The following information was posted with the attached DRAFT document:

Feb. 17, 2016

NIST IR 8103

DRAFT Advanced Identity Workshop on Applying Measurement Science in the Identity Ecosystem: Summary and Next Steps

On January 12-13, 2016 the Applied Cybersecurity Division (ACD) in the National Institute of Standards and Technology's (NIST) Information Technology Laboratory hosted the "Applying Measurement Science in the Identity Ecosystem" workshop to discuss the application of measurement science to digital identity management. Draft NISTIR 8103 summarizes the concepts and ideas presented at the workshop and serves as a platform to receive feedback on the major themes discussed at that event.

Comments on Draft NISTIR 8103 should be emailed to NSTICworkshop <at> nist.gov.
The comment period closes on **March 31st, 2016**.

NISTIR 8103 (Draft)

Advanced Identity Workshop on Applying Measurement Science in the Identity Ecosystem: Summary and Next Steps

Mike Garcia
Paul Grassi
Ryan Galluzzo
Walter McClean

This publication is available free of charge from:

National Institute of Standards and Technology
U.S. Department of Commerce

NISTIR 8103 (Draft)

Advanced Identity Workshop on Applying Measurement Science in the Identity Ecosystem: Summary and Next Steps

Mike Garcia
Paul Grassi
Applied Cybersecurity Division
Information Technology Laboratory

Ryan Galluzzo
Walter McClean
Deloitte & Touche LLP
Rosslyn, VA

This publication is available free of charge from:

February 2016

U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Under Secretary of Commerce for Standards and Technology and Director

National Institute of Standards and Technology Internal Report 8103
14 pages (February 2016)

National Institute of Standards and Technology
Attn: Applied Cybersecurity Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 2000) Gaithersburg, MD 20899-8930

Abstract

On January 12-13, 2016 the National Institute of Standards and Technology's (NIST) Applied Cybersecurity Division (ACD) hosted the "Applying Measurement Science in the Identity Ecosystem" workshop to discuss the application of measurement science to digital identity management. This document summarizes the concepts and ideas presented at the workshop and serves as a platform to receive feedback on the major themes discussed at that event.

Keywords

Identity; NSTIC; authentication; biometric authentication; biometrics; identity proofing; attributes; metadata; identity management; cybersecurity; security; information security.

Disclaimer

Any mention of commercial products or reference to commercial organizations is for information only; it does not imply recommendation or endorsement by NIST, nor does it imply that the products mentioned are necessarily the best available for the purpose. Content was derived from workshop participant discussions captured by note takers and aggregated for the purposes of summarizing the event. Any misrepresentation of comments or concepts is unintentional. Corrections or clarifications can be provided through the open comment period.

Workshop Comments

Organizations are encouraged to review this draft publication during the public comment periods and provide feedback.

Comments on this publication may be submitted to: NSTICworkshop@nist.gov.

Public comment period: February 17, 2016 – March 31, 2016.

Acknowledgements

The authors would like to thank panelists Ian Glazer, Kim Little-Sutherland, David Kelts, Dario Berini, Brent Williams–Verato, Julian White, Brett McDowell, Stephanie Shuckers, Vance Bjorn, Cathy Tilton, Liz Votaw, LaChelle Levan, Darran Rolls, Gerry Gebel, Ryan Disraeli, and Robin Wilton and facilitators Kirk Brafford, Mike Wyatt, Roger Cressey, Kiersten Todt, JR Reagan, and Colin Soutar, as well as the workshop participants who provided valuable input to this report.

Table of Contents

Introduction

On January 12 and 13, 2016, the Applied Cybersecurity Division (ACD) in the National Institute of Standards and Technology's (NIST) Information Technology Laboratory hosted the "Applying Measurement Science in the Identity Ecosystem" workshop in Gaithersburg, Maryland. The two-day workshop brought together security practitioners, identity solution providers, subject matter experts, and policy makers from across the public and private sectors to discuss the application of metrics and measurement science to common identity management practices.

The identity ecosystem has matured to the point where it is appropriate to undertake the work of building measurement science for application in the market—a critical step in further aiding expansion and innovation of the identity ecosystem. This workshop was held to obtain feedback from stakeholders on the feasibility of, and approaches necessary to, measure and compare three disciplines of digital identity management:

1. Strength of identity proofing;
2. Strength of authentication; and
3. Attribute metadata and confidence scoring.

NIST's ultimate goal is to establish frameworks that enable objective measurement of identity solutions, so that their ability to mitigate risk is more quantitatively measurable, they can more easily be compared, and, ultimately, measured when combined. NIST believes making progress in this space will achieve greater alignment of identity solutions and technology with risk assessment and management practices. This document provides a summary of the proceedings to ensure NIST captured stakeholder feedback accurately as it executes the next steps in its broad effort towards improved digital identity.

