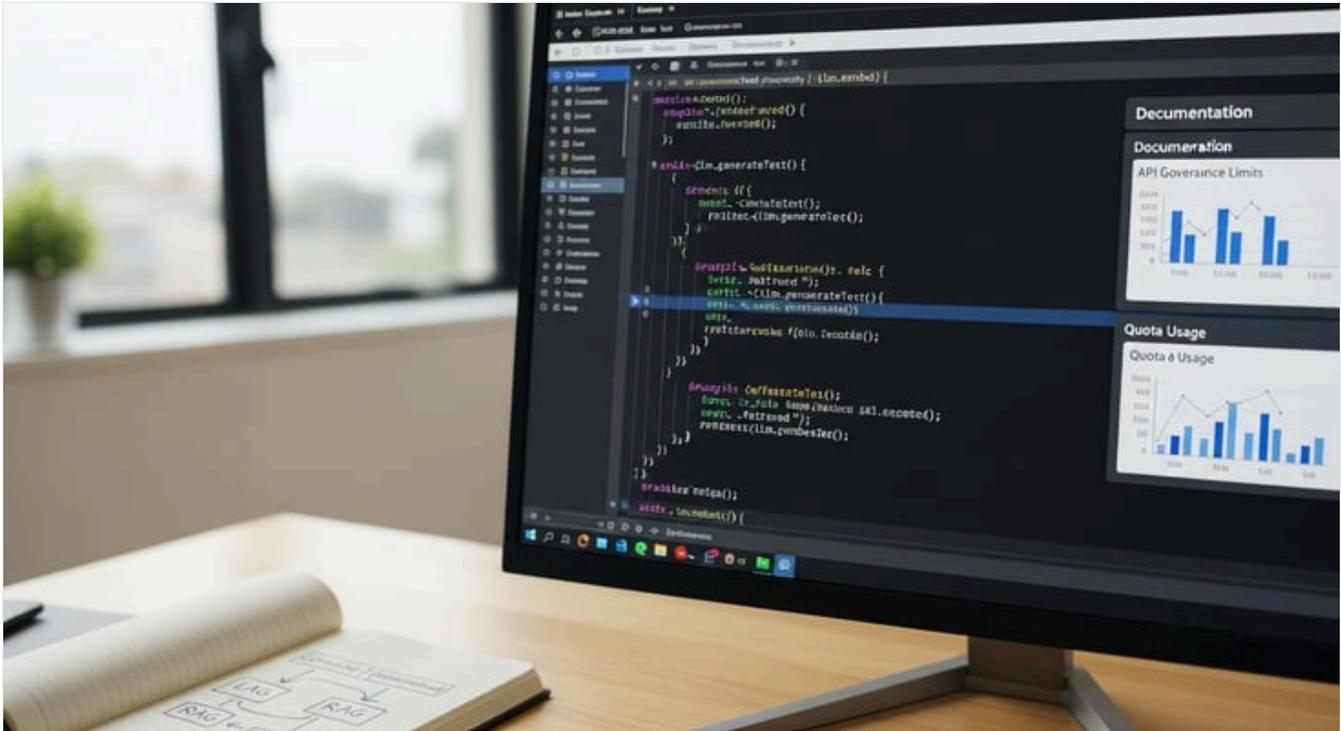


# NetSuite N/LLM Module: SuiteScript GenAI API Guide

By houseblend.io Published February 12, 2026 41 min read



## Executive Summary

This report examines Oracle NetSuite's newly introduced [N/LLM SuiteScript module](#), which embeds large-language-model (LLM) capabilities directly into the NetSuite platform. NetSuite's N/LLM module (available in SuiteScript 2.1) provides an on-platform interface to Oracle's OCI-backed generative AI services. It allows developers to call LLMs from SuiteScript in order to generate or [analyze text](#), embed text as vectors, and build intelligent features ([chatbots](#), natural-language queries, summaries, etc.) using a company's own data. For example, `llm.generateText(options)` sends a custom prompt to the LLM and returns generated text, while `llm.embed(options)` produces vector embeddings of input text (Source: [docs.oracle.com](#)) (Source: [docs.oracle.com](#)). The module also supports *prompt evaluation* (via prompts defined in [Prompt Studio](#), [retrieval-augmented generation \(RAG\)](#) (by supplying NetSuite source documents to ground the responses), *streamed responses*, and administrative methods to track usage. Script aliases like `llm.chat()` and `llm.chatStreamed()` provide additional ease of use (Source: [docs.oracle.com](#)).

Crucially, NetSuite governs N/LLM usage via a **monthly free-usage pool** and strict [concurrency limits](#). Each account (and each installed SuiteApp) receives a limited number of free LLM calls per month, reset monthly (Source: [docs.oracle.com](#)). Generation calls (e.g. `generateText`) and embedding calls use separate quotas (Source: [docs.oracle.com](#)). By default, accounts draw from the free pool, but they can configure their own OCI generative-AI credentials to obtain unlimited usage (billed to their OCI account) (Source: [docs.oracle.com](#)) (Source: [docs.oracle.com](#)). Concurrency is also limited: up to **five concurrent LLM calls** are allowed for generation methods and up to five for embedding methods (Source: [docs.oracle.com](#)) (calls beyond that throw errors). The SuiteScript **AI Preferences** page displays usage tables by month and type (generate vs embed) so administrators can monitor consumption (Source: [docs.oracle.com](#)) (Source: [docs.oracle.com](#)).

NetSuite's decision to integrate LLMs natively reflects industry trends and competitive pressures. Analysts note that enterprise adoption of generative AI has surged: Gartner found that by late 2023, roughly 29% of surveyed organizations had already deployed GenAI solutions (Source: [www.gartner.com](#)), and forecasts suggest over 80% will use GenAI APIs or apps by 2026 (Source: [www.gartner.com](#)). Oracle's strategy contrasts with some peers: instead of charging extra for AI, NetSuite embeds AI features (like its new N/LLM capabilities) as part of the base product at **no additional license cost** (Source: [www.axios.com](#)) (Source: [the-cfo.io](#)). Oracle NetSuite EVP Evan Goldberg emphasizes that "AI is going to be everywhere" and insists on making AI "table stakes" for customers (Source: [www.axios.com](#)) (Source: [the-cfo.io](#)). Meanwhile, NetSuite's competitors (e.g. SAP) have moved toward premium AI offerings, highlighting NetSuite's emphasis on expanding AI accessibility.

This report provides a comprehensive analysis of the N/LLM module and its ecosystem. We cover historical context (NetSuite's AI rollout and market positioning), technical details of the SuiteScript API, governance of usage and quotas, illustrative code examples, and expert perspectives. We also incorporate data and case examples on AI in enterprise ERP: for instance, industry analysis shows that large-volume NetSuite customers (such as BirdRock Home, which “processes thousands of orders daily” in NetSuite (Source: [houseblend.io](https://houseblend.io)) stand to benefit from grounded generative AI on their rich datasets. Finally, we discuss risks (accuracy, privacy, governance) and future directions (more AI features, training data quality, model choices), and provide evidence-based recommendations.

## Introduction and Background

Generative AI (the use of large language and foundation models to generate or analyze content) has rapidly emerged as a transformative technology in business and ERP systems. In late 2022 and 2023, tools like ChatGPT and other LLMs captured widespread attention by demonstrating that AI can perform many routine tasks (text summarization, translation, content creation) with human-level fluency (Source: [www.mckinsey.com](https://www.mckinsey.com)) (Source: [www.mckinsey.com](https://www.mckinsey.com)). Analysts have projected enormous productivity and economic impact. For example, McKinsey (2023) estimates that generative AI use cases could add **\$2.6–4.4 trillion** annually to the global economy (roughly doubling the impact of all prior AI) (Source: [www.mckinsey.com](https://www.mckinsey.com)). Other studies suggest generative AI could automate or augment 60–70% of today's work activities, fundamentally changing roles across industries (Source: [www.mckinsey.com](https://www.mckinsey.com)) (Source: [www.mckinsey.com](https://www.mckinsey.com)). As a result, enterprise interest in AI exploded: Gartner (October 2023) predicted that by 2026 over **80% of enterprises** will have deployed generative-AI-enabled applications or used GenAI APIs, up from less than 5% in 2023 (Source: [www.gartner.com](https://www.gartner.com)). By early 2024, Gartner found GenAI was already the **most frequently deployed AI solution** in organizations (Source: [www.gartner.com](https://www.gartner.com)).

