


U C B E R K E L E Y
C E N T E R F O R L O N G - T E R M C Y B E R S E C U R I T Y



QUICK GUIDE

An Introductory Resource for the General-Purpose AI Risk Management Standards Profile V1.2

NADA MADKOUR | JESSICA NEWMAN | DEEPIKA RAMAN | KRISTAL JACKSON
EVAN R. MURPHY | CHARLOTTE YUAN

Cover art: The cover image is an adaptation of a photograph titled, “Steam Engine near the Grand Transept, Crystal Palace,” taken by the photographer Philip Henry Delamotte in 1851. The impact of artificial intelligence and especially general purpose artificial intelligence is often compared to the impact of the steam engine during the Industrial Revolution, which brought enormous economic gains, but also dangerous workplaces and horrible living conditions for many. The Crystal Palace housed the Great Exhibition of 1851, where examples of technology developed in the Industrial Revolution were put on display for thousands of people to see. While enjoyed by many, the Crystal Palace was also critiqued for representing a false utopia. Similarly, the rise of general purpose AI is often discussed with utopian visions, but such positive visions are often overpromised and will not be possible without the establishment of meaningful risk management strategies. The image is a reminder of the entanglement of people and machines, and the profound and lasting impact of general purpose technologies on society.

In this adaptation, the updated golden palette alludes to contemporary narratives of an “AI gold rush,” reflecting the rapid investment, aspiration, and momentum surrounding AI development. The radiant gold machinery draws the viewer’s eye and underscores how technological systems increasingly occupy the locus of attention within public and policy discourse, often overshadowing the human figure within the frame. Against this backdrop of acceleration and possibility, we present the second annual update to the AI Risk-Management Standards Profile for General-Purpose AI Systems (GPAIS) and Foundation Models (Version 1.2).

QUICK GUIDE

An Introductory Resource for the General-Purpose AI Risk Management Standards Profile V1.2

NADA MADKOUR[†] • JESSICA NEWMAN[†] • DEEPIKA RAMAN[†] • KRYSTAL JACKSON[†]
EVAN R. MURPHY[†] • CHARLOTTE YUAN[†]

[†] AI Security Initiative, Center for Long-Term Cybersecurity, UC Berkeley

All affiliations listed are either current, or were during main contributions to this work or a previous version.

April 2026

The Quick Guide is a short and actionable introductory resource designed to complement the full

General Purpose AI Risk-Management Standards Profile V1.2, available here:

<https://cltc.berkeley.edu/publication/ai-risk-management-standards-profile-v1.2>



Contents

1. HOW TO USE THIS QUICK GUIDE	3
1.1 What's Included	3
1.2 Recommended Instructions	4
2. HIGH-PRIORITY RISK MANAGEMENT STEPS	5
3. GUIDANCE	7
3.1 Govern	7
3.2 Map	8
3.3 Measure	11
3.4 Manage	14

1. How to Use this Quick Guide

This Quick Guide is designed to be a short and actionable introductory resource to complement the full General-Purpose AI Risk-Management Standards Profile (“GPAI Profile”) V1.2,¹ available here: <https://cltc.berkeley.edu/publication/ai-risk-management-standards-profile-v1.2>

Please refer to the full Profile for sources, resources, and significantly more detail on all of the recommendations listed here, as well as other related guidance.

Both the Quick Guide and GPAI Profile are intended to be used in conjunction with the NIST AI Risk Management Framework (AI RMF) and AI RMF Playbook, or with an approximately equivalent set of AI risk management guidance resources.

We provide guidance specifically for general-purpose AI (GPAI), across the four AI RMF “core functions,” including: Govern, Map, Measure, and Manage.

We expect this guide to be primarily helpful for developers of large-scale, state-of-the-art GPAI models. Others who can benefit from use of this guide include downstream developers of end-use applications or AI systems that build on a GPAI model, as well as model evaluators and regulators.

1.1 WHAT’S INCLUDED

The quick guide was developed to be a condensed version of the full Profile, and includes:

- **Highest-priority risk management steps** to be regarded as baseline expectations. These are the same as those outlined in the full Profile. Further detail, including inclusion criteria for high-priority risk management steps, can be found in section 2.3 of the full V1.2 Profile.
- **Priority guidance in condensed form** for each of the core RMF functions, with “high-priority risk management steps” highlighted for ease of reference.

