Response to Request for Information
National Privacy Research Strategy
NCO for Networking and Information Technology Research & Development (NITRD)
17 October 2014

Submitted via Email to nprs@nitrd.gov

Erin Kenneally, MFS, JD
CEO and Founder – Elchemy, Inc.
8677 Villa La Jolla Dr., #1133
La Jolla, CA 92037
(858) 232-6255
erin@elchemy.org

Scott E. Coull, PhD
Research Scientist – RedJack, LLC
8484 Georgia Ave., Suite 940
Silver Spring, MD 20910
(301) 886-8785
scott.coull@redjack.com

In response to the RFI issued by the CSIA R&D SSG, we provide the following input to all four presented questions to inform the development of the National Privacy Research Strategy to guide federally-funded privacy research and provide a framework for coordinating research and development in privacy-enhancing technologies.

1. Privacy Objectives:

The broad scenario addressed by our responses is privacy and organizational data sharing. We discuss the scenario mostly as it pertains to privacy in network layer data sharing (e.g., metadata, machine communications) since it is less understood yet no less of a concern, owing largely to the Snowden revelations. Therefore, the problem and suggested solution impact any human layer domain that runs atop the network—from industrial and national security, to healthcare, social media and energy sectors. Despite the rampant liberation of data flowing from online social networking and Internet-enabled technologies, many initiatives to encourage persistent data sharing at an institutional level (whether it be government, industry and/or research) have either stalled or been 'driven under the radar' due to data privacy and confidentiality concerns.

These data sharing efforts are hindered by a lack of a sustainable model that concomitantly assesses privacy risks, articulates utility goals, chooses appropriate disclosure controls, and operationalizes those controls to mitigate risks while maintaining utility. The current state of practice encourages the application of ad-hoc policy and technical approaches that often fail to appropriately balance between privacy risks and utility of the shared data. Moreover, the inherent variability in the disclosure control process among data sharing efforts makes it difficult to re-use legal or technological infrastructure, resulting in excessive labor costs and an inability to properly audit the process. To achieve effective data sharing, a "standard"[1] organizing framework is necessary to guide any data provider to synchronously consider privacy risks and pursue utility benefits using a range of disclosure controls, without adding unreasonable costs to existing business process.

- <u>Understanding the Privacy Need and Barriers</u>.

We are currently playing a zero-sum game when it comes to achieving privacy and availing ourselves of the benefits of technology advancement. The pervasive nature of modern technology, rapid expansion of Internet based services, and collective migration of nearly all aspects of civil society to the web make it easier than ever to generate, collect, use and share information. Data shared about and from individuals helps in the development and delivery of new services and products. These benefits, however, often come with a cost to

[1] The use of "standard" contemplates both informal, de facto notions as well as formal processes giving rise to industry or regulatory standards.

© 2014 Elchemy and RedJack

individuals' privacy and thereby create the need to strike a balance between the risks and utility associated with disclosing data. This disclosure tension exists similarly at the organization level with data collected about and from computer networks and devices such as packet traces, network flow logs, or intrusion detection alerts. For organizations, the intended operational, security, legal, and economic benefits that are derived from the shared data by all data recipients must be considered in parallel with the probable risks that disclosure presents to the data provider.

Widespread sharing is tempered by real and perceived privacy risks associated with the data to all participants, although they are born primarily by the *provider* of the data. Since network and security data encompasses information about a range of human and device communications, it may contain directly sensitive information about individuals such as passwords, mailing addresses, financial transactions, and behavioral information. The data may reveal information considered sensitive to the disclosing organization itself such as intellectual property, security procedures, and business relationships restricted by non-disclosure agreements and other policy restrictions. Complications arise when the sensitive information is encoded within the statistical properties of the data itself. This can add uncertainty about its sensitivity and make it difficult to control for the risks associated with the information disclosure. While managing these privacy and confidentiality-related challenges falls on a spectrum of difficulty, all effective approaches require understanding the risks and the available methods for controlling those risks.

There is express and tacit dissatisfaction with the prevailing models for considering and managing the risks and benefits of information sharing. Discourse revolves around course-grained criteria for the disclosure of data that necessarily shoehorns data *providers* and *recipients* to accept disjointed and binary trade-offs between utility and risk – for ex- ample, can we disclose network traces to entities from China? Can we disclose flow records if we scrub all Internet Protocol Addresses? This 'state-of-the-art' reflects a demand for a more sophisticated approach to disclosing data that allows participants to address data sharing according to the reality of nuanced disclosure contexts.