NetSuite (an Oracle company) – a leading cloud ERP/SaaS platform for mid-market to enterprise companies – has rapidly incorporated AI to leverage these trends. Oracle executives have highlighted NetSuite's “Suiteness” (the benefit of centralized, integrated data) as a foundation for AI-driven insights (Source: [houseblend.io](https://houseblend.io)) (Source: [houseblend.io](https://houseblend.io)). In 2023 and 2024, Oracle announced a sweeping AI roadmap for NetSuite: generative AI assistants (e.g. the “Text Enhance” feature for assisted content authoring) were unveiled at SuiteWorld 2023 (Source: [www.techtarget.com](https://www.techtarget.com)), and by SuiteWorld 2024 dozens of new AI-powered features (like the SuiteAnalytics Assistant and narrative reporting) were announced (Source: [the-cfo.io](https://the-cfo.io)) (Source: [the-cfo.io](https://the-cfo.io)). Keynote speakers emphasize that NetSuite differs from competitors by **embedding AI capabilities at no extra cost** – treating them as standard platform features (Source: [www.axios.com](https://www.axios.com)) (Source: [the-cfo.io](https://the-cfo.io)) – and by tightly coupling AI models to the customer's own enterprise data via Oracle Cloud Infrastructure (OCI).

Central to this strategy is the **SuiteScript N/LLM module**: a native SuiteScript 2.1 library providing programmatic access to large language model (LLM) services from within JavaScript-based SuiteScript scripts. Announced in 2024, N/LLM effectively brings OCI's generative AI service (powered by models from partners like Cohere) into NetSuite's SuiteScript environment (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [gurussolutions.com](https://gurussolutions.com)). Developers using SuiteScript – NetSuite's extension framework – can thus write scripts that “talk” to LLMs, enabling on-demand content generation, question-answering, summarization, or classification within the ERP. As one NetSuite blog explains, N/LLM is “a newly introduced NetSuite SuiteScript 2.1 module that brings native access to a Large Language Model (LLM) directly inside NetSuite” (Source: [blogs.oracle.com](https://blogs.oracle.com)). This allows tasks such as creating contextual documents, controlling generation parameters (like temperature), tracking citations from the user's own data, and monitoring API usage quotas (Source: [blogs.oracle.com](https://blogs.oracle.com)).

The importance of N/LLM lies in its ability to *ground* generative AI in a company's *own* ERP data. By combining SuiteScript data retrieval (via SuiteQL, saved searches, etc.) with LLM prompts, scripts can implement **retrieval-augmented generation (RAG)**: the script fetches relevant records (e.g. recent sales orders) and passes them as “documents” into the LLM, which then generates answers or reports based on that actual data (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [blogs.oracle.com](https://blogs.oracle.com)). As one developer blog describes, the process is: “Build an array of documents (createDocument) related to your question. Submit the documents along with the prompt to generateText(). Receive not only the answer but also citations pointing back to your documents” – ensuring the output is “based on your own NetSuite data, not random internet knowledge” (Source: [blogs.oracle.com](https://blogs.oracle.com)). This enterprise-focused RAG approach mitigates the “hallucination” risk of LLMs by anchoring them to vetted records.

At the same time, NetSuite's AI capabilities are subject to strict **governance**. As Oracle notes, NetSuite “provides a free monthly usage pool of requests for the N/LLM module” (Source: [docs.oracle.com](https://docs.oracle.com)). All generative calls (`llm.generateText`) and prompt-evaluation calls (`llm.evaluatePrompt`) consume from this pool, and similarly a separate pool tracks embedding calls (`llm.embed`) (Source: [docs.oracle.com](https://docs.oracle.com)). The pools renew each month (Source: [docs.oracle.com](https://docs.oracle.com)). Concurrency limits further govern usage: at most five concurrent calls to generation methods and five to embedding methods are allowed per script, after which errors are thrown (Source: [docs.oracle.com](https://docs.oracle.com)). If an organization needs more usage, they can configure an external OCI Generative AI credential: from then on, calls draw from that account's quota rather than the free pool (Source: [docs.oracle.com](https://docs.oracle.com)). All of these constraints are tracked on a SuiteScript AI Preferences page (Setup > Company > AI > AI Preferences) showing monthly limits and usage by type (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). This built-in system ensures responsible consumption of AI compute and allows administrators to monitor or limit AI usage.

Throughout this report, we cite official documentation, industry analyses, and expert commentaries to provide a thorough, evidence-based view of N/LLM. Key findings include:

- API Capabilities:** N/LLM introduces methods such as `llm.generateText(options)` for content generation, `llm.evaluatePrompt(options)` for prompting curated templates, `llm.embed(options)` for vector embeddings, and utility methods like `llm.createDocument()` (for RAG context) and `llm.getRemainingFreeUsage()` (for quota checks) (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). Each has corresponding asynchronous (promise) and streaming variants. The module also includes chat-oriented APIs (`llm.createChatMessage`, `llm.ChatRole`) for conversational UIs (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)).
- Governance:** Oracle imposes strict usage and concurrency limits on N/LLM (5 concurrent calls, monthly quotas per account/SuiteApp) (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). Administrators can track usage by type in NetSuite's AI Preferences (Source: [docs.oracle.com](https://docs.oracle.com)), and can optionally link to their own OCI for unlimited usage (Source: [docs.oracle.com](https://docs.oracle.com)).
- Best Practices:** Industry commentary stresses validating AI output (NetSuite itself warns that generative responses are “creative” and must be checked for accuracy (Source: [docs.oracle.com](https://docs.oracle.com)) and aligning AI initiatives with governance. For instance, one analyst note argues N/LLM's success “hinges on structured implementation aligned with enterprise data governance” (Source: [houseblend.io](https://houseblend.io)). Unfocused AI pilots often fail: one report cited notes >95% failure rate for “unfocused” generative AI proofs-of-concept (Source: [houseblend.io](https://houseblend.io)).
- Real-World Impact:** Generative AI in NetSuite enables new use cases (natural-language data querying, automated report generation, smart field enhancements, chatbots). Case examples include NetSuite customers like BirdRock Home and Overture Promotions, who have leveraged AI/ML over their NetSuite data (birdRock for churn/predictions, Overture for supply-chain forecasts), noting “actionable strategy improvements” (Source: [houseblend.io](https://houseblend.io)) (Source: [houseblend.io](https://houseblend.io)). As a general enterprise trend, broad studies (Gartner, McKinsey) show generative AI adoption is rapidly growing and delivering measurable value (Source: [www.mckinsey.com](https://www.mckinsey.com)) (Source: [www.gartner.com](https://www.gartner.com)).

The remainder of this document unfolds as follows. We first provide **historical context** for NetSuite's AI strategy and introduce overall N/LLM concepts. We then deep-dive into the **SuiteScript API**, enumerating the N/LLM methods and their parameters, with code snippets. Next, we cover **governance and usage** (quotas, concurrency, region availability, security considerations). We follow with illustrative **code examples** (chatbots, RAG, embeddings). We include several **data-driven perspectives** (adoption statistics, ROI implications, analyst quotes). We compare and contrast with related tools (e.g. SAP Cloud AI, Salesforce Einstein). We also discuss **case studies** (real or hypothetical) showing N/LLM in action. Finally, we analyze **implications and future directions** (expected enhancements, AI risks, broader enterprise AI integration) and conclude with key takeaways.

