

STUDY ON DYNAMIC VALUE CHAIN OPTIMIZATION FOR THE  
ENGINEERING-TO-ORDER INDUSTRY

By

David Y. Kim, B.S., M.S.

DISSERTATION

Presented to the Swiss School of Business and Management Geneva

In Partial Fulfillment

Of the Requirements

For the Degree

DOCTOR OF BUSINESS ADMINISTRATION

SWISS SCHOOL OF BUSINESS AND MANAGEMENT GENEVA

DECEMBER, 2025

STUDY ON DYNAMIC VALUE CHAIN OPTIMIZATION FOR THE  
ENGINEERING-TO-ORDER INDUSTRY

by

David Y Kim

APPROVED BY

*Iva Buljubic*

---

Dr. Dejan Curovic, Chair

RECEIVED/APPROVED BY:

*Renee Goldstein Osmic*

---

Admissions Director

# ABSTRACT

Engineer-to-Order (ETO) environments are characterised by high variability, frequent engineering change, supplier uncertainty, and late-stage field constraints, making traditional linear value-chain and integration-centric execution models increasingly inadequate. While Industry 4.0 initiatives have improved digital connectivity through BIM, ERP, and MES integration, persistent coordination failures and lifecycle blind spots remain unresolved.

This study proposes and evaluates a Dynamic Value Chain Optimisation (DVCO) framework designed to address these limitations through semantic integration, knowledge-graph-based dependency modelling, and AI-enabled reasoning. Using Symbotic as a reference case, the research develops a formal ETO ontology, instantiates a cross-domain knowledge graph, and integrates large language models (LLMs) as semantically constrained decision-support agents. A structured, simulation-driven evaluation compares baseline ETO execution behaviour with a DVCO-enabled environment across engineering, supply chain, manufacturing, installation, and commissioning stages.

The findings demonstrate that DVCO significantly improves lifecycle traceability, change-impact visibility, predictive risk awareness, and cross-functional coordination. Results further show that engineering change and supplier variability are the dominant drivers of non-linear risk propagation, and that DVCO delivers its greatest performance gains in downstream installation and commissioning, where recovery from coordination failure is most

difficult. The study concludes that DVCO represents a scalable, technology-agnostic orchestration paradigm capable of transforming ETO execution from a reactive, siloed process into a semantically aligned, model-driven, and AI-supported value chain.

## ACKNOWLEDGEMENTS

I wish to express my deepest appreciation to Dr. Dejan Curovic, whose intellectual guidance, critical insight, and patient mentorship were invaluable throughout the development of this dissertation. His scholarly discipline and thoughtful direction have profoundly shaped both the methodological rigor and conceptual framing of this research.

I also extend my sincere gratitude to Professor Jeongsik Lee, my colleagues and collaborators at Symbotic and AI&S LLC, whose deep expertise in SAP, robotics, data integration, and analytics informed and enriched the framework presented herein. Their practical experience and technical excellence provided meaningful grounding between theory and real-world application.

This work is also the culmination of over 40 years of global industry experience as a techno-functional architect across diverse sectors including manufacturing, automation, aerospace, energy, pharmaceuticals, and digital transformation. I am grateful to the many engineers, IT leaders, executives, and operational specialists I have worked with over the decades — each of whom contributed in some measure to the knowledge embodied in this study.

To my family, I offer heartfelt thanks for their unwavering support, encouragement, and belief in my calling and purpose. Finally, and most importantly, I give thanks to God, whose wisdom, guidance, and grace sustained me throughout this journey and made the completion of this work possible.

# Table of Contents

ABSTRACT.....	3
ACKNOWLEDGEMENTS.....	5
Table of Contents.....	6
LIST OF TABLES.....	14
LIST OF FIGURES.....	15
CHAPTER I.....	16
Introduction.....	16
1.1 Background and Context.....	16
1.2 Problem Statement.....	17
1.3 Research Aim and Objectives.....	19
1.4 Research Questions.....	22
1.5 Significance of the Study.....	23
1.5.1. A new semantic and model-driven architecture for ETO execution.....	24
1.5.2. Demonstration of AI-assisted reasoning in complex project environments.....	24
1.5.3. Bridging engineering, planning, manufacturing and operational domains.....	24
1.5.4. A scalable model for next-generation warehouse-automation deployments.....	25
1.5.5. Actionable insights for digital transformation in the ETO industry.....	25
1.6 Dissertation Structure.....	26
CHAPTER II.....	29
LITERATURE REVIEW.....	29
2.1 Introduction.....	30
2.2 Engineer-to-Order (ETO) Value Chain Literature.....	31
2.2.1 Characteristics and Structural Complexity of ETO Systems.....	32
2.2.2 Common Bottlenecks Identified in ETO Literature.....	33
2.2.3 Limitations of Traditional ETO Planning and Control.....	34
Relationship to ETO Literature Limitations.....	35
2.2.4 Emergence of Digital and Model-Based Solutions in ETO Literature.....	37
2.2.5 Literature Gap: Lack of a Unified, Semantic, AI-Enabled Framework.....	38
2.3 Dynamic Nature of ETO Value Chains.....	39
2.4 Limitations of Traditional Linear Planning Models.....	45
2.4 Model-Based Systems Engineering (MBSE) and the Model-Based Enterprise (MBE)....	47

2.4.1 From Document-Based to Model-Centric Engineering.....	47
2.4.2 Principles of Model-Based Systems Engineering.....	48
2.4.3 Transition from MBSE to the Model-Based Enterprise (MBE).....	49
2.4.4 Role of BIM within the Model-Based Enterprise.....	50
2.4.5 MBSE and MBE in Engineer-to-Order (ETO) Environments.....	51
2.4.6 Contributions of MBSE and MBE to Integrated Digital Engineering.....	53
2.5 Digital Twins and Cyber-Physical Systems.....	53
2.5.1 Conceptual Foundations of Digital Twins.....	54
2.5.2 Cyber-Physical Systems (CPS) as the Digital Twin Substrate.....	56
2.5.3 Digital Twin Applications in Engineering and Operations.....	58
2.5.4 Limitations of Current Digital Twin Implementations.....	60
2.6 Ontologies, Semantics, and Knowledge Graphs in Engineering and Operations.....	62
2.6.1 The Role of Semantics in Complex Engineering Environments.....	63
2.6.2 Ontologies in Engineering, Manufacturing and Supply Chain Domains.....	65
2.6.3 Knowledge Graphs as Lifecycle Integration Structures.....	67
2.6.4 Ontology + KG as an AI Enabler.....	69
2.6.5 Applications of Ontologies and KGs in ETO and Industry 4.0 Literature.....	71
2.6.5.1 Change-Impact Analysis.....	71
2.6.5.2 Engineering–Manufacturing Alignment.....	71
2.6.5.3 Installation and Commissioning Management.....	72
2.6.5.4 Supply Chain Risk Analysis.....	72
2.6.5.5 Digital Twin Synchronisation.....	72
2.6.5.6 AI Reasoning and Natural-Language Querying.....	73
2.6.6 Current Limitations in Semantic and KG Adoption.....	73
2.6.7 Contribution of Semantic Technologies to DVCO.....	75
2.7 Artificial Intelligence and Large Language Models in Operations and Value-Chain Management.....	77
2.7.1 Evolution of AI in Operations and Manufacturing.....	78
2.7.2 Emergence of Large Language Models in Industrial Applications.....	79
2.7.3 LLMs and Knowledge Graphs: A Synergistic Architecture.....	81
2.7.4 AI-Enabled Reasoning in Engineering and Operations.....	82
2.7.5 Limitations of Current AI Approaches in ETO Contexts.....	85
2.7.6 AI as a Catalyst for DVCO.....	86
2.7.7 Summary.....	88

