JANUARY 2022

Move Over JARVIS, Meet OSCAR

Open-Source, Cloud-Based, AI-Enabled Reporting for the Intelligence Community

AUTHOR
Emily Harding

A Report of the CSIS International Security Program

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

JANUARY 2022

Move Over JARVIS, Meet OSCAR

Open-Source, Cloud-Based, AI-Enabled Reporting for the Intelligence Community

AUTHOR
Emily Harding

A Report of the CSIS International Security Program

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

About CSIS

The Center for Strategic and International Studies (CSIS) is a bipartisan, nonprofit policy research organization dedicated to advancing practical ideas to address the world's greatest challenges.

Thomas J. Pritzker was named chairman of the CSIS Board of Trustees in 2015, succeeding former U.S. senator Sam Nunn (D-GA). Founded in 1962, CSIS is led by John J. Hamre, who has served as president and chief executive officer since 2000.

CSIS's purpose is to define the future of national security. We are guided by a distinct set of values—nonpartisanship, independent thought, innovative thinking, cross-disciplinary scholarship, integrity and professionalism, and talent development. CSIS's values work in concert toward the goal of making real-world impact.

CSIS scholars bring their policy expertise, judgment, and robust networks to their research, analysis, and recommendations. We organize conferences, publish, lecture, and make media appearances that aim to increase the knowledge, awareness, and salience of policy issues with relevant stakeholders and the interested public.

CSIS has impact when our research helps to inform the decisionmaking of key policymakers and the thinking of key influencers. We work toward a vision of a safer and more prosperous world.

CSIS does not take specific policy positions; accordingly, all views expressed herein should be understood to be solely those of the author(s).

Acknowledgments

This report is made possible by support from Amazon Web Services.

Center for Strategic & International Studies
1616 Rhode Island Avenue, NW
Washington, DC 20036
202-887-0200 | www.csis.org

Contents

The Issue **1**

Introduction: OSINT Opportunities for the IC **3**

1 | What Is OSINT, and What Should It Be? **6**

2 | OSCAR'S Grandfather: IC Use of the Cloud, AI/ML, and OSINT Today **8**

3 | Obstacles to OSCAR **10**

The Culture Problem *11*

The Security Problem *14*

The Policy Problem *17*

When Absence of Policy Is the Problem *19*

4 | Possibilities Ahead **21**

OSCAR 1.0: Near-Term Possibilities *22*

OSCAR 2.0: Mid-Term Possibilities *24*

OSCAR 3.0: High-Powered Compute and Quantum-Enabled Possibilities *25*

5 | Recommendations: Training OSCAR and Training Us **27**

Culture *28*

Security *29*

Policy *30*

People *31*

Bold Steps *32*

Conclusion **33**

About the Author **35**

The Issue

"The OSINT [open-source intelligence] world outside the IC [intelligence community] is so far advanced and we're not keeping up with it. We've already been eclipsed, and we run the risk of becoming irrelevant. We need to catch up quickly so that we don't become a very expensive irrelevance."—workshop participant[1]

"In future information environments of ubiquitous sensing and continual awareness, the commercial sector's faster technology adoption rates and superior facility with OSINT could give it the advantage over the IC in assessing fast-moving global events . . . [IC] analysts may fall behind and outside of policymakers' information and decision cycles . . . The task force has concluded that the IC must fundamentally reconceptualize OSINT as a cornerstone of U.S. intelligence."—Maintaining the Intelligence Edge[2]

Systems using artificial intelligence (AI) could save an intelligence community (IC) analyst as much as 364 hours, or more than 45 working days, a year.[3] Universities run 100,000 core AI or machine learning (ML) models in partnership with an unclassified cloud provider, and high-level IC policymakers have publicly stated the necessity of embracing open-source intelligence (OSINT) as a

[1] Workshop participant #1.

[2] Brian Katz et al., *Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence through Innovation* (Washington, DC: Center for Strategic and International Studies, January 2021), https://www.csis.org/analysis/maintaining-intelligence-edge-reimagining-and-reinventing-intelligence-through-innovation.

[3] Kwasi Mitchell et al., "The Future of Intelligence Analysis: A Task-Level View of the Impact of Artificial Intelligence on Intel Analysis," Deloitte Insights, December 11, 2019, https://www2.deloitte.com/us/en/insights/industry/public-sector/artificial-intelligence-impact-on-future-intelligence-analysis.html.

core analytic discipline.[4] One can easily imagine a near-term future in which something like Tony Stark's JARVIS can assist analysts and operators in a wide range of duties. This paper dubs a hypothesized open-source, cloud-based, AI-enabled reporting capability for the intelligence community as "OSCAR." It may be years before OSCAR can recognize analysts' sarcasm as JARVIS can with Tony, but the combination of unclassified cloud capability, vast new sources of publicly available information, and AI/ML tools could accelerate intelligence work and transform capabilities in the near term.

Despite these potential advantages, the intelligence community has been slow to adopt AI/ML capabilities to make sense of masses of untapped OSINT data. This report explores the reasons for the delay, building on the CSIS report *Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence through Innovation*, which clearly laid out the benefits of incorporating unclassified, cloud-based, open-source intelligence. This follow-on report redefines OSINT and looks at how the IC is currently using the key technologies that will underpin the OSINT revolution. It then discusses obstacles to that revolution, breaking down the elements of ingrained culture, security practices, and policy decisions that hold back adoption of AI/ML applications, and the mechanisms that would greatly assist OSINT integration into the larger IC digital workflow. In the "Possibilities Ahead" section, it looks to near-, mid-, and long-term opportunities for the IC to adopt OSCAR and accelerate intelligence work. Finally, it puts forth a slate of actionable recommendations that will break the logjam and allow the IC to obtain and deliver intelligence from anywhere, taking advantage of cloud and edge-cloud computing, AI/ML tools, and OSINT. In particular, this report argues that the intelligence community should stop "recreating the internet in a classified environment, which is highly expensive and time consuming," as one interviewee put it, and instead accept a small amount of risk to run applications on an unclassified cloud, taking advantage of increasingly sophisticated and automated cloud security and obfuscation capabilities.[5]

Much has been written advocating for the IC to make urgent strides in both OSINT and AI/ML capabilities. To dig deeper into the root causes behind its reluctance to adopt these capabilities and to provide actionable recommendations, researchers engaged in more than 20 interviews with experts in the fields of intelligence, commercial use of publicly available information, AI/ML, and cloud computing. We also held a workshop with experts in these areas, as well as in government contracting and innovation, testing hypotheses regarding which hurdles are the most significant obstructions to progress. Experts who recently departed from government service discussed their first-hand experience addressing sources of friction and candidly described both the challenges and why previous attempts did not succeed. The workshop followed Chatham House rules, which capture the deep expertise of participants but grant them the anonymity to speak freely; thus, people are cited by number rather than name in the footnotes.

[4] Interviewee #2.

[5] Interviewee #3.

Introduction

OSINT Opportunities for the IC

"The story of the past five years is that the underlying math and finances of trying to build bespoke widgets for the high-side environment got absurd."—interviewee[6]

Industry is gleaning astonishing insights from a combination of publicly available information and AI/ML tools. Organizations such as Bellingcat have built a reputation as OSINT wizards. Using publicly available information,[7] they have discovered illegal shipping of chemical weapons precursors, identified a high-ranking Russian intelligence officer as a key suspect in the shooting down of Malaysian Airlines Flight 17, and identified Russian intelligence officers as suspects in the poisoning of Sergei and Yulia Skripal.[8] Unclassified data sources hosted in cloud environments have mapped supply chain issues for weapons systems and critical infrastructure.

One interviewee cited a company that had bought and effectively curated anonymized cell phone data, then used the geolocation data and public records to identify the phones' owners—the kind of

[6] Interviewee #5.

[7] PAI is defined by the Department of Defense's (DOD) Manual 5240.01 as "Information that has been published or broadcast for public consumption, is available on request to the public, is accessible on-line or otherwise to the public, is available to the public by subscription or purchase, could be seen or heard by any casual observer, is made available at a meeting open to the public, or is obtained by visiting any place or attending any event that is open to the public. Publicly available information includes information generally available to persons in a military community even though the military community is not open to the civilian general public."

[8] "About," Bellingcat, September 29, 2021, https://www.bellingcat.com/about/.

work the IC has used classified information to do for decades.[9] The *New York Times* Privacy Project in 2019 tracked President Trump's entourage using location-tracking data from apps.[10] In space, private companies are replicating what only nation-states could do; Planet can image any point on Earth up to 12 times a day. That data is uploaded to the cloud, where computer-vision algorithms can characterize, baseline, and then automatically detect changes, extracting insights and building trend data.[11] In August 2021, *The Economist* remarked on the possibilities from OSINT: "There are websites which track all sorts of useful goings-on, including the routes taken by aircraft and ships. There are vast searchable databases. Terabytes of footage from phones are uploaded to social media sites every day, much of it handily tagged."[12] *The Economist* pointed to using Denis Rodman to measure the size of a hydrogen bomb. Rodman's height is a known quantity; photographs of Rodman standing next to Kim Jong Un give a reliable measure of the latter's height, and thus a reliable estimate of Kim's head; therefore the many photographs of Kim standing next to North Korea's nuclear arsenal allow estimates of the sizes of these weapons. Taking the capability a step further, one could task an AI/ML system with finding pictures of Kim's face near objects that the machine could be trained to recognize as weapons, then giving an estimate of their size. Perhaps the machine could even alert a human when there is a sudden miniaturization or expansion in size of the arsenal.

Taking advantage of OSINT and an unclassified cloud, analysts can deliver fast, attributable, relevant reporting from anywhere, including austere field environments and their homes. The Covid-19 pandemic forced many agencies to reevaluate what was truly a necessary security posture and what precautions were being taken by default. The National Geospatial-Intelligence Agency (NGA), for example, sent its workforce home and discovered they were able to conduct much of their business on an unclassified level. While exquisite, highly classified reporting can be invaluable, OSINT can provide the quick, detailed, and deep reporting that builds trust between intelligence professionals and policymakers.

