FISD Alternative Data Council: Compliance Considerations for Generative AI

Executive Summary –These materials are intended to address the financial information industry's need for a common understanding of compliance concerns with respect to the use of Generative Artificial Intelligence "GenAI". These Considerations address key terms, risks, governance, controls, vendor diligence and onboarding GenAI models, and contractual GenAI models. The ultimate goals of the Considerations are to promote good corporate citizenship and sustainable behaviors for the financial information industry, and to serve as a guide to financial firms regarding compliance practices while taking into account individual use cases and risk tolerances.

Key Terms

GenAI	Refers to a specific branch of artificial intelligence called Generative Artificial Intelligence. GenAI solutions use multi-layered computational abstractions of neural networks to generate content in an attempt to replicate human logic and commonsense reasoning. GenAI is distinct from other branches of artificial intelligence such as traditional machine learning (**ML**), which is focused on identifying patterns in specific data to make accurate predictions.
GenAI Model	A third-party algorithm or similar tool that is capable of inferring/reasoning from the prompts it receives from users (**inputs**) to create content, including text, images, audio, or video (**outputs**), based on underlying training data. GenAI Models are capable of producing unique and novel outputs not directly represented in their training data. GenAI Models may be procured directly or via a third-party platform, system or tool that utilizes GenAI Model(s) (e.g., Microsoft Azure).
LLM	**Large Language Model** is a text-based artificial intelligence model that powers GenAI Models.
Fine-Tune	The process of providing additional data to a pre-trained GenAI Model to augment its capability and be more effective for a specific use case (e.g., incorporating firm data to be leveraged in outputs).
RAG	**Retrieval Augmented Generation** allows GenAI Models to access specific pieces of data as part of its process of creating outputs (e.g., firm documents) which have been optimized for LLMs without Fine-Tuning.

Risks

MNPI	Firms may face regulatory scrutiny if they have reason to know that the output generated by a GenAI Model contains **material non-public information** obtained in breach of a duty.
Confidentiality & IP Leakage	Depending on their terms of service and cybersecurity protocols, GenAI Models may have access to firm inputs/outputs (e.g., confidential information) and/or be vulnerable to cyber attacks. Additionally, if firm inputs/outputs are used as training data, such information may be exposed to others outside of the firm, thereby waiving trade secret protections and/or undermining the confidential status of the inputs/outputs.
Infringement	GenAI relies on training data from a vast set of resources, which may be subject to certain laws (e.g., copyright). Firms may face risk depending on the format of the output they receive and whether the data used to develop the GenAI Model violates such laws.
Privacy	Information associated with natural persons, such as Personally identifiable information (**PII**) may be subject to certain privacy laws and require adherence to specific notice and consent protocols prior to use with GenAI.
Accuracy	GenAI can produce inaccurate results and does not often provide source citations to verify the accuracy of those results by default. GenAI Models can "hallucinate" whereby they generate inaccurate results with confidence.
Bias	Due to biases in training data, GenAI may produce outputs that reflect biases. Federal, state, and international regulators have warned about the discriminatory potential of GenAI in making decisions having legal effect, such as employment or credit decisions.
Consumer Protection	Offering products to the market that do not disclose the products' non-human origin may run afoul of requirements issues by various federal and state regulatory authorities, e.g., FTC, CFPB and state attorneys general.
Regulatory	These are risks associated with failing to use GenAI in adherence with laws, rules and regulations, such as failing to enact reasonably designed policies and procedures or providing sufficient disclosures.

