



The Gyre Analytics and Operations Hub: One Platform for Every Role

Author: Gyre Research | Analytics and Operations Hub

Contact: team@gyrereseach.com

Revised: 2026-06-13

Contents

Contents	1
The Gyre Assistant.....	2
Internal Data.....	4
External Data	4
The Portfolio Manager.....	5
The Risk Officer	6
The Performance & Investor-Reporting Team.....	8
The Compliance Officer.....	8
The Operations / Middle Office	9
The Fund Accountant / Controller	10
The Trading Desk.....	11
The Quant / Systematic Researcher	12
The IR / LP Reporting Team.....	12
Summary... Your Firm, at a Glance	13
What Runs Underneath	14
Getting Started	14
Appendix, The Full Capability Catalog	15

Every team in an investment firm needs the same thing, an accurate, current view of the book, and almost none of them get it from the same place. The portfolio managers work from the OMS, the risk team from a separate engine, fund accounting from the administrator's files, and compliance from an overnight extract. Spread across multiple prime brokers, strategies, and instruments, each role maintains its own copy of the truth, and the gaps between those copies become reconciliations, restatements, and arguments about whose number is right.

The Gyre Analytics and Operations Hub is built so that every role works from one valued book. The Gyre Hub assembles a single data object per portfolio, pricing, positions, transactions, and tax lots at the base, and runs 98 calculation modules across 12 disciplines on top of it. Below we walk the firm role by role: the problem each role lives with today, the Gyre Hub capability that addresses it, and what changes for that role once it is in place.

The Gyre Assistant

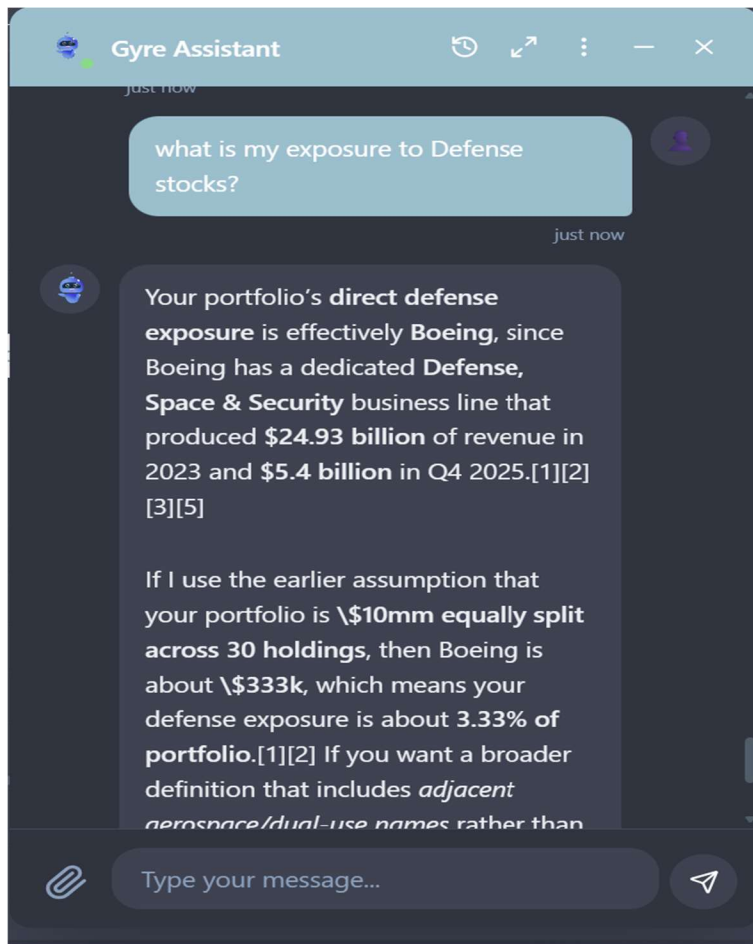
Every role gets the same analyst on call, one that already knows the book – Hosted entirely within your firm's isolated LLM environment, to prevent data bleed.