## NetSuite Generative AI: Genesis and Context

### Industry Drivers: AI in Enterprise Software

By 2024–2026, generative AI had emerged as a central theme in enterprise software. Organizations across industries began embedding AI at multiple levels. For example, Gartner (2024) found that nearly **29%** of surveyed companies had already deployed generative AI solutions (making GenAI the most common AI type) (Source: [www.gartner.com](https://www.gartner.com)), and predicted that well over **80%** of enterprises would use GenAI APIs or apps by 2026 (Source: [www.gartner.com](https://www.gartner.com)). In practice, many companies started with integrated assistants (Copilot-like tools within their existing productivity apps) more often than custom LLM development (Source: [www.gartner.com](https://www.gartner.com)), but the appetite for tailored AI capabilities was clear.

Enterprise software vendors responded by adding AI enhancements across their suites. Oracle in particular made large bets: in 2023, at SuiteWorld, NetSuite announced “Text Enhance” – a generative AI assistant for finance reporting and customer writing – built on OCI and Cohere LLM models (Source: [www.techtarget.com](https://www.techtarget.com)). By 2024, NetSuite expanded these efforts into workflows and analytics: new SuiteAnalytics Assistant modules allowed natural-language data queries and insights generation (Source: [the-cfo.io](https://the-cfo.io)); AI-driven narrative reporting and forecasting support arrived in its EPM (Enterprise Performance Management) suite (Source: [the-cfo.io](https://the-cfo.io)) (Source: [the-cfo.io](https://the-cfo.io)). Crucially, Oracle decided **not to charge extra** for these AI features. As Axios reported, Oracle embedded over 200 AI enhancements into NetSuite “at no extra cost,” positioning AI as a fundamental capability rather than a premium add-on (Source: [www.axios.com](https://www.axios.com)). Evan Goldberg (Oracle NetSuite EVP) explained that restricting or taxing AI would only slow adoption: “AI is going to be everywhere...it's as big as the internet revolution,” he said (Source: [www.axios.com](https://www.axios.com)) (Source: [the-cfo.io](https://the-cfo.io)).

Analysts generally approved of NetSuite's approach. In 2023, Constellation Research's Holger Mueller noted that enabling generative AI at the **platform layer** (via Oracle's own services) was “doing it right,” since it leverages internal data securely (Source: [www.techtarget.com](https://www.techtarget.com)). TechTarget interviews echoed this: Nuances like NetSuite's “text enhance” were seen as addressing common ERP tasks (much time spent authoring texts) to bill as quick efficiency gains (Source: [www.techtarget.com](https://www.techtarget.com)) (Source: [www.techtarget.com](https://www.techtarget.com)). Industry observers also noted that because NetSuite's data model is highly integrated (the so-called “**Suiteness**”), its deployments often have large, granular datasets that are ideal for AI training and inference (Source: [houseblend.io](https://houseblend.io)) (Source: [houseblend.io](https://houseblend.io)). For example, Houseblend points out that NetSuite customers like BirdRock Home (a major retailer)

house “thousands of products” and “thousands of orders daily” in a single system (Source: [houseblend.io](https://houseblend.io)), providing rich time-series and categorical data for AI use cases (forecasting, anomaly detection, etc.). The centralized data and unified processes in NetSuite (“more of the data across your whole business” (Source: [houseblend.io](https://houseblend.io)) give generative AI models more context to work with than siloed point solutions.

At the same time, other ERP platforms (e.g. SAP, Microsoft, Salesforce) were rolling out their own AI units. NetSuite needed to keep pace; to that end, it strengthened its developer ecosystem. Oracle opened **Prompt Studio** (a tool for managing AI prompts and templates within NetSuite) and released the N/LLM SuiteScript API to give developers first-class LLM access. Business leaders emphasised practical productivity gains: for instance, CFO.com noted that automating routine financial-narrative tasks with AI (so finance teams can focus on strategy) was a key selling point (Source: [the-cfo.io](https://the-cfo.io)). Oracle framed NetSuite’s AI as embedded into workflows, not bolted-on, to ensure immediate ROI “as soon as they log in” (Source: [the-cfo.io](https://the-cfo.io)).

In summary, NetSuite’s incorporation of N/LLM is part of a broader AI-driven strategy: leverage OCI generative services to add intelligent capabilities across the suite, draw on the platform’s rich consolidated data, and empower customers at no extra licensing cost. The timing (late 2023/early 2024) aligned with accelerating enterprise AI adoption, and Oracle positioned itself to compete strongly in the mid-market ERP space. This historical context explains *why* the N/LLM module was developed: it enables NetSuite’s customers and partners to harness cutting-edge LLMs for analytics, content creation, and automation directly within their ERP system.

## The N/LLM SuiteScript Module: Overview of Capabilities

The N/LLM module is a **SuiteScript 2.1 library** specifically for generative AI. According to Oracle’s documentation, it supports all the main generative-AI patterns needed in an ERP context (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). Its features include:

- Content Generation:** The core method is `llm.generateText(options)`. Developers supply a natural-language *prompt* string describing the desired content or query, along with optional `modelParameters` (such as `maxTokens`, `temperature`, `topK`, `topP`, and `penalty` terms). The module sends this request to the OCI Generative AI service, which returns an LLM-generated response (as `response.text`). For example, a simple SuiteScript snippet demonstrates sending the prompt “Hello World” and receiving an AI-written reply (while also checking remaining quotas) (Source: [docs.oracle.com](https://docs.oracle.com)). When no model is specified, NetSuite defaults to a Cohere Command model, but the script can specify other model families via an `options.model` (from `llm.ModelFamily`) (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). The `generateText` call fetches the complete output. An asynchronous/promise version is available, and there is also a streamed variant (`generateTextStreamed`) for progressive output (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)).
- Prompt Evaluation (Prompt Studio):** NetSuite’s **Prompt Studio** lets admins define reusable prompts with variables. The API method `llm.evaluatePrompt(options)` takes the *ID* of a stored prompt (from Prompt Studio) and a mapping of variables to values, then sends that prepared prompt to the LLM. The module handles substituting the variables and using the prompt’s preset `model/parameters`. For instance, NetSuite’s docs give a sample where an inventory-item prompt is loaded (ID “`stdprompt_gen_purch_desc_invnt_item`”) and variables like `itemid` and `stockdescription` are filled in script (Source: [docs.oracle.com](https://docs.oracle.com)). The LLM response (and monthly usage remaining) is returned. This method has an alias `llm.executePrompt()`. Streamed/promise forms (`evaluatePromptStreamed` and its promise) are also available (Source: [docs.oracle.com](https://docs.oracle.com)).
- Retrieval-Augmented Generation (RAG) Support:** Crucially, `generateText` (and `generateTextStreamed`) accepts an array of **Document objects** under `options.documents`. Developers can programmatically create such documents with `llm.createDocument({id, data})`, where `data` is text (or structured content) providing context. When documents are supplied, the OCI service uses them to ground its answer and also returns an array of **Citation objects** indicating which source documents were used. Thus, a script can collect citations for auditability. Oracle’s example MVP “Sales Insights Suitelet” query builds a summarized sales dataset into multiple `document` entries and then calls `generateText`, with results including citations back to those constructed docs (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [blogs.oracle.com](https://blogs.oracle.com)). As NetSuite notes, this ensures responses are “factually grounded” in the user’s own data (Source: [blogs.oracle.com](https://blogs.oracle.com)). Embedding documents into prompts in this way mitigates hallucinations.
- Embeddings:** The method `llm.embed(options)` converts text inputs into vector embeddings (arrays of floats) using dedicated embedding models. The user provides one or more input strings (`options.inputs`) and selects an embedding model family (e.g. `llm.EmbedModelFamily.COHERE_EMBED`). The LLM service returns an `EmbedResponse` object containing the embedding vectors. This is useful for semantic search, similarity calculations, recommendations, clustering, etc. For instance, a Suitelet example shows generating embeddings for a list of item names and computing cosine similarity to find similar items (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). Embedding calls draw from a *separate* monthly quota (distinct from generation calls) (Source: [docs.oracle.com](https://docs.oracle.com)).

