

PNNL-39552

Modeling the Reference Coordination Architecture: An Exercise in Architectural Validation

July 2026

Joshua J Yip
Shiva Poudel
Orestis Vasios
D Jonathan Sebastian-Cardenas
Alexandre B Nassif
Rabayet Sadnan
Md Touhiduzzaman
Todd Wall
Jaime T Kolln
Chris Klasen

Pacific Northwest
NATIONAL LABORATORY

U.S. DEPARTMENT of ENERGY

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights**. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <https://www.ntis.gov/about>
Online ordering: http://www.ntis.gov

Modeling the Reference Coordination Architecture: An Exercise in Architectural Validation

July 2026

Joshua J Yip
Shiva Poudel
Orestis Vasios
D Jonathan Sebastian-Cardenas
Alexandre B Nassif
Rabayet Sadnan
Md Touhiduzzaman
Todd Wall
Jaime T Kolln
Chris Klasen

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Acknowledgments

This project was supported by the Department of Energy, Office of Electricity. The authors would like to thank Joe Paladino, William Parks, and Sandra Jenkins for their support and contributions to shaping the scope and direction of this work.

Acronyms and Abbreviations

AMI	advanced metering infrastructure
API	application program interface
BESS	battery energy storage system
DAMO	distribution area market operators
DER	distributed energy resource
DMO	distribution market operator
DO	distribution operator
kVA	kilovolt-amperes
kvar	kilovolt-amperes reactive
kW	kilowatt
kWh	kilowatt-hour
LDDL	large dynamic digital load
MO	market operator
MW	megawatt
MW/h	megawatt-hour
OPF	optimal power flow
PV	photovoltaic
RCA	reference coordination architecture
SCADA	supervisory control and data acquisition
T&D	transmission and distribution
TO	transmission operator

Contents

Acknowledgments ii
Acronyms and Abbreviations iii
1.0 Introduction 7
2.0 Overview of Architecture 8
2.1 DER Registration and Coordination 8
2.1.1 DER Registry 9
2.2 Market Coordination 11
2.2.1 Market Data Warehouse 12
2.3 Operational Coordination 13
3.0 Coordination Modules 15
3.1 DER Registration and Coordination 15
3.1.1 Bid Crafting 15
3.1.2 Procedure and Output of Bid Crafting 18
3.2 Market Coordination 21
3.2.1 Market Validation 21
3.2.2 Market Monitoring 22
3.3 Operational Coordination 22
3.3.1 Operational Validation 22
3.3.2 T&D Node Visibility 25
3.3.3 Delivery Optimization 25
3.3.4 Distribution System State Estimation 26
4.0 RCA Modeling and Validation 27
4.1 Testbed 27
4.2 Case Studies 28
4.2.1 Large Dynamic Digital Load 28
4.2.2 Extreme Cold Weather Event 38
4.2.3 Cascading Blackout 45
4.3 Summary of Case Studies Validation 56
5.0 Conclusions 58
6.0 References 59

Figures

Figure 1. Modules, Processes, and Information Exchange for DER Registration and Coordination 8
Figure 2. Module, Process, and Information Exchange for Market Coordination 12

Figure 3. T&D Coordination 14
Figure 4. Operational Validation 14
Figure 5. Procedure for Bid Crafting by an Aggregator 20
Figure 6. Iterative Procedure Between MO and DOs to Achieve Market Clearing 24
Figure 7. Swimlane Diagram Depicting Workflow for Walk Phase of LDDL Case Study 29
Figure 8. Swimlane Diagram Depicting Workflow for Jog Phase of LDDL Case Study; Only Distribution Part is Shown as Bulk Grid Sequence Remains the Same 32
Figure 9. Swimlane Diagram Depicting Workflow for Run Phase of LDDL Case Study; Only Distribution Part is Shown as Bulk Grid Sequence Remains the Same 35
Figure 10. Distributed Optimization and Local Market for Distribution Services 36
Figure 11. Swimlane Diagram Depicting Workflow for Walk Phase of Extreme Temperature Case Study 39
Figure 12. Swimlane Diagram Depicting Workflow for Jog Phase of Extreme Temperature Case Study 41
Figure 13. Swimlane Diagram Depicting Workflow for Run Phase of Extreme Temperature Case Study 44
Figure 14. Swimlane Diagram Depicting Workflow for Walk Phase of Blackout Case Study 46
Figure 15. Swimlane Diagram Depicting Workflow for Jog Phase of Blackout Case Study 48
Figure 16. Swimlane Diagram Depicting Workflow for Run Phase of Blackout Case Study 51
Figure 17. Distributed Optimization Across Area Sub-Circuits for Blackout Restoration 52
Figure 18. Wholesale Market and Centralized Coordination 56
Figure 19. Distribution Market and Distributed Coordination 57

Tables

Table 1. Schema for Only Table in Type-Blind Model of DER Registry 9
Table 2. Schema for Main Table in Type-Aware Model of DER Registry 10
Table 3. Schema for BESS Table in Type-Aware Model of DER Registry 11
Table 4. Schema for PV Table in Type-Aware Model of DER Registry 11
Table 5. Market Data Warehouse Structure and the Role of Each Subfolder in the Bid Life Cycle 13
Table 6. Data-Class Schema for Market Participation Ruleset 15
Table 7. Data-Class Schema for Market Conflict and Dual Participation Identification 16
Table 8. Data-Class Schema for Service Opportunity 16

Table 9. Data-Class Schema for Operating Envelope....17
Table 10. Data-Class Schema for Contract....17
Table 11. Data-Class Schema for Qualification....17
Table 12. Data-Class Schema for Service Bid....18
Table 13. Data-Class Schema for Product v. Value Curve (ProdValueCurve)....19
Table 14. Data-Class Schema for Product v. Value Point (ProdValuePoint)....19
Table 15. Data-Class Schema for Service Award....24

1.0 Introduction

This document provides a modeling guide to support implementation and validation of the proposed Reference Coordination Architecture (RCA) (Touhiduzzaman, Poudel et al. 2026). The RCA was developed to provide principles, requirements, and structure for a coordination framework composed of modular functions that support distributed energy resource (DER) coordination, market coordination, and operational coordination. The purpose of the modeling effort is twofold. First, it establishes a structured guide for implementing modeling exercises that evaluate whether the RCA can support coordinated grid and market operations across institutionally distinct actors. Second, it applies three representative use cases—Large Dynamic Digital Load (LDDL), Extreme Cold Weather Event, and Cascading Blackout (Touhiduzzaman, Wall et al. 2026)—across the RCA Modeling and Validation Testbed, referred to hereafter as the testbed, to provide evidence that the architecture is scalable, generic, and extensible across a range of operating conditions and levels of system complexity.

The RCA defines a coordination framework through which transmission operators (TOs), distribution operators (DOs), market operators (MOs), and aggregators (or customers) exchange information, validate decisions, and coordinate market and operational activities. The architecture includes coordination modules such as bid crafting, market validation, and operational validation, which define how information, models, and decisions are exchanged among actors to ensure awarded services remain feasible and reliable from both transmission and distribution (T&D) perspectives. The modeling effort described in this document is intended to translate those architectural concepts into executable workflows, actor-specific data environments, and validation procedures.

A central principle of this modeling approach is that grid participants do not have universal access to one another's internal information. Instead, each actor retains responsibility for its own models, forecasts, datasets, and operational processes. For example, TOs maintain transmission network models, while DOs maintain feeder models and customer- or sensor-related data. These data assets are treated as locally owned and are not assumed to be visible to other actors unless they are intentionally shared through defined coordination or market processes. This approach supports data sovereignty, explicit information exchange, separation of data and logic, and modularity across the coordination framework.

The three case studies are organized using *Walk-Jog-Run* progression. In this document, Walk, Jog, and Run represent increasing levels of architectural maturity and operational complexity, while also reflecting the expected evolution of coordination capability over time. The Walk stage establishes the initial validation basis using more centralized coordination and limited resource participation. The Jog stage introduces greater visibility, expanded coordination, and more flexible participation of DERs. The Run stage represents a more advanced future-grid condition in which distributed operations, localized needs, market interactions, and automated coordination functions become more complex and dynamic.

The testbed evaluates whether the RCA can accommodate these increasing levels of complexity while remaining sufficiently generic to support all three case studies. The validation process focuses on whether the information specified by the architecture is sufficient for coordination modules to execute, whether data and models can be communicated to the appropriate actors while respecting ownership and confidentiality constraints, whether decision-making sequences can be executed as intended, and whether market and operational coordination processes can interact effectively.

2.0 Overview of Architecture

The RCA decomposes a coordination framework according to functional requirements and then associates them with modules. Adopting with minimal change the framework, requirements, and modules of the RCA, the present document suggests preliminary designs for realizing the modules, particularly to fulfill energy services.

This section provides an overview of the coordination architecture introduced in the RCA. It consists of three main compartments: DER Registration and Coordination, Market Coordination, and Operational Coordination.

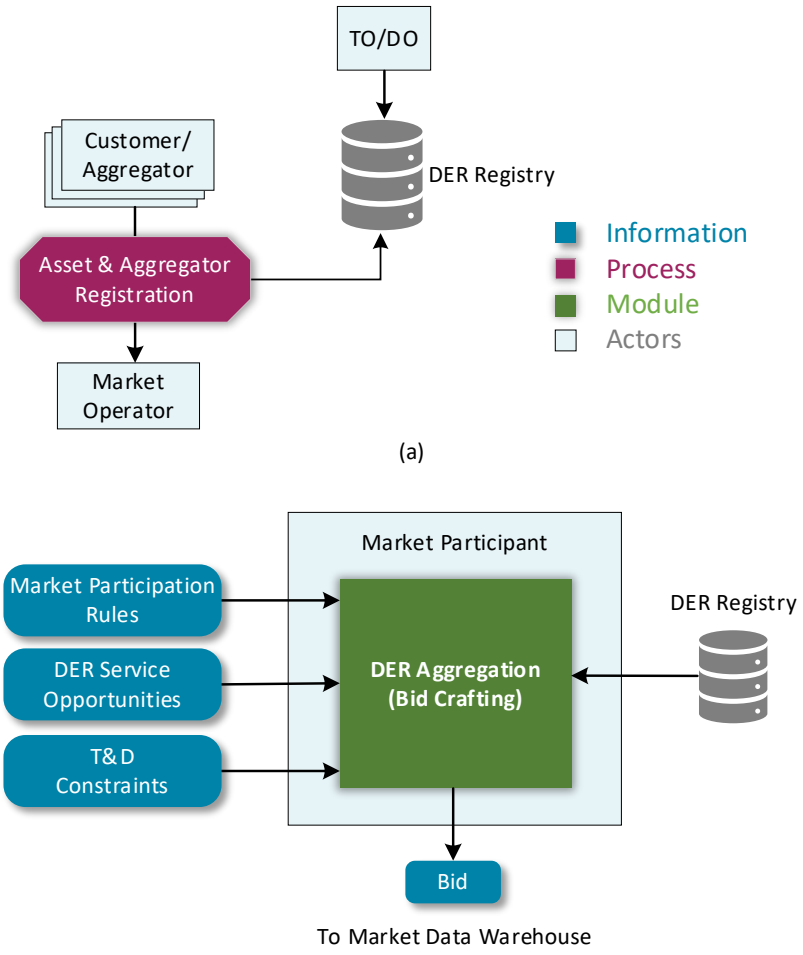


Figure 1. Modules, Processes, and Information Exchange for DER Registration and Coordination

2.1 DER Registration and Coordination

Figure 1, drawn from the RCA without modification, depicts the compartment of the coordination architecture responsible for registration of DER assets and coordination of the bid-crafting process.

The Asset and Aggregator Registration process serves as the entry point for participation in the coordination architecture. In this process, DER assets are enrolled into the market through an aggregator, which acts as the representative entity for one or more customer-owned resources. During registration, the aggregator submits asset information such as resource characteristics, operational capabilities, ownership information, location, and any other attributes required for market participation and operational coordination. The registration process updates the DER Registry, creating or modifying records that describe the registered resources and their relationship to the aggregator. The DER Registry therefore becomes the authoritative repository of resource information that can later be used by validation, market participation, operational coordination, and dispatch functions.

After the underlying DER assets have been registered, the aggregator itself must also register with the MO. This registration establishes the aggregator as an authorized market participant and associates the aggregator with the portfolio of registered DERs maintained in the DER Registry. Once both asset registration and aggregator registration are completed, the aggregator becomes eligible to participate in subsequent market activities on behalf of its registered resources. The TO/DO may also access information from the DER Registry to obtain visibility into registered resources within their operational area and to support future coordination and validation activities.

2.1.1 DER Registry

The DER Registry serves as a shared reference accessible to multiple actors. The aggregator queries it to identify which DERs are eligible for a given service call (for bid crafting). The DO queries it to understand which DERs are proposed for dispatch and where they connect (for operational validation). The MO references it for eligibility verification. In a real deployment, this registry would be maintained by a neutral authority (e.g., the utility or an independent system operator) and queried via secure application program interfaces (APIs) rather than direct file access.

Because the DER Registry does not currently need to preserve historical changes to DER attributes, a single-table structure is the simplest viable logical model. Table 1 presents the proposed schema. The schema does not record or disclose the technology type of each DER. The value each DER has in the Premise ID column is assumed to be sufficient information for the relevant DO to associate the DER with a position in its distribution grid without compromising privacy.

Table 1. Schema for Only Table in Type-Blind Model of DER Registry

Column	Data Type	Description
DER ID	Integer	An ID designating the DER uniquely within the transmission network of the MO.
Real Power Capacity (kW)	Float	A floating-point value specifying the nameplate real power capacity, in kilowatts, of the DER.
Energy Capacity (kWh)	Float	A floating-point value specifying the nameplate energy capacity, in kilowatt-hours, of the DER for the purposes of charging and discharging.
Apparent Power Capacity (kVA)	Float	A floating-point value specifying the nameplate apparent power capacity, in kilovolt-amperes, of the DER.

Column	Data Type	Description
Ramp Rate (%/min)	Float	A floating-point value specifying the maximum rate, in percent of nameplate real power capacity per minute, at which the DER can change its output.
Number of Phases	Integer	The number of phases (e.g., 1 for single-phase) through which the DER is connected to the distribution grid.
Service Area ID	Integer	The ID of the DO managing the service area (i.e., distribution grid) that contains the premise at which the DER resides.
Premise ID	String	A readable sequence of characters designating, uniquely within the geographic information system of the relevant DO, the position of the DER.
Responsible Entity ID	Integer	The ID of the entity (e.g., aggregator) managing the DER.

If knowledge and consideration of technology type are desired, then an alternative model for the DER Registry is given by the schemas in Table 2 through Table 4. For DERs belonging to the battery energy storage system (BESS) and photovoltaic (PV) technology types, specifications particular to their types are represented by Table 3 and Table 4, respectively. To enable joining in database queries, the DER ID of a BESS or PV asset in these two tables is intended to have the same value as in Table 2.

Table 2. Schema for Main Table in Type-Aware Model of DER Registry

Column	Data Type	Description
DER ID	Integer	An ID designating the DER uniquely within the transmission network of the MO.
Technology	String	A readable sequence of characters indicating the technology type (e.g., natural gas) of the DER.
Real Power Capacity (kW)	Float	A floating-point value specifying the nameplate real power capacity, in kW, of the DER.
Apparent Power Capacity (kVA)	Float	A floating-point value specifying the nameplate apparent power capacity, in kVA, of the DER.
Ramp Rate (%/min)	Float	A floating-point value specifying the maximum rate, in percent of nameplate real power capacity per minute, at which the DER can change its output.
Number of Phases	Integer	The number of phases (e.g., 1 for single-phase) through which the DER is connected to the distribution grid.
Service Area ID	Integer	The ID of the DO managing the service area (i.e., distribution grid) that contains the premise at which the DER resides.
Premise ID	String	A readable sequence of characters designating, uniquely within the geographic information system of the relevant DO, the position of the DER.
Responsible Entity ID	Integer	The ID of the entity (e.g., aggregator) managing the DER.

Table 3. Schema for BESS Table in Type-Aware Model of DER Registry

Column	Data Type	Description
DER ID	Integer	An ID designating the BESS uniquely as a DER within the transmission network of the MO.
Energy Capacity (kWh)	Float	A floating-point value specifying the nameplate energy capacity, in kWh, of the BESS for the purposes of charging and discharging.
Charging Efficiency (%)	Float	A floating-point value specifying the percent efficiency of the BESS when charging.
Discharging Efficiency (%)	Float	A floating-point value specifying the percent efficiency of the BESS when discharging.

Table 4. Schema for PV Table in Type-Aware Model of DER Registry

Column	Data Type	Description
DER ID	Integer	An ID designating the PV uniquely as a DER within the transmission network of the MO.
Maximum Reactive Power Injection (kvar)	Float	A floating-point value specifying the maximum reactive power, in kilovolt-amperes reactive, the PV can inject.
Minimum Reactive Power Injection (kvar)	Float	A floating-point value specifying the maximum reactive power, in kilovolt-amperes reactive, the PV can absorb.

Regardless of whether the type-blind model or the type-aware one is adopted for the DER Registry, it is also anticipated that the attributes that can be retrieved (i.e., queried) from the DER Registry may vary between entities.

2.2 Market Coordination

Figure 2, drawn from the RCA, depicts the compartment of the coordination architecture responsible for market coordination.

Once a bid is crafted, it undergoes evaluation by the Market Validation module to ensure compliance with market requirements. Specifically, this module ensures that the bid respects a list of rules provided by the relevant MO. Additionally, this module identifies market conflicts (where the same DER would be committed to two or more mutually exclusive markets) and potential dual participation (where the same action by a DER would lead to double compensation). Issues that the market validation detects within bids result in notifications that are then stored in the Market Data Warehouse, which is accessible to MOs, DOs, and market participants (e.g., aggregators). Likewise, bids that pass market validation are also stored in the Market Data Warehouse.

The Market Monitoring module affords passive oversight and long-term analysis of the behavior of the rest of Market Coordination compartment of the RCA, with the aim of detecting market manipulation or flaws in the market design. In contrast, the Operational Validation module within the Operational Coordination compartment is managed by the DO to ensure that, in the process of market clearing, service awards granted by the MO to market participants in recognition of their bids do not cause operational issues in the circuit of the DO. Although operational

envelopes are issued to aggregators for incorporation into their bid crafting, the bids once crafted and validated for market considerations still need to be validated again with power-flow simulations that assume forecasted states. Therefore, the MO and the Operational Validation module of the DO orchestrate in an iterative process to mold validated bids from the Market Data Warehouse into service awards that are operationally valid. With this process entrusted to the MO and the DO, the need for multiple rounds of bidding in response to the same service opportunity is obviated.

However, it should be emphasized that the set of designs (i.e., data models and procedures) given in Section 3.3 to realize such an iterative process is merely one possible set of designs to accomplish market clearing. An alternative set of designs may suppose a different iterative process in which, over the course of multiple rounds of bidding, any bids that Operational Validation finds to cause issues are returned to the market participants for reformulation.