Controls

To address and mitigate the GenAI risks identified above, firms may consider implementing the following controls:

Guardrails	Technical and administrative controls that re-enforce the firm's stance and risk tolerance on GenAI (e.g., policies and procedures; limiting use cases; blocking certain types of inputs).
Training	Specialized training better ensures that users are aware of the firm's GenAI Guardrails and may also assist users' understanding of the limitations and capabilities of GenAI (e.g., providing guidance on proper prompt construction to generate more accurate results).
Testing	Benchmarking and sampling GenAI outputs may help identify errors and hallucinations and assess how well the firm's GenAI Models perform on data they have never seen before (e.g., test GenAI output against validated, out-of-sample data to confirm accuracy and identify bias).
Monitoring	As the technology, use cases, and regulations evolve, GenAI controls may need to be reassessed.
Oversight	Endeavor to inject human oversight into the GenAI process when reasonably practicable (e.g., code review, periodic human review/validation/manual intervention)

Governance

Discussing the risks outlined above with the firm's stakeholders will help determine how the firm will address and mitigate them. They can be considered in the firm's core compliance policy, contract negotiation playbook, GenAI approval procedures, Guardrails, and elsewhere.

Policies and Oversight are important corporate governance features that seek to ensure that firms and their employees have clear guidance and consistent views on GenAI and how it may be utilized. Consider beginning the firm's policy with a definition of GenAI, which may include reference to specific tools that are in-scope of the policy (e.g., OpenAI's ChatGPT) and/or a recitation of out-of-scope tools (e.g., ML/NLP models). The possible span of use cases for GenAI can be very wide and while it is important to acknowledge that GenAI has enabled unique use cases that may present novel risks (e.g., using a GenAI chatbot to communicate with investors), much of GenAI's most immediate and widespread impact has been accelerating existing and well accepted capabilities (e.g., data cleansing, information extraction). To streamline GenAI adoption and maintain consistency with existing firm practices, careful thought should be given as to where the firm's GenAI policy will fit into its overall compliance structure. Where possible, firms may leverage and/or enhance existing policies and procedures to address the novel risks and **use cases presented by GenAI**. For instance, data that is input into a GenAI Model should follow the firm's Data Policy requirements to better ensure that the firm doesn't access or receive MNPI in the form of GenAI output. Firms may also wish to have their policy address information subject to confidentiality obligations, IP considerations, and under what circumstances there will be a human review overlay (i.e., review of outputs). These risks will evolve over time and policies should be designed to adapt and address change.

It is important to recognize that there are different **versions of GenAI Models** that firms may choose to utilize. These include public versions (e.g., free and API versions of ChatGPT); enterprise-licensed versions hosted on a private platform (e.g., OpenAI offered through Microsoft Azure); and private or proprietary versions (e.g., open-source or firm-developed GenAI Models hosted on a firm cloud or firm hardware). Different versions of GenAI Models carry different risk profiles and firms may consider having a list of versioning requirements and corresponding use cases that have been vetted and deemed acceptable. Articulating the types of information that should or should not be input into a GenAI Model may also be a prudent exercise, which may depend on the use case and the version utilized.

GenAI is increasingly incorporated into everyday **commercial tools and services** (e.g., DocuSign, Slack) and firms may want to consider how to address those issues. Commercial tools incorporating GenAI may use one or more versions of GenAI Models (i.e., a Microsoft Word plug-in may query a public version of ChatGPT). When assessing the permissibility of a commercial tool with GenAI functionality, consider prioritizing business critical applications and conducting a risk analysis related to the importance, scale, information security and sharing practices of the tool regarding firm inputs and. Guardrails may be tailored to account for specific risks identified. For example, some use cases may involve ingestion of PII or PHI that could raise privacy, bias or discrimination concerns. In sensitive cases, for instance in the healthcare space, firms may decide to seek out tools that integrate HIPAA compliance considerations into their GenAI Models.

Another corporate governance consideration is implementing a **GenAI approval process** that specifies the internal personnel/committees that must provide input and approval. This approval process may address legal/security/business review requirements; use cases and firm information that may be utilized with GenAI; the internal Testing and Monitoring process for GenAI solutions; the workflow for reviewing and enhancing GenAI policies and procedures as the environment evolves. This approval process may function as a complementary Guardrail to the firm's policies and procedures and may also better ensure that business stakeholders and control functions have a clear Oversight and understanding of specific projects where GenAI is deployed throughout the firm.