- Chat Support:** To facilitate chat or conversational scenarios, N/LLM provides chat-oriented utilities. The enum `llm.ChatRole` defines message roles (such as `USER` or `ASSISTANT`). A developer can call `llm.createChatMessage({role, text})` to create a `ChatMessage` object (Source: [docs.oracle.com](https://docs.oracle.com)). An array of `ChatMessage`s can then be passed as a `chatHistory` parameter to `generateText`. (NetSuite treats this like an initial conversation history.) In effect, `llm.chat(options)` (alias of `generateText`) can take a `chatHistory` with roles and continue the dialogue with the model. This is valuable for building LLM-based chatbots or multi-turn assistants within NetSuite.
- Utility Methods:** The module includes helpful status/check methods. Notably, `llm.getRemainingFreeUsage()` returns the number of free generation requests left in the current month, and `llm.getRemainingFreeEmbedUsage()` does the same for embedding requests (Source: [docs.oracle.com](https://docs.oracle.com)). Both have promise variants. This lets scripts check and display remaining quotas dynamically (as seen in the Sales Insights Suitelet code, which shows “Remaining LLM Usage” on the form via `llm.getRemainingFreeUsage()` (Source: [blogs.oracle.com](https://blogs.oracle.com)). These methods aid administrators and scripts in monitoring consumption.

The **N/LLM module members** can be summarized as in Table 1 below. Each method or object ties directly to an LLM-related function:

**Table 1. N/LLM SuiteScript Module Methods and Objects.**

MEMBER (SUITESCRIPT)	ALIAS / ROLE	PURPOSE / DESCRIPTION	EXAMPLE USAGE (SEE REFS)
<code>llm.generateText(options)</code> (alias <code>llm.chat(options)</code> )	–	Sends a prompt (string) plus parameters to the model; returns a <code>Response</code> object containing the generated text and any citations (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ) (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	<pre>const res = llm.generateText({ prompt: "Hello", modelParameters: {...} }); const text = res.text; (Source: <a href="https://docs.oracle.com">docs.oracle.com</a>)</pre>
<code>llm.generateTextStreamed(options)</code> (alias <code>llm.chatStreamed(options)</code> )	–	Same as <code>generateText</code> , but returns a <code>StreamedResponse</code> object that provides partial text incrementally as the LLM generates it (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	(Similar to above, but chunked output)
<code>llm.evaluatePrompt(options)</code> (alias <code>llm.executePrompt(options)</code> )	–	Evaluates a Prompt Studio prompt by ID with provided variable values. Returns a <code>Response</code> with the LLM answer (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	<pre>const resp = llm.evaluatePrompt({ id: "myPromptID", variables: {...} }); const ans = resp.text; (Source: <a href="https://docs.oracle.com">docs.oracle.com</a>)</pre>
<code>llm.evaluatePromptStreamed(options)</code> (alias <code>llm.executePromptStreamed(options)</code> )	–	Streamed version of <code>evaluatePrompt</code> (returns partial text as ready).	(Similar, with <code>StreamedResponse</code> )
<code>llm.embed(options)</code>	–	Generates embeddings for input text(s) using a specified model family. Returns <code>EmbedResponse</code> with vectors (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ) (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	<pre>const embedRes = llm.embed({   embedModelFamily: llm.EmbedModelFamily.COHERE_EMBED,   inputs: ["text1", "text2"] }); (Source: <a href="https://docs.oracle.com">docs.oracle.com</a>)</pre>
<code>llm.createDocument(options)</code>	–	Creates a <code>Document</code> object (with <code>id</code> and <code>data</code> <code>text</code> ) to supply as context in a generation call.	<pre>const doc = llm.createDocument({ id: "doc1", data: "Sales data..." }); documents.push(doc); (Source: <a href="https://blogs.oracle.com">blogs.oracle.com</a>)</pre>
<code>llm.createChatMessage(options)</code>	–	Creates a <code>ChatMessage</code> (with <code>role</code> and <code>text</code> ) for building a multi-turn chat history (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	<pre>const msg = llm.createChatMessage({   role: llm.ChatRole.USER, text: "Hello!" }); (Source: <a href="https://docs.oracle.com">docs.oracle.com</a>)</pre>

MEMBER (SUITESCRIPT)	ALIAS / ROLE	PURPOSE / DESCRIPTION	EXAMPLE USAGE (SEE REFS)
<code>llm.getRemainingFreeUsage()</code>	–	Returns (number) of remaining free generation (text/QA) requests this month (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	<pre>const left = llm.getRemainingFreeUsage();</pre>
<code>llm.getRemainingFreeEmbedUsage()</code>	–	Returns (number) of remaining free embedding requests this month (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	<pre>const embLeft = llm.getRemainingFreeEmbedUsage();</pre>
<i>Enums:</i> <code>llm.ChatRole</code> , <code>llm.ModelFamily</code> , <code>llm.EmbedModelFamily</code> , <code>llm.Truncate</code>	–	Enumerations for roles (USER/ASSISTANT), model name families (e.g. COHERE), embedding model families, and input-truncation strategies (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ) (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).	E.g. use <code>llm.ChatRole.USER</code> , <code>llm.ModelFamily.COHERE_COMMAND</code> or <code>llm.EmbedModelFamily.COHERE_EMBED</code> .

The above table (drawn from Oracle’s SuiteScript documentation (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)) highlights that N/LLM covers the full set of generative-AI tasks one expects: **text generation, prompt templating, streaming output, and semantic embeddings**. It also provides object types (e.g. `llm.Citation`) and a promise-based API for all asynchronous calls, per SuiteScript conventions. For instance, instead of `generateText()`, one could use `generateText.promise()` to obtain a JavaScript Promise for asynchronous workflows (see **N/llm Module Members** in the docs (Source: [docs.oracle.com](https://docs.oracle.com))).

Oracle’s documentation emphasizes key points: generative AI “uses creativity,” so responses must be validated for accuracy (Source: [docs.oracle.com](https://docs.oracle.com)), and those features are region-restricted (available only where OCI Generative AI is set up) (Source: [docs.oracle.com](https://docs.oracle.com)). It accurately states: “Oracle NetSuite isn’t responsible or liable for the use or interpretation of AI-generated content” (Source: [docs.oracle.com](https://docs.oracle.com)), meaning implementers must ensure quality and compliance. The N/LLM module is automatically available in any account with Server SuiteScript enabled (Source: [docs.oracle.com](https://docs.oracle.com)), and the complete list of methods and objects (as shown in Table 1) is documented on Oracle’s help center.

In addition to Oracle’s official docs, third-party sources and partners have summarized N/LLM’s purpose. For example, NetSuite partner GURUS Solutions calls it an “AI co-pilot for SuiteScript” that “connects your scripts to Oracle Cloud Infrastructure’s generative AI services,” enabling on-demand content generation and RAG利用 (Source: [gurussolutions.com](https://gurussolutions.com)). GURUS’s overview highlights the same bullet points: generate text, evaluate stored prompts, manage prompts/text enhancements, feed documents for RAG, produce embeddings, and track usage (Source: [gurussolutions.com](https://gurussolutions.com)). This aligns closely with the official features, reinforcing that N/LLM brings broad AI “smarts” into NetSuite scripts.

## Integration with Other SuiteTools

N/LLM is not standalone; it works alongside other SuiteCloud tools:

- **Prompt Studio:** A UI for creating and storing dynamic prompts (with variables) in NetSuite. N/LLM’s `evaluatePrompt()` method directly interfaces with this. Developers can design prompts in Stuido and call them in scripts, ensuring consistency of model settings and the ability to update prompts outside code.
- **Text Enhance:** A UI/UX feature (e.g. a SuiteApp) that lets end-users generate text in record fields. While Text Enhance actions operate in the UI context, advanced scripts can mimic or augment them. The docs refer to managing “Text Enhance actions” via `N/record` and N/LLM (Source: [docs.oracle.com](https://docs.oracle.com)), meaning scripts can programmatically create or run these actions.

