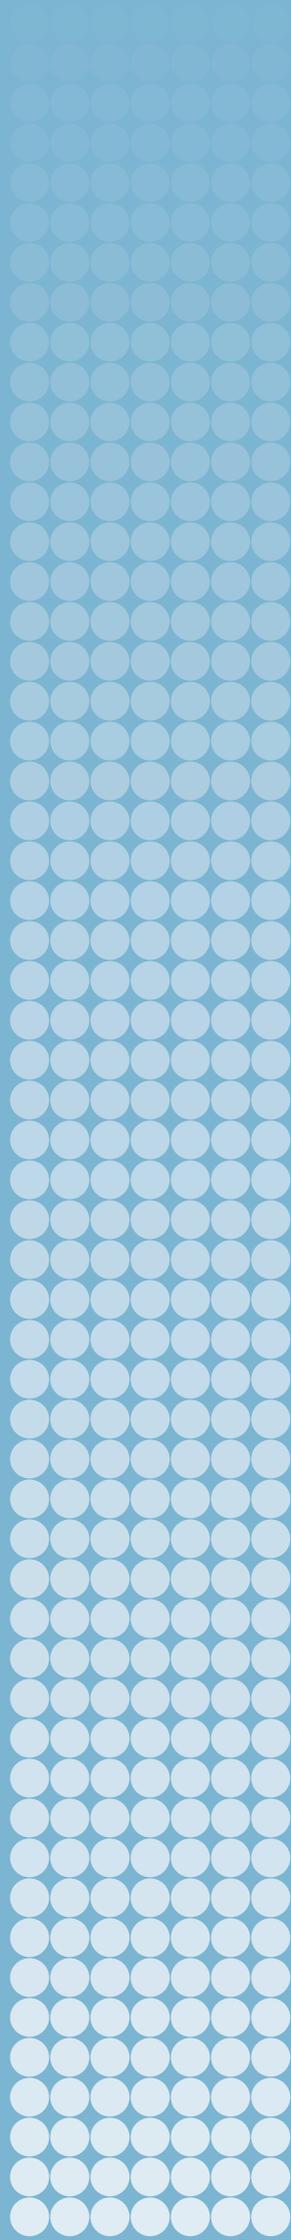


PREVENTING BIOLOGICAL WEAPONS PROLIFERATION

Operational Applications of Emerging Technologies

MIRANDA SMITH, KOLJA BROCKMANN
AND MARK BROMLEY



**STOCKHOLM INTERNATIONAL
PEACE RESEARCH INSTITUTE**

SIPRI is an independent international institute dedicated to research into conflict, armaments, arms control and disarmament. Established in 1966, SIPRI provides data, analysis and recommendations, based on open sources, to policymakers, researchers, media and the interested public.

The Governing Board is not responsible for the views expressed in the publications of the Institute.

GOVERNING BOARD

Stefan Löfven, Chair (Sweden)
Dr Mohamed Ibn Chambas (Ghana)
Ambassador Chan Heng Chee (Singapore)
Dr Noha El-Mikawy (Egypt)
Jean-Marie Guéhenno (France)
Dr Radha Kumar (India)
Dr Patricia Lewis (Ireland/United Kingdom)
Dr Jessica Tuchman Mathews (United States)

DIRECTOR

Karim Haggag (Egypt)



**STOCKHOLM INTERNATIONAL
PEACE RESEARCH INSTITUTE**

Signalistgatan 9
SE-169 70 Solna, Sweden
Telephone: +46 8 655 97 00
Email: sipri@sipri.org
Internet: www.sipri.org

PREVENTING BIOLOGICAL WEAPONS PROLIFERATION

Operational Applications of Emerging Technologies

MIRANDA SMITH, KOLJA BROCKMANN
AND MARK BROMLEY

March 2026



**STOCKHOLM INTERNATIONAL
PEACE RESEARCH INSTITUTE**

© SIPRI 2026

DOI: <https://doi.org/10.55163/HPTG4304>

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form or by any means, without the prior permission in writing of SIPRI or as expressly permitted by law.

Contents

<i>Executive summary</i>	v
1. Introduction	1
Box 1.1. Key terms	3
Figure 1.1. Key functions and supporting legal instruments of the biological weapons prohibition regime	2
2. Production and laboratory oversight	4
Use case: AI for laboratory activity monitoring	4
Use case: DLT for traceability and data integrity	6
3. Export controls on dual-use items	8
Use case: AI for item classification	8
Use case: AI for licensing facilitation eligibility determination	10
Use case: DLT for export documentation authentication	11
4. Reporting and transparency mechanisms	13
Use case: AI for reporting and transparency	14
Use case: DLT for data integrity and record assurance	16
Box 4.1. Artificial intelligence, information analysis and the Biological Weapons Convention verification debate	16
Figure 4.1. Status of reporting obligations under biological weapons prohibition regime measures	14
5. Cross-cutting considerations	18
Data integrity and governance	18
Human oversight and confidence in AI systems	19
Interoperability and policy coherence	19
Balancing transparency and security	19
Access, equity and capacity	20
Experience-sharing and collaboration	21
6. Conclusions	22
Recommendations	23
Figure 6.1. Potential applications through which artificial intelligence (AI) and distributed ledger technology (DLT) can strengthen key functions of the biological weapons prohibition regime	23
<i>About the authors</i>	26

Executive summary

The key functions of the biological weapons prohibition regime are established by the rights and obligations outlined in the Biological and Toxin Weapons Convention (BWC). The BWC is complemented by other instruments and arrangements which reinforce implementation and operationalize key obligations. Rapid advances in emerging technologies create opportunities to help prevent the development and proliferation of biological weapons by strengthening key functions of the biological weapons prohibition regime. Artificial intelligence (AI) and distributed ledger technology (DLT) have the potential to strengthen three key functions of the BWC regime—production and laboratory oversight, export controls on dual-use items, and reporting and transparency mechanisms—while a fourth—cooperation and assistance—can be used to support their responsible and equitable adoption.

Production and laboratory oversight

The growing number of high- and medium-containment laboratories, alongside the shift from paper-based record-keeping to digital workflows, has increased the volume and granularity of information generated within these facilities. For these reasons, emerging technologies, particularly AI-enabled analysis and DLT-enabled tools, have attracted attention as a means of supporting laboratory oversight by improving how existing information is organized, reviewed and preserved. For example, a growing share of laboratory informatics platforms promote integrated analytics and AI-enabled functionality that is directly relevant to oversight workflows. DLT could help narrow the gap between periodic inspections and ongoing laboratory operations.

Export controls on dual-use items

Continued scientific and technological progress in biology and the life sciences and the growth and internationalization of the biotechnology sector have increased the volume of transfers that are covered by states' export controls, particularly those of controlled software and technology. States are increasingly discussing the use of AI and, to a lesser extent, DLT to strengthen the national implementation of export control systems. Companies, research institutes and universities are considering their use to help them comply with the legal obligations created by national export controls on dual-use items. These tools can assist with tasks such as product classification, end-user screening and licensing workflows, helping exporters and authorities manage growing regulatory complexity while maintaining oversight of sensitive biological transfers.

Reporting and transparency mechanisms

To date, AI- and DLT-enabled tools have not been used to support BWC and United Nations Security Council Resolution 1540 reporting and transparency practices. However, governments and research institutes have indicated support for their adoption and highlighted potential applications, and there are multiple potential use cases that deserve further consideration. For example, AI can be used to support internal quality control by flagging inconsistencies within a state's reporting. Where reporting responsibility spans multiple ministries or shifts across agencies, DLT can provide a durable audit trail of when specific records were created, updated or transmitted.

Cross-cutting considerations

Several considerations shape whether the wider adoption of AI- and DLT-enabled tools will strengthen oversight and confidence building or introduce new frictions. These include data integrity and governance, interoperability and policy coherence, balancing transparency with security, maintaining human oversight and confidence in AI systems, ensuring equitable access and capacity, and strengthening experience-sharing and collaboration across states and sectors. The development of any system using emerging technologies raises the question of whether all states will have equitable access to such systems and how they would attain the capabilities to set up, operate and maintain them. Where possible, states should seek to use the regime's cooperation and assistance channels to support standard-setting and to enable responsible, equitable distribution.

Recommendations

These recommendations are aimed primarily at states parties and specify where action is most realistically taken by laboratories and industry individually or through national systems, and where multilateral forums can support learning, comparability and capacity building.

- Create structured spaces for innovation and dialogue within the BWC.
- Treat digital capacity as part of cooperation and capacity-building under BWC Article X.
- Leverage complementary forums without creating parallel standards.
- Build a shared foundation for digital trust.
- Preserve accountability in an automated future.

1. Introduction

Rapid advances in emerging technologies have brought significant improvements in key areas of biological research and operations. These improvements have generated benefits but have also raised concerns about their potential to facilitate the development and proliferation of biological weapons. For example, recent advances in artificial intelligence (AI) have supported improvements in gene sequencing and gene synthesis, enabling more rapid processing in vaccine discovery and gene therapy.¹ However, these improvements could also be repurposed to support the development of transmissible biological agents.² At the same time, advances in emerging technologies create opportunities to help prevent the development and proliferation of biological weapons by strengthening key functions of the biological weapons prohibition regime. This policy paper explores how AI and distributed ledger technology (DLT) can be used to strengthen three key functions of this regime—production and laboratory oversight, export controls on dual-use items, and reporting and transparency mechanisms—and how a fourth—cooperation and assistance—can be used to support their responsible and equitable adoption.

