Locational Marginal Emissions

Technical Documentation

JUNE 2026

Highlights

- MISO's Locational Marginal Emissions (LME) method provides nodal, 5-minute data to estimate how system level emissions change in response to changes in load.

- The LME dataset published by MISO provides a robust characterization of locational emissions estimates through the addition of several novel enhancements

- Transparent methodology and real-time data validation provide an impartial, credible, and useful LME signal to meet the needs of MISOs stakeholders

MISO

misoenergy.org

Table of Contents

LME calculation Overview

Introduction to LMEs

Locational Marginal Emissions, or LMEs, describe how emissions would change in response to increases or decreases in electrical demand at different points on the grid. The basic idea is simple: if demand rises at a specific location, generators that are currently on the margin must adjust their output to keep the system balanced. Some generators may ramp up, some may ramp down, and each of these changes has an associated emissions impact.

To understand LMEs, it helps to look at two main pieces of information:

1. How marginal generators would adjust their output when demand changes at a given location.
2. How much each of those generators emits for every unit of power they produce.

If the response of each generator and its emissions rate are known, the total change in emissions resulting from a small change in demand at each location can be estimated.

In practice, determining how much each generator would adjust output is more complicated than simply matching generator output to load demand. The power grid must also obey physical limits. For example, some transmission lines may be operating at their maximum allowable flow. These are called binding constraints. When a line is binding, any change in demand that shifts power flows across that line must be offset with generator adjustments that keep the flow from exceeding its limit. Different locations on the system affect power flows in different ways, so the system must choose a specific combination of generator redispatch to keep the grid within its limits.

Another factor is system losses. Delivering power is not perfectly efficient, and the amount of loss changes depending on where power is produced and consumed. A small increase in demand at one location may slightly raise or lower total system losses, because it would result in power flow changing along lines with different loss factors. These changes must also be offset by generator adjustments so that supply and demand remain balanced.

All of these requirements, balancing load, holding transmission constraints constant, and accounting for changes in system losses, can be described through a set of relationships that determine how marginal units must respond. By organizing these relationships into a linear system of equations, it becomes possible to solve for the unique set of marginal generator adjustments associated with a change in demand at any location on the system (described by the redispatch matrix).

Redispatch matrix calculation details

From a mathematical perspective, the system of equations can be expressed in terms of several matrices. The binding constraint sensitivities can be described by a sensitivity matrix (sometimes also called a shift factor matrix), which describes the sensitivity of the power flows along each binding constraint to a shift in load at every node in the system. Marginal loss sensitivity factors (MLSFs) describe the sensitivity of system losses to changes in load at each node. The goal is to solve for the redispatch matrix, which describes the amount that each marginal unit is redispatched in response to a change in load at any node in the system.

The resulting equation is:

$$\begin{bmatrix} s_L \\ S_N \end{bmatrix} = \begin{bmatrix} s_L \\ S_G \end{bmatrix} [R]$$

Where s_L is the marginal loss sensitivity factor vector, S_N is the nodal sensitivity matrix, S_G is the marginal generator sensitivity matrix, and R is the redispatch matrix. This can also be described as: Loss-augmented nodal sensitivity matrix = loss-augmented marginal generator sensitivity matrix * redispatch matrix.

These matrices have dimensions $A_{(BC+1)\times N} = B_{(BC+1)\times G} R_{G\times N}$, where BC is the number of binding constraints, N is the number of nodes for which LMEs are being calculated, and G is the number of marginal units. Thus, the redispatch matrix $R_{G\times N} = B^{-1}{}_{(BC+1)\times G}\, A_{(BC+1)\times N}$

From a mathematical perspective if G is greater than BC+1, then the system of equations is underconstrained, and there is no unique solution. In real world operations, G is often greater than BC+1, so to find a solution the Moore–Penrose pseudoinverse of the loss-augmented marginal unit sensitivity matrix is used.

Based on how the redispatch matrix is calculated, it becomes evident that an accurate solution (and thus an accurate LME value) depends highly on:

1. Accurately identifying the set of binding constraints whose sensitivities must be considered
2. Accurately calculating the marginal loss sensitivity factors
3. Accurately identifying the set of marginal units that would respond to changes in load

None of the three above factors are directly available as inputs to this calculation. Thus, the following sections describe the approach to identifying the correct set of binding constraints and marginal units and calculating the loss factors.

Input parsing and validation

The inputs into the LME calculation come from the MISO's ex-ante, real-time market clearing solution from its security-constrained economic dispatch (SCED) engine.

Use of ex-ante vs. ex-post results

MISO's ex-ante market run is the solution that is used to send dispatch instructions to units in the real time market. The ex-post run is used for market settlements to determine the price that is actually paid to units. Historically, the ex-ante and ex-post solutions have been nearly identical, but in recent years these solutions have started to diverge more often. Generally, they are still similar except during capacity emergencies. MISO's LME calculation uses ex-ante results for several reasons:

- Ex-ante results better align with the dispatch instructions sent to generators, making them a more accurate reflection of real operational redispatch
- The ex-post market run is generally more volatile
- Block-loaded units cannot be price-setting in the ex-ante run. However, they can appear price-setting ex-post, even though they would not actually redispatch if load changed.

Case validation

During each market run, MISO may run several possible cases (scenarios) with operators normally approving a single case. However, when grid conditions change rapidly, the originally approved case may no longer represent real-time operating conditions, and additional cases may be approved.

