IETF 118 HotRFC

Brought to you by
Gorry Fairhurst and **Liz Flynn**

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (https://www.ietf.org/contact/ombudsteam/) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- BCP 9 (Internet Standards Process)
- BCP 25 (Working Group processes)
- BCP 25 (Anti-Harassment Procedures)
- BCP 54 (Code of Conduct)
- BCP 78 (Copyright)
- BCP 79 (Patents, Participation)
- https://www.ietf.org/privacy-policy/(Privacy Policy)

Note Really Well

- IETF meetings, virtual meetings, and mailing lists are intended for professional collaboration and networking, as defined in the IETF Guidelines for Conduct (RFC 7154), the IETF Anti-Harassment Policy, and the IETF Anti-Harassment Procedures (RFC 7776). If you have any concerns about observed behavior, please talk to the Ombudsteam, who are available if you need confidentiality to raise concerns confident about harassment or other conduct in the IETF.
- The IETF strives to create and maintain an environment in which people of many different backgrounds and identities are treated with dignity, decency, and respect. Those who participate in the IETF are expected to behave according to professional standards and demonstrate appropriate workplace behavior.
- IETF participants must not engage in harassment while at IETF meetings, virtual meetings, social events, or on mailing lists. Harassment is unwelcome hostile or intimidating behavior—in particular, speech or behavior that is aggressive or intimidates.
- If you believe you have been harassed, notice that someone else is being harassed, or have any other concerns, you are encouraged to raise your concern in confidence with one of the Ombudspersons.

This session is being recorded

IETF 118 Meeting Tips

In-person participants

- Make sure to sign into the session using the Meetecho (usually the “Meetecho lite” client) from the Datatracker agenda
- Use Meetecho to join the mic queue
- *Keep audio and video off if not using the onsite version*

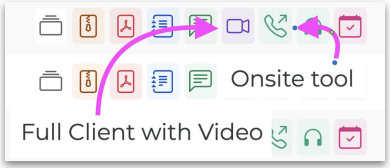

Remote participants

- Make sure your audio and video are off unless you are chairing or presenting during a session
- Use of a headset is strongly recommended

The Ground Rules

- **HotRFC is how you make a Request For Conversation**
  - It's a good way to find IETF people to talk to, for various reasons
- Each person gets four minutes from "Go" to "Please Applaud"
  - At four minutes, we start applauding (see next slide)
  - When you hear applause, please hand the microphone over 😇
- We don't do questions here - each person provides follow-up info
  - (in-person attendees can follow presenters to the bar, of course)
- So you can follow along, we're using the datatracker for all slides
  - Let the conversations begin!

Please Applaud!!! (and the crowd goes wild)

IPv6 Traffic% and Packet Loss Rate – An Update

HotRFC Talk, IETF 118

XiPeng Xiao, Huawei Germany & v6ops co-chair
xipengxiao@huawei.com

What We've Learned after the Previous Talk

Then

- Reported 2 problems in HotRFC Talk 117
  - IPv6 traffic% = IPv6 user% * IPv6 content% * IPv6 connectivity%
    = 41% * 67% * 100% = **27%**
    greatly exceeding reported traffic% below

IETF 117	Traffic %	Date	Source
AMS-IX	5.0%	2023 07	https://stats.ams-ix.net/sflow/ether_type.html
Akamai	16.4%	2022 06	Value derived combining two independent posts

- IPv6 PLR (packet loss rate) much higher than IPv4 PLR

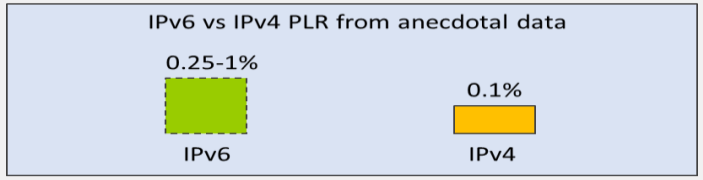

Now

- IXP IPv6 traffic% not including CDN → not representative
- Traffic% available: FB, China-IPv6, Cloudflare
- User% available: Google, APNIC, Akamai
- Stats may not be what are claimed (i.e. is it traffic% or user% or connection%)
- What matters is IPv6-usability% = IPv6 traffic% / (IPv6 user% * IPv6 content%)
  - If IPv6 equally usable as IPv4, then IPv6 traffic% = IPv6 user% * IPv6 content% → IPv6-usability%=100%
  - If IPv6-usability% > 100%, IPv6 more usable than IPv4
  - If IPv6-usability% < 100%, IPv6 less usable than IPv4
  - Anecdotal IPv6-usability calculated from a single company stats for Dual-Stack content (with caveats, please take with grain of salt, contact author for full disclosure)

	World	USA	Canada	Germany	UK	Australia	NZ	India	Indonesia	South Africa	Egypt	Argentina	Brazil
Traffic%	37%	60%	41%	56%	36%	35%	31%	69%	13%	1.70%	4%	18%	48%
User%	41%	48%	37%	73%	44%	29%	20%	71%	15%	1.50%	5%	20%	48%
Content%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
IPv6-usability%	90%	125%	111%	77%	82%	121%	155%	97%	87%	113%	80%	90%	100%

- IPv6 PLR (packet loss rate) ~0.6%, derived from TCP failure rate of 1.27% (involving 2 packets) reported by APNIC, 6x of IPv4 PLR, clearer indicator of connectivity problem

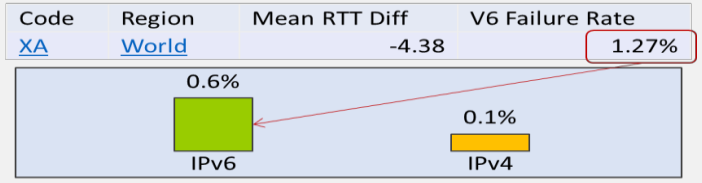

Big thanks to: Weitong Li, Luca Niccolini, Paul Saab, Jordi Palet, Geoff Huston, Erik Nygren, Eduard Vasilenko, Paolo Volpato, Nalini Elkins, Mike Ackermann, Ryo Yanagida, Tim Winters, Gabor Lencse, Ted Lomon, Ole Troan and many others for their input

More on What We Learned, and What Problems Remain

Learned

- IPv6 requirements on residential CPEs dated – v6ops will update CPE requirement RFC 7084
- ISP deployment of IPv6 is mostly in overlay not so much in underlay
- BCP for enterprise IPv6 deployment is needed – please contribute
- Networks in vertical industries are far from using IPv6 – many applications & devices not supporting IPv6
- Maybe more productive to focus on “Converting Residual IPv4” than on IPv6

Problems

- At 0.6%, IPv6 PLR is 6x of IPv4’s
- The reasons for high IPv6 PLR (presented in IETF 117) are still relevant
  - Packet drop with EHs,
  - NCE exhaustion causing packet drop
  - Rate limiting to prevent /64 scanning causing NCE exhaustion
  - Long headers causing congestion/drop at mobile backhaul links
  - Fragmentation-related drops
  - Flash renumbering-related drops
  - Note: Firewall/middleboxes may create PLR asymmetricity between clients/servers
- Does high IPv6 PLR imply some IPv6 issues not yet known?
  - Possibly. Please join Nalini Elkins’ talk at v6 side meeting (Thur Nov. 9, 9:30-11:00)

We will Continuously Improve IPv6 Operations. Please Contribute

- Do you agree: IPv6-usability% = IPv6 traffic% / (IPv6 user% * IPv6 content%) is good for comparing IPv6 with IPv4?
- Provide IPv6 traffic stats from operators & enterprises
- Measure IPv6 PLR in various scenarios
  - Inside enterprise & operator's AS, at IXPs, at content providers
  - Identify root causes of high IPv6 PLR
- Co-author drafts about issues and solutions
  - One theory: is IPv6 PLR mostly from transit points & FW/middle boxes? What can be done to prove/disprove that?
- Help to convert "Residual IPv4"
  - Residual IPv4 users & content – we know where they are but what can be done?
  - Many IPv4 residuals in vertical industries (e.g. railways). Call for people with vertical domain knowledge to contribute

Disclaimer: IPv6 has shorter latency and other benefits over IPv4, but this talk focuses on the issues so as to improve

Please Applaud!!! (and the crowd goes wild)

IETF HotRFC
Sustainability/Energy

IETF 118

mpalmero@cisco.com
emile.stephan@orange.com

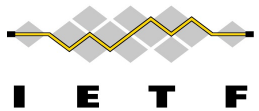

Orange Vélodrome

Smart Power Delivery - Powered by Orange & Cisco

10k wifi guest / event
22 events / year

PoE energy for each Access Point

1 AP		1041 AP		1041 AP 24/7
6.3 Wh	→	7 kWh	→	57 MWh/year

This wifi infrastructure is turned on 24/7.
The energy consumption is extrapolated.

119 MWh/year

24%

Based on the Orange Velodrome case study, the simulation for full seasons 2023 events show that we can save almost 25% of energy consumption.

Solution

A simple device connected to the current infrastructure monitoring all PoE devices. Based on usage, we schedule the shutdown and power up of the energy distribution on PoE Devices.

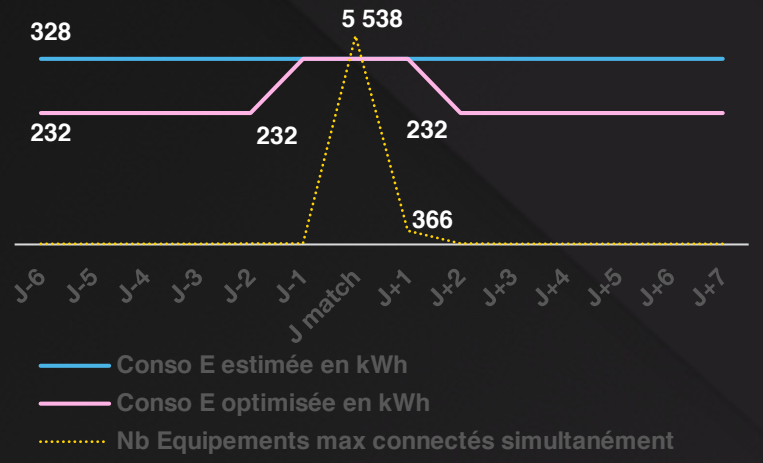

Total Savings

119 MWh/year → 91 MWh/year → 24%

2T CO_2eq

Sustainability Insights Current Challenges

Circular Economy

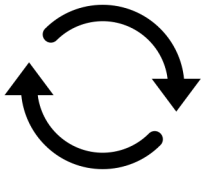

Normalization

Optimization

Accuracy & Granularity

Cross Domain

Sustainability Insights Current Challenges

Circular Economy

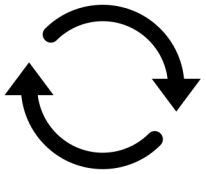

Normalization

Optimization

Accuracy & Granularity

Cross Domain

IETF Draft(s)

