RECOMMENDATIONS FOR INTEGRATING A DIGITAL TWIN (DT) OF THE ELECTRIC PROPULSION DRIVE SYSTEM (EPDS) INTO SOFTWARE-DEFINED ELECTRIC VEHICLES (SDEV)

PRG2532

Prof Anton Rassõlkin

Dr Mahmoud Ibrahim

Funder Sihtasutus Eesti Teadusagentuur

Grant number PRG2532

Grant title Advanced Digital Tools to Accelerate the Development of Software-Defined Electric Vehicles

Project abstract

The project aims to advance Electric Propulsion Drive System (EPDS) Digital Twin (DT) technology for Software-Defined Electric Vehicles (SDEVs), with a focus on achieving DT adaptive and intelligent levels. It addresses the need for efficient testing and evaluation of electric propulsion systems in line with EU clean energy transition goals. Leveraging the rapid development of DT technology, the project seeks to contribute to SDV technology through enhanced modeling, data gathering, IoT integration, and system optimization. Key challenges include lifecycle management, data processing, and real-time communication between physical and virtual systems. The project encompasses advanced modeling, data gathering, IoT, and communication infrastructure, system integration, optimization, and technology demonstration.

This guidance synthesizes the prior project PSG453 “Digital twin for propulsion drive of autonomous electric vehicle” (supported by the Estonian Research Council) research outputs into actionable recommendations for integrating an EPDS DT across the SDEV design, manufacturing, and maintenance/operations lifecycle. It targets engineering leadership, systems/controls engineers, production engineering, diagnostics, and after-sales.

Principal Investigator

Name Anton Rassõlkin

ETIS https://www.etis.ee/cv/Anton_Rassolkin

ORCID iD 0000-0001-8035-3970

Email anton.rassolkin@taltech.ee

Phone +372 53 919 587

Version: 1.0

Date: November 19, 2025

TABLE OF CONTENTS

Table of Contents 3
List of Abbreviations 4
Key Insights and Objectives 5
Reference Architecture (Twin-of-Twins) 6
Twin Instances 9
Integration into SDEV Design 11
Integration into Manufacturing 13
Integration into Maintenance & Operations 15
Methods & Tooling Recommendations 17
KPIs and Acceptance Criteria 19
Possible Implementation Roadmap (Phased) 22
Safety, Cybersecurity, and Compliance 23
List of Publications that Underpin the Recommendations 25

LIST OF ABBREVIATIONS

Abbreviation	Meaning
AI	Artificial Intelligence
ALM	Application Lifecycle Management
ANN	Artificial Neural Network
CAN	Controller Area Network
CBM	Condition-Based Maintenance
CI/CD	Continuous Integration / Continuous Deployment
DDS	Data Distribution Service
DoE	Design of Experiments
DT	Digital Twin
DTC	Direct Torque Control
EoL	End-of-Line
EPDS	Electric Propulsion Drive System
EV	Electric Vehicle
FMEDA	Failure Modes, Effects, and Diagnostic Analysis
FOC	Field-Oriented Control
GDPR	General Data Protection Regulation
HIL	Hardware-in-the-Loop
IoT	Internet of Things
KPI	Key Performance Indicator
LSTM	Long Short-Term Memory (neural network)
MIL	Model-in-the-Loop
ML	Machine Learning
MLOps	Machine Learning Operations
MQTT	Message Queuing Telemetry Transport
NRMSE	Normalized Root Mean Square Error
NTF	No Trouble Found
OPC UA	Open Platform Communications – Unified Architecture
OTA	Over-the-Air
PMSM	Permanent Magnet Synchronous Motor
PLM	Product Lifecycle Management
QoS	Quality of Service
ROS2	Robot Operating System version 2
RUL	Remaining Useful Life
SDEV / SDV	Software-Defined Electric Vehicle / Software-Defined Vehicle
SIL	Software-in-the-Loop
SoH	State of Health
SOME/IP	Scalable Service-Oriented Middleware over IP

KEY INSIGHTS AND OBJECTIVES

Why now?

The integration of a Digital Twin (DT) for the Electric Propulsion Drive System (EPDS) is particularly relevant at this stage of Software-Defined Electric Vehicles (SDEV) development. Propulsion represents the central value-creation mechanism in an SDEV, linking software functionality, performance, and sustainability targets. A well-implemented DT enables early validation of control strategies and power electronics while supporting system-level design optimization, fault diagnosis, predictive maintenance, and lifecycle management.

What to build?

The recommended architecture is a tiered, hybrid EPDS DT that combines physics-based models and data-driven approaches. This twin spans the entire propulsion system, from inverter, electrical machine, and transmission to high- and low-voltage interfaces, forming a comprehensive virtual representation of the propulsion system. Physics-based layers capture electromechanical, thermal, and energy conversion dynamics, while data-driven modules enhance adaptivity and capture nonlinear system behavior from operational data. All elements are linked through a digital thread that connects the lifecycle stages from initial requirements, design models, and validation tests to manufacturing configurations and in-service fleet data. This structure ensures traceability, model reuse, and cross-domain feedback, enabling continuous improvement and rapid innovation cycles.

How to run?

The implementation strategy should include three twin operational instances, each serving a distinct lifecycle phase. An offline *design twin* supports system architecture optimization, control algorithm development, and HIL/SIL validation. A *production twin* connects to test benches and End-of-Line (EoL) systems, validating components and assemblies during manufacturing by comparing their signatures with expected values. Finally, a *runtime twin* operates within the vehicle and cloud infrastructure, using real-time telemetry for performance monitoring, fault detection, and predictive analytics. Continuous calibration loops between these twins maintain fidelity and enable adaptive updates as the physical system evolves throughout its lifecycle.

Expected impact?

Such framework integration reduces reliance on physical prototypes, accelerates development cycles, and facilitates continuous improvement through real-time data feedback and adaptive software updates, ultimately contributing to more resilient and energy-efficient mobility solutions. In manufacturing, it enables right-first-time assembly, automatic fault localization, and process optimization by comparing virtual models with actual test results in real time. During vehicle operation, the DT supports predictive maintenance, real-time diagnostics, and performance tuning while providing the framework for Over-the-Air (OTA) feature upgrades, which are fundamental to the SDEV business model.