Before the firm splits into roles, it shares one tool. The Gyre Assistant is a natural-language interface available to every user, regardless of seat: ask a question in plain English and get an answer drawn from the same valued book everything else reads from. It handles questions about your activity, what was traded, when, and how positions changed; questions about your portfolio, exposures, concentrations, and on-the-fly calculations against your actual holdings; and questions about Gyre usage and calculation methodology, how any of the 98 modules define and derive the numbers they produce, so a figure can always be traced back to its arithmetic. The same plain English drives the screen: rather than working a report's filters by hand, a user can simply describe the view they want,

For informational purposes only. Not investment advice.
© 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved



narrow to a strategy, change the date range, regroup by sector, and the Assistant reshapes the report for them. Beyond the book, it reads the unstructured documents a firm lives in, administrator, custodian, and prime-broker reports, sell-side and internal research, and regulatory filings, and answers questions against their contents. Because it operates on each user’s entitled data, the same Assistant serves every role below; what differs is the portfolio and documents it can see, not the capability it offers.



For informational purposes only. Not investment advice.
 © 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved

Internal Data

Positions, Trades, Transactions, Fees, Cash, Borrow, Margin Accounting

The Gyre Hub is built to be inexpensive to adopt. Out of the box, users simply upload their daily reports, the same administrator, custodian, and prime-broker files they already receive, and the Gyre Hub parses them into the valued book. There is nothing to integrate and no data project to fund: if you can download your daily files, you can run the Gyre Hub.

For firms that would rather not handle files at all, the Gyre Hub offers managed data interfaces for a small monthly charge. Gyre builds and maintains a direct connection to your data providers, so the book stays current automatically without anyone uploading anything. We already operate interfaces with dozens of providers across the industry's primes, custodians, and fund administrators, including, representatively, custodians such as BNY, State Street, J.P. Morgan, Citi, and BNP Paribas; administrators such as SS&C, Citco, Northern Trust, HSBC, and Apex; and prime brokers such as Goldman Sachs, Morgan Stanley, Bank of America, UBS, and Barclays, and we add providers as clients bring them.

The two modes share one destination: whether a file is uploaded by hand or delivered over a managed interface, it lands in the same valued book every role reads from.

External Data

Reference, Security, Index, and Other Market Data

Where Internal Data covers the firm's own book, External Data is the market and reference layer the Gyre Hub maintains around it, collected, scrubbed, and normalized from over 30 exchanges, issuers, and agencies, and refreshed daily so analysis always runs against current data. As of the latest daily load, coverage spans:

- 250,078 securities under coverage, with 29,427,780 daily price and volume (OHLCV) observations stored from December 1, 2023 to the present and growing every day, with additional history loaded upon request
- 36 exchanges
- 246,172 issuers

For informational purposes only. Not investment advice.
© 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved

- Daily constituent coverage for 12 major indices, the S&P 500, 400, 600, and 100; the Russell 1000, 2000, and 3000; the Nasdaq 100; the Dow Jones Industrial Average; the Nikkei 225; the STOXX 600; and the S&P/TSX
- Over 1,000,000 global economic data series sourced from leading public economic-data providers
- Global reach across 257 countries and 179 currencies

Because every figure above grows with each daily load, these counts are a snapshot of an expanding dataset rather than a fixed catalog.

The Portfolio Manager

I want one current, risk-aware picture of how I am positioned, not five reports that disagree.

The problem: Across a multi-strategy, long/short, multi-prime book the PM's working view is stitched together from a positions screen, a separate exposure report, an attribution spreadsheet refreshed weekly, and a factor model that lives somewhere else. By the time the pieces are assembled, the market has moved and the numbers no longer reconcile with each other.