2.8 Literature Gap and Summary.....	89
2.8.1 Identified Gaps Across the Literature.....	89
2.8.3 Summary of Chapter 2.....	98
CHAPTER III.....	101
Conceptual Framework and DVCO Architectural Model.....	101
3.1 Introduction.....	101
3.2 Conceptual Foundations of the DVCO Framework.....	106
3.3 DVCO Architectural Model.....	108
3.3.1 Architectural Overview.....	108
3.3.2 Layered Architecture Description.....	110
3.3.3 Architectural Validation and Traceability to Findings.....	111
3.3.4 Architectural Validation and Traceability to Chapter 6 Findings.....	113
3.4 Model-Based Engineering (MBE).....	113
3.4.1 Requirements Modelling and MBSE.....	114
3.4.2 EBOM, MBOM, Routing, and Configuration Rules.....	114
3.4.3 BIM and 3D/4D/5D Models.....	114
3.4.4 Semantic and Ontology Layers.....	114
3.5 Software-Defined Execution Layer.....	115
3.5.1 Model-Linked Workflows and Orchestration Engines.....	115
3.5.2 Integration Across ERP, PLM, MES, and WMS Systems.....	115
3.5.3 Dynamic Resequencing and Constraint Management.....	116
3.5.4 Exception and Variance Handling.....	116
3.6 Real-Time Feedback Loop: IoT, Data Infrastructure, and Digital Twins.....	116
3.6.1 IoT and Operational Telemetry.....	116
3.6.2 Digital Twins for Assets and Processes.....	117
3.6.3 Closed-Loop Updates.....	117
3.6.4 Data Infrastructure.....	117
3.7 Decision Intelligence Layer.....	117
3.7.1 Predictive Analytics.....	118
3.7.2 Prescriptive Optimisation.....	118
3.7.3 Knowledge Graph Reasoning.....	118
3.7.4 AI-Enabled Control Towers.....	118
3.8 Integration Across the D2O Lifecycle.....	118
3.9 Summary.....	119

CHAPTER IV.....	121
TECHNOLOGIES AND ARCHITECTURE ENABLING DVCO.....	121
4.1 Introduction.....	121
4.2 Building Information Modelling (BIM) as the Structural and Spatial Backbone.....	124
4.2.1 BIM–SAP Integration and the Digital Thread.....	125
4.2.1.1 BIM as the Engineering Source of Truth for SAP.....	125
4.2.1.2 Synchronisation Across Engineering, Procurement and Installation.....	125
4.2.1.3 Establishing the Digital Thread.....	126
4.2.1.4 Evidence from Industry Practice.....	127
4.2.2 Extending BIM Beyond SAP: MES, IoT/CPS and the DVCO Semantic Layer....	127
4.2.2.1 The Cross-Domain Interoperability Challenge.....	127
4.2.2.2. DVCO’s Semantic Layer as the Integration Mechanism.....	129
4.2.2.3. Enabling a Fully Connected Digital Thread.....	129
4.3 Ontologies and Knowledge Graphs: The Semantic Backbone of DVCO.....	131
4.3.1 The Role of Ontologies in Lifecycle Semantic Alignment.....	131
4.3.3 Why Ontologies and KGs Are Essential in ETO Environments.....	134
4.4 IoT, CPS and Telemetry Layer.....	136
4.4.1 IoT Telemetry as the Foundation of Real-Time Visibility.....	138
4.4.2 Digital Twin Synchronisation and Behavioural Alignment.....	139
4.4.2.1 Geometric and Spatial Synchronisation (BIM–IoT Integration).....	140
4.4.2.2. Operational Synchronisation (MES–IoT–SAP Integration).....	140
4.4.2.3. Behavioural and Predictive Synchronisation (AI-Enabled Digital Twins)...	141
4.4.2.4 Role of DVCO in Synchronising the Digital Twin Ecosystem.....	141
4.4.3 Telemetry-Driven Anomaly Detection and Commissioning Assurance.....	142
4.4.3.1. Detection of Equipment and Automation Anomalies.....	143
4.4.3.2. Spatially Contextualised Anomaly Interpretation.....	143
4.4.3.3. Commissioning Validation and Readiness Assurance.....	144
4.4.3.4. Early-Warning and Predictive Assurance.....	145
4.4.3.5. Lifecycle Impact and Feedback Integration.....	145
4.4.4 Closed-Loop Feedback into Engineering, Planning and Execution Systems.....	146
4.4.4.1. Engineering Feedback: Updating Models, Constraints and Designs.....	146
4.4.4.2. Planning Feedback: Adjusting Schedules, Milestones and Dependencies.....	147
4.4.4.3. Execution Feedback: Orchestrating Production, Installation and Commissioning..	148
4.4.4.4. Knowledge Graph Integration: Persisting Traceability and Context.....	148

4.4.4.5. System-Level Outcome: A Self-Updating, Orchestrated D2O Environment.....	149
4.4.5 The Role of the Telemetry Layer in DVCO.....	149
4.5 Decision Intelligence Layer.....	150
4.5.1 Predictive Analytics and Machine-Learning Models.....	151
4.5.2 Semantic Reasoning Using Ontologies and Knowledge Graphs.....	152
4.5.3 LLM-Based Agents for Orchestration, Insight and Human–AI Collaboration.....	152
4.5.4 Human-in-the-Loop and Organisational Decision Enablement.....	153
4.6 Multi-Layer DVCO Architecture.....	155
4.6.1 Layered Architectural View.....	156
4.6.2 Cross-Cutting Digital Thread and Feedback Loops.....	158
4.6.3 Figure Description.....	159
4.7 LLM-Enabled Reasoning and Decision Intelligence.....	160
4.7.1 LLMs as Cognitive Interfaces Across the Value Chain.....	161
4.7.2 Hybrid Reasoning: Integrating LLMs with Ontologies and Knowledge Graphs...	162
4.7.3 LLM-Enabled Use Cases in DVCO.....	163
4.7.3.1 Natural-Language Querying of Multi-System Data.....	164
4.7.3.2 Change-Impact Analysis.....	164
4.7.3.3 Automated Technical Documentation and Reporting.....	165
4.7.3.4 Predictive and Prescriptive Reasoning.....	165
4.7.3.5 Multi-Scenario Simulation Support.....	166
4.7.3.2. Change-Impact Analysis.....	167
4.7.3.3. Automated Technical Documentation and Reporting.....	167
4.7.3.4. Predictive and Prescriptive Reasoning.....	168
4.7.3.5. Multi-Scenario Simulation Support.....	168
4.7.4 Human–AI Collaboration in DVCO Decision Workflows.....	169
4.8 Integrated DVCO Architecture.....	170
4.8.1 Overview of the Multi-Layer Architectural Model.....	171
4.8.2 Cross-Layer Data Flow and Lifecycle Synchronisation.....	172
4.8.3 Architectural Integration Principles.....	173
4.8.4 Value-Chain-Level Digital Twin.....	174
4.8.5 DVCO Orchestration Logic.....	175
4.8.6 Architectural Significance for ETO and Symbiotic.....	176
4.9 Summary.....	177
CHAPTER V.....	179

