

ServiceNow Large Language Model (LLM) Model Card (v2.0)

ServiceNow Large Language Model (LLM) Model Card (v2.0)

Model Name

gpt-oss-120b

Model Developer

OpenAI

Version

Initial public release (August 5, 2025)

Source

<https://openai.com/index/introducing-gpt-oss/>

License

[Apache 2.0](#)

Usage in Our System

gpt-oss-120b is leveraged within ServiceNow AI-driven solutions to support natural language understanding, automation, and decision support across enterprise workflows.

Built with OpenAI's gpt-oss family, the model is used as-is (no ServiceNow fine-tuning applied) and is selected for its strong reasoning, tool-use capabilities, and efficient deployment profile (MoE; designed to run efficiently on a single 80 GB GPU), making it suitable for on-prem and privacy-sensitive use cases.