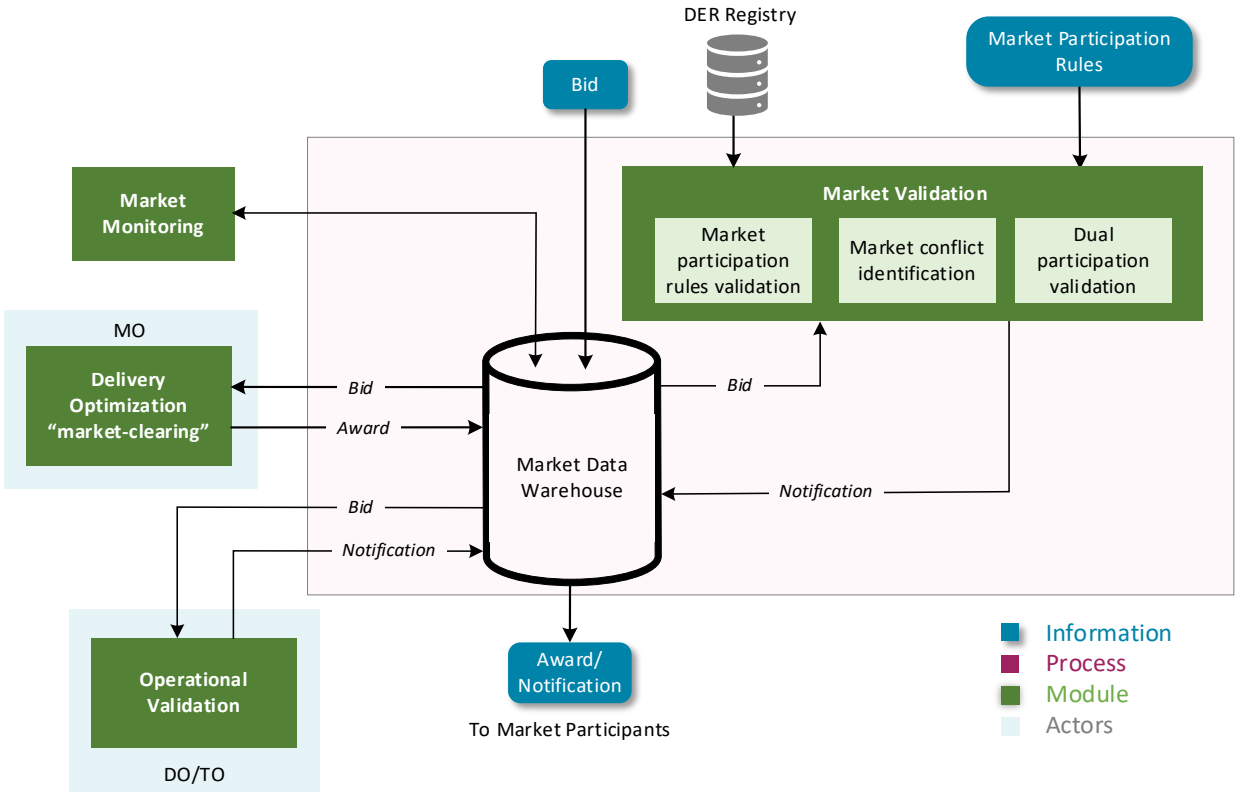


Figure 2. Module, Process, and Information Exchange for Market Coordination

2.2.1 Market Data Warehouse

The Market Data Warehouse is a shared, neutral data store that functions as the communication backbone between market participants, grid operators, and MO. It is organized into life cycle stages that mirror a market transaction as shown in Table 5.

Table 5. Market Data Warehouse Structure and the Role of Each Subfolder in the Bid Life Cycle

Compartments	Purpose
Service	Contains service call signals. When the delivery optimizer detects congestion, it places a service request file here, broadcasting the need for resources to all market participants.
Submissions	Receives bid files (JSON) from participants (e.g., aggregators) responding to a service call. Each bid includes product-value curves, duration, DER IDs, and market metadata.
Market-validated	Holds bids that have passed market validation rules, conflicts, and dual participation rules. These are forwarded to the DO/TO for operational feasibility assessment.
Operation-validated	Holds bids that have additionally passed operational validation by the DO/TO (e.g., no grid constraint violations from dispatching the proposed DERs).
Market-rejected	Bids that failed market validation.
Operation-rejected	Bids that passed market rules but were rejected on operational grounds (e.g., dispatching the DERs would cause distribution feeder overloads).
Award	Final awarded bids, signaling to participants that their resources should be dispatched. The cleared quantity and award are included.

2.3 Operational Coordination

Figure 3, drawn from the RCA without modification, depicts the T&D Coordination portion that is continually executing within the Operational Coordination compartment of the coordination architecture.

Using data from field sensors, supervisory control and data acquisition (SCADA) systems, and Advanced Metering Infrastructure (AMI) meters, the Distribution Grid State Estimation module attempts to recover the state of the distribution circuit to identify abnormal conditions. Then, to achieve enough mutual visibility without intruding on privacy, the TO and DO exchange real-time information on states and power deliverability near the interchange points between their circuits. The information shared may include voltage measurements, power measurements, power-flow limits (present or upcoming), and other forecasted information.

With T&D node visibility thus established, the T&D Optimization module coordinates the optimization of the objectives of the MO, TO, and DO such that, subject to both market-based and operating constraints as well as forecasting, the most locationally relevant and cost-

effective resources are used system-wide. This co-optimization is what produces the service opportunities that are processed by the rest of the coordination architecture.

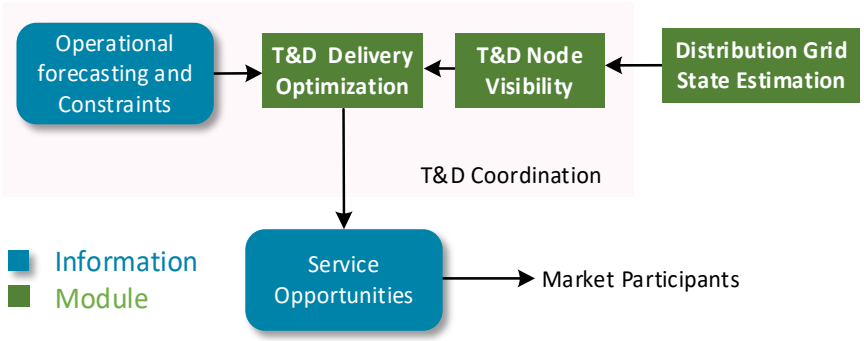


Figure 3. T&D Coordination

Figure 4, drawn from the RCA without modification, depicts the Operational Validation portion that is applied to bids that have passed market validation. With specifications from the DER Registry as well as dynamic states (such as from operational forecasting), Operational Validation executes power-flow simulations to assess whether the DER injection values of bids are likely to cause issues in the proper operation of distribution grids. Such issues may include voltage violations or thermal overloads. If no operational issues are found, then Operational Validation yields approval. If some issues are found, then to aid market participants in reformulating bids, Operational Validation produces either “constraints” on DER injection values or contextual information.

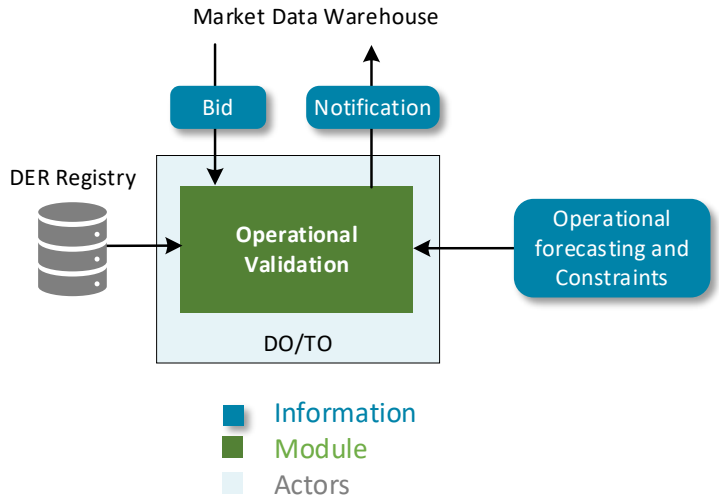


Figure 4. Operational Validation

3.0 Coordination Modules

This section presents the coordination modules associated with the selected validation use cases and describes their roles, data models, and operational procedures. The objective is to illustrate how the RCA enables information exchange and coordination among participating actors. For each module, representative data models and workflows are provided to demonstrate a possible realization of the architecture, with data structures illustrated using SQL schemas, CSV, and Python data classes where appropriate. The modules described in this section are not intended to be exhaustive; rather, they represent the subset of coordination functions required to support the case studies considered in this work and to demonstrate the communication and coordination mechanisms enabled by the RCA.

3.1 DER Registration and Coordination

3.1.1 Bid Crafting

The three main kinds of inputs to the bid-crafting process are market participation rules, DER service opportunities, and T&D constraints, as shown in Figure 1. Service opportunities represent the constantly evolving operational needs in the transmission network; in the case of energy services, an opportunity is defined chiefly by the range of real-power injection/absorption desired system-wide but also other qualifications, such as the permitted locations of participating DERs. Market participation rules specify characteristics required of aggregators or individual DERs for participation in each of the various types of service opportunities. On the basis of power-flow simulations that assume nameplate capacities and forecasted system states (e.g., available capacities for variable generation resources), T&D constraints impose operational limits (e.g., envelopes) that DERs must satisfy during participation in opportunities. The limits only permit certain amounts of injection that DERs can sustain for the duration of the opportunity.

Market participation rules altogether may be conceived as a collection of rulesets. A potential data-class schema for a market participation ruleset is given in Table 6. Additionally, each ruleset describes a service market, where a service opportunity in a market may be subject to one of multiple possible rulesets defined for the market. A potential data-class schema for a service market is given in Table 7.

Table 6. Data-Class Schema for Market Participation Ruleset

Field	Data Type	Description
Ruleset ID	Integer	An ID for uniquely designating the ruleset.
List of Rules	List[String]	A list of rules that bids must respect.
Accepted Locations	Set[String]	The unordered set of locations permitted for DERs participating under the ruleset. Each location can be associated by the relevant DO with a position within its distribution grid. If unspecified, then all locations are acceptable.
Market ID	Integer	The ID of the service market (e.g., energy) that the ruleset is describing.

Table 7. Data-Class Schema for Market Conflict and Dual Participation Identification

Field	Data Type	Description
Conflicting Market IDs	Set[(Integer, Integer)]	A set of pairs of Market IDs that are considered to be conflicting by regulators.
Dual Market IDs	Set[(Integer, Integer)]	A set of pairs of Market IDs that are similar so potential participation in both would lead to double counting of services.

A potential data-class schema for a service opportunity is given in Table 8. In Python, Date-Time and Time-Delta values can be specified with types defined within the built-in date-time module.

Table 8. Data-Class Schema for Service Opportunity

Field	Data Type	Description
Opportunity ID	Integer	An ID for uniquely designating the service opportunity.
Requestor ID	Integer	The ID of the entity requesting the service.
Contract ID	Optional[Integer]	The ID of the contract to which the service opportunity belongs. Unspecified if the opportunity does not belong to any contract.
Start Time	Date-Time	An object specifying the date and time at which the interval of the service begins.
Duration	Time-Delta	An object specifying the interval duration for which the service is sustained.
Qualification	Qualification	An instance of the Qualification class, which describes the nature of the service being performed.
Market ID	Integer	The ID of the market (e.g., energy) to which the service opportunity belongs.
Ruleset ID	Integer	The ID of the market participation ruleset by which bids in response to the service opportunity must abide.

T&D constraints broadly concern the transmission network, a distribution grid, or the interface between the two. Focusing on constraints concerning a distribution grid, the present designs envision them as operating envelopes where, to avert disproportionate impact, each envelope pertains to only the DERs under the control of a certain aggregator within a particular distribution grid. Compliance with an operating envelope ensures that, if no other aggregator in the same distribution grid controls DERs, no operational issues arise (e.g., voltage violations, thermal overloads).

It is supposed that, to each aggregator controlling DERs within a distribution grid, a set of operating envelopes is issued concurrently with the posting of the service opportunity. From the set issued to it, each aggregator selects a single envelope at a time to constrain its DER scheduling for participation in the opportunity.

A potential data-class schema for an operating envelope is given in Table 9.

Table 9. Data-Class Schema for Operating Envelope

Field	Data Type	Description
Envelope ID	Integer	An ID for uniquely designating the operating envelope.
Provider ID	Integer	The ID of the entity (e.g., aggregator) providing the service.
Service Area ID	Integer	The ID of the DO managing the service area (i.e., distribution grid) in which the envelope limits the entity's DER scheduling.
DER Injection Upper Bounds [kW]	Dict[Integer,Float]	A dictionary specifying as a floating-point value the maximum real-power injection, in kW, that the envelope permits from each individual DER (specified by ID) controlled by the entity in the service area.
DER Injection Lower Bounds [kW]	Dict[Integer,Float]	A dictionary specifying as a floating-point value the minimum real-power injection, in kW, that the envelope permits from each individual DER (specified by ID) controlled by the entity in the service area. (A negative value can be used to permit real-power absorption.)

The data-class schema proposed earlier for service opportunities relies on the notions of a contract and a qualification. Not previously defined, a contract is understood to refer to an agreement, between a service requestor and a service provider, that the provider is to furnish a particular service on a particular recurring schedule. Potential data-class schemas for a contract and a qualification are given in Table 10 and Table 11, respectively.

Table 10. Data-Class Schema for Contract

Field	Data Type	Description
Contract ID	Integer	An ID for uniquely designating the contract.
Requestor ID	Integer	The ID of the entity requesting the service.
Provider ID	Integer	The ID of the entity (e.g., aggregator) providing the service.
Service Terms	Object	An object specifying the service terms. A placeholder to be expanded in future designs.
Service Guarantees	Object	An object specifying the service guarantees. A placeholder to be expanded in future designs. At a minimum, expected to include a penalty for failure to provide service.
Qualification	Qualification	An instance of the Qualification class, which describes the nature of the service being performed.
Recurring Schedule	Object	An object specifying the recurring schedule on which the service is to be provided. A placeholder to be expanded in future designs.

Table 11. Data-Class Schema for Qualification

Field	Data Type	Description
System-Wide Injection Range [kW]	Tuple[Float,Float]	A pair of floating-point values specifying the lower and upper limits of the range of real-power injection, in kW,

Field	Data Type	Description
		desired in total across bids in response to a service opportunity. (A negative injection value indicates absorption.)
Minimum Ramp Rate [%/min]	Float	A floating-point value specifying the minimum ramping capability, in percent of nameplate real-power capacity per minute, required of DERs participating in a service opportunity.

3.1.2 Procedure and Output of Bid Crafting

The output of the bid-crafting process for a service opportunity is a bid from each provider (e.g., aggregator) that is eligible to participate in the opportunity. As stated before, each eligible provider must further ensure that its bid neither involves any DER that is not eligible to participate nor involves DER scheduling that no operating envelope supports in the set of envelopes issued to the provider.

A potential data-class schema for a service bid is given in Table 12. The schema relies on the notion of a product v. value curve, which is defined as a collection of discrete points whose numerical pairs relate discrete real-power injection levels to offer costs. Each aggregator crafting a bid determines, for each discrete injection level in the distribution grid of each DO, how to furnish the injection in the grid at lowest cost to the service requestor while abiding by at least one operating envelope issued to the aggregator. The injection level and the lowest cost then forms the numerical pair of a point in a product v. value curve. Nevertheless, the operating envelope allowing the lowest cost may differ between injection levels. Therefore, although the bid crafted by an aggregator contains a single product v. value curve for each distribution grid, each point in the curve for a given grid may be associated with a different envelope.

A potential data-class schema for a product v. value curve is given in Table 13. The schema further relies on the notion of a product v. value point, which is characterized not only by a numerical pair of injection level and offer cost but also by an envelope ID and the IDs of the DERs that the aggregator intends to command to accomplish the injection. A potential data-class schema for a product v. value point is given in Table 14. The schemas in Table 13 and Table 14 do not prescribe fields for curve and point IDs, respectively, because it is not anticipated that curves and points will need to be uniquely designated.

Table 12. Data-Class Schema for Service Bid

Field	Data Type	Description
Bid ID	Integer	An ID for uniquely designating the service bid.
Requestor ID	Integer	The ID of the entity requesting the service. Copied from the service opportunity.
Provider ID	Integer	The ID of the entity (e.g., aggregator) providing the service.
Opportunity ID	Integer	The ID of the service opportunity to which the bid is crafted in response. Copied from the service opportunity.
Contract ID	Optional[Integer]	The ID of the contract to which the service opportunity belongs. Copied from the service opportunity.

Field	Data Type	Description
Market ID	Integer	The ID of the market (e.g., energy) to which the service opportunity belongs. Copied from the service opportunity.
Ruleset ID	Integer	The ID of the market participation ruleset by which bids in response to the service opportunity must abide. Copied from the service opportunity.
Product v. Value Curves	Dict[Integer] Set[ProdValueCurve]	A dictionary mapping the ID of each DO to the unordered set of product v. value curves that the entity (e.g., aggregator) determines in the distribution grid of the DO.
Involved DER IDs	Set[Integer]	The union of the sets of IDs of DERs that each curve point involves, across all curves. Thus, the unordered set of IDs of DERs that can at all possibly be involved in the award (if any) resulting from the MO's eventual adjustment of the bid to become operationally valid.
Start Time	Date-Time	An object specifying the date and time at which the interval of the service begins.
Duration	Time-Delta	An object specifying the interval duration for which the service is sustained.
Service Terms	Object	An object specifying the service terms. A placeholder to be expanded in future designs.
Service Guarantees	Object	An object specifying the service guarantees. A placeholder to be expanded in future designs. At a minimum, expected to include a penalty for failure to provide service.

Table 13. Data-Class Schema for Product v. Value Curve (ProdValueCurve)

Field	Data Type	Description
Curve Points	List[ProdValuePoint]	The list of product v. value points comprising the curve.

Table 14. Data-Class Schema for Product v. Value Point (ProdValuePoint)

Field	Data Type	Description
Injection Level [kW]	Float	A floating-point value specifying the injection level, in kW, of the curve point.
Offer Cost [$]	Float	A floating-point value specifying the offer cost, in USD, for which the injection can be furnished.
Involved DER IDs	Set[Integer]	The unordered set of IDs of DERs that the aggregator intends to command to accomplish the injection.
Envelope ID	Integer	The ID of the operating envelope assumed in DER scheduling to furnish the injection level.

A desired uniform granularity (e.g., 25 kW) of discrete real-power injection levels is assumed. In response to a service opportunity and after receiving a different set of operating envelopes for the distribution grid of each DO, each aggregator considers the characteristics of its DERs to craft a bid by the following procedure, illustrated in Figure 5.

1. If the aggregator at large is not eligible to participate in the service opportunity, then simply do not craft a bid for the aggregator. The rest of the steps in the procedure are not followed. Reasons for aggregator ineligibility may include insufficient entity capacity.
2. Within the distribution grid of each DO, identify the DERs that are controlled by the aggregator and that are eligible to participate in the service opportunity. Reasons for DER ineligibility may include insufficient ramping capability or a location not among those accepted.
3. If the aggregator does not control any eligible DERs in any distribution grid, then again do not craft a bid for the aggregator. The rest of the steps in the procedure are not followed.
4. For each distribution grid, consider the aggregator's eligible DERs in the grid and a range of discrete injection levels starting from 0 kW and incrementing by the uniform granularity. Then, perform the following steps:
    a. For each injection level, determine which operating envelope within the grid permits the injection to be furnished at the lowest cost by the aggregator's eligible DERs in the grid. Take note also of which eligible DERs would be involved to achieve the lowest cost subject to the envelope. (If a DER stores energy, then the aggregator may need to weigh short-term benefits from its usage against mid-term benefits of maintaining its availability for future intervals within some time horizon.)
    b. Construct a product v. value curve spanning the range of injection levels. The curve is a collection of discrete points, each characterized by not only a numerical pair of injection level and offer cost but also by an envelope ID and a set of IDs of involved DERs.
5. For each distribution grid, the aggregator should now have a single product v. value curve, each of whose points may be associated with a different operating envelope within the grid. Include these curves in the bid, along with any auxiliary information (metadata) required for the bid as a whole. In particular, take the union of the sets of IDs of DERs that the individual curve points involve, across all curves, and include this union as a field value in the bid. Publish the bid in the Market Clearinghouse to subject the bid to market validation (followed by operational validation, and last, awarding, which both occur outside the Market Clearinghouse).