Decisions by data providers to share network data are ultimately anchored in trust, the nature and extent of which is a confluence of: the capacity to satisfy relevant legal requirements; the value placed on potential benefits; and confidence that the data recipient will not increase the provider's risk of disclosing the data. What is most notable about the current state of practice is that most would-be data providers have a relatively weak understanding of these individual issues and their interplay, particularly in the context of computer network data. Risk is fueled by uncertainty about the application and interpretation of legal restrictions and obligations (e.g., regulations and privacy laws) related to network data disclosure, and exacerbated by unfamiliarity with disclosure control methods. Trust then is largely a function of ad hoc, bilateral, and interpersonal relationships between individuals within the organizations providing and receiving the data. This is substantiated by the fact that network data providers and recipients lack a scalable and transparent process by which providers– be they researchers, operators or technology vendors– can make data available for operational or research purposes in a sustainable manner. Consequently, what exists is a self-perpetuating cycle that does nothing to improve the lack of trust in organizational data sharing, instead reinforcing a risk-averse posture that precludes all but the most restrictive forms of sharing between organizations. At present most efforts related to data sharing focus on defining common data formats and exchange procedures for information interoperability, but very little work has been done to address the need for generalizable and scalable guidance that helps data providers understand and reason about these data sharing issues, thereby enabling risk-sensitive data disclosures that consider both legal constraints and utility needs.

- <u>Proposed Capabilities to Meet the Privacy Need</u>.

We propose a risk-utility model – *the Disclosure Control Framework* – tailored to help data providers understand and take action in the face of this tension, thereby enabling risk-sensitive data sharing that considers both risk constraints and utility needs. While this framework is born from deficiencies in cyber security data sharing, its underlying motivations and core components are generalizable across data sharing scenarios in other domains

such as healthcare, education and social media.

Our proposed risk-utility framework addresses this systemic deficiency by conceptualizing risk, articulating utility needs, and helping to operationalize data disclosure controls. This approach allows data providers to more efficiently manage risk and achieve desired utility in the face of legal speculation and indeterminate outcomes. The significant merit of this framework is that it arms data providers with an objective and defensible *process* for considering, designing and communicating how to balance risk and utility when making data disclosure decisions. This is the prerequisite to engendering the discourse and shared reasoning that is fundamental to raising the quality and scope of trust that drives data sharing. Importantly, the impact of sustainable and defensible organizational sharing can be an evolution of collective norms about 'risks of first impression' and the relative value of sharing, thereby lowering the real and perceived privacy barriers.

2. Assessment Capabilities:

The concepts, methods, and capabilities needed to address the Privacy Objectives discussed in Question 1 are contemplated in our Disclosure Control Framework[2]. It outlines a process for both understanding the objective factors (the "what") that define the myriad of privacy risks and utility objectives, and also, for applying them according to the subjective choices (the"how") of the participants in any *data sharing scenario*. Our framework is broken into three phases– analysis, application and assessment. Each phase in the framework discusses, in necessarily general terms, the primary considerations essential to developing a data provider's data sharing scenario: the desired utility objectives, relevant risks, options for disclosure controls, and the impact of those controls on the chosen risk and utility determinations. This framework is intended to address our normative and empirical expectations about data sharing risks and utility by embedding a process for decision-making that forces the provider to:

- consider risks drawn from authoritative sources (i.e., laws, private agreements, standards, ethics);
- alongside conscious choices about data utility;
- given a comprehensive set of options that can be tailored to its appetite for accepting certain risks and altered outcomes;
- tied together by an evaluative loopback that reinforces the evidence-based deliberateness and of the providers chosen actions.

The initial phase examines the two components around which the sharing scenario is anchored– the risks involved with sharing the data, and the intended utility for the provider and/or recipient. The input for this phase is the data that the provider wishes to make available to the recipient(s), such as network flow records, packet traces, or application logs. Risk considerations and utility determinations are applied to the data, thereby producing a data *risk profile* and a *utility profile*, respectively. The risk profile helps the provider understand what common statutory, contractual, proprietary, ethical, policy, and best practices obligations and restrictions are implicated when releasing network data. Importantly, this phase considers another aspect of risk, the *threat factor*, which addresses the capability and motivation of the intended recipient(s) to enhance the provider's risk.

