

An Overview of Security & Privacy in Content-Centric Networking (NDN and CCN)

Gene Tsudik

CS Department
University of California, Irvine (UCI)
www.ics.uci.edu/~gts

1

Current Research Topics

- Security of Embedded Devices (IoT?)
- Private Set Operations
 - Cloud/DB apps
 - Genomic S&P
 - Input size-hiding
- Privacy in Social Networks
- Usable Security
- Weird Biometrics
- S&P in CCN/NDN

For more info see: sprout.ics.uci.edu

OUTLINE

- Internet
- NDN/CCN Overview
- NDN Security & Privacy
- Anonymous Retrieval
- Cache Privacy
- Denial of Service
- Trust Management
- Optional Topics, e.g.,
 - Access Control, Accounting, Fragmentation, NACKs

3

NEED TO KNOW

- Basic networking & Internet concepts
- Network security principles
 - Protocols
- Basic knowledge of applied cryptography
 - Basic crypto primitives

4

Today's Internet

- Tremendous, unexpected and long-lasting global success story
- 35-year-old design: architecture defined in RFC 791/793 (1981 and earlier)
- Enables any host to talk to any other host
 - Names boxes and interfaces
 - Supports end-to-end conversations
 - Provides unreliable packet delivery via IP datagrams
 - Compensates for simplicity of IP via complexity of TCP

5

IP-Based Internet

- Helped facilitate today's rich global-scale communication
- But, was not designed for it
- Fundamental communication model: point-to-point conversation between two hosts (IP interfaces)
- The central abstraction is a host identifier corresponding to an IP address

6

Recent Decades

- Last 20 years – profound change in nature of Internet communication
 - From email/ftp/telnet to what?
 - From a few thousands of users to that?
 - From static wired nodes (computers, terminals) to what?
 - From friendly, clubby, trusting ambience, to what?
- Massive amounts of data constantly produced and consumed
 - Web (esp. media sharing and social networking),
 - Audio-/video-conferencing
 - Email, etc.

7

Key Aspects of Internet Change

- Multimedia
- Mobility / Wireless-ness
 - Delays and Disruptions
- Distribution Scale
- Cloud

8

Internet Security & Privacy

- S&P in the current Internet are certainly **NOT** a success story
- Retrofitted, incremental, band-aid-style solutions, e.g.:
 - SSH,
 - SSL/TLS,
 - IPSec + IKE,
 - DNSSec,
 - sBGP,
 - AAA, etc.

9

NSF Future Internet Architectures (FIA) Program

- Targeted NSF-funded program, 2-tiered competition
- Major goals:
 - Design comprehensive next-generation Internet architectures
 - Accommodate current and emerging comm-n paradigms
 - Security and privacy from the outset (by design)
- Started in 2010
 - Phase I: 2010-2014
 - Phase II: 2014-2018
- Projects:
 - Nebula (Phase I)
 - MobilityFirst (Phases I and II)
 - XIA: eXpressive Internet Architecture (Phases I and II)
 - NDN: Named-Data Networking (Phases I and II)
 - ChoiceNet (started in 2012, not strictly speaking FIA)

10

Caveat Emptor

- I was part of the NDN FIA project 2010-2014
- Worked on S&P in NDN (and CCN)
- Was funded by the NSF ('till 09/15)
- Thus... take everything with a grain of salt, draw your own conclusions, and explore further

Also:

- I focus on NDN and CCN
- There are other ICN efforts

11

The image shows a height chart with markings from 4'4" to 7'0" in increments of 2 inches. A hand is holding a dark green sign with the word 'DATA' written in white capital letters. The sign is positioned at the 4'8" mark. The text 'NDN & CCNx' is overlaid in large, bold, black letters across the top of the chart.

- "Named data networking project (NDN)", <http://named-data.org>
- "Content centric networking (CCNx) project", <http://www.ccnx.org>
- "Networking named content", ACM CoNEXT, 2009.

12

Communication

- For almost 150 years, communication meant:
A wire connecting two devices



- The Web forever changed that:
What matters is content, not the host it came from



13

DN vs. CN

	Communication	Distribution
Naming	Endpoints	Content
Memory	Invisible, Limited	Explicit; Storage = Wires
Security	Communication process	Content

Today's Internet: a communication network, used as a distribution network

14

Named-Data Networking (NDN)

✧ An instance of ICN and CCN

NDN focuses on:
Scalable Content
Distribution which is poorly
served by today's Internet