In such circumstances, data from the most recently-approved case is used. This is one reason why input data is uploaded after the start of the market run, to ensure that the most recently-approved case is used. Upon upload, the system validates that all input data use the same case ID.

Input data files

For the LME calculation, five different types of inputs are used:

- Binding transmission constraint identities
- Binding reserve constraint identities
- Binding constraint sensitivities
- Unit dispatch results
- LMP values

Identifying binding constraints

Section 5.1.1.2 of MISO's Business Practice Manual 002 describes multiple types of constraints that contribute to the calculation of the marginal congestion component (MCC) of the LMP:

- Transmission constraints
- Sub-regional power balance constraint
- Reserve procurement regulation-up (RPRU) deployment constraint
- Reserve procurement regulation-down (RPRD) deployment constraint
- Zonal Reserve procurement contingency reserve (RPRCR) deployment constraint
- Zonal reserve procurement short-term reserve (RPRSTR) deployment constraint

The current LME calculation only considers data regarding transmission constraints (including market-to-market constraints), sub-regional power balance constraints, and zonal reserve procurement short-term reserve deployment constraints. Data for RPRU, RPRD, and RPRCR constraints are not currently available.

Although MISO identifies the set of constraints that are binding, a binding constraint is not considered in this calculation if it is violated. A constraint is considered violated when its limit cannot be respected or when the limit has been overridden. While violated constraints affect the calculation of the Marginal Cost Component (MCC) of the LMP, once a constraint is violated, it no longer affects the redispatch solution and is therefore excluded from the marginal unit response calculation.

To determine whether a binding constraint is violated, its line flow and shadow price are compared to the structure of its Transmission Constraint Demand Curve (TCDC). The TCDC defines penalty breakpoints and corresponding shadow price levels. If the observed flow and shadow price place the constraint on a part of the curve associated with a violation (for example, between or beyond penalty breakpoints), then the constraint is treated as violated rather than binding.

As shown in Figure 1, the TCDC can have different breakpoints and different penalty values, but one common example is a curve with a $1,000 penalty if the line exceeds 100% of its rated capacity, and a $2,000 penalty if the line exceeds 102% of its rated capacity. The constraints identified by MISO as "binding" include any constraints at or above 100% of their rated capacity. If the flow on the line exactly equals either of the breakpoint values, the constraint is considered binding for these calculations. If the flow on the constraint is between the first and second breakpoints, or greater than the second breakpoint, the binding constraint is considered to be violated and is not included in the LME calculation.

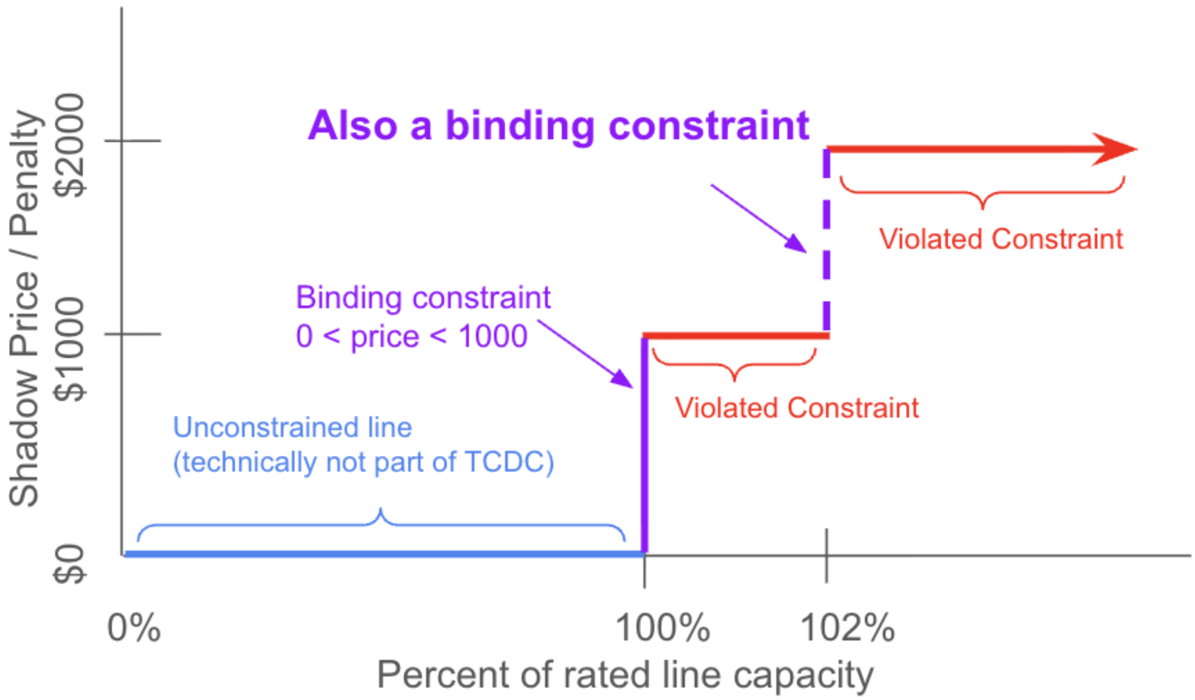

Figure 1: An example TCDC with breakpoints at 100% and 102% of the rated line capacity. The constraint is binding if the line flow is exactly 100% or 102% of the capacity, or if the shadow price is not equal to $0, $1,000, or $2,000.

