Minutes IRTF ICNRG Interim meeting in Berlin Sunday July 17th, 2016

Agenda

9:00-10:30 Session 1 - Routing & Caching -- (Minutes taker: Chris Wood)

- Welcome, Agenda Bashing, Minutes takers, Bluesheets - Chairs (10 min)
- Native Content Distribution through Off-Path Content Discovery - Ioannis Psaras (30 min)
- CDN pain points and TLS deployment (a summary of interviews with companies [Fastly, CloudFlare, Netflix, EFF] in the bay area) - Chris Wood (20 min)
- Routing/Locators discussion - Marc Mosko (30 min)

11:00-12:30 Session 2 - Security Topics -- (Minutes taker: Ioannis Psaras)

- Secure on-boarding - Ralph Droms (15 min)
- Object security and protected object manipulation in caches - Jörg Ott (15 min)
- Transport offload, or how to separate data encryption and authentication for encrypted PEPs - Chris Wood (15 min)
- Broadcast and group key encryption - Chris Wood (15 min)

14:00-15:00 Session 3 -- (Minutes taker: Ravi Ravindran)

- POINT project - Dirk Trossen (30 min)
- Name Resolution
- Requirements for Name Resolution Service in ICN - Jungha Hong (20 min)
- draft-dong-nrs-requirements-00.txt - Cedric Westphal (10 mins)
- draft available at the ICNRG Wiki (below)

15:30-17:00 Session 4 -- (Minutes taker: Marc Mosko)

- Planning for future Interop / Hackathon (20 min)
- CCNx test rig (design and running snapshot) - Chris Wood
- Terminology Draft Update - Bastiaan Wissingh (20 min)
- https://datatracker.ietf.org/doc/draft-wissingh-icnrg-terminology/
- https://www.ietf.org/internet-drafts/draft-zhang-icnrg-icniot-architecture-00.txt : this is the ICN-IoT architecture draft focusing on a middleware design and solutions - Ravi Ravindran (15 min)
- https://www.ietf.org/internet-drafts/draft-ravi-icnrg-ccn-notification-00.txt : Proposes PUSH in CCN - Ravi Ravindran (15 min)
- Dagstuhl 16251 report and takeaways - Chris Wood (15 min)
- ACM ICN-2016 conference accepted papers, http://conferences2.sigcomm.org/acm-icn/2016/index.php
- Authors of papers can give 2-3 minute summaries (3 confirmed so far)
- The presentations are available at: https://datatracker.ietf.org/meeting/interim-2016-icnrg-03/session/icnrg/

Agenda items not presented:

- Secure replicas and nomad sessions with CCNxKE - Chris Wood (15 min)

Minutes

Session 1 - Routing & Caching

Native Content Distribution through Off-Path Content Discovery - Ioannis Psaras (30 min)

- DaveO: Does the D-FIB have any notion of the distance to an off-path cached copy?
    - Not yet -- working on it
    - One could use interest hop counts as a rough estimate of this distance
- Nacho: You have an extra table per node such that all requests are logged? Help routers forward requests to other downstream nodes, i.e., other edge caches
- Nacho: D-FIBsize = 20*cache size in terms of objects?
    - Yes
- Nacho: Why is impact of router traffic going up?
    - Interests are forked and multple responses may be received

- Nacho: NDN has less traffic overhead because cache is hit more?
    - Yes
- Paul: summing up overhead on all of the paths
- Ravi: overhead is in terms of interest packets or content packets?
    - Content (data) packets
- Marc: Nodes are spending more memory on D-FIB than cache -- why not use that memory for cache?
    - It only stores names, not the content itself, so it's small(er) than the cache
- Dave: D-FIB is significantly smaller if content sizes are small. Overhead is related to the size of items in the cache
- Dave: No aggregation in the D-FIB, right? No compression?
    - Right. The size is O(# objects fetched).
- Nacho: Wasteful without caching at the edge?
    - Yes
- Chris: An attacker can fill up DFIB with random stuff.
    - Cache pollution will lead to DFIB pollution
    - Similar to PIT table, it's per packet state