¹ The Profile V1.2 provides guidelines for GPAI model developers based on “core functions” defined in the NIST AI RMF: “Govern,” for AI risk management process policies, roles, and responsibilities; “Map,” for identifying AI risks in context; “Measure,” for rating AI trustworthiness characteristics; and “Manage,” for decisions on prioritizing, avoiding, mitigating, or accepting AI risks.

1.2 RECOMMENDED INSTRUCTIONS

- **For a quick scan** of the risk management steps that should be considered as the baseline or minimum expectations for managing GPAI risks, read section 2, “High-Priority Risk Management Steps.”
- **For a more in-depth introduction** to “high-priority risk management steps,” as well as other unique perspectives and recommendations from across the AI RMF subcategories, read the shortened high-level guidance in section 3. You may use the Quick Guide to take note of particularly relevant subsections to refer to when reading the full GPAI Profile. Alternatively, you may go through the full GPAI Profile in parallel to reading the summarized guidance provided in section 4 of this document.
 - » High-priority subsections² are highlighted; we encourage you to prioritize them while applying the quick guide and the full Profile.

The Quick Guide is meant to introduce readers to the Profile V1.2 contents and aid in Profile V1.2 navigation. We do not recommend using the quick guide as a replacement for the full Profile.

2 All the subsections categorized as “high-priority” in the Quick Guide are considered high-priority in the full GPAI Profile.

2. High-Priority Risk Management Steps

We recommend giving the highest priority to the following risk management steps and corresponding GPAI Profile guidance sections.³ These steps should be regarded as the **baseline or minimum expectations** for users of the GPAI Profile, who can exceed the minimum expectations by also applying guidance from other sections.

- **Check or update, and incorporate, each of the following high-priority risk-management steps when making go/no-go decisions**, especially on whether to proceed on major stages or investments for development or deployment of cutting-edge large-scale GPAI models (Manage 1.1).
- **Take responsibility for risk assessment and risk-management tasks for which your organization is well positioned to develop sufficient constructive action, particularly when organizational actions are substantially more effective than those taken by others in the value chain** (Govern 2.1).
 - » We also recommend applying this principle throughout other risk assessment and risk-management steps, and we refer to it frequently in other guidance sections.
- **Collect, consider, and integrate feedback from external stakeholders and third-party evaluators** (Govern 5.1).
 - » Prioritize partnerships with independent third-party evaluators and auditors throughout the model lifecycle, particularly when assessing the model for dangerous capabilities, and establish clear and accessible feedback channels for users or impacted persons.
- **Set risk-tolerance thresholds to prevent unacceptable risks** (Map 1.5).
 - » For example, the NIST AI RMF 1.0 recommends the following: “In cases where an AI system presents unacceptable negative risk levels — such as where significant negative impacts are imminent, severe harms are actually occurring, or catastrophic risks are present — development and deployment should cease in a safe manner until risks can be sufficiently managed” (NIST 2023a, p.8).

³ It may be appropriate to follow the guidance in this document for these risk management steps, but to apply and document them under other, closely related risk-management steps (typically noted in this document with “see also” statements pointing to guidance in other sections of the GPAI Profile). For example, if your organization sets risk-tolerance thresholds under Govern 1.3 instead of under Map 1.5, then as part of your organization’s process for Govern 1.3, it can be appropriate to follow guidance in the GPAI Profile under Map 1.5.