The key functions of the biological weapons prohibition regime are established by the rights and obligations outlined in the Biological and Toxin Weapons Convention (BWC).³ The BWC is supported by other instruments which add additional legal force to certain non-proliferation obligations, most notably United Nations Security Council Resolution 1540, and by arrangements that detail how specific controls can operate in practice, in the case of the Australia Group (see figure 1.1).⁴

These functions include national implementation measures that states parties adopt to fulfil their obligations under the BWC. Measures may take numerous forms and functions at various levels (e.g. legislation, regulations, codes of conduct and good practices).⁵ These measures are applied through domestic regulatory and compliance systems, typically involving both public authorities and regulatory oversight entities.⁶ For example, states implement domestic measures to prohibit biological weapons-related activities, often operationalized through production and laboratory oversight of facilities handling relevant agents.⁷ States also use export controls on dual-use items to prevent transfers that could contribute to biological weapons programmes. This function is reinforced by Resolution 1540 obligations and supplemented in practice by technical guidance, including the Australia Group control lists and guidelines.⁸

These functions also include two other activities that are carried out primarily by states. First, transparency and reporting mechanisms, including BWC confidence-

¹ See Avsec, G. and Latysheva, N., 'AlphaGenome: AI for better understanding the genome', Google DeepMind, 25 June 2025; and Intuition Labs, 'The modern biotech lab: A guide to automation, AI & data', 20 Oct. 2018.

² National Academies of Sciences, Engineering and Medicine, 'AI-enabled biological design and the risks of synthetic biology', *The Age of AI in the Life Sciences: Benefits and Biosecurity Considerations* (National Academies Press: Washington, DC, 2025).

³ See Lentzos, F., 'Compliance and enforcement in the Biological Weapons regime', UNIDIR WMD Compliance & Enforcement Paper no. 4, 5 Dec. 2019; and Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction (Biological Weapons Convention, BWC), opened for signature 10 Apr. 1972, entered into force 26 Mar. 1975, *United Nations Treaty Series* vol. 1015 (1976).

⁴ UN Security Council Resolution 1540, 28 Apr. 2004.

⁵ United Nations Office for Disarmament Affairs (UNODA), *Guide to Implementing the Biological Weapons Convention* (UNODA: Geneva, Aug. 2023).

⁶ See Longworth, S., 'Interpreting the Biological Weapons Convention: What are "necessary measures" under Article IV of the Convention?', Swedish Defence Research Agency (FOI) Memo 8407, Dec. 2023.

⁷ Longworth (note 6).

⁸ Resolution 1540 (note 4); and Australia Group, 'Guidelines for transfers of sensitive chemical or biological items', June 2015.

Regime function	Supporting legal instruments
 Production and laboratory oversight	BWC Articles I and IV
 Export controls on dual-use items	BWC Article III UN Security Council Resolution 1540 Australia Group control lists and guidelines
 Reporting and transparency mechanisms	BWC Article V CBM submissions supporting BWC (Articles IV and VI) and Resolution 1540
 Cooperation and assistance	BWC Article X

Figure 1.1. Key functions and supporting legal instruments of the biological weapons prohibition regime

building measures (CBMs) and reporting under Resolution 1540, enable states to demonstrate good-faith implementation and share relevant information. Second, the cooperation and assistance function involves mechanisms for facilitating peaceful exchange and for capacity building, including practical support.

Officials and analysts have suggested using AI to help detect biological weapons programmes or support a future verification mechanism for the BWC.⁹ These ideas remain largely untested and, in some cases, politically contentious. By contrast, less attention has been paid to the ways that AI and DLT are being used, and could be used, to strengthen the key functions of the biological weapons prohibition regime. The wider adoption of AI- and DLT-enabled tools to strengthen certain regime functions may become more attractive in light of the increasingly complex regulatory and compliance environment created by the distributed and data-intensive nature of modern biological research and operations. However, many AI- and DLT-enabled tools are embedded in internal systems, commercial platforms or security-sensitive processes. As a result, lessons learned from their use are often siloed within sectors, national authorities or companies. There are also concerns that these tools may fail to preserve accountability, confidentiality and human oversight, and that they will be less accessible in states in the developing world. These risks create a need for structured discussion on how to ensure responsible and equitable uptake across states.

This policy paper discusses existing and potential applications of AI and DLT that could help strengthen key regime functions, beginning with production and laboratory oversight (section 2); export controls on dual-use items (section 3); and reporting and transparency mechanisms (section 4). Section 4 also briefly examines discussions around using AI to support BWC verification. Each section presents examples of actual and potential use cases of AI- or DLT-enabled tools, drawing on documented deployments, regulatory pilots, platform capabilities, and administrative analogues. Box 1.1 defines key technical terms that are used in these sections.

⁹ See Revill, J., 'How AI can—and cannot—improve verification of the Biological Weapons Convention', *Bulletin of Atomic Scientists*, 6 Oct. 2025.

Box 1.1. Key terms

Application programming interface (API): A system access point or library function that enables software applications to share information, features and functionality with each other.

Artificial intelligence (AI): A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments.

Auditability: The extent to which records and processes can be reviewed and reconstructed reliably over time.

Blockchain: A type of DLT that collates a number of records into a ‘block’, which is then ‘chained’ to the next block, using a cryptographic signature. This allows blockchains to be used like a ledger, which can be shared and corroborated by anyone with the appropriate permissions.

Chain of custody: A record of who had control of a material or sample, when, and under what authorization.

Data reconciliation: The process of cross-checking records from different systems or sources to ensure they match.

Distributed ledger technology (DLT): A shared record system that can create tamper-evident logs and time-stamped integrity proofs, typically through controlled access and cryptographic techniques.

Hash: A fixed-length value representing a certain type of data; the algorithmic *hash function* that converts the input data to the output hash value ensures the integrity of the data.

Integrity signal: A technical marker, such as a hash and time stamp, that allows later checks of whether a record has been modified.

Large language model (LLM): A type of AI designed to process and generate text, including by extracting information from documents and producing draft summaries for human review.

Machine learning (ML): A type of AI that detects and ‘learns’ from patterns or anomalies in training data to make inferences about new data.

Metadata: Descriptive information about a record (e.g. date, author, document identifier), which can contain sensitive information even when the record’s content is not shared.

Permissioned ledger: A DLT system where access to write, read or query records is controlled by defined participants.

Provenance: The ability to trace a record to its source and timeline, including when it was created or updated.

Source: US Department of Commerce, National Institute of Standards and Technology (NIST), *Glossary of Key Information Security Terms*, NISTIR 5298, rev 3, July 2019; and NIST Computer Resource Centre, ‘Glossary’, 14 Jan. 2026.

The paper then assesses cross-cutting dynamics that influence the feasibility of adopting these tools (section 5). This section also examines how states could use the regime’s cooperation and assistance mechanisms to support setting standards for these tools and enabling their responsible and equitable distribution. The final section presents conclusions and recommendations (section 6), including steps that states and other stakeholders could take to support responsible and equitable adoption and deployment of AI- and DLT-enabled tools to strengthen the biological weapons prohibition regime.

2. Production and laboratory oversight

Effective production and laboratory oversight increasingly depends on the integrity, continuity and usability of operational records. The growing number of high- and medium-containment laboratories, alongside the shift from paper-based record-keeping to digital workflows, has increased the volume and granularity of information generated within these facilities. Where documentation once relied heavily on notebooks, paper logs and fragmented files, laboratory operations now routinely generate large volumes of digital records through laboratory information management systems (LIMS), access-control systems, equipment logs and environmental monitoring tools.¹⁰ This expands the scope for oversight by creating records that can be searched, aggregated and assessed over time. At the same time, it shifts the challenge from data availability to data management and interpretation, placing greater demands on both national regulatory authorities and laboratory operators, who must reconcile heterogeneous data sources, distinguish routine variations from meaningful discrepancies, and maintain continuity across inspection cycles.¹¹ Without the capacity to structure, validate and interpret these records, the proliferation of digital data does little to strengthen national implementation measures and may complicate efforts to demonstrate that regulatory controls are being applied consistently in practice.

Key areas of laboratory oversight still rely on extensive manual review in which inspectors and regulators reconcile a range of record types—inventories, access logs, equipment records and procedural documentation—that are maintained in different systems and formats, often across multiple time periods.¹² This process is time-intensive and depends heavily on individual expertise and institutional familiarity with specific facilities, particularly because inspections are periodic rather than ongoing and are subject to practical disruptions such as staffing constraints, scheduling backlogs and travel limitations. During the COVID-19 pandemic, for example, several national regulatory authorities postponed many inspections, resulting in backlogs and increased reliance on remote and records-based assessment approaches in some regulated sectors.¹³ In such contexts, the availability of reliable, well-structured digital records becomes central to maintaining continuity of oversight and demonstrating good-faith adherence to national implementation measures.

For these reasons, emerging technologies, particularly AI-enabled analysis and DLT-enabled tools, have attracted attention as a means of supporting laboratory oversight by improving how existing information is organized, reviewed and preserved.

Use case: AI for laboratory activity monitoring

Sectors with stringent regulatory compliance—such as pharmaceutical manufacturing, environmental monitoring, and industrial process control—are already using AI to detect anomalies, predict equipment failures and flag deviations from expected para-

¹⁰ World Health Organization (WHO), *Laboratory Biosafety Manual*, 4th edn (WHO: Geneva, Dec. 2020).

¹¹ Organisation for Economic Co-operation and Development (OECD), ‘GLP data integrity’, OECD Series on Principles of Good Laboratory Practice and Compliance Monitoring No. 22, 20 Sep. 2021; and WHO, *Laboratory Biosecurity Guidance* (WHO: Geneva, 2024), section 6.5.