Workshop Summary and Key Takeaways

Workshop attendees represented diverse public and private sector stakeholders. In total, 224 people attended the event: 67% from the private sector, 26% from government organizations, and 7% from academia and non-profits. The workshop included moderated panels and facilitated working sessions for each workshop topic. Throughout the event, participants shared risk management practices, security evaluation approaches, and testing processes that they utilized within their organizations. Additionally, participants identified barriers, evaluated solutions, and specified implementation considerations to enable greater quantification of strength within each digital identity management discipline the workshop covered.

The summary below identifies takeaways and observations from the event. These do not necessarily indicate items that were unanimously supported by those in attendance, but rather frequently voiced ideas and input among panelists, audience questions, and the breakout teams during the course of workshop.

Overall Observations

NIST heard several recurring themes that transcended the individual workshop topics. These typically involved NIST's overall effort to apply measurement science to digital identity, the efficacy of measurement within each topic, and how the relationships, or lack thereof, between topics could influence a future direction.

- **Application of Metrology to Digital Identity and Access Management.** Many participants saw value in NIST's effort to establish measurement science to communicating the strength, and ability to mitigate risk, of identity management practices and solutions. Furthermore, most expressed willingness to remain engaged as those efforts develop and mature. Some attendees expressed a view that mandatory metrics and measurements may place an undue burden on vendors. Overall, attendees felt the three focal areas of the workshop were appropriate to evolve and enhance the identity ecosystem, and supported NIST's efforts to produce measurement-based guidance associated with each.

  While the idea of producing additional guidance regarding measurements and metrics was generally well received, there was no consensus on any specific approach to apply measurement science to digital identity, nor how to develop such approaches. Likewise, there was no consensus on the metrics that should be measured and reported within systems and federations. A few participants felt scoring digital identity processes and technologies was neither feasible nor appropriate.

- **Improved Transparency and Standardization.** Most participants expressed a desire to see increased transparency and standardization across identity practices—particularly in the realm of remote identity proofing practices. Many attendees expressed a desire to better understand the way identity solutions operate and to overcome a lack of visibility, whether real or perceived, into how proprietary scoring works within existing remote identity proofing solutions. Many participants also saw a need to better understand the processes that contribute to data they leverage and trust to remote proof identities. Many felt that standardized processes for evaluating solutions and communicating the efficacy of these solutions would provide greater interoperability and trust on a broad scale.

- **Flexibility and Extensibility.** Participants broadly encouraged NIST to ensure that any future guidance is both flexible and extensible to support the diverse needs of different communities, trust frameworks, and sectors. Notably, most participants wished to ensure that any guidance was reflective of the need to address risks associated with NIST's primary constituents—federal agencies—while also acknowledging the needs and concerns of the private sector. This reflected the view that many, if not all, digital identity solutions will come from

the private sector vendor community, so NIST must attempt to develop guidance that does not create an environment where cross-sector solutions will no longer be viable within the Federal Enterprise. Many expressed a strong desire for NIST to craft documentation in a manner that could subsequently be submitted as a work product in open, consensus-based international standards development organizations.

- **Topic Area Relationships:** Participants acknowledged the pre-event white papers as thoughtful starting points for much of the workshop discussion, but sought greater insight on how the measurement of authenticator strength, remote and in-person identity proofing, and attribute confidence would or could impact each other in future NIST deliverables. Many suggested NIST should explore an overarching model of identity measurement to help clarify the role of measurement science in digital identity and the interplay between these, and potentially other, components.