CDN pain points and TLS deployment - Chris Wood

- Marc: TLS terminator impersonates the Origin server? Yes. Everything in the POP is cleartext in the cache service.
- DaveO: Producer's private keys exposed in the 100s of physical locations (POPs), as opposed to the normal "locked up in one place" private keys.
- DaveO: HTTPS everywhere. What do they think they are getting? No data in the clear? Protect the producer's services & data? Preserve anonymity of clients? THink most businesses care about 1st two, not 3rd, whereas IETF is concerned about privacy of client.
    - Netflix, for example, does it because everyone else does it. It's cheap. They are already transfering encrypted content (the movies).
- What is EFF? Electronic frontier foundation.
- Paul: The data itself is not https? They are in transition phase, because data already encrypted.
- DaveO: HTTPS w/o crypto? No, full HTTPS. It's minimal overhead, AES & chacha is pretty quick.
- Nacho: It's not trivial to add HTTPS everywhere.
- DaveO: Distributing long form content, setup is small and data phase is giant majority. Not like transactional where setup is high overhead.
- Nacho: Setting up HTTPs internally is a lot of work.
- DaveO: pre- or post- eDNS? pre.
- Dirk: multicast, you can get addresses that don't work any more.
- Why deltas better in ICN? Becaue we're transferring structured data.
- Marc: Caching API? Named function networking
- DaveO: (1) replacing content with new content whose value is different, (2) purging content so it is no longer fetchable. They are different. The latter is extrememly hard.
- Chris: In Fastly no versioning, so just have to keep hitting same name.
- DaveO: Difference between getting latest data and never getting old data -- Purging means you cannot get old data.
- Chris: In CDN world, they are confounded. My understanding is they are the same.
- Cedric: Do today's CDNs talk about federations?
- DaveO: stuck at IESG
- Ravi: Akamei? How does it compare? Chris: Didn't talk with them.
- ??: How did you chose providers? Chris: Picked companies near me and asked them. These are the ones I got responses.
- ??: Pain points? Who defined? Chris: They defined the pain points.
- Dirk: You will continue talking with them? Chris: Yes.

Routing/Locators discussion - Marc Mosko

- Lixia: the earlier efforts (8+8 included) are not really about locator/identifier issue --the problem they tried to address is routing scalability
- Lixia: RFC1955 maps locater-to-another locater (IP address of a host), not identifiers
    - This was the first documented map-encap scheme, outcome of IAB ROAD study group
- Ravi: Which of these are used?
    - LISP is in limited use, none others
- Lixia: background of this work is routing scalability -- the notion of identifier has been continually evolving. It used to be using an IP address to identify *host* (box), HIP started the notion of supporting mobility and introduced host IDs that are not IP addresses, nonetheless host IDs. in ICN we mean data identifiers.
- DaveO: that's right. There were a number of routing scalability problems: size scalability for Internet, multi-homing and mobility meant that topological identifiers would fall apart. Even non-topological things were intended to identify boxes, not data. ICN reformulates the meaning of an identifier to data.