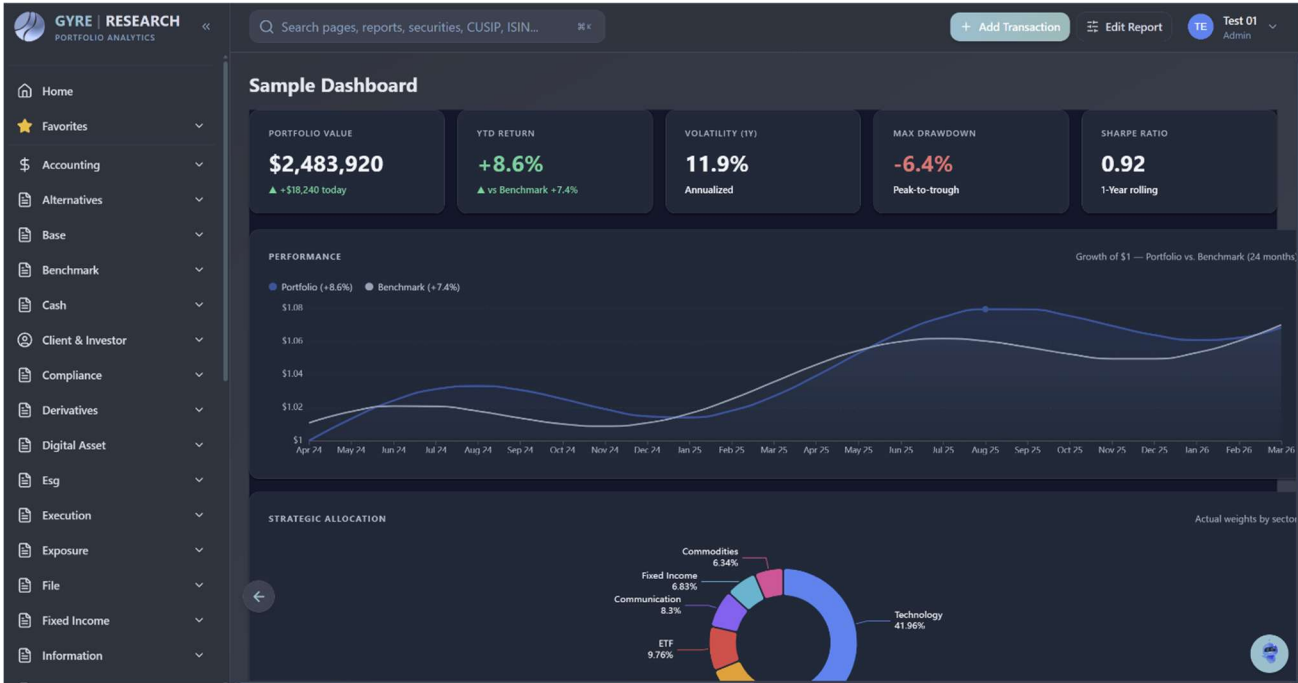
What the Gyre Hub does: The Gyre Hub gives the PM a single book with positions, gross/net and currency exposure, performance and contribution, Brinson attribution, factor and custom-beta exposures, and a mean-variance, risk-parity, or CVaR optimizer, all derived from the same valuation. Position, exposure, performance, factor_model, custom_betas, and portfolio_opt read the one book, so the exposure that drives risk is the exposure that explains return.

The outcome: The PM stops assembling the picture and starts using it: positioning, the risk it implies, and the return it produced are one consistent view, current as of the latest mark.

Under the hood: *position (position_calculate), exposure (exposure_calculate, allocation_attach, strategy_enrich), fx_exposure (decompose_position/hedge/pnl, hedge_effectiveness), return_calc (daily_chain_twr, modified_dietz, mwr_newton, growth_of_dollar), perf_attribution (brinson_bhb, brinson_fachler, multi_period_carino), factor_model (compute_value/momentum/quality/size/low_vol/carry,*



compute_exposures, compute_risk_attribution), custom_betas (ols_regression, rolling_betas, vasicek_shrinkage), portfolio_opt (optimize_mean_variance/risk_parity/cvar/kelly).



The Risk Officer

My risk numbers are only as good as last night's extract, and they miss the leverage and counterparty exposure that move a levered book.

The problem: Risk runs on a snapshot imported into a standalone engine. When VaR or tracking error moves, tracing it back through the import to the underlying holdings is slow, and a stale or unreconciled position quietly corrupts the entire risk picture.