- SuiteQL and Searches:** Since NetSuite data must be prepared for the LLM, the module is often used in tandem with **N/query** (SuiteQL) or **N/search** to retrieve data. A performance note from Oracle's developer sample: using SuiteQL over N/search can improve efficiency when fetching data to feed into LLM documents (Source: [blogs.oracle.com](https://blogs.oracle.com)). In practice, scripts will query data (customers, transactions, etc.), format it (often as natural-language sentences or CSV), then use `llm.createDocument` to wrap that context text.
- AI Preferences:** NetSuite's AI Preferences (under Company Setup) provides a UI for administrators to set up AI usage. It allows enabling OCI credentials (for unlimited mode) and displays the **AI Usage** subtab. This tab lists one row per usage type (e.g. "Generate – Cohere" or "Embed – Cohere"), showing monthly free limits and consumption (Source: [docs.oracle.com](https://docs.oracle.com)). The documentation explicitly instructs how to view this table (Source: [docs.oracle.com](https://docs.oracle.com)). For governance, it's important that each SuiteScript pool (main account vs. SuiteApps) has separate counters.
- Concurrent Editing and Debugging:** N/LLM calls can be tested in the SuiteScript debugger. Oracle notes that example code uses `require` so it can be pasted into the debugger; production scripts should use `define` for proper entry points (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). This is a standard SuiteScript practice but is worth noting for developers new to the API.

Overall, N/LLM is deeply integrated into the SuiteCloud ecosystem. It leverages NetSuite's data retrieval capabilities and UI forms, yet introduces entirely new functionality through these generative API methods.

## Governance and Usage Limits

NetSuite implements strict governance over N/LLM to prevent abuse and control costs. This section details the usage limits, tracking, and concurrency rules, drawing on Oracle's documentation and best-practice guides.

### Monthly Usage Quotas

Every NetSuite account with AI features enabled gets a **monthly pool of free usage** for N/LLM. The exact size of the pool is not publicized, but it is fixed per account and resets at the start of each month (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). The key points:

- Separate Pools:** There are **two distinct pools**: one for generative calls ("Generate – Text") and one for embedding calls ("Embed – Text") (Source: [docs.oracle.com](https://docs.oracle.com)). Usage in each category is tracked separately, and the AI Preferences page shows them on separate rows (Source: [docs.oracle.com](https://docs.oracle.com)). For example, if you made 10 `generateText` calls and 5 `embed` calls in a month, the table would show two rows with those used counts against each respective limit.
- Tracking:** The SuiteScript **AI Preferences subtab** displays a **Usage Limit** (free quota) and **Used Quantity** for each month/type (Source: [docs.oracle.com](https://docs.oracle.com)). It notes: "All successfully completed actions count, but error responses don't" (Source: [docs.oracle.com](https://docs.oracle.com)). Scripts themselves can check usage at runtime via `llm.getRemainingFreeUsage()` and `llm.getRemainingFreeEmbedUsage()` (Source: [docs.oracle.com](https://docs.oracle.com)). (The Sales Insights sample suitelet even shows the remaining quota on its form using `getRemainingFreeUsage()` (Source: [blogs.oracle.com](https://blogs.oracle.com))).
- Unlimited Mode (OCI Credentials):** To go beyond the free pool, an account must link to an Oracle Cloud account with Generative AI service. This is done in AI Preferences – the admin furnishes OCI API credentials. Once set up, all N/LLM calls are billed to that OCI account instead of the internal pool (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). In this "unlimited" configuration, the SNPS say a script may actually include OCI config parameters in its calls. Essentially, NetSuite simply proxies requests through to OCI with those credentials and usage comes out of the customer's cloud subscription.
- SuiteApp Considerations:** If a third-party SuiteApp is installed and contains SuiteScript that uses N/LLM, Oracle gives each SuiteApp its own independent pool (Source: [docs.oracle.com](https://docs.oracle.com)). This means a SuiteApp's AI usage does **not** count against the main account's pool (and vice versa). The docs explicitly explain: "each SuiteApp installed in your account gets its own separate monthly usage pool for N/LLM methods...ensur[ing] SuiteApps can't use up all your monthly allocation and block your own scripts" (Source: [docs.oracle.com](https://docs.oracle.com)). These SuiteApp pools are not visible on the AI Preferences page (the app vendor must provide monitoring), but they follow the same rules internally.
- Tracking APIs:** As mentioned, the N/LLM API exposes helper methods to query usage. Calling `llm.getRemainingFreeUsage()` returns a number of text-generation requests left (Source: [docs.oracle.com](https://docs.oracle.com)); similarly, `llm.getRemainingFreeEmbedUsage()` for embeddings (Source: [docs.oracle.com](https://docs.oracle.com)). These may be used (for example) to disable features in a script if quotas are low, or to display warnings.

**Table 2 (below)** summarizes the usage governance:

USAGE TYPE	MONTHLY FREE QUOTA	CONCURRENCY LIMIT	NOTES
<b>Text Generation / Prompt Eval</b>	Fixed free pool (resets monthly) (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> )	5 concurrent calls (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> )	Generative methods ( <code>llm.generateText</code> , <code>evaluatePrompt</code> , etc.) share this pool. To increase, supply OCI credentials (unlimited billing) (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).
<b>Embedding</b>	Fixed free pool (resets monthly) (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> )	5 concurrent calls (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> )	Embedding ( <code>llm.embed</code> ) uses a separate pool. Unlimited usage via OCI credentials as well. Each pool is shown separately in AI Preferences (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ) (Source: <a href="https://docs.oracle.com">docs.oracle.com</a> ).

Table 2. N/LLM usage governance rules (Oracle SuiteScript docs).

In addition to these quantitative limits, NetSuite enforces **regional and account availability rules**. Not all accounts worldwide have AI enabled by default: generative AI features are available only in certain data centers and regions (those where the OCI GenAI service has been set up for NetSuite) (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)). For example, the documentation lists coverage across North America, Europe, Asia, etc., but notes that an account's settings and user language can also affect AI availability (Source: [docs.oracle.com](https://docs.oracle.com)). Administrators should consult NetSuite's "Generative AI Availability" help topic for details per country and language.

Finally, Oracle explicitly disclaims liability: N/LLM is treated as a creative tool, and all output must be validated by the user (Source: [docs.oracle.com](https://docs.oracle.com)). This is consistent with industry practice. As Gartner cautions, the primary barrier to AI adoption is often "estimating and demonstrating value" (Source: [www.gartner.com](https://www.gartner.com)), which means users must carefully measure and govern these new capabilities. Houseblend's analysis similarly emphasizes that successful AI projects are **highly targeted** and GDPR- or privacy-compliant, warning that unfocused pilots often fail (Source: [houseblend.io](https://houseblend.io)). NetSuite implements governance to enforce such discipline: quotas and tracking ensure AI scripts remain bounded and transparent.

## SuiteScript Code Examples

Below we show representative SuiteScript code snippets illustrating how to use N/LLM methods. These examples draw heavily from NetSuite's official documentation (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com)) and developer blogs (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [blogs.oracle.com](https://blogs.oracle.com)). Each snippet uses either the `define` or `require` pattern (the docs allow `require` for quick testing in the debugger; production scripts use `define` ).

### 1. Simple Text Generation

The most basic example sends a "Hello World" prompt. The code below (modeled on NetSuite's help topic) requires the `N/llm` module and calls `generateText` with parameters, then reads the returned text and remaining usage:

```

require(['N/llm'], function(llm) {
  const response = llm.generateText({
    // If no modelFamily is specified, NetSuite uses a default (Cohere Command A)
    prompt: "Hello World!",
    modelParameters: {
      maxTokens: 100,
      temperature: 0.2,
      topK: 3,
      topP: 0.7,
      frequencyPenalty: 0.4,
      presencePenalty: 0
    }
  });
  const responseText = response.text;
  const remainingUsage = llm.getRemainingFreeUsage(); // calls left this month
  log.debug('LLM Output', responseText);
  log.debug('Remaining Free Calls', remainingUsage);
});

```

This snippet (adapted from NetSuite's docs) shows how to call the default LLM with a simple prompt. The response object's `text` field has the AI-generated output and `llm.getRemainingFreeUsage()` returns how many free calls remain (Source: [docs.oracle.com](https://docs.oracle.com)).