REFERENCE ARCHITECTURE (TWIN-OF-TWINS)

The DT architecture for the EPDS within SDEV provides a comprehensive framework for modeling, monitoring, and optimizing system performance across the vehicle lifecycle. This architecture is structured in multiple interconnected layers, each performing a distinct function while contributing to the overall digital thread. The layers - *Physical, Virtual Models, Data and Services*, *Applications,* and *Communication* create a cohesive ecosystem that links the real and virtual domains. Figure 1 visually summarizes the four layers of the DT architecture and their main components. At its core, this architecture establishes how data flows between hardware and software environments, ensuring real-time synchronization, accurate representation of system behavior, and continuous feedback for improvement. The *Physical layer* captures real-world measurements and testing data; the *Virtual Model*s layer builds digital representations of the propulsion system; the *Data and Services* layer manages the information pipeline and infrastructure; the *Application* layer enables the use of these models for design, validation, manufacturing, and operational decision-making; and the *Communication* layer provide real-time data exschange and analytical services by implementing commonly used protocols. Together, they form a robust foundation for achieving high reliability, efficiency, and adaptability in SDEV propulsion systems.

1. Physical enity

Lab rigs,
EoL testers,
In-vehicle telemetry.

2. Virtual entity (multi-fidelity)

Control/behavioral models,
Circuit models,
Data-driven models.

3. Communication

Real-time data exchange,
Analytical services using automotive and industrial protocols.

4. Data & services

Time-series ingestion,
Labeling,
Synchronization,
Machine Learning Operations (MLOps)
Model management,
Twin registry.

5. Applications

Design verification,
Controller evaluation,
EoL pass/fail analytics,
Fault detection and prognostics,
Over-the-Air (OTA) optimization.

Figure 1 Digital Twin (DT) Architecture Layers for the Electric Propulsion Drive System (EPDS).

The **physical layer** represents the tangible systems, test environments, and vehicles that provide real-world data for the DT. This includes laboratory test rigs composed of inverters, Permanent Magnet Synchronous Machines (PMSM), dynamometers, and battery emulators. These setups allow controlled experimentation across a wide range of torque-speed operating points. EoL testing stations in manufacturing use the same principles to validate every propulsion assembly against its expected behavior. In addition, in-vehicle telemetry streams real-time data during normal operation and testing campaigns. Collectively, these physical assets form the data foundation that drives model calibration, fault detection, and performance benchmarking in the DT ecosystem.

The **virtual model layer** provides a digital representation of the propulsion system with multiple levels of fidelity, each serving a specific engineering purpose. The models can be categorized into three main types:

- Control/behavioral models. These include simplified d–q domain representations of the PMSM and averaged inverter models. They are optimized for high-speed simulation and for the development of control algorithms, including testing Field-Oriented Control (FOC), Direct Torque Control (DTC), and sensorless estimation strategies. Driveline dynamics are often included at this level to study vehicle responses.
- Circuit models. At a more detailed level, circuit models capture electrical (and thermal) interactions within the inverter and motor. They predict current harmonics, switching losses, and temperature profiles at the semiconductor junctions and windings. These models are essential for design validation and durability analysis.
- Data-driven models. To complement the physics-based approaches, Artificial Neural Networks (ANN) and Long Short-Term Memory (LSTM) models are introduced. These models learn relationships between inputs and latent states, such as torque, speed, component temperatures, and State of Health (SoH). They provide accurate predictions even when system nonlinearity or parameter uncertainty is significant.

The combination of these three model types forms a hybrid architecture that balances interpretability, computational efficiency, and accuracy.

The **data and services layer** is responsible for the collection, storage, processing, and management of all information used by the DT. It handles time-series data ingestion from physical rigs, manufacturing systems, and vehicles. Each dataset is labeled, synchronized, and versioned to maintain traceability. Machine Learning Operations (MLOps) processes support automated training, validation, and deployment of data-driven models. A centralized twin registry ensures each digital instance is uniquely identifiable, with configuration metadata linking it to its physical counterpart. This layer serves as the backbone of the digital thread, connecting engineering, manufacturing, and operational data throughout the product lifecycle.

The **application layer** defines how the DT adds tangible value to engineering, production, and service processes. It supports several use cases, including:

- Design verification and optimization. Engineers can perform what-if analyses, evaluate new control algorithms, and optimize inverter or motor parameters without physical prototypes.

- Controller evaluation. The twin enables testing of FOC, DTC, and sensorless control methods under variable conditions before hardware deployment.
- Manufacturing analytics. At the EoL stage, the twin compares real test data against expected signatures to identify deviations and predict potential faults.
- Fault detection and prognostics. Continuous monitoring and comparison between predicted and measured states allow early identification of faults such as stator inter-turn short circuits and inverter phase imbalance.
- OTA optimization. In operational vehicles, the twin supports OTA calibration updates, adaptive control tuning, and predictive maintenance strategies.

This top layer transforms data and models into actionable insights, closing the loop from design to operation and back, ensuring the propulsion system remains optimized throughout its life cycle.

The **communication layer** serves as the interface that enables continuous, reliable data flow between the physical components, virtual models, and data infrastructure of the DT. It establishes the connectivity framework that synchronizes measurements, commands, and model outputs across distributed environments. This layer incorporates multiple communication technologies tailored to the system's requirements for bandwidth, latency, and security. Within test benches and production systems, real-time communication protocols such as CAN, Ethernet, and OPC UA ensure deterministic data exchange between sensors, controllers, and local servers. In manufacturing and plant networks, MQTT and REST APIs enable efficient data transfer to centralized data lakes and cloud-based analytics. For in-vehicle and cloud-edge synchronization, Data Distribution Service (DDS) or SOME/IP provides scalable, event-driven communication that is well-suited to the dynamic nature of SDEV. Security measures, encryption, authentication, and digital signatures, are embedded in the layer to comply with ISO 21434 cybersecurity standards. By managing data routing, quality of service (QoS), and timing synchronization, the communication layer ensures that DT operations remain accurate, up-to-date, and secure throughout the design, manufacturing, and runtime phases.