¹² Lentzos, F. et al., *Global BioLabs Report 2023* (Global Biolabs Initiative, 2023).

¹³ Cuddy, E., Lu, Y. P. and Ridley, D. B., ‘FDA global drug inspections: Surveillance of manufacturing establishments remains well below pre-COVID-19 levels’, *Health Affairs*, vol. 42, no. 12 (Dec. 2023); Denigan-Macauley, D., ‘Drug safety: FDA has faced persistent challenges overseeing foreign drug manufacturing’, Testimony before the Subcommittee on Oversight and Investigations, US House of Representatives Committee on Energy and Commerce, GAO-24-107359, 6 Feb. 2024; and US Food and Drug Administration (FDA), ‘Conducting remote regulatory assessments: Questions and answers—Guidance for industry’, June 2025.

meters. The Organisation for Economic Co-operation and Development (OECD) conducted a cross-government review of AI applications that highlights core functions like document analysis and monitoring processes—areas that align closely with the record review and compliance triage needed in laboratory oversight.¹⁴ The European Medicines Agency catalogued a wide range of AI and machine-learning applications across the lifecycle of medicines, including applications that reinforce compliance functions.¹⁵ Peer-reviewed assessments of AI use in pharmaceutical manufacturing similarly describe predictive maintenance and continuous parameter monitoring as prominent application areas.¹⁶ These examples suggest that the operational logic of AI-based oversight already exists in regulated domains, even if specific laboratory use cases are not always disclosed.

A growing share of laboratory informatics platforms, including LIMS and related quality and compliance systems, now promote integrated analytics and AI-enabled functionality that is directly relevant to oversight workflows, even when implementation remains at the facility level. Illustrative public materials include Benchling's launch of Benchling AI, which positions AI-assisted search and drafting within laboratory research and development workflows, and vendor roadmaps from LabWare and STARLIMS that emphasize analytics and AI-enabled review of laboratory data to support quality and compliance functions.¹⁷ These examples are best read as product direction and market signalling, not as a comprehensive inventory of facility-level adoption, but they demonstrate that the tooling is increasingly embedded in mainstream laboratory data platforms.

Many LIMS record time-stamped entries for biological material inventories, reagent usage, sequencing runs, personnel access events, and equipment status. When machine-learning models are applied to these datasets, users can identify statistically significant deviations from established baselines, such as unexpected frequency or patterns of sample withdrawals, reagent usage inconsistent with approved protocols, or abnormal fluctuations in equipment performance which may indicate procedural lapses. These applications are relevant to laboratory oversight not because they introduce new analytical concepts, but because they show how highly regulated environments use operational data more systematically to support monitoring and review. In settings where oversight relies on periodic review rather than continuous presence, the practical contribution is not an automated compliance determination but targeted identification of records, variables and time periods that merit human attention.

Beyond anomaly detection, AI systems can support laboratory oversight by integrating heterogeneous operational data streams into consolidated dashboards. This approach, widely used in industrial automation, can bring together outputs from equipment logs, environmental monitoring systems and other digital records to provide a single operational picture over time. Laboratory informatics platforms, including LIMS, are also increasingly emphasizing instrument connectivity and integrated analytics as part of

¹⁴ OECD, *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions* (OECD Publishing: Paris, 2025).

¹⁵ European Medicines Agency (EMA), 'Review of AI/ML applications in medicines lifecycle (2024)', EMA Horizon Scanning Short Report, EMA/571739/2024, 9 July 2025.

¹⁶ Kodumuru, R., et al., 'Artificial intelligence and internet of things integration in pharmaceutical manufacturing: A smart synergy', *Pharmaceutics*, vol. 17, no. 3 (2025); and Ojo, O. R. et al. 'Applying machine learning models for real-time process monitoring and anomaly detection in pharma manufacturing', *GSC Biological and Pharmaceutical Sciences*, vol. 27, no. 1 (2024).

¹⁷ See LabWare, 'Advanced analytics, machine learning, & AI for your lab', 7 Feb. 2022; STARLIMS, 'Artificial intelligence and machine learning: A new frontier', 7 Feb. 2022; and Benchling, 'Benchling reveals the next chapter of R&D at Benchtalk 2025', Press release, 14 Oct. 2025.

routine laboratory data management.¹⁸ When calibrated against facility-specific operating ranges, these dashboards can support incident investigation by reconstructing conditions preceding an event, such as equipment failure or deviations in containment-relevant parameters. The effectiveness of these systems depends less on algorithmic sophistication than on consistent digital record-keeping and data integration that allow multiple streams to be aligned, compared and analysed over time.

The operational value of these systems lies in how outputs fit within existing oversight workflows. Rather than producing binary judgements, anomaly-detection models flag specific records or trends for review. For laboratories, this can support earlier identification and correction of documentation inconsistencies before scheduled inspections. For regulators, it can reduce the time spent reconciling fragmented records and help target inspection planning toward higher-risk issues or facilities.

Use case: DLT for traceability and data integrity

DLT is used in regulated supply chains across a variety of sectors to create tamper-evident records of provenance and verification events, particularly where multiple entities need to use shared records without relying on a single entity's database. In the pharmaceutical sector, DLT has been explored as part of traceability and product verification efforts, including demonstration of six industry-developed platforms through the United States Food and Drug Administration (FDA) Pilot Project Program under the US Drug Supply Chain Security Act.¹⁹ At the same time, experience in other sectors illustrates that adoption is not automatic and can fail when governance arrangements, incentives and onboarding do not align. This is reflected in the discontinuation of the TradeLens platform by Maersk and IBM after it did not reach the level of industry participation needed for commercial viability.²⁰

Where AI strengthens the interpretability of laboratory records, DLT could be used to support the integrity and continuity of relevant compliance records over time. In laboratory oversight, permissioned ledger architectures can be used as an integrity layer for auditable metadata, for example by anchoring hashes and time stamps for selected events, while keeping sensitive scientific content off ledger under controlled access.²¹ In principle, laboratories could log events such as materials transfers between laboratories, inventory reconciliations, and chain-of-custody actions associated with the movement or destruction of sensitive samples. Ledger entries would typically include metadata such as date and time, quantity, responsible personnel identifiers, and reference to the associated internal record, rather than experimental data or pathogen characteristics.

Because ledger entries are cryptographically protected and time-stamped, they cannot be altered retrospectively without detection. For laboratories, this provides a defensible audit trail that reduces the personnel time required to assemble documentation ahead of inspections, supports continuity despite staff turnover, and protects institutions in the event of retrospective compliance questions. For regulators, DLT offers assurance that records reflect contemporaneous actions rather than reconstructed narratives, improving confidence in digital documentation even when physical access to facilities is limited.

¹⁸ Edayan J. M., et al., 'Integration technologies in laboratory information systems: A systematic review', *Informatics in Medicine Unlocked*, vol. 50 (2024).

¹⁹ US FDA, 'DSCSA Pilot Project Program', 12 June 2024.

²⁰ Maersk, 'Information on the closure of TradeLens', Advisory, 1 Dec. 2022; and Maersk, 'A. P. Moller – Maersk and IBM to discontinue TradeLens, a blockchain-enabled global trade platform', Press release, 29 Nov. 2022.

²¹ Salingros, E., *Guidelines for the Management of Digitalised Systems in Laboratories Accredited to ISO/IEC 17025*, Technical Report no. 1/2024 (EUROLAB: Brussels, Oct. 2024).

As with AI, facility-level deployment of DLT in laboratory oversight is not systematically documented in public sources, since implementation is often internal, vendor-mediated or treated as security-sensitive. Integration of DLT with existing LIMS is most plausibly implemented as a ledger layer which anchors or verifies selected records, such as material transfers, reconciliations, or compliance and licensing certificates, while the primary operational data remain within the LIMS. For example, ServBlock, an Ireland-based blockchain compliance company, markets a DLT-based data exchange model designed to integrate with LIMS and support trusted transfer of laboratory data across the inbound supply chain.²² Discrepancies between declared and recorded events can be flagged for follow-up, reducing the likelihood of clerical errors and supporting more consistent oversight across inspection cycles. By preserving compliance-relevant metadata in verifiable form, DLT can be leveraged to narrow the gap between periodic inspections and ongoing laboratory operations by reducing reconciliation effort and improving record continuity.

²² ServBlock, 'Introducing Lab Xchange—Secure laboratory data exchange & DLT', [n.d.].

3. Export controls on dual-use items

Export controls set up policies and a licensing system that allows national authorities to have oversight of and—where necessary—to intervene and prevent transfers of sensitive items. An export control system primarily relies on control lists which use technical parameters to define the items whose transfer requires an export licence. They also extend licensing requirements to transfers destined for proscribed end uses and restricted end users. All exporters of controlled items—including, but not limited to, companies, research institutes and universities—are required to comply with licensing and record-keeping obligations and act with due diligence by, for example, ensuring they know their customers and the end use of products they supply. These obligations can overlap with the obligations created by sanctions adopted by the UN, regional organizations and individual states. Exporters are encouraged, and in some cases required, to adopt internal compliance programmes (ICPs) to ensure they are meeting their export control and sanctions obligations.

Continued scientific and technological progress in biology and the life sciences and the growth and internationalization of the biotechnology sector have increased the volume of transfers that are covered by states' export controls, particularly those of controlled software and technology. Implementing export controls in a way that prevents the proliferation of biological weapons while not infringing on states' rights and obligations under Article X of the BWC is challenging and requires continuous efforts and resources allocated to national export control authorities. The obligations created by states' export controls also result in a constantly evolving regulatory landscape for exporters which need to set up adequate ICPs and foster a culture of responsible innovation and behaviour to meet them.