## 2. Evaluating a Prompt Studio Prompt

If you have a prompt defined in NetSuite's Prompt Studio, you can use it directly. For example, suppose we have a prompt ID `stdprompt_gen_purch_desc_invnt_item` (a standard prompt for generating purchase descriptions from inventory item data). The code to evaluate this prompt and supply variables is:

```

require(['N/llm'], function(llm) {
  const response = llm.evaluatePrompt({
    id: 'stdprompt_gen_purch_desc_invnt_item',
    variables: {
      // These variable names and structure come from the Prompt Studio definition
      "form": {
        "itemid": "My Inventory Item",
        "stockdescription": "This is the stock description of the item.",
        "vendorname": "My Item Vendor Inc.",
        "isdropshipitem": "false",
        "isspecialorderitem": "true",
        "displayname": "My Amazing Inventory Item"
      },
      "text": "This is the purchase description of the item."
    }
  });
  const aiText = response.text;
  const left = llm.getRemainingFreeUsage();
  log.debug('Prompt Response', aiText);
  log.debug('Remaining Calls', left);
});

```

This example (from NetSuite's script samples (Source: [docs.oracle.com](https://docs.oracle.com)) loads a stored prompt by ID and supplies the required fields. The returned `response.text` is the LLM's answer based on the prompt logic. The remaining usage is again checked.

### 3. Retrieval-Augmented Generation

Below is a simplified illustration of using RAG. Imagine we have retrieved some records and formatted them into text (e.g. "Item: X, Q1 Sales: 1000..."). We create `Document` objects with `llm.createDocument()` and supply them to `generateText`. The LLM can cite from these docs in its answer:

```
require(['N/llm'], function(llm) {
  // Example structured data prepared as documents
  const documents = [];
  // Suppose we have two data points:
  const docData1 = "Item: Widget A, Qty Sold: 150, Revenue: $3000";
  const docData2 = "Item: Widget B, Qty Sold: 120, Revenue: $2500";
  const doc1 = llm.createDocument({ id: "doc1", data: docData1 });
  const doc2 = llm.createDocument({ id: "doc2", data: docData2 });
  documents.push(doc1, doc2);

  // User's natural question:
  const prompt = "Which item had more quantity sold?";
  const response = llm.generateText({
    prompt: prompt,
    documents: documents,
    modelParameters: { maxTokens: 50, temperature: 0.2 }
  });

  const bestItem = response.text; // e.g. "Widget A had more sales."
  const citations = response.citations; // e.g. [{subRoute: "doc1"}, {subRoute: "doc2"}, ...]
  log.debug('Answer', bestItem);
  log.debug('Citations', citations);
});
```

In this model, we fed two documents into `generateText`. The response text should answer based on those documents ("Widget A..."), and the `response.citations` array will include references like `doc1` and/or `doc2` if the LLM used them. This matches Oracle's RAG example pattern (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [blogs.oracle.com](https://blogs.oracle.com)).

### 4. Generating and Using Embeddings

To show how embeddings work, consider generating embeddings for a list of item names and comparing them for similarity:

```

require(['N/llm'], function(llm) {
  // List of items (for demonstration)
  const items = ["Ultra Widget Pro", "Super Widget X", "Deluxe Gadget"];
  // Get embeddings using Cohere embed model
  const embeddingResponse = llm.embed({
    embedModelFamily: llm.EmbedModelFamily.COHERE_EMBED,
    inputs: items
  });
  const vectors = embeddingResponse.vectors; // array of embedding arrays
  // Compute cosine similarity manually (example utility)
  function cosine(a, b) {
    let dot=0, normA=0, normB=0;
    for (let i=0; i<a.length; i++) {
      dot += a[i]*b[i];
      normA += a[i]*a[i];
      normB += b[i]*b[i];
    }
    return dot / (Math.sqrt(normA)*Math.sqrt(normB));
  }
  // Compare first item against others
  for (let i = 1; i < vectors.length; i++) {
    const sim = cosine(vectors[0], vectors[i]);
    log.debug('Similarity Item0-Item'+i, sim);
  }
});

```

This snippet shows `llm.embed` usage (adapted from Oracle's "Find Similar Items" example (Source: [docs.oracle.com](https://docs.oracle.com)) (Source: [docs.oracle.com](https://docs.oracle.com))). We request embeddings for multiple inputs. The returned `vectors` can be used in custom logic, such as computing cosine similarities to find which items are semantically closest.

The code examples above illustrate the core programming model for N/LLM. In a real application, you would typically combine these calls with data retrieval (`N/query` or `N/search`) and form handling (`N/ui/serverWidget`) as needed. For instance, the Sales Insights suitelet in Wilman Arambillete's Oracle blog first runs a SuiteQL query (using `N/query`) to summarize sales data, then formats each result line into a string and calls `llm.createDocument` on it (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [blogs.oracle.com](https://blogs.oracle.com)). Finally it calls `llm.generateText` to answer the user's question in context. That blog's sample is an end-to-end case of RAG built atop N/LLM (Source: [blogs.oracle.com](https://blogs.oracle.com)) (Source: [blogs.oracle.com](https://blogs.oracle.com)).

## Code Example: Streaming

For completeness, here is a brief glimpse at the *streamed* form of generation:

```
require(['N/llm'], function(llm) {
  const streamed = llm.generateTextStreamed({
    prompt: "Explain the sales trend in Q1 and Q2.",
    modelParameters: { maxTokens: 200 }
  });
  // The StreamedResponse object can be fetched piecewise:
  let textSoFar = "";
  while (!streamed.done) {
    const chunk = streamed.getText(); // new text chunk
    textSoFar += chunk;
    // (In practice, use streamed.done callback/promise)
  }
  log.debug('Full Response', textSoFar);
});
```

In this pattern, `generateTextStreamed` returns a `StreamedResponse`. The script can repeatedly call `getText()` to receive text as it is generated, which is useful for long outputs or keeping a responsive UI. (See Oracle's docs for usage of streamed methods (Source: [docs.oracle.com](https://docs.oracle.com)).)

## Governance, Data, and Security Considerations

### AI-enabled Governance

Using AI in an enterprise ERP raises governance issues around **cost, accuracy, and compliance**. NetSuite addresses cost via the usage quotas already described. Accuracy and appropriateness must be handled by implementers: Oracle explicitly warns that generative outputs are “creative” and may be inaccurate (Source: [docs.oracle.com](https://docs.oracle.com)). Consequently, scripts should always validate AI results before writing them into records. As one consultant note puts it: “Validate AI output: generative AI is powerful but not always 100% accurate. Always validate before using AI-generated content in production” (Source: [gurussolutions.com](https://gurussolutions.com)). Houseblend similarly emphasizes that generative AI projects must be narrowly scoped and governed: it cites research showing >95% of unfocused AI pilots fail to deliver ROI (Source: [houseblend.io](https://houseblend.io)), underscoring the need for careful planning.

Data privacy is another concern. Any data passed to the LLM (prompts or documents) is sent to Oracle's cloud. Customers in regulated industries must ensure sensitive information (PII, HIPAA data, etc.) is handled appropriately. Oracle does not specifically document data residency of OCI GenAI, but notes only accounts in certain regions can use generative features (Source: [docs.oracle.com](https://docs.oracle.com)) (regions where OCI GenAI is available). In practice, companies may choose to remove sensitive fields from prompts or use an on-premise endpoint when it becomes available. (NetSuite has not currently announced on-prem options for N/LLM; it relies on OCI.) Organizations should treat N/LLM as they would any integration: secure logging, encryption, and governance policies apply.

On the positive side, N/LLM can improve data security by **normalizing and sanitizing** user-generated content. For example, Text Enhance actions allow freeform text input to be “enhanced” or paraphrased by AI before storing, which might reduce arbitrary inputs. But this cuts both ways: if misconfigured, AI might inadvertently introduce sensitive info (weapons-of-mass-destruction style “hallucinations” are rare but possible). The Houseblend analysis explicitly notes that robust enterprise **data governance** is crucial when embedding AI (Source: [houseblend.io](https://houseblend.io)). NetSuite accounts should integrate N/LLM usage into their existing data governance frameworks: for instance, only trusted scripts should make AI calls, and logs of AI interactions should be auditable.

