AVALANCHEALLEY.COM
A PROGRAM DATABASE & RISK COMMUNICATION TOOL

Richard W. Steiner, Jr.[1*], Adam M. Clark[1], and Calvin Hill[2]

1 Snowslip LLC, Whitefish, MT, USA
2 Allied Computing, Kalispell, MT, USA

ABSTRACT: In 2010 avalanche forecasters with the BNSF Avalanche Safety Program (Program) began developing the Interactive Avalanche Information System (IAIS). BNSF Avalanche Safety was founded in January 2005 to assist the Railway with reducing avalanche risk. Its operational area is along the southern boundary of Glacier National Park, John F. Stevens Canyon, located in northwest Montana, USA. The IAIS was originally developed as an in-house, computer-based avalanche atlas and database where forecasters could enter and manage data. An oral presentation explaining the initial development of the IAIS was given at the 2012 International Snow Science Workshop in Anchorage, Alaska. The proceedings paper can be found at https://arc.lib.montana.edu/snow-science/item.php?id=1703. Program staff transitioned away from the IAIS beginning in 2012 and began developing a user-interactive, web-based platform (website), www.avalanchealley.com. This transition allowed forecasters to manage data remotely on various electronic devices while also utilizing the website for risk communication. Development of the website had three main objectives: (1) to allow forecasters to electronically record industry standard avalanche, snowpack, and weather data, (2) display these data, free of charge, for online viewer dissemination, (3) to integrate these data with a user-interactive, web-based avalanche atlas. Now, 15 years later, the Avalanche Alley website continues to operate efficiently and has become a robust Program risk communication tool. Keeping the website running has required enduring software glitches, programming issues, a web developer switch, and a completely new platform for managing both the website and database.

KEYWORDS: Avalanche Alley, IAIS, Data, Data Management, Web Atlas.

1. BACKGROUND

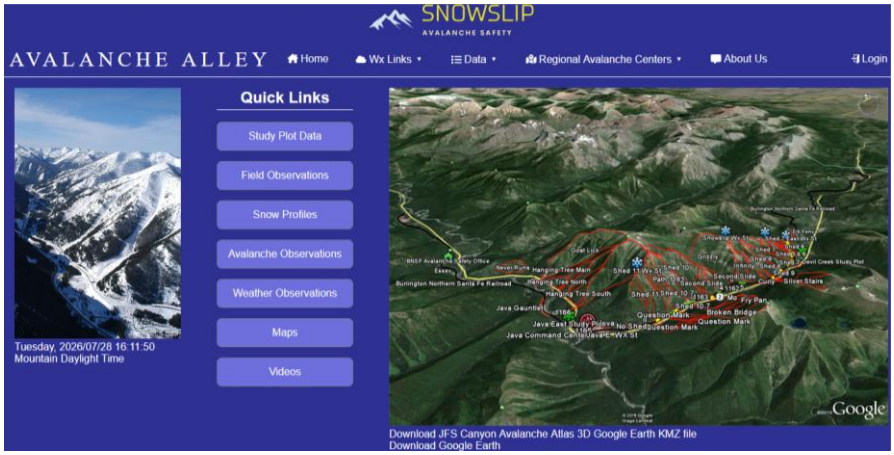
In January 2005 BNSF Railway (Railway) established an avalanche safety program into its operations. The Program is located in Essex, Montana, USA and its purpose is to reduce avalanche risk for Railway elements during winter months. The Program's operational area is along the southern boundary of Glacier National Park in John F. Stevens Canyon, of northwest Montana, USA. **(Hamre, 2004)**. Program education, forecasting, and mitigation is conducted by a contractor, currently Snowslip LLC, that works closely with Railway management.

Retrieval and recording of Program data associated with avalanche snowpack and weather observations began in 2006 and parameters utilized for collection were based on avalanche industry standards derived from the United States Snow, Weather & Avalanche Guidelines – SWAG **(Greene, 2022)**.

Not long after Program establishment, it became apparent that electronically storing and geo-referencing the Program's data was critically important. Beginning in 2010, and with the assistance of contractor, Mobile Logistic Mapping (MLM), forecasters developed an in-house electronic database that stored data and incorporated that into a web-based atlas utilizing Goggle Earth (GE).