- **Existing Work and Fora:** The digital identity community is one of constant innovation and many complementary efforts. Participants repeatedly sought to ensure that NIST recognize and collaborate with existing initiatives focused on similar outcomes. Participants identified multiple existing efforts as places where NIST could leverage synergies to advance the community's collective interests. Efforts consistently mentioned include:
    - UK Cabinet Office and the Good Practices Guides (GPGs) : https://www.gov.uk/government/collections/identity-assurance-enabling-trusted-transactions
    - The Kantara Initiative Identity Assurance Working Group: https://kantarainitiative.org/confluence/display/idassurance/Home
    - The IETF Vectors of Trust internet draft: https://tools.ietf.org/html/draft-richer-vectors-of-trust-00
    - OASIS Trust Elevation Group: https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=trust-el
    - ISO/IEC SC 27 Working Group 5: http://www.iso.org/iso/iso_technical_committee?commid=45306
    - ISO/IEC SC 37 biometric activities: http://www.iso.org/iso/home/standards_development/list_of_iso_technical_committees/iso_technical_committee.htm?commid=313770
    - FIDO Certification Working Group: https://fidoalliance.org/working-groups/
    - FIDO Biometric Assurance Sub-working Group

Strength of Identity Proofing

Strength of identity proofing was the first topic of the workshop. Participants discussed existing and potential identity proofing methods and ways to measure strength of each

individual process, as well as the establishment of a scoring framework to communicate common results of digital identity proofing for the purposes of risk management. Across the discussion groups, several major themes emerged.

- **Develop a common lexicon**. Many participants identified a lack of standardized terminology regarding identity proofing processes and functions. For example, some attendees used the term "verification" while others preferred "validation" for the same process. For the purposes of NIST's work, attendees suggested a common vocabulary should be developed to help ensure consistency in the framework and across communities, and that the taxonomy be aligned to the best extent possible with existing schemes.

- **Identify functional components of proofing.** Attendees in most sessions came to the conclusion that proofing could be broken down into a set of component functions or actions that could potentially be evaluated to provide a greater understanding of the processes and the results associated with verifying a claimed identity. Each component could potentially serve as the basis for a scoring structure.
    - Participants suggested additional functional components that could be added to those currently explored in the white papers. Specific suggestions included: ongoing maintenance of an identity (i.e., how a provider manages the identity, updates it, and supports necessary modifications when needed); fraud and compromise detection; document authentication; activity history of an identity; biometric collection to support the binding of proofing to a credential; and processes for binding proofing data to an identity.

- **Avoid a single score**. Many participants expressed the belief that any scoring of the processes associated with identity proofing should not be aggregated into a single score. Instead, many felt it more appropriate to provide individual scores for the processes that could be considered and weighted by RPs to meet their needs. In some instances, more fine-grained knowledge of the processes an individual underwent to confirm a claim of identity would be just as valuable as a score.

- **Consider existing standards and practices**. Several participants referenced UK GPG 45 as an example of combining high-level scoring with desired outcomes. Participants discussed the potential to draw lessons from the UK GPGs and apply them to a US based identity proofing framework.

- **Define scope of proofing**. Participants also discussed the scope of identity proofing, specifically that the goal of identity proofing guidelines should be scoped to proving a valid identity exists. Proofing should not, for example, validate an individual's rights and privilege to obtain specific entitlements. Determining entitlements and eligibility is an RP decision that goes beyond confirming that and identity is associated with a specific individual.

Strength of Authentication

The second workshop session addressed the strength underlying various authentication methods. The session explored measuring mitigations of known threats to an authentication system as a method to determine an overall score for authenticator performance as well as an overall construct that would enable the assessment and comparison of distinct authentication mechanisms. While the strength of authentication white paper identified biometric authentication as a starting point for an overall authentication framework, several of the workshop sessions ended up extensively discussing the broader concept of evaluating various authenticator technologies. While biometrics was selected first due to its increasing consumer and commercial adoption rates, the framework envisioned by NIST would support the evaluation of authentication strength regardless of form or factor, making these broader discussions extremely valuable. Across the groups, several major themes of discussion emerged.

- **Emphasize importance of addressing usability**. Most participants felt that user experience with a chosen authentication method is one of the most important factors in selecting technologies that are not only secure, but also likely to be successfully adopted. Many pointed out that the largest driver behind the adoption of mobile biometric solutions is market demand and the ease with which users are able to access services. As a result of this ease-of-use focus by consumers, many participants noted that security may not be the primary objective of many relying parties (RP) when instituting authentication solutions. Therefore, inclusion of user experience in any evaluation scheme may have a benefit to both security personnel required to assist in risk management and mitigation, and to business decision makers.

- **Consider user experience—or a poor user experience—as a vulnerability**. This led some participants to suggest incorporating a usability score into the framework. Participants also considered the possibility that poor UX could be considered a system "vulnerability" and weighted, evaluated, and scored much as the other components of the score. However, poor user experience should not be the only metric. Rather, the result of poor user experience will be the users themselves trying to exploit workarounds to improve their individual experience with the technology. These workarounds would be considered the vulnerability.

