



Dynamic stability twin for monitoring and control of TSO-DSO operations

Deliverable No: D8.4
Work package: WP8

Official delivery date: 30.06.2026
Actual delivery date: 01.07.2026 (v1), 25.09.2026 (v2)

Dissemination level: Public



Co-funded by
the European Union

Project: 101136119 | HORIZON-CL5-2023-D3-01 | www.twineu.net

Disclaimer

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

Document Authors	
Jonathan Avilés Cedeño (TUD)	José Rueda Torres (TUD)
Peter Palensky (TUD)	Alexandru Stefanov (TUD)

Version	Date	Author(s)	Notes
0.1	2025-04-30	Jonathan Avilés Cedeño / José Rueda Torres	First draft
1.0	2026-06-30	Jonathan Avilés Cedeño / José Rueda Torres	First release of the deliverable
2.0	2026-09-25	Jonathan Avilés Cedeño / José Rueda Torres	Revised release after consortium review

Responsible Partner	Delft University of Technology (TU Delft)
Checked by WP leader	Esteban Pastor Calatayud (ETRA), 01.07.2026 (v1) 24.09.2026 (v2)
Verified by the appointed Reviewers	Daniele Porcu (ENEL), 16.09.2026
Approved by Project Coordinator	Padraic McKeever (Fraunhofer), 01.07.2026 (v1) 25.09.2026 (v2)

Table of Contents

Executive Summary	9
1 Introduction	11
1.1 Scope of Task 8.2	11
1.2 Structure of the Deliverable	11
2 Use Case NL-02: DT-based dynamic stability assessment under power flow changing events	12
2.1 Case description	12
2.2 Use-case KPIs for the demonstrator	14
2.3 Use case conditions	15
2.4 Use case diagrams	16
2.5 Actors/Systems	17
2.6 Overview of the use case scenarios	17
2.7 Methodology for model development	18
3 Development of the Onshore Dutch Power System Model	20
3.1 Model Characteristics	20
3.2 Network Topology	21
3.3 Multi-Energy Scenario Analysis	21
3.4 Case Study Results	24
3.4.1 Case A - Only Wind Generation	24
3.4.2 Case B - Only Solar Generation	24
3.4.3 Case C - Mix of Solar and Wind	25
4 Federated State Estimation for Integrated Transmission–Distribution Networks	26
4.1 Role of State Estimation in the Digital-Twin Workflow	26
4.2 Extension of the Dutch Model to Distribution Networks	27
4.3 Scenario Generation and Measurement Datasets	29
4.4 Federated State Estimation Framework	30
4.5 State Estimation Results	32
4.6 Connection to EMT Model Initialization	34
5 Real-Time EMT Model Development	35
5.1 Randstad Region’s Network	35
5.2 Model Development in RSCAD	36
5.2.1 Power and Control Connections of a Wind Turbine in RSCAD	36
5.2.2 Small Time-Step Block (Small dt Block)	36

5.2.3	Grid VSC Controls	37
5.2.4	Permanent Magnet Synchronous Machine (PMSM) VSC Controls.....	37
5.2.5	Wind Turbine and Measuring Devices	38
5.2.6	Equivalent Models of Neighboring Networks	38
5.3	Simulation Results using the Real-Time Digital Simulator (RTDS).....	39
5.3.1	Real-Time Simulation Performance	39
5.3.2	Signal Selection	39
5.3.3	Steady-State Simulation	39
5.3.4	Behaviour under Three-Phase Short-Circuit Fault	40
6	National-Scale EMT Stability Analysis under High-RES Scenarios	43
6.1	Extension to the Full National Transmission Network	43
6.2	Renewable Penetration Scenarios and Disturbance Scenario	44
6.3	Stability Indicator Extraction Methodology	45
6.4	Simulation Results	47
6.4.1	Voltage Response	47
6.4.2	Frequency and RoCoF Response	50
6.4.3	Summary of Stability Indicator Results	52
7	Conclusions.....	53
8	References	54

List of Figures

Figure 1. Integrated scheme of the Task 8.2 digital-twin demonstrator for stability assessment.	16
Figure 2. Randstad region's transmission network single-line diagram (SLD).....	16
Figure 3. Synthetic model development flowchart for the Dutch EHV network of 2030.....	20
Figure 4. Single-line diagram of the synthetic digital model of the Dutch EHV system.	22
Figure 5. Illustration of the electrolyzer locations. Grid map taken from TenneT's website.	23
Figure 6. Geographic diagram of contingency analysis for n-1 security.....	25
Figure 7. Single-line diagram of the integrated transmission–distribution system.....	28
Figure 8. Single-line diagram of the synthetic distribution networks.....	29
Figure 9. Composition of input features and predicted voltage states for the residual MLP estimator.	31
Figure 10. Validation RMSE of (a) the centralized and (b) the federated estimators, in per unit (p.u.).	33
Figure 11. Test performance of the centralized and federated estimators, expressed as RMSE and MAE in per unit (p.u.).	34
Figure 12. Randstad region's electrical network. Adapted from TenneT's website.....	35
Figure 13. Power and control connections of a wind farm in RSCAD.....	36
Figure 14. Small time-step block (small dt block).	37
Figure 15. PMSM current controller.	38
Figure 16. Wind turbine aerodynamic model.....	38
Figure 17. Power output of generators connected to Substation 6 during normal start-up to reach steady state in RSCAD.	39
Figure 18. Power output of Hollandse Kust wind farms (1-Alpha and 2-Beta) and BritNed during normal start-up to reach steady-state in RSCAD.....	40
Figure 19. Behaviour of generators connected to Substation 6. Short-circuit fault.....	41
Figure 20. Behaviour of Hollandse Kust Noord wind farm. Short-circuit fault.	41
Figure 21. Behaviour of Hollandse Kust wind farms (1-Alpha and 2-Beta). Short-circuit fault.	42
Figure 22. Annotated view of the RSCAD implementation around Substation 1, showing the main busbar selected for the three-phase short-circuit fault and the modelled elements, including synchronous machines, transformers, and dynamic PQ sources for offshore wind and HVDC infeeds. Internal labels are anonymized for confidentiality.....	44
Figure 23. Phase-to-neutral voltages at Substation 1 for the three RES penetration scenarios.	48
Figure 24. Positive-sequence RMS voltage at Substation 1 for the three RES penetration scenarios.	49
Figure 25. Instantaneous frequency at Substation 1 for the three RES penetration scenarios.	51
Figure 26. Rate of change of frequency (RoCoF) at Substation 1 for the three RES penetration scenarios.	52

List of Tables

Table 1. Summary of demonstrator-level key performance indicators (KPIs).....	14
Table 2. Summary of case scenarios with solar/wind capacity percentages and presence of conventional generation, interconnections, and electrolyzer demand.	24
Table 3. Description of cases. Active power generation, demand, electrolyzer usage, and power surplus.....	24
Table 4. Federated learning hyperparameters for each strategy.....	32
Table 5. Stability indicators per scenario.....	52

List of Abbreviations and Acronyms

Acronym	Meaning
AC	Alternating Current
CRoF	Control Room of the Future
DC	Direct Current
DER	Distributed Energy Resources
DFT	Discrete Fourier Transform
DSO	Distribution System Operator
DT	Digital Twin
EHV	Extra-High Voltage
EMT	Electromagnetic Transient
FL	Federated Learning
HV	High Voltage
HVDC	High-Voltage Direct Current
KPI	Key Performance Indicator
MAE	Mean Absolute Error
MLP	Multilayer Perceptron
MPPT	Maximum Power Point Tracking
PI	Proportional-Integral
PLL	Phase-Locked Loop
PMSM	Permanent Magnet Synchronous Machine
PV	Photovoltaic
PWM	Pulse-Width Modulation
P2x	Power-to-X
RES	Renewable Energy Sources
RMS	Root Mean Square
RMSE	Root Mean Square Error
RoCoF	Rate of Change of Frequency
RSCAD	Real-Time Simulation Software (RTDS companion software)
RTDS	Real-Time Digital Simulator
SCADA	Supervisory Control and Data Acquisition
SE	State Estimation
SGD	Stochastic Gradient Descent

SLD	Single-Line Diagram
T-D	Transmission–Distribution
TSO	Transmission System Operator
UC	Use Case
VRES	Variable Renewable Energy Sources
VSC	Voltage Source Converter
WLS	Weighted Least Squares

Executive Summary

This deliverable presents the final status of the Dutch pilot demonstrator for use case NL-02, “DT-based dynamic stability assessment under power flow changing events,” developed within TwinEU Task T8.2, “Digital twin for stability of power systems with high penetration of renewable energy sources.” This work aimed to develop and assess a digital-twin demonstrator for examining network constraints and post-fault dynamic behaviour of the Dutch power system under future high-renewable operating conditions.

The work comprises four complementary developments organized around a common modeling and stability-assessment workflow. First, a steady-state synthetic model of the Dutch extra-high-voltage (EHV) transmission network was built in DigSILENT PowerFactory from publicly available data and used to screen 2030 multi-energy scenarios combining variable renewable energy sources (VRES) and electrolyzers. The analysis showed that transformer capacity between the HV and EHV networks, rather than generation headroom, is the binding constraint for accommodating high renewable shares: solar-dominated scenarios with limited local consumption relative to solar production, despite high total national demand, drove voltages down to about 0.9 p.u. in certain regions, and contingency screening left several EHV/HV transformers loaded to 100%. Electrolyzer siting materially relieved these constraints: an even 6.5 GW distribution across three locations was most effective in the wind-only case, while only specific placements kept all component loadings below 70% in the mixed solar-wind case.

Second, the steady-state model was extended to the distribution level with three synthetic distribution networks and used to generate 30,000 feasible operating points for a federated state estimator. The purpose of this development is to reconstruct consistent operating conditions across both voltage levels from measurements that remain local to each operator, so that no centralized pooling of operator data is required. A residual multilayer perceptron was trained across one TSO and three DSO clients, and the best-performing federated strategy, FedAdam, reached a test RMSE of 0.0158 p.u. and an MAE of 0.0104 p.u. on the predicted voltage components, against 0.0033 p.u. and 0.0023 p.u. for a centralized benchmark with full data access. These errors correspond to voltage magnitude deviations of around 1 to 2%, comparable to measurement uncertainty in operational systems.

Third, a real-time electromagnetic transient (EMT) model of the Randstad region, including a detailed Type-4 wind-turbine representation scaled to 700 MW per farm for the Hollandse Kust wind farms, was developed in RSCAD and run on the Real-Time Digital Simulator (RTDS) at TU Delft. The model sustained continuous real-time execution at a 50 μ s network time step (1.4-2.5 μ s for converter components) across three NovaCor racks with no latency, establishing that EMT-level fidelity for this part of the EHV network is achievable in real time. The fault studies proved more revealing than the steady-state checks: wind farms near the fault rode through with reactive-power support and recovered once it cleared, but a three-phase fault near the Hollandse Kust Zuid connection produced active- and reactive-power oscillations that did not settle, indicating that converter-dominated corridors may require coordinated synchronous/converter control and fast protection to maintain stable post-fault behaviour.

