

IAB E-Impact Workshop Session 3: Improvements
2022-12-09
Webex Chat Log

from Hosein Badran ISOC to Everyone: 6:10 AM
Good morning/afternoon from Ottawa, Canada

from Noa Zilberman to Everyone: 6:13 AM
Would the slides from all the talks be uploaded or shared?

from Colin Perkins to Everyone: 6:13 AM
Slides are being uploaded to <https://datatracker.ietf.org/group/eimpactws/meetings/> as we get them

from Noa Zilberman to Everyone: 6:15 AM
Great

from Pascal Thubert to Everyone: 6:15 AM
+1 on protocol designs; add to the list use of broadcast in reactive protocols

from Russ White to Everyone: 6:17 AM
one thing to think about here is that optimizing for multiple metrics is np(complete) ... so somehow any metrics used must be "merged" into a single metric for any sort of shortest/loop free path mechanism to work correctly

from Operations and Management Area to Everyone: 6:19 AM
This sounds like it could be documented perhaps as experimental YANG data models. This could be done both as device models and also augmentations to the network wide YANG modules. For the device ones, it would probably also be helpful if they could also be defined in a somewhat unified way with the OpenConfig YANG models.

from Louise Krug to Everyone: 6:20 AM
with network as current I can't imagine that routing based on metrics will help - the stuff draws power once it is there

from Martin Flack to Everyone: 6:20 AM
Discarding consumer and industrial equipment is huge... really an area for government to help by incentivizing or requiring mandatory return of used equipment to manufacturer or their agent in order to most responsibly reuse / recycle / refurbish. Need to consider end of life up front.

from Qin Wu to Everyone: 6:22 AM
The hardware yang model has already had some sensor related attributes which can be reused or completely refactored.

from Russ White to Everyone: 6:22 AM
there is, I think, little difference in power usage when forwarding versus "at rest" for a router ... though I've not seen research on this for a long while

from Russ White to Everyone: 6:23 AM
there are other techniques, though ... I have some of these in our presentation

from Louise Krug to Everyone: 6:24 AM
Russ - we don't see much variability on the latest devices in our network

from Jari Arkko to Everyone: 6:24 AM
Good point about metrics being actionable at the IETF

from Operations and Management Area to Everyone: 6:25 AM
@Loiuse, if you can move all traffic off an interfaces, or all interfaces on a linecards then you could potentially power down those interfaces and linecards, which would save power (until an increase in traffic requires them to be enabled again). Whether it would save any meaningful amount of power compared to other places that power is being used in the network is more questionable.

from Colin Perkins to Everyone: 6:26 AM
Sorry - who is "Operations and Management Area"?

from Operations and Management Area to Everyone: 6:26 AM
Rob Wilton - sorry, I didn't know that Webex had picked up my title from somewhere

from Colin Perkins to Everyone: 6:26 AM
Of course - thanks! WebEx does that...

from Louise Krug to Everyone: 6:27 AM
the question is if there are enough linecards that you can do that with - with modern routers and large link data rates you may have few links. You typically need a couple of links up for resilience - rapid response to an event like the manchester bombing - bringing a link back up from deactivated is quite slow - the line cards are quite power hungry though

from Vesna Manojlovic to Everyone: 6:27 AM
I had to look up OPSAWG, so here's the link: <https://datatracker.ietf.org/wg/opsawg/about/>

from Chris Adams to Everyone: 6:27 AM
Is Ramp up / ramp down much of a common a common metric designers are actively optimising for? That was one we discussed earlier this week, and I might have missed it in the presentation if it was discussed

from Louise Krug to Everyone: 6:28 AM
and then we need to understand the carbon impact of fibre (say you have more smaller links) vs the embodied energy in the routers- so many unknowns

from Daniel Schien to Everyone: 6:28 AM
What knowledge is there about powering down networking devices. Anecdotally, routers or line cards are not powered off for fear they might not come back online without manual intervention. Is there any data on this?

from Chris Adams to Everyone: 6:28 AM
In other sectors where we see folks needing to accomodate a high peak of demand, when the average is usually lower, I think there might be some less used approaches that might be helpful.