RESEARCH METHODOLOGY.....	179
5.1 Research Framework and Phases.....	179
5.1.1. Phase 1: Exploration – Mapping the ETO Value Chain and Identifying Bottlenecks. 179	
5.1.2. Phase 2: Modelling – Ontology and Knowledge Graph Construction.....	180
5.1.3. Phase 3: Simulation and Validation – Testing Improvements Using AI-Enabled Scenarios.....	181
5.2 Research Design.....	182
5.2.1 Qualitative Systems Analysis.....	183
5.2.2 Model-Based Design Using Ontologies and Knowledge Graphs.....	184
5.2.3 AI-Enabled Simulation and Scenario Analysis.....	184
5.2.4 Justification for the Hybrid Research Design.....	185
5.3 Data Sources and Data Collection.....	186
5.3.1 Documentary Sources.....	187
5.3.2 Synthetic Data Construction.....	193
5.3.3 Expert Knowledge and Domain Understanding.....	196
5.3.4 Synthetic Data for Ontology and Simulation.....	197
5.3.5 Data for AI Reasoning and LLM Interaction.....	198
5.3.6 Data Collection Approach and Rationale.....	198
5.4 Data Analysis Procedures.....	201
5.4.1 Qualitative and Documentary Analysis.....	202
5.4.2 Ontology Structuring and Semantic Alignment.....	202
5.4.3 Knowledge Graph Construction and Consistency Checking.....	203
5.4.4 Simulation-Based Evaluation and Decision Analysis.....	204
5.5 Ontology Construction Process.....	205
5.5.2 Conceptual Modelling and Entity Definition.....	207
5.5.3 Formal Representation Using OWL/RDF.....	208
5.5.4 Iterative Validation and Refinement.....	209
5.6 Reliability, Validity and Bias Mitigation.....	210
5.6.1 Reliability.....	210
5.6.2 Validity.....	212
5.6.3 Bias Mitigation.....	213
5.6.4 Summary.....	213
5.7 Ethical Considerations.....	214

5.7.1 Data Privacy and Confidentiality.....	214
5.7.2 Responsible Use of Artificial Intelligence.....	215
5.7.3 Avoidance of Harm and Misuse.....	215
5.7.4 Transparency and Integrity in Research Reporting.....	216
5.7.5 Ethical Considerations in Multi-Party Value Chains.....	217
5.7.6 Compliance with Academic Ethical Standards.....	217
5.7.7 Summary.....	218
Chapter VI.....	219
Findings and Discussion.....	219
6.1 Introduction.....	219
6.2 Findings from Ontology and Semantic Integration.....	223
6.2.1 Enhanced End-to-End Lifecycle Traceability.....	223
6.2.2 Resolution of Cross-System Semantic Inconsistencies.....	224
6.2.3 Improved Quality and Reliability of Downstream Models.....	224
6.3 Findings from Knowledge Graph Evaluation.....	225
6.3.1 Accurate and Complete Change-Impact Propagation.....	225
6.3.2 Bidirectional Lifecycle Traceability.....	225
6.3.3 Structural Insights Beyond Traditional Systems.....	226
6.3.4 Foundation for LLM-Based Reasoning.....	226
6.4 Findings from LLM-Enabled Reasoning.....	226
6.5 Findings from Integrated DVCO Simulation.....	227
6.7 Summary.....	228
6.8 Transition to Strategic Recommendations.....	230
CHAPTER VII.....	233
CONCLUSION AND RECOMMENDATIONS.....	233
7.1 Introduction.....	233
7.2 Summary of the Research.....	233
7.3 Theoretical Contributions.....	235
7.4 Practical Implications for Symbiotic.....	237
7.4.1 Improved deployment repeatability.....	237
7.4.4 Enhanced cross-functional collaboration	
The semantic layer creates a shared operational language across engineering, manufacturing, installation and commissioning teams. By aligning terminology and lifecycle semantics, DVCO reduces misunderstandings at handover points and improves cross-functional coordination (Främling et al., 2013; Sheth, 2020).....	238

7.5 Recommendations for Implementation.....	238
7.5.1 Establish a Semantic Integration Programme.....	239
7.5.2 Deploy a Knowledge Graph Platform.....	239
7.5.3 Integrate LLMs Through a Hybrid Architecture.....	239
7.5.4 Pilot DVCO at a Selected Customer Deployment.....	240
7.5.5 Build a Digital Twin Governance Workflow.....	240
7.5.6 Expand Cross-Functional Training Programmes.....	241
7.6 Limitations of the Study.....	242
7.7 Directions for Future Research.....	243
7.8 Final Conclusion.....	244
BIBLIOGRAPHY.....	248
GLOSSARY OF KEY TERMS.....	254
INDEX.....	258

# LIST OF TABLES

Table 3.1 Comparison of DVC (Cubreath) and DVCO (This Dissertation).....	104
Table 3.2 Architectural Validation to Chapter 6 Findings.....	105
Table 4.1 Revit(BIM) and SAP Mapping table.....	145
Table 6.1 Mapping of Key Findings to Theoretical and Managerial Implications.....	212
Table 6.2 Alignment Between Dynamic Value Chain Figure and Chapter VI Findings.....	215

# LIST OF FIGURES

Figure 1.1 The structure of Automation warehouse by Robotics.....	18
Figure 2.1 The end-to-end ETO/D2O value chain .....	38
Figure 3.1 Dynamic Value Chain Optimisation (DVCO) framework.....	97
Figure 3.2 DVCO Architecture - ETO/E2O Value Chain.....	112
Figure 4.2 Multi-Layer DVCO Architecture.....	149
Figure 5.1 DVCO ARCHITECTURE - THREADS VIEW .....	176
Figure 6.1 Dynamic value chain in Engineer-to-Order (ETO) deployments and.....	205

# CHAPTER I

## Introduction

### 1.1 Background and Context

Engineer-to-Order (ETO) industries are characterised by high variability, extensive engineering dependencies and the coordination of multiple technical and organisational domains over long and complex project lifecycles (Gosling and Naim, 2009). Symbotic, a leading provider of AI-enabled warehouse-automation systems, operates within one of the most sophisticated ETO environments in contemporary industry. Each Symbotic deployment constitutes a substantial cyber-physical infrastructure consisting of autonomous robots, sensor-mediated control systems, integrated mechanical structures and tightly sequenced commissioning procedures.

Traditional project-management instruments—such as ERP-based material planning, static Gantt scheduling, and spreadsheet-driven resource tracking—struggle to accommodate the structural, temporal and behavioural complexity intrinsic to such deployments (Hicks and McGovern, 2009; Olhager and Wikner, 2014). As Symbotic’s global operations continue to expand, the need for real-time integration across design, engineering, manufacturing, logistics and installation domains becomes increasingly critical.

Simultaneously, technological advances in model-based engineering, semantic data modelling, knowledge-graph reasoning and digital-twin simulation have introduced new opportunities for operational transformation (Eigner et al., 2017; Stark and Damerau, 2019). These developments enable organisations to move beyond isolated information silos toward fully integrated value-chain intelligence spanning the complete Design-to-Operate (D2O) lifecycle.

Accordingly, this dissertation investigates how these technologies may be systematically combined within a unified Dynamic Value Chain Optimization (DVCO) framework to improve Symbotic’s orchestration of high-complexity automation deployments.

## 1.2 Problem Statement

Despite Symbotic’s advanced automation capabilities and its leadership in modular, AI-driven warehouse-robotics systems, its ETO deployments confront challenges common to engineering-intensive production ecosystems (Hicks, McGovern and Earl, 2000). These difficulties arise largely from the fragmented nature of information architectures underpinning the value chain.

First, data is distributed across heterogeneous systems—including BIM, SAP S/4HANA, MES platforms and IoT-derived telemetry—without unified semantics or lifecycle alignment. Second, engineering-design artefacts and real-world site execution frequently diverge, with operational feedback rarely propagated into upstream planning models. Third, change-impact analysis is

predominantly expert-driven and manually executed, limiting the capacity for scalable responsiveness.