Fourth, the real-time EMT platform was extended to the full national Dutch transmission network to enable stability assessment under projected 2030 renewable penetration levels of 50%, 65%, and 80%. Voltage- and frequency-based stability indicators were extracted following a three-phase short-circuit fault and analyzed systematically. The maximum rate of change of frequency (RoCoF) emerged

as the most penetration-sensitive indicator, increasing from 2.04 Hz/s to 14.54 Hz/s across the three scenarios, consistent with the expected loss of synchronous inertia. The frequency nadir and post-fault maximum also degraded monotonically, whereas voltage indicators were less sensitive to penetration level, with sag durations dominated by the imposed fault clearance time and both recovery times peaking in the intermediate 65% case.

The central contribution is a real-time EMT digital-twin core that extends from a regional to a national-scale high-voltage model while preserving real-time execution, giving a reusable basis for stability screening of high-renewable scenarios. Around that core, the steady-state model supplies consistent operating points, and the federated estimator provides a path for reconstructing conditions across transmission and distribution networks from operator-local measurements. Together, these developments bring network modeling, operating-point reconstruction, dynamic simulation, and indicator extraction into a coherent workflow for assessing stability in power systems. The findings show why high-renewable operation must be examined through both network constraints and disturbance response, where the effects of declining synchronous inertia become particularly apparent.

1 Introduction

The TwinEU Dutch pilot for stability studies is implemented in Task T8.2 “Digital twin for stability of power systems with high penetration of renewable energy sources.”

This report presents the developments of use case NL-02, “DT-based dynamic stability assessment under power flow changing events.”

1.1 Scope of Task 8.2

Task T8.2 is led by Delft University of Technology, in collaboration with TenneT, Alliander, Stedin, and Jedlix, and runs from M3 to M33 (30 months).

The task demonstrates how digital replicas [6] can support in-depth stability analyses, covering frequency behaviour, voltage integrity, and overall system resilience, under evolving operational conditions and high shares of inverter-based generation [17]. The terminology used in this deliverable follows ISO/IEC 30173:2023, which provides a standardized conceptual and terminological basis for digital-twin systems [14].

The Dutch pilot demonstrator combines two complementary components: a synthetic data generator and a real-time simulation core for digital-twin-based stability assessment. The synthetic data generator consists of a steady-state model that produces responses resembling data from actual power systems, an approach adopted in response to industry partners’ confidentiality requirements and limitations on data sharing. The same steady-state model generates the datasets used to develop the federated state estimator presented in Section 4. This capability provides a route from distributed measurements to the operating conditions required for dynamic simulation.

For this task, a digital twin demonstrator of the Dutch transmission network has been implemented using real-time simulation hardware from RTDS Technologies and its companion software RSCAD [24], both hosted at the Control Room of the Future (CRoF) at TU Delft. The demonstrator reported in this deliverable consolidates and extends modeling and simulation work previously reported by the authors in conference publications [3], [4], together with the national-scale EMT stability assessment presented at IFAC 2026 [5]. The deliverable itself provides the TwinEU Task T8.2 integration of these developments into a DT-based stability-assessment demonstrator. The federated state-estimation framework adapted in Section 4 is reported in a companion publication [1].

1.2 Structure of the Deliverable

The remainder of this report is organized as follows. Section 2 provides a detailed description of use case NL-02 following the IEC 62559-2 template. Section 3 presents the development and results of the steady-state synthetic digital model of the Dutch EHV network used for multi-energy scenario analysis. Section 4 introduces the role of state estimation in the digital-twin workflow, extends the steady-state model to the distribution level, and reports the federated state-estimation framework together with its results. Section 5 describes the development and evaluation of the real-time EMT digital model for the Randstad region. Section 6 describes the extension of the real-time EMT model to the full national transmission network and the stability-indicator-based assessment under high-RES penetration scenarios. Section 7 provides conclusions and summarizes how the demonstrator illustrates the use of future DT-based environments for stability monitoring and assessment.

2 Use Case NL-02: DT-based dynamic stability assessment under power flow changing events

2.1 Case description

The increasing complexity of modern cyber-physical power systems, with the high integration of renewable energy sources and dynamic loads, necessitates advanced tools for stability assessment [2]. This use case focuses on demonstrating the use of Digital Twin (DT) technology to assess the stability of power systems under power-flow-changing conditions associated with high renewable penetration, changing demand patterns, and transmission-network disturbances [9]. In this deliverable, these conditions are represented through synthetic steady-state scenarios, real-time EMT simulation, and fault-based voltage and frequency stability assessment. The increasing penetration of renewable generation is also aligned with broader climate-policy commitments and national decarbonization objectives [7], [18].

Power-system stability is understood as the ability of a power system, for a given initial operating condition, to regain a state of operating equilibrium after being subjected to a physical disturbance, with most system variables bounded so that practically the entire system remains intact [10]. Stability is not a single property but a set of interacting phenomena, conventionally grouped into rotor-angle, voltage, frequency, and, more recently, converter-driven and resonance-related dynamics, each associated with distinct physical mechanisms and characteristic time frames [10]. The work reported in this deliverable concentrates on the dynamic response that follows a transmission-network disturbance, where several of these phenomena appear together in the measured quantities: the depression of voltage during a fault and its subsequent recovery, the excursion of system frequency away from its nominal value, and the active- and reactive-power oscillations exchanged between generating units and the network as the system settles.

The character of that response is shaped by the interaction between synchronous machines, converter-interfaced generation, and the prevailing network conditions. Synchronous machines contribute to rotating inertia and sustained fault current, and their electromechanical dynamics govern the initial exchange of power after a disturbance. Converter-interfaced units respond instead through their control systems and within the current limits of the power electronics, so that their contribution to voltage support and power balance is determined largely by control design and settings rather than by machine physics. Because these behaviours combine over short time frames, real-time EMT simulation offers a detailed environment in which the resulting interactions can be examined for selected operating conditions and disturbances [4], [24]. On that basis, this demonstrator compares simulated responses across operating conditions and extracts voltage- and frequency-based indicators that characterize post-fault dynamic behaviour.

Properly managing power-flow changes is crucial for the seamless interaction between the distribution and transmission networks. Effective coordination between Distribution System Operators (DSOs) and Transmission System Operators (TSOs) supports the reliability and stability of the entire power system, especially during disturbance events that affect voltage and frequency behaviour. This coordination depends on a shared view of the operating state across both voltage levels, which in practice must be reconstructed from measurements that remain under separate ownership. The use case, therefore, also covers the reconstruction of network operating states under TSO–DSO data separation, so that the operating condition used for dynamic assessment reflects distribution-level conditions without requiring operators to exchange raw measurement data.

The digital models are configured to reproduce relevant static and dynamic behaviour of the power system using simulation-generated signals, accessible published databases, public technical reports, and representative operating assumptions.

Critical aspects of the study include:

- **Model configuration and consistency checks:** Using simulation-generated signals, publicly available data, technical reports, and representative assumptions to configure the digital models and verify that their behaviour is numerically consistent with expected Dutch transmission-system behaviour.
- **Scenario analysis:** Simulating representative high-renewable operating conditions and disturbance scenarios to study their impact on system stability, including frequency deviations, voltage response, post-fault recovery, and power oscillations.
- **State estimation:** Developing a federated estimator that reconstructs transmission and distribution voltage states from operator-local measurements and could support the preparation of operating conditions for dynamic assessment.
- **Real-time EMT simulation:** Demonstrating the feasibility of using RTDS/RSCAD as the real-time simulation core of the DT demonstrator, first for the Randstad region and subsequently for the national-scale Dutch transmission network.
- **Stability assessment:** Extracting voltage- and frequency-based indicators from simulated waveforms to evaluate the dynamic response of the system under high-renewable operating conditions and representative transmission-network faults.

Power-flow-changing events considered in this demonstrator include:

- **High-renewable operating scenarios:** Different wind and solar generation configurations used to represent projected 2030 operating conditions.
- **Power-to-X demand placement:** Different electrolyzer demand distributions near offshore wind connection areas, used to assess their impact on power flows, loading, and voltage profiles.
- **Transmission-network contingency conditions:** Network configurations and contingency cases used to evaluate the behaviour of the synthetic steady-state model.
- **Regional three-phase short-circuit faults:** Fault cases applied in the Randstad EMT model to evaluate real-time dynamic behaviour of synchronous generation, offshore wind farms, and DC interconnection responses.
- **National high-RES fault scenarios:** A common three-phase short-circuit fault applied under 50%, 65%, and 80% renewable penetration to compare voltage recovery, frequency response, and RoCoF behaviour.

The objective of the use case is to demonstrate how DT-based environments can support stability assessment of the Dutch transmission system under representative high-renewable and disturbance conditions. The demonstrator provides a controlled simulation environment for analyzing power-flow patterns, voltage behaviour, frequency behaviour, and post-fault recovery, thereby illustrating how future digital twins can support stability screening, scenario comparison, and stability-oriented decision-making.

2.2 Use-case KPIs for the demonstrator

A set of use-case key performance indicators (KPIs), summarized in Table 1, has been established to describe the assessment framework for NL-02 within the present demonstrator. These KPIs are defined at the demonstrator level and are aligned with the implemented workflow: steady-state scenario analysis, federated state estimation across the transmission and distribution levels, regional real-time EMT simulation, and national-scale EMT stability assessment. They quantify the capability of the demonstrator to reproduce relevant operating conditions, execute representative disturbances, and extract voltage- and frequency-based stability information from the simulated response.

Table 1. Summary of demonstrator-level key performance indicators (KPIs).

ID	Name	Description	Reference to use case objectives
KPI-01	Scenario Analysis Coverage Index	Number and scope of representative operating conditions and disturbance scenarios implemented in the demonstrator, including steady-state high-RES scenarios, integrated transmission–distribution datasets for state estimation, regional EMT tests, and national high-RES EMT scenarios.	Supports the objective of demonstrating scenario-based stability assessment under future operating conditions.
KPI-02	Voltage Response Index	Voltage deviation, voltage sag duration, and voltage recovery behaviour extracted from simulated disturbance responses.	Supports the objective of assessing voltage behaviour and post-fault recovery under representative disturbance conditions.
KPI-03	Frequency Response Index	Frequency nadir, post-fault maximum frequency, and frequency recovery behaviour extracted from simulated disturbance responses.	Supports the objective of assessing frequency behaviour under high-renewable operating conditions.
KPI-04	RoCoF Sensitivity Index	Maximum absolute RoCoF extracted from the simulated frequency response under different renewable penetration scenarios.	Supports the objective of evaluating the sensitivity of system dynamics to increasing inverter-based generation.
KPI-05	Real-Time Simulation Capability Index	Demonstration of stable real-time execution of the EMT models using RTDS/RSCAD for regional and national-scale stability-assessment workflows.	Supports the objective of demonstrating real-time digital-twin capabilities for stability studies.
KPI-06	State-Estimation Accuracy Index	Root-mean-squared error and mean absolute error of the estimated busbar voltage components, reported in per unit, obtained from the federated estimator under TSO–DSO data separation and compared against a centralized benchmark trained with full data access.	Supports the objective of reconstructing consistent network operating states across voltage levels without centralized data sharing.