from Daniel Schien to Everyone: 6:29 AM
@Chris - what examples come to your mind that are outside of the electricity sector?

from Louise Krug to Everyone: 6:29 AM
Data no - they just don't want to go there - although we have sponsored some research that is looking at different router architecture - so if you lose the odd link it wouldn't matter so much

from Louise Krug to Everyone: 6:30 AM
That was to Dan sorry

from Daniel Schien to Everyone: 6:31 AM
Does latency really stack up as a driver along throughput and availability?
Yes, it's important, but are there substantial penalties of electricity
consumption incurred to provide latency levels?

from Russ White to Everyone: 6:32 AM
yes, there can be

from Daniel Schien to Everyone: 6:32 AM
@Russ - example?

from Russ White to Everyone: 6:32 AM
I'll give one in my presentation ... :-)

from Daniel Schien to Everyone: 6:32 AM
:)

from Chris Adams to Everyone: 6:32 AM
@dan - TBH I was thinking of the electricity sector specifically. there you
have to deal with over provisioning power, because it needs to be balanced.
So, we have peaker plants, and batteries that are only on a few hours a
year, just to meet the peaks

from Jari Arkko to Everyone: 6:33 AM
@daniel re: powering down network devices. There's a lot of talk about this,
but I think it would be helpful if we saw that not as on/off but more of a
continuum... can i shut down a core, slow down CPU, shut down one network
card in a multi-link bundle, do wake-on-packet, shut down cards, shut down
for a period, etc. I find the smaller grade shutdowns more useful personally

from Louise Krug to Everyone: 6:33 AM
interesting question Dan - latency drives the need for fewer electronic hops
-> lower energy but also fast processing -> higher energy

from Toerless Eckert to Everyone: 6:34 AM
@jari: ideally we could "shut down" link/interfaces for data but keep them
up/running at low power just for management traffic. pure optical equipment
can do this. Alas, router can not

from Louise Krug to Everyone: 6:35 AM
what are the standards used in data centres for reporting energy consumption
etc? is it standardised?

from Michael Welzl to Everyone: 6:35 AM
Faster processing may not need higher energy: low latency transfer can mean
finishing faster, which can mean longer sleep times. This is in my paper.

from Daniel Schien to Everyone: 6:36 AM
@Chris - yes, the electricity sector is useful as a model. That is why we
point out the TNUoS charges when we think about representing increases to
peak use in the electricity intensity metric. What we don't spend much time
on in the paper: In contrast to our smooth transformation function, the TNUoS
chargers exclusively penalise peak use. The equivalent would be to let all
data use outside of peak go for free and burden all electricity use to just
peak use (for the purpose of carbon accounting).

from Jari Arkko to Everyone: 6:36 AM

@louise adjustable link speed? maybe. But I think CPUs can also shut down or at least significantly slow power usage, and then wake upon memory write from network card. (Eve?)

from Vesna Manojlovic to Everyone: 6:37 AM

@Luise : Data Centers energy consumption is covered by, for example, SDIA: <https://knowledge.sdialliance.org>

from Louise Krug to Everyone: 6:38 AM

Adjusting rates – its all about how quickly we can get the rates back up – which I suspect needs both ends to agree to – which adds delay?. WOULD be good to at least have the ability to cost it so that design choices can be more informed than they are today!

from Franco Davoli to Everyone: 6:38 AM

@Jari: I very much agree with your comment. All this goes under the category of "dynamic adaptation", which can be used to trade-off energy/performance under different traffic conditions and multiple time scales.

from Toerless Eckert to Everyone: 6:39 AM

@Michael: the energy for transceiver to send/receive 100Gbps is a lot more than to send/receive e.g.: 10 Mbps over the same distance. Just a matter of clock frequency and DSP speed needed. Thats why optical equipment has such low bitrate management channels.

