



Project Falcon: Self Sovereign Consent DNA Bitmap



If you'd like to follow along the webinar as you read this document [click here for the video](#) and start reading in the document at [The Grammar of Consent](#).

Thank you for reading this doc!

NOTE: We are hosting weeking meetings, and there are groups working on this project from different perspectives.

Feel free to request an invite here:

falcon@privacyco-op.com

Drop us a line and enjoy the read!

List of Author(s) and Contact Information

Include names and email address of author(s), along with other platforms/contributors that worked on this proposal

J. Oliver Glasgow — jay@privacyco-op.com

Alan Nekhom — alan.nekhom@gmail.com

W3C WICG Group -

<https://discourse.wicg.io/t/proposal-project-falcon-global-consent-dna-bitmaps-and-consent-management-service-cms/4634>

More to come...


Document Map

User stories will be lettered with capitals: A, B, C.

API will be lettered in lower-case and are bookmarked for easy location: a), b), c).

Response to Call for Proposals

This next section complies with the template, “GPWG: Call for Proposals on a Global Technical Architecture for Transparency and Control”.



1. Problem Statement

See [Executive Overview](#)

2. Key Considerations

See [Commercial Consideration](#)

3. Goals

A Consent Name System component deployed as a system that solves all

“crazy consent corner cases” (or C-4) for any platform (so that all other proposed Rearc solutions can focus on their own core functionality), and that can be operated by multiple entities using common standards of interoperability and rules to provide:

- A simple and free “yes or no” election responses to use data of any subject (person, place, or thing) based on a request containing a combination of the following:
 - One to many businesses
 - One to many legal uses
 - One to many jurisdictions
- Available as a free commons (just like DNS is today)
- Support for “lazy provisioning”
- Support for externally defined AuthN+AuthZ
- Works well transactionally as well as *en masse* for Big Data
- Interfaces for any transport/protocol (including signals)
- Can adopt data set from TCF
- Can be addressed by pointers and privacy strings
- Can be initially set by jurisdiction, modified per business, and ultimately set by authorized subject
- Supports licensing for each subject by authorized agents
 - Gains affirmative express consent compliant with various regulations where applicable
 - Timeboxed values
- Versioned by Date
 - Data collected today used in accord with regulations 50 years from now
 - Data collected 50 years from now used in accord with regulations today

4. High-level proposed approach

Including diagrams and text descriptions

See Solution starting at [The Grammar of Consent](#)

5. Open questions

See [Open Questions](#)

This next section complies with the template, “Addressability Proposal Template” and “Accountability Proposal Template”

Identify Problem Space(s) This Proposal Addresses

A micro-component for Self Sovereign Consent (SSC), as outlined in this proposal can help other platform and product solutions in ANY space accomplish the use of data for “secondary purposes” for all of the following situations...

1. When there is no identifier - Project Falcon relies on lazy provisioning information by default strata of consent settings for any business, legal use of data, and jurisdiction. When no identifier is provided, the best yes/no answer is still provided.
2. When only on-device identification can be used - Project Falcon is geo-distributed, geo-redundant, and can be cached on a device. As the bitmap for each Person, Place, or Thing identifier is projected to be less than 84kb, it can easily be cached even on a device for offline use.
3. When there is an identified user via a user-provided identifier - Identified user can be an extremely complex entity, as it can be a person, a device used exclusively by that user, an unverified attribute, or even a persona. In all of these cases and more, Project Falcon can resolve consent on a spectrum of request types from single transaction to even very massively large insights (multi-petabyte Big Data map reductions).
4. When there is an identified user via a pseudonymous ID not provided directly by the user - Project Falcon supports interfaces for Self Sovereign Identity, Federated Identity, and more, as it is a component that resides between platforms and products as a marketplace, not a product itself. A single SSC DNA Bitmap can be associated with such a pseudonymous ID, which in turn can be related to other identifiers (IDs as well as attributes) by other platforms/products. In fact, just by having a SSC DNA available might make such relationships possible and permissible in a myriad of competing rights and regulations for such platforms/products.

Considerations or Applicability by Media Channel?

Different media channels present unique identifier and technology environments. This approach provides SSC DNA addressability solutions within:

- Web/Browser environments
- Mobile App environments
- OTT environments

But we will also demonstrate addressability in the following:

- IOT
- Big Data

- Devices
- ...and more

Business Use Case Support

We believe at this stage that Project Falcon can support all of the following use cases. We will happily remove any that are found to fall out as we develop the spec:

Publishers

Publisher Life-Cycle Stages and Distinct Business Activities		Business Impact	Proposal Support?
Content/Service Development			
	Audience Analytics	High	X
	Content Management	None	X
Consumer/Customer Acquisition			
	SEO / SEM	Moderate	X
	Advertising	High	X
	Social Media Promotion	High	X
Personalization and Delivery			
	Basic Content Delivery	None	X
	Content Personalization	High	X
	Consumer Privacy	High	X
	Pay Walls	None	X
	Registration Walls	None	X
	Blocker Detection & Recovery	Moderate	X
	Fraud Detection	High	X
Monetization			
	Advertising - Direct Sales	High	X
	Advertising - Programmatic Sales	High	X
	Affiliate Links	High	X
	Audience Data Sales	High	X
	Audience Extension	High	X
	Donations	None	X

	Sponsored Content	High	X
	Subscriptions	None	X
	Ecommerce Goods and Services	High	X
	Virtual Goods and Services	High	X
Re-Engagement / Retention			
	Email Promotion	None	X
	Advertising	High	X
	Notifications	Slight	X

Marketers

Marketing Life-Cycle Stages and Distinct Business Activities		Business Impact	Proposal Support?
Consumer Research	Closed Loop ROI Analysis	Moderate	X
	Lift Studies (Awareness, PI, etc)	High	X
	Look-a-like Audience Modeling	High	X
Communications Strategy		None	
Creative Development		None	
Media Channel Execution + Optimization	Affiliate Marketing	Moderate	X
	Algorithmic Optimization	High	X
	Audience Targeting - Prospecting	High	X
	Audience Targeting - Retargeting	High	X
	Audience Targeting - Exclusion	High	X
	Contextual Targeting	Slight	X
	Conversion Measurement (view and click-through)	High	X
	Dynamic Creative Optimization (DCO)	High	X
	Frequency Capping	High	X
	Impression Counting (inc.	High	X

Post Campaign Analysis	Viewability, IVT)		
	Multi-touch Attribution (MTA)	High	X
	Pacing	High	X
	Sequential Messaging	High	X
	Conversion Measurement (view and click-through)	High	X
	Multi-touch Attribution (MTA)	High	X


Short Outline of the Solution:

1. The Grammar of Consent--what it is we are actually trying to solve.
 - a. Singular Subject--a Person, Place, or Thing
 - b. Compound Predicate--one to many combinations of a Business (or Gov Agency), a Legal Use of Data (sometimes called a "Program"--we know...we wish lawyers hadn't picked that word), and a Jurisdiction. This can be a single X+Y+Z or an array of [xyz1, xyz2, etc].
2. Building the Self Sovereign Consent (SSC) DNA Bitmap.
 - a. Put the X + Y on a two dimensional plane and understand its properties and attributes and how they work.
 - i. Need for standardized business/gov agency ledger additions. There are competing standards that can be considered.
 - ii. Also similar need for Program ledger additions.
 - b. Replicate the X+Y into various copies on a Z axis and modify their default setting for a single jurisdiction per layer.
3. SSC DNA Bitmap Consent Resolution. How the Bitmap resolves singular, complex, compound, or complex/compound requests in any language and in any technology.
 - a. 3D drill down into one Z axis intersection to see how we arrive very quickly at a "0 or 1" response.
4. Compression
 - a. Using various native bitmap libraries to minimize not only footprint but also resulting functionality for any technical realization.
 - b. Example: hex code.
5. SSC DNA metadata
 - a. Storage Address
 - b. Assignability
 - c. Provisioning
 - d. Relationships
 - e. Architectural distribution/deployment
6. Lazy Provisioning
 - a. How a response can be given even before a user makes their election known.

- b. Living Logs
- 7. Consistent Addressability
 - a. Supports all use cases from transactional to Big Data in a decentralized + cacheable distribution.
 - b. Browser
 - c. Apps in a Browser
 - d. Big Data
 - e. 3rd Party Use of Data
 - f. Compound/Complex use of Data
 - g. Contract Snippets
 - h. Masks (bitmasks)
- 8. The Pointer - a focus of additional standardization work in W3C
 - a. Open for interpretation and even supports existing IP claims, etc.
 - b. Prependability
 - c. Pointer Resolution Rules
 - d. Referential Integrity
 - e. Organization Identifier
 - f. Privacy String--competing standards need to be considered.
 - g. Optional Subject ID
- 9. Pointer Resolution Rules, a subcomponent of the pointer, can be singular, compound, complex, or compound/complex.
 - a. Puts customization in the hands of those legally on the hook to adhere to regulations for their particular use.
 - b. Can handle crazy corner cases.
- 10. Standardized updates
 - a. Low threshold for changing 1 to 0.
 - b. High threshold for changing 0 to 1.
 - c. Supports externality of Affirmative Express Consent.
 - d. Supports Authorized Agency.
- 11. Proposed re-purposing of DNS server open-source code and architecture
- 12. Supports platforms below and products above the marketplace
 - a. A role for SSI
 - b. A role for Blockchain
 - c. A role for AuthN+AuthZ
- 13. Potentially removes the data plane from simple tech stack

Proposed Solution:

Executive Overview



Cookies are going away. Privacy Regulations are increasing and tightening. Businesses in all sectors (not just tech) are feeling pressure to grow revenue from secondary uses of data, all while consumers are growing increasingly upset by perceptions of data exploitation.