- draft-almprs-sustainability-insights
- draft-lindblad-tlm-philatelist
- draft-kll-yang-label-tsdb
- draft-opsawg-poweff
- draft-petra-path-energy-api

More on the
Side Meeting
“Sustainability Insights”
Monday@8:45, Karlin 4

Please Applaud!!! (and the crowd goes wild)

Path-aware Interdomain Architecture Updates

HotRFC - IETF 118
05/11/2023

Nicola Rustignoli (nic@scion.org)
Corine de Kater (cdk@scion.org)
SCION Association

SCION – context

- SCION is an Internet architecture that fills gaps in several areas:
  - *Inter-domain* path-aware networking
  - Trust-enhanced networking (e.g. geofencing)
  - Routing security (path authorization/path validation)
- There is deployment experience:
  - Productive finance, healthcare network in Switzerland
  - More being evaluated (government, energy, DLTs)
- There are multiple implementations of SCION
- SCION can get even better with community feedback

Current drafts

We want to find a space for SCION at the IETF:

- Now documenting the current protocol as implemented in existing deployments
- In the long term use the initial specification for future IETF work to evolve the protocol

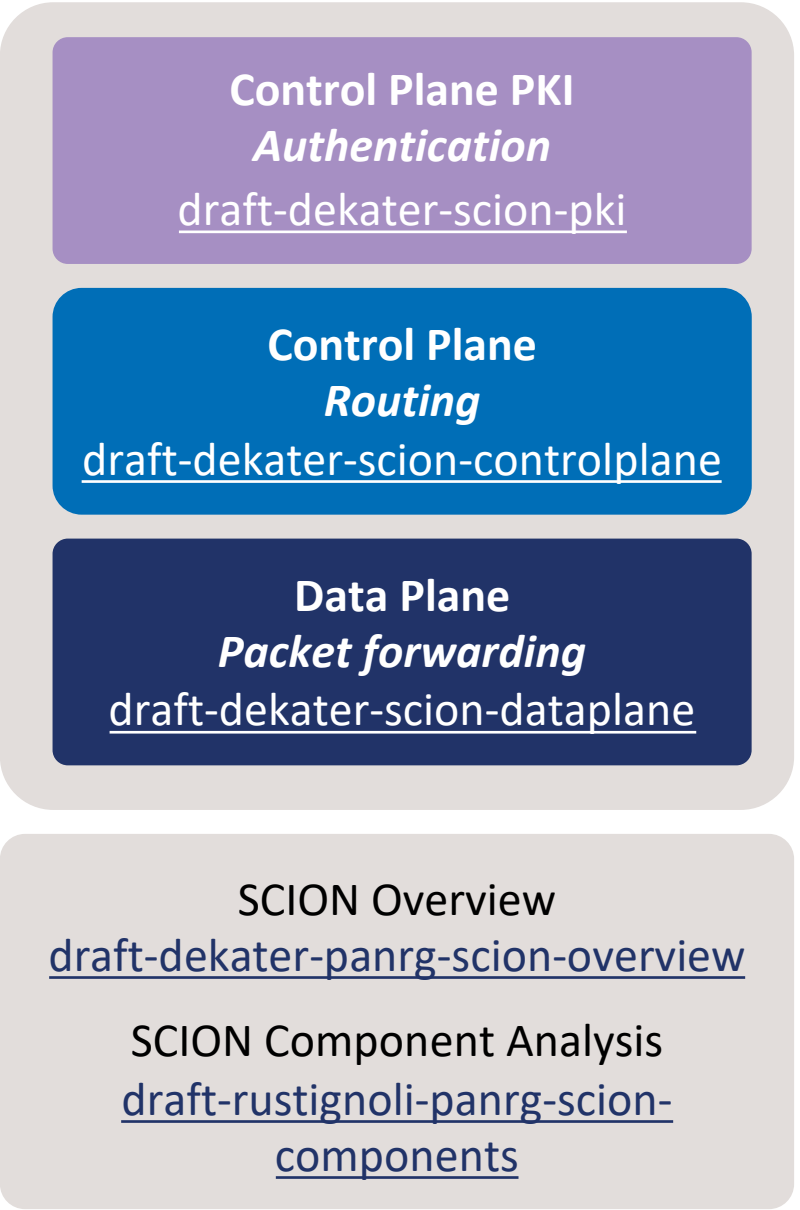

SCION at this IETF

- Hackdemo happy hour – Tomorrow 18:30-19:30
- Extended PANRG discussion – Tomorrow 13:00-15:00
  - Drafts
  - Deployment experience by Anapaya (SCION vendor)
  - User experience in the finance industry (SIX Swiss Exchange)
  - Discussion on the next steps within IETF/IRTF
- Path validation side meeting – Tuesday 18:30-20:00

We need you

- Comment and review on our drafts
- Support us in getting the right space within the IETF/IRTF

Nicola Rustignoli (nic@scion.org)
Corine de Kater (cdk@scion.org)

Feedback & Collaboration Welcome!

Please Applaud!!! (and the crowd goes wild)

Collective Communication Optimization(CCO): Use cases, Problems, and Requirements

Personal I-Ds:
【1】https://datatracker.ietf.org/doc/draft-yao-tsvwg-cco-problem-statement-and-usecases/
【2】https://datatracker.ietf.org/doc/draft-yao-tsvwg-cco-requirement-and-analysis/

Kehan Yao, China Mobile

Shiping Xu, China Mobile

Yizhou Li, Huawei

Hongyi Huang, Huawei

Dirk Kutscher, HKUST(GZ)

IETF 118 hotRFC

Concept:

Collective communication is a communication model which plays a key role in high performance computing and modern distributed AI model training workloads such as recommender systems and natural language processing. It involves a group or groups of processes participating in collective operations like AllReduce or AllGather. The communication model can be one-to-all, all-to-one or all-to-all and is usually realized by a sequence of unicast messages.

Use cases:

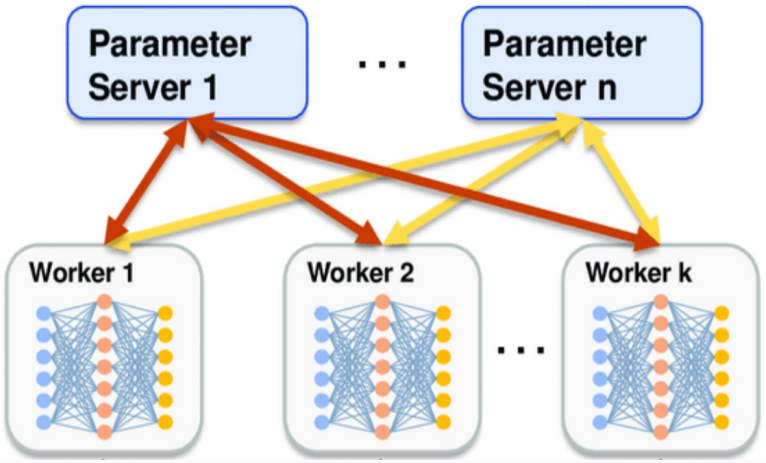

Distributed AI Model Training

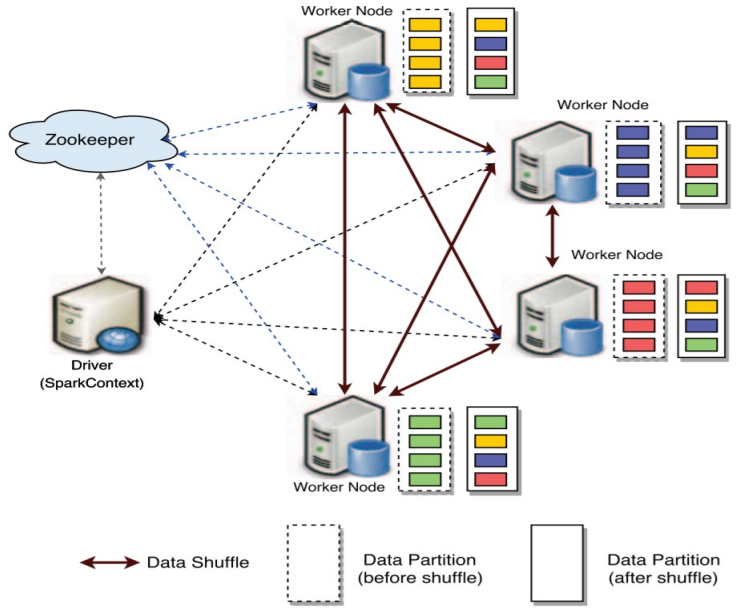

Spark Shuffle
in Big Data Anlaysis

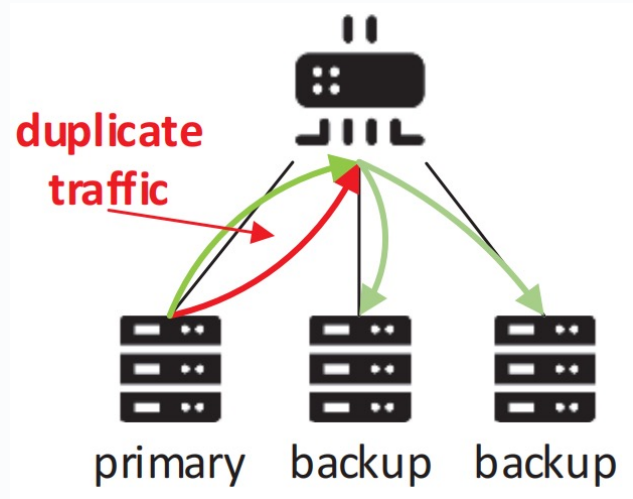

Distributed Storage

Major Problems & Observation:

- P2P implementation of Collective Communication incurs much overhead, reflected in:
  - **large bandwidth occupancy(duplications & redundancy)**
  - **much data movement(end-to-end transmission)**
  - **large number of data copies at endpoints(sending one pkt needs to copy at least one time).**

Communication bottleneck & performance degradation

- It should
  - save bandwidth(This is extremely important for BW-sensitive Apps like distributed AI model training workloads, **since BW is the new oil**).

    --- "***The metaphor is not from me, but I think it is quite impressive.*** ☺"
  - reduce data movement.
  - decrease data copies.

- Offloading collective operations to the network is important for achieving benefits above and very necessary, especially for these performance-driven Apps.

