


HORIZON CREDIT ADVISORY · THE WHOLE PICTURE · NO. 1

Counterparty credit risk, in one picture

The perimeter fails at the joins, not within any one panel

Version	v1.15
Date	7 September 2026
Authors	Seamus Maher and Michał Falkowski
Status	Signed off for publication

*Horizon Credit Advisory, counterparty credit risk in one picture*

This note accompanies the picture above: Horizon's one-page map of counterparty credit risk in an energy and commodity trading firm. The picture shows six parts of the discipline working as one system: before a trade, the pre-deal check; at the centre, the engine that measures exposure; to the right, the collateral that reduces it, and the cash that collateral demands; to the left, the stress test that asks what breaks; underneath, the trading book that drives every number. We call the whole of it the perimeter, because that is what it is: everything a firm's credit risk lives inside.

Why draw it as one picture? Because in most trading firms counterparty credit risk is owned in fragments, split across functions in a pattern that varies from firm to firm: a credit analyst typically owns the limits, a quant team the exposure engine, treasury or credit the collateral and the margin calls it triggers, risk the stress tests, and IT the systems underneath all of it. The exact ownership map differs by firm; the fragmentation does not. Each fragment can be individually sound while the whole is quietly broken, because the failures that matter happen at the joins, in the handoffs between teams rather than within any one of them.

This note walks through the six panels in turn. For each one it describes what the panel does, how it breaks in real firms, what the public record says that breakage costs, and one question worth asking about your own firm.

A note on what this is not. The Committee of Chief Risk Officers published the industry's original best-practices reference for energy credit risk back in 2002, and is now assembling a new, more comprehensive edition chapter by chapter, with the first completed chapter (on credit-risk technology) already published. That is the atlas, complete or not. This is the map: short enough to read in one sitting, specific enough to be uncomfortable.

One honest boundary. This picture covers the margined and marginable book, the derivatives and structured trades that dominate exposure at most European gas and power desks. The open-account physical book, with its receivables, letters of credit, country risk and title questions, deserves its own picture, and will get one in a later note.

A note on the figures. The figures in the accompanying picture describe an illustrative portfolio. They are worked examples, not client data. The figures in this note are from the public record, and each carries its source.

1. Before the trade: the limit that does not bind

The textbook failure in limit design is not the absence of a limit. It is a limit that exists, is breached, and binds nobody. The collapse of Archegos Capital Management in March 2021 remains the defining case. Archegos, the family office of Bill Hwang, had amassed vast, concentrated equity positions through total return swaps with a small circle of prime brokers, among them Credit Suisse, which carried the largest exposure. When Archegos failed to meet its margin calls and its positions were liquidated, Credit Suisse absorbed a loss of roughly \$5.5 billion, the heaviest of any counterparty. What elevates the episode from an expensive misjudgement to a study in control failure is the Special Committee report the bank's own board commissioned afterwards, which found something worse than a blind spot: a limit framework that saw the danger clearly and still did nothing about it. In the months preceding the default, the bank's scenario limits on the exposure were breached week after week; each breach was tolerated, escalated nowhere, and the position was allowed to grow, while the response the limit existed to force never came. That pattern of breach without consequence explains why the eventual cost ran to billions rather than millions, and why the Bank of England's Prudential Regulation Authority, on investigation, fined Credit Suisse's UK entities a record £87 million. The fine was not for weak models. It was for failing to act on breaches the bank's own systems had dutifully, and repeatedly, reported.

The energy and commodities version of this failure is quieter, but it takes the same shape. A pre-deal check that takes longer than the deal itself teaches the desk to execute first and ask afterwards, because waiting carries a cost nobody wishes to pay twice. The incremental exposure of a structured transaction often cannot be computed at the point of decision, and so it books at zero, not because anyone judged it worthless, but because nothing compelled a better answer. A limit framework that speaks only after execution is not controlling anything. It is documenting what has already happened.

What good looks like. The incremental exposure of the trade, after netting and collateral, priced into the ticket before execution, with breach escalation that names a person and a deadline. In practice the pre-deal answer comes from approximations, grids and sensitivities rather than a full re-simulation, and that is fine, provided the approximation carries a stated accuracy budget, meaning an agreed limit on how far the fast answer may differ from the full overnight run, and someone owns the reconciliation between the two. The Basel Committee's 2024 guidelines on counterparty credit risk management say much the same, at greater length.

Ask yourself: when your desk last wanted a fast answer on a large structured deal, did credit give one, or did the deal go ahead without it?

2. The exposure engine: the right number, ignored

An exposure engine is the computational core of a counterparty credit risk system: the simulation that projects, path by path, what a counterparty could come to owe as markets move. Such engines fail in two ways, and firms tend to prepare for only one of them.