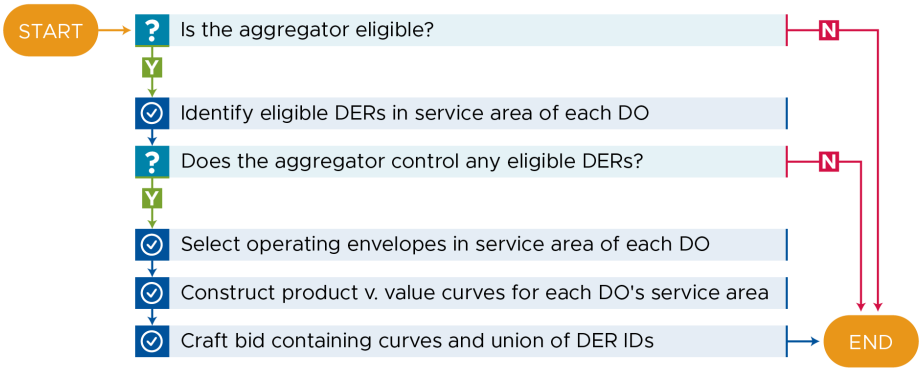


Figure 5. Procedure for Bid Crafting by an Aggregator

Each eligible aggregator submits a bid containing at least one separate product v. value curve for every distribution grid in which it controls eligible DERs. Although such granularity of curves may be viewed as an excessive level of disclosure, it is seen in Section 3.3 to greatly aid the

MO in its impartial role of facilitating cooperation between aggregators and DOs. If DER participation in wholesale markets through aggregators were permitted to cause frequent operational issues within distribution grids, then over time, DOs would likely grow less inclined to accommodate DERs even if they lacked the authority to ban DERs outright.

3.2 Market Coordination

Market coordination includes Market Validation and Market Monitoring modules as shown in Figure 2.

3.2.1 Market Validation

The Market Validation module receives all bids once they are formed and performs a list of checks to ensure compliance with market regulations.

The first step is to validate the fact that the bid conforms with market participation rules. Specifically, the Market Validation module uses the Market ID field to identify the relevant rules for the bid. In this phase, we assume that the Market Validation module implements one simple rule: the sum of capacities of the DERs that are part of the bid (given by the Involved DER IDs field) can be no less than the offered quantity of power (given by the Product v. Value Curves). In future phases of work, the number of rules will increase.

The next step is to identify potential market conflicts or dual participation. Bids are used to create pairs of DER IDs with Market IDs (using the Market ID and Involved DER IDs fields) that can be checked against the pairs of both conflicting and dual market IDs (from the Conflicting Market IDs and Dual Market IDs of the appropriate schema, respectively) to identify whether any DER might be committed to either provide services that create conflict or provide services that are similar and create the risk of double counting of services, if different bids in which it participates are accepted.

This is implemented in the following algorithm:

For every *bid*:

- Retrieve *Market ID*.
  - Retrieve rules corresponding to *Market ID*. If a rule is violated, reject and terminate.
- Retrieve DER IDs from Involved DER IDs.
  - Form all (DER ID, Market ID) pairs.
  - Retrieve Conflicting Market IDs. If conflicting Market IDs share the same DER ID, reject and terminate.
- Retrieve DER IDs from Involved DER IDs.
  - Form all (DER ID, Market ID) pairs.
  - Retrieve Dual Market IDs. If dual Market IDs share the same DER ID, reject and terminate.
- Accept *bid*.

3.2.2 Market Monitoring

The Market Monitoring module leverages data accessible through the i) Market Data Warehouse, such as cleared bids, dispatch notifications, market validation logs, ii) DER Registry, and iii) meter-based measurements to evaluate how well market participants adhered to their awarded schedules. In the modeling activity, this functionality is implemented by systematically comparing the cleared dispatch instructions against the actual delivered output for each participant and time interval. The model also incorporates registration and contractual information to verify whether participants provided the committed services for the required duration. By processing these datasets together, the module identifies deviations, quantifies performance, and flags potential compliance or operational concerns.

3.3 Operational Coordination

Operational Coordination encompasses the processes required to validate bids from an operational standpoint and to optimize the grid for efficient and reliable service delivery. For instance, the DO must verify aggregator bids before DERs are dispatched, ensuring that proposed actions are feasible within distribution system constraints. Similarly, the TO and/or MO conducts system-level optimization to ensure that the bulk power system operates economically and efficiently. The following sections describe how these Operational Coordination modules are modeled to support validation of the coordination framework.

3.3.1 Operational Validation

Once all bids submitted in response to a service opportunity have completed market validation, the remaining steps are to assess their operational feasibility and ultimately clear the market to generate service awards. The sequencing of these activities is a design choice. For example, operational validation may be performed prior to the market clearing, after the market clearing, or through an iterative process in which market clearing and operational validation are repeatedly coordinated. This report describes one such procedure in detail, where the MO coordinates iteratively with multiple DOs to determine feasible service awards.

The MO retrieves from the Market Data Warehouse the bids that passed market validation. With these validated bids, the MO then orchestrates with the relevant DOs in the following iterative procedure, illustrated in Figure 6, to mold bids into operationally valid service awards. The iterative procedure is intended to avoid requiring multiple rounds of bidding by aggregators and, for comprehension, may be viewed as a black box taking the place customarily occupied by market clearing.

1. From the product v. value curves of the bids, the MO selects points such that, at lowest total cost, the sum of their injection levels falls within the system-wide range of the service opportunity (but preferably is at or near the upper range limit). The total cost accounts for the need to compensate each aggregator at the same clearing price in $/megawatt-hour (MWh), such that the clearing price is given by the ratio of the total cost to the sum of the injection levels. Within each aggregator's bid, at most one point may be selected from the aggregator's curve for each distribution grid. However, a point may be selected from each multiple of the aggregator's set of curves, such that altogether multiple product v. value curve points are selected from within the aggregator's bid.
2. Given the results of step 1, to each DO the MO issues the injection levels of the selected curve points, along with their operating envelope IDs—with which DOs can retrieve the full envelopes on their own—and the IDs of the involved DERs. (These injection levels

correspond to the “Cleared Bid” information sent by the MO into Operational Coordination in Figure 2.) Internally, the MO also tracks the offer costs of the selected points.

3. The injection level of a curve point represents a total injection by an aggregator across its DERs within the distribution grid of a DO. Therefore, each DO receiving any injection levels and auxiliary information in step 1 performs the following steps:
    a. Because it is possible for the DER-wise limits of an operating envelope not to be saturated, it is still not generally possible to deduce from the injection level of a point the exact allocation across the involved DERs of the aggregator within the distribution circuit. Consequently, given the selected envelopes as well as nameplate capacities and forecasted system states, the DO solves an optimal power flow (OPF) that finds the worst-case allocations of the injection levels of aggregators across their involved DERs such that the allocations together cause the most severe operational issues, particularly in terms of voltage violations or thermal overloads. (Optionally, the schema in Table 14 may be modified to contain not merely the IDs of involved DERs but the allocation of the injection level across them. Then, with the forecasted states and the injections of individual DERs, the DO simply executes a conventional power-flow simulation such as with OpenDSS.)
    b. If any operational issues are detected in the OPF, then the DO downscales each aggregator’s selected injection level in the distribution grid by a consistent factor (e.g., 90 percent), again to avert disproportionate impact. For example, a downscaling factor of 90 percent applied in two iterations of this step would translate to a cumulative downscaling factor of 81 percent from the original injection levels of the points. Also, the DO downscales the DER-wise limits of each aggregator’s operating envelope by the same factor; the downscaled envelope will again be enforced on the DER operations of the aggregator within the grid of the DO. Then, with the downscaled injection levels and envelopes, the DO returns to step 3.a for another iteration.
4. Each DO sends back to the MO the downscaling factor—which may be unity (i.e., 1) in the case of no downscaling—needed to regain valid operation in its distribution grid. (This downscaling factor corresponds to the “Constraint/Approval” information returned by Operational Coordination to the MO in Figure 2.) The next action taken by the MO depends on whether a non-unity downscaling factor is received from any DO.
    a. If there is no DO for which downscaling is detected, then the MO proceeds to step 5.
    b. If downscaling is detected from at least one DO, then with each non-unity factor, the MO internally downscales each aggregator’s selected curve-point numerical pair—both injection level and offer cost—in the circuit of the DO, and the procedure returns to step 1 for another iteration. However, in the new and every subsequent iteration, all selected curve points downscaled in the present iteration remain selected but with their downscaled injection levels and offer costs. Additionally, in the new and every subsequent iteration, each distribution grid with any downscaled curve points in the present iteration cannot select curve points offered by aggregators other than those offering the downscaled curve points, and different curve points within the grid instead of the downscaled curve points cannot be selected from the aggregators offering the downscaled curve points. The MO considers the downscaled points to be frozen as the selected points in the circuit of the DO, and the MO does not engage in any further orchestration—steps 2 through 4—with the DO in the new and every subsequent iteration. It is also observed that, after the downscaling in the present iteration, the MO needs to ensure the curve points it selects in the new iteration bring

the sum of injection levels back up to within the system-wide range of the service opportunity.

5. The MO translates the now-selected curve points into service awards that are then stored in the Market Data Warehouse, and the relevant market participants are notified. A potential data-class schema for a service award is given in Table 15.

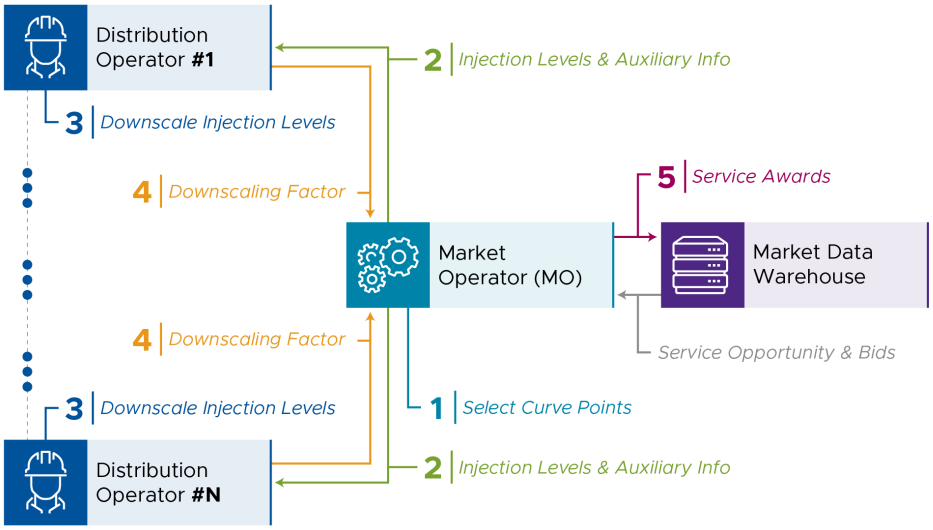


Figure 6. Iterative Procedure Between MO and DOs to Achieve Market Clearing

Table 15. Data-Class Schema for Service Award

Field	Data Type	Description
Award ID	Integer	An ID for uniquely designating the service award.
Bid ID	Integer	The ID of the service bid on which the award is based. Copied from the service bid.
Requestor ID	Integer	The ID of the entity requesting the service. Copied from the service bid.
Provider ID	Integer	The ID of the entity (e.g., aggregator) providing the service. Copied from the service bid.
Opportunity ID	Integer	The ID of the service opportunity that is the basis of the bid (and now award). Copied from the service bid.
Contract ID	Optional[Integer]	The ID of the contract to which the service opportunity belongs. Copied from the service bid.
Market ID	Integer	The ID of the market (e.g., energy) to which the service opportunity belongs. Copied from the service bid.
Injection Level [kW]	Float	A floating-point value specifying the injection level, in kW, of the award.
Offer Cost [$]	Float	A floating-point value specifying the offer cost, in USD, that the aggregator would consider sufficient compensation for the injection. Should reflect clearing price.
Clearing Price [$/MWh]	Float	A floating-point value specifying the clearing price, in USD per MWh, ultimately granted to all aggregators receiving service awards for the opportunity.

Field	Data Type	Description
Involved DER IDs	Set[Integer]	The unordered set of IDs of DERs that the aggregator intends to command to fulfill the award.
Start Time	Date-Time	An object specifying the date and time at which the interval of the service begins. Copied from the service bid.
Duration	Time-Delta	An object specifying the interval duration for which the service is sustained. Copied from the service bid.
Service Terms	Object	An object specifying the service terms. A placeholder to be expanded in future designs. Copied from the service bid.
Service Guarantees	Object	An object specifying the service guarantees. A placeholder to be expanded in future designs. At a minimum, expected to include a penalty for failure to provide service. Copied from the service bid.
Downscaling Factors	Dict[Integer, Float]	A dictionary mapping the ID of each DO to the downscaling factor, as a floating-point value, determined to be needed for regaining valid operation within the circuit of the DO.

3.3.2 T&D Node Visibility

The T&D Node Visibility module facilitates information exchange between the DO and TO regarding the electrical conditions and operational constraints at the T&D interface nodes. The module may be executed by either the DO or the TO, depending on the coordination arrangement. To support interoperability, the exchanged information should follow a standardized data model that defines the attributes associated with interface nodes, such as power injections, voltage measurements, forecasts, operating limits, and flexibility availability. By providing a consistent representation of interface-node information independent of the underlying systems used by the DO and TO, the module enables seamless information sharing and serves as a foundational input to other coordination modules, including Delivery Optimization and State Estimation.

3.3.3 Delivery Optimization

The Delivery Optimization module operates on the DO, TO, and MO sides (both bulk and distribution), allowing each side to optimize power delivery subject to constraints and available resources. For example, a DO may run optimization routines focused on voltage management or feeder loading, whereas a TO may optimize for broader bulk-system objectives such as contingency analyses, voltage stability assessments, and interchange scheduling. To model these optimization functions, each network will be represented using appropriate simulation tools. Distribution networks can be modeled using tools such as OpenDSS or GridLAB-D, while transmission networks will be modeled using PSSE, MATPOWER, or PowerWorld. These models provide the electrical topology, equipment parameters, and system constraints that can be used to run the load flow and feed relevant information into the optimization engine.

Additionally, the MO will also execute the Delivery Optimization, which is referred to as “market clearing.” While market clearing can be a separate module, we use the same name in our modeling exercise as per Figure 2. For instance, given DO-provided load forecasts, SCADA measurements, generator bids, aggregator bids, and a transmission network model, a security-constrained economic dispatch formulation can be constructed by the MO. In our modeling

activity, DO's load forecast is provided through a CSV file with proper tags for feeder or substation identifiers, time-stamped demand values, and associated metadata, while the aggregator bid follows the structured schema defined in Table 12. SCADA measurements from field devices will be linked to their corresponding equipment models in the simulator using unique device identifiers. Infrastructure limits—such as line thermal ratings, ramp rates, voltage limits, reserve requirements, and N-1 security conditions—will be encoded as constraints in the mathematical model. Depending on operator needs, optimization can be implemented using rules-based logic, heuristic algorithms, or mathematical optimization using off-the-shelf solvers.

3.3.4 Distribution System State Estimation

The Distribution System State Estimation module is intended to fuse information received from multiple actors to develop an enhanced understanding of the current operating condition of the distribution system. Inputs to the module may include measurements, forecasts, outage information, operational constraints, interface-node information exchanged with the TO, and resource-related information provided by aggregators, such as availability, schedules, telemetry, and historical service-delivery performance. By combining these diverse information sources, the module can estimate not only electrical states such as voltages, power flows, and equipment loading, but also operational conditions that influence DER participation and service delivery. In addition to estimating system states, the module can identify conditions that may affect the deliverability of DER services, including planned outages, feeder reconfigurations, curtailments, equipment limitations, communication failures, and other operational events. The resulting outputs provide situational awareness regarding both network conditions and resource availability.

This module can be modeled as a network-aware situational assessment problem. In addition to a distribution network model and measurement model, the module requires an event model that represents planned outages, switching actions, feeder reconfigurations, DER curtailments, communication failures, and other conditions that may affect DER deliverability. The module can then estimate system states while also identifying whether DERs remain electrically deliverable under current or expected network conditions.

The outputs of the module are intended to serve as inputs to other coordination modules, including Operational Validation, Delivery Optimization, and T&D Node Visibility. To facilitate interoperability among actors and systems, the estimated states and inferred operational conditions can be communicated using standardized data models and information exchange formats.

4.0 RCA Modeling and Validation

The modeling and validation exercise is designed to reflect the organizational separation that exists among actors in a real-world grid ecosystem. Rather than assuming that all system information is centrally available, the testbed represents each participant as maintaining its own local data environment and operating only on the information it would realistically possess. In this modeling approach, each actor, such as TO, DO, aggregator, and MO, has responsibility for its own models, forecasts, and operational datasets. For example, a TO maintains transmission network models, and a DO maintains feeder models and customer or AMI-related data. These data assets are treated as locally owned and potentially proprietary. They are not assumed to be visible to other actors unless they are intentionally shared through defined coordination or market processes.

4.1 Testbed

The testbed used in the modeling and validation process distinguishes between an actor's data environment and the computational processes that operate on that data. Actor-specific functions, such as bid formulation, delivery optimization, feasibility assessment, or market validation, are implemented as separate executable modules. These modules use the relevant actor's available information as input and produce outputs that may either remain local or be shared with other actors depending on the modeled process.

This testbed supports several important principles:

- **Data sovereignty:** Each actor retains control over its own models and datasets, reflecting the practical reality that grid participants do not have universal access to one another's internal information.
- **Explicit information exchange:** Communication between actors occurs only through defined interfaces or shared artifacts. For example, if an MO requires a load forecast from another participant, that forecast must be intentionally provided as part of the modeled workflow. This makes information exchange transparent, traceable, and auditable.
- **Separation of data and logic:** The testbed separates the information each actor owns from the algorithms each actor runs. This allows analytical methods, optimization routines, and validation processes to be updated independently of the underlying datasets.
- **Modularity and extensibility:** Because actor roles, datasets, and computational processes are represented independently, individual modules can be replaced, improved, or tested without restructuring the broader coordination framework.

Overall, the testbed allows modeling of the grid ecosystem as a set of interacting but institutionally distinct actors. This design allows the validation process to better capture real-world constraints around data ownership, limited visibility, and structured coordination among market and grid participants.

We developed several modules representing different actors that coordinate through file-based communication using a centralized Market Data Warehouse and a DER Registry. Each module operates independently and continuously monitors designated locations for the data and models required from other actors. When executed, the modules enter a listening state and remain idle until all prerequisite information becomes available. Once the required inputs are detected, the corresponding module proceeds with its computation and writes its outputs to predefined

locations, making the results available for downstream actors or other coordination modules. This asynchronous execution emulates the information exchange and coordination process among multiple entities involved in managing transmission congestion caused by large data center loads. Certain coordination modules may exist at multiple operational levels. For example, a market validation module may be implemented within both wholesale and distribution market contexts. To clearly distinguish between these instances throughout this report, wholesale-level modules are denoted using the prefix W- (e.g., `W-Market Validation`), while distribution-level modules are denoted using the prefix D- (e.g., `D-Market Validation`). The same naming convention is applied to shared data repositories and databases. For example, wholesale and distribution market data warehouses are referred to as `W-Market Data Warehouse` and `D-Market Data Warehouse`, respectively. Similarly, to have each use case be as self-sufficient as possible, the descriptions of data to be stored or modules to be executed may be repeated.