As new innovations and solutions are being defined, they are starting their documents with a list of assumptions that unnecessarily straightjacket imagination. These include: Consumers can't know what 3rd Parties are using data, solutions can't adhere to all regulations, we can't support asynchronous Identity, stored data can't be used the same as transactional data. And there are more.

Information Rights define the abilities to use information as a scarce resource with competing rights. It can be defined in sentence form. The Subject (a Person, Place OR Thing) which relates to a Predicate (Organization, Program (legal use of data), AND Jurisdiction). As a constraint for this proposal, we will treat it as the subject (OR) must be singular and the predicate (AND) must contain at least one of all three elements.

While a great deal of logic and circumstance can go into whether or not an enterprise can use data for a particular purpose in a particular jurisdiction (legal nexus) for a particular subject, it ultimately results in a binary conclusion of yes or no, or in computer-speak, 1 or 0.

These types of elections, and the applicable Program(s) (a legal use of data), can apply to various corrals of consent, such as **accessing** data from a device, selling or **sharing** data with 3rd parties, and **sending** communications to the subject. But today there is no consistent ability to resolve all of those corrals for a single use case, which leads to confusion by enterprises and complexity for users so great that there is today no meaningful end-to-end consent understanding for most significant uses of information (data). This hurts legal implementations of the phase space between privacy and publicity and impedes enterprise profitability.

Default settings for consent for any such use can be set by policy or regulations. They can be modified by businesses on advice from their counsel, but ultimately get cast by a subject as an election. Again, it results in a 1 or a 0.

The one thing consistent for all complexities mentioned is the ultimate distillation down to a 1 or a 0.

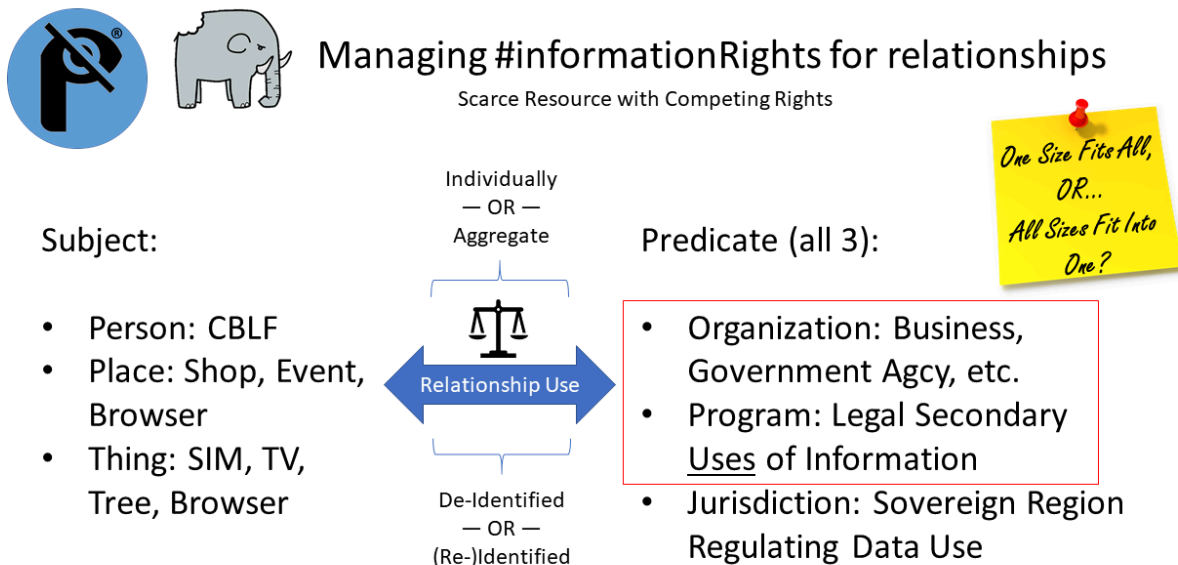
This election, then, can reside at the intersection of the three components of the Predicate (Organization, Program, and Jurisdiction) as defined by an address: X, Y, and Z. Therefore, a collection of consents can be combined neatly, without collision or conflict, into a 3D bitmap that can be utilized by any programming language using native libraries and can be applied to any data use sentence from transactional on one end of the technical spectrum, all the way to Big Data on the other end, and utilize all types of data from anonymous to re-identified data.

This 3D cube can then be related to any Subject as a stand-alone copy that can then be updated over time by the Subject (or authorized agent) and dated/versioned off. This supports historical use of data over time against historical defaults per various jurisdictions.

The goal of Project Falcon is to normalize and standardize the approach for defining, updating, consuming, and distributing this bitmap, and by doing so, providing a new source of record for all relying systems, thereby providing the missing end-to-end consent understanding for significant uses of information (data).

To this end, the project has already built a working version of the CNS, deployed to AWS and Azure, and onboarding early affiliates and users through a beta program by request: betaProgram@privacyco-op.com.

1. The Grammar of Consent



On the left side, the subject, we must constrain a request to one single object--a person, place or thing. Some people get confused on what we mean by a person, so we have adopted the CBLF (Carbon-Based Life Form) monicar. The goal here is to stretch your thinking beyond a

browser or a device. In today's world, trees put out data that can be used by various relying parties using tech such as Augmented Reality. Consent applies to all.

Therefore, this solution starts with a subject that can be all of these things and more... but only one at a time.

On the right side, the predicate, must be a combination of at least one of the following: x=organization, y=program, and z=jurisdiction. We will show later in the spec how you can have more than one of any of these things in request that can be singular, compound, complex, or even compound/complex. Spoiler alert: we use an array of x+y+z.

It's important to note that the relationship of the subject to a given predicate can reference, and must support, data that can be individual, aggregate, identified, de-identified, or re-identified.

To examine how we are going to build the SSC DNA Bitmap, let's start with only the X+Y and save the Z for a little bit later.

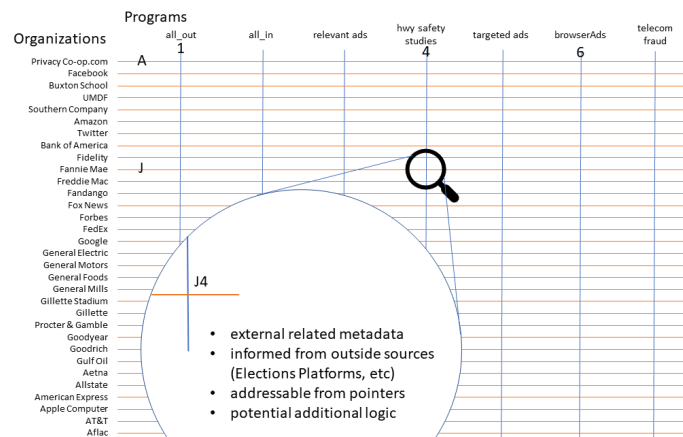
2. Building the Self Sovereign Consent (SSC) DNA Bitmap

Predicate Grid Organizations	Programs (Legal Secondary Uses of Information)						
	all_out	all_in	relevant ads	hwy safety studies	targeted ads	browerAds	telecom fraud
Privacy Co-op.com	A	1				6	
Facebook							
Buxton School							
UNDP							
Southern Company							
Amazon							
Twitter					0		
Bank of America							
Fidelity	J	0	0				
Fannie Mae							
Freddie Mac							
Fandango							
Fox News							
Forbes							
FedEx							1
Google							
General Electric	R			1			
General Motors							
General Foods							
General Mills							
Gillette Stadium							
Gillette							
Procter & Gamble							
Goodyear							
Goodrich							
Gulf Oil							
Aetna							
Allstate							
American Express							
Apple Computer	AE		1				
AT&T							
Aflac							
Albertson's							

Default values from initial Regs
Adopting businesses change by Policy
Based on Regs + Privacy Policy
User changes by
business, program, intersection
Users can All In/Out change

The initial grid is easy to approach. The goal is to distil down a consent to use information for secondary purposes to a single bit that resides at an X+Y coordinate.

Project Falcon is agnostic to the work leading up to the intersection value as well as what happens with it afterward. However, the subject of these specs largely focus on the CRUD for the value (Create, Read, Update, and Delete).



There's room for improving the CRUD requirements, but this effort will attempt to descope more external requirements. The primary objective is to identify interfaces for the CRUD to make the SSC DNA Bitmap extremely usable for a variety of platforms, applications, transactions and massive data use.

Work can be captured here for ideas on a standard way to approach the organization ledger (X), which can leverage some existing capabilities such as IAB's TCF, but any additional functionality brought with it will be de-scoped. Primarily, we are interested in the following user stories:

- A. As a company authority, I would like to create with AuthN+AuthZ my company uniquely in the SSC DNA Bitmap so that we can manage our own default values for the various programs listed, and later leverage the values set on that same record across multiple bitmaps for our customers in order to help keep us compliant with various regulations.
- B. As a relying party using data for a secondary purpose, I want to be able to addressably request permission to use data for purpose Y from a record for company X in order to help us honor a subject's (person, place, or thing) elections and to help keep us compliant with various regulations.

No delete functionality is planned for the X axis at this time. A dead company results in a permanent record for many ongoing purposes including nonrepudiation.

a) createOrg(name: string, uniqueID: string), Goal: add a new company, Response: {result: number = -1, 0, 1; err?: ERROR}

b) consent(bitmapName: string, intersection: string[XYZ]), Goal: understand consent for specified use, Response: {result: number = -1, 0, 1; err?: ERROR}

Work can be captured here for ideas on a standard way to approach the program ledger (Y), which can leverage some existing capabilities such as existing Consent Management Platforms,

but any additional functionality brought with them will be de-scoped. Primarily, we are interested in the following user stories:

- C. As a legal authority, I would like to create with AuthN+AuthZ my own unique program (legal secondary use of data) in the SSC DNA Bitmap so that we can manage our own default values to be applied as foundational defaults to all organizations (X), knowing that they can later update as they see fit.
- D. (Initially a replica of B) As a relying party using data for a secondary purpose, I want to be able to addressably request permission to use data for purpose Y from a record X in order to help us honor a subject's (person, place, or thing) elections and to help keep us compliant with various regulations in order to help keep us compliant with various regulations.