Reserve constraints do not have TCDCs, so if the flow is greater than the limit, the reserve constraint is violated.

Validating binding constraint sensitivities

The binding constraint sensitivity data contains the sensitivity values for each binding constraint identified for each node in the system. To reduce the size of the data, sensitivity values are only provided for CPnodes with non-zero sensitivities. Zero values are filled for all CPnodes for which a sensitivity is not provided.

Because the sensitivity data is such a critical component of the LME calculation, a number of validation checks are performed on the data provided.

First, the sign convention of the sensitivity data is standardized based on the direction of flow on the binding constraint. If a constraint is at its limit but the flow is negative (indicating flow in the opposite direction of the sign convention), the sign of the sensitivity values for that constraint are also flipped.

The sensitivity values that are provided are validated by reconstructing the marginal congestion cost (MCC) component of the LMP from the sensitivity values and the shadow price of the constraint, and comparing that value to the actual MCC. If the reconstructed MCC value is not within $0.01 of the actual MCC value, the sensitivity value for that node is flagged as a “bad” value and removed from the dataset. This validation also ensures that the missing sensitivity values are actually zero values.

For any bad sensitivities, there are certain circumstances when the correct value for that CPnode can be derived:

1. If there is a single binding constraint in the interval, the sensitivity value for a node is calculated as (actual MCC at node i) / (binding constraint shadow price)
2. If there is only a single missing sensitivity value for a binding constraint, the sensitivity for that missing node is calculated as (actual MCC - reconstructed MCC) / (binding constraint shadow price)
3. If there are multiple sensitivities missing, but they are all for regional directional transfer (RDT) constraints, which should be only 0 or 1, all (0,1) combinations of the missing nodes are enumerated, and if exactly one combination reproduces the actual MCC, that combination of sensitivities is used.

For these circumstances the calculated MCC value replaces the missing MCC value.

Identifying locational marginal units

While MISO’s dispatch results include a “marginal unit” flag, this flag indicates the set of “energy marginal” units, which are the set of units that set prices in the absence of congestion. Energy marginal units are those units where their marginal cost equals the MEC component of the LMP. Even in the absence of congestion, it is frequently the case that multiple units are “energy marginal” if they have the same offer price for the marginal unit of energy. These energy marginal units form the basis of MISO’s existing “fuel on the margin” public report.

In an interval without congestion (no binding constraints), the set of energy marginal units is the same as the set of “locational marginal” units (LMUs). However, when congestion exists, the energy marginal units are not necessarily a subset of the locational marginal units. For the purpose of the LME calculation, it is necessary to consider the set of “locational marginal” units (LMUs), which are marginal considering congestion.

This identification requires using other data from the unit dispatch results. While in theory, LMUs would be the set of units whose marginal cost is equal to the LMP at their node, in reality, the complexity of the SCED calculation, rounding errors, and other factors means that this approach is not certain to identify the correct or complete set of LMUs. The uncertainty around the identity of the true set of LMUs is one source of potential error in the LME calculation.

This uncertainty is mitigated using a two-step approach to LMU identification:

1. Identify the full set of possible LMUs using a set of theoretical filters and analytical approaches
2. Identify the subset of these units that result in the redispatch matrix with the best overall validation metrics

Identifying possible LMUs

To identify the full set of possible LMUs from the set of dispatched units, a series of data filters and analytical approaches are used.

Removing non-marginal units
The first step is to remove any units that cannot be marginal from a theoretical perspective.

1. Remove units with a control status or "offline" or "online but off control." These units cannot be dispatched by the market clearing engine. This leaves units that have a control status of "load following" or "regulating"
2. Remove block-loaded units. Block loaded units are those that cannot be dispatched incrementally: they can only be off or operating at a fixed setpoint. These units are identified as units whose minimum dispatch level equals its maximum dispatch level.

Identifying units that match the economic definition of LMU
Once the above set of units are removed, locational marginal units are identified based on their price. As introduced above, LMUs are defined from a theoretical perspective as the set of units whose marginal cost equals the LMP value at the node. Because the SCED calculation is complex and subject to rounding errors, this theoretical relationship may not always hold exactly. To account for this uncertainty, each unit's "LMP differential" is calculated. This value represents the absolute difference between the unit's nodal LMP and its marginal cost:

$$LMP\ Differential_n = |LMP_n - MC_n|\ \ for\ each\ node\ n$$

Since LMUs should theoretically have an LMP differential of $0.00, all units are ranked from lowest to highest differential. If multiple units have the same LMP differential, they are secondarily ranked in order from lowest to highest nodal LMP. Once ranked, the first BC+1 units with the lowest differential are flagged as potentially marginal. "BC + 1" is the number of unique non-violated binding constraints + 1, and represents the number of units necessary to calculate a unique redispatch solution. In the case that multiple units have the same LMP differential and nodal LMP, these units are treated as a single group, and if any one unit in the group would be flagged, all units will be flagged.

While in theory, the true set of LMUs should be represented by the units with the lowest LMP differential, in practice this is often not the case. Often, the redispatch matrix calculated with this set of units is unable to reconstruct the sensitivity matrix, MLSF vector or LMP values with zero error, or requires extreme redispatch magnitudes. Most of this error and extreme behavior was the result of not having the appropriate unit in the set of LMUs that could be used to effectively manage congestion based on its location on the grid relative to the binding constraint. To address this, a set of analytical approaches to identify additional units that would be most likely able to manage congestion were developed.