Fourth, cross-domain dependencies remain obscured due to incompatible data schemas, inconsistent terminology and the absence of semantic interconnectivity. Fifth, existing enterprise tools fail to provide dynamic constraint modelling that can holistically evaluate spatial, temporal, logistical and behavioural constraints.

Furthermore, variability in supplier performance and engineering-change frequency introduces cascading disruption effects that are difficult to anticipate without integrated and adaptive reasoning systems (Ballard and Thomas, 2021).

Collectively, these challenges lead to rework, delays, sub-optimal resource deployment, cost fluctuations and inconsistent commissioning outcomes. Incremental improvements achieved through ERP extensions or isolated BIM adoption remain insufficient, as they do not constitute a unified semantic representation of the value chain.

This indicates the need for a novel integrated methodology that connects engineering, manufacturing, logistics, functional execution and site-level operational data through a model-driven, semantically grounded and AI-assisted execution framework.

Figure 1.1 illustrates the structural architecture of an automation warehouse enabled by robotics, representing a core execution layer within a modern Engineering-to-Order (ETO) driven logistics and fulfillment environment. The figure demonstrates how robotic systems are orchestrated to

integrate physical material handling with digital control layers, forming a cyber-physical warehouse system.

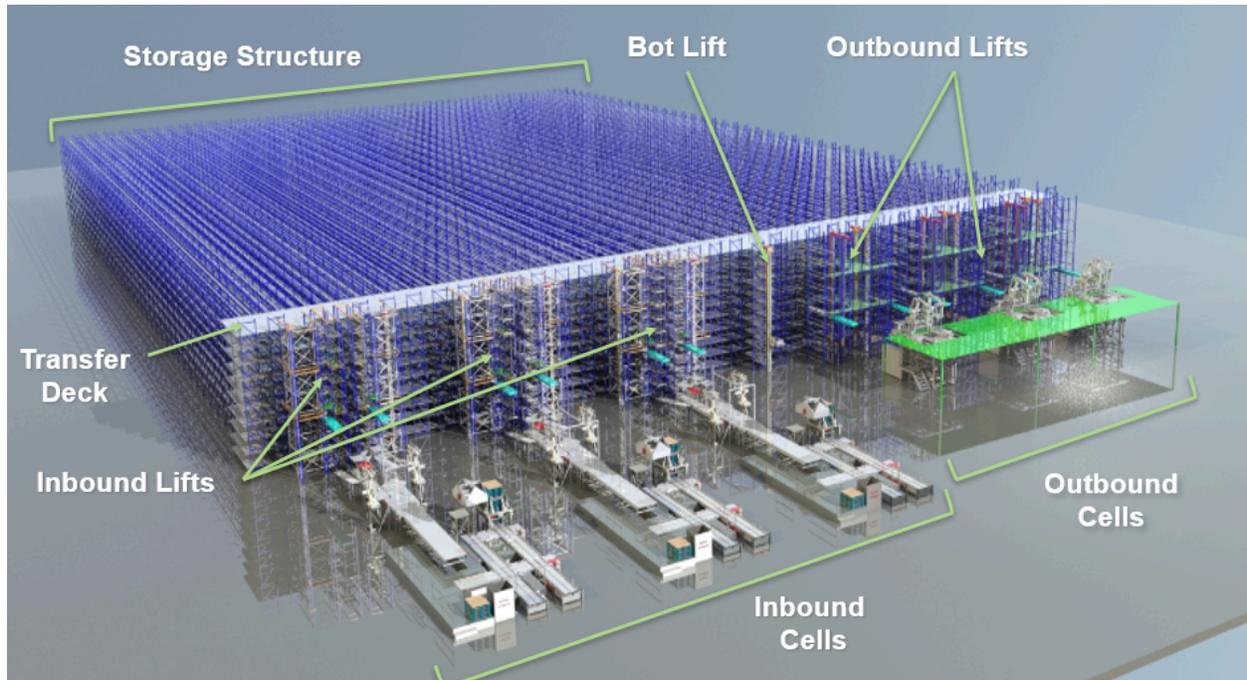


Figure 1.1 The structure of Automation warehouse by Robotics

### 1.3 Research Aim and Objectives

#### 1.3.1 Research Aim

The aim of this dissertation is to design and evaluate a Dynamic Value Chain Optimization (DVCO) framework that directly addresses the coordination inefficiencies, limited end-to-end

predictability, and weak lifecycle alignment identified in Symbotic's Design-to-Operate (D2O) execution model.

Specifically, the framework is intended to resolve the following core problems identified in Section 1.2:

1. Fragmented coordination across design, engineering, manufacturing, and logistics, caused by disconnected systems, document-centric workflows, and late-stage information handoffs.
2. Low operational predictability in Engineering-to-Order (ETO) execution, where frequent engineering changes propagate inconsistently into planning, production, and warehouse automation activities.
3. Insufficient lifecycle feedback mechanisms, resulting in operational and performance data from execution and field operations being weakly reintegrated into upstream design and planning decisions.

To address these problems, the proposed DVCO framework integrates:

- Model-Based Engineering (MBE) to establish a single, authoritative digital representation of product, process, and system intent;
- Semantic integration to ensure consistent interpretation and traceability of engineering, operational, and logistical data across heterogeneous enterprise systems; and

- AI-enabled reasoning to dynamically assess change impacts, support decision-making, and optimize value chain coordination in near real time.

By doing so, the framework aims to enhance coordination efficiency, improve execution predictability, and achieve tighter lifecycle alignment across Symbotic's D2O execution, thereby mitigating the systemic limitations observed in current ETO-oriented value chain operations.

### 1.3.2 Research Objectives

To accomplish the research aim, the study is guided by six objectives:

1. Map Symbotic's ETO/D2O value chain to identify bottlenecks, fragmentation points and cross-functional dependencies across engineering, supply chain, manufacturing, logistics, installation and commissioning.
2. Develop a cross-domain ontology that formalizes MBE, engineering, manufacturing, logistics, installation and operational concepts, including their attributes, lifecycle states and semantic relationships.
3. Construct a knowledge graph capable of linking BIM, SAP/MES, IoT and installation-task data into a unified, lifecycle-connected digital structure supporting reasoning and propagation analysis.

4. Evaluate AI-enabled reasoning, integrating LLM capabilities with the knowledge graph to anticipate risks, identify hidden dependencies and provide explainable decision support for project managers and engineers.
5. Demonstrate the value of MBE-enabled DVCO by simulating how unified engineering models, semantic data, and AI reasoning improve visibility, coordination efficiency, impact analysis and commissioning readiness.
6. Provide recommendations for Symbotic and the wider ETO industry regarding the adoption of MBE, semantic integration and AI-assisted operational intelligence as foundations for next-generation D2O execution models.

#### 1.4 Research Questions

The study is guided by three primary research questions, each addressing a critical dimension of Symbotic's Engineer-to-Order (ETO) and Design-to-Operate (D2O) value chain:

Research Question 1:

How can Model-Based Engineering (MBE), together with semantic modelling through ontologies and knowledge graphs, improve visibility, traceability and lifecycle coherence across Symbotic's ETO value chain?

This question investigates how unified engineering models, formal semantics and graph-based data structures resolve fragmentation between BIM, ERP, MES, IoT and field-execution systems.

Research Question 2:

How can AI-enabled reasoning—particularly large language models (LLMs) grounded in knowledge graph semantics—enhance early risk detection, change-impact analysis, and decision-making throughout the D2O lifecycle?

This explores the value of combining semantic intelligence with predictive analytics to detect risks sooner and support more informed, explainable decision processes.