4.2 Case Studies

Three case studies are applied across the testbed to support modeling and validation of the RCA: LDDL, Extreme Cold Weather Event, and Cascading Blackout. For each case study, the following sections describe the associated data, models, workflows, and modeling details. The case studies are organized according to the *Walk-Jog-Run* progression to represent increasing levels of coordination maturity, operational complexity, and validation scope.

4.2.1 Large Dynamic Digital Load

The objective of this case study is to demonstrate how the RCA enables opportunities to monetize customer-sited resources on the sub-T&D systems while mitigating the operational and reliability challenges posed by the operation of LDDLs. The relevant models and data for this case study are summarized below:

- The TO owns multiple transmission test systems, and for the demonstration purposes, we used IEEE 39-bus in PSS/E and PowerWorld formats.
- The DO maintains the IEEE 123-Bus distribution model in OpenDSS format (lines, loads, PVs, batteries, regulators, switches) and load forecast data.
- The MO or any third-party entity manages the Market Data Warehouse with structured queues (submissions, market-validated, operation-validated, rejected, award, service) containing JSON bid files that flow through validation stages.
- The LDDL actor (representing a data center or industrial facility) maintains 24-hour demand profiles for transmission studies. If this is not possible, the TO can forecast based on its historical performance.
- The aggregator performs DER registration and updates aggregator-DER relationships in the DER Registry.
- The DER Registry maintains a centralized database with information on DER ID, Real Power Capacity, DER Type, Premise ID, and Responsible Entity ID, enabling both market validation queries and distribution model lookups.

4.2.1.1 Walk

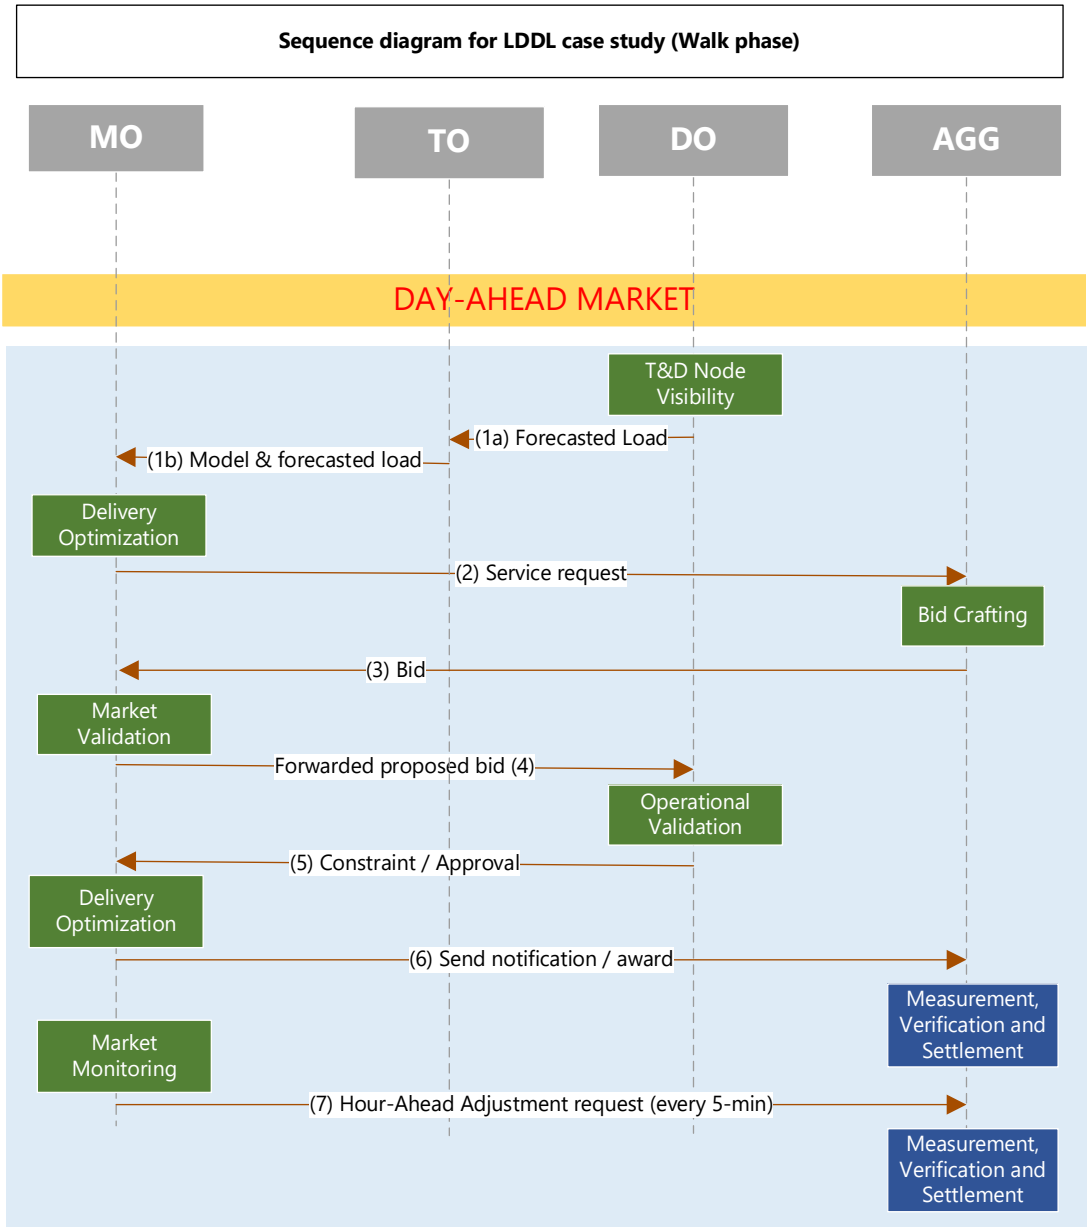


Figure 7. Swimlane Diagram Depicting Workflow for Walk Phase of LDDL Case Study

The LDDL case study, in the Walk phase, consists of several modules representing different actors that coordinate through file-based communication using a centralized Market Data Warehouse. Each module operates independently and continuously monitors data and/or models provided by other actors. When executed, the modules enter a listening state and remain idle until all prerequisite information becomes available. Once the required inputs are detected, the corresponding module proceeds with its computation and writes its outputs to predefined locations, making the results available for downstream actors. This asynchronous execution emulates the information exchange and coordination process among multiple entities involved in managing transmission congestion caused by large data center loads.

A. The DO executes the `T&D Node Visibility` function, yielding information on DO system status relevant to demand such as near real-time peak (1a). The same function may also be executed on the transmission side to share the TO status to the DO such as system states (e.g., voltage at the interface).

B. The `W-Delivery Optimization` module acts as the bulk grid's delivery optimization engine within a multi-actor market environment. On startup, the module polls for forecasted load from TO and DO as well as a network model from the TO (1a and 1b). Once all inputs are received, it executes a load flow study (or optimization) using the network model and forecasted loads to identify potential line congestion (overloads). If any lines are forecasted to exceed their thermal ratings, it flags congestion and issues a "call for service" by writing a service request file to the Market Data Warehouse (2), signaling market participants to submit bids. Note that this module is not exclusively executed by the MO as shown in Figure 7. Depending on the operational context, it may also be invoked and operated by the TO, such as when the TO is executing the `T-Delivery Optimization` process and identifies a need for additional services. After calling for service, the `W-Delivery Optimization` waits (polls) for validated or rejected bid files to appear from the DO's operational validation process. Based on the validation result, it either awards the bid (writes to the award folder, notifying participants to dispatch resources) or rejects it with constraints.

C. The `Bid Crafting` module is executed by an aggregator that assembles bids from its portfolio to participate in markets (3). On startup, the aggregator reads the DER Registry to identify the resources it manages that are applicable to the given service opportunity. For example, the service might be needed in a specific jurisdiction, and it is the responsibility of an aggregator to identify the relevant resources using information from the DER Registry. The module polls the Market Data Warehouse's service folder, waiting for a "call for service" to appear. Once the service call is detected, the module aggregates the capacity of its eligible DERs, computes product-value curve, and packages a bid with other relevant information. Finally, the drafted bid is written as a JSON file to the Market Data Warehouse submission queue, making it available for downstream market validation.

D. The `W-Market Validation` sits downstream of the `Bid Crafting` module. After an aggregator submits a bid JSON to the Market Data Warehouse's submissions queue, this module picks it up and runs a series of validation checks before either accepting or rejecting it (3). On startup, it loads the DER Registry and configures the expected bid schema and market rules (including which markets conflict and which carry dual-compensation risk). The module polls the Market Data Warehouse's submissions compartment, waiting for bid JSON files to appear. Once a bid is detected, the module runs a sequential validation pipeline that checks for compliance with market participation rules, dual compensation, and market conflicts. If any check fails, the bid is moved to the market-rejected queue. If all checks pass, the module records the DER-to-market commitments (to inform future conflict checks) and moves the validated bid to the market-validated queue, making it available for downstream validation and clearing (4).

E. The `D-Operational Validation` module is executed by the DO that determines whether a market-validated bid is physically feasible on the distribution grid (5). At the Walk maturity level, the DO has limited visibility into grid conditions and relies on

simplified checks to screen for obvious operational concerns. On startup, it loads the DER Registry and a distribution feeder model. The module polls the Market Data Warehouse's market-validated queue, waiting for bids that have already passed market validation. Once a bid is detected, the module runs a basic power-flow-based screening:

a. DER dispatch modeling: each DER in the bid is dispatched at full-rated capacity to approximate a worst-case scenario.

b. Voltage feasibility under peak and low load: power flow is solved under two representative loading conditions, checking whether bus voltages remain within acceptable bounds.

c. This is not a comprehensive operational study but rather a coarse check meant to catch bids that would clearly violate grid constraints given the DO's limited observability. If either scenario produces a violation, the bid is written to the Market Data Warehouse's operation-rejected queue. If both pass, the bid is written to the operation-validated queue, signaling to the MO that the bid is operationally acceptable for clearing.

F. The `W-Delivery Optimization` module performs the market clearing process and communicates the outcome to participating aggregators (6). For bids that successfully pass all market and operational validation checks, the module issues an award notification, enabling aggregators to dispatch their resources in accordance with the awarded service. For bids that are not selected or fail validation, the module provides notification along with contextual feedback to help aggregators understand the reasons for the outcome and improve future bidding and operational strategies.

G. The `Market Monitoring` module is currently a placeholder. Its intended purpose is to perform retrospective analysis of market operations, examining historical data for discrepancies in payments, dispatch instructions, or other aspects of market activity. Once implemented, it will serve as an oversight function to identify irregularities and ensure the integrity of market outcomes over time.

4.2.1.2 Jog

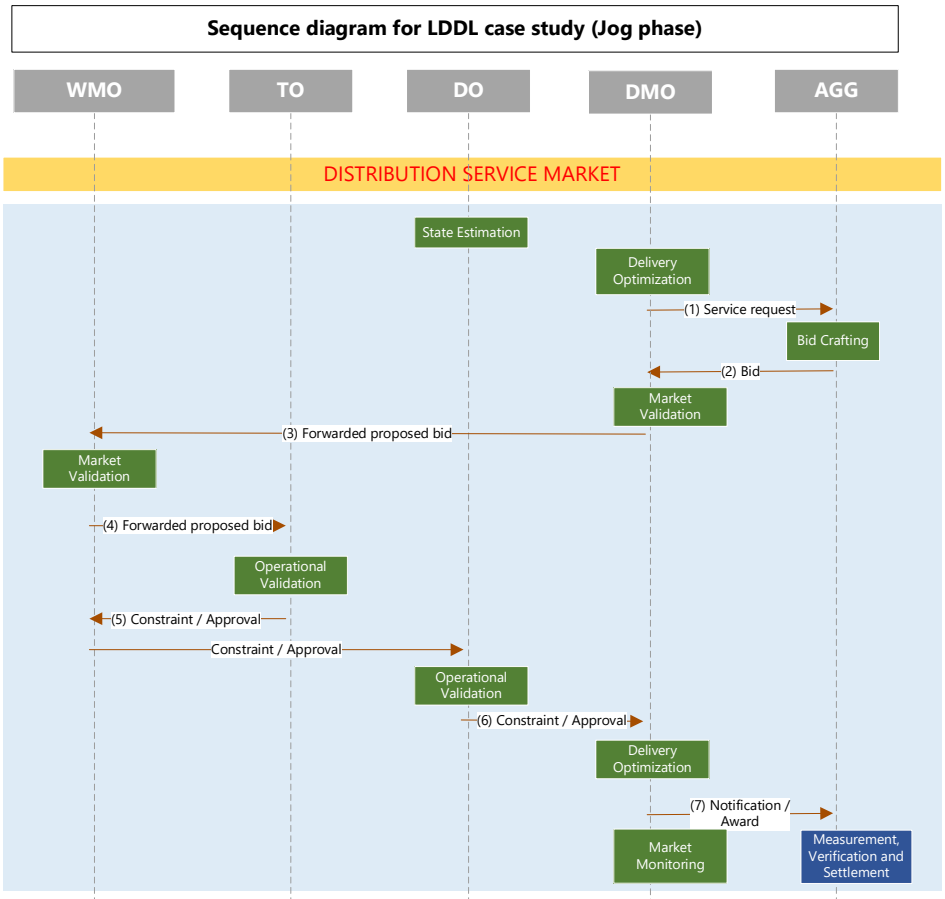


Figure 8. Swimlane Diagram Depicting Workflow for Jog Phase of LDDL Case Study; Only Distribution Part is Shown as Bulk Grid Sequence Remains the Same

The key distinction in the Jog phase is the introduction of a distribution-level market, which necessitates the creation of a dedicated Distribution Market Data Warehouse (D-Market Data Warehouse) separate from the Wholesale Market Data Warehouse (W-Market Data Warehouse). This evolution also introduces a DMO responsible for managing distribution-level market activities, along with a delivery optimization module that ensures market outcomes remain feasible and reliable within distribution system operational constraints. The module advancements in the Jog phase are explained in detail below:

A. The `D-Delivery Optimization` module is a new capability introduced in the Jog phase, enabling a distribution-level market. It acts as the distribution grid's engine, operating independently of—but coordinated with—the `W/T-Delivery Optimization`. On startup, the module polls for a distribution network model and AMI measurement data from the DO. Once received, the module runs a distribution operations simulation (such as OpenDSS-based power-flow analysis) to identify local congestion or constraint violations on the distribution feeder. If congestion is detected, it issues a "call for service" by writing a service request file to the Distribution Market Data Warehouse (D-Market Data Warehouse), opening a local market opportunity for DER participation. After calling for service, it waits (polls) for market-validated or rejected bid files in the D-Market Data Warehouse (1). Based on the validation result, it either awards

the bid (writes to the D-Market Data Warehouse award queue) or rejects it. If awarded, it adjusts the distribution load flow accordingly and re-invokes the optimization loop to confirm the local constraint is resolved. This dual-market structure allows DERs to participate in both wholesale and distribution markets, with coordination handled through cross-market validation checks.

B. The `Bid Crafting` module is executed by an aggregator that assembles bids from its portfolio to participate in markets. In the Jog phase, the Bid Crafting module from the Walk phase is extended to support both the wholesale and distribution market service opportunities. On startup, the aggregator reads the DER Registry to identify the resources it manages, filtering by its eligibility based on service call. It then polls the appropriate Market Data Warehouse's service queue (W-Market or D-Market), waiting for a "call for service" to appear (1). Once the service call is detected, the module selects eligible DERs from its portfolio until the required capacity is met, computes a product-value curve (megawatt [MW] capacity and total bid price), and packages a bid with relevant metadata (duration, DER IDs, market ID, ruleset ID, provider ID, timestamps). The drafted bid is written as a JSON file to the appropriate Market Data Warehouse submission queue, making it available for downstream market validation (2). This dual-market extension allows the same aggregator to respond to service calls from either the wholesale market operator (WMO) or DMO.

C. The `D-Market Validation` module is a new addition in the Jog phase, providing market validation for bids submitted to the D-Market Data Warehouse (2). This module operates similarly to the `W-Market Validation` and follows the same validation workflow. In addition to being validated against information maintained in the D-Market Data Warehouse, the bid must also be validated using data available in the W-Market Data Warehouse. Accordingly, the bid is forwarded to the W-Market Data Warehouse for `W-Market Validation`. Upon successful completion of this validation, the bid proceeds to the next stage in the coordination process (4), namely `T-Operational Validation`, for downstream operational feasibility assessment.

D. The `Operational Validation` module is executed by both the TO and DO to determine whether a market-validated bid is physically feasible on the grid. First, the bid gets validated at the TO level (4) and the notification is sent back to the DO for further validation (5). At the Jog maturity level, the DO has enhanced visibility into grid conditions through AMI data (state estimation or real-time measurements made available due to increased sensing and analytics capability). On startup, it loads the DER Registry and the full distribution feeder model (123-Bus OpenDSS model). The module polls the W-Market Data Warehouse's market-validated queue, waiting for bids that have passed market validation. Once a bid is detected, the module performs a rigorous power-flow-based validation:

    a. Load modeling: AMI data (forecasted/estimated load at each bus) is applied to the feeder model, establishing the expected operating point rather than relying on generic peak/low-load proxies as in the Walk phase.

    b. DER dispatch modeling: Each DER in the bid is dispatched proportionally based on its rated capacity relative to the total available capacity, matching the bid quantity. This proportional allocation reflects a more realistic dispatch assumption than the Walk phase's full-rated-capacity worst-case approach. This assumes

that the DO and aggregator are coordinated for DO to know the actual dispatch strategies used by the aggregator.

c. Capacity feasibility check: Before running the power flow, the module verifies that the aggregate DER capacity meets the bid quantity; if insufficient, the bid is immediately rejected.

d. Voltage feasibility: Power flow is solved with the dispatched DERs and estimated injections, checking whether bus voltages remain within acceptable bounds (0.95–1.05 p.u.).

e. If any violation is detected, the bid is written to the W-Market Data Warehouse rejected/operation queue with a False validation flag. If all checks pass, the bid is written to the operation-validated queue with True validation flag, signaling to the MO that the bid is operationally acceptable for clearing and award (6).

E. The `D-Delivery Optimization` module performs the market clearing process and communicates the outcome to participating aggregators (7). For bids that successfully pass all market and operational validation checks, the module issues an award notification, enabling aggregators to dispatch their resources in accordance with the awarded service. For bids that are not selected or fail validation, the module provides notification along with contextual feedback to help aggregators understand the reasons for the outcome and improve future bidding and operational strategies.

F. The `Market Monitoring` module is currently a placeholder in our modeling activity. Its intended purpose is to perform retrospective analysis of market operations, examining historical data for discrepancies in payments, dispatch instructions, or other aspects of market activity. Once implemented, it will serve as an oversight function to identify irregularities and ensure the integrity of market outcomes over time.

