# Report | Cultural Policy Hub on Content Provenance, the Media, and Democracy

The Cultural Policy Hub at OCAD University is a national, bilingual platform that aims to build informed, inclusive and integrated cultural policy. It contributes to policy issues in collaboration with researchers, policymakers, artists and creators from across Canada's academic, government, non-profit and private spheres.

*Interested in learning more about the Hub?*
*Visit our website, subscribe to our monthly newsletter, or follow us on LinkedIn*

## Executive Summary

The rise of artificial intelligence (AI), combined with the popularity of social sharing platforms, has made it easier than ever to create and distribute fake images, videos, and news online. This is contributing to eroding public trust: when anyone can create synthetic content that is perceptually indistinguishable from "real" content, people start to doubt the authenticity of all the content they are consuming. This leads to what is called the "liar's dividend," where bad actors can dismiss real information as fake, making it harder for the public to know what to believe.

Technical solutions exist for identifying the authenticity of digital content; One such solution is the open, interoperable standard developed by the Coalition for Content Provenance and Authenticity (C2PA). This approach deals with the entire information ecosystem, from capture through production, distribution and consumption. Its adoption does, however, require the participation of a wide range of unrelated parties. Coordinating and motivating these interest holders is where the federal government can accelerate progress.

To address this challenge, the Cultural Policy Hub convened experts from media, government, tech, and civil society for a workshop on "Content Provenance, the Media, and Democracy" at the annual DemocracyXChange (DXC) summit in April 2026. The workshop was designed with support from colleagues at Neural Transform and the CBC and facilitated by Aaron Williamson.

The goals of the workshop were:

- To explore how a voluntary Code of Practice could be structured and adopted to set standards for future regulation; and
- To ensure that provenance systems remain financially and technically accessible and are adopted by a wide range of users, including small media organizations, independent creators, and civil society groups.

To achieve this, participants first examined how provenance data could be tracked from content creation to consumption and identified the specific challenges for content labelling at each stage of the supply chain. They then discussed what the key elements of a voluntary Code of Practice for Canada might include. To translate these insights into action, the workshop identified the following priority actions to guide Canada's approach toward immediate, scalable adoption:

- Canada should prioritize piloting the Code of Practice in high-impact areas, such as major newsrooms and federal/provincial government data and to create a baseline of trust and demonstrate feasibility.
- Federal and provincial departments can lead by example by applying content provenance and authenticity labels to their own content, such as reports and data releases.
- Governments should act as a catalyst by using procurement power (RFPs) to require secure media provenance labels as a part of all new public facing systems. Compliance should be tied to market access, meaning non-compliant entities could be excluded from government procurement opportunities or deemed ineligible for grants.
- Newsroom workflows will need to be upgraded to incorporate AI and Media Provenance technologies. This comes at a time when financially strained newsrooms "make do" with current tools. Journalism grants to incentivize hardware, software, and workflow upgrades will speed the adoption of the media provenance ecosystem. Much of this work is business-to-business system integration and is a productivity investment.
- Citizens must understand how to interpret provenance signals. This requires integrating media literacy into school curricula (drawing on existing models in other jurisdictions) and launching targeted public awareness campaigns tailored to diverse demographics.
- Strict protections must be built into the Code of Practice to shield vulnerable sources—such as journalists and activists—from surveillance, ensuring that provenance data cannot be weaponized against individuals.

## Background and Context

The rapid spread of high-quality AI-generated content makes it essential to trace its origins and preserve origin information. Deepfakes and manipulated media—images, videos, and even news reporting—have become tools for disinformation and election interference. The urgency is clear, and technical solutions alone are insufficient. Public concern is high, with 70% of Canadians seeing AI as a serious risk, and 88% demanding stronger governance. Incidents like the Kirkland Lake bot campaign and AI-generated ads mimicking news outlets during the 2025 federal election highlight the threat AI-generated content poses to Canada's democratic integrity. The provenance of online content is often invisible, and the ecosystem remains fragmented. Without visible and reliable verification details, the public struggles to distinguish between authentic and manipulated content, undermining trust in the information ecosystem.

The Hub's workshop at the DemocracyXChange posed a key question: What baseline expectations should guide the adoption of content provenance in Canada? Participants addressed this question by aiming to develop the contours of a voluntary framework and Code of Practice for media provenance in Canada. This framework would ensure that provenance systems are accessible to all while establishing secure, traceable digital content flows. It would align interest holders on shared norms and maintain Canada's competitive position as a global leader in fostering trust and resilience in the information ecosystem.

