

Interim ICNRG meeting @ IETF-89 in London Sunday March 2nd, 2014

The focus of this meeting will be our working documents and the new work item "ICN and video distribution" and the proposed new work item "ICN and IoT", other agenda proposals are also welcome.

Time: Sunday March 2nd, 2014, 9.00-17.00 (Confirmed)

Location: Hilton London Metropole, Room: Park Suite (Confirmed)

Etherpad for notes: <https://neclab.titanpad.com/ICNRG-IETF89-2014-03-02>

Presentation material: <http://www.ietf.org/proceedings/interim/2014/03/02/icnrg/proceedings.html>

Attendance: see end of the minutes

Agenda

ICN and video distribution

- Adaptive Video Streaming over ICN (60 min)
- draft-video-streaming-over-ICN (Cedric Westphal)
- Video optimisation management (Daniel Corujo)

ICN and IoT (90 min)

- ICN based Architecture for IoT
- draft-zhang-iot-icn-architecture-00
- Contextualized Information-centric Home Networking (Ravi Ravindran)
- draft-ravindran-cibus-homenet-01
- Consumer Driven Information Freshness Approach for Content Centric Networking (Daniel Corujo)

Web traffic and CCN/NDN (Mark Stapp/Dave? Oran) (30 min)

Transport schemes for ICN (Fei Song) (30 min)

Discussion on Evaluation Methodology (30 min)

draft-irtf-icnrg-evaluation-methodology

Scalable Container-based Routing Scheme for ICN (Joo-Chul Kevin Lee)

Editing sessions for working documents

ICN Baseline Scenarios

ICN Research Challenges

Slides can be found at

<http://www.ietf.org/proceedings/interim/2014/03/02/icnrg/proceedings.html>

Meeting Notes

Video over ICN

(3 sets of slides used - from Westphal/Lederer/Corujo)

Cedric.. Slides

New version of draft (-01)

Dirk Trossen: Is the point to make ICN over P2P because we can? Thought it was effectively P2P anyway.. Isn't it putting big load on routing plane?

Cedric: Partly because don't know what every user has in video... use of swarms

DT: might be useful to reengineer existing protocols over ICN but a big challenge... so what's the objective?

Dave Oran: May be tied up in terminology problems... conflict between ICN and P2P worlds... P2P is trying to solve some problems that don't exist in ICN (e.g. multipath). Is the mapping just a (ICN-P2P object) name mapping problem i.e. in the name resolution

DT: maybe should be a concept mapping issue rather than (re)building P2P protocol

Spiros Spirou: To do DASH over ICN using URI over CCN.. taking ICN for granted.. no changes in ICN.. is this really in scope 'cos this is application layer... maybe this isn't ICN

Cedric: Disagrees.. need to redesign apps to take advantage of ICN capabilities.. may be important to have idea of location... Is this necessary?

ss: Notes that changes are in app layer rather than 'transport layer' i.e. ICN

Jan Seedorf: P2P is an overlay.. surveys say it is a stupid idea.. so why do P2P over ICN? Name resolution has a different granularity in P2P as compared with ICN. How would you move to a totally distributed chunk distribution scheme.

Nacho Solis: Comment on general term (not doc): Should be careful that we don't assume that we *do* have a well defined base to work with rather than working on requirements for that base. Nacho would be happy if we could get the requirements for video distribution and then decide on the 'lines' of architecture to actually deliver these requirements.

DT: Have we already accepted that ICN is better engineered HTTP or something better? If we do better in ICN (dynamic routing etc) then we could maybe do a better job than just reengineering P2P over ICN.. Would like to see less protocol engineering? Do we care about provenance? Educational video for example... provenance is very important there.. Also personalisation. can ICN deliver this sort of thing?

CW: Should it support this sort of thing? Narrows down to discussing requirements,

Dirk K: Discussing mapping to p2p may be a bit mechanistic but do need to think about how to do things like video.

DO: Good to not live entirely on architecture level.. for practicality need to do some protocol engineering. / We probably need to think about 'catalogue operations'... streaming may a good point to start thinking about catalogue operations.. / Issue of object immutability.. drafts don't think

about immutability .. binding between name and video

Ralph Droms: Sensor networking will also need catalogues/ Need to think about a tie up between partial names and catalogues...

(Back to slides)

Service quality assurance for IPTV..