TWIN INSTANCES

The DT ecosystem for the EPDS comprises several instances, each tailored to a specific phase of the vehicle lifecycle. These instances, Development DT, Production DT, and Maintenance/Usage DT, as shown in Figure 2, are interconnected through a continuous digital thread that ensures information consistency and feedback across domains. Each DT instance plays a unique yet complementary role. *Development DT* enables virtual prototyping and system optimization, *Production DT* ensures product quality and process efficiency, and the *Maintenance/Usage DT* delivers in-service monitoring and continuous improvement. Together, they establish a fully integrated, lifecycle-spanning framework for data-driven decision-making in SDEV.

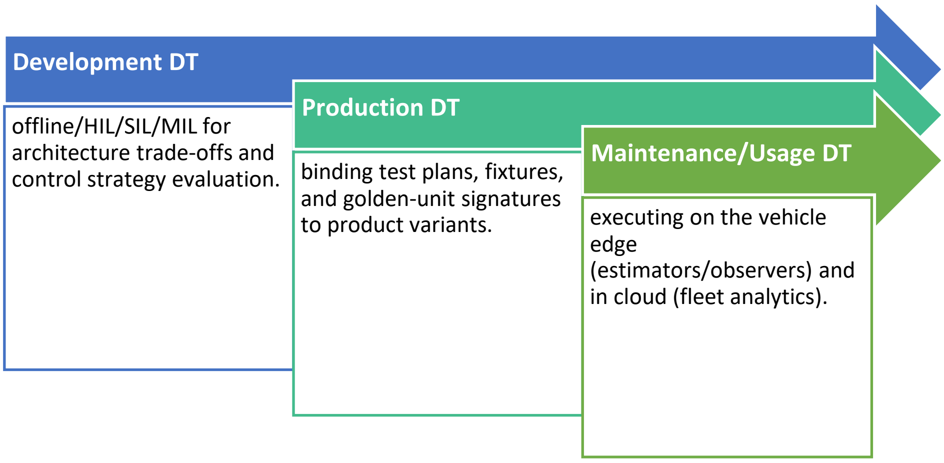

Figure 2 Digital Twin (DT) Instances Across the Vehicle Lifecycle.

The **Development DT** operates primarily in offline environments and supports Model-in-the-Loop (MIL), Software-in-the-Loop (SIL), and Hardware-in-the-Loop (HIL) configurations. It is used during the concept and development stages to explore architectural trade-offs, evaluate control strategies, and optimize powertrain parameters before physical prototypes are available. Through simulation, engineers can validate the performance of inverter and motor models under various operating conditions, assess the robustness of control algorithms such as FOC and DTC, and predict system behavior during transient events. The Design DT provides a low-risk environment for innovation and accelerates calibration by integrating real test data progressively as the system matures.

The **Production DT** bridges the virtual and physical domains in the production environment. It binds detailed test plans, calibration parameters, and equipment configurations, such as EoL testers, fixtures, and instrumentation, to individual product variants. Each production unit is validated against a digital reference model, often called the "golden unit," that represents ideal electrical and thermal signatures. By comparing live EoL data with the twin's predicted values, manufacturing engineers can detect deviations, optimize test duration, and identify potential faults early in the process. This instance also supports statistical process control and continuous improvement by feeding production data back into design and process optimization loops.

The **Maintenance/Usage DT** functions once the vehicle enters service, executing across both the edge (in-vehicle systems) and cloud environments. At the vehicle edge, lightweight estimators and observers use real-time sensor data to estimate unmeasured variables, such as torque, speed, temperature, and component degradation. In parallel, the cloud-based twin aggregates fleet data to perform advanced analytics, enabling predictive maintenance, fault diagnosis, and performance benchmarking across vehicle populations. The Maintenance/Usage DT is continuously updated via field data feedback and supports OTA updates, enabling adaptive control tuning and model recalibration without disrupting vehicle operation. This capability transforms the propulsion system into an intelligent, self-optimizing entity throughout its operational life.

INTEGRATION INTO SDEV DESIGN

Integrating the DT of the EPDS into the SDEV design phase enables a seamless transition from concept to implementation. This stage focuses on building accurate models, developing and validating controllers, planning experiments, and ensuring data traceability and model fidelity. DT acts as both a predictive simulation tool and a validation reference, ensuring that design decisions are data-driven, verifiable, and reproducible across the engineering lifecycle. Figure 3 illustrates the continuous and cyclic workflow used to integrate the DT of the EPDS into the SDEV design process.

Figure 3 Iterative Digital Twin (DT) Development Process in Electric Propulsion Drive System (EPDS) Design. (image generated by AI)

A well-defined **modeling strategy** is fundamental to the success of the DT framework. The recommended approach is hybrid modeling, combining physics-based and data-driven methods, to capture both fundamental dynamics and system-specific nonlinearities. Baseline physics models, such as PMSM representations in the d–q reference frame and averaged inverter models. These models are accurate enough for most control and performance studies while remaining computationally efficient. To enhance realism and address system behaviors not fully described by analytical equations, learned surrogate models based on ANNs or LSTM networks are incorporated. These data-driven components handle complex nonlinear relationships, parameter uncertainties, and proprietary control features that component suppliers may not disclose. To ensure scalability and efficiency, model fidelity must be adapted to each use case. For early concept evaluations, low-fidelity models (with millisecond-level time steps) enable fast trade-off studies and architecture optimization. For control design, high-fidelity models (microsecond-level) capture pulse-width modulation (PWM) switching and transient behavior. Thermal models, operating on slower timescales (seconds to minutes), support thermal management and derating assessments. The drivability and system-level models run at the vehicle level to evaluate user experience, energy efficiency, and system integration effects. This allows engineers to balance computational load and accuracy according to the needs.

DT serves as a virtual testbed for **controller development**, enabling safe, cost-effective evaluation of control strategies across a wide range of operating conditions. It can simulate both FOC and DTC algorithms, enabling direct comparison in terms of torque ripple, efficiency, and dynamic response. Flux-weakening regions, torque-speed transitions, and system responses to voltage sags or thermal drift can all be explored in simulation long before hardware is available. Additionally, the DT facilitates the validation of sensorless observers, assessing their robustness to noise, load variations, and parameter drift. From this foundation, engineers can derive virtual sensors—software-based estimators for parameters such as torque, speed, and temperature—that complement or replace physical sensors. These virtual sensors enhance redundancy and resilience, reducing sensor costs and improving system fault tolerance. Because these estimators can be updated and refined using field-collected data, they also form part of the OTA deliverable feature set, ensuring continuous improvement and adaptability of control functions throughout the vehicle's lifespan.