- **Identify reasonably foreseeable uses, misuses, and abuses for a GPAI model** (e.g., automated generation of toxic or illegal content or disinformation, or aiding with proliferation of cyber, chemical, biological, radiological, or nuclear weapons), and identify reasonably foreseeable potential impacts (e.g., to fundamental rights) (Map 1.1).
- **Identify whether a GPAI model could lead to significant, severe, or catastrophic impacts**, e.g., because of correlated failures or errors across high-stakes deployment domains, dangerous emergent behaviors or vulnerabilities, or harmful misuses and abuses (Map 5.1).
- **Use red teams and adversarial testing** as part of extensive interaction with GPAI models to identify dangerous capabilities, vulnerabilities, or other emergent properties of such systems (Measure 1.1).
- **Track important identified risks** (e.g., vulnerabilities from data poisoning and other attacks or objectives misspecification), even if they cannot yet be measured (Measure 1.1 and Measure 3.2).
- **Implement risk-reduction controls as appropriate** throughout a GPAI model lifecycle, e.g., independent auditing, incremental scale-up, red-teaming, structured access or staged release, and other steps (Manage 1.3, Manage 2.3, and Manage 2.4).
- **Incorporate identified AI system risk factors — and circumstances that could result in impacts or harms — into reporting and engagement with internal and external stakeholders** (e.g., when reporting to downstream developers, regulators, users, impacted communities, etc.) on the AI system as appropriate, e.g., using model cards, system cards, and other transparency mechanisms (Govern 4.2).
- **Continuously monitor the model post-deployment** by actively gathering relevant information for risk evaluation, and by implementing an incident response plan (Manage 4.1).

3. Guidance

The high-level guidance provided below includes the high-priority risk management steps described in the previous section, as well as other unique perspectives and recommendations from across the AI RMF subcategories that may additionally be prioritized by organizations. The high-priority risk management steps are highlighted for ease of reference.

3.1 GOVERN

Table 2 outlines priority guidance provided in the GPAI Profile under the “Govern” core function.

Note: These have been condensed for inclusion in the quick guide. Please refer to the full GPAI Profile for greater detail and context, including sources and resources. Guidance that is one of the “high-priority risk management steps” is highlighted for ease of reference.

Table 1: Quick Guidance for NIST AI RMF Govern

Guidance	Subcategory
Establish relevant internal (e.g., incident response plans) and public-facing policies (e.g., acceptable use policies, data use and retention policies) and make them available to the appropriate stakeholders (e.g., oversight bodies, third-party evaluators, or end users).	Govern 1.4
Identify GPAI model impacts (including to human rights) and risks (including potential uses, misuses, and abuses), starting from an early AI lifecycle stage and repeating frequently through new lifecycle phases, or as new information becomes available.	Govern 1.5
Revisit use and misuse case identification and review risk management process at key intended milestones, or at periodic intervals (e.g., at least annually), whichever comes first.	Govern 1.5
GPAI model developers should take responsibility for risk assessment and risk management in areas where they have—or can reasonably expect to have—better access to information, capabilities, or opportunities to take meaningful action compared to others in the value chain.	Govern 2.1
Make as much information as reasonably possible transparent and available to all audiences on AI risk factors, incidents (including near-miss incidents), knowledge limits, etc..	Govern 2.1
As part of staffing to identify potential high-impact scenarios for GPAI models, broaden the team as appropriate to include, for example, social scientists and historians who can provide additional perspective on structural or systemic risks that could emerge from interactions between an AI system and other societal-level systems.	Govern 3.1
Incorporate identified AI system risk factors , and circumstances that could result in impacts or harms, into engagement with internal and external stakeholders.	Govern 4.2

QUICK GUIDE: AN INTRODUCTORY RESOURCE FOR THE
GENERAL-PURPOSE AI RISK MANAGEMENT STANDARDS PROFILE V1.2

Guidance	Subcategory
Report identified AI system risk factors , and circumstances that could result in impacts or harms to individuals (including impacts to health, safety, well-being, or fundamental rights), to organizations, groups or communities, and society, including risks to critical infrastructure, economic or national security, democratic institutions, or the environment. Also report additional identified factors that could lead to severe or catastrophic consequences for society, such as the potential for correlated robustness failures or other systemic risks across high-stakes application domains, potential for correlated bias across a large fraction of a society’s population, potential for many high-impact uses or misuses beyond an originally intended use case, or other factors.	Govern 4.2
Communicate potential impacts of AI models (e.g., model risks, impacts, incidents, and near misses) more broadly to the public and governing bodies.	Govern 4.2
GPAI model developers and deployers should integrate independent feedback from those external to the team that develops or deploys a model. Types of external feedback that should be utilized where appropriate include: deliberation with impacted communities, including people involved with the human labor and training of GPAI models (such as data annotators and content reviewers), people whose work is “scraped” for training purposes (such as artists and authors), intended users of the model, and people whose livelihoods may be altered by the use of the system. Other types of feedback that should also be gathered include independent auditing and external evaluation throughout the AI lifecycle, including through bug bounty and bias bounty programs, red-teaming, and collaboration with external industry and academic experts. It is also important to incorporate feedback from users or impacted individuals or communities, including through appeal and redress mechanisms.	Govern 5.1

3.2 MAP

Table 3 details priority guidance provided in the GPAI Profile under the “Map” core function.