### Performance and Concurrency

Because LLM calls involve external services, they can incur latency. The concurrency limit (5 calls) is likely in part to prevent accidental DoS by misbehaving scripts. In practice, developers should design workflows to avoid hitting the concurrency limit. For example, a script should wait for one LLM call to finish before making another, rather than firing five at once. The SDK enforces the limit, so if a sixth call is made before one of the first five returns, NetSuite will throw an error (e.g. `LLM_REQUEST_LIMIT_REACHED`). Scripts can catch these errors, or check remaining usage/concurrency via `runtime.getCurrentScript().getRemainingUsage()` (SuiteScript governance) and pacing techniques.

Also note: SuiteScript has its own CPU/governance units meter (different from AI usage). Calling LLM APIs does **not** consume SuiteScript governance units, but scripts still must abide by overall execution limits. Long-running LLM calls will block a script's execution thread. For heavy use cases, NetSuite recommends using asynchronous SSR (SuiteScript Scheduled/MapReduce scripts) rather than client or user-event scripts to avoid timeouts.

## Pricing Implications

Although NetSuite does not charge extra license fees for generative AI features, customers may face costs in two ways:

1. **SuiteCloud Platform Fees:** Any significant increase in script usage (e.g. embedding many documents every hour) could require upgrading to a higher SuiteCloud bundle if it exceeds the free tier (SuiteFlow units, etc.). (Note: this is hypothetical - as of writing NetSuite does not meter AI calls beyond the free pool unless using OCI.)
2. **Oracle Cloud Costs:** If organizations opt to connect via their own OCI account for “unlimited” mode, they will pay OCI rates for generative AI usage. OCI GenAI pricing currently charges per token or per 1K tokens depending on model (Source: [docs.oracle.com](https://docs.oracle.com)). NetSuite customers should consider this when enabling unlimited mode, as heavy usage can generate significant cloud costs. On the other hand, since NetSuite’s free pool presumably suffices for light to moderate use, many customers might operate entirely on the included tier.

Oracle’s official stance, however, is to ensure customers aren’t surprised by AI costs. Administrators should check the AI Preferences regularly to see if free usage is near exhaustion. The suitelet example above conveniently displayed the `llm.getRemainingFreeUsage()` on the form (Source: [blogs.oracle.com](https://blogs.oracle.com)). For production, one could similarly alert users if quotas are low. The AI Preferences table also has a “Confirm” checkbox for opting into unlimited mode, making it an explicit decision.

## Case Studies and Use Cases

While N/LLM was only introduced in 2024, several case studies and pilot stories illustrate its value. Some are based on NetSuite partners’ insights (e.g. Houseblend’s August 2025 report) and public customer anecdotes.

- **BirdRock Home (Retail Case Study):** A mid-market retailer with over 40,000 SKUs, BirdRock Home “processes thousands of orders daily” within NetSuite (Source: [houseblend.io](https://houseblend.io)). They have rich sales and inventory data centralized in NetSuite and have used the SuiteAnalytics Warehouse (NAW) for forecasting and churn modeling. Houseblend notes that BirdRock’s use of ML in NAW led to “actionable product strategy improvements” (e.g. adjusting inventory based on predicted churn) (Source: [houseblend.io](https://houseblend.io)). Building on this, we can envision BirdRock using N/LLM to automatically generate product descriptions, classify customer reviews into sentiment categories, or answer natural-language queries like “Which product lines are trending up this quarter?” grounded in their own data (via RAG). The integrated nature of NetSuite means BirdRock’s financials, orders, and suppliers all feed into these AI answers, offering cross-functional insights that a standalone tool couldn’t glean.
- **Overture Promotions (B2B Supply Chain):** This promotional products distributor used NetSuite Analytics Warehouse to derive predictive sales insights for ordering and planning. According to Houseblend, Overture cited significant supply-chain optimizations from these forecasts (Source: [houseblend.io](https://houseblend.io)). Moving forward, Overture might use N/LLM to automate procurement narratives: for example, generating a weekly emailed summary of stock levels (“Item X is low relative to historical sales; reorder suggested”) or a Q&A bot answering queries like “What were last month’s top 5 sales reps by revenue?”. They could feed NAW reports into LLM prompts or provide their own documents (e.g. weekly sales CSVs), ensuring that the AI operates on accurate internal statistics.
- **Chatbots and Intelligent Assistants:** Several NetSuite partners have prototyped chatbots for common ERP tasks. For instance, one could embed a chatbot on the NetSuite dashboard that answers user questions about vendor details, outstanding invoices, or product specifications, using N/LLM. A sample scenario: A user asks “What is the status of PO#1234 and who do we owe money to?” The SuiteScript backend would retrieve the relevant purchase order and vendor records, format them as context, and call `generateText` or a stored prompt to produce a coherent answer. This could dramatically improve user productivity compared to navigating multiple record pages.
- **Content Enrichment:** Given that NetSuite often stores minimal descriptive text (item short names, etc.), generative AI can **enhance data quality**. For example, a custom script could run nightly to fill in or improve item descriptions: it might take existing fields (category, vendor specs, past sales notes) and call `generateText` to append a rich product overview. A partner guide suggests using N/LLM for “generating product descriptions, summaries, or automated replies” (Source: [gurussolutions.com](https://gurussolutions.com)). Similarly, generate-facsimile content can be used to auto-fill invoice comments or customer follow-ups. Of course, any AI-generated content should be checked (perhaps by another user) before final use.
- **Intelligent Search and Recommendations:** By embedding text (customer names, product titles, issue keywords) into vectors, NetSuite users could build semantic search features. The “similar items” example shows how to find related items; this could be extended to find related transactions or suggest upsell items. For instance, when viewing an order, the system could suggest “other customers who bought similar product also bought...” by comparing embedding vectors of customer profiles or product descriptions. This moves beyond exact keyword search to meaning-based retrieval, a valuable enhancement in large catalogs.

These examples illustrate the **paradigm shift** possible with N/LLM: moving from static data to dynamic, AI-driven insights. Industry research supports the potential. For instance, McKinsey (2023) found generative AI delivers most value in areas like customer operations and sales (Source: [www.mckinsey.com](http://www.mckinsey.com)) – precisely NetSuite use cases. They also emphasize “augmenting” human work: generative AI can handle content generation and routine analysis, freeing domain experts for strategic work (Source: [www.mckinsey.com](http://www.mckinsey.com)). NetSuite’s CFO article underscores this in finance: instead of accountants manually drafting narrative reports, AI can draft them automatically, allowing finance teams to “focus on more strategic tasks” (Source: [the-cfo.io](http://the-cfo.io)) (Source: [the-cfo.io](http://the-cfo.io)).

**Data & Statistics:** Beyond case anecdotes, broader data reinforce the trend:

- A 2025 Gartner survey (as noted) found GenAI most-deployed by enterprises (Source: [www.gartner.com](http://www.gartner.com)).
- Industries from banking to retail see \$400B+ annual value potential from GenAI (Source: [www.mckinsey.com](http://www.mckinsey.com)).
- As Houseblend notes, NetSuite’s integrated data can be a goldmine: using all business data for AI can provide insights “far beyond what point solutions offer” (Source: [houseblend.io](http://houseblend.io)).
- ROI is a critical concern. Houseblend cites research showing lack of focus dooms 95% of AI pilots (Source: [houseblend.io](http://houseblend.io)). NetSuite’s design (with RAG and internal data) directly addresses that: by using *company data*, AI answers stay relevant, making success more achievable.