- DaveO: goal: try to make it seem that some "state" existed in a specific box (?)
    - Not as part of the IP architecture, but as something above and behind the scenes, hence the addtional complexity to handle it.
- NDN Link concept: Lixia: you're using an old reference
- Lixia: instead of enumerating mechanisms, more interesting to talk about *what maps what, and for what reason*. In MobilityFirst, there are two levels of mapping: application names to ID --they need this step (instead of app name directly to location) for security purpose, then ID to location.
- DaveO: mapping to IP address only because IP was the underlay
- DaveO: MF did some interesting work with scope-based GUID mapping
- Ravi: GNRS is contextual -- not just a name lookup.
- Ravi: MF addressed security by self-certifying names.
- Lixia: Yes, all functions are supported by lookups from servers.
- Marc: Some stuff is (obviously) omitted from the slides -- see publication list for more details. The point was that there are two lookups to go from application names to IP addresses.
- Why NDN has LINKS: Lixia: requesting by name would not scale (names are not necessarily aggregatable) -- so we need to handle scalability, and this is what LINKs and NDNS try do
- Dirk: who puts LINKs in interests?
- Lixia: yes, at the client -- but not necessarily nailed down to a single point
- DaveO: loop detection based on nonces in the presence on nonces?
- Lixia: LINK is not involved in loop detection
- DaveO: so there's still a belief that loop detection works with nonces (right). Adding and removing LINKs from interests does not change the nonce?
    - No
- Lixia: LINK only tries to help guide interest forwarding
- Dirk: a LINK can be a name in the same namespace
    - Yes -- everything is a name
- Ralph: is the PIT entry based on the name or LINK?
    - Name -- links are just forwarding hints
- Clarifying concept in fetching nameless content:
- Lixia: if a content does not have a name, then what is the name carried in the interest that wants to fetch that content?
    - the interest carries no name, but carries the hash
- Ralph: with nameless objects, interest is just a locator
- Nacho: only requirement in interest name is that it's forwardable to *something*
- Ralph: these are roughly the same thing, i.e., where are things placed in packet fields
    - In CCN: there's only one locator
    - In NDN: there can be multiple links
- DaveO: forwarded on name and matched on hash
- Nacho: the purpose of the name is to forward stuff
- DaveO: with LINKs, there can be multiple locators. In the other mindset, the object has a single name -- the hash
- Ralph: Without a hash, then the name is both a locator and identifier. With the hash, the name is a locator.
- Nacho: All names are locators, but sometimes they are also identifiers.
- DaveO: locators and identifiers just confuse people
- Ralph: but we need a differentiation between the unique name and how to find it. Since we sometimes overload locators/IDs, it can be confusing
- DaveO: let's ban the "locator and ID" terminology
- ??: Please explain
- DaveO:
    - (1) everything has a hash
    - (2) if I know the hash, I can always ask by hash but the routing system might not be able to deliver it
    - (3) in CCN world we have multiple names
- Paul: do you mean the name is dependent on where it's stored?
- DaveO: no, the point is that we don't care. The point is that we can have as many names as we need in order to have it in the network
- Lixia: this needs a writeup for clarification -- there is confusion between locator and identifier. Maybe this is not necessarily the right model (given history)?
- Lixia: hashing is an optimization for static content -- but it does not work for dynamic content
- Lixia: all names represent some entities, so name in a link represents an entity (att.com), and once requests get into the entity the name carried in the interest can be directly used for further forwarding
- Ravi: there should be a distinction between the application and the network layer
- e.g., applications work with identifiers, network layer announces locators
- Lixia: why?
- Ravi: application determines name, right?
- Marc: there is no clean cut between a locator and identifier in a name
    - This is only known when a request reaches a place where the data is stored

- DirkT: identification is an app layer task, locating is a network layer task. PURSUIT should be added to the list. We've had this discussion before.
- DirkK: is the NDN/CCN difference that NDN locators are explicit, whereas they are implicit in CCN
- Marc: how one receives LINKs/locators is orthogonal to how they're used and where they're put
- Alex: no security context for nameless content object?
- Nacho: not true, it inherits context from the manifest
- Nacho: in theory you could have a CCN packet with multiple names, which would be fine since security for nameless objects is about the hash, not the name (there is none)
- Ravi: a fundamental challenge here is that of routing
- Nacho: if you want to use non-topological names, but for what?
- Nacho: do we want the flexibility of application-specific names in CCN, and do we want to reflect that in the network? @Ravi: do you want any name to be chosen? Put the locator in the application.
- Ravi: are applications only working with one name? how do you distribute data around the network in this case?
    - Nacho: no one said this was the case
- Ravi: how do we deal with mobility in this case? do apps just keep flooding new names as they move around?
- DirkT: instead of talking about locator/ID splits, what is the *type* of name we want the network to operate on? what is the granularity and semantic richness that we want the network to operate on? Is it topological?
- DirkK: one view is that we don't need the term "locator" (we have alternative names), another view is [missed]
- Nacho: applications will give content names however they want -- we don't want these to be reflected in the network
- Lixia: locator in IP context is a slightly inaccurate term. Old telephone numbers are strictly locators , i.e., in US they have area codes. even though IP inherited point-to-point communication model from telephone networks, IP addresses are no longer tied to specific locations. The address-to-location binding was done through routing announcements. IP prefixes are not locators as in the telephone sense. Point: what is in the routing table is not a locator, but it's just info on how to reach a prefix.
- Lixia: IP addresses are not inherently tied to specific *location*, but only through (dynamic) routing announcements.
- Cedric: we need to define our terms and try to come to some agreement (or at least clearly define the different perspectives)

Session 2 - Security Topics

Secure On-boarding - Ralph Droms

Talk Conclusions:
- secure authentication and authorisation
- resilient to outsider and insider attacks

Ravi: from routing perspective, how do you reach the destination?
Ralph: there is a 2-byte ID in the packet
Ravi: there is some security computation - what is the meaning of this
Ralph: encryption and mac computation/verification (and decryption).
Ravi: (missed the question)
Ralph: all nodes that have joined, they have joined using this scheme. Starting f
Ravi: New users joining, do they depend on the leaf nodes to do the authentication?
Ralph: yes, this is what is happening.