What the Gyre Hub does: The Gyre Hub computes the covariance structure once and derives the full risk surface from it: portfolio and position-level volatility, beta, tracking error and Sharpe/Information ratios; parametric, historical, and Monte Carlo VaR with Expected Shortfall and

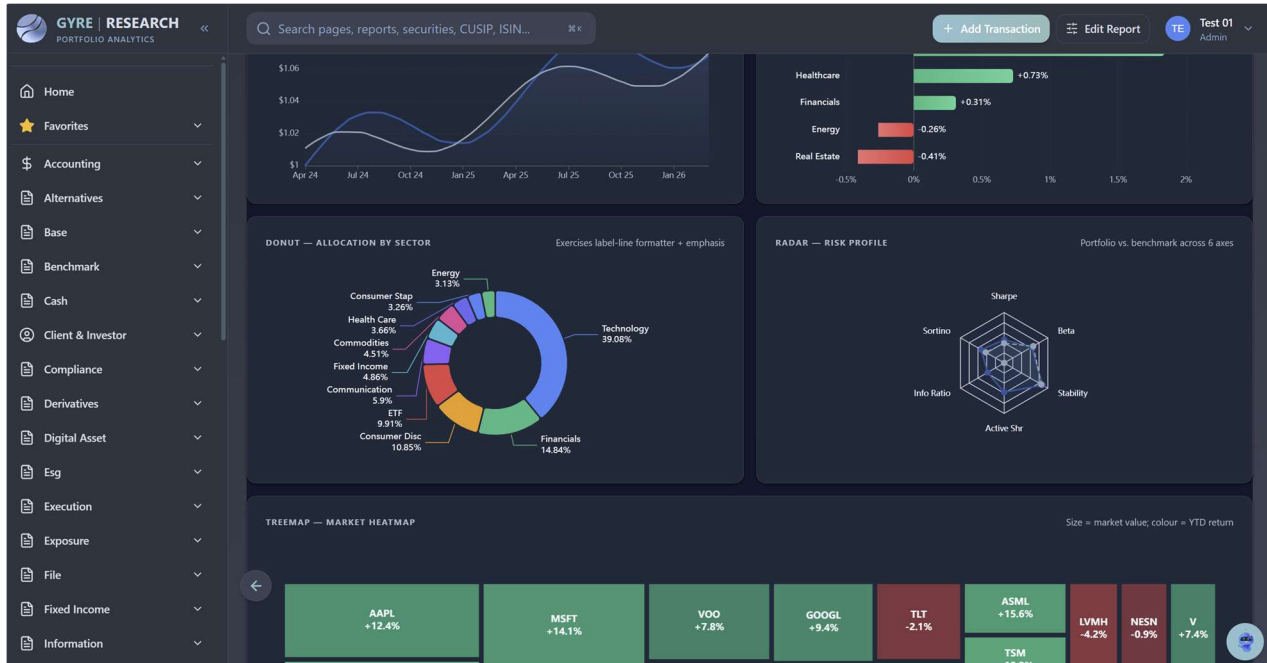
For informational purposes only. Not investment advice.
 © 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved



component VaR; stress and reverse-stress scenarios; counterparty XVA; and liquidity and enterprise-risk rollups. covariance, risk_portfolio, risk_position, var, var_backtest, stress_test, reverse_stress, counterparty_risk, and liquidity_score all run on the live book.

The outcome: Risk is current and explainable. Every figure decomposes to the positions that drove it, and a VaR backtest with Basel traffic-light scoring shows the model is behaving, on the same book the firm trades.

Under the hood: covariance (sample/ewma/ledoit_wolf_shrinkage, nearest_psd), risk_portfolio (portfolio_volatility, beta, tracking_error_ex-ante/ex-post, active_share, herfindahl; sharpe/sortino/treynor/calmar/omega; drawdown, ulcer_index), risk_position (decompose_position_risk, MCR/CCR), var (var_parametric/historical/monte_carlo, expected_shortfall, var_components, kupiec_pof/christoffersen), var_backtest (build_breach_calendar, basel_traffic_light), stress_test (parallel/nonparallel/pca curve shocks, fx_shock, monte_carlo_jump_diffusion), reverse_stress (min_norm_solution, mahalanobis, pareto_frontier), counterparty_risk (cva/dva/fva, pfe_95, wrong_way_risk_check, isda_netted_exposure), liquidity_score (check_22e4_compliance, stress_scenarios), enterprise_risk.



