

risk-specific
methodology report

CRITICAL MINERALS

A Structured Framework for Assessing Climate-Driven Risks in ETF Portfolios

CRITICAL MINERALS RISK METHODOLOGY

Document version: 1.0 | Last updated: May 4, 2026

Introduction

Critical minerals such as lithium, cobalt, and copper face climate-driven pressures that are expected to reduce availability and increase price volatility. For companies that depend on these inputs, supply shocks could weaken financial performance and, in turn, weigh on ETF valuations and returns relative to the market expectations that prevailed at the time of initial investment.

This assessment focuses exclusively on forward-looking drivers not yet reflected in market prices. Drivers already in force are assumed to be at least partly priced in and are therefore excluded.

Critical minerals mapping

In this initial release, Kilda's mapping covers five critical minerals::

- **Lithium**
- **Cobalt**
- **Nickel**
- **Copper**
- **Rare Earth Elements** (specifically neodymium Nd, praseodymium Pr, dysprosium Dy, and terbium Tb)

These minerals were selected against two criteria.

- First, each is consumed across multiple industries rather than confined to a single sector.
- Second, global supply is concentrated in a small number of producing countries, many of which carry material exposure to climate-driven drivers.

Coverage will expand as authoritative data for additional minerals matures.

Note: Bulk commodities such as iron ore fall outside the framework. They are mined in many countries and sourced through diversified supply chains, which makes them far less vulnerable to climate-driven disruptions. Cobalt and rare earth elements, by contrast, depend on a handful of producers and face pronounced upstream bottlenecks.

Policy drivers

Policy drivers arise from the global shift toward a low-carbon economy. The broader TCFD transition-risk taxonomy also covers legal, market, technological, and reputational categories; in this initial release, only policy drivers are assessed, since they can be evaluated with defensible precision using authoritative sources such as government agencies, legislative notices, and consultation papers. Three types of policy drivers are considered:

1. Export Restrictions

Export controls on critical minerals are often tied to national strategies for clean energy and advanced manufacturing. Governments may reserve minerals for domestic use in renewable-energy, electric-vehicle, and battery-storage supply chains, reinforcing local value creation while limiting foreign access. These restrictions typically persist or tighten over time, creating lasting pressure on global pricing and availability.

2. Biodiversity and land-use restrictions

Conservation measures such as expanded protected areas, deforestation bans, and no-go zones can restrict or fragment access to mineral deposits, either physically or through extended permitting. While these policies serve biodiversity and climate objectives, they also introduce long-term barriers to mineral supply.

3. Community and Indigenous Rights Enforcement

Requirements for free, prior, and informed consent (FPIC), combined with litigation and local mobilisation, raise the threshold for project approval and extend the time required to clear it. These processes can delay or cancel projects, creating supply-timing uncertainty across affected supply chains.

Note: This category records the supply-chain implications of FPIC and community processes. No value judgment regarding Indigenous peoples or affected communities is implied; only the potential effects on project timelines and dependent supply chains are assessed.

Extreme weather drivers

Extreme weather drivers arise from acute events and chronic shifts in climate patterns that could disrupt the supply and pricing of critical minerals. In this initial release, Kilda's assessment uses three extreme weather drivers:

1. Annual maximum 5-day precipitation

Annual maximum 5-day precipitation measures the maximum accumulated precipitation (both rainfall and snowfall) over any five consecutive days in a given area and year. It serves as a proxy for extreme downpour severity.

2. Daily maximum wet-bulb temperature

Daily maximum wet-bulb temperature captures the combined effects of temperature and humidity. Derived from daily maximum temperature and relative humidity, it represents the lowest temperature achievable through evaporative cooling. This makes it a direct measure of heat-stress conditions relevant to outdoor mining operations and worker safety.

3. Area under extreme drought (SPEI < -2)

The Standardised Precipitation Evapotranspiration Index (SPEI) measures meteorological drought by accounting for both precipitation and potential evapotranspiration. A SPEI value below -2 is classified as extreme drought, while a value near zero represents near-normal conditions. This indicator reports the area affected by extreme drought in at least one month of the year.

Note: In this initial release, the selected drivers are restricted to climate variables that are measurable, projection-ready, and available from the Climate Impact Explorer under the NGFS 'Current Policies' scenario.