Units with non-zero sensitivities to binding constraints

To balance how a change in load affects a constraint, there must be at least one LMU that influences that constraint (meaning it has a non-zero sensitivity). However, sometimes a single marginal unit affects multiple constraints (non-zero sensitivity for multiple constraints), so adjusting it would also change flows on others. For units that affect only one constraint, the one with the lowest LMP differential is flagged because it can balance that constraint on its own. If no unit is unique to a single constraint, a step-by-step process is used to identify which marginal units to flag.

1. Rank all units with a non-zero sensitivity to the given constraint from lowest to highest LMP differential
2. If there are any units that have the same sensitivity to a constraint, all units with duplicate sensitivities are dropped, only keeping the first (lowest LMP differential). This is because from the perspective of balancing power flow on the constraint, units with identical sensitivities are equivalent and do not give additional degrees of freedom in the calculation.
3. Flag the unit with the lowest LMP differential.
4. If a uniquely sensitive unit was not already flagged (i.e. there are not yet 2 total units flagged), flag a second non-uniquely sensitive unit with the next lowest LMP differential. Selecting 2 units gives us enough degrees of freedom to independently manage each constraint.

In addition to flagging up to 2 units as described above, the unit that has the largest magnitude sensitivity to that constraint is also flagged (if not already flagged). This unit is flagged because with a larger sensitivity, the unit needs to be redispatched less to achieve the same impact on the binding constraint flow, which potentially helps keep the norm of the redispatch matrix low.

Unit in each "congestion cluster"

The marginal congestion cost can indicate which regions of the grid react in similar ways to the constraints that are binding. When several nodes have similar MCC values, they can be thought of as belonging to the same congested pocket of the system. In principle, having at least one marginal unit in each pocket helps us manage congestion more effectively. To find these pockets, a K-Means clustering algorithm is used to group all nodes into BC+1 clusters based on their MCC values. Nodes with an MCC of zero are placed in their own group because an MCC of zero means that node is not experiencing congestion. After these clusters are created, the unit in each cluster with the lowest LMP differential is identified.

Unit in each lost load zone for reserve constraints
In intervals with reserve constraints, each constraint may apply to a specific local resource zone (LRZ) in the footprint, based on a loss of load scenario in that LRZ. To ensure that there is a unit that can be used to manage that reserve constraint, the lowest LMP differential unit located in the LRZ to which the reserve constraint applies is flagged.

Unit in each subregion for power balance constraints
Subregional power balance constraints, also called regional directional transfer constraints, represent limits on how power can flow between MISO's two major subregions, MISO North and MISO South. In the case of a binding power balance constraint, the lowest LMP differential unit in MISO North and the lowest differential unit in MISO South are flagged.

In many cases, a unit will be flagged by more than one of the above filters.

Selecting the best set of LMUs

While the above filters ideally identify the full set of possible LMUs, they may often have false positive identifications, identifying units that are not actually LMUs. To identify the "correct" set of LMUs from the list of possible candidates, different combinations of units are iteratively tested to identify the combination that results in the best validation metrics for the run.

For example, if 10 marginal units and 5 binding constraints were present, the process is to:

1. Test the full set of 10 units
2. Test all combinations of 9 units by removing one unit at a time from the list of 10 units
3. Identify which run performed the best, and the unit that was removed to get that result
4. Remove the worst-performing unit from the list of 10 units, and save the validation metrics for the best run with 9 units
5. Test all combinations of 8 units by removing one unit at a time from the list of 9 remaining units
6. Repeat steps 3-5 until there have 6 units (BC+1) left
7. Identify the run with the best overall metrics, and get a list of all of the units that were removed to get that result.
8. Remove those units from the list of total units, and use that to then calculate the final result.

The “best” result is defined as the result that has the lowest error, considering the lowest redispatch norm, and the number of units. For more details on these metrics, see the “validation” section. Specifically, the results are ranked from lowest to highest based on the following metrics:

1. Combined Mean Absolute Percent Error (MAPE) for sensitivity and MLSF reconstruction. These two metrics are averaged to weigh them equally. Generally, if these MAPEs are not zero, something is wrong with the set of units that were selected.
2. MAPE for LMP reconstruction: the LMP values at all nodes should be able to be reconstructed, but because LMPs are not directly considered in the redispatch matrix calculation, this is only considered in the case of a tie between the combined MAPE above.
3. Redispatch matrix infinity norm
4. MAPE for LMP reconstruction from the Marginal costs

All MAPE metrics are rounded to 2 decimal places (nearest percent), and the norm is rounded to the nearest whole number. This is done to ensure that if the preferred metric is only slightly below an alternative, this does not over-select on that metric and ignore all other metrics. Effectively, if two MAPEs are within 1% of each other, the next metric is considered.

A unit in the initial set of LMU candidates is only a “false positive” if removing it strictly improves the rounded reconstruction/norm metrics. If removing a unit merely ties those metrics, the unit is likely co-marginal (e.g. two units with equal marginal cost but different emission rates) with another unit. Co-marginal units should be retained because removing one of them would be arbitrary, and if the co-marginal units have different emission rates, this would arbitrarily impact the LME results. Keeping both co-marginal units means that the LME will reflect a blend of both unit’s emission rates.

Thus, if all the above metrics tie, the results are then ranked in descending order by the number of marginal units, which will prefer the tied result with the largest number of marginal units.