- **Consider framework usability**. In addition to the importance of usability with respect to authentication solutions, many identified a need for any scoring or evaluation framework to be usable as well. Attendees indicated that some RPs struggle to balance the need to deliver cutting edge solutions to the market with the needs of security and privacy. For a measurement based framework to have broad adoption it must enable rapid evaluation of solutions to allow users to maintain pace with markets and customer demands.

- **Evaluate the complete authentication system**. Many attendees agreed with the perspective in the white paper that a scoring framework should look at the authentication solution as a system, regardless of whether it is deployed as one

single system on a local device, or a system of systems distributed across a network and security boundaries. It was a common observation that authentication technologies face many different attack vectors, threats, and vulnerabilities, which target different aspects of the authentication process. The strength of a solution can only be effectively assessed and determined by identifying how the layers of security combine to protect the entire process. While a common practice in the public and private sectors, the fact that locally matched biometrics on a mobile device are often used to unlock a cryptographic token was introduced as an example that a scoring framework cannot just evaluate the performance of the biometric sensor and matching algorithm. The framework also needs to evaluate additional security functions such as the way keys are generated, the algorithms used, how keys are stored, and how keys are transmitted by the device.

- **Incorporate consideration of multiple biometric modalities.** Much like the previous comment focused on evaluating the authentication method as a complete system, many attendees expressed the belief that any measurement framework needs to take into account the ability of a biometric authentication schema to incorporate multiple different modalities (e.g., fingerprint, iris, voice) to increase assurance or security of the overall system. There was also discussion of how the layering of biometric modalities could be compared to other authentication methods. For example, some participants asked if the incorporation of more than one modality is comparable to adding more than one factor to an authentication. A measurement framework geared towards biometrics should be able to address multiple modes and support comparison of these modes and how they can be combined to address vulnerabilities as part of a complete authentication system.

- **Incorporate liveness testing.** Participants expressed the importance of liveness testing—the ability for a biometric system to determine if the biometric being presented is from a live authorized person—to understanding the strength of biometric authentication methods. This is a particular challenge in the case of biometrics on consumer-owned mobile devices as the biometric event is unattended. In some venues, such as the FIDO alliance, there have been some initial discussions of how liveness detection can be incorporated, tested, and evaluated on mobile devices, but NIST should place particular emphasis on ensuring this is properly incorporated into any measurement framework.

- **Consider testing and evaluation**. There was substantial conversation of the challenges associated with testing biometric devices, from the need to understand the difference between a laboratory and operational environments, to the infrastructure needed to support testing, and the challenges of ensuring conformance in a production environment. Many participants saw value, but substantial challenges, associated with standardizing an evaluation process. Participants recommended that any framework to determine strength should consider the requisite testing procedures before being finalized. In other words, if the testing procedures are difficult, costly, and time consuming, the framework itself may be deficient.

- **Consider existing testing models**. Participants raised comparisons to Federal Information Processing Standard (FIPS) 140-2 testing and certification processes, though stakeholders had mixed feelings over the prospect of a similar concept for commercial authentication technologies. While some participants were in favor of an "approved product list" for authenticators to ease acquisition, many others felt that a testing and certification program based on the FIPS 140-2 model would not be agile nor flexible enough to support commercial innovation or maintain pace with emerging threats.

Attribute Metadata and Confidence Scoring

This topic centered on defining standardized attribute metadata to assist relying parties in making risk based decisions about the efficacy of using an attribute when evaluating access control policies. In addition, the whitepaper on this topic also proposed a scoring framework for attribute confidence scores based on the metadata. Across the discussion groups, the following major themes emerged.

- **Consider cost and performance**. Some attendees saw value in the development of standardized metadata to support greater RP understanding of an attribute's trustworthiness—so long as the development and implementation of this metadata is done in a way that considers the impact to cost of integration and system performance. Participants noted that without careful consideration of these factors, the metadata is unlikely to see broad adoption or achieve its desired impact. The focus should be on metadata in identity assertions, not necessarily the requirement for data storage technologies to adopt attribute metadata confidence schemes.