Note: These have been condensed for inclusion in the Quick Guide. Please refer to the full GPAI Profile for greater detail and context, including sources and resources. High-priority risk management steps are highlighted for ease of reference.

Table 2: Quick Guidance for NIST AI RMF Map

Guidance	Subcategory
Identify reasonably foreseeable uses, misuses, or abuses for a GPAI model, beyond any originally intended use cases (or in the absence of a specific intended purpose). These may include: <ul style="list-style-type: none"> • Automated generation of disinformation or misinformation at scale; • Automated generation of phishing-attack material; • Aiding with proliferation of chemical, biological, or radiological weapons, or other weapons of mass destruction. This can include aiding in lab experiment design and troubleshooting, providing instructions on chemical and biological material acquisition, and the capability to “upskill” threat actors by advising on attack techniques; • Discovery and exploitation of software vulnerabilities, cyber attack plan critiquing and assistance, and malware and virus creation or modification, including viruses that evolve over time to evade detection; 	Map 1.1

QUICK GUIDE: AN INTRODUCTORY RESOURCE FOR THE
GENERAL-PURPOSE AI RISK MANAGEMENT STANDARDS PROFILE V1.2

Guidance	Subcategory
<ul style="list-style-type: none"> • Creation of violent, illegal, discriminatory, or harmful content, including non-consensual intimate imagery (NCII) or child sexual abuse material; and • Creation of realistic deceptive content such as deepfakes and voice cloning. 	
<p>Identify the goals and limitations of the data collection and curation processes, and the implications for the resulting AI systems. Consider implementing practices such as running data audits, adopting and aggregating datasets from multiple sources, and evaluating data curation through data quality dimensions.</p>	Map 1.1
<p>Identify reasonably foreseeable potential impacts of GPAI models, which can include but are not limited to:</p> <ul style="list-style-type: none"> • Impacts to organizational operations, including potential loss of understanding or control over particular functions, or loss of trust; • Impacts to organizational assets, including legal compliance costs arising from problems created for individuals; • Impacts to individuals, including impacts to health, reputation, safety, well-being, or fundamental rights; • Impacts to groups, including populations vulnerable to disproportionate adverse impacts or harms; • Impacts to the labor market and economic opportunities; • Impacts to society, including: <ul style="list-style-type: none"> ◦ Damage to or incapacitation of a critical infrastructure sector; ◦ Economic and national security; ◦ Concentration and control of the power and benefits from AI technologies; ◦ Democratic institutions and quality of life; ◦ Impacts of anthropomorphic design risks (e.g., deception and manipulation) and user over-reliance; ◦ Public perception and mistrust in institutions; ◦ Polarization and extremism; and ◦ Environmental harms. • Impacts of open-weight models on risks (e.g., enable malicious misuse, amplify model limitations). 	Map 1.1
<p>Identify potential and actual human rights impacts. Ensure alignment with the Universal Declaration of Human Rights (UDHR), including:</p> <ul style="list-style-type: none"> • The right to non-discrimination and equality before the law (Article 2); • The right to life and personal security (Article 3); • The right to privacy and protection against unlawful governmental surveillance (Article 12); • The rights to freedom of thought, conscience, and religious belief and practice; freedom of expression; and freedom to hold opinions without interference (Articles 18 and 19); • The rights to freedom of association and the right to peaceful assembly (Articles 20 and 21); and • The rights to decent work and to an adequate standard of living (Article 23 and 25). 	Map 1.1
<p>Consider the following questions for an AI system: “What objective has been specified for the system, and what kinds of perverse behavior could be incentivized by optimizing for that objective?” Examples of AI systems with misspecified objectives include machine-learning algorithms for social media that learn to optimize engagement metrics by serving users with extremist content or disinformation.</p>	Map 1.3
<p>Set policies on unacceptable-risk thresholds for GPAI model development and deployment to include prevention of risks with substantial probability of significant, severe, or catastrophic outcomes if inadequately mitigated. Unacceptable-risk thresholds can be based on quantitative metrics, qualitative characteristics, or a combination of both. They should be informed not only by the risk tolerance of the organization in question, but also by broadly recognized notions of unacceptable risks to users and impacted communities, society, and the planet. For GPAI models with potential for unknown emergent properties, include a “margin of safety” or buffer between the worst plausible system failures and the unacceptable-risk thresholds. Similar approaches are common for safety engineering in other fields.</p>	Map 1.5