When considering different combinations of the same marginal units, the number of marginal units will be the same, and thus this metric cannot serve as a tiebreaker in the case that all of the primary error metrics tie. In this case, the non-rounded versions of the above metrics are used to break any ties.

Calculation of Marginal Loss Sensitivity Factors

Marginal Loss Sensitivity Factors (MLSFs) are not directly output by MISO's market clearing engine so must be calculated. LMP data for each node, including the LMP's Marginal Energy Component (MEC), Marginal Loss Component (MLC), and Marginal Congestion Component (MCC) is used to derive the MLSF. Per MISO's BPM 5.1.1.1, the MLC is calculated as:

$$MLC = -MLSF \times MEC$$

Thus, using the LMP data, the MLSF for each node can be calculated as $MLSF = -\frac{MLC}{MEC}$.

However, because the LMP data is only available to the nearest whole cent (2 decimal places), the accuracy of the MLSF value will be impacted by rounding.

This approach also has limitations when the MEC is zero or close to zero, which can happen when a zero marginal cost unit is on the margin. The three scenarios that are challenging are:

- MEC = 0, MLC = 0. In this case, the MLSF could be any number
- MEC = 0, MLC is not zero. This is mathematically impossible but does occur.
- Both MEC and MLC are close to zero (e.g. 0.01). This results in very high (>0.5) MLSF values which are likely not correct and result in extreme values in the LME calculation.

In the above cases, the MLSF value for these nodes is set to 0, effectively assuming no impact on system losses.

Locational Marginal Emissions

Assigning heat rate curve-based emission rates

One criticism of previous marginal emissions estimates is that they utilize annual-average emission rates to quantify the emissions of marginal units, and these annual averages do not reflect the marginal emission rate of the unit at its current output level. While the level of error introduced by annual-average emissions factors is likely low relative to other sources of error and uncertainty in marginal emissions estimates, the criticism is still valid.

To address this, Singularity uses its emissions calculation engine, called GRETA (Generator REal-Time emissions Assignment). For each interval, GRETA assigns a time-specific fuel type and emission rate to each generator based on the generator's modeled heat rate curve. These heat rate curves are modeled using public data that each generator reports to the EPA and/or EIA. Public data is the best source of consistent, high-quality data across the MISO network model because all large fossil generators report emissions and generation data.

The use of GRETA, rather than static average emissions factors, enables a greater level of accuracy because:

- Certain generators cofire multiple fuels or switch fuels throughout the year. GRETA estimates a time-specific fuel type based on past generator behavior.
- The efficiency (heat rate) of generators varies based on their output level (i.e., a generator operating at full capacity may be more efficient than when it is operating at 50% capacity). GRETA uses heat rate curves to assign output-specific emission rates to generators.
- The heat rate of a generator can also vary by season based on ambient temperatures and seasonal requirements to operate pollution controls. GRETA includes season-specific heat rate models for each generator.
- Generators generally are less efficient during start-up and can also burn different fuel during startup. GRETA attempts to detect startup events and assign startup-specific fuel types and heat rate curves to the data.

Full documentation for this Singularity methodology can be found in the GRETA documentation, available as a separate document in the documentation menu of MISO's marginal emissions dashboard.

Emission rates for carbon-free resources

Generators with carbon-free fuel types (hydro, nuclear, wind, solar) are assigned an emission rate of 0.

Reservoir hydro power is distinct from other carbon-free resources in that in certain hydrologic zones, there is a finite amount of water (and thus power) available on a seasonal basis. In these cases, the use of hydropower at one time has a marginal opportunity cost because that power cannot then be generated at a later time, which has a secondary marginal impact on emissions. This problem has been widely acknowledged in the academic research on marginal emissions for at least a decade but has yet to be solved for this style of marginal emissions estimation method. Furthermore, this opportunity cost only applies to certain types of hydropower in certain locations under certain conditions, and so would be challenging to quantify with any degree of accuracy. Thus, the convention of assigning a 0 emission rate to marginal hydropower is followed.

Energy Storage Emissions

Energy storage technologies such as batteries and pumped hydro do not produce emissions on their own. However, a battery must charge before it can discharge, and it only holds a limited amount of energy at any time. Because of this, *when* a marginal battery discharges its energy can influence when it needs to charge later, and it also reduces how much energy it can provide in the future. For example, if a battery discharges in one hour, it may either need to charge in the next hour to restore its energy level or it will have less energy available to discharge in the hour after that. These shifts in charging and discharging can create secondary emissions effects at different times. A marginal storage unit will generally have some secondary consequential emissions impact, but it is very difficult to measure these impacts in a reliable way. This challenge is similar to hydropower and is well known in academic research but does not yet have a clear solution.

Assigning an emission rate of 0 to marginal storage units would be inappropriate because it may incentivize behavior that actually results in non-zero emissions impacts on the grid. Likewise, assigning a high emissions rate to storage may encourage behavior that results in less impact than expected. Thus, to be conservative, storage discharge is assigned a regional, annual-average emission rate with the idea that use of such a factor would not push the LME rate to either extreme.

While quantifying secondary, counterfactual, marginal impacts of energy storage is not currently well researched, and may not ever be possible for this style of marginal emissions methodology, it may be possible to refine the conservative average approach in the future. For example, instead of using an annual-average, regional-average value, a more temporally and spatially-granular average emissions rate that may be more relevant to the specific time and location in which the storage unit is operating could be used.