Communication Pattern:

Collective Operations (Allreduce,bcast,Allgather)

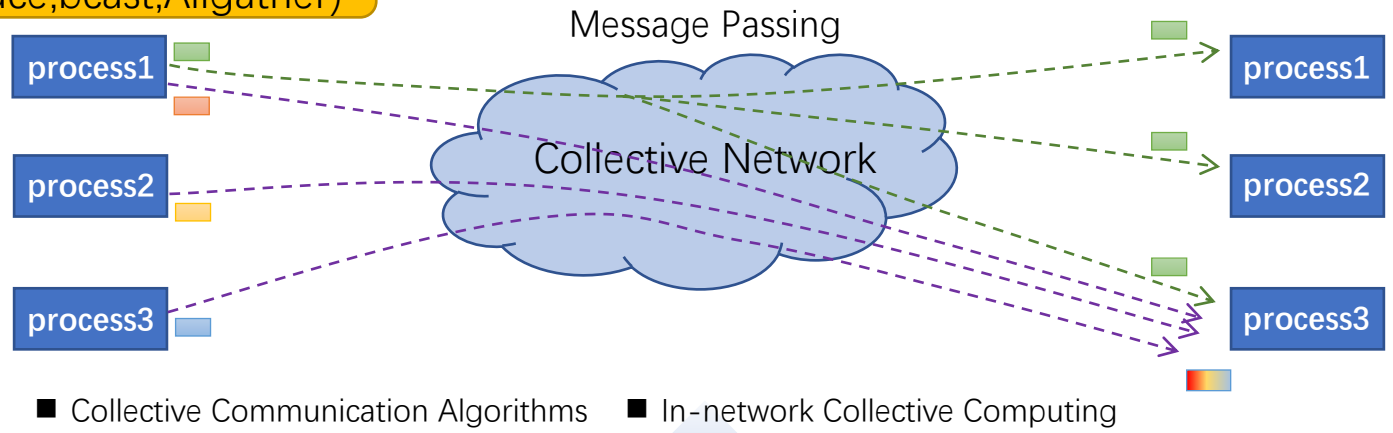

■ Collective Communication Algorithms ■ In-network Collective Computing

Underlying Network for Collective Communication

How to design ?

Underlying network (Multicast, P2P)

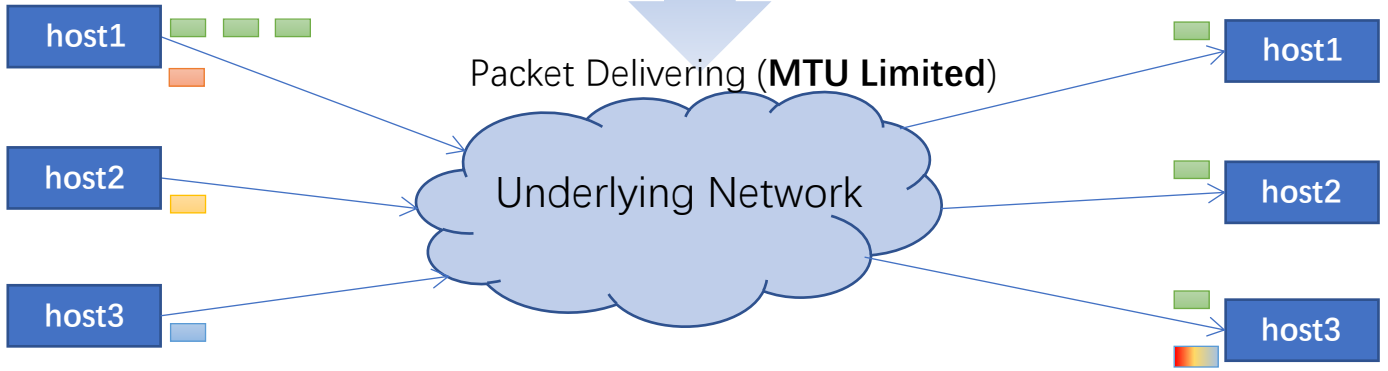

■ P2P Reliability ■ Transport Congestion Control ■ IP Multicast

Design Issues[1],[2]:

- ***Transport Issues:***
  - *Reliability*
    underlying network lacks collective communication reliability
  - *Semantic Gap*
    message passing vs packet delivering
  - *Blocking & Non-blocking*
    different optimizations for different communication modes
- ***One-to-Group Transmission:***
  - *IP Multicast for Message Bcast/AlltoAll/…*
    IP multicast is the most direct way, perhaps there is a better way
- ***Data & Control & Management:***
  - *In-network Primitives*
    collective operations based on unified In-network primitives
  - *Topology Awareness*
    to improve existing topology aware algorithms to support in-network computing

More in our I-Ds

【1】https://datatracker.ietf.org/doc/draft-yao-tsvwg-cco-problem-statement-and-usecases/
【2】https://datatracker.ietf.org/doc/draft-yao-tsvwg-cco-requirement-and-analysis/

Related Side Meeting in IETF118:

- https://wiki.ietf.org/meeting/118/sidemeetings

Title : Collective Communication Optimization(CCO),

Time Schedule: 9th, Nov, Thursday, 14:30 -- 16:00, Palmovka ½

Agenda: ***https://github.com/CCO-IETF/ietf118-side-meeting***

Looking for collaborators to seek for potential standardization opportunity of the work in IETF, and welcome for more discussions and contributions.

Thanks!

Please Applaud!!! (and the crowd goes wild)

Universal Name System (UNS)

Global and unified cryptographic namespace
for person and non-person entities

Universal Certificate Authority (UCA)

Global and automated cryptographic infrastructure
to solve code and data provenance, integrity,
authenticity, confidentiality and privacy

DNS + CAs	UNS + UCA
Domain-centric	**Entity**-centric
Entities are **mere resources** in domains	Entities have **independent existence**
Bolted-on cryptographic security	**Built-in** cryptographic security
Administered	**Automated and generated**
Global “Zero Trust” **impossible**	Global “Zero Trust” **possible**

Enabled by Confidential Computing (TEEs)

1. **Cryptographic protection of data in use**, even from system admins
2. **Remote attestation and verification** of HW and SW authenticity
3. **Verifiable confidential entropy** for decentralized cryptography

World Wide Weaknesses

#1 – Privileged insiders

#2 – Weak relational links

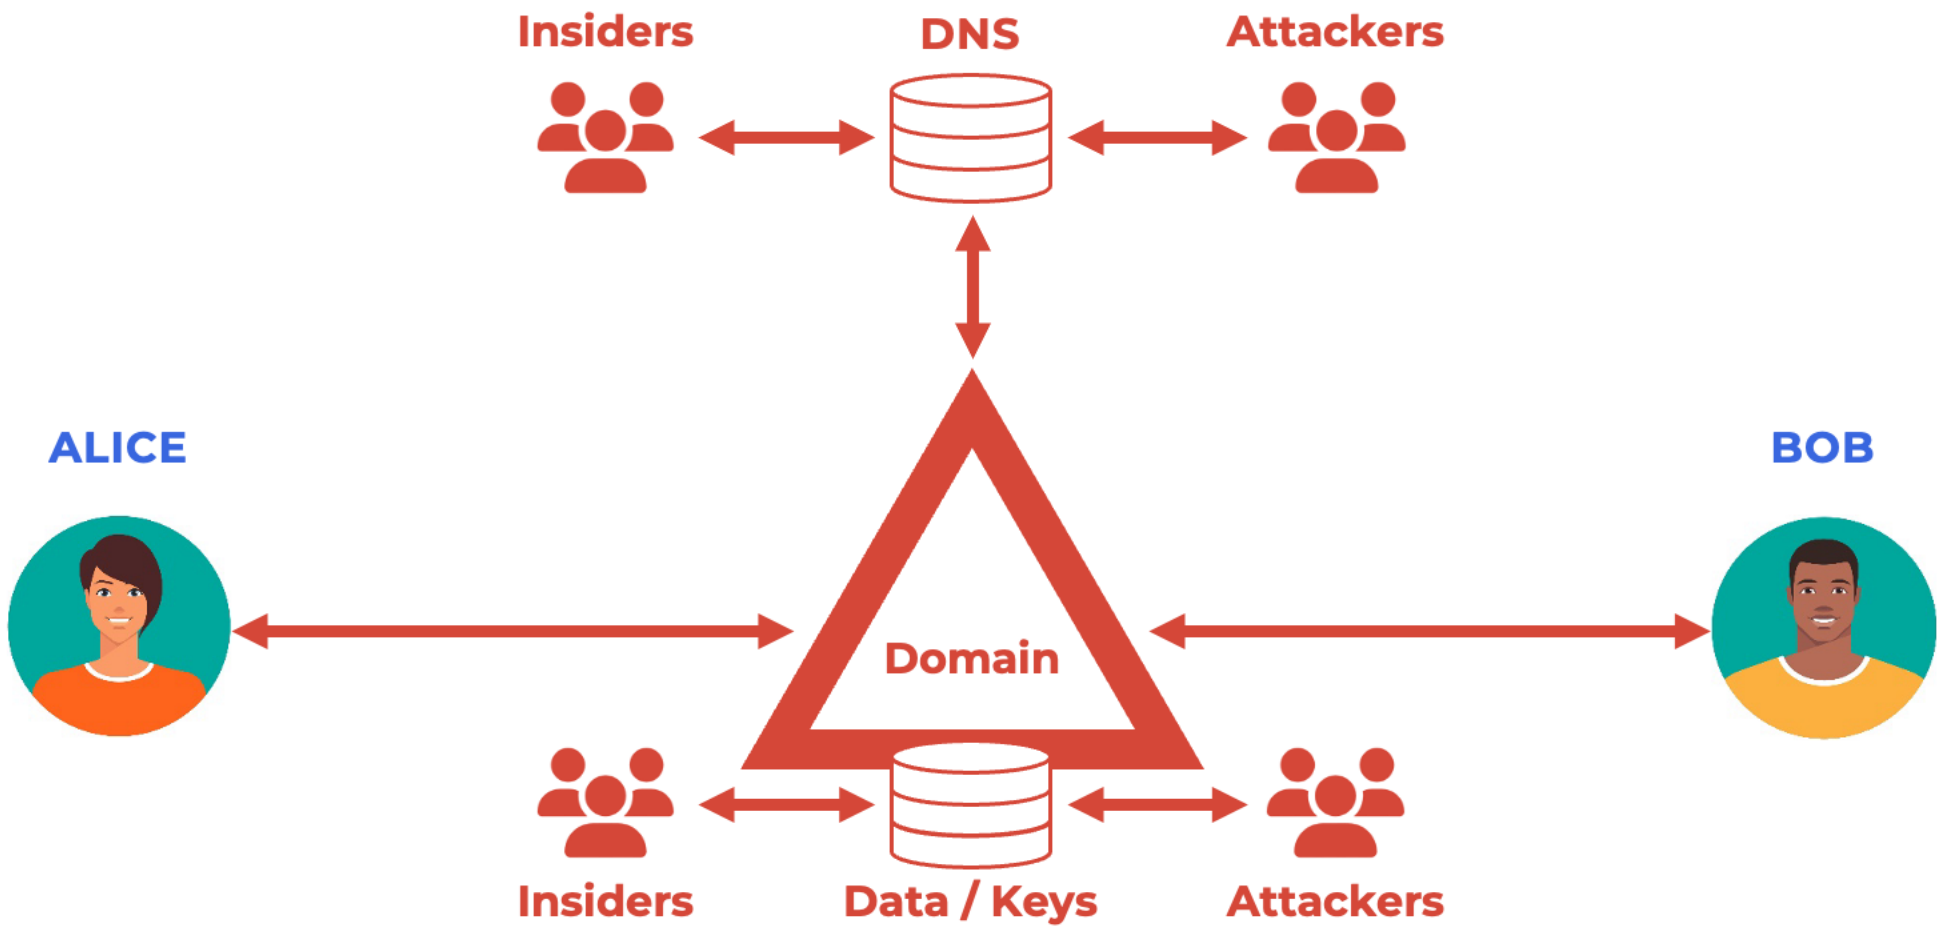
Insiders
DNS
Attackers
ALICE
BOB
Domain
Insiders
Data / Keys
Attackers