DT: Asks about problems of interest aggregation in IPTV.. problem exists in ip world in trying to aggregate multicast groups. Are lots of sessions a problem?

SS: Problem addressed in draft is that client has to ask for every chunk.. in existing IPTV systems this doesn't exist.

CW: Depends on granularity of chunks.. overhead not an issue if chunks are large.. but then there is a responsibility issue.

DT: Have we agreed on interest aggregation as a capability of ICN? Pull based architecture creates this problem... Not all architectures have this issue... need to be clear what architectural story we are taking about.

DT: (Aside Keshav suggested...)

Georgios K: Question about challenge of size of requests/chunks and rate of requests...

CW: Notes that this is (as DT said) architecture specific.

GK: Need to differentiate between types of content.. generates different types of requirements.. e.g. streaming versus static.

DT: Backward compatibility is a lower priority target: More important to see if we can get advantages from (say) doing IPTV using ICN

Marie-Jose Montpetit: IPTV is *not* multicast most of the time. Need to think more carefully about how IPTV *really* works.. think about the pause button

WC: Thinks that authors of draft *are* pretty knowledgeable

Daniel Corujo : Agrees with M-J

NS: Back to overhead... How do we address the overhead of sending interests in CCN? The CCN people *are* thinking about overheads.. and acknowledges that is being thought about!

(Future Research Challenges)

CW: Request for additional challenges

Move to slide set by Stefan Lederer et al (Klagenfurt) who are doing DASH over CCN prototype
DASC = DASH + CCN

Implementation available (URL in slides) based on CCNX from PARC
DASH VLC plugin available

Notes that 01 version of draft talks about DRM issues but very little visible work so far
Qn: Can we build a Video on Demand service based on ICN technology without DRM???
Depends on business model!

M-JM: W3C has a WG on this subject.. netflix contributor... may go back to end point model rather than content model like cable model partly because DRM systems are incompatible
This would build on using device authentication which is more lightweight than heavy encryption.

May go to some sort of automatic content recognition scheme.

EBD: Issue with what is meant by term 'server' in broadcast encryption proposal.. there is some discussion of whether you have to go back to the originator of content in ICN RG evaluation draft security section. Also notes revocation issue that isn't mentioned.

DT: One could view a situation where caches are able to distribute keys as a form of distributed content delivery network.

NS: (Not a DRM expert) but we may be confusing DRM with network requirements...

DO: Need to separate what is going on with DRM today and the wish list of things people want to be able to do with protected content. Which things *are* we going to do?

DO: e.g. deterministic cache eviction

BO: Next steps discussion will continue at meeting on Monday

DK: Should we discuss DRM more deeply?

JS/SS: IETF doesn't discuss DRM (There is an RFC)

DK: But we are the IRTF so we can
[Break]

Daniel Corujo's slides ICN and Video
Looking at how link layer could be impacted by ICN and application requirements

How to optimize delivery for both requester and server/cache from their requirements.. and how to reflect that in ICN requirements/implementation

M-JM: The sorts of choice mentioned here are not ICN specific. Question: Will you be sending these requests about in a protocol or just doing them locally? Will there be signalling?

DO: Recent work indicates that putting caches in the middle of adaptive video may make delivery worse!

DO: Are you looking at a fixed set of objects and let the network decide which one to deliver or specifying things at a higher level and letting the ICN system decide how to do things in a much more complex way.

DT: What role does transcoding play?

DO: The fundamental question is 'immutability'... The way that calculation in intermediate nodes is involved is vital.. transcoding is just one example.

DC: Hope all these aspects will be covered and both of DT's options will be considered.

(slides again: Content reception Optimization)

DK: (comment) We have had some of these architectural discussions.. is there going to some condensation of layers - terminal vs end point control

DC: Group haven't decided on the solution as yet

(Slides: Content Optimization.Adaptation)

DO: (comment) Mentioned at end: Are ICN mechanisms useful for managing the network? What do RG think about this possibilities

DC: Video as a use case for what we have to do for management

DO: be good if we could do (ICN) network management using ICN (mechanisms)
Programming languages not considered mature until you can write a self-compiler in the language.. [Aside EBD: interesting question for interpretive languages!]

DO's definition of network management: "Everything you haven't yet been able to solve" (or something like that) - here you go: "Network Management is everything you don't know how to do, because if you did know, you would have put in the base system or application and then you would not have called it Network Management" - right!