15

NDN: key participants




16

Content Distribution over IP



17

Content Distribution over NDN



18

NDN Basic Concepts

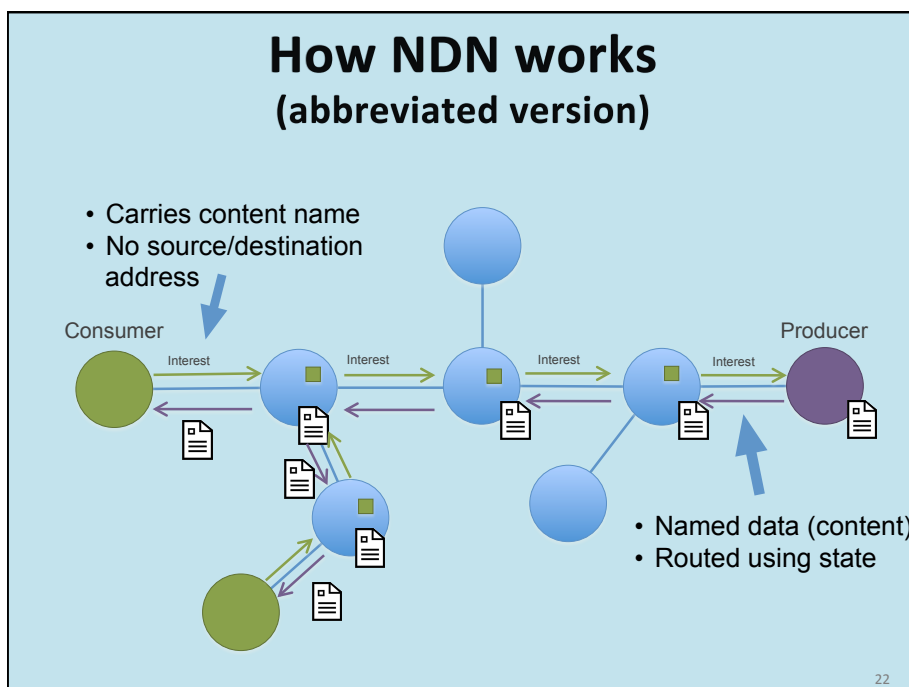
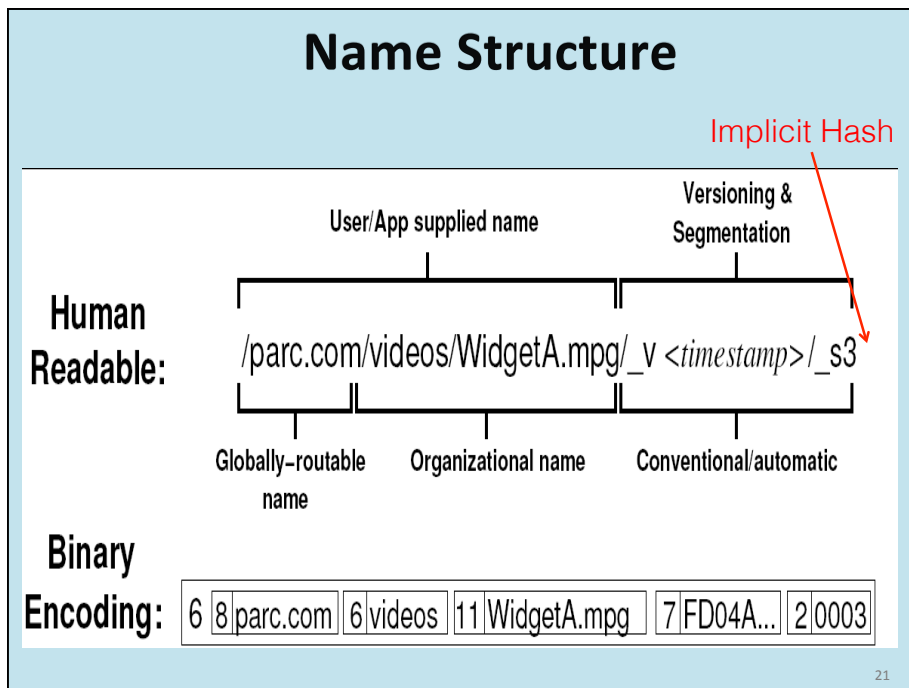
- Name
 - Human-readable, path/url - like
- Roles:
 - Consumer
 - Producer
 - Router
- Objects:
 - Content
 - Interest

19

As opposed to IP

- Host
- Interface address (IP address)
- Datagram/Packet
- Router

20



Zooming In:

Interest	Incoming face
/ndn/uci/content	face0, face3

Interest: /ndn/uci/content

Every router has a:


- PIT: Pending Interest Table
- CS: Content Store (Cache)
- FIB: Forwarding Information Base

23

Forwarding

- Main operation is prefix-based longest match lookup, like IP
- Interests forwarded according to routing table, but multipoint forwarding, broadcast, local flooding are all okay
- Data follows interest path in reverse

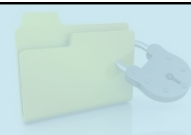
24



Security

- **Now:** secure the pipe
 - Data is authentic because it emanates from the right box (which is an end-point of the right secure pipe)
- **NDN:** Integrity and trust as properties of content
 - Should be inferred from content itself

27




Securing Content: how?

Current SSL/TLS 3-way handshake model is not a good fit for NDN:

- Secures channel, not data
- Authentic content can come from anywhere
- But, access control (and accounting) is difficult
- After content retrieved from origin, it's served by the network (from caches)

28




Authenticity of Content

Content can be retrieved from anywhere by any consumer

- How can it be trusted?
- How do we know who produced it?
- How do we know it is the right content?

29



Securing Content

NDN Content object:

Name
Data
Signature

- **Integrity:** is data intact and complete?
- **Origin:** who asserts this data is an answer?
- **Correctness:** is this (content) an answer to my question (interest)?
- **Bonus feature:** routers can choose to verify content (with caveats)

30




Private Content

Access to content can be restricted, e.g.:

- Encrypt once with a symmetric key
- Symmetric key distributed using “standard” techniques (pigeons?)
- Access control on key rather than content
 - This can make long-term secrecy problematic


31



Trust Model?

- All content is signed
- **Interests are not...**
- NDN is PKI-agnostic
- Application-specific vs network-layer trust

32



NDN: Privacy Benefits


- Interest has no source address/identifier
- Content can be routed without knowing consumer identity and/or location
- One observed interest may correspond to multiple consumers at various locations
- Router caches reduce effectiveness of observers close to producers

33

NDN: Privacy Challenges

- Name privacy in interests
 - **/ndn/us/wikipedia/STDs/herpes**
- Name privacy in content
 - **/ndn/zimbabwe/piratebay/XSOQW(#E@UED\$%.mp3**
- Signature privacy
 - Leaks content publisher identity
 - Classical privacy vs. security conflict
- Cache privacy
 - Detectable hits/misses


34



NDN: Security Benefits

- Simplicity
- All content is signed
- No need for security handshakes in real time
- A producer's public key is a type of content
 - Pull PKC first, then request content

35



NDN: Security Challenges

- State in routers is both a blessing and a curse
- Any such state can be abused
- DoS attacks:
 - Interest Flooding
 - Content Poisoning: proactive & reactive
- Covert Channels & Geo-location
- Content Access Control
- Trust management at the network layer

36

NDN: quick recap

PRODUCER

- Announces name prefixes
- Names and signs content packets
- Injects content by answering interests

CONSUMER

- Generates interest packets referring to content by name
- Receives content, verifies signature, decrypts if necessary

ROUTER

- Routes interests based on (hierarchical) name prefixes – inherently multicast
- Remembers where Interests came from (PIT), returns content along same path
- Optionally caches content (in CS)
- May verify content signatures

37



Some Recent & Ongoing Work

- Anonymous content retrieval
- DoS/DDoS defense:
 - Content poisoning countermeasures
 - Interest flooding mitigation
- Privacy in Router Caching
- Covert channels and Geolocation
- Secure content fragmentation
- NDN security in non-distributive settings
 - Instrumented Environments (actuation/control)
 - Sensor Networks
 - Bidirectional low-latency communication
- Trust Management
- Fragmentation
- Accounting
- Content Deletion
- Negative Acknowledgments
- Access Control
- Key Name Service (PK Discovery)
- Private Content Retrieval

38

Why Name Privacy?

