PRELIMINARY DISCUSSION DRAFT

9/10/18

Using International Export Controls to Bolster Cyber Defenses

Herb Lin

Joel P. Trachtman

Abstract

Export controls have traditionally been used to restrict the international availability of technology, information, and software for national security, economic and/or foreign policy reasons. In 2013, parties to the Wassenaar Arrangement agreed to export control provisions related to intrusion software, which cyber attackers could use to the detriment of national security and the economic welfare of a nation. In this article, we (i) describe the history of the problem, (ii) explain some of the difficulties involved in defining the scope of controlled software, including the need to allow transfers for legitimate and defensive reasons, (iii) describe the existing regime for export control relating to intrusion software under the Wassenaar Arrangement and in the U.S. and E.U., (iv) develop a proposal for a “verified end-user” regime that would involve agreements among private sector entities, as well as agreements of host governments to respect the non-disclosure obligations of those entities, and (v) describe some of the possible institutional features of such a regime.

1. Introduction

In the past decade, cybersecurity has become an increasingly important problem of public policy for nations around the world as the frequency and sophistication of cyberattacks has skyrocketed. Criminals are behind many of these cyberattacks, but national governments are culpable as well. For example, targeted attacks on entities such as Sony and the Democratic National Committee, and broader attacks under names like Wannacry and Petya, have been carried out by foreign governments or their agents. These governments utilize technical means to achieve access without consent to the computer networks of these civilian entities, to extract information, to modify information, or to make information unavailable. Private entities such as Hacking Team have assisted repressive governments in hacking civil society organizations within and outside the territories of those governments.

Some nations have sought to use export control mechanisms to reduce the availability to governments of the instruments needed to conduct such attacks. In 2013, members of the Wassenaar Arrangement ("WA"), a significant group of 42 (then 41) mostly western states, adopted a definition of "intrusion software." The definition of intrusion software currently used by the WA is as follows:[1]

> "Software" specially designed or modified to avoid detection by 'monitoring tools', or to defeat 'protective countermeasures', of a computer or network-capable device, and performing any of the following:
>
> a. The extraction of data or information, from a computer or network-capable device, or the modification of system or user data; or
>
> b. The modification of the standard execution path of a program or process in order to allow the execution of externally provided instructions.

Note that the WA definition does not call for controls on intrusion software per se (if it did, such software would be less available for "white hat" uses that support cybersecurity efforts); rather, it controls "'software' specially designed or modified for the generation, command and control or delivery of 'intrusion software.'"[2] For convenience, we refer to this WA-controlled software as "intrusion-related software."

The WA definitions of intrusion software and intrusion-related software are subject to important criticisms, which we address below. Nevertheless, they are a starting point for exploring the relevance of export controls to reducing the availability of such software to parties that might use it for destructive purposes.

There are a number of important questions about whether a system of controls on intrusion software would work and whether it would be excessively costly in relation to its benefits. First, would intrusion software remain available to the likely attacking governments

[1] Wassenaar Arrangement on Export Controls for Conventional Arms and Dual Use Goods and Technologies, List of Dual Use Goods and Technologies and Munitions List, December 2017, available at https://www.wassenaar.org/control-lists/.

[2] Id.

despite controls? Second, can export controls be designed to reduce access by likely attacking governments without being excessively costly? Third, can export controls be designed to reduce access by likely attacking governments without excessively compromising the ability to protect against and respond to attacks? Fourth, can intrusion technology be kept from collaborators or partners of governments that may wish to conduct attack?

We start by noting that physical equipment is significantly easier to observe and control. Export controls were first devised for physical weapons, for which geographic destination was congruent with the relevant policy goal of keeping weapons out of the hands of the governments that control those destinations. Because software moves across borders at essentially zero cost, and because software development, management, and use occur throughout an effectively borderless world, it is challenging to apply territorial border-based export controls. Furthermore, current export control law in the U.S. and elsewhere also includes restrictions not just on the transfer of certain commodities or technologies to particular nations, but also to particular persons. For example, software and technical data disclosed to a foreign national, even if present within the U.S., can be a "deemed export."[3]

Second, greater concerns have been expressed in the past several years regarding the use of intrusion software by governments to attack and to undermine the human rights of their own citizens.[4] Export controls regarding intrusion software have largely focused on these types of human rights violations, with the implicit assumption that human rights obligations are owed by governments to their own citizens, and not "extraterritorially" (or more precisely "extra-nationality") to citizens of other states. This focuses on restricting the availability of intrusion software to foreign state attackers rather than human rights violations by governments against their own citizens. Still, as a practical matter, a nation with access to intrusion software can use it against its own citizens or against other nations, and the core question addressed in export controls is which states may be permitted access to these technologies.[5]

A third key point is that in seeking to bolster cyber defenses, the availability of intrusion software plays two key roles. First, the availability of intrusion software to malicious parties increases the likelihood that they can use it for nefarious purposes. Second, the availability of intrusion software to legitimate security researchers and multinational firms increases their cyber defensive capabilities. To the extent possible, any use of export controls to restrict the availability of intrusion software should support legitimate uses while restricting access for malicious uses.