from Romain Jacob to Everyone: 6:39 AM

@jari, toerless and co: We are currently performing measurements on gradual powering/connecting on/off parts of routers and switches. We are just getting started, but that's the idea

from Louise Krug to Everyone: 6:40 AM

would be good to have similar tools/interfaces to data centres as there is such a blur with NFV etc

from Wim Vanderbauwhede to Everyone: 6:40 AM

@Jari yes, CPUs have dynamic voltage and frequency scaling

from Colin Perkins to Everyone: 6:41 AM

The trade-off between hyper-optimised to save power for a particular use case, and general-purpose, flexible, less efficient but hence perhaps longer life, is an interesting one.

from Michael Welzl to Everyone: 6:41 AM

@Toerless: you have to factor in the time – 100 Gbps for a shorter time. The example in my paper involved WiFi sleep times. So, this depends on equipment being able to sleep.

from Toerless Eckert to Everyone: 6:42 AM

@Michael: Look at how ≥ 100 Gbps works. you can not short-term enable/disable them due to keeping track of the frequency response of the channel, aka: DSP state.

from Michael Welzl to Everyone: 6:43 AM

@Toerless: I believe you. That's wired high-speed stuff, just a different case from the one I meant.

from Carsten Bormann to Everyone: 6:44 AM

(40 years ago starting a car engine was a big deal. Now this is done automatically all the time. You "just" have to design things to sleep.)

from Toerless Eckert to Everyone: 6:44 AM
@Michael: also, i am coming from the RFC8994 idea where you'd have a small BMC style CPU connected to the management channel on interfaces so you can reliably control witching on/off the rest of the system without killing yourself (linecards, forwarding plane engines, etc. pp.)

from Hosein Badran ISOC to Everyone: 6:45 AM
Q: When eliminating redundant paths, could this adversely affect network resiliency?

from Rob Wilton to Everyone: 6:45 AM
Yes

from Toerless Eckert to Everyone: 6:46 AM
Without doing the separation between data plane and control/management Yes.. (IMHO).

from Rob Wilton to Everyone: 6:46 AM
But if you have 3 paths then possibly could go down to 2 and still keep some level of resiliency

from Eve Schooler to Everyone: 6:46 AM
@Jari @Wim Techniques for sleep states and dynamics apply to CPUs, so could they apply to other components of the system architecture

from Alex Clemm to Everyone: 6:47 AM
The question is how to best mitigate those goals – constraint-based optimization – how to minimize footprint by maintaining just the level of resiliency that is required (and being aware of additional resiliency cost)

from Toerless Eckert to Everyone: 6:47 AM
In data-centers you always have an out-of-band network you can keep up running to manage the rest of the network and switch on/off network components/paths there... Only in WAN we have given up on that separation.

from Pernilla Bergmark Ericsson to Everyone: 6:48 AM
@Dan and others – not sure what you refer to as networking devices but for cellular networks this is something that has become in focus and improved overall energy performance since 5G. There is also systems out whihc demonstrate the potential of such as micro-sleep. Important as systems has traditionally been optimized for top performance not for typical performance. Insights developed during the H2020 project Earth have been instrumental in improving 3GPP standards.

from Colin Perkins to Everyone: 6:49 AM
A reminder to presenters to send your slides to Jari and I, for the website

from Carsten Bormann to Everyone: 6:50 AM
(Slides are gone?)

from Alex Clemm to Everyone: 6:51 AM
I can see slides here: <https://datatracker.ietf.org/meeting/interim-2022-eimpactws-03/session/eimpactws>

from Chris Adams to Everyone: 6:52 AM
so one question that came up from Ross's presentation is the idea of nodes in a network being able to communicate how much longer they are able to sustain a given amount of throughput or latency. Do any protocols includes this and allow for this?

from Louise Krug to Everyone: 6:52 AM
experience on mobile – reluctance to use the sleep modes as it affects
rootmetrics performance assessment of the network

from Louise Krug to Everyone: 6:53 AM
root-metrics measures things like speed on all the UK mobile networks and
publishes the results

from Wim Vanderbauwhede to Everyone: 6:54 AM
@Eve, there is research e.g. on NetFPGA Gb routers with frequency scaling

from Russ White to Everyone: 6:55 AM
we tried doing this in EIGRP many years ago, and discovered you can end up in
a situation where changing metrics drives utilization, which then drives
metrics, which then drives utilization, which ... it was not happy

from Russ White to Everyone: 6:56 AM
BABEL solves this with its energy metrics by holding convergence to defined
periods -- by slowing down convergence through hysteresis (sp?) ... this is a
good technique, but the tradeoff is slower convergence ...

