

FINANCE FEATURE · RISK & CAPITAL

Stress-Testing a \$600 Billion Bank: Inside the GCC Engine Room

How an applied research project built an agentic-AI-enabled, Federal Reserve-style credit risk stress testing prototype inside the operating model of a banking Global Capability Centre.

Based on the IMT Ghaziabad (CFM) applied research programme on BFSI Global Capability Centres, April–June 2026

WHAT IS A GLOBAL CAPABILITY CENTRE?

A **Global Capability Centre (GCC)** is an offshore hub, frequently located in India, through which a multinational bank operates specialised functions such as data management, quantitative risk analytics, model validation, compliance and regulatory reporting. GCCs originated as shared service centres and have since evolved into Centres of Excellence (CoEs) for enterprise risk management, with teams in India supporting the global risk functions of banks such as HSBC, J.P. Morgan, Barclays and NatWest through centralised delivery and standardised methodologies.

Their strategic value lies in centralised risk analytics, standardised methods, automation and scalability, independent governance, faster regulatory reporting and better capital-planning support.

THE PROJECT AND ITS OBJECTIVE

Banks operate amid macroeconomic variables such as inflation, interest-rate movements, geopolitical shocks and market disruptions – forces that directly erode borrowers' repayment capacity. Even a moderate downturn raises default rates and expected credit losses; a severe downturn additionally depresses collateral values and places material pressure on regulatory capital. Regulators therefore mandate stress testing: a structured assessment of whether a bank holds sufficient capital to absorb losses under hypothetical yet plausible adverse scenarios.

This project analyses a **PD–LGD–EAD credit risk stress testing framework** that estimates the impact of adverse macroeconomic scenarios on portfolio-level expected credit losses and capital adequacy and, crucially, embeds the model within the real-world operating structure of a banking GCC.

“How can a multinational bank translate adverse macroeconomic scenarios into portfolio-level credit losses and capital impact through a transparent, governance-driven stress testing framework that reflects the operating model of a Global Capability Centre?”

WHY BANKS STRESS TEST

Need	Outcome
Capital planning	Assess CET1 resilience
Regulation	Meet supervisory expectations
Portfolio risk	Identify vulnerable sectors
Management	Support strategic decisions
Early warning	Detect hidden fragilities
Governance	Improve transparency

INSIDE THE GCC: WHO DOES WHAT

Enterprise stress testing in a GCC is a sequential, multi-team discipline in which each function hands validated work to the next:

Function	Role	Contribution
Enterprise Data Operations	Data sourcing	Reliable portfolio data
Data Engineering	ETL, transformation	Model-ready datasets
Credit Risk Analytics	PD/LGD/EAD estimation	Loss estimation
Portfolio Analytics	Aggregation	Sector insights

Function	Role	Contribution
Treasury & Capital Mgmt.	CET1 assessment	Capital planning
Model Validation	Independent review	Model robustness
Business Oversight & Compliance	Governance	Control effectiveness
Regulatory Reporting	Submission packs	Regulator communication

THE WORKFLOW, END TO END

Each stage passes validated work to the next, and every result is independently reviewed before submission to the regulator:




Fig. 1 - The GCC stress testing workflow, from regulatory scenario to reporting

PART II - THE MODEL

The Stress Testing Model: From Macro Shock to Credit Risk

A \$600BN PROXY PORTFOLIO

The Excel model converts a JPMC-inspired loan portfolio into a \$600 billion proxy portfolio across six loan categories, attaches Basel-style LGD and risk-weight assumptions, and links Federal Reserve 2026 stress scenario variables into credit risk through normalised macro drivers. The core arithmetic is deliberately simple and auditable:

$$\text{Expected Loss} = \text{EAD} \times \text{PD} \times \text{LGD}$$

PD (probability of default) and LGD (loss given default) are stressed separately, because defaults and recoveries respond to different economic channels: unemployment drives credit-card defaults, while falling house prices erode mortgage recoveries.




