Use of Third-Party Natural Language Generators Guideline

Cybersecurity Services Branch and Privacy, Policy and Governance Branch
Technology and Innovation

Version: 1.0

Approved by: ARMC/ADM IMT	**Owner:** Executive Director, Cybersecurity Services Branch, Cybersecurity Services Division, Technology and Innovation Executive Director, Privacy, Policy and Governance Branch, Innovation, Privacy and Policy Division, Technology and Innovation	
Approval date: 13 October 2023	**Last reviewed:** 13 October 2023	**Review date:** 13 October 2025
Contact: cybersecurity@gov.ab.ca Sa.InformationManagement@gov.ab.ca	**Policy Instrument type:** Guideline	

Contents

Guideline Statement 3
Authority 3
Application and Scope 3
Definitions 3
Guideline Specification 3
General Considerations 3
Formal Business Process Integration 4
Acceptable Use Cases 4
Compliance 6
References and Supporting Resources 6
Contact 6
Appendix A – Definitions 7

Guideline Statement

This guideline is intended to educate Government of Alberta (GoA) staff on the acceptable use of third-party Natural Language Generator (NLG) applications by providing acceptable use cases for the technology.

Authority

This guideline was developed in conjunction and alignment with the Acceptable Use of Third-Party Natural Language Generators Directive.

Application

For more information on the application and scope of this guideline, please see the Acceptable Use of Third-Party Natural Language Generators Directive.

Definitions

For definitions of the terms used in this guideline, please refer to **Appendix A - Definitions**

Guideline Specification

General Considerations

Use of third-party NLG applications by GoA staff must align with the Acceptable Use of Third-Party Natural Language Generators Directive. In addition to the requirements detailed in the directive, it is strongly recommended that business areas:

- refresh their familiarity with the Official and Transitory Records Directive, as inputs provided to third-party NLGs and/or third-party NLG outputs may be classified as official and/or transitory records;
- refresh their familiarity with the Data and Information Security Classification Standard, and ensure that all content in the custody and/or under the control of the GoA (e.g., outputs generated by a third-party NLG and modified by GoA staff) has a security classification applied;
- review all privacy statements/policies and/or End User License Agreements prior to using a third-party NLG, paying particular attention to what information is being collected and how the information is being stored, used, and disclosed—collaboration with Privacy Services may be necessary in some instances;
- assume that any content provided to a third-party NLG (including, but not limited to, user-provided inputs and information used to create an account) is collected, retained, and potentially used for different purposes by the application**;**
- use generalized prompts wherever and whenever possible (e.g., asking a third-party NLG how “an organization” should respond to an issue, rather than how “the GoA” should respond);
- consider not only accuracy, bias, and legality, but also tone, clarity, consistency, and context (e.g., necessity of ensuring alignment with the GoA voice, tone and style guide) when validating and/or verifying third-party NLG outputs;
- recognize the possibility of negative stakeholder responses when determining which use cases are considered “acceptable”;
- establish processes for recording user-provided inputs and third-party NLG outputs, if necessary; and
- consider how and when third-party NLGs are used by contractors, as contract provisions may need to be established or reviewed.

Formal Business Process Integration

Business areas may leverage third-party NLGs on an ad hoc basis under the requirements established in the Acceptable Use of Third-Party Natural Language Generators Directive; however, business areas may be interested in formally integrating third-party NLGs into business processes. Formal integration of third-party NLGs into business area processes may be subject to requirements beyond those detailed in the Acceptable Use of Third-Party Natural Language Generators Directive.

NOTE: Third-party NLGs that have been subject to assessment and approval prior to the publication of the Acceptable Use of Third-Party Natural Language Generators Directive may not require immediate reassessment. Business areas with questions regarding assessment of third-party NLGs are encouraged to contact Cybersecurity Services.

The process for formal integration is as follows:

- Business areas will develop a business case for integrating a third-party NLG into a business process—this may involve collaboration with Cybersecurity, Privacy Services, legal counsel, and content management subject matter experts.
- Cybersecurity will require the completion of a Security Threat and Risk Assessment (STRA).
    - An STRA must be performed for each third-party NLG platform that a business area intends to integrate into their business processes.
    - The STRA may require the completion of an assessment of privacy risk in collaboration with Privacy Services.
    - Once completed, the STRA may be leveraged by other business areas with similar business processes and/or intended use cases.
- Depending on the specifics of the business area and the business process, the STRA may result in:
    - additional process requirements (e.g., mandatory recording of inputs provided to third-party NLGs, revised reporting structures, etc.);
    - additional usage requirements (e.g., requiring the use of GoA email addresses and/or identities for account creation);
    - prescribed use cases (e.g., enabling NLG use in scenarios that may fall outside the parameters of the Acceptable Use of Third-Party Natural Language Generators Directive); and/or
    - the ability to use content in the custody and/or under the control of the Government of Alberta with applied security classifications of Protected A, Protected B, and/or Protected C as third-party NLG inputs.