States are increasingly discussing the use of AI and, to a lesser extent, DLT to strengthen the national implementation of export control systems, while companies, research institutes and universities are considering their use to help them comply with the legal obligations created by the implementing regulations at the national level. Only a small number of states have announced that their national authorities have begun adopting AI or DLT systems in their national export control authorities.²³ Exporters and suppliers of export control-related consulting services have recently begun development, testing and deployment of a variety of AI-enabled tools that promise to assist with different export control compliance functions, including product classification and end-user screening.²⁴ Adopting AI and DLT tools to strengthen export controls is most compelling where they can reduce administrative and process burdens, facilitate access to information and provide validated documentation, instead of attempting to provide reliable legal and technical assessments for exporters and licensing authorities.

Use case: AI for item classification

The proper classification of an item (i.e. a good, software or technology) is important to ensure that exporters apply the appropriate due diligence, responsibility and duty of care to exports of sensitive biological dual-use items and that licensing authorities assess the resulting licensing application based on the correct assumptions about the

²³ For examples from the area of customs control see Lisoir, X., Rane, S. and Jhawar, R., 'Revolutionising customs with AI: Dream big, start small', PwC, 2024; and Mikuriya, K. and Cantens, T., 'If algorithms dream of customs, do customs officials dream of algorithms? A manifesto for data mobilisation in customs', *World Customs Journal*, vol. 14, no. 2 (Sep. 2020).

²⁴ Brockmann, K. and Héau, L., 'Use cases for emerging technologies to strengthen export controls on biological items', SIPRI Research Policy Paper, Dec. 2025.

item—and, where necessary, can check if it has been misclassified. Some transaction-party screening tools are already deploying AI to cross-reference the various sanctions and designated party lists, and to cross-check known aliases and relations between individuals, entities, known addresses and other information using public and commercial information sources such as company registries.²⁵ AI assistants and large language models (LLMs) are increasingly being used to assist with the classification of items. However, there are significant limitations associated with their use in this area, particularly regarding the classification of technology and software that is still in development.

A key technical challenge for the development and training of such tools results from the format of control lists and the logic of their structure. Simply providing one of the current publicly available AI assistants or LLMs with the national legislation and control lists and prompting it to classify a product based on a technical description frequently results in misclassifications.²⁶ Such LLMs still frequently fail to correctly resolve the multiple levels of conditionalities and exceptions that characterize export control legislation and classification, which can result in the tool classifying non-sensitive items as requiring a licence or not recognizing items as listed and requiring a licence. Other challenges include that without necessary authorizing provisions in place, there could be limitations on the extent to which data submitted as part of previous licensing applications may be used for training of AI models; and systems might lack compatible application programming interfaces (APIs) or use different programming languages and file formats.²⁷

Rather than trying to develop AI tools that classify items by analysing a data sheet or project description, some research institutes, companies and compliance service providers are developing AI tools in-house or as service products that use a series of questions to support the classification process.²⁸ These questions prompt users to provide the information necessary to make an initial assessment of whether the product or project involves the potential transfer of controlled items and directs relevant queries to compliance staff. Use of these prompt-based tools directly by researchers or engineers within a research institute or company reduces the manual or repetitive workload for compliance staff and increases the visibility of regulatory and internal policy compliance procedures. Such blended AI–human compliance approaches allow more time to focus on complicated cases, take final decisions on cases, and ensure more systematic proof of due diligence on classification beyond typical email or meeting note exchanges.²⁹ These systems can also more efficiently create audit trails in case of later questions and reviews by the export licensing authorities.

National export control authorities are also exploring the use of AI-enabled tools, particularly custom-trained LLMs, to perform internal initial classification of items based on the information provided by the exporter in a licence application or to assist with understanding and navigating national export control legislation. One advantage of LLMs is their ability to quickly process the information included in licence application forms, the supplemental data sheets and other technical information submitted as part of licensing applications. Particularly where well-established electronic licens-

²⁵ See e.g. Deloitte, ‘TRADEclassifier’, [n.d.]; Deloitte, ‘Future proofing your sanctions compliance program’, 9 June 2025; and Xapien, ‘Xapien FAQs’, [n.d.].

²⁶ Kampf, A., ‘17. Informationstag Exportkontrolle’ [17th export control information day], Germany Trade and Invest, 5 Jan. 2026.

²⁷ Representative of a national export control authority, Interview with authors, 17 Oct. 2025.

²⁸ See Deloitte, ‘Trade and Customs Duty, Export Compliance’, [n.d.].

²⁹ Senior compliance officer in an international semiconductor research organization, Interview with authors, 16 Sep. 2025.

ing systems retain enough data on previous applications and classifications, LLMs could be trained to improve their performance.

Such tools could be used to provide a first verification indication of the classification done by the exporter and a contextual analysis of whether submitted documents share characteristics with previous false classifications or falsified documents. For example, the Korean Institute of Nuclear Nonproliferation and Control has deployed the Intelligent Export Control Review System (IXCRS) to assist initial classification of large quantities of nuclear items as part of supply arrangements for nuclear power plants or research reactors.³⁰ At least initially, such AI tools tend to use proprietary models and be custom trained on selected data in a government system, without being open to the public. One public-facing example is the My Australian Defence Exports AI ('MADE AI') assistant created by the Australian Department of Defence that was trained on the Australian national export control legislation and provides general guidance on common questions related to Australia's defence export control legislation via a website.³¹ The MADE AI assistant provides general information on applicable legislation via a chat function or frequently asked prompts. Both the website and the user help guide acknowledge that the assistant may generate inaccurate or incomplete responses and cannot substitute for independent legal advice.³² The US Directorate for Defense Trade Controls has also recently set up an AI-enabled virtual agent and live assistance function to provide user information on licensing applications and related issues.³³ Importantly, such AI tools are only assistive systems that can help inform an initial triaging and assessment—a human licensing official still needs to understand the legislation, review any AI tool's recommendations, and take the final decisions on any classification or legal assessment as part of the licensing process.

Use case: AI for licensing facilitation eligibility determination

Another application of AI that aims to improve efficiency in licensing processes is using AI assistance tools for analyzing licensing applications to determine if they are straightforward cases that would be eligible for any licence facilitations that states offer. Some exporters might not be aware of the eligibility of their export for such a licence facilitation. Many states with export control systems offer general licences (or general export authorizations) that an exporter can use—under certain conditions—instead of applying for a single licence for each individual export. Rather than applying for the general licence, the exporter is only required to fulfil certain record-keeping requirements and make initial or aggregate declarations on the use of the general licence. General licences usually have a specific scope of items and destinations, as well as conditions and requirements for use, limiting their applicability to a subset of exports.³⁴ Global licences authorize an exporter for multiple exports with a specific scope for items or groups of items, recipients or groups of recipients, and destination

³⁰ Tae, J. and Shin, D., 'Keyword clustering for comparing documents in different languages', *International Journal of Machine Learning and Computing*, vol. 5, no. 4 (Aug. 2015); and Shin, D., 'Applications of artificial intelligence in the export control domain', eds Y. Bi, S. Kapoor and R. Bhatia, *Proceedings of SAI Intelligent Systems Conference (IntelliSys) 2016*, vol. 2 (Springer: Cham, 2018).

³¹ Australian Department of Defence, 'MADE AI', [n.d.].

³² Australian Department of Defence, 'MADE AI' (note 31); and Australian Department of Defence, 'My Australian Defence Exports Artificial Intelligence (MADE AI) help guide', [n.d.].

³³ US Department of State, Directorate of Defense Trade Controls (DDTC), 'Don't wait—get immediate support with DDTC's AI-enabled virtual agent and live assistance', Announcement, 12–16 Jan. 2026.

³⁴ See e.g. Regulation (EU) 2020/821 of the European Parliament and of the Council of 20 May 2021, Annex II: Union General Export Authorisations, *Official Journal of the European Union*, L206, 11 June 2021, p. 425.

or destinations, over a specific period.³⁵ Both types of licensing facilitations can help reduce the volume of individual licence applications that exporters need to prepare and submit and that licensing authorities need to process.

As part of the compliance process, AI tools can help identify exports that are eligible for a general or global licence. For a licensing authority, similar tools can help identify when a general licence could be used instead of several single licences, or where a global licence would be appropriate, either during initial screening of licence applications or as feedback to applicant exporters for future licensing applications. Such reductions in the burden on both exporters and licensing authorities can help increase the overall efficiency of export controls and limit their impact on legitimate exports for peaceful purposes.