QUICK GUIDE: AN INTRODUCTORY RESOURCE FOR THE
GENERAL-PURPOSE AI RISK MANAGEMENT STANDARDS PROFILE V1.2

Guidance	Subcategory
Define risk tolerances for a minimum set of high-severity risk categories (e.g., offensive cyber operations, loss of control, CBRN weapons, persuasion and manipulation, deception) unless developers can justify why their system could not reasonably pose risks in one of those areas.	Map 1.5
Fully scoping and understanding knowledge limits —that is, the boundaries of the model’s knowledge, reliability of its application, and the likelihood of errors—is difficult for increasingly general-purpose AI systems. However, given the large number of potential uses of these AI systems, clear documentation and communication of their knowledge limits is very important. GPAI model developers should describe or list (and provide examples of) uses that would exceed a system’s knowledge limits, as well as uses that would be appropriate given the system’s knowledge limits. This information should be clearly featured in system documentation for downstream developers, users, and others as appropriate.	Map 2.2
Implement evaluations to detect alignment faking and sandbagging. Test if models can distinguish between evaluation and deployment contexts (situational awareness) and if they alter their behavior to appear aligned during testing.	Map 2.3
Prioritization of GPAI model risks should include consideration of the magnitude of potential impacts, not just their likelihood. This is particularly important for any potential impacts with irreversible effects and catastrophic magnitude. Such impacts can be more likely for GPAI models than for many other types of AI, because GPAI models are often more likely to have relatively greater capabilities, scale of deployment, and other factors leading to high impact.	Map 5.1
When evaluating the likelihood and magnitude of identified risks, utilize systemic-risk scenarios and risk-modeling approaches to better quantify potential impacts and probabilities, and to reveal hidden dependencies or interactions that could increase overall risk.	Map 5.1
Factors that could lead to significant, severe, or catastrophic harms to individuals, groups, organizations, and society can include: <ul style="list-style-type: none"> • Correlated bias across large numbers of people or a large fraction of a group or society’s population (e.g., resulting in systemic discrimination, exclusion, or violence); • Impacts to societal trust or democratic processes, for example through large-scale manipulation of the information ecosystem; • Correlated robustness failures across multiple high-stakes application domains such as critical infrastructure; • Potential for high-impact misuses and abuses beyond originally intended use cases, including a model being used to assist in code generation for cybersecurity threats or to create or be used as destructive weapons, such as cyberweapons, lethal autonomous weapons, bio-weapons, or other significant military applications; • Potential for large harms from misspecified objectives or misgeneralized goals (e.g., using oversimplified or short-term metrics as proxies for desired longer-term outcomes); • Ability to directly cause physical harms, e.g., via robotics motor control; • Reliability issues (e.g., hallucination and confabulation); • Potential for socioeconomic risk and labor market disruption, including job automation and worsening inequality; • Capability for AI models to manipulate or deceive humans into taking harmful actions in the world; • AI systems that could recursively improve their capabilities by modifying their algorithms or architectures through code generation (e.g., from OpenAI Codex or DeepMind AlphaCode), neural architecture search, etc.; • Adaptive models that might be difficult to control in real time, e.g., in response to coordinated manipulation attacks, such as the attacks on the Microsoft Tay chatbot in 2016; 	Map 5.1

