IMPROVING ACCESSIBILITY OF OPERATIONAL AVALANCHE BULLETINS THROUGH A MODEL CONTEXT PROTOCOL (MCP) SERVER

Norbert Lanzanasto[1]*, Christoph Stanger[2], Bengt Haunerland[1], and Simon Legner[3]

[1] *Avalanche Warning Service Tyrol, Innsbruck, Austria*
[2] *Google, Zirl, Austria*
[3] *TBBM, Innsbruck, Austria*

ABSTRACT: Timely access to official avalanche bulletins is critical for safe decision-making in avalanche terrain. While avalanche warning services publish structured and text-based bulletins, these formats are primarily optimized for direct human consumption. Large language model (LLM) based tools offer new opportunities to improve accessibility by enabling interactive, multilingual, and user-adapted safety assessments. However, a key technical challenge is ensuring LLM responses are strictly grounded in real-time authoritative forecasts to prevent safety-critical hallucinations. To address this, we developed the Albina Model Context Protocol (MCP) server, which exposes published avalanche bulletins directly to LLMs as secure, machine-readable resources and tools. Built upon the Canadian Avalanche Association Markup Language (CAAML) standard, our server acts as a standardized middleware layer that retrieves the latest forecasts at query time. We propose a unified global aggregation model where a single centrally operated server aggregates CAAML-compliant feeds from multiple regional services. This model reduces technical friction for developers, minimizes operational costs for forecast organizations, and preserves strict authority control under the official warning services. Prototype integrations demonstrate that conversational assistants can reliably summarize hazard levels, resolve location queries, and retrieve operational blog posts while remaining strictly grounded in the official, currently active bulletin.

KEYWORDS: Avalanche Bulletin, Model Context Protocol, Large Language Models, CAAML, Information Accessibility

1. INTRODUCTION

Operational avalanche warning services publish critical safety bulletins daily to help backcountry recreationists and professionals make informed choices in avalanche terrain. While these forecasts are highly detailed, their traditional dissemination channels (such as web portals and custom mobile applications) are designed primarily for direct human reading and manual retrieval. This presentation format limits rapid integration with automated trip-planning platforms and modern artificial intelligence (AI) decision-support systems.

Large Language Models (LLMs) and conversational assistants present a powerful opportunity to bridge this gap. By enabling interactive, user-adapted, and multilingual access to avalanche safety bulletins, LLMs can summarize complex warning structures and answer specific user terrain questions. However, deploying conversational agents in safety-critical domains poses a severe risk of model hallucinations or reliance on outdated training datasets. To ensure public safety in avalanche terrain, AI applications must be strictly grounded in real-time, official forecasts issued by authorized warning agencies.

2. SYSTEM ARCHITECTURE & EXPOSED TOOLS

To establish a secure, real-time connection between operational avalanche forecast data and conversational AI, we developed the Albina Model Context Protocol (MCP) server (accessible at https://extra.avalanche.report/mcp). The Model Context Protocol is an open standard designed to structure how LLM applications interact with external data sources and tools. The Albina MCP server serves as standardized middleware that converts structured forecasts into LLM-readable contexts at query time.

* *Corresponding author address:*
Norbert Lanzanasto, Avalanche Warning Service Tyrol, Wilhelm-Greil-Strasse 17, 6020 Innsbruck, Austria;
tel: +43 512 508 7730
email: norbert.lanzanasto@tirol.gv.at

Developed in TypeScript, the server exposes the following specialized tools and prompts to connected LLM clients:

- get_bulletin: Retrieves the latest detailed avalanche report for a specified micro-region.
- get_bulletin_micro_region: Resolves geographic locations by mapping mountain names or regions to their corresponding avalanche bulletin micro-regions.
- get_blog_posts: Retrieves recent operational blog posts from warning services. This tool provides LLMs with access to detailed qualitative descriptions of snowpack development, forecaster field reports, and weather event analyses, supplementing the structured bulletin with vital context.

3. GEOGRAPHIC RESOLUTION AND COORDINATE TRANSLATION

To support natural language queries regarding specific geographic locations or coordinate inputs, the Albina MCP server implements a robust geographic resolution workflow. When a user provides a place or mountain name, the server integrates with the OpenStreetMap Nominatim API for geocoding. Host systems must specify the NOMINATIM_USER_AGENT environment variable to authenticate with the Nominatim API; if omitted, the program reports an error and refuses to start.

For direct GPS coordinate inputs, the server leverages an coordinate-to-region loader component. This component ingests precise latitude and longitude coordinates to map the user's location directly to the corresponding avalanche warning micro-region. This localized mapping ensures that backcountry recreationists receive precise safety advisories matching their exact geographic position.

4. CENTRALIZED GLOBAL AGGREGATION MODEL

A primary challenge for developers integrating avalanche safety data into AI assistants is the technical fragmentation of endpoints across different countries and regional services. To overcome this, the Albina MCP architecture leverages the internationally established Canadian Avalanche Association Markup Language (CAAML) standard. CAAML provides XML schema definitions designed to facilitate the uniform exchange of public avalanche bulletin information.

Rather than requiring each individual forecasting service to build, host, and maintain its own AI-compatible middleware, we propose a centralized global aggregation model. A single, centrally operated MCP server instance aggregates CAAML-compliant bulletins from participating regional warning services, presenting a single, unified access point for LLM applications. This approach minimizes integration friction for developers, eliminates operational overhead for individual warning services, and preserves absolute authority control by retrieving unaltered, official forecasts directly from the responsible agencies.

5. DISCUSSION AND CONCLUSION

Standardizing access to operational avalanche bulletins via the Albina MCP server marks a significant advancement in safety communication. By integrating both structured bulletins (via get_bulletin) and descriptive, event-based forecaster blogs (via get_blog_posts), the server enables conversational AI assistants to deliver a highly synthesized, yet completely grounded, safety assessment. The incorporation of coordinate-based mapping guarantees precise localization in the field. Ultimately, this framework establishes a secure, reliable, and accessible interface that bridges the gap between complex snow science data and public safety needs.