2.3 Use case conditions

The Digital Twin (DT) demonstrator for voltage and frequency stability assessment is based on specific assumptions and prerequisites that define the scope of the implemented simulation environment. These conditions support the generation of relevant insights into the dynamic behaviour of the Dutch transmission system under representative high-renewable and disturbance scenarios.

Assumptions:

1. The demonstrator uses synthetic and publicly derived models to represent relevant characteristics of the Dutch transmission system, in line with confidentiality and data-access constraints. PowerFactory is used for steady-state scenario analysis, while RTDS/RSCAD is used for real-time EMT stability studies.
2. The data used to configure and test the demonstrator are obtained from simulation-generated signals, accessible published databases, public technical reports, and representative assumptions consistent with the studied operating conditions.
3. The demonstrator focuses on the Dutch transmission system and uses equivalent representations where detailed external or neighboring systems are not explicitly modelled.
4. Communication between the simulation components and analysis workflows is assumed to be reliable within the controlled demonstrator environment, enabling consistent execution and post-processing of the stability studies.

Prerequisites:

1. Secure access to the required PowerFactory and RSCAD/RTDS environments, ensuring proper configuration and licensing for steady-state scenario analysis and real-time EMT simulations.
2. Proper setup and configuration of the tools required for the demonstrator environment, including PowerFactory, RTDS/RSCAD, data-processing scripts, and post-processing routines for stability indicator extraction.
3. Availability of the physical and digital infrastructure of the Control Room of the Future (CRoF) at TU Delft, including the RTDS hardware and supporting computing resources required for real-time EMT simulation.
4. Availability of the required datasets and assumptions for configuring, testing, and analyzing the demonstrator, including public grid information, published technical reports, simulation-generated signals, and scenario definitions.

2.4 Use case diagrams





Figure 1. Integrated scheme of the Task 8.2 digital-twin demonstrator for stability assessment.

Note: This single-line diagram was generated with a larger resolution; thus, it is possible to zoom in to look at component details



2.5 Actors/Systems

The primary actors in this use case are the Transmission System Operator (TSO) and Delft University of Technology (TU Delft). The TSO provides the system operator's perspective needed to define relevant stability-assessment needs and to contextualize the demonstrator results. TU Delft, as the owner of the Control Room of the Future (CRoF), provides the real-time digital simulation capabilities required for the demonstration. Specifically, TU Delft hosts the Real-Time Digital Simulator (RTDS) hardware and companion software RSCAD, where the Dutch digital-twin demonstrator for stability assessment is implemented and evaluated. The federated state-estimation framework of Section 4 represents these operator roles structurally rather than organizationally: one client holds the transmission-level measurements and coordinates the shared model, in the role of the TSO, while three further clients hold the measurements of separate distribution regions, in the role of DSOs. The partitioning reproduces the data boundaries between operators inside the research environment, and no measurement data from actual operators is used, nor is any deployment at an operator implied.

2.6 Overview of the use case scenarios

The following six scenario categories summarize how the NL-02 demonstrator supports dynamic stability assessment of the Dutch transmission system:

Scenario 1. High-RES Steady-State Operating Conditions.

This scenario category represents projected 2030 operating conditions with high wind and solar penetration. The synthetic PowerFactory model is used to assess power-flow patterns, voltage profiles, and component loading under different renewable generation and demand configurations.

Scenario 2. Power-to-X Demand Placement.

This scenario category evaluates the effect of electrolyzer demand placement near offshore wind connection areas. The demonstrator compares different distributions of electrolyzer demand and assesses their impact on network loading, voltage behaviour, and contingency performance.

Scenario 3. Integrated Transmission–Distribution Operating Conditions.

This scenario category represents operating conditions of the Dutch network with the three synthetic distribution networks connected at the 150 kV level. Demand, photovoltaic, and wind allocations are swept across 25 demand bins between 12 GW and 26 GW to produce 30,000 feasible operating points, which are used as the measurement and state datasets for the federated estimator described in Section 4.

Scenario 4. Regional Real-Time EMT Operation.

This scenario category demonstrates real-time EMT simulation of the Randstad transmission network in RSCAD/RTDS. It includes initialization to steady-state operation and monitoring of active and reactive power responses from synchronous generation, offshore wind farms, and the BritNed DC link.

Scenario 5. Regional Three-Phase Fault Response.

This scenario category evaluates the dynamic behaviour of the Randstad EMT model under representative three-phase short-circuit faults. The response of generators and offshore wind farms is analyzed to assess real-time model behaviour under severe disturbances.

Scenario 6. National High-RES EMT Stability Assessment.

This scenario category extends the real-time EMT workflow to the national Dutch transmission network. Three renewable penetration levels, 50%, 65%, and 80%, are compared under a common three-phase short-circuit fault, and voltage- and frequency-based stability indicators are extracted from the simulated waveforms.

2.7 Methodology for model development

A power system can be represented at several levels of detail, and the choice of representation determines which phenomena a study is able to resolve. A steady-state power-flow solution describes a single operating condition: it returns the bus voltage magnitudes and angles, the active and reactive flows on each branch, and the balance between generation and demand once losses are accounted for. It says nothing about how the system arrived at that condition or how it would behave if the condition were disturbed, but it establishes the consistent starting point that any dynamic study requires, and it is efficient enough to be repeated across large numbers of scenarios.

Dynamic studies extend this picture in time. RMS, or phasor-domain, simulation tracks the evolution of system states using a fundamental-frequency representation of the network, in which voltages and currents are described by slowly varying phasors and the network itself is treated as algebraic. This representation is well suited to electromechanical phenomena such as rotor-angle swings and frequency regulation, and it allows long simulation horizons over extensive networks at moderate computational cost. EMT simulation instead reproduces instantaneous electrical waveforms and the dynamic behaviour of network elements and converter controls, with switching represented explicitly where required by the selected model. This enables the examination of fast network and converter interactions beyond the scope of conventional fundamental-frequency representations [34].

Model detail is consequently selected according to the phenomena of interest and the scale of the network under study, since the additional fidelity of an EMT representation carries a correspondingly higher computational cost and data requirement [34]. A separate characteristic is whether the simulation is executed in real time: real-time operation concerns the rate at which a simulation advances relative to wall-clock time and applies to phasor-domain and EMT representations alike, so it should not be read as a property of either modeling approach [24]. The developments reported in this deliverable combine a steady-state representation for operating-condition analysis with a real-time EMT representation for disturbance response, and the following paragraphs describe how the two were built and used.

The initial phase of this work involved modeling the system using DigSILENT PowerFactory to perform power flow calculations as a foundation for steady-state analysis. The digital model was implemented and fine-tuned using publicly available data from credible sources. This model was then extended to include multi-energy scenario analyses, incorporating varying levels of renewable generation, converter-based loads, and system contingencies. The model anticipates future challenges by thoroughly examining the integration of VRES and electrolyzers and provides strategic insights into renewable energy distribution under different operating conditions. The same model serves a second purpose in the workflow: because each solved case is a consistent operating point of the network, sweeping demand and renewable injections produces the labelled datasets required to train and evaluate a data-driven state estimator. Scenario analysis and estimator training, therefore, draw on one common modeling basis.

Building upon this steady-state foundation, the second phase focused on developing a real-time EMT model using RSCAD, starting with the Randstad region. This model supports detailed stability

studies for scenarios involving high levels of renewable penetration and converter-interfaced resources. The workflow was then extended to the national-scale EMT model, where voltage- and frequency-based stability indicators were extracted under high-RES scenarios. Operating-point preparation links these developments: the steady-state model supplies power-flow solutions for dynamic studies, while the federated estimator of Section 4 could support the reconstruction of corresponding operating conditions from operator-local measurements. Together, these stages demonstrate how DT-based environments can support stability assessment of future Dutch transmission-system operating conditions.

3 Development of the Onshore Dutch Power System Model

The purpose of this model is to provide the demonstrator with a consistent, steady-state representation of the Dutch network: it solves power flows for defined operating conditions, screens future multi-energy scenarios, and supplies the operating points used by the subsequent state-estimation and EMT work.

3.1 Model Characteristics

This model has been developed based on data available from public sources [11], [26], [28]. The generation/demand profile is theoretically defined based on TenneT's public technical reports, where specific regional values and VRES installation capacities are extensive [27]. The transmission lines and power transformer parameters were derived from previous research conducted at Delft University of Technology. The steady-state model development and multi-energy scenario analysis summarized in this section build on the work reported in [3].

The transmission system was modelled in PowerFactory, and Python scripting was used to automate data reading and perform scenario simulations. A flowchart illustrating the complete process of developing the synthetic digital model is shown in Figure 3.




Figure 3. Synthetic model development flowchart for the Dutch EHV network of 2030

After constructing the base model, several verification steps were conducted. First, the model's results were compared to system performance data and studies available from TenneT's technical reports to verify consistency with historical and expected system behaviours. Second, the parameters of transmission lines, transformers, and generation capacities were cross-referenced with values reported in technical documents from different public data sources. Although real-time operational data was unavailable, these comparisons ensured that the synthetic model aligned with the expected

performance of the Dutch EHV grid. Adjustments were made where discrepancies were found, particularly concerning generator settings and regional demand values.

Subsequently, the model was updated with projections for 2030, including adjusted installed capacities and topology configurations. The dispatch configuration and testing for these future scenarios were conducted as part of the case studies described in Sections 3.3 and 3.4.

3.2 Network Topology

The network was expanded in compliance with different ongoing or planned projects mentioned in technical reports available on TenneT's website [27], [29]. Various network configurations were tested, including both typical operating conditions and more extreme scenarios designed to push the system to its operational limits. These tests considered factors such as high demand, solar and wind renewable energy integration, installation of electrolyzers, and potential grid contingencies, allowing us to assess the network's resilience under diverse conditions.

The single-line diagram of the final system for 2030 is given in Figure 4. The example diagram uses colors to indicate voltage levels and loading conditions: green indicates normal operating conditions, light blue to blue indicates increasingly lower voltage ranges, and yellow to red indicates increasingly higher levels of loading and potential overloads. Additionally, highly loaded components (in this case, transformers) are circled in magenta, and a large area with undervoltages of around 0.95 p.u. is coloured in light blue.

3.3 Multi-Energy Scenario Analysis

Simulations were performed for the 2030 scenario, which was selected as a reference because it is based on installed capacities from the Dutch climate agreement [18]. For each scenario, the power flow within the transmission network was analyzed with the introduction of electrolyzers at locations near new offshore wind connections. The chosen locations are illustrated in Figure 5, using TenneT grid-map information [30].

Different distributions of the electrolyzers were assumed for different cases. Different combinations of VRES generation, demand, and electrolyzer distribution are considered for the operational scenarios. In all cases, the Netherlands is assumed to be self-sufficient. This assumption was made as an extreme condition, since it is expected that 70% of the main transmission network transfer capacity in the future should become available for the market [8]. The different cases that were chosen to be investigated are the following:

- **Case A:** Wind energy as the only source of electricity generation for a high demand.
- **Case B:** Solar energy as the only source of electricity generation for a high demand.
- **Case C:** A mix of high wind and some solar generation during peak demand.