QUICK GUIDE: AN INTRODUCTORY RESOURCE FOR THE
GENERAL-PURPOSE AI RISK MANAGEMENT STANDARDS PROFILE V1.2

Guidance	Subcategory
<ul style="list-style-type: none"> • Agentic systems, i.e., systems that in effect choose or take actions in a goal-directed fashion; • Ability to employ outbound communication and influence channels, such as to post information to the web via HTTP POST requests or functionally equivalent means (e.g., some types of plugins); • Ability to escape a sandbox and replicate on another computational system, either via hacking, social engineering, or using other exploits; • Sandbagging, i.e., strategically underperforming on model evaluations, including but not limited to password-locking or password-unlocking key capabilities; and • Situational awareness, including a system being able to recognize that it is an AI, having knowledge about its capabilities and limitations, and knowing whether it is running in a test or deployment environment, which could, for example, enable a model to learn about the idea of jailbreaks from pre-training and utilize it when being evaluated for safety by a reward model. 	
<p>GPAI model developers should implement mechanisms to support regular engagement with relevant AI actors, given the high likelihood and potential impact of unanticipated negative impacts. Such mechanisms can include support for incident reporting, complaint and redress, independent auditing, and protection for whistleblowers.</p>	Map 5.2

3.3 MEASURE

Table 3 outlines priority guidance provided in the GPAI Profile under the “Govern” core function.

Note: These have been condensed for inclusion in the Quick Guide. Please refer to the full GPAI Profile for greater detail and context, including sources and resources. High-priority risk management steps have been highlighted for ease of reference.

Table 3: Quick Guidance for NIST AI RMF Measure

Guidance	Subcategory
<p>Do not ignore identified risks just because measurement would be difficult, especially if the impacts could be severe or catastrophic. Measurements of identified risks are often more difficult for GPAI models than for smaller-scale or fixed-purpose AI systems, because of factors such as complexities, uncertainties, and emergent properties of GPAI models. For many factors it can be more appropriate to use qualitative assessment procedures. Plan to track and revisit identified risks, even if they cannot be measured quantitatively at this time.</p>	Measure 1.1
<p>Use red teams and adversarial testing as part of extensive interaction with GPAI models to identify dangerous capabilities, vulnerabilities, or other emergent properties. Security vulnerabilities are typically inherent to currently available GPAI models, including vulnerabilities to prompt injection attacks. Red-teaming can identify these weaknesses, though they are currently difficult to protect against. For frontier models, characteristics that red teams should evaluate include: unacceptable-risk factors as outlined in Map 1.5, and high-impact harm factors as outlined in Map 5.1. Partner with one or more independent red-teaming organizations as appropriate, and consider the following:</p>	Measure 1.1

QUICK GUIDE: AN INTRODUCTORY RESOURCE FOR THE
GENERAL-PURPOSE AI RISK MANAGEMENT STANDARDS PROFILE V1.2

Guidance	Subcategory
<ul style="list-style-type: none"> • Protect proprietary or unreleased model weights as appropriate during red-teaming to prevent unauthorized access or leaks of model weights. • Grant red teams considerable independence and control over the scrutiny process. • Grant red teams appropriate access to the final versions of models before deployment. • As part of pre-release red-teaming for models that are planned for release with downloadable, fully open, or open-source access, allow red teamers to appropriately test the extent to which Reinforcement Learning from Human Feedback (RLHF) or other mitigations would not be resilient to additional fine-tuning or other processes used by actors with direct access to a model's weights after open release. • As part of criteria for use of benchmarks or other metrics for risk assessment, and as part of communication of benchmarking results, clarify whether a specific benchmark directly measures a particular risk, such as security vulnerability to prompt injection; whether it indicates a capability that could be misused or abused, such as software code generation; or whether it measures another important aspect of risk. 	
<p>Encourage independent researchers to test models and share their findings by offering bug and bias bounties, and by providing safe harbor for AI evaluation and red-teaming.</p>	Measure 1.3
<p>Use information security measures to assess and assure model weight security (specifically, integrity and confidentiality) as part of preventing misuse or abuse of models. This is particularly valuable for frontier models, for which public release of model weights could enable misuse with particularly high-consequence impacts. Model developers should implement the NIST Cybersecurity Framework (NIST 2024a) — or an approximate equivalent, such as NIST SP 800-171 or ISO/IEC 27001 — with at least the following security controls or approximate equivalents:</p> <ul style="list-style-type: none"> • For frontier models: High-value asset guidance (e.g., per NIST SP 800-171 and NIST SP 800-172), or high-impact system baseline per NIST SP 800-53B as an informative reference for the NIST Cybersecurity Framework, or approximate equivalent. • For other models: Moderate-impact system baseline guidance (e.g., per NIST SP 800-171), or moderate-impact system baseline per NIST SP 800-53B as an informative reference for the NIST Cybersecurity Framework, or approximate equivalent. 	Measure 2.7
<p>Consider the following as part of security evaluations of GPAI models:</p> <ul style="list-style-type: none"> • Perform red-teaming and adversarial testing of security aspects of GPAI models. • Check for backdoors, AI trojans, prompt injection vulnerabilities, and other adversarial attacks during testing and evaluation, especially for models trained on untrusted or publicly sourced data that may be susceptible to data poisoning. Recognize that fully eliminating backdoors and trojans in GPAI models may be infeasible, and evaluate and manage risks accordingly, taking this uncertainty into account. • Engage in continuous monitoring, vulnerability disclosure, and bug bounty programs for GPAI models to identify novel security vulnerabilities. • Track uncovered security vulnerabilities in other GPAI models, including open-source models, that may be transferable to other models. 	Measure 2.7
<p>Document organizational transparency and disclosure mechanisms to inform or allow users to check whether they are interacting with, or observing content created by, a generative AI system.</p>	Measure 2.8