Backstop emission rates

In the case that unit-specific, heat rate curve-based emissions estimates are not available, various backstop methods are applied in the following order:

- Unit-specific, annual-average emission rates from the most recently available year of data in the Open Grid Emissions dataset
- If the fuel type of the resource is known, but not the specific identity of the unit, fleet-average, annual-average emission rates from the most recently available year of data in the Open Grid Emissions dataset are used. Fleet-average refers to the average emission rate of all units of the same fuel type (e.g. natural gas) in the same grid region (e.g. MISO).
- Any other units, including units with an "unspecified" fuel type, are assigned the regional-average, annual-average emission rate from the most recently available year of data in the Open Grid Emissions dataset.

CO_2-equivalent Values

All emissions estimates are reported in CO_2 equivalent (CO_2e).

The combustion of fuel for power generation results in emission of various greenhouse gases (GHGs) including Carbon Dioxide (CO_2), Methane (CH_4), and Nitrous Oxide (N_2O). Each of these GHGs contributes to global warming differently, as described by its Global Warming Potential (GWP). These GWPs can be used to calculate a CO_2-equivalent (CO_2e) value. The Intergovernmental Panel on Climate Change (IPCC) regularly updates these values over time in published Assessment Reports (AR). Emissions occurring in 2021 and later use the AR6 values, so CO_2e estimates are currently calculated in OGE using AR6 IPCC weights.

Locational Marginal Fuel

Locational Marginal Fuel (LMF) data describes the fuel mix of the resources that are being redispatched. These data can be useful for contextualizing the LME value at any node by revealing what type of units were contributing to the LME value being high or low.

Assigning fuel categories

To determine the LMF value, a fuel category must first be assigned to each marginal unit. Fuel category assignment, like emissions assignment, can be variable over time, and uses a combination of GRETA models and static fuel category assignments where an EIA mapping is unavailable.

Each category groups multiple energy sources, which are assigned at the plant or generator level according to EIA data. In each interval, a fuel type is assigned based on sub-plant level attributes (if sub-plant level mapping is available). If only plant-level mapping is available, the fuel type is assigned based on the primary attribute of the entire plant. The fuel type is assigned as follows:

- If the (sub)plant is a single fuel (sub)plant, the “primary fuel” identified in the Open Grid Emissions Dataset is assigned. This is based on the fuel type reported in EIA-860 and actual fuels consumed as reported to EIA-923 (see the methodology here).
- If the generator fuel switches or co-fires multiple fuels at once, the fuel type is assigned based on a regression model trained on historical fuel consumption data that the generator reports to EIA-923. This model identifies the probability of a certain fuel type being consumed at the plant in a given month. The model then assigns the fuel type with the greatest probability of consumption in that month.

The fuel categories that can be used for LMF data include:

- Biomass: includes wood, biogas, landfill gas, and other biomass fuels
- Coal
- Hydro
- Natural Gas
- Solar

- Storage: includes both battery storage and pumped hydro
- Wind
- Unspecified: No information is available about the fuel type of this resource. These are often distributed energy resources (DERs)

Net Redispatch

As described above, for a hypothetical +1MW load shift, some combination of marginal units will be redispatched up or down so that the net redispatch sums to approximately +1MW (the impact of marginal losses means that this net redispatch often does not sum directly to +1.0MW.

For each fuel type represented by the marginal units, the net redispatch for each node is calculated. For example, if the redispatch for a given node was:

- Unit 1 (Gas): +0.3
- Unit 2 (Gas): -0.2
- Unit 3 (Coal): +0.9
- Unit 4 (Wind): +0.05
- Unit 5 (Wind): -0.15

Then the net redispatch reflected in the LMF values would be:

- Gas: +0.2
- Coal: +0.9
- Wind: -0.1

The precise unit of measurement for these values is "MW of redispatch per MW of load shift." However, for the sake of clarity, and because this essentially represents a rate of MW per MW, the data is presented as a percentage value. So in the above example, the Gas value of +0.2 would be presented as +20%.

Because the redispatch values for units of the same fuel type are calculated on a net basis, this could potentially mask information about negative redispatch. For example, if there was one gas unit that had +1.5 redispatch, and another that had -0.5 redispatch, this would just show up as +1.0 redispatch. In the case that these two gas units have dramatically different emission rates, the net redispatch value may mask information about why a marginal emission rate is higher or lower than expected, even if there is only a single marginal fuel type.

LME Validation

For every LME calculation, a set of validation metrics are calculated to evaluate how close the derived redispatch matrix matches the actual (but unknown) redispatch matrix.

Redispatch matrix validation metrics

Although the “true” redispatch matrix is unknown, there are multiple matrix identities that suggest that if the true redispatch matrix has been identified, it should be able to used to exactly re-construct the nodal sensitivity matrix, the marginal loss sensitivity factors, and the LMP values at each node. Because the redispatch matrix is what propagates costs/sensitivities from the marginal nodes to the rest of the nodes in the system, it should be able to be used to reconstruct these nodal values.

