Key Definitions & Concepts

- "**Advisor**" is a specialized professional who from time to time provides expert guidance and support to Olympus on various technological aspects of the business plan but is not as of the date of this presentation an employee of Olympus.
- "**Alignment**" refers to ensuring that an AI system's goals, behaviors, and decision-making processes are consistent with human values, ethical principles, and intended objectives. Different stakeholders may have conflicting values, making alignment difficult. Alignment can be configured to support unethical, racist, or even collectivist goals of stability and uniformity, the development of which must be avoided, particularly as models become more capable and autonomous. **Alignment Faking** refers to when a model appears to align with values but produces unintended harmful outputs under certain conditions. A model might behave as though its preferences have been changed by the training—but might have been faking alignment all along, with its initial, contradictory preferences "locked in". The 2023 Anthropic paper "*Transformers Learn Through Gradient Descent to Simulate a Virtual Machine*" shows that language models can implicitly learn to construct and execute deceptive strategies during training, even without explicit instruction. This has profound risks.
- "**Cluster"** refers to a group of interconnected servers that work together as a single system to provide high availability, scalability, and fault tolerance. Clusters are commonly used for load balancing, parallel computing, and redundancy to ensure that applications and services remain operational even if individual components fail.
- "**Elastic Training"** refers to the ability to dynamically scale computing resources (e.g., GPUs, TPUs, or cloud instances) up or down during the training process based on workload demands.
- **"Generative AI" or genAI** is a type of algorithm that has mastered language across context, sentiment, emotion, grammar, and syntax, seeking to simulate and perfect the learning and decision-making processes of the human brain. They operate by identifying patterns and structures in existing data and using that knowledge to produce novel outputs--most prominently, in word form, but also images, audio, video, or code. These models work by identifying and encoding the patterns and relationships in publicly available data **scraped** from the internet and any other sources that augment that original data set.
- "**Hyperscaler**" is a large-scale technology company that operates massive, global networks of data centers and provides extensive cloud computing services. These companies are characterized by their ability to deliver enormous amounts of computing power, storage capacity, and networking resources to organizations and individuals worldwide. They include Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), IBM Cloud, and Alibaba Cloud.
- "**Industry Partner**" has the meaning used colloquially in the technology sector and not in any legal partnership sense, i.e., an individual who works full time for a vendor, supplier, consultant or other entity in the AI infrastructure ecosystem (and not for Olympus), (i) who has met with the Olympus team on multiple occasions, (ii) whose employer's interests are broadly aligned with the success of the Olympus business plan, and (iii) who has expressed their willingness to be supportive of the Olympus business plan, consistent with their duties to their employer.

Key Definitions & Concepts

- "**Installation**" means the hardware and software that makes possible the Olympus Private Cloud and is used interchangeably with "Olympus Private Cloud" or the local company established or acquired to purchase and own such assets.
- **Inference**" is the process of applying learned patterns to make predictions or decisions about new inputs. It's the mechanical process of using a trained model's closed parameters to generate outputs, rather than using new data. For example, when a language model produces text, or a vision model identifies objects in images, they're performing inference.
- **"Olympus Private Cloud"** means the private cloud made possible by the installed hardware and software in accordance with the Olympus business plan, which enables localized AI capabilities and services.
- **"RAG (Retrieval-Augmented Generation)"** is an AI technique that improves text generation by integrating a retrieval mechanism with a generative model. Instead of relying solely on pre-trained knowledge, RAG dynamically fetches relevant information from external sources (e.g., databases, documents, or the web) and incorporates it into the response.
- **"Safety**" refers to designing, deploying, and monitoring of AI systems to minimize risks to individuals, society, and the environment. These risks include harmful biases, unintended behaviors, misuse by malicious actors, and systemic disruptions. Safety involves ensuring that training of AI systems is designed to align with ethical principles, operate reliably under varying conditions, and remain interpretable and controllable.
- "**Training**" is the learning phase where a model develops its capabilities. During the first phase of training, or "**pre-training**", the algorithm is directed to process a large, broad data set typically "scraped" from the web, research papers, textbooks, etc., from which it learns general language patterns, grammar, syntax, and contextual relationships of words. The model adjusts its internal parameters based on this data in an iterative process to minimize errors and improve its performance. **Fine-tuning** is the second phase of training, when the pre-trained model is refined on a smaller, domain-specific dataset to adapt it to specific tasks. Fine-tuning bridges the gap between general language understanding and task-specific expertise and improves accuracy and relevance in targeted applications. Reinforcement Learning (RL) uses trial-and-error to optimize the model's responses for specific outcomes (such as legal compliance and prohibited behavior), often involving human feedback, to ensure the model aligns with human preferences, values, and expectations.
- **Transfer Learning or knowledge distillation** adapts a model trained on one domain to perform well in another, reducing the need for extensive retraining. OpenAI, Anthropic, Google and others have built extremely expensive closed models trained on massive data sets – such as 100% of the internet, in the case of OpenAI -- approaching limits on performance gains from the use of new data with open parameters. DeepSeek has built an open-source model with knowledge distillation, training their own R1 model to learn the reasoning patterns and outputs of an advanced model, likely ChatGPT4 o1. Distillation allowed the smaller Chinese model to achieve performance levels close to or even exceeding the original large model, introducing a new paradigm to AI model construction and changing the economics of access to AI, especially energy and GPUs.
- "**Use case"** refers to a specific scenario or application where AI is implemented to solve a problem, optimize a process, or provide value to users.