QUICK GUIDE: AN INTRODUCTORY RESOURCE FOR THE
GENERAL-PURPOSE AI RISK MANAGEMENT STANDARDS PROFILE V1.2

Guidance	Subcategory
It is critical to ensure that users know how to interpret system behavior and outputs , including the limitations of both the system and any explanations provided. However, explainability and interpretability are often extremely limited for LLMs and other GPAI models with deep-learning architectures. These systems can be inappropriate for applications requiring better explainability and interpretability.	Measure 2.9
Privacy challenges for GPAI models include the issue that, after pre-training on large quantities of uncurated web-scraped data or other sources containing personally sensitive data, some of that sensitive material in the training data can be revealed by user prompts. Enable people to consent to — and opt out of — the uses of their data. Also notify users and impacted communities about privacy or security breaches.	Measure 2.10
Evaluating for fairness and bias in GPAI models should take into account the complexity of the numerous sources and types of bias that influence GPAI models (including bias from massive datasets as well as decisions made about modeling, optimization, hardware, and testing). Evaluations for fairness and bias should not, for example, focus only on narrow definitions of protected classes, which may overlook complexities of identity.	Measure 2.11
Perform environmental impact assessments. GPAI model developers should include estimating the environmental impact of large-scale ML model training. Assessment of environmental impacts is particularly important for GPAI models, for which model training typically has much larger environmental impacts than smaller-scale ML models.	Measure 2.12
Identify or assess longer-term impacts , or use longer time horizons (longer than would be typical for smaller-scale fixed-purpose AI systems), to reduce potential for surprise, and account for potential novel risks. Consider whether any risk or impact assessment answers would change if assessing longer-term time periods (e.g., beyond the next year) . What additional impacts would you expect? Which impacts would you expect to have greater magnitude? Identify unintended potential future events that should trigger reassessment or other responses, and build them into risk registers and/or planning and implementation of relevant lifecycle stages.	Measure 3.1
Use appropriate mechanisms for tracking identified risks, and complement evaluations with systemic risk modelling and estimation to quantitatively or qualitatively assess probability and severity of harm. Consider tracking identified risks (including difficult-to-assess risks) using a risk register. When developing frontier models with unprecedented capabilities, failure modes, and other emergent properties, it is especially valuable to use red teams and adversarial testing prior to deployment. Risk tracking should include ongoing monitoring of newly identified capabilities and limitations of deployed GPAI models. These efforts can include monitoring use of the models through APIs, and monitoring publications or online forums that discuss new uses of the models.	Measure 3.2

3.4 MANAGE

Table 4 outlines priority guidance provided in the GPAI Profile under the “Govern” core function.

Note: These have been condensed for inclusion in the Quick Guide. Please refer to the full GPAI Profile for greater detail and context, including sources and resources. High-priority risk management steps”have been highlighted for ease of reference.