In 2012, Program forecasters decided the configuration of the Program's database needed to be updated where data would be stored in "the cloud" rather than on specific computer(s). To do this, again with the assistance of MLM, a Program website, **www.avalanchealley.com** (AA), was developed. This configuration allowed forecaster to enter data remotely and data entered became accessible to all website visitors, free of charge. Data was stored on a Network Solutions server and redundantly backed-up by MLM. The AA web atlas also became accessible on the website via a website-linked GE KMZ file **(Steiner, 2012)**.

*Corresponding author address:
Richard W. Steiner Jr., Snowslip LLC,
PO Box 4203, Whitefish, MT 59937-4203;
tel: +1 406-212-0588
email: ted.steiner@avalanchealley.com



AvalancheAlley.com

2. WEBSITE & WEB ATLAS

2.1 *Data entry and utilization*

Data on AA has been utilized by Program forecasters for in-house forecasting products and statistical analysis. Data has also been regularly accessed and evaluated for regional avalanche forecast products produced by forecasters with the Flathead Avalanche Center. Public users have utilized AA data for general information, decision-making, and winter activity planning. Avalanche researchers have also used AA data, especially now that Program records extend back over 20 years. The United States Geological Survey (USGS) has utilized AA data for several research projects and published papers, see website: USGS Northern Rocky Mountain Science Center.

Observational data is updated regularly by Program forecasters and displayed chronologically and includes:

- **Avalanche Observations**. Avalanche characteristics are tabular defined and listed in columnar format.
- **Snowpack Profiles**. JPEG images generated from *Snow Pilot* data entry. Images are chronically listed and defined by individual rows.
- **Manual Weather Observations** (Snow Study Plot Data). Weather and snow-related data is tabular defined and listed in columnar format.
- **Field Observations**. Columnar defied and text format.

At the end of the 2025/26 season, Program observations totaled:

- Manual Weather Observations: 3521
- Avalanche Observations: 622
- Field Observations: 258
- Snow Profiles: 241

Data is displayed on the website and can be accessed by navigating to the AA homepage and then scrolling over the “Data” link located in the homepage header links. Once scrolled over, a drop-down menu is displayed and allows users to access the various databases previously referenced.

Additional data available on the AA website are links to:

- Hourly data of seven (7) automated weather stations in and near the Program Area.
- USGS records of historic avalanche events and dendrochronology studies.
- Regional web resources related to avalanche, snowpack and weather conditions.

2.2 *Website search filter*

When the data tab is scrolled over on the homepage, a dropdown menu is activated and allows users to access a variety of datasets. Once a set is selected, a table displays the associated data. If a user is looking for more detailed information, they can choose the displayed data and search for specifics. This is a powerful database search tool and has become increasingly useful as the AA database keeps increasing in size.

2.3 *Web atlas*

The AA web atlas (atlas), which utilizes GE as a base layer, allows users to access three-dimensional imagery of the Program Area terrain and provides a graphical overlay of GE generated shape files that identify Program Area attributes. A link to activate the Program's atlas (KMZ File) is located at the bottom of the AA homepage. Access to the atlas is available to all that visit the AA website.

To access the web atlas requires that:

- GE software (available online) is already downloaded on a computer.
- The Program Area KMZ file is activated on the AA website.

Some attributes in the atlas are user-interactive and allow users to mouse over and click on them for related attribute information. For example, all identified Program Area avalanche paths are named and represented by distinct polygons. Each polygon can be activated with a computer mouse click. When activated, a particular path's atlas data is displayed along with links to that path's associated avalanche, snowpack, and weather data that have been previously entered into the AA database.

Additional interactive atlas attributes include:

- Program area weather station icons: Mouse click displays hourly weather collected at these sites.
- Snow study plot icons- Mouse click displays manual snow and weather observations collected on a regular basis.
- Rescue cache and Helicopter Landing Zone location icons- Mouse click displays coordinates and contents.

2.4 *Web atlas utility*

The web atlas has proven to be a powerful presentation tool that allows atlas users to interact with the graphically displayed Program area terrain and attributes. This ability has been essential when giving presentations and has also been beneficial when utilizing the atlas in conjunction with conference calls, meetings, and Railway operational briefings.

For example, if presenting on a conference call regarding Program area avalanche conditions, forecasters can reference attributes in the atlas to convey pertinent information with those on the call. As long as attendees have GE and the atlas loaded on their computer, they can utilize that to simultaneously view information being referred to and have a graphical representation of the area in question. Also included with each GE avalanche path polygon are files that display Program Area avalanche atlas data along with links to recorded avalanche, snowpack, and weather observations.