Three base cases were created where no electrolyzers are used, and the national system's generation and demand are balanced: A0 for only wind generation, B0 for only solar generation, and C0 for a mix of solar and wind generation. The Dutch transmission system is assumed to be self-sufficient, so interconnections are disabled.




Figure 4. Single-line diagram of the synthetic digital model of the Dutch EHV system.




Figure 5. Illustration of the electrolyzer locations. Grid map taken from TenneT's website.

Subsequently, 3 cases (A, B, and C) were created, each with seven scenarios in which different electrolyzer distributions are considered. This includes cases locating electrolyzers only at one location, two locations, and all three locations. For these scenarios, the generation of VRES exceeds the demand, so electrolyzers are used for Power-to-X (P2x) generation. Table 2 and Table 3 present an overview of the different cases' characteristics.

Table 2. Summary of case scenarios with solar/wind capacity percentages and presence of conventional generation, interconnections, and electrolyzer demand.

Case Type	Synthetic Case Description	Solar (%)	Wind (%)	Conv. Gen.	IC	P2x. Dem.
A0	Base case, only wind, no IC	0	87	N	N	N
B0	Base case, only solar, no IC	65.5	0	N	N	N
C0	Base case, mix solar and wind, no IC	18	90	N	N	N
A1-7	Electrolyzer cases, only wind, no IC	0	100	N	N	Y
B1-7	Electrolyzer cases, only solar, no IC	91	0	N	N	Y
C1-7	Electrolyzer cases, mix solar and wind, no IC	52	100	N	N	Y

Notes: Conv. Gen.: Conventional Generation, IC: Interconnections, P2x. Dem.: Electrolyzer Demand.

Table 3. Description of cases. Active power generation, demand, electrolyzer usage, and power surplus.

Case Type	Solar (GW)	Wind (GW)	Conv. Gen. (GW)	Regional Dem. (GW)	P2x. Dem. (GW)	Import (-) / Export (+) (GW)	Gen. - Dem. (GW)	Active Power Surplus (GW)
A0	0.0	16.2	0.0	16.0	0.0	0.0	0.2	0.2
A1-7	0.0	18.6	0.0	12.0	6.5	0.0	6.6	0.1
B0	16.2	0.0	0.0	16.0	0.0	0.0	0.2	0.2
B1-7	22.5	0.0	0.0	12.0	10.0	0.0	10.5	0.5
C0	4.5	16.7	0.0	21.0	0.0	0.0	0.2	0.2
C1-7	12.9	18.6	0.0	21.0	10.0	0.0	10.5	0.5

Notes: Conv. Gen.: Conventional Generation, Import/Export: Interconnections, P2x. Dem.: Electrolyzer Demand.

3.4 Case Study Results

The synthetic digital model proved to be a valuable tool for testing multi-energy scenarios, with several notable observations arising from the different cases.

3.4.1 Case A - Only Wind Generation

For scenarios A0-7, large electricity transport can be seen from offshore wind farms towards inland connections. Large exports can especially be seen from the regions of Zeeland to Brabant in the Netherlands. As most wind generation is concentrated at the 380 kV and 220 kV levels, considerable power flows were also directed toward the 150 kV and 110 kV regions. Contingency analysis revealed that the most effective distribution of electrolyzers occurs when they are evenly distributed across the three locations, with a total electrolyzer demand of 6.5 GW.

3.4.2 Case B - Only Solar Generation

In the only-solar-generation scenarios B0-7, results differed notably from those in Case A. For the base scenario B0, significant power transport was observed from the northern to the southern regions of the Netherlands due to the higher installed solar generation capacity in the North. A total electrolyzer demand of 10 GW was considered. In these scenarios, the power exported from the HV to the Groningen and Drenthe EHV regions led to undervoltages, reaching approximately 0.9 p.u. This issue arose due to the combination of high solar capacity, limited local consumption relative to solar production despite high total national demand, and insufficient transformer capacity between the HV and EHV sections of the network.

3.4.3 Case C - Mix of Solar and Wind

In this case, the optimal locations for electrolyzers were identified as Locations 1 and 3 (case C5) and Locations 1 and 2 (case C6), as only these two cases displayed no component loading above 70%. Results from the contingency analysis for cases C5 and C6 are shown in Figure 6. In both cases, transformers between the EHV and HV regions were critical for n-1 security in several areas, with some transformers showing loadings up to 100%.




Figure 6. Geographic diagram of contingency analysis for n-1 security

The following section builds on this steady-state foundation by introducing three synthetic distribution networks and using the extended model to generate datasets for federated state estimation.

4 Federated State Estimation for Integrated Transmission–Distribution Networks

4.1 Role of State Estimation in the Digital-Twin Workflow

Dynamic stability assessment always begins with an operating condition. The steady-state model of the Dutch network described in Section 3 supplies that condition: it computes consistent power-flow solutions for the national system, supports the multi-energy scenario analysis used to select representative future conditions, and generates the operating-point datasets on which a data-driven estimator can be trained. State estimation is the step that turns measured quantities into such an operating condition, and it therefore sits between network modeling and dynamic simulation in the demonstrator workflow.

State estimation (SE) is the reconstruction of the electrical state of a network, represented by the voltage magnitudes and angles at every busbar, from a redundant and noisy set of field measurements, so that downstream operational functions all act on a single consistent network solution. In transmission practice, SE is performed by the TSO using a centralized weighted least squares (WLS) formulation. WLS does not require full measurement coverage, but it assumes an observable transmission network and represents each distribution grid as a static net active and reactive power injection at the transmission–distribution (T–D) interface. DSOs run their own estimators independently and exchange information with the TSO only at the interconnection points.

This separation does not hinder transmission-level SE, but it does limit the accuracy of operational studies that rely on boundary sensitivities, or that must predict how the T–D exchange varies under contingencies and redispatch actions, because the internal behaviour of DER-rich feeders, including renewable variability, voltage-dependent loads, and local control actions, remains invisible to the transmission operator. A coordinated TSO–DSO estimator addresses this by learning how internal distribution conditions influence the T–D exchange. A unified model reduces reliance on static equivalents and pseudo-measurements, improves voltage-security assessment through a better representation of reactive-power responses and reverse flows, and enables more coherent congestion and flexibility management across voltage levels.

Machine-learning estimators offer a fast alternative by learning nonlinear mappings between measurements and states. Neural and ensemble models can approach WLS accuracy with lower inference cost and good robustness to noise [12], and physics-informed or graph-based architectures embed structural constraints to improve generalization [32]. Centralized machine learning, however, requires the aggregation of operator data, which raises concerns regarding privacy, ownership, and scalability.

Federated learning (FL) is a distributed training scheme in which each participant trains a local copy of a shared model on its own data and exchanges only model parameters with a coordinating server, so that raw measurements never leave the operator domain. Applied to SE, this approach aligns collaborative training with the operational separation of TSOs and DSOs and distributes computational effort among participants. FL has previously been applied to SE [31] and to anomaly detection on state-estimation data [20], and personalized or hierarchical variants improve adaptation to heterogeneous networks [33]. Existing studies, however, largely rely on small benchmark networks, which motivates the assessment carried out here on a network of realistic size and heterogeneity [1].

4.2 Extension of the Dutch Model to Distribution Networks

The model of Section 3 represents the 380, 220, 150, and 110 kV backbone of the national system and accounts for distribution-level effects by attaching, at the lowest transmission voltage levels, aggregated models of solar generation, wind generation, and equivalent loads that stand in for the downstream power flows [3]. Coordinated estimation across the T–D boundary requires those downstream systems to be represented explicitly, so three synthetic distribution networks are connected at representative 150 kV busbars located in Noord-Holland (DN1), Gelderland–Flevoland–Utrecht (DN2), and Brabant (DN3). Together, these areas account for a significant share of national demand and renewable capacity. The resulting integrated system is shown in Figure 7.

Critical demand values and nominal renewable capacities at the three interfaces are redistributed within the associated distribution networks. The original aggregated blocks are kept in the model as disabled elements, and new elements are instantiated inside the synthetic distribution networks, which preserves comparability with the transmission-only results of Section 3.

Each distribution network is connected through a 150/33 kV transformer, followed by a 33/10 kV stage that supplies two feeders serving industrial and residential consumers. The single-line diagram common to the three networks is shown in Figure 8. The aggregated 150 kV loads of the original model are disabled and represented at 10 kV within each distribution network, maintaining the regional active-power totals with a 60%/40% industrial/residential split.

Onshore wind capacity originally connected at the selected 150 kV nodes is partially relocated to 33 kV within the corresponding distribution network, reflecting the empirical mix of high- and medium-voltage connections of onshore wind generators. Utility-scale photovoltaic (PV) generation, previously represented at 150 kV, is de-aggregated to the distribution level and split between 33 kV for large plants and 10 kV for small plants and rooftop installations, using regional deployment fractions derived from national statistics. Although small-scale rooftop PV physically operates at low voltage, approximately 0.4 to 0.7 kV, it is aggregated at 10 kV to limit network depth while preserving its dominant influence on feeder voltages and local power balance.






Figure 8. Single-line diagram of the synthetic distribution networks.

4.3 Scenario Generation and Measurement Datasets

A dataset of steady-state operating points is generated in DIgSILENT PowerFactory by varying net demand and renewable injections across the integrated T–D system. All inverter-based resources operate at a maximum power factor between 0.9 and 0.95 and follow a voltage-reactive-power droop control scheme. The network is simulated under fully renewable operating conditions, with the reactive effects of lines, transformers, and shunt elements fully represented. Each operating point is solved through an AC power flow, and the resulting bus voltages and branch flows constitute the data used in the federated experiments. Transmission- and distribution-level measurements remain partitioned so as to emulate the data boundaries between system operators.

Operating points are generated through structured parameter sweeps around stable base-case values. Total national demand is divided into 25 bins between 12 GW and 26 GW, and within each bin, the renewable generation is scaled to match the aggregated demand plus transmission losses. The allocations of PV, wind, and load are then perturbed by small random variations to obtain diverse yet feasible operating points. Solar generation spans 0 to 0.80 p.u. and wind generation spans 0 to 1.00 p.u., representing a wide range of Dutch renewable conditions.

To emulate measurement uncertainty, bounded uniform noise of $\pm 2\%$ is added to all input measurements. This choice is consistent with SCADA accuracy classes, which specify a maximum percentage error rather than a specific probability distribution.

The estimator operates on 243 input features and predicts 270 voltage-state components. The input features comprise busbar voltage magnitudes, generator and load active and reactive powers, and the active and reactive power exchanges on the tie lines at the T–D interface. The prediction targets are the real and imaginary components of the voltage at every busbar in both the transmission and the distribution systems, which together define the complex network state. Each operator accesses only the features and predictions belonging to its own system, consistent with the federated setting described in Section 4.4.

A total of 30,000 feasible operating points is generated. Of these, 10,000 are first used to warm start the neural estimator, reflecting the fact that practical SE systems are trained continually and never begin from an uninitialized model. The remaining 20,000 scenarios are used during the federated training rounds.