Domain name / path / resource

Entity StemID # relationship # keyname

is the HMAC key derivation function

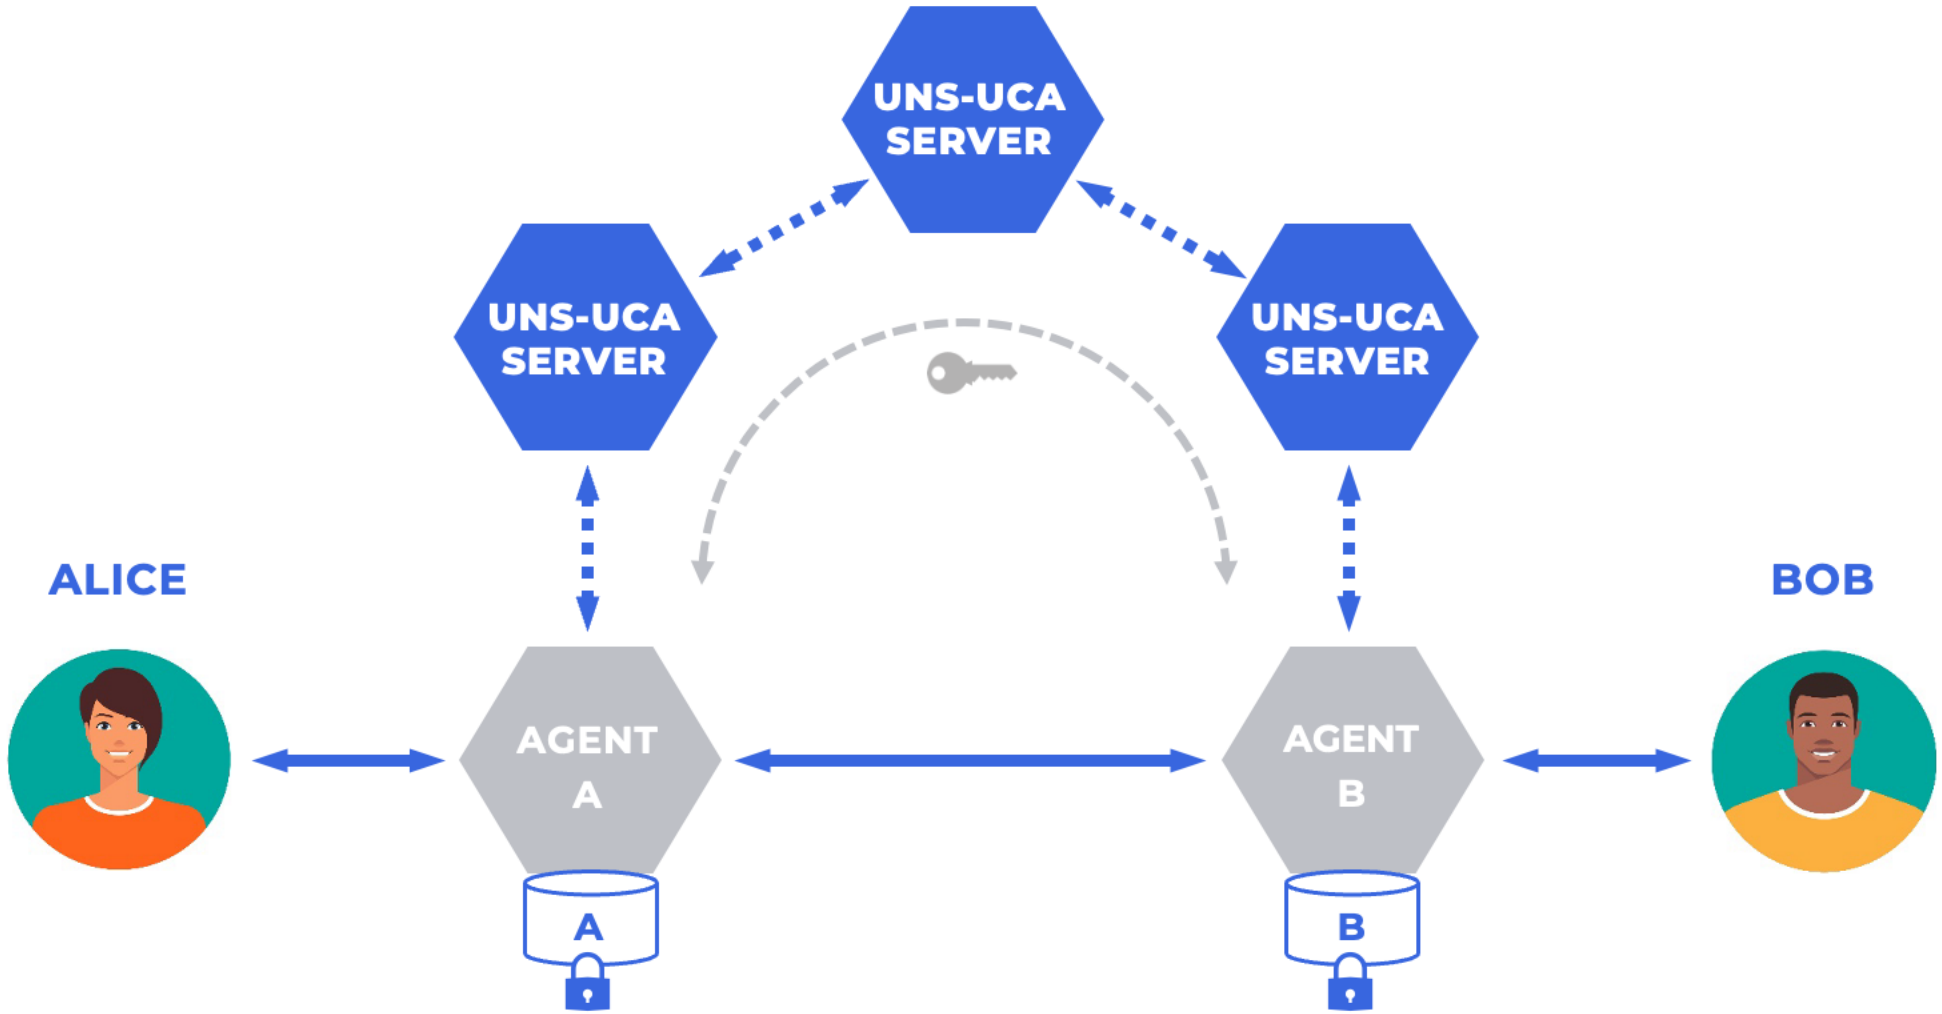
UNS-UCA SERVER
UNS-UCA SERVER
UNS-UCA SERVER
ALICE
BOB
AGENT A
AGENT B
A
B

Chaining entities (not blocks)