from Russ White to Everyone: 6:56 AM
it might be fine to converge more slowly sometimes, and it might be other
times ... we don't have a good handle on this today, I think ... we just
assume "faster is better"

from Russ White to Everyone: 6:57 AM
we could also make the routing protocols themselves MUCH more efficient in
terms of the number of packets sent to converge, etc. ... I don't know if
this would have enough of an impact on energy usage to matter, though ...
again, we don't really have measurements, so we don't really know

from Louise Krug to Everyone: 6:58 AM
you can get a feel for the amount of energy for routing by looking at the
power of the routing engine within a whole fully configured router –
from Toerless Eckert to Everyone: 6:58 AM
It's not as if closed loop control systems are new...

from Toerless Eckert to Everyone: 6:59 AM
we just never put a lot more effort into the matter in routing at least not
in the industry

from Toerless Eckert to Everyone: 6:59 AM
every room thermostat has a more advanced control loop as routing
protocols ;-)

from Fieke to Everyone: 7:00 AM
Unfortunately I have to go, again great session and thanks for all the input.
See you all next week. Happy weekend.

from Dom Robinson Greening of Streaming to Everyone: 7:00 AM
The key challenge to multicast is economic: There are CDNs and large
operators making a lot of money from unicast.

from Russ White to Everyone: 7:01 AM
it's also hard to do multicast properly, even from the perspective of a
cdn ...

from Jari Arkko to Everyone: 7:01 AM
Thanks Fieke and have a nice weekend you as well!

from Louise Krug to Everyone: 7:01 AM
i'm all for multicast

from Russ White to Everyone: 7:01 AM
multicast, done right, is brilliant, though ... its just not easy to do :-)

from Suresh Krishnan to Everyone: 7:02 AM
Multicast is pretty expensive in wireless

from Dom Robinson Greening of Streaming to Everyone: 7:02 AM
i ran a multicast CDN (pioneered M-BGP) 2001-2010 - we got to scale but then grew our vod / unicast business, and at that point multicast efficiency crippled our unicast and associated traffic commits and so everything unicast (the dominant traffic) became more expensive

from Russ White to Everyone: 7:02 AM
for instance ... suppose I use anycast to reach the closest instance of a service ... what's the tradeoff of carrying the traffic farther versus holding the data closer? do we know how to measure this, or characterize it? I don't know

from Colin Perkins to Everyone: 7:03 AM
What's the trade-off between network multicast and application-layer multicast overlays?

from Toerless Eckert to Everyone: 7:03 AM
Suresh: Cell Broadcast uses radio "multicast" ;-)

from Dom Robinson Greening of Streaming to Everyone: 7:03 AM
@colin - 'QoS'

from Martin Flack to Everyone: 7:03 AM
More multicast would be useful for traffic where everyone wants to download at the same time (think World Cup), or a download can be scheduled and subscribed (think Windows update). Unpredicted, on-demand, user-based access is a driver of unicast usage. In defense of CDN's (I work for one) they also do a lot of fan-out in heirarchical trees that save lots of byte transit.

from Russ White to Everyone: 7:03 AM
application layer overlay generally means you're duplicating traffic closer to the source versus closer to the destination ... to get efficiency, you want to replicate as close as possible to the destination

from Toerless Eckert to Everyone: 7:04 AM
Colin: Anything doing replication on non-router devices such as PCs requires the energy consumption of that additional equipment. Energy wise, multicast is a no-brainer. Just all those oter ecosystems repeated here in the chat stand in the way

from Jari Arkko to Everyone: 7:04 AM
It is a great observation that CDNs might help with the inter-domain issues we previously had. But isn't there also a different impact? If we have a CDN node in a DSL ISP for instance, we have limited savings potential, because the last mile links still need separate packets

from Suresh Krishnan to Everyone: 7:04 AM
@Toerless: I was thinking more of 802.11 but I should have been more specific :)

from Dom Robinson Greening of Streaming to Everyone: 7:05 AM
Why would CDNs reduce their revenue from unicast?