By embedding consideration of the threat environment as part of the risk assessment, this framework allows all stakeholders– providers, recipients, and third- party oversight authorities– to assess and defend the reasonableness of a provider's choice of disclosure controls based on adherence to applicable performance criteria – utility and risk. The risk profile is scorecard of risk factors for common data types based on sensitivity and context. The risk profile is manifest qualitatively according to a High - Medium - Low scale for each data type

[2] For a detailed description of the framework, see, Coull, Scott E. and Kenneally, Erin, A Qualitative Risk Assessment Framework for Sharing Computer Network Data (March 31, 2012). 2012 TRPC. Available at SSRN: http://ssrn.com/abstract=2032315 or http://dx.doi.org/10.2139/ssrn.2032315

contained within the disclosed dataset, accumulating into a Phase 1 output that indicates risk on any type of ordinal scale, prior to the application of any disclosure controls. Similarly, the utility profile is a confluence of the choices for six (6) utility choices (audience, timeliness, duration, level of detail, functionality, output) for the entire dataset that also exist on a continuum from highly restrictive data uses (5) to completely unhindered use (1), where use is proportional to chosen utility. The output similarly a qualitative, ordinal depiction of the initially-desired outcome for the data sharing scenario.

The second phase of the framework considers the Phase 1 risk-assessed data (*risk profile*) and utility-ascribed (*utility profile*) data in parallel and offers a menu of disclosure control options for the provider to apply to the shared data to achieve appropriate risk and utility since each option has measurable impact on each of the two components. These controls are organized as: (a) Operational– how the recipient may interact with the shared data; (b) Technical– how the data can be altered to prevent sensitive data leakage; and (c) Policy– how a provider can address the identified data risks ex post via contractual and policy-oriented agreements concerning the access, use and secondary disclosure of the data by the recipient. The notional application of the disclosure control(s) allows the provider and recipient to foresee how a chosen control will enhance or decrease both risk and utility. The output of this application phase is a *modified risk profile* and a *modified utility profile* that reflects the provider's acceptable risk and expected outcome.

The final phase of the framework addresses the impact of the choices made in the previous phase by mapping back to the identified risks and stated utility objectives from the first phase to determine how the *risk profile* and a *utility profile* has changed. This Phase directs the provider to compare the *modified risk profile* to the original *risk profile* and assess whether the recalibrated risk is acceptable, or whether the provider needs to loopback to Phase 2 and apply different disclosure control(s) that will pre-empt unacceptable risk exposure. In parallel, the provider is instructed to map the *modified utility profile* back to the original *utility profile* in order to evaluate whether the desired properties of the disclosure-controlled data (*modified utility profile*) satisfy the original utility objectives beyond the point where it undermines the purpose of sharing the data.

This Disclosure Control Framework uses qualitative metrics to enable data providers to practically assess and communicate privacy risk and utility decisions related to the data disclosure. The component parts– identification of utility and risk, application of operational and technical disclosure controls, and assessment of their impact on utility and risk– can be described using a simple rating scale, where 1 indicates low risk, 3 indicates medium risk, and 5 indicates high risk. The numbers are not used in a quantitative sense, but rather, as common means to symbolize the risk continuum. While true quantitative methods provide the data provider with a specific value indicating the relative severity of the risk of leaking sensitive information, these methods are effective when applied to a very specific threat scenario with a well-defined notion of what the sensitive information is. In the case of network data, however, sensitivity defies precise definition, as it is often context and fact-specific. Instead we are more concerned about general qualities of the data and what they might imply based upon our knowledge of how computer networks operate. In particular, this approach reflects the philosophy that the data provider is better served by understanding general properties of the disclosure control process and then applying that knowledge to the intricacies of the situation at hand. Doing so ensures that the framework is general enough to be applied to new technologies and data sharing scenarios, and avoids the problem of trusting the results of quantitative analysis when they may not be applicable.

3. Multi-Disciplinary Approach:

It is imperative that any practicable and effective approach to identifying privacy risks and beneficial objectives of data sharing, to addressing a solution that contemplates both components, and to assessing/evaluating that solution, enable a range of stakeholders' input. To do otherwise is to ignore the value that each perspective can provide to a problem space that necessarily involves technical and legal expertise at a minimum, and can very well benefit from socio-economic proficiency.

