

Hari Prasath Ulaganathan

Data Analyst | Analytics Engineering, BI & Operations Analytics

Silverthorne, CO | (857) 468-8793 | hari8059@gmail.com | US work authorization: EAD (I-485)

linkedin.com/in/hari-prasath-u- | hari-94.github.io | github.com/hari-94

Summary

Data analyst with 4+ years of turning messy operational data into tools people use every day. Most recently replaced manual housekeeping scheduling and budget tracking at Grand Colorado on Peak 8 with Python applications, including a scheduling engine 41 colleagues use daily, saving about 990 hours and \$36,250 a year. Previously built a tested dbt and Snowflake pipeline behind a project-delivery dashboard at ThoughtSpot, and ran KPI reporting for the IT infrastructure team at CMA-CGM. Equally at home designing the data model and sitting with the supervisors who use the output. Seeking data analyst, BI or analytics engineering roles; open to relocation. Case studies at hari-94.github.io.

Skills

Languages & Databases: SQL, Python, R, PL/SQL, PostgreSQL, MySQL, Microsoft SQL Server

Analytics Engineering: dbt, Snowflake, Databricks, Airflow, Azure Data Factory, Azure Data Lake, AWS, ETL, dimensional modeling, data quality testing

BI & Reporting: Power BI, Tableau, Looker (LookML), ThoughtSpot, Qlik Sense, QlikView, Excel (Power Query, pivot tables), KPI design

Python & Statistics: Pandas, NumPy, SciPy, scikit-learn, Statsmodels, Matplotlib, Seaborn, regression, time series forecasting, clustering, NLP

Automation & AI: Google OR-Tools, Streamlit, Supabase, Power Automate, REST APIs, LLM agents, MCP, Git/GitHub, ReportLab, openpyxl, SharePoint

Experience

Internal Tools Developer (Housekeeping Operations), Grand Colorado on Peak 8 Nov 2025 – present

- Built a Python scheduling engine (custom bin-packing plus a Google OR-Tools solver) that assigns every Full Clean and Daily Service room, then automated it with Power Automate: it builds at 5 AM and lands in HotSOS and Teams with no manual steps. Used by 41 colleagues; saves 2+ hours and about \$50 a day (~\$18,000 a year).
- Added a three-week staffing forecast built from live bookings, plus per-person performance and room-timing reports, giving managers headcount and productivity figures they never had.
- Re-sequenced inspector assignments floor by floor after analyzing travel patterns, cutting unnecessary walking between floors during quality inspections.
- Built a budget application (Python, Streamlit, Supabase/PostgreSQL) that replaced an hour a day of spreadsheet entry and blocks duplicate expenses at the database level.
- Added a live budget dashboard with month-over-month trends, variance tracking, over-budget alerts and automated PDF reports (ReportLab), which helped cut spending by at least \$1,500 a month (~\$18,000 a year).
- Turned an 85-sheet, multi-year Excel scheduling workbook into 30,401 clean shift records for 181 employees with a Python ETL (openpyxl), then used trend, clustering and workload-balance analysis to inform scheduling policy.
- Deployed and maintain all three tools in production (GitHub, Streamlit Cloud).

Solutions Analyst, ThoughtSpot Aug 2022 – Jun 2023

- Turned client requirements into KPI definitions, report specs and on-schedule dashboards that client teams used for self-service analysis.
- Wrote SQL against Snowflake for ad-hoc analysis and dashboard metrics.
- Wrote the data analysis behind proposals that won new client engagements, and coached teammates on SQL and data-quality investigation.

ThoughtSpot ETL project: delivery analytics pipeline *Databricks · Snowflake · dbt · Python · SQL*

Replaced untested ad-hoc SQL with a layered, tested pipeline feeding the project-delivery dashboard.

- **Raw:** Python ingestion on Databricks pulling four operational sources (projects, resources, resource allocations, support tickets) into raw tables in Snowflake.
- **Staging (dbt views):** cleaning only. Types cast, free-typed status values mapped to a controlled vocabulary, Y/N flags turned into booleans.
- **Intermediate (ephemeral dbt models):** business rules defined once, pricing each allocation at the resource's rate and measuring ticket resolution hours against P1/P2/P3 SLA targets held in a macro.

- **Marts (dbt tables):** project and resource dimensions, allocation and ticket facts at a declared grain, and a pre-aggregated project-health table the dashboard reads as is, so every report shows the same number.
 - **Tests and reports:** 51 data tests (uniqueness, referential integrity, accepted values, ranges) and two singular tests stop bad data at build time; Python scripts on the marts generate the recurring KPI reports.
- PMO Senior Executive, CMA-CGM** Jun 2020 – Nov 2021
- Built the IT infrastructure team’s KPI dashboards and published weekly and monthly status reports on resource allocation and delivery.
 - Spotted delivery risks in project and resource trends and escalated them early, so blockers were resolved faster.
 - Ran contractor invoicing, SharePoint reporting and documentation, and first-round hiring interviews.
- Management Trainee, Risk Management, CMA-CGM** Jul 2019 – Jun 2020
- Handled ad-hoc analysis, project reporting and asset tracking for leadership in the risk management team.

Projects

- AI Job-Search Agent** *LLM agent · MCP · Python · REST API · Task Scheduler*
- Designed an LLM agent that runs 10+ nationwide searches each morning through Indeed and ZipRecruiter MCP connectors, screens each posting against written rules (experience, pay, recency, clearance), skips repeats and sends ranked matches with a fit-and-gap note to WhatsApp: 40 vetted matches in its first six days.
 - Moved 186 scheduled reminders from GitHub Actions, whose cron ran hours late, to Windows Task Scheduler, and traced intermittent 403 errors to the message gateway’s command-injection filter, fixing them and verifying the fix with a dry run over every message.
- Self-Service BI, CMA-CGM** *Looker · Snowflake · SQL*
- Built LookML Explores and metric definitions on curated Snowflake tables so PMO and infrastructure teams could pull ticketing and resource KPIs themselves, and profiled source data to fix upstream quality issues.
- Stock Data Pipeline and Strategy Backtest** *Python · Pandas · Snowflake* github.com/hari-94/ETL
- Python ETL from the Yahoo Finance API into Snowflake: 7-day moving average in Pandas, table created on first run, a suspended warehouse resumed automatically, credentials kept in environment variables.
 - Backtested a moving-average crossover strategy against buy-and-hold on six stocks (notebook): buy-and-hold won on five of six, NVDA the exception at 21,349% against 18,864%.
- BI Dashboards** *Power BI · Tableau · Snowflake · DirectQuery* github.com/hari-94/PowerBI-Dashboard
- Power BI star-schema mart on Snowflake’s TPC-H data over DirectQuery, comparing spending and savings by country with Sankey and waterfall views, and a Tableau Public sales dashboard with year-over-year KPIs.
- Case studies:** scheduling engine, budget tracker and dbt warehouse, with animated walkthroughs, at hari-94.github.io

Education

- Colorado State University**, MS in Management and Computer Information Systems (MCIS) Aug 2023 – Dec 2024
- Great Learning**, PG in Data Science and Engineering Jun 2021 – May 2022
- SRM Institute of Science and Technology, Chennai**, MBA in Systems and Marketing 2017 – 2019
- Rajiv Gandhi College of Engineering and Technology**, BE in Computer Science 2012 – 2016

Certifications

- ThoughtSpot Business Analyst Certification (Mar 2025)
- Databricks Academy: Lakehouse Fundamentals (Dec 2024), Generative AI Fundamentals (Feb 2025)