Yet, despite these advances, the intelligence community has not yet warmed to OSINT, with potentially disastrous effects. For the intelligence community to meet its mission of "all-source" analysis, it cannot afford to ignore a wealth of available data solely because it is unclassified. In a best-case scenario, the IC will lose policymaker attention and trust as they compete with private intelligence. But the worst-case scenario is more problematic: U.S. adversaries are pursuing this same technology aggressively and outstripping IC capabilities. While China has an estimated 100,000 science and technology intelligence officials, who do most of their work on open sources, the United States has an estimated 100.[13] Part

[9] Interviewee #4.

[10] Stuart A. Thompson and Charlie Warzel, "How to Track President Trump" *New York Times*, December 20, 2019, https://www.nytimes.com/interactive/2019/12/20/opinion/location-data-national-security.html.

[11] Interviewee #4; Martin Van Ryswyk, "Planet's New Rapid Revisit Platform to Capture up to 12 Images Per Day," Planet , June 9, 2020, https://www.planet.com/pulse/12x-rapid-revisit-announcement/; and "Data Empowering the World's Decision Makers," Orbital Insight, https://orbitalinsight.com/.

[12] "Open-Source Intelligence Challenges State Monopolies on Information," *The Economist*, August 7, 2021, https://www.economist.com/briefing/2021/08/07/open-source-intelligence-challenges-state-monopolies-on-information?giftId=7f75b253-f29f-4827-9121-8736a3cbd5fe.

[13] Wm. C. Hannas and Huey-Meei Chang, China's STI Operations: Monitoring Foreign Science and Technology through Open Sources (Washington, DC: Center for Security and Emerging Technology, January 2021), https://cset.georgetown.edu/wp-content/uploads/CSET-Chinas-STI-Operations.pdf.

of the reason for the IC's resistance to developing OSINT is a dependence on an old definition of what makes up "open-source" intelligence. This report will posit several old definitions and suggest a new approach to shift the IC's thinking on what is a valuable intelligence use of OSINT.

1

What Is OSINT, and What Should It Be?

"Data is the lifeblood of intel."—interviewee[14]

The intelligence community has long struggled with the idea of OSINT. Intelligence agencies, after all, were created to steal secrets and to discover what the enemy is attempting to hide. The Department of State gathers information that a particular nation wants Washington to see; the news media gathers information that the public wants to know. For decades, OSINT meant press reporting, which analysts were supposed to be aware of and incorporate into their work, but mostly as supporting information for highly classified reports. In the CNN era, OSINT was also "CNN-INT." CNN was live on the ground and reporting the first edges of a story before signals intelligence (SIGINT) could be processed or human-intelligence (HUMINT) sources could report back. Instructions for Central Intelligence Agency (CIA) analysts were always, "Don't try to be CNN"—in other words, writing anything in a President's Daily Brief (PDB) or World Intelligence Review (WIRe) publication that CNN had already broadcast to the world 12 hours earlier was going to be unhelpful for policymakers and embarrassing for the agency. Instead, analysts should seek to add value to the breaking news, such as by adding expert commentary stemming from a deep understanding of the political context or providing classified reporting that illuminates—or even contradicts—CNN's original story.

Today, researchers are swimming in a sea of publicly available data that reaches far beyond cable news reporting, and this definition of OSINT is outdated. A revised definition now needs to encompass a wide range of information—from the VKontakte feed of a Russian operative to the photos a tourist posted on Twitter of a square just before a bombing to advertising data collected by a Chinese company.

[14] Interviewee #10.

A review of classic definitions of OSINT provides a starting point. Some of these definitions prove forward-looking, and their key elements can inform a new framework:

- The Office of the Director of National Intelligence (ODNI) defines OSINT as "publicly available information appearing in print or electronic form including radio, television, newspapers, journals, the Internet, commercial databases, and videos, graphics, and drawings."[15] While comprehensive, this definition reads like it was written to cover all eventualities in anticipation of a future legal challenge—particularly its reference to anything appearing on the internet.
- The Department of Defense (DOD) Strategy for Open-Source Intelligence highlights the purpose of the intelligence produced: "Open-source intelligence (OSINT) is intelligence that is produced from publicly available information and is collected, exploited, and disseminated in a timely manner to an appropriate audience for the purpose of addressing a specific intelligence requirement."[16]
- The RAND Corporation defines OSINT as "publicly available information that has been discovered, determined to be of intelligence value, and disseminated by a member of the IC."[17] This definition skips over the processing that must take place to fully capitalize on publicly available data such as sound data.

All these definitions grasp that intelligence information can come from non-clandestine sources, but few acknowledge that OSINT goes beyond press reporting. Companies such as Planet Labs and HawkEye 360 have stretched capably into territory that once belonged only to wealthy nation-states. Other companies advertise private analysis services, and some, such as BP, claim to incorporate human sources. NSO Group, out of Israel, recently made headlines for selling foreign intelligence services the capability to hack into any phone with nary a click from the target.[18] This kind of capability blurs the line between clandestine and open information. We may assume, however, that if data is available to legally purchase, most likely the United States' adversaries have access to it, as do any number of corporate entities. It should be counted as open source, which would require expanding the definition further.

For the purposes of this paper, OSINT is *intelligence collected from publicly available or available-for-purchase information, obtained for addressing a specific intelligence requirement, and processed to derive new insights.*

[15] "What Is Intelligence," Office of the Director of National Intelligence, https://www.dni.gov/index.php/what-we-do/what-is-intelligence.

[16] "Responsibilities of Secretary of Defense Pertaining to National Intelligence Program," United States Code, 2011 Edition (U.S. Government Publishing Office, 2011), https://www.govinfo.gov/content/pkg/USCODE-2011-title50/html/USCODE-2011-title50-chap15-subchapI-sec403-5.htm.

[17] Heather J. Williams and Ilana Blum, *Defining Second Generation Open Source Intelligence (OSINT) for the Defense Enterprise* (Santa Monica, CA: RAND Corporation, 2018), https://www.rand.org/pubs/research_reports/RR1964.html.

[18] Patrick Howell O'Neill, "NSO Was about to Sell Hacking Tools to France. Now It's in Crisis.," MIT Technology Review, November 24, 2021, https://www.technologyreview.com/2021/11/23/1040509/france-macron-nso-in-crisis-sanctions/.

2

OSCAR'S Grandfather

IC Use of the Cloud, AI/ML, and OSINT Today

The intelligence community has taken initial steps toward adopting each of the technologies required to make OSCAR a reality. While the IC has had an on-again, off-again relationship with OSINT and its adoption of AI/ML is uneven, it *has* committed to using cloud resources. It has used commercial cloud services (C2S) on the Secret and Unclassified fabrics since 2017.[19] In 2019, *ODNI's Strategic Plan to Advance Cloud Computing in the Intelligence Community* said:

> The IC's cloud capabilities will support a diverse set of users to include disconnected or edge operations. These capabilities will provide innovative and contemporary technologies such as artificial intelligence (AI), machine learning (ML), and high-performance computing to meet current and future needs. These capabilities will require unified security processes and acceptance that enable quick adoption and portability of applications, data, and code. The IC will leverage these capabilities in an approach that favors vendor flexibility, simplifies use and adoption of new and cloud-native technologies, and promotes necessary culture changes.

Regarding AI/ML, elements of the intelligence community are using computer vision to tip and cue collection, in particular at NGA.[20] While OSINT has always been used in intelligence analysis, many within the IC recognize the value of publicly available data sets and have run the gauntlet of lawyers and contracts to try to procure the most critical ones.

[19] Office of the Director of National Intelligence, *Strategic Plan to Advance Cloud Computing in the Intelligence Community* (Washington, DC: Office of the Director of National Intelligence, June 2019), https://www.dni.gov/files/documents/CIO/Cloud_Computing_Strategy.pdf.

[20] Workshop participant #6.

The Biden administration has also made clear that it sees data and data analysis as strategically important. Deputy Secretary of Defense Kathleen Hicks, in a speech at the DOD's Artificial Intelligence Symposium and Tech Exchange in June and in a memo to the workforce on "Creating Data Advantage" in May 2021, said the department will forward-deploy operational data teams (ODTs) to each combatant command's data office in order to scale existing efforts and make data "visible, accessible, understandable, linked, trustworthy, interoperable, and secure."[21]

[21] Terri Moon Cronk, "Hicks Announces New Artificial Intelligence Initiative," U.S. Department of Defense, June 22, 2021, https://www.defense.gov/News/News-Stories/Article/Article/2667212/hicks-announces-new-artificial-intelligence-initiative/; and Kathleen Hicks, "Memorandum for Senior Pentagon Leadership," issued on June 21, 2021, https://media.defense.gov/2021/May/10/2002638551/-1/-1/0/DEPUTY-SECRETARY-OF-DEFENSE-MEMORANDUM.PDF.

3

Obstacles to OSCAR

Despite several successful pilot programs and high-level commitment to adoption, the IC has still lagged in adopting the combination of OSINT, cloud computing, and AI/ML tools that will propel intelligence work in the coming decade. During CSIS's literature review and expert interviews, seven hypotheses emerged regarding the biggest obstacles to general acceptance of this technology:

1. **It's a policy problem.**
   - Acquisition policy is too slow, and requirements are too strict. Laws, regulations, and requirements focus on technological specifications instead of desired outcomes.
   - Budget cycles are ponderous and too rigid to account for innovation.
2. **It's a messy-data problem.**
   - IC data is not standardized, and much of it is not usable in its current form. The up-front expense and time required to update, model, and move data is a major impediment to adopting new technology that could capitalize on legacy data.
3. **It's a security problem.**
   - IC security officers have unresolved concerns about unclassified cloud and off-the-shelf AI/ML applications—despite cloud providers' assertions that their platforms are more secure than agency-owned servers because of their ability to rapidly identify and patch problems and their large and up-to-date security teams.
4. **It's a bespoke-requirements problem.**
   - The IC says it wants to adopt off-the-shelf cloud and AI/ML applications but then adds many

specific, additional requirements. Those additional needs not only wipe out much of the cost savings and limit the functionality of the cloud environment, but they also make it difficult for smaller tech firms to adapt commercial off-the-shelf (COTS) products to compete.