Mutable Data in ICNs - Object security and protected object manipulation - Jorg Ott

DaveO: how do you deal with order (partial, full etc.)
Joerg: deliberately not talking about order here. There is no absolute order in the network, application would have to deal with this - work in progress still
DaveO: what's the meaning of authorised edit/merge?
Joerg: someone who is authorised to edit the doc/project
DirkK: you have built the application but have not done the merging yet
Joerg: yes, still to work on the rest
MarcM: the hash tree that git has got will be helpful here.

Secure Transport Offload with Encrypted PEPs - Chris Wood

DaveO: MCTLS can do two things that this cannot
1) you can chain the middleboxes with MCTiOS - with this you can't
ChrisW: yes you cannot exactly do that
2) MCTLS can support partial encryption/decryption
ChrisW: yes, this can be done

Joerg: since this is focusing at transport, the interesting arch question: how can you build an architecture that integrates that? Will it be a failure if ICN needs to redesign PEPs?
MarcM:
KostasP: the arch that you're describing is 2006.
DirkK: the question is whether the arch changes can support this - it doesn't matter if it's 2G, 3G or 5G.
MarcM: Bring the content as close as to the basestation as possible. You're always going to have to fetch the content from far away.
Ravi: who owns the PEPs?
MarcM: it doesn't matter. There are lots of different deployment scenarios.

Group Key Encryption - Chris Wood

DaveO: it seems like for some applications providers want to change keys quite rapidly. so with this, you'd have to create manifests again and again.
MarcM: you'd have to do this within the manifest tree

DaveO: I think I understand why you think there's no threat in having the nameless root manifests unencrypted. If you encrypted that, it would be much more difficult for someone to see inside it.
ChrisW: that's correct.
DaveO: it's easier to intercept requests for manifests than requests of everything
ChrisW: agreed

Session 3 -- (Minutes taker: Ravi Ravindran)

FLIPS : Flexible IP Services (over ICN) - POINT project
Dirk Trossen - InterDigital, Europe

- European efforts, POINT and RIFE

- Practical SDN approach for Flexible Routing
Challenge - HTTP unicasting
- HTTP/ICN gateway
- Surrogate Service, edge caching, near to zero second integration with routing fabric

DK- how is HTTPS handled ?
DT: Handled at the NAP, certifiacates handled between the client and the NAP. Currently uses HTTP.

Demo of this work was shown in MWC, Barcelona

Resolution optimization using cache state in the cNAP and sNAP

Multicasting using OR of pathIDs at the sNAP

sdn integration - opendayligh
Mobility, HTTPS, resilience, COAP support for IoT

Planed at ITU-T IMT2020 ICN PoC

ICN based mluticast overlay - BIER WG

Georgios : Relation between MEC, SDN and ICN, What the complexity with HTTPS ?
DT: : MEC is a separate use case, Services come up dynmically, SDN is used for forwarding, no growing flow tables. Easy integration with SDN, BIER similarly uses bit field based on link notion., Networking monitoring, underlying ICN substrate. Surrogate management, reacts on delay, base don user QoE.

Ravi : does the surrogate service affect the DNS name registration
DT: No, that is only exposed to the ICN routing.

Q. Which one of the surroagate has the content, how do you know this ?
DT: This is part of the Randezvou server state

Name Resolution System,
Requirements for name resolution service in ICN
- Jungha Hong

Goal of the document : motivates the use of NRS in ICN

Why we need NRS ..
- to support Flat name (ID) - self-certifiying IDs..
- Mobility - provider, host mobility

Use case 1:

Name to locators

MM: In using flat name in NRS, there needs to some trust for a client to register to the NRS ?
JH: Security issue
BO: Use the name authority, to generate a certificate for the client registration

Nacho: What DNS cannot do ?
JH: Flat naming handling and mobility

Nacho : can we use flat names in DNS ?, or what is the part that DNS if it doesnt support ?
BO: DNS is a viable solution, we may also be able to mobility
JH : Mobility has a challenge using DNS
DK: Also depends on how the locators are defined.

Nacho : Is it missing a feature in DNS to support the NRS document ?

Use case 2: Name to name (alias)

Use Case 3: Name to IP

Requirements : Scalability, low latency, fast update, low maintannce cost, locality deployability, Resilienc, fault isolation, security challenges

Questions:
Dk: More on the mailing list

Requirmemtns fo NRS in ICN : Cedric

MM : today, we combine, name resolution to replic a ocation, a using DNS, how should ICN operate, in terms of getting the best record ?
CW: No guarantees..