Effective **test planning and structured data management** are essential to ensure that model development, calibration, and validation processes remain consistent and traceable. The DT supports this by defining comprehensive Design-of-Experiments (DoE) matrices that specify operating points and test conditions across the torque–speed map. These experiments are executed on dynamometer setups or battery emulators, covering operating conditions from no-load to approximately 80% of full load. During each test, key variables, including DC bus voltage and current, three-phase currents and voltages, resolver signals, and thermal node temperatures, are recorded. The data must be time-aligned and annotated with configuration metadata, including model version, software revision, and environmental parameters. To ensure traceability, every artifact, from requirements through test cases, models, and validation results, should be version-controlled within Product Lifecycle Management (PLM) or Application Lifecycle Management (ALM) systems. Each dataset and model is assigned a unique identifier, forming a digital thread that links simulation, testing, and validation results across teams and domains.

The **validation and calibration** phase closes the loop between simulation and reality. DT's predictions are compared with bench-test results using quantitative metrics such as torque, speed, efficiency, and temperature error bands. Validation ensures that each model behaves within acceptable tolerances, confirming both the accuracy and reliability of the simulation environment. Once validated, the models undergo continuous calibration as new bench and in-vehicle data become available. This iterative process keeps the twin synchronized with its physical counterpart, reflecting component aging, software updates, and evolving operating conditions. Model release gates should be governed by confidence metrics, such as the Normalized Root Mean Square Error (NRMSE) between predicted and measured torque, speed, and junction temperature. Only models that meet or exceed the defined thresholds are approved for deployment within design or control workflows. This feedback-driven calibration process establishes the DT as a living, evolving representation of the propulsion system, one that continuously improves with every design iteration and real-world data update.

INTEGRATION INTO MANUFACTURING

Integrating the DT of the EPDS into the manufacturing environment ensures that every production unit meets its design intent and quality requirements. The DT becomes an operational tool for validating system performance, optimizing production processes, and maintaining traceability from design to assembly. This integration is structured around three key domains: production twin, process optimization, and data operations.

Production DT serves as the virtual counterpart of the physical testing environment in manufacturing, particularly during the EoL verification stage. For each propulsion system variant, a dedicated twin configuration is created, containing reference models and expected performance characteristics. These include reference waveforms, torque–speed signatures, inverter self-test sequences, and thermal soak profiles, all derived from validated design and prototype data. During EoL testing, the measured signals from sensors, such as current, voltage, torque, and temperature, are continuously compared against predictions generated by the twin. This comparison produces residuals, which represent the difference between actual and expected values. Real-time anomaly scoring algorithms evaluate these residuals to identify deviations from nominal behavior. By analyzing these deviations, the Production DT can automatically flag potential defects, such as rotor or stator imbalances, inverter phase asymmetries, and encoder or resolver misalignments. This capability allows rapid root-cause identification and classification of faults, reducing the number of false fails and ensuring that only genuine issues trigger corrective actions. Production DT therefore acts as both a quality assurance system and a diagnostic assistant at the manufacturing stage.

Beyond product validation, the DT also contributes to *process efficiency*. Through continuous learning from production data, it can identify patterns and optimize test execution strategies. One key application is reducing test cycle time through adaptive test sequencing. Instead of running all tests sequentially, the DT can estimate an early pass/fail probability based on partial data and dynamically skip redundant checks when a unit clearly meets specifications. This approach significantly improves throughput without compromising test integrity. The DT also facilitates automated rework instructions. When anomalies are detected, the twin can trace the source of deviation, whether it originates from the product, test fixture, or cabling, and automatically generate step-by-step corrective actions. This not only accelerates troubleshooting but also ensures consistency in maintenance procedures across production lines. Moreover, the DT monitors station health by observing long-term drifts in equipment behavior. For example, changes in load motor drive performance, sensor calibration, or power analyzer accuracy can be detected through persistent deviations in twin residuals. By introducing canary parts, known reference units periodically tested on each station, the system can differentiate between product faults and equipment degradation, maintaining production reliability and test repeatability.

Efficient *data management* is critical for sustaining the integrity of the manufacturing DT. All measurement signals and test outcomes must adhere to a standardized data schema to ensure compatibility and traceability across systems. Each record includes absolute timestamps, per-cycle counters, and metadata describing the product variant, configuration, and testing conditions.

Raw signals are stored alongside engineered features (such as harmonic content, RMS values, and temperature deltas) and labels derived from test verdicts and anomaly classifications. This structured data foundation enables downstream analytics, fleet correlation, and the training of predictive models for future designs. Moreover, each dataset is linked to its corresponding serial number and DT instance, ensuring complete traceability across the digital thread, from design simulations and prototype testing to production and in-service monitoring. This closed-loop data structure supports continuous improvement, allowing insights from manufacturing and field performance to refine design models, test procedures, and operational strategies.

By deploying the Production DT within manufacturing, the organization gains a powerful mechanism for quality assurance, process efficiency, and data-driven feedback. The combination of real-time anomaly detection, adaptive testing, and standardized data management establishes a seamless connection between the physical factory floor and the virtual engineering environment. As a result, manufacturing operations become more predictive, efficient, and resilient, core characteristics of the SDEV production paradigm.

INTEGRATION INTO MAINTENANCE & OPERATIONS

The integration of the DT into the maintenance and operational phases of the SDEV lifecycle transforms vehicle support from reactive servicing to a predictive and data-driven process. The DT operates in real time, processing both on-board sensor data and fleet-level analytics to monitor component health, detect anomalies, and guide service actions. This integration is realized through three key components: the edge/runtime twin, prognostics, and after-sales workflows.