5. **It's a business-processes problem.**
   - The U.S. government is not designed to change quickly, and quick change is essential to keeping up with AI/ML. By the time the process of granting authority to operate (ATO) is complete, the state of the art has moved on.
6. **It's a congressional-oversight problem.**
   - Congress does not allow "wasting" money on "failure," and the IC assesses that incorporating unclassified cloud and AI/ML applications is high-risk.
7. **It's an amorphous culture problem, and none of the above adequately express the cause.**

Researchers convened a workshop to test these hypotheses further. At the outset, participants voted on which of the seven hypotheses represented the most important problem to tackle. They selected the culture problem as the most critical issue, followed by security, then a tie between policy and speed, which we combined into an exploration of policy, including those policies that slow down the acquisition process. Researchers broke the participants into three discussion groups, and each group tackled one of these problems.[22] Groups then reconvened to share findings with the rest of the participants and continue the conversation.

One overarching takeaway emerged from the discussions: The U.S. government is adopting only the technology that fits within the current system, but the current system discourages risk-taking and rapid innovation. The government likely needs to rework the acquisition and adoption process—or even develop a parallel one—so it can circumvent many legacy requirements and adapt to revolutionary technologies.

The Culture Problem

"What's missing is the discussion about people. It's all about the humans: empowering them and incentivizing them to transform public servants from rule-bound to problem-solving."[23]*—Stephen Goldsmith*

Interviewees and workshop participants indicated the culture obstacle is rooted in a lack of demand signal from analysts and operators and in an incentive structure that rewards narrowly defined success, disincentivizes risk-taking, and encourages normal-course-of-business work.

The lack of a demand signal is itself made up of two components. First, at the working level, there is a lack of recognition of what OSINT and AI/ML can accomplish and how they can revolutionize intelligence work; second, there has yet to be an attention-focusing crisis that creates the sense of urgency necessary to break through organizational barriers.

[22] The third group focused largely on policy, viewing speed as a symptom of the policies in place to acquire and adopt technology.

[23] Stephen Goldsmith and Neil Kleiman, *A New City O/S: The Power of Open, Collaborative, and Distributed Governance* (Washington, DC: Brookings Institution Press, 2017), https://www.brookings.edu/book/a-new-city-os/.

RECOGNITION OF THE NEED: HOW CAN OSCAR HELP?

"No one can do perfect like we can."—former IC member, discussing the "not made here" mindset[24]

Intelligence agencies have exquisite toys, from covert communications capabilities to disguised widgets for intelligence collection. These capabilities are often created to fulfill a particular mission, and as such, are perfect for the moment. This bespoke perfection can lead to a "not made here" bias and a lack of recognition that competition exists for some IC functions. One workshop participant said, "[There is] not a widespread recognition that intel agencies' capabilities in AI/ML and large-scale analysis have been outstripped by [the] private sector. In the mid-2000s, [the IC] were the best, but not anymore."[25]

In general, line analysts and operators do not understand the labor-saving possibilities of AI/ML, so there is limited to no demand signal from the front lines or middle management. One workshop participant said, "Demand for these capabilities can come from different places: top-down, which is less effective; but it's more effective to come from the mission level ('Damn it, I need this to do my job')."[26] A major constraint is the very human limitation of not having enough time to learn a new skill and develop the habit of using it, in particular, when under pressure to produce "current intelligence" products.[27] Analysts have no time to figure out a new AI/ML software system when a PDB is due in a few hours.

Ideally, a data scientist would be embedded in each analytical and operational unit to demonstrate how AI/ML can be a force multiplier when collecting and processing OSINT, but a dearth of available talent means analysts and operators rarely get to see the proof of concept. Instead, large quantities of data and information technology (IT) talent are occupied maintaining legacy systems, some of which have outlived their usefulness, because of a combination of familiarity and inertia on reallocating those resources. "You typically can't retire those systems until legacy workers leave," said one workshop participant.[28] A stopgap measure, further explored in the recommendations section, would be to train incoming intelligence officers on the benefits of OSCAR and cycle existing intelligence officers through additional training to create a new demand signal. Reaching for OSCAR should be second nature when writing a PDB or asking a new in-depth research question.

SENSE OF URGENCY: WHY DO ANALYSTS NEED OSCAR NOW?

"Unless there is someone senior, there is no way to get it done. [We] attempted to do it at the 05/06 [Lt. Colonel/Colonel] level, but nothing can get done unless someone with stars on their shoulder forces it through the system."—workshop participant[29]

Senior IC and DOD leadership have given speeches and repeatedly said that it is critically important to adopt low-side cloud-computing capabilities to make better use of AI/ML, but then initiatives

[24] Workshop participant #8.

[25] Workshop participant #1.

[26] Workshop participant #9.

[27] IC professionals refer to "current intel" when discussing short-fuse pieces like the PDB or WIRe.

[28] Workshop participant #10.

[29] Workshop participant #11.

either never make it out of the pilot phase or never scale up. These leaders have the best of intentions, but they cannot devote the consistent attention required to push initiatives over the goal line. The lack of attention from on high and demand from below result in no sense of urgency to solve the problem, and the middle managers who are generally engines of productivity also lack the technical understanding to push the change needed.

Several interviewees and workshop participants highlighted the difficulty of accomplishing even initial-phase incorporation of new technologies without either the consistent attention and support of senior officers or a forcing event such as a national security crisis. One participant said, "When something is new and paradigm-breaking, you have to protect it and nurture it, not bury it somewhere. If anything, we are seeing centers of excellence bury [AI/ML projects] further down in their organizations and not putting any more budget in than when they started."[30] Often, adding funding would require removing funding from existing, proven systems—anathema to a government bureaucracy. The press of day-to-day business is too great for OSCAR to garner high-level attention, and, luckily, an acute crisis has not emerged to drive it to the top of the agenda.[31] There is, however, a slow-fuse crisis bubbling, wherein the United States' competitors are outstripping the intelligence community in adopting such technology.

ACCEPTING RISK

"The perpetuation of legacy is the greatest barrier to change and progress."—workshop participant[32]

The U.S. defense and intelligence communities are focused on being the first line of defense and making truly life-or-death decisions, resulting in an ethos of a no-fail mission. The perception—although the reality should be tested—is that promotions are awarded based on who can minimize risk and succeed in mission or project execution. Risk-taking and failure are not conducive to rising in the ranks, so getting a flag officer or deputy director to devote attention to a moonshot is less likely; those officers likely rose to their positions by making careful risk calculations.

To take risks, those executing a project need top cover—preferably an explicitly stated acceptance of risk—from every level of their supervisory chain. For contract officers, in particular, flexibility and risk-taking are not part of the job description. Judging from interviewees' and workshop participants' first-hand views, implementers perceive that cover is lacking. One workshop participant said that "lots of good ideas die because of comments from budgeting officers that Congress is going to hate an idea."[33] Contracts that finish under budget and on time, with deliverables that match exactly the specifications in the initial contract, count as successes for contracting officers. Creating a new mindset of contracts that can adapt to cutting-edge developments and take risks on new technology may require breaking old, bad habits and debunking perceptions about risk acceptance.

[30] Interviewee #12. Also see: Melissa Flagg and Jack Corrigan, "Ending Innovation Tourism: Rethinking the U.S. Military's Approach to Emerging Technology Adoption," *CSET Policy Brief,* Center for Security and Emerging Technology, July 2021, https://cset.georgetown.edu/publication/ending-innovation-tourism/.

[31] One participant said that Russian interference in the 2016 election was our "9/11 moment" for open-source intelligence, but "we missed it" because of the political maelstrom surrounding the allegations.

[32] Workshop participant #10.

[33] Workshop participant #19.

Analysts, too, have an incentive to operate in the realm of the normal and familiar. A former intelligence officer who participated in the workshop recalled their supervisors explicitly releasing their team from the normal set of performance requirements in order to make space to pioneer activity-based intelligence, which at the time was revolutionary for its tipping and cueing potential: "We were empowered to do our jobs [because we] did not have the burden of production numbers . . . , empowered to create relationships across the IC, to have our own deployments. When leadership removes the burden of production, that's when . . . analysts are empowered to think about hard problems."[34]

The current IC and DOD ecosystem—which includes acquisition professionals, security standards, and industries—creates a self-reinforcing culture. As one workshop participant put it, "There is a security infrastructure that jealously protects what it does. The acquisition community is built on that and is resistant to change because it could impact their bottom line."[35] In other words, everyone currently in the system benefits from the stability of the system. Industry insiders have cracked the code and are profiting from the current situation, so they have no incentive to change. Security teams and acquisition chains gravitate toward the proven and the familiar. Thus, the pressure to change is from the outside, particularly from smaller companies yelling in frustration at navigating the gates to entry. Put bluntly, "Antibodies come into the system for a new thing."[36]

The Security Problem

IC elements have been burned by public disclosures of sensitive information—from Edward Snowden's leak of National Security Agency (NSA) data to WikiLeaks' publishing of the Vault 7 files on CIA cyber activities—leaving security officers highly reluctant to expand access to IC activities. They are also constantly facing aggressive, broad-scope espionage attempts from nation-states—including the Russian hack of SolarWinds and thousands of cyberattacks by Chinese hacking campaign Hafnium, to offer only two recent examples.[37] Security officers are deeply committed to protecting the intelligence community, and their incentive structures are all geared toward extreme caution. This approach can be jarring for those in industry, or even former IC officers who transitioned into industry. One participant said, "Despite being partners, the IC forces industry to go through an antiquated [vetting] process because there is a trust deficit. AI cannot be built in the government, and there needs to be better communication between industry and [the] IC based on a world of trust."[38] Another participant described the hurdles for obtaining authorization to operate (ATO) as "too high, too slow, too complicated;" achieving an ATO can take 18 months in the CIA's Directorate of Analysis, and there

[34] Workshop participant #7.