4.4 Federated State Estimation Framework

Let the set of participating operators be denoted by K , comprising one TSO and three DSOs in the present configuration. Each client k holds a private dataset of n_k measurement-state pairs, and the total number of data points across all clients is n . Federated learning seeks the model parameters w that minimize the weighted empirical risk

$$F(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w) \quad (1)$$

with local objectives

$$F_k(w) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell \left(f_w \left(x_i^{(k)} \right), y_i^{(k)} \right) \quad (2)$$

where f denotes the neural estimator and ℓ the mean-squared error between predicted and reference voltages. Each global communication round proceeds in three stages. The TSO first broadcasts the current global model to all DSOs. Each client then initializes its local model from the broadcast weights and performs a fixed number of local training epochs on its own measurements. The TSO finally aggregates the returned client models into an updated global model.

In the local-update stage, client k initializes its local model as the broadcast global model and performs E stochastic gradient descent (SGD) steps indexed by j . FedAvg [16] and FedAdam [21] use the standard local SGD rule

$$w_{t,j+1}^{(k)} = w_{t,j}^{(k)} - \eta \nabla F_k \left(w_{t,j}^{(k)} \right) \quad (3)$$

which is effective when client datasets are not strongly heterogeneous. FedProx [15] instead modifies the client-side objective by adding a proximal term that limits drift from the broadcast model,

$$F_k^{prox}(w; w_t) = F_k(w) + \frac{\mu}{2} \|w - w_t\|^2 \quad (4)$$

where μ controls the strength of the regularization and recovers the FedAvg objective when set to zero. The resulting local SGD step becomes

$$w_{t,j+1}^{(k)} = w_{t,j}^{(k)} - \eta \left(\nabla F_k \left(w_{t,j}^{(k)} \right) + \mu \left(w_{t,j}^{(k)} - w_t \right) \right) \quad (5)$$

After completing the E local steps, the client returns its updated model to the TSO.

In the aggregation stage, the TSO updates the global model using one of the following strategies. FedAvg and FedProx form a data-size-weighted average of the returned client models,

$$w_{t+1} = \sum_{k=1}^K \frac{n_k}{n} w_{t+1}^{(k)} \quad (6)$$

whereas FedAdam accelerates convergence under heterogeneous updates by applying adaptive moment estimation on the server side. The TSO first forms the aggregated model delta,

$$\Delta_t = \sum_{k=1}^K \frac{n_k}{n} (w_{t+1}^{(k)} - w_t) \quad (7)$$

then maintains exponential moving averages of that delta and of its element-wise square,

$$m_{t+1} = \beta_1 m_t + (1 - \beta_1) \Delta_t \quad (8)$$

$$v_{t+1} = \beta_2 v_t + (1 - \beta_2) \Delta_t \odot \Delta_t \quad (9)$$

and finally updates the global model as

$$w_{t+1} = w_t + \eta_s \frac{m_{t+1}}{\sqrt{v_{t+1}} + \tau} \quad (10)$$

where β_1 and β_2 are the decay factors of the first and second moments, $\odot$ denotes the element-wise product, τ is a small constant that ensures numerical stability, and η_s is the server learning rate.




Figure 9. Composition of input features and predicted voltage states for the residual MLP estimator.

All clients employ the same residual multilayer perceptron (MLP), which maps the 243 input features to the 270 predicted voltage components. The network uses an input width of 64, two residual blocks of width 64, and a final linear output layer, totalling approximately 5×10^4 trainable parameters. The hidden layers apply GELU activations, dropout, and layer normalization, which help maintain robustness and good gradient flow during learning. The 20,000 scenarios used in the federated rounds give a sample-to-parameter ratio appropriate for a model of this size and support

stable training across heterogeneous clients. The composition of the input features and of the predicted voltage states is shown in Figure 9.

The federated process reflects the operational hierarchy of the Dutch grid, in which the TSO maintains the global model and the DSOs train locally on their private measurements. The implementation uses the open-source Flower framework with synchronous communication rounds. The 20,000 available scenarios are split globally into 80% training, 10% validation, and 10% testing. Feature normalization uses the global training statistics, and each client receives only the standardized features corresponding to its own measurement set, which preserves TSO–DSO data separation while ensuring consistent scaling across all local datasets. Although the formulation above is written in terms of gradient-descent updates, the implementation uses AdamW for local optimization, which empirically improves stability with heterogeneous data. Each DSO trains for two local epochs per round with a batch size of 128. The hyperparameter values used for the three strategies are summarized in Table 4.

Table 4. Federated learning hyperparameters for each strategy.

Parameter	FedAvg	FedProx	FedAdam
Warm start	Yes	Yes	Yes
Rounds	300	300	300
Clients (K)	4	4	4
Batch size	128	128	128
Local epochs	2	2	2
Client learning rate	10^{-4}	10^{-4}	10^{-4}
Client weight decay	10^{-4}	10^{-4}	10^{-4}
Proximal term μ	–	0.005	–
Server learning rate	–	–	10^{-3}
Server (β_1, β_2)	–	–	(0.9, 0.999)
Server τ	–	–	10^{-9}

Estimation performance is quantified using the root-mean-squared error (RMSE) and the mean absolute error (MAE) between the predicted voltage components $\hat{y}_i$ and their reference values y_i ,

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2} \quad (11)$$

$$MAE = \frac{1}{N} \sum_{i=1}^N |\hat{y}_i - y_i| \quad (12)$$

where N is the number of predicted components. Both metrics are computed over the 270 predicted real and imaginary voltage components and are reported in per unit, and the global metrics average these errors across all clients.

4.5 State Estimation Results

A centralized implementation of the estimator was trained to provide a reference for the federated strategies. It uses the same warm-start initialization and the same pool of 20,000 scenarios used in

the federated experiments, providing a centralized reference with access to measurements from all operator domains. Centralized aggregation of operator data is generally incompatible with multi-operator data-ownership constraints, which is the reason the federated configuration is pursued.

Convergence is assessed on the validation split. Figure 10 shows the validation RMSE of the centralized baseline and of the three federated strategies. The centralized model remains stable throughout training, with validation RMSE between 3.35×10^{-3} p.u. and 5.07×10^{-3} p.u. All three federated methods converge rapidly. FedAvg and FedProx reduce their validation error smoothly during the first 30 global rounds, with FedProx slightly more stable owing to the proximal term that limits client drift. FedAdam oscillates initially but quickly benefits from the adaptive server-side moment updates and reaches the lowest steady-state validation error. Convergence stabilizes after roughly 120 global rounds for all three strategies.




Figure 10. Validation RMSE of (a) the centralized and (b) the federated estimators, in per unit (p.u.).

Accuracy is reported on the held-out test split, which comprises 10% of the federated scenarios and is disjoint from the data used for training and validation. Figure 11 summarizes the test errors. The centralized model achieves an RMSE of 0.0033 p.u. and an MAE of 0.0023 p.u. FedAdam obtains the best federated performance, with an RMSE of 0.0158 p.u. and an MAE of 0.0104 p.u., followed by FedProx with 0.0167 p.u. and 0.0110 p.u., and FedAvg with 0.0170 p.u. and 0.0112 p.u. An RMSE on

the order of 10^{-2} p.u. corresponds to voltage magnitude errors of around 1 to 2%, which is comparable to typical measurement uncertainty levels in operational power systems. FedProx improves robustness under imbalanced client datasets, while FedAdam benefits from the adaptive server-side updates.




Figure 11. Test performance of the centralized and federated estimators, expressed as RMSE and MAE in per unit (p.u.).

The reduced distribution feeders provide a manageable test environment for examining differences between operator datasets. Extending the framework to larger distribution networks would allow its performance to be assessed with richer measurement sets and a broader range of operating conditions.

4.6 Connection to EMT Model Initialization

The estimator returns the real and imaginary voltage components of every busbar across both voltage levels, which is a complete network solution rather than a set of interface quantities. This is the same class of information that defines the starting point of a dynamic study: a consistent set of nodal voltages and the corresponding branch flows and injections, from which the initial condition of a time-domain simulation is built. Estimated states can therefore support the preparation of operating conditions for EMT studies directly from measurements distributed across operators, which complements the scenario-based operating points generated in Section 3.

A complete dynamic initialization involves more than the network solution. Synchronous machine field and rotor quantities, converter DC-link voltages and inner-loop current references, and the integrator states of the associated controllers must all be consistent with the network condition and with one another, and each is derived from the network solution together with the corresponding device models. The estimated voltage state provides the network-level input to that process, and the estimator developed here is the component that makes this input available under operator-level data separation [1].

The following section presents the regional real-time EMT model, which establishes the dynamic simulation capabilities subsequently extended to the national transmission network in Section 6.

5 Real-Time EMT Model Development

Building on the steady-state modeling foundation, this section develops the regional EMT representation required to study the dynamic response of synchronous machines, converters, and their controls under transmission-network disturbances.

5.1 Randstad Region's Network

Offshore wind connection areas are informative locations for a detailed dynamic study, because they concentrate in a small part of the network the elements whose interaction determines the disturbance response: large converter-interfaced generation, the remaining synchronous units, the transformers and cable circuits that carry the infeed inland, and the connections to neighboring networks. The manner in which the offshore capacity is connected, together with the response of the converter controls, governs how active and reactive power are exchanged with the network while voltage is depressed and during the recovery that follows, so the electrical layout and the control behaviour cannot usefully be examined in isolation from one another [34].

Beginning with a regional representation allows these components and their interactions to be examined individually before network coverage is extended. At this scale, the response of each converter, machine, and circuit can be inspected directly, model behaviour can be checked against expectations, and the computational resources required for an EMT representation remain modest enough to permit iteration. The regional model reported here, therefore, serves both as a study object in its own right and as the development basis for the national-scale extension described in Section 6 [4]. The topology of the Randstad region is described in the following paragraphs.




Figure 12. Randstad region's electrical network. Adapted from TenneT's website.

The first step in building the extra-high-voltage (EHV) transmission network corresponding to the Randstad region was establishing its corresponding topology. The transmission lines and power transformer parameters were sourced from previous research conducted at Delft University of

Technology. Additional grid data was obtained from public reports from the Transmission System Operator TenneT. Figure 12 shows the Randstad electrical grid and the offshore wind farms considered for the model [30]. The regional real-time EMT model described in this section builds on the implementation reported in [4].

5.2 Model Development in RSCAD

5.2.1 Power and Control Connections of a Wind Turbine in RSCAD

Wind farms' power outputs were supported by data from the Netherlands Enterprise Agency (RVO) [23]. This model includes a 6-MW type-4 wind turbine and its power electronic converters. The model is scaled up to 700 MW to match the power output of Hollandse Kust (Noord), Hollandse Kust (Zuid) Alpha, and Hollandse Kust (Zuid) Beta wind farms. Figure 13 illustrates the components of the wind farm model, comprising:

- **Small dt Block (6 MVA):** This smaller time-step block contains a permanent magnet synchronous machine (PMSM), voltage source converters (VSC), and the DC link between the converters.
- **Grid VSC Controls:** This hierarchy box manages the controls for the grid-side converter.
- **PMSM VSC Controls:** This hierarchy box governs the machine-side converter.
- **Wind turbine and measuring devices:** A hierarchy box that includes measuring devices and functional blocks to simulate the dynamic behaviour of the wind turbine.

Each component will be examined in further detail in the subsequent sections.




Figure 13. Power and control connections of a wind farm in RSCAD.