NDN names are expressive and meaningful, but...

- Leak information about requested content
- Easy to filter/censor content, e.g., block everything like:

/ndn/cnn/world-news/russia

However:

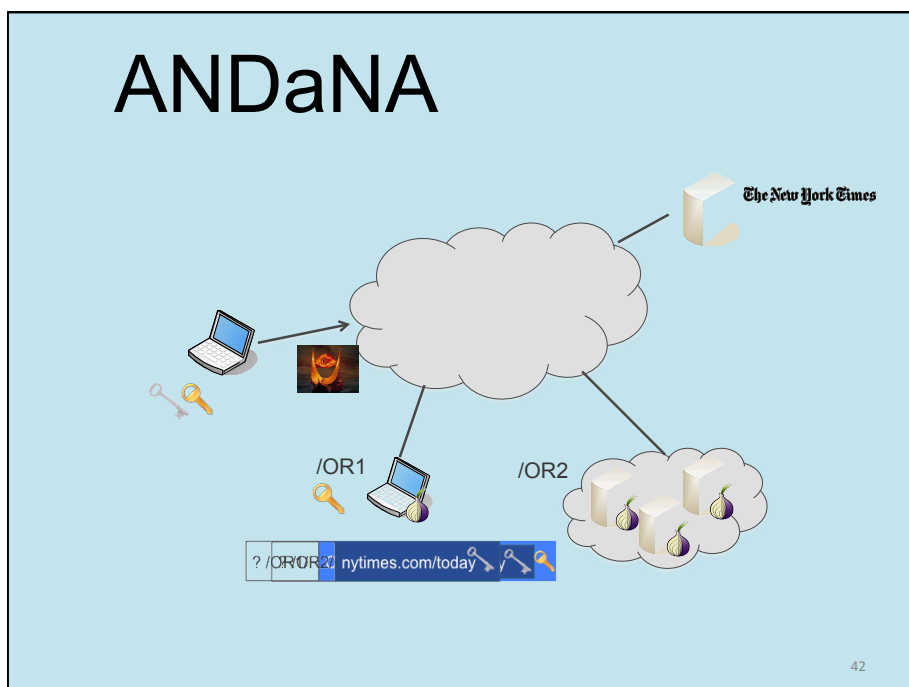
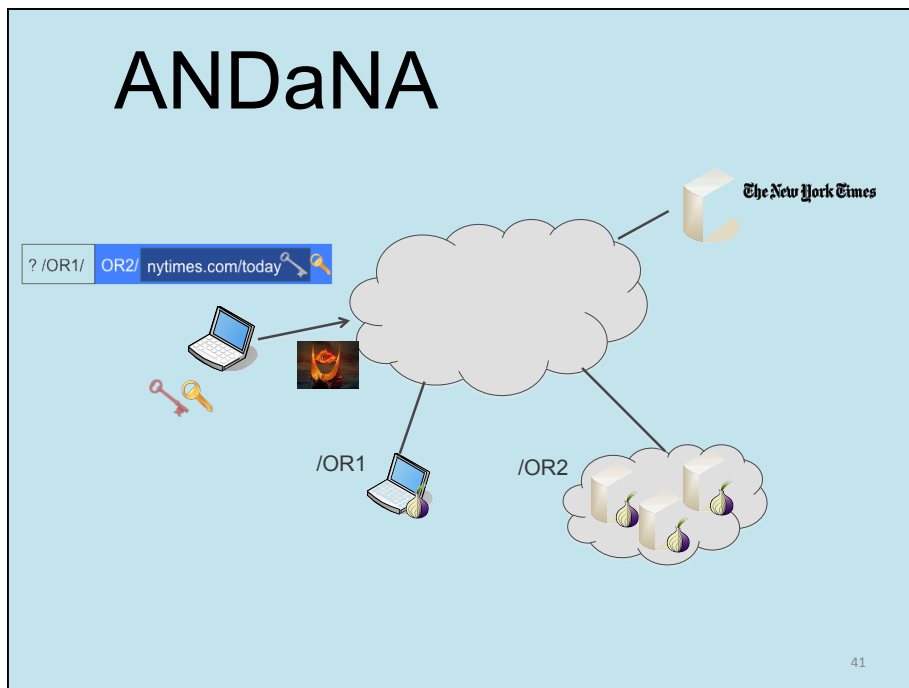
- NDN names are opaque to the network
- Routers only need to know name component boundaries – “/”
- Names can carry binary data