[35] Workshop participant #2.

[36] Workshop participant #1.

[37] Isabella Jibilian and Katie Canales, "The US is readying sanctions against Russia over the SolarWinds cyber attack. Here's a simple explanation of how the massive hack happened and why it's such a big deal," Business Insider, April 15, 2021, https://www.businessinsider.com/solarwinds-hack-explained-government-agencies-cyber-security-2020-12; and "30,000 U.S. Organizations Breached by Cyber Espionage Group Hafnium," *Security Magazine*, March 9, 2021, https://www.securitymagazine.com/articles/94781-000-us-organizations-breached-by-cyber-espionage-group-hafnium.

[38] Workshop participant #18.

is very little oversight of security officers' decisions to grant or deny it.[39] A third participant said his company is on month 36 of an ATO process.[40] Running this security gauntlet to gain trust can be lucrative for a large company that can afford to spend two years in the initial process, but for smaller companies or those with cutting-edge technology, navigating this process makes no business sense.

The security obstacle is composed of two main issues: comfort in classifying data and processes and a lack of understanding of the security features of new technology.

THE CLASSIFICATION SECURITY BLANKET

A major obstacle to using unclassified cloud architecture is concern about adversaries being able to discern IC priorities from queries of the data. China, for example, could surveil analysts' activities on the open side and understand what questions they are asking, thus deriving their intelligence gaps, which would inform their denial and deception efforts and make U.S. analysts' collection more difficult. Industry partners, however, report that this perception of a security weakness is overstated, given the security features of the cloud, the possibility of submitting obfuscating queries, and the likelihood that the United States' adversaries already know or can guess most intelligence priorities without this additional surveillance.[41] Put differently, an interviewee said, "The IC is uncomfortable working with open-source vendors because, by targeting certain data sets, it becomes clear what their interests are. But is that much of a secret anyway? Do images of Russian bases need to be classified or not? Does their collection in open-source [platforms] hinder what the IC can do with them?"[42] Another workshop participant relayed his company's experience: the IC funded an unclassified program, but security vetoed the implementation because of concerns that outside actors could see the queries of the data.[43]

Classifying analysts' questions is comfortable and easy and minimizes risk, and security officers rarely see the downside. To shift work into the unclassified realm, it will be essential for top leadership to change the conversation about risk and priorities. General Robert Ashley, the former director of the Defense Intelligence Agency, wrote in Defense One about one of his efforts to make battlefield data more accessible: "Operational commanders made it policy that captured material would remain unclassified until such time as it was combined with other sensitive data."[44] Adopting this mindset would help the intelligence community use publicly available information much more effectively.

The sheer volume of data running into these security restrictions may be the tipping point for the acceptance of unclassified cloud architecture. One workshop participant said, "It's impossible to bring all PAI [publicly available information] behind the firewall for processing, and even doing large portions of it is cost-prohibitive. We must figure out how to work with open data in the open . . .

[39] Workshop participant #8.

[40] Workshop participant #11.

[41] Workshop participant #2.

[42] Interviewee #4.

[43] Workshop participant #11.

[44] Robert P. Ashley, "Ten Years after Bin Laden, We Still Need Better Intelligence Sharing," Defense One, May 1, 2021, https://www.defenseone.com/ideas/2021/05/ten-years-after-bin-laden-we-still-need-better-intelligence-sharing/173748/.

It's much easier to apply tradecraft solutions than invent new math and computer science for this challenge."45 In other words, a wealth of data exists in the world that will help the United States confront its toughest security challenges. The U.S. intelligence community can either take a small risk and operate in the unclassified space or pursue a fool's errand in trying to move the entire internet onto the high side, then search it with antiquated tools.

CLOUD QUESTIONS: OSCAR IS MY COPILOT

The trust deficit particularly emerges when attempting to convey the security features incorporated in cloud architecture. One interviewee used an apt analogy for the difference between using legacy, on-premises, classified servers and adopting the cloud: "[It's like] flying in a small, single-engine Cessna with someone you personally know and trust versus flying in an Airbus 380 and having instant trust with someone you haven't met. The difference with [the] cloud is that you just need to understand the cockpit. Once you do, you realize you actually have much more control and safety than you used to."[46] The same participant used another apt example: "The security processes being used today [to evaluate] cloud services and features in the IC are akin to attempting to perform a tune-up on a 2021 Tesla using the same tools designed for a 1968 Camaro."[47]

Overcoming skepticism of the cloud will largely entail education at both senior and middle-management levels, but those conversations take time and need to be done in private because, as one interviewee put it, "leaders are embarrassed by how much they don't know."[48] A former senior intelligence official said, "They're still working on the assumption that the internet is insecure but don't have a good understanding of the security investments that have gone into the cloud. Everyone needs to better understand the walls and barriers available. A lot more education needs to go on regarding the security side of the cloud and teaching how it works."[49]

Security Features of the Cloud

The security features available in the cloud are better than the perimeter-based security that exists in most of the IC and DOD today. Traditional on-premises systems are a combination of products and services in which one weakness can be exploited and corrupt the whole system—often serviced by traditionally understaffed IT departments.[50] The cloud, on the other hand, is built and monitored by large staffs dedicated exclusively to security using a wide array of automated services and sophisticated features. An interviewee from a cloud services provider said his company's cloud is "by far the most secure system I've ever worked on. It's all zero-trust architecture and device-level. Perimeter defense is foolish today."[51]

[45] Workshop participant #13.

[46] Interviewee #10.

[47] Ibid. Note that the on-premises, high-side cloud referenced is for classified information. This paper focuses on the opportunities in using unclassified cloud, but the analogy and security reassurances still hold.

[48] Ibid.

[49] Interviewee #14.

[50] Interviewee #15.

[51] Interviewee #4. The service provider was Google.

Several interviewees pointed to the ability to use AI/ML on the unclassified cloud to look for anomalies in access patterns and learn from previous attacks, further bolstering security measures. Then, once security vulnerabilities are discovered, applications and services running on cloud platforms can be patched quickly: "Security models are achieved by patching at speed, and the only organizations who can't do that are the U.S. DoD and IC."[52]

Industry, however, should be prepared for a slew of security requirements, if it is not already. "As the government moves to a service model, all proxy vendors must be considered a legitimate target of our adversaries. Companies should get used to more intrusive expectations from the government."[53]

The Policy Problem

Culture has its roots in policy. Agency policies elucidate what is possible and create friction for decisionmakers, guiding their decisions and behavior. As organizations add policies over time, often as a reaction to a problem, they can interlock in unanticipated ways, like closely packed gears preventing each other from turning. For example, combining policies designed to ensure measurable competition with those attempting to prevent security breaches can result in overly specific contract requirements, which prevent new entrants to the market. Policies exist to encourage pilot programs, but crossing the "valley of death" from the pilot phase to full availability means navigating a maze of requirements and institutional inertia, including locating new funding streams.[54] As one interviewee said, "What typically happens is that someone runs a pilot, users love it, then it dies quickly because it can't be hosted, and incumbent embedded legacy companies resist it and seek additional scope in the form of full-time equivalent staff to build a customer high side app."[55]

Workshop participants and interviewees pointed to specific policies as obstacles to incorporating new technologies, but they also described dead-end outcomes where the cause was less clear, perhaps pointing to a systematic failure. Specific problematic policies are described immediately below. However, the second category is indicative of a larger ill, and more research will be required inside the DOD and IC acquisition communities to identify obstacles and opportunities to remove them. For now, such outcomes are outlined below as symptoms, rather than root causes.

WELL-INTENTIONED POLICIES THAT BECOME BARRIERS

Interviewees identified the following set of policies as barriers to rapid incorporation of technology such as OSCAR in the intelligence community:

- **Intractable contracts**. The U.S. government errs on the side of specificity in contracts, with an eye toward ensuring accountability and fairness in awarding them. However, exacting requirements and delivery schedules do not allow for incorporating off-the-shelf capabilities or for rapidly updating when the state of the art evolves. Prescriptive requirements also mean each contract creates a one-off deliverable for a client, rather than meeting a larger demand, disincentivizing

[52] Interviewee #12.

[53] Interviewee #15.

[54] Interviewee #11.

[55] Interviewee #10.

commercial providers from working with the government. For example, the NSA's WildandStormy cloud computing contract had intensive requirements, requiring significant adaptions to an otherwise off-the-shelf product.[56]

- **Rigid budget processes.** DOD's planning, programming, budgeting, and execution (PPBE) process takes place years in advance.[57] Corin Stone, in her series of articles on AI/ML in the intelligence community, described in detail the challenges rigid budgeting presents to acquiring AI/ML.[58] In her words, the PPBE is "a methodical and deliberate three-year budget cycle that calls for defined and steady requirements at the beginning of the cycle." This type of planning may be helpful for procuring items with long lead times, such as key durable components of submarines and aircraft carriers, but predicting cutting-edge developments for a young technology like AI/ML several years in advance is nearly impossible. She explains that the semi-permanent "color of money" is an additional challenge: "The IC's budgeting processes require that IC spending fit into a series of discrete sequential steps, represented by budget categories like research, development, procurement, or sustainment." These linear buckets assume a slow, predictable progression of the technology, rather than rapid advancements. In other words, AI/ML "development" actually looks very similar to AI/ML software "sustainment." Product-based statements of work issued at the earliest stages of soliciting bids will be quickly outstripped by advancements in this rapidly changing area. Together, setting standards early in the process and strictly adhering to budget categories handcuff innovation.
- **Rigid budget definitions.** Contract officers struggle to create contracts with "as a service" deliverables that can be purchased as subscriptions, preferring instead to enter into long-term "per head" licensing agreements, which perpetuates level-of-effort contracts for general systems engineering and technical assistance with limited measures of performance. Similarly, because the cloud model of "pay only for what you use" can provide detail on actual expenses, AI/ML to advance specific missions can be budgeted and tracked as operations and maintenance instead of capital expenses, an approach that would allow flexibility to quickly meet mission need but does not comport with acceptable practice for budget officers.[59]
- **The authorization to operate (ATO) maze**. Vendors routinely lament the slow, confusing process involved in securing an ATO. According to the General Services Administration (GSA), it takes a minimum of 18 weeks and a myriad of reviews to go from request to ATO for new software; but industry partners report a more realistic estimate is years. The Federal Risk and Authorization Management Program's (FedRAMP) 21-page guide to securing an ATO includes 27 separate acronyms. A 2017 study found the average cost to obtain a FedRAMP certification was $350,000–

[56] Workshop participant #16.