Analysis

In this section, each mineral is carried from producer mapping, through driver tagging and severity-likelihood scoring, to an aggregated mineral-country colour, and finally to the production-weighted risk score that links it to company exposures.

Mapping minerals to producing countries

For each mineral, the smallest set of producing countries whose cumulative share reaches or exceeds 80% of global mine production is identified. Three rules govern the selection:

1. The list is capped at five countries.
2. If the threshold is reached with fewer than three countries, the top three are listed regardless.
3. If two or more countries are tied at the cutoff rank, all tied countries are included, even if this adds one beyond the cap.

The primary source is the USGS Mineral Commodity Summaries. National geological surveys (e.g., Geoscience Australia, Natural Resources Canada) are used for country-level verification.

Example: Lithium producers covering $\geq 80\%$: Australia 37.39%, Chile 20.82%, China 17.42%, Zimbabwe 9.35%.

Tagging future policy drivers

For each mineral-country combination, all relevant instances of future policy drivers are identified and tagged. The three categories defined in the Policy drivers section apply:

1. **Export Restrictions**
2. **Biodiversity and Land-Use Constraints**
3. **Community and Indigenous Rights Enforcement**

Identification and tagging rely exclusively on free, authoritative sources such as government legal notices, legislative records, and official consultation papers. A standardised source list and search protocol for each driver category are maintained separately and updated quarterly.

Only forward-looking drivers are recorded. A driver qualifies as forward-looking if it represents a new measure or a material tightening of an existing measure. Existing statutes already in force are treated as current operational conditions, assumed to be at least partly priced in, and excluded.

Example: Lithium - Chile - Community and indigenous rights enforcement.

Chile will apply ILO Convention 169 indigenous consultation systematically to all new lithium Special Operation Contract (CEOL) assignments under the National Lithium Strategy, conditioning the pace and viability of new lithium project licensing.

Tagging future extreme weather drivers

For each mineral-country combination, all relevant future extreme weather drivers are identified and tagged. The three indicators defined in the Extreme weather drivers section apply:

1. **Annual Maximum 5-day Precipitation**
2. **Daily Maximum Wet Bulb Temperature**
3. **Area under extreme drought (SPEI < -2)**

The analysis uses median projected changes for 2030 under the NGFS Current Policies scenario, sourced from the [Climate Impact Explorer \(CIE\)](#). For each indicator, the median projection and the 5th-to-95th-percentile range are extracted.

Full precision is retained internally. Reported values are rounded to one decimal place.

Example: Lithium - Chile - Area under extreme drought (SPEI < -2) (2030, NGFS Current Policies).

Median projected change: 37.0%. Uncertainty range (5th-to-95th percentile): 27.9% to 46.3%.

Scoring framework for future policy drivers

For every mineral-country-driver combination, two scores are assigned:

1. **Severity:** The degree to which the driver, if realised, would disrupt future supply of the mineral or materially raise its costs.
2. **Likelihood:** The directional confidence that the driver will materialise or intensify. This is an ordered classification (Green, Yellow, Red), not a probabilistic estimate.

The rubric for policy drivers is set out below. The rubric for extreme weather drivers follows in the next section.

	Green (low)	Yellow (medium)	Red (high)
Severity	Narrow, non-binding, or low-coverage constraint. Limited scope, easy carve-outs, or measures that don't minerally hinder cross-border supply or permitting.	Meaningful but partial constraint. Could hinder supply in some producing geographies or stages (mine/refine/ship), or introduce recurring delays/uncertainty.	Broad and binding constraint with clear potential to reshape access or timelines in one or more systemically important producing countries or supply stages.
Likelihood	Low probability that a binding constraint will materialize.	Discernible pathway but not locked in. Drafts/consultations exist; intent is visible, but timing, scope, or enforcement remain uncertain.	High probability—legal/institutional machinery is in motion toward enforcement or tightening.

Example: Lithium - Chile - Community and indigenous rights enforcement.

Severity: Yellow. Systematic indigenous consultation is required before CEOL allocation under the National Lithium Strategy, adding a binding procedural step that can delay or condition projects, but not an outright prohibition.

Likelihood: Red. The government has already led 12 consultations, with 8 completed in key salares, showing entrenched practice that will continue shaping CEOL licensing pace.