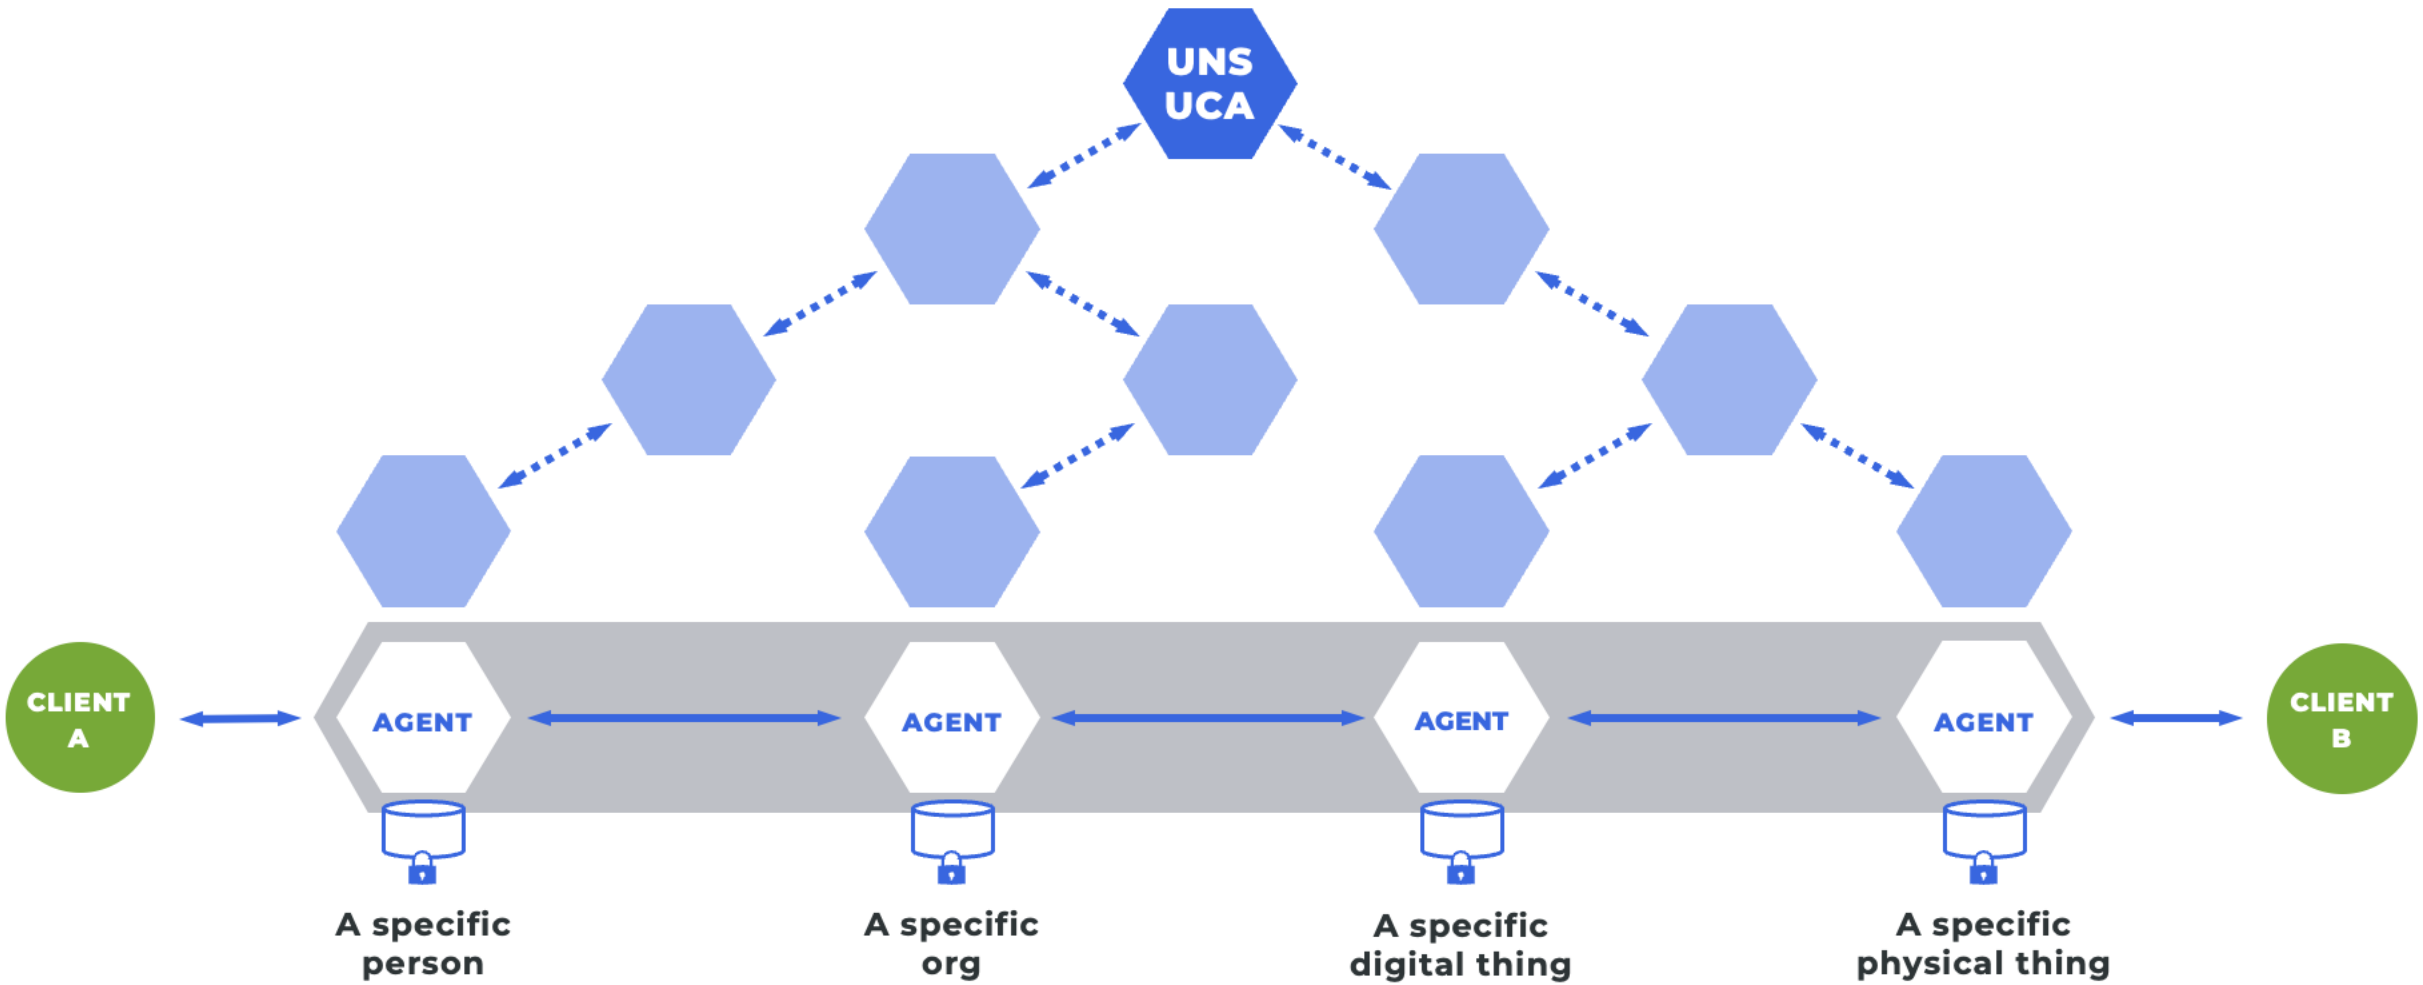

Fully automated, no humans in the middle

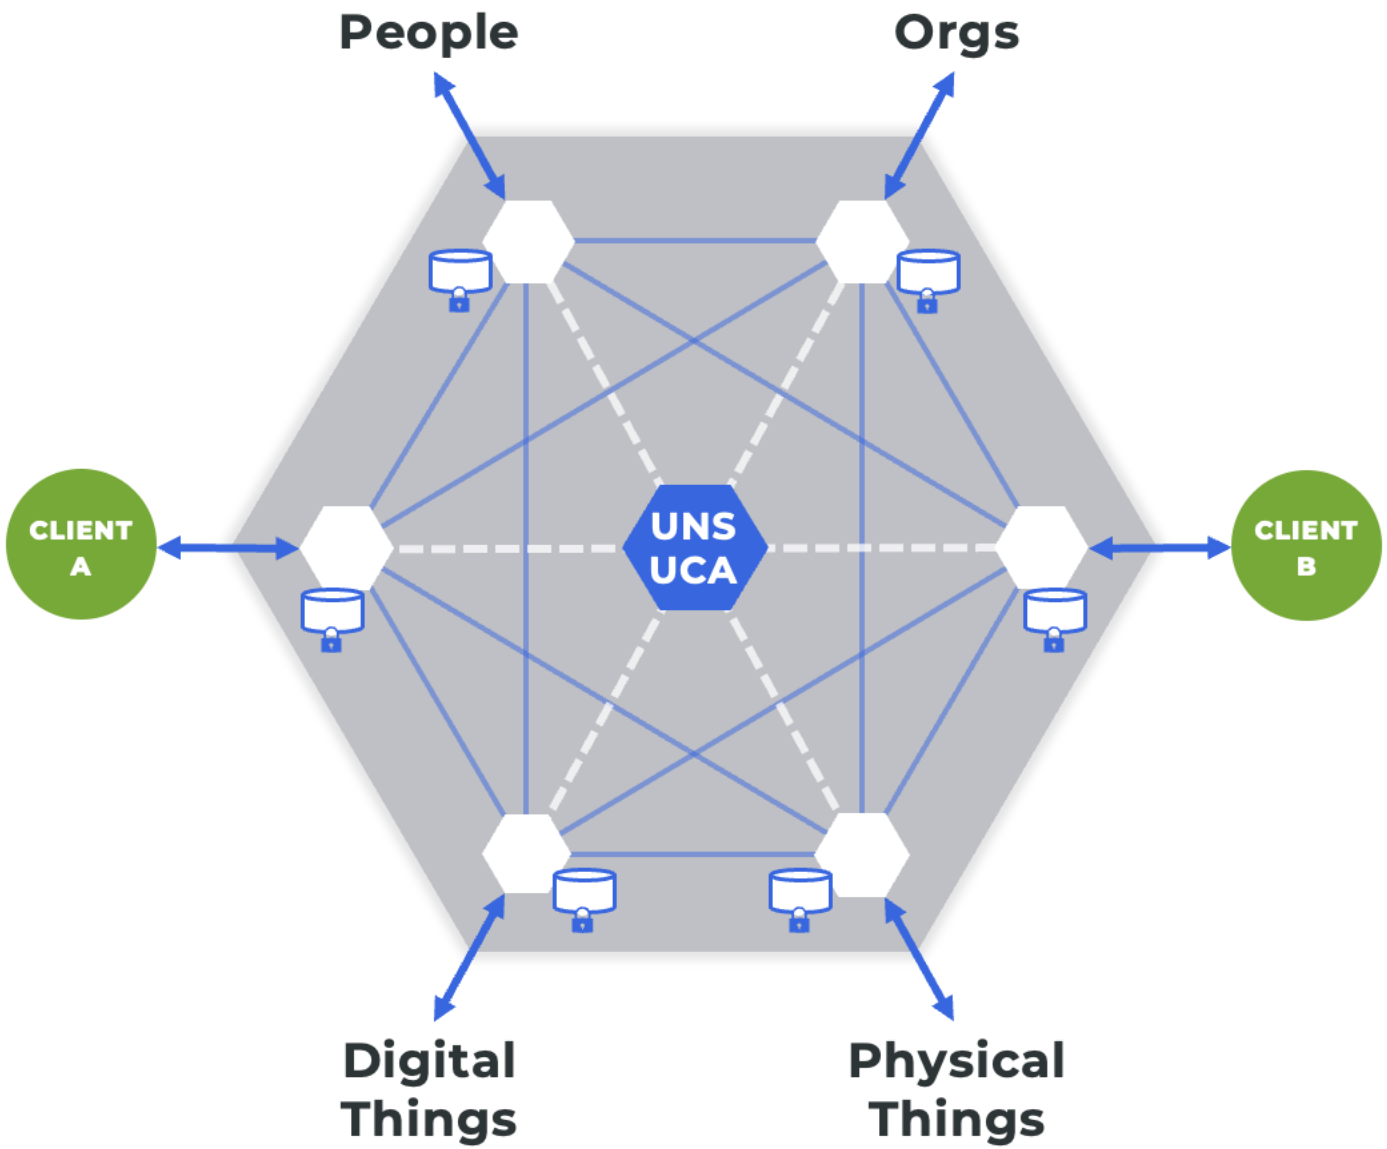

Present and discuss this work

Figure out where in IETF this work should go

Starting coalition for **global and neutral governance**

Side Meeting in **Palmovka 1/2** on **Monday @ 14:00**

manu@hushmesh.com

Please Applaud!!! (and the crowd goes wild)

KIRA – Scalable Zero-Touch Routing

KIRA: Kademlia-directed ID-based Routing Architecture

Karlsruhe Institute of Technology

Contact:
bless@kit.edu

- Scalability: 100,000s of nodes (in a single domain)
- Zero-touch: no configuration required
- Goal: provide highly resilient autonomous control plane connectivity
  - Running code provides zero-touch IPv6 connectivity
- First Internet-Draft: https://datatracker.ietf.org/doc/draft-bless-rtgwg-kira/
  - Want to standardize it
  - Many practical options → need broader IETF expertise
- Scheduled presentations @IETF118
  - RTGWG Tuesday Session I (Nov 7th), 09.30h
  - NMRG Friday (Nov 10th), 13.00h–15.00h
- Side meeting/BarBOF Wednesday Nov 8th, 19.00–20.00h, Karlin 4
  - KIRA use cases, Q&A, collaboration, next steps towards standardization

More Info

https://s.kit.edu/KIRA

Please Applaud!!! (and the crowd goes wild)

Will Post Quantum Crypto make Constrained IoT Devices and Networks obsolete?