Fig. 2 - Portfolio composition by loan category (share of \$600Bn proxy EAD)

Interpreting the chart: the two largest segments are 'Other Loans' (26.8%), a residual bucket of mixed lending, and first-lien mortgages (25.5%), which are collateralised by residential property. Unsecured credit cards account for only 16.5%. Composition matters because stress losses depend on risk intensity rather than exposure size alone: a large, well-collateralised segment can lose less in a crisis than a smaller unsecured one, a pattern the loss results later confirm. The capital view differs again: mortgages attract only a 30% Basel risk weight while the corporate, CRE and residual buckets attract 100%, so within the \$459.95Bn of risk-weighted assets the capital burden concentrates in the corporate and residual segments rather than in mortgages.

THE SEVERE SCENARIO

The Severely Adverse case represents a synchronised macro-financial deterioration in which repayment capacity, refinancing access, collateral values and market liquidity weaken simultaneously. All eleven Federal Reserve 2026 variables enter the model:

Macro variable	Baseline	Severe	Stress
Real GDP growth	1.0%	-5.4%	-6.4 pp
Unemployment rate	4.5%	10.0%	+5.5 pp
CPI inflation	2.8%	1.0%	-1.8 pp
3-month Treasury	3.7%	0.1%	-3.6 pp
10-year Treasury	4.1%	2.3%	-1.8 pp
BBB corporate yield	5.1%	8.2%	spread +4.7 pp
Mortgage rate	6.2%	6.0%	spread +1.3 pp
House prices (HPI)	323	227	≈ -30%
CRE prices	306	187	≈ -39%
Equity market	67,502	28,490	≈ -58%
VIX volatility	26.4	72.0	+45.6 pts

The scenario is internally coherent - the first property a GCC analyst verifies before relying on any credit estimate. Double-digit unemployment is consistent with consumer credit

stress; a 58% equity decline and a VIX spike to 72 are consistent with a collapse in risk appetite; and declines of 30% and 39% in house and CRE prices are consistent with weakened collateral. Because labour-market, credit-spread and asset-price stresses reinforce one another, defaults rise at precisely the moment recoveries weaken - the environment in which credit losses become nonlinear.

HOW MACRO SHOCKS BECOME CREDIT RISK

Two parameters underpin the model. **PD (probability of default)** measures the likelihood that a borrower ceases repayment; it responds to pressures on income and cash flow such as unemployment, recession and widening credit spreads. **LGD (loss given default)** measures the proportion of an exposure actually lost once a borrower defaults; this depends on what can be recovered from the borrower and from collateral, so its drivers are collateral-value and recovery channels such as house prices, CRE prices and distressed-market conditions. Each macro shock is scaled relative to the severe scenario, and category-wise driver weights convert the shocks into separate PD and LGD stress scores. The Moderate Adverse case is an intermediate scenario between baseline and severe, and reappears in the loss and capital tables later in the report.

The category-wise driver weights below convert the normalised macro shocks into separate PD and LGD stress scores for each loan category:

Category	PD drivers (income & cash-flow)	LGD drivers (collateral & recovery)
Mortgages	Unemployment 35%, mortgage spread 25%; GDP, HPI, volatility 10% each	HPI 70%; GDP, spread, volatility 10% each
C&I	Corporate spread 30%, GDP 25%, volatility 15%	Spread 30%, equity 30%, volatility 20%
CRE	CRE prices 25%; corporate & mortgage spreads 20% each	CRE prices 70%
Credit cards	Unemployment 45%, GDP 15%, CPI 15%	Volatility 30%; GDP, unemployment, spread 20% each
Other consumer	Unemployment 40%; GDP, CPI, volatility 15% each	Unemployment 25%, volatility 25%, GDP 20%
Other loans	GDP 25%, spread 20%; unemployment, equity, volatility 15% each	Spread 25%, equity 25%, volatility 20%

Interpreting the table: each weight measures the strength with which a macro shock affects that category, and the PD and LGD columns deliberately differ. Consider mortgages: unemployment carries the largest PD weight (35%) because loss of employment is the principal reason households cease repayment, yet 70% of the LGD weight rests on house prices, because the dwelling is the collateral and a forced sale into a falling market reduces recovery. For C&I loans the default channel operates through business cash flow (GDP, 25%) and refinancing cost (corporate spreads, 30%), while recovery depends on the value of a distressed business and its assets; equity values (30%) and credit spreads (30%) therefore drive LGD. Credit cards concentrate PD weight on unemployment (45%), and their LGD weights rest on broad market-stress indicators because no collateral is available for recovery.