[57] "Planning, Programming, Budgeting & Execution Process (PPBE)," Defense Acquisition University, https://www.dau.edu/acquipedia/pages/ArticleContent.aspx?itemid=154.

[58] Corin Stone, "Artificial Intelligence in the Intelligence Community: The Tangled Web of Budget & Acquisition," Just Security, October 8, 2021, https://www.justsecurity.org/78362/artificial-intelligence-in-the-intelligence-community-the-tangled-web-of-budget-acquisition/.

[59] Interviewee #10.

$865,000, which it described as good news, given previous estimates of $1 million or more.[60] Even the low end of that estimate is an eye-popping price tag for a startup. For government employees sponsoring a cloud service provider's candidacy, FedRAMP's slides include this foreboding caution: "Confirm a CSP's dedication to taking on the FedRAMP authorization process. Clearly outline the level of effort involved."[61] The clear subtext is that this is an arduous process, and all parties need to be clear-eyed about the level of work and commitment required over the next four to five months. Even after the FedRAMP process is completed, the provider may also need to obtain an agency-specific ATO. The GSA also says it prefers a "do once, use many" approach—that is, to reuse security assessments to save time and resources. This seems a sound policy, given the investment required to secure ATOs, but has the side effect of funneling opportunities to well-established players and away from smaller startups.

- **A lack of interoperability.** Contractors usually prefer to build support systems from scratch, be they computers, data, or apps.[62] There can be good reasons, such as a lack of insight into the systems an agency is already using or a desire to protect proprietary information. However, too often the result is siloed data. Breaking those silos into a cloud-based data lake exploitable by AI/ML will be costly and labor-intensive.[63]
- **Outdated destruction rules.** Controls and standards regarding media destruction were designed for physical hardware in the control of an agency or department. In a cloud environment, cryptographic destruction renders information inaccessible more effectively.[64] IC policies, many of which were written with the security of physical systems in mind, should be reviewed and updated to account for cloud infrastructure.

When Absence of Policy Is the Problem

While participants were able to pinpoint a collection of policies that definitively obstruct AI/ML adoption, a key takeaway was the problem of a lack of a policy. As one participant put it, "When we say we have an interoperability problem, it's typically a policy or data-model problem, nothing to do with the technology."[65] Different agencies, or even different elements within the same agency, use different data models. Making that data usable for AI/ML applications means cleaning and sorting data

[60] Jason Miller, "New Report Tries to Bust Fedramp Myths about Cost, Usage," Federal News Network, May 8, 2017, https://federalnewsnetwork.com/cloud-computing/2017/05/new-report-tries-to-bust-fedramp-myths-about-cost-usage/.

[61] FedRAMP, *FedRAMP Agency Authorization Playbook (Version 2.0)* (Washington, DC: FedRAMP, October 2021), https://www.fedramp.gov/assets/resources/documents/Agency_Authorization_Playbook.pdf.

[62] Workshop participant #10.

[63] For more on interoperability, see Emily Harding, McKenzie Richardson, and Matt Strohmeyer, "From Data to Insight: Making Sense out of Data Collected in the Gray Zone," Center for Strategic and International Studies, *Commentary*, October 21, 2021, https://www.csis.org/analysis/data-insight-making-sense-out-data-collected-gray-zone.

[64] Interviewee #16. Cryptographic destruction of data relies on the idea that the data has been encrypted and owners have, in effect, thrown away the key, rendering it impossible to access. However, given quantum computing's theorized capability to undo much of today's encryption, it is worth exploring whether this form of data destruction will stand the test of time.

[65] Interviewee #15.

to a common standard.[66] This will be expensive and time-consuming, but starting such an undertaking without a policy on data standards is a recipe for wasting money. Similarly, there is no test, evaluation, validation, and verification policy within DOD or the IC for AI/ML. Many pilot projects have been prevented from effectively scaling because there is no process to assure their functionality.[67]

Given the challenges, expense, and cultural change needed to achieve a robust, AI/ML-enabled OSINT capability, some have argued the IC should leave OSINT to the private sector, and whatever the market will bear. However, this mindset ignores two key factors: the IC's all-source mission and the trust relationship between intelligence agencies and the policy community. First, the IC seeks facts, wherever they may exist. The mission is to gather those facts and then to use deep expertise to put them in context to create insight, like indications and warnings. The IC should not discount facts because they are unclassified. Second, for the same reason that the IC does not take the *New York Times* at face value and republish it as an intelligence product, the IC should not let off-the-shelf commercial analysis be a permanent stand-in for an in-house OSINT capability. The IC has a reputation for high analytic standards, unbiased and tailored to policymaker questions, and should bring those high standards to OSINT analysis as well. Those high standards and ability to show the work is what builds the trust relationship with policymakers.

[66] Interviewee #16.

[67] Interviewee #10.

4

Possibilities Ahead

"Game changer."—John Edwards, chief information officer at the CIA, discussing adopting commercial cloud services at the AWS Public Sector Summit in 2019.[68]

Cultural, security, and policy challenges are myriad but surpassable, particularly if the intelligence community decides AI/ML capabilities are mission-critical and firmly establishes a vision for how a capability like OSCAR can be a game-changing partner for IC officers. We asked workshop participants to imagine the horizon of cloud-based, AI-enabled intelligence—how OSCAR might become reality. Hollywood always reflects the leading edge of imagination, so the model is Tony Stark's JARVIS: a computer people can talk to, ask questions of, and get instantaneous and easily understandable results.

OSCAR would likely start as adoption of "scrape and search" COTS capabilities installed on an unclassified cloud. Once the U.S. government puts enough money and labor into grooming and training data sets, OSCAR can evolve into an AI/ML application able to "read" and sort data, make basic decisions, and—most importantly—tip and cue analysts to focus on the most concerning data points. In the long run, OSCAR may surpass JARVIS, incorporating quantum computing to crunch massive data sets instantaneously.

[68] John Edwards, "AWS Public Sector Summit 2017 Customer Keynote: John Edwards, Central Intelligence Agency," Amazon Web Services, YouTube, June 21, 2017, 12:20 to 13:01, https://www.youtube.com/watch?v=DEc6kVAXSs8&t=744s&ab_channel=AmazonWebServices.

Time Horizon of AI/ML Capabilities on an Unclassified Cloud

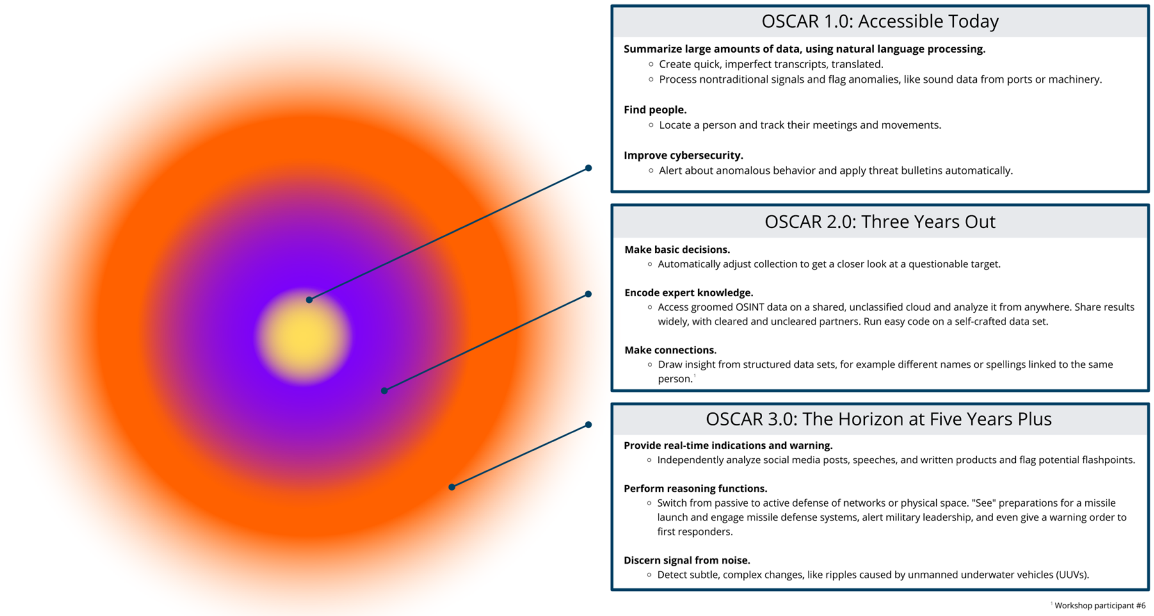

OSCAR 1.0: Accessible Today

Summarize large amounts of data, using natural language processing.

- Create quick, imperfect transcripts, translated.
- Process nontraditional signals and flag anomalies, like sound data from ports or machinery.

Find people.

- Locate a person and track their meetings and movements.

Improve cybersecurity.

- Alert about anomalous behavior and apply threat bulletins automatically.