The **edge** or **runtime** DT is the in-service counterpart of the vehicle's propulsion system. It functions directly on the vehicle's embedded hardware or at the edge computing layer, providing real-time monitoring, fault detection, and performance estimation. This twin deploys lightweight observer models, often implemented as ANNs or LSTM estimators, which run efficiently on vehicle control units or connected edge devices. These observers act as virtual sensors, estimating parameters that are difficult or costly to measure directly, such as torque, rotor speed, and component temperatures. By continuously comparing model predictions to real sensor readings, the twin identifies discrepancies, known as residuals, that signal potential degradation or faults. Runtime DT also detects specific fault modes, including stator inter-turn short circuits, phase current imbalances, and resolver or encoder misalignments, which might not be easily visible in traditional diagnostic systems. When anomalies are identified, the system evaluates their severity and trend over time, enabling condition-based maintenance (CBM) rather than scheduled or reactive servicing. If residuals exceed defined thresholds or follow concerning trends, DT automatically triggers OTA diagnostics or sends service bulletins to fleet management systems. This allows maintenance actions to be planned before failures occur, minimizing downtime and optimizing service scheduling across the fleet. Over time, continuous synchronization between edge devices and cloud twins ensures that models remain accurate and adaptive to changing vehicle behavior, component aging, and environmental conditions.

The **prognostics layer** of the DT extends beyond fault detection to predict the Remaining Useful Life (RUL) of critical components within the propulsion system. Using fleet-wide operational data collected from runtime twins, data-driven RUL models are trained to forecast degradation patterns and failure probabilities. For example, power semiconductor components in the inverter experience stress from repetitive junction-temperature cycling, which can lead to wear-out over time. The DT analyzes temperature histories, switching patterns, and thermal gradients to estimate semiconductor lifespan. Similarly, bearings and insulation systems within electric machines are monitored using vibration, temperature, and current harmonics, which serve as indirect indicators of health. By combining physics-based degradation models with data-driven predictions, the DT can provide highly accurate RUL estimates. These predictions are integrated into maintenance planning systems, enabling proactive scheduling of service windows. This ensures that parts are replaced or refurbished before failures occur, improving vehicle uptime, reducing warranty costs, and enhancing customer trust. Furthermore, aggregated RUL data across the fleet allows manufacturers to identify systemic reliability issues, improve component design, and refine control strategies to extend component life in future vehicle generations.

The DT also reshapes **after-sales** and **service operations**, empowering technicians and service centers with advanced diagnostic tools. One of the most practical implementations is the *Technician DT*, a localized, portable instance of DT that can be used in workshops or field environments. Technician DT provide guided diagnostic scripts that replay recorded DT scenarios, simulating how the propulsion system behaves during fault conditions. This helps technicians reproduce issues accurately and test hypotheses before performing physical repairs. Each diagnostic session can automatically generate a case report, summarizing key residuals, fault classifications, and probable root causes. This level of automation streamlines troubleshooting, minimizes human error, and ensures consistent maintenance decisions across service locations. It also creates a feedback channel from service operations back to engineering, where real-world fault data and diagnostic outcomes are used to refine models, update prognostic thresholds, and improve future vehicle releases. Ultimately, integrating DT technology into after-sales workflows turns every service event into a learning opportunity, thereby strengthening the overall accuracy and resilience of the DT ecosystem.

By combining the abovementioned workflows, DT enables a closed-loop maintenance system that continuously evolves with vehicle use. Real-time sensing, predictive analytics, and automated diagnostics replace traditional static maintenance schedules with intelligent, adaptive servicing strategies. The result is a propulsion system that remains efficient, reliable, and continuously optimized throughout the operational life of the SDEV.

METHODS & TOOLING RECOMMENDATIONS

Table 1 summarizes the recommended software tools, frameworks, and standards required to support the lifecycle deployment of the DT for the EPDS based on prior project PSG453 and will be implemented in PRG2532. It highlights the integration of modeling, data-driven analytics, middleware, and governance mechanisms to ensure the DT remains accurate, secure, and maintainable from design through runtime.

Table 1 Methods and Tooling Recommendations for Digital Twin (DT) Implementation into Electrical Propulsion Drive System (EPDS) of Electric Vehicle (EV)

Category	Recommended Tools and Frameworks	Purpose and Application in the EPDS DT
Modeling	MATLAB/Simulink for control algorithm development and d–q domain machine modeling;	Enables creation of multi-fidelity models for control, electrical, and mechanical domains. These tools form the foundation for model-based design, control validation, and system-level integration of the EPDS.
Data-Driven Methods	Python environment with ANN and LSTM frameworks; supported by MLOps pipelines with model registry and continuous integration (CI).	Facilitates the development of surrogate models and estimators for torque, speed, temperature, and degradation prediction. MLOps ensures consistent version control, retraining, and deployment of data-driven models across edge and cloud environments.
Platforms	Cloud/edge-based DT platforms capable of managing twin graphs, time-series data, and OTA updates. Simulation-based tools such as Ansys Twin Builder are recommended for physical and hybrid twin integration.	Provides the computational and data infrastructure to host, synchronize, and scale DT instances across design, manufacturing, and runtime stages.
Middleware	ROS2, UDP, and CAN for test rig streaming; OPC UA and MQTT for plant-level communication; in-vehicle protocols such as SOME/IP or DDS for runtime data exchange.	Ensures reliable, real-time communication between the physical system, virtual models, and cloud services. These middleware solutions standardize data flow, enabling synchronization across heterogeneous systems.
Governance and Standards	ISO 26262 for functional safety (ASIL assignment for estimator functions), ISO 21434 for cybersecurity, data privacy policy adherence, and	Establishes compliance, safety, and traceability across the DT lifecycle, ensuring that models, data, and

Category	Recommended Tools and Frameworks	Purpose and Application in the EPDS DT
	calibration/versioning management rules.	software updates meet industry and regulatory standards.

KPIS AND ACCEPTANCE CRITERIA

Establishing clear, measurable Key Performance Indicators (KPIs) and acceptance criteria is essential for evaluating the effectiveness of DT integration throughout the SDEV lifecycle, as shown in Figure 4. These metrics ensure that the DT delivers tangible improvements in design validation, manufacturing efficiency, and operational performance. KPIs should be continuously tracked and reviewed as part of the digital thread governance framework to maintain system reliability, quality, and value creation.

