

Technical Support Center: Optimizing Historical BGP Routing Data Storage

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: AS1928370

Cat. No.: B12371349

[Get Quote](#)

This guide provides researchers, scientists, and drug development professionals with best practices and troubleshooting for the storage and analysis of historical Border Gateway Protocol (BGP) routing data, notionally for an autonomous system like **AS1928370**.

Frequently Asked Questions (FAQs)

Q1: What is the standard format for historical BGP data, and where can I source it?

Historical BGP data is typically stored in the Multi-threaded Routing Toolkit (MRT) format, as defined in RFC6396.^[1] This format captures BGP routing information, including routing table dumps (RIBs) and BGP update messages. Prominent public archives where you can obtain this data include the University of Oregon's Route Views project and the RIPE NCC's Routing Information Service (RIS).^{[1][2]} These archives collect data from various BGP peers worldwide, offering a comprehensive view of the internet's routing landscape.

Q2: My raw BGP data files are enormous. What is the first step to make them more manageable?

The primary challenge with historical BGP data is its volume. A crucial first step is compression. Due to the high redundancy in BGP update messages, where often only timestamps differ between consecutive messages, the data is highly compressible.^[2] Standard compression utilities like Gzip or Bzip2 are effective. For more advanced, query-friendly storage, consider chunked compression, where data is compressed in smaller blocks to allow for faster access to specific time ranges without decompressing the entire file.^[2]

Q3: I'm trying to query prefix announcements for a specific time range, but it's extremely slow. How can I optimize this?

Querying raw, compressed MRT files for specific time ranges or prefixes is inefficient because it requires a full sequential scan. To optimize this, you need to preprocess the data and build an indexed database. A common approach involves creating a back-end database with indices for prefixes and AS paths.[2] Tools like BGP-Inspect utilize B+ tree indices to significantly speed up queries.[2] Open-source frameworks like BGPStream are also designed for efficient processing of both live and historical BGP data, abstracting away the complexities of data parsing and filtering.[3][4]

Q4: What are the key differences between a RIB dump and a BGP update message?

A Routing Information Base (RIB) dump is a snapshot of a BGP router's complete routing table at a specific point in time. It provides a comprehensive view of all the preferred routes to reach various destinations. BGP update messages, on the other hand, capture the dynamics of routing changes. They announce new routes, withdraw old ones, or update path attributes. For historical analysis, using both is crucial: RIB dumps provide the baseline state, while update messages provide the granular changes between these snapshots.

Q5: How can I detect routing anomalies like hijacks or leaks in the historical data?

Detecting anomalies involves analyzing patterns in BGP announcements. For instance, a BGP hijack can be identified when an AS unexpectedly starts originating a prefix it doesn't own. A route leak might be suspected if an AS starts announcing routes learned from one peer to another, violating routing policies. To detect these, you can monitor for changes in the origin AS for a given prefix or sudden, significant increases in the number of unique origin ASes for a set of prefixes.[3] Tools that process historical data can be scripted to look for these specific patterns.

Data Storage Optimization Strategies

Choosing the right storage strategy depends on the scale of your data and the types of queries you intend to run. Below is a comparison of common approaches.

Storage Strategy	Pros	Cons	Best For
Compressed MRT Files	Simple to store; preserves original data format.	Extremely slow for analytical queries; requires full-file decompression.	Archival purposes; initial data staging.
Time-Series Database (e.g., InfluxDB, TimescaleDB)	Optimized for time-based queries; good compression for time-series data.	May not be ideal for complex graph-based queries (e.g., AS path analysis).	Monitoring routing stability over time; analyzing update frequency.
Columnar Database (e.g., ClickHouse, Apache Druid)	High-performance analytical queries; excellent data compression.	Can be complex to set up and maintain; requires a well-defined schema.	Large-scale, fast aggregations and statistical analysis of routing data.
Custom Indexed Database	Highly optimized for specific BGP query patterns (prefix, AS path).[2]	Requires significant development effort; less flexible for ad-hoc queries.	Building specialized BGP analysis tools with high-performance requirements.

Experimental Protocols

Protocol 1: BGP Data Ingestion and Indexing Pipeline

This protocol outlines the steps to process raw MRT files into an indexed database for efficient querying.

Objective: To create a queryable database of historical BGP data.

Methodology:

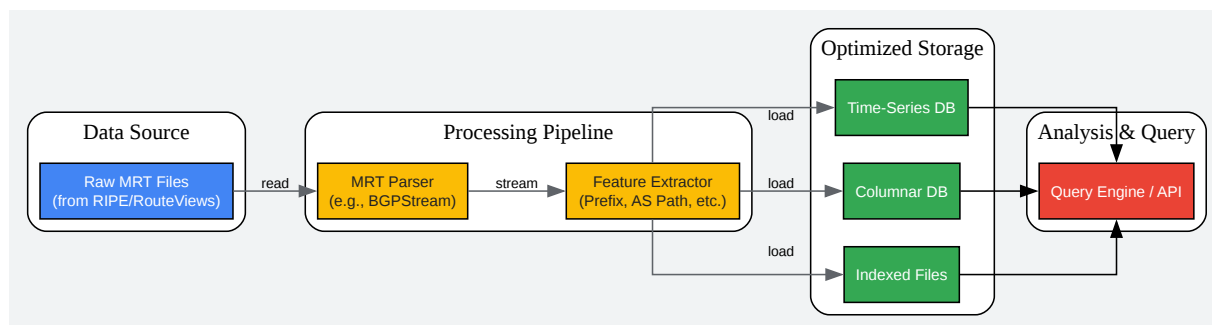
- **Data Acquisition:** Download historical BGP update and RIB dump files in MRT format from a public repository (e.g., RIPE RIS or Route Views).
- **Data Parsing:** Use a library capable of parsing MRT files (e.g., the one provided with BGPStream) to extract individual BGP messages.

- Data Extraction: For each BGP message, extract key fields:
 - Timestamp
 - Peer AS
 - Prefix (and prefix length)
 - Origin AS
 - AS Path
 - Community Attributes
 - Message Type (Announcement or Withdrawal)
- Database Loading:
 - Choose a suitable database backend (e.g., a columnar or time-series database).
 - Define a schema that maps the extracted fields to database columns.
 - Load the parsed data into the database in batches.
- Indexing:
 - Create indices on the most frequently queried columns: timestamp, prefix, origin_as, and as_path.
 - For AS path queries, consider specialized indexing techniques if your database does not support array indexing efficiently.
- Verification: Run a set of predefined queries (e.g., find all announcements for a specific prefix in a given month) to verify the integrity of the data and the performance of the indices.

Visualizations

BGP Data Processing Workflow

The following diagram illustrates the pipeline for ingesting and processing raw BGP data for analysis.



[Click to download full resolution via product page](#)

Caption: Workflow for ingesting and processing historical BGP data.

Decision Tree for Storage Optimization

This diagram presents a logical decision-making process for selecting an appropriate data storage solution.

Caption: Decision tree for choosing a BGP data storage strategy.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. LACNIC's BGP Collector [lacnic.net]
- 2. user.eng.umd.edu [user.eng.umd.edu]
- 3. researchgate.net [researchgate.net]

- 4. bigdama.ait.ac.at [bigdama.ait.ac.at]
- To cite this document: BenchChem. [Technical Support Center: Optimizing Historical BGP Routing Data Storage]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b12371349#optimizing-data-storage-for-historical-as1928370-routing-data]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com