Independent Model Rating

The Triumph Visualizer is the proprietary investment system that powers our active equity strategy. It is a complete platform: research, scoring, portfolio construction, position management, and live performance tracking all run through a single application built and operated in house. This report is an internal audit of where the system stands today against the institutional bar Triumph holds itself to.

98/100

INSTITUTIONAL GRADE

as of May 14, 2026

The Visualizer combines a backtest validated signal across more than ten years of monthly Bloomberg fundamentals with a live, monthly rebalanced top twenty basket and a deterministic state machine that records the reason for every entry and exit. The methodology has been formally stress tested across the four robustness dimensions a serious allocator examines — beta sensitivity, universe survivorship, forward factor mix, and pillar weight optimisation — and the production parameters survive every test. A live Hedge Lab prices six structured option strategies for any holding off real time Bloomberg data. The codebase runs at the institutional baseline of zero type errors, zero lint warnings, a green test suite, and a weekly automated audit. Every number a portfolio manager sees on the screen is reproducible from the data that was actually visible at that moment.

WHAT BACKS THE GRADE

SIGNAL DEPTH

Stocks scored every day	**1207**
Quality factors across four pillars	**40**
Valuation factors	**14**
Forward versus trailing data mix	**66 / 34**
Primary data source	**Bloomberg SAPI**

BACKTEST HISTORY

Annual vintages tested	**12 (2014 to 2025)**
Stock-years observed	**~5,691**
Sharpe (production, survivorship-corrected)	**1.31**
Annual volatility	**15.7%**
Median forward return (Alpha Zone)	**~23%**
Robustness tests passed (May 2026)	**4 / 4**

LIVE PERFORMANCE

Inception	**2026-05-01**
Basket size, monthly rebalance	**10 names**
Basket return since inception	**–**
SPY return same window	**–**
Edge versus SPY	**–**
Hit rate (positions positive vs SPY)	**–**

FORWARD-LOOKING RISK PROFILE

Expected forward return	**9.5%**
Expected volatility	**33.4%**
Expected Sharpe	**0.29**
Beta versus SPY	**1.93**
Top sector	**Technology 59%**
Top five names share of book	**59%**

Methodology. The grade is an internal audit against the criteria above. The model itself, the code, the live basket, the backtest, and the audit script are all visible inside the Triumph Visualizer. Live numbers are pulled at the moment this report is generated.

A complete investment system

Most quantitative tools do one thing: a screen, a backtest, a risk report. The Visualizer does all of them, in one application, against one source of truth. A portfolio manager moves through six surfaces in a single workflow — research, scoring, portfolio review, position management, operational health, and external presentation — with every number on every screen reconciling end to end. The same composite that ranks a stock in the morning research is the same composite that drives the basket rebalance, the same composite shown to a client on the website, and the same composite the manager checks before sizing a new position.

ACTIVE BOOK

Portfolio

Live tracking of every holding with weighted return today, month to date, and year to date, alongside the same window for the SPY benchmark and the resulting excess. Side by side: sector exposure, concentration, covariance based volatility, beta, expected forward return, and a three layer fee analysis from gross fully invested down to net of management and performance fees. Click any name to see its contribution to risk and to return.

LIVE MODEL BASKET

Alpha Zone

The current top 10 names by composite, conviction weighted, with a live since inception versus SPY card, a rebalance log that captures every entry and exit with the locked in return for the leg, and a 3D scatter that lets the manager scrub through a hundred and thirty seven months of history to see how today's basket would have looked at any prior point in the cycle.

DAILY RESEARCH NOTE

Research Desk

A daily idea note published only when a name clears strict incremental Sharpe, expected return, and correlation gates. A deterministic state machine confirms the rating before the supporting language is written; the prose is constrained to numbers the desk has already computed, so it can never fabricate a fact. Quiet days are the right answer when nothing clears the bar.

POSITION MANAGEMENT

Hedge Lab

For any holding the manager opens the Hedge Lab and sees six structured option strategies — trim, covered call, collar, cash secured put, stock replacement, hold — priced live from Bloomberg option chains. Each candidate carries a transparent fit score, a payoff diagram showing where the structure protects you and where it caps you, and a tax aware net of cost view.

OPERATIONAL HEALTH

Pipeline

A rolling five day status grid for every external pull and key compute step, with one click recovery for any failed cell. Cross file integrity checks verify scores, the basket, and the price cache are mutually consistent on every refresh. A site wide banner appears immediately if any check fails.

EXTERNAL

Presentation and Website

A live sixteen slide investor deck and the public Triumph website both read directly from the same data files the model produces. When the basket changes in the morning, the deck and the site reflect it before lunch. There is no parallel set of marketing numbers; what a client sees is what the model just decided.

One application, one source of truth. Each surface is a different lens on the same underlying state. Reconcile any number on any page and the same input file produces it.

How the model thinks

The Visualizer scores every name on two independent axes. The Quality axis blends 40 fundamental factors that measure margins, returns on capital, growth trajectory, balance sheet strength, and shareholder payout. The Valuation axis blends 14 factors, eleven peer-relative multiples plus three universe-scaled absolute factors covering asymmetry, discounted cash flow, and today's discount to fair value. Roughly two thirds of the inputs on each axis are forward-looking estimates rather than trailing data, which is the design choice that lets the model anticipate rather than react.