4.2.1.3 Run

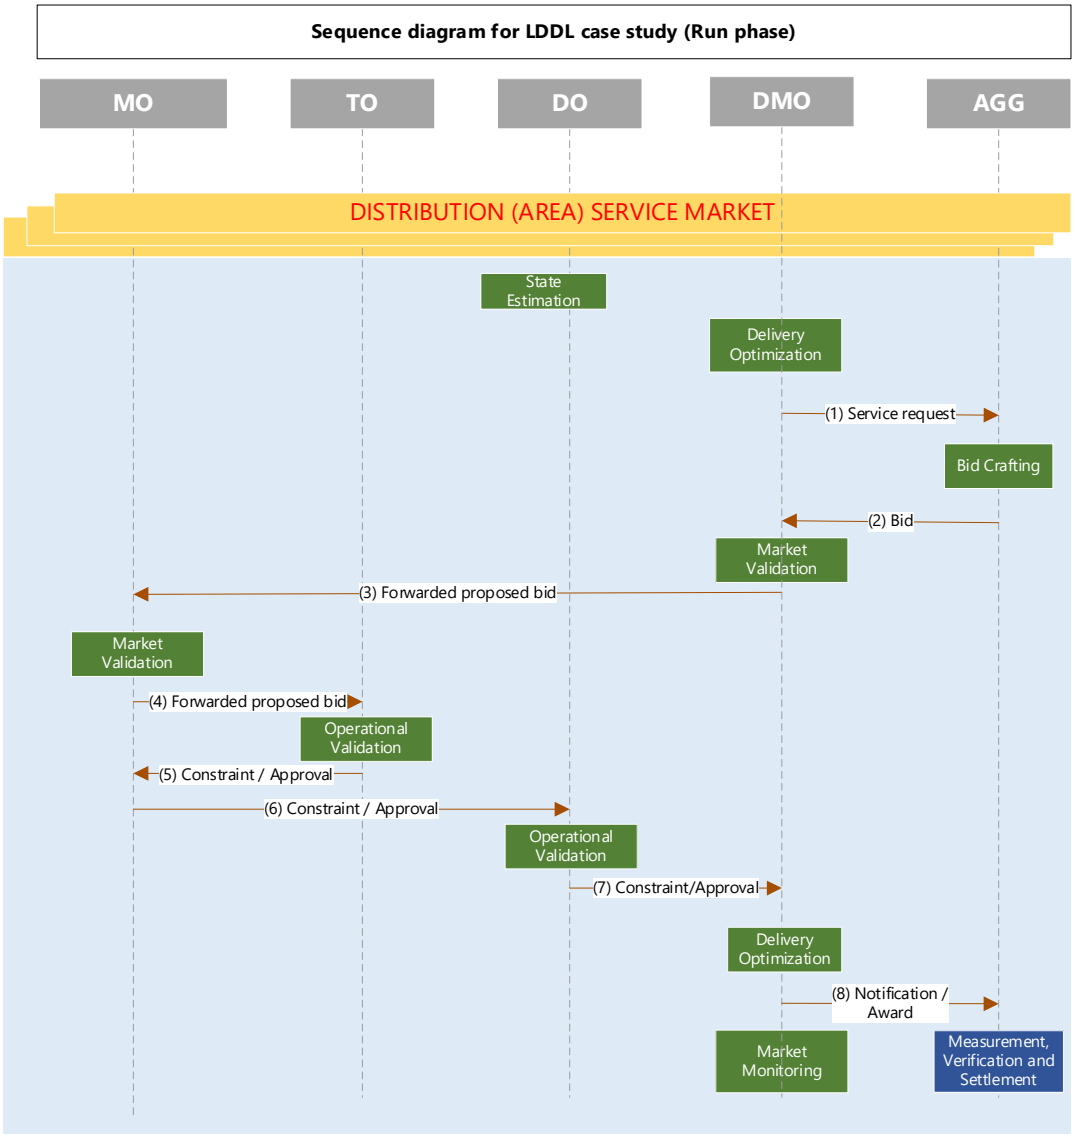


Figure 9. Swimlane Diagram Depicting Workflow for Run Phase of LDDL Case Study; Only Distribution Part is Shown as Bulk Grid Sequence Remains the Same

The Jog and Run phases share the same modular design; the primary distinction lies in how the distribution `Delivery Optimization` is modeled. In the Jog phase, the distribution network is treated as a single circuit, and the DO runs a centralized power flow to detect congestion and call for service. In the Run phase, the distribution network is decomposed into multiple local zones (e.g., distribution areas) as shown in Figure 10, and the distributed agent performs a distributed optimization where each area's sub-circuit is solved independently with boundary power flows and voltages exchanged iteratively until there is system-wide convergence. This area decomposition introduces the need for greater coordination: service requirements must be estimated and procured on a per-zone basis, bids from multiple aggregators across different zones must be aggregated into a unified merit-order supply curve, and market clearing must account for the spatial distribution of resources relative to the location of the constraint.

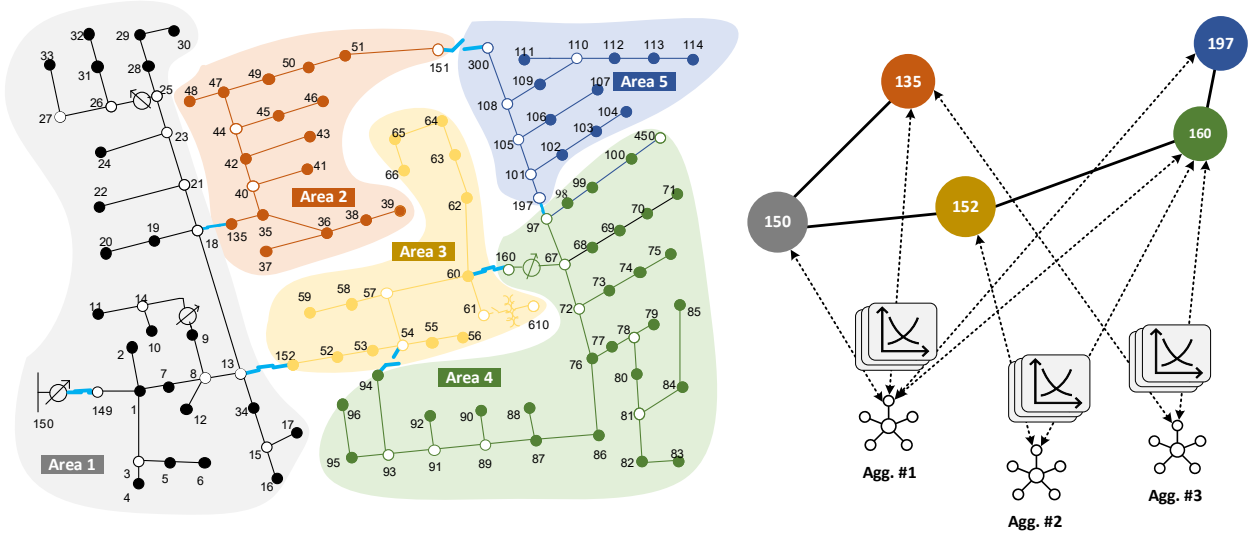


Figure 10. Distributed Optimization and Local Market for Distribution Services

The module advancements in the Run phase are explained in detail below:

A. The `D-Delivery Optimization` module is updated in the Run phase, enabling a distribution-level local market for distribution services and is referred to as `DA-Delivery Optimization`. It acts as the optimization engine for DAMO, performing area-decomposed power flow analysis and market analytics across the feeder. On startup, the module loads a distribution network model that has been partitioned into multiple distribution areas (e.g., five areas of the IEEE 123-Bus feeder) and the associated area topology information from the DO. The module runs a distributed power-flow simulation: it solves each area's sub-circuit independently using OpenDSS, passing boundary power flows and voltage information between adjacent areas iteratively until convergence. If any violation is detected (such as undervoltage in this case) and all direct DO controls are exhausted, the module estimates the service requirement by issuing a "call for service" by writing a service request file—including the estimated service quantity—to the Distribution Area Market Data Warehouse (DA-Market Data Warehouse), opening a local market opportunity for DER participation (1). In this case study, reactive power (kVAR) service is used to mitigate the undervoltage scenario. After calling for service, the module waits (polls) for validated bid files to appear in the DA-Market Data Warehouse. Once bids arrive, it performs market clearing.

A. The `Bid Crafting` module executed by the aggregator is mostly unchanged from the Jog phase. On startup, the module waits for service opportunities to be posted by, rather than a single DMO, the multiple DAMOs in the service queue of their respective DA-Market Data Warehouse instances. Whenever a service opportunity is found, the module has each aggregator post its bid (2) in the submissions queue if eligible for the corresponding DAMO.

B. The `DA-Market Validation` module instance of each DAMO acts in a manner similar to the `D-Market Validation` module of the Jog phase. The module instance ensures that bids posted in the DA-Market Data Warehouse instance of the area meet market requirements before allowing them to proceed in the workflow. On startup, the instance waits for bids in the submissions queue of the DAMO's instance of DA-Market Data

Warehouse. Whenever a bid is detected (2), the module instance processes the bid. Besides the customary checks, the module instance verifies that each DER listed in the bid actually resides at a location permitted by the DAMO. Depending on whether a bid passes all checks, it is moved to either (3) the market-validated queue or the market-rejected queue of the DA-Market Data Warehouse instance. This bid also goes to the W-Market Data Warehouse (3) for validation using the `W-Market Validation`.

C. First, the bid undergoes operational validation (4) using the `T-Operation Validation`. Once the bid is validated, it goes into the respective queue of W-Market Data Warehouse (5). The `D-Operational Validation` module of the DO is updated from the Jog phase, and it polls the bid from W-Market Data Warehouse (6). On startup, it loads the DER Registry and its corresponding area network model. The module polls the market-validated queue of the corresponding DA-Market Data Warehouse instances. Despite being read from different instances, the bids are evaluated simultaneously in the same distributed optimization. Because this is a local problem in a distribution context, this is coordinated with `DA-Delivery Optimization` module (7). Therefore, coordination is required among the distributed agents before finally delivering an award.

    a. Individual aggregator bid curves (price–quantity pairs) are merged and sorted by ascending price, creating a single area-level supply stack for dispatch.

    b. The DAMO walks the merged merit-order curve for the target area, dispatching DERs from cheapest to most expensive until the required kVAR is met. The last accepted offer sets the marginal clearing price, and each aggregator's cleared quantity is tracked (`DA-Delivery Optimization`).

    c. After clearing, the distributed power flow (`Operational Validation`) is re-executed with the awarded reactive injections dispatched at each DER's bus across all five areas, iteratively solving area-by-area with boundary voltage/flow exchange to confirm the system voltage is restored above 0.95 p.u.

    Depending on whether the simulation results improve the voltage profile, a bid is moved to either the operation-validated queue or the operation-rejected queue of its corresponding DA-Market Data Warehouse instance. For bids that successfully pass all market and operational validation checks, the `DA-Delivery Optimization` module issues an award notification (8), enabling aggregators to dispatch their resources in accordance with the awarded service. For bids that are not selected or fail validation, the module provides notification along with contextual feedback to help aggregators understand the reasons for the outcome and improve future bidding strategies.

D. The `Market Monitoring` module is currently a placeholder in our modeling activity. Its intended purpose is to perform retrospective analysis of market operations, examining historical data for discrepancies in payments, dispatch instructions, or other aspects of market activity. Once implemented, it will serve as an oversight function to identify irregularities and ensure the integrity of market outcomes over time.

4.2.2 Extreme Cold Weather Event

The goal of the Extreme Temperature case study is to demonstrate the operation and evaluate the performance of the coordination framework during an extreme weather event. Such events challenge the system and can lead to its collapse if not mitigated properly. The data and models necessary to execute according to the framework are as follows:

- The TO has access to a model of the transmission system under its control, which in this case is the IEEE 8-bus system. The relevant data can be provided in a variety of formats such as PSS/E or PowerWorld files or any other format appropriate for data sharing, as long as the data can be used to solve an OPF problem. The TO also receives information about outages related to the extreme weather event. Such information is necessary for an effective response.
- The DO also has access to a model of the system under its control, which is the IEEE 123-Bus system here. A distribution model may also be available in multiple formats such as OpenDSS. The DO also has access to load forecast data, which could be in a format such as CSV or JSON or any other appropriate data sharing format.
- The MO has all the relevant information necessary to operate a market organized in a Market Data Warehouse. A Market Data Warehouse should contain all information related to submitted bids and their processing. It needs to track service calls, and bids as they are submitted, validated (and classified as either accepted or rejected) for market and operations, and finally either awarded or rejected.
- The aggregator is responsible for registering the DERs it controls into a DER Registry and for updating the DER Registry for relevant changes.
- The DER Registry maintains information on DERs installed within its area of responsibility according to the schemas presented with the coordination framework.

4.2.2.1 Walk

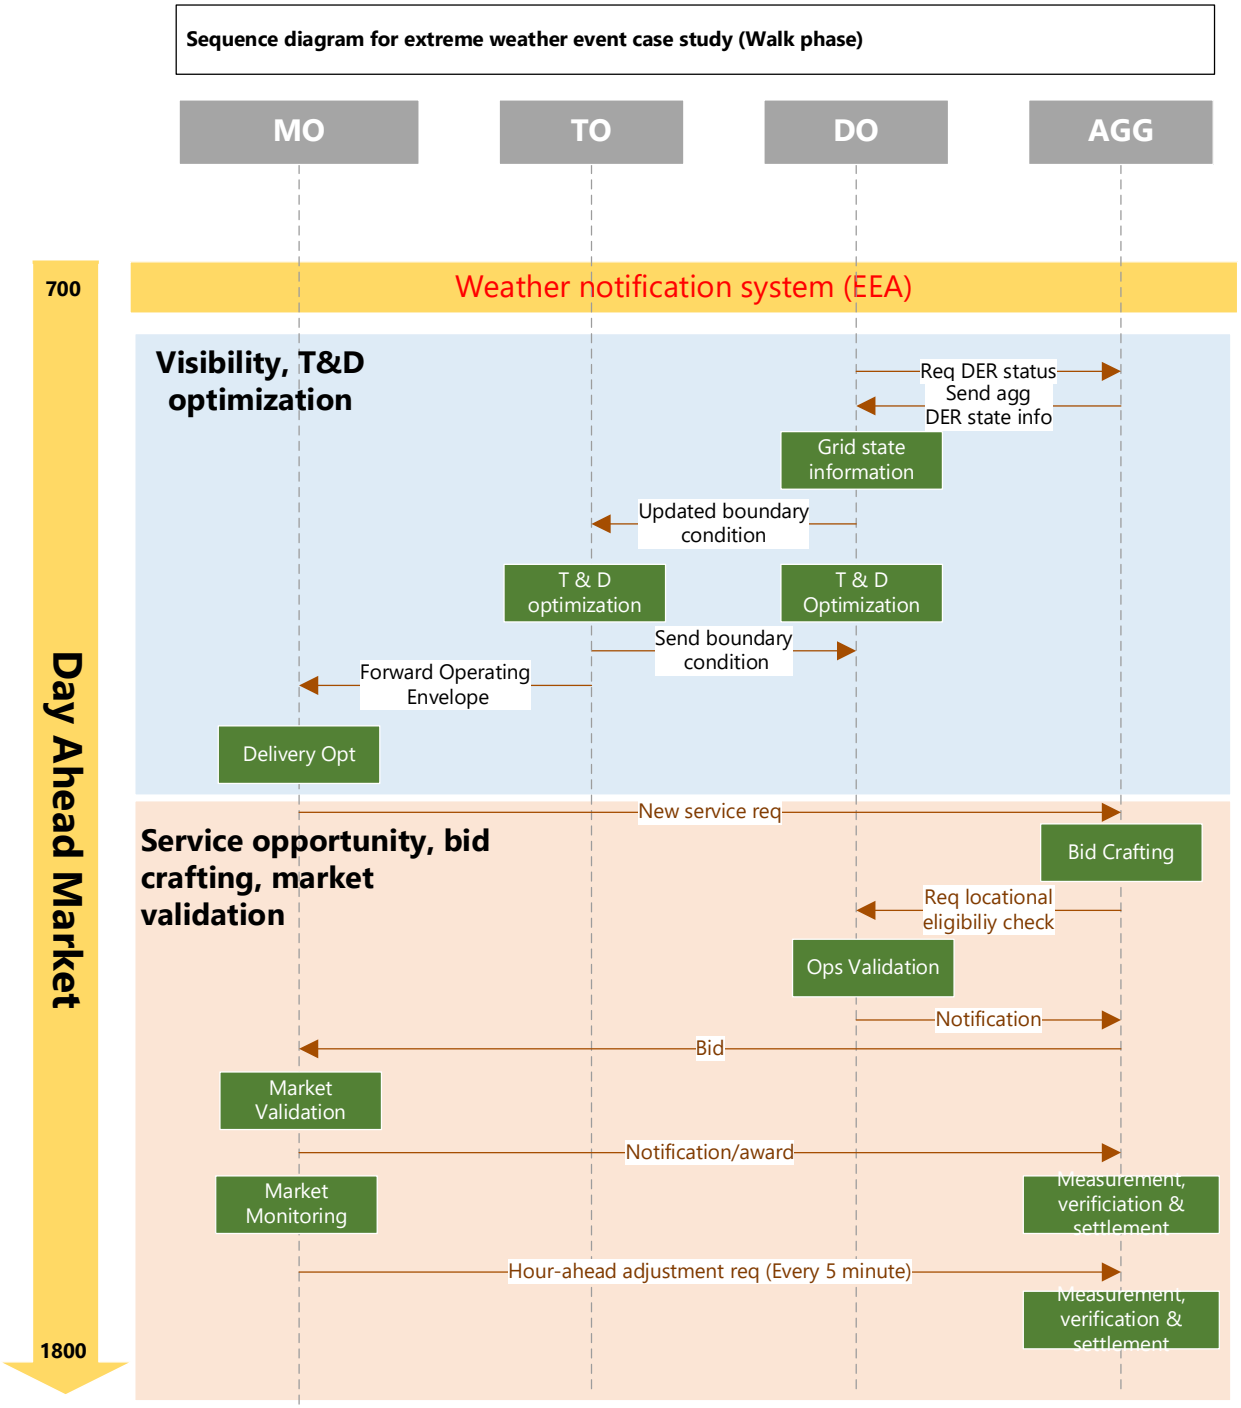


Figure 11. Swimlane Diagram Depicting Workflow for Walk Phase of Extreme Temperature Case Study

As was the case with the LDDL case study, the various actors communicate and coordinate through a centralized Market Data Warehouse. In the Walk phase, there is only one Market

Data Warehouse. Its responsibility is at the wholesale level, so it is denoted as `W-Market Data Warehouse` despite its uniqueness. Once more, each module begins execution through a listening state and remains idle until all necessary information appears in the `W-Market Data Warehouse`. Each module's outputs are stored in the `W-Market Data Warehouse` at the end of execution. In the Walk phase, the `Market Monitoring` module is a placeholder for future analyses and market irregularity detection. The rest of the modules operate as follows.

A. The DO executes the `T&D Node Visibility` function to inform the TO about the status of the distribution system.

B. The `W-Delivery Optimization` module begins execution by receiving a load forecast as well as the expected generation shortfall at the wholesale level. It also receives the model of the transmission network from the TO. Then, it performs a load flow to detect whether the grid needs more generation or not. If a need for more generation is indeed detected, the module issues a call for service to the `W-Market Data Warehouse`.

C. The `Bid Crafting` module enables market participation by an aggregator in response to a call for service. Specifically, the aggregator assembles a set of resources that can respond to a received call for service, produces a price versus quantity curve, and submits it to the `W-Market Data Warehouse`. An aggregator can poll the DER Registry for relevant information that may be necessary to form a bid.