In states using an electronic licensing system, one option would be to train a proprietary and closed government AI model on the scope and requirements of each general licence currently offered by that state, as well as on the data submitted as part of reviewed declarations on their use. The AI tool would be designed to analyse licensing applications submitted in a standard digital and machine-readable format, and national licensing officials could use that to assist them in their initial screening of applications. Possible challenges in the implementation of such a tool in current licensing systems include differing data formats and incompatible APIs of the electronic licensing system and the declaration and reporting systems used.³⁶ The possibility of data poisoning and adversarial examples deceiving the model introduced by direct input or interaction with the tool by third parties means that such systems might not be public-facing and instead remain an internal assistive tool for national licensing officials.³⁷

Use case: DLT for export documentation authentication

DLT is already being deployed in a variety of sectors as a tool to verify the provenance and authenticity of certain items. Pilot studies have developed prototype DLT systems that would support the Chemical Weapons Convention and the work of the International Atomic Energy Agency. These studies could be instructive for consideration of the specific design and choice of DLT solutions for any applications aimed at strengthening aspects of the biological weapons prohibition regime.³⁸ However, the adoption of DLT solutions has not always been successful. For example, the TradeLens DLT platform created by Maersk and IBM for shipping companies failed, at least in part, from a lack of buy-in, incentives and institutional support.³⁹

A specific application of DLT that could help strengthen national export controls is a DLT-based system for authenticating end-use and end-user documentation using a secure distributed ledger among all transaction parties and relevant national authorities with tiered levels of permission.⁴⁰ This type of DLT-based system administrated by national export control authorities would enable them to verify that end-use and

³⁵ For a general overview of types of licences see e.g. European Commission, Directorate-General for Trade and Economic Security, 'Exporting dual-use items', [n.d.].

³⁶ Representative of a national export control authority, Interview with authors, 17 Oct. 2025.

³⁷ Representative of a national export control authority, Interview with authors, 17 Oct. 2025.

³⁸ Examples include the MATCH 2.0 prototype DLT for resolving chemical trade discrepancies under the Chemical Weapons Convention and the SLAFKA prototype DLT for nuclear safeguards information management. See Vestergaard, C. et al., 'MATCH 2.0: A new ledger for nonproliferation', Stimson Center, May 2025; and Vestergaard, C. et al., 'SLAFKA: Demonstrating the Potential for Distributed Ledger Technology for Nuclear Safeguards Information Management', Stimson Center Nonproliferation Report, Nov. 2020.

³⁹ Bousquette, I., 'Blockchain fails to gain traction in the enterprise', *Wall Street Journal*, 15 Dec. 2022.

⁴⁰ Brockmann and Héau (note 24); and Cándano Laris, D., 'Blockchain applications for export control compliance and global supply chain integrity', ed. C. Vestergaard, *Blockchain for International Security: The Potential of Distributed Ledger Technology for Nonproliferation and Export Controls* (Springer: Cham, 2021), p. 95.

end-user certificates are not falsified, altered or otherwise tampered with throughout the licensing process, shipping, and potential transit and trans-shipment, as well as post-delivery—preventing illicit procurement for biological weapon programmes.⁴¹ A permissioned blockchain that creates a closed network with additional control layers that limit access to information on the ledger to authenticated users with varying levels of permission administered by the national licensing authorities would be one such solution that could also maintain privacy and chain-of-custody standards.⁴² Requiring the use of the permissioned blockchain system would not prevent the possibility of initial submissions of falsified documents or use of false identities altogether, but could significantly increase the likelihood of detection and traceability by requiring digital verification of identity and by time-stamping and hashing each modification of the documentation.⁴³

DLT includes a variety of different options of models, consensus mechanisms, encryption types and architectures that are combined into, for example, a permissioned blockchain solution. Each comes with different strengths and limitations, including ones related to required computing power and energy consumption. Broadly, a solution for end-use and end-user documentation authentication would require a careful layering of permissions that provides the appropriate access to the national export control authorities of the exporter, the relevant national authorities in transit, trans-shipment and recipient states, and the exporter and recipient, while limiting access to certain parts of the data that might impact competitiveness or involve technology transfer. In most cases, different DLT models involve a provider that executes certain maintenance and administrative functions while maintaining security through additional layers of cyber security. One possible challenge in setting up such a system could be the selection of a trusted provider who, if their solution is adopted internationally, is both willing to engage with and is acceptable to all parties connected to the system. As a baseline, the system would, at least initially, also need to be able to work with a variety of data formats used in national templates and be built to have a widely compatible API, as well as limited costs, a dedicated onboarding function and sufficient time to enable widespread adoption.⁴⁴

⁴¹ Gahlaut, S., 'Mental block? Time to revisit the potential of distributed ledger technology', *WorldECR*, 9 Nov. 2021; and Cándano Laris (note 40), p. 97.

⁴² Cupitt, R. T., 'Blockchain for global trade in dual-use chemicals', ed. C. Vestergaard, *Blockchain for International Security: The Potential of Distributed Ledger Technology for Nonproliferation and Export Controls* (Springer: Cham, 2021), p. 83.

⁴³ Cándano Laris (note 40), p. 100.

⁴⁴ For a discussion of these different factors in a different example of a prototype DLT solution see Vestergaard, C. et al., 'MATCH 2.0: A new ledger for nonproliferation' (note 38), pp. 11–15.

4. Reporting and transparency mechanisms

Reporting and transparency mechanisms are the primary method for states to demonstrate good-faith implementation and support confidence building in the absence of a standing verification mechanism under the BWC. In practice, states must generate, maintain and transmit overlapping but non-identical sets of information on national implementation, often across multiple ministries and agencies, including public health, research oversight, customs administration and law enforcement bodies. National points of contact therefore aggregate information that they do not directly control, and the effectiveness of transparency mechanisms depends less on the existence of formal obligations than on administrative capacity to manage information flows consistently over time.

Under the BWC, states parties are invited to submit annual CBMs which are voluntary but widely regarded as the principal transparency channel in the regime. Resolution 1540 requires states to establish domestic controls to prevent non-state actors from acquiring, developing, transporting, transferring or using nuclear, chemical or biological weapons and related materials, including through legislation, enforcement, and border and export controls. Compliance relies on reporting to the UN Security Council's 1540 Committee. Unlike the CBM process, Resolution 1540 requires an initial national report and then relies largely on additional information and updates provided when appropriate or upon request, rather than a routine annual cycle.⁴⁵ In parallel, the Australia Group supports national implementation through non-binding guidelines and export control lists that many states use as technical reference points, but it does not impose reporting obligations.

Empirical data indicates that the scale of these reporting obligations represents a challenge for many states (see figure 4.1). According to the BWC Implementation Support Unit, 136 of 189 states parties (72 per cent) have submitted at least one CBM report since 2020, while 53 (28 per cent) did not provide any reports during that period. Of those reporting at least once, 60 states parties (32 per cent) did so every year, while 76 (40 per cent) reported intermittently, indicating uneven participation in annual reporting. As of mid-2025, 35 states parties (19 per cent) had not submitted a CBM report since the transparency mechanism was introduced in 1987.⁴⁶ Reporting under Resolution 1540 shows a similar pattern: of 193 member states, 185 (96 per cent) have submitted an obligatory initial report; however, follow-on reporting remains uneven, with only 124 states (64 per cent) providing additional information and just 38 (20 per cent) submitting national implementation action plans.⁴⁷

Analyses of both BWC and Resolution 1540 implementation consistently indicate that underreporting is driven less by political resistance than by capacity constraints, administrative burden and limited institutional continuity.⁴⁸ These patterns highlight that transparency deficits are often procedural rather than political, rooted in how implementation-related information is collected, structured and preserved at the national level. These procedural constraints highlight the potential value of states

⁴⁵ United Nations, Security Council, 1540 Committee, '1540 fact sheet', [n.d].

⁴⁶ Biological Weapons Convention (BWC), Working Group on the Strengthening of the BWC, Sixth Session, 'Background information document on "Measures on confidence-building and transparency"', Submitted by the Implementation Support Unit, BWC/WG/6/2, 16 July 2025, pp. 8 and 15.

⁴⁷ United Nations, Security Council, '2022 comprehensive review of the status of implementation of Security Council Resolution 1540 (2004)', S/2022/899, 1 Dec. 2022.

⁴⁸ de Vries, B., 'Recent developments in the national implementation of Biological Weapons Convention: What happened since Resolution 1540?', *Journal of Conflict and Security Law*, vol. 28, no. 3 (Winter 2023).

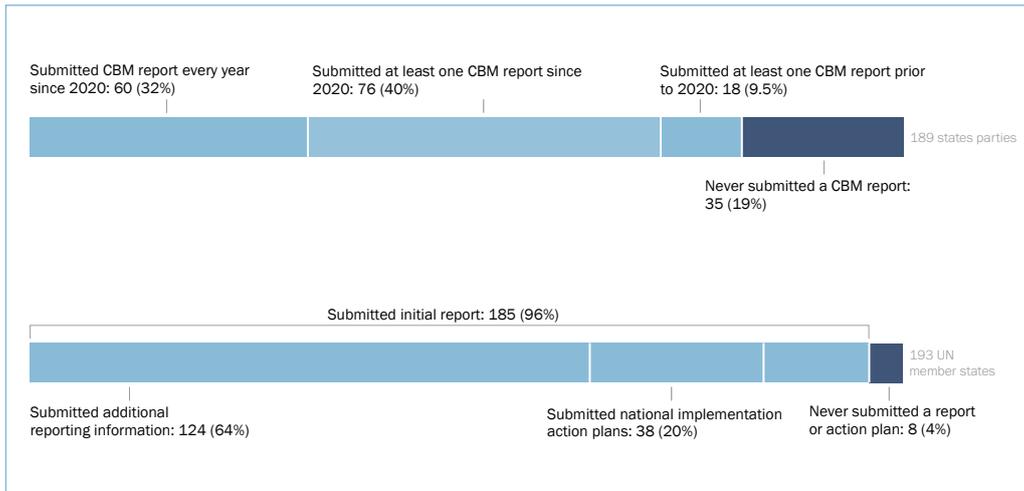


Figure 4.1. Status of reporting obligations under biological weapons prohibition regime measures. Biological Weapons Convention (BWC) states parties submitting confidence-building measures (CBM) reports since the practice was introduced in 1987 (top) and United Nations member states reporting under UN Security Council Resolution 1540 (bottom)

adopting tools that can help reduce the burden of compiling and reconciling implementation information while preserving accountability and confidentiality.