The initial risk and utility profile phase and the final impact assessment phase of the Disclosure Control Framework are designed to engage the frontline technical data provider (e.g., security administrator, network operator) and the business layer risk manager (e.g., legal counsel, CISO) in an exchange of knowledge that is necessary to effectively translate network layer concepts to business layer notions of sensitivity and threat. The sheer complexity of many data sharing situations undoubtedly allow for gaps in the disjointed efforts by engineers, the legal community, and privacy experts in addressing the utility, policy, and privacy concerns arising from data releases. The framework considers the impracticality of an approach that expects personnel who have intimate technical familiarity with the data will be capable of rendering ultimate determinations related to the risk sources (e.g., law, regulation, contract). Therefore, the development of the risk profile and impact involves an iterative exchange of information between the provider's technical and risk management personnel.

For instance, once a provider has engaged the three phases of the disclosure framework, the provider may evaluate whether the overall residual risk and utility is acceptable. This may involve another round of consultation with legal counsel whereby s/he engages the *modified utility profile* has appropriately reduced the risk identified in the initial *risk profile* to an acceptable level for the organization. If the data risk is not acceptable, the provider can notionally apply additional and/or alternate operational, technical, or policy controls to determine how the risk may be modified, and iterate with the risk manager. In general, the quantity and quality of the disclosure controls have a direct relation to the data risk, whereas they are inversely proportional to utility objectives– high data risk demands more restrictive controls, and vast utility calls for lower disclosure encumbrances. So, striking a balance -- where the disclosure controls lowers the data risk to acceptable levels and does not modify the utility objectives such that the purpose of the disclosure is rendered unattainable—requires that the persons who speak "utility" engage with those who speak "risk" using a process that considers both. This framework enables a provider to craft a disclosure strategy along a spectrum of risk thresholds and utility needs according to a repeatable, transparent and evidence-based process. As such, this design embeds a risk management approach to data disclosure that allows technical providers to demonstrate a documented and reasoned process that can be audited, measured and compared over time by legal and business overseers, both within and across provider organizations.

4. Privacy Architecture:

The Disclosure Control Framework described in response to the previous three questions is a concrete architecture, which, because of the design goals described previously, naturally produces a "responsible use framework" as contemplated in this RFI's reference to the Big Data Report. Below we describe this architecture[3], necessarily repeating some previous content, but doing so in terms of general specification in order to demonstrate how the conceptual principles can be operationalized. It is that operational implementation that effectuates a "responsible use framework."[4]

---

[3] Coull, Scott E. and Kenneally, Erin E., "Toward a Comprehensive Disclosure Control Framework for Shared Data", IEEE International Conference on Technologies for Homeland Security (November 2013, Boston, MA). Available at SSRN: http://ssrn.com/abstract=2326264 or http://dx.doi.org/10.2139/ssrn.2326264 .

[4] The disclosure control framework presented in this paper has been used, in part, to develop privacy and disclosure policy recommendations for the FCC's Measuring Mobile Broadband (MMB) project. The main goal of the project is to gather information about the speed, performance, and coverage of current mobile carriers within the United States from the mobile devices of volunteers. To accomplish this goal and make the results available to the public, it is important to protect sensitive information about the volunteer's location information while still ensuring that the level of specificity for each measurement provides useful data for distinguishing performance characteristics among different geographic areas. The situation is complicated by the

