



CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

TOWARD A NUCLEAR FIREWALL

Bridging the NPT's Three Pillars

MARCH 2017

Toby Dalton, Wyatt Hoffman, Ariel E. Levite,
Li Bin, George Perkovich, and Tong Zhao



CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

TOWARD A NUCLEAR FIREWALL

Bridging the NPT's Three Pillars

Toby Dalton, Wyatt Hoffman, Ariel E. Levite,
Li Bin, George Perkovich, and Tong Zhao

© 2017 Carnegie Endowment for International Peace. All rights reserved.

Carnegie does not take institutional positions on public policy issues; the views represented herein are the authors' own and do not necessarily reflect the views of Carnegie, its staff, or its trustees.

No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Carnegie Endowment. Please direct inquiries to:

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue, NW
Washington, DC 20036
P: +1 202 483 7600
F: +1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost
at CarnegieEndowment.org/pubs.

Contents

About the Authors	v
Acknowledgments	vii
Summary	1
Introduction	5
The Firewall Framework	11
Possible Uses of the Firewall	28
Conclusion	44
Notes	47
Carnegie Endowment for International Peace	48

About the Authors

Toby Dalton is co-director of the Nuclear Policy Program at the Carnegie Endowment for International Peace, where his research focuses on regional security challenges and the evolution of the global nuclear order. He is a co-author with George Perkovich of *Not War, Not Peace?* (Oxford University Press, 2016). Before joining Carnegie, he served in several positions at the U.S. Department of Energy's National Nuclear Security Administration, including as senior policy adviser in the Office of Nonproliferation and International Security and director of the Department of Energy office at the U.S. embassy in Islamabad, Pakistan. As a Luce scholar in 2001–2002, he worked as a visiting fellow at the Institute for Far Eastern Studies in Seoul, South Korea.

Wyatt Hoffman is a nonresident research analyst with the Nuclear Policy Program and the Cyber Policy Initiative at the Carnegie Endowment for International Peace. He is also a Rotary Global Grant scholar in the Department of War Studies at King's College London. His research focuses on active cyber defense in the private sector, norms for cyberwarfare, and the intersection of nuclear weapons and cybersecurity. He previously worked as a project manager and junior fellow with the Nuclear Policy Program at Carnegie and as a research intern at the Arms Control Association.

Ariel E. Levite is a nonresident senior fellow with the Nuclear Policy Program and the Cyber Policy Initiative at the Carnegie Endowment for International Peace. Prior to joining Carnegie in 2008, Levite was the principal deputy director general for policy at the Israeli Atomic Energy Commission from 2002 to 2007. He earlier served as the deputy national security adviser for defense policy and was head of the Bureau of International Security and Arms Control (an assistant secretary position) in the Israeli Ministry of Defense.

Li Bin is a senior fellow with the Nuclear Policy Program and the Asia Program at the Carnegie Endowment for International Peace. A physicist and expert on nuclear disarmament, he conducts research on China's nuclear and arms control policies and on U.S.-China nuclear relations. Li is also a professor of international relations at Tsinghua University. He previously directed the Arms Control Division at the Institute of Applied Physics and Computational

Mathematics, where he served as executive director of the Program for Science and National Security Studies. Li was a Social Science Research Council–MacArthur Foundation Peace and Security fellow at the Massachusetts Institute of Technology and Princeton University and a member of China’s delegation to the Comprehensive Nuclear-Test-Ban Treaty negotiations. He is the author of *Arms Control Theories and Analysis* and co-editor of *Strategy and Security: A Technical View* and *Understanding Chinese Nuclear Thinking*.

George Perkovich is vice president for studies at the Carnegie Endowment for International Peace. He works primarily on nuclear strategy and nonproliferation issues. On Iran and South Asian security affairs, he has advised many agencies of the U.S. government and testified before both houses of Congress. He has been a member of the National Academy of Sciences’ Committee on International Security and Arms Control as well as the Council on Foreign Relations’ Independent Task Force on U.S. Nuclear Weapons Policy, and he was a principal adviser to the International Commission on Nuclear Nonproliferation and Disarmament, a joint initiative of the governments of Japan and Australia. Perkovich is the author of the prizewinning book, *India’s Nuclear Bomb*; co-author with James Acton of a 2008 Adelphi Paper, *Abolishing Nuclear Weapons*; and co-author with Toby Dalton of *Not War, Not Peace?* (Oxford University Press, 2016).

Tong Zhao is a fellow with the Nuclear Policy Program at the Carnegie Endowment for International Peace, based at the Carnegie–Tsinghua Center for Global Policy in Beijing. His research focuses on strategic security issues, including nuclear arms control, nonproliferation, missile defense, space security, strategic stability, and China’s security and foreign policy. Before joining Carnegie, he was a Stanton Nuclear Security fellow at the Belfer Center at Harvard University. He is a member of the Asia-Pacific Leadership Network for Nuclear Non-Proliferation and Disarmament.

Acknowledgments

The Nuclear Policy Program of the Carnegie Endowment for International Peace wishes to acknowledge the contributions of a number of individuals whose insights and efforts are reflected in this paper. Over the course of six years, the authors consulted dozens of experts and officials from around the world to solicit their views on nuclear energy and nuclear weapons technologies and policies. We are indebted to them for their ideas, comments, and criticisms, which we have sought to capture as faithfully as possible, albeit without attribution. Brian Radzinsky, while a project manager at Carnegie, helped immensely in launching this effort. Chelsea Green, a Herbert Scoville Jr. Peace Fellow, provided excellent research and editing assistance. The Carnegie communications team, in particular senior editor Lori Merritt, helped us improve and clarify the text. And the project would not have been possible without generous support from the MacArthur Foundation, the Hewlett Foundation, the Edgerton Foundation, the Sloan Foundation, the Prospect Hill Foundation, and the Skoll Global Threats Fund. The authors alone bear responsibility for the paper's contents.

Summary

There is no clear, internationally accepted definition of what activities or technologies constitute a nuclear weapons program. This lack of definition encumbers nuclear energy cooperation and complicates peaceful resolution of proliferation disputes. A “nuclear firewall” could enhance the distinction between nuclear weapons–related activities and other non-weapons uses of nuclear technology. Applying a firewall framework for analyzing nuclear programs could improve international governance of nuclear technology and facilitate peaceful nuclear cooperation and disarmament. It could also expand the time and means available to key states and international bodies, such as the International Atomic Energy Agency and United Nations Security Council, to diplomatically resolve impending proliferation crises.

Defining Nuclear Weapons

The Nuclear Non-Proliferation Treaty (NPT), which establishes the norms and rules that guide the international management of nuclear technology, does not define the term nuclear weapon. Nor does it identify the evidence that would determine whether a state is seeking to manufacture or acquire nuclear weapons.

Such definitional and analytic ambiguity exacerbates the task of distinguishing whether components, equipment, nuclear materials, and facilities are related to nuclear weapons programs or, instead, are for purely peaceful applications of nuclear technology. It also complicates national and international deliberations over the legitimate boundaries for peaceful civil nuclear applications, as well as the handling of proliferation risks and responses.

The Carnegie Endowment for International Peace has developed an analytic approach—a nuclear firewall—to help distinguish between activities and programs that are purely peaceful and those that merit definition as nuclear weapons–related. Like an information-system firewall, the nuclear firewall would

- identify those activities, materials, and equipment that should be inhibited because they are purely or strongly associated with nuclear weapons programs;
- distinguish activities that should be facilitated because they are fully consistent with peaceful applications of nuclear technology and know-how; and
- assess in-between activities depending on transparency and reassurance measures that states would undertake.

The Firewall Framework

To develop the firewall, the Carnegie team worked with leading technical and policy experts from nuclear weapon and non-nuclear-weapon states to identify pertinent indicators and contextual factors that demarcate peaceful from nuclear weapons activities. Tested against historical case studies, the endeavor yielded a framework that can withstand challenges posed by the lack of complete and accurate information about a country's activities and the possibility of deliberate efforts at deception.

The process predictably revealed that (1) technical parameters alone are rarely sufficient to draw reliable assessments about the overall orientation of a country's nuclear program; and (2) in some cases, the *absence* of certain activities, items, or policies can better indicate whether a program's stated purpose is indeed peaceful.

The resulting multidimensional framework—designed to be country-neutral, transparent, and easily employed—has several features:

- **Evaluates the presence *and* absence of activities, equipment, materials, patterns of behavior, and the broader context.** The firewall assesses whether these elements—individually and collectively—are compatible with the purposes states proffer for them. The identity of a country should not prejudice analysis of its nuclear program.
- **Provides insights into the nature and direction of nuclear programs and helps assessment of potential proliferation concerns.** The firewall can suggest which combinations of particular actions and other indicators should, over time, increase or decrease the sense of assurance or level of concern about a given state's nuclear-related activities, which could inform discussion of such concerns in national or multilateral settings.
- **Augments effective implementation of all three pillars of the NPT, namely nonproliferation, the peaceful use of nuclear energy, and disarmament.** By helping users to systematically identify weapons-oriented activities and enhance policy options for redressing them—the firewall can facilitate legitimate applications of peaceful nuclear energy and highlight ways that states can reassure others of the peaceful orientation of their programs. Further, by identifying comprehensive indicators of nuclear weapons programs, it addresses a necessary condition for progress toward disarmament. For without a technologically detailed template for defining how to turn military nuclear programs into purely peaceful ones, nuclear disarmament will not fully enhance security and therefore will not be politically achievable.

While elements of the framework have been vetted with distinguished international experts, governments and civil society should study them and consider which might merit further technical development and which might be ripe for implementation in multiple institutional settings. A firewall application could be employed by individual states assessing proliferation risks and making decisions about strategic trade controls, by international organizations and in multilateral forums, and by academic institutions and the nonproliferation community more broadly. Carnegie stands ready to assist with these efforts.

Introduction

Atomic energy, with its potential for peaceful and military applications, emerged internationally in the 1940s. In subsequent years, states incrementally established institutions and rules for managing this potential, such as the International Atomic Energy Agency (IAEA), created in 1957. In 1968, the Nuclear Non-Proliferation Treaty (NPT) opened for signature, and it came into force in 1970.¹ To motivate states to join the treaty, its principal authors, the United States and the Soviet Union, agreed that the obligations to stem weapons proliferation needed to be matched with obligations to facilitate access to peaceful applications of atomic energy and progress toward nuclear disarmament. These obligations were reflected in Articles I, II, and III of the treaty (nonproliferation), Article IV (peaceful nuclear cooperation), and Article VI (disarmament). By 1995, when the terms of the NPT called for an international conference of parties to decide whether, and under what conditions, to extend the treaty, 178 states had joined. At the 1995 Review and Extension Conference, the parties decided to extend the treaty indefinitely. Today, there are 190 NPT parties.

Inevitably, much has changed in the world since the treaty's enactment and extension. The circumstances in which the nuclear order must be maintained in the coming decades have changed dramatically and in ways that the NPT and other multilateral instruments could not anticipate.

The NPT and related agreements, instruments, and arrangements—such as the IAEA, the Zangger Committee, and the Nuclear Suppliers Group—remain invaluable foundations of nuclear order. However, innovations and new developments in technology sometimes arise and, inevitably, take time for the existing order to comprehend and manage. Similarly, institutional innovations take time to become widely implemented. For example, while progress has been made over the past two decades to introduce and promote the IAEA Additional Protocol, the goal of making it universal remains elusive. Other efforts to enhance these arrangements, such as the IAEA's advancement of the State-Level Concept, have proven contentious. In this situation, some states have been tempted to act individually or react to address what they perceived to be challenges and/or opportunities for gain in the nuclear order. Yet such uncoordinated action frustrates efforts to build consensus and may undermine the goals of the NPT.

Leading states in the international system could significantly strengthen the prospects of fulfilling the NPT's bargains on peaceful nuclear cooperation and disarmament if they cooperated in devising more comprehensive means for

assessing and addressing proliferation risks. For such cooperation to be feasible, useful, and internationally accepted, it must be based on solid scientific and engineering knowledge and experience, as well as transparent analytic tools that could be applied consistently and generically to all states. Less transparent and comprehensive assessments of risks, and/or favoritism in addressing them, will over time not only fail to manage emerging challenges but also erode the foundations of the existing order and further exacerbate stresses in the regime.