OSCAR 2.0: Three Years Out

Make basic decisions.

- Automatically adjust collection to get a closer look at a questionable target.

Encode expert knowledge.

- Access groomed OSINT data on a shared, unclassified cloud and analyze it from anywhere. Share results widely, with cleared and uncleared partners. Run easy code on a self-crafted data set.

Make connections.

- Draw insight from structured data sets, for example different names or spellings linked to the same person.[1]

OSCAR 3.0: The Horizon at Five Years Plus

Provide real-time indications and warning.

- Independently analyze social media posts, speeches, and written products and flag potential flashpoints.

Perform reasoning functions.

- Switch from passive to active defense of networks or physical space. "See" preparations for a missile launch and engage missile defense systems, alert military leadership, and even give a warning order to first responders.

Discern signal from noise.

- Detect subtle, complex changes, like ripples caused by unmanned underwater vehicles (UUVs).

[1] Workshop participant #6

Source: Author's own analysis.

OSCAR 1.0: Near-Term Possibilities

"How do you do more on the low side at scale? That's the next big step."—interviewee[69]

Today, machines can read, see, write, and think in limited ways. Image-recognition algorithms can be trained to identify cats in YouTube videos, as well as known terrorists captured on video by unmanned aerial vehicles (UAVs) over Syria or Afghanistan.[70] Personal digital assistants such as Siri, Alexa, and Cortana can read and respond to emails, suggest grammar and substantive changes, and sort important emails from spam. They get these tasks right a substantial portion of the time, but not always. One workshop participant described the current state of machine learning as follows: "If you have imagery and 9 out of 10 people can look [at the image] and make a similar answer, then it is currently accessible to create an AI system that can do that today."[71]

With this capability, the intelligence community could immediately begin widely using machines to comb through unstructured data such as images, articles, and reports; build a structured database; and classify the data into categories created by analysts or by the machine itself.[72] The most basic

[69] Interviewee #5.

[70] Greg Allen and Taniel Chan, *Artificial Intelligence and National Security* (Cambridge, MA: Belfer Center for Science and International Affairs, July 2017), https://www.belfercenter.org/publication/artificial-intelligence-and-national-security.

[71] Workshop participant #6.

[72] Mandeep Kaur, "Top 10 Real-Life Examples of Machine Learning," Big Data Made Simple, May 13, 2019, https://bigdata-madesimple.com/top-10-real-life-examples-of-machine-learning/.

application of this capability is tipping and cueing. For example, a satellite collects far more data in one pass around the planet than a human can capably review. Human eyes get tired and bored; AI/ML, on the other hand, can classify every speck in a satellite image of the Pacific Ocean, for example, as either an island, fishing boat, container ship, warship, or anomaly. Even when the machine is confused or the image is unclear, AI/ML can tell human analysts which images to evaluate more closely, saving hundreds of labor hours.

A host of COTS capabilities could create OSCAR 1.0 tomorrow. If widely adopted, it could:

- **Automate summaries of large amounts of data using natural language processing.** AI/ML can create quick, imperfect transcripts, translate them, summarize main points, and search for keywords automatically. Letting AI/ML point out the most relevant content can save analysts hundreds of hours of combing manually through videos and documents. AI/ML can also passively monitor the written word, agnostic of format.[73] Similarly, AI/ML could automate a variety of relatively rote IC tasks, like creation of cable templates for the field, along with some auto-population of data fields. Many HR and logistics-related applications could be 95 percent automated, saving time and resources for tasks that need human judgement.[74]
- **Find people in the smart-city haystack.** China and other countries have installed extensive surveillance capabilities, building "smart cities." Paired with facial or body recognition, an intelligence service can locate a person—friend or foe—and track their meetings and movements. The counterintelligence implications are immense. Starting with the face of a suspected case officer, for example, an intelligence service can track their movements across a city, including meetings with potential assets, allowing that service to identify a mole in its ranks. Conversely, an intelligence service can use the same capability in real time to see if a case officer is being followed to meetings by a person or set of people, allowing them to abort a meeting at the last minute. This capability, which the intelligence community calls "ubiquitous technical surveillance" (UTS), will change the face of HUMINT work.[75]
- **Improve cybersecurity defense.** AI/ML applications are already proving effective at cybersecurity defense, alerting humans to anomalous behavior and effectively applying threat intelligence alerts to networks. Such applications can more accurately diagnose and isolate a breach when the attacker and their methods are known, but recent advances have allowed AI/ML to detect more advanced threats by focusing on anomalies rather than known attackers.[76] In summer

[73] Workshop participant #11.

[74] Workshop participant #11.

[75] According to a *Foreign Policy* article, "Cities in at least 56 countries worldwide have deployed surveillance technologies powered by automatic data mining, facial recognition, and other forms of artificial intelligence. Urban surveillance is a multibillion-dollar industry, with Chinese and U.S.-based companies such as Axis, Dahua, Hikvision, Huawei, and ZTE leading the charge." See Robert Muggah and Greg Walton, "'Smart' Cities Are Surveilled Cities," *Foreign Policy*, April 17, 2021, https://foreignpolicy.com/2021/04/17/smart-cities-surveillance-privacy-digital-threats-internet-of-things-5g/.

[76] Chuck Brooks and Frederic Lemieux, "Three Key Artificial Intelligence Applications for Cybersecurity by Chuck Brooks and Dr. Frederic Lemieux," *Forbes*, September 24, 2021, https://www.forbes.com/sites/chuckbrooks/2021/09/24/three-key-artificial-intelligence-applications-for-cybersecurity/?sh=4714489f7b7e.

2021, Darktrace used AI/ML monitoring tools to neutralize an attack on the Tokyo Olympics.[77] Meticulous Research has estimated that the market for AI/ML in cybersecurity will grow to $46.3 billion by 2027.[78]

Given the current possibilities in COTS AI/ML, the intelligence community could effectively supercharge its capabilities. It has already invested heavily in the necessary cloud capability; using an unclassified cloud to scale up a common platform that could serve as a test bed for low-side commercial AI will allow for these advances. The last sections of this paper make recommendations for how to take that leap.

OSCAR 2.0: MID-TERM POSSIBILITIES

In three years, enough training data will have been curated and entered for AI/ML systems that they can accomplish a broader range of intelligence tasks with greater precision and recall. Today, one of the limiting factors in AI/ML capability is the paucity of cleaned and labeled training data to prepare systems to analyze information properly. For example, if an AI/ML system has not processed enough satellite images of various ballistic missile sites, it cannot differentiate a silo from a radar tower or distinguish between types of missiles. Three years from now, this issue will be solved.

With this "knowledge" on board, AI/ML systems will be able to make basic decisions. For example, an AI/ML system will be able to examine a series of images from a satellite or UAV and decide whether it is routine or requires a more thorough human examination.

Improvements in AI-enabled intelligence collection and analysis, however, will not be limited to government agencies. Over the next three years, commercial intelligence firms will continue to expand and adopt similar technologies to the point that the private sector will have many of the same capabilities. Companies like HawkEye 360 and the investigative group Bellingcat have already begun replicating IC capabilities, and such organizations will only continue to proliferate in the next few years. This will pose a challenge for the U.S. government, as there are currently no guidelines for how and when the intelligence community can access data from commercial firms due to privacy concerns. The United States needs to develop protocols for accessing this information or risk closing out a major source of intelligence.

In three to five years, AI/ML capabilities might include the following:

- **Decisionmaking.** Enough training data should exist by then that machines can consistently make a basic decision, such as adjusting collection to get a closer look at a questionable target. For example, AI/ML systems mounted on semi-autonomous UAVs monitoring activity in the Pacific Ocean could "see" a ship, identify it as a fishing vessel but note some anomalous features, then "decide" to loiter over the ship to collect additional information.

[77] Oakley Cox, "AI Neutralizes IoT Attack That Threatened to Disrupt the Tokyo Olympics," Darktrace, September 20, 2021, https://www.darktrace.com/en/blog/ai-neutralizes-io-t-attack-that-threatened-to-disrupt-the-tokyo-olympics/.

[78] Meticulous Research, *Artificial Intelligence (AI) In Cybersecurity Market by Technology (ML, NLP), Security (Endpoint, Cloud, Network), Application (DLP, UTM, Encryption, IAM, Antivirus, IDP), Industry (Retail, Government, Automotive, BFSI, IT Healthcare, Education), Geography - Global Forecast to 2027*, (Meticulous Research, June 2020), https://www.meticulousresearch.com/product/artificial-intelligence-in-cybersecurity-market-5101.

- **Encoding expert knowledge.** A combination of robust U.S. government training data, the capacity to analyze that data in-house on a shared cloud, and tools that encode expert knowledge for better identification will reduce costs and make intelligence available from anywhere.
- **Simplifying data science.** Within five years, knowledge of data-science techniques will spread through the IC workforce to an extent where a political analyst—whether at home or in Abuja—can run simple code on a self-crafted data set. A small number of highly skilled data scientists can assist those with general knowledge, serving as force multipliers.
- **Making connections.** Once an unclassified cloud capability has enough structured data, AI/ML could make sense of the connections between that data, including by finding linkages and resolving similarities such as next-order combination/fusion—e.g., different names or spellings for the same person.[79]

More importantly, the intelligence community can anticipate additional pressure from private intelligence services. Workshop participants expected to see a surge in high-quality AI/ML intelligence products in the next three to five years, incorporating unprecedented capability to conduct high-resolution commercial collection of IMINT, SIGINT, and even HUMINT from publicly available information. Without the weight of government acquisition procedures—or the security restrictions that come with operating in a 90 percent classified environment—commercial entities will have the freedom to innovate and create alternative streams of high-quality intelligence. While no private entity is likely to ever supplant the IC's vast reach and exquisite capabilities, the community will need to anticipate that its audience will read not only the *New York Times* but also intelligence streams from private-sector competitors before they read the WIRe.