from Toerless Eckert to Everyone: 7:05 AM
@Suresh: always depends on how many receivers you want to reach. 802.11 could do automatic switchover from replicated unicast to multicast if it wanted. We just never had enough use-cases for that optimization.

from Colin Perkins to Everyone: 7:06 AM
@toerless could those non-router devices cache and save transmission that way? Multicast is great if you have simultaneous receivers, it's not clear it helps for time-shifted

from Russ White to Everyone: 7:06 AM
it's chicken-and-egg, isn't it? the more multicast business there is, given multicast replication at the edge rather than source, cdn's could make money off multicast ... I don't know if these are "solved problems" though

from Toerless Eckert to Everyone: 7:07 AM
@colin: "Deigital FOuntain", ca. 2002 - all the schemes to do on-demand video via multicast segments. I mentioned on the mailing lists that it would be fun to do an energy comparison of bandwidth used for that vs. energy for PC+storage

from Russ White to Everyone: 7:07 AM
wireless efficiency often depends on hidden speakers, I'd think ... ?? I've not spent time thinking about these problems in a long'ish while, though

from Dom Robinson Greening of Streaming to Everyone: 7:08 AM
in our model we did get the access networks running IGMPv2 and so we really did a very pure multicast delivery - and we charged for bytes received rather than 'sent' so our revenue was similar to unicast, but (as i mentioned above) the problem was our wholesale buying which adversely affected the economic of delivering of the dominant VoD unicast.

from Jari Arkko to Everyone: 7:08 AM
I don't think there's a business reason why CDNs could sell multicast-based solutions. But isn't there also a dependency on the ISP supporting BIER etc?

from Russ White to Everyone: 7:09 AM
you can do something just in the cdn overlay ... it just won't be as efficient

from Toerless Eckert to Everyone: 7:09 AM
@Jari: Look at the AMT work in mboned driven by folks from Akamai

from Colin Perkins to Everyone: 7:10 AM
@toerless the digital fountain work was interesting, certainly. As you say, it would be interesting to compare energy use of the different approaches

from Martin Flack to Everyone: 7:10 AM
(1) "Gigabytes Delivered" is a key CDN billing metric to customers. If it could all be done with less power and greener, i.e. multicast, Gigabytes Delivered to each listening client would still be the same. It's a question of making it work for all business requirements. (2) At a high level, contracts are based on outcomes, not strictly technical metrics.

from Toerless Eckert to Everyone: 7:10 AM
@Jari: The BIER story is orthogonal, that's another enhancement.

from Toerless Eckert to Everyone: 7:10 AM

The story of why CDNs would like to see multicast is for peak shaving:

from Toerless Eckert to Everyone: 7:10 AM
average CDN utilization now is 30% of peak.

from Toerless Eckert to Everyone: 7:11 AM
aka: peak of live events such as FIFA world cup are what you need to design your CDN against.

from Louise Krug to Everyone: 7:12 AM
ISPs are def motivated to get more multicasting to work - it saves energy but also reduces costs. I think some of the problems is in the vast array of end devices?

from Dom Robinson Greening of Streaming to Everyone: 7:13 AM
@Toerless - Problem is the background VoD video traffic still dwarfs the live peaks

from Brendan Moran to Everyone: 7:13 AM
@Louise; I wonder if that can be solved by the ISP-provided routers/modems?

from Louise Krug to Everyone: 7:13 AM
live can be multicast

from Toerless Eckert to Everyone: 7:13 AM
Multicast is currently decommissioned from more and more services in ISP. IPTV being the latest one. Nobody wants to afford multicast troubleshooting in home networks.

from Louise Krug to Everyone: 7:13 AM
we are looking into that Brendon I think (slightly out my area)

from Brendan Moran to Everyone: 7:14 AM
I suspect DRM is the big deal here.