DC: Looking at management issues in DC's current project (some of it in management draft)

ICN and IoT
=====

(Notetaker Ralph Droms)

ICN based Architecture for IoT, draft-zhang-iot-icn-architecture-00

Review of work previously discussed in Vancouver
Question: can ICN be applied to IoT as an emerging technology?

Slides contain list of popular scenarios and summary of motivation

Architectural requirements from IoT: naming, scalability, resource constraints, traffic characteristics (local/wide-area, unicast/multicast/anycast), contextual communication, mobility, storage and caching, security and privacy, communication reliability, self-organization

DO: Discovery is usually farther down in the architecture:

Ravi Ravindran: We are referring to a more complicated discovery that includes policy evaluation, etc.

Dirk Kutscher: What about actuators?

RR: Push events into the network to control actuators.

DK: Does this require changes to object properties

RR: IoT requires both push and pull models; push required for actuators.

DK: When requesting content, no interest in where content comes from

RR: IOT requires name resolution system to associate content of interest with name; also required for actuators

DT: confusion is still there about discovery; reformulate to be specific about "what is being discovered". In this case, "discovery" is searching information space.

RR: "Discovery" is contextual as in "what information is nearby"

DR: Discovering names, contextualized services, ???

RR: Anything that the application is interested in

DO: If you use "service" you're not talking about ICN; ICN is about data produced by services, not the services themselves

NS: I disagree

Georgios Karagiannis: Might be interested in object or the function

DO: No, object is ICN...not interested in how object is generated

DO: IoT looks at "things" not "services"

RR: Tied to environment, so context is important

RR: Diagram shows an example of distributing the service itself; ICN integrates cache and computing

DT: Feedback on DO's radical proposition: mostly agree, but...(1) can't break what we haven't agreed on, (2) what is a "network" "application" is unknown yet, (3) what is it that could be the architecture and then how do we break it

DK: is context naming really ICN?

ED: perhaps this thing described as a "service" is something that, when you poke it, puts out a different data object; this definition may be a stretch and not terribly useful

DO: when people use the word "service", things get squishy fast

NS: "services" are part of something; maybe overloaded now but we need to discuss it as part of

ICN

NS: how does discovery, services and cabs all come together

RR: this scenario is an example of a general problem

DK: two issues: IoT has a communication model when an actuator gets talked to by a controller; is semantic-level discovery part of ICN? a search of the namespace give better information about where the data is.

GK: agreed; what part of ICN architecture will be used in this IoT architecture

Spiros Spirou: often a prepackaged "black box" which is ICN; this architecture seems to assume some set of functions in ICN. What does IoT need that ICN doesn't have?

RR: Depends on the particular protocol

RD: what are abstractions to pull out?

Contextualized Information-centric Home Networking (Ravi Ravindran), draft-ravindran-cibus-homenet-01

Current home networking:

- single subnet
- heterogeneous systems that can't be interconnected
- non-interoperable standards
- routing requirement for mulitcast, etc.

DO: objects to characterization of homenet as "enabling IPv6"; solving every problem on that slide

RR: disagrees; draft does not claim to be solving

DO: homenet users *do not* want to configure policies

RR: sure - program your wine cellar; need security to authorize access

Marie-José Montpetit: What, exactly, does ICN bring to my 80-year old mother?

RR: "Why ICN is more apt" slide should address this question.

MM: Are these benefits to the home network operator or the home owner?

RR: E.g., "Configuration-less" means home owner doesn't need to know about multiple routers in the homenet.

NS: Are we talking about application designers or network designers?

RR: Moving Nest to ICN would be transparent.

DK: some of these features may give a broad benefit across all networks; what is specific to home networking

RR: We're just setting the agenda to motivate deeper discussion

DT: "ICN can elevate home networking to a new level of experience" - "configuration-less" is *really dangerous"; ICN requires a high degree of configuration; what are proper examples of "new levels of experience"

MM: experience for whom

DT: "home networking" is not an ideal application because the gains are not obvious

DO: home networks are interesting because they are easy:

- no hard scaling requirement
- one administrator
- robustness and reliability bar is low
- limited range of type of devices
- so...make it work here and then we might be able to make it work in harder setting

RR: I agree

ED: Are the functions in the APIs written down?

RR: No. Some ideas in the draft.