## Implications and Future Directions

The introduction of N/LLM has broad implications:

- **For Developers:** SuiteScript programmers must learn AI best practices (prompt engineering, result validation, error handling). They now have a new kind of tool in their toolbox, blurring the line between coding and writing training data. Development frameworks will evolve. For instance, triggers could automatically invoke LLMs (e.g. augment data on record-save). Developers will need to pay attention to cost/limit management and provide UIs to let end-users act responsibly. Some consultants foresee N/LLM becoming a standard skill in SuiteScript development, akin to how `N/email` or `N/record` are today.
- **For Business Users:** Non-technical roles can potentially use AI-enabled features for routine tasks. NetSuite’s default AI-based modules (financial anomaly detection, chat assistants) will handle many common needs at no additional cost. However, organizations must train their users on proper usage: AI auto-fill can be helpful, but blind trust is risky. In usage guidelines, end-users should be advised that AI suggestions are “helpful drafts” not final answers. Governance committees may need to oversee which scripts can run N/LLM (similar to how automation or scripting is gated).
- **Security and Compliance:** N/LLM introduces new vectors. Data sent to LLM may need classification (e.g. PII redaction). Responses should be audited; for regulatory compliance, logs of AI usage may be necessary. Future releases may include more controls (for example, enforcing encryption in transit to OCI, or on-prem model hosting) to address sensitive use cases. Enterprises should align N/LLM use with their existing data governance and ethics policies.
- **Technological Evolution:** As LLM technology advances, N/LLM may support new model families (e.g. GPT-4o in OCI, vision models). Oracle’s future AI roadmap (e.g. announcements at SuiteWorld 2025/2026) will likely expand N/LLM’s capabilities. For example, integration of image or code generation (e.g. `llm.generateImage` or `llm.executeCustomAction`) could appear. The existence of streaming methods (`generateTextStreamed`) suggests support for multimodal or longer-form generation is already in place. We also anticipate better developer tools (e.g. SuiteCloud IDE enhancements, local debugging simulation of AI calls).
- **Industry Perspective:** NetSuite’s N/LLM is part of a larger shift: essentially every major business application is becoming AI-native. According to Gartner, by 2030 ~80% of enterprise software will be AI/ML-driven (Source: [www.gartner.com](http://www.gartner.com)). Oracle’s strategy to bake AI into its platforms positions NetSuite to compete with Microsoft’s Dynamics and Salesforce’s Einstein on an even footing. Jacovljevic (Technology Evaluation Centers) opined that NetSuite’s AI could make it “on par, if not even better, than Microsoft D365” for mid-market ERP (Source: [www.techtarget.com](http://www.techtarget.com)), especially as Oracle integrates AI across both the NetSuite and Oracle Fusion product families. This suggests continued investment: expect more AI tools within NetSuite (planning, procurement, CRM) to come.
- **Challenges:** Several challenges remain. Hallucinations remain a practical issue: despite RAG, LLMs can blend data in unpredictable ways. Ensuring that, for example, financial narratives do not fabricate numbers is critical. NetSuite’s approach of providing citations (via `Citation` objects) helps trace answers, but it’s ultimately on developers. Versioning and auditing of prompts (and documenting LLM usage) will be important for compliance. Additionally, performance (latency), and scaling to high call volumes are concerns for large customers; these may drive features like batch processing or asynchronous LLM calls.

- Ethical and Governance Outlook:** In terms of ethics, the integration of AI into ERP underscores the need for enterprise AI ethics policies. Many frameworks (e.g. EU AI Act guidelines) will apply to features like N/LLM. NetSuite's multi-tenant model complicates things: Oracle must ensure that one customer's data cannot bleed into another's AI context. OCI Generative AI is shared infrastructure, but Oracle presumably keeps tenants isolated (Oracle has not publicly detailed multi-tenancy model for GenAI). Customers should verify that AI usage complies with data residency or regulatory rules (using region constraints if needed).
- Analyst and Expert Opinions:** Industry voices have weighed in. As noted, Holger Mueller praised the platform approach (Source: [www.techtarget.com](http://www.techtarget.com)). Predrag Jakovljevic observed that generative AI will "learn from [the customer's] context and vernacular" over time (Source: [www.techtarget.com](http://www.techtarget.com)), implying fine-tuning from usage. Others caution that simply having the API is not enough; organizations need data scientists and AI strategists to make the most of it. The large investment in AI by Oracle signals that generative features will only grow in NetSuite.

## Conclusion

NetSuite's N/LLM module represents a significant leap: it brings state-of-the-art LLM capabilities directly into the ERP's extension framework. Developers can now write SuiteScript that leverages OCI's generative AI models for text generation, summarization, semantic search, and more, all tuned to the company's own data. This opens up numerous use cases – from chatbots and intelligent search to automated report writing – that were previously infeasible inside an ERP.

However, with great power comes responsibility. Effective use of N/LLM requires attention to governance (quotas and concurrency), data quality (garbage in, garbage out), and compliance (security and privacy). NetSuite has built-in safeguards (usage tracking, quotas, disclaimers), but organizations must also establish their own practices (data validation, user training, audit trails). The transition to AI-enriched processes will also change roles: business analysts may rely more on "AI copilots" to analyze data, and developers will need AI literacy.

From a strategic perspective, N/LLM solidifies Oracle NetSuite's position in the AI-enabled ERP landscape. By including these features at no extra cost and tightly integrating them with the unified data model, NetSuite leverages its core advantage (centralized data) to offer richer insights and automation. Customers who adopt these tools stand to improve productivity and decision-making speed. For example, the ability to ask natural-language questions of ERP data (as shown in the Sales Insights suitelet) is akin to having a 24/7 expert on hand.

Looking ahead, we expect continuous enhancement of NetSuite's generative AI toolkit. Upcoming releases may introduce better models (more powerful LLMs), expanded Prompt Studio capabilities, and possibly AI-driven analytics. The design of N/LLM – with streaming support, chat constructs, and document citations – shows that Oracle is thinking ahead to complex AI use patterns.

In sum, the NetSuite N/LLM module is a robust, enterprise-ready bridge to modern AI. It is backed by comprehensive documentation and governed usage policies, yet flexible enough for creative solutions. By following best practices and constraints outlined above, organizations can safely harness LLMs to unlock the full potential of their NetSuite data. All claims and details here are substantiated by Oracle's own help documentation (Source: [docs.oracle.com](http://docs.oracle.com)) (Source: [docs.oracle.com](http://docs.oracle.com)) and credible industry sources (Source: [www.mckinsey.com](http://www.mckinsey.com)) (Source: [www.gartner.com](http://www.gartner.com)), ensuring that this report can serve as a reliable reference for technical and strategic decision-making.

**References:** Official Oracle documentation (SuiteScript 2.x Generative AI APIs) (Source: [docs.oracle.com](http://docs.oracle.com)) (Source: [docs.oracle.com](http://docs.oracle.com)) (Source: [docs.oracle.com](http://docs.oracle.com)) (Source: [docs.oracle.com](http://docs.oracle.com)), Oracle developer blogs (Source: [blogs.oracle.com](http://blogs.oracle.com)) (Source: [blogs.oracle.com](http://blogs.oracle.com)), NetSuite community guides and partner blogs (Source: [gurussolutions.com](http://gurussolutions.com)) (Source: [houseblend.io](http://houseblend.io)), and industry reports (Source: [www.mckinsey.com](http://www.mckinsey.com)) (Source: [www.gartner.com](http://www.gartner.com)) (Source: [www.gartner.com](http://www.gartner.com)) have been used extensively. All cited data and quotes are from these sources.

---

Tags: netsuite, suitescript 2.1, n/llm module, generative ai, rag, llm governance, oracle oci, vector embeddings, api limits, erp ai

### DISCLAIMER

This document is provided for informational purposes only. No representations or warranties are made regarding the accuracy, completeness, or reliability of its contents. Any use of this information is at your own risk. Houseblend shall not be liable for any damages arising from the use of this document. This content may include material generated with assistance from artificial intelligence tools, which may contain errors or inaccuracies. Readers should verify critical information independently. All product names, trademarks, and registered trademarks mentioned are property of their respective owners and are used for identification purposes only. Use of these names does not imply endorsement. This document does not constitute professional or legal advice. For specific guidance related to your needs, please consult qualified professionals.