[3] 15 C.F.R. 734.13(b)

[4] The Commission sought to add human rights to the coverage of these export controls in 2017. http://www.osborneclarke.com/insights/cyber-surveillance-technology-and-export-control-changes-on-the-horizon-part-1/.

[5] Note also that intrusion software is used by national law enforcement agencies, and by intelligence agencies, for legitimate purposes. For example, the intrusion may be carried out pursuant to a procedurally satisfactory search warrant. So not all use of intrusion software against citizens is illegitimate, and export controls must distinguish between exports to governments that will use intrusion software only for legitimate purposes and those that will not.

Fourth, we observe that cybersecurity is a growing field of economic activity, and demand for cybersecurity products and services is steadily on the rise. Moreover, cybersecurity is a field that is knowledge-intensive rather than capital-intensive, a fact that puts the development of a cybersecurity industry within reach of many capital-poor nations. Under the current export control regime, the transfer of intrusion software to such nations would be inhibited, putting into place a significant hurdle for a nascent industry to overcome. Again, a tailored export control regime might be able to minimize such a hurdle.

Given the instantaneous, costless, and often difficult to observe nature of disclosures of software, as well as the need to be able to transfer software across borders for legitimate, or even protective, purposes, whether within firms, between firms, or from firms to governments, we believe that an approach to export controls on intrusion software emphasizing the bona fides of the transferee, rather than the host nation of the transferee, may be of some value. We elaborate below on the reasons for this conclusion.

2. Structure of the Analysis

In order to evaluate the need for revised international law and international organization to address the issue of export controls for intrusion software, we will structure our analysis as follows.

We first address the core technical, policy, and legal problem of defining the technology to be controlled, and providing exceptions from control. It is difficult to define the software to be controlled without restricting important development and security activities.

We then describe the existing regime for controlling intrusion software. This regime includes both national governments and an international organization, the Wassenaar Arrangement. This regime has been in a state of flux, and raises important concerns as it stands. We also will describe the U.S. and E.U. rules, which follow, to some extent, the Wassenaar rules.

Third, we examine some existing regimes that focus on the bona fide nature of the recipient, or end-user, rather than the geographic destination per se. We explore the possibility of developing a regime for combining agreement on controls with agreement on verified end-users ("VEUs"). Since the goal, as expressed above, is to keep certain tools out of state hands, transfer to these VEUs must be conditioned on agreement by the states that have legal or physical authority over the VEU to respect the VEU's responsibility to hold the controlled technology in confidence, without governmental interference. Once we develop these substantive rules, we focus on the structure of an international legal regime, including possibly a revised Wassenaar regime, to administer, modify, and enforce the rules.

3. Problems in Definition of Controlled Software

We refer to the 2013 WA definition of "intrusion software" set forth above. As noted above, under the WA, intrusion software itself is not controlled. Rather, "intrusion-related software," as defined above to include software "specially designed or modified for the generation, command and control or delivery of 'intrusion software'" is controlled. This structure was intended to avoid excessive constraint on innocent research into flaws, prevention, and remediation. It was proposed by the government of the United Kingdom to address human rights and security concerns.[6]

A useful table of the types of software tools that could be included within an “intrusion software” category is contained in Goodwin, Griffin, Peltier & Walton (2016). These different types of tools raise a variety of issues regarding definitions, rationales for export, and appropriate exceptions. However, there are a number of difficult definitional issues, as well as conceptual issues. On the definitional side, there are questions about the meaning of the terms “monitoring tools” and “protective countermeasures.”

More intractably, certain system administrative tools or penetration testing software designed to discover vulnerabilities might fall within this definition. It turns out that the good guys use the same tools as the bad guys. And the good guys often involve multinational firms with offices in multiple states, as well as foreign nationals, working with handoffs around the clock. So, innocent intracorporate sharing, not to mention innocent sharing with clients or colleagues, will require export licenses that may take weeks to obtain. This may delay response to attacks. One expert concludes: “The Wassenaar Arrangement as written would have required export control licenses for nearly anyone involved in defensive security activities involving an export of, for example, command and control software and technology shared in taking down a botnet attack in real time.”[7]

> Definitions of intrusion software may focus on capabilities of components, or may focus on more complete “intrusion delivery platforms”, in which overall design reflects an intent to intrude, not just a component of capacity to intrude. Similarly, export controls for dual use goods would not ordinarily restrict a truck tire that could be used in a mobile missile launcher. However, if components may be assembled with sufficient ease, then a focus on platforms may be insufficient.
>
> Intent-based definitions can be underinclusive, or difficult to apply, and capability-based definitions that focus on more complete intrusion delivery platforms may also be underinclusive. Thus, the question in determining whether it is sufficient to focus on intrusion delivery platforms will turn on the extent to which the particular component is (i) critical to the intrusion delivery platform, and (ii) not so readily available or substitutable to make restriction futile. Intent-based definitions that infer intent from design may be under-inclusive.