No delete functionality is planned for the Y axis at this time. A dead program results in a permanent column for many ongoing purposes including nonrepudiation.

c) createProg(name: string, uniqueID: string), Goal: add a new Program (legal use of data), Response: {result: number = -1, 0, 1; err?: ERROR}

With this approach, we honor the ongoing integrity of a growing upper-right quadrant and can depend on its consistency long term.

Organizations	Programs						telecom fraud
	all_out	all_in	relevant ads	hwy safety studies	targeted ads	browserAds	
Privacy Co-op.com	A	1		4		6	
Facebook							
Buxton School							
UMDF							
Southern Company							
Amazon							
Twitter							
Bank of America							
Fidelity							
Fannie Mae	J						
Freddie Mac							
Fandango							
Fox News							
Forbes							
FedEx							
Google							
General Electric							
General Motors	R						
General Foods							
General Mills							
Gillette Stadium							
Gillette							
Procter & Gamble							
Goodyear							
Goodrich							
Gulf Oil							
Aetna							
Allstate							
American Express							
Apple Computer							
AT&T							
AT&T							
Albertson's							

In this way, data collected 50 years from now can be used in compliance with a program that exists today, and data collected today can be used in compliance with a program that gets added 50 years from now.

[illegible]

With the default, “best-guess” effort of the first grid placed on the floor, we can consider it layer zero (0), and make a copy of it and add that as layer one (1).

© copyright 2020, Privacy Co-op - Distribute with Permission Only

That new layer can now be assigned to a jurisdiction (Z). For example, GDPR

The diagram illustrates the GDPR impact on data processing, showing a 3D representation of a table with 'Businesses' as rows and 'Programs' as columns. The 'Programs' are categorized into 'at rest', 'relevant ads', 'heavy safety warnings', 'targeted ads', and 'systemic tracking'. The 'Businesses' listed are Privacy Co-op.com, Facebook, Amazon, Twitter, Bank of America, Fidelity, Facebook Messenger, Freddie Mac, Fandango, Fox News, Forbes, FedEx, Google, General Electric, General Motors, General Foods, General Mills, and Gillette Stadium. The diagram is divided into two horizontal sections: 'Z = 1' (top) and 'Z = 0' (bottom). The 'Z = 1' section is labeled 'GDPR' and shows a green line for 'at rest' and a red line for 'relevant ads'. The 'Z = 0' section is labeled 'Default' and shows a green line for 'at rest' and a red line for 'relevant ads'. The 'Z = 1' section is shaded blue, and the 'Z = 0' section is shaded yellow.

Businesses	Programs at rest	relevant ads	heavy safety warnings	targeted ads	systemic tracking
Privacy Co-op.com	A				
Facebook					
Amazon					
Twitter					
Bank of America	J				
Fidelity					
Facebook Messenger					
Freddie Mac					
Fandango					
Fox News					
Forbes					
FedEx					
Google	R				
General Electric					
General Motors					
General Foods					
General Mills					
Gillette Stadium					
Gillette					

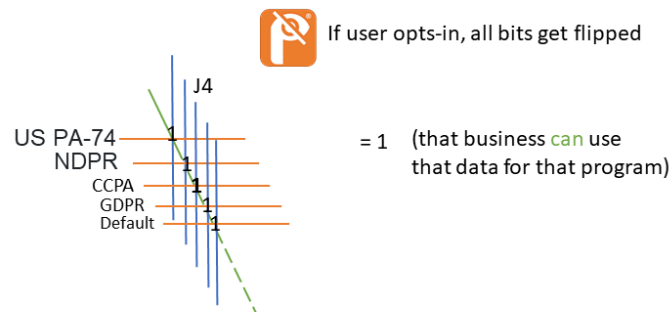

© copyright 2020, Privacy Co-op - Distribute with Permission Only

Work can be captured here for ideas on a standard way to approach the jurisdiction layers (Z), which can leverage some existing capabilities such as WorldLii.org, but any additional functionality brought with it will be de-scoped. Primarily, we are interested in the following user stories:


- E. As a nation state or sovereign government authority, I would like to create with AuthN+AuthZ my own unique jurisdiction layer in the SSC DNA Bitmap so that we can manage our own default values to be applied as foundational defaults to all organizations (X) and all programs (Y) knowing that they organizations can later update as they see fit so that we can provide observable guidance for our guardrails of data usage when subject to our laws.
- F. As a relying party using data for a secondary purpose, I want to be able to addressably request permission to use data for purpose Y from/by organization X by jurisdiction Z in order to help us honor a subject's (person, place, or thing) elections and to help keep us compliant with various regulations in order to help keep us compliant with various regulations.


Place, or Thing associated with this particular bitmap cast a consent election (this idea to be covered soon in an upcoming section). To be candid, we don't care. All we are after is to support all those required user stories, and then to support the Relaying Party user story that wants an answer other than an acknowledgement, null, or error. Even if it's lazy-provisioned at the moment of a first request, it will return a "0 or 1".

With this 3D grid in place, we can support any number of related configurations. Here are some samples...



This happens to be the way Privacy Co-op handles it.





Work can be captured here for ideas on a standard way to approach CRUD for XYZ intersection, which can leverage some existing capabilities values in various stores of record or directly impacted by things like the CCPA requirement for an “Opt-Out” button on privacy policies, but any additional functionality brought from those will be de-scoped. Primarily, we are interested in the following user stories:

- G. As a subject (person, place, or thing), I want a simple way for my user interface to set the value for my election for a given X+Y+Z based upon my election and carry through to all applicable subsequent uses of that data so that I can have peace of mind that my election matters in a subsequent reliable and verifiable way.
- H. As an authorized agent for a subject (person, place, or thing), I want a simple way for my user interface to set the value for my subject's election across all Z layers for any given X+Y based upon their election and carry through to all applicable subsequent uses of that data so that I or the subject I represent can have peace of mind that their election and my stewardship matters in a subsequent reliable and verifiable way.

e) `updateConsent(bitmapName: string, intersection: string[XYZ], vals: number[0/1])`,
 Goal: update consent for specified use, Response: {result: number = -1, 0, 1; err?: ERROR}

4. Compression

Every programming language we have been able to identify has a native bitmap library of functions. These libraries are, by design, the most efficient functions in that language, as this is how the bare metal of the hardware processors work. This is literally no more efficient approach to information representation and compression than this.

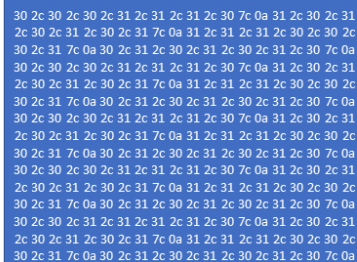
Project Falcon is not attempting to pick any one language or OS, but rather set forth a component approach that can be applied or federated with by all. In the following example, we are visually representing it as hex, and the Privacy Co-op already does a variation of this, but this is illustrative and for some uses such as Python, other approaches to visualize and treat the data is fine.

SSC DNA Metadata on bitmap

Located externally, or innate

- Storage Address
- Assigned to:
 - 1) a subject
 - 2) an attribute
- Can be partial (lazy provisioned)
- Can be related to other bitmaps
- Copied/dated at last change
- Can identify deltas
- Can be cached/pushed/reconciled (think architecture) – supports lightweight Geo-Dist/Geo-Redund

SSC DNA



```
30 2c 30 2c 30 2c 31 2c 31 2c 31 2c 30 7c 0a 31 2c 30 2c 31
2c 30 2c 31 2c 30 2c 31 7c 0a 31 2c 31 2c 30 2c 30 2c
30 2c 31 7c 0a 30 2c 31 2c 30 2c 31 2c 30 2c 31 2c 30 7c 0a
30 2c 30 2c 30 2c 31 2c 31 2c 31 2c 30 7c 0a 31 2c 30 2c 31
2c 30 2c 31 2c 30 2c 31 7c 0a 31 2c 31 2c 31 2c 30 2c 30 2c
30 2c 31 7c 0a 30 2c 31 2c 30 2c 31 2c 30 2c 31 2c 30 7c 0a
30 2c 30 2c 30 2c 31 2c 31 2c 31 2c 30 7c 0a 31 2c 30 2c 31
2c 30 2c 31 2c 30 2c 31 7c 0a 31 2c 31 2c 31 2c 30 2c 30 2c
30 2c 31 7c 0a 30 2c 31 2c 30 2c 31 2c 30 2c 31 2c 30 7c 0a
30 2c 30 2c 31 2c 31 2c 31 2c 31 2c 30 7c 0a 31 2c 30 2c 31
2c 30 2c 31 2c 30 2c 31 7c 0a 31 2c 31 2c 31 2c 30 2c 30 2c
30 2c 31 7c 0a 30 2c 31 2c 30 2c 31 2c 30 2c 31 2c 30 7c 0a
```


- L. As an SSI (Self Sovereign Identity) developer, after I have validated a person's claims so that they can subsequently be verified by others, I want to write a function to associate a customer's work persona consent elections with their Company Only Equipment device, and while that relationship is stored in our data store and the ledger record in our blockchain, I'd like a consistent pointer that I can rely upon in the future verifications of my validation so that I can support our business plan.
- M. As a telco developer of our welcome app, I want to ask permission to onboard a customer to various opt-in and opt-out products even before they have created their login credentials, as their device will register to the 5G gateway before that registration step, and I don't want to handle a null condition when a 0 or a 1 would be fine in order to support our shorter welcome cycle, and I want them to be able to later review those values and make changes.