The Nuclear Firewall Concept

With these premises, the Nuclear Policy Program of the Carnegie Endowment for International Peace engaged external experts in a multiyear, low-profile brainstorming exercise to explore whether we could reliably identify generic activities or programs aimed at the development of nuclear weapons—which is prohibited in principle by the NPT—and offer options for employing such analysis in a menu of policy applications. The effort was inspired by the conviction that a clearer divide between weapons and non-weapons nuclear activities would simultaneously facilitate more effective implementation of the NPT's three pillars (nonproliferation, the peaceful use of nuclear energy, and disarmament).

This project was stimulated in part by recognition that the NPT does not define what a nuclear weapon is and therefore what constitutes nuclear weapons development activity prohibited under the treaty. The resulting concept of the project—a nuclear firewall—suggests a way to address this omission by helping to draw a line between activities that are exclusively weapons-oriented and those that are peaceful or for dual use. The firewall first seeks to identify activities or combinations of activities that experience suggests are either uniquely useful for the purposes of developing, acquiring, and sustaining an operational nuclear weapon capability, or whose most plausible use is in support of such an aim. Such activities are in the domains of research and development, testing and engineering, acquisition of materials, components and systems, construction of facilities for these purposes, and development of a military infrastructure (both physical and human) to employ nuclear weapons. This focus on weapons-related activities derives from the NPT's clear requirement that non-nuclear-weapon states refrain from manufacturing or acquiring nuclear weapons. Though individual indicators by themselves are rarely definitive, a program whose multiple activities occurring together reflect a cohesive effort to acquire a nuclear weapon would be in clear violation of this requirement.

For dual-use activities (that could serve military and/or civilian purposes), the firewall seeks to identify observable or otherwise measurable technical and contextual means and trends to distinguish activities that may confidently be regarded as having a nuclear weapons orientation from those that appear to be compatible with their stated non-nuclear-weapons purposes. In cases where the exclusivity of an activity for legitimate (non-nuclear-weapons) purposes is questionable, it identifies lines of inquiry about the alternative scientific,

technological, and commercial explanations that could help clarify whether an application other than nuclear weapons development is plausible.

Two Analyses: The Parts and the Whole

These lines of inquiry into alternative explanations address two levels of analysis (see figure 1 on page 10). The first is to assess whether individual activities have rationales that can plausibly justify, for example, their technical and/or commercial parameters, the materials employed, and the quantities involved. The second is to assess whether the program or enterprise to which a specific activity is related has the attributes of a credible program. To illustrate the second level of analysis in the domain of nuclear energy, the firewall project has delineated a standard set of traits, or a template, for a nuclear power program. That is, we identify activities and undertakings that comprise the common practices (technical, legal, institutional, and regulatory) of virtually all genuinely peaceful nuclear programs. This template of activities and behaviors provides a basis for comparison with any particular program. Consistency with this template, or deviation from it, permits assessment of the plausibility of alternative explanations when assessing other countries' nuclear activities and programs.

The first level of analysis—assessing how much concern exists about a certain item or activity in connection with a nuclear weapons application—is routinely conducted by the five nuclear weapon states under the NPT (as well as some other states), usually on a national basis and privately. This concern may actually be indicative of a clandestine weapons program. The second level of analysis—assessing the peaceful credibility of an entire program—however, is an innovation of the nuclear firewall concept and could be developed further as a standard practice for proliferation assessment.

A central, related innovation is the identification of categories of rationales that might be offered to justify (but also provide cover for) certain activities indispensable for a nuclear weapons program. In addition to the template of a peaceful nuclear energy program, the firewall concept highlights the utility of further expanding analysis in the areas of nuclear energy; ballistic missiles; and nuclear weapons research, development, and deployment. This effort, which requires the expertise and resources of states, would include the mapping of plausible legitimate explanations for activities in each domain and the identification of the requirements and conditions for such applications to be deemed credible. The most immediate benefit of developing this analysis is diagnostic. When and where indicators of concern emerge, the provision of rationales could either generate reassurance that the activities are indeed legitimate and deserving of support, or conversely, reinforce concerns that they are inconsistent with legitimate applications and thus merit further investigation. The aim is to suggest an internationally acceptable basis for engaging in this sort of analysis and inquiry in both intragovernmental settings and intergovernmental dialogues between the IAEA and its member states, within the Zangger

Committee and the Nuclear Suppliers Group, or between states and nongovernmental organizations (NGOs) in other institutions and contexts.

This analytical approach of assessing not only a state's activities but also possible omissions in its technical development, activities, or practices (compared with the templates of legitimate applications of nuclear technology) is a significant value of the firewall. The complementary lines of inquiry described above permit a more comprehensive analysis of nuclear programs than traditional proliferation assessment techniques.

The Five Cs Approach

Consistent with the IAEA's obligations pursuant to comprehensive safeguards agreements, the IAEA assesses the *completeness* and *correctness* of declarations regarding a country's nuclear activities. The firewall project seeks to provide an additional analytic framework to complement the IAEA's safeguards work by increasing the breadth and depth of proliferation analysis from two to five Cs (hereafter, the "5Cs"). We add considerations of *compatibility*, *cohesion*, and *consistency*. *Compatibility* refers to the degree to which individual activities undertaken are indeed commensurate with their stated peaceful purposes, including the sequencing and scale of activities and economic rationale underlying them. Activities that are incompatible register higher concern and may indicate weapons orientation. *Cohesion* refers to the extent that diverse, individual activities undertaken by a state are interconnected and reflect most of the known elements and pathways of a nuclear weapons program. Finally, *consistency* refers to the degree to which attributes of a nuclear program are consistent with the established hallmarks of credible civil or nonexplosive applications of nuclear technology (for example, in safety, security, environment, and liability).

The 5Cs describe the firewall project's holistic approach to assessing nuclear programs. Such an approach can help establish a legitimate basis for evaluating concerns regarding any given nuclear program, as well as provide reassurances about its nature. The 5C approach could aid discussion on potential cases of proliferation and help develop policy options to assuage concerns that might arise from looking at individual activities. Furthermore, it provides an opportunity to do so as a complement to national or multinational proliferation assessments.

Potential Implications and Applications

The project's comprehensive approach to proliferation analysis has several important implications for efforts to strengthen nuclear governance. By clarifying the interdependence of Articles II and IV, this approach could advance the objectives of all three pillars of the NPT. For example, we considered whether a program that closely conforms to the patterns, scope, synchronicity, and scale of established peaceful programs and exhibits no indicators of nuclear weapons could attract increased international cooperation (for as long as this

characterization holds). On the other hand, a state conducting activities that depart from standard practices of peaceful programs—including, especially, activities associated with weaponization—would be more readily assessed as transgressing or at a minimum encroaching on its obligation not to acquire or manufacture nuclear weapons. In such cases, states undertaking more sensitive activities might voluntarily reassure the international community of their commitment to peaceful obligations. The insights developed through this iterative analysis were subsequently tested using historical and contemporary case studies of proliferation and peaceful nuclear development in order to refine their conclusions.

Naturally, extensive discussions were required to establish how the firewall could be used to assess both indicators of proliferation warning and of peaceful orientation. The resulting analytic framework accounts for and weighs such indicators statically and over time to establish whether any given country is moving closer to or further away from nuclear weapons development and acquisition. Whether utilized in institutional or diplomatic contexts, this framework could facilitate a process of discussion or clarification.

The firewall project consciously refrained from prescribing how states individually or collectively should respond to concerns that may arise from this analysis. The linkage between assessment and policy prescription is left open for further discussion in the context of how the firewall could be applied in the future. However, to be effective, the analysis must be timely, recognized as technically valid, country-neutral, reliable (it would yield broadly similar results based on the same data set or information regardless of who used it), and robust (it employs all the indicators necessary to reach a valid conclusion). The analysis should also have built-in resiliency to defeat efforts to game it by determined proliferators, as well as to protect against biased or undependable results. In a complementary way, this analysis could inform states' efforts to reassure others that their nuclear programs are peaceful.

Governments, NGOs, academic institutions, and even private enterprises could use the firewall in a range of applications. Some potential analytic and diagnostic applications include informing multilateral deliberations on issues of proliferation concern and helping to guide national and multilateral strategic trade implementation and outreach. Some potential constructive or normative applications could include informing future actions that would strengthen the delineation between nuclear weapons activities and the rest.

Such applications, all of which would require further explication of the requirements, as well as development and testing, are discussed later in the paper. To enhance the political and technical acceptability of these applications, an innovation such as the firewall should have the following attributes:

- *Scope:* cover all three pillars of the NPT.
- *Application:* consistently and uniformly apply to all NPT non-nuclear-weapon states, with an explicit effort to enhance the legitimacy and political acceptance of the firewall concept and tool.

- *Country neutrality*: utilize to the maximum possible degree established, generic criteria, indicators, procedures, and institutions without prejudice or favor to the state involved.
- *Clarity*: explicitly delineate the boundaries between rights and obligations, and in practice, impose limitations and conditions on certain types of activities.
- *Transparency*: make the general principles public, even if some sensitive technical details and applications are withheld exclusively for official governmental use.
- *Robustness*: include all relevant known indicators to make the firewall as resilient to the evolution of technology as possible, as well as to prevent efforts to game it by potential proliferators. This will require periodic reassessment and refinement.
- *Parsimony and user-friendliness*: employ only those parameters necessary to generate valid conclusions and provide user manuals that enable them to be easily and effectively employed.
- *Modularity*: allow for the application(s) to be employed at different levels of resolution, with varying degrees of access to information.

Figure 1: The Firewall Concept

<p>The firewall evaluates the presence and absence of activities, equipment, materials, patterns of behavior, and the broader context to establish a comprehensive assessment of a nuclear program.</p>
<p>Analyzing the Parts and the Whole</p>
<ol style="list-style-type: none"> 1. Assessing whether individual activities have rationales that can plausibly justify, for example, their technical and/or commercial parameters, the materials employed, and the quantities involved. 2. Assessing whether the program or enterprise to which a specific activity is related has the attributes of a credible program.
<p>Using a 5Cs Approach</p>
<p><i>Completeness and correctness</i>: The attributes of a state’s initial declaration assessed by the IAEA under comprehensive safeguards agreements.¹</p> <p><i>Compatibility</i>: The degree to which individual activities undertaken are commensurate with their stated peaceful purposes—for example, related to the sequencing and scale of activities and the economic rationale underlying them. Lower compatibility registers higher concern and may indicate a weapons orientation.</p> <p><i>Cohesion</i>: The extent that diverse individual activities undertaken are interconnected and reflect many/all of the known elements and pathways of a nuclear weapons program, thus presenting more like a nuclear weapons enterprise than a peaceful nuclear program.</p> <p><i>Consistency</i>: The degree to which a program has the attributes of credible civil or nonexplosive applications of nuclear technology (for example, in the safety, security, environment, and liability realms).</p>

¹ See O.J. Heinonen, "Verification of the Completeness and Correctness of Initial Declarations," IAEA, <http://www-pub.iaea.org/MTCD/publications/PDF/ss-2001/PDF%20files/Session%202/Paper%202-02.pdf>.

The Firewall Framework

The firewall distinguishes activities and patterns of behavior in non-nuclear-weapon states that raise degrees of concern about nuclear proliferation, as compared with activities and patterns that are unambiguously peaceful. The former would obviously warrant closer monitoring, possibly alongside further steps by the state to reassure the international community about the peaceful purpose of its activities.

The firewall looks at both the presence and absence of activities and their evolution over time as indicators of purpose. It systematically identifies and weighs indicators and patterns of behavior. This approach diminishes challenges associated with the availability of comprehensive and valid information. It provides a holistic picture of a weapons program (as opposed to, for example, looking only at fissile material production). This approach also avoids the pitfall of technological determinism, which may arise in using methodologies focused on critical path analysis. Historically, many proliferators have pursued multiple parallel approaches to nuclear weapons development, thus underscoring the need to assess a broad range of activities over time. Augmenting this descriptive analysis, the firewall framework also utilizes a template of a credible nuclear power program to assess the compatibility and consistency of a state's nuclear activities. (As discussed later as an area of potential future development, a thorough assessment of all plausible non-nuclear-weapons applications of relevant technologies and activities and the incorporation of such templates would greatly enhance this framework.)