5.2.2 Small Time-Step Block (Small dt Block)

Simulations run in the Real Time Digital Simulator (RTDS) via RSCAD require high-power processing to achieve real-time performance [24]. To optimize processing demands, RSCAD allows the

adjustment of simulation step times based on the specific requirements of different components. A dedicated step-time of 50 μ s is established for the main network operations to balance performance and computational load.

However, power electronics components involve complexities that require a finer time resolution for accurate simulation. The small time-step block addresses this by allowing a reduced step time, ranging from 1.4 μs to 2.5 μs , thereby ensuring numerical stability.

Figure 14 shows the configuration within this block. It includes a PMSM connected through a three-level voltage source converter to a DC link. The other end of the DC link is connected with another three-level converter linked to a transformer, which acts as the interface between the small step-time subnetwork (small dt block) and the main network.




Figure 14. Small time-step block (small dt block).

This section also hosts triangular wave generators and firing pulse generators for the operation of the three-level converters. These components utilize a comparison method where sinusoidal waves, generated from either the grid side or the PMSM, are compared with a triangular signal. This comparison generates the necessary logic to trigger the converter valves accurately.

5.2.3 Grid VSC Controls

This hierarchy box contains the controls of the grid-side VSC, which also allows the wind turbine to recover after a fault occurs. A synchronous reference frame control, or dq-control, provides adequate inputs to the controllers.

A phase-locked loop (PLL) will set the q-axis to zero. Then, this block's output angle will be used to calculate the dq components via the Park transformation. This transformation will be performed for the grid voltages as well as for the currents.

The PLL outputs are also used to create the shape of the triangular wave and its rate, which, in this case, is 19 times higher than the power system frequency.

Direct- and quadrature-voltage components are converted to the abc frame. Then, they are used to build the sinusoidal waves to compare with the triangular wave and give the firing logic for the valves on the grid side of the three-level converter.

5.2.4 Permanent Magnet Synchronous Machine (PMSM) VSC Controls

On the machine side, the currents that need to be controlled are the stator currents. As illustrated in Figure 15, these currents are transformed into the dq reference frame using the rotor angle in the Park transformation. After transformation, the direct current i_d is compared with the machine speed and then introduced into the current controller.

i_q is then compared to a reference current derived from a proportional-integral (PI) controller within the maximum power point tracking (MPPT) module, optimizing power generation across varying wind speeds.

The adjusted current is fed into the current controller, whose output comprises voltages in the dq reference frame. Subsequently, these voltages are converted to the abc reference frame to generate sinusoidal waves. These waves are then compared with a triangular wave to establish the firing logic for the valves on the synchronous machine side.




Figure 15. PMSM current controller.

5.2.5 Wind Turbine and Measuring Devices

This hierarchy box encompasses the aerodynamic model of the wind turbine, as depicted in Figure 16. The power coefficient (C_p) vs. tip-speed ratio (λ) curve is central to this model, which illustrates how efficiently the turbine converts wind kinetic energy into electrical power across different blade speeds. Additionally, the model integrates a control mechanism for adjusting the pitch angle, which is regulated by a slider that modifies the PMSM speed to optimize the turbine's response to changing wind conditions.




Figure 16. Wind turbine aerodynamic model.

5.2.6 Equivalent Models of Neighboring Networks

Since only the electrical network of the Randstad region is under study, connections to neighboring networks are modelled based on their active and reactive power flow extraction. These equivalents are represented using dynamic load blocks with active and reactive power consumption equal to the corresponding power flow through the network interconnection.

5.3 Simulation Results using the Real-Time Digital Simulator (RTDS)

5.3.1 Real-Time Simulation Performance

In a real-time simulation, simulated time advances at the same rate as wall-clock time. Every calculation belonging to a given simulation step, including the solution of the network, the component models, and the associated controls, must therefore be completed within the interval that the step represents. Whether this condition can be met depends on the computational resources available, on how the network is partitioned across processors, and on the level of detail with which individual components are represented, since each of these determines the work to be performed in each interval [24]. Maintaining synchronization with real time is what allows a simulation to exchange signals with equipment outside the solver, so a model executing in this manner could subsequently be coupled to external controllers, protection, or converter hardware, or operator-facing interfaces. The performance obtained for the regional model is reported below.

This study conducted real-time simulations using three NovaCor racks within the RTDS system at Delft University of Technology. The model operates continuously in real-time, with a consistent simulation time step of 50 μ s for the main network and a time step ranging from 1.4 μ s to 2.5 μ s for high-frequency components. The system demonstrated stable performance throughout, maintaining real-time operation without latency or delays.

5.3.2 Signal Selection

To effectively test the system's behaviour, the following signals were selected:

- Active power from three generators connected at Substation 6.
- Active and reactive power from the Hollandse Kust Noord and Zuid wind farms.
- Active power of the BritNed DC link.

5.3.3 Steady-State Simulation




Figure 17. Power output of generators connected to Substation 6 during normal start-up to reach steady state in RSCAD.

Figure 17 illustrates the active power response of the generators connected to Substation 6. The generators reach a steady state within 6 seconds.

Figure 18 depicts the response of the wind farms and the BritNed DC connection. The three wind farms achieve a steady-state power output of 700 MW approximately 15 seconds after initialization. Their reactive power remains stable at around zero, indicating no significant fluctuations. Additionally, the BritNed DC connection attains its nominal active power in about 6 seconds and maintains it throughout the simulation.




Figure 18. Power output of Hollandse Kust wind farms (1-Alpha and 2-Beta) and BritNed during normal start-up to reach steady-state in RSCAD.

5.3.4 Behaviour under Three-Phase Short-Circuit Fault

The system resilience was tested by introducing a 150 ms three-phase short-circuit fault at the main busbar of Substation 1, near the connection point of Hollandse Kust Noord. Exceptional settings were used as an exercise to see how unprecedented oscillations could theoretically be caused. Hence, the duration of this disturbance is longer than typical real-world clearance times, and it was chosen to assess the system dynamics under unprecedented challenging conditions. Figure 19 shows the active power of the generators, which oscillates after the fault but stabilizes once the fault is cleared. In practice, faster fault clearance would likely result in less pronounced oscillations.




Figure 19. Behaviour of generators connected to Substation 6. Short-circuit fault.

Figure 20 illustrates the performance of the northern wind farm. During the fault, there is a noticeable drop in active power, which stabilizes post-fault. Reactive power is employed during the fault to support voltage levels and returns to zero after the fault is cleared. The fault did not significantly affect the southern wind farms, maintaining a steady-state active power of 700 MW. Therefore, to see the response of these wind farms, a second location was chosen to apply the fault.




Figure 20. Behaviour of Hollandse Kust Noord wind farm. Short-circuit fault.

The second fault was applied at Substation 6's main busbar to evaluate the Hollandse Kust Zuid connection. As seen in Figure 21, the wind farms' active and reactive power show significant oscillations, which do not stabilize. In reality, coordinated control between synchronous generators and wind farms and faster protection schemes would mitigate such issues.




Figure 21. Behaviour of Hollandse Kust wind farms (1-Alpha and 2-Beta). Short-circuit fault.

6 National-Scale EMT Stability Analysis under High-RES Scenarios

6.1 Extension to the Full National Transmission Network

The national-scale EMT stability assessment developed within TwinEU Task T8.2 is also reported in an accepted IFAC 2026 conference paper [5].

The Randstad model used small-time-step subnetworks ($1.4 \mu\text{s}$ to $2.5 \mu\text{s}$) to host switched VSC models for the offshore wind farms. Extending this approach to the full national network is not feasible in real-time due to the computational cost of multiple such subnetworks. The national model is therefore solved entirely at the $50 \mu\text{s}$ distribution-mode time step. This choice preserves real-time execution at the larger network scale while remaining well below the dominant electromagnetic time constants of the overhead lines, transformers, and synchronous machines that drive the short-term post-fault dynamics of interest.

The model represents the 380 kV, 220 kV, 150 kV, and 110 kV voltage levels in detail, including synchronous machines, transformers, overhead lines, and dynamic loads. Network topology and substation data are based on publicly available information from TenneT. Interconnections with neighboring areas are represented by Thévenin equivalents tuned to reproduce the short-circuit levels of the boundary buses under the operating conditions studied.

Converter-interfaced renewable sources are represented using the three-phase dynamic PQ source available in the RSCAD distribution library [24]. This component captures the network-interaction behaviour of grid-following converters, including a phase-locked loop, abc-dq transformation, and outer-loop PI regulators with anti-windup and rate limits, while remaining compatible with the $50 \mu\text{s}$ solver. The inner current loop and PWM switching are abstracted; the outer-loop d- and q-axis current references are injected directly through controlled current sources. This approach enables national-scale real-time simulation while retaining the grid-following PLL tracking and outer-loop recovery dynamics that are relevant for the voltage- and frequency-based stability indicators studied here. As an example, Figure 22 shows this implementation around Substation 1, where dynamic PQ sources represent the offshore wind and HVDC infeeds.




Figure 22. Annotated view of the RSCAD implementation around Substation 1, showing the main busbar selected for the three-phase short-circuit fault and the modelled elements, including synchronous machines, transformers, and dynamic PQ sources for offshore wind and HVDC infeeds. Internal labels are anonymized for confidentiality.

6.2 Renewable Penetration Scenarios and Disturbance Scenario

Renewable penetration is a useful dimension along which to compare dynamic operating conditions, because it captures in a single parameter a change in the composition of the generating fleet that has direct consequences for the disturbance response. As online synchronous generation is displaced by converter-interfaced sources, the synchronous inertia available to the system is reduced, the fault current supplied to the network changes in magnitude and character, and a larger share of the response to a disturbance is determined by converter control action rather than by the electromechanical behaviour of rotating machines [10], [34].

It is important to distinguish installed renewable capacity from the share of generation actually supplied by renewable sources at a given operating point. Installed capacity describes the available generating fleet, whereas the renewable generation share depends on resource availability, demand,

and dispatch at a particular operating point. The associated dynamic response also depends on which synchronous units remain connected and on the operating modes of converter-interfaced sources. Penetration is nevertheless not the only influence: demand level and location, the dispatch of the remaining synchronous units, network topology and outage conditions, and the settings of converter and protection controls all shape the observed response. Comparing operating points, therefore, requires the disturbance itself to be held constant, so that differences in the recorded response can be attributed to differences in the operating condition rather than to differences in the event. The three renewable-generation shares examined in this section are constructed on that basis.

Three operating cases were constructed at renewable shares of 50%, 65%, and 80% of total generation, bracketing a plausible range of 2030 Dutch operating conditions [19], [22], [25]. The active-power outputs of converter-interfaced sources and the dispatch of conventional synchronous generation were scaled across scenarios to meet these target shares. Individual unit ratings are omitted for confidentiality. Each operating point is initialized from a steady-state power-flow solution.

To assess system performance under the three scenarios using a common disturbance, a three-phase short-circuit was applied at the extra-high-voltage busbar of a representative substation (referred to as Substation 1 for confidentiality). The fault duration was set to 50 ms, corresponding to typical primary-protection clearance times in the 380 kV transmission network (approximately 2.5 cycles). Applying the same disturbance across all scenarios enables a clean comparative assessment of the system response under varying RES penetration.