However, both intent-based definitions and capability-based definitions can be overbroad, as well as underinclusive.[8] Bratus, Locasto and Shubina (2014) argue that these

[6] Wassenaar Arrangement, ‘Public Statement 2013. Plenary meeting of the Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods and Technologies’, 4 Dec. 2013.

[7] Katie Moussouris, Serious Progress Made on the Wassenaar Arrangement for Global Cybersecurity, The Hill, December 17, 2017.

[8] See Sergey Bratus, Michael Locasto, and Anna Shubina, Why Wassenaar Arrangement’s Definitions of “Intrusion Software”and “Controlled Items” Put Security Research and Defense At Risk, July 23, 2014.

categories of controlled software restrict “the primary known means through which research and engineering progress has been made in all known aspects of software, including security.” These means are “automation of generation and operation of software elements.” Often legitimate program features will need to be designed to “defeat protective countermeasures,” so that these legitimate features will be caught up in the definition of intrusion software. Similarly, the control of software that automatically generates vulnerabilities of the kind utilized by intrusion software, or “exploits,” would ordinarily be included in software verification programs.

So, according to these commentators, the definition of controlled items in this field remains unacceptably, and hopelessly, overbroad. The conclusion: “people who defend and protect computer networks need access to the exact same tools and information that attackers use.”[9] The problem: those people that defend are not concentrated in a single country.

Some areas of overbreadth may be addressed through specified exceptions. What are appropriate reasons for exceptions from control? One group of researchers has proposed that exceptions from licensing requirements be considered for certain categories of software, including the following:

i. Non-commercial use under a monetary threshold
ii. Hypervisors, debuggers or software reverse engineering tools
iii. Sweep patch validation or assessment tools
iv. Digital rights management
v. Software designed for asset tracking, incident response, and recovery
vi. Software for adding to, testing, or modifying the security or functionality of systems, in some cases with informed consent, for securing data, systems and networks, and for creating reports and analysis for customers to remediate security issues, and
vii. Network capable devices.[10]

These categories need to be evaluated in order to determine the extent to which they may allow transfers of dangerous capabilities, and their definitions must be modified, in order to exclude such transfers. The purpose of this article is not to provide specific definitions and exceptions, but to suggest the parameters that might be considered in constructing definitions and exceptions.

In 2015, after the U.S. Bureau of Industry and Security published proposed rules for implementing the 2013 revised WA controls,[11] industry groups and civil society groups objected to the apparent unintended limits on cross-border vulnerability research. The cybersecurity industry community argued that these export controls were overbroad, and would stifle both research leading to improved security, and coordination of response to attack. In

[9] Tom Cross, New Changes to Wassenaar Arrangement Export Controls will Benefit Cybersecurity, Forbes, January 16, 2018.

[10] Goodwin, Griffin, Peltier & Walton (2016).

[11] U.S. Department of Commerce, Bureau of Industry and Security, Proposed Rule, Wassenaar Arrangement 2013 Plenary Agreements Implementation: Intrusion and Surveillance Items, 80 FR 28853 (May 20, 2015).

addition, while the WA rules excluded “zero day flaws”—actual vulnerabilities in systems—the U.S. proposal included them.[12] Restrictions on transmission of zero day flaws will impede security coordination, as well as “bug bounties” by which legitimate companies pay bounties for disclosure of zero day flaws in their systems.

In response to these objections, at the December 2017 WA meeting, the U.S. sought exceptions to export controls on intrusion software for use in research. These modifications were designed to clarify that “technology exchanged for vulnerability disclosure or cyber incident response purposes are not controlled, and updates or upgrades are not controlled,” so long as they themselves are not intrusion software.[13] These modifications address important elements of the concerns expressed by the software community about allowing legitimate defensive operations, both before and after an incident.

Legal categories are constructed to serve a substantive purpose. In the current context, the substantive purpose is to deprive specified governments of tools that can be used to intrude without consent on civilian computer networks. The problem is that this purpose cannot readily be translated into an elegant category of software to be controlled, but will require continued adjustment and updating pursuant to a global partnership between the defensive industry and the governments establishing the export controls. The WA will need to be revised to provide a satisfactory structure for this partnership.

4. The Existing Regimes

Assuming that there are alternative sources of intrusion technology, then export controls present a cooperation problem. Cooperation is implicit in the idea of export controls: government intervenes to cause producers to restrict their sales in order to promote the common good. National government solves the cooperation problem at the national level; international law or informal regimes solve cooperation problems at the international level.[14]

The idea of an export control assumes governmental restriction of private activity: governments prohibit export of restricted technology without a license. Private persons could, in theory, develop their own set of export controls, and if those controls could operate effectively, they might pre-empt the need for governmentally-imposed controls. Controls could evolve based on an industry standard, similar to the “traffic light protocol” that has developed to secure confidential information, or similar to ISO 27001 or 27032.[15]

[12] Mailyn Fidler, Proposed U.S. Export Controls: Implications for Zero-Day Vulnerabilities and Exploits, Lawfare, June 10, 2015.