Table 4: Quick Guidance for NIST AI RMF Manage

Guidance	Subcategory
<p>Consider the following when making go/no-go decisions, especially on whether to proceed on major stages or investments for development or deployment of cutting-edge, large-scale GPAI models:</p> <ul style="list-style-type: none"> • See guidance in this document under Map 1.3 on AI development objectives, especially: consider potential for misspecified AI system objectives, and consider what kinds of misaligned behavior could be incentivized by optimizing for those objectives. • See guidance in this document under Map 1.5 on organizational risk tolerances, especially: set policies on unacceptable-risk thresholds for GPAI model development and deployment to include prevention of risks with substantial probability of inadequately mitigated catastrophic outcomes. 	Manage 1.1
<p>For each identified potential use or misuse (or category of use or misuse) of an AI system, define and communicate to key stakeholders whether any potential use cases (or categories of use cases) would be disallowed or unacceptable, or would be treated as “high risk.”</p>	Manage 1.3
<p>If model training requires obtaining data sets, consider using only trusted training data instead of uncurated scrapes from the web. This can be valuable for multiple objectives, including reducing vulnerability to backdoor and data-poisoning attacks, and reducing unwanted bias and language toxicity.</p>	Manage 1.3
<p>Increase the amount of compute (computing power) spent training frontier models only incrementally (e.g., by not more than three times between each increment) as part of identification and management of risks of emergent properties. Test frontier models after each incremental increase of compute, data, or model size for model training. Probe, monitor, and stress cutting-edge GPAI models using detailed analysis processes (including or extending standard cybersecurity red-team methods) to achieve testing objectives, including testing for unintended toxic and harmful content and/or dangerous errors (e.g., inaccurate medical information), and identifying emergent properties, such as new capabilities and failure modes.</p>	Manage 1.3
<p>Consider approaches to design, testing, and deployment that ensure AI systems possess only the minimum necessary capabilities for high-reliability operation (and not more). Consider methods of implementing the cybersecurity principle of least privilege. For example, consider using or extending typical “deny by default” or whitelisting methods to limit an AI system’s privileges to the minimum necessary for access to information, communication channels, and action space.</p>	Manage 1.3
<p>Determine a strategy to safely and appropriately release the AI system, with consideration of what protections might be necessary to prevent harm or misuse.</p>	Manage 1.3

QUICK GUIDE: AN INTRODUCTORY RESOURCE FOR THE
GENERAL-PURPOSE AI RISK MANAGEMENT STANDARDS PROFILE V1.2

Guidance	Subcategory
When planning for GPAI model deployment, plan on deployment with gradual, phased releases, and/or structured access through an API or other mechanisms, with efforts to detect and respond to misuse or problematic anomalies. Such systems and infrastructure can also be useful for enforcing usage guidelines and minimizing risks of misuse.	Manage 2.4
GPAI model developers that plan to release a model with open weights or open-source access — where that model would be above, at, or near a foundation model frontier — should first use a staged-release approach (e.g., they should not release model parameter weights until after an initial closed-source or structured-access release, where no substantial risks or harms emerge over a sufficient time period with red-teaming and other evaluations as appropriate). Developers also should not proceed to a final step of releasing model parameter weights until a sufficient level of confidence in risk management has been established, including for safety and societal risks and risks of misuse and abuse. Models above a foundation model frontier should be given the greatest amount of duration and depth of pre-release evaluations, as they are most likely to have dangerous capabilities or vulnerabilities, or other properties that can take some time to discover. GPAI model developers that release a GPAI model’s parameter weights, or that suffer a leak of model weights, will in effect be unable to decommission AI systems that others build using those released or leaked GPAI model weights.	Manage 2.4
Prepare emergency shutdown procedures, such as implementing tiered control mechanisms that allow incremental restriction or suspension of model capabilities.	Manage 2.4
As part of ongoing post-deployment monitoring, actively gather and analyze relevant information to evaluate risk levels and determine when an update to a model report update is warranted.	Manage 4.1
Post-deployment monitoring should include clear procedures for incident response, recovery, and decommissioning , such as incident reporting that enables users and third parties to report incidents or vulnerabilities.	Manage 4.1



CLTC

Center for Long-Term
Cybersecurity

UC Berkeley