- **Revisit proposed metadata elements**. There was substantial feedback on the individual metadata elements. Examples include:
    - Attribute currency and specifically how concepts such as decay rate, freshness, and date since last verification could affect confidence scoring for attributes. Many participants acknowledged that these factors may not apply equally to all attributes, but should be explored in greater detail none the less.
    - Complications around the term consent in "individual consented," and how privacy enhancing requirements could be better instantiated in the metadata elements.
    - Concerns about terminology, particularly with respect to "provenance" and the types of values allowable under "verification."

- **Determine when and how metadata is communicated**. Attendees reinforced the need to understand how and where this metadata would fit into transactions. Participants discussed if these details would be agreed to or established through contracts or existing structures such as trust frameworks and federations. Some participants questioned if the metadata would be included in transactions during

run-time or sent on a regular, but less frequent basis. There was concern about which elements would be mandatory and which optional.

- **Discuss differences between provider and individual attribute scores**. Participants questioned the granularity of the attribute score and if it would be best applied as a single provider score or individual attribute scores. If single provider, steps would need to be in place to ensure which attributes the provider was capable of confidently asserting, or if the provider score would be applicable to any or all of the attributes of which it is a steward.

- **Maintain flexibility and extensibility**. A large share of participants stated that any metadata and scoring structure for attributes would need to be flexible and extensible enough to support a broad set of different use cases. A common assertion was that NIST could not plan to produce a single universal scoring structure or metadata standard that addresses every community or potential application. Participants in several breakouts discussed the concept of a creating a core framework, which could be tailored or "profiled" to support specific sectors.

Next Steps

The workshop and this document, designed to present the major themes NIST identified from stakeholder input, are just the first steps in establishing new resources and guidelines for digital identity management.

Based on the feedback collected during the workshop and other interactions with stakeholders, NIST intends to further develop documentation to support advancement of the identified topics. In the coming months, NIST will focus on three steps: 1) determining the type of material that will need to be developed to most effectively forward these efforts; 2) establishing new processes to foster greater collaboration, as well as early and frequent community interaction in the development of NIST documents; and 3) determining the best fora for advancing these efforts. Each is essential to ensuring the establishment of viable and broadly accepted framework for employing metrology in the identity ecosystem.

In addition to the general steps listed above, NIST has determined a specific course of action across the three topic areas to make immediate progress and remain transparent in our development activities. NIST will:

- Publicly post "project charters" which will outline the goals, objectives, stakeholders, milestones, and development/outreach methods for each topic area.
- Transition the attribute metadata and confidence white paper to a NIST Internal Report (NISTIR). This NISTIR will cover metadata only and serve as an "implementer draft" to determine scope of market adoption, identify lessons learned, and determine a path forward post-NISTIR. NIST will research a scoring framework in parallel efforts.
- Upon completion of the initial draft of the attribute metadata NISTIR, commence a series of iterative public comment and development periods utilizing Github

public repositories to manage contributions and comments. This process will follow similar processes in government, such as OMB's comment period to Circular A-130, Managing Information as a Strategic Resource. NIST will also consider offering traditional methods of submitting comments for those not comfortable or willing to use Github. NIST will use these relatively short, iterative public comment periods in lieu of an extended comment period. However, based on NIST feedback to this approach, NIST may also provide an extended comment period.

- Solicit stakeholder feedback to determine the scope and path for measuring the strength of identity proofing.
- Transition the strength of authentication white paper to a NISTIR focused on a framework for evaluating biometric authentication systems. This document will focus on leveraging the proposed vulnerability based model to better understand how biometric systems can be scored and compared to mitigate authentication risk. As part of this effort, NIST will also identify elements of the framework for scoring a biometric authentication system that can be used in future work on other authentication systems.

Stakeholders who wish to comment on this document, the workshop white papers, or provide written contributions to any of the digital identity topics discussed at the workshop can do so by emailing NSTICworkshop@nist.gov. NIST encourages respondents to consider the following questions:

1. Are there other takeaways you heard over the course of the event that have not been captured in this report?
2. Are there any additional initiatives, venues, or fora where similar or complementary work is being conducted?
3. For each topic, is there a recommended development approach that you would like to see NIST undertake?

Tackling the challenges associated with managing digital identities in a privacy preserving, secure, interoperable, user friendly, and economical manner will take time, effort, and cooperation. It requires an adaptable, transparent, and iterative approach capable of driving the innovation necessary to improve and evolve the Identity Ecosystem. NIST believes that this flexibility lies at the nexus of measurement, metrics and risk management, and intends to work with our stakeholder community to establish an improved foundation that aligns each of these within the identity ecosystem.