JS how do we get sometign without the client location ?
CW: is it similar to CDN ..

Session 4

CCNx Test Rig - Chris Wood

- Policy-specific check slide has typo should be "Link B, Link C"
- DK: is this useful, does icnrg understand purpose? I think it's useful.
- Chris: You can download it right now and check your forwarder.
- BO: If we arranged an interop would people be interested in participating?
- DK: Need to be clear on what to achieve. If we had new feature or new idea, maybe help make a technical conclusion if it's a good idea.
- BO: If you have ideas for this, please take it to mailing list.
- Ravi: This is assuming you are using CCNx 1.0 code. It shows that you are aligning to the spec.
- Chris: Yes.

Terminology Draft

- DK: Long awaited terminology draft posted to the list.
- DO: a lot of discussion on the mailing list now that we have a draft. Best thing right now is for everybody to read the

draft and then respin it once we have feedback.

- DO: Anyone have something to say about it?
- DK: More people could have opinion on it, express it on the mailing list

ICN IoT draft - Ravi Ravindran

- Marc:The title sounds like it's "the" ICN IoT architecture. Is it a specific protocol?
- Ravi: Protocol independent description, then described how it could be done in MF, NDN, CCN.
- Marc: Normally when one does an architecture document for a protocol its called "foo architecture" then you do something like "foo over ipv6", etc. The way this is presented sounds like its the only ICN IoT architecture.
- Ravi: Based on certain in paradymn and principles.
- LZ: ICNRG a research group so it's ok
- DK: What's the maturity level?
- Ravi: We have a CCN-Lite implementation. Security for device discovery not implemented. Service discovery from aggregation point onwards. Pieces validated. A lot of the details, like the constraint part, not done yet. The procedure for discovery similar to NDN, like Jeff Burke's work.
- DK: Yes, there's lots of work happening. For this group, report on what's been built, experience report.
- Ravi: We've not pushed in to papers yet.
- DK: Not sure how much benefit we have for architectures, we need more experience reports.
- Ravi: There is published work on the prototype systems (or we are planning to). We can limit discussion if we need to. We can collect different results from different groups.
- Ravi: Motivating people to think IoT in ICN and motivate challenges.
- LZ: Most useful thing not specific architecture, most useful thing is build it and tell us what works.
- Ravi: Yes, we will try for that.
- Vince Park: Describe the issues you are trying to solve with this design. Connected car, drones, home authomation, agricultural, etc. Some of your decisions, middleware functions, are based on assumptions you are trying to solve. Outline the rationale.
- Ravi: We do prototype. We don't have speicifc numbers like Ralph's presentation. We have basic prototypes and we need to go to next step.

Push for CCN - Ravi Ravindran

- DK: Draft specifies new message type and discusses these options?
- Ravi: yes
- DK: How do you plan to recommend specific approaches? What's future plan?
- Ravi: Just the notification spec. The idea is to motivate that you need this thing. Then address the routing and forwarding.
- DK: Would also be good to hear more about implementation.
- Ravi: we don't have specific simulation results yet.

Report on Dagstuhl Seminar on ICN & Security - Chris Wood

- DO: producer independent from the data it produces?
- CW: Think it was tied up to the data it produces
- Ravi: What about access control
- DO: Not arguing against consume identity, but just asking about producer identity separate from its data
- DO: Do you mean "can" or "may"? (application advertise any name)
- Chris: "can" There hsould not be a mechanism to inject just any name.
- LZ: What do you mean by "ideal name privacy: no more than IP address and port"?
- DO: no more information than is revealed by IP and port
- LZ: What does that mean? what does IP/port reveal?
- Ravi: Name is allowing things to be correlated. Isn't that the major concern?
- DO: Go all the way to PIR.
- Marc: To clarify: generally post-quantum is pluggable, so long as your protocol is not based on things like generating EC keypairs is cheap.
- RD: are we done in we should stop or because we have a complete solution?
- RD: Anything about ICN for IoT?
- Chris: there was a breakout, but don't know
- CW: What IoT would look like, what sort of densities, scalability, etc.
- Marc: data in motion is different than data at rest, different threat models. Data on hard drive has different attacks than data on the wire.

ACM ICN Conference paper summary

- DO: congestion control paper: why rate-based congension control in ICN
- DO: TCP-ICP Carrying TCP over an ICN infrastructure
- Chris: Network names and how to encode them in the network

Conclusion

- Will there be a discussion on the BoF?
- Yes, on Thursday