Scoring framework for future extreme weather drivers

For each mineral-country-driver combination, the same two dimensions used for policy drivers apply: severity and likelihood. The definitions are adapted for climate projections:

1. **Severity:** The median projected change in the indicator by 2030. Larger changes signal greater potential to disrupt supply or raise costs.
2. **Likelihood:** The directional confidence that the hazard will worsen by 2030, evaluated from the CIE 5th-to-95th-percentile projection range. Values are oriented so that higher values equal worse outcomes. This is an ordered classification, not a probabilistic estimate.

The likelihood rubric is the same across all three indicators:

	Green (low)	Yellow (medium)	Red (high)
Likelihood	Entire 5th-to-95th range \leq 0 (no projected worsening)	5th-to-95th range crosses 0 (worsening possible but not certain)	Entire 5th-to-95th range $>$ 0 (worsening projected with high confidence)

Severity thresholds differ by indicator and are set out below.

Annual maximum 5-day precipitation

Thresholds are expressed as percentage change in the median projection.

	Green (low)	Yellow (medium)	Red (high)
Severity	\leq 1 %	$>$ 1 and \leq 5 %	$>$ 5 %

Daily maximum wet-bulb temperature

The Climate Impact Explorer does not provide explicit scoring thresholds for this variable. The severity cutoffs below were derived from research literature and observed conditions. Thresholds are expressed as absolute wet-bulb temperature values ($^{\circ}$ C).

	Green (low)	Yellow (medium)	Red (high)
Severity	\leq 10 $^{\circ}$ C	$>$ 10 and \leq 21 $^{\circ}$ C	$>$ 21 $^{\circ}$ C

Area under extreme drought (SPEI $<$ -2)

The Climate Impact Explorer does not provide explicit scoring thresholds for this variable. The severity cutoffs below were derived from research literature and observed conditions.

Thresholds are expressed as a percentage of area affected by extreme drought in at least one month of the year.

	Green (low)	Yellow (medium)	Red (high)
Severity	\leq 20 %	$>$ 20 and \leq 40 %	$>$ 40 %

Example: Lithium - Chile - Area under extreme drought (SPEI $<$ -2) (2030, NGFS Current Policies).

Severity: Yellow. The median projected area under extreme drought is 37.0%, which falls above the 20% threshold and at or below the 40% threshold.

Likelihood: Red. The 5th-to-95th-percentile range is 27.9% to 46.3%, which falls entirely above zero.

Aggregate scoring

Aggregation proceeds in three steps.

1. **First**, severity and likelihood are combined into a single colour for each individual driver.
2. **Second**, all drivers within each category (policy drivers and extreme weather drivers) are averaged into one category-level colour.
3. **Third**, the two category-level colours are combined into a final mineral-country colour.

Step 1: Combining severity and likelihood

For each mineral-country-driver combination, the severity and likelihood scores assigned in the previous sections are combined into a single colour using the matrix below:

Severity \ Likelihood	Green	Yellow	Red
Green	Green	Yellow	Yellow
Yellow	Yellow	Yellow	Red
Red	Yellow	Red	Red

Example: Lithium - Chile - Community and indigenous rights enforcement.

Severity: Yellow, Likelihood: Red → Score: Red

Step 2: Aggregating within each driver category

Each combined colour from Step 1 is encoded numerically: Green = 1, Yellow = 2, Red = 3. For policy drivers and extreme weather drivers separately, the arithmetic mean of all driver scores within the category is computed. The mean is then mapped back to a colour:

	Green	Yellow	Red
Threshold	Mean < 1.5	1.5 ≤ mean < 2.5	Mean ≥ 2.5

These thresholds apply identically to both categories.

Example: Lithium - Chile - extreme weather drivers.

- Annual maximum 5-day precipitation: Yellow (2)
- Area under extreme drought (SPEI < -2): Red (3)
- Daily maximum wet-bulb temperature: Red (3)

Mean: $(2 + 3 + 3) / 3 = 2.67$. Since $2.67 \geq 2.5$, the aggregated extreme-weather colour is Red.

Handling climatically irrelevant indicators. Where an extreme weather indicator is not climatically relevant for a producing country, a Green (1) score is assigned. This reflects negligible exposure rather than missing data and keeps the averaging denominator consistent across all mineral-country combinations.