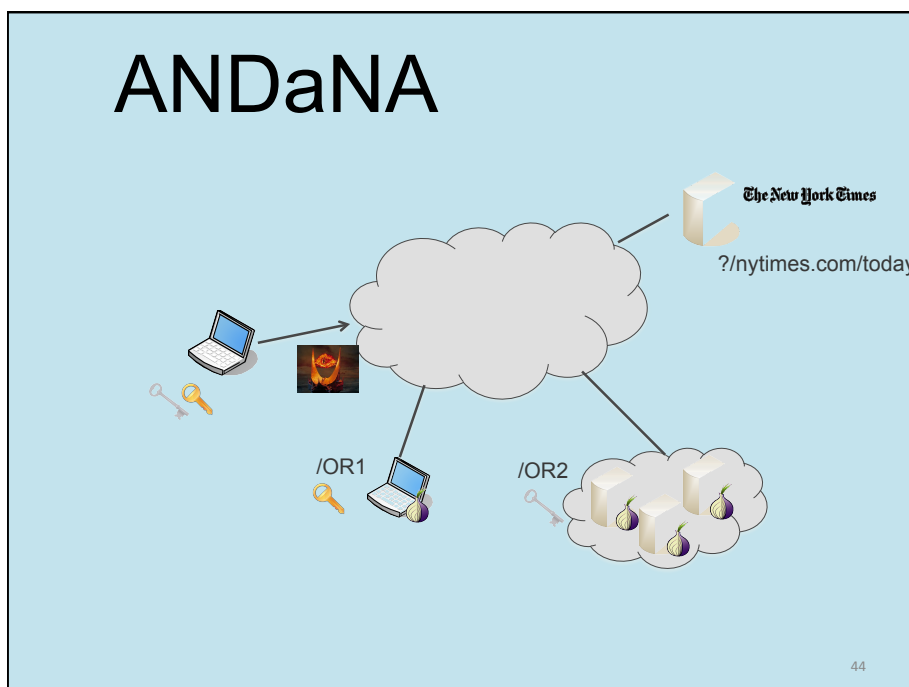
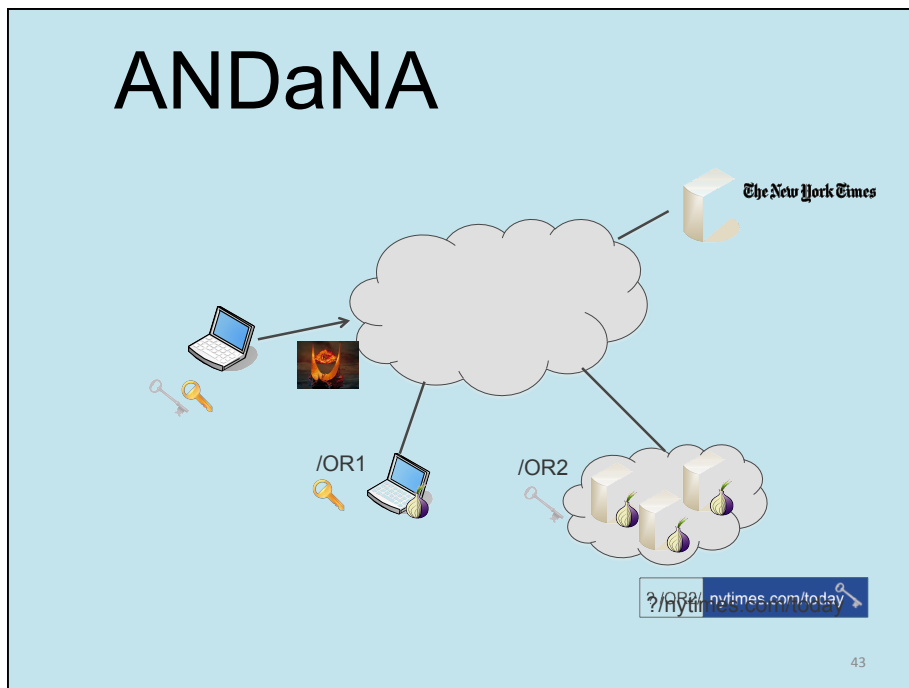
39

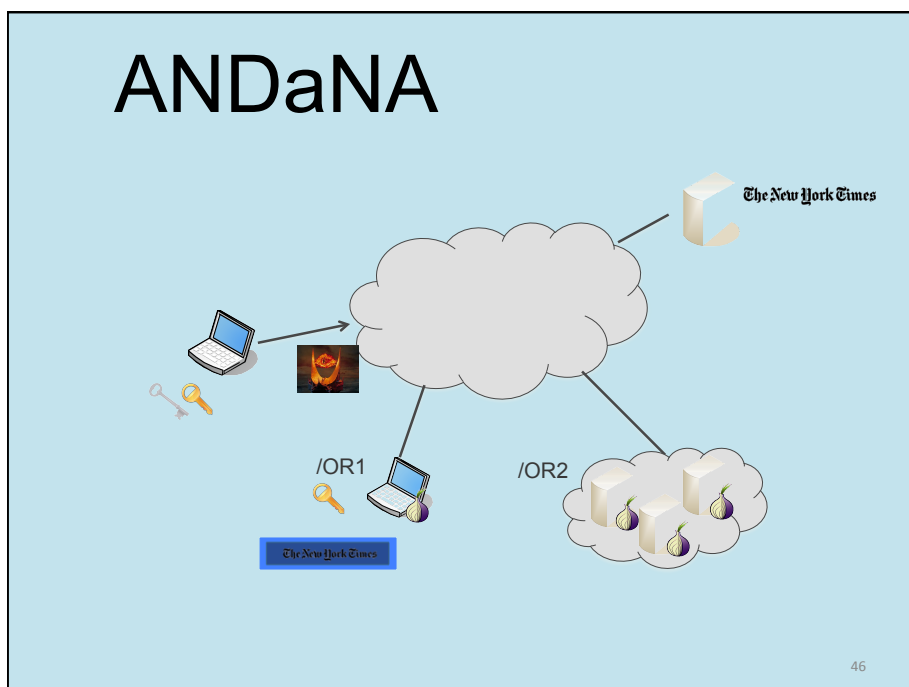
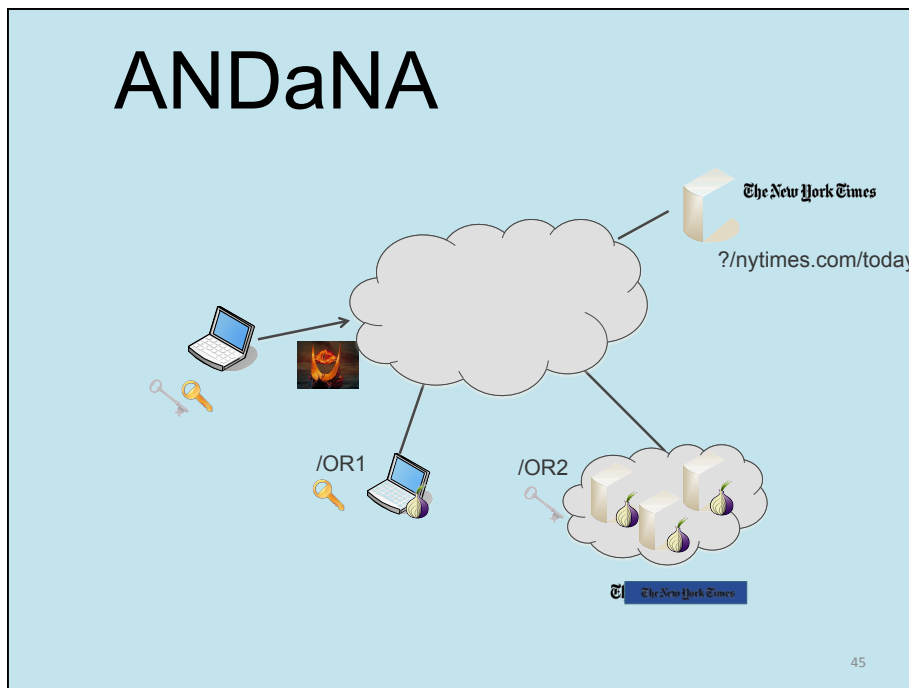
ANDaNA: Anonymous Named Data Networking Application