D. The `W-Market Validation` module is responsible for the first validation process of aggregator-submitted bids. It begins execution by checking whether a bid follows the appropriate data schema or is malformed. Then, it validates the bid against market rules. In the Walk phase, the only applicable rules are market conflict and dual participation. Finally, the `W-Market Validation` module places each investigated bid into either the accepted or rejected location of the `W-Market Data Warehouse`.

E. The `D-Operational Validation` module is responsible for determining whether a bid is physically feasible. To perform this role, this module uses its knowledge of distribution system topology to perform load flows and check whether there are any operational limit violations in case of bid acceptance. Because this module has limited visibility into the distribution grid conditions in the Walk phase, it executes its load flow simulations assuming worst-case grid conditions, which includes peak-load and low-load conditions.

F. Finally, control returns to the `W-Delivery Optimization` module, which waits for the operational validation stage to conclude, checks the validated bids, and awards or rejects bids through a market-based process.

4.2.2.2 Jog

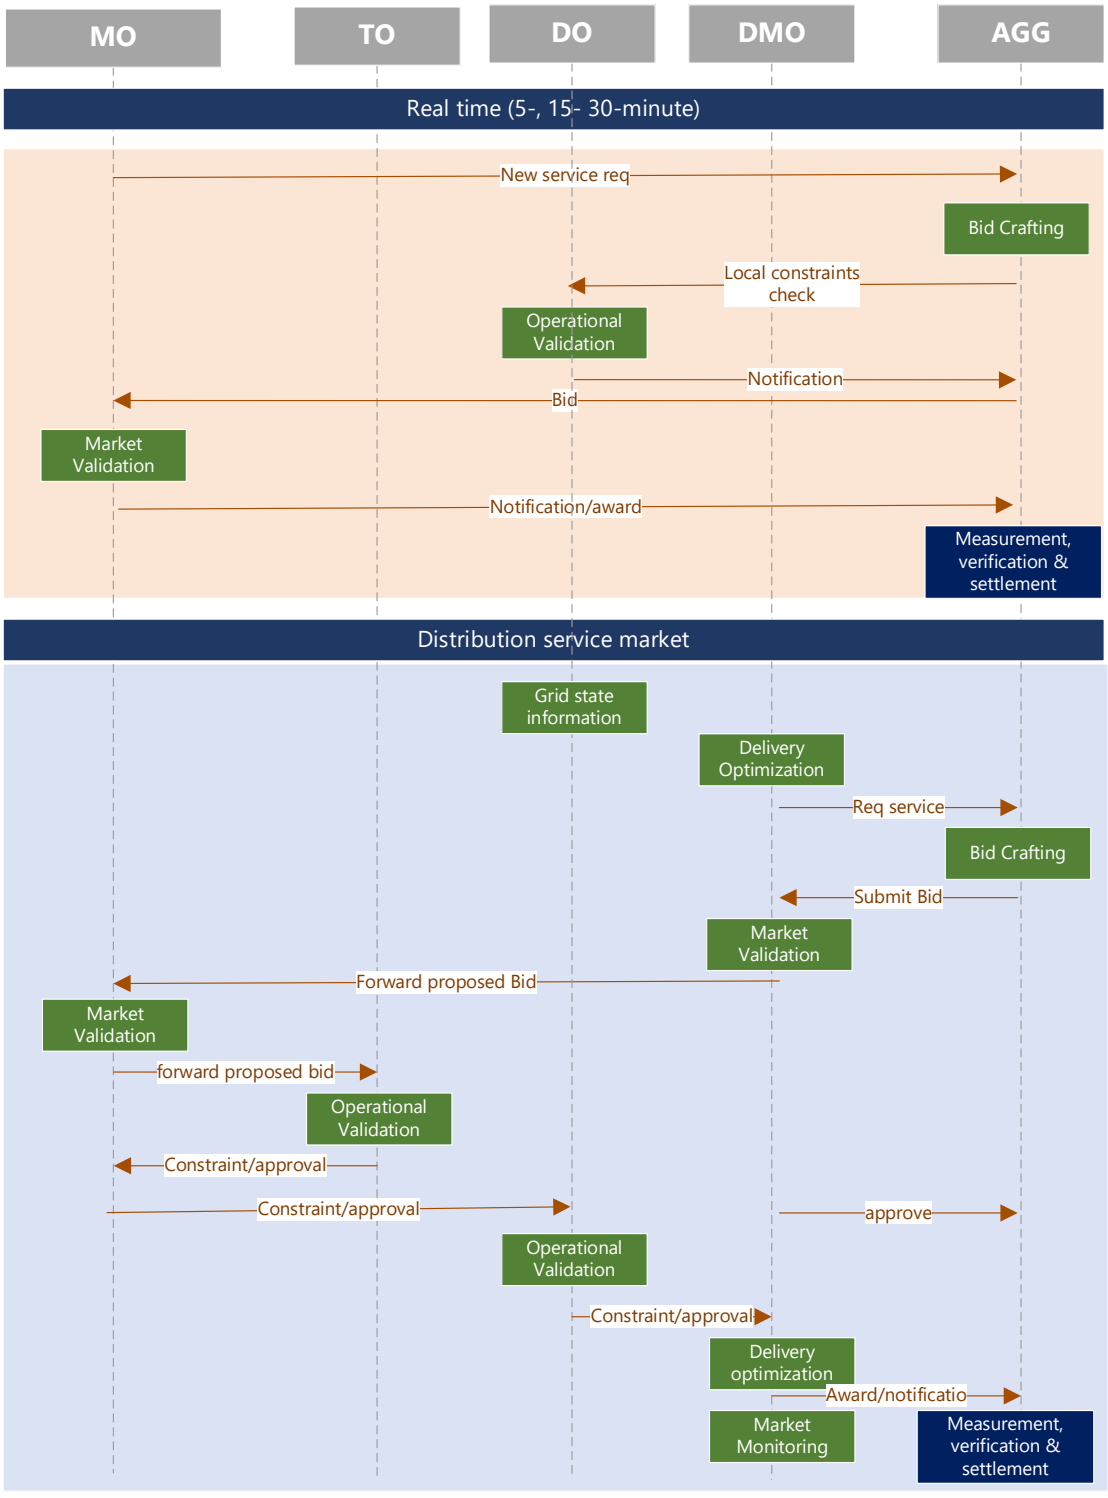


Figure 12. Swimlane Diagram Depicting Workflow for Jog Phase of Extreme Temperature Case Study

The Jog phase adds a distribution-level market, which is represented by a `D-Market Data Warehouse` and a DMO. The additional market means that market-related modules should also be added to the distribution side. In the Jog phase, the `Market Monitoring` module is also left as a placeholder for future development. The specific modules are as follows.

A. The `D-Delivery Optimization` mirrors the functionality of its wholesale counterpart. Its first responsibility upon entering execution is to receive a distribution grid model (in a format such as OpenDSS) and AMI measurement data from the DO. It uses this information to perform a distribution load flow to detect whether grid operation is feasible with the available generation. If that is not the case, a call for service is issued through the corresponding `D-Market Data Warehouse`. The `D-Delivery Optimization` module then waits for the output of operational validation. Once it receives it, it checks the operationally validated bids and issues awards for the lowest-cost bids that cover the anticipated generation shortfall.

B. The `Bid Crafting` module is again utilized by an aggregator. However, in this phase it can respond to both wholesale and distribution calls for service. To achieve this, the `Bid Crafting` module now needs to know the correspondence between available resources and available markets. Otherwise, its behavior is the same as in the Walk phase. Specifically, it identifies both a set of resources and a price versus quantity curve and uses them to form a bid and submit to either the `W-Market Data Warehouse` or the `D-Market Data Warehouse`. Once again, an aggregator can poll the DER Registry for any necessary information.

C. The `D-Market Validation` module is the distribution-level counterpart to `W-Market Validation` and performs similar checks on each received bid. Each bid needs to be validated for both wholesale and distribution market operations, and the two modules work serially to achieve this objective. This collaboration is made possible through seamless communications between `W-Market Data Warehouse` and `D-Market Data Warehouse`.

D. The `T-Operational Validation` module performs load flow at the transmission level for each bid it receives to determine whether acceptance of the bid would violate transmission limits. If no such violation is detected, the bid is marked as accepted and sent to the `W-Market Data Warehouse` for further processing. Otherwise, the bid is marked as rejected in the `W-Market Data Warehouse`.

E. The `D-Operational Validation` module in the Jog phase has enhanced visibility into the distribution grid as it can access AMI data. Therefore, it can combine these data with the distribution-level grid model it receives from the DO to execute distribution load flow and determine whether a bid can endanger grid operations if accepted. It should be noted that this module processes only bids that have already been marked as acceptable at the market and transmission operations levels. In the Jog phase, the `D-Operational Validation` module has accurate information on the load at each distribution node unlike the Walk phase, where it had to assume worst-case conditions. Each DER included in the bid is dispatched proportionally to its rated capacity for load flow simulations. Any violation of operational limits leads this module to reject the bid; otherwise, the bid is marked as accepted.

F. Finally, the `D-Delivery Optimization` module assumes control again, uses all bids that have passed all validation checks until this stage to clear the market, and issues dispatch signals to the aggregators with awarded bids.

4.2.2.3 Run

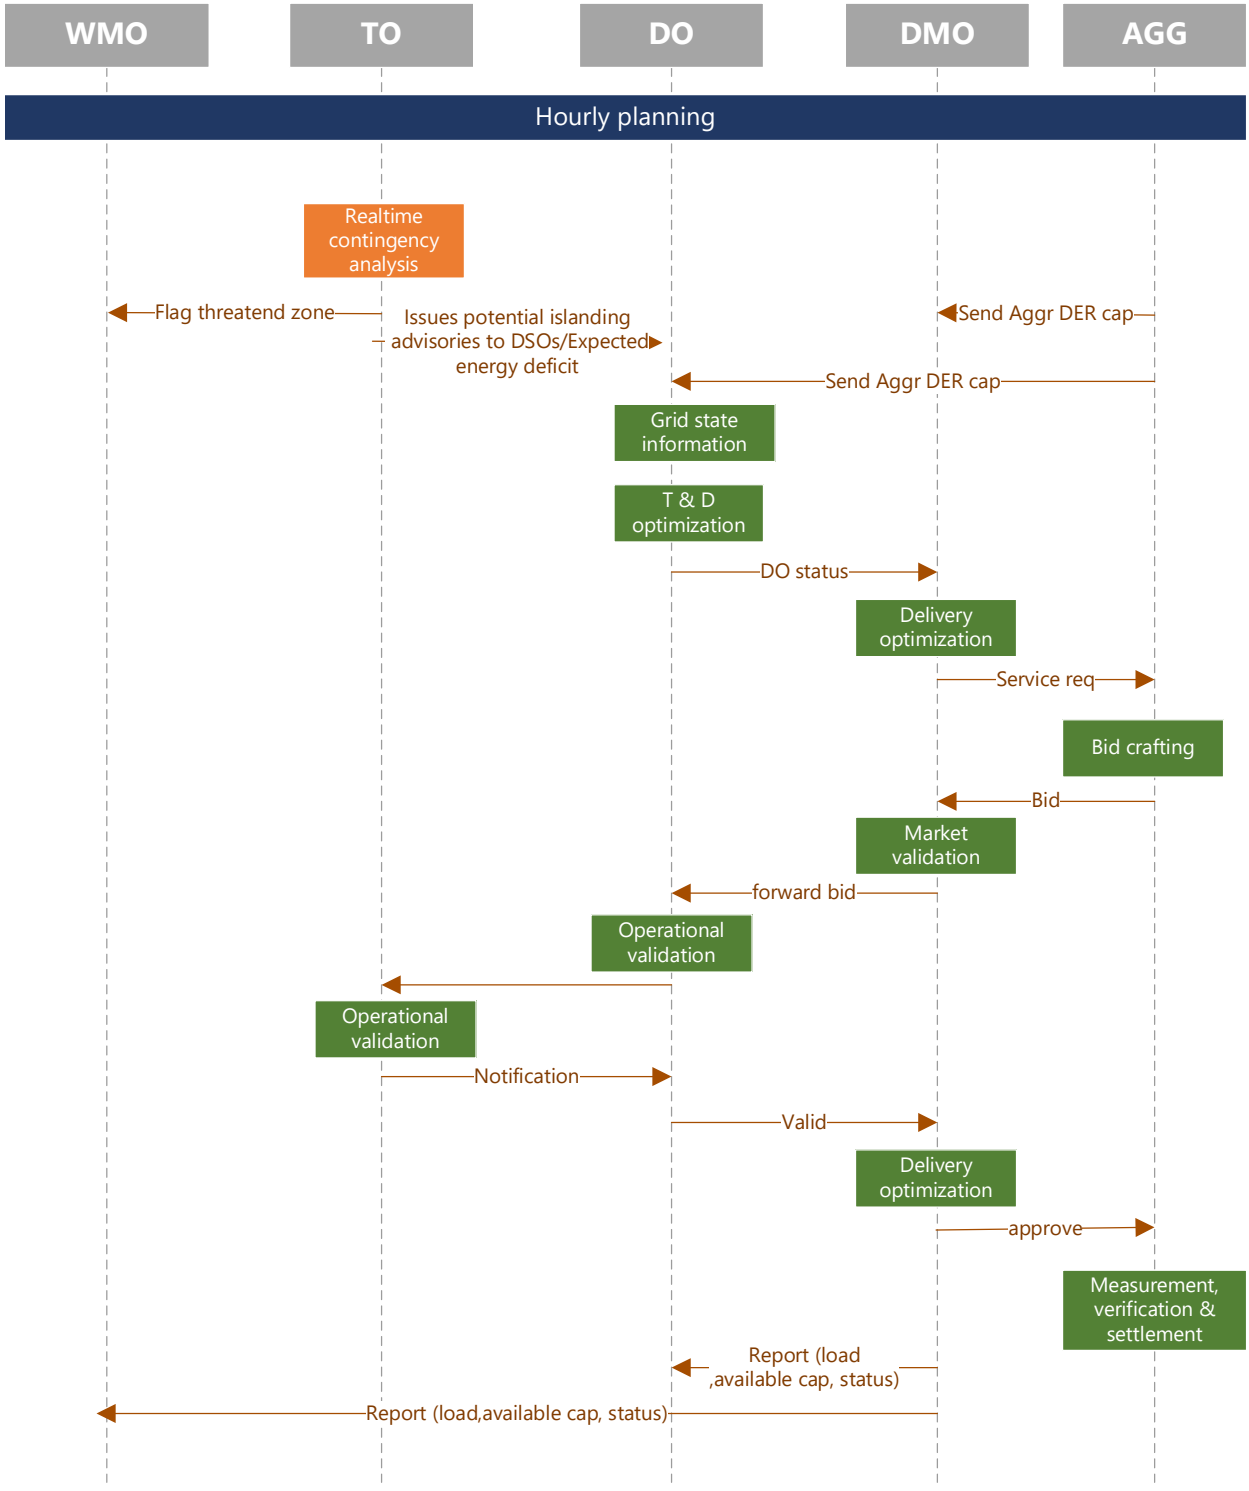


Figure 13. Swimlane Diagram Depicting Workflow for Run Phase of Extreme Temperature Case Study

The main novelty of the Run phase is the splitting of the distribution-level grid, which was assumed to be under the control of a single entity in Walk and Jog, into independent areas (or zones). To maintain coordination, new distributed control agents are introduced, one for each distribution-level independent area. Their operation will be explained in more detail in the module explainer below. The introduction of these new distributed agents has two main effects on the RCA. The first is the need for distributed optimization and load flow techniques and the second is the need for spatial awareness for bids and resources. Once more, the `Market Monitoring` module is not implemented for simulation, as its focus is on retrospective analyses of system operations. The modules in the Run phase operate as follows.

E. The new `DA-Delivery Optimization` module performs the functions of the `D-Delivery Optimization` module with the main difference being that now the module scope is an independent distribution-level area rather than the whole distribution grid. A new distribution area MO (DAMO) is also introduced to replace the distribution market operator (DMO) of the Jog phase. The new module performs optimization and load flow only for the distribution area under control, passes information on boundary voltages and flows to neighboring areas and continues iterating until convergence (i.e., agreement between all neighboring areas about flows and voltages). Call for service requests are now also issued on a per-area basis and are stored to a `DA-Market Data Warehouse`. The main goal of the `DA-Delivery Optimization` module is to properly estimate the amount of necessary generation required to cover any shortfalls.

F. The `Bid Crafting` module is similar to its Jog counterpart, with the main difference being that instead of forming bids for a single wholesale or distribution market, it instead creates per-area bids. Therefore, increased awareness about the exact location of resources is now necessary. The module can poll the DER Registry once again for any relevant information.

G. The new `DA-Market Validation` module is also a per-area version of the `D-Market Validation` module, performing similar checks as in the Walk and Jog phases.

H. The `D-Operational Validation` module converts the checks performed for distribution-level operational validation in the Jog phase into a distributed form. Bids that have been cleared from a market perspective are evaluated here simultaneously across different distribution-level areas using distributed load flow.

4.2.3 Cascading Blackout

The intent of the Blackout case study is to demonstrate how the coordination framework might be leveraged to facilitate restoration of the transmission system (as well as distribution systems within) following a blackout. The relevant models and data organized by an actor within the case study are summarized below:

- The TO manages the model for a 9-bus transmission network. The model may be represented in PowerWorld or PSS/E formats, as was done for the IEEE 39-Bus System in the LDDL case study of Section 4.2.2. Alternatively, because the Blackout case study entails the analysis of the 9-bus network primarily through OPF rather than conventional power

flow, the specifications of the network may be adequately represented in a format suitable for that purpose.

- The DO manages the IEEE 123-Bus Feeder model in OpenDSS format. Although the transmission network is considered to encompass four DOs—one at each of Buses 6 and 7 and two at Bus 8—solely DO #1 at Bus 6 is selected for demonstration of the Jog and Run phases of the Blackout case study. In the Jog phase, DO #1 entrusts a single DMO with control of the DERs in the IEEE 123-Bus Feeder. In the Run phase, the DMO distributes control among four DAMOs, each of which maintains the OpenDSS model for its area sub-circuit within the IEEE 123-Bus Feeder.
- As in the other two case studies, the Jog phase has the DMO or a third-party entity maintain the D-Market Data Warehouse with structured queues (submissions, market-validated, operation-validated, market-rejected, operation-rejected, award, service) containing JSON bid files that flow through validation stages. In the Run phase, each DAMO has its own DA-Market Data Warehouse instance with these queues.

