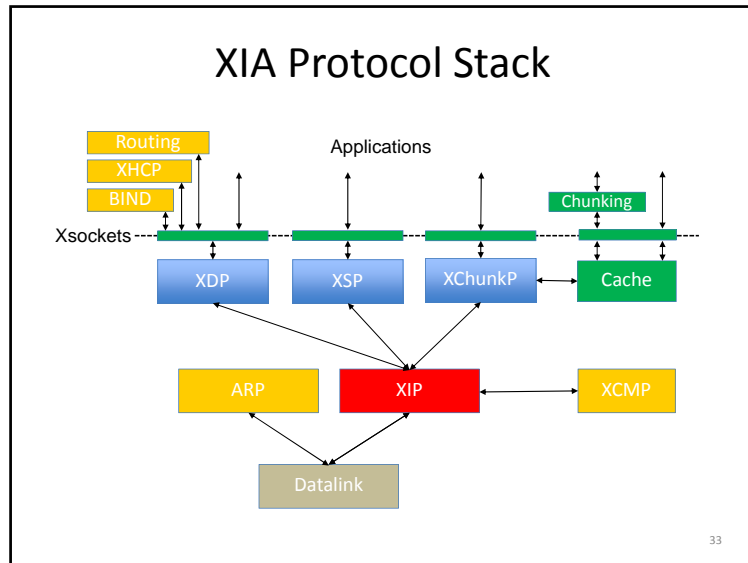


Look-aside cache of 1024 entries

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

XIA Prototype

- Completed Click-based prototype
 - Full stack for routers, caches, and end-points
 - Basic transport, routing, naming, diagnostics, ...
 - User-level/in-kernel, native/overlay
- Open-source release of prototype on GitHub
 - More info: <http://www.cs.cmu.edu/~xia>
 - Support for GENI and VM-based experiments
 - Good starting point for collaboration
- Extend over time by integrating research



- ### A Broad Research Agenda
- Applications and services
 - Web, CDNs, video delivery, teleconferencing, games, mobility services, ...
 - Key to evaluating and improving the architecture
 - Protocols and network infrastructure
 - Security, transport protocols, naming, mobility, routing, service deployment, principal types, ..
 - Targeted deployments
 - Use XIA to optimize unique networks, e.g. wireless access, Scada, sensors, “ad hoc”, data center, ...

- ### Conclusion
- XIA supports evolution, expressiveness, and trustworthy operation.
 - Multiple principal types, intrinsic security, and flexible addressing
 - <http://www.cs.cmu.edu/~xia>
 - Broad research agenda applications, protocols, and deployments
 - Use XIA to fundamentally improve the network: transport protocols, trust management, applications, services, ...
 - Use flexibility to target demanding network deployments
 - Customize without giving up interoperability

Extra slides

37

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!

38