[13] BIS, FAQs. See also Cross, supra note ***Error! Bookmark not defined.***; Shaun Waterman, The Wassenaar Arrangement’s Latest Language is Making Security Researchers Very Happy, CyberScoop, December 20, 2017.

[14] See, e.g., Joel P. Trachtman, The Future of International Law: Global Government (Cambridge 2013).

[15] If we understand network security as a human right, the supply of intrusion software to government violators may be understood as aiding or abetting a human rights violation. The

However, even if the software industry as a whole would benefit from export controls on intrusion software, individual software companies would have incentives to defect from an industry-agreed rule, and it does not appear that there are sufficient market-based incentives to comply. While a contractual agreement among capable suppliers could reduce problematic transfers of software, it may be difficult to induce a sufficient number of capable suppliers to participate, it may be difficult to enforce, and it may raise competition law issues in some jurisdictions. At the national level, this is a public goods cooperation problem, and a main role of national government is to solve it.

Of course, different governments will have different incentives in relation to export controls on intrusion software. Governments of states that are incapable of producing effective intrusion software would generally oppose controls, and governments of states less likely to be harmed by intrusion software, or more likely to benefit from transfers of intrusion software, would have less incentive to accept and comply with controls. It should be noted, and highlighted, that export controls may not be able to address indigenous intrusion software capabilities in countries like China, Russia, Iran, and North Korea. However, export controls may reduce these capabilities, and would have greater effectiveness on countries with less robust cyber capabilities.

Assuming that it would be useful to reduce trade in intrusion software, the international "intrusion software export control game" is a weakest-link public goods game,[16] in which unless all of the potential producers of intrusion software comply, the public good of restriction is not produced.

The Wassenaar Arrangement is designed to partially address this cooperation problem. We describe it, and its limitations, below, and then briefly describe the U.S. and E.U. export control regimes for intrusion software. In this Part, we do not address the difficult definitional problems in designing an export control regime; that is reserved for Part 4.

a. Wassenaar Arrangement

Under the Wassenaar Arrangement ("WA"), participating states have agreed, but do not accept formal treaty-based international legal obligations, to maintain, through national rules, export controls on items included in the WA control lists. Each participating state retains formal discretion to restrict exports or to allow them. To be clear, the WA is a coordinated list of items that its members plan to, and on a non-binding legal basis agree to, subject to national export controls. In addition, members are required, on a non-binding legal basis, to report transfers or denials of transfers of certain controlled dual-use items.[17]

U.N. Guiding Principles on Business and Human Rights suggests responsibilities of companies to avoid aiding or abetting human rights violations.

[16] See Hirshleifer, J., 1983, From weakest-link to best-shot: The voluntary provision of public goods, Public Choice, 41, 371-386.

[17] https://www.wassenaar.org/the-wassenaar-arrangement/

The WA is not a treaty, but instead operates as a “soft law” commitment among its member states.[18] In that way, it is similar to the Basle Committee bank capital accords, the Codex Alimentarius, and other non-legal rules. The fact that these rules do not impose formal legal requirements under international law does not mean that they do not have effects on state behavior. But the effects operate at the political or informal level in international relations.

Because the WA is not binding international law, and because it does not have a formal process of interpretation and dispute settlement, different states apply it with different scopes and degrees of effectiveness. The predecessor regime of the WA, COCOM, included a rule that exports of certain sensitive items by any member state would require prior notification to the other members, and were subject to veto by any member. The U.S. regime is one of the strictest of the WA export control regimes. However, for example, the U.S. has not yet, as of July 2018, adopted the WA restriction of intrusion-related software.

The WA includes 42 states,[19] including all of the member states of the E.U. The WA notably includes Russia, but, also notably, excludes a number of states with strong software capabilities, including Brazil, China, Iran, Israel, North Korea, Pakistan, South Korea, and Taiwan. There are also concerns that Russia does not apply the Wassenaar rules as reliably as western states do. The WA also encourages voluntary adherence to its standards by non-member states.

The WA control lists were first established in 1996 and have been revised annually thereafter, by negotiation among the members. Decisions regarding what to include on the control list are made by consensus.

While much of the world’s intrusion software capability is covered by the WA member states, and other states that adhere to its standards, certain non-members would have the capability to produce effective intrusion software. Indeed, the more successful a “cartel” arrangement such as the WA is, the greater the incentives to defect, or simply to avoid accepting its obligations. So, in order to have an effective regime, movement towards universal membership, or at least universal compliance, will be attractive. And as the regime grows more effective, it will need greater inducements to comply and to remain part of the regime, in order to overcome the increasing attractions of defection.