Use the existing [consent\(\)](#) function.

6. Lazy Provisioning

Lazy provisioning by design is light on "should we do this?" post-authorization. This may result in multiple bitmaps representing a single subject by different identifiers, and that's OK. The trade-off for the potential need for a subsequent [Consolidating Bitmaps](#) functionality is balanced well by the most-efficient responsiveness gained by immediate consent resolution.

Subject



From right to left, this diagram visualized lazy provisioning. When a request came in on 3/15 for a bitmap that didn't exist, the system touched the data store with the germain intersections to the request as it asked for a value. In this case, it responded with a 0.

Three months later, the user became aware of the data use and reviewed a more complete set of elections for that package of elections and made a different choice on one.

Over the next 16 months, various products asked for the user's consent.

On 2/14/20, an authorized entity added a new program (secondary use of data) to the default bitmap, which propagated across all root jurisdictions and then across all existing bitmaps. While this propagation took place, if any requests were made for that program to any bitmaps that had not been so updated, it would have lazy provisioned the value from the default set. If a request is made that isn't in the default set, with proper AuthN+AuthZ, it will lazy provision the default set and then lazy provision the specified addressed bitmap.

- N. As a consumer driving a rental car across Europe, I'd like to see a history of the various consent elections I encountered when I get home so that I can understand what companies accessed and used my data even if I was not made aware of it at the time, so that I can follow up with my express consent afterward.
- O. As a Big Data engineer, I want to be able to pull and cache copies of SSC DNA Bitmaps by device IDs so that I can include them in my next Hadoop map reduction run and I don't want my code to handle gaps in the bitmap array due to users that haven't provided affirmative express consent yet, so that I can run my insights daily for a growing list of devices over time without having to refit my approach in order to meet our business goals.

f) consentHistory([bitmapName]: string[]), Goal: understand the history of a SSC DNA Bitmap, Response: {result: [bitmapName: string[{dateModified: Date, SSCDNAB: bitmap}]; err?: ERROR}

g) consentBitmap([bitmapName]: string[]), Goal: get the full bitmap in its current state for a specified SSC DNA Bitmap, Response: {result: [bitmap]; err?: ERROR}

Living Logs

As these lazy provisioning events occur, the facts per subject are easily understood and reportable. It is the aim of early development prototypes for the CNS to turn logging into a publish and subscribe service.

Verified Subject owners, or their authorized agents, could subscribe to all events registered for their various Subjects. In this way, simple logging would become transparent knowledge, placing urgency on judicious data use by organizations.

Moreover, verified auditors might be given anonymous, random logging by intersection, x, y, or z axis to get aggregate and anonymous sampling information. For example, they may be interested in a certain business and how they use data, or in a certain use of data many businesses may be participating in, or behavior across a certain jurisdiction.

Instead of logs being viewed as dead artifacts, they would become living logs supplying information to reasonably verified parties, and even trigger for other behavior.

For example, an authorized agent for a Subject with the technical wherewithal could subscribe to a certain intersection and upon reaching some agreed upon amount, could change the value of the election on behalf of the Subject. For example, a consumer may agree to 10 unsolicited direct advertisement emails a week, but no more. After the 10th request, the agent could change a 1 to a 0. This means that frequency capping could conceivably move directly into the control of each Subject.

Frequency Capping

In production today, we have added Frequency Capping that is set by the subject or the subject's Authorized Agent. Our early use case for this is not intended to be the only use of this functionality, but it has turned out to be very popular with beta testers. A subject or their agent can set a frequency of use for an affirmative express consented and validated end-point, currently supporting email, postal address, and mobile phone numbers. For example, you can agree to receive up to 300 transactions/messages per month and then the 1 gets flipped to a 0. In this case, we are leveraging existing affiliates through AWS to pay per transaction to the subject's authorized agent. This "royalty" is then paid to the subject, successfully parsing publicity law from privacy law and finding a phase-space of opportunity in between.

We believe royalties may ultimately fund the global commons intended for this CNS.

7. Consistent Addressability

A word about Transports and Signalling...

There are low level protocols called TCP and UDP. What these protocols do is allow sending arbitrary data from one point in the network to another.

In our browsers, transport protocols that allow sending arbitrary data from both the browser to the web server and vice versa include XHR, SSE and Websocket.

Signaling protocols go one step higher. They express some mechanism – a way to tell the other end something. In our case it can be the need to gain consent to open a call, using availability, and identification. To that end, we can either invent a protocol to do that or use a predefined protocol – something that people have already agreed upon in the past. This protocol is a signaling protocol.

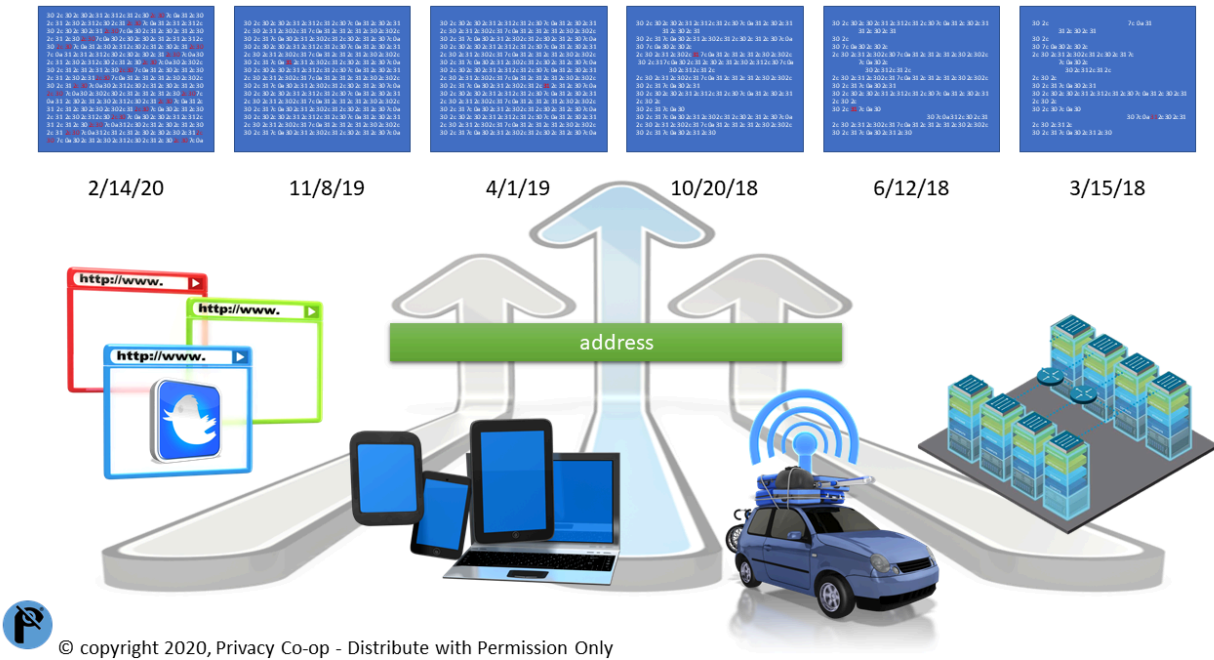
The predefined ones? H.323, SIP and XMPP. There are more, but these are the main ones used for things like VoIP and instant messaging.

The SIP signaling protocol uses TCP or UDP for its transport. For WebRTC, there is an adaptation of SIP over Websocket.

Project Falcon is not attempting to build a signalling solution, nor is it attempting to build a transport solution, but the bitmaps can reside in data stores accessible by potentially any one of these or adjacent to them.

While this spec assumes a TCP/IP use case, and specifically the application layer of HTTP, it's important that the ultimate solutions are mindful of the likely need to support other protocols.

Subject to Predicate

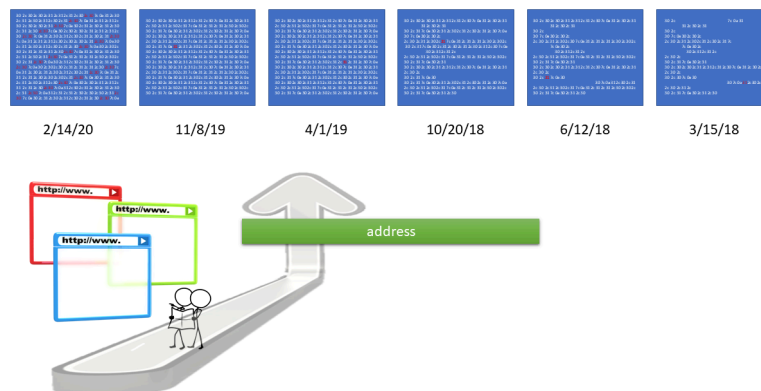


It's critical that the SSC DNA Bitmap is broadly adopted. This means that it must support a single transaction in a millisecond while it must also support Big Data map reductions. It must work while residing in a decentralized framework or cashed by a pull request.

With these in place...

Browsers

Browsers can leverage whatever identity they wish or wish not. We have already seen how an unprovisioned bitmap for a specified identity can be lazy provisioned, but by the same token, when no identifier is provided, this solution will simply respond using the default bitmap. In this way, each browser manufacturer is free to decide which root address, subject ID, and [XYZ] coordinates they would like to use as they build a better browser than their competition.