The results of a firewall analysis as it is applied to each case can be represented along a color spectrum ranging from green to yellow to red. This coloring scheme can be used in two ways: to delineate and describe the character of particular activities, either singularly or in combination; and to illustrate the sum output of all of the indicators and patterns, which provides an overall assessment of the state of a nuclear program at any given time. An activity or capability that registers green is peaceful. Dual-use activities are generally considered yellow, and a program whose capabilities and activities could be utilized for peaceful and/or military purposes is coded yellow. Finally, an activity that has exclusively weapons-oriented applications is coded red. A program registers as red when its collective capabilities and patterns of activities are uniquely, or strongly, associated with a nuclear weapons program. When looked upon comparatively over time, this coloring scheme portrays not merely the status of any given nuclear program but also its trajectory toward or away from nuclear weapons. Pictographically, clusters of activities and other patterns of behavior as viewed through the firewall are captured in a pyramid (see figure 2).

Figure 2. The Firewall Color Spectrum



The green-yellow-red representation of the output signifies three of the firewall's major policy contributions: (1) to provide a basis for determining that particular capabilities and activities—alone or in combination—are inconsistent with the obligation not to seek to acquire nuclear weapons; (2) to provide warning that a nuclear program is trending away from legitimate peaceful purposes toward illegitimate military ones—that is, from green to yellow or red; and (3) to identify patterns of behavior that could be modified to enhance confidence in the peaceful intentions underlying any given nuclear program.

The purpose of the firewall is not simply to assess where states fall on this color spectrum, but also to reveal concerning trends and patterns in a state's activities well in advance of reaching the red zone. To this end, a proliferation warning and reassurance application of the firewall could be built on the basis of a simple assessment scheme. A nuclear program whose collective attributes are coded as green, arguably, should be allowed to continue uninterrupted (while still being monitored). This judgment could favorably inform national decisions on issues of international nuclear cooperation or transfers. An observed increase in the pace or scope of disparate yellow activities in a state should result in heightened concern and vigilance and could prompt increased measures toward clarification and assurance. Finally, red activities or a program that is collectively coded red should elicit relatively high proliferation concerns. Such cases would warrant further investigation into activities and consideration of possible policy responses. A negative or alarming trend identified early on in a program would be enough to warrant discussion of a state's activities and the potential measures it could take to reassure others of its peaceful intentions.

Because the firewall aims to invite such discussions or clarification rather than to serve as the singular basis for punitive measures and seeks to extend

the time horizon to deal with such delicate issues diplomatically, the approach should be more tolerant of false positive interpretations in the absence of complete information on a state's activities. For the same reason, a proliferation assessment would not have to be wholly precise as to the exact status of a program at any given time, as this would require comprehensive, high-quality, and unambiguously accurate data. Instead, for the firewall to prove its worth, it should suffice to utilize information that, once integrated into the framework, could reliably and robustly indicate trends and overall patterns.

The comprehensiveness of the firewall derives from its use of multiple indicators and patterns, its accounting for the absence of indicators normally found in peaceful programs, and its registering of trends over time. The resultant framework specifically addresses three challenges: (1) the unavailability of complete and accurate information; (2) the complexity inherent in the dual-use nature of much nuclear technology, compounded by the existence of numerous (and changing over time) paths leading to nuclear weapons; and (3) deliberate concealment and deception by a state designed to throw off those organizations monitoring its activities (gaming the firewall). The multi-expert elicitation approach employed in developing the firewall also addresses a fourth challenge, which stems from variation in interpretation based on technical, cultural, or other differences in the backgrounds of individual experts. While the features of the firewall make it inherently difficult to game, concern over gaming could be further addressed by introducing degrees of transparency into parts of the framework or making it transparent only at a more general level. This, however, may come with some tradeoff in terms of weakening its value as a potential deterrent of cheating.

Firewall Indicators

A country may conduct many activities that are sensitive from the perspective of nuclear proliferation. Some of these activities may not be fully consistent with the country's stated purposes, and some may enhance the inherent (latent) capability of the country to acquire nuclear weapons whether the country plans to or not. Some activities, especially those that augment transparency, may be useful for redressing suspicions, assuaging concern, or attenuating the inherent nuclear weapons capability. The presence or absence of combinations of some important activities in a country can present a much clearer picture of the capability and possible future direction of the nuclear program than any one activity can. This clarity can emerge particularly when activities are measured over time and against certain contexts.

The firewall framework includes a group of technical indicators and a group of contextual factors (see figure 3 on page 27). The indicators are categorized under four technical domains: (1) fuel-cycle and reactor operations, (2) nuclear weapons research and development (R&D) and weaponization, (3) weapons delivery and systems (payload) integration, and (4) militarization. The contextual factors reflect state policy and conduct, mostly in aggregate; and they

include, for example, a state's track record of compliance and cooperation with the IAEA, sequencing of nuclear decisions, implementation of international norms, and preparation for military utilization of nuclear capabilities.

Fuel-Cycle and Reactor Operations

Naturally, fuel-cycle and reactor operations constitute one of the most important and relevant aspects of the firewall, consistent with the focus of IAEA safeguards on fissile material. Fuel-cycle activities are paramount and indispensable—without them, activities in other domains (such as weaponization and systems integration) have more modest meaning. However, for the purpose of the firewall, absence of evidence on fuel-cycle activity should not be interpreted to mean that no nuclear weapons program exists. While nobody doubts that possession of fissile material is a necessary condition for attaining a nuclear weapon capability, its mere existence is not sufficient to make a determination. Work in other domains would be necessary and should still register concern even if conducted in the absence of indications of fissile material work. For example, an implosive test involving depleted uranium or other simulant metal is inherently important. The rationale for this analytic perspective is threefold: the information available on some other activities may be incomplete (especially given the relatively low signature of centrifuge enrichment); some of the activities associated with fissile material acquisition and production could, validly or not, lend themselves to alternative (and more benign) explanations; and the fissile material required for bomb production may also be produced or acquired at a late stage, including from a foreign source.

In a country possessing (or seriously embarking on) a nuclear power program, the firewall project concluded that most basic reactor and fuel-cycle R&D activities would probably register as green. The most notable exceptions are some laboratory activities that involve facilities that have the capacity to produce high-purity fissile materials. These activities should be considered yellow. The scale of fuel-cycle activities, especially the quantity of materials involved, would be an important consideration in assessing the compatibility of such activities with stated nuclear energy research purposes.

The firewall project explored six areas relevant to fuel-cycle and reactor operations:

- Reactor design and operation
- Metallic fuels
- Uranium enrichment
- Fuel production and fabrication
- Reprocessing
- Fissile material storage

In the area of reactor design and operation, accepted international practice has evolved toward a preference for power plants that are inherently less well

suitable to weapons-grade plutonium production (for example, light-water moderated, low-enriched uranium [LEU] fuel for nuclear power plants [NPPs], as well as zero power research reactors). The design and construction of reactors that depart from these practices are indicators for consideration in the firewall (for example, thermal reactors with continuous online refueling, fast reactors with a uranium blanket, or research reactors with a thermal output of several tens of megawatts or more). In the area of reactor operations, similar deviations from evolved and common international practices are potential indicators of nonpeaceful nuclear objectives. For instance, short irradiations of large amounts of natural uranium, an unscheduled shutdown after a short fuel-irradiation period, and the unloading of some fuel after a short operation may be indicators of nonpeaceful objectives.

The issue of how to handle uranium metal in the firewall naturally arises. Conventional military applications of uranium metal are not proscribed, which complicates assessment in this area. Still, uranium metal has little practical utility in common commercial reactor and fuel-cycle operations. Few reactors still use metallic fuel (though future fast reactors might, albeit in highly specific processes). Uranium metal could be used in a laser enrichment program (which is a separate indicator). Thus, uranium metal is an indicator, yet one that has to be assessed in combination with others.

Uranium enrichment is a yellow activity and one for which context is important. For any state, enrichment greater than 20 percent uranium-235 (U235) is a particularly significant indicator. (This is not to say that enrichment beyond the more standard range of 3 percent to 5 percent U235, for example, should not raise concern.) For states that do not have an associated nuclear power program (operating or planned), any shift from laboratory-scale basic research and development to production would be a potential indicator of nonpeaceful objectives. Construction of facilities for enrichment beyond laboratory scale and production of uranium hexafluoride would constitute additional indicators. Two caveats present themselves: states beyond the P5 (China, France, Russia, the United Kingdom, and the United States) that have or purport to have naval propulsion programs said to require highly enriched uranium (HEU) fuel and states that claim to require HEU for medical isotope production. Low-enriched uranium alternatives are available for both applications; and therefore, apart from legacy systems, these caveats should get maximum scrutiny. The weight assigned to both of these alternative explanations for the intended production of HEU would hinge on evidence (including transparency) that such requirements are real enough and whether the pattern, quantities produced, and sequence and form match such requirements.

For states with operating nuclear power programs, a contextual indicator is whether production and storage capacities are appropriately scaled and sequenced to planned fuel requirements and configured solely for LEU production and mated with fuel fabrication (and related qualification), or instead appear purposed to accumulate fissile material. Another indicator is the development of enrichment technology long preceding actual fuel requirement, if

the aim is to produce more than around 100 grams of LEU. All legitimate enrichment R&D activity should be possible without accumulation or use of larger quantities of LEU.

Fuel fabrication is seen as an area with fewer potential indicators of use for the firewall. However, we considered one indicator to be the construction of fabrication facilities capable of handling high criticality, which could potentially suggest weapons-related purposes. A related indicator is preparation of a facility for work on metallic fuel and facilities to produce targets involving lithium or bismuth, which could be useful in boosting a weapon or making initiators. Another indicator is the existence of enrichment capability purportedly serving a civilian nuclear energy activity, but without any fuel fabrication activity to convert it into reactor-usable fuel.

Like enrichment, reprocessing is a heavily scrutinized capability. Scaling of R&D facilities (hot cells) for processing irradiated fuel larger than gram quantities would be a potential indicator, as would construction of a facility intended to result in pure elements.

In terms of material and waste storage, for most programs, there is no requirement to store large quantities of pure plutonium or enriched uranium. The accumulation of substantial quantities of pure elements or easily processed compounds is an indicator.

Nuclear Weapons R&D and Weaponization

A second category of indicators involves weapons R&D and weaponization activity. Whereas fuel-cycle issues are relatively well understood from the perspective of both nuclear governance and international proliferation risk assessment, weapons R&D and weaponization are relatively poorly described and therefore signify a substantial gap in understanding nuclear weapons programs. Evidence of engagement in this category of activities may not only provide more definitive indication of weapons intent than fuel-cycle activity, but may also generate such indications earlier in the life of a weapons program. By offering an additional vantage point on activities necessary to attain nuclear weapons, these indicators also enhance the probability of timely identification and correct interpretation of purpose.

Consequently, creation of a logically consistent and comprehensive typology of the activities required to turn fissile material into a nuclear explosive device or weapon is required. These activities range from theory and computation to fabrication and construction of nonfissile components. From the standpoint of robust indicators of proliferation warning, the firewall concept relies on an expansive definition of nuclear weapons. In this context, it is useful to distinguish between a nuclear *device* that can produce an explosion and a nuclear *weapon* that can be stored and moved safely, delivered reliably to a target, and detonated when intended within the design parameters. Nuclear devices are often large, have delicate components, and may require manual input prior to detonation. Nuclear weapons, on the other hand, are small and rugged enough

to be carried on a delivery vehicle (aircraft or missile). Weapons include fusing, firing, and safety systems that permit them to explode at a desired time and/or location. Development and acquisition of some types of nuclear weapons (for example, of the gun type) pose much lower requirements for technical sophistication, infrastructure, and testing.

The firewall project explored eight areas relevant to nuclear weapons development:

- Detonators
- High explosives
- Plutonium/uranium shells
- Neutron generators
- Tritium technology
- Hydrodynamic codes coupled with neutron transport calculation
- Hydrodynamic experiments
- Weaponization (in other words, preparation of the physics package for systems integration)

These eight areas comprise a range of activities and the equipment or technology necessary to carry them out, including on a simulated basis. R&D activity is obviously only one indicator of a nuclear weapons program, and no single activity in this domain clearly delineates a peaceful program from a nonpeaceful one. Many potential activities and technologies necessary for nuclear weapons R&D have civilian or at least non-nuclear-weapons applications, although not necessarily with the same specifications or materials or requiring the same tests or production techniques. Moreover, several categories of such technology are now practically ubiquitous or at least available in technologically advanced states.