Research Question 3:

What operational improvements can be achieved through the integration of the DVCO framework, and how does this integration affect coordination, predictability, installation performance and commissioning readiness?

This evaluates the tangible operational benefits of DVCO and its impact on reducing delays, rework, cost variability and cross-functional misalignment.

### 1.5 Significance of the Study

This study is significant for both academic inquiry and industrial practice, as it addresses one of the most persistent and complex challenges in Engineer-to-Order (ETO) environments: the fragmentation of engineering, manufacturing, supply-chain and operational data across the project lifecycle. By proposing and evaluating a Dynamic Value Chain Optimization (DVCO)

framework grounded in Model-Based Engineering (MBE), semantic integration and AI-enabled reasoning, the research offers several important contributions.

#### 1.5.1. A new semantic and model-driven architecture for ETO execution

The study provides a unified conceptual and technical architecture that connects BIM, ERP, MES, IoT and field-execution data through ontologies and knowledge graphs. This represents a significant advancement over traditional system-integration approaches, which often rely on point-to-point interfaces without shared semantics or lifecycle coherence.

#### 1.5.2. Demonstration of AI-assisted reasoning in complex project environments

By integrating large language models (LLMs) with knowledge-graph semantics, the research demonstrates how AI can support impact analysis, risk prediction, natural-language querying and decision support in high-variability, low-repeatability ETO deployments—an application area historically underserved by conventional machine learning.

#### 1.5.3. Bridging engineering, planning, manufacturing and operational domains

The research shows that DVCO creates a continuous digital thread spanning:

- engineering and spatial modelling (BIM, MBE),
- planning and procurement (SAP S/4HANA),
- production execution (MES/PEO),
- real-time sensing and site behaviour (IoT).

This holistic digital integration improves lifecycle traceability, cross-functional coordination and the accuracy of change-impact propagation.

#### 1.5.4. A scalable model for next-generation warehouse-automation deployments

Symbolic's deployments represent a new generation of cyber-physical infrastructure. The DVCO framework offers a scalable blueprint for coordinating engineering changes, supply-chain uncertainties, robotic-system integration, installation sequencing and commissioning readiness across global deployment programs.

#### 1.5.5. Actionable insights for digital transformation in the ETO industry

The study provides practical guidance for organisations seeking to:

- adopt MBE as a lifecycle backbone,
- implement semantic data models,
- operationalise digital twins,
- embed AI-enabled reasoning into value-chain orchestration,
- move beyond siloed ERP/MES/BIM workflows toward integrated execution models.

These insights have broad applicability across aerospace, robotics, industrial machinery, construction automation, shipbuilding and other advanced ETO sectors.

## 1.6 Dissertation Structure

This dissertation is organised into seven chapters, each contributing to the development, application and evaluation of the Dynamic Value Chain Optimization (DVCO) framework within Symbolic's Engineer-to-Order (ETO) and Design-to-Operate (D2O) environment.

### Chapter I – Introduction

Establishes the background, problem context and motivation for the study. It articulates the research aim, objectives and questions, highlights the significance of integrating MBE, semantic modelling and AI reasoning, and outlines the structure of the dissertation.

### Chapter II – Literature Review

Examines scholarly research and industry studies related to:

- ETO value-chain complexity
- Industry 4.0 technologies and integration challenges
- Model-Based Systems Engineering (MBSE)
- Digital twin concepts and limitations

- Semantic modelling through ontologies and knowledge graphs
- AI and LLM applications in manufacturing and engineering

The review identifies major gaps and establishes the foundation for the DVCO framework.

### Chapter III – Conceptual Framework and DVCO Architectural Model

Presents the proposed Dynamic Value Chain Optimization framework, grounded in Dynamic Capabilities Theory, digital-thread principles and value-chain analysis. It defines the core components of DVCO—model-based engineering, software-defined execution, real-time sensing, semantic integration and AI-enabled decision intelligence—and explains their interactions across the D2O lifecycle.

### Chapter IV – Enabling Technologies and Integrated Architecture

Provides a detailed analysis of the technologies required to operationalise DVCO, including:

- BIM and digital-twin modelling
- Ontologies and knowledge graphs
- SAP S/4HANA, MES and process integration
- IoT telemetry and site-level sensing

- Palantir Foundry as an ontology-based integration layer
- LLM-driven predictive and semantic reasoning

This chapter forms the technical foundation for subsequent modelling and simulation.

#### Chapter V – Research Methodology

Describes the multi-phase research design, including exploration, ontology and knowledge-graph modelling, and AI-enabled simulation. It addresses data sources, development procedures, reliability, validity and ethical considerations.

#### Chapter VI– Findings and Discussion

Presents the results of the semantic models and simulations.

It discusses insights regarding engineering-change propagation, dependency behaviours, AI-enhanced risk detection, digital-twin dynamics and the operational impacts of DVCO.

The discussion positions these findings within existing literature.

#### Chapter VII – Conclusion and Recommendations

Summarises the contributions of the study, outlines strategic and operational recommendations for Symbotic and the wider ETO industry, and identifies areas for future research including agentic AI, probabilistic modelling and scalable graph deployment.

## CHAPTER II

### LITERATURE REVIEW

## 2.1 Introduction

The purpose of this literature review is to establish a comprehensive and interdisciplinary foundation that connects management theory, systems engineering principles, digital-transformation frameworks, and emerging developments in artificial intelligence into a unified narrative of value-chain evolution (Porter, 1985; Teece, Pisano and Shuen, 1997; Jarke et al., 2020). By examining the progression from classical value-chain concepts to contemporary cyber-physical, semantic, and AI-enabled paradigms, the review positions Dynamic Value Chain Optimisation (DVCO) within the broader academic and industrial landscape (Lee, Bagheri and Kao, 2015; Schuh et al., 2017).

This chapter is organised chronologically and thematically. It begins by reviewing the foundational principles of ETO value-chain research and operations-management theory. It then examines the rise of Industry 4.0 and cyber-physical systems, highlighting how these movements reshaped understanding of manufacturing connectivity, system interoperability, and real-time orchestration (Culot et al., 2020; Wang, Törngren and Onori, 2015). The review subsequently explores Model-Based Enterprise (MBE) and digital-twin methodologies, which shift organisations from document-centric to model-centric execution (Hedberg and Feeney, 2017; Grieves, 2014). Attention is also given to semantic technologies—including ontologies and knowledge graphs—that address long-standing issues of cross-functional fragmentation and data inconsistency in engineering-intensive environments (Beetz, van Leeuwen and de Vries, 2009;

Yu and Liu, 2020; Sheth, 2020).

Finally, the review surveys the emerging role of artificial intelligence, with particular emphasis on large language models (LLMs) and their potential to augment reasoning, forecasting, and complex decision workflows (Davenport and Ronanki, 2018; Kaplan and Haenlein, 2019). Across these domains, persistent challenges are identified: misaligned engineering and operational processes, incomplete lifecycle integration, limited adaptability of traditional optimisation approaches, and the absence of unified semantic structures capable of supporting intelligent, cross-domain decision-making (Gosling and Naim, 2009; Schuh et al., 2017).

These gaps highlight the need for a new integrative framework. DVCO emerges as a response to the limitations identified in the literature, offering a semantically coherent, AI-enabled, model-centric approach to orchestrating ETO value chains (Lin, Lee and Ma, 2020). This literature review therefore serves as the conceptual scaffolding for the DVCO framework developed in Chapter 3 and empirically evaluated in Chapters 5 and 6.