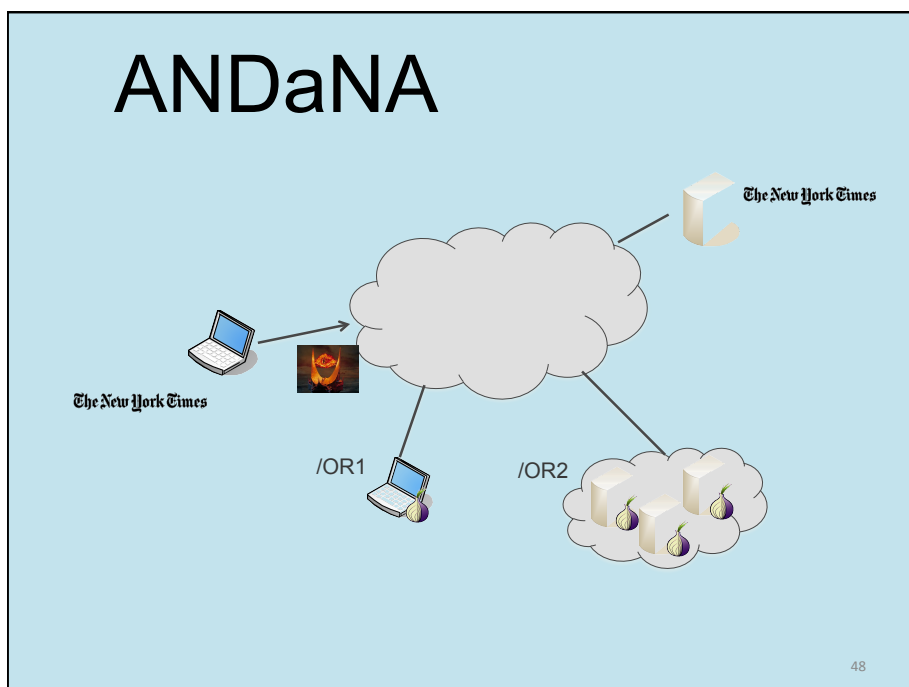
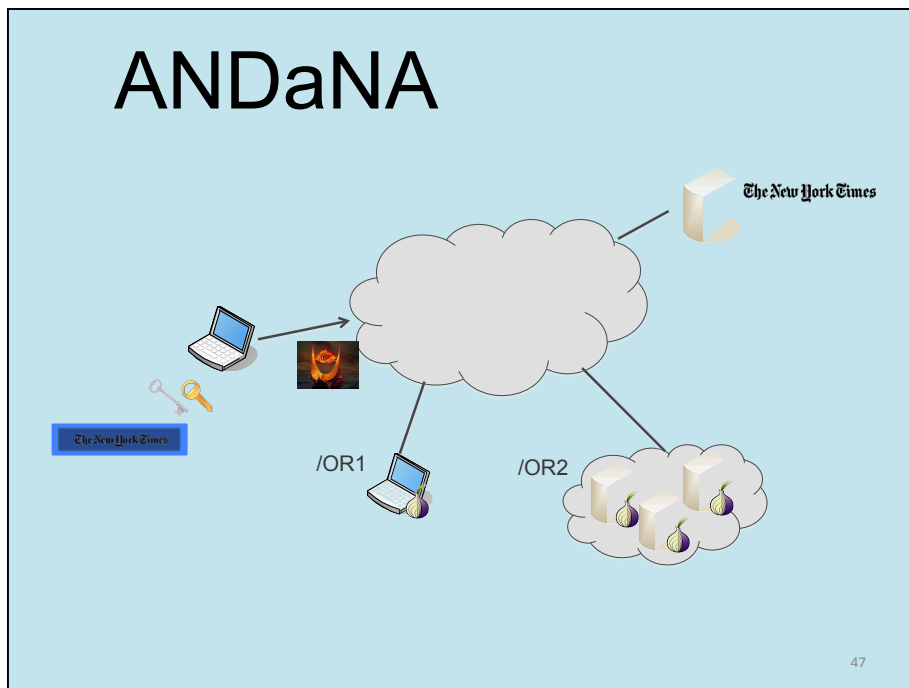
- Observers close to consumer should not learn what content is being requested
- Target: low-to-medium-volume interactive communication
- Producers might not be aware of ANDaNA


[DGTU-NDSS⁴⁰2012]













ANDaNA

Privacy with 2 hops comparable to Tor with 3

- Why? Lack of source address in interests
- Anonymizing routers do not learn origin of traffic (only the previous hop)
- Lower overhead


49



NDN Cache Privacy

- Router Caching is good for performance
 - Better bandwidth utilization
 - Lower latency
- But... bad for privacy
 - Timing attacks
 - Cache harvesting attacks