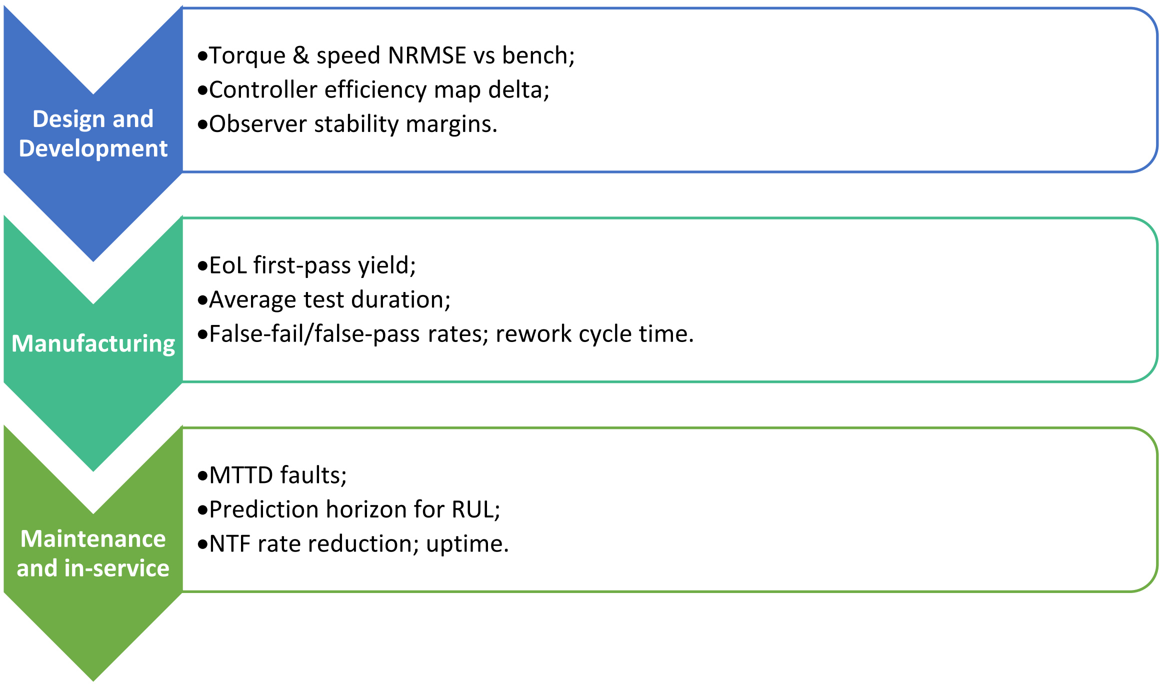

Figure 4 Key Performance Indicators (KPIs) Across the Digital Twin (DT) Lifecycle.

During the **design and development phase**, KPIs focus on model accuracy, controller performance, and system stability.

- Torque and Speed NRMSE vs. Bench Measurements. The NRMSE between simulation outputs and physical test bench data serves as a quantitative measure of model fidelity. Low NRMSE values indicate that the DT accurately reproduces physical system behavior, thereby increasing confidence in virtual validation and reducing reliance on extensive prototyping.
- Controller Efficiency Map Delta. This metric compares the efficiency maps derived from DT simulations to those measured on physical test rigs. The smaller the deviation (delta), the more reliable the DT becomes as a tool for optimizing control strategies such as FOC or DTC.
- Observer Stability Margins. For virtual sensors and estimators used in control, stability margins, such as phase and gain margins, ensure robustness under varying conditions. Acceptance criteria should verify that observers maintain stable performance across temperature ranges, speed transitions, and component tolerances.

Together, these KPIs validate the DT as a credible simulation and design-optimization environment capable of accelerating innovation while maintaining engineering rigor.

In the **manufacturing phase**, KPIs measure production quality, testing efficiency, and diagnostic accuracy.

- EoL First-Pass Yield. This metric reflects the percentage of propulsion systems that pass EoL testing on their first attempt. A high first-pass yield indicates effective calibration of the DT models and accurate fault detection thresholds, leading to reduced rework and improved throughput.
- Average Test Duration. DT enables optimization of test procedures through adaptive sequencing. A measurable reduction in average test duration demonstrates that the DT is successfully reducing cycle time without compromising quality or coverage.
- False-Fail and False-Pass Rates. These rates quantify the accuracy of the DT-driven anomaly detection system. Low false-fail rates prevent unnecessary rework, while low false-pass rates ensure that defective units are not mistakenly released.
- Rework Cycle Time. When rework is required, this metric measures the time from defect detection to resolution. A shorter rework cycle time indicates effective fault localization and improved technician guidance, often supported by insights from the production twin.

Collectively, these metrics confirm DT’s role in enhancing manufacturing precision, process efficiency, and product reliability.

For **maintenance and in-service operations**, KPIs evaluate fault-detection capability, predictive accuracy, and fleet performance.

- Mean Time to Detect (MTTD) Faults. This measures how quickly the DT identifies emerging faults compared to traditional diagnostics. A reduced MTTD indicates earlier fault detection and prevention of secondary failures.
- Prediction Horizon for RUL. The DT’s predictive models estimate how far in advance it can forecast a component’s degradation or end-of-life event. Longer and more accurate prediction horizons improve maintenance scheduling and resource allocation.
- No-Trouble-Found (NTF) Rate Reduction. NTF cases—when a reported fault cannot be replicated during service—are costly and time-consuming. The DT reduces NTF rates by improving fault traceability and providing data-driven root cause insights.
- System Uptime. Uptime reflects the overall availability of the vehicle or fleet. A higher uptime percentage demonstrates DT’s effectiveness in maintaining operational reliability through proactive diagnostics and predictive maintenance.

These metrics collectively validate the Runtime as a cornerstone of CBM and lifecycle optimization.

By monitoring KPIs across design, manufacturing, and maintenance, organizations can quantitatively assess the maturity and effectiveness of their DT implementation. Success is measured not only by model accuracy or process efficiency but also by tangible business outcomes, faster development cycles, reduced operational costs, and improved vehicle reliability. These acceptance criteria ensure that the DT continues to deliver measurable value throughout the full lifecycle of the SDEV.

