760 Market Street
Suite 400
San Francisco, CA 94013

September 15, 2021

Unity appreciates the opportunity to share our perspectives on managing risks associated with artificial intelligence ("AI") with the National Institute of Standards and Technology ("NIST") in response to NIST's Request for Information ("RFI") concerning the development of an Artificial Intelligence Risk Management Framework ("AI RMF"). Unity welcomes NIST's commitment to develop the AI RMF consistent with a consensus-driven, open, and collaborative process, including public workshops and other opportunities for stakeholders to provide input. As an organization that utilizes the NIST Cybersecurity and Privacy Frameworks in our internal governance, Unity views the AI RMF initiative as an important step towards strengthening enterprise risk management practices surrounding the design, development, and deployment of AI systems. We look forward to further dialogue with NIST and its stakeholders throughout the AI RMF development process.

About Unity

Unity is the world's leading platform for creating and operating interactive, real-time 3D ("RT3D") content. We believe the world is a better place with more creators in it, and this belief is at the core of our efforts to democratize innovation for all. Our products unlock opportunities for creators to entertain and educate people, and deliver groundbreaking RT3D experiences that add value across industries.

Unity is a comprehensive RT3D platform. Developers can enter the Unity ecosystem at any stage in their process and utilize our platform to advance their unique business goals. Our platform has two central pillars -- our Create and Operate solutions -- with the Unity Runtime Engine being the connective tissue between these pillars.

Our Create Solutions are used to create, edit, run, and deploy real-time 2D, as well as high definition, RT3D content. Create Solutions includes custom scripting tools and a high-definition render pipeline for developers; graphics, animation and audio tools for artists; and navigation, networking and user interface tools for designers. In addition, businesses can leverage our complete, end-to-end Operate Solutions to realize the value of their creations. We provide a complete, full stack solution for development, engagement, and ongoing management that is heavily utilized in mobile gaming applications and an increasing number of industries.

AI plays a critical role in Unity's current offerings and our future growth, as well as the success of developers who utilize our platform. For example, Unity Simulation is a suite of products that harnesses the power of cloud computing to run millions of simulations to test, train and validate AI algorithms at scale. Available on Google Cloud Platform and Amazon Web Services, Unity Simulation lowers the barrier for anyone to train, test, or validate new products and services in a high-fidelity 3D environment, so creators can bring their ideas to the world faster. Initial use

cases for Unity Simulation have focused on computer vision-oriented applications for the automotive, gaming, and robotics industries, however, the product can be used for any project, prototype, or concept created in Unity.

Foundational Principles for the AI RMF

Unity is committed to the democratization of AI by making AI systems widely available to developers and organizations of all sizes. Unity also recognizes broad adoption of AI will depend on the trustworthiness of AI systems, including the underlying principles that guide AI designers and developers. Accordingly, Unity supports NIST's focus on transparency, fairness, and accountability as the foundational principles underlying the AI RMF. These concepts align well with Unity's six guiding AI principles, as described below.

- Be Unbiased. Design AI tools to complement the human experience in a positive way. Consider all types of human experiences in this pursuit. Diversity of perspective will lead to AI complementing experiences for everybody, as opposed to a select few.

- Be Accountable. Consider the potential negative consequences of the AI tools we build. Anticipate what might cause potential direct or indirect harm and engineer to avoid and minimize these problems.

- Be Fair. Do not knowingly develop AI tools and experiences that interfere with normal, functioning democratic systems of government. This means saying no to product development aimed at the suppression of human rights, as defined by the Universal Declaration of Human Rights, such as the right to free expression.

- Be Responsible. Develop products responsibly and do not take advantage of your products' users by manipulating them through AI's vastly more predictive capabilities derived from user data.

- Be Honest. Trust the users of the technology to understand the product's purpose so they can make informed decisions about whether to use the product. Be clear and be transparent.

- Be Trustworthy. Guard the AI derived data as if it were handed to you by your customer directly in trust to only be used as directed under the other principles found in this guide.

The first version of the AI RMF should focus on practical approaches to implementing NIST's proposed foundational principles -- transparency, fairness, and accountability -- as a baseline for organizations to develop and iterate further on principles that address their needs. Unity recognizes the scope and implementation methodologies for AI principles will differ for organizations based on their role in the AI ecosystem, compliance obligations and governance structure, societal commitments and concerns, among other factors. With the AI RMF envisioned as a living document, later versions could provide greater breadth and depth on the development and implementation of AI principles.

We support NIST's intention to consider risks from unintentional, unanticipated, or harmful outcomes that arise from intended uses, secondary uses, and misuses of the AI. Indeed, AI systems designed and developed for one purpose can be adapted to serve different purposes. Organizations benefit from considering not only whether they can create an AI system for a

given purpose, but whether they should develop the system when taking into account different use cases and potential abuse.