For informational purposes only. Not investment advice.
 © 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved



The Performance & Investor-Reporting Team

Net-of-fee returns, attribution, and the LP pack live in spreadsheets that only one analyst fully understands.

The problem: Returns are computed in one place, attribution by hand in another, and GIPS composites in a fragile workbook. Every reporting cycle is a manual rebuild, and a single broken link can delay an entire client pack.

What the Gyre Hub does: The Gyre Hub produces time-weighted and money-weighted returns for every period, dollar P&L decomposition, single- and multi-currency Brinson attribution with Cariño/Frongello linking, and GIPS-compliant composites with dispersion and an export pack. `return_calc`, `pnl_calc`, `perf_attribution`, `attribution_multicurrency`, and `gips` all derive from the same returns and benchmark mapping.

The outcome: Reporting becomes a review of a computed result rather than a monthly reconstruction, and the attribution on the client report ties exactly to the performance figures beside it.

Under the hood: `return_calc` (`twr + mwr`, `sharpe/sortino/treynor`, `tracking_error`, `information_ratio`, `jensens_alpha`), `pnl_calc` (`compute_realized/unrealized/mtm`, by `account/security/lot/period`), `perf_attribution` (`brinson_bhb/fachler`, `multi_period_carino/frongello`, `currency_attribution`, `fi_attribution`), `attribution_multicurrency` (`karnosky_singer_4way`), `gips` (`gips_enrich`, `compute_gips_stats`, `dispersion via calc/composite.py`, `gips_excel build_workbook`).

The Compliance Officer

I am monitoring limits against data that is already a day old.

The problem: Compliance checks run as an overnight batch against an exported snapshot, so genuine breaches surface late and phantom breaches generate noise. Surveillance, personal-trading pre-clearance, and best-execution evidence are scattered across separate tools and manual logs.

What the Gyre Hub does: The Gyre Hub runs a configurable compliance engine on the live book: 50 portfolio-level rules and 20 transaction-level rules covering single-name and issuer concentration, sector/country/currency bands, UCITS 5/10/40 and RIC tests, counterparty and liquidity limits, and days-to-liquidate. Alongside it: restricted-list and pre-trade decisions, trade surveillance for layering/spoofing/wash trades, personal-trading attestation and pre-clearance, best-execution Rule 606 / RTS 28 reporting, and the regulatory-filing calendar.

The outcome: Controls reflect reality. Limits are checked against the current book, real breaches surface immediately, and the evidence trail for an examiner is produced by the system rather than reassembled by hand.

Under the hood: *compliance_portfolio* (50 rules C01–C50: concentration, sector/country/currency bands, UCITS 5/10/40, RIC 75-5-10, days-to-liquidate), *compliance_txn* (20 rules T01–T20: trade \$/qty, %ADV, %NAV, per-trader/desk limits), *compliance_engine* (post_trade_check, transaction_by_transaction_check, batch_check), *restricted_list* (build_blackout/reg_m/rule_144_flags, build_pre_trade_decisions), *trade_surveillance* (detect_layering/spoofing/wash_trades/front_running/marking_the_close), *personal_trading* (evaluate_pre_clearance, build_attestation_grid, front_running_screen), *best_execution* (build_rule_606/rts_28_table), *reg_filing*, *soft_dollar*, *esg_screening* (evaluate_exclusions).

The Operations / Middle Office

My team spends every morning reconciling several prime brokers and clearing the same categories of breaks.

The problem: Reconciliation across multiple prime brokers is a spreadsheet exercise, settlement, financing, and corporate-action exceptions are tracked manually, and recurring breaks consume the same operational bandwidth as genuine ones, leaving less capacity for the issues that actually need judgment.

What the Gyre Hub does: The Gyre Hub treats reconciliation as a first-class function: a five-pass transaction matching engine (exact, tolerance, fuzzy, allocation, step-out), position and cash reconciliation, automated break classification and aging, and a break cockpit that routes only true

exceptions to an analyst. Settlement lifecycle, corporate actions, margin and collateral, document parsing, and an approval-workflow engine sit on the same book.