POSSIBLE IMPLEMENTATION ROADMAP (PHASED)

Figure 5 outlines the sequential phases for implementing the DT of the EPDS within the SDEV lifecycle. Each phase builds progressively on the previous one, ensuring a structured deployment of modeling, data, and analytics capabilities.

- Phase 0 – Foundations. Establish core infrastructure, including data pipelines, standardized metadata schemas for test benches, and a version-controlled model repository. Select the DT platform and edge runtime environment.
- Phase 1 – Design Twin. Develop machine and inverter models, collect Design of Experiments (DoE) datasets, and train torque and speed estimators. Validate results against bench data and introduce controller performance comparisons.
- Phase 2 – Production Twin. Integrate EoL testers with the twin, create golden reference signatures, and implement residual-based anomaly detection. Optimize test sequences through adaptive logic to improve efficiency.
- Phase 3 – Runtime Twin. Deploy virtual sensors and fault monitoring algorithms in vehicles, initiate fleet learning programs, and enable OTA model calibration updates.
- Phase 4 – Prognostics & Optimization. Incorporate RUL models for power electronics and electric machines. Close the feedback loop by translating field issues into new design requirements and refining manufacturing test procedures.

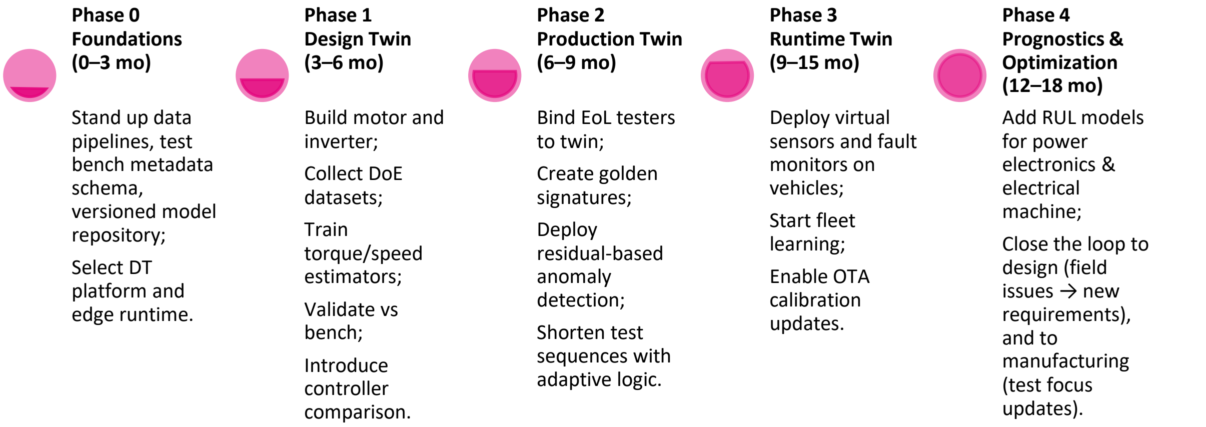

Figure 5 Phased Implementation Roadmap for the Electric Propulsion Drive System (EPDS) Digital Twin (DT).

This phased roadmap provides a clear timeline and structure for DT integration, from infrastructure setup and design modeling to production validation, fleet analytics, and lifecycle optimization. It ensures that each stage contributes measurable value while maintaining alignment between physical systems and their digital counterparts.

SAFETY, CYBERSECURITY, AND COMPLIANCE

Incorporating functional safety and cybersecurity into the DT framework is not optional, but it is a prerequisite for trust and adoption in production-grade SDEV. By aligning with ISO 26262, ISO/SAE 21434, and the lifecycle governance principles. DT interacts closely with the control systems of the EPDS, it must be treated as a safety-relevant software element and governed under the same functional safety and cybersecurity frameworks that apply to embedded automotive systems. Compliance must extend across the full lifecycle, from model design and validation to deployment, operation, and OTA updates.

Functional Safety and Software Integrity. The DT outputs that influence or support control decisions, such as torque estimations, temperature predictions, or fault detection, should be classified as Software Safety Elements under ISO 26262. This requires performing an Automotive Safety Integrity Level (ASIL) assessment and decomposition to determine the safety level each DT function must achieve based on its potential impact on vehicle operation and occupant safety. A structured Failure Modes, Effects, and Diagnostic Analysis (FMEDA) should be conducted to identify and quantify the risks associated with DT estimators, observers, or virtual sensors. Each potential failure mode—such as an incorrect torque estimate, delayed thermal response, or corrupted model output—must be analyzed for severity, occurrence, and detectability. The DT architecture should include safety mechanisms such as range monitoring, plausibility checks, and model confidence metrics to detect and isolate unsafe outputs.

Redundancy and fault containment should be designed into the DT environment. For example:

- Independent physical sensor readings should periodically cross-validate virtual sensor outputs.
- Fallback strategies should be in place to revert to conservative control modes if the DT detects inconsistency or model degradation.
- Safety-critical estimations (e.g., torque or speed) should be separated from non-critical analytics functions using ASIL partitioning or software sandboxing.

This layered safety approach ensures that DT-assisted control operations do not compromise vehicle integrity, even in the presence of degraded data or model faults.

Cybersecurity and Data Protection. From a cybersecurity perspective, DT is part of a connected data ecosystem spanning the vehicle, edge, and cloud. Consequently, it must comply with ISO/SAE 21434 for automotive cybersecurity management. All data flows, from real-time vehicle telemetry to cloud analytics, should be secured end-to-end using encryption, mutual authentication, and integrity verification. The communication layer between the DT instances (vehicle, plant, and cloud) should implement secure transport protocols (e.g., TLS over MQTT or DDS Secure) to protect against eavesdropping and tampering. Each DT instance, whether local or cloud-based, should possess a unique digital identity and operate within a zero-trust architecture that verifies access before every interaction. To protect model integrity, all DT models and configuration files must be cryptographically signed before deployment. This prevents unauthorized or corrupted models from being loaded into runtime environments. For OTA updates, rollback mechanisms must be implemented to restore previous stable versions of the twin or its models in case of unexpected behavior after an update. Continuous monitoring of software and firmware integrity should be conducted via secure boot mechanisms and runtime attestation. Data collected from vehicles, test benches, and production systems should be anonymized or pseudonymized before external storage or analysis to comply with data privacy regulations such as the General Data Protection Regulation (GDPR). Access control policies must ensure that sensitive data, including operational parameters, fleet identifiers, and proprietary model structures, are available only to authorized personnel and systems.