## 2.2 Engineer-to-Order (ETO) Value Chain Literature

Engineer-to-Order (ETO) production systems represent one of the most complex forms of industrial value-chain execution (Gosling and Naim, 2009; Bolton and Hannon, 2016). Unlike

Make-to-Stock (MTS) or Assemble-to-Order (ATO) environments—where product structures are largely predefined—ETO organisations begin engineering work after customer requirements are confirmed. This creates a value chain in which design, procurement, manufacturing, installation, and commissioning are not strictly sequential but deeply interdependent, fluid and frequently disrupted. The literature consistently highlights that ETO environments display high variability, long lead times, and a high degree of coordination uncertainty (Gosling and Naim, 2009; Gosling and Towill, 2011).

### 2.2.1 Characteristics and Structural Complexity of ETO Systems

ETO systems are distinguished by several structural characteristics:

- Customer-specific engineering that triggers bespoke design, often requiring new components or configurations.
- Interdependence between engineering and procurement, where design freeze is often delayed due to iterative negotiation or evolving customer specifications.
- Long and uncertain lead times, particularly for critical or custom-manufactured components.
- High product and process complexity, creating an environment where engineering changes can propagate widely across the value chain.
- Reliance on specialist suppliers and subcontractors, with varying degrees of integration and data-sharing capability.

- Late-stage site installation activities, which depend on both engineering accuracy and logistical reliability.

These characteristics create what researchers term structurally coupled value chains, where deviations in one part of the system quickly ripple across others (Gosling and Naim, 2009; Stevenson and Spring, 2007).

### 2.2.2 Common Bottlenecks Identified in ETO Literature

#### Engineering Change Propagation

Engineering changes are a dominant disruption mechanism. Studies show that ETO organisations frequently face cascading delays because engineering modifications do not propagate effectively across downstream processes (Terwiesch and Loch, 1999). This is exacerbated by inconsistent data structures across engineering, planning and manufacturing systems (Hicks, McGovern and Earl, 2000).

#### Cross-Functional Fragmentation

Research points to chronic fragmentation between engineering, supply chain and production functions. Each operates within its own system environment—CAD/BIM, ERP, MES—resulting in poor synchronization (Hicks, McGovern and Earl, 2000; Gosling and Naim, 2009).

### Poor Visibility into Multi-Domain Dependencies

ETO projects typically lack tools that provide end-to-end transparency across design, procurement, manufacturing and installation (Gosling and Towill, 2011). Dependencies remain hidden within disconnected systems or within the tacit knowledge of experienced engineers.

### High Vulnerability to Supplier Variability

Long-lead or custom components introduce supply-chain uncertainty. Literature notes that ETO performance is often constrained by vendor reliability rather than internal capability (Stevenson and Spring, 2007).

### Installation and Commissioning Challenges

ETO deployments—particularly those involving automation or infrastructure-intensive systems—experience delays due to site conditions, coordination complexity and limited integration between engineering models and field execution (Whyte et al., 2016).

## 2.2.3 Limitations of Traditional ETO Planning and Control

### Limitations of Traditional ETO Planning and Control

The limitations discussed in Section 2.2.3 are related to, but analytically distinct from, the limitations identified in the preceding ETO literature review sections. While the earlier literature-based limitations are derived from established academic studies on ETO supply chains

and project-based manufacturing, the limitations in Section 2.2.3 are identified through a practice-oriented and system-level analysis of traditional ETO planning and control mechanisms.

### Relationship to ETO Literature Limitations

The limitations identified in the ETO literature (Section 2.2.1–2.2.2) primarily emphasize:

- Structural characteristics of ETO environments (e.g., high customization, uncertainty, and long lead times),
- Organizational and supply chain coordination challenges, and

Conceptual shortcomings of classical planning models when applied to ETO contexts. In contrast, Section 2.2.3 focuses specifically on the operational and technological limitations of traditional ETO planning and control approaches, such as:

- MRP- and project-centric planning logic,
- Document-driven engineering change processes,
- Static scheduling and rule-based control mechanisms, and
- Limited real-time feedback from execution systems.

### How the Limitations in Section 2.2.3 Are Identified

The limitations in Section 2.2.3 are identified through a synthesis of three complementary sources:

1. Analytical Decomposition of Traditional Planning and Control Models

A systematic analysis of how conventional ETO planning approaches (e.g., ERP-based MRP, project scheduling, and siloed MES/WMS control) behave under conditions of high engineering variability and frequent change.

2. Industry Practice and Practitioner Evidence

Observations derived from industry implementations in ETO and automation-intensive environments, including robotics-based warehouses and D2O execution models, where traditional planning and control mechanisms fail to maintain alignment across design, engineering, and execution.

3. Gap Analysis Between Literature Assumptions and Digital Execution Reality

A comparison between the assumptions embedded in traditional ETO planning frameworks (e.g., stable BOMs, frozen designs, sequential handoffs) and the realities of digitally enabled, dynamically reconfigurable production and logistics systems.

Positioning Within the Overall Literature Review

Accordingly, Section 2.2.3 does not duplicate the limitations identified in the ETO literature, but rather:

- Operationalizes them at the planning and control level,
- Translates conceptual limitations into system and process constraints, and
- Establishes the direct motivation for introducing a Dynamic Value Chain Optimization (DVCO) framework in subsequent chapters.

This distinction is essential, as it bridges the gap between theoretical limitations discussed in prior research and the practical execution challenges that arise when traditional ETO planning and control models are applied in modern, digitally integrated D2O environments.

#### 2.2.4 Emergence of Digital and Model-Based Solutions in ETO Literature

Recent literature highlights several emerging approaches to address ETO challenges:

- Digital twins, enabling simulation and verification before physical execution (Grieves, 2014; Boschert and Rosen, 2016).
- BIM, improving spatial coordination and engineering accuracy (Eastman et al., 2011; Bryde, Broquetas and Volm, 2013).

- MBE (Model-Based Engineering), supporting consistent engineering representation across disciplines (Hedberg and Feeney, 2017).
- Knowledge graphs and ontologies, addressing semantic inconsistencies and enabling dependency modelling (Beetz, van Leeuwen and de Vries, 2009; Yu and Liu, 2020; Sheth, 2020).
- AI-enabled reasoning, offering predictive and interpretive support in environments with sparse structured data (Davenport and Ronanki, 2018; Zhang et al., 2021).

However, research also acknowledges that these technologies are often implemented in isolation, and the absence of an integrated framework limits their value (Culot et al., 2020).

#### 2.2.5 Literature Gap: Lack of a Unified, Semantic, AI-Enabled Framework

Despite substantial research into ETO challenges and digital-transformation tools, the literature consistently identifies several unmet needs:

- Lack of cross-domain semantic integration connecting engineering, planning, manufacturing, logistics and site execution (Yu and Liu, 2020; Beetz, van Leeuwen and de Vries, 2009).
- Limited adoption of MBE beyond the engineering function (Hedberg and Feeney, 2017).

- Insufficient modelling of lifecycle dependencies across BIM, ERP, MES and IoT systems (Eastman et al., 2011; Wang, Törngren and Onori, 2015).
- No widely adopted framework for AI-enabled reasoning in ETO coordination and risk prediction (Davenport and Ronanki, 2018; Sheth, 2020).
- Absence of a holistic value-chain-level digital twin (Grieves, 2014; Boschert and Rosen, 2016).

These gaps highlight a critical opportunity to develop a unified conceptual and technical model—such as DVCO—that synthesizes ontology-driven integration, digital twins, IoT telemetry and AI reasoning.