CALIBRATING TO REAL-WORLD BENCHMARKS

Calibration denotes tuning the model so that its severe outputs align with trusted external benchmarks rather than floating freely. In the table below, **PD floor** is each category's default rate in a normal year (FDIC 2013-19 averages); **Base LGD** is the normal-times loss severity from Basel-style assumptions; **Target LLR** is the severe loan loss rate the scenario must reproduce, taken from CCAR benchmarks; and the **Multiplier** indicates how many times worse than baseline the severe loss rate is, ranging from 6.4x for collateralised CRE to 34.2x for unsecured credit cards:

Category	PD floor	Base LGD	Target LLR	Multiplier	Status
Mortgages	2.33%	5%	2.00%	17.2x	OK
C&I	0.79%	45%	11.50%	32.3x	Check
CRE	2.33%	20%	3.00%	6.4x	OK
Credit cards	0.96%	50%	16.40%	34.2x	Check
Other consumer	0.96%	30%	3.10%	10.8x	OK
Other loans	1.52%	45%	4.50%	6.6x	OK

LGD cannot exceed 100%, since a bank cannot lose more than it has lent. For C&I loans and credit cards the severe scenario pushes modelled LGD to this ceiling; the model therefore cannot fully attain its benchmark loss rate and flags the gap ('Check') for review rather than concealing it. Outputs breaching **PD > 20%, LGD > 70% or loss rate > 15%** raise automatic review flags.

FROM MACRO SHOCK TO LOSSES

Category	Base EL	Mod. EL	Sev. EL	Incr.	Share	LLR
Mortgages	\$0.18	\$1.34	\$3.06	\$2.88	8.6%	2.00%
C&I	\$0.28	\$2.66	\$5.93	\$5.65	16.7%	7.57%
CRE	\$0.35	\$1.21	\$2.28	\$1.93	6.4%	3.00%
Credit cards	\$0.47	\$7.15	\$16.01	\$15.53	45.1%	16.19%
Other cons.	\$0.10	\$0.51	\$1.03	\$0.93	2.9%	3.10%
Other loans	\$1.10	\$3.86	\$7.23	\$6.13	20.4%	4.50%
Total (\$Bn)	\$2.48	\$16.74	\$35.53	\$33.05	100%	-

Expected credit loss by category and scenario (\$Bn)

The three scenario columns are the baseline, Moderate Adverse and Severely Adverse cases defined on page 2. Aggregate expected loss rises **14.3 times** from baseline to severe; the incremental severe shock of \$33.05Bn represents the loss that capital must absorb beyond routine provisioning. The clearest pattern in the table concerns **the presence or absence of collateral**. Credit cards hold only 16.5% of exposure yet generate 45% of severe losses: unsecured borrowers default rapidly when unemployment rises, and with no collateral to recover, loss severities remain structurally high. Mortgages and CRE, secured on property to which the bank holds a claim, show far lower severe loss rates (2.00% and 3.00%) despite heavy collateral-price shocks. 'Other Loans', the largest exposure bucket, is the second-largest loss contributor - a governance concern, since residual buckets can conceal heterogeneous risks.




Fig. 3 - Exposure share vs. share of severe losses

TWO MODELS, ONE CROSS-CHECK

The project developed two models with distinct roles. **Model A establishes the model and its assumptions:** starting from published loss outcomes - Fed/JPM CCAR loss rates, FDIC 2013-19 baseline default rates, Basel III LGDs - it rebuilds the portfolio backwards (balance = loan loss ÷ loss rate, scaled to \$600Bn) and back-solves the stressed PD each benchmark implies, reaching **\$38.40Bn** of stressed loss and a post-stress CET1 of 10.56%. **Model B simulates that established model across variables and sensitivities:** eleven Fed 2026 macro shocks flow through driver weights into PD and LGD, producing losses for three scenarios plus six sensitivity overlays.