Compliance and Lifecycle Governance. Compliance throughout the DT lifecycle requires an integrated governance framework that covers functional safety, cybersecurity, and data management. It is important to emphasize the importance of traceability, ensuring that every software element, model version, and dataset can be linked to its verification evidence, configuration history, and release approval. Key practices include:

- Maintaining a DT Safety Case, documenting the rationale, validation results, and verification evidence supporting the safe operation of DT components.
- Applying change control procedures for model and data updates, ensuring that each modification undergoes hazard assessment and regression testing before deployment.
- Conducting periodic audits and penetration tests on cloud and OTA infrastructures to identify vulnerabilities.
- Integrating safety and cybersecurity assessments into the Continuous Integration/Continuous Deployment (CI/CD) pipeline used for DT model delivery.

Moreover, it is important to establish cross-functional safety and cybersecurity boards that oversee risk analysis, incident response, and regulatory compliance. These boards should clearly define ownership of DT safety and security across engineering, IT, and operations teams.

LIST OF PUBLICATIONS THAT UNDERPIN THE RECOMMENDATIONS

No.	Reference	Keywords / Core Topics
1.	Rassõlkin, A.; Rjabtšikov, V.; Vaimann, T.; Kallaste, A.; Kuts, V.; Demidova, G. (2020). Digital Twin Data Handling for Propulsion Drive System of Autonomous Electric Vehicle: Case Study. IEEE RTUCON.	Digital twin, propulsion drive, data handling, autonomous EV
2.	Ibrahim, M.; Rassõlkin, A. (2025). Hybrid-Driven Digital Twin Modelling Framework for an EV Propulsion Drive System. IET Intelligent Transport Systems.	Hybrid digital twin, EV propulsion, modelling framework
3.	Rjabtšikov, V.; Rassõlkin, A.; Vaimann, T.; Kallaste, A.; Kuts, V. (2020). Concept of the Test Bench for Electrical Vehicle Propulsion Drive Data Acquisition. IEEE ICEPDS2020.	Test bench, EV propulsion, data acquisition
4.	Rjabtšikov, V.; Rassõlkin, A.; Kudelina, K.; Kallaste, A.; Vaimann, T. (2023). Review of EV Testing Procedures for Digital Twin Development. Energies.	EV testing, digital twin, procedures review
5.	Ibrahim, M.; Järg, O.; Seppago, R.; Rassõlkin, A. (2025). Performance Optimization of a PMSM Drive for Formula EV. Sensors.	PMSM optimization, Formula EV, performance modelling
6.	Rjabtšikov, V.; Ibrahim, M.; Asad, B.; et al. (2023). Digital Twin Service Unit for EV Induction Motor Fault Detection. IEEE IEMDC2023.	Digital twin, induction motor, fault detection
7.	Naseer, M.U.; Kallaste, A.; Asad, B.; Vaimann, T.; Rassõlkin, A. (2022). Analytical Modelling of Synchronous Reluctance Motor. IET Electric Power Applications.	SynRM modelling, analytical model
8.	Rassõlkin, A.; Rjabtšikov, V.; Vaimann, T.; Kallaste, A. (2020). Digital Twin of an Electrical Motor Based on Empirical Model. IEEE ICEPDS2020.	Digital twin, empirical motor model
9.	Jegorov, S.; Rassõlkin, A.; Rjabtšikov, V.; Ibrahim, M.; Kuts, V. (2022). Novel Digital Twin Concept for Industrial Propulsion Drive. ASME IMECE2022.	DT concept, industrial propulsion

No.	Reference	Keywords / Core Topics
10.	Ibrahim, M.; Rassõlkin, A.; Vaimann, T.; Kallaste, A. (2024). Propulsion Drives and Control Algorithms of EVs. Elsevier Handbook.	EV propulsion, control algorithms
11.	Vaimann, T.; Rassõlkin, A.; Belahcen, A.; et al. (2020). AI in Monitoring and Diagnostics of Energy Conversion Systems. IEEE IWED2020.	AI monitoring, diagnostics
12.	Gilbert Zequera, R.A.; Rassõlkin, A.; Vaimann, T.; Kallaste, A. (2023). Data Science-based Techniques for Battery Modelling and Diagnostics. IEEE EDPE2023.	Battery diagnostics, data science
13.	Zequera, R.A.G.; Rassõlkin, A.; Belolipetskaja, D. (2025). Remote Data Transfer through PyBaMM in Battery Applications. IEEE ICDCM2025.	Battery modelling, PyBaMM, remote data
14.	Kudelina, K.; Vaimann, T.; Rassõlkin, A.; Kallaste, A. (2021). Impact of Bearing Faults on Vibration of BLDC Motor. IEEE IECON2021.	BLDC motor, bearing faults
15.	Valme, D.; Rassõlkin, A.; Liyanage, D.C. (2025). Review of Hyperspectral Imaging on Mobile Robots. Sensors.	Hyperspectral imaging, robotics
16.	Orosz, T.; Pánek, D.; Rassõlkin, A.; Kuczmann, M. (2022). Robust Design Optimization of Electrical Machines. Electronics.	Optimization, electrical machines
17.	Wang, R.; Sell, R.; Rassõlkin, A.; Otto, T. (2020). Intelligent Functions Development on Autonomous EV Platform. Journal of Machine Engineering.	Autonomous EV, intelligent functions
18.	Ibrahim, M.; Raja, H. A.; Rassõlkin, A.; Vaimann, T.; Kallaste, A. (2023). An EV-Traction Inverter Data-Driven Modelling for Digital Twin Development. IEEE EPE2023.	EV inverter, data-driven modelling, DT
19.	Ibrahim, M.; Rjabtšikov, V.; Jegorov, S.; Rassõlkin, A.; Vaimann, T.; Kallaste, A. (2022). Conceptual Modelling of an EV-PMSM Digital Twin. IEEE PEMC2022.	PMSM digital twin, conceptual modelling