DK: last thing homeowner wants to do is change batteries every 2 weeks

RR: energy efficiency is in scope but not in prototype

BO: Everyone agrees policies are needed; DO says users don't want to configure them. Are policies built into the system

RR: Homeowner wants to decide about new devices

BO: No, wants it to "just work"

DT: Nail down the specific requirements and characteristics; what are the bits and pieces of ICN that make homenets easier to run? homenet is multi-stakeholder for policies

RR: Asks for helps in making those changes to draft

Consumer Driven Information Freshness Approach for Content Centric Networking (Daniel Corujo)

(agenda bashed; see below)

Web traffic and CCN/NDN (Mark Stapp/Dave? Oran) (30 min)

Questions/motivation:

1. Naive faith ICN-style networking will make web applications "better": information-fetching applications
2. Lack of traffic models

Sufficiently rigorous first intuition about traffic and impact on ICN

What are the traffic patterns? What actually comes back in response to synthesized queries?

Jan Seedorf: We ran similar tests; results differ a lot depending on where the tests are run

DO: Yeah, we know. Remembred we're counting returned objects and sizes, not performance [missing notes; network failure]

...see slides for several interesting observations

Multiple CNAMEs could be good news or bad news:

bad: have to recreate CNAMEs in ICN

good: "brilliant routing" might eliminate need for CNAMEs

Queries are all gzipped to get down to fundamental amount of information in the query; compresses <domain>/a/a/a/a/a/a/a/a/a/a

Message example slide: shows distribution of total bytes in Interest messages for fetching large objects; query includes whole RESTful state many times for a large (multi-segment) object

Kostas Pentikousis: Might generate a whole new mindset for page development; current structure is a result of what we have now; might not want to extrapolate ICN web from today's web

DO: agrees 100%

KP: You will need 10 years; JavaScript needed 10 years to be taken seriously

DO: You're right; let's not get into that political discussion. One conclusion: web might not be place for market insertion

NDN Advantages slide

Didn't capture certificate fetches for SSL; no intuition about what effect key fetch might have

2 things in progress:

1. what Jan pointed out: issuing same request from multiple places and multiple times to measure spacial and temporal repeatability
2. issue same logical request from multiple clients; how much spacial caching will be available

DO: Is this useful?

DK: work does a good job of measuring today; shows how ICN maps to today's web architecture

DO (aside): could rerun scripts on other ICN architectures

DK: try to measure what client really needs in an Interest

DO: all the cookies get sent; can ICN figure out which cookies really need to be sent

DO: this work assumes "don't re-engineer the web"; will ICN "save our bacon"

Marc Mosko: a lot of sites going to HTTPS; trying to give same privacy wont allow caching

DO: all streaming video depnd on HTTPS only for control; video streaming uses broadcast encryption keying, etc.

Steven Farrell: HTTPS WG did similar work comparing SPDY to HTTPS; ask Mark Nottingham

KP: can researchers on this work review (some other work)

(from offline input) Steven Farrell - here is the HTTPS work:

https://github.com/http2/http_samples

Consumer Driven Information Freshness Approach for Content Centric Networking (Daniel Corujo)

IoT can provide requirements on content providers and consumers
resource restricted devices as content generators

How does ICN "information freshness" impact applciation to IoT

information freshness - difference between time of generation and time of access for some content

simulated via NDNsim

Big Idea: Consumers assist producer in determining "freshness value"

DO: immutable objects?

DC: immutable names, mutable objects

NS: prefix match naming?

DC: complete match

RD: Why not timestamp data?

DC: Avoid need for knowing timestamp syntax.

DO: Why not "Don't give me information more than n seconds old"?

DC: Agreement gives better results

DO: But 1 out of 100,000,000 consumers might break caching for all other consumers

MarcM: May not have time on all sensors; "freshness" doesn't behave like freshness in HTTP, it moves through the network with the object, gets restarted with every move to a new cache

DO: If you don't have wall clock time you don't have wall clock time. With no wall clock, you have "how many revisions back"

Notetaker: Nacho Solis

Evaluation Methodology - Kostas Pentikousis

History of document ... document grew large, decided to split it into 2.

Not a lot of recent activity. Only a few of the original proposers left.

Not a lot of interest from IPPM

Does anybody want to work on it?.... <silence>

What are people doing for evaluations?

Lorenzo Saino is doing work using a previous methodology. Kostas wants to talk about this methodology.

Spiros: What is the way forward?

Kostas: More discussion, maybe on mailing list

BO: Intention of document is having a common way to evaluate ICN.