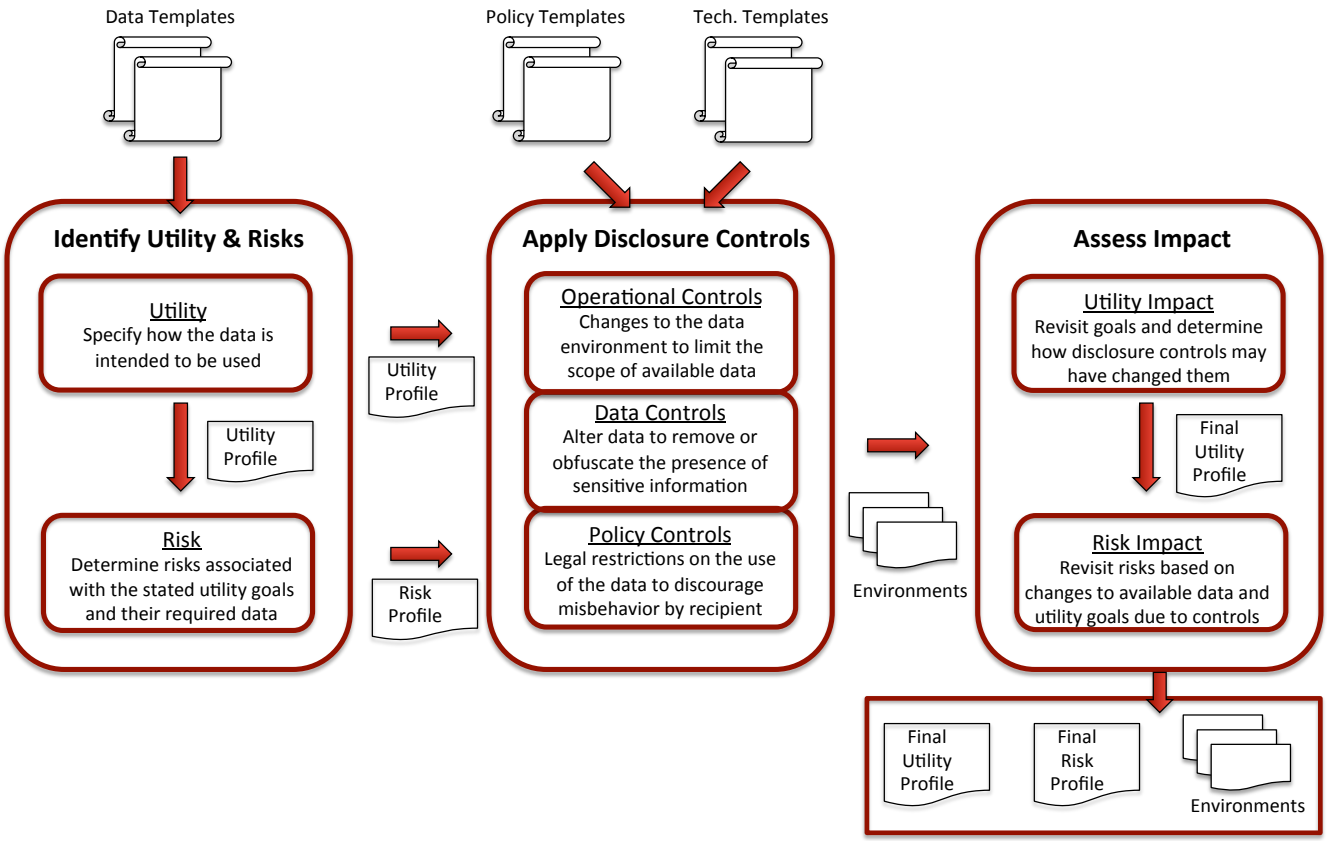

Figure 1.

The disclosure control framework (Fig. 1.) is made up of three basic components: templates, environments, and a risk assessment methodology. Generally speaking, a *template* is a data structure that encodes information about disclosure control components that transcend individual data releases or sharing scenarios. An *environment*, on the other hand, is a concrete instantiation of the disclosure controls and related data sharing infrastructure chosen by the provider for a particular data release or scenario (*e.g.,* access controls, server software, etc.).

Finally, the *risk assessment framework* guides the provider through a decision- making process that essentially transforms a set of templates describing available controls and sharing options into fully- specified environments that can (given the right infrastructure) be automatically implemented and guaranteed to reflect the risk and utility goals of the release. A high-level overview of the workflow is shown in Figure 1. The provider begins by establishing the primary utility goals of the release, which then inform the identification of associated risks to be mitigated. In the second phase of the workflow, the provider chooses a set of disclosure controls to apply to the data, which may include controls to change the data or establish penalties for misusing it. The process ends by having the provider describe how the utility goals and identified risks are impacted by the chosen controls. Throughout this process, the templates restrict the questions and options given to the provider, and connect those options to the available environment configurations. Due to space restrictions, we provide only general specifications for the components and a high-level description of the assessment methodology itself, and refer interested readers to the full version of the paper for more details.

---

sheer number of entities involved with the data sharing effort – mobile carriers, volunteers, the FCC, independent researchers, and curious members of the general public.