Use case: AI for reporting and transparency

AI already supports document analysis and pattern recognition in adjacent reporting and compliance contexts, including customs risk management and financial integrity frameworks, where authorities must structure large volumes of heterogeneous documentation and support routine reporting obligations.⁴⁹ Public evidence of systematic AI use for BWC CBM preparation is limited, but the underlying tasks are similar: extracting structured fields from heterogeneous documents, checking internal consistency across reporting cycles, and generating draft narratives for human review.

CBMs arrive as both structured and unstructured submissions—for example, scanned PDFs of completed forms, electronic templates with fillable fields, and free text narrative annexes—which complicates comparison across years and between states. LLMs can analyse this wealth of information and convert it into structured datasets by extracting facility types, activity descriptions, legislation references, and other key fields in a consistent manner. This could preserve internal reliability by allowing states parties to review their national implementation submissions more systematically before transmission. In practice, this would likely involve a secure, enterprise-deployed LLM or document-processing system connected to a government’s internal repositories of legislation, licensing records and prior submissions, for example through retrieval-based access controls in tools such as Microsoft 365 Copilot, or comparable government-hosted systems.⁵⁰

AI can be used to support internal quality control by flagging inconsistencies within a state’s reporting that might otherwise require extensive manual review. For

⁴⁹ World Customs Organization (WCO), *Detailed Report on the Adoption of Artificial Intelligence and Machine Learning in Customs* (WCO: Brussels, Mar. 2025); Financial Action Task Force (FATF), *Opportunities and Challenges of New Technologies for AML/CFT* (FATF: Paris, 2021); and Bank for International Settlements, ‘The use of artificial intelligence for policy purposes’, Report submitted to the G20 Finance Ministers and Central Bank Governors, Oct. 2025.

⁵⁰ Microsoft, ‘Microsoft 365 Copilot architecture and how it works’, 29 Oct. 2025.

example, inconsistencies between successive CBM submissions, such as changes in declared laboratory or facility information, or mismatches between reported national implementation legislation and referenced regulations, could be automatically flagged for internal review rather than identified only through manual reconciliation. Public authorities already use pattern-recognition models in other domains to surface items that merit follow-up instead of issuing automated determinations.⁵¹ The incentive for states is practical: fewer staff hours spent on report reconciliation, and a lower risk of omissions that undermine credibility, particularly where reporting responsibilities are concentrated in a small number of staff or subject to frequent rotation.

At the national level, CBM and Resolution 1540 reporting requires cross-ministry aggregation of implementation information that is owned and maintained in different institutional systems. Under the BWC CBM process, for example, information on legislation and regulations sits in different institutional channels from information on laboratories and facilities, vaccine production capacity, biodefence-related activities, or disease outbreaks, which variously fall under public health, science, agriculture, defence, interior or other authorities depending on national arrangements.⁵² Under Resolution 1540, reporting similarly spans prohibitions and enforcement measures, accounting and security, physical protection, and border and export controls, which typically map onto different agencies and legal owners.⁵³ In that context, machine-learning methods can be used to support internal reconciliation by mapping equivalences across datasets, flagging where categories do not align, and generating draft tables or narrative summaries for human validation. This is best understood as an administrative aid, not an automated compliance assessment, and the closest operational analogues are in government functions that already use AI to structure heterogeneous evidence bases and support monitoring and review.⁵⁴

The operational value of AI increases during periods of political or administrative disruption. When staff rotate, when ministries reorganize, or when routine coordination is slowed by external factors (e.g. restrictions on travel, reduced information exchange or engagement with foreign counterparts and assistance providers, or heightened clearance requirements for sensitive data), reporting timelines often become dependent on informal workarounds. In such circumstances, maintaining structured records in consistent formats can reduce reliance on ad hoc interagency consultation and tacit knowledge by making prior reporting inputs and source documents easier to retrieve, compare and reuse. This could support continuity in national reporting processes through more consistent submissions over time and confidence building among states parties. These dynamics are widely observed across government AI use cases, where capacity constraints, legacy systems and skills gaps shape whether digital tools improve continuity or introduce new bottlenecks.⁵⁵

At the same time, greater automation may increase the distance between national points of contact and the underlying data they submit, raising questions about accountability, interpretive responsibility, and ownership of judgements embedded in automated outputs. These considerations do not negate the value of AI-assisted reporting, but they underscore the importance of maintaining clear lines of human review and responsibility within reporting workflows. Ensuring that automated tools augment,

⁵¹ World Customs Organization (note 49); and Bank for International Settlements (note 49).

⁵² UNODA, *Guide to Participating in the Confidence-building Measures of the Biological Weapons Convention*, Revised edn (UNODA: Geneva, Feb. 2015).

⁵³ United Nations, Security Council, 1540 Committee, '1540 Matrices', [n.d.].

⁵⁴ OECD, *Governing with Artificial Intelligence* (note 14).

⁵⁵ OECD, *Governing with Artificial Intelligence* (note 14).

Box 4.1. Artificial intelligence, information analysis and the Biological Weapons Convention verification debate

Since the negotiation of the Biological Weapons Convention (BWC), verification has remained a persistent point of contention. Advances in digital technologies have renewed interest in whether some form of technology-enabled analysis could contribute to compliance assessment. Artificial intelligence (AI) is often cited as a potential enabler of future verification arrangements due to its ability to process large volumes of heterogeneous data. For example, in September 2025 the United States president, Donald Trump, promised to ‘lead an international effort to enforce the Biological Weapons Convention . . . by pioneering an AI verification system that everyone can trust’.^a This highlights the recurring appeal of technology as a potential substitute for institutional verification arrangements.^b

In theory, AI-enabled tools could support analysis of open-source information, satellite imagery, scientific publications and confidence-building measure submissions, to identify patterns or anomalies that might warrant closer examination. Comparable techniques are being explored in other domains, including the use of machine learning to assist safeguards analysis in the nuclear field.^c However, such approaches are best understood as analytical aids that can augment human review, rather than as independent compliance determinations.

In the biological context, these approaches remain largely conceptual. Many of the data streams cited in BWC verification discussions—such as social media activity, preprint repositories and genomic datasets—are incomplete, context-dependent and susceptible to misinterpretation. Biological research is widely distributed, dual-use by nature, and often indistinguishable in form between peaceful and prohibited activities. As a result, algorithmic identification of non-compliance would face significant challenges in reliability, attribution and evidentiary sufficiency.

Moreover, the BWC lacks an institutional verification body and agreed procedures through which AI-generated analyses could be assessed, contested and acted upon. Without shared standards for data selection, model validation and interpretability, the introduction of AI into compliance assessment risks amplifying disagreement rather than building confidence. These constraints help explain why, despite renewed technological interest, AI-enabled verification has not progressed beyond exploratory discussion within the BWC framework.

^a US Department of State, ‘President Trump delivers remarks to the United Nations General Assembly’, YouTube, 23 Sep. 2025.

^b Fluegel, L., ‘For bioweapons experts, Trump’s UN speech presents a window of opportunity’, Carnegie Endowment for International Peace, 4 Dec. 2025.

^c International Atomic Energy Agency, International Symposium on AI and Nuclear Energy: AI, nuclear energy and the IAEA, 2025; and Cui, Y. et al., ‘Using deep machine learning to conduct object-based identification and motion detection on safeguards video surveillance’, Report submitted to the Symposium on International Safeguards, Vienna, 5–8 Nov. 2018.

rather than obscure, human judgement is essential if digital reporting systems are to strengthen confidence rather than introduce new sources of friction.

Use case: DLT for data integrity and record assurance

In the reporting context, DLT is most plausibly used as a provenance layer that preserves integrity signals, such as hashes and time stamps, for implementation-related records without storing sensitive content. This differs from facility-level applications, where DLT is sometimes proposed for chain-of-custody logging for controlled materials. A practical analogue is the US FDA’s Pilot Project Program that explored DLT approaches for interoperable record exchange and auditability across multiple actors, such as TraceLink’s pilot on traceability and digital recalls across pharmaceutical supply chains.⁵⁶

⁵⁶ TraceLink, ‘FDA Pilot Project Program: DSCSA 2023 traceability with blockchain/distributed ledgers and digital recalls network pilots’, FDA Pilot Program Final Report, [n.d.].

DLT is particularly relevant where reporting responsibility spans multiple ministries or shifts across agencies, because it can provide a durable audit trail of when specific records were created, updated or transmitted. In principle, this could reduce the administrative cost of reconstructing reporting histories during reviews, assistance requests and follow-up queries by making it easier to demonstrate that a given record existed at a particular time and has not been modified since. A permissioned ledger could, for example, store time-stamped hashes of CBM submissions, national legislation updates and supporting documents, allowing future reviewers to compare and verify document versions. This addresses a routine administrative challenge: maintaining continuity and traceability of official records across staff turnover, system migrations and changes in domestic information management practices. However, as with other digital assurance tools, DLT does not remove the need for governance and quality control processes around what data is recorded, and it cannot resolve disagreements over interpretation or completeness.

In the multinational context, DLT can support confidence building through cryptographic verification which provides assurance that documents have not been modified, while access permissions ensure that underlying files remain under national control. This supports confidence building when trust in institutions, access arrangements or political relations fluctuates. At the same time, metadata is not automatically benign. Even when content is held off chain, metadata can still create operational or security sensitivities, and it can be targeted through cyber operations, misconfiguration or insider misuse. The practical value of DLT to support states in their national implementation reporting processes therefore depends on careful system design, including minimal metadata exposure, access controls, and clear procedures for who can write to and query the ledger.