Step 3: Combining policy and extreme weather scores

The aggregated policy-driver colour and the aggregated extreme-weather colour from Step 2 are combined into a single mineral-country colour using the same matrix as Step 1:

Policy \ Extreme weather	Green	Yellow	Red
Green	Green	Yellow	Yellow
Yellow	Yellow	Yellow	Red
Red	Yellow	Red	Red

Linking companies to mineral-producing countries

This section links company-level mineral exposures to the producing countries that supply those minerals, and assigns a production-weighted risk score to each company-mineral pair. The process has three steps.

Step 1: Mineral to producing-country mapping (supply side)

From the preceding sections, the following inputs are established:

- The current list of critical minerals (lithium, cobalt, nickel, copper, and rare earth elements).
- For each mineral, the latest USGS Mineral Commodity Summaries (MCS) production table, used as the primary source.

The coverage rule from the Mapping minerals to producing countries section applies ($\geq 80\%$ of global mine production, minimum three countries, capped at five countries if no ties occur)

The mapped producer list is cross-checked against the British Geological Survey (BGS) World Mineral Statistics. Material discrepancies are flagged. USGS remains the tie-breaker, given its standardised annual coverage.

Example: Lithium producers covering $\geq 80\%$: Australia 37.39%, Chile 20.82%, China 17.42%, Zimbabwe 9.35%.

Step 2: Company to mineral exposure (demand side)

Each company's sector, industry, and primary geographical market are obtained from free, authoritative disclosures, APIs, or registries (e.g., Yahoo Finance). Sector and industry are mapped to GICS codes where needed.

A standardised industry-mineral exposure map is then applied. This map uses open, sectoral evidence (IEA technology-demand chapters, World Bank minerals-for-climate reports, BGS/USGS technical notes) to define which minerals each industry structurally depends on. No company-specific private data is used.

Example: Battery and EV manufacturers depend on lithium, nickel, and cobalt (battery cathodes).

Note: Evidence links are documented at the sector level, ensuring the mapping remains source-based and reproducible.

Step 3: Company to mineral and producing-country risk

Reliable company-specific trade-flow data for critical minerals are not publicly available. Supply chains are also heavily concentrated upstream in a small number of producing countries. This step therefore links each company's mineral exposure to primary producing countries rather than to inferred import origins.

For each company-mineral pair, three inputs are combined:

Industry dependency: the company's industry is mapped to the minerals on which it structurally depends, using the exposure map from Step 2.

Example: Consumer Cyclical, Auto Manufacturers → Rare Earth Elements (Nd, Pr, Dy, Tb) used in EV traction motors.

Producer concentration: the top producing countries covering ≥ 80% of global mine output for that mineral are taken from Step 1.

Example: Rare Earth Elements: China 68.5%, United States 11.4%, Burma 7.9%.

Risk assignment: each producing country already carries an aggregate forward-looking colour for the mineral in question, scored as Green = 1, Yellow = 2, Red = 3 (from the aggregate scoring section).

These scores are combined using production-share weights to create a single, production-weighted mineral risk:

$$\text{Country + Mineral Risk Score}_m = \sum_{c \in C_m} p_{m,c} \times R_{m,c}$$

where

C_m = set of producing countries for mineral m ,

$p_{m,c}$ = country c 's share of global production for mineral m ,

$R_{m,c}$ = that country's aggregate risk score for mineral m (Green = 1, Yellow = 2, Red = 3).

The resulting weighted score is mapped back to a colour using the same thresholds as the category-level aggregation: Green if < 1.5, Yellow if $1.5 \leq \text{score} < 2.5$, Red if ≥ 2.5 .

Example: Rare Earth Elements. China: Red (3); United States: Yellow (2); Burma: Green (1).

$$\text{Country + Mineral Risk Score}_m = (0.685 \times 3) + (0.114 \times 2) + (0.079 \times 1) = 2.36 = \text{Yellow}$$

The weighted score of 2.36 falls in the Yellow range ($1.5 \leq 2.36 < 2.5$).

Portfolio aggregation

Per-mineral risk is aggregated to portfolio-level outputs in two stages: each company is first assigned a single 0–100 risk score derived from its mineral-level colours, and ETFs are then summarised through paired coverage and intensity figures.