This approach offers a number of benefits over current ad- hoc methods. For one, we are able to reuse the basic components of past data releases through established templates; even going so far as to enable community-wide sharing of common templates. These templates describe the standard language of legal documents and basic functionality of technical disclosure controls methods, and then provide a common interface for customizing them. Additionally, the risk assessment method-ology itself forces the provider to consider utility and risk together while providing a unified set of both policy and technical options to achieve the goals of the data release, which enables a more refined balance between utility and risk mitigation. The outputs of each phase in the process also provide a standardized way of describing and justifying the data release so that it can be easily audited by third parties, like Institutional Review Boards (IRBs) or regulatory agencies. In some cases, this standardized information can be used to automatically configure data collection and sharing environments that can be verified to meet the goals set forth in the framework outputs. Overall, the framework provides a way to minimize the long-term costs of the disclosure control and data release process, while simultaneously providing data that is more useful and where the risks are more well-understood by all parties involved.

Creating Reusable Templates & Environments

The foundation of the framework is built upon reusable components called templates and environments. The purpose of the templates is to clearly separate the baseline information and procedures that exists across all data releases from the information about a specific data sharing scenario that must be extracted from the provider. Based on the specifics of the scenario at hand (*e.g.,* type of recipient, data, etc.), there may be many such templates to describe relevant legal and policy documents, as well as the format of the data being released and the applicable technical disclosure controls in use. All templates contain a distinguished name, a user-friendly description of the functionality of the template, and categorical information used to organize the templates for easy examination by the provider. The key functionality here is that the template must be able to facilitate the translation of the provider's choices in the risk assessment and data sharing process into a specification for a concrete implementation. In this paper, we consider document, data, and technology templates, though other types may be added as the need arises (See Figure 2. for examples in JSON).

- A *document template* is meant to encode the boilerplate text of legal and policy documents commonly used when collecting and sharing data, along with a series of questions that must be answered by the provider to customize that text to the current scenario. The boilerplate text contains variables associated with each of the questions such that answering the questions fills in the blanks, so to speak, and allows us to create a complete document that can be used when collecting or sharing data. This is similar to the way privacy impact assessment templates are currently used, though on a much broader scale. The template also contains category information about the types of policy controls (*i.e.,* clauses) contained within the documents.

- To describe the data being released, we use a *data template* that contains information on the syntax of the data type and how disclosure controls may be applied. More specifically, the data template consists of parsing rules (or a pointer to a parser implementation) and a data schema that breaks the data into individual fields that we may apply disclosure controls to. Each of these fields is associated with a type that is used to determine which disclosure controls may be used on that field.

- The *technology templates* describe a single implementation of a disclosure control or supporting technology, such as server software or data collection utilities. In practice, these technology templates will often be abstractions of specific parts of much larger technologies or software implementations, such as a specific type of data filter in the collection software. The template includes information about the field and data types that it may be applied to, its disclosure control categories (dis- cussed in Section IV), available parameters including default settings, and a pointer to its implementation. This information is enough to guide the application of controls to the appropriate types of data, and to the appropriate fields within that data.

The environments are simply sets of templates that have been chosen and configured by the provider during the risk assessment and data sharing process. The primary purpose of these environments is to provide a concise and standardized description of the data sharing scenario, including specific instantiations of policy documents, data sharing software, and disclosure control parameter settings. Like the templates, we limit the scope of the environments we examine to only include collection and sharing environments, which are the two areas where disclosure controls are most often used for network and security data. Other situations where there is no control over the data being collected or which have intermediary processing steps may use a different number of environments.

The *collection environment* is made up of the completed templates that govern how the data is received from up- stream providers, whether they are individual users or large organizations. This environment may include privacy policies, collection filters, and the storage format for the collected data. The *sharing environment* specifies the set of controls used when providing the data to a downstream recipient, such as a researcher, data repository, or the general public, and may include non-disclosure agreements (NDAs), authentication and access control mechanisms, and the disclosure controls to be applied to the data itself. With the right infrastructure, it may be possible to automatically generate the implementation of these scenarios using, for example, a set of baseline virtual machines. At the very least, the environments output by the framework provide a reasonable basis for auditing the provider's data collection and sharing practices.