UNIVERSE	QUALITY FACTORS	VALUATION FACTORS	FORWARD / TRAILING
1207	40	14	66 / 34
US, Canadian, global ADR	Across four pillars	11 peer-relative + 3 absolute	Anticipate, not react

THE COMPOSITE FORMULA

The two axes resolve into a single composite score per stock per day through a deliberately asymmetric formula. When valuation is favourable, quality and cheapness compound geometrically and a great company at a reasonable price is rewarded quadratically. When valuation is unfavourable, the same stock is penalised linearly. This shape encodes the philosophy of the strategy: the largest opportunities sit where excellence and value reinforce each other, and the largest mistakes sit where price has run ahead of fundamentals. The formula is smooth, monotonic, and ranges from roughly minus six to plus seven.

Per-position upside and downside math anchors on each company's own five year price-earnings history, with a stress floor below the empirical low and a capped multiple above the empirical high. This replaced an earlier approach that pulled from universe-wide percentiles and produced results that did not match how individual companies actually trade. The current method gives downside and upside numbers a portfolio manager can act on rather than statistical artefacts that need translation.

FROM SIGNAL TO BASKET

From the scored universe, the Alpha Zone selects the top 10 names by composite, gated to the BUY tier so that only high-conviction companies qualify. Position weights tilt toward conviction, with the highest scoring names receiving the largest weights up to a hard cap of fifteen percent per name. The basket rebalances on a monthly cadence chosen empirically from a walk-forward backtest that tested daily, weekly, monthly, and annual cadences across twelve annual vintages from 2014 through 2025.

> *Across twelve annual vintages from 2014 through 2025, the Alpha Zone outperformed the rest of the scored universe by an average of 7.7 percentage points per year, with positive performance in nine of those twelve years. Production rules deliver a survivorship-corrected Sharpe of 1.31 at 15.7% annual volatility, validated against the full Russell 3000.*

THE DISPLAY

Every stock in the universe lives as a sphere in a three dimensional scene where Quality is one axis, Valuation is another, and the Z axis shows a momentum signal. The Alpha Zone appears as a translucent ellipsoid

around the names that currently qualify. A monthly history slider lets the manager rewind the entire universe to any past month and watch how the spheres move. The same scene exists on the public website, the investor deck, the portfolio page, and the alpha zone page, scoped to the relevant subset on each surface. It is not a chart of the model. It is the model.

An asymmetric reward, by design. Cheap quality compounds. Expensive quality is taxed.

Engineered to institutional standards

The signal is only as good as the data behind it and the engineering that keeps that data trustworthy. Triumph runs the entire pipeline on hardware we own and maintain, against data sources we pay for directly, with checks that surface any drift the moment it happens.

WHERE THE DATA COMES FROM

Every fundamental and forward consensus field used in the score is sourced through a Bloomberg API service running on a Triumph operated host, with user mode authentication that allows the same field coverage as the terminal itself. A second commercial data feed serves as a fallback for non-Bloomberg fields and a cross check on price history. Every fetch is logged, time stamped, and counted; the daily morning refresh writes a versioned snapshot of the entire scored universe to a versioned archive, which over time becomes the substrate for forward factor research that cannot be tainted by hindsight.

THE MORNING CHAIN

A scheduled job runs every weekday morning before six. It refreshes fundamentals, refreshes forward consensus estimates, recomputes the scored universe, recomputes the Alpha Zone, repaints the daily snapshot archive, regenerates the morning research note, refreshes option chains for every holding, repaints the public website, and runs cross file integrity checks. The chain is sequenced so that downstream steps see only the upstream data they actually depend on. By the time the manager sits down at the desk, the workbench is current; nothing waits on a human.

HOW ERRORS ARE CAUGHT

Three independent layers protect the data. A daily integrity check verifies that scores, the live basket, the price cache, and the portfolio book are mutually consistent on every refresh; if any check fails, a red banner appears site wide before anyone reads a number. A weekly automated code health audit checks build state, security advisories, and architectural drift and writes the result to a file the team reads on Sunday morning. A pipeline page shows a rolling five day status grid for every external pull and gives one click recovery for any failed cell. Nothing about the model, the basket, or this rating depends on a black box.

WHY WE CALL IT INSTITUTIONAL GRADE

Three things, in order. First, the signal is real. Twelve annual vintages, almost six thousand stock-years of forward returns, a walk-forward Sharpe of 1.36, and an average annual edge of 7.7 percentage points against the rest of the universe. The number was produced by walk-forward methodology on data the model could only have seen at the time, not by an in-sample fit. Second, the engineering is real. The system runs at the same baseline an institutional client would expect from a public software vendor: zero type errors, zero lint warnings, full test coverage of the scoring code, automated daily integrity checks, automated weekly code health audits. Third, the transparency is real. The model itself, the code, the live basket, the backtest, and the audit script are all visible inside the application that produced this report. There is no version of the system Triumph keeps that a client cannot see.

The result is a workbench that does for the Triumph portfolio manager what the desks of the largest institutional shops do for theirs, run by a focused team that can actually act on what it sees. That is the model behind the rating.

Internal Triumph audit, not a third-party rating. Live numbers are computed at generation time and may differ slightly from the values shown on the live application after the next refresh. Past performance is not a guarantee of future results.