6.3 Stability Indicator Extraction Methodology

To quantify the dynamic response of the national-scale EMT model, voltage- and frequency-based stability indicators are extracted from the simulated waveforms. These indicators implement the measurable voltage, frequency, RoCoF, and recovery-time dimensions of the demonstrator-level KPI framework defined in Section 2.2.

Indicators of this kind summarize selected characteristics of a disturbance response as scalar quantities, making it easier to compare, tabulate, and rank scenarios alongside inspection of the underlying waveforms. The voltage indicators address three distinct questions. The minimum voltage reached describes the severity of the sag, the durations spent below defined thresholds describe its persistence, and the recovery time describes how long the system takes to restore an acceptable voltage once the fault has been cleared. The frequency indicators are organized in the same way. The nadir and the post-fault maximum bound the excursion away from nominal, the rate of change of frequency characterizes how quickly that excursion develops, and the frequency recovery time describes the return toward nominal conditions.

Several indicators are needed together because similar extreme values can accompany quite different behaviour. Two scenarios may reach the same frequency nadir while differing markedly in how rapidly the deviation develops or in how long the system takes to settle, and a shallow sag that persists can be more demanding than a deeper one that clears promptly. Reading the indicators as a set, therefore, gives a fuller picture than any single value. For the comparison to be meaningful, the quantities must also be derived consistently, since each indicator depends on the conventions used to obtain it: the voltage base against which per-unit values are expressed, the measurement window over which rates of change are evaluated, the filtering applied before the derived quantities are computed, and the reference instants for fault inception and clearing. The procedure set out below defines these conventions and the signal processing applied to the simulated waveforms.

The fundamental-frequency phasor for each phase $\varphi \in \{a, b, c\}$ is estimated by a one-cycle sliding Discrete Fourier Transform (DFT) applied to the real-valued sampled instantaneous phase-to-neutral voltage $v_\varphi[n]$ [kV]:

$$\bar{V}_\varphi[n] = \frac{2}{N} \sum_{k=0}^{N-1} v_\varphi[n-k] e^{+j\frac{2\pi k}{N}} \quad (13)$$

where $N = \left\lfloor \frac{1}{f_0 T_s} \right\rfloor = 400$ is the number of samples per nominal cycle at $f_0 = 50$ Hz and $T_s = 50$ μ s, and $\bar{V}_\varphi[n] \in \mathbb{C}$ [kV] is the resulting complex peak phasor. The factor $\frac{2}{N}$ ensures that $|\bar{V}_\varphi[n]|$ equals the peak amplitude of the fundamental component for a pure sinusoid.

The positive-sequence phasor is then obtained by applying the symmetrical-components rotation to the three per-phase phasors:

$$\bar{V}_1[n] = \frac{1}{3} (\bar{V}_a[n] + \alpha \bar{V}_b[n] + \alpha^2 \bar{V}_c[n]), \alpha = e^{j\frac{2\pi}{3}} \quad (14)$$

The positive-sequence RMS voltage, a real non-negative scalar [kV], is

$$V_{1,\text{RMS}}[n] = \frac{|\bar{V}_1[n]|}{\sqrt{2}} \quad (15)$$

The instantaneous frequency [Hz] is derived from the phasor angle $\theta[n] = \angle \bar{V}_1[n]$ [rad], unwrapped to remove the 2π discontinuities introduced by the principal-value branch of the inverse tangent, via a centered difference:

$$f[n] = \frac{\theta[n+1] - \theta[n-1]}{4\pi T_s} \quad (16)$$

After validity gating and 20 ms moving-average smoothing (described below), the RoCoF [Hz/s] is estimated by a central difference over a 100 ms measurement window:

$$\text{RoCoF}[n] = \frac{f[n+M] - f[n-M]}{2 M T_s} \quad (17)$$

where $M = \left\lfloor \frac{50 \text{ ms}}{T_s} \right\rfloor = 1,000$ samples, so that $2 M T_s = 100$ ms.

The faulted bus is a generation substation hosting multiple synchronous machines and step-up transformers, and consequently, operates with a pre-fault voltage above the nominal phase value. Using the system nominal $\frac{380}{\sqrt{3}}$ kV as the base would yield pre-fault values above 1 p.u. and distort the interpretation of the voltage stability indicators. The chosen base is therefore the pre-fault steady-state magnitude $V_{ss} = \langle V_{1,\text{RMS}}[n] \rangle_{t < t_{\text{fault}} - \frac{T}{2}}$, where $T = \frac{1}{f_0}$ and the average is taken over phasor windows lying fully before the fault.

Voltage-based stability indicators are then computed from $\frac{V_{1,RMS}[n]}{V_{ss}}$ and include the minimum value V_{min} , the sag durations below 0.90 and 0.80 p.u., and the voltage recovery time T_{rec} (the elapsed time from fault clearing until $V_{1,RMS}$ enters and remains within $\pm 5\%$ of V_{ss} for at least one fundamental cycle).


During deep voltage sags ($|\bar{V}_1[n]| \rightarrow 0$), the angle increment in Eq. (16) becomes ill-conditioned and produces large non-physical spikes that contaminate $f[n]$ and $RoCoF[n]$. To avoid reporting these as physical effects, samples for which $\frac{V_{1,RMS}[n]}{V_{ss}} < 0.30$ are flagged as invalid, and the invalid mask is dilated by one full DFT cycle to account for window smearing of the sliding DFT. The frequency $f[n]$ is then smoothed with a 20 ms moving average over the valid samples before $RoCoF[n]$ is computed via Eq. (17). The 100 ms window in Eq. (17) is consistent with the functional requirements for $RoCoF$ measurement defined in IEC 60255-181 [13].

Frequency-based stability indicators include the nadir, post-fault maximum, maximum $|RoCoF|$, and the frequency recovery time $T_{f,rec}$ (the first post-clearing instant at which the smoothed valid frequency enters and remains within ± 0.1 Hz of nominal for at least one fundamental cycle).


6.4 Simulation Results

6.4.1 Voltage Response


Figures 23 and 24 show the three-phase voltages and the positive-sequence RMS voltage envelope at Substation 1 for the three RES penetration scenarios. The minimum RMS voltage, evaluated from the positive-sequence phasor of Eqs. (13)–(15) and normalized to the pre-fault steady-state base defined in Section 6.3, decreases only slightly with renewable share, from 0.0347 p.u. at 50% to 0.0330 p.u. at 80%, confirming a near-complete voltage sag at the fault location in all cases. The sag durations below 0.90 p.u. and 0.80 p.u. remain in the narrow ranges 73–75 ms and 68–70 ms respectively, dominated by the imposed 50 ms fault duration rather than by penetration level. Voltage recovery time shows a non-monotonic pattern, with the 65% case requiring the longest time due to an oscillatory recovery transient.



(a) 50% RES penetration




(b) 65% RES penetration




(c) 80% RES penetration


Figure 23. Phase-to-neutral voltages at Substation 1 for the three RES penetration scenarios.



(a) 50% RES penetration



(b) 65% RES penetration




(c) 80% RES penetration


Figure 24. Positive-sequence RMS voltage at Substation 1 for the three RES penetration scenarios.

6.4.2 Frequency and RoCoF Response


Figure 25 shows the post-clearing instantaneous frequency excursion at Substation 1, derived from the positive-sequence phasor angle via Eq. (16). The frequency nadir deepens from 49.83 Hz at 50% RES to 49.28 Hz at 80%, and the post-fault peak rises from 50.17 Hz to 51.42 Hz, reflecting the larger transient swing as inertia is reduced. Figure 26 shows the RoCoF computed via Eq. (17), which is the most penetration-sensitive indicator: the maximum absolute RoCoF increases monotonically from 2.04 Hz/s at 50% to 5.31 Hz/s at 65% and 14.54 Hz/s at 80%, spanning nearly an order of magnitude. This constitutes a textbook signature of reduced equivalent inertia and weaker electromechanical coupling. The frequency recovery time follows a non-monotonic trend similar to the voltage recovery time, peaking at 65% due to an oscillatory return attributable to the interplay between residual synchronous inertia and converter-driven damping.



(a) 50% RES penetration




(b) 65% RES penetration




(c) 80% RES penetration


Figure 25. Instantaneous frequency at Substation 1 for the three RES penetration scenarios.



(a) 50% RES penetration



(b) 65% RES penetration



(c) 80% RES penetration

Figure 26. Rate of change of frequency (RoCoF) at Substation 1 for the three RES penetration scenarios.

6.4.3 Summary of Stability Indicator Results

Table 5. Stability indicators per scenario.

Stability Indicator	50% RES	65% RES	80% RES
$V_{\min}$ [p.u.]	0.0347	0.0342	0.0330
Sag < 0.90 p.u. [s]	0.0736	0.0733	0.0752
Sag < 0.80 p.u. [s]	0.0680	0.0678	0.0699
T_{rec} [s]	0.0174	0.0373	0.0188
f_{nadir} [Hz]	49.83	49.79	49.28
f_{max} [Hz]	50.17	50.46	51.42
$ \text{RoCoF} _{\text{max}}$ [Hz/s]	2.04	5.31	14.54
$T_{f,\text{rec}}$ [s]	0.0683	0.1259	0.0970

The indicators in Table 5 summarize the dynamic response of the national EMT demonstrator for the selected three-phase fault scenario. The frequency-side indicators all degrade monotonically with renewable share: the maximum RoCoF increases from 2.04 Hz/s to 14.54 Hz/s, the nadir deepens from 49.83 Hz to 49.28 Hz, and the post-fault maximum rises from 50.17 Hz to 51.42 Hz, a signature of reduced equivalent inertia at the faulted bus. The voltage-side indicators change only marginally: $V_{\min}$ also decreases monotonically, but by less than 5%, and the sag durations below 0.90/0.80 p.u. remain within 73-75/68-70 ms, set by the imposed 50 ms fault duration. This is consistent with the fault current being dominated by the synchronous machines in service, since the converter models provide no reactive-current support during the sag. Both recovery times, T_{rec} and $T_{f,\text{rec}}$, are non-monotonic and peak in the 65% case, where partial synchronous inertia coexists with substantial converter-driven dynamics, possibly producing an oscillatory recovery. As the converter models include PLL and outer-loop dynamics but no fault-ride-through or inner-loop current control, the indicators represent a conservative, no-fault-support bound, obtained while preserving real-time execution at national scale.

7 Conclusions

This deliverable reports the work performed within TwinEU Task T8.2 on the development of a digital-twin demonstrator for stability assessment of the Dutch power system. By combining steady-state scenario modeling, federated state estimation, and real-time EMT simulation, this study establishes a research workflow for relating future operating conditions to their response under disturbances [6], [10].

The steady-state synthetic model gave the demonstrator a foundation for examining 2030 high-VRES, electrolyzer-heavy operating conditions without access to confidential network data. It generated consistent operating points for dynamic studies and the labelled datasets used to train and evaluate the state estimator. The scenario results also showed where voltage and transformer-loading constraints could arise and how electrolyzer placement affected them.