from Louise Krug to Everyone: 7:14 AM
yes the rights management is a good one - games releases

from Romain Jacob to Everyone: 7:14 AM
What's DRM?

from Martin Flack to Everyone: 7:14 AM
digital rights mgmt

from Louise Krug to Everyone: 7:14 AM
digital rights management

from Toerless Eckert to Everyone: 7:14 AM
DRM AFAIK is not an issue. Satellite networks had the DRM for content worked out way before IP multicast came into the picture

from Russ White to Everyone: 7:14 AM
and advertising ... you have to have a backchannel to "see" advertising has been placed correctly, check on viewers watching, etc.

from Carsten Bormann to Everyone: 7:15 AM
Multicast works best if the CDN actually extends down into the end-users device...

from Dom Robinson Greening of Streaming to Everyone: 7:15 AM

good points on mic Martin

from Toerless Eckert to Everyone: 7:16 AM
@russ targeted advertisement with multicast was also done a decade ago by some company we worked for ;-) check out patents against it ;-)

from Louise Krug to Everyone: 7:16 AM
security encodings is nasty

from Louise Krug to Everyone: 7:19 AM
the difference between 2 research home hub devices many years ago was 8W or 24W , and the main reason apparently use of IPSEC -

from Toerless Eckert to Everyone: 7:20 AM
AES support seems still not ubiquitous on ARM, whereas it is on x86 since 6? year.

from Jari Arkko to Everyone: 7:20 AM
Agree that points from Martin were really good. Always great to have these meetings were you have the actual players and get a definite opinion.

from Alex Clemm to Everyone: 7:21 AM
How about the energy impact of REST (vs non-REST) and the need to continuously transfer all that state? (tradeoff with need to maintain state at the server, of course)

from Russ White to Everyone: 7:22 AM
have to drop for another meeting ... have a good weekend

from Michael Welzl to Everyone: 7:22 AM
@Alex: a very interesting question, IMO ... the REST overhead could affect MANY things

from Louise Krug to Everyone: 7:22 AM
I did some back of envelope sums for circuit and packet switching (state in router vs state carried) and the circuit could be factor 10 more efficient for longer hold traffic - so the REST question could be interesting

from Jari Arkko to Everyone: 7:23 AM
Actually, I always have to install tools to look at JSON ... pretty-print at least

from Toerless Eckert to Everyone: 7:26 AM
@Brendan I would place what you say the other way around: We need to build the whole diagnostic framework out a lot more for constrained and secured protocols before we can sunset all the non-constrained/non-secure protocols easily.

from Carsten Bormann to Everyone: 7:26 AM
So it would be interesting to hear where REST contributes to message size in actual interchanges

from Toerless Eckert to Everyone: 7:29 AM
@carsten thanks. i think that was the answer to the REST question

from Toerless Eckert to Everyone: 7:31 AM
@brendan I've not seen IETF to define the necessary wireshark definitions for binary protocols ;-)

from Jari Arkko to Everyone: 7:32 AM
This is kind of related to the REST question perhaps, but I was curious if we

had any data on where in the Internet tech & apps we might have biggest issues re: bloated formats. Do we know? And can we quantify how much we could save, not as a percentage of the specific protocol exchange but on a more global scale?

from Brendan Moran to Everyone: 7:32 AM
@Toerless, that's fair, but that's also not the whole story

from Chris Adams to Everyone: 7:32 AM
if every place that could adopt CBOR instead of JSON, did adopt CBOR instead of JSON – is there any rough indicator of what impact that might represent? does any data exist that could give at least a back of a napkin figure of what kind of savings in absolute terms?

from Chris Adams to Everyone: 7:33 AM
You see the figures from cisco's annual reports for machine to machine communication – would that give an idea?

from Brendan Moran to Everyone: 7:33 AM
@Chris, probably difficult to quantify. The 33% number from SenML might be helpful

from Toerless Eckert to Everyone: 7:33 AM
@chris how do you pick between CoAP and CBOR ?

from Chris Adams to Everyone: 7:34 AM
CoAP?