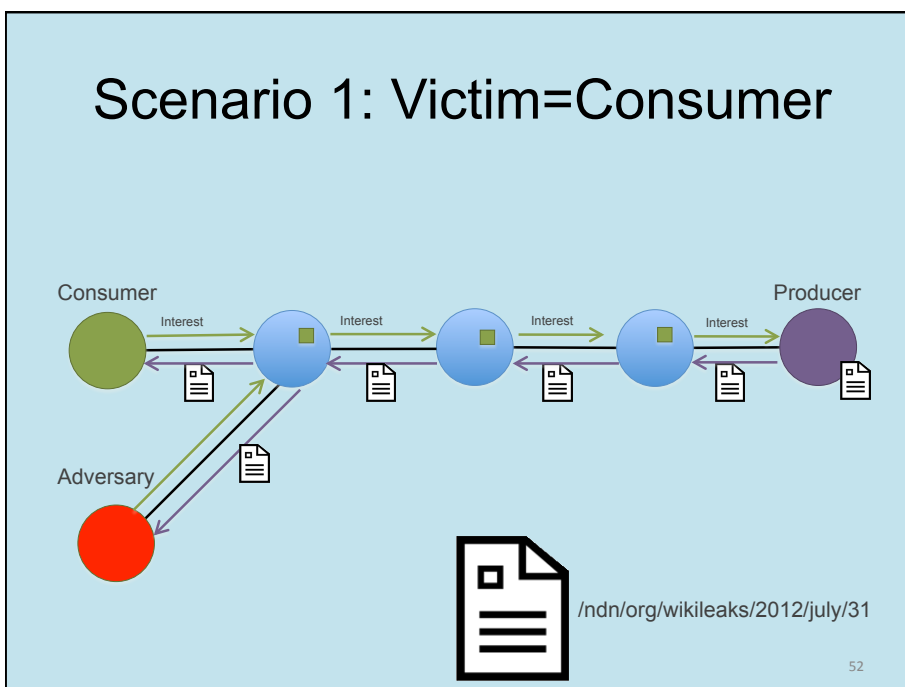
50




Cache Privacy

- Who could the adversary be?
 - Another host or router
 - A malicious application on victim's device
- Where could the adversary be?
 - Near consumer, e.g., on the same LAN/WLAN segment
 - Near producer (opposite sides of first hop router)
 - In both places at once

51




Scenario 2: Victim=Producer



/ndn/org/wikileaks/2012/july/31

53

Scenario 3: Victims=Both



Are Alice and Bob talking?

54

Countermeasures

- Do not cache content at all
 - Bad idea...
- Cache and delay
 - Which content? Who decides?
 - How long to delay?

55

Countermeasures

- Two types of traffic:
 - Private
 - Non-private
- Two communication types:
 - Low-latency (interactive) traffic
 - Use unpredictable content names
 - Content distribution traffic; details in paper, IEEE ICDCS'13
 - Random delay
 - Content-specific delay
- Introduce a privacy bit in interests and/or content?

56

DoS/DDoS in NDN

57

DoD/DDoS Resistance?

Some current DoS + DDoS attacks become irrelevant in the NDN architecture

- Content caching mitigates targeted DoS
- Content is **not** forwarded without prior state set up by interest(s)
- Multiple interests for same content are collapsed
- Only one copy of content per “interested” interface is returned
- Consumer can’t be “hosed” with unsolicited content

58

DoS/DDoS

- Attacks on infrastructure
 - Loop-holing/black-holing
 - Interest flooding
 - Router resource exhaustion
- Attacks on Consumers + router caches
 - Content flooding
 - Cache pollution
 - Content/cache poisoning

59

Interest Flooding

Adversary generates numerous non-sensical interests, e.g.:

`/ndn/us/ca/uc/uci/cs/gene.tsudik/random-string`

Any legitimate producer prefix

- Guaranteed to reach the producer
- Consumes precious router resources (PIT entries)
- IF attack affects both routers and producers

60

Interest Flooding

Potential countermeasures:

1. Unilateral rate limiting/throttling

- Resource allocation determined by router state

2. Collaborative rate limiting/throttling

- Routers push back attacks by interacting with neighbors

61

Content Poisoning

1. Adversary is on the path to producer (e.g., a router)

- Intercepts genuine interest, replies with fake content
- Content settles in routers

2. Adversary is NOT on the path to producer

- Anticipates demand for content
- Issues own interest(s), replies with fake content
- Content settles in routers

62

Content Poisoning

Potential countermeasures:

- Signature verification in routers?
- Consumer feedback?
- AS egress router verification only?

BTW: what is "fake" content?

- Bad signature (fails verification) ,
- Bad signing key

63

Content Poisoning Mitigation

- NDN objective is content distribution
- Facilitated by caches + PITs in routers
- Consumer must verify content signatures
- But ... how to flush fake content from router caches?
- NDN allows exclusion filters in interests (by hash)
 - Can be used, with very limited efficacy
 - Immediate flush: DoS
 - Verifying signatures: expensive + another DoS type
- Consumer authentication contradicts interest opacity

64

Public Keys in NDN

- A public key is a type of content, i.e., a certificate
- Contains authorized name prefixe(s):

- For example:

`/cnn/usa/web/key`

OR

`/verisign/europe/key`

65

Content Poisoning

Two reasons:

- Ambiguous interests
- No unified trust model: applications are diverse and dynamic

AXIOM: Network-layer trust and content poisoning are inseparable

Routers should do minimal work:

- **Not** verify/fetch public keys (except for routing)
- Do bounded, fixed amount of work per content
 - e.g., verify at most one signature

66

Interest-Key Binding Rule (IKB)

IKB (general): An interest must reflect the trust context of the consumer's application, thus making it (easily) enforceable at the network layer

IKB (NDN/CCN): An interest must reflect the public key of the content producer

67

Interest-Key Binding Rule (contd.)

IKB (NDN/CCN): An interest must reflect the public key of the content producer

- Make `PublisherPublicKeyDigest` (PPKD) field mandatory in every interest
- Consumers obtain and validate keys, using
 - Pre-installed root keys
 - Key Name Service (KNS)
 - Global search-based service

68

Interest-Key Binding Rule (contd.)

- Producer:
 - Includes public key in each content's `KeyLocator` field
- Router:
 - Matches `KeyLocator` digest to PPKD in PIT
 - Verifies signature using `KeyLocator`
 - **No fetching, storing, parsing of public keys**
 - Note: PIT entry collapsing takes PPKD into account

69

Is this Secure?

CLAIM:

Adherence to IKB → security against content poisoning

- Assume:
 - All nodes abide by IKB
 - Consumer not malicious
 - Consumer-facing routers – not malicious
 - Consumer ← → first-hop router link not compromised

70

Is this Secure?