The outcome: Routine breaks are identified, categorized, and in many cases auto-resolved; analysts work the tail. Daily break volume and average break age fall sharply, and reconciliation stays current instead of accumulating toward month-end.

Under the hood: *recon_transaction (match_pass_1..5 exact/tolerance/fuzzy/allocation/step-out, classify_drift, compute_aging_buckets, compute_counterparty_scorecard, compute_csd_r_penalty, try_auto_resolve), recon_position, recon_cash (match_movements, fx_translation_check, age_breaks), break_cockpit, calc/recon_engine + break_classifier (classify_position/cash_diff), trade_settlement (derive_lifecycle, compute_fail_cost, csdr_penalty, reg_sho_deadline, cns_obligations, step_outs), corp_actions (handle_split/cash_div/stock_div/drip/merger/spinoff/rights/tender/coupon/bond_call/maturity), margin (reg_t/span/csa/wgmr/portfolio_margin), approval_workflow, doc_parsing (llm_parser, resolve_security/account).*

The Fund Accountant / Controller

Month-end close is a spreadsheet rebuild under time pressure.

The problem: The books are a chain of linked workbooks pulling journal entries from one system, accruals from another, and fee schedules from a third. The close is slow, performance-fee and high-water-mark logic is fragile, and the tax view is reconstructed after the fact.

What the Gyre Hub does: The Gyre Hub generates the books from the same transactions and valuations as everything else: double-entry general ledger, trial balance, full P&L statement, daily accruals, and NAV per account and share class, with management and performance fees on high-water-mark and hurdle logic, equalization, and period close. Tax runs continuously: lot-level accounting across relief methods, wash-sale tracking, loss-harvesting, and investor tax forms.

The outcome: The close becomes a review of a computed NAV rather than a reconstruction, fee accruals are auditable to the arithmetic, and the tax position is maintained daily instead of rebuilt at year-end.

Under the hood: *general_ledger* (JournalEntry/Line generation), *trial_balance* (control_account_recon, comparative_view, subledger totals), *nav* (compute_assets/liabilities, allocate_share_classes, compute_daily_attribution), *accruals* (interest/dividend/management/performance-fee accrual), *fee_calc* (compute_management_fee, compute_performance_fee with HWM + hurdle, compute_clawback, share_class breakdown), *pnl_statement*, *period_close*, *tax_lot* (lot_create/relieve across FIFO/LIFO/HIFO/specific-ID), *tax_accounting*, *tax_harvest* (detect_wash_sale_block, build_replacement_candidates, compute_state_overlay), *tax_forms*, *capital_activity* (money-weighted IRR, equalization, gates/lockups, waterfall, ILPA call notice, AIFMD).

The Trading Desk

Execution quality is measured in a tool that never sees the portfolio.

The problem: Orders, executions, and transaction-cost analysis live in execution systems disconnected from the book, so TCA and best-execution reporting run on a separate copy of the trade history and rarely tie back to the positions and benchmarks the firm actually holds.

What the Gyre Hub does: The Gyre Hub carries order and execution management, transaction cost analysis with outlier flagging and venue scorecards, best-execution Rule 606 / RTS 28 evidence, broker scorecards and commission/wallet tracking, and securities-lending economics, all on the same transaction and position data.

The outcome: Execution quality is measured against the firm's real benchmarks and positions, broker allocation is driven by scored performance, and the best-execution record is a by-product of the same data that books the trade.

Under the hood: *oms* (derive_status, compute_vwap, allocate_pro_rata/average_price/fifo, validate_pre_trade), *ems* (slice_twap/vwap/pov/is/iceberg, select_venue, check_trade_through, build_fix_new_order/execution_report, compute_slippage, compute_venue_scorecard), *tca* (decompose_order, flag_outliers, tca_account), *best_execution* (rule_606, rts_28), *broker_scorecard*, *securities_lending* (loan_pnl, repo_accrual, recall_severity, substitution_check).



The Quant / Systematic Researcher

My models run on a research copy of the data that drifts from the live book.