Hannes Tschofenig

<Hannes.Tschofenig@gmx.net>

Lots of Progress over the last 10+ years

- Constrained IoT devices have traditionally not be blessed with great security capabilities.
- Work in the IETF and other bodies to tailor security and protocols to those devices.
- Today, they can run public key crypto well.

Picture References

Optimizations

- A lot of energy was spent with “compressing” protocols to reduce every byte of protocol exchanges for low-power networks like LoRaWAN.

Examples	Bytes
DTLS 1.3 - RPKs, ECDHE	880
cTLS-08 - X.509s by reference, ECDHE	405
EDHOC - Signature X.509s, x5t, ECDHE	242

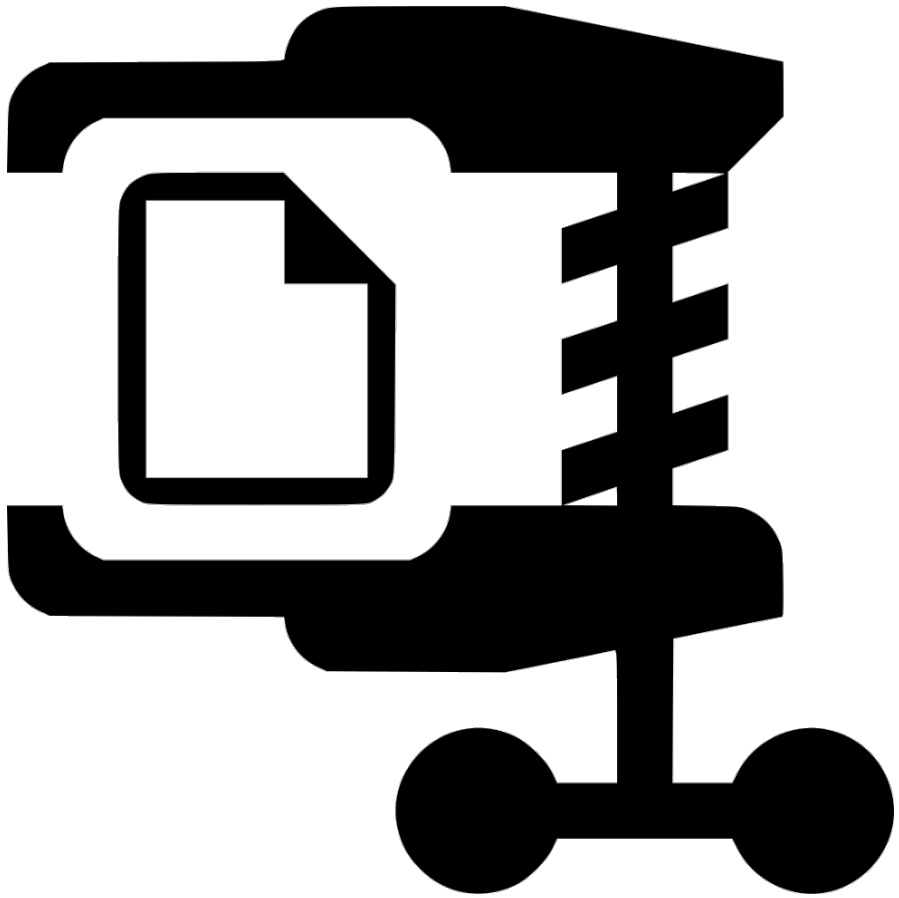

Reference: https://datatracker.ietf.org/doc/draft-ietf-iotops-security-protocol-comparison/

Picture reference

… then PQC algorithms came along!

- Performance and key sizes not great.
- Uncertainty about
  - the timeframe for Cryptographically Relevant Quantum Computer
  - the security of the new algorithms
  - the transition
  - operational aspects
  - ….

What does this mean for constrained IoT devices and networks?

- Preconditions are not great: Long lifetime and limited resources
- Options:
  - Switch to general purpose hardware
  - Use symmetric key cryptography (with longer key sizes) / Kerberos
  - Invent new cryptographic algorithms
  - ???

I am interested in your view.

Reach out to me at Hannes.Tschofenig@gmx.net

Please Applaud!!! (and the crowd goes wild)

Can we improve certificate/JWT/CWT revocation?

Hannes Tschofenig

<Hannes.Tschofenig@siemens.com>

A lot has been said about certificate revocation already

- **Long-lived certificates** require a story for revocation. Solutions are available but usage remains “mixed”.

- Certificate Revocation Lists (RFC 5280)
- Online Certificate Status Protocol (OCSP) + extensions for stapling in TLS (see RFC 6961) and other protocols
- CRLite (Mozilla)
- CRLSets (Google)

- Reducing the lifetime of certificates is also frequently being proposed (and used).

JSON Web Tokens (JWTs) have now become certificates as well

- With the proof-of-possession extension (see RFC 7800) JWTs (RFC 7519) have effectively become certificates.
- With OAuth, these JWTs (when used as access tokens) are generally short-lived and created for use with specific relying parties.

- The work on **Verifiable Credentials** turns them into long-lived certificates.

  → JWTs require revocation.

- Same is true for CBOR Web Tokens (CWTs, RFC 8392) and the proof-of-possession extension defined in RFC 8747.

Status Lists

- New work in OAuth WG to define **new revocation mechanism**.
- Mimics the “Let’s Revoke” concept (academic publication, focused on X.509 certificates)
- Prior work also in the W3C on Verifiable Credentials Status List v2021 (w3.org)

What are Status Lists?

- Issuer adds a URL to the status list and an index to the JWT (or CWT)
- Issuer maintains information about revoked JWTs/CWTs in a bit string – called status list.
- Verifier fetches this status list by
  - Downloading the status list from the URL provided in the JWT/CWT
  - Retrieves the bit position based on the index value.
- The status list (containing the bit string) itself is again a JWT/CWT.
- To reduce the size of the bit string, apply GZIP.
- To make it bigger again then apply base64encoding ;-)

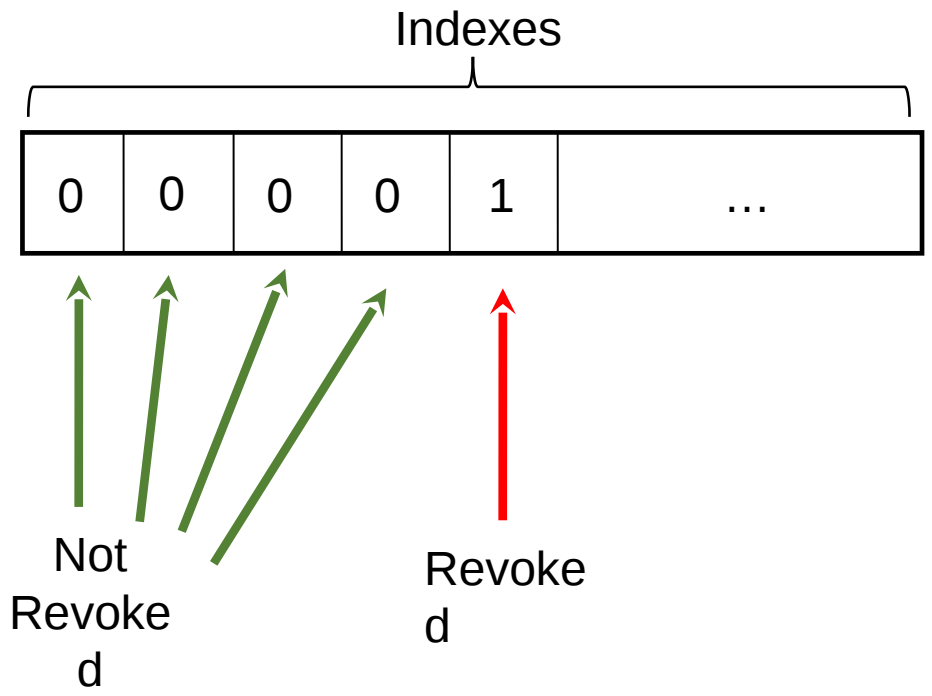

Your experience is needed!

- Is this a useful concept?
- If it is useful for JWTs/CWTs, should it be applied to X.509 certificates as well?
- Is there room for improvement?

Come to the OAuth WG and tell us!

Please Applaud!!! (and the crowd goes wild)

Merkle Tree Ladder Mode (MTL) Signatures

Joe Harvey
jsharvey@verisign.com
IETF-118

What is MTL Mode?

MTL Mode is a method for reducing a signature scheme's operational impact on an expanding message series.

- MTL mode is a technique for using a signature scheme to authenticate an evolving series of messages
- Rather than signing individual messages, MTL mode signs Merkle Tree Ladders
- Ladder = subset of nodes or "rungs" in generalized Merkle tree construction (not just a single root)
- Messages are authenticated with Merkle proofs relative to ladders
- Ladders provide backward compatibility since they can potentially verify Merkle proofs constructed relative to future ladders too
- Signature on message can be just Merkle proof + reference to signed ladder
- MTL mode operations can be aligned with the underlying signature scheme to ensure proper cryptographic separation (I-D proposes alignment with SPHINCS+)
- Presented by Burt Kaliski (Verisign) at NIST 4th PQC Standardization Conference and CT-RSA 2023

Benefits of MTL Mode

- Hash-based scheme ➔ quantum-safe design
  - “Stateful” hash-based (if evolving Merkle tree is considered to be state), but graceful degradation of security instead of key compromise if state is reused
- Hash functions are already available in many hardware platforms, making MTL mode performant
- Merkle proofs are typically much shorter than PQC signatures ➔ reduces size of messages that are transmitted across the wire or stored in memory/cache
- Batching can reduce the number of underlying signatures computed
- Hybrid signatures can be applied to ladders rather than individual messages

Outline of draft-harvey-cfrg-mtl-mode

1.	Introduction	Introduces MTL Mode.
2.	Preliminaries	Gives preliminaries including definitions, operators, functions and algorithm style.
3.	General Model	Presents the general model for authenticating messages in MTL mode.
4.	Security Parameter, Cryptographic Functions, Address Scheme	Introduces the security parameter, abstract cryptographic functions and address scheme used in the document, which are based on SPHINCS+.
5.	Computing Data Values from Messages	Shows how to compute data values for the Merkle node set from messages.
6.	MTL Node Sets	Describes the various concepts behind MTL mode operations including seeds and series identifiers, node sets, leaf nodes, internal nodes, ladders, authentication paths and backward compatibility.
7.	Data Structures	Defines the data structures for ladders, rungs and authentication paths.
8.	MTL Node Set Operations	Provides interoperable specifications for MTL node set operations.
9.	Signing and Verifying Messages in MTL Mode	Discusses how to sign and verify messages in MTL mode, including the concepts of "full" and "condensed" signatures.
10.	SPHINCS+ in MTL Mode	Proposes instantiations of SPHINCS+ in MTL mode using the SHAKE and SHA2 hash function families.
11.	Related Work	Discussion on related work.
12.	IANA Considerations	Comments on IANA considerations.
13.	Security Considerations	Covers the security considerations.
14.	References	Lists the references.