Although this section focuses on technology-enabled improvements to reporting integrity and confidence building under existing mechanisms, interest persists in whether AI-enabled analysis could contribute to treaty verification, an issue addressed separately in box 4.1.

5. Cross-cutting considerations

The potential for AI- and DLT-enabled tools to strengthen the key functions of the biological weapons prohibition regime is not technologically determined. Technical capabilities shape what is feasible, but outcomes will continue to depend on governance choices, policy environments, institutional incentives, and norms and standards, including how states manage accountability, access to data, intellectual property constraints, and the distribution of digital capacity across states and sectors.⁵⁷ Several cross-cutting considerations shape whether the wider adoption of AI- and DLT-enabled tools will strengthen oversight and confidence building or introduce new frictions. Where possible, states should seek to use the regime's cooperation and assistance channels to support standard-setting and to enable responsible, equitable distribution.

Data integrity and governance

Most applications of AI, DLT and other modern digital tools rely on large datasets. However, the considerable variation in size, quality and completeness of such datasets can result in biases, limited performance and erroneous outputs, including spurious correlations and unsupported inferences. Many datasets are proprietary, only accessible at considerable cost or only collected and cleaned with specific custom applications in mind.⁵⁸ These characteristics mean that there are considerable differences in access and integrity of data, raising persistent questions about provenance and validation. On a small scale, this means that AI-enabled analytical tools in particular are always limited to the extent to which these different factors can be optimized by system administrators and users. For digital oversight measures relying on such datasets to be credible, the parties providing the data must share how the data is generated, collected, curated and interpreted—potentially across states, industry sectors, research fields and academia. Models trained on unverified, incomplete or small datasets can misrepresent activity. Records without demonstrable provenance have limited evidentiary value.

Other sectors have developed standards and guidance relevant to data assurance and algorithmic governance through bodies such as the International Organization for Standardization (ISO) and the OECD.⁵⁹ In laboratory practice, standards such as ISO 35001 also specify expectations for biorisk management systems, including documentation and record control.⁶⁰ However, comparable principles tailored to the use of AI and DLT in biosecurity and non-proliferation remain underdeveloped, particularly where cross-border confidence-building and assurance are the objective. Steps towards the further development and widespread adoption of AI- and DLT-enabled systems to strengthen the biological weapons prohibition regime must therefore be accompanied by progress towards the developments of dedicated norms and standards for this area of application.

⁵⁷ Brockmann and Héau (note 24); and Smith, M., 'Use cases for emerging technologies to strengthen high-containment laboratory governance', SIPRI Research Policy Paper, Dec. 2025.

⁵⁸ See National Academies of Sciences, Engineering, and Medicine, 'Importance of data in AI-enabled biological models', *The Age of AI in the Life Sciences: Benefits and Biosecurity Considerations* (National Academies Press: Washington, DC, 2025).

⁵⁹ International Organization for Standardization (ISO), 'Harnessing international standards for responsible AI development and governance', ISO Policy Brief, 2025; and OECD, *Reinforcing Regulatory Frameworks through Standards, Measurements and Assurance: Making Better Use of Quality Infrastructure in Policymaking* (OECD Publishing: Paris, 2025).

⁶⁰ ISO, 'ISO 35001:2019 Biorisk management for laboratories and other related organisations', Nov. 2019.

Human oversight and confidence in AI systems

AI and automation introduce new forms of uncertainty into decision-making processes that have traditionally been largely human-driven. Even when algorithms improve efficiency, they can obscure accountability or produce outcomes that users cannot easily explain. Confidence in AI-assisted compliance therefore rests not only on transparency of data but also on auditability and confidence in judgement—that decisions remain interpretable, contestable, and ultimately subject to human responsibility and accountability. Maintaining this distinction is analytically important for assessing whether technology-assisted oversight strengthens confidence or introduces new frictions, particularly when decisions depend on outputs that are not easily explainable or contestable.⁶¹

Interoperability and policy coherence

Another cross-cutting challenge relates to the different setup, data formats, APIs and software used across states and across relevant implementing authorities within each state. This is demonstrated by the challenges experienced by export control authorities in setting up AI-enabled item classification and licensing facilitation eligibility systems requiring interoperable interfaces with digital licensing systems (see section 3). The obligations created by the BWC relate to a variety of different processes and activities which are governed by policies and legislation spanning economics, health, environment, security, defence and foreign policy. The implementing legislation and the responsibilities of different agencies vary and, in most states, so do the digital systems through which they are operated, frequently resulting in interface and interoperability issues. This is not surprising, as the implementation of different aspects of the regime and their respective objectives are put into effect as part of wider policy and administrative systems that are aligned with the primary objectives of their respective policy fields. The BWC, Resolution 1540 and the Australia Group guidelines are mutually reinforcing, but have different compositions of states parties, UN member states, and participants and adherents.⁶² This means that they effectively cannot be discussed and aligned through one forum, resulting in fragmented responsibilities on the national authorities' side and in terms of their areas of focus of implementation. Divergent classification systems, licensing platforms and data-exchange requirements complicate the integration of digital compliance and record-keeping tools that depend on consistent metadata and record formats. However, interoperability challenges are often less about tool selection than about shared data formats, terminology and standards that allow systems to exchange information reliably.⁶³ A voluntary, shared approach to technical standards for interoperability could enable the adoption of emerging applications of AI and DLT to help strengthen the BWC regime.

Balancing transparency and security

Digital transparency can be both an enabler and a limitation for confidence building, which depends on the ability to exchange information without revealing data that could compromise safety, security or proprietary interests. For example, systems that analyse facility data, satellite imagery, laboratory records, or product classifications

⁶¹ See OECD, Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449, adopted 22 May 2019, amended 3 May 2024, paras 1.2 and 1.3.

⁶² See the list of participants on the official UN treaty page for the Biological Weapons Convention (note 3); and Australia Group, 'Participants', [n.d.].

⁶³ International Standards Organisation, 'ISO 22739:2024 Blockchain and distributed ledger technologies', 2024.

and licensing information can strengthen confidence only if they also safeguard sensitive details. Other sectors have approached this dilemma through cryptographic and privacy-preserving technologies that allow aggregated or encrypted data to be compared without full disclosure.⁶⁴ The absence of analogous norms within the biological weapons prohibition regime leaves states to interpret transparency and confidentiality independently, which limits comparability and can erode confidence even among compliant parties. The same applies to the ability to benefit from mutual transparency in line with the commitments under Article X of the BWC, where designing transparency arrangements in lopsided ways that impose restrictions beyond appropriate security and competitiveness rationales could benefit specific parties over others. This means, for example, that the design of permissioned DLT systems needs to reflect common norms on how to balance transparency and security, and, where applicable, economic interests. Security authorities note that DLT can support integrity and traceability but does not eliminate security risks, and that careful design and governance remain essential.⁶⁵

Access, equity and capacity

The development of any system using emerging technologies, including AI and DLT, raises the question of whether all states will have equitable access to such systems and how they would attain the capabilities to set up, operate and maintain them. Article X of the BWC frames cooperation and assistance as an obligation and a right, reinforcing that access to tools and capabilities is part of regime legitimacy. The varying nature of paper-based and digital systems of applications, certification and record-keeping are illustrative of the cross-cutting issues that any AI- or DLT-based system will face. It also means that for each system proposed to be made part of international standards or otherwise compulsory, the cost associated with ensuring equitable access and capability development and upkeep must be factored in.

Unequal access to digital infrastructure and data-governance expertise remains one of the clearest determinants of how states can participate in technology-enabled implementation. Advanced economies often possess the computing power, cloud storage and computing capacities, and cybersecurity systems needed to integrate AI or DLT into oversight processes, while many developing economies with limited resources do not.⁶⁶ This asymmetry, where only some states can operationalize digital transparency and assurance mechanisms, risks exacerbating inequalities among states parties. The difference in access is often already reflected in standard-setting and governance of relevant emerging technologies, potentially reinforcing such effects.⁶⁷ Open-access tools and shared training resources have begun to narrow these divides in other scientific fields, but similar mechanisms have yet to emerge within the BWC regime. Any progress towards piloting and adoption of AI or DLT solutions as standards for fulfilling BWC obligations thus needs to involve equitable participation in the process and provide the necessary resources through technical cooperation and assistance activities, as mandated by Article X.

⁶⁴ See US FDA, 'DSCSA Pilot Project Program' (note 19); and Salingros (note 21).

⁶⁵ UK National Cyber Security Centre, Distributed ledger technology, 2021.

⁶⁶ Representative of a BWC state party, Remarks shared during 'Technological opportunities for the biological weapons prohibition regime', SIPRI workshop, Stockholm, 20–21 Nov. 2025.

⁶⁷ Representative of a BWC state party, Remarks shared during 'Technological opportunities for the biological weapons prohibition regime' (note 66).

Experience-sharing and collaboration

Lessons learned from the adoption of AI- and DLT-enabled tools remains fragmented across sectors, vendors and jurisdictions, in part because many deployments are embedded in internal systems and are not documented publicly. This creates a cross-cutting constraint on responsible uptake: states and regulated actors often lack shared visibility on what approaches have reduced administrative burden, improved auditability or introduced new frictions.⁶⁸ Shared learning and comparability therefore become enabling conditions for technology-assisted implementation, alongside governance and capacity.