To guide their thinking on the potential solutions, participants were introduced to one of the possible technical solutions which has gathered global momentum: the Coalition for Content Provenance and Authenticity (C2PA). C2PA provides an open interoperable technical standard

for attaching tamper-evident origin information (or “provenance signals”) to digital content such as images, videos, and documents, enabling verifiable content credentials. It functions like a digital receipt, tracking digital content from its origins—from the lens of a camera, a word processor, a sound recorder, etc.—all the way to its final destination: our phones and computer screens. This "glass to glass" approach creates a complete, unbroken record of where the content came from, how it was edited, and provides the consumer with reliable information on which to make a trust decision. It can be thought of as the equivalent of a nutrition label for media: consumers may choose to ignore it, but its presence allows for informed decisions. However, its effectiveness grows with wider adoption, visible display, and user literacy.

Prior to the workshop, participants received a scene-setter briefing prepared by the Cultural Policy Hub, which offers a detailed overview of the landscape and key challenges. To access this briefing document, click here.

## Provenance Tracking: From Creation to Consumption

The workshop used a supply chain framework to map and trace the digital information lifecycle from Creation through Verification, Packaging, Distribution, and finally, Consumption. Participants, organized into working groups, identified actors, threats, and workflow challenges at each stage. This approach moved discussions from abstract policy to actionable insights, revealing where provenance signals are lost or compromised.

Participants mapped a host of good actors and bad actors influencing the adoption of provenance standards across the information supply chain, including journalists, editors, creators, archivists, social media platforms, AI companies, advertisers, fact-checkers, researchers, government bodies, bot farms, non-human agents, and users.

Participants identified a convergence of threats and challenges across the information supply chain. They highlighted the following interconnected systemic barriers:

1. The Erosion of Truth and Trust

    - Bad actors can dismiss real information as fake and vice versa. Even government officials admitted their institutions have been "deceived many times" by sophisticated fakes, creating a risk where the public doubts the veracity of all content. This is causing an overall decline in trust in the news and information ecosystem.
    - "Troll factories" and bot farms flood the internet with propaganda. Furthermore, AI creates content based on old data, which new AI then trains on, creating a loop where the original truth is diluted or lost forever.
    - AI tools often summarize news without citing sources, hiding the original creator and making it impossible for consumers to distinguish between valid reporting and AI-generated noise.

2. Operational and Economic Pressures

    - The pressure to publish "breaking news" can leave journalists with limited resources for fact checking. The process for verification needs to be adapted to be augmented with machine-to-machine verification. There is a fundamental tension between the

need for speed, and the time and effort required to undertake rigorous verification.

- Large networks may have the budgets to adopt provenance standards, whereas local newspapers and independent journalists may not. Without support, local community information could soon be invisible or overlooked due to the lack of recognized provenance signals.
- Some journalists can no longer publish with confidence, as they lack the resources, time, or technical tools to distinguish real content from AI fakes. The credibility of their work is threatened by a constant flow of fake news and misinformation, and they are often ill-equipped to address the problem.

3. Technical Fragility and Data Integrity

- The core standards for embedding provenance like C2PA already exist and are now being deployed as a feature of cameras, edit suites, generative AI tools, and browsers. The functional gap is with system integration, workflow adaptation, and training. News outlets hesitate to invest because they do not yet see an immediate financial benefit or trust the reliability of these systems.
- As AI agents read and share news automatically, verification data must be readable by machines, not just humans, adding technical complexity.
- Platforms routinely remove provenance signals. Accidental stripping—where the digital proof of a content's origins is destroyed—can also occur when files are shared from one source to another. Technologies to restore the provenance data exists, but it shifts the cost burden from the platforms to the financially strained newsrooms.

4. Privacy, Safety, and Surveillance

- Provenance signals, while optional, could amplify mass surveillance, rather than just verification, and could be used by governments or bad actors to track down and punish people.
- Attaching too much data to a post could inadvertently reveal the location or identity of a source, creating a dangerous trade-off between verification and personal or political safety. Redaction of metadata is an important function.

5. Legal and Structural Conflicts

- Provenance signals can become permanent records if read by AI training systems, which conflict with laws allowing individuals to request data deletion. Resolving how to get AI models to "unlearn" data is a massive hurdle.
- There is no agreement on whether news content can be used to train AI models without permission or payment, creating a messy legal battlefield. Provenance signals can contain information on consent conditions, which would remove some doubt about a creator's intent.
- Even with provenance in place, AI Models could ignore consent terms and could "capture" content by keeping it on closed platforms, cutting creators off from their audience and revenue.