Company risk score

Once each company is mapped to its relevant critical minerals and each mineral carries an aggregate colour (Green, Yellow, Red) from the preceding steps, a single risk score is constructed per company.

The score satisfies four design requirements:

1. **Colour dominance.** Any Red outweighs any number of Yellows and Greens. Any Yellow outweighs any number of Greens.
2. **Breadth sensitivity.** The score reflects the number of flagged minerals, not just the worst colour.
3. **Fixed range.** The score spans 0 (no mapped exposure) to 1 (exposure to all tracked minerals, all Red).
4. **Transparency.** The score is a structured coding of colours and counts, not a calibrated estimate of financial loss.

Definitions

Let:

- M = total number of minerals tracked in this framework.

For company i :

- r_i = number of minerals flagged Red (high risk).
- y_i = number of minerals flagged Yellow (medium risk).
- g_i = number of minerals flagged Green (low risk).
- $K_i = r_i + y_i + g_i$ = number of mapped minerals for the company.

Step 1: Assigning a colour band

Each company is placed in one of four bands based on its worst colour:

Condition	Band (L_i)
No mapped exposure ($K_i = 0$)	0
Green only ($r_i = 0, y_i = 0, g_i > 0$)	1
At least one Yellow, no Red ($r_i = 0, y_i > 0$)	2
At least one Red ($r_i > 0$)	3

- **No mapped exposure:** $K_i = 0 \rightarrow$ band $L_i = 0$.
- **Green-only:** $r_i = 0, y_i = 0, g_i > 0 \rightarrow$ band $L_i = 1$.
- **Yellow (no Red):** $r_i = 0, y_i > 0, \rightarrow$ band $L_i = 2$.
- **Red:** $r_i > 0 \rightarrow$ band $L_i = 3$.

Step 2: Assigning a colour band

Computing the normalised score

$$S_i = \frac{L_i}{4} + \frac{1}{4} \times \frac{3r_i + 2y_i + g_i}{3M} \in [0,1]$$

where:

- S_i = normalised company risk score for company i on a 0–1 scale.
- M = total number of minerals tracked in the framework (currently 5).
- L_i = worst-colour band assigned to company i (0, 1, 2, or 3 as defined above).

Properties

The construction guarantees the following:

- Companies are placed in one of four ordered bands: no exposure, Green-only, Yellow, Red.
- All Red-exposed companies score above all Yellow-only companies, and all Yellow-only companies score above all Green-only companies, regardless of mineral counts.
- Within a band, the score rises smoothly with the total weighted number of flagged minerals. Red contributes more than Yellow, and Yellow more than Green.
- A score of 0 represents no mapped critical-mineral exposure. A score of 1 represents the framework's maximum: exposure to all tracked minerals, all of them Red.

Worked example

Company X depends on three critical minerals: cobalt (Red), lithium (Yellow), and copper (Green). With $M = 5$:

- $r = 1, y = 1, g = 1, K = 3$
- Worst colour is Red, so $L = 3$
- Band Position: $\frac{3}{4} = 0,75$
- Weighted count: $(3 \times 1) + (2 \times 1) + (1 \times 1) = 6$
- Within-band adjustment: $\frac{1}{4} \times \frac{6}{15} = 0.10$
- $S = 0.75 + 0.10 = 0.85$
- Investor-facing score: $100 \times 0.85 = 85$

Investor-facing scale

For reporting, the normalised score is rescaled to 0–100:

$$\text{Company Risk Score}_i = 100 \times S_i$$

The four bands map to the following score ranges:

Band	Score Range	Interpretation
No mapped exposure	0	No critical-mineral risk identified within this framework.
Green only	25-50	Exposed only to Green-flagged minerals. Score rises with the number of Green dependencies.
Yellow (no red)	50 to 75	Exposed to at least one Yellow-flagged mineral. Score rises as Yellow and Green dependencies accumulate.
Red	75 to 100	Exposed to at least one Red-flagged mineral. With five tracked minerals

The distances between bands (for example, between 25 and 50) do not represent calibrated differences in financial loss. They encode the ordinal ordering of colours and a weighted count of flagged minerals in a transparent and consistent manner.

ETF score construction

For each ETF, company-level scores are aggregated in two steps: coverage and intensity.