3. WEBSITE MAINTENANCE, REPAIR, REBUILD

3.1 *Maintenance and repair*

To keep the AA web atlas, website, and database up to date and functioning properly since 2010 has required regular monitoring of the web-based products and function of the operating software. This has been accomplished by regularly monitoring the website, ensuring data entry and displays were functioning properly, and making sure data was redundantly backed up in independent servers.

Operational issues that have occurred throughout the tenure of both the website and web atlas have been associated with outdated operational software and associated coding. These issues have oc-

curred due to changes and updates on the server site and web host, while the software and associated code used to build the site remained static. The resulting web-based problems were related to the inability to upload and store data in the database. In the winter of 2022, this situation was frustratingly obvious as data related to weather observations could not be entered into the database from the web portal. After a substantial amount of time and effort in trying to troubleshoot the cause, it was found to be related to outdated development software that no longer was recognized by the website host server.

It became clear throughout 2023 that the operational software used to build the AA platform was completely outdated and it needed to be rebuilt. However, doing so was a logistical complexity as MLM was no longer in business and it was imperative to find a web developer that understood the original software and coding utilized to develop the AA website, database, and the GE integration of the web atlas.

Thankfully, in the fall of 2023, a web development company, Allied Computing, was identified as a potential business that could assist with an AA platform rebuild and associated updates. In the meantime, Peter Petri, the owner of MLM graciously put together a detailed list of everything related to the development of the original AA platform and its current state. Peter made this list available for Allied Computing to utilize as necessary for rebuild assistance. In 2025, Calvin Hill, owner of Allied Computing, accepted the rebuild project and AA's new platform began taking shape.

3.2 *Website rebuild*

Allied Computing modernized the AA operating platform end-to-end. The application, website, now runs on current, supported versions of MySQL and PHP, replacing a stack that had fallen years behind and was increasingly a liability for security patches and long-term support.

Every database query in the application was rewritten rather than patched in place, and the front-end HTML / CSS was optimized for a faster, more consistent experience across pages.

The site itself now runs on a custom-built MVC framework that cleanly separates routing, business logic, and display — making future updates faster to build, easier to test, and far less likely to introduce regressions elsewhere in the site.

Alongside this, Allied Computing built a new administration system that gives Program forecasters a straightforward way to update website content and data directly, without needing a developer to make routine changes.

Allied Computing also completely rebuilt the Google Earth KMZ map (atlas). The previous version had accumulated duplicate and inconsistent entries over time; the new atlas consolidates that data into a single clean source, pulls live data directly from the website instead of relying on static exports, and adds several newly mapped avalanche paths, linked directly to the corresponding website records.

The most substantial change in this update is to the underlying database structure itself. Rather than a surface-level refresh, the database was restructured to fix long-standing data-integrity issues, remove a fragile and risky dynamic report-generation system, and lay a cleaner foundation for future growth. Details regarding all AA updates and changes are documented in detail on the following page.

3.3 *Rebuild and update details*

Database Updates:

1. Inventory-Level Changes
    1. Tables: +9 new, −2 removed, 45 carried over
    2. Views: −10 (all removed)
    3. Stored procedures: −25 (all removed)

    New tables: about_info, bad_access, failed_logins, lut_avy_rads, lut_boolean, regional_centers, reports, site_admins, weather_links
    Removed tables: media_files, reports_reportslist

Removed views: vw_images_all, vw_obs_av_all, vw_profiles_all, vw_reports_all, vw_reports_file, vw_reports_pub, vw_wx_sum_1_hnswe, vw_wx_sum_hn, vw_wx_sum_swe, vw_wx_summary_base
Removed stored procedures: all 25, including a large cluster of "dynamic SQL" helpers (lut_listMaker, lut_listTables_avy, lut_avy_CodeLabel, wx_ss_hn, wx_ss_swe, wx_SeasonTotals_inVar, reports_admin_list, reports_public_list, etc.)

2. **New Tables — What They Add**
   site_admins: Admin accounts with login / password and an active / access flag
   failed_logins: Records failed login attempts (username, IP, timestamp)
   bad_access: Rate-limiting / abuse tracking — repeated bad requests per IP, with attempt counts and indexed lookup
   reports: Replaces reports_reportslist— leaner report-definition table
   regional_centers: Ordered list of linked regional avalanche centers
   weather_links: Ordered, categorized list of external weather links
   about_info: Single free-text "about" block for the site
   lut_boolean: Generic Yes / No lookup values
   lut_avy_rads: New avalanche lookup list (code / label), referenced by a new rads column added to data_obs_avy

   All nine new tables use code / label / listorder-style or purpose-built columns consistent with the rest of the schema's lookup-table convention. Eight of the nine — all except lut_boolean — use InnoDB rather than MyISAM.