Intellectual Property

- Verisign announced a public, royalty-free license to certain intellectual property related to the Internet-Draft
- The license provides a “Standards Development Grant” for the purpose of facilitating standardization of the Internet-Draft
- The license also provides a “Grant on Standardization” for MTL Mode Implemented Using Binary Rung Strategy
- IPR declarations 6170-6176 give the official language (datatracker link)

Next Steps

- Please review the draft and provide feedback
- We have released an open-source library that combines MTL Mode with SPHINCS+
  - https://github.com/verisign/MTL
- We also plan to publish an I-D on using MTL Mode with DNSSEC (DNSOP?).
- Pseudo-code for data structures and algorithms in current draft is runnable Python code (see Appendix A).
  - Test code (see Appendix B) shows examples of how to do operations like sign or verify a message.

Please Applaud!!! (and the crowd goes wild)

Personal Digital Agent Protocol (pdap)

HotRFC
November 5, 2023

pdap@ietf.org

Enabling a shift
from
proprietary platforms **to personal agents**

- Replace forced platform association with one’s choice of community.
- My agent is **interoperable** by vendors and service providers.
- I can switch the host of my agent anytime. **No lock-in.**
- My agent’s **policies are portable** across host communities.
    - Swarm and federated personal AI agents can be supported.

Foundational Principle

Universal Human Right of Freedom of Association and Assembly

- Individual choice of hosting and support communities for one's digital agent.
- Self-hosting is supported for those that have the skill and interest.
- Research perspective: https://datatracker.ietf.org/doc/draft-irtf-hrpc-association/

Platform Issues and Regulatory Responses

- US: FTC vs. Amazon
  - Bundling of Search and Logistics - drives up cost for vendors
  - Most-favored-pricing clauses drive up cost for customers
- EU: Payment Services Directive PSD-2
  - Opens interface between payment services and banks
  - Increased competition for both payment services and banks
- EU: Digital Services Act
  - Customer lock-in and lack of transparency and customer agency
  - Very Large Online Platforms (VLOP) and Very Large Online Search Engines (VLOSE)
- India: Aadhaar, UPI, India Stack, beckn protocol
  - Avoid US / EU style platform oligopolies
  - Support essential digital tools as a public good (Aadhaar, India Stack)
  - Enable customers and vendors to choose their agents (UPI, beckn protocol)

Personal Digital Agent

- Self-Sovereign (independent of any jurisdiction or federation)
- Community-Hosted (self-hosted option for pure self-sovereignty)
- Semi-autonomous
- Intelligent (adaptable, learning, context-aware, conversational)
- General across all types of vendors, service providers and jurisdictions
- Accepted by most service providers and resource servers
    - Standardized
    - Fair
    - Cost-effective
    - Secure

Scope of a Personal Digital Agent Protocol

- Separate choice of Authorization Server (as agent) from Resource Server (as vendor)
  - IETF GNAP - **Last Call** - multiple implementations available
  - https://datatracker.ietf.org/doc/draft-ietf-gnap-core-protocol/
- Standardize Request Presentation and Authorization Tokens
  - OAuth 2.0 Rich Authorization Requests https://datatracker.ietf.org/doc/html/rfc9396
  - Message Signature https://datatracker.ietf.org/doc/draft-ietf-httpbis-message-signatures/
  - Authorization and Revocation Capabilities - TBD - Consider https://github.com/ucan-wg/spec
- Standardize the Service Endpoint for a personal digital agent Authorization Server
  - Support both URIs and DIDs
- Scope of the Authorization Server Policy Management Interface
  - TBD - Consider CEDAR https://www.cedarpolicy.com/en

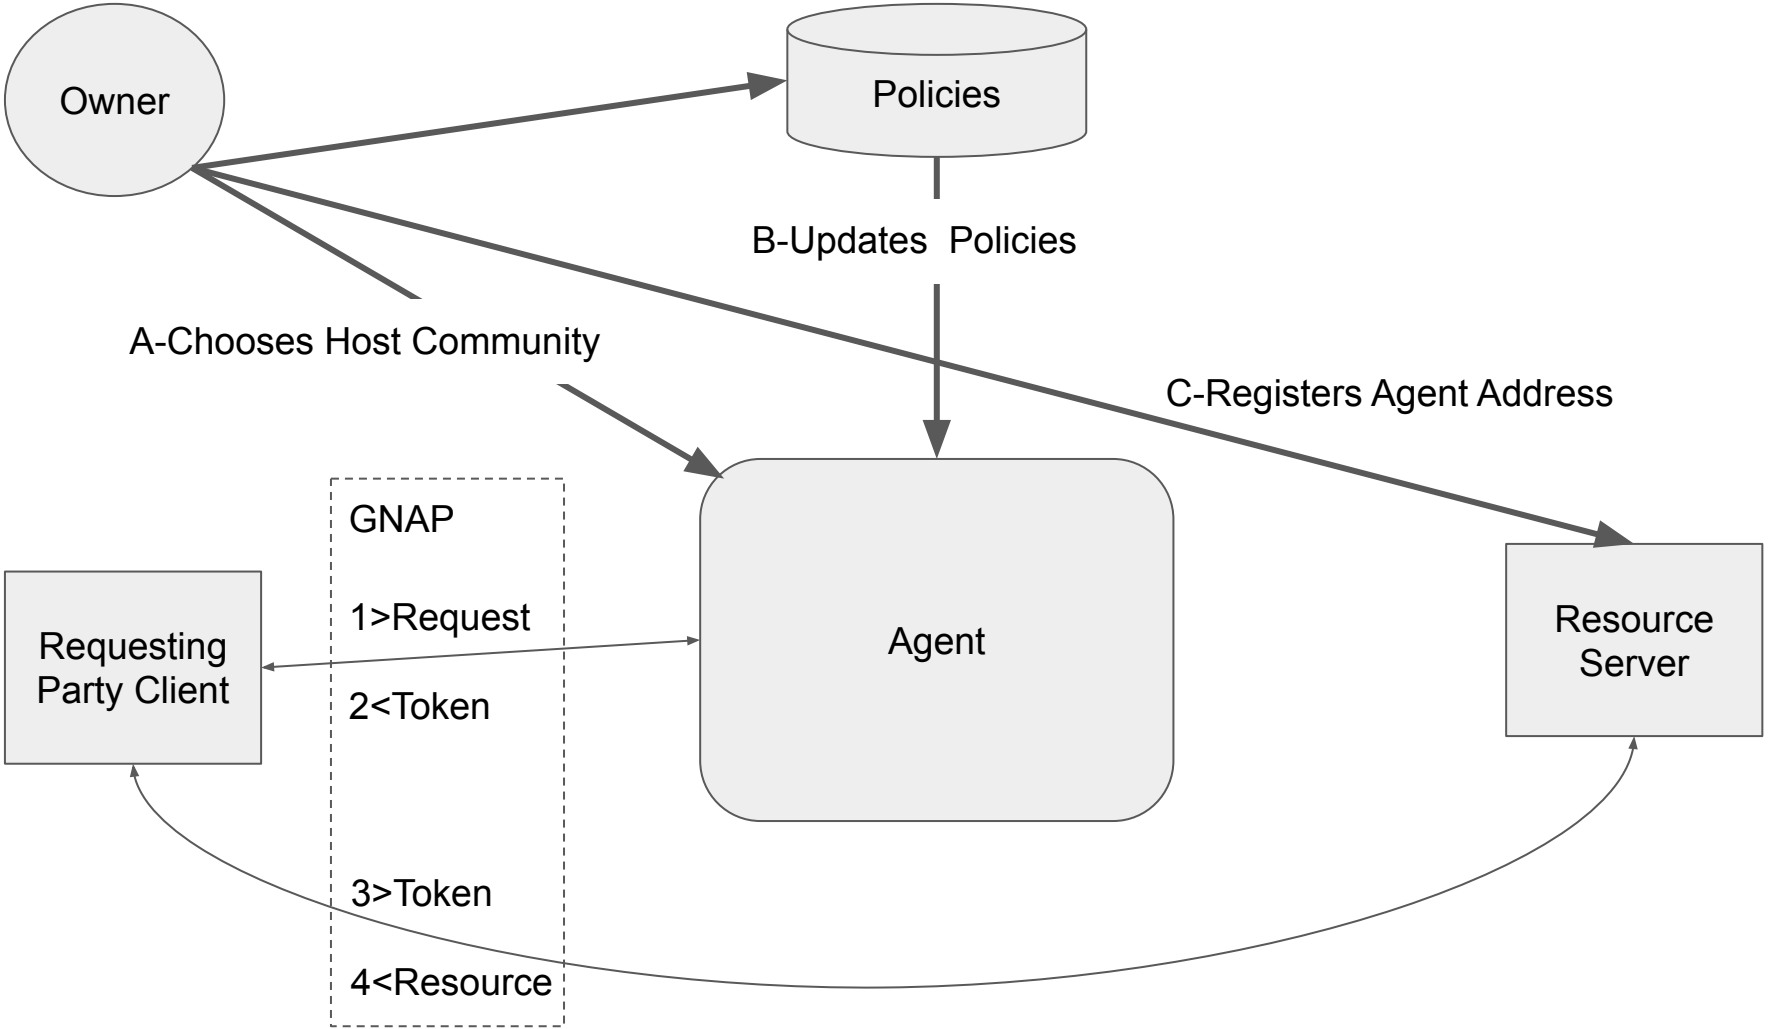
Owner
Policies
B-Updates Policies
A-Chooses Host Community
C-Registers Agent Address
GNAP
1>Request
2<Token
3>Token
4<Resource
Requesting Party Client
Agent
Resource Server

Steps toward a Personal Digital Agent Standard

- 2015-on: Kantara UMA 2 Protocol Standard
  - OAuth-based
  - “Open World” Authorization Server
- 2020-on: IETF GNAP Protocol Standard
  - Not OAuth-based
  - Open world user request state machine
- 2022: W3C Verifiable Credentials Data Model Standard
  - Validation and Verification as a commodity service
  - Supports open world, decentralized, and self-sovereign flows
- 2023: Personal Digital Agent Discussion Group
  - Define the scope of a vendor standard for interoperable personal agents
  - Excellent vendor and customer experience without platform lock-in
  - https://docs.google.com/document/d/19GU6L1QxaVsIfm9iBKg9T2qV9y0zttRmtlQuPireMMU/edit
- 2024: IETF Standard Workgroup Established

Join our mail list

https://www.ietf.org/mailman/listinfo/pdap

These slides and older notes at:

https://docs.google.com/document/d/19GU6L1QxaVsIfm9iBKg9T2qV9y0zttRmtlQuPireMMU/edit

Suggest vendors and implementers to participate.

agropper@healthurl.com

Thank you.