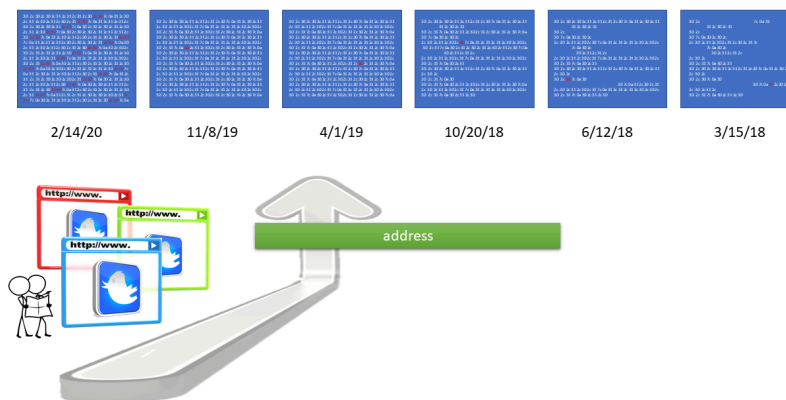
Because the XYZ is defined by the people that are actually using the data and are accountable for any potential exploitation, then they are the ones who should be empowered with the ability to ask the right questions for them. For example, in the diagram above, the developer that is deciding on the right XYZ question to ask for their browser is consulting with his company's legal counsel so that they can make the choices that are right for them.

- P. As a browser developer, I want to distinguish our code by supporting natively both GDPR and CCPA consent positions for content allowance by using an internal anonymous identifier each installed instance will have in order to increase adoption by my target users.

Use the array parameter of the existing [consent\(\)](#) function.

Websites

Meanwhile, static app developers using some approach such as Angular code compiled and stored in AWS S3 buckets can, within that same browser, either use the features offered by that browser, which in many cases requires different code logic per, or they can elect to develop their own approach using the exact same resources the browser do.



Just like the browser example above, it's incumbent on this app developer to work with their own legal counsel to decide which XYZ question is the right one for them to set and ask.

- Q. As a website developer, I'm tired of modifying my code over and over based on the browser in which we may be found, and I want a clear way to tap into a public commons solution for consent that I don't have to worry about provisioning, and not based upon the current logged-in state, so that my time to market and updates are faster and more profitable.

Use existing [consent\(\)](#) function.

Devices

Devices have always brought additional Identity complexities, namely because they can be shared. We made the choice that subjects can be a person, place, or thing -- but only one. This frees up different device manufacturers, transport providers, or carriers to decide on their own identifiers.

Work can be captured here for ideas on a standard way to approach the identifier ledger, which can leverage some existing capabilities such as various SKU systems, but any additional functionality brought with it will be de-scoped. Primarily, we are interested in the following user stories:

- R. As a developer for a mobile provider, I want to use geo-fenced proximity to direct the pizza delivery person to them while they are roaming in a different country, while staying compliant with a nexus of legal regulations in different jurisdictions.

Use the array parameter, specifically multiple Z values, with the existing [consent\(\)](#) function.


IOT

This is a great example of potentially different signalling or transports using this solution. For example, Zigbee and Z-Wave are both mesh networks – meaning the signals can hop from gadget to gadget around the home and each device or sensor doesn't need to connect to Wi-Fi – but they usually have a central hub which connects to the internet. It's conceivable that while the hub may request a pull of the bitmap, there may be a need for a device to provide some functionality, but only if there is proper consent.

- S. As an IOT owner, I want my automated freezer to dispense a fudgesicle to each of my children, but only if that requestor has my consent, in order for me to maximize my household budget and help them manage their nutrition health.

Use the existing [consent\(\)](#) function.

Big Data



In Big Data, there are many time costs that are just accepted as “costs of doing business”.

1. As you hydrate a data lake, you must scrub (remove) opted-out records for the next insights run--let's say it's for targeted advertising insights. This is processor intensive but worse, takes a lot of time as for each record you code identifies who owns the rights to it, does some form of lookup for consent, and takes the time to erase that record or otherwise not use it.
2. A map reduction is executed using logic rules in the Hadoop lake.
3. After that run, your next run may be for a different program like relevant advertising. Because you want to maximize the profit, you don't want to re-use the same data because it was scrubbed for a different reason. You likely drain the lake and re-hydrate, this time scrubbing for a different program (legal use of data).



These costs are staggering. It can take 3 times longer or more to set up a run than actually doing the map reduction. But with SSC DNA Bitmaps and their consistent addressability, the above process can look like this...

1. As you hydrate the lake for targeted advertising insights, for each record, identify who owns the rights to it, and prepend that record with a pointer to the consistently addressable SSC DNA Bitmap and write it to the disc.
2. A map reduction is executed using the logic rules in Hadoop lake, which include the pointer resolution as a mathematical additional feature to all the other features.

- After that run, your next run may be for a different program. Just run it and include the different XYZ addresses in the Hadoop map reduction.

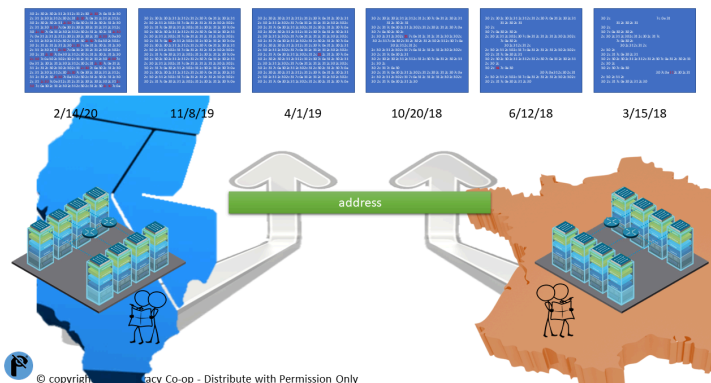
An independent test of this approach 5 years ago in an Israeli foundry found that in step 1, they cut the write to disc time by 60%. In step 2, they added 4 minutes and 13 seconds to a 4 hour map reduction of over 6 petabytes of data. And in step 3, they saved an average of 12 hours for each subsequent run...forever.

User Story -- see story O. above for Big Data.



© copyright 2020, Privacy Co-op - Distribute with Permission Only

Just like the above examples for browsers and apps, it's important that device manufacturers, IOT developers, and Big Data architects all consult with their own legal counsel to make sure they're asking the right questions.



© copyright 2020, Privacy Co-op - Distribute with Permission Only

In this way, businesses in France will likely ask a different set of XYZ questions than businesses in California. This is the power that Project Falcon brings Consent Enforcement.

3rd Party Use of Data

How valuable would the above Big Data solutioning be to a large enterprise? Well...it could save millions and at the same time help them immediately become compliant with all global regulations. But now when they sell data or insights to a third party, will they delete the pointers?

Why would they do that? The pointers add value.

Would a 3rd party delete the pointers? Why would they do that? It would cost them more to do the work and then even more to condition the data for their own needs. In fact..it would be to their advantage to not only use the pointers, but to add their own segment so they could specify their own XYZ for their own legal uses of the same data.

What would that data then look like? It would begin to speak for itself.

What would the bitmaps tell us about the usage? As much as the lazy provisioning and the associated metadata could tell us.

External Address

For now, until W3C finalizes their approach, let's assume that the pointer will include a minimal byt.ly type URL with assumed prefix info (<https://etc>). Keep in mind, other signalling and other transport needs will require pointers to include other protocol approaches. Therefore, this segment of the pointer payload will likely need to be an array. ['byt.ly/ssc/'].

SSC DNA Bitmap Address

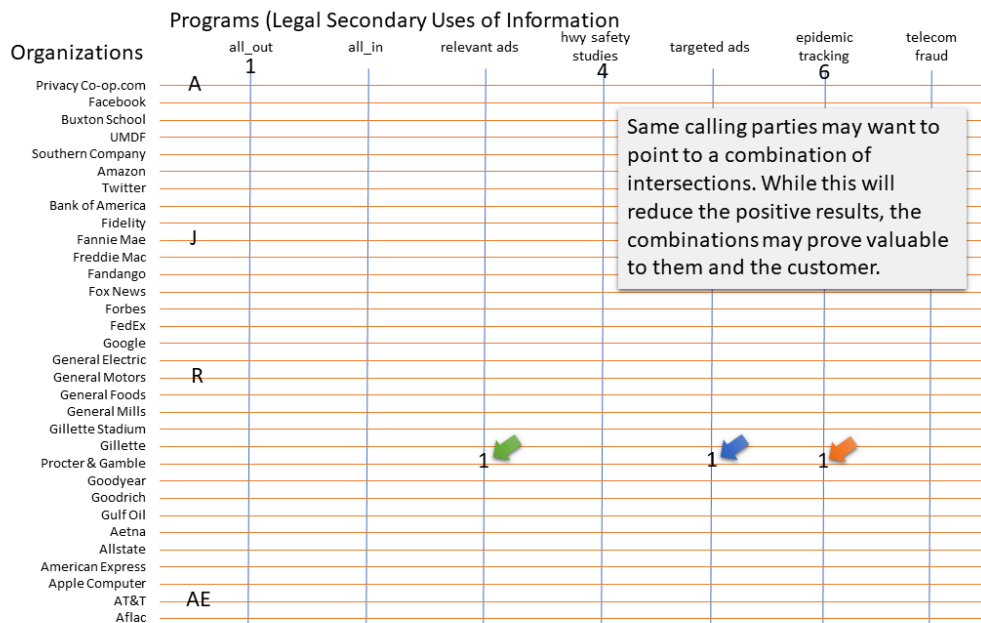
Once arriving at the service, an optional subject identifier can further point to a specific bitmap contained within. If not provided, the default SSC DNA Bitmap will be used. Because this resolution requires a singular subject, for now, the assumption is that this segment of the pointer payload will likely be a string. ['byt.ly/ssc/'].'+12055551212023'

Internal Address

Once the process arrives at the designated bitmap, the X+Y+Z coordinate is used. Because it may request more than one intersection, this segment of the pointer payload will likely be an array. ['byt.ly/ssc/'].'+12055551212023'.[XYZ].