Document Template

```
"name": "Privacy Notice",
"description": "A privacy notice and terms of use.",
"categories": ["upstream", "terms", "covenants"],
"text": "We collect [#(COLLECTED_DATA)] kinds of
        information to measure the performance of
        your mobile broadband service.
        [COLLECTED_DATA]
        This data is protected using
        [PROTECTIONS]. You can find more detail in
        the FCC's technical summary of this program.",
"questions": [
        {"question": "Enumerate data items collected.",
        "answer": "COLLECTED_DATA"},

        { "question": "Enumerate protections for raw
                data after collection.",
        "answer": "PROTECTIONS"}
]
```

Technology Template

```
"name": "Quantize_Location",
"description": "Aggregate location data into blocks.",
"categories": ["aggregation", "location_data"],
"pointer": "http://example.com/quantize_loc"
"parameters":
{
        "k": "10",
        "granularity": "0.5"
}
```

Data Template

```
"name": "GPS Data",
"description": "Lat. and long. data",
"categories": ["location_data"],
"parser": "http://example.com/gps"
"schema": "
{
        "accuracy": "float",
        "latitude": "float",
        "longitude": "float",
        "timestamp": "datetime"
}"
```

FIgure 2.

Utility & Risk Assessment

Utility and risk are often inextricably linked in data sharing efforts. Despite this close relationship, most risk assessment methodologies rarely, if ever, explicitly consider utility even though it is the driving force behind the effort. Rather than ignoring utility, our methodology makes it the central focus. The first step in our framework, therefore, is to assess the utility goals of the data release and the risks related to achieving those goals. These key factors are captured in *utility and risk profiles*, respectively. The profiles allow the data provider to create a concise and standardized audit trail of the decision-making process underlying the data release.

Choosing Disclosure Controls

Once the utility and risk profiles have been created in the assessment phase, the data provider's next task is to create one or more environments to sufficiently mitigate the identified risks and uphold the utility goals. This is accomplished by choosing from among the available document and technology templates, then configuring their properties and applying them to the data as a whole, or in part. In many cases, there are several ways to reach the same end state by applying different combinations of controls, such as when we limit data collection via filters or delete parts of the data after it is collected. Obviously, the chosen controls may achieve the same level of risk mitigation, but often at different costs. The category associated with each template is used to organize the available options, while the name and description provide an understanding of the specific usage, benefits, and drawbacks for the template. There are three broad classes of that span operational, data, and policy controls, along with several sub- categories for each.

Operational controls restrict different aspects of the sup- porting data collection and sharing infrastructure in an effort to minimize the exposure of the sensitive data, either before it is collected or after it is made available to the recipients. Some examples include access controls, use of specific data formats, and timing restrictions on data availability. The operational controls are broken into the following six sub-categories:

- Filtering: Limits the data to a specific sub-population as it is being collected and stored. In some cases, the most risky data population can simply be ignored during collection to mitigate its potential risks.
- Duration: Specifies the amount of time the data is available to recipients, which limits exposure of the data to potential abuse.
- Timeliness: Controls how long the data is retained before it is made available to the recipient. Based on the time-sensitivity of the risk factors involved in the use cases, it may be possible to enforce long waiting periods before the data can be accessed.
- Length: Collecting data over short periods of time often provides more limited exposure for potentially risky data, while longitudinal collection often leads to information that is deeply rooted in the patterns that emerge over time, which can make mitigation more difficult.
- Format: Some data formats naturally encode less de- tailed information than others, or may naturally restrict the data to only a small number of pertinent fields. This may help to focus the data collection effort to only the most basic information necessary.
- Access: There are several methods that can be used to control and audit access to the data itself, including limited query interfaces or other mitigated environments. When the data is accessed within a controlled environment, it may be possible to offset any potential risk factors of data exposure with stronger authentication and auditing mechanisms to recover from malicious activities.

Data controls alter the data itself after it has already been collected and stored. Here, the data controls may be applied to the dataset as a whole, to specific rows or columns, or to very specific pieces of information (*i.e.,* data cells). These data abstractions are dictated by the data templates associated with this data release.

There are six sub-categories of data controls:

- Deletion: Simply removes a row (record), column (field), or specific set of cells in the data. The difference between deletion and the operational limitations above is that using deletion allows the provider to examine the data and make more dynamic disclosure control choices based on the results of the collection itself.
- Aggregation: Takes the values of a field over several records and aggregates them into a single record value. For instance, the age of participants in a survey may be aggregated by taking the average of their ages. These methods attempt to blend the impact of any one record in with others while still providing useful trending information.

- Generalization: Uses the semantics of a field to generalize several related classes of values into a single large class. An example of this would be truncating zip codes or Social Security numbers, which effectively generalizes the values into groups based on larger geographic areas.
- Pseudonymization: Replaces identifiers with a linkable or partially-unlinkable pseudonym to hide the real identity associated with the record, but maintain the ability to group those records together.
- Perturbation: Changes the value of a field and combining it with noise, such as adding noise taken randomly from a Laplace distribution to a number.
- Synthetic Data: Given a set of specific statistical properties to maintain, generative models can be trained to produce data that is guaranteed to match those proper- ties, but which has no connection to the original data for any other property. Data imputation techniques are also considered to be a type of synthetic data generation method.