The visible failure is being wrong: calibrations anchored to a gentle history colliding with a violent present. The European gas market of 2022 supplied the canonical example. Futures on TTF, the Dutch hub price that serves as Europe's gas benchmark, settled around €227 per megawatt-hour on 7 March 2022, in the turmoil that followed Russia's invasion of Ukraine, having traded near €345 intraday, roughly twelve times their early-2021 level; by late August the settlement itself stood near €340. An engine calibrated to the preceding five years of European gas prices was not measuring that world. It was remembering a different one.

The subtler failure is being right and unheeded, or right and misread. A simulation produces not one number but a distribution of future exposure, and different statistics of that distribution answer different questions. Peak potential future exposure, the high quantile of the distribution at a stated confidence and horizon, is what a limit should see, because a limit exists to cap what could plausibly be lost. The expected exposure profile, the probability-weighted average across paths, is a pricing and provisioning input: it drives what you charge the counterparty for the risk, and what you set aside against the losses you expect to incur; regulatory capital, where it applies, consumes yet another read-out of the same simulation. The distinction is easy to state and surprisingly easy to lose in practice, and the point is not the alphabet of acronyms: it is that every decision must name the measure it consumes. A firm that manages the book to the expected profile while the limit framework nominally watches the peak is running two views of the same portfolio, and acting consistently on neither. On the illustrative portfolio in the accompanying chart the difference is material: against a

€50m limit, peak PFE of €35.2m leaves €14.8m of headroom, while the expected profile peaks at €19.1m and leaves €30.9m of headroom, more than twice as much. Neither of those is the dashed line on the picture, which is the fast add-on approximation the pre-deal check actually uses, and which tracks the full simulation to within €2.8m across the profile. A desk steering by the wrong curve believes it has room the limit framework never granted; and because the confusion is seldom written down anywhere, nobody notices until it matters.

What good looks like. One simulation, two read-outs, and everyone able to say which number governs which decision. And when the quick approximation (the add-on grid in the picture, which answers in milliseconds at the point of trade) disagrees with the full simulation, someone must own the reconciliation.

Ask yourself: could your credit committee say, without preparation, whether your limits bind on peak or expected exposure, and at what confidence?

3. Collateral and margining: the protection you do not hold

Collateral reduces exposure by exactly as much as its terms deliver, and no more. Three contractual parameters determine how much of the gross number a firm actually wears: the threshold, which is the exposure permitted to run uncollateralised before any margin is called; the minimum transfer amount, beneath which no call is worth making; and the margin period of risk, the assumed interval between a counterparty's final posting of collateral and the close-out of its positions. On the last of these the academic literature has been clear for years: the textbook assumption of ten business days flatters reality, because disputed calls and illiquid positions stretch the close-out, and each added day of the margin period raises collateralised exposure materially above its textbook value. Archegos, once more, supplies the operational lesson. Its margins were set once, at inception, as static terms, and were never recalibrated as the exposure beneath them transformed in scale and concentration, even as Credit Suisse's own dynamic margining tools flagged the problem a month before the default.

The European energy crisis of 2022 added a system-scale lesson to the same point. As Russia's invasion of Ukraine drove gas prices to unprecedented levels, the Bank for International Settlements documented initial margin on benchmark TTF futures rising from about €5 per megawatt-hour in early 2021 to peaks of €79 on ICE and €145 on EEX, the two principal exchanges, while average daily variation-margin calls rose sixteen-fold, touching €3.4 billion on the single most volatile day. Collateral terms negotiated in calm markets were thus tested in violent ones, which is, of course, the only time they ever really are.

What good looks like. Margin terms modelled as they are written, not as intended; margin-period assumptions that survive a dispute; and the residual, uncollateralised number reported beside the gross one wherever decisions are made.

Ask yourself: if your largest counterparty disputed a call tomorrow, how many days of market moves would you wear, and does your exposure number already assume that answer?

4. Liquidity under margin calls: the mitigation that calls cash

Everything in the third panel has a mirror image. Collateral, viewed from the credit desk, is protection; viewed from the treasury, it is a contingent claim on somebody's cash, and half the time that somebody is you. The year 2022 made this duality the best-documented failure mode in the entire picture, because the sums involved were simply too large to miss. As margin requirements surged in the weeks after Russia's invasion of Ukraine, the European Federation of Energy Traders, the trade body for the continent's wholesale energy market, wrote to governments on 8 March 2022 requesting time-limited emergency liquidity support, and supplied the arithmetic for a single hedged producer: initial margin of about €1 billion in the summer of 2021, rising to €4 billion by October and €6 billion by March. Uniper, Germany's largest importer of Russian gas, had seen the strain coming earlier still: it arranged €10 billion of precautionary credit facilities in the first days of January 2022, before the invasion, explicitly to cover margining needs of exactly this kind.