Compound/Complex use of Data

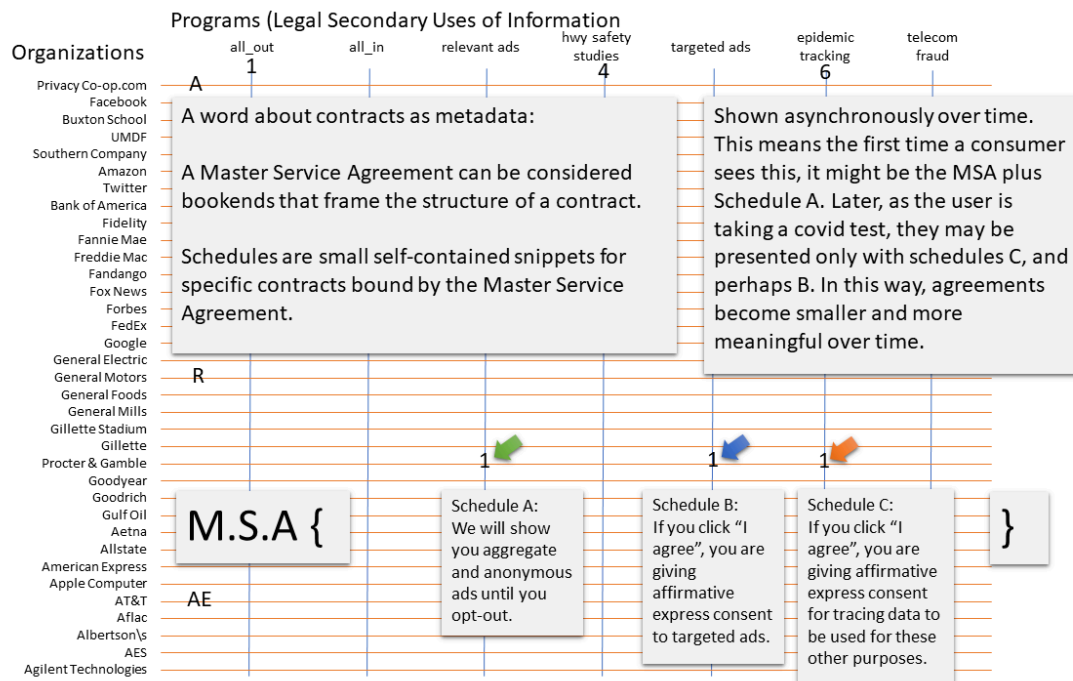
Let's say that a company wants to do something that requires a combination of multiple companies, and/or programs, and/or jurisdictions. For this example, let's say it's a medical company wanting to use data from Covid Tracing to offer relevant or even targeted advertising.



In this case, the XYZ pointer payload segment might look something like this: [X1Y1Z1, X1Y2Z1]. You'll notice the only change is the second Y. But the same approach could be used for different jurisdictions or even companies. Now, it's expected the resulting number of 1's would go down, but the legal and appropriate use of the data may be so valuable that it would warrant such a reduction in pool size.

- T. As a Theme Park operator, I want to combine data from a child's smart bracelet that we provide them with their guardian's cell phone info so that we can tailor their experience and suggest the rides with the shorter lines so that we can improve our customer satisfaction numbers.

Contract Snippets



A word about contracts as metadata: A Master Service Agreement (MSA) is a foundational contract that can be considered bookends that frame the structure of a contract, which may contain one or many schedules. Schedules are small self-contained snippets of language that are themselves small specific contracts bound by the Master Service Agreement in which they are contained.

Each intersection of Organization and Program can be assigned to a small schedule in any/all pertinent languages, tailored to specific Jurisdictions.

This MSA legal structure lends itself well to this Self Sovereign Consent DNA Bitmap, as an organization can establish a Master Service Agreement with any subject and then later add, present and acquire schedules for consent elections as they become needed.

For example, a company such as Procter & Gamble in the above illustration may initiate a relationship with a patient by providing them with an MSA and the Schedule A from the illustration so that they can use informed consent to show that patient aggregate and anonymous ads until the subject opts-out. Later when an epidemic hits, and perhaps the organization is contracted to do testing, they and the patient may find benefit in direct notification constituting advertising for products specific to the disease. At the time of the test, the patient wouldn't have to read the whole MSA or Schedule A again. They would only have to be provided Schedule B and/or C with an ability for affirmative express consent.

Can you imagine a world where huge notice and consent agreements are replaced with ever tighter, more consistent legal language across all organizations with application to specific jurisdictions? People would start reading the things presented to them and could trace them directly to the represented value at any time.

Duty of Care - a New Standard

As time passes, these initially bespoke schedules will be copied and reused, normalized. Some will address contractual duties, others criminal duties. This makes it more likely that criminal and contractual duties language will come closer and closer together, ultimately forming a commons of risks.

Masks (bitmasks)

Since the underlying tech will be bitmaps, that means we instantly have “bitmasks” at our disposal. A bitmask is data that is used for bitwise operations, particularly in a bit field. Using a mask, multiple bits in a byte, nibble, word etc. can be set either on, off or inverted from on to off (or vice versa) in a single bitwise operation. Bitmasks can, and do serve many purpose from the platform standpoint.

Locking Masks

A locking mask (or block) can prevent bits from being changed. This could be termed, “immutable”. While it may not be likely that any Program (legal use of data) ever be considered immutable across ALL Jurisdictions, it certainly may be considered such within a jurisdiction. For example, in the US, the FCC makes provisions for telcos and other regulated industries to use data for fraud. In fact, it’s considered a duty. Within its jurisdiction, the US may determine that a subject may not change the cast election for fraud from a “1” to a “0”. By using a bitmask applied to that jurisdiction, the Program of Fraud could be so blocked. This would require no (0, zero) additional logic to be introduced, and keep the performance of the CNS at its most efficient.

As you can see, masks play an essential and even critical enforcement and performance role in the overall implementation of the CNS.

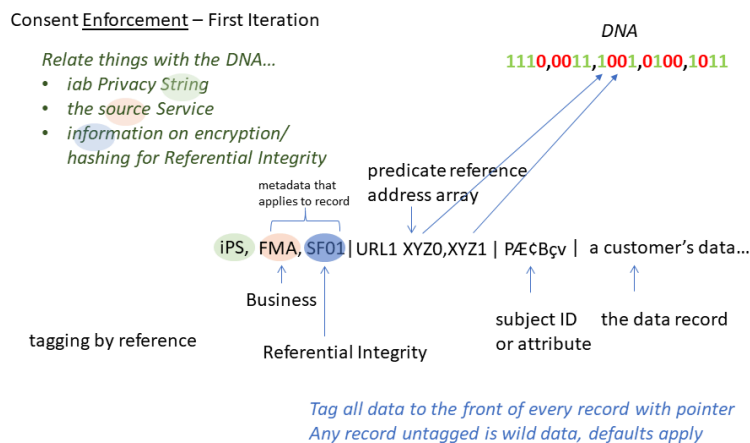
8. The Pointer

We have already addressed the pointer payload segments above starting in the [External Address](#) section, but there is so much more we can include here. Let’s dig in...

There are some businesses that own patents on various uses of pointers, and while there's no intellectual property contained in Project Falcon (trust me... many attorneys have culled this approach for anything they could find), you can't get a patent on bitmaps, pointers, or consent. However, there are some very clever technical implementations that further define ways to manipulate functionality of things like pointers.

This is the perfect time to make sure we support these innovations and invite the owners to help shape this new standard for the Duty of Care. If this approach can be leveraged by their specific product and further help their customers - super!

The writers of this document are aware of at least three companies that have patents for various additional features that very nicely bolts on to Project Falcon as value adds. They have been invited to co-write these standards and have shown positive feedback towards that as they consider joining.



Prependability

While developing the standards for the pointer, one requirement is that it remains prependable to various types of data records.

Pointer Resolution Rules

We have already covered the details in [Internal Address](#) and [Compound/Complex Use of Data](#)

Referential Integrity

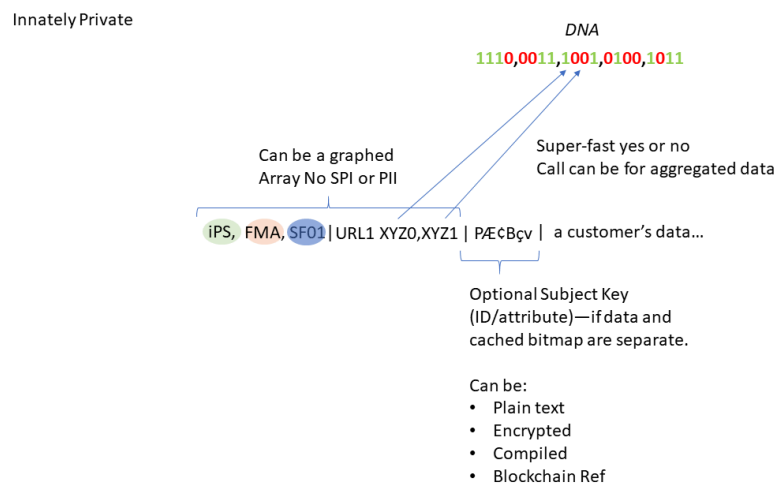
In the section of the pointer that deals with metadata that applies to the record, there's an opportunity for the business that is co-managing the information rights for their customers to include valuable resources such as security information or perhaps referential integrity data.

Organization Identifier

Work can be captured here for ideas on a standard way to approach the unique organization identifiers, which can leverage some existing capabilities such as those used by the NY Stock Exchange, but any additional functionality brought with it will be de-scoped.