Please Applaud!!! (and the crowd goes wild)

HOW TO ENSURE TECHNOLOGY DOESN'T DO CERTAIN THINGS BY DESIGN

AN EXPLORATION FOR AN IRTF RG

THE IDEA

“I’d love to change the world
but
they won’t give me the source code.”

Some genius

BUT WHY?

SOME EXAMPLES

DPLs

CROSS-BORDER Xchange

AI
ML

NEURAL IMPLANTS and others

POSSIBLE/LIKELY AREAS

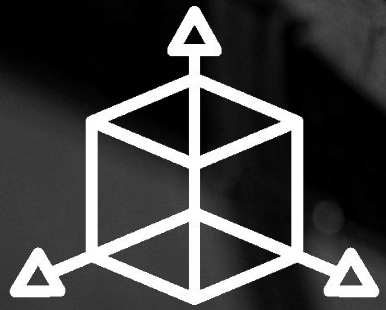

MODELING

TAXONOMIES

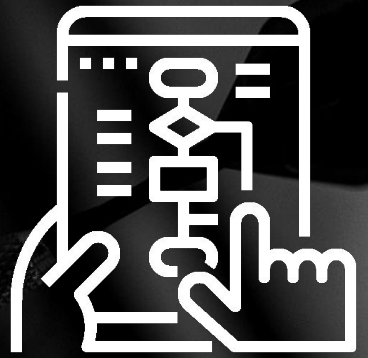

USE CASE SCENARIOS

THE ASKS

As this been attempted before?

What worked?

What didn't?

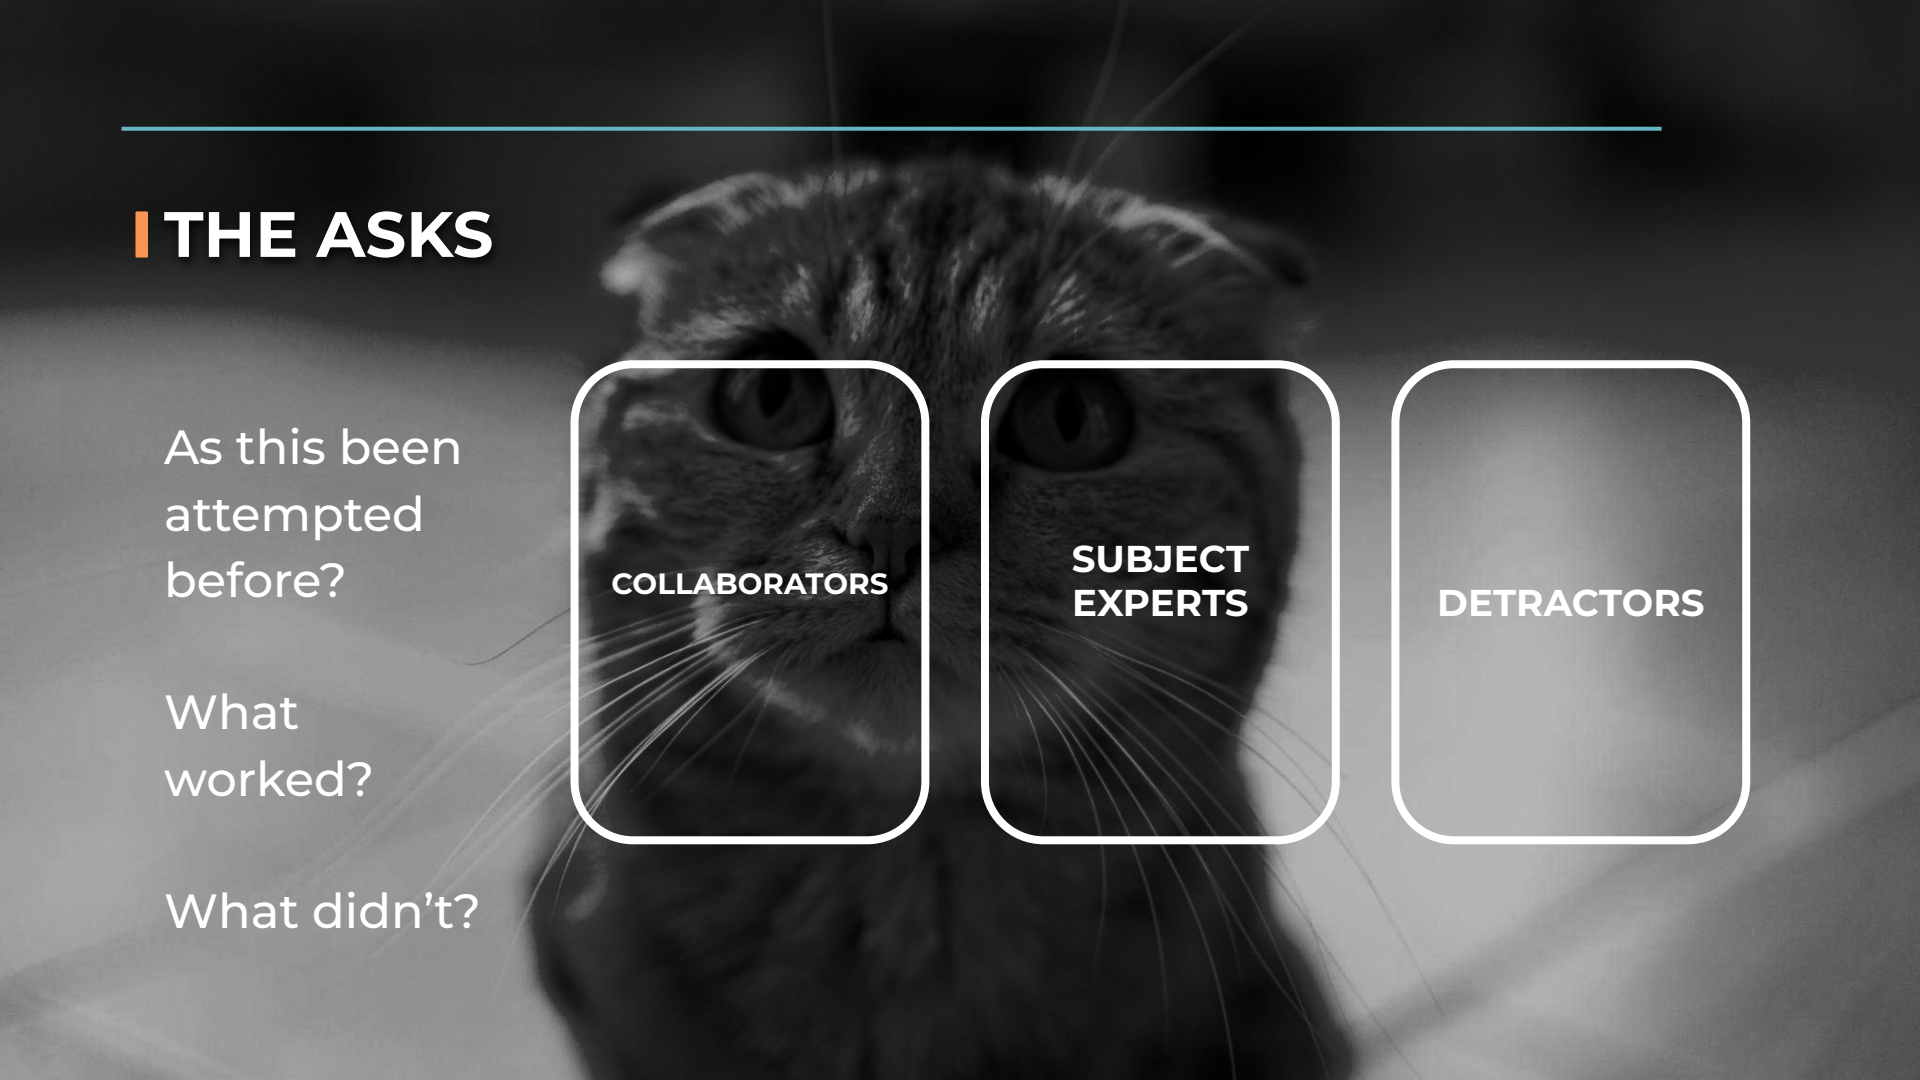

CONTACT ME

JFQueralt@TheIOFoundation.org

OR FIND ME AROUND THE VENUE

Image credits: Unsplash - Pixabay - Burst


Thanks!
From Unsplash

Please Applaud!!! (and the crowd goes wild)

draft-janfred-eap-fido
A new EAP method based on FIDO keys

IETF 118 in Prague – HotRFC | 05.11.2023

Janfred Rieckers

Why do we need something new?

- Eduroam configuration is not trivial
  - It’s too easy to misconfigure eduroam
- Certificate check is hard. There is no implicit way to derive the expected certificate names, users need to configure it manually
- Bootstrapping problem especially in BYOD environment

What are the new requirements?

- Trivial configuration
  - Security properties must not depend on the ability of users to configure their security parameters correctly
  - Users have no idea what the parameters mean and what implications they have
- Don’t use passwords
  - Authentication by Knowledge with the requirement to reveal the knowledge is a bad idea.

How do we solve this?

- Use asymmetric crypto that is easy to provision:
  - FIDO keys \o/
  - Should be available on most new devices in Software.
  - Can also be used with hardware-token
- Provisioning via web frontend (currently out of scope for EAP-FIDO spec)
  - EAP server just needs access to the DB of known FIDO Public Keys
- TLS certificate parameters are implicit through realm configuration
  - Users do not need to configure anything security related.
  - We just use WebPKI, it’s used for WebAuthn during registration anyway.

What now?

- See the EAP-FIDO Draft (draft-janfred-eap-fido)
- We have a side-meeting on Monday (tomorrow) 18:00 in Karlin 4
  - We are looking for
    - Feedback on Design
    - Experience of EAP operators
    - Input from EAP/RADIUS server and/or supplicant implementers
    - Input from people with FIDO/WebAuthn/CTAP experience
- The work will also be presented in the emu WG session
- Or find me in the hallways.

- Contact
  - Jan-Frederik Rieckers

Mail: rieckers@dfn.de
Phone: 0049 30 884299-339
Fax: 0049 30 884299-370

Address:
DFN-Verein, Geschäftsstelle
Alexanderplatz1
10178 Berlin

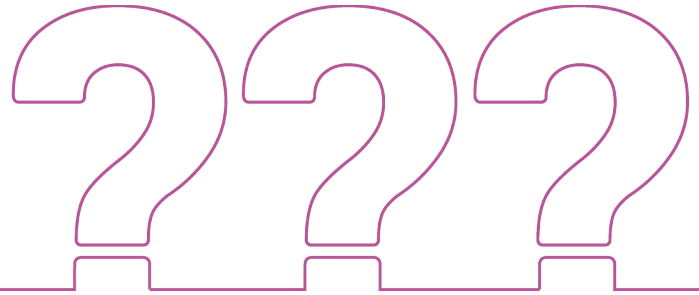

Please Applaud!!! (and the crowd goes wild)

Thank you to the presenters!