Generally speaking, countries working on activities in all eight weapons R&D areas should invite much more scrutiny than countries working in only one or two. These activities are not equivalent in the exclusivity of their utility for nuclear weapons. Thus, their diagnostic value also varies. Some activities have several credible alternative purposes, while others are narrowly tailored to nuclear weapons or have unique technical specifications that make their use difficult to justify with a commercial, scientific, or any other rationale. Secrecy and concealment surrounding these activities (even if some technical specifications are withheld, say, for proprietary reasons) should only serve as another source of concern. By assessing the plausibility of alternative rationales for each area of activity, we can measure the compatibility of these activities with their stated purpose and thus roughly determine the strength of each area as an indicator of the nuclear weapons orientation of a program. (Note that this may not necessarily correlate with the nuclear weapons utility of each particular activity.)

In addition to these areas, there are other downstream activities that may or may not be conducted as part of a nuclear weapons development effort. For example, a state may prepare for and conduct an underground or underwater nuclear explosion. Depending on the scale of the effort, states may also build dedicated facilities for the design, testing, and manufacturing of non-nuclear-weapon components, as well as for weapons assembly. These types of indicators are perhaps more closely associated with weaponization and parallel work in the domain of weapons delivery and systems (payload) integration.

Weapons Delivery and Systems (Payload) Integration

States developing nuclear weapons confront a series of decisions about how to turn an R&D effort into an operational one. Most of these decisions relate to how to pair a nuclear warhead with a delivery vehicle. This category of activity, termed systems integration, is similar to weapons R&D in that no single activity indicates nuclear weapons application, but multiple activities that have dubious plausibility for other applications can establish a compelling assessment of incompatibility with the stated non-nuclear purposes.

Though many states that have tested and deployed nuclear weapons have done so first through gravity bombs, all have at least subsequently developed ballistic missiles, and much of the analytic effort is focused in this area. Moreover, this pattern may be changing. Many states that may consider developing nuclear weapons in the future already possess or are in the process of developing or acquiring potentially nuclear-capable ballistic missiles and may also be working on space launch vehicle (SLV) programs. Under the cover of an SLV program, a state may seek to develop and qualify technologies for ballistic missiles.

Much of the required technology for delivery vehicles is common for commercial, conventional, and ballistic missiles. For instance, software is widely available for the launcher design, trajectory, simulation, and guidance—all of which can be used for work on nuclear weapons delivery. In some areas, though, there are technical standards uniquely required for nuclear weapon system integration and delivery. It seems possible to identify requirements imposed by genuine space launch and related programs that would be missing if the SLV activity is merely intended as cover for a ballistic missile program. These would relate both to the overall scope of the program—for example, an array of satellite receiving stations, without which space activity looks unreasonable—and to its footprint and associated transparency. Civil space programs would also reasonably entail, for example, expectations regarding reliability, commercial aspects, and insurance arrangements. If possible, further work could be undertaken to define specificities with respect to civil/conventional applications and associated red lines for nuclear weapon activities.

Visible indicators of nuclear weapons orientation in seven key areas of delivery vehicle development include the following (certain technical specifications are deliberately withheld):

- Main functional missile specifications exceeding commercial requirements (for example, acceleration faster than 5 g).
- Main functional missile specifications with questionable conventional military uses, such as a range exceeding 2,000 kilometers.
- Software designed for reentry simulation, and development of a high-performance guidance, navigation, and control subsystem.
- Equipment for ground and flight testing with certain specifications.
- Certain features of reentry vehicles, such as the ability to detonate at high altitudes.
- Absence of an array of ground receiving stations.
- Absence of certain aspects of propellant technologies that are requirements for SLVs but not ballistic missiles.

Having noted these technical parameters distinguishing SLVs and potential ballistic missiles, any delivery vehicle, including cruise missiles and other unmanned aerial vehicles, capable of exceeding the 500 kilogram payload/300 kilometer range (and thus one that would be defined as Category I by the Missile Technology Control Regime) would raise the level of concern surrounding a program.

Militarization

The militarization of nuclear activity comprises steps intended to induct nuclear weapons for use by armed forces or to make them militarily useful. In this context, the firewall project considered four major questions:

- Should indicators of militarization be incorporated in the firewall?
- How unambiguous are they in terms of indicating nuclear weaponization?
- How relevant are they in terms of warning time for the identification of a nuclear weapons program compared to other indicators?
- Do technological advances in the last forty years or variations among nations in the way they go about the induction of nuclear weapons into their militaries mean that we might require different indicators today?

Upon looking closely at possible indicators of nuclear militarization activities, it became evident that while such individual indicators are not easy to detect and are somewhat ambiguous, they nevertheless have significant contextual diagnostic value. This is in part because they may indicate a purpose that goes beyond a latent and even hedging nuclear posture and in part because they are compelling as evidence of concrete intent and a timeline. Naturally, these indicators are especially potent when several are evident and when combined with activities of concern in other domains pertinent to the development of a nuclear weapon capability. Consequently, we concluded that militarization

indicators should indeed be incorporated into the firewall, along with other contextual factors to help amplify their utility. Given that some or all of the militarization functions might be performed in some countries by quasi-military or paramilitary organizations, the indicators mentioned below refer to any state organs performing the relevant military functions.

Possible indicators of militarization include, but are not limited to, the following (certain details are deliberately withheld):

- Direct military involvement in the administration and coordination of a nuclear program, or the employment of administrators with military backgrounds.
- Military responsibility for the construction of nuclear facilities and test sites.
- Military responsibility for the security and safety of nuclear facilities.
- Training, equipment, and preparation for the storage and maintenance of nuclear weapons.
- Involvement of nuclear experts in SLV activity.
- Establishment and training of special missile or bomber forces and operational bases.
- Dual-capable military aircraft deployed in an order of battle.
- Existence of ballistic missile forces in an order of battle.
- Training exercises for regular armed forces involving nuclear-capable forces.
- Incorporation of nuclear weapons into military doctrine or operational planning (for example, the necessary route, target information, and authorization codes).
- Hardening of command, control, and communications infrastructure.
- Exceptional arrangements accompanying certain missile activities and their missions.
- Special handling of certain (classified) military missions.

Other potential indicators of militarization could relate to certain engineering, testing, communications, and training capabilities and activities (details withheld).

The above listed indicators may vary widely in terms of visibility and in their timeframe relative to nuclear weapons development. But they all are potentially suggestive of a nuclear weapons orientation. Furthermore, the visibility of some of these measures may be greater than intuitively thought, precisely because they involve detection of anomalies in the routine practices common to all other missile and military SLV activity.

The logic underlying this conclusion is that a nuclear weapons program is much more than just the production of fissile material and weapons system

manufacture and integration. The development, deployment, and maintenance of fully engineered weapons that are incorporated into the armed forces of a state require considerable resources (human and financial) and a whole host of other activities. In particular, these include the organizational and bureaucratic structures needed to build and then sustain any real operational capability, the training of personnel on equipment and procedures, and the development and maintenance of a supporting infrastructure. Operational capabilities take time to develop and mature. If not constantly sustained, they will quickly atrophy. Furthermore, a historical review of several cases of nuclear militarization suggest that such activities have often started several years before a nuclear weapon capability was actually attained, in order to have an operational military capability ready when the technical preparations mature. Thus, they are doubly useful as warning indicators.

We underscore that various technical activities associated with nuclear weapons development and militarization activities run parallel and/or are contingent on progress being made in other directly and indirectly linked areas. This multidimensional process involves multiple activities running at various speeds and levels (from experimental to in-service). Some run continuously, others on a stop-go basis, and yet others episodically.

Finally, the relevance of this category of indicators has not diminished over time. While considerable variation occurs in the way various nuclear weapon states and aspirants approach the militarization challenge, they all basically aim to address a core set of requirements and dilemmas necessary to acquire and maintain an operational capability.

Contextual Factors

Many nuclear technologies—as well as other technologies required to produce nuclear weapons—are recognized to have dual uses. They may be used for peaceful nuclear applications, for military applications other than nuclear weapons (from naval propulsion and shaped charges to depleted uranium shells), for various commercial applications (such as oil drilling), and even for basic science. Given this challenge, the project concluded that the firewall framework should draw on factors related to the broader context surrounding a state's nuclear program to complement the technical analysis of individual activities. Context can help facilitate and clarify an interpretation of purpose in the event that a non-nuclear-weapon state makes technical and/or operational progress in areas relevant to the acquisition of a nuclear weapon capability.

It is useful analytically to distinguish two clusters of contextual factors characterizing a state's nuclear program and policies: features of and anomalies in the state's activities and programs that may cause alarm, and anomalies created by the *absence* of attributes that are indispensable for a credible peaceful nuclear energy program. These clusters provide an aggregate sense of a state's track record and purpose. Combining the aggregated presence *and* absence of key activities and technologies yields a compelling picture regarding the overall

orientation of a state regarding nuclear weapons aspiration. This helps in interpreting particular technical and operational indicators that by themselves may yield an ambiguous or inconclusive assessment of proliferation risk. The contextual factors are akin to a checklist. The greater the number and diversity of the alarming boxes checked, the higher the level of concern it ought to trigger, and vice versa. Just as the contextual factors could heighten concerns over a program, they could also reassure the international community about the purpose of a state's activities.

Four important considerations went into the development of the lists of contextual factors. First, the factors should have diagnostic utility and be objective; expert analysts from a range of states need to agree on their merit regardless of which state is being analyzed. Second, the factors should be parsimonious yet add to the robustness of the analysis to help defuse concerns that a state seeking nuclear weapons could not easily game the firewall by avoiding designated indicators and still succeed in producing a nuclear weapon. Third, the factors should be evaluated with relative ease and reliability. They should be easy to monitor and assess so that no massive clandestine intelligence collection effort would be required to use them for a proliferation analysis. By implication, every user of the criteria should reach the same conclusions based on easily obtainable data. Fourth, considering challenges associated with information collection and diplomatic utility, the contextual factors must be predominantly based on publicly observable behavior and practice. Scoping the factors in this way ensures not only easy data collection but also easy data use, without sensitivity to sources and methods (which often precludes the public or diplomatic use of intelligence information).

The contextual factors are divided into two categories: those that could easily be operationalized (binary factors) and those that are more difficult to operationalize (for the moment, qualitative factors). The totality of these factors is more important than whether any one of them is apparent.

Binary Factors

- NPT membership.
- Full-scope safeguards agreed upon with the IAEA and implemented.
- Additional Protocol ratified and implemented with the IAEA.
- Clean track record of IAEA safeguards' implementation (for instance, no issue of safeguards implementation brought before the Board of Governors in the preceding decade).
- No referral to the United Nations (UN) Security Council by the IAEA Board of Governors for failure to address an open safeguards issue in the preceding decade.
- Open and transparent nuclear activities (for example, details of the program are regularly published).

- No meaningful indigenous fuel-cycle activities.²
- Fuel-cycle activities entirely run by commercial entities.
- Independent national nuclear regulator in place for health, safety, environmental monitoring, and licensing.
- Nuclear program subjected to international safety, security, and liability standards (Convention on Nuclear Safety, Convention on the Physical Protection of Nuclear Material, Joint Convention on the Safety of Spent Fuel Management and on the Safety of Radioactive Waste Management, Convention on Supplementary Compensation for Nuclear Damage, or Vienna or Paris conventions).
- No parallel chemical and/or biological weapons programs.
- Membership in good standing with other leading nonproliferation (disarmament) treaties (Biological and Toxin Weapons Convention, Chemical Weapons Convention, and Comprehensive Test Ban Treaty) and weapons-free zones.
- No development of intermediate-range ballistic missiles, intercontinental ballistic missiles, SLVs, and large or long-distance cruise missiles (exceeding Missile Technology Control Regime guidelines).
- No unresolved information indicating clandestine or dubious nuclear-related procurements (to include export denial notifications) over the past decade.
- No information indicating use of space activity as false cover for ballistic missile activity.

Qualitative Factors

- Non-implementation of UN Security Council Resolution 1540.
- Challenging security environment (having one major or sustained adversarial interstate relationship).
- Concealment of nuclear activities (that in established peaceful programs are not concealed).
- Incompatibility or asynchronicity between fuel-cycle activity and nuclear power requirements for civilian (or naval propulsion) purposes:
 - Temporal—preceding substantial requirement for fuel.
 - Scope—quantities involved exceed reasonable requirements or are produced in absence of technology or a license to make associated fuel.
 - Safety culture—activity occurs in the absence of a transparent and formal regulatory safety review and formal licensing arrangements.
 - Priority—fuel production takes precedence over power.
 - Types—numbers and arrangement of centrifuges in cascades are not aligned with stated purpose.