Firms may wish to provide guidance in their policies or training on the use of GenAI tools to prevent the generation of inaccurate or biased information. Training on "prompt engineering" may result in more useful and accurate outputs from GenAI models. Whether or not training is provided, consider a policy requirement that users of GenAI tools remain ultimately

responsible for ensuring the accuracy of outputs. You may also wish to formalize approval of the uses of GenAI in the partial or full automation of your business processes.

Investment advisers are subject to **recordkeeping requirements** under the Investment Advisers Act Rule 204-2 (the "Books and Records Rule") and certain uses of GenAI Models may implicate aspects of the Books and Records Rule. For instance, use of a GenAI Model to create "Advertisements", may require an adviser to maintain books and records. Although the state of the law around recordkeeping and GenAI Models is in flux, it seems unlikely that an Investment Adviser would be required to retain *all* inputs and outputs of a GenAI Model. While the user interface to such tools is conversational in nature, prompts and outputs may not be true "communications" that would be required under section 204-2(a)(7) of the Books and Records Rule. But note, if the adviser is using a GenAI chatbot to directly interact with clients or communicate the placement of orders, those communications would be subject to the recordkeeping requirements. Notwithstanding the requirements of the Books and Records Rule, an adviser may want to consider retaining certain prompts submitted to GenAI Models and/or the outputs of such tools. Factors to consider in making this determination include: how useful would the records be for business, compliance, or other purposes; what costs and technical requirements are there in maintaining such records (e.g., cloud storage, proxy intercept costs). However an adviser lands on the above determinations, it may desire to consider whether the length of the GenAI providers' retention period for prompts and input (e.g., 30 days) conflicts with the adviser's recordkeeping requirements.

Diligencing and Onboarding GenAI Models

To address some of the novel issues presented by GenAI, firms may consider implementing different vetting factors and Guardrails depending on the use case being considered and which version(s) of GenAI Models are used. As previously mentioned, to better ensure consistency and transparency, firms may consider documenting risk assessment, legal review and approval standards, as well as maintaining list(s) of approved GenAI Models and use cases, along with any applicable limitations.

Firms may consider reviewing the GenAI Model's **cyber security** terms, product statements and disclosures (e.g., ToS, FAQs) to confirm that they do not provide the GenAI provider with use or access rights to the firm's inputs or outputs (or Fine-Tuned models, where applicable) that contravene the firm's security and/or confidentiality requirements. GenAI Model documentation may set out the degree to which firm information is subject to data access and security parameters, such as whether data will be confined within the firm's IT environment, or will be transferred to, stored, or processed within the GenAI provider's IT environment or physical location. The use of cloud-based GenAI Models in an otherwise on-premises data processing model may also pose heightened concerns around data leakage. Nearly all enterprise GenAI providers offer solutions subject to SOC2 and other relevant information security frameworks. Firms may desire to review these certifications and similar terms to confirm that they meet the firm's requirements.

GenAI Models are **trained** on incredibly large corpuses of data, much of it collected from the web by GenAI providers (or taken from public, open-source data collections) or licensed from various data partners. While it may be challenging to understand the full scope of training data incorporated into a GenAI Model, firms may consider reviewing white papers, research papers, or other documents describing the training corpus, along with techniques used by the GenAI provider to cleanse the training data of sensitive, confidential, or otherwise objectionable material. With respect to MNPI, the cutoff or "knowledge date" of the training data may be relevant, as a GenAI Model trained on stale data is less likely to reproduce MNPI.