The London Metal Exchange's nickel suspension of the same day, 8 March 2022, remains the extreme case. The covering of a vast short position drove the price to levels the exchange judged disorderly, meaning trading had stopped reflecting genuine supply and demand, and the LME faced \$19.7 billion of margin calls that its own independent review concluded would have put multiple clearing members into default; rather than make the calls, it cancelled the morning's executed trades, unwinding hedges that firms believed they held. That is a credit lesson in its own right, distinct from the liquidity one: protection that depends on an exchange's discretion is less final than it appears. By the autumn, the official backstops put a number on the scale of the problem: a £40 billion UK

guarantee scheme for energy firms, little of it drawn in the end, alongside Danish, Swedish and Finnish schemes worth some €46 billion between them. Backstops are measured by what they render unnecessary as much as by what they pay out.

Precision matters in the best-known case. Margin strain drove Uniper's January facilities; the later rescue, ultimately exceeding €30 billion, was driven principally by the failure of physical supply, as the company bought replacement gas at spot prices to honour contracts Russia no longer served. The two channels compounded one another: credit mitigation became a liquidity event at precisely the moment credit risk itself was peaking.

Regulators drew the conclusion in due course. The Financial Stability Board's 2024 recommendations on liquidity preparedness for margin and collateral calls apply explicitly to commodity traders, and the expectation is now written down plainly: a firm that posts or receives margin holds, and stresses, a buffer sized to the stressed call, not the average one.

Ask yourself: does your liquidity buffer reflect what your credit function actually agreed in the CSAs, and has anyone modelled the day a large margin call and a credit stress hit at the same time?

5. Stress and scenario: when correlations go to one

A stress-testing framework fails in two distinct ways: through scenarios calibrated to recent history, and through scenarios nobody acts upon. The year 2022 supplied the definitive example of the first. Price moves in European energy ran far outside any scenario set built on the preceding decade, a decade of ample gas and low volatility, and the firms that coped were those asking a different question of their scenarios. Not "what has happened before?" but "what breaks us?": a reversal of the usual direction of stress design, in which the analyst reasons backwards from ruin rather than forwards from history.

The deeper failure lies in wrong-way risk: the tendency of exposure and counterparty quality to deteriorate together, so that a firm's claim grows precisely as its counterparty's capacity to honour it shrinks. The European Central Bank's November 2022 financial-stability analysis described exactly this dynamic at work in energy markets, where banks extending credit lines to traders to fund margin calls were doing so at the very moment the same price moves were weakening those traders. Yet within this category, one distinction deserves to be written into the framework explicitly, because 2022 punished the firms that missed it. A hedged producer facing rising prices is strained for liquidity but often economically stronger, its physical output worth more even as its hedges bleed cash. A counterparty whose creditworthiness itself falls with the same move that raises your exposure is the true wrong-way case. The first calls for a funding answer, the second for a credit one, and a framework that conflates the two will misjudge both.

In the accompanying panel, the stress engine re-runs the entire book under shock, and its essential feature is that it re-runs the collateral, the margin calls and the limits alongside the exposure, because in a genuine stress all three move together, and a scenario that shocks prices while holding the mitigants still is only half a scenario. On the illustrative portfolio the shock takes peak exposure from €35.2m to €57.2m, a 63 per cent rise that carries it €7.2m through the same €50m limit the base case cleared with €14.8m to spare. The simulation behind the panel adds one further lesson worth retaining: under stressed parameters, the exposure peak arrives earlier as well as higher, at around five months instead of eleven. The limit does not merely break. It breaks sooner than the base profile suggests, and a framework watching the base-case timeline is watching the wrong time horizon.

What good looks like. Joint market-and-credit scenarios, reverse stress tests that identify the moves that break the book, and a documented action that follows a breach. A stress test whose result changes nothing is a report, not a control.

Ask yourself: name the scenario that breaks your book. If the answer is a percentile rather than a story, the framework is measuring, not stressing.

6. The foundation: the book nobody can state

Every number in the five panels above is computed on the book: positions, netting sets, legal terms, tenors, physical and financial legs. The least glamorous failure in credit risk is that this foundation is wrong, or right but scattered across systems nobody can reconcile, and the banking sector's own report card makes the point at scale. The Basel Committee's risk-data aggregation principles were published in 2013; a decade on, in 2023, the Basel Committee's own progress report found only two of thirty-one global systemically important banks fully compliant, citing fragmented IT estates and manual workarounds. The relevance is the comparison, not the age of the finding: if banks with the largest compliance budgets in finance are still mostly failing this test after ten years, trading firms with a fraction of that budget have no reasonable claim to be ahead of them.