Let w_i be the portfolio weight of the company i in the ETF, with $\sum_i w_i = 1$.

Step 1: Computing coverage

Coverage measures the share of the ETF's total weight that is mapped to a company score (that is, holdings with at least one mapped mineral, $K_i > 0$):

$$C_{ETF} = \sum_{i: K_i > 0} w_i$$

C_{ETF} in $[0,1]$

$$\text{Coverage \%} = 100 \times C_{ETF}$$

Holdings with no mapped exposure or with insufficient data are excluded from the intensity calculation and reduce overall coverage.

Step 2: Computing intensity

Intensity measures the average risk level of the covered portion of the ETF.

First, portfolio weights are renormalised over covered holdings only:

$$\tilde{w}_i = \frac{w_j}{\sum_{j: K_j > 0} w_j} = \frac{w_i}{C_{ETF}} \text{ for } K_i > 0$$

Then intensity is computed as the weighted average of company scores:

$$I_{ETF} = \sum_{i:K_i>0} \tilde{w}_i S_i$$

$$I_{ETF} \in [0,1].$$

$$Intensity \% = 100 \times I_{ETF}$$

If $C_{ETF} = 0$ (no holdings carry a mapped company score), I_{ETF} is undefined by construction and is set to 0 by convention.

Worked example

An ETF holds three companies with the following weights and scores:

Company	Weight (w_i)	Company score (S_i)	Mapped minerals (K_i)
Tesla	0.50	0.85	3
Samsung	0.30	0.52	2
JPMorgan	0.20	-	0

- JPMorgan has no mapped critical-mineral exposure, so it is excluded from the intensity calculation
- Coverage $C = 0.50 + 0.30 = 0.80$ (80%). The remaining 20% of the ETF's weight is unmapped.
- Renormalised weights: $\tilde{w}_{Tesla} = \frac{0.50}{0.80} = 0.625$; $\tilde{w}_{Samsung} = \frac{0.30}{0.80} = 0.375$
- Intensity: $(0.625 \times 0.85) + (0.375 \times 0.52) = 0.531 + 0.195 = 0.726$

Reading this result: 80% of the ETF's weight carries mapped mineral exposure, and that exposed portion has a high average risk intensity of 72.6 out of 100.

Limitations and future development

This methodology is in its initial release. The constraints below apply at this stage and will be addressed in subsequent versions:

- Mineral coverage is limited to five entries. The set will expand as authoritative data for additional minerals matures.
- Policy drivers cover three categories only. The legal, market, technological, and reputational categories of the TCFD transition-risk taxonomy fall outside the critical-minerals scope and will be addressed elsewhere in Kilda.

- Extreme weather analysis relies on a single climate scenario (NGFS Current Policies) and a single time horizon (2030). Additional scenarios and longer horizons will be incorporated in subsequent iterations.
- The industry-mineral exposure map operates at sector level. Precision will improve as reliable company-specific supply-chain data becomes available.
- Severity and likelihood rubrics for policy drivers are qualitative; inter-analyst reliability therefore depends on a maintained library of reference cases.

Conclusions

This framework provides a transparent assessment of forward-looking critical-minerals exposure for companies and ETFs.

Country-level mineral risks are expressed as ordinal colours, company exposures as a 0-100 score, and ETF outputs as paired coverage and intensity figures, with the pairing preserving the distinction between portfolio breadth and the average risk of the covered portion.

The framework is intended for use as one input within a wider investment due-diligence process.

Notice of Confidentiality and Limitation of Liability

This document is confidential and proprietary, and is furnished solely for informational purposes to the intended recipient. No part of this document may be distributed, reproduced, disclosed, or otherwise shared, in whole or in part, by any means or in any form, without the prior written consent of the author or issuing organization. The information contained herein is provided on an "as is" basis, without representations or warranties of any kind, whether express or implied, including but not limited to warranties of accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Neither the author nor the issuing organization shall be liable for any direct, indirect, incidental, consequential, or special damages arising out of or in connection with the use of, or reliance upon, this document or its contents. Recipients are advised that no reliance should be placed on any statement, methodology, or recommendation set forth herein for decision-making purposes without independent verification and, where appropriate, the advice of qualified professional advisors. All rights not expressly granted are reserved.