Why do the two models produce different figures - \$38.4Bn against \$35.5Bn? The difference arises by design. Model A reproduces its benchmark loss rates exactly, on a portfolio mix rebuilt from those same benchmarks, whereas Model B recalculates losses from macro drivers on the JPMC-mapped portfolio, where binding LGD ceilings hold modelled losses slightly below benchmark for two categories. A divergence of roughly 8% between an independent anchor and a simulation constitutes corroboration rather than contradiction: it confirms that the macro engine is calibrated within the correct range while behaving as a genuine model rather than a replica - the cross-check on which the capital verdict rests.

WHAT REVIEWERS SHOULD CHALLENGE

The workbook embeds its own challenge agenda. Before any result is communicated, a GCC review team should confirm that the macro variables form an internally coherent narrative (unemployment, spreads, asset prices and volatility deteriorating together); disaggregate the 'Other Loans' residual bucket or document why aggregation is acceptable; prepare a detailed explanation of unsecured credit-card loss behaviour; investigate the C&I calibration variance caused by the binding LGD cap rather than forcing the benchmark; and avoid simplistic 'pass' language where any capital view is marginal. These checkpoints exist because the numbers invite behavioural traps: anchoring on the comfortable 15.70% starting ratio, over-weighting the headline pass, and over-trusting generated explanations.

SIDE-BY-SIDE: HOW INPUTS, CALCULATIONS AND OUTPUTS INTERACT

Dimension	Model A - Back-Calculation	Model B - Macro-Driven
Loss rates are...	Inputs - benchmarks reproduced by design	Outputs - generated from macro shocks
Macro variables	None enter directly (implicit in benchmarks)	11 Fed 2026 variables drive PD & LGD
Scenarios	Two: baseline + one severe	Three + six sensitivity overlays
PD / LGD	PD back-solved (loss rate ÷ LGD); LGD static	Both stressed via driver indices; capped at 100%
Severe loss / verdict	\$38.40Bn; CET1 10.56% - Pass	\$35.53Bn; CET1 10.93% - narrow Pass
Governance	In-sheet notes and explanations	Validation flags, audit trail, AI evidence artifact

PART III - GOVERNANCE & AI

Capital Verdict, Agentic AI and What the Model Cannot Do

THE CAPITAL VERDICT: A BORDERLINE PASS

The waterfall starts from a **15.70% starting CET1 ratio**, deducts the baseline expected loss and the incremental stress loss net of the **21% tax shield**, and tests the result against a **9.50% floor** (4.5% Basel minimum + 2.5% conservation buffer + 2.5% G-SIB surcharge):

Scenario	Scenario EL	After-tax hit	Post-stress CET1	Verdict
Baseline	\$2.48Bn	-	15.29%	Pass
Moderate	\$16.74Bn	\$11.27Bn	13.41%	Pass
Severe	\$35.53Bn	\$26.11Bn	10.93%	Pass

Capital outcomes vs. the 9.50% floor




Fig. 4 - Post-stress CET1 ratio vs. the 9.50% regulatory floor

The bank passes in every scenario; however, the severe case clears the floor by only 1.43 percentage points. So thin a buffer offers little comfort: modest changes to key assumptions consume most of it, and the result should therefore be read as borderline and assumption-sensitive:

Sensitivity case	Add'l EL	Post-stress CET1	Buffer
Unemployment +100 bps	\$3.55Bn	10.47%	0.97 pp
HPI decline 10% worse	\$0.31Bn	10.89%	1.39 pp
CRE decline 10% worse	\$0.23Bn	10.90%	1.40 pp
BBB spread +100 bps	\$1.32Bn	10.76%	1.26 pp
LGD +10 pp	\$3.55Bn	10.47%	0.97 pp
Starting CET1 -100 bps	-	9.93%	0.43 pp

Modest changes cut the margin below 1 pp; weaker starting capital leaves 0.43 pp.