Structure of the AI RMF

Unity utilizes both the Privacy and Cybersecurity Frameworks in our operations, and we have found both frameworks to be especially useful as we work to mature and continually improve our risk governance operations. We would encourage NIST to harmonize the AI RMF with NIST's existing frameworks as much as possible, including the inclusion of a maturity model. For example, the AI RMF should emulate the existing frameworks' focus on the maturity of organizational programs through critical analysis documentation, repeatability of processes and feedback loops, and other relevant indicators of maturity,

Attributes of the AI RMF

Unity supports NIST's proposed attributes of the AI RMF, and we encourage NIST to consider its application of these proposed attributes throughout the initial AI RMF development process as well as its ongoing renewal as a living document. Rather than a sequential response, our comments are organized to address NIST's proposed attributes where those attributes reflect complementary goals.

Development of the AI RMF
AI risk management will require multi stakeholder and multilateral engagement over many years to come. We appreciate NIST's commitment to developing the AI RMF through an open and consensus-driven process, similar to the processes utilized in the Cybersecurity and Privacy Frameworks. We encourage NIST to leverage its relationships with national and international standards organizations around the world to raise their awareness about the AI RMF, and to participate in NIST's process where appropriate. This dialogue will be critical as the European Union and other jurisdictions address AI risk management through their own policy development processes.

Characteristics of the AI RMF and its integration with risk management programs
Given that many AI use cases will fall outside of regulated sectors and activities, it is critical that the AI RMF provide guidance that is accessible to organizations that may not have well developed governance practices. We support NIST's goal of creating a risk-based, outcome-focused, voluntary, and non-prescriptive approach in the AI RMF, which will be critical for organizations that do have sector-specific regulators. Likewise, we support NIST's intention that the AI RMF should be usable as part of an organization's broader risk management strategy and processes. Indeed, the AI RMF may serve as the only AI-specific risk management framework for many organizations in the AI ecosystem.

Use of plain language and adaptability across multiple organizations
Broad adoption of the AI RMF will depend on its ability to communicate complex topics using plain language. We support NIST's focus on ensuring that the AI RMF is comprehensible for non-technical stakeholders, and we encourage NIST to build the AI RMF with an eye towards modular use of its components. For example, the AI RMF could distinguish between higher versus lower risk use cases for AI, while at the same time identifying practices that are so fundamental that any organization deploying AI should take them into account.

3

Common definitions and consistency with other approaches to managing AI risk
Developing common definitions for the AI RMF is critical and those definitions should mesh with definitions under consideration in related workstreams, especially definitions that are emerging outside of the United States. We encourage NIST to drive for alignment between definitions provided in the AI RMF with definitions proposed in the [European Commission's proposed AI Act](#) as well as definitions under development at the International Standards Organization.

Maintenance of the AI RMF as a living document
The sophistication of AI systems, and associated risk mitigations, are both advancing rapidly. We support NIST's intention to maintain the AI RMF as a living document, and we encourage NIST to articulate its update schedule. In contrast with the Privacy and Cybersecurity Frameworks, the AI RMF would be NIST's first framework focused on a technology (AI) that is still relatively new and evolving at the speed of innovation. Stakeholders would benefit from the ability to anticipate and prepare for updates to the AI RMF.

Encouraging Workforce Development Through the AI RMF

Unity would support a strong emphasis on workforce development in the AI RMF, including measures that encourage AI providers to offer training and learning opportunities to the current and future workforce. This is driven by our belief that learning should be accessible to everyone. We know there is workforce demand for technical programming skills which is why we contribute our products, technology, and expertise to enable over 300,000 [students](#) and [educators](#) every year to learn how to use Unity's products and services.

The AI RMF should encourage organizations to share their AI expertise as a means of advancing AI skills in the current and future workforce. At Unity, we make two AI courses available to the public online; [Artificial Intelligence for Beginners](#), which has had over 15,000 viewers in 2021 alone, and [Machine Learning (ML) Agents: Hummingbirds](#), which has had nearly 6,000 viewers in 2021. This level of engagement demonstrates there is strong interest in learning about AI and its practical applications, and expanding the breadth and depth of similar offerings across the AI ecosystem can only benefit the workforce.

Conclusion

Unity commends NIST for its initiative to develop a risk management framework that is broadly relevant to organizations throughout the AI ecosystem. We are grateful for the opportunity to share our perspectives, and we welcome dialogue about our input. We look forward to further collaboration with NIST and its stakeholders to develop the AI RMF.