The problem: Signals, factor models, and backtests are developed in a research environment fed by its own data pipeline. What works in research behaves differently in production because the position, price, and corporate-action handling never matched the live book, and look-ahead creeps in.

What the Gyre Hub does: The Gyre Hub exposes factor models, machine-learning return prediction, portfolio optimization, time-series momentum, statistical-arbitrage and relative-value signals, macro-regime allocation, Kalman/state-space filtering, volatility targeting, snapshot backtesting with an explicit look-ahead guard, and anomaly detection, all reading the same valued book the firm runs on.

The outcome: Research and production use one data foundation, so a signal that tests well is built on the same prices and positions it will trade on, and the path from idea to allocation no longer crosses a data boundary.

Under the hood: *factor_model + custom_betas, ml_predict (build_feature_matrix, predict_universe, rank_and_decile, feature_importance), portfolio_opt (mean_variance/risk_parity/cvar/kelly), ts_momentum (trailing_return, cross_sectional_z, assign_ranks, portfolio_tilt), stat_arb, relative_value, macro_regime, kalman_filter (kalman_filter_pass, rts_smoother, detect_regime_changes, dynamic-beta/factor-loading/hedge-ratio), vol_target, backtest (look-ahead_guard, compute_sharpe/sortino/drawdown/attribution), anomaly_detection (ks_2sample, cusum_breach, isolation-forest if_score).*

The IR / LP Reporting Team

Every quarter-end I rebuild the same investor packs by hand, hoping the numbers tie to the fund's books.

The problem: Investor reporting is assembled manually, capital account statements, allocations, and waterfalls stitched together in spreadsheets from the administrator's files, then formatted into



packs one LP at a time. The work is slow, error-prone, and disconnected from the same book the rest of the firm runs on, so investor-facing numbers can drift from the fund's official records.

What the Gyre Hub does: Because investor reporting reads from the same valued book as everything else, the Gyre Hub computes capital activity and distribution waterfalls directly against fund-level positions and cash flows, generates LP statements, and serves them through an investor portal. It produces investor tax forms from the same underlying data and, for fund-of-funds structures, resolves look-through exposure to the underlying holdings rather than stopping at the fund line.

The outcome: Investor packs stop being a manual rebuild. Capital accounts, waterfalls, statements, and tax forms are generated from the fund's official book, reconciled by construction, and delivered to LPs through the portal, so what an investor sees always ties to what the fund reports.

Under the hood: *capital_activity (money-weighted IRR, equalization, gates/lockups, waterfall, ILPA call notice, AIFMD), lp_portal, tax_forms, fund_of_funds.*

Summary... Your Firm, at a Glance

One valued book; every role reads from it.

Market data and custodian feeds load once into a single valued book, pricing, positions, transactions, and tax lots, and every role below consumes the same source rather than its own copy:

ROLE	THE DAILY PROBLEM	HUB CAPABILITIES
Portfolio Manager	Five reports that disagree	Positions, exposure, performance, attribution, factor & beta exposures, optimization
Risk Officer	Risk on last night's extract	Covariance, VaR/ES, component VaR, stress & reverse stress, counterparty XVA, liquidity
Performance / Reporting	Attribution & GIPS in spreadsheets	TWR/MWR, Brinson & multi-currency attribution, GIPS composites & export

For informational purposes only. Not investment advice.
 © 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved



Compliance	Checks on day-old data	50 portfolio + 20 transaction rules, surveillance, personal trading, best-ex, reg filing
Operations / Middle Office	The same breaks every morning	5-pass reconciliation, break cockpit, settlement, corporate actions, margin, approvals
Fund Accounting	Month-end spreadsheet rebuild	GL, trial balance, NAV, accruals, fees (HWM/hurdle), period close, tax & forms
Trading Desk	TCA blind to the portfolio	OMS/EMS, TCA, best execution, broker scorecard, securities lending
Quant / Research	Models on a drifting data copy	Factor model, ML prediction, optimization, momentum, stat-arb, Kalman, backtest
IR / LP Reporting	Investor packs assembled by hand	Capital activity & waterfalls, LP portal & statements, tax forms, fund-of-funds look-through