OSCAR 3.0: HIGH-POWERED COMPUTE AND QUANTUM-ENABLED POSSIBILITIES

In five years, if the intelligence community has made the necessary investments in unclassified cloud computing, AI/ML, and OSINT, systems operating on curated government data will be seamless and scalable to a variety of challenges. The emphasis will also shift to being able to quickly deploy customized programs for specific challenges.

Some of the greatest leaps in AI-enabled analysis will be in natural language processing. In five years, AI/ML could begin to independently analyze, not just summarize, social media posts, speeches, telephone conversations, signals intercepts, and written messages, leading to intelligence and warning in something much closer to real time.

AI/ML applications could perform many of the reasoning functions humans do today. In the national-security context, AI/ML-enabled systems could switch from passive to active defense of networks or physical space. An AI/ML system could "see" preparations for a missile launch and engage missile-defense systems, alert military leadership, and even give a warning order to first responders and hospitals—detecting and responding faster than a human to new technology such as hypersonic missiles. AI/ML systems with advanced computing power could detect undersea changes as well, from tsunami-generating earthquakes to subtle ripples caused by unmanned underwater vehicles (UUVs). Validating human assets is also vital to intelligence analysis—in the future, AI/ML could use facial recognition, query publicly available data holdings, and deliver a report on whom a person has met, where they have lived, and what jobs they have had in the last 10 years.

[79] Workshop participant #6.

A significant increase in computational power would also lead to a massive acceleration in AI/ML capabilities. Quantum computing, once only a theory, is slowly evolving into reality. In 2019, a Google team claimed its quantum computer carried out a specific calculation that it estimated would take the most advanced "classical" computer more than 10,000 years.[80] *Nature* quoted Scott Aaronson, a theoretical computer scientist at the University of Texas–Austin: "The scientific achievement is huge, assuming it stands, and I'm guessing it will."[81] The enhanced speed of quantum computing, along with its ability to process complex if/then calculations, will transform AI/ML applications.

[80] Madhumita Murgia and Richard Waters, "Google Claims to Have Reached Quantum Supremacy," *Financial Times*, September 20, 2019, https://www.ft.com/content/b9bb4e54-dbc1-11e9-8f9b-77216ebe1f17; and Vidar, "Google's Quantum Computer Is about 158 Million Times Faster than the World's Fastest Supercomputer," Medium, February 28, 2021, https://medium.com/predict/googles-quantum-computer-is-about-158-million-times-faster-than-the-world-s-fastest-supercomputer-36df56747f7f.

[81] Elizabeth Gibney, "Hello Quantum World! Google Publishes Landmark Quantum Supremacy Claim," *Nature*, October 23, 2019, https://www.nature.com/articles/d41586-019-03213-z.

5

Recommendations

Training OSCAR and Training Us

"Focus on concrete, actionable things that individuals can do. Changing policy is a multi-year effort, but changing behaviors around policy can be faster. Encourage new efforts to take risks."—workshop participant[82]

Imagine an analyst is tasked with evaluating whether China has laid a new undersea cable off the coast of a third country. Today, that analyst would manually search a myriad of databases for classified reporting on China's intentions regarding cable lines and capabilities to deploy them, Beijing's interactions with the third country's capital, and perhaps publicly available databases on ship movements. They would also examine any available imagery of that area in the last few months and ask NGA colleagues to help look for evidence of cable-laying or on-shore disruptions that indicate building a ground station. If they do not find evidence, they could say, "We see no indications of a new undersea cable," but would also explain the collection gaps and task assets to follow up. Maybe in the following weeks or months, a new piece of intelligence would illuminate China's activities—likely too late for policymaker engagement with the targeted country.

If that same analyst had OSCAR, their timeline would be dramatically shorter and the presented product more complete. They could ask OSCAR to run a set of queries in a data lake on an unclassified cloud, using computer vision to search commercial imagery to identify ships likely to have cable-laying capability. OSCAR could notice changes in the foliage on the shore, suggesting landfall for a cable or the location of a new ground station. OSCAR could uncover and translate social media posts by Chinese sailors or find and translate a speech in which a local politician thanks a Chinese businessman for coming to visit his coastal town and expresses excitement about a new construction project. OSCAR

[82] Workshop participant #8.

could query commercial radio-frequency collection databases, finding push-to-talk communications between a Chinese shipping vessel and a trailing vessel, indicating that the onboard cargo is not what was reported on the manifest. Even if the analyst finds none of these things, they can ask never-blinking OSCAR to issue an alert if a ship matching a description enters the area. Finally, they could task OSCAR to run a bunch of dummy queries to confuse any adversary who might have access to their activity on the unclassified cloud.

So how does the U.S. intelligence community get from here to there? What policies should the United States change, and what elements of culture need to shift in order to make OSCAR a reality?

The sections below offer actionable recommendations in the key areas identified as the biggest obstacles to this vision. Acknowledging the urgency and the immensity of the task, the final section presents dramatic steps the intelligence community should take if more modest steps do not achieve results.

Culture

Culture is the most difficult part of an organization to change. To improve uptake of AI/ML in the intelligence community, a culture shift needs to take place. This should start with changing the demand signal:

- **Reframe OSINT to emphasize the value added by the intelligence community**. IC officers think of OSINT as press—from CNN to ITAR-TASS—and the value added as its translation. Today, OSINT is far more, and the real value extracted from the information is in the processing. Agencies should start teaching analyst cadres to think of OSINT as publicly available information that has been collected and processed to provide insights available nowhere else.
- **Involve analysts, operators, and other users in the acquisition.** There should be a close link between what analysts need and what the IC buys and incorporates. For example, some agencies are leveraging natural language processing through a new data-gathering capability called InsightAI.[83] The Defense Intelligence Agency (DIA), when developing the new Machine-assisted Analytic Rapid-repository System (MARS) program, used this tool to understand the needs of analysts so they could spend time, energy, and resources on the most impactful capabilities, rather than deploy a new system that had no real use case. The only way to cross the "valley of death" and have widespread incorporation is if analysts create the demand.[84]
- **Explicitly tie OSCAR to missions.** IC leadership has heard many overblown sales pitches about the promise of AI. Overselling a capability will only backfire. IC contracting officers should ask critical questions about professed capabilities, particularly regarding scalability and data training, and those selling AI/ML for OSINT should explain the limitations of the solution, how the state of the art is evolving, and explicitly tie the capability to the mission.
- **Create an OSCAR habit.** A benchmark indicating the IC has fully adopted OSCAR is when analysts

[83] See InsightAI at https://content.bmnt.ai/insightai.

[84] See BMNT Inc. at https://www.bmnt.com/.

turn to these tools as part of their normal routine. Creating that routine will start with training—IC elements should stand down groups of analysts for a week at a time to learn the necessary new skills. That will create a demand signal, at least among early adopters, for easy access to the toolset. For a subset of these analysts, the IC should remove the expectation of producing current intel for a few months to allow them to perfect their skills and become OSCAR ambassadors. Having upper leadership recognize and celebrate when these analysts use OSCAR effectively in PDBs and other current production will further normalize and drive demand.

- **Get to know the data scientist next door.** A week of training will soon wear off if there are no resources close at hand to help analysts continue to use the tool. IC agencies should embed one, or ideally several, data scientists with every mission center to champion the technology and serve as a resource for analysts seeking to learn how to use OSCAR on a daily basis.

Security

The security argument against having analysts and operators query unclassified databases is that adversaries may be able to discern U.S. intelligence priorities and even map who is accessing the databases. These are legitimate, but solvable, concerns. The conversation needs to shift from what happens if the IC implements OSCAR to what happens if it does not—in which case the United States will rapidly be outstripped by adversaries such as China, which will conduct intelligence collection and warfighting at high speed, with the United States falling days or weeks behind in indications and warning.

- **Issue clear orders from leadership on accepting risk**. ODNI and the leadership of each agency should explicitly declare adoption of OSCAR a priority and, in repeated internally public statements, talk about how the risk is acceptable and necessary. Agencies should also seek to "experiment when failure is cheap," by testing several versions of a minimally viable capability before investing heavily in development.[85]
- **Bring security officers into the discussions from the beginning**. Their participation can be framed as helping the IC safely get to yes. For every complication and obstacle raised, the IC should ask them to come up with two potential solutions, no matter how expensive or bureaucratically painful. It should be made explicit that they are encouraged to take risks, and that leadership will support them if security issues arise later.
- **Embed security professionals with cloud providers for a month to learn cloud security features.** Security officers should sit with security professionals at Amazon, Oracle, and Microsoft to learn about available security features and get comfortable with the idea of unclassified cloud access.
- **Train intelligence professionals on security practices in the unclassified realm.** As analysts learn about OSCAR, they should learn about the mechanisms that can obscure their activity as they play on the unclassified cloud, like VPNs, and on how to preserve that anonymity.[86]

[85] Workshop participant #17.

[86] Workshop participant #7.

Policy

Policy changes will speed the uptake of OSCAR, but none as much as reforming acquisition of software. Interviewees and participants in the workshop agreed that total reform of the acquisition process is desperately needed, with some suggesting creating a new, parallel acquisition process because the current one is too slow.