The WA does not directly include private sector participation, and there seems to be wide agreement that the U.S. administration failed adequately to obtain private sector input before

[18] The constitutive document for the Wassenaar Arrangement is the Guidelines and Procedures, including the Initial Elements, the latest version of which is December 2016 (the “Initial Elements”).

[19] As of April 14, 2018, the 42 members of the Wassenaar Arrangement were: Argentina, Australia, Austria, Belgium, Bulgaria, Canada, Croatia, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, India, Ireland, Italy, Japan, Latvia, Lithuania, Luxembourg, Malta, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, Republic of Korea, Romania, Russian Federation, Slovakia, Slovenia, South Africa, Spain, Sweden, Switzerland, Turkey, Ukraine, United Kingdom and United States.

accepting the 2013 Wassenaar intrusion software provisions. We discuss this issue in Part 4 below.

b. U.S.

The U.S. regime for export controls is administered by the Bureau of Industry and Security ("BIS") of the Department of Commerce, under the Export Administration Regulations, including the Commerce Control List of Dual-Use Items ("CCL"). Items on the CCL are assigned an "export control classification number" or "ECCN". These regulations have the force of law.

The CCL has not yet been revised to reflect the 2013 WA agreement or subsequent WA revisions. As noted above, in 2015, the BIS proposed to incorporate the 2013 WA agreement by adding the relevant references to the CCL. These categories were proposed to be controlled for national security (NS), regional stability (RS), and anti-terrorism (AT) reasons to all destinations, except Canada. The proposal, while requiring licenses for all other destinations, proposed "favorable review" if "destined to a U.S. company or subsidiary not located in Country Group D:1 or E:1 [specified countries less trusted or subject to embargo], foreign commercial partners located in Country Group A:5 [countries more trusted], or government end users in Australia, Canada, New Zealand or the United Kingdom [with the U.S., the "Five Eyes"] . . . ." "Note that there is a policy of presumptive denial for items that have or support rootkit or zero-day exploit capabilities."[20]

Because of private sector criticism in response to the 2015 implementation proposal, as discussed in Part 4, the U.S. has not yet implemented WA intrusion related software restrictions, nor has it implemented more recent WA exceptions for "vulnerability disclosure" and "cyber incident response."

c. E.U.

The E.U. export control regime governs export controls for all E.U. member states, pursuant to Regulation 428/2009, amended by Regulation 2016/1969.[21] Member states are permitted to impose more stringent restrictions, and Germany has done so with respect to intrusion software. The existing E.U. controls generally track the WA definitions.

The E.U. controls are proposed to be revised.[22] [check status] "The draft regulation introduces the new concept of 'human security' to export controls, to prevent the human rights

[20] Wassenaar Arrangement 2013 Plenary Agreements Implementation: Intrusion and Surveillance Items, 80 FR 28853, May 20, 2015.

[21] Commission Delegated Regulation (EU) 2016/1969 of 12 September 2016, amending Council Regulation (EC) No 428/2009 setting up a Community regime for the control of exports, transfer, brokering and transit of dual-use items.

[22] Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL setting up a Union regime for the control of exports, transfer, brokering, technical assistance and transit of dual-use items (recast), COM/2016/0616 final - 2016/0295 (COD), September 28, 2016 ("E.U. 2016 Proposal").

violations associated with certain cyber-surveillance technologies."[23] "The proposal sets out a two-fold approach, combining detailed controls of a few specific listed items with a 'targeted catch-all clause' to act as an 'emergency brake' in cases where there is evidence of a risk of misuse."[24] "The targeted catch-all control applies where there is evidence that the items may be misused by the proposed end-user for directing or implementing serious violations of human rights or international humanitarian law in situations of armed conflict or internal repression in the country of final destination."[25] This approach attempts to mitigate some of the potential underinclusiveness in the definition of intrusion-related software subject to control.

d. Limitations of the Existing Regime

The existing regime has a number of limitations that make it unlikely to be effective in avoiding transfers of intrusion software capabilities. First, it only covers a limited number of countries. Second, the countries that are covered have different interpretations of the controls, and different levels of enforcement rigor. Third, the definitions of controlled software are, as discussed below, overbroad and under-inclusive in important respects. Fourth, the institutional structure does not provide for effective international enforcement.

5. Toward Exclusions of Verified End-Users

Of course, if all transferees could be trusted to refrain from malicious use of intrusion software, there would be no need for export controls at all. The definition of controlled software thus interacts with the definition of permitted transferee. That is, with a broad group of permitted transferees—a broad group that has been vetted for reliability—it is less of a problem to have a broad definition of controlled software.

a. Intra-Company Transfers and Transfers to Private Sector End-Users

Under U.S. export controls, license exception ENC[26] may authorize export without a license to any country (except certain countries designated as terrorism-supporting or embargoed countries), if the item is being exported either (i) to a subsidiary of a U.S. company, including to foreign nationals who are employees, contractors or interns of a U.S. company or its subsidiaries, for internal company use; or (ii) to private sector end users, headquartered in what is defined as a "Favorable Treatment Country" (NATO countries and certain other closely allied countries)[27] for internal development or production of new products. This exception serves as an example of an end-user based exception to export controls.