Privacy String

Work can be captured here for ideas on a standard way to approach additional privacy information, which can leverage some existing capabilities such as those used by the IAB Privacy String, but any additional functionality brought with it will be de-scoped. Potentially competing standards need to be considered.



Optional Subject ID

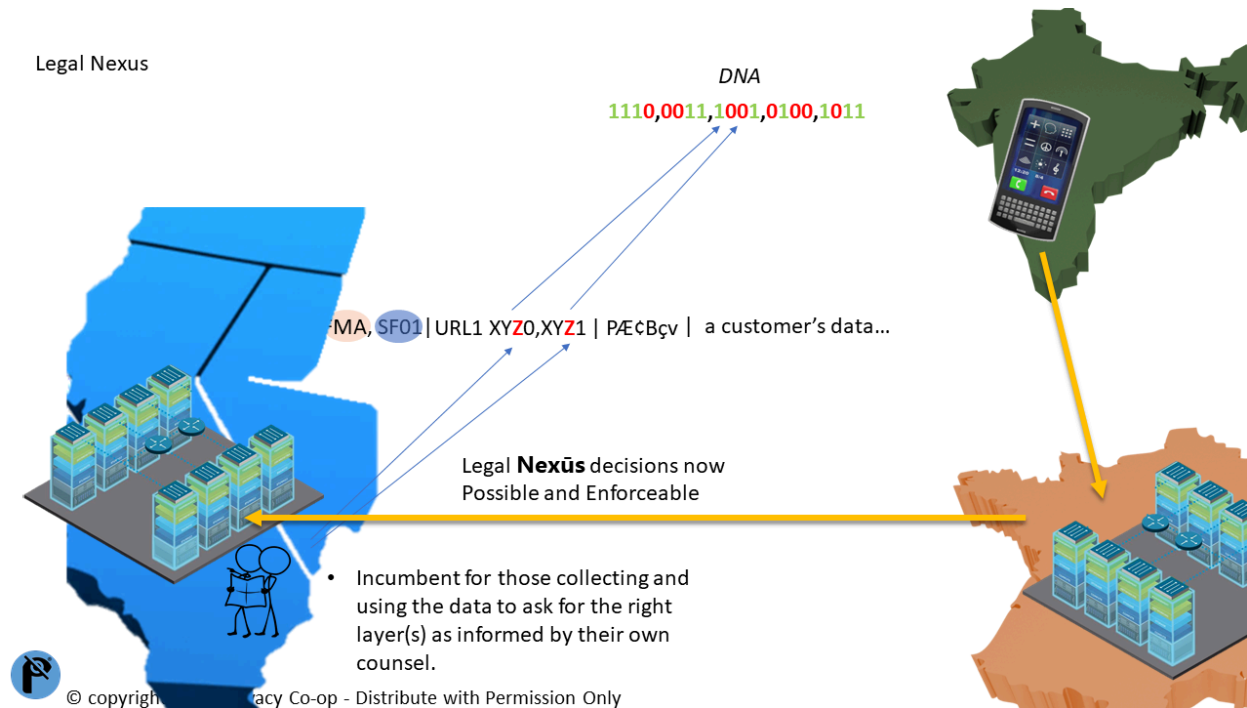
As previously stated in the [SCC DNA Bitmap Address](#), an ID specifier is optional. As these specifications mature, it's likely that the identifiers will need to support a number of formats including: plain text, encrypted, compiled, or a reference such as to a blockchain distributed ledger or verifiable container.

Interoperability

In production today, the CNS accepts any unique identifier. For convenience as well as performance, we separate the subject arrays in the JSON into email, phone, and account ID. In a recent update, we are working with custom JSON to accept specified subject types. In this trial, we are working with Unified ID 2.0 as we understand that spec. Current assumption is that the subject value will remain unique across all of these convenience groupings.

9. Pointer Resolution Rules

As a subcomponent of the pointer, the pointer payload segment can be singular, compound, complex, or compound/complex. This allows for the most powerful aspect of Project Falcon...it's flexibility for Consent Enforcement based on the decisions by those that will be held accountable.



In this diagram, a device is used in India and data is collected by a company in France, where it is packaged and sold to a company in California. Before Project Falcon, developing the legal Nexus understanding of the jurisdictions that might apply might be arbitrary and different case by case. But now the counsel in France can make informed decisions of the XYZ settings, while the counsel in California may decide on a different setting.

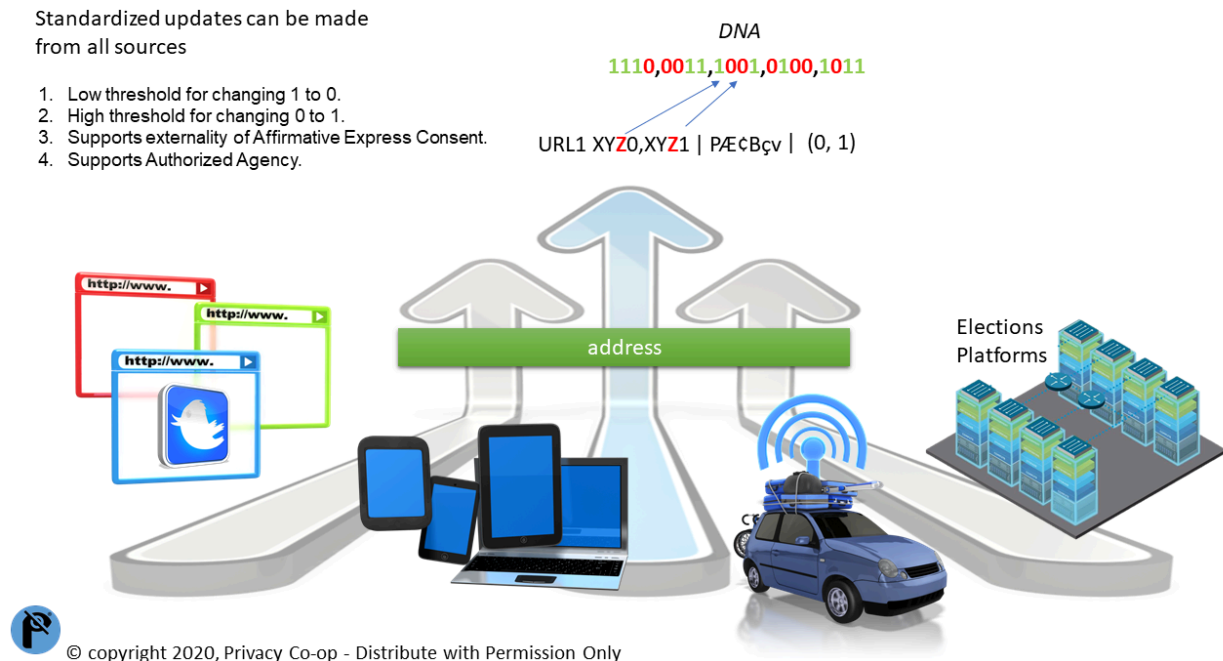
For Project Falcon, we don't have to think through all the possible crazy consent corner cases (C-4, which blows projects up) and now neither does any of the platforms and systems that may leverage the SSC DNA Bitmap. How big would it be for this component to be included in all the plans to move away from cookies?!

10. Standardized Updates

Ultimately, updates to default settings for each intersection must be allowed from the individual or authorized agent for the subject (person, place, or thing). While this must be secure, the

security layer we define must provide thresholds that are reasonably lax or strict depending on the need.

For example, if an individual is responding to a query from their smart dashboard while they are driving and they want to opt-out (effectively changing an existing 1 to a 0), then the security threshold could be reasonably low. The potential harm from an opt-in nefariously being removed are less impactful and convenience can be more highly weighted.



Conversely, changing an opt-out to an opt-in (effectively changing an existing 0 to a 1) should require a heightened AuthN + AuthZ.

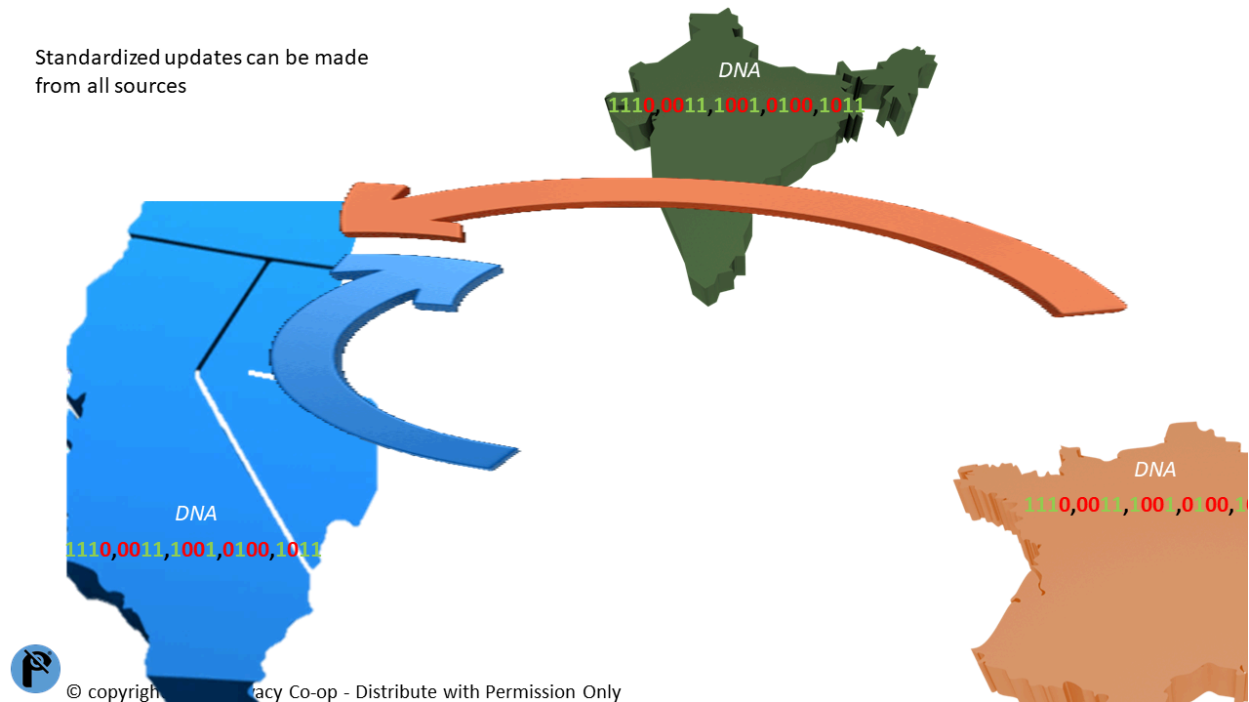
Remember, concepts like Affirmative Express Consent may be required by some laws before a 0 can be set to a 1, but the bitmap itself is not responsible for that rigor. But the security structure placed immediately around it will need to apply AuthN + AuthZ to ensure the right individual or authorized agent is making the change.