The version we meet most often in energy and commodities looks like this: the only place the methodology actually lives is inside the code that implements it, nowhere else. That has a specific, recognisable shape: the model is the SQL, because nothing outside the query defines what the model does. The person who wrote it has moved on, and the requirements were never written down, only discussed in meetings that nobody minuted properly. When the firm then decides to rebuild or replace the system, it discovers it cannot state what the current one does, and the project inherits an archaeology phase nobody budgeted for.

What good looks like. Netting sets held as data somebody owns, meaning a definitive record of which trades can be set off against which, and under which signed master agreement. Methodology documented separately from the code that implements it. The architecture and the systems the model runs on recorded and actively maintained, not documented once at build and left to drift. One person should be able to say, of every number on the picture, which system produced it, from what inputs, and on what infrastructure.

Ask yourself: if your PFE engine's author left tomorrow, how long before you could change the model with confidence? If the answer is measured in quarters, the documentation is the code.

A fair objection

A reasonable reader could stop here and object: none of this is new. Each of the six failures above is already the subject of extensive guidance, from the Basel Committee's counterparty credit risk guidelines to the FSB's work on margin preparedness to the CCRO's own best-practice reference. That is true, and it is exactly the point rather than a rebuttal of it. The guidance addresses each panel on its own terms, because each sits with a different owner and answers to a different part of the regulatory apparatus. What none of it addresses is the seventh failure, the one this note is actually about: the handoff between panels, where the credit analyst's limit, the quant team's simulation, treasury's buffer and IT's data model are each individually defensible and collectively unreconciled. The perimeter picture exists because the joins have no natural owner, and no single piece of existing guidance is written to cover a boundary rather than a discipline.

Where the perimeter goes next

The picture is drawn from the world 2022 tested, and the next tests are already visible in a specific direction: exposure moving onto books the current toolkit was not built to see. Power purchase agreements and data-centre offtakes are writing decade-long exposure to unrated special-purpose vehicles onto books built for rated utilities, with no CSA to soften it. The post-2022 drift from cleared markets back to bilateral terms has quietly re-created the exposure clearing had removed, while intraday and algorithmic trading is compressing the gap between the overnight exposure run and the speed of the book itself. Credit insurance capacity, having repriced, now carries exclusions that will be read closely for the first time in the next stress. None of these breaks the perimeter on its own. Together, they are the same lesson 2022 taught, arriving early: a firm that waits for the next crisis to examine its systems will find out what those systems cannot measure at the moment the answer matters most.

The question that finds the joins

One question locates the joins faster than any review can. Take a number that crosses a boundary, and the stressed margin call is the cleanest of them: it is computed in the stress framework, it is governed by collateral call a different team negotiated, and it arrives as a cash requirement on a third desk. Ask what that number is. Then ask who owns it.

A firm whose perimeter works returns one figure and one name. Many firms return two figures, or one figure and a pause, and that gap is the subject of this note. It is not a modelling failure, because no individual model is wrong. It is a failure of the handoff, which is why no single team's improvement plan will close it, and why it survives audits that examine each panel on its own terms.

The perimeter is one system. It is drawn as one picture because it fails as one system, at the joins, where the fragments meet.

Sources

Figures cited: Credit Suisse Special Committee report on Archegos (2021); PRA final notice, July 2023; ECB Financial Stability Review, November 2022; BIS Bulletin No. 77 (2023); Fortum/Uniper announcements, January 2022; EFET letter, 8 March 2022; LME independent review by Oliver Wyman (2023); FSB, Liquidity Preparedness for Margin and Collateral Calls (2024); BCBS progress report on risk data aggregation (2023); BCBS Guidelines for counterparty credit risk management (2024); CCRO, Credit Risk Management Practice Standards (in progress) and CCRO Technology for Credit Risk (2024). The portfolio shown in the accompanying picture is illustrative; its figures are worked examples, not client data.

About Horizon Credit Advisory

Horizon Credit Advisory is a specialist counterparty credit risk advisory for energy and commodity trading, working where credit policy meets the systems that enforce it. We hold no ratings business and have no platform to sell.

Seamus Maher has spent over 25 years delivering credit risk programmes across energy, commodities and banking, including credit risk rollout at BP, CCAR revaluation at Citi and market risk scenario tools at Barclays. Michał Falkowski, PhD, has fifteen years in exposure and PFE analytics, pre-deal limits and stress testing, at BP and Anglo American; his doctoral research was on commodity financialisation, and the perimeter picture this series accompanies is his work.

The observations in these notes come from that delivery experience, not from survey data.

Practice bases in London and Warsaw. horizoncreditadvisory.com