3. **Removed Tables — What They Took With Them**
   media_files — a generic file-metadata table (filename, path_sys, path_web, filesize, mime_type, uploader user / IP / time) that media_photos, media_videos, and data_obs_profiles all joined against via a fileID column. That column has been dropped from all three tables in the new schema, and media_obs_av no longer stores a separate media_id — it now stores a single avrid referencing the avalanche observation directly (previously it stored obs_id + media_id as a two-sided link).
   reports_reportslist — a table that stored not just report metadata but full hand-built SQL fragments as text (sql_select, sql_from, sql_where, sql_group, sql_ordby, report_sql), assembled by PHP or the stored procedures into ad-hoc queries. Replaced by reports.

4. **Reports Table Redesign: reports_reportslist → reports:**
   The old design, built queries by concatenating up to six stored text fragments at runtime (visible in the removed stored procedures and views, which did exactly this via PREPARE / EXECUTE). The new table collapses that into two SQL fields plus a JSON config field, removing the PHP-template path columns and the fragment-reassembly step entirely.

5. **Changes to Tables That Exist in Both:**

   A. **Systemic changes (repeated identically across ~40 tables)**
      Audit timestamp columns — every edcr_time / eded_time column changed:
      Type: timestamp → datetime
      Default: DEFAULT '0000-00-00 00:00:00' → DEFAULT CURRENT_TIMESTAMP
      eded_time additionally gained ON UPDATE CURRENT_TIMESTAMP in most tables (auto-refreshes on edit)
      Relaxed NOT NULL constraints — many columns that were NOT NULL with no default (edcr_user, edcr_ip, eded_user, eded_ip, and various data columns) became nullable (DEFAULT NULL).

   B. **Table-specific changes**
      The new database structure makes several targeted improvements to simplify the schema, reduce redundancy, and improve data integrity. Key changes include widening fields where needed, removing obsolete columns and duplicate audit information, and replacing manually maintained calculated values with stored generated columns for avalanche vertical elevation, crown-depth difference, and weather snow density, ensuring these values are always calculated consistently. Boolean-style avalanche

flags were reduced from int to tinyint, new avalanche rating support was added, and several fields were made nullable where the data is not always available.

The lut_avy_paths table was significantly redesigned to eliminate duplicated geographic path information and follow the consistent code / label lookup-table structure used elsewhere in the database. Media and profile tables were also simplified by removing dependencies on the obsolete media_files structure, while relationships in media_obs_av were consolidated from separate obs_id and media_id fields into a single avrid. Finally, several new tables were added to support administration, reporting, access logging, regional centers, weather links, and avalanche ratings, providing a cleaner and more scalable structure for the application.

C. **Tables with no real changes**
   The remaining ~30 tables (mostly lut_avy_*, lut_wx_*, and other lookup tables) needed no updates and are structurally identical apart from the cosmetic dump-format differences noted above.