### 2.3 Dynamic Nature of ETO Value Chains

Engineer-to-Order (ETO) value chains are fundamentally dynamic in nature, exhibiting levels of variability, interdependence, and uncertainty that significantly exceed those found in Make-to-Stock (MTS) or Assemble-to-Order (ATO) environments. While ETO value chains are often represented as sequential lifecycle stages—from customer engagement through engineering, planning, manufacturing, installation, and service—such representations can obscure the reality that these stages are neither linear nor temporally isolated. Instead, ETO value

chains operate as continuously evolving systems characterized by frequent feedback loops, iterative rework, and real-time adjustments across organizational and functional boundaries.

A defining characteristic of ETO environments is the volatility of customer-driven requirements. Unlike standardized products, ETO solutions are typically customized to specific operational, spatial, regulatory, and performance constraints. As a result, customer requirements often evolve throughout the project lifecycle due to changes in business objectives, market conditions, regulatory interpretations, or site-specific discoveries. These changes may occur even after contractual agreements have been established, creating downstream impacts on engineering designs, procurement plans, production schedules, and installation strategies. Consequently, requirement definition in ETO projects cannot be treated as a one-time, front-loaded activity but must instead be continuously managed as a dynamic process.

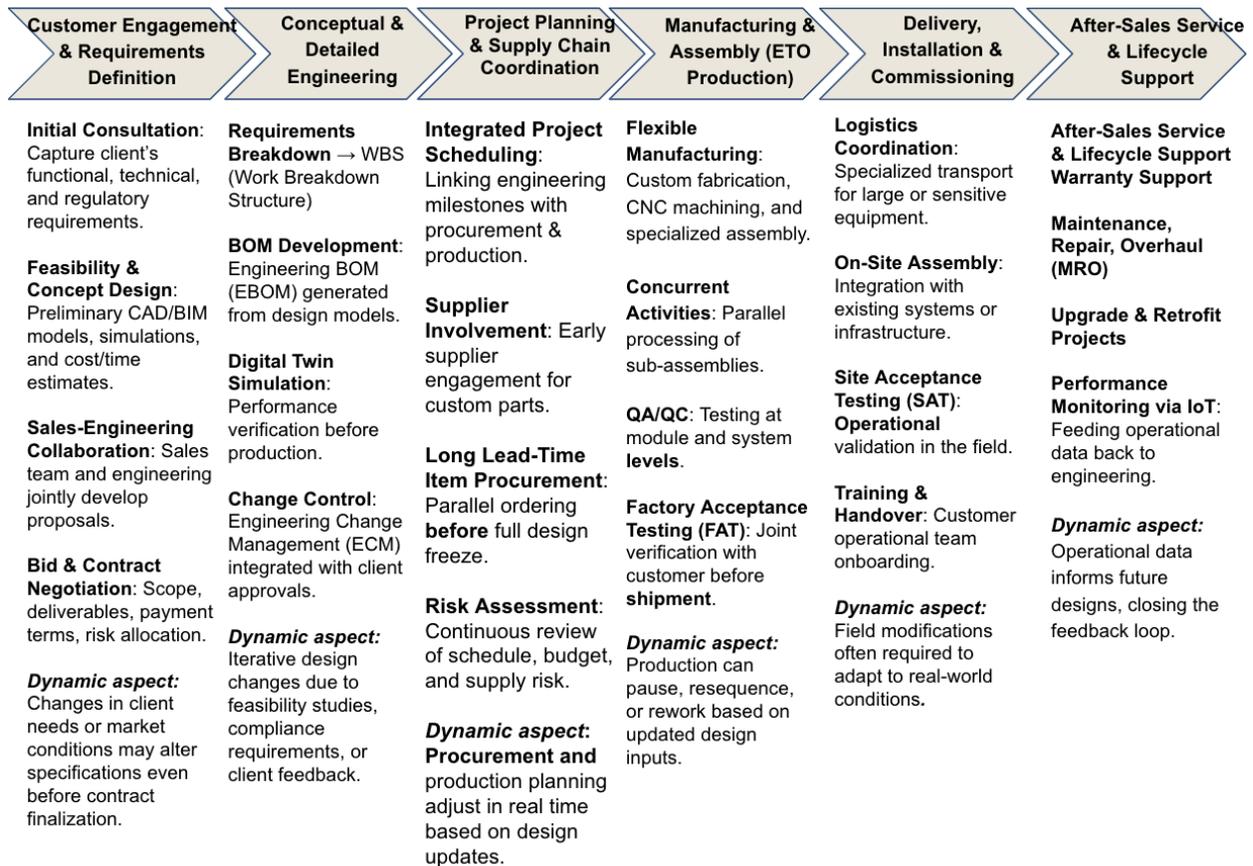


Figure 2.1 The end-to-end ETO/D2O value chain

Figure 2.1 illustrates the end-to-end ETO/D2O value chain and highlights the inherently dynamic nature of each lifecycle phase, characterized by frequent design changes, feedback loops, and real-time adjustments.

The engineering phase further amplifies this dynamic behavior. Conceptual and detailed engineering activities in ETO projects rely heavily on iterative design exploration, simulation, and validation. Preliminary CAD and BIM models are frequently refined based on feasibility assessments, performance simulations, compliance requirements, and stakeholder feedback. Engineering change management (ECM) becomes a central operational activity rather than an exception, as design modifications propagate across engineering bills of materials (EBOMs), routing structures, and configuration rules. These engineering changes rarely remain confined to the design domain; instead, they exert cascading effects on procurement, manufacturing, and installation processes, often requiring rapid replanning and cross-functional coordination.

Project planning and supply chain coordination in ETO value chains must therefore contend with a high degree of temporal overlap between engineering, sourcing, and production activities. To mitigate long lead times, procurement of custom or critical components often begins before final design freeze, exposing projects to increased risk when engineering changes occur. Supplier involvement is typically initiated early to support manufacturability and cost optimization, yet suppliers themselves become participants in the dynamic change process, responding to revised specifications, delivery schedules, and quality requirements. As a result, project schedules in ETO environments are inherently provisional, requiring continuous risk assessment and adjustment rather than static baseline adherence.

Manufacturing and assembly activities in ETO contexts further illustrate the non-linear nature of the value chain. Production processes are frequently characterized by flexible manufacturing,

custom fabrication, specialized assembly, and parallel processing of subassemblies. Unlike repetitive manufacturing environments, ETO production must accommodate interruptions, resequencing, and rework triggered by late engineering updates or supply disruptions. Quality assurance and testing activities, including Factory Acceptance Testing (FAT), often reveal issues that necessitate design clarification or modification, thereby reintroducing engineering and planning activities into what would traditionally be considered downstream phases.

The delivery, installation, and commissioning stages introduce additional layers of dynamism, particularly due to site-specific conditions that cannot be fully anticipated during design. On-site assembly frequently requires adaptation to existing infrastructure, spatial constraints, and environmental factors, leading to field modifications that diverge from original engineering assumptions. Site Acceptance Testing (SAT) and commissioning activities may uncover integration issues or performance deviations that require rapid coordination between engineering, manufacturing, and field teams. These field-driven feedback loops further challenge linear lifecycle assumptions and highlight the necessity of real-time information exchange across the value chain.

The dynamic nature of ETO value chains extends beyond project delivery into after-sales service and lifecycle support. Maintenance, repair, and overhaul (MRO) activities, upgrade projects, and performance optimization initiatives generate valuable operational data that can inform future engineering designs and planning assumptions. The increasing adoption of IoT-enabled monitoring technologies accelerates this feedback loop, enabling continuous performance

tracking and condition-based insights. In mature ETO organizations, operational data increasingly influences not only service strategies but also upstream design decisions, effectively closing the loop between operation and engineering.