11. Re-purpose DNS

There are many open source DNS servers that already have well developed functionality that directly supports nearly all of the requirements to implement SSC DNS Bitmaps. If you think about it, you can use DNS. You do all the time. You enter something like "<https://amazon.com>" in a browser and DNS resolves the IP address for the browser.

The “name:value” pair for DNS is “domain name”: IP Address, or <https://mydomain.com:255.255.255.255>.

But you can’t change Amazon’s DNS settings. So, there’s already the right kind of security. Also, DNS is built to be Geo-redundant and Geo-distributed, effectively placing resolution at the edge of the network for the most efficient results.



Changes in one place rapidly get propagated around the world. That’s handy and much needed!

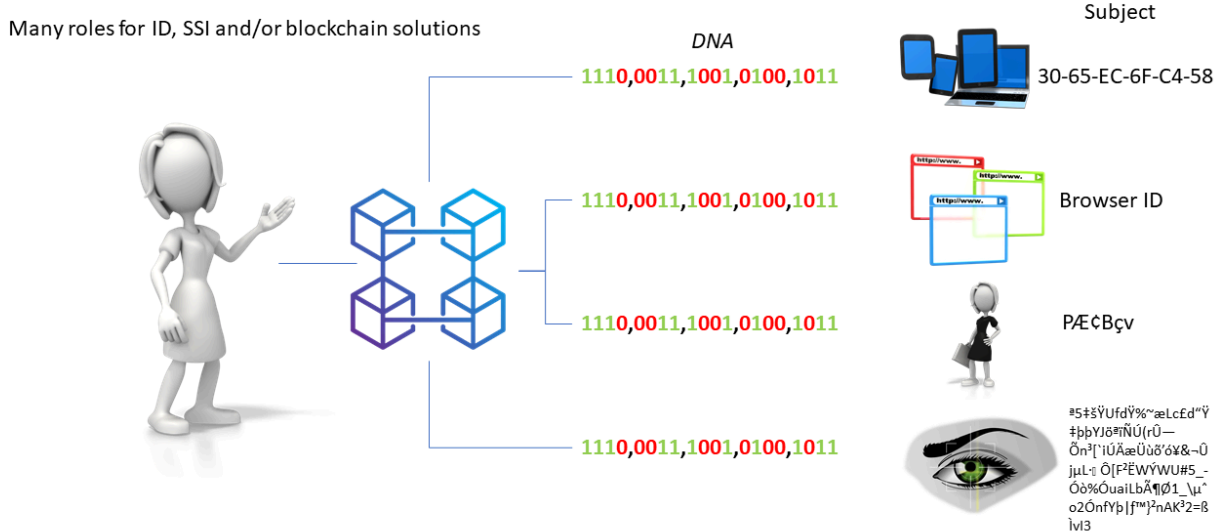
So what’s the difference? Not much. The name:value pair of “URL:IP Address” would need to change to “consent name: bitmap”. Instead of the Domain Name System (DNS), it would be the Consent Name System (CNS).

Yep... it’s pretty much like that. Everything else is very helpful.

12. Platforms Below, Products Above (the marketplace)

There are many existing players in this space. There are all sorts of Identity products, security platforms, Self Sovereign Identity frameworks, decentralized ledgers, etc. They all are very, very good at connecting the dots of Identity to attributes and properties of a person, place, or things.

Project Falcon is not good at connecting those dots. But a bitmap can be assigned to things as varied as a MAC ID, a browser attribute, a business persona, and a biometric.



None of those things are a person. All of those things are managed by a person. Below the SSC DNA Bitmap can be a layer of API that allows platforms to leverage, or even manage relationships between various bitmaps, while a layer of API above can support products that produce or consume data.

Consolidating Bitmaps

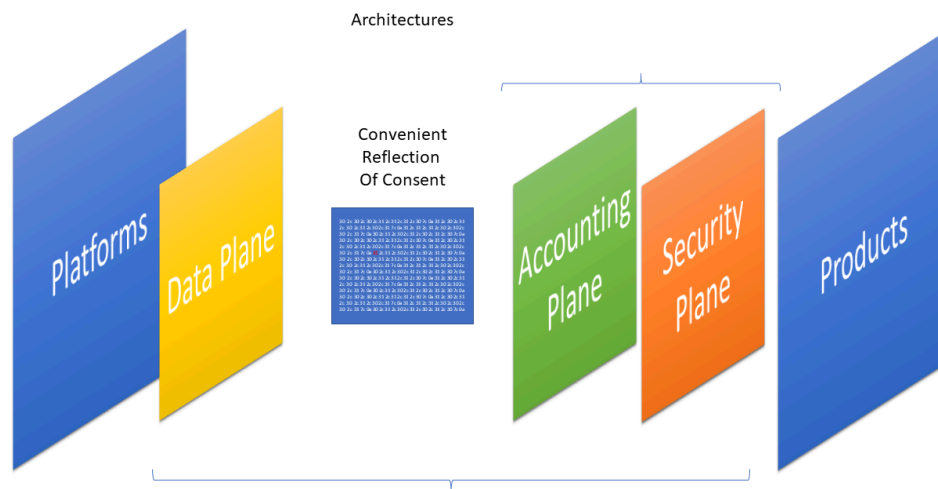
Bitmaps of the same subject, but stored in with names/addresses generated by different attributes of the same subject (like a SSN or a driver's license number), will need to remain in their states in the Consent Naming Service (CNS). However, relying parties may wish to consolidate their values for a singular answer. Bitmaps totally support this kind of combinative functionality natively in all programming languages. For this purpose, the relying party will use the existing API to pull a copy of the entire bitmap, [consentBitmap\(\)](#), for each attribute represented.

- U. As a developer for a Self Sovereign ID platform, when my software identifies two bitmaps that represent two attributes for the same subject, I want the ability to consolidate the two bitmaps to arrive at a resolution based on my own rule set in order to honor the best guess at the user's desired state.

Use the existing [consentBitmap\(\)](#) function.

13. The Disappearing Data Plane

Many businesses that maintain actual consent election values, either as a direct service provider or as a business that provides consent as a service, must maintain a similar architecture stack. Here's a simple marketecture diagram...



The bracket below shows how most of these providers must maintain a data plane. But if they had their druthers, they may jettison that layer in lieu of a common service like the SSC DNA Bitmap version of DNS called CNS.

Think of it as the Consent Name System (CNS). Until this becomes available, the concepts in this document are so simple and amenable to any technology, that there's no reason why any solution couldn't adopt and develop their own version of it and contribute to the larger evolution of a public common CNS.

The main point here is to be an assumption buster as various platforms attempt to fill in the gaps when cookies are removed from the landscape. The assumptions we feel can be removed are: Consumers can't know what 3rd Parties are using data, solutions can't adhere to all regulations, we can't support asynchronous Identity, stored data can't be used the same as transactional data. And there are more.

Project Falcon busts all of those assumptions, freeing up platforms and products to imagine broader, more exciting solutions.

The goal of Project Falcon is to normalize and standardize the approach for defining, updating, consuming, and distributing this bitmap, and by doing so, providing a new source of record for

all relying systems, thereby providing the missing end-to-end consent understanding for significant uses of information (data).

Commercial Considerations

Does the proposal require a commercial structure?

This proposal is already available in a private implementation hosted by the Privacy Co-op, who is seeking to deploy a public commons version of the same thing to support broader adoption.

If so, please describe the suggested funding and cost model.

For any authorized agent (nonprofit privacy cooperative, of which there are more than one), they can license the use of opt-ins via this model. Funding would be generated from multiple affiliates so licensing the use of data they already collect, gaining them their desired opt-ins, thus bringing them into compliance with various global privacy regulations and helping them mitigate risk inherent in secondary data use.

Additionally, how does the proposal contemplate the role of open standards - in terms of supporting industry transparency/accountability - while also supporting open innovation and market competition?

This proposal is grounded in deploying this as a commons leveraging existing open standards and open source from the DNS stack.

Critical Dependencies

Are there critical dependencies worth noting, like consumer adoption dependencies, browser dependencies, commercial resource dependencies, or other cost implications?

All additional critical dependencies will be addressed through open standard interface agreements, such as:

- AuthN+AuthZ
- Metadata
- Pre and Post Logic required by silo'd relying parties

Open Questions:

What does this not address, or require additional feedback / input from the working group?

This component is not meant to be an end-to-end solution for projects such as IAB's Rearc. However, with this component, many such proposals will stand a better chance to succeed since they will avoid crippling scope creep for crazy consent corner cases (C-4, that blows projects up). We need to work with all of them to see how they can benefit from its adoption.