Firms may consider assessing the GenAI provider's technical ability to access **firm inputs** and use them to train GenAI Model(s) and/or reproduce the inputs to third parties. Absent technical controls or contractual terms preventing this, firms should be aware that information used to train a GenAI Model and, to a lesser extent, information entered as part of an input even if not used for training (particularly for public versions of GenAI Models), carry a risk of disclosure to the GenAI provider and to future users of the GenAI Model. Such disclosure can negatively impact a firm's maintenance of **trade secret protections**. Although trade secret law varies from one jurisdiction to another, it will protect certain valuable and generally

unknown information from disclosure and use, so long as efforts that are reasonable under the circumstances were made to protect its secrecy. Purposeful or inadvertent "leaks" of proprietary, secret information can result in loss of protection. The disclosure described above may result in making information protected by trade secret law available to GenAI in a way that destroys trade secret protection. In some cases, bespoke settings that lower disclosure risks may be available, such as opting-out of human review and monitoring. Many providers now offer Zero-Data Retention which can be favorable to obtain in business contexts.

Firms may also consider reviewing the GenAI Model's contractual use rights and restrictions regarding **firm outputs** (e.g., terms-of-use, acceptable use policy, data retention) to confirm that they provide, and/or do not preclude, the firm's required use rights (for example, the right to create derived data from the output), and that they do not provide the GenAI provider with any use or access rights to the output that the firm does not wish to provide. Users intending to produce materials for which **copyright** and authorship protection may be sought should be aware that pursuant to US Copyright Office guidance, a work is not copyrightable if it was produced using a GenAI prompt alone, with no additional human authorship. The extent to which human authorship of GenAI output is required for the purpose of copyright and authorship protection is not yet a settled area of the law. GenAI Models that disclose outputs to the GenAI provider and/or that do not limit such provider's rights in the output may result in the firm failing to establish new trade secret protections for the output it generates.

In addition to diligencing GenAI Models and providers, firms may also choose to assess how **data vendors** are incorporating GenAI into the development of their data products. Areas that may be relevant for GenAI due diligence include: the versions of GenAI Models used by the data vendor (i.e., public vs. private); the role that GenAI played in the production of the data product; the inputs the data vendor used; the Guardrails, Testing and Monitoring practices implemented by the data vendor to check the data product's accuracy; the method for Fine-Tuning the data vendor's GenAI Model(s), etc. In an effort to address these issues, FISD has added GenAI-related questions to its standard Compliance Due Diligence Questionnaire and firms should consider using this document to better identify and address the issues above.

Contractual Considerations for GenAI Models

While many of the contractual considerations set forth by FISD guidelines will be applicable when onboarding GenAI Models, there are a few areas that may warrant particular focus given the novelty of GenAI. To the extent a GenAI provider offers its product subject to non-negotiable terms that are posted online and subject to change, firms may consider periodically reviewing the terms for material changes.

To preserve the **confidentiality** of the firms' information (and protect the firms' trade secrets), firms should pay close attention to the contractual language related to GenAI providers' ability to access, store, or utilize firms' inputs and outputs . Firms may seek to restrict such capabilities to limited circumstances (e.g., only to resolve security or abuse violations). Relatedly, firms may seek to include provisions confirming the firms' **ownership** rights to the inputs/outputs and that such information may not be used to train the GenAI provider's GenAI Models, nor be included in any output provided to other users of the GenAI Models. If a firm will Fine-Tune a GenAI Model, the firm should consider including provisions that address ownership or long-term license of the Fine-Tuned model.

Because GenAI Models are trained on large amounts of web data collected across the Internet, a number of high-profile copyright infringement actions have been filed against GenAI providers alleging that some of the training data is copyrighted and used without permission. To address this issue, firms may consider seeking enhanced **indemnification** obligations for IP infringement and, where available, using **protective settings** offered by GenAI providers to mitigate IP infringement risks (e.g., enabling content match filtering). Lastly, firms may consider the contractual distinctions between licensed GenAI Models versus early access/beta versions of GenAI Models. Oftentimes, early access/beta versions of GenAI Models offer fewer **representations and warranties** than licensed GenAI Model.