6. **Engine and Charset**
   A. New tables created in the new schema use InnoDB.
   B. Default character set differs (utf8mb4 / utf8mb3 in the original dump vs utf8 in the new dump)
7. **Why the New Structure Is Better**
   A. **Removed a SQL-injection-prone report engine.** The old reports_reportslist table stored raw, hand-built SQL fragments (sql_select, sql_from, sql_where, …) that were reassembled and executed via dynamic PREPARE / EXECUTE in several stored procedures and views. Building queries by concatenating stored strings is inherently fragile and risky. The new reports table replaces that with two bounded SQL fields plus a JSON config column — far less surface area for malformed or malicious SQL, and no PHP-template path columns baked into the database.
   B. **Eliminated ~25 stored procedures built around dynamic SQL and information_schema cursors.** Many of the removed procedures existed purely to loop over tables matching a name pattern (lut_avy_%, lut_wx_%) and PREPARE / EXECUTE a constructed string — brittle logic that breaks silently if a table is renamed, and that's much easier to write, test, and secure in application code than in MySQL procedural SQL. Removing it moves that logic out of the database, where it's harder to review and version-control.
   C. **Removed 10 views that depended on the two tables being retired.** Since media_files and reports_reportslist are gone, the views built on them (vw_images_all, vw_reports_all, etc.) would have broken anyway; retiring them along with their source tables keeps the schema consistent.
   D. **Fixed broken audit timestamps.** DEFAULT '0000-00-00 00:00:00' is an invalid, sentinel-style date that MySQL only accepts in a legacy / relaxed SQL mode and that many drivers and ORMs mishandle. Every audit column now defaults to CURRENT_TIMESTAMP, and edit timestamps auto-update via ON UPDATE CURRENT_TIMESTAMP — so "last edited" is now reliable without application code having to set it explicitly.
   E. **Added computed columns for derived values.** av_vert_m, ch_diff_cm (in data_obs_avy), and wxnsdensity (in data_obs_wx) are now GENERATED ALWAYS AS (...) STORED instead of manually entered numbers. This guarantees the derived value can never drift out of sync with its source columns — previously a data-entry error or missed update could leave them inconsistent.
   F. **Normalized lut_avy_paths and simplified media_obs_av.** lut_avy_paths now matches every other lut_* table's simple code / label / listorder shape, removing duplicated data that already lived in data_avy_atlas. Similarly, media_obs_av now links to a photo / video directly through a single avrid column instead of a two-part obs_id / media_id pair, giving a clearer, more direct relationship and a simpler join for the most common query pattern (find media for an observation).
   G. **Trimmed a number of dead or duplicate columns.** Beyond the tables above, the new schema drops several columns that were unused or redundant with columns already present elsewhere in the same table, including data_usg_avoccraw.PathID, data_usg_dendro.d_rain / d_warming, and media_videos.EdUser / EdTime (a duplicate of that table's own edcr_user / edcr_time). Fewer columns per row means less data MySQL

has to scan on a full table read, which matters most on the larger observation tables as they continue to grow.

H. **Added operational / security tables that didn't exist before.** failed_logins and bad_access give the application a place to track and rate-limit abusive access, and site_admins formalizes admin accounts — none of which existed as first-class tables in the original schema. regional_centers, weather_links, and about_info similarly give newer site features (linked centers, external weather links, an about page) a dedicated table instead of being bolted onto an existing one, which keeps future changes to one feature from risking another.

I. **New tables use InnoDB, which is a better fit for that workload.** InnoDB supports row-level locking, transactions, and crash recovery; MyISAM only offers table-level locking and no transactions. Concretely, that means a write to failed_logins or bad_access (e.g. logging a login attempt) no longer locks out other readers / writers of that whole table the way a MyISAM write would, which matters for tables that are written to frequently and potentially concurrently. The tables best suited to that pattern — login / security tracking, admin accounts, report definitions — were built on InnoDB from the start.

4. SUMMARY

Avalanchealley.com, or what had been referred to as the IAIS, has now been in existence and functioning for 16 years. The web atlas is currently functioning as intended and data in the database continues to be updated on a regular basis. Attributes on the web atlas have been updated throughout the years to include additional resources, Program infrastructure, and avalanche paths. All data on the website remains available for dissemination and is available free of charge. There is no username or password needed to access any AA attributes.

5. ACKNOWLEDGEMENTS

Thank you to all who have assisted and supported AA throughout the years. In particular, thank you to management with BNSF Railway who have supported the creation and utilization of AA as a Program information tool and public education product. And, thank you to David Hamre, who supported the original creation of AA and provided objective feedback regarding the utility of AA as part of the Program. Huge gratitude to Peter Petri of MLM who was dedicated to creating the original AA platform and managed the web-based products for nearly 15 years. Finally, thank you to Calvin Hill of Allied Computing for taking on the task of the AA platform rebuild. The updates and rebuilding done by Calvin were logistically seamless, allowing AA to be utilized and remain online while the work was completed. Calvin’s work greatly improved AA dated functionality and its overall operational speed.

REFERENCES

Greene, E., et al., 2022. Snow, Weather, and Avalanche. 4th ed, American Avalanche Association, (pp. 11-85).

Hamre, D., Overcast, M., 2004. Avalanche Risk Analysis, John F. Stevens Canyon, Essex, MT., (pp.1-24).

Steiner, R.W., Dundas, M., Petri, P., 2012. Interactive Avalanche Information System (IAIS) in: Proceedings of the International Snow Science Workshop, Anchorage, AK, 16-24 September, 375-380.