

Brightside Daily Good News - Current Implementation Specification

Product: Brightside Daily Good News

Live path: 'https://mykindlyevents.com/dailynews'

Captured: September 3, 2026

1. Purpose

Brightside is a standalone, public-facing positive-news experience within Kindly Events. It gathers recent, family-safe good-news stories, creates original summaries, and presents them in a searchable and browsable feed. Its visual system, header, footer, and theme preferences are separate from the Kindly Events website chrome.

2. Public experience

Routes

Route - Purpose - Indexing

'/dailynews' - Main rolling news feed with search and category filters. - Indexable

'/dailynews/:slug' - Individual article detail page, source attribution, stats, related stories, bookmarking, and sharing. - Indexable

'/dailynews/explore' - Latest-batch stories grouped by geographic region in collapsed accordions. - Indexable

'/dailynews/bookmarks' - Browser-local saved-story list. - 'noindex, nofollow'

Feed behavior

- The feed calls 'GET /api/dailynews/articles?date=recent&limit=100'.
- The API anchors the feed to the newest stored 'batch_date' and returns stories in newest-first 'fetched_at' order.
- 'date=recent' includes articles whose batch date is within five days of the latest batch date.
- The server enforces the configured 'stories_per_day' value across pagination; the client cannot exceed it by requesting a larger limit.
- Category filtering, title search, offsets, and direct batch-date filtering are supported by the same endpoint.
- The displayed total is the number of articles returned in the current response, rather than the total matching records in storage.

Article detail behavior

- Each detail page fetches the full article and up to four related articles from the same category and batch.
- Body text is stored as paragraph-separated text and rendered as paragraphs with a decorative first-letter treatment.
- Article cards and detail pages use an optional stored image URL. When none is present, they render a category-specific gradient and icon instead.
- Source attribution is shown when source data exists. A verified article permalink uses the 'Read Article ->' call to action; a site-level source uses 'Visit Site ->'.
- Bookmarks are stored only in the visitor's browser under 'bs-bookmarks'; they are not associated with a user account and can disappear when browser storage is cleared. Bookmarks to deleted articles simply no longer resolve.

- Share uses the native sharing sheet when available and otherwise copies the current URL to the clipboard.

Theme, navigation, and legal notice

- Brightside has a scoped deep-violet, peach, and category-color design system using '--bs-*' variables in 'src/styles/globals.css'.
- The selected light/dark theme is stored in browser local storage under 'bs-dark'.
- The sticky Brightside header includes Feed, Explore, and Bookmarks navigation, a theme toggle, and a mobile navigation menu.
- The header uses the uploaded Brightside wordmark asset at a maximum rendered height of 40px.
- The footer identifies Brightside as a product of Infynity Labs LLC and includes a disclaimer dialog. The dialog supports click-outside close, Escape close, and explicit close/confirmation controls.

3. Content generation and refresh pipeline

Triggering

A refresh can be started manually through 'POST /api/dailynews/refresh' or automatically through the scheduler. The workflow requires the configured 'OPENAI_API_KEY' server secret.

Generation model

- The system uses the OpenAI Responses API with 'web_search_preview'.
- The configured model is supplied from Brightside settings; the currently saved model is 'gpt-4o' with temperature '0.7'.
- The prompt instructs the model to search for recent stories from the preceding seven days, open each selected direct article permalink, and write an original summary based on the full article.
- The response is expected as a raw JSON array. The parser removes markdown fences, extracts an array when surrounded by extra text, and attempts recovery of complete top-level JSON objects if the response is truncated.

Editorial rules

The prompt requires stories to be uplifting, constructive, verifiable, recent, original in wording, and appropriate for all ages. It excludes violence, conflict, crime, accidents, injury, adult or romantic content, alcohol/drugs/tobacco, politics, religion, tragedy, disaster, death, promotional content, opinion, and unverified claims.

The prompt assigns geographic quotas for the United States, India, Asia excluding India, and the rest of the world. It requests coverage across configured categories.

Validation and de-duplication

Before any retention cleanup, a response must produce at least one row with a title, snippet, and body. Empty or unusable results preserve all existing articles.

For each candidate article, the pipeline:

1. Normalizes and bounds text fields before storage.
1. Accepts an article source URL only when it is not a homepage or single-segment section path.
1. Uses a source-publication homepage or Google source search as a fallback when a usable direct article URL is not supplied.
1. Persists the source type as 'article', 'site', or 'null'.
1. Rejects rows with a duplicate source URL already stored.
1. Rejects headlines that share at least 60% of meaningful words with an existing headline.
1. Inserts unique rows individually and records failures without stopping the entire batch.

Refresh completion and purge order

Retention purge is the final operation of a refresh. It runs only when at least one new article has been successfully inserted.

After successful insertion, rows with a 'batch_date' older than the configured 'article_retention_days' cutoff are deleted. A refresh that yields only duplicates, has no valid content, or cannot save any rows skips purge and leaves the existing feed intact.