- Weak commercial rationale for fuel-cycle activity (fewer than ten NPPs requiring such fuel, inadequacy of available uranium supply, lack of reliable and cost-effective enrichment technology).
- International isolation (lack of integration into the international community).

Firewall Algorithm

In developing the firewall, the project aimed to translate the framework, technical indicators, and contextual factors into a set of calculation rules—a firewall algorithm—that could assess the individual and combined impact of proliferation-sensitive activities and patterns of behavior on advancing the capacity of a nation to produce a nuclear explosive device and subsequently produce, deploy, and use a nuclear weapon. The algorithm offers a transparent methodology to measure latent nuclear weapon capability, or alternatively, the progress toward nuclear disarmament in a country.

The project developed a proof-of-concept algorithm to demonstrate its usefulness and feasibility.

The algorithm is built around two core concepts, defined as output variables, to assess the proliferation situation in a country: (1) latent nuclear weapon capability and (2) compatibility with stated purpose. The first describes the progress made by a country toward acquiring nuclear weapons. The second describes how and to what extent a country's nuclear activities are consistent with their stated (civilian or conventional military) purposes.

Both variables have implications for the level of concern in proliferation assessments, but they are important at different points in time. When a country is still a few years away from acquiring nuclear weapons, incompatibility with stated purpose is more useful for detecting hidden efforts toward a nuclear weapon program early and giving meaning to activities that are conducted visibly but whose purpose is concealed. When a country needs less than a few years to acquire nuclear weapons, both variables assume salience and policy relevance. Since the NPT does not explicitly preclude many activities that contribute to nuclear latency, the use of both of these variables (latency and compatibility) in a complementary fashion is indispensable.

A country may or may not acquire capabilities as the result of a deliberate effort to acquire nuclear weapons. When the capability is the unintended side effect of a country's peaceful nuclear and other pursuits, it is referred to as latent capability or simply latency, following a common practice in the academic literature.³ The intentional acquisition of capabilities could be for the development of a nuclear option (in other words, hedging posture) or the pursuit of a full-fledged nuclear weapon capability. Regardless of the purpose, a country could accumulate alarmingly significant capabilities as a by-product of extensive dual-use nuclear and other activities while remaining in good standing with the NPT. Nevertheless, a country that acquires an advanced latent capability (providing it with the capability to acquire nuclear weapons

relatively quickly were it to embark on such a path) would merit special attention to ascertain that it does not at some point change its orientation toward nuclear weapons.

For this reason, latent nuclear weapon capability is a technical measure defined in the algorithm as the accumulation of capabilities that facilitate the acquisition of nuclear weapons. Latent nuclear weapon capability can be measured by the ratio of proliferation-sensitive activities that a country has been working on to all the activities required for a nuclear weapon capability as measured in the algorithm. The threshold for warning of a nuclear weapon capability should be the ability to produce a nuclear explosive device (rather than a full-fledged arsenal). But the algorithm could be useful at higher thresholds as well (for example, in assessing vertical proliferation). We focused on designing an algorithm to warn of the ability to produce a nuclear explosive device; an algorithm better suited to assess risks of vertical proliferation could be developed.

In contrast to latency as it is used more generally, latent nuclear weapon capability incorporates all activities—nuclear and non-nuclear, both dual-use and exclusively weapons-oriented. It does not presuppose the nature or intent of a program. It reflects the completeness of a nuclear weapons program irrespective of its purpose. This measure is thus merely a technical assessment of a country's capabilities and is insufficient by itself to discern whether such capabilities are designed for a nuclear weapons program.

Compatibility with the stated peaceful purpose is based on observable criteria that indicates whether proliferation-sensitive activities can be reasonably justified for non-nuclear-weapons use, based on technical, economic, and other measures of plausibility of its alternative use. (Non-nuclear-weapons use could be for peaceful scientific, medical, agricultural, and energy purposes or for conventional military applications.)

Compatibility with the stated peaceful purpose of an individual activity depends on the extent to which the activity (or absence thereof) conforms with reasonable indicators associated with the purpose, or non-nuclear-weapons military application. For an individual, proliferation-sensitive activity, compatibility with the stated purpose is assessed mainly on the technical nature of the activity. Depending on the activity and purpose in question, it may raise concern regardless of the presence or absence of other indicators.

The algorithm also yields an assessment of the overall compatibility of a country's program based on the combination of individual activities or patterns of behavior. The result may also be evaluated according to technical, economical, and legal parameters constituting the context surrounding a program's activities. This context could indicate, based on established international practices, that a program closely shares characteristics of those in nuclear-armed states or conversely of a peacefully oriented nuclear program. During the fire-wall development process, the project evaluated qualitative and quantitative

criteria for assessing compatibility for most important activities. The results laid a foundation for future calibration of the core compatibility indicator.

Combinations of different activities may change the assessment of general compatibility with the stated purpose of a country's program, with both linear and nonlinear effects. A linear effect is a situation in which two or more activities contribute their individual compatibilities to the general compatibility score according to a weighting based on the nuclear weapon-related utility of an activity. This is a measure of how much that particular activity advances a country's nuclear weapon capability. If a specific activity has a large nuclear weapon utility, it will have a larger weighted contribution to the general scoring of compatibility with a stated purpose.

Proliferation-sensitive activities by their technical nature may have different relations among each other in a nuclear weapons program. They could be parallel (for example, fissile material production and missile development). They could be alternative (for example, the production of plutonium and of highly enriched uranium, though this particular example could also constitute parallel activities). Or sets of activities could be sequential to each other (for example, the development of gas centrifuge machines and the production of enriched uranium). An effective algorithm must account for various ways to measure the implications—for warning or reassurance—of parallel, alternative, and sequential activities.

In particular, the algorithm must consider the role of nonlinear effects involving the combination of different proliferation-sensitive activities. Many individual dual-use activities might be ongoing in a country that does not have a nuclear weapons program. Furthermore, these activities may not be technically or institutionally coordinated. In this case, the achievement of latent weapon capability would not be pursued through clandestine actions. However, some activities must be integrated at certain stages if they are designed to contribute to a nuclear weapon capability. To that end, the project identified some pairs of activities that are particularly relevant to a warning of clandestine nuclear weapons programs (details withheld).

Some important proliferation-sensitive activities in different domains of nuclear weapons development have been identified, as have the linear and nonlinear effects of the combination of these important proliferation-sensitive activities. Available information for a proliferation assessment may be incomplete because potential proliferators may hide elements of their nuclear programs. The firewall methodology attempts to mitigate the problem of incomplete information through the scope and robustness of its indicators across all relevant domains, applied over an extended period of time. The algorithm should be able to identify urgent or alarming situations if a few important proliferation-sensitive activities are known. We leave open the possibility for not only a single firewall algorithm but several algorithms scoped for different applications. Each algorithm would require an appropriate assessment of requirements to meet the desired parameters for the intended application,

as well as a user manual that provides relevant instructions and guidance on understanding any uncertainty inherent in the model.

Figure 3: The Firewall Framework

<p>The firewall systematically identifies and weighs patterns of behavior using indicators under four technical domains and various contextual factors.</p>	
<p>Four Technical Domains: The necessary technical and organizational elements whose presence or absence is most diagnostically useful in assessing nuclear programs.</p>	
<p><i>Fuel-cycle and reactor operations:</i> Activities that could contribute to the production of fissile material, such as reactor design and operation, uranium enrichment, and fuel production and fabrication.</p>	<p><i>Weapons research and development:</i> Activities intended to turn fissile material into a nuclear explosive device or weapon, including detonators, plutonium/uranium shells, and hydrodynamic experiments.</p>
<p><i>Delivery vehicles and systems (payload) integration:</i> Activities intended to mate a warhead to a delivery vehicle, including the design of software for reentry simulation; the development of a high-performance guidance, navigation, and control subsystem; and the advancement of certain features for reentry vehicles, such as the ability to detonate at high altitudes.</p>	<p><i>Militarization:</i> Activities intended and required to induct nuclear explosive devices into military systems and structures and to make them militarily reliable and useful over time, including direct military involvement in the administration and coordination of a nuclear program, or the employment of administrators with military backgrounds, and involvement of nuclear experts in SLV activity.</p>
<p>Contextual Factors: The various economic, political, military, or other characteristics of the conduct of a country that shed light on the purpose behind its nuclear and related activities.</p>	
<p><i>Binary factors</i> include full-scope safeguards agreed upon with the IAEA and implemented, no referral to the UN Security Council by the IAEA Board of Governors for failure to address an open safeguards issue in the preceding decade, and fuel-cycle activities entirely run by commercial entities.</p>	<p><i>Qualitative factors</i> include concealment of nuclear activities (that in established peaceful programs are not concealed) and incompatibility or asynchronicity between fuel-cycle activity and nuclear power requirements for civilian (or naval propulsion) purposes.</p>
<p>Firewall Algorithm(s): A set of indicators and calculation rules to assess the individual and combined impact of proliferation-sensitive activities on advancing a nation’s capacity to produce a nuclear explosive device or weapon. The algorithm offers a country-neutral, transparent methodology to measure how alarming a proliferation situation is or, alternatively, how much progress has been made toward nuclear disarmament. Different algorithms varying in function and level of detail could be employed in different contexts. The algorithm uses two main variables.</p>	
<p><i>Latent nuclear weapon capability:</i> The accumulation of capabilities that facilitates the acquisition of a nuclear weapon. This is based on the product of nuclear weapons utility (a measure of how much an activity advances a country’s nuclear weapon capability) and the status of the activity in a given state.</p>	<p><i>Compatibility with stated purpose:</i> A judgment, based on observable criteria, of whether proliferation-sensitive activities can be reasonably justified as intended for non-nuclear-weapons use (be it for peaceful nuclear, other civilian, scientific, or conventional military purposes).</p>

Possible Uses of the Firewall

The firewall can have several potential applications for governments (individually and collectively), NGOs, academic institutions, and companies. These applications fall into two broad categories: (1) analytic and diagnostic—facilitating assessments of current and past nuclear programs and related activities and (2) constructive or normative—potentially informing future actions and policies to delineate purely peaceful nuclear programs from military ones. This forward-looking application of the framework could facilitate development of peaceful nuclear energy programs and, in other settings, could inform nuclear disarmament measures at the national, regional, and global levels. Applying the firewall in multiple ways could synergistically strengthen each and, therefore, all of the central pillars of the NPT.

The most obvious category of applications is analytic: to provide a framework for states, international organizations and forums, and NGOs to identify and assess indicators of proliferation. The firewall helps integrate and assess the relative importance of disparate pieces of information on a country's known and suspected activities. It helps weigh this knowledge and assessment against standardized models of nuclear weapons programs and peaceful applications of nuclear technology. This would be useful not only for highlighting proliferation threats and singling out activities or patterns of behavior that raise concern, but also for providing guidance on how a country can reassure others of the peaceful nature of its activities. Such an application of the firewall could be employed by (1) individual states assessing proliferation risks and making decisions about strategic trade controls, (2) international organizations and multilateral forums, and (3) academic institutions and the nonproliferation community more broadly. Moving beyond the current, proof-of-concept phase of firewall development into validation of the framework through extensive, independent testing using hypothetical or real-world scenarios would greatly enhance the relevance of the firewall in these contexts.

The firewall could be used in multiple analytic and diagnostic applications:

- In national and multilateral deliberations on issues of proliferation concern.
- In an algorithm tool for the technical assessment of proliferation concerns and reassurance.
- In national and multilateral strategic trade control consultations, implementation, and outreach.
- As a complement to IAEA safeguards conclusions for assessing possible proliferation concerns about a state's nuclear program.
- For identifying activities that could be deemed inconsistent with peaceful nuclear programs, with a view toward supporting the roll-back of nuclear weapons programs and informing the design of future disarmament regimes.

- For benchmarking non-nuclear-weapons applications of dual-use activities.
- For facilitating the reassurance and verification of states' compliance with obligations.

It could also be used in constructive or normative applications that are forward looking:

- For strengthening the buffer between peaceful and weapons-oriented activities and enhancing the credibility of peaceful nuclear energy programs.
- For designing national, regional, and/or global disarmament agreements.