[23] European Parliament, Review of Dual Use Export Controls, January 12, 2018.

[24] E.U. 2016 Proposal, supra note 22, p. 5.

[25] Id., at 9. See Fabian Bohnenberger, The Proliferation of Cyber-Surveillance Technologies: Challenges and Prospects for Strengthened Export Controls, 3:4 Strategic Trade Review 81 (2017).

[26] For "encryption." 15 CFR 740.17.

[27] Supplement No. 3 to Part 740 CFR.

Of course, private sector end-users may not always refrain from disclosing the relevant software to governments or other prohibited users, and exporters have some obligations, at least under U.S. law, to “know your customer.” The U.S. export control regulations provide a set of “red flags” that are intended to make exporters aware of when they are “on notice” that a purported end-user cannot be trusted.[28] Red flags are defined as “any abnormal circumstances in a transaction that indicate that the export may be destined for an inappropriate end-use, end-user, or destination.”[29]

In addition, the U.S. BIS maintains a “validated end user” (“VEU”) facility, under which a transferee in India or China may be approved in advance.[30] It is worth quoting in full the regulatory standard for approval of a VEU:

> In evaluating an end user for eligibility under authorization VEU, the ERC will consider a range of information, including such factors as: the entity's record of exclusive engagement in appropriate end-use activities; the entity's compliance with U.S. export controls; the need for an on-site review prior to approval; the entity's capability of complying with the requirements of authorization VEU; the entity's agreement to on-site reviews by representatives of the U.S. Government to ensure adherence to the conditions of the VEU authorization; and the entity's relationships with U.S. and foreign companies. In addition, when evaluating the eligibility of an end user, the ERC will consider the status of export controls and the support and adherence to multilateral export control regimes of the government of the eligible destination.[31]

While this program has been subject to criticism,[32] and has not expanded beyond India and China, it is a model for expansion of the scope for international regulatory cooperation between exporting states and importing states, in order to facilitate controlled exports. We discuss how a modified program might address some of the concerns about intrusion software below.

As described above, there are several problems with the existing export control regime for intrusion software. One problem is that the specification of the controlled software is overbroad, but this problem would be ameliorated substantially if the universe of license-free transferees, or transferees that could be licensed generally and in advance, could be expanded. The ENC and VEU programs provides some models for such expansion. What if affiliated companies, other companies that cooperate on software development, and even customers,

[28] These “red flags” formally apply only to nuclear, chemical, and biological weapons, and missile technology, but they may be applied on other bases to other areas of control. See BIS “Best Practices” for Industry to Guard Against Unlawful Diversion Through Transshipment Trade, August 31, 2011.

[29] 15 CFR Appendix Supplement No 3 to Part 732.

[30] 15 CFR 748.15.

[31] 15 CFR 748.15.

[32] See Government Accountability Office, Challenges with Commerce’s Validated End-User Program May Limit Its Ability to Ensure That Semiconductor Equipment Exported to China Is Used as Intended, September 2008.

including law enforcement and intelligence agencies, could be approved in advance as transferees of the controlled intrusion software? This could be done on three conditions.

- First, it would be necessary to perform a due diligence investigation of these transferees, including their internal safeguards and end-use of the products, and for the transferees and their employees to contract not to disclose the software in violation of the exporting country's export control laws.
- Second, these transferees would be required to agree to monitoring and auditing of their activities in connection with the transferred intrusion software, by (a) the exporting country government authorities, (b) private sector agents approved by exporting country government authorities, or (c) importing country government authorities approved by exporting country government authorities.
- Third, the importing country government would be required to agree to enforce and not to interfere with the transferee's compliance with obligations not to disclose the software in violation of the exporting country's export control laws.

This expanded VEU regime would be designed to provide appropriate assurances that intrusion software would not be used for purposes of attacks on civilians in other states, and presumably in the transferee state as well. Why would transferee states accept this regime? As discussed below, they would be likely to do so in order to allow greater dissemination into their country of intrusion software products utilized for legitimate purposes, as well as software development expertise. This would bring commercial and developmental benefits.

We might ask, can a non-state entity really resist orders or inducements from host governments or other governments to disclose intrusion software? Much would depend on the ability to induce host governments to make binding commitments to comply with the relevant regime, and to construct a set of punishments for both the disclosing private entity and the government to which the software is disclosed. To the extent that software can be prepared in a way that makes its origin identifiable after an attack, it would be easier to attribute a failure of export controls to a firm and to a government, and to impose punishments. This is a dual attribution problem: the attack needs to be attributed to a government, and the source of the software needs to be attribute*d to* a software firm (the VEU).