Acceptable Use Cases

In general, third-party NLGs can be leveraged by GoA users for the following purposes without requiring formal business process integration:

- **creating summaries or overviews** of lengthy publicly available documents (such as whitepapers, research papers, etc.);
- **generating ideas or brainstorming** (e.g., prompting a third-party NLG for potential solutions to an issue), provided that the prompts are general, and the outputs are further developed and refined by staff; and
- **within licensed systems** (e.g., third-party NLG functionality added to existing systems), provided the systems have been subject to review by Cybersecurity and approval from the appropriate business area(s).

The example use cases highlighted in this guideline are not intended to be prescriptive nor exhaustive; as NLGs advance, acceptable use cases will evolve. Depending on the desired use case, business areas may want to collaborate with relevant subject matter experts (e.g., Cybersecurity, privacy, legal counsel, etc.) before considering the use of third-party NLGs.

	Examples
Permitted without authorization	• Using data from the Government of Alberta's Open Data Portal to generate a summary. • Conducting an initial jurisdictional scan on published policies from other jurisdictions. • Summarizing a publicly available scientific paper or report. • Using internet search engines that have NLG integration by default, given that the outputs are publicly accessible. • Generating ideas or brainstorming, provided that the prompts are general, and the outputs are further developed and refined by staff.
May be permitted upon completion of an STRA	• Inputting content that is not publicly accessible, so long as the process has been approved by the Information Controller and/or an STRA has been completed with Cybersecurity Services. • Supporting language translation, so long as outputs are verified and validated by qualified staff. • Supporting proofreading and copyediting of documents (e.g., addressing grammatical errors, improving clarity and/or readability), so long as the inputs do not contain prohibited information and the outputs are verified and validated. • Drafting communications to stakeholders, provided that proper approval processes (i.e., an STRA) have been completed, and that all third-party NLG outputs are sufficiently validated, verified, and/or identified. • Any uses determined to be acceptable upon completion of necessary approval processes (i.e., an STRA).
Prohibited	• Using third-party NLGs in a manner contrary to the rules, requirements, and obligations laid out in legislation and relevant content and privacy management policy instruments (e.g., neglecting to record both inputs and outputs in instances where recording is mandatory). • Inputting content that is not publicly accessible without approval of the Information Controller and/or the completion of an STRA with Cybersecurity Services. • Integrating third-party NLGs into business area processes without completing an STRA. • Using third-party automated transcription services (e.g., Otter.AI) to transcribe internal government meetings. • Writing responses to stakeholders using content that is not publicly accessible as a third-party NLG input and/or without validating and verifying third-party NLG outputs. • Drafting briefing materials for a minister or Cabinet.

	Examples
	• Using third-party NLG outputs verbatim without clearly citing, referencing, or otherwise indicating that the content is a third-party NLG output. • Generating, reviewing, or in any other way interacting with GoA proprietary computer code. • Generating outputs that impersonate another individual or entity. • Using an NLG for any illegal, unethical, or otherwise inappropriate activities.

Compliance

For more information on compliance with this guideline, please see the Acceptable Use of Third-Party Natural Language Generators Directive.

References and Supporting Resources

- Acceptable Use of GoA IT Assets Directive
- Information Controller and Information Custodian Directive
- Use of Third-Party Natural Language Generators Directive
- Code of conduct and ethics for the Public Service of Alberta
- Data and Information Security Classification Standard

Contact

Types of Questions	Contact
Content Management	Sa.InformationManagement@gov.ab.ca
Cybersecurity	cybersecurity@gov.ab.ca
Privacy	privacy@gov.ab.ca

Appendix A – Definitions

Application – a collection of computer programs, databases, and procedures designed to help the GoA perform particular tasks or handle particular types of IT problems by automating a business process or function. Also known as a business application.
Source: Information Security Management Directives

Contractor - a person or company that undertakes a contract to provide materials or labor to perform a service or do a job.

End-User License Agreement (EULA) - a legal contract between a software supplier and a customer. An EULA specifies in detail the rights and restrictions which apply to the use of the software.

Input - refers to any content that is provided to an application for processing. In the context of NLGs, inputs may take the form of prompts (e.g., specific user instructions, commands, questions and/or discussion topics provided to an NLG to elicit an output response).

Natural Language Generators: Natural Language Generators (sometimes referred to as "Natural Language Processors") are computer programs that use linguistic and/or non-linguistic representations of information (such as large language models) to generate coherent and contextually appropriate human-like outputs in response to a human-provided input. Common examples of third-party Natural Language Generators include (but are not limited to) ChatGPT, Bard, and LLaMA.
Source: NLG DIRECTIVE

Licensed systems - all products and systems developed by or for the GoA which are licensed, sold, distributed, or otherwise transferred to the GoA by a third party.

Output – the products generated as the result of processing by an application. In the context of NLGs, outputs may take the form of music notation, poetry, plain text, computer code, etc.

Personal information – means recorded information about an identifiable individual, including (but not limited to) age, name, income, opinions, biometric information, and ethnicity. For a comprehensive list of content considered personal information, please refer to the *Freedom of Information and Protection of Privacy Act* (FOIP Act).
Source: FOIP Act

Third party - includes external parties and persons outside the direct reporting structure of the Information Controller or Information Custodian (e.g., an individual, a business or organization, personnel from another branch of government, or another level of government). Third party also include vendors, service delivery agents, businesses, and citizens.
Source: Information Security Management Directives