Policy controls mitigate risk not by trying to hide or limit access to the risky data, but instead by providing strong incentives for appropriate behavior and penalties for abuse. In addition to risks arising from the exposure of the data itself, there are often other risks related to various policy aspects of the collection and sharing process, such as the need for informed consent from users or transitive application of agreements from upstream providers to downstream recipients. Examples of these policy controls include privacy policies, memorandums of agreement, data licenses, and non-disclosure agreements associated with upstream data providers. The basic functionalities of the policy controls are categorized as follows:

- Performance: Describes the bargain or exchange, such as the scope of data that is protected, and license grants or restrictions.
- Consideration: Restrictions related to required fees or necessary services that are related to the collected data or its source.
- Covenants & Conditions: Requirements or obligations placed on the parties for use of the data, such as consent notices, confidentiality obligations, and destruction of data after a specified period of time.
- Accountability & Enforcement: Guarantees and mechanisms to enforce or police them, including penalties and auditing rights.
- Terms & Termination: Specific termination conditions for the use of the data, time period of use, conditions under which can it be ended.
- General: Basic requirements imposed by governing law or other third-party governing law, interpretation and adjudication; binding effects, third party beneficiaries.

The process of choosing the disclosure controls using this framework is guided by the categorization of the policy and technology templates, and their applicability to the chosen data templates. Once a control is selected by the provider, the list of questions or parameters found in the associated template are presented so that they may be customized to the risk level of the current data sharing scenario. The provider also chooses which environment (if more than one exists) the control should be associated to. The output of this phase of the framework is a set of environments containing the configured templates chosen by the data provider. While the framework does not currently consider the notion of completeness (*i.e.,* the idea that there is a necessary set of templates), it is possible that in the future it may be extended to establish requirements for certain types of controls. For instance, sharing obviously cannot occur without choosing some type of server technology, and so that may end up becoming a requirement in future iterations of the framework's implementation.

<u>Evaluating Disclosure Control Impact & Using Framework Outputs</u>

The final phase of the framework comes full circle to determine how the choices of disclosure controls has changed the original utility goals and risk factors identified at the start of the process. The evaluation proceeds by adding an additional column to the utility and risk profiles, called the *impact statement*. The provider uses the field to specify how they believe things have changed, either quantitatively or qualitatively. As with the initial profiles, we cannot rely on a one-size-fits-all approach when talking about quantitatively measuring

change in inherently qualitative utility properties and risk factors. We can, however, make some quantitative measurements where they naturally occur, such as an increase in lag time between the time data is collected and when it is made available to recipients. In general, though, we believe that qualitative impact statements will be the most generally useful approach. Again, there is no claim that the data is guaranteed to be safe, but these final profiles help encourage defensible and pragmatic solutions. In fact, use of the development of community-driven templates and use of the framework itself helps to set a standard for what can be considered to be a “reasonable” level of due diligence on the part of the provider.

Once the profiles have been completed, the framework outputs the final utility and risk profiles, as well as the set of environments created during the disclosure control phase of the framework. There are several uses for these outputs that greatly improve the current state-of-practice. Probably the most obvious use is to provide the profiles and environments to third-party auditors to review the decisions made in choosing the controls. Since the profiles provide direct support for the environment configurations and those environments are standardized, it is much easier to have a data privacy expert or attorney verify that the controls meet the necessary requirements. By comparison, the current approach would be to engage the experts on an ad-hoc basis with little or no information about the complete data sharing scenario, instead receiving only piecemeal verification of the controls and data sharing policies.

Another, more ambitious use of the output is to use it to automatically assemble implementation artifacts for each of the environments. As mentioned earlier, it is possible to create one baseline virtual machines (VMs) for each environment in the output, with all of the available tools pre-installed. Then, the applications within the VMs are configured according to the templates within their respective environments. Such a system would remove almost all technology and implementation costs involved with new data sharing efforts, and enable simple verification procedures for compliance with stated policies. At the moment, compliance checking of any sort is actually impossible because the disclosure control policies are not formalized and the implementations are not standardized.