The problem of attributing software to the VEU is an issue of software provenance.[33] To be more precise, the problem assumes that the restricted software is made available to the VEU in source code form (which we call G_s (G for oriGinal, s for source)).[34] Another party (named R) improperly obtains G_s from the VEU and incorporates it into another body of source code (called

[33] A useful perspective on software provenance can be found on the Software Engineering Institute Blog at https://insights.sei.cmu.edu/sei_blog/2014/02/provenance-inference-in-software.html.

[34] The distinction between source code and object code is important. Source code is the medium in which humans program computers. Source code is readable by humans. Through a process knowns as compilation, source code is turned into object code, a particular sequence of ones and zeros that are meaningful to the computer and instruct the computer about what to do step-by-step. Of particular importance is the fact that during the compilation process, information that is meaningful to humans but not to computers is lost.

N_s (N for new, s for source). To use N_s in an attack, R compiles the new source code into object code (called N_o (N for new, o for object)), and runs N_o on the targeted computer. Forensic investigators F can usually obtain N_o from the targeted computer; in special cases, N_s may be available to F. The operative question is whether F can determine whether N_o (or in special cases, N_s) is associated with G_s, the software originally provided to the VEU. Note also that the original code in object form G_o is also available to F.

Software can be associated with a given party in two different ways. One set of techniques is based on the identification of characteristic features or aspects of the original software, much as a painting might be associated with a given artist because of a similarity between the pattern of brushstrokes used on that painting and other paintings known to be done by that artist.

For example, software bertillonage is an approach that uses the presence of various software features to reduce the effort of trying to locate a software object (in either source or object form) within a large corpus of possibilities.[35] Once the entity is determined with high probability to be amongst a more limited set of known software objects, other techniques can be used to make a more precise identification.

Another approach based on code stylometry[36] analyzes the style with which source code has been written and seeks to associate a software entity of unknown provenance with another specific entity of known provenance.[37] Recent research has also suggested that, using machine learning techniques, it is more feasible than previously believed to use stylometric techniques on object code.[38] The significance of the latter research is that in the wake of an actual attack, object code may be recoverable while source code will be unavailable, barring very unusual circumstances.

A second way of associating software with a given party is to introduce into the original code G_s certain features (here called watermarks) that would be preserved in any re-use of that

[35] Julius Davies, Daniel M. German, Michael W. Godfrey, and Abram Hindle, “Software Bertillonage: Finding the Provenance of an Entity”, 8th Working Conference on Mining Software Repositories, May 21-22, 2011, Waikiki, Honolulu, HI, USA, http://softwareprocess.es/pubs/davies2011MSR-bertillonage.pdf

[36] Stylometry is the generic name given to techniques that have been used to identify previously unknown works of Shakespeare—these techniques examine the style of an unknown work and determine that it is highly similar in style to those of known works of Shakespeare.

[37] Aylin Caliskan-Islam et al, “De-anonymizing Programmers via Code Stylometry”, Proceedings of the 24th USENIX Security Symposium, Washington, D.C., August 12–14, 2015, https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-caliskan-islam.pdf.

[38] Aylin Caliskan et al, “When Coding Style Survives Compilation: De-anonymizing Programmers from Executable Binaries,” Network and Distributed Systems Security (NDSS) Symposium 2018, 18-21 February 2018, San Diego, CA, USA, http://dx.doi.org/10.14722/ndss.2018.23304.

software.[39] If R incorporates G_s into its own software N_s, examination of R's software should reveal the watermark, thus indicating the true origin of that code, and thus a violation of the VEU agreement could be identified. On the other hand, R's illicit use of code obtained from the VEU would almost surely be accompanied by an effort by R to remove the watermark from the body of code in question, and if R knew about the watermark, it would be able to do so (easily if R knew the details of the watermark, and with more difficulty if R only knew of its presence). If only G_o –the object version of the original software— is made available to the VEU, R is likely to have a harder time removing the watermark; this would be a very good reason for the original creator of G only providing object code to the VEU.

Still, the watermarking is a cat-and-mouse game. Watermarkers constantly strive for better watermarks—those that will resist attempts at program transformation and other removal techniques—while anti-watermarkers will strive for better ways to remove watermarks.[40]

6. Need for Broad Agreement: The Carrot-Stick

If a regime were designed around the enhanced VEU approach described above, it would have to address a number of issues. First, should it remain, like the WA, in the form of soft rules, or be converted to a legally binding international treaty? Second how would this regime incorporate the views of the private sector, in order to avoid the kinds of errors made in the 2013 WA revisions? Third, how would intrusion software-competent states be induced to adhere and comply, and how would other states be induced to adhere and comply? Finally, what organizational features, in terms of decision-making, adjudication, and executive functions, including research, surveillance, and enforcement, should a revised organization have?

a. Toward Export Control Rules for Intrusion Software

A cooperation regime can utilize formal law or informal rules. Informal rules may be easier for states to adopt, and may provide attractive flexibility. Formal law has the advantage that states may take it more seriously, and it can more readily be subjected to adjudication in order to definitively interpret the definitions, exceptions, and thus the obligations. It can therefore be a basis for greater trust. Perhaps counterintuitively, some states may be willing to enter into law of this nature because they can rely more on the performance of other states. Whether a regime is composed of formal or informal rules, it would require some of the same basic features.

b. Role of Private Sector

[39] Mila Dalla Preda and Michele Pasqua, "Software Watermarking: A Semantics-based Approach", Electronic Notes in Theoretical Computer Science 331:71-85, 2017, https://www.sciencedirect.com/science/article/pii/S1571066117300075.