Analytic and Diagnostic Applications

National and Multilateral Deliberations

The most obvious application of the firewall is to provide a platform and common language for enhanced interstate collaboration on nonproliferation matters. More than ever, such collaboration is necessary to successfully disincentivize and counter nuclear proliferation. Yet, collaboration in this domain is impeded by suspicions that states are biased in their approaches to particular cases of proliferation concern. Within the IAEA Board of Governors or the UN Security Council, for example, one or more states may display greater or lesser vigilance against an adversary state than against a friendly one. A generic, transparent basis for evaluating nuclear programs, such as the one afforded by the firewall, could facilitate more effective deliberation and collaboration in this realm, particularly among the P5. The firewall could also be valuable in facilitating diplomacy at the IAEA Board of Governors or other diplomatic consultations for cases in which questions arise about the nature of a state's nuclear activities and whether the international community and/or the state should take reassurance action to lessen associated proliferation concerns.

The firewall could provide similar benefits for deliberations *within* states. The ability to provide (share) classified analysis of proliferation risks within the ranks of government can become entangled in politically charged controversy, which in turn complicates effective handling of the issue. A firewall framework (co-produced generically by acknowledged experts) could facilitate the communication of proliferation concerns in more easily understood terms to nontechnically oriented decisionmakers.

Algorithm for Technical Assessment of Proliferation Concerns and Reassurance

The differing, often discordant approaches to assessing proliferation threats by individual states, international organizations, and the nonproliferation community often results in conflicting interpretations and highly politicized debates. The firewall aims to deliver a transparent, technically valid means to help distinguish between peaceful nuclear programs and illegitimate nuclear

weapons programs. The algorithm proof-of-concept is one manifestation of this methodology. The algorithm—like other proliferation models developed by government, laboratory, and academic experts—is not designed to supplant technical expertise or national or international authorities. Rather, it seeks to provide a common analytical framework that would help facilitate and guide discussions within and between expert communities.

The project began by laying out a set of criteria identified by experts that would guide development of the algorithm. The proof-of-concept was evaluated and refined through application to historical case studies of nuclear programs. This process in turn informed the development of the firewall more broadly. However, the creation of a robust, polished algorithm ready for application by states exceeded the resources of this project. A fully developed firewall algorithm would require additional, rigorous statistical modeling to ensure it is technically validated. Further refinement or modification of the methodology would be required for use in certain settings. The resolution or level of detail required for national proliferation risk-assessment activities, for example, may greatly exceed what would be required to utilize the firewall in diplomatic outreach. Thus, in addition to the proof-of-concept, we explored technical requirements to guide further development, whether undertaken by governments or other actors. A complementary approach would involve an open-source firewall platform that governments, NGOs, and others could test and improve for their desired application. Should NGOs or academic institutions take up this approach, one could imagine the algorithm being used to facilitate comparative open-source analysis. These approaches are not mutually exclusive.

*National and Multilateral Strategic Trade Control
Consultations, Implementation, and Outreach*

One important potential application of the firewall is in the area of nuclear trade regulation. Here, the focus is on utilizing the firewall framework to strengthen (and help streamline) national approaches for assessing proliferation risk involving Trigger List items or dual-use nuclear commodities, as well as international coordination, outreach, and information sharing in venues such as the Nuclear Suppliers Group.

The firewall could help improve the formulation and implementation of national nuclear trade policy, particularly nuclear and dual-use export controls. All supplier states struggle to ensure industry compliance with national export guidelines. States develop rules and guidelines but often find it difficult to educate (and update) myriad technology suppliers about them and, relatedly, to explain the rationales behind the rules. The firewall could facilitate government decisionmaking—and improved unclassified communication between government regulators and industry—to help industry (as well as financiers and lenders) better assess proliferation risk in their commercial spheres and strengthen compliance with national export controls. A comprehensive

open-source risk assessment tool such as the firewall could also help government and industry adapt more quickly to changes in procurement patterns and proliferation threats.

Internationally, the firewall could serve multiple purposes, including in venues such as the Nuclear Suppliers Group, where participating governments discuss and coordinate multilateral export control policies. The firewall could provide a more common basis for assessing proliferation risk among participating governments, as well as a means to structure and facilitate shared understanding of proliferation risks through established information exchange procedures. Again, because this is an open-source product, it obviates concerns that individual governments may have regarding classification issues. More ambitiously, the firewall could also improve coordination among states to assess and strengthen existing supply lists and guidelines beyond the current operational baseline.

Export control regimes, such as the Nuclear Suppliers Group and (to a lesser extent) the Zangger Committee, encourage robust export control procedures to be implemented in all countries. However, the denial of licenses and the bans on transfers of particular technologies can be perceived as unfair by those affected. The rationale underlying such decisions often cannot be explained (in part because of the sensitive information on which it is based), and even then it may not be easily understood given the complicated technologies involved. The firewall's open-source indicators could be used to help explain some of the reasoning behind export licensing policies and decisions and could thus help improve the understanding of and collaboration on such decisions. The firewall therefore has a high outreach potential among national industries, universities, and research organizations working on sensitive technologies, as well as state governments that have significant transshipment and/or dual-use commodity businesses.

Finally, the firewall could inform procedures for monitoring procurement channels such as those incorporated on an ad-hoc basis in the Joint Comprehensive Plan of Action for Iran, as part of international nonproliferation agreements. By weighing indicators of proliferation concern and drawing attention to activities that seem out of sync with stated peaceful purposes, the firewall can hone analysis and monitoring of procurement.

Complementing IAEA Safeguards Analysis

IAEA safeguards analysis has evolved along with the agency's mandate, tools, and ability to utilize disparate streams of information. A robust debate continues at the agency over further development of safeguards concepts. State-level analysis in particular has provoked considerable discussion.⁴ Recognizing these sensitivities, as well as the agency's special role in verifying the absence of undeclared activities that might contribute to a nuclear weapon, the project considered whether and how the technical elements of the firewall could help inform

state-level analysis. We believe it could be particularly useful in complementing and helping to prioritize safeguards coverage in a state that is developing pathways to acquire fissile material for nuclear weapons, whether such analysis is performed by IAEA member states or the agency's technical staff.

IAEA safeguards analysis currently relies in part on a Physical Model of the nuclear fuel-cycle and related activities.⁵ The IAEA developed the model (with the assistance of governmental experts) in response to the problematic safeguards experience in Iraq prior to 1991. After periodic updates to the model over the years, it now comprises a multivolume, internal IAEA set of documents, which analyzes in some detail the various nuclear materials and related development paths to nuclear weapons. The model provides IAEA analysts with a cognitive map to help inform state-level safeguards analysis and implementation planning, detect indicators (with designated strengths) of nuclear weapons-oriented activities in the nuclear fuel-cycle and related domains, and put in context information the agency collects as part of its safeguards analysis.⁶ It has become an integral part of the ongoing state evaluation process in the IAEA Department of Safeguards.

The Physical Model, relying primarily on critical path analysis, looks exclusively at nuclear materials to detect inconsistency between state declarations and other available information about the state. Indicators are identified as strong, medium, or weak. The strength attributed to each indicator, however, is not only a matter of subjective judgment but also is to a degree dependent on the economic circumstances prevailing in the country. The limitation inherent in critical path analysis is that these indicators do not change in combination; there is no methodology for attributing a weight or combined strength to a *series* of disparate indicators. The model can help the IAEA detect activities relevant to nuclear weapons development, but the presence of certain strong indicators can only suggest reason for concern. The firewall can supplement this analysis with an assessment based on complementary methodologies as well as quantitative and transparent criteria. Such complementary analysis, for example, could help emphasize safeguards coverage of certain fissile material acquisition paths over others (if the indicators suggest a state may be moving along one or more particular paths) or seek in accordance with its mandate supplementary access or information to address any concerns.

The firewall's inclusion of indicators beyond those related directly to fissile materials enables it to more effectively cue analysis and investigation of the entirety of a state's nuclear-related activities. This includes analysis of the presence or absence of activities that when viewed individually would not be alarming, but when considered collectively may be an indicator of capabilities that warrant a higher level of IAEA verification—to maintain the necessary safeguards assurances regarding one or more possible acquisition pathways in that state for obtaining nuclear materials. The firewall could thus inform the agency's state-level assessments and its specific approach to implementing safeguards in a given country.

Benchmarking Non-Nuclear-Weapons Applications of Dual-Use Activities

The firewall project underscores the potential benefit of a complementary examination of not only the characteristics and parameters of proliferation-sensitive activities being conducted by a state, but also the credibility of peaceful and other non-nuclear-weapon applications offered to justify such activities. Analyzing an activity or program's consistency and compatibility with normal peaceful nuclear activities and programs can reveal as much about the purpose of an activity as focusing on its application to nuclear weapons. Non-weapon applications often impose much different and at times even more stringent requirements, relating to matters such as financing, reliability, quality assurance, certification, product liability, and marketing in the case of commercial products. If an activity is intended for weapons use, it is likely that many of these civilian requirements would not be met. Omission of these characteristics thus becomes an important indicator. Persistent opacity surrounding an activity and the parameters under which it is conducted would give further reason to question its purpose.

A template of a normal peaceful nuclear energy program can highlight aberrations that should raise concerns over the ultimate purpose of activities. Such a template can be particularly useful in benchmarking plausible applications of proliferation-sensitive, dual-use activities. These applications could include scientific research, commercial use, technological development, and conventional military or security applications. Several examples illustrate how such a template could be created and applied to activities in the fuel-cycle (other than nuclear power) and nuclear weapons R&D domains (certain details withheld):

Reactor Technology

- If a case is made that research reactors are to be used for training and instruction, that logic should manifest itself inter alia in transparency surrounding the project and its purpose; the selection of the type and features of the research reactor (for example, zero or low power); and the associated arrangements for training. There also must be a realistic need and concrete plans for leveraging this knowledge and personnel into a program for the peaceful application of nuclear technology.
- Research reactors for medical, agricultural, or industrial applications should similarly be tailored for these purposes and refrain from significant product accumulation in excess of demonstrated need. Even more importantly, they should be accompanied by extensive arrangements to endow such products with quality assurance, other safeguards and certification, and associated efforts to market and distribute the product to end users.

Nuclear Weapons R&D

- To justify development of neutron generators that could be used for weapons, a country might argue that these are intended for industrial use in oil-well logging, security diagnosis, or reactor startup. Closer scrutiny

could reveal whether the generators have costs and capabilities for precision timing that exceed requirements for civilian applications. The inquiry could also establish the presence or absence of the investment, logistics, and economic rationale for commercial oil development. What application requires the neutron generator to have these specific output requirements? And if there is an ostensible application, are there other indicators that this is a serious, commercial enterprise?

- Actinide machining is required for some advanced reactor work and basic science of actinide metals. But the manufacturing of plutonium or uranium shells that might be used in an implosion device requires special technology and methods. If the explanation offered would suggest a material science application, the natural rejoinder would be to probe deeper into the features of the materials science program and ask for the scientific articles it produces.
- Explanations for work with hydrodynamic codes could be offered based on the development of conventional warheads, astrophysics, and other basic research. But these would necessarily raise questions on the specific materials and testing parameters employed, depending on which purpose is suggested.
- Similarly, hydrodynamic experiments may plausibly be used for the development of conventional warheads (especially shaped charges). The inquiry would center on the implosion aspects and its purpose, sponsor, and background.

The requirements to fully develop this line of reasoning—by identifying all plausible explanations for dual-use activities and integrating templates of each into a robust methodology—would be considerable. The effort would demand not only expertise regarding the nuclear weapons applications of these technologies but also expertise in any given area of peaceful activity—from oil drilling to satellite launch programs to nuclear medicine.

*Facilitating Reassurance and Verification
of States' Compliance With Obligations*

The firewall project has found that traditional verification could be enhanced significantly by mechanisms for reassurance. In this regard, discussions of the reassurance and verification implications of the firewall drew lessons from several existing verification regimes including the Chemical Weapons Convention, Biological and Toxin Weapons Convention (BTWC), and the Comprehensive Test Ban Treaty. These verification regimes share many procedural commonalities and face similar verification challenges.