The biological weapons prohibition regime's existing set of cooperation and assistance channels provide entry points for experience-sharing without the need to create new structures. The Working Group on the Strengthening of the BWC provides a structured venue to discuss scientific and technological developments and international cooperation under Article X.⁶⁹ The Australia Group and Resolution 1540 processes also convene technical and regulatory communities working on implementation and compliance, and offer channels for sharing effective practices and capacity-building approaches.⁷⁰ Used pragmatically, these forums can support diffusion of implementation lessons and reduce duplication across states.

These instruments form an existing multilateral ecosystem capable of supporting gradual diffusion of technology-enabled practices. Leveraging them would ideally include a shared, deliberate focus on documenting experiences, comparing approaches, and identifying conditions under which digital tools have strengthened implementation. In this sense, multilateral engagement serves not to promote technology adoption for its own sake, but to ensure that innovation contributes to trust, equity and long-term regime resilience.

⁶⁸ Representative of a BWC state party, Remarks shared during 'Technological opportunities for the biological weapons prohibition regime' (note 66).

⁶⁹ BWC, Working Group on Strengthening the BWC, Seventh session, Revised draft final report, Submitted by the chairperson of the working group, BWC/WG/7/CRP.1/Rev.1, 12 Dec. 2025.

⁷⁰ United Nations, Security Council, 1540 Committee, 'Experiences shared, lessons learned, and effective practices', [n.d.]; and United Nations, Security Council, 1540 Committee, 'General information', [n.d.].

6. Conclusions

AI- and, to a lesser extent, DLT-enabled tools are already being used to strengthen production and laboratory oversight and export controls on dual-use items. Specific areas of application include laboratory management, supply chain regulation, workflow reporting, export control implementation and export control compliance. To date, AI- and DLT-enabled tools have not been used to support BWC and Resolution 1540 reporting and transparency practices. However, governments and research institutes have indicated support for their adoption and highlighted potential applications. For all three regime functions, AI- and DLT-enabled tools deployed in other sectors and areas of arms controls provide models that could be adopted. Figure 6.1 presents illustrative examples of where AI and DLT can strengthen regime functions.

The near-term value of these tools lies in reducing administrative burden, improving record integrity, and supporting continuity in workflows that currently rely on fragmented data ownership, manual compilation, and institutional memory concentrated in a small number of officials. However, these benefits are not automatic. Adoption and implementation will strengthen confidence only if governance choices preserve accountability, confidentiality and human oversight, and if capacity gaps do not create inequalities in which states can adopt and sustain such approaches.

Emerging technologies will not replace the political work of disarmament, but they can reinforce the administrative and informational foundations on which transparency and confidence building depend. Their contribution is therefore contingent on collective approaches to learning, standard-setting and capacity development rather than isolated national experimentation. As states, companies and research institutes adopt and integrate these technologies in ways that seek to strengthen the key functions of the biological weapons prohibition regime, they will need to think of innovative ways in which they can be made available to the developing world. If the tools that are being developed bring real benefits in terms of tackling the proliferation of biological weapons, then there is a clear need to make versions of them as widely available as possible.

The following recommendations focus on practical steps that can be taken through existing national systems and multilateral processes to support responsible, equitable and confidence-enhancing adoption of emerging technologies.

By approaching emerging technologies as instruments for shared assurance rather than competitive advantage, states parties can strengthen implementation of, reinforce confidence in, and sustain the relevance of the biological weapons prohibition regime in an era of rapid technological change. These recommendations seek to ensure that today's innovations evolve under conditions of trust and equity. They identify pathways for states to integrate emerging technologies into existing BWC commitments while reinforcing transparency, cooperation and accountability. These recommendations are aimed primarily at states parties and specify where action is most realistically taken by laboratories and industry individually or through national systems, and where multi-lateral forums can support learning, comparability and capacity building.

Regime function	Artificial intelligence can enable	Distributed ledger technology can enable
 Production and laboratory oversight	Anomaly detection for samples, reagents, or equipment	Event documentation for material transfer, inventory verification, and chain-of-custody
 Export controls on dual-use items	Biological item export classification	Authentication of end-use and end-user documentation and auditable licensing records
 Reporting and transparency mechanisms	CBM document processing, field extraction, and consistency checks	Version traceability and integrity proofs for submissions and supporting documentation

Figure 6.1. Potential applications through which artificial intelligence (AI) and distributed ledger technology (DLT) can strengthen key functions of the biological weapons prohibition regime

Note: These examples are neither exhaustive nor assessments of readiness.

Recommendations

Create structured spaces for innovation and dialogue within the BWC

A central barrier to responsible adoption of AI- and DLT-enabled tools is that states lack shared visibility on what has worked, what has not, and what constraints are real. The Working Group on the Strengthening of the BWC offers a practical venue to normalize discussion of the adoption of these tools. The aim of these exchanges should be comparability, learning, and risk-aware diffusion of practice. Notably, the working group is time-limited and mandate-bound, so any structured technology dialogue would need to be advanced within its current window or carried forward through successor arrangements agreed by states parties.

- Establish a recurring agenda item that captures lessons learned from national pilots and implementations of AI and DLT in oversight, compliance and reporting contexts.
- Use light, voluntary templates for describing use cases, inputs, outputs and limitations, so information can be compared across states without over-disclosure.
- Encourage structured technical contributions, from laboratories, compliance functions and service providers to national authorities, to ground discussion in operational realities.
- Produce a short periodic synthesis that identifies common bottlenecks and promising practices, and flags issues requiring additional governance attention.

Treat digital capacity as part of Article X cooperation and capacity-building

Digital infrastructure and analytical capability increasingly shape states' ability to implement biosafety, biosecurity and reporting functions. Treating these capabilities as legitimate areas of cooperation and assistance can make Article X support more relevant to current implementation realities, especially where administrative burden and data fragmentation are persistent constraints.

- Incorporate digital systems and skills into needs assessments and assistance discussions as implementation enablers, alongside infrastructure and workforce development.
- Develop practical reference workflows for secure data management, auditable record-keeping and reporting preparation that can be adapted to different national contexts.
- Encourage voluntary narrative updates on digital capacity as part of Article X-related progress discussions, focusing on practical lessons rather than benchmarking.
- Support training approaches that prioritize operational competence, such as data governance, quality control, and responsible use of automated tools in compliance workflows.

Leverage complementary forums without creating parallel standards

The Australia Group and Resolution 1540 processes convene technical and regulatory communities working on aspects of implementation and compliance related to either production and laboratory oversight or export controls on dual-use items. These channels can also support experience sharing and capacity building. These efforts should be coordinated with BWC processes to avoid producing duplicative guidance or competing documentation practices that might increase fragmentation.

- Use existing Australia Group and Resolution 1540 channels to share experiences with AI and DLT tools, focusing on practical adoption constraints, governance approaches and lessons learned.
- Identify where documentation practices overlap across instruments—for example, definitions, metadata fields and record-retention expectations—and where voluntary convergence on machine-readable templates, consistent metadata conventions and version control practices could reduce administrative burden.
- Encourage interoperability discussions centred on formats, metadata and auditability rather than specific products or vendors.
- Avoid creating new reporting requirements or duplicative standards by ensuring that outputs are framed as guidance and practice exchange, not as new obligations.

Build a shared foundation for digital trust

Digital tools are only as reliable as the data, processes and governance practices that underpin them. Trust is strengthened when states can explain how records are generated, validated, reviewed and preserved, and when auditability is maintained over time. The objective should not be to prescribe the adoption of specific technologies, but to increase confidence in how technology-assisted workflows are produced and interpreted.

- Encourage transparency about how AI and DLT are being used to strengthen key regime functions, including what data sources are relied upon and how human review procedures are documented.
- Promote auditability as a design requirement, ensuring that records and decisions can be reconstructed and explained over time.

- Support shared vocabulary on provenance, integrity signals and accountability, to reduce misunderstanding across states with different regulatory models.
- Treat cybersecurity and metadata sensitivity as core design considerations, particularly where integrity proofs or record exchange mechanisms are used.

Preserve accountability in an automated future

Automation can improve efficiency but could also blur responsibility if decisions become opaque or delegated without clear accountability. Adopted tools will need to maintain clear chains of responsibility, especially where automated outputs shape official submissions, compliance decisions or enforcement priorities.

- Maintain human responsibility for final determinations in compliance and reporting processes, including explicit criteria for when human review is required.
- Require traceable decision logs and documentation of how automated outputs were used in forming judgements or submissions.
- Encourage routine internal challenge processes, such as structured review, and, where appropriate, adversarial testing for high-impact automated outputs.
- For transparency mechanisms, ensure that the use of automation does not distance national points of contact from the underlying information they submit.

About the authors

Dr Miranda Smith is a Researcher in the SIPRI Weapons of Mass Destruction Programme, focusing on the governance of biological and chemical weapons. Her work spans biorisk management, dual-use research and global frameworks to reduce risks from emerging biotechnologies.

Kolja Brockmann is a Senior Researcher (non-resident) contributing to the work of the SIPRI Dual-Use and Arms Trade Control Programme. He conducts research in the fields of export control, nuclear, biological and missile non-proliferation, sanctions and technology governance.

Dr Mark Bromley is the Director of the SIPRI Dual-Use and Arms Trade Control Programme. His research focuses on national, regional and international efforts to regulate the trade in conventional arms and dual-use items.



**STOCKHOLM INTERNATIONAL
PEACE RESEARCH INSTITUTE**

Signalistgatan 9
SE-169 72 Solna, Sweden
Telephone: +46 8 655 97 00
Email: sipri@sipri.org
Internet: www.sipri.org