- Consumer sends interest containing PPKD
- Router ensures that:
 - Valid content signature using key in `KeyLocator`
 - Digest of `KeyLocator` matches PPKD in PIT
- Consumer-facing router not malicious → only possibility of poisoned content is **hash collision**
- If upstream malicious routers send fake content:
 - Consumer-facing router detects and drops it

71

Optimizations

- Include keys in interest:
 - ✓ Save storage
 - x Requires changes to interest & content structure
- Only AS border routers implement IKB
 - ✓ Better performance
 - x Possible attacks within AS
 - But ... detectable by border routers

NOTE: each router must at least do a PPKD match

72

Optimizations (contd.)

- Self-Certifying Name (SCN)
 - Hash of content (including name) as last component of name
- Benign consumers use SCN → network delivers “valid” content
- **No** signature verification by routers:
 - Only one hash re-computation
- How to get content hash in the first place?

73

Catalogs and SCN-s

A catalog:

- An authenticated (signed) data structure
 - Contains one or more SCN-s, nesting arbitrary
 - Any authenticated data structure
 - Hash chains, MHTs, skip-lists, etc.
 - Structure is application-specific
 - Use IKB to bootstrap (fetch catalogs)
- SCN obtained from a catalog:
 - ✓ **No** signature verification by routers/consumers
 - ✓ **No** need to sign content by producers

74

Two types of traffic

1. Content Distribution, e.g.:

- Video streaming:
 - One big catalog containing SCNs of all segments
 - Or, hash chains (with data), or MHT, etc.
- For example, Web browsing:
 - HTML file as a catalog
 - Contains SCN of sub-pages/components
 - Works only for static content

75

Two types of Traffic (contd.)

2. Interactive Traffic

- Content generated on demand (real-time), e.g., audio/video conferencing,
- Catalogs not viable
- Content must be requested by setting PPKD in interest

76

Content NACKs: what if?

- Consumer obtains hash H of content C from P's catalog
- Consumer generates interest for C referring to H
- But, C is no longer available at P
- P receives interest and ???
 - Drops it – bad for Consumer
- Or:
 - NACK-s it – routers will drop the NACK since a NACK's hash doesn't match H

Bottom-line: need to augment iKB and interest format to allow for SCN-carrying interests to still refer to P's public key.

77

Fragmentation

- Internet connects heterogeneous devices over heterogeneous links, with different:
 - Physical layers (copper, fiber, radio, laser)
 - MAC layers
 - Maximum Transmission Unit (MTUs)
 - Determined by MAC layer

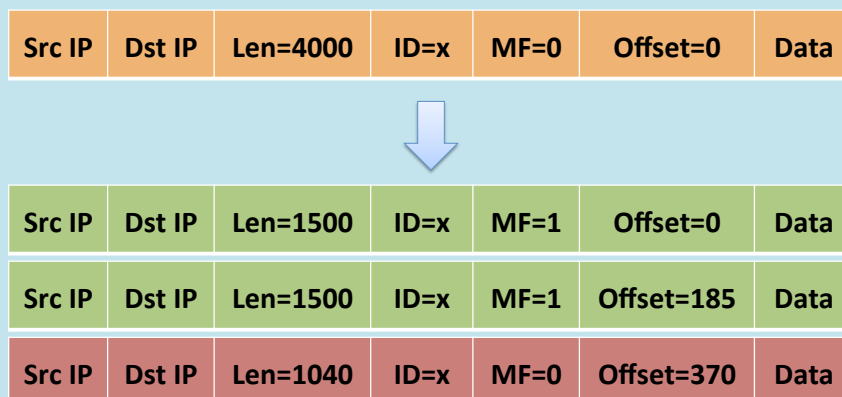
78

Fragmentation

- Fragmentation: splitting a packet into fragments that fit into outgoing link MTU
 - Fragment header encodes ordering of related fragments
 - Re-fragmentation can occur if smaller MTU is encountered

79

Fragmentation – IPv4



80

Fragmentation – IPv4


- **Issues:**
 - Several attacks
 - Ping of death
 - Tiny fragment
 - Router overhead and code complexity
- **Results:**
 - Deprecated in IPv6 and limited to source-based fragmentation

C. Kent and J. Mogul, Fragmentation considered harmful, SIGCOMM 1987.


81

Fragmentation – NDN/CCN

- **Two messages types:**
 - Interest message
 - Content object



Interest



Content

82

Interest Fragmentation

- Recall that NDN interests are processed using referenced content name
 - Names can be of arbitrary size
 - Longest-prefix match on a name requires the entire name before performing a search
- Intermediate fragmentation & reassembly for interests is unavoidable

83


Content Segmentation

- Segmentation (at source) can avoid fragmentation
 - Data segmented by application
 - Signature computed per segment
- Segments are numbered
 - /youtube/dancingcats/s0
 - /youtube/dancingcats/s1

84

Content Segmentation


- Use path MTU discovery
 - Mark interests with smallest transmit MTU in a path



85

Content Segmentation


- Use path MTU discovery
 - Mark interests with smallest transmit MTU in a path



86

Content Segmentation


- Use path MTU discovery
 - Mark interests with smallest transmit MTU in a path



87

Content Segmentation


- Use path MTU discovery
 - Mark interests with smallest transmit MTU in a path



88

Content Segmentation


- Use path MTU discovery
 - Mark interests with smallest transmit MTU in a path



89

Content Segmentation

- Use path MTU discovery
 - Mark interests with smallest transmit MTU in a path



90

Content Segmentation

- Problem
 - Producer cannot segment for all MTUs

91

Content Segmentation

- Problem
 - Producer cannot segment for all MTUs

**Content Intermediate re-fragmentation is
unavoidable**

92

Content Segmentation

- In CCN/NDN:
 - Routers are not required to verify signatures
 - But... they might

93

Content Segmentation

- In CCN/NDN:
 - Routers are not required to verify signatures
 - But... they might

**Provide content authentication without
intermediate reassembly**

94

Content Fragmentation

- In today's Internet
 - Packet fragments might not follow same path
- In CCN/NDN:
 - All content fragments follow the same path
 - But... out of order delivery is possible, even between adjacent routers
 - Parallel links with different speeds and/or loss/error

95

Content Fragmentation

FIGOA: Fragmentation with
Integrity **G**uarantees and **O**ptional
Authentication