[40] For a discussion of this struggle (but with the watermarkers winning at this time), see Zhe Chen, Zhi Wang, and Chunfu Jia, "Semantic-integrated software watermarking with tamper-proofing", Multimedia Tools and Applications, 77:11159–11178, 2018, https://link.springer.com/content/pdf/10.1007%2Fs11042-017-5373-7.pdf.

Any new regime must provide appropriate transparency, notice and comment, and probably a formal role for private sector representatives. This will be important in crafting and in interpreting commitments and exceptions in a way that will not have unintended or excessive adverse consequences. One model for private sector participation is purely consultative, and this may be sufficient. Another model would provide for formal private sector participation, along the lines of the International Labor Organization's inclusion of employer and employee representatives. [expand; also refer to Codex Alimentarius, Basle, and others]

c. Membership Inducements

Given the weakest link public good nature of the cooperation problem in connection with intrusion software, it would be useful to procure participation by all potential source countries. Indeed, the very idea of territoriality is misleading in this context, because software development may be highly mobile due to its technological character. So, broader participation beyond those countries presently enjoying robust intrusion software capabilities may be appropriate.

One way to achieve broad participation would be to link this cooperation with, or incorporate it in, an existing more or less universal organization, such as the United Nations or the World Trade Organization. Of course, these organizations would have to approve such link or incorporation through a unanimous decision, which may be difficult to achieve. However, log-rolling has allowed these organizations to make effective changes in the past.

There is a natural and elegant punishment for non-participation: refusal to transfer intrusion software to the non-participating state. There is precedent for this in the Basle Convention on Transboundary Movement of Hazardous Waste, which provides in Article 4(5) that "a Party shall not permit hazardous wastes or other wastes to be exported to a non-Party or to be imported from a non-Party."[41] However, some states may find that they would rather maintain freedom to export, while accepting this punishment. So, further evaluation of the effectiveness of this mechanism will be necessary in order to determine whether it will be sufficiently effective.

Another approach would be to punish non-participation through other means, including reputational sanctions, as the Organization for Economic Cooperation and Development (OECD) has done in the context of its Harmful Tax Practices program, inducing most tax haven countries to cease some of their worst tax haven abuse practices.[42] This type of linkage-based punishment is at the core of William Nordhaus' proposal for "climate clubs" using trade sanctions to induce states to join carbon reduction regimes.[43] This structure is not reliant on hard law, or on membership in an international organization, but it has used exposure and international pressure to cause changes in practices. This model, including its surveillance and reporting functions,

[41] Basle Convention on Transboundary Movement of Hazardous Waste and Their Disposal, Mar. 22, 1989, 28 I.L.M. 657, 1673 U.N.T.S. 57, Art. 4(5)

[42] See OECD Global Forum on Transparency, Tax Transparency 2017: Report on Progress (2017).

[43] William Nordhaus, Climate Clubs: Overcoming Free-Riding in International Climate Policy, 105 American Economic Review 1339 (2015). Available at https://www.aeaweb.org/articles?id=10.1257/aer.15000001

may be sufficient to support adherence and compliance to export control obligations in connection with intrusion software.

d. Organizational Functions

In order to provide for periodic revisions to the "control list" and other aspects of the relevant obligations, in order definitively to interpret them, in order to engage in research, negotiation support, monitoring, reporting, dispute settlement, and enforcement functions, it would be useful to have an international organization. These types of functions could be taken over by the WA, although it is not clear that these functions are needed in other areas of export controls. However, these functional additions could be confined to intrusion software export controls. Alternatively, a new U.N. specialized agency, or the WTO, or some other organization, could house these functions.

One important organizational function would involve adjudicating determinations of what types of software is covered, and whether relevant exceptions are available.

7. Conclusion

With respect to software or other data that is intangible, geographic location is not necessarily an accurate proxy for government control: it is both over-broad and under-inclusive. Government control can occur with respect to software that exists on computers outside the physical borders of the state, for example through foreign investment, migration, or subornment, and government control may not exist with respect to software that exists on computers within the physical borders of the state. A VEU regime may help to address this problem by focusing on the rationale for control, rather than the geographic location of the recipient. But the many difficulties in enforcing such a regime may minimize its practical import. In particular, better techniques for ascertaining software provenance (where "better" refers to techniques that are difficult to circumvent) would greatly assist in the establishment of a VEU regime.

[more to be provided]