In the context of the BTWC, for example, relevant equipment, materials, facilities, and knowledge generally have both peaceful and weapons applications. This is true also for some nuclear activities, but with biological

infrastructure, the problem is much more complex and in many ways much less tractable. Furthermore, experience with these regimes has shown that assumptions about the purpose of state parties' actions cannot easily be made based on the presence or absence of single, or even multiple, factors. Rarely can the nature of a state's program be described in black and white terms—instead the picture is often obscure and multidimensional. At the operational level, therefore, some lessons and techniques can be borrowed from existing regimes. For example, the draft BTWC verification protocol (though never concluded) contained clarification procedures short of challenge inspections to deal with situations of ambiguity, and more recently, parties to the BTWC have explored the use of a voluntary peer review system.⁷

The emphasis given here to the mechanisms built into these regimes to handle implementation concerns is not intended to suggest that their success or failure is exclusively determined by their technical merit. Indeed, historical successes and failures have made it clear that the effectiveness of verification systems often depends on factors external to the structure of the mechanism itself. These factors include the influence of major powers, the level of trust or hostility toward inspectors, and the integration of the inspected state into the international system. Psychological and personal factors may also play an important role, such as the ability of inspectors to deal with uncooperative states or practices. That said, the relevance of technical and procedural issues and their impact on the outcome should not be underestimated.

Analysis of the biological and chemical weapons regimes suggests that the fire-wall could help with reassurance and verification for the following reasons:

- A broad range of activities and facilities must be accounted for across multiple levels.
- Not just quantitative, but qualitative indicators are important and should be factored into conclusions (for example, not just what is done but the nature of how activities are conducted).
- The broader context within which activities take place is important. This includes temporal considerations: confidence and assurance may grow or diminish over time. The overall trend may reveal the changing nature of a program.
- The temporal dimension is not only important to ascertain intention, but also to allow sufficient space for peaceful resolution of concerns if and when they arise. Hence, the ability to flag concerns in a timely manner (without stigmatizing the states involved) is germane to the success of any verification regime.

For purposes of reassurance and verification, experience with the biological and chemical weapons conventions suggests that effective use of a nuclear fire-wall would require the following:

- Mechanisms in place to deal with concerns short of outright noncompliance, probably based on a mandate from the IAEA Board of Governors, UN Security Council, or a regional organization. Evidence of suspicious activity should not only trigger greater clarification or investigation measures for that activity but also broaden the information search to include other possible related activities. Algorithms can be used for directing and informing this search.
- The cooperation of the inspected state party, as well as the availability of well-established procedures and parameters governing the investigation of concerns, are necessary conditions for peaceful resolutions to challenging situations.
- Challenge inspections and special investigations, to the extent that they could be feasibly pursued, must be technically and procedurally proficient and sustainable to be credible. On-site investigations require whole-system integration of technical and nontechnical aspects to be effective. Equipment needs to be technically sound and reliable.
- The 5Cs approach discussed earlier would considerably enhance the chances of timely detection of a proliferation concern and, at the same time, would significantly enhance the confidence in a conclusion that no proliferation-related concerns exist in a given state.

Given that the firewall's assessment methodology has comprehensive, transparent, and generic characteristics, it could augment effective communication between national stakeholders and help identify diplomatic approaches and tools (export controls, safeguards) to build confidence and reassurance about activities of concern. As such, the firewall could prove useful for structuring a dialogue between relevant states aiming to defuse any proliferation crisis diplomatically and expeditiously. The country-neutral, normative aspect of the firewall should create an expectation that when it flags a proliferation concern, the state in question should constructively engage the P5, the IAEA, or other appropriate states or bodies to provide reassurances about its purposes and resolve all such concerns without delay. The firewall, in concept and detail, can provide a technically and historically vetted, country-neutral basis for such dialogue, avoiding controversy generated by selective, case-by-case approaches.

Should such dialogue prove inconclusive while the matter(s) of proliferation concern are not resolved, the firewall could ease the challenge of referring the matter to the UN Security Council. Council members could then draw on the firewall framework to address the possible proliferation concern at hand. Based on past experience, this could include enhancing and expanding the IAEA's authority to investigate the matter and keep the council abreast of its findings. The UN Security Council could also authorize use of other tools to selectively and fairly address the concern. Application of the firewall in such ways should in turn enhance the prospects of deterring non-nuclear-weapon states from

coming close to the nuclear threshold while facilitating, in fact encouraging, constructive engagement over matters of concern when they arise.

Constructive or Normative Applications

To support all three pillars of the NPT, the firewall is designed to provide a persuasive basis for defining activities that should be considered inconsistent with peaceful and nonproscribed military applications. This delineation serves to enhance the warning of proliferation risks; guide steps to reassure the international community that current and future programs are peaceful; and define possible future disarmament arrangements at the national, regional, and global levels.

Template for Enhancing the Credibility of Peaceful Nuclear Energy Programs and for Strengthening the Buffer Between Peaceful and Weapons-Oriented Activities

The IAEA serves as the international clearinghouse for helping member states pursue peaceful nuclear activities and providing guidelines on how to do so. The firewall could help maintain international support of and confidence in IAEA efforts, particularly those that provide member states aid and best practices related to peaceful nuclear programs, including as delineated in the IAEA's document, *Milestones in the Development of a National Infrastructure for Nuclear Power*.⁸

By enhancing the credibility of peaceful nuclear energy programs, the firewall could reduce ambiguities and disputes in interpreting rights under the NPT. This could help make nuclear cooperation more predictable, which in turn would enhance the credibility, sustainability, and economic prospects of nuclear energy. Insights from the firewall could inform state export control policies related to nuclear and non-nuclear technologies and facilities and help identify measures that potential recipients could undertake to reinforce the peaceful credentials of a nuclear program and obviate proliferation concerns.

We do not underestimate the controversy that arises today from any effort to refine and apply criteria for managing access to nuclear-related technology. The purpose here is to provide a transparent framework that states could utilize to begin building a constructive or normative template for a peaceful nuclear program, including measures that could be taken to reassure others that dual-use technologies and activities will not lead to the development of nuclear weapons. Such a template would help clarify ambiguities in Articles II and IV of the NPT. The treaty does not define what constitutes the manufacturing and acquisition of nuclear weapons. Nor does it enumerate specific technologies, materials, or activities to which states have inalienable rights and/or reasonable expectations of international cooperation to obtain. (We later discuss how the firewall could give practical meaning to Article VI of the NPT and help define parameters of nuclear disarmament.)

Recognizing that states will not consider how they could begin building such a norm unless and until they conclude when it would be desirable and feasible, this report limits itself to substantive considerations, not procedural or political ones. The core principle of such a template would be to single out those (relatively few) individual activities and clusters thereof that should be deemed illegitimate—namely where a reasonable presumption can be made that they serve a nuclear weapons program. Such activities would be either precluded categorically or subject to certain conditions, without which they would be impermissible in all non-nuclear-weapon states.

Utilizing the coloring scheme depicted earlier (and elaborated in figure 4), the constructive or normative template would highlight red capabilities and activities, which are uniquely or strongly associated with nuclear weapons programs or which, in combination, are more plausibly suited for nuclear weapons development than any other purpose. That is, these activities should generally be regarded as inconsistent with peaceful nuclear programs—certainly unless and until reassurance is given, on a case-by-case basis, to ease the concerns associated with those activities. This would de facto make them yellow activities. Dual-use (yellow) activities should only be carried out after a state has taken the actions necessary to reassure observers of its peaceful objectives for the aforementioned work.

Figure 4: Interpreting the Firewall Color Classifications

The firewall uses color to evaluate and classify individual activities or a program, as well as the overall output from the firewall analysis. The color scheme has both descriptive (informing proliferation analysis) and prescriptive (suggesting possible reassurance) connotations.

Green: Activities or capabilities that are peaceful, or a program whose capabilities and activities are for peaceful purposes. Programs that register green would offer opportunities for further international nuclear cooperation or transfers, but they should still be monitored.

Yellow: Activities or capabilities that could be used for dual purposes (peaceful and/or military), or a program with elements that could be indicative of nuclear weapons development. Programs that register yellow would elicit heightened vigilance and concern and could prompt increased measures toward clarification and assurance.

Red: Activities or capabilities that are exclusively weapons-oriented, or a program whose collective attributes are uniquely associated with a nuclear weapons program. Programs that register red would elicit heightened concern about proliferation and should prompt a necessary investigation into activities and the consideration of possible policy responses.

Similarly, guidelines would define those numerous activities that reasonably can be presumed, by default, to serve one or more legitimate purposes, hence deemed green. Activities or capabilities registering as green could be considered peaceful. This implies they may be undertaken without additional restriction on access to peaceful nuclear technology, and indeed perhaps with some enhanced cooperation and assistance, provided there remains an absence

of other indications that would cast doubt about their true purpose. Naturally, there are a range of potential mechanisms—from safeguards to export controls and confidence-building steps—that could be used selectively for green and yellow activities to reassure the international community that they are and will remain peaceful.

This template for distinguishing purely peaceful nuclear programs from illegitimate military ones would be applied according to two core principles. First, the firewall concept would apply to all NPT non-nuclear-weapon states. Second, a distinction ought to be made between applying the guidelines to future systems and facilities on the one hand and their application to legacy systems and facilities on the other. Guidelines could more readily be applied to the former, while legacy systems and facilities would be phased out over a longer period (including in P5 states). Unavoidably, some states currently undertake activities that, in the future, could be inconsistent with the template for a peaceful program. As with the application of new rules or norms in other fields, legacy activities that would not be acceptable under new or amended rules should ultimately be phased out in a timely manner.

The following sections illustrate, for heuristic purposes, parameters and characteristics of fuel-cycle and reactor operations and of military R&D and operations that should be precluded or limited in states with purely peaceful nuclear programs.

Fuel-Cycle and Reactor Operations

In addition to those fuel-cycle activities undertaken by the P5, such activities may presently take place in states with NPPs in operation or under construction and in states possessing no more than research reactors. States without NPPs should have only green activities, with the exception of some basic research (yellow). States with NPPs could undertake some yellow activities (for example, reactor design, fuel-cycle facilities), and there could be some transition time for phasing out existing reactors or facilities (for example, Purex reprocessing or use of heavy water reactors with continuous online reloading).

In line with this logic, the firewall project envisages the following approaches to handling all new sensitive fuel-cycle activities by non-nuclear-weapon states:

- Commercial activity involving enriching uranium or processing plutonium would be conditioned on demonstrating (to the supplying state[s]) a valid commercial and technical rationale for such activity—specifically, established requirements for fresh fuel to sustain a nuclear power program of specified and planned capacities or a plan for managing nuclear waste and handling irradiated fuel.
- Once undertaken, such commercial-scale activities would be further subject to the provision that they not yield materials in forms or quantities that are easily usable for a weapons program (or tempting targets of theft).

Specific measures to that effect could include no meaningful accumulation of enriched uranium other than in fuel rods or as part of international commercial activity and no production or accumulation of pure plutonium metal. Nor should such material be a product of any interim stage of the recycling process. No conversion of uranium or plutonium oxides to metallic form should be undertaken, and use of metallic fuel in research reactors should henceforth be phased out. Research and development activity involving plutonium or uranium would be limited to a single declared site. Furthermore, these activities would be quantitatively capped at 1 gram for plutonium and 10 grams for enriched uranium.

- Additionally, no enrichment of uranium for any purpose would be permissible above 20 percent other than very small production solely for the purposes of irradiation targets intended to produce isotopes for medical applications.
- The aforementioned regime would apply to all forms of enrichment and processing of plutonium regardless of the technology and facilities employed.

Some enrichment of uranium is and will remain necessary to fuel most power reactors and all research reactors and to facilitate the production of targets for medical isotope production. However, HEU production presently constitutes the easiest proliferation route, as it basically only requires possession of natural uranium and an enrichment facility and can be done independently of any civilian application of nuclear power. With relatively little requirement for major infrastructure, this means of production constitutes a relatively easy path by which a state might clandestinely seek to acquire fissile material. Converting part of a commercial enrichment facility to produce HEU would require only a modest effort, while the covert construction and operation of an enrichment facility would present a more difficult detection challenge than the production of plutonium. Furthermore, handling enriched uranium does not require heavy protection and formidable safety measures. Thus, covert production and/or diversion of HEU could provide a simple way to develop a nuclear weapon or impede the firewall with lower probability of detection.

With regard to HEU fuel for research reactors and for naval propulsion (or land-based prototypes of these reactors), certain technical solutions allowing for the use of 20 percent or less enriched fuel already exist or are under development (for both high-performance research reactors and naval propulsion). These might eventually allow the phasing out of HEU fuel for these reactors. The same could also apply for high-temperature reactors. This would eliminate all requirements for HEU other than for irradiation targets (medical applications), which in turn warrants production of only small quantities of HEU. Moving toward the future elimination of the use of HEU for civilian and military purposes might thus both prove feasible (assuming an adequate transition period and exception for legacy systems) and expedient for mitigating the

risk of diversion, thereby defining a much stronger firewall than the one that exists to date.