6. The Adoption Paradox

    - If the provenance verification process is too difficult, newsrooms will ignore it; if it is too strict, valuable stories may never get published. This needs to be baked into the technical processes to minimize burden on reporters and to get uniform adoption.
    - Major platforms may slow progress because their business models prioritize engagement and speed over truth and content origin. Provenance does not solve this fundamental conflict of interest, but it does provide hooks for resolution.

Key Insights:

Participants identified several principles that will guide the development of the Voluntary Code of Practice:

- Provenance standards require more than technology and workflow upgrades. They demand parallel investment in policy, education, and human capacity. To bridge the resource gap, governments, industry collaborations, and public-private partnerships must step in to fund and deliver the training needed for equitable adoption.
- The technology to solve provenance challenges already exists. For widespread adoption, tools must be so simple that anyone can use them without technical expertise. The goal is to ensure everyone can trace content origins, regardless of their stance on the content itself.
- Trust is personal and contextual. Users need transparency: the ability to see who created and shared the content, and to assess their credibility. Trust is the choice of an informed recipient of information, not imposed by governments or big companies.
- A major realization was that we need to distinguish between "distribution" (sharing a story) and "training" (using a story to teach an AI); publishers should have the right and technical ability to say "no" to the latter, without losing access to the search listings and advertising revenue of the former.
- Participants agreed that the old way of getting news (creator to reader) is broken. Now, AI acts as a middleman that often hides the source. The group realized that without a way to see the "life cycle" of a story, consumers cannot tell if what they are reading is real or just a remix of old data.

The workshop revealed that technical solutions for provenance tracking already exist, but their adoption is hindered by economic, legal, and operational barriers. A collaborative, community-driven approach—supported by open standards and hybrid verification—is essential to preserve trust, accuracy, and fairness in the digital information ecosystem. This can be achieved by developing a voluntary Code of Practice for Canada. The following section outlines the basic principles of this code, operationalizing the insights above into baseline expectations for media organizations, government bodies, and technology vendors.

## Developing a Voluntary Code of Practice for Canada

To address the fragmentation of the information ecosystem, the workshop participants moved

beyond problem identification to prototype a Voluntary Code of Practice tailored to the Canadian context. This framework is designed not as a rigid mandate, but as an industry-led set of baseline expectations that media organizations, government bodies, technology vendors, and end-users can adopt to establish secure, traceable digital content flows. The core philosophy driving this code is flexibility: recognizing that a large national broadcaster and an independent community journalist face vastly different resource constraints, the framework allows for tiered compliance based on an organization's capacity and risk profile. This approach aims to lower barriers to entry while ensuring that the essential principles of provenance are upheld across the spectrum of content creation.

## Defining Scope and Governance

A central pillar of the Code of Practice is the establishment of clear, shared definitions for content types and the governance of the framework itself. Participants agreed that the framework must distinguish between human-made content, AI-assisted creation, and fully AI-generated material, grounding these distinctions in intent, perception, and process. To ensure the Code of Practice is practical and legally sound, the workshop established several critical structural elements:

- It must explicitly include end-users in its scope, ensuring that the public is not just a passive recipient but a recognized interest holder in the provenance ecosystem.
- It needs to clearly separate legal requirements (such as Canadian privacy laws) from global technical standards (like ISO or C2PA specifications), allowing for local adaptation without compromising international interoperability.
- It must clearly define what constitutes a "material change" to content, distinguishing between a technical edit (like fixing the brightness of a photo) and a material editorial change (like using AI to put someone in a place they never visited). If a significant change occurs, the content must be relabeled to show it has been altered. This should be done with a reference to international standards, such as the IPTC media type codes.
- It must outline the consequences for non-compliance, providing a clear understanding of the risks for organizations that choose not to adopt the standards, even within a voluntary framework. They may still face repercussions, such as loss of grants eligibility or exclusion from government procurement opportunities (RFPs).

## Embedding Privacy and Safety

Privacy and safety are woven into the fabric of the proposed Code of Practice, reflecting a critical concern raised throughout the workshop regarding the potential for provenance data to be weaponized against vulnerable individuals. The Code of Practice needs to explicitly include provisions to protect the identities of sources, like journalists and activists operating under oppressive governments.