The federated estimator extended that foundation across the transmission–distribution boundary. Trained across one TSO and three DSO clients with all measurements held locally, it reconstructed the voltage state of both systems with test errors of 0.0158 p.u. RMSE and 0.0104 p.u. MAE for the best-performing strategy, against 0.0033 p.u. and 0.0023 p.u. for a centralized benchmark with full data access. The practical value lies in achieving this estimation performance while retaining measurements within each operator’s domain. The estimator, therefore, supplies the component that connects distributed measurements to the operating conditions used for dynamic assessment, and the remaining step toward an automated initialization path is identified as further development.

The regional real-time EMT model established RTDS/RSCAD as a viable simulation core for the digital twin and, more importantly, opened the door to hardware-in-the-loop, controller-oriented, and operator-oriented studies that the steady-state model cannot support. This is the capability the demonstrator adds beyond conventional offline analysis.

At the national scale, the selected fault scenario revealed a clear trend in the frequency indicators as renewable penetration increased, consistent with the reduced synchronous inertia of the simulated operating points. The voltage indicators changed less systematically, and both recovery times peaked in the case of intermediate renewable penetration. The key lesson is that these responses must be interpreted together: a clear trend in one set of indicators does not capture the full post-fault behaviour.

Together, these developments establish a common framework for preparing operating conditions, examining disturbance responses, and interpreting their implications for system stability. The demonstrator therefore shows how future digital twins can support large-scale power system stability monitoring and assessment. In future operator-facing implementations, similar DT-based environments could combine updated operating conditions, real-time or near-real-time simulation, and automated extraction of stability indicators to support voltage and frequency stability screening, scenario comparison, contingency assessment, and operator training. The work reported in this deliverable provides a practical demonstration of this concept within a controlled research environment and establishes a technical basis for further development of digital-twin-based stability-assessment tools for high-renewable power systems.

8 References

- [1] Akaouche, A., Aviles-Cedeno, J., and Rueda-Torres, J. (2026). Federated Learning-Based State Estimation for Integrated Transmission–Distribution Networks. Presented at the IEEE PES General Meeting 2026. Proceedings forthcoming at the time of this deliverable.
- [2] Alvarez-Alvarado, M.S., Apolo-Tinoco, C., Ramirez-Prado, M.J., Alban-Chacón, F.E., Pico, N., Aviles-Cedeno, J., Recalde, A.A., Moncayo-Rea, F., Velasquez, W., and Rengifo, J. (2024). Cyber-physical power systems: A comprehensive review about technologies drivers, standards, and future perspectives. *Computers and Electrical Engineering*, 116, 109149. DOI: [10.1016/j.compeleceng.2024.109149](https://doi.org/10.1016/j.compeleceng.2024.109149).
- [3] Aviles-Cedeno, J., Rueda, J., and de Roos, A. (2024). Toward a Digital Twin of the Dutch EHV Network: Analyzing Future Multi-Energy Scenarios. In 2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), 1–5. DOI: [10.1109/ISGTEUROPE62998.2024.10863095](https://doi.org/10.1109/ISGTEUROPE62998.2024.10863095).
- [4] Aviles-Cedeno, J., Rueda, J., Perilla, A., and Palensky, P. (2024). A Real-Time EMT Digital Model for a Dutch Regional EHV Network: Integrating Offshore Power. In IECON 2024 - 50th Annual Conference of the IEEE Industrial Electronics Society, 1–6. DOI: [10.1109/IECON55916.2024.10905203](https://doi.org/10.1109/IECON55916.2024.10905203).
- [5] Aviles-Cedeno, J., and Rueda-Torres, J. (2026). Real-Time EMT Stability Analysis of the 2030 High-RES Dutch Power System. Presented at the IFAC World Congress 2026. Proceedings forthcoming at the time of this deliverable.
- [6] Brosinsky, C., Naglič, M., Lehnhoff, S., Krebs, R., and Westermann, D. (2024). A Fortunate Decision That You Can Trust: Digital Twins as Enablers for the Next Generation of Energy Management Systems and Sophisticated Operator Assistance Systems. *IEEE Power and Energy Magazine*, 22(1), 24–34. DOI: [10.1109/MPE.2023.3330120](https://doi.org/10.1109/MPE.2023.3330120).
- [7] Dimitrov, R.S. (2016). The Paris Agreement on climate change: Behind closed doors. *Global Environmental Politics*, 16(3), 1–11. DOI: [10.1162/GLEP_a_00361](https://doi.org/10.1162/GLEP_a_00361).
- [8] ENTSO-E. (2019). Bidding zone review methodology, assumptions and configurations - submitted to NRAs. Available at: <https://www.entsoe.eu/news/2019/10/07/bidding-zone-review-methodology-assumptions-and-configurations-submitted-to-nras/>. Accessed on 2024-06-28.
- [9] Fuller, A., Fan, Z., Day, C., and Barlow, C. (2020). Digital Twin: Enabling Technologies, Challenges and Open Research. *IEEE Access*, 8, 108952–108971. DOI: [10.1109/ACCESS.2020.2998358](https://doi.org/10.1109/ACCESS.2020.2998358).
- [10] Hatziargyriou, N., Milanovic, J., Rahmann, C., Ajarapu, V., Canizares, C., Erlich, I., Hill, D., Hiskens, I., Kamwa, I., Pal, B., Pourbeik, P., Sanchez-Gasca, J., Stankovic, A., Van Cutsem, T., Vittal, V., and Vournas, C. (2021). Definition and Classification of Power System Stability - Revisited & Extended. *IEEE Transactions on Power Systems*, 36(4), 3271–3281. DOI: [10.1109/TPWRS.2020.3041774](https://doi.org/10.1109/TPWRS.2020.3041774).
- [11] HoogspanningsNet. (2024). HoogspanningsNet. Available at: <https://www.hoogspanningsnet.com/>. Accessed on 2024-06-28.
- [12] Huy, T.H.B., Vo, D.N., Nguyen, H.D., Truong, H.P., Dang, K.T., and Truong, K.H. (2023). Enhanced power system state estimation using machine learning algorithms. In 2023 International Conference on System Science and Engineering (ICSSE), 401–405.
- [13] International Electrotechnical Commission. (2019). IEC 60255-181:2019 - Measuring relays and protection equipment, Part 181: Functional requirements for frequency protection. IEC, Geneva, Switzerland. Available at: <https://webstore.iec.ch/en/publication/30962>.

- [14] ISO/IEC 30173:2023. Digital twin - Concepts and terminology. International Organization for Standardization and International Electrotechnical Commission, 2023. Available at: <https://www.iso.org/standard/81442.html>. Accessed on 2024-06-28.
- [15] Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A., and Smith, V. (2020). Federated optimization in heterogeneous networks. In Proceedings of Machine Learning and Systems, 2, 429–450.
- [16] McMahan, B., Moore, E., Ramage, D., Hampson, S., and Agüera y Arcas, B. (2017). Communication-Efficient Learning of Deep Networks from Decentralized Data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, PMLR, 1273–1282.
- [17] MIGRATE Consortium. (2023). MIGRATE Project: Massive InteGRation of power Electronic devices. Horizon 2020 research and innovation programme. Available at: <https://cordis.europa.eu/project/id/691800>.
- [18] Ministerie van Economische Zaken en Klimaat. (2019). Klimaatakkoord hoofdstuk elektriciteit. Available at: <https://www.rijksoverheid.nl/documenten/2019/06/28/klimaatakkoord>. Accessed on 2024-06-28.
- [19] PBL Netherlands Environmental Assessment Agency. (2021). Netherlands Climate and Energy Outlook 2021: Summary. Technical report, PBL, The Hague, The Netherlands. Available at: <https://www.pbl.nl/uploads/default/downloads/pbl-2021-netherlands-climate-and-energy-outlook-2021-summary-4709.pdf>. Accessed on 2026-05-10.
- [20] Raghuvamsi, R., Teeparthi, Y.T., Mao, B., and Ukil, A. (2022). Incentive edge-based federated learning for false data injection attack detection on power grid state estimation. Applied Energy, 314, 118828.
- [21] Reddi, S.J., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., and McMahan, H.B. (2021). Adaptive federated optimization. In Proceedings of the 9th International Conference on Learning Representations (ICLR).
- [22] Reis, F., Torres, J.R., Palensky, P., and Gonzalez-Longatt, F. (2023). Future Dutch electricity grid: Assessing the potential of overplanting in photovoltaic systems. In IECON 2023 - 49th Annual Conference of the IEEE Industrial Electronics Society, 1–7. DOI: [10.1109/IECON51785.2023.10312436](https://doi.org/10.1109/IECON51785.2023.10312436).
- [23] Rijksdienst voor Ondernemend Nederland (RVO). (2024). Offshore Wind Energy. Available at: <https://offshorewind.rvo.nl>. Accessed on 2024-05-13.
- [24] RTDS Technologies Inc. (2024). RSCAD Power System Manual. RSCAD 7.5 edition. RTDS Technologies Inc., Winnipeg, Manitoba, Canada. Available at: <https://knowledge.rtds.com>.
- [25] Statistics Netherlands (CBS). (2024). Over Half of Electricity Production Now Comes from Renewable Sources. Available at: <https://www.cbs.nl/en-gb/news/2024/39/over-half-of-electricity-production-now-comes-from-renewable-sources>. Accessed on 2026-05-10.
- [26] TenneT. (2019). System & transmission data - explanatory documents - grid diagram.
- [27] TenneT. (2020). Investeringsplannen, net op land 2020-2029. TenneT technical report.
- [28] TenneT. (2024a). TenneT: Leading European Grid Operator. Available at: <https://www.tennet.eu>. Accessed on 2024-06-28.
- [29] TenneT. (2024b). Onshore Projects in the Netherlands, December 2023. Available at: https://tennet-drupal.s3.eu-central-1.amazonaws.com/default/2024-02/GB_DEC2023_Onshore_Netherlands.pdf. Accessed on 2024-05-13.
- [30] TenneT. (2026). Grid maps: high-voltage grid, substations and interconnectors in the Netherlands. Available at: <https://www.tennet.eu/nl-en/our-high-voltage-grid/grid-maps>. Accessed on 2026-05-10.

- [31] Veena, C.H., Jithendar, P., Anandhi, A.R.J., Singla, A., Bansal, M., and Ali, R.R. (2023). Federated learning for real-time state estimation and dynamic phasor analysis in wide-area power grids. In 2023 International Conference on Power Energy, Environment & Intelligent Control (PEEIC), 678–682.
- [32] Wang, L., Zhou, Q., and Jin, S. (2020). Physics-guided deep learning for power system state estimation. *Journal of Modern Power Systems and Clean Energy*, 8(4), 607–615.
- [33] Wu, H., Xu, Z., and Ruan, J. (2024). PFL-DSSE: A personalized federated learning approach for distribution system state estimation. *CSEE Journal of Power and Energy Systems*, 10(5), 2265–2270.
- [34] CIGRE Working Group C4.56. (2022). Electromagnetic transient simulation models for large-scale system impact studies in power systems having a high penetration of inverter-connected generation. Technical Brochure 881. CIGRE. Available at: <https://www.e-cigre.org/publications/detail/881-electromagnetic-transient-simulation-models-for-large-scale-system-impact-studies-in-power-systems-having-a-high-penetration-of-inverter-connected-generation.html>