- **Sensitivity matrix reconstruction error**: because the linear system of equations requires generator sensitivities to balance load sensitivities, the load sensitivities should be able to be reconstructed by multiplying the redispatch matrix by the sensitivities at the generator node. This metric helps measure numerical error in the matrix inversion (which may result from ill-conditioning or use of a pseudoinverse). If the error is 100% for a given constraint, this can also indicate that the set of LMUs is missing a unit that has a non-zero sensitivity to that constraint.
- **MLSF vector reconstruction error**: this metric again validates an identity that is enforced in the redispatch matrix construction itself. Any error here may be the result of the use of pseudo-inverse calculations, numerical precision, or missing MLSF values at marginal generator nodes.
- **LMP vector reconstruction error**: This metric is calculated two ways: once by multiplying the LMP by the redispatch matrix, as well as by multiplying the LMU marginal costs by the redispatch matrix. In theory, the LMPs at any node should be a linear combination of the marginal costs of the marginal units. This validation provides a sanity check on the consistency of the redispatch matrix with the economic theory of the market. When calculating these metrics, adjusted LMP values, which subtract out the congestion cost of violated binding constraints, are used since the calculation does not consider violated constraints. These error metrics are generally higher than the previous two, which may be a result of non-SCED pricing effects on the LMPs.
- **Redispatch matrix norm**: the norm indicates the magnitude of the generator redispatch relative to a perturbation. Specifically, the infinity norm asks what is the largest absolute redispatch in response to a 1MW shift in load at any node? Generally, these values should be small, indicating that a small load shift requires a small overall response. However, during highly constrained intervals, the norm may be higher. High norms could also indicate an incorrect LMU set, or a nearly-singular matrix that results in large redispatch values. Intervals with larger norms are more likely to have extreme LME values.

- **Redispatch matrix condition number**: this metric tells us about the numerical stability of the redispatch matrix solution, or how sensitive the results are to small errors in the inputs. Generally, exactly constrained problems will have a lower condition number than underconstrained problems. A high condition number can indicate that the LME values are potentially not well-identified from the available SCED data.

Marginal condition context
Generally, a well-constrained system of equations will have exactly 1 more LMUs than there are binding constraints. However, in reality, this relationship often does not hold. These metrics help us identify information about the system conditions in this interval and the resulting marginal state.

- **LMU:BC relationship**: this metric tells us if there are exactly BC + 1 locational marginal units identified, or if the system of equations is underconstrained.
- **Number of selected marginal units**
- **Number of binding constraints:** the total number of non-violated binding constraints of all types.
- **Number of unique binding constraints:** Sometimes, a single constraint can appear as both a reserve constraint and a binding constraint. This metric tells us the number of uniquely-named binding constraints.

LMU selection process context
The following metrics provide useful diagnostics regarding the iterative LMU selection process and the set of marginal units that were selected.

- **List of selected marginal units**: this is the list of marginal CPnodes that were identified as marginal by the LMU selection algorithm.
- **List of marginal units removed by the iterative process**: This is the set of CPnodes that were identified as potentially marginal by the selection filters, but were ultimately removed as candidates during the iterative selection process.
- **Maximum LMP differential**: This metric describes the maximum "LMP differential" value of any single selected LMU. Ideally this metric would be close to zero.
- **Maximum LMU marginal cost and maximum LMP for the interval**: These metrics indicate when shortage pricing may be in effect, which can impact the accuracy of the LMU selection.

LME Solution performance
These metrics describe how complete the LME solution is, and whether the solution contains any extreme or unusual values.

- **Total number of LME values**: The total number of valid CPnodes for which an LME value should be produced.
- **Number of extreme LME values**: The number of LMEs outside of the range of +/- 20,000 $lbCO_2e$/MWh.

- **Number of missing LME values:** The number of nodes for which an LME value could not be calculated, either due to missing or invalid input data.
- **Min/median/max LME values**: These values provide a snapshot of the range of LME values across all nodes in an interval.

Data Outputs

Data Release Timing

LME values are generally expected to be published within 5 minutes of the start of each market interval. On occasion, there will be missing data, which could either result from the unavailability of market results for that interval or the data failing data quality or validation checks.

Confidentiality

With any new data product, MISO assesses whether the data release presents a risk of revealing market-sensitive or confidential information. At a high-level, MISO believes that publication of LME data presents no additional risks on top of any risks associated with the publication of LMP data.

As with LMPs, LME data reveals no information about the identity of any specific marginal units, and due to the complexity of the calculation, it would not be practical to reverse engineer or back-calculate any operational characteristics (output levels, heat rates, etc.) of any specific unit.

LMEs describe system-level changes, so the location of the marginal unit(s) that affect the LME may not be anywhere near a given node, even in the case of congestion.

Future Enhancements

Shortage pricing
During shortage conditions, MISO can use emergency resources, including emergency range of resources, emergency DR, load modifying resources, external resources, and emergency energy purchases. During times of Operating Reserve and/or other reserve scarcity, Ex Ante LMPs will be impacted by Scarcity Prices determined by Reserve Demand Curves and will be capped at the Value of Lost Load ("VOLL"). In the unlikely event of an Energy deficiency, all LMPs and MCPs will be set equal to the VOLL. Under emergency pricing, emergency resources as described above are cleared based on their Proxy Offer that is established as the maximum of the Emergency Offer Floor and the resource's offer up to a maximum of the Energy Offer Hard Price Cap ($2,000).

During times of shortage pricing, the current method may struggle to identify the correct set of marginal units, which is observed as higher error rates for these intervals. This is because a key element of the locational marginal unit identification methodology relies on comparing the unit's marginal cost to the nodal LMP, and because data about reserve demand curves or proxy offers are not currently used.