from Carsten Bormann to Everyone: 7:34 AM
Toerless: You pick both, of course :-) (Well, CoAP has a certain domain of applicability)

from Toerless Eckert to Everyone: 7:34 AM
rfc7252

from Rob Wilton to Everyone: 7:34 AM
@Chris, it goes back to the key question, if most traffic is video then moving to a binary encoding would only make a small difference. But I still prefer binary encodings generally – why waste the space.

from Vesna Manojlovic to Everyone: 7:35 AM
My "social engineering" suggestion is aligned with this (from Carsten) -- similar to "flying shame" concept for making people to D0 take a train instead of flyign -- can we come up with some kind of social pressure for people / engineers being pressured AGAINST using the wasteful encoding, wasteful protocols, wasteful equipemnt... and TOWARDS the more sustainable option.

from Toerless Eckert to Everyone: 7:35 AM
@carsten with anima i am just seeing a very small slice of the IoT problem of having too many alternatives at all protocol layers... lot of "job security" for developers to do the same work in IoT multiple times *sigh*

from Louise Krug to Everyone: 7:35 AM
traffic peaks today mostly video & game downloads

from Carsten Bormann to Everyone: 7:36 AM
CoAP is a transfer layer, like HTTP. CBOR is a data representation format, like JSON or XML (or, implicitly, XDR or ASN.1)

from Toerless Eckert to Everyone: 7:36 AM

@vesna This is already happening today. Vendors who have constrained and unconstrained products recognize they can re-use the constrained, but not the unconstrained protocols.

from Michael Welzl to Everyone: 7:36 AM
Argh... we're back at email, yet even in a reasonable context!

from Colin Perkins to Everyone: 7:37 AM
We need stats for the traffic mix in different networks, to see where this will save significant bandwidth.

from Carsten Bormann to Everyone: 7:37 AM
Answer: To an appropriate level

from Carsten Bormann to Everyone: 7:38 AM
In YANG, the answer is YANG SIDs.

from Romain Jacob to Everyone: 7:38 AM
@Colin: this recent survey provide some high-level numbers: <https://dl.acm.org/doi/pdf/10.1145/3485447.3512283>

from Romain Jacob to Everyone: 7:39 AM
Not broken up by network types, but's that's already something

from Toerless Eckert to Everyone: 7:40 AM
@carsten there are probably more than one CBOR based protocols in IETF that could have reused CoAP instead of using CBOR. Likewise, CoAP could use CBOR today... price of evolution and broad markets, but its a cost we carry forward unless we converge more in the future

from Colin Perkins to Everyone: 7:41 AM
Sure, that sort of data is helpful. But to see the energy savings, maybe the question is whether switching formats saves enough to allow certain links to be shutdown temporarily? Very much depends on the mix in particular networks.

from Louise Krug to Everyone: 7:42 AM
you dont just have to close links to save power. if we can slow the rate of growth of peak traffic then network energy consumption has a chance to decline as the efficicny improvemnets that occur over time will more than account for the growth in peak traffic

from Carsten Bormann to Everyone: 7:45 AM
@toerless: CoAP does use CBOR, e.g., in OSCORE.

from Toerless Eckert to Everyone: 7:45 AM
@carsten hmm. but ietf spec of coap predates all the cbor encoding rules, not ?

from Chris Adams to Everyone: 7:45 AM
Thanks Eve! My thinking was knowign this makes is possible to build system level abilities to respond to spikes in demand

from Carsten Bormann to Everyone: 7:46 AM
@toerless: Not by date of publication, but by date of design. CBOR essentially came on the agenda when CoAP was done and we were moving up the stack.

from Toerless Eckert to Everyone: 7:47 AM
@carsten so oscore spec is different from ietf one, using cbor encoding ?

from Rob Wilton to Everyone: 7:47 AM
@Louise, I think that the point that Dom make previously is pertinent here – how can we help consumers make a choice between my audio/video has to be at the highest possible rate vs what is good enough (e.g., will I really see/hear a difference for 1.Mb/s FLAC vs 160 Kbp/s ACC/Opus). Can we make it easier for consumers to understand the energy impacts of their choices? It seems that most folks pay a flat rate for bandwidth and hence don't have much financial incentive to minimize it.