Key privacy safeguards embedded in the voluntary Code of Practice include:

- Collecting only the necessary information required for verification.
- Extending provenance data beyond the moment of capture to include archiving, reuse, and rights management.

- Allowing creators to control and edit what provenance information they optionally provide to be included with the media.
- Addressing the tension between immutable records and legal requirements for data deletion where applicable.

## Driving Adoption and Market Incentives

While the code is voluntary, the workshop identified several strategic mechanisms to encourage widespread adoption and ensure its effectiveness. A major focus was on securing vendor buy-in. Major hardware and software providers and platform owners (like Canon, Google, Adobe, Microsoft, and Meta) are already deploying features based on C2PA provenance standards due to global demand. The group agreed that the Code of Practice must demonstrate the Canadian benefits of compliance and clarify the risks of non-adoption to continue to motivate these and smaller providers.

To facilitate collaboration without violating competition laws, the Code of Practice would require 'house rules' for partnerships. These rules will allow organizations to share non-proprietary, technical, and operational knowledge to collectively promote the use of provenance tools, while strictly prohibiting the exchange of commercially sensitive information or any coordination that could restrict competition.

## The Role of Government

Government bodies were recognized as uniquely positioned to lead by example, though a key concern raised was that centralized control of provenance systems, especially by government, could be misused to silence voices or control narratives. With this safeguard in mind, participants suggested that federal and provincial departments should begin by signing their own data releases and research reports to establish a baseline of trust. The code also envisions the use of procurement power and other levers to drive adoption. These include:

- Mandating compliance with provenance standards in government Requests for Proposals (RFPs) to create market incentives;
- Attaching standards compliance to journalism grants and support programs;
- Supporting media literacy campaigns to help citizens understand and value provenance signals; and
- Clarifying the consequences of non-compliance to encourage voluntary participation.

## The Path Forward

The path forward for the Voluntary Code of Practice prioritizes scaling adoption of multiple reinforcing sub-systems above waiting for a perfect solution. The workshop concluded with a focus on a Minimum Viable Provenance (MVP) strategy, focusing first on high-impact use cases within major newsrooms and government data releases. This means tackling specific, high-value slices of the news and information supply chain now, rather than waiting for a perfect system that may never materialize. Crucially, this phase will involve active interest holder consultation to gather feedback from media organizations, government, and civil society on the draft framework, ensuring broad acceptance.

Much of this work has already been pioneered by CBC/Radio-Canada in conjunction with international partners. It would be helpful if this know were to be shared with other members of the Canadian news ecosystem.

## Conclusion

The DemocracyXChange 2026 workshop successfully mapped the complex ecosystem of digital content provenance, revealing that the challenge of synthetic media is a multi-dimensional issue affecting our social fabric, political stability, and economic health. The consensus is clear: the "liar's dividend" and the fragmentation of our information supply chain pose a direct threat to the resilience of Canadian democracy. The emergence of AI based summaries as a primary information source exacerbates the challenge.

Moving forward requires a coordinated approach that brings together government, industry, and civil society. Rather than waiting for rigid regulations to impose change, Canada can join a global movement to securely label digital content by establishing a voluntary Code of Practice and adopting a "Minimum Viable Provenance" strategy. This approach allows us to start verifying the origin of content immediately, without waiting for perfect technology.

The workshop concluded with a commitment to share findings and continue this critical dialogue. The goal is to ensure that technology serves the truth, protecting both creators and citizens. Canadian public institutions and tech providers must act proactively now to align interest holders and set norms. By doing so, we can safeguard the integrity of our information ecosystem before external pressures or crises force our hand.

We already have a blueprint for success: EBU and CBC/Radio-Canada won the 2026 NAB Technology Innovation Award for their end-to-end workflow capturing video on a C2PA equipped Sony Camera, editing it using C2PA tools from Adobe, signing it using an internally developed CBC/Radio Canada tool, and playing it on open-source EBU C2PA video player. This achievement demonstrates that, with the right leadership and collaboration, Canadian innovators can not only meet these challenges, but are poised to play their part in setting the global standard for digital trust.

International standards for open secure digital media provenance now exist and are rapidly being incorporated into existing products and services. The Canadian challenge is for businesses to prioritize the integration of this capability into established news workflows and train staff to interpret the results. Once these foundations are in place, the effort can expand to include public education. The creation of a Code of Practice will add momentum to this effort.