Over time, work will be done to incorporate shortage pricing data into the calculations to ensure that locational marginal units can be correctly identified during these periods.

Imports

The current methodology only considers market-participating unit dispatch results to identify marginal units, and does not consider any data about imports into MISO. Currently, the only data about external resources that is considered in the calculation are external asynchronous resources (EARs).

Loss factors

Currently, marginal loss sensitivity factors (MLSFs), which are part of the redispatch calculation, are not directly provided as data inputs. Instead, these MLSFs are derived from the marginal loss component (MLC) and marginal energy component (MEC) of the LMP, since the MLC = -MLSF*MEC. However, the MLC and MEC values are only available to the nearest cent (two decimal places), which means that the derived MLSF values may be subject to rounding error. Generally, these rounding errors are small, but in periods with MECs near zero, these errors will be potentially larger. Furthermore, in intervals with an MEC of $0.00, the MLSF cannot be calculated due to division by zero.

Missing or invalid binding constraint sensitivities

In some intervals, input data is missing binding constraint sensitivity values for certain nodes, or contains sensitivity values that fail the validation process described above. In these cases, LME values for the affected nodes cannot be calculated. If an affected node happens to be an LMU node, then the redispatch matrix cannot be calculated.

Furthermore, sub-regional power balance constraints are understood to have binary sensitivities of either 0 or 1. However, there are instances where such constraints have non-binary sensitivities. It is unclear what impact such values have on the redispatch calculation.

Storage emissions

As discussed in depth above, storage assets can have secondary marginal emissions impacts, which are not reflected in the current methodology. In the future, further refinements to this method could be considered

Glossary

AR: Assessment Report (from the IPCC). Each successive version of the report has a corresponding number.

AR6: The sixth major synthesis of global climate science findings published by the Intergovernmental Panel on Climate Change.

BC: Binding Constraint. A transmission limit actively managed in market operations that restricts flows and influences dispatch and prices.

CH_4: Methane, a potent greenhouse gas emitted from natural and human sources, with a high global warming impact per molecule.

CO_2: Carbon dioxide. A primary greenhouse gas produced by combustion and industrial processes.

CO_2e: Carbon dioxide Equivalent. A metric expressing the climate impact of different greenhouse gases in terms of the amount of CO2 with the same warming effect.

CP Node: Commercial pricing node. A MISO-market construct representing a market participant.

EARs: External Asynchronous Resources. Resources outside a balancing authority that can provide energy or grid services without synchronized connection.

EIA: Energy Information Administration. This agency collects, analyzes, and disseminates energy information to inform policy making, efficient markets, and public understanding of energy and its interaction with the economy and the environment.

EPA: Environmental Protection Agency. This agency develops and enforces regulations, gives grants, studies environmental issues, and publishes information for the purpose of protecting human health and the environment.

GHGs: Greenhouse Gases. Gases in the atmosphere that trap heat.

GRETA: Generator REal-Time emissions Assignment. A method or tool for assigning real-time emissions to specific generators.

GWP: Global Warming Potential. Measure of how much infrared thermal radiation a greenhouse gas added to the atmosphere would absorb over a given time frame, as a multiple of the radiation that would be absorbed by the same mass of added carbon dioxide. GWP is 1 for CO_2.

IPCC: Intergovernmental Panel on Climate Change. This group is the United Nations' body for assessing science related to climate change.

LME: Locational Marginal Emissions. A measure of the incremental change in emissions resulting from generating one additional unit of electricity at a specific location.

LMF: Locational Marginal Fuel. The fuel type(s) of the marginal units serving an incremental change from generating one additional unit of electricity at a specific location.

LMU: Locational Marginal Unit. The generator(s) on the margin whose dispatch would change in response to a small change in load at a given node in MISO's real-time market.

LMP: Locational Marginal Price. The cost of supplying the next increment of electricity at a specific node, reflecting energy, congestion, and losses.

LRZ: Local Resource Zone. A defined geographic area used for resource adequacy and planning, with zonal requirements such as LCR and PRMR assessed per LRZ.

MCC: Marginal Congestion Cost. The congestion component of LMP attributable to active transmission constraints in MISO's markets.

MEC: Marginal Energy Cost. The energy component of LMP reflecting the marginal cost of producing the next unit of energy absent congestion and losses.

MLC: Marginal Loss Cost. The loss component of LMP reflecting transmission energy losses.

MLSF: Marginal Loss Sensitivity Factor. A factor indicating how incremental injections or withdrawals affect losses; in MISO's LME method, MLSFs are derived from the LMP loss component.

MISO: Midcontinent Independent System Operator

MW: Megawatt

N_2O: Nitrous oxide. A long-lived greenhouse gas emitted from agriculture, combustion, and industrial processes.

RDT: Regional Directional Transfer. A limit on directional power transfers between MISO sub-regions (i.e., North/Central and South).

SCED: Security Constrained Economic Dispatch. The optimization process that dispatches generation at least cost while respecting reliability constraints.

TCDC: Transmission Constraint Demand Curve. A pricing mechanism that reflects the cost of relaxing a transmission constraint in energy markets.

VOLL: Value of Lost Load. An estimate of the economic cost to consumers when electricity is not supplied.

For a complete glossary of emissions terms, see the "MISO Emissions Glossary," available as a separate document in the documentation menu of MISO's marginal emissions dashboard.

Documentation last updated: July 1, 2026