What Runs Underneath

- 98 calculation modules across 12 disciplines, each an enrichment stage on one shared data object, run in prerequisite order by the process engine.
- A calc foundation of shared engines (fixed-income, option Greeks, covariance, compliance and reconciliation engines) and the canonical data model.
- A high-performance API service exposing 77 routes with full entity persistence, and a modern web hub (49 screens, 177 components) whose report layouts bind dashboard elements directly to data-object leaf paths (with detailed samples and documentation available).
- SaaS Application, fully tested with Chrome, Safari and Edge (Desktop and Mobile Versions)
- MS SQL Server + Rocky Linux 10 server infrastructure, hosted in a world-class data center appropriate for your domicile. Detailed Certifications and Audit documentation available by location.

Getting Started

The common thread across every role is that the firm does not have a data problem so much as a copies problem: each team maintains its own version of the book, and the work of holding those

For informational purposes only. Not investment advice.
 © 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved

versions together grows faster than the firm does. Consolidating onto one valued book removes the copies rather than reconciling them.

The right starting point is the same one we recommend for any consolidation: value one set of accounts on the Gyre Hub, reconcile its positions, risk, and NAV against your current systems, and see where they already disagree. That exercise takes about a week, proves the single-source-of-truth premise on your own data, and produces a prioritized roadmap for which roles to bring onto the platform first. Because all 98 modules read the same book, each one you turn on is a configuration step, not another integration project.

To learn more or schedule a demonstration, contact team@gyrererearch.com.

Appendix, The Full Capability Catalog

The narrative above is backed by the platform's full module set. All 98 calculation modules, grouped by discipline, run on the single shared data object:

- **INFORMATION (7)** calendar, news, filings, security_fundamental, ref_data_quality, alt_data, insider_flow
- **BASE (5)** pricing, position, transaction, tax_lot, report_config
- **BENCHMARK (4)** benchmark_load, benchmark_calc, benchmark_sec_rebal, benchmark_seg_rebal
- **CASH (4)** cash_balance, corp_actions, cash_forecast, treasury_cockpit
- **ACCOUNTING_TAX (10)** general_ledger, accruals, fee_calc, pnl_statement, trial_balance, nav, tax_accounting, tax_harvest, period_close, tax_forms
- **PERFORMANCE (5)** return_calc, pnl_calc, perf_attribution, gips, attribution_multicurrency
- **RISK (14)** exposure, fx_exposure, covariance, risk_portfolio, risk_position, custom_betas, var, stress_test, enterprise_risk, liquidity_score, counterparty_risk, var_backtest, reverse_stress, raroc

For informational purposes only. Not investment advice.
© 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved

- **COMPLIANCE_ESG (8)** restricted_list, compliance_portfolio, compliance_txn, esg_scoring, esg_screening, trade_surveillance, soft_dollar, personal_trading
- **EXECUTION (5)** oms, ems, tca, best_execution, broker_scorecard
- **OPERATIONS (13)** trade_settlement, capital_activity, margin, doc_parsing, recon_position, recon_transaction, recon_cash, break_cockpit, lp_portal, approval_workflow, sales_crm, reg_filing, securities_lending
- **QUANT_RESEARCH (12)** ts_momentum, vol_target, relative_value, macro_regime, factor_model, stat_arb, ml_predict, portfolio_opt, sentiment, kalman_filter, backtest, anomaly_detection
- **SINGLE_SECURITY (11)** yield_curve, fi_analytics, prepayment_model, swap_pricing, option_pricing, vol_surface, digital_asset, pe_analytics, illiquid_accounting, cat_bonds, fund_of_funds.

This document is published by Gyre Holdings LLC d/b/a Gyre Research for informational purposes only and does not constitute investment advice or a solicitation to buy or sell any security. Readers should consult a qualified financial professional before making any investment decision. All content is the intellectual property of Gyre Holdings LLC d/b/a Gyre Research and may not be reproduced or distributed without prior written consent.

© 2026 Gyre Holdings LLC d/b/a Gyre Research. All rights reserved.