from Louise Krug to Everyone: 7:48 AM
yes that would be good

from Henk Birkholz he/him to Everyone: 7:48 AM
@toerless: the intuitive approach was and remains to use CBOR via CoAP. rfc dates might not reflect that well

from Carsten Bormann to Everyone: 7:48 AM
@toerless: RFC 8613 abstract: This document defines Object Security for Constrained RESTful Environments (OSCORE), a method for application-layer protection of the Constrained Application Protocol (CoAP), using CBOR Object Signing and Encryption (COSE). OSCORE provides end-to-end protection

from Jan Lindblad to Everyone: 7:48 AM
Returning to earlier presentations; One of the really power hungry parts of the ICT industry is how we do device& network monitoring. Operators are often setting up mechanisms that are polling huge amounts of state data every few minutes. The global power usage for this somewhat crude approach corresponds to multiple nuclear powerplants full output 24/7/365

from Suresh Krishnan to Everyone: 7:49 AM
<https://datatracker.ietf.org/wg/tvr/about/>

from Suresh Krishnan to Everyone: 7:49 AM
is what Toerless was talking about

from Toerless Eckert to Everyone: 7:49 AM
@carsten: ack. so future CoAP work including recent ones uses CBOR instead of repeating it's own "one-off" encoding... Fair ?

from Henk Birkholz he/him to Everyone: 7:50 AM
http2 opens new options while coap was also augmented to work well with bigger message sizes

from Carsten Bormann to Everyone: 7:51 AM
@toerless: CoAP's option structure stays intact. What you put into the options and in the payload often is CBOR these days. E.g., RFC 9290...

from Toerless Eckert to Everyone: 7:51 AM
@carsten "soft convergence ;-)"

from Rob Wilton to Everyone: 7:52 AM
The main alternative is streaming telemetry, which is more efficient, but then that may just allow or encourage the data to be pushed off at higher frequency. E.g., the saving may just get sucked up elsewhere.

from Rob Wilton to Everyone: 7:52 AM
But, Jan, I am in complete agreement that we should do better.

from Dom Robinson Greening of Streaming to Everyone: 7:52 AM

nice @Chris

from Rob Wilton to Everyone: 7:55 AM
@Suresh, I like this idea, and this would seem to be something that could be worked on straight quite quickly (without waiting for more data before we do anything).

from Chris Adams to Everyone: 7:55 AM
at my end - i don't understand enough how people pay for capacity - I know you might pay in terms of Xmbts of capacity, but being able to specify how, much ramp up and how fast it happens would really help with knowing how much capacity needs to be kept ready

from Louise Krug to Everyone: 7:56 AM
great fun!

from Suresh Krishnan to Everyone: 7:56 AM
Sounds good Rob. I will see if I can put something together after the workshop to collect some of the things in this space.

from Chris Adams to Everyone: 7:56 AM
i really enjoyed this folks, thanks

from Michael Welzl to Everyone: 7:57 AM
+1, very interesting and inspirational !

from Dom Robinson Greening of Streaming to Everyone: 7:57 AM
excellent week - looking forward to #4

from Suresh Krishnan to Everyone: 7:57 AM
Thanks Jari, Colin and the PC for putting this together! This has been excellent so far.

from Hosein Badran ISOC to Everyone: 7:58 AM
Thank you for great sessions!

from Cedric Westphal to Everyone: 7:59 AM
thank you all!

from Wim Vanderbauwhede to Everyone: 7:59 AM
Thank you everyone!

from Dom Robinson Greening of Streaming to Everyone: 7:59 AM
Have a great weekend!

from Daniel Schien to Everyone: 7:59 AM
thanks all!

from Martin Flack to Everyone: 7:59 AM
thanks!

from Esther Roure Vila to Everyone: 7:59 AM
enjoy!

from Vesna Manojlovic to Everyone: 7:59 AM
Thank you! Bye!

from Henk Birkholz he/him to Everyone: 7:59 AM
thx to all! bye!

from Louis Navarre to Everyone: 7:59 AM
Thank you! Bye!