THE AGENTIC AI LAYER

The workbook was designed to be AI-ready. A dedicated **Evidence Artifact** sheet structures every output - scenario, loan category, macro drivers, PD, LGD, loss rate, expected loss, warning flags and human-review notes - so that a Streamlit-based agentic AI assistant can consume it directly. Within the companion prototype, this layer operates as a **governed reasoning companion above the deterministic stress engine**, converting a spreadsheet-intensive exercise into a guided analytical workflow whose modular Python components handle validation, shock extraction, driver mapping, loss calculation, the capital waterfall, sensitivities, exports and audit logs:

User Input → Agentic AI Layer → Scenario Engine → Stress Model → Dashboard Output

From first principles, effective stress testing requires **three linked judgments**: the nature of the shock being applied, the manner in which it transmits into borrower risk, and whether the resulting capital position is decision-useful. The table below maps each judgment to the support the AI layer provides. Unlike traditional static stress testing, in which analysts move between spreadsheets, policy notes, charts and commentary, scenario design, model execution, visualisation and governance are here reviewed in one place; although prototype-scale today, the module boundaries support scaling into enterprise workflows.

Judgment	How the AI layer supports it
What shock is being applied	Narrates scenario severity and coherence directly from the normalised macro shocks in the evidence artifact
How it transmits into borrower risk	Explains driver weights and the resulting PD and LGD movements; surfaces warning flags (PD > 20%, LGD > 70%, loss rate > 15%)
Whether the capital position is decision-useful	Summarises the capital waterfall and remaining buffer, frames the verdict candidly and drafts human-review questions

The division of labour is strict: the AI may **explain** deterministic outputs, surface warnings, highlight drivers, frame human-review questions and generate stakeholder-friendly narratives, but it **must not calculate, approve, validate or override the model**.

THE TOOL IN PRACTICE




Fig. 5 - The Streamlit prototype in operation: evidence artifact preview, explicit consent and the governed Gemini explanation (top); governance boundaries and the human approval gate with named approver (bottom)

The prototype renders these boundaries tangible. Model controls - the starting CET1 ratio, corporate tax rate and regulatory floor - are exposed as adjustable inputs; the evidence artifact is previewed before anything is transmitted to the AI; an explicit consent step records that the AI cannot alter model numbers; and every generated narrative enters a human approval gate, with a named owner, before disclosure. Notably, the AI-generated executive summary frames the severe result exactly as the governance design intends - as a narrow pass demanding attention to credit-card losses rather than a comfortable clearance.

WHAT THE LAYER DELIVERS

Benefit	How it is delivered
Faster interpretation	Narratives generated from the evidence artifact replace manual commentary drafting
Visible assumptions	Inputs, warning flags and review notes travel with every explanation
Sanitised communication	Unsupported regulatory claims are removed before stakeholders see them
Consistent control	Consent step, approval gate and audit logs govern every AI output
Analyst leverage	Effort shifts from mechanical assembly to judgment, challenge and scenario quality

The strategic consequence completes the arc with which this report began: a GCC that owns the analytics product, its governance framework and its AI tooling is no longer a shared service centre executing delegated tasks but a **Centre of Excellence** operating, and continuously improving, an enterprise risk capability. The destination is not autonomous decisioning but **governed intelligence around deterministic models**.

THE HUMAN LOOP

A nine-point review checklist assigns each control - portfolio mapping, macro scenario capture, driver weights and calibration, the capital waterfall, sensitivities, governance documentation, AI readiness and stakeholder traceability - to a named owner before results are finalised. Together with the governance audit trail, this arrangement counters automation bias, anchoring on the headline pass, and the temptation to favour assumptions that preserve a pass.