Processing of plutonium is far less commonly practiced today for any purpose and also requires reactors and equipment and tools for handling plutonium. Because some states might choose reprocessing for various reasons (for example, to make most efficient use of uranium, reduce waste volume, or address long-term disposal of spent fuel), a total ban is neither presently possible nor *prime facie* desirable—at least with respect to legacy programs. The challenge rests with the limitation of any further expansion of plutonium processing activity by all states until such processing can be modified so as not to yield (at any stage) material for nuclear weapons and until alternative means of disposal can be made workable at an industrial scale. (For the foreseeable future, the need for this requirement would be in non-nuclear-weapon states.)

One option worth considering would be to condition the acceptability of any fuel-cycle activity on transparent implementation of the above-mentioned fundamental nuclear benchmarks for credible nuclear energy programs (in the domains of safety, security, environment, liability, and transparency). That is, states that do not subscribe to and implement the pertinent international rules and norms in these areas should be considered unsuited to conduct activities that result in the acquisition of weapons-usable nuclear materials. In addition, further consideration should be given to the desirability of safeguards-by-design, especially for bulk handling facilities, and the incorporation of proliferation-resistant features in sensitive facilities.

Weaponization, Delivery Systems, Systems Integration, and Militarization

Unlike the fuel cycle, around which much of the IAEA safeguards and export control regimes have been constructed, there are few governance mechanisms (indeed norms) that exist for the domains of weapons development, systems integration, delivery vehicles, and militarization (as related to nuclear weapons programs). Therefore, there is no basis for allowing non-nuclear-weapon states to conduct activities oriented toward the acquisition and maintenance of nuclear weapons. Activity germane to weaponization should be constrained or conditioned in ways that reassure and demonstrate they will not lead to the production of nuclear weapons.⁹ Operationalizing these requirements is nonetheless easier said than done because of the ambiguity of many activities in these domains and because of the challenge in detecting them if a state sought to pursue them covertly. But in the final analysis, we must consider pursuit of activities in several areas—some by themselves and others in combination (for example, neutronic and hydrodynamic codes)—to be red and hence out of bounds (certain details withheld).

- Milling of plutonium or uranium shells, spheres, or hemispheres.
- Neutron generators (especially neutron initiators).

- Tritium technology (especially production), apart from heavy water detritiation or as part of a credible commercial operation or an international fusion technology program.
- Hydrodynamic codes (especially when linked with neutronic codes).
- Hydrodynamic experiments, especially with a focus on implosive tests involving heavy metals.
- Preparations for conducting an underground or underwater nuclear explosion (or a test of a device with fissile material components substituted for inert materials).
- Modification of a delivery vehicle to carry a nuclear weapon payload.
- Development of a high-velocity reentry vehicle.
- Weaponization.

Activities considered yellow include, but are not limited to, the following:

- Design, testing, and manufacturing of non-nuclear components and weapon integration.
- Design, construction, and operation of specialized production facilities (for criticality safety) for fissile material and high explosives.

Possible Implications Beyond Non-Nuclear-Weapon States

Much of the paper has focused on firewall applications pertaining to non-nuclear-weapon states and their programs. The template devised for building over time a greater buffer between peaceful and military-oriented activities is most obviously applicable to these states. But the five recognized nuclear weapon states (the P5) could also apply many of the template's guidelines to strengthen the distinction between peaceful and military nuclear programs in practice and underscore the importance of this distinction while building toward a universal norm. (Naturally, some of this logic could also be applicable to India, Pakistan, and Israel, as well as the Democratic People's Republic of Korea, notwithstanding the particular circumstances of each state.)

Distinguishing peaceful from military nuclear programs would serve the strong nonproliferation interest of the P5 without in any way undermining the NPT-furnished basis of their nuclear weapons activities. The immediate benefit of movement by P5 states in this direction would be to underscore their commitment to nonproliferation, by making demands not only of others but also of themselves. The firewall framework could guide how the P5 might wish to conduct future peaceful and nonproscribed nuclear activities and, over an extended period of time, replace (or phase out) legacy systems that do not meet this constructive or normative template and encourage others to follow suit.

Some ideas on how the P5 could buffer their civilian programs from their military ones are self-evident. For example, in the domain of weapons-usable

fissile materials, nuclear weapon states could phase out the use of HEU in non-weapons applications such as naval propulsion. But we believe the benefits of movement in this direction greatly exceed its direct nonproliferation gains, insofar as this could also be beneficial in demonstrating that the P5 states have taken important, albeit measured and paced, steps in illustrating what eventual disarmament could require.

A Guide for Disarmament Agreements

A framework for distinguishing purely peaceful nuclear activities and programs from nuclear weapons programs can clearly augment nonproliferation. In a complementary fashion, the firewall can inform the design of nuclear disarmament regimes. Such disarmament could occur within an individual state or as part of a bilateral, multilateral, regional, or global arrangement (for example, as envisioned in the NPT process). As a general principle, activities that alone and/or in combination elicit warning that nuclear weapons are being pursued should not appear in states that have completed a nuclear disarmament process. Of course, real-world exigencies likely would dictate that any disarmament process—national, regional, or global—would proceed in phases. Complete nuclear disarmament would be an end point of successive disarming moves over time. The firewall, or something like it, could inform the prioritization of de-escalatory and disarming steps, as well as the ultimate desired end state.

The initial firewall methodology seeks to warn of efforts to acquire a nuclear explosive device. In a disarmament context, the focus is on advanced arsenals; thus, the methodology would prioritize activities that contribute to vertical proliferation. It is obvious that this would include in the first instance cessation of fissile material and delivery vehicle development, testing and production, nuclear weapons testing (or simulation thereof), and military deployment and preparation for use. But the firewall could helpfully address the more difficult questions of how to prioritize, sequence, and identify rollback endpoints for activities across the range of indicators described in this paper (in the domains of fuel-cycle, nuclear weapons R&D and weaponization, weapons delivery and systems integration, militarization, and context). For instance, if a given state was to retain a nuclear energy program following disarmament, the firewall could help identify activities beyond only those related to the fuel cycle that the state should agree to forego.

To our knowledge, no similar detailed basis for defining and designing nuclear disarmament has been developed. The UN General Assembly in December 2016 authorized negotiations of a treaty to prohibit nuclear weapons, leading to their complete elimination. This negotiation is to commence in March 2017.¹⁰ The NPT process has, since 1995, prioritized movement toward creating a zone free of weapons of mass destruction in the Middle East. The UN Security Council has mandated the denuclearization of the Democratic

People's Republic of Korea. To realize the objectives of any of these and other efforts to achieve nuclear disarmament, extremely complicated work must be undertaken to define the terms that would make these processes meaningful and verifiable. We believe the firewall framework offers an approach that could constructively inform such efforts.

Beyond Initial Applications

The aforementioned diagnostic and analytic applications, as well as constructive and normative applications, do not exhaust the range of possibilities but demonstrate the potential value of the firewall in strengthening implementation of the NPT. These applications provide a good foundation for further discussion. The applications of greatest potential to address existing challenges in the nuclear order also pose difficult fidelity and sensitivity requirements. As with any attempt to add definition to a regime, there exists a tension between binding requirements and the desire by participating states to maintain space for sovereign decisions. The firewall project does not presuppose how states might seek to strike that balance in this context. Rather, it offers a range of options that may allow for fine-tuning.

Conclusion

This report represents the culmination of a six-year effort by the Nuclear Policy Program of the Carnegie Endowment for International Peace. The goal was to address a core problem in working to strengthen the NPT's three pillars—a lack of clear definitions that would help distinguish legitimate from illegitimate nuclear activity under the terms of the treaty. Informed by consultations with outstanding policy and technical experts, Carnegie has attempted to develop a robust and methodologically consistent approach to ameliorate this problem. This approach centers on a framework and indicators for demarcating activities that should be inhibited because they are purely or strongly associated with nuclear weapons and conversely those that should be facilitated because they are fully consistent with peaceful applications of nuclear technology and know-how. This paper has explored the feasibility of operationalizing such an approach and its utility in a number of potential applications and contexts. Carnegie obviously bears sole responsibility for the outcome.

The firewall project has succeeded in constructing, refining, and vetting the framework through a series of structured, in-depth discussions among a distinguished group of international experts. Based on these discussions and those with officials who contributed to or encouraged the project's efforts and the progress made, we recommend that governments study the ideas contained herein and consider which applications of the firewall might merit further technical development and which might be ripe for implementation in multiple institutional settings. Such efforts may be carried out nationally or

multilaterally, with or without the participation of Carnegie or other nongovernmental experts. Carnegie stands ready to pursue further development of the firewall concept, the methodology, or specific policy applications in parallel or to pursue outreach more broadly should such an approach be deemed most appropriate.

Notes

1. For background on the treaty, see “Treaty on the Non-Proliferation of Nuclear Weapons,” United Nations Office for Disarmament Affairs, <https://www.un.org/disarmament/wmd/nuclear/npt/>.
2. Meaningful is defined here as exceeding 100 grams of enriched uranium or 1 gram of separated plutonium.
3. For instance, John Carlson, “‘Peaceful’ Nuclear Programs and the Problem of Nuclear Latency,” Nuclear Threat Initiative, November 2015, http://www.nti.org/media/pdfs/Peaceful_Nuclear_Programs_and_the_Problem_of_Nuclear_Latency.pdf; and John Carlson, “Scientific, Technical and Industrial Potential as a Precondition for Nuclear Weapons Development,” in *Limits to Secure Nuclear Tolerance*, ed. Viatcheslav Kantor (Moscow: International Luxembourg Forum, 2014), http://www.luxembourgforum.org/media/documents/Book_by_VVK_2014_eng.pdf.
4. See Mark Hibbs, “Iran and the Evolution of Safeguards,” in *Verification and Implementation*, 2015, ed. Larry MacPaul (London: VERTIC, 2015), <http://www.vertic.org/media/assets/Publications/Verification%20and%20Implementation%202015.pdf>.
5. See Z. Liu and S. Morsy, “Development of the Physical Model,” International Atomic Energy Agency (IAEA), 2007, <http://www-pub.iaea.org/MTCD/publications/PDF/ss-2001/PDF%20files/Session%2013/Paper%2013-07.pdf>.
6. See IAEA, *IAEA Safeguards Glossary 2001 Edition* (Vienna: IAEA, 2002).
7. Jez Littlewood, *The Biological Weapons Convention: A Failed Revolution* (Aldershot, UK: Ashgate, 2005); and United Nations Biological Weapons Convention Ad Hoc Group, “Protocol to the Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction,” May 31, 2001, [http://www.unog.ch/80256EDD006B8954/\(httpAssets\)/EAC25C464A88DBB7C12577C8002B4B24/\\$file/BWC_AHG_CRP.08.pdf](http://www.unog.ch/80256EDD006B8954/(httpAssets)/EAC25C464A88DBB7C12577C8002B4B24/$file/BWC_AHG_CRP.08.pdf).
8. See “The IAEA Milestones Approach,” IAEA, November 12, 2015, <https://www.iaea.org/NuclearPower/Infrastructure/milestone/index.html>.
9. In this respect, the recent Joint Comprehensive Plan of Action with Iran, which contains proscriptions of this nature in perpetuity, affirms the utility of such an approach as well as its potential negotiability.
10. United Nations General Assembly, Resolution 71/258, “Taking Forward Multilateral Nuclear Disarmament Negotiations,” December 23, 2016, http://www.un.org/en/ga/search/view_doc.asp?symbol=A/RES/71/258.

Carnegie Endowment for International Peace

The **Carnegie Endowment for International Peace** is a unique global network of policy research centers in Russia, China, Europe, the Middle East, India, and the United States. Our mission, dating back more than a century, is to advance the cause of peace through analysis and development of fresh policy ideas and direct engagement and collaboration with decisionmakers in government, business, and civil society. Working together, our centers bring the inestimable benefit of multiple national viewpoints to bilateral, regional, and global issues.

The **Carnegie Nuclear Policy Program** is an internationally acclaimed source of expertise and policy thinking on nuclear industry, nonproliferation, security, and disarmament. Its multinational staff stays at the forefront of nuclear policy issues in the United States, Russia, China, Northeast Asia, South Asia, and the Middle East.

BEIJING

BEIRUT

BRUSSELS

MOSCOW

NEW DELHI

WASHINGTON

THE GLOBAL THINK TANK



CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

CarnegieEndowment.org