4.2.3.1 Walk

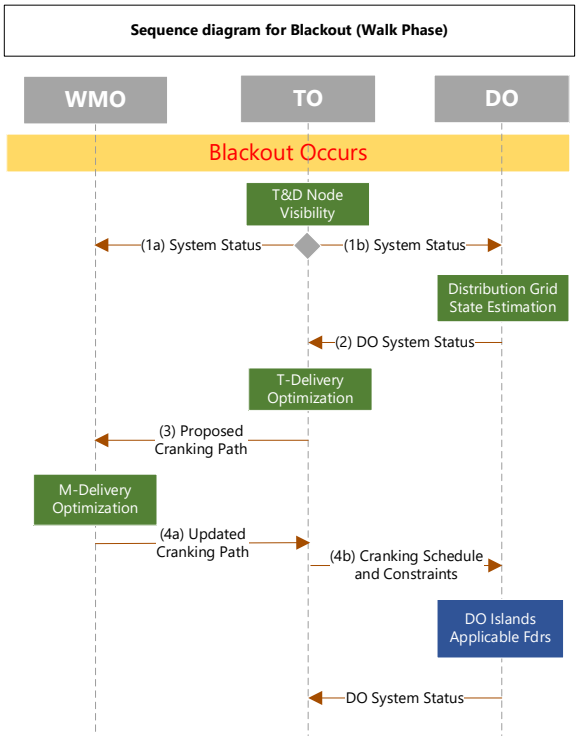


Figure 14. Swimlane Diagram Depicting Workflow for Walk Phase of Blackout Case Study

The Walk phase, whose workflow is depicted in Figure 14, initiates the Blackout case study as primarily consisting of two modules: `Distribution Grid State Estimation` and `Delivery Optimization`. The latter is actually expressed as two sub-modules, managed by the TO and WMO. While not technically a module, the `T&D Node Visibility` function also serves a role. The modules and function coordinate through file-based communication as described in Section 4.1.

A. The `T&D Node Visibility` function of the TO yields information on TO system status relevant to blackout restoration. Execution of the function initiates the workflow for the Blackout case study. The function issues (1a) to the WMO a version of the transmission network model updated to reflect impacts of the blackout event. It also issues (1b) to the DOs a notice that black-start restoration is commencing at the transmission level.

B. The `Distribution Grid State Estimation` module instance of each DO shares information related to DO system status. On startup, the module instance waits for a black-start restoration notice from the TO. Upon receipt (1b), the module instance sends back (2) to the TO the MW value of total peak load within the distribution circuit of the DO.

C. The Delivery Optimization sub-module belonging to the TO is designated as `T-Delivery Optimization`. On startup, the sub-module waits for the peak load data from all the DOs. Upon receipt from all of them (2), the sub-module solves an optimal black-start restoration problem (Maloney et al. 2025). The proposed cranking path—in terms of time-series values concerning the amount of load supplied at each transmission bus and the energization state of each transmission line—encoded by the solution values is then passed (3) to the WMO.

D. The Delivery Optimization sub-module belonging to the WMO is designated as `M-Delivery Optimization`. On startup, the sub-module waits for the proposed cranking path from the TO. Upon receipt (3), the sub-module performs any required adjustments to the cranking path and then sends back (4) the updated cranking path. In the case of multiple TOs (perhaps managing different areas of the bulk power system) under the WMO, the intent of `M-Delivery Optimization` is to determine and send back to the TOs edited cranking paths reconciling any discrepancies or conflicts identified between the originally proposed cranking paths. In the case of a single TO under the WMO, `M-Delivery Optimization` may simply allow the proposed cranking path to pass through and be sent back to the TO practically unmodified. In either case, the TO then takes the updated cranking path from the WMO, extracts for each DO the time series of the amount of load supplied by the transmission network, and issues (4b) to each DO its corresponding time series.

4.2.3.2 Jog

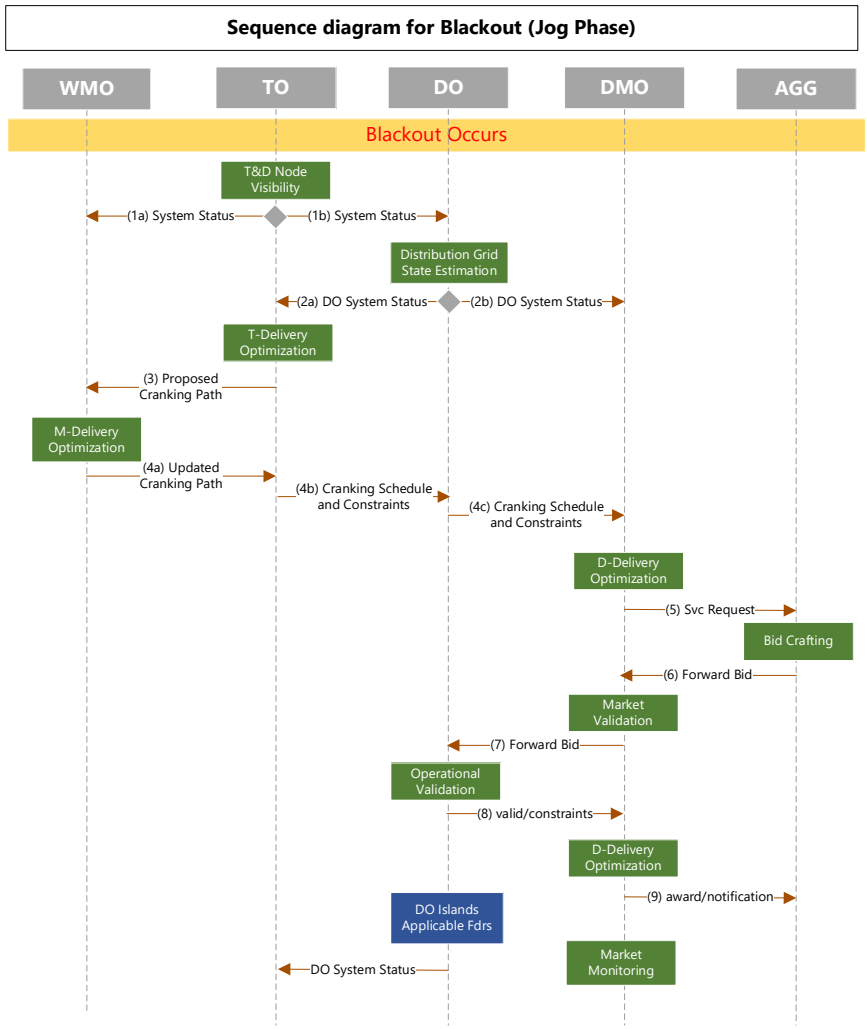


Figure 15. Swimlane Diagram Depicting Workflow for Jog Phase of Blackout Case Study

The key development in the Jog phase, whose workflow is depicted in Figure 15, is the enabling of DER participation in the blackout restoration of distribution grids. Although black-start restoration at the level of the bulk power system still does not originate at the distribution level, DERs are able to restore supply to some loads in a distribution circuit even before energization arrives along cranking paths to the transmission bus hosting the distribution circuit. Compensation to aggregators and DER owners for participation in blackout restoration is determined beforehand by contract; however, the trappings of a distribution-level "market" are nonetheless employed to orchestrate their contributions. The Jog phase introduces a DMO and DER aggregators as new entities, as well as a D-Market Data Warehouse for coordination of DER activities.

The Jog phase of the Blackout case study consists mainly of five Python modules: `Distribution Grid State Estimation`, `Bid Crafting`, `D-Market Validation`, `Operational Validation`, and `Delivery Optimization`, the last of which is expressed

as three sub-modules managed by various entities (i.e., TO, WMO, and DMO). The addition of new modules and evolution of modules existing in the Walk phase are as follows:

A. The `T&D Node Visibility` function of the TO is effectively unchanged from the Walk phase. It still yields information on TO system status relevant to blackout restoration: an updated transmission network model to the WMO (1a) and a black-start restoration notice to the DOs (1b).

B. The `Distribution Grid State Estimation` module instance of each DO is mostly unchanged from the Walk phase. Upon receipt of a black-start restoration notice from the TO (1b), the module instance sends back (2a) to the TO the MW value of total peak load within the distribution circuit of the DO. Additionally, the Jog phase has the instance for DO #1 issue (2b) to DMO #1—the DMO of DO #1—a notice that distribution-level blackout restoration is starting.

C. The `Bid Crafting` module of the aggregators is newly added to essentially report the availability of DERs to participate in the distribution-level "market" of DMO #1 for blackout restoration. On startup, the module waits for service opportunities to be posted (5) by DMO #1 in the service/ queue of the D-Market Data Warehouse. When a service opportunity is detected, the module has each aggregator with any DERs eligible for the opportunity post (6) a bid in the submissions/ queue. The total MW amount of a bid simply includes the full available capacity of each eligible DER. The bid does not include any pricing information because compensation is determined beforehand by contract.

D. The `D-Market Validation` module is newly added to ensure that bids posted in the D-Market Data Warehouse meet market requirements before allowing them to proceed in the workflow. On startup, the instance waits for bids in the submissions/ queue of the D-Market Data Warehouse. Whenever a bid is detected (6), the module instance performs on the bid the customary checks—as in the other two case studies—for issues such as DER capacity, market conflict, and dual participation. If a bid passes all checks, then the bid is moved (7) from the submissions/ queue to the market-validated queue. Otherwise, if a bid fails any check, then it is moved to the rejected/market/ queue.

E. The `Operational Validation` module is newly added to verify that market-validated bids cause no violations of operational limits in the distribution circuit. On startup, the module waits for some time. Once that period has elapsed, the module reads the bids that have arrived (7) in the market-validated queue of the D-Market Data Warehouse and processes all these market-validated bids simultaneously. Specifically, the module executes an OpenDSS simulation assuming each DER listed in any of the bids injects real power at its nameplate capacity. Details on simulation setup and violation assessment may be found in the Jog-phase description of the LDDL case study, in Section 4.2.2.2. If no violation of operational limits is detected within the simulation results, then the module moves (8) all the bids in the market-validated queue of the D-Market Data Warehouse to the operation-validated queue. Otherwise, the module moves all the bids to the operation-rejected queue.

F. The `T-Delivery Optimization` sub-module is effectively unchanged from the Walk phase. It still waits for the peak load data from all the DOs. Upon receipt from all of them (2a), it solves an optimal black-start restoration problem and issues (3) to the WMO the proposed cranking path.

G. The `M-Delivery Optimization` sub-module is mostly unchanged from the Walk phase. It still waits for the proposed cranking path from the TO. Upon receipt (3), the sub-module sends back (4a) to the TO the updated cranking path. The TO then issues (4b) to each DO its extracted time series of load supplied. However, in the Jog phase, DO #1 then in turn forwards (4c) the time series it receives to DMO #1.

H. The Delivery Optimization sub-module belonging to DMO #1 is newly added with the designation of `D-Delivery Optimization`. Compensation for participation in blackout restoration is pre-determined by contract in the distribution-level "market" of DMO #1. Consequently, unlike in a typical electricity market, the aggregators do not concern themselves with competitive pricing, as a merit order is not considered in dispatching DERs.

   On startup, `D-Delivery Optimization` waits for the cranking schedule and constraints from DO #1. Upon receipt (4c), the sub-module requests (5) DER participation in blackout restoration (not black-start, which concerns cranking paths for energization of the transmission network) by posting a service opportunity in the service/ queue of the D-Market Data Warehouse. Indicating blackout restoration as its purpose, this service opportunity contains such details as the start time and duration, as well as characteristics required of DERs participating in the opportunity.

   The `D-Delivery Optimization` module then waits for some time. Once that period has elapsed, if the instance finds any bids (8) in the operation-validated queue of the D-Market Data Warehouse, it reads all these bids that have passed both market and operational validation. The module solves a LinDistFlow-based OPF that aims to supply as much load as possible in the distribution circuit not already supplied by the bulk power system, while limiting the injection of each participating DER to its available capacity.

   With the OPF solution in hand, `D-Delivery Optimization` takes each bid from the operation-validated queue of the D-Market Data Warehouse, appends to the bid an MW amount to reflect what portion was actually scheduled by the solution, and moves (9) the bid to the award queue. The module also issues (9) notifications to the aggregators that had submitted the validated bids.

4.2.3.3 Run

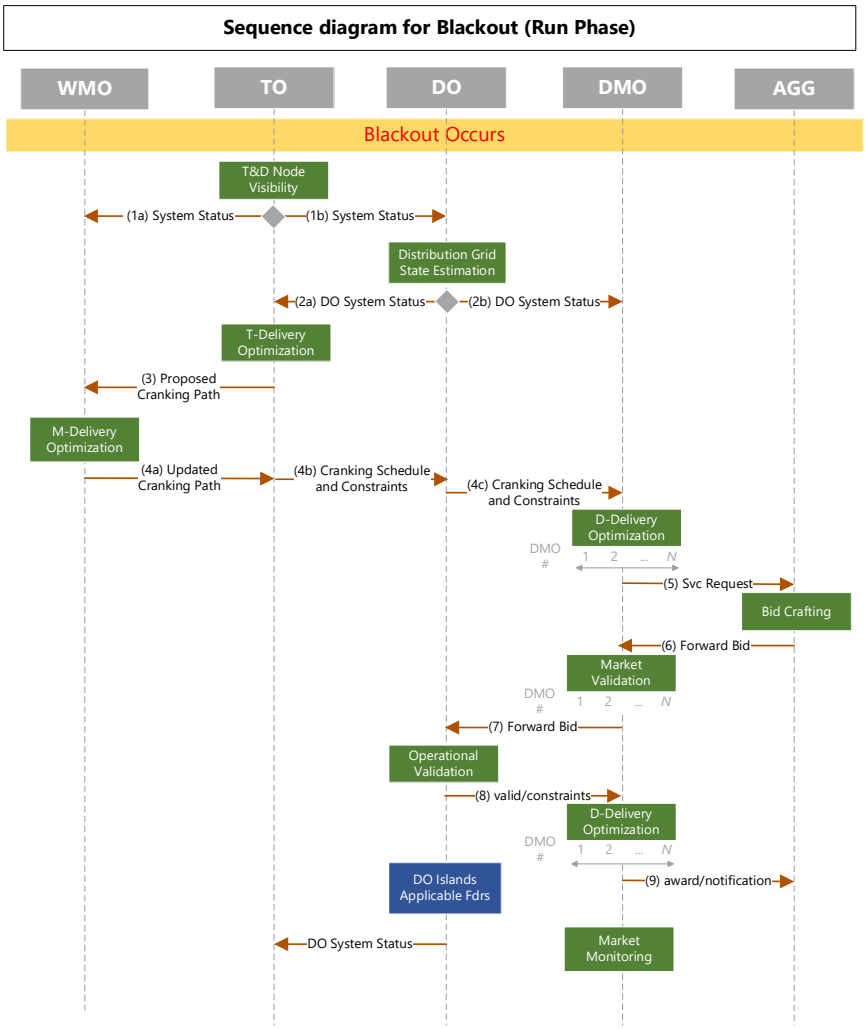


Figure 16. Swimlane Diagram Depicting Workflow for Run Phase of Blackout Case Study

The key development in the Run phase, whose workflow is depicted in Figure 16, is the division of DER control among four DAMOs. A new kind of entity, each DAMO operates a local "market" to orchestrate the contributions of DERs to blackout restoration in its area sub-circuit. The DAMO of each area coordinates with the DAMOs of the neighboring areas in a paradigm of distributed optimization until all DAMOs have experienced convergence of the subproblems for their respective areas. The division of the IEEE 123-Bus Feeder into four area sub-circuits and an islanded microgrid—which is omitted from the distributed optimization—is depicted in Figure 17. DER aggregators now interface with the multiple DAMOs through their corresponding DA-Market Data Warehouse instances—and respective service opportunities—rather than a single, unified D-Market Data Warehouse (Yip et al. 2026).

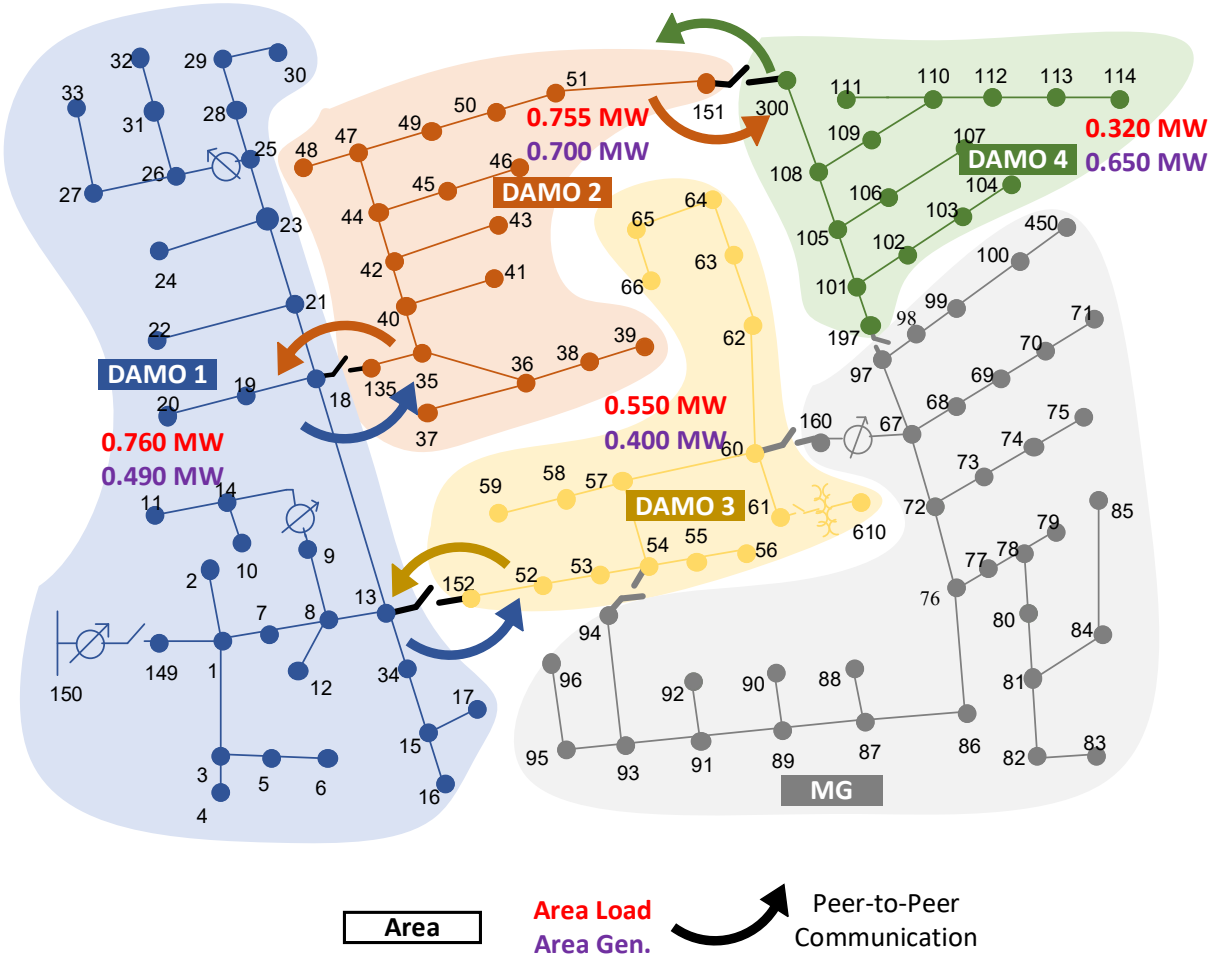


Figure 17. Distributed Optimization Across Area Sub-Circuits for Blackout Restoration

The Run phase of the Blackout case study consists mainly of five Python modules: `Distribution Grid State Estimation`, `Bid Crafting`, `D-Market Validation`, `Operational Validation`, and `Delivery Optimization`, the last of which is expressed as four sub-modules managed by various entities (i.e., TO, WMO, DMO, and DAMO). The addition of new modules and evolution of modules existing in the Jog phase are as follows:

A. The `T&D Node Visibility` function of the TO is effectively unchanged from the Walk and Jog phases. It still yields information on TO system status relevant to blackout restoration: an updated transmission network model to the WMO (1a) and a black-start restoration notice to the DOs (1b).