96

Content Fragmentation

- **FIGOA** supports:
 - Cut-through switching & optional intermediate reassembly
 - Security via *Delayed Authentication*
 - Also supports integrity with optional authenticity

97


Content Fragmentation

- **FIGOA** supports:
 - Cut-through switching & optional intermediate reassembly
 - Security via *Delayed Authentication*
 - Also supports integrity with optional authenticity
- Not CCN/NDN-specific
- Works with any network architecture with path consistency guarantees

98

Content Fragmentation


- Hash-and-sign



99

Content Fragmentation


- Merkle-Damgard construction



100

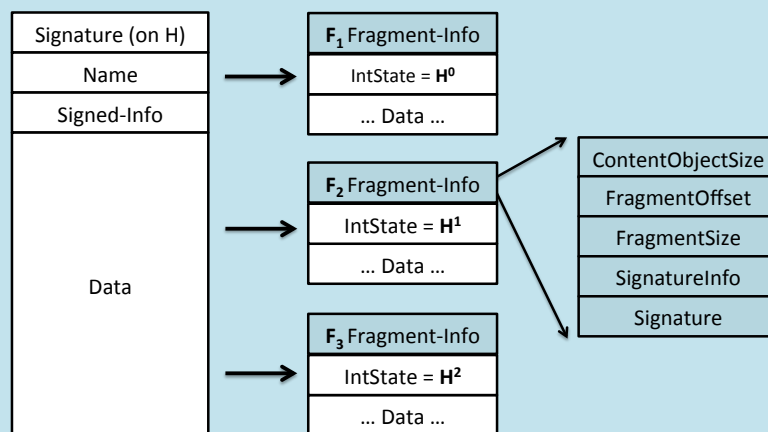
Content Fragmentation

- Merkle-Damgard construction




101

Content Fragmentation




102

Content Fragmentation




103

Content Fragmentation



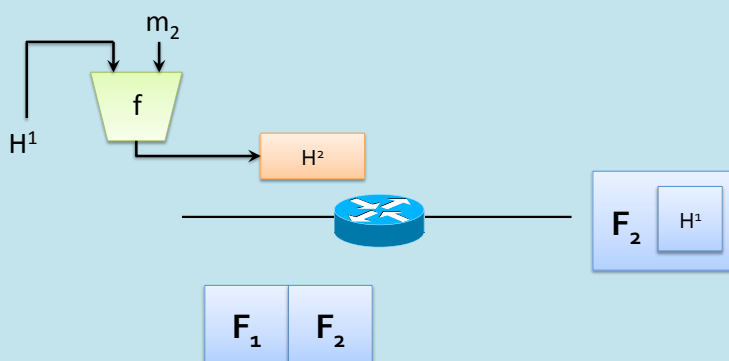
104

Content Fragmentation




105

Content Fragmentation



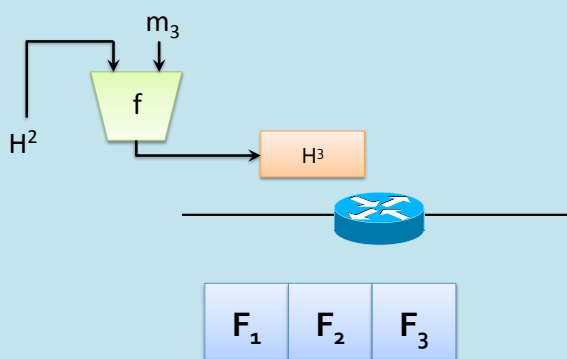
106

Content Fragmentation



107


Content Fragmentation



Verify signature (contained in F_3) using H^3


108

Content Fragmentation



109

Content Fragmentation



110

Content Fragmentation

- Hash is computed gradually
- FIGOA works with out-of-order fragments
 - F_1 is received \rightarrow calculate H^1
 - F_3 is received \rightarrow calculate H^3
 - F_2 is received \rightarrow calculate H^2
 - Ensure calculated IntState = received ones
 - Verify signature

111

Content Fragmentation

- If content is cached:
 - Routers store fragment info, including data
 - Content can be cached fragmented or assembled
- If content is not cached, routers store:
 - Fragments offsets
 - Intermediate state

112

Conclusion

- Fragmentation is a must in CCN/NDN
 - Interest
 - Content
- Segmentation does not avoid fragmentation
- Neither does MTU discovery

FIGOA: Fragmentation with **I**ntegrity **G**uarantees and **O**ptional **A**uthentication

113

NDN S&P References

- S. DiBenedetto, P. Gasti, G. Tsudik and E. Uzun,
ANDaNA: Anonymous Named Data Networking Application, NDSS 2012.
- J. Burke, P. Gasti, N. Nathan and G. Tsudik,
Securing Instrumented Environments over Content-Centric Networking:
the Case of Lighting Control via Named-Data Networking, IEEE NOMEN 2013.
- G. Acs, M. Conti, C. Ghali, P. Gasti and G. Tsudik,
Cache Privacy in Name-Data Networking, IEEE ICDCS 2013.
- P. Gasti, G. Tsudik, E. Uzun, and L. Zhang,
DoS & DDoS in Named-Data Networking, IEEE ICCCN 2013.
- A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun and L. Zhang,
Interest Flooding Attack and Countermeasures in Named Data Networking, IFIP Networking 2013.
- A. Compagno, M. Conti, P. Gasti and G. Tsudik
Poseidon: Mitigating Interest Flooding DDoS Attacks in Named Data Networking, IEEE LCN 2013.
- C. Ghali, G. Tsudik and E. Uzun,
Elements of Trust in Named-Data Networking, ACM SIGCOMM CCR, October 2014.
- M. Conti, P. Gasti and G. Tsudik,
Exploring Covert Channels in Named Data Networking, AsiaCCS 2014.
- A. Compagno, M. Conti, P. Gasti, L. Mancini and G. Tsudik
"Violating Consumer Anonymity: Geo-locating Nodes in Named Data Networking", ACNS 2015.
- A. Compagno, M. Conti, C. Ghali and G. Tsudik,
To NACK or not to NACK? Negative Acknowledgments in Information-Centric Networking, IEEE ICCCN 2015.

114