- **Conduct a zero-based review for the software acquisition process**. Acquisition officers should sit down with agency leadership and lay out the entire process, including why each step is necessary. Is it a legal requirement? A long-held policy? A relic of the past? The group should then brainstorm how to cut the number of steps by 75 percent.
- **Prepare to mitigate.** The IC should explicitly acknowledge the risk it is accepting by changing this process and think through what might mitigate that risk. The Intelligence Committees should be briefed on the plan before it is implemented.
- **Shift to statements of objectives (SOOs).** The standard procedure—requiring statements of work (SOWs)—establishes defined timelines and activities, which are ill-suited to rapid advancements in software. As Corin Stone describes it,

  > SOWs make sense when the government understands with precision exactly what is needed from the contractor and how it should be achieved. SOOs, on the other hand, are more appropriate when the strategic outcome and objectives are clear, but the steps to achieve them are less so. They describe 'what' without dictating 'how,' thereby encouraging and empowering industry to propose innovative solutions.[87]

  As an interviewee put it, "The way you ask the question of the market matters. If you over-specify capabilities, [you] will drive out good product. Instead, provide end-state goals."[88]

- **Retrain acquisition officers on buying cycles for emerging tech.** These officers are professionals committed to protecting the government, but sticking with past practice is easy and adopting new procedures is hard on a human level. Retraining should include incentives for accepting risk and creating efficiencies to support the mission.
- **Implement Other Transaction Authority (OTA) and Commercial Solutions Openings (CSOs).**[89] Stone advocates for these vehicles in her recent series: "[OTA] allows specific types of transactions to be completed outside of the traditional federal laws and regulations that apply to standard government procurement contracts, providing significantly more speed, flexibility, and accessibility than traditional contracts." CSOs, meanwhile, are a "relatively quick solicitation method to award firm fixed price contracts up to $100 million. CSOs can be used to . . . close capability gaps or provide technological advances through an open call for proposals that provide offerors the opportunity to respond with

[87] Stone, "Artificial Intelligence in the Intelligence Community."

[88] Interviewee #15.

[89] "10 U.S. Code § 2371B - Authority of the Department of Defense to Carry out Certain Prototype Projects," Legal Information Institute, https://www.law.cornell.edu/uscode/text/10/2371b; "Contracting Cone," DAU Adaptive Acquisition Framework, https://aaf.dau.edu/aaf/contracting-cone/defense-cso/.

technical solutions of their own choosing to a broadly defined area of government interest."[90]

- **Stretch Indefinite Delivery/Indefinite Quantity (IDIQ) contracts.** IDIQ allows for changing technical details but not the nature of the product. Agencies should work with Congress to allow these contracts to also make space for product updates, such as advancements in AI/ML.
- **Force explicit accounting for decisions to build rather than buy.** In order to overcome the "not built here" bias against buying COTS software, agencies should have to prove their "build" is obviously better based on cost, efficiency, and security.[91]

A secondary, but no less critical, policy change will be to clarify data standards. The IC needs to develop rules for cataloging and training data for AI/ML use. Ideally, this would be another opportunity to make use of COTS—companies would compete to curate publicly available information for the IC, eventually building trust that will grant them access to clean up already collected data.[92]

- **Request detailed data models.** The IC should request detailed data models from vendors and partners. While an added expense, it will be critical to analyzing and storing data, as well as sharing with partners such as the Five Eyes intelligence alliance.[93]

People

While the technological advances are substantial, they mean nothing if the professionals in the IC do not know how to use them and have no interest in trying. Leadership will need to create both opportunities to learn and excitement about the available possibilities in order to spark this evolution in intelligence work.

- **Create rotations for operational data teams**. In partnership with tech firms who do business with (or want to do business with) the federal government, the IC should create opportunities for data scientists to do a two-year rotation into federal service. The U.S. government could reimburse part of their salaries, commensurate with government service, and private organizations could subsidize the rest. Tech firms would gain a greater understanding of how the federal government thinks and operates, while the government would benefit from their expertise. Ideally, members of these operational data teams would be cleared for full access and understanding of their assigned IC mission—but, if not, they could still work on the difficult data problems and advise analysts on how to tackle questions.
- **Use OSCAR as a training ground.** While this paper does not argue for creating a new Open Source Center, opting instead for integrating OSINT analysis as a discipline into every all-source analyst's toolkit, ODNI could create an open-source intelligence temporary position to which those with a conditional offer of employment could contribute while awaiting their final security clearance. The IC loses diverse talent in part because not everyone can afford to wait months or years for clearance. At an OSCAR-based training ground, those in the pipeline can learn how to use AI/ML to exploit publicly available information, ready to carry those new skills into their line-analytical jobs, and the IC can draw on that talent in the meantime.

[90] Stone, "Artificial Intelligence in the Intelligence Community."

[91] Workshop participant #8.

[92] Interviewee #12.

[93] Interviewee #15.

Bold Steps

All these recommendations are achievable and should bring the United States closer to keeping pace with China's advancements in AI/ML. However, changing how the IC conceives of intelligence capability is both urgent and vital—and the IC should be prepared to take bold steps if, in the next year, significant progress is elusive. Bold steps should include the following:

- **Create a parallel acquisition process for AI/ML and cloud systems.** The current acquisition process is a hindrance to progress; the U.S. intelligence community should be prepared to take the risk of casting it aside and starting from scratch for quickly changing products like AI/ML and the cloud services on which they run. At first, these could be approved only on a limited basis, but they should be stress tested and placed under consideration for full adoption.
- **Create a new Indefinite Delivery/Outcome Oriented (IDOO) contract category.** This will combine the strengths of IDIQ with the idea of defining outcomes, rather than products, to allow for ultimate flexibility in updating the AI/ML software and cloud capabilities as the state of the art evolves.
- **Create an IC innovation incubator to nurture and launch OSCAR.** If the cultural changes above do not take hold, creating an incubator like Lockheed Martin's Skunk Works would allow OSCAR to develop and demonstrate its worth.
- **Give ODNI total budget authority over AI/ML and cloud acquisition.** ODNI should handpick an interagency team of specialists to purchase, contract, and distribute software and other capabilities. Agencies and departments will likely object to this centralization, but it may be needed to prompt real change.
- **Shift personnel incentives.** Congress should initiate a resource shift for this mission by mandating (and funding) big bonuses for retraining on OSCAR capabilities and requiring demonstrated proficiency in OSCAR before, for example, promoting an analyst to GS-12 or sending them on an overseas tour. IC agencies should also send analysts on rotation to cloud and AI/ML companies to learn the products' capabilities and how industry builds them.

Conclusion

Like any organizational change, widespread IC adoption of OSCAR will not come from any one initiative. It will need champions, pushes from several angles, and the amplification of success stories worth emulating. Tony Stark and JARVIS rarely succeed on their own—for example, in *The Avengers*, JARVIS and Tony use quick analytics to fix the helicarrier, but Captain America's analog support is essential to complete the mission.

Culture is always the hardest thing to change. Organizational inertia is powerful, and the urge to keep doing what we have been doing remains attractive as long as there is no real pressure to upend established processes. But the IC is perfectly capable of rapid shifts in the name of mission. After 9/11, it reorganized itself overnight and retrained analysts, operators, and support staff to focus on a new mission. Russian attempts to influence the 2016 election were another wake-up call, which resulted in a new election threats executive and other reforms.

Some have speculated, however, that the 2016 elections should have been the intelligence community's 9/11 moment for OSINT; after all, Moscow used proxies and social media platforms to spread misinformation and disinformation, and the U.S. government was largely dependent on Twitter, Facebook, and other social media companies to trace and report on that activity. Until faced with intense pressure from Congress, these platforms did not self-report, and the IC (and its lawyers) have still not figured out what the rules should be for using their publicly available information, which includes data on Americans. Reorganization and retraining were limited and lost momentum amid a fraught political environment. What could have been a crisis-turned-opportunity simply passed.

A workshop participant who specializes in driving innovation described the challenges to IC adoption of new tech: "The three main risks for any innovation project are user desirability (users want it and key stakeholders will adopt it), technical feasibility (can we build the solution in the time available), and organizational viability (is there a pathway through the bureaucracy to transition it)."[94] The intelligence community can create user demand by showing IC officers what is possible—how the mission can be served by open-minded adoption of OSINT—and explaining key tools to make it better. The second risk, technical feasibility, is the easiest to address—provided the U.S. government focuses on the third pillar: creating a pathway through the bureaucracy. Stakeholders at several levels need to be involved early and often in the development of these new tools—from users to senior leaders to skeptics. The IC needs to shake off antiquated contract practices and take on a modicum of risk if it is to continue to serve the nation. Put bluntly by a former IC member, "Where there's a will, there's a waiver." The IC has always prided itself on being the first line of the nation's defense. OSCAR is coming to help.

[94] Workshop participant #17.

About the Author

Emily Harding is deputy director and senior fellow with the International Security Program at the Center for Strategic and International Studies (CSIS). She joined CSIS from the Senate Select Committee on Intelligence (SSCI), where she was deputy staff director. In her nearly 20 years of government service, she has served in a series of high-profile national security positions at critical moments. While working for SSCI, she led the Committee's multiyear investigation into Russian interference in the 2016 elections. The five-volume, 1,300-page report reshaped the way the United States defends itself against foreign adversaries seeking to manipulate elections, and it was lauded for its rigor, its thoroughness, and as the only bipartisan effort on election interference. During her tenure on the Committee, she also served as the subject matter expert on election security, counterintelligence and associated cybersecurity issues, and the Middle East. She oversaw the activities of 18 intelligence agencies and led SSCI staff in drafting legislation, conducting oversight of the intelligence community, and developing their expertise in intelligence community matters.

She began her career as a leadership analyst at CIA, and then became a manager of analysts and analytic programs. She led the Iraq Group during the attempted Islamic State takeover of Iraq and Syria and led a multidisciplinary group of analysts working crises worldwide, drawing from many perspectives to provide rich analysis to policymakers. During a tour at the National Security Council, she served as executive assistant to the deputy national security adviser for global democracy strategy and then as director for Iran, where she led interagency efforts to create innovative policies drawing on all elements of national power. After leaving the White House, she served on a team running the first Office of the Director of National Intelligence-led presidential transition, where she was responsible for liaising with both campaigns and briefing the incoming administration on a wide range of intelligence topics. Harding holds a master's degree in public policy from Harvard University's Kennedy School of Government and a bachelor's degree in foreign affairs from the University of Virginia.

COVER PHOTO BLUE PLANET STUDIO/ADOBE STOCK

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

1616 Rhode Island Avenue NW
Washington, DC 20036
202 887 0200 | **www.csis.org**