B. The `Distribution Grid State Estimation` module instance of each DO is effectively unchanged from the Jog phase. Upon receipt of a black-start restoration notice from the TO (1b), the module instance sends back (2a) to the TO the MW value of total peak load within the distribution circuit of the DO. Additionally, the instance for DO #1 issues (2b) to DMO #1 a notice of distribution-level blackout restoration.

C. The `Bid Crafting` module of the DER aggregators is mostly unchanged from the Jog phase. On startup, the module waits for service opportunities to be posted (5) by, rather than a single DMO, the multiple DAMOs in the service/ queues of their respective DA-Market Data Warehouse instances. Whenever a service opportunity is found, the module has each aggregator post (6) its bid in the submissions/ queue if eligible for the corresponding DAMO.

D. The `DA-Market Validation` module instance of each DAMO acts in a manner similar to the `D-Market Validation` module of the Jog phase. The module instance ensures that bids posted in the DA-Market Data Warehouse instance of the area meet market requirements before allowing them to proceed in the workflow. On startup, the instance waits for bids in the submissions/ queue of the DAMO's instance of DA-Market Data Warehouse. Whenever a bid is detected (6), the module instance processes the bid. Besides the customary checks, the module instance verifies that each DER listed in the bid actually resides at a location permitted by the DAMO. Depending on whether a bid passes all checks, it is moved to either (7) the market-validated queue or the market-rejected queue of the DA-Market Data Warehouse instance.

E. The `Operational Validation` module of the DO is mostly unchanged from the Jog phase. On startup, the module waits for some time but then reads the bids that have arrived (7) in the market-validated queues of the multiple DA-Market Data Warehouse instances. Despite being read from different instances, the bids are evaluated simultaneously in the same simulation. Depending on whether the simulation results violate operational limits, a bid is moved to either (8) the operation-validated queue or the operation-rejected queue of its corresponding DA-Market Data Warehouse instance.

F. The `T-Delivery Optimization` sub-module is effectively unchanged from the Walk and Jog phases. It still waits for the peak load data from all the DOs. Upon receipt from all of them (2a), it solves an optimal black-start restoration problem and issues (3) to the WMO the proposed cranking path.

G. The `M-Delivery Optimization` sub-module is effectively unchanged from the Jog phase. It still waits for the proposed cranking path from the TO. Upon receipt (3), the sub-module sends back (4a) to the TO the updated cranking path. The TO then issues (4b) to each DO its extracted time series of load supplied, and DO #1 in turn forwards (4c) the time series it receives to DMO #1.

H. The `D-Delivery Optimization` sub-module is quite different from how it was in the Jog phase. Its main purpose now is to delegate DER control to four DAMOs, each operating a local "market" to orchestrate DER contributions in its area sub-circuit. On startup, the sub-module waits for the cranking schedule and constraints from DO #1. Upon receipt (4c), the sub-module issues coordination specifications for LinDistFlow-based distributed optimization to each DAMO. In particular, each DAMO is informed of its immediately neighboring areas—at most one upstream but potentially multiple downstream. For each immediately downstream area, a DAMO is moreover informed of the boundary bus with that area and provided with initial estimates of the phase-wise MW and Mvar values of total peak load of the area and any areas farther downstream from it.

I. The Delivery Optimization sub-module instance belonging to a DAMO is newly added with the designation of `DA-Delivery Optimization`. The sub-module instance for an area operates a local "market" to orchestrate DER contributions to blackout restoration of the area sub-circuit. Unlike in a typical electricity market, the aggregators do not concern themselves with competitive pricing because compensation is pre-determined by contract. Aside from considering a mere sub-circuit, a major distinction of `DA-Delivery Optimization` from `D-Delivery Optimization` is the need to solve a series of LinDistFlow-based OPF subproblems rather than just a single problem.

On startup, each instance of `DA-Delivery Optimization` waits for coordination specifications from DMO #1. Upon receipt, the instance requests (5) DER participation in blackout restoration (not black-start, which concerns cranking paths for energization of the transmission network) by posting a service opportunity in the service/ queue of its area's D-Market Data Warehouse. Indicating blackout restoration as its purpose, this service opportunity contains such details as the start time, the duration, and some characterization of permitted locations (e.g., buses) for DERs participating in the opportunity. Clearly, a participating DER must reside at a location within the area sub-circuit of the DAMO.

The `DA-Delivery Optimization` instance then waits for some time. Once that period has elapsed, if the instance finds any bids (8) in the operation-validated queue of its area's D-Market Data Warehouse, it reads all these bids that have passed both market and operational validation. A macro-iteration index is initialized to $k = 0$ and, until convergence of subproblems for its area within the distributed optimization is detected, the `DA-Delivery Optimization` instance performs the following steps so that the instances collectively accomplish distributed optimization (Sadnan and Dubey 2021):

a. The DAMO sub-module instance solves the Iteration k subproblem for its area. The subproblem permits each DER participating in the service opportunity of the DAMO to inject up to its available capacity. If the area is the farthest upstream, then the target total real power across the phases at the substation is given by the scheduled load supplied by the transmission network to the distribution circuit at large of DMO #1. If the area is not the farthest upstream, then the target total real power at the head bus of the area sub-circuit is given by the sum of total net load (load minus DER capacity) within the area itself and the present estimates the area has for total net loads of areas downstream from it. Other assumptions of the subproblem differ depending on whether it is for the first macro-iteration (i.e., $k = 0$):

    i. Iteration $k = 0$: If the area is the farthest upstream, then the per-unit voltage magnitude on each phase at the substation bus is fixed to the regulated value (e.g., 1.05 p.u.). Otherwise, the per-unit voltage magnitude on each phase at the head bus of the area is fixed to a naive value of 1.00 p.u. Regardless of whether an area is the farthest upstream, the MW and Mvar consumption on each phase at each boundary bus the area has with an immediately downstream area are fixed to their initial estimates (total peak load without subtracting any DER capacity)

furnished in the coordination specifications that DMO #1 issued to the DAMO.

ii. Iteration $k > 0$: If the area is the farthest upstream, then the per-unit voltage magnitude on each phase at the substation bus is still fixed to the regulated value. Otherwise, the per-unit voltage magnitude on each phase at the head bus of the area is fixed to the estimate most recently shared by the DAMO of the immediately upstream area. Regardless of whether an area is the farthest upstream, the MW and Mvar consumption on each phase at each boundary bus with an immediately downstream area are fixed to the estimates (essentially total net load that subtracts DER capacity) most recently shared by the DAMO of that area.

b. From the solution values of the Iteration k subproblem, the DAMO sub-module instance extracts the MW and Mvar flows on each phase at the head bus of the area sub-circuit. The instance shares these phase-wise flows with the immediately upstream DAMO, which subsequently interprets them as updated estimates of the MW and Mvar consumption at its boundary bus with the area of the DAMO module instance. Likewise, from the solution values of the Iteration k subproblem, the instance extracts the per-unit voltage magnitude on each phase at each boundary bus with an immediately downstream area. The instance shares these phase-wise voltage magnitudes with the corresponding immediately downstream DAMO, which subsequently interprets them as updated estimates of the per-unit voltage magnitudes at its head bus.

c. For each boundary value shared in the present macro-iteration with an immediately neighboring DAMO, determine whether—based on the kind of value—the value differs by more than some threshold from the value of the same kind shared in the previous macro-iteration. If no boundary value shared with any immediately neighboring DAMO differs by more than its intended threshold, then the DAMO sub-module instance deems its subproblems to have converged, and it will merely proceed to issue the same values to immediately neighboring DAMOs in subsequent iterations.

d. Otherwise, if the DAMO sub-module instance does not detect convergence of its subproblems, then it waits to receive the updated estimates of boundary values from its immediately neighboring DAMOs. Upon receipt of the updated estimates, the sub-module instance advances the macro-iteration index k for its area by one, and its execution returns to step a.

Upon convergence of its area subproblems, the DAMO sub-module instance takes each bid from the operation-validated queue of the DA-Market Data Warehouse, appends to the bid an MW amount to reflect what portion was actually scheduled by the solution to the final subproblem, and moves (9) the bid to the award/ queue. The instance also issues (9) notifications to the aggregators that had submitted the validated bids.

4.3 Summary of Case Studies Validation

During the modeling and validation process, all relevant RCA compartments were represented together within the testbed to evaluate whether the architecture could function as an integrated coordination structure rather than as separate case-specific implementations. This integrated view included the relevant actors, shared databases, coordination modules, market functions, operational validation steps, and information exchanges defined by the RCA.

As shown in Figure 18, the initial configuration represents a wholesale-market-focused coordination structure, where transmission operations, distribution operations, market participants, and the MO interact through defined RCA functions such as T&D Node Visibility, Bid Crafting, Market Validation, Operational Validation, and Delivery Optimization. This configuration supports the lower-maturity implementation of the testbed, where the primary market interaction occurs through the wholesale market and coordination remains more centralized.

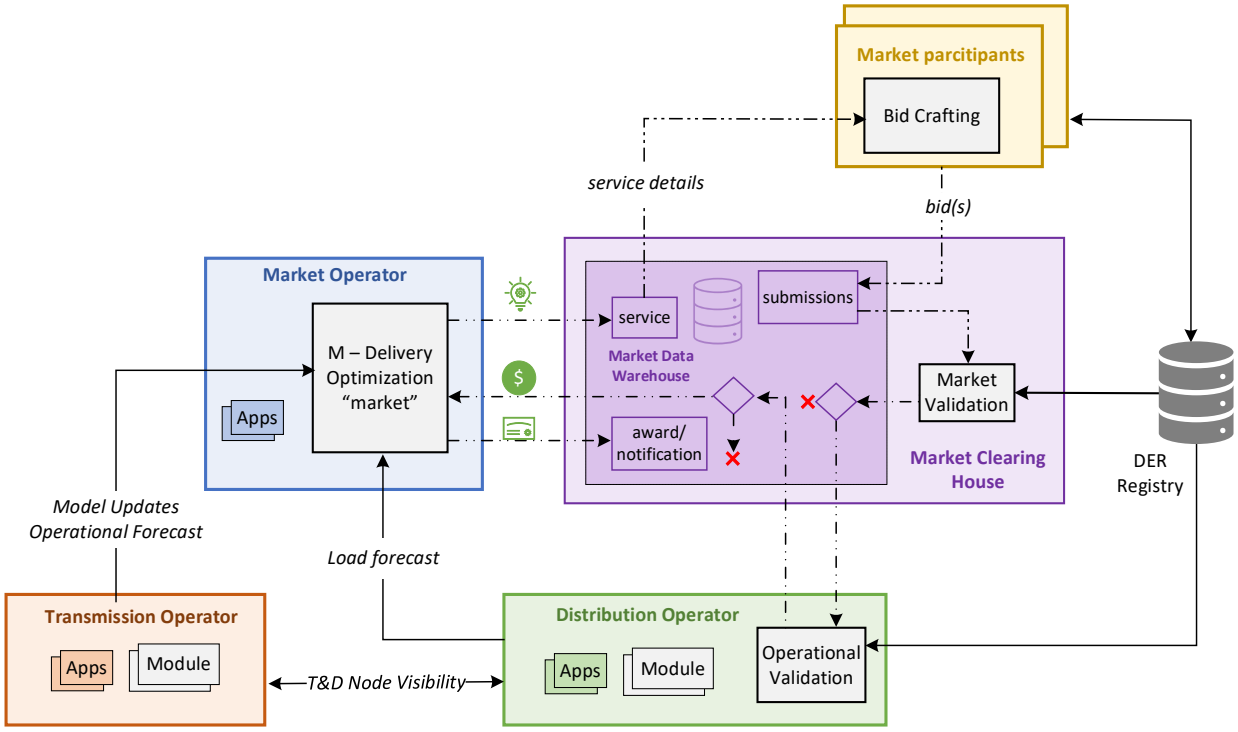


Figure 18. Wholesale Market and Centralized Coordination

Figure 19 illustrates the evolution of the testbed toward a more advanced configuration that includes distribution market functions (in addition to wholesale) and distributed operational coordination. In this configuration, the RCA compartments are expanded to include both wholesale and distribution MOs, distribution-level market clearing functions, bid and service exchanges at multiple levels, and greater interaction among TOs, DOs, MOs, DER Registry, and market participants. This progression reflects the RCA *Walk-Jog-Run* approach, where successive phases introduce greater situational awareness, richer market interactions, expanded DER participation, and more distributed control.

Across the LDDLs, Extreme Cold Weather Event, and Cascading Blackout case studies, the same RCA structure was able to support different data needs, models, workflows, and operational objectives. The ability to model all *Walk-Jog-Run* phases across the testbed provides evidence that the RCA is scalable, modular, and sufficiently generic to support a range of grid coordination scenarios. The validation effort also showed that RCA modules can be composed and expanded as needed, allowing the architecture to transition from a wholesale-market-centric current state toward a future state with distribution markets, distributed operations, and more automated coordination capabilities. This supports the broader RCA objective of enabling coordinated market and operational activities among TOs, DOs, MOs, customers, and aggregators while maintaining clear actor responsibilities and interfaces.

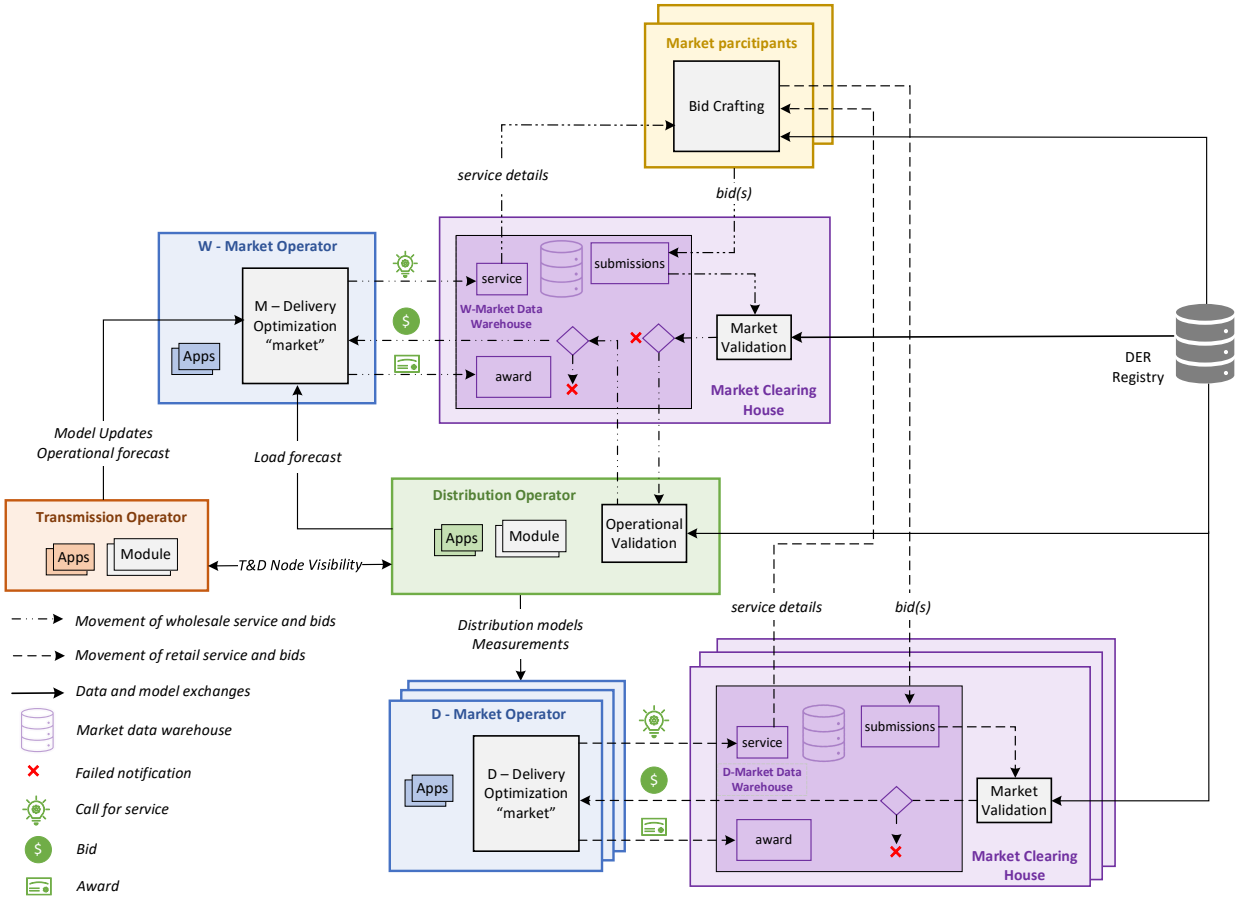


Figure 19. Distribution Market and Distributed Coordination

5.0 Conclusions

This document provides a modeling guide for translating the proposed RCA into practical workflows, actor interactions, data exchanges, and validation processes that can be adapted to different user needs. The RCA supports both the current state and future state of grid operations by allowing users to begin with today's realities—limited visibility, actor-specific data ownership, centralized coordination, and emerging DER participation—and progressively evolve toward more distributed, high-observable, and market-integrated operations. The three case studies, LDDLs, Extreme Cold Weather Event, and Cascading Blackout, provide evidence that the architecture is scalable and generic enough to represent different operating conditions and levels of system complexity.

The *Walk-Jog-Run* progression provides a practical transition path for how operations and markets can adapt over time. In the Walk phase, the RCA supports foundational coordination among DER, market, and operational functions. In the Jog phase, it enables expanded DER participation, distribution-level market functions, improved visibility, and more flexible coordination. In the Run phase, it supports highly localized services, distributed optimization, peer-to-peer coordination, and near-real-time operational adjustments. However, realizing this future state will require complementary technologies, including advanced sensing, communications, forecasting, cybersecurity, interoperability standards, DER controls, data platforms, and distributed computing. The RCA provides the coordination structure, while these enabling technologies support its practical implementation and long-term scalability.

6.0 References

Touhiduzzaman, M. d., S. Poudel, J. Kolln, A. B. Nassif, D. J. Sebastian Cardenas, O. Vasios, and C. Klasen. 2026. Reference Coordination Architecture. Pacific Northwest National Laboratory, PNNL-37902.

Touhiduzzaman, M. d., T. Wall, A. B. Nassif, S. Poudel, J. Kolln, and C. Klasen. 2026. Case Studies Applying the Reference Coordination Architecture. Pacific Northwest National Laboratory, PNNL-38459.

Maloney, P., M. Zhao, X. Ke, P. R. Mendez-Curbelo, R. Sadnan, A. Nassif, K. P. Guddanit, K. Mahapatra, O. Vasios, X. Li, M. Abdelmalak, F. B. dos Reis, J. C. Bedoya Ceballos, V. Chalishazar, X. Fan, and M. Elizondo. 2025. *Recovery Simulator and Analysis Formulation: Mathematical Framework for Enhanced Resilience and Resource Allocation*. PNNL-38549. Richland, WA: Pacific Northwest National Laboratory.

Sadnan, R., and A. Dubey. 2021. “Distributed optimization using reduced network equivalents for radial power distribution systems.” *IEEE Trans. Power Syst.* 36 (4): 3645–3656.

Yip, J. J., R. Sadnan, S. Poudel, A. B. Nassif, L. T. Wall, J.T. Kolln, and C. P. Klasen. 2026. "Reference Coordination Architecture for Market and Operations Coordination: A Blackout Case Study." In *2027 IEEE PES Grid Edge Technologies Conference & Exposition*. PNNL-SA-225106.



Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99354

1-888-375-PNNL (7665)

www.pnnl.gov

Pacific Northwest
NATIONAL LABORATORY

U.S. DEPARTMENT of ENERGY