Scalable Container-based Routing Scheme for ICN - by Joo-Chul Kevin Lee

Nacho: Are you using IP? Where?

JCKL: Our scheme is an overlay.

Nacho: Scalable on routing table?

JCKL: Yes

Nacho: But the name resolution service has to scale to all objects:

JCKL: Yes

DirkKutcher: How does it relate to similar architectures? (NetInf routing scheme... etc)

RaviRavindran: Yes, presented in Orlando

DaveOran: Some similar paper submitted to ICNP 2013

Nacho: Hawei's proposal on the ICNRG page.

DirkKutcher: Good if you can state what is the delta to previous presentations.

JCKL: I wanted to check if this approach makes sense

<multiple>: It's hard to tell

Nacho: How does the evaluation methodology apply to this?

Kostas: This is what I said, we're doing evaluation innovation.

Baseline Scenarios - by Kostas

If you have a scenario in your mind but don't have a peer-reviewed paper, look somewhere else.

Scenarios document complete in terms of coverage.

Current version addresses concerns from reviewers.

Georgios K: I wanted to contribute but didn't have a chance.

Kostas: There will be more opportunities to participate.

DirkKutcher: Talked to DirkTrossen?

Kostas: Yes, he wants to send a scenario. Depending on scenario it can be on this document.

Challenges Document - by Dirk Kutcher

One open issue: Link and impact to IETF technologies, section still empty.

Nacho: who proposed that section?

Dirk: I think me. Document would be valuable without it.

Georgios K: Catalogs, best-effort, etc. Do you think there is a challenge there?

Kostas: The IETF section might be a draft by itself.

wrapup

Attendance:

#	Name	Email	Comment
1	Börje Ohlman	Borje.Ohlman at ericsson.com	
2	Dirk Kutscher	the usual one	
3	George Xylomenos	xgeorge at aueb.gr	
4	Bartosz Balazinski	bartosz.balazinski at interdigital.com	
5	Thomas Schmidt	t.schmidt at ieee.org	
6	Dave Oran	daveoran at orandom.net	
7	Spiros Spirou	spis at intracom-telecom.com	
8	Jussi Kangasharju	jussi.kangasharju at helsinki.fi	
9	Daniel Corujo	dcorujo at av.it.pt	
10	Cedric Westphal	cwestphal at gmail.com	
11	Ravi Ravindran	ravi.ravindran at huawei.com	
12	Jan Sedorf	lastname at neclab dot eu	
13	Anders Lindgren	andersl at sics dot se	
14	Fehmi Ben Abdesslem	fehmi at sics dot se	
15	Gareth Tyson	gareth.tyson at qmul ac uk	
16	Matthias Wählisch	waehlich at ieee dot org	
17	Damien Saucez	damien.saucez at inria.fr	
18	Nacho (Ignacio) Solis	Ignacio.Solis at parc.com	
19	Marc Mosko	Marc.Mosko at parc.com	
20	Glenn Edens	Glenn.Edens at parc.com	
21	Marie-José Montpetit	marie at mjmontpetit.com	
22	Diego Perino	diego.perino at alcatel-lucent.com	
23	Eiichi Muramoto	muramoto.eiichi at jp.panasonic.com	
24	Georgios Karagiannis	karagian at cs.utwente.nl	
25	Dirk Trossen	Dirk.Trossen at InterDigital?.com	
26	Pasquale Cataldi	pasquale.cataldi at nominet.org.uk	
27	Alfredo Grieco	alfredo.grieco at poliba.it	
28	Gennaro Boggia	gennaro.boggia at poliba.it	
29	Joo-Chul Kevin Lee	rune at etri.re.kr	
30	Wan-Seon Lim	wslim at etri.re.kr	
31	Elwyn Davies	davieseb at scss.tcd.ie	
32	Hongke Zhang	hkzhang at bjtu.edu.cn	
33	Shuai Gao	shgao at bjtu.edu.cn	
34	Fei Song	fsong at bjtu.edu.cn	
35	Wassim Haddad	wassim.haddad at ericsson.com	

36	Ralph Droms	rdroms.ietf at gmail.com	
37	Kostas Pentikousis	k.pentikousis at eict.de	
38	Dessalegn Mequanint	dessalegn_m at yahoo.com	
39	Konstantinos Katsaros	k.katsaros at ucl.ac.uk	