The current code uses the date-based retention setting in this final purge. Although the settings UI includes a 'retention_cap', the current refresh implementation does not enforce that maximum-total-articles setting.

4. Scheduled refresh

- Scheduler implementation: 'node-cron' in 'src/server/lib/brightside-scheduler.ts'.
- The scheduler starts after the Express server begins listening and reads saved Brightside settings from 'app_settings'.
- It supports daily or weekly jobs and schedules them in UTC.
- Saving settings calls 'applySchedulerSettings', which starts, stops, or reschedules the task without a server restart.
- Current saved schedule: enabled, daily, at '06:00 UTC'.

5. Administration

The Brightside control page is available at '/admin/brightside' within the application admin area. It provides:

- Latest-batch status and article count.
- Manual refresh.
- Two-step flush-all deletion.
- Category add/remove controls.
- Article-length selection.
- Stories-per-day selection.
- A retention-cap field.
- Geographic-distribution controls with a 100% total indicator.
- OpenAI model and temperature controls.
- Automatic refresh enable/disable, daily/weekly frequency, weekly day, and UTC hour controls.

The page is behind the client-side admin route gate, which checks '/api/admin/me' after authentication.

Current source observation: the Brightside mutation handlers themselves ('POST /api/dailynews/refresh', 'POST /api/dailynews/flush', and 'POST /api/dailynews/settings') do not currently apply the available 'requireAdmin' middleware in their route registration or handler code. The admin page is protected, but the API endpoints should be treated as needing server-side authorization hardening before relying on them as admin-only controls.

6. API contract

Method and path - Description - Key behavior

'GET /api/dailynews/articles' - Read feed articles. - Supports 'date', 'category', 'region', 'search', 'limit', and 'offset'. 'date=recent' is a five-day rolling window from the newest batch; default is newest batch only.

'GET /api/dailynews/articles/:slug' - Read one article and related stories. - Returns '404' when absent; related stories match category and batch.

'GET /api/dailynews/regions' - Return region counts for the newest batch. - Used by Explore; returns an empty list when no content exists.

'GET /api/dailynews/settings' - Read configured settings merged with defaults. - Publicly reachable; includes categories, model, schedule, retention, and geography settings.

'POST /api/dailynews/settings' - Save settings and apply scheduler changes. - Merges request values with stored and default values.

'POST /api/dailynews/refresh' - Run a content refresh. - Returns batch date and inserted count; needs the OpenAI secret.

'POST /api/dailynews/flush' - Delete all news articles. - Permanently removes every 'news_articles' row.

7. Data model

'news_articles'

Field - Purpose

'id' - Auto-increment primary key.

'slug' - Unique route identifier derived from title plus timestamp-based suffix.

'title', 'snippet', 'body', 'why_it_matters' - Display copy and long-form article content.

'category', 'region', 'country' - Classification and browsing metadata.

'source_name', 'source_url', 'source_url_type' - Attribution data and CTA behavior.

'stats_json' - JSON-encoded list of up to three displayed stat tiles.

'image_ref' - Optional source image reference.

'published_at', 'fetched_at', 'batch_date' - Original publication timing, ingestion time, and refresh-batch anchor.

Indexes support batch-date, category, region, and slug queries.

'app_settings'

Brightside configuration is stored as a JSON string under the 'brightside_settings' key. The stored values are merged with defaults at read time.

Current saved configuration:

- Categories: Environment, Wildlife, Community, Health, Science, Innovation, Sports, Charity, Pets, Inspirational.
- Article length: medium.
- Stories per day: 50.
- Retention days: 7.
- Retention cap: 100 (present in settings but not currently enforced by refresh logic).
- Geographic distribution: 25% US, 25% India, 25% Asia excluding India, 25% rest of world.
- Automatic schedule: enabled, daily, 06:00 UTC.

At the time this specification was captured, storage contained 56 articles with batch dates ranging from September 1 to September 3, 2026.

8. SEO

- Feed, Explore, and article-detail pages use server-compatible Helmet metadata for titles, descriptions, canonical URLs, and basic Open Graph/Twitter card values.
- Bookmarks is intentionally excluded from indexing.

Brightside Daily Good News - Current Implementation Specification

- Brightside content is currently client-fetched after hydration rather than supplied through route loaders, so article content is not present in the initial server-rendered HTML.
- The feed and article pages declare canonical URLs under 'https://mykindlyevents.com/dailynews'.

9. Current implementation boundaries

- Brightside article content is persisted in MySQL; bookmarks and theme preferences are local to each browser.
- Article images are optional source references, not a dedicated image-ingestion pipeline.
- Article visibility uses a fixed five-day rolling feed window, while data deletion uses the configurable date-based retention window. These are intentionally separate behaviors in the current code.
- The configured retention cap is currently a UI/persisted setting only; it has no deletion logic.
- The public settings endpoint exposes configuration metadata.
- Server-side protection should be added to Brightside mutation routes before considering manual refresh, settings changes, and flush operations secure.