LIMITATIONS & CONSTRAINTS

This is an academic workflow prototype, not a validated regulatory model. The most important assumption is that the economic

variables themselves are proxies: macro paths are taken from the published Federal Reserve 2026 scenario, and the portfolio, LGD and risk-weight inputs are public benchmarks rather than bank-internal data. In addition, the \$600Bn portfolio is a scaled proxy; the RWA capital view is directional only, since losses are computed on EAD; the 21% tax shield requires jurisdiction-specific validation; the Moderate Adverse scenario is interpolated at 55% of the severe shocks rather than independently specified; the 'Other Loans' residual bucket aggregates mixed exposures; and where PD/LGD caps bind, calibration can no longer attain benchmark loss rates.

CONCLUSION

This project set out to answer a single question: how can a multinational bank translate adverse macroeconomic scenarios into portfolio-level credit losses and capital impact through a transparent, governance-driven framework that reflects the operating model of a Global Capability Centre? The answer developed here is an integrated one. A PD-LGD-EAD stress testing framework, applied to a \$600Bn six-category proxy portfolio under three Federal Reserve 2026 scenarios, was embedded within the full GCC operating chain - data engineering, portfolio mapping, modelling, capital assessment, independent validation, business oversight and regulatory reporting - and instrumented for governed AI explanation.

Analytically, the framework demonstrates how quickly credit costs escalate under coordinated macro-financial stress: expected losses rise from \$2.48Bn at baseline to \$16.74Bn under the Moderate Adverse scenario and \$35.53Bn under the Severely Adverse scenario - a 14.3-fold amplification. The distribution of those losses is the central economic finding. Unsecured credit cards, holding only 16.5% of exposure, generate \$16.01Bn (45.1%) of severe losses, while property-secured mortgages and CRE remain contained at 2.00% and 3.00% loss rates; the presence or absence of collateral, not exposure size, governs loss concentration. On capital, the post-stress CET1 ratio of 10.93% clears the 9.50% regulatory floor in every scenario, but by a margin of only 1.43 percentage points - and the sensitivity analysis shows that a 100-basis-point rise in unemployment, a 10-point LGD overlay, or a weaker starting capital position each compress that buffer below one percentage point, to as little as 0.43. The verdict is a pass, but a borderline and assumption-sensitive one.

Operationally, the project confirms that stress testing is an enterprise capability rather than a standalone quantitative exercise. Every number in this report is traceable through a controlled chain: benchmark-anchored assumptions established by the back-calculation model (whose independent \$38.40Bn severe-loss

estimate corroborates the macro engine within roughly 8%), a nine-point human review checklist with named owners, warning thresholds that flag extreme outputs automatically, and a governance audit trail designed to counter anchoring, confirmation and automation bias.

Methodologically, the prototype demonstrates that the credible next step for the discipline is not autonomous decisioning but governed intelligence around deterministic models. The agentic AI layer compresses interpretation time, keeps assumptions visible and sanitises unsupported claims - yet every narrative it produces is generated from a structured evidence artifact, passes through an explicit consent step, and terminates at a human approval gate. The division of labour holds throughout: the deterministic engine produces the numbers; the AI explains them; a named human approves them.

For Global Capability Centres, the implication completes the arc with which this report began. A GCC that owns the analytics product, its governance framework and its AI tooling is no longer a shared service centre executing delegated tasks but a Centre of Excellence operating and continuously improving an enterprise risk capability. Before any production use, the recognised boundaries must be closed: bank-internal data in place of proxy benchmarks, a formally specified adverse scenario in place of interpolation, validated tax and RWA treatment, and a disaggregated residual loan bucket. Within those boundaries, the model 'passes'; its more important achievement is demonstrating how banking GCCs convert regulatory stress testing into strategic decision-making - humans firmly in the loop, and AI as explainer, never approver.

Key result	Value
Baseline / severe expected loss	\$2.48Bn / \$35.53Bn (14.3x)
Largest severe-loss contributor	Credit cards - 45.1% of losses from 16.5% of exposure
Post-stress CET1 vs floor	10.93% vs 9.50% (buffer 1.43 pp)
Thinnest sensitivity buffer	0.43 pp (starting CET1 -100 bps)
Independent cross-check (Model A)	\$38.40Bn severe loss; within ~8% of the macro engine
AI governance boundary	Explains and flags; cannot calculate, approve or override

The report in six numbers