Collectively, these characteristics demonstrate that ETO value chains function as adaptive, feedback-driven systems rather than deterministic, sequential processes. Changes originating in any lifecycle phase—whether driven by customer requirements, engineering constraints, supply chain disruptions, or operational feedback—can propagate rapidly across the entire value chain. Traditional representations that assume stable handoffs between discrete phases fail to capture this reality and risk underestimating the coordination complexity inherent in ETO environments.

This intrinsic dynamism has profound implications for planning, execution, and decision-making. Static planning models, document-centric change management processes, and siloed enterprise systems struggle to maintain alignment in the face of continuous change. As ETO value chains grow in scale and complexity, the limitations of linear and deterministic approaches become increasingly pronounced. These limitations, examined in the following section, provide the critical motivation for the development of dynamic, model-driven, and intelligence-enabled frameworks such as Dynamic Value Chain Optimization (DVCO).

## 2.4 Limitations of Traditional Linear Planning Models

Despite the well-documented dynamic characteristics of Engineer-to-Order (ETO) value chains, many industrial planning and execution approaches continue to rely on linear, deterministic models that were originally developed for stable and repetitive production environments. These models typically assume clearly defined phase boundaries, fixed requirements, and predictable information handoffs between engineering, planning, manufacturing, and installation. While such assumptions may hold in Make-to-Stock (MTS) or Assemble-to-Order (ATO) contexts, they are fundamentally misaligned with the complexity and volatility inherent in ETO environments.

Traditional ETO planning practices are often structured around static Work Breakdown Structures (WBS) and milestone-driven project schedules. These approaches presuppose that engineering outputs can be finalized early and subsequently “handed off” to downstream functions with minimal iteration. However, empirical studies have shown that ETO projects experience frequent design changes that propagate across multiple lifecycle stages, rendering static baselines rapidly obsolete (Gosling and Naim, 2009). As a result, planning systems that depend on frozen structures struggle to accommodate late-stage modifications without extensive manual intervention and schedule disruption.

Enterprise Resource Planning (ERP) and Material Requirements Planning (MRP) systems further reinforce linear assumptions by embedding deterministic logic into production planning and procurement processes. These systems typically require stable bills of materials, routings, and lead times to generate feasible plans. In ETO contexts, where engineering bills of materials

(EBOMs) and manufacturing structures evolve iteratively, such stability rarely exists. Hicks and McGovern (2009) highlight that conventional ERP-based planning mechanisms lack the flexibility required to manage concurrent engineering and production activities, leading to frequent replanning cycles and diminished decision confidence.

Change management processes in traditional ETO environments are similarly constrained by document-centric and siloed organizational structures. Engineering change management (ECM) is often treated as an exception-handling function rather than a core operational capability. Changes are communicated through manual notifications, spreadsheets, and meetings, with limited automated assessment of downstream impacts. This fragmented approach delays decision-making and increases the risk of inconsistencies between engineering intent and execution reality. Gosling and Naim (2009) describe this phenomenon as a key contributor to coordination failures in ETO supply chains, particularly when changes originate from multiple sources simultaneously. Another critical limitation of linear planning models lies in their inability to integrate real-time feedback from manufacturing, installation, and operation. Traditional planning frameworks are predominantly forward-looking, relying on predefined assumptions rather than continuously updated operational data. As a result, deviations identified during production, Factory Acceptance Testing (FAT), or Site Acceptance Testing (SAT) are often addressed through reactive measures rather than proactive optimization. This reactive posture not only increases cost and lead-time variability but also undermines the potential value of operational learning across projects (Hicks and McGovern, 2009). Furthermore, linear models inadequately support cross-functional and cross-organizational coordination, which is essential in

ETO environments characterized by early supplier involvement and distributed execution. Suppliers and subcontractors frequently operate with partial or outdated information, as traditional systems lack mechanisms for synchronized, model-driven collaboration. The absence of a shared, authoritative representation of requirements, designs, and execution states exacerbates misalignment and amplifies the impact of change across the value chain. Collectively, these limitations demonstrate a fundamental mismatch between the dynamic nature of ETO value chains and the static assumptions embedded in traditional planning and execution models. Linear approaches are not merely inefficient in this context; they are structurally incapable of supporting the continuous adaptation required in ETO/D2O environments. As ETO projects increase in scale, customization, and integration complexity, the shortcomings of deterministic, document-driven systems become increasingly pronounced.

## 2.4 Model-Based Systems Engineering (MBSE) and the Model-Based Enterprise (MBE)

### 2.4.1 From Document-Based to Model-Centric Engineering

Historically, engineering information has been managed through distributed documents such as CAD files, spreadsheets, specifications and PDF reports. Such artefacts are typically fragmented across organisational silos, creating persistent issues with version control, inconsistent interpretations and slow propagation of design updates (Hedberg et al., 2016).

MBSE replaces this paradigm by positioning the model as the authoritative representation of system requirements, behaviour, structure and interfaces. In this approach, the model does not merely describe the system—it *is* the system representation, providing a consistent and analysable basis for design, verification and lifecycle management (Hedberg and Feeney, 2017).

#### 2.4.2 Principles of Model-Based Systems Engineering

The literature describes MBSE as grounded in a set of principles that define how engineering knowledge is formalised and managed across complex systems. Four principles are widely emphasised:

Formalisation.

Engineering knowledge is expressed through structured, machine-interpretable models with explicit syntax and semantics. This reduces ambiguity and provides a basis for automated analysis and validation.

Integration.

MBSE brings together mechanical, electrical, software, controls and operational domains within a unified modelling framework, improving cross-disciplinary alignment and supporting integrated verification.

Consistency.

Model changes propagate across dependent views and artefacts, preventing divergence and ensuring that structural and behavioural representations remain synchronised throughout development.

Lifecycle Connectivity.

MBSE models persist from early requirements and conceptual design through simulation, detailed engineering, manufacturing, installation and commissioning, strengthening traceability across the lifecycle (Hedberg et al., 2016).

The literature consistently identifies these principles as important for managing the tightly coupled evolution of complex socio-technical systems.

### 2.4.3 Transition from MBSE to the Model-Based Enterprise (MBE)

While MBSE provides a rigorous foundation for system definition, its scope often remains limited to the engineering domain. The Model-Based Enterprise (MBE) extends these principles across the wider organisation, enabling model-driven execution across manufacturing, logistics, quality and operations (Främling et al., 2013).

Key elements characterising an MBE-enabled organisation include:

- Authoritative digital twins that maintain alignment between engineering intent and operational reality.
- Semantic interoperability across engineering, ERP, MES, PLM and quality systems, reducing ambiguity at system interfaces.
- Integrated models of product, process, resource and facility information, supporting coordinated planning across disciplines.
- End-to-end digital workflows that replace document-based handovers with model-based exchanges.
- Closed-loop feedback where operational data flows back into engineering models to support continual improvement (Hedberg and Feeney, 2017; Hu et al., 2021).

These capabilities position MBE as an organisational extension of MBSE, aiming to maintain fidelity between engineering models and downstream enterprise processes.

#### 2.4.4 Role of BIM within the Model-Based Enterprise

Building Information Modelling (BIM) represents a domain-specific instantiation of model-based principles within architectural, structural and facilities engineering contexts. BIM provides a parametric, geometry-rich representation of built-environment assets that supports multidisciplinary coordination and lifecycle management (Eastman et al., 2011).

Within an MBE context, the literature identifies several contributions of BIM:

- Parametric, geometry-rich modelling of building, structural and mechanical elements.
- Spatial coordination across architectural, structural, mechanical and automation disciplines.
- Clash detection and design-for-assembly analysis.
- Integration of geometry with cost and schedule data for construction planning.
- 4D sequencing for installation and construction processes.
- 6D operations integration linking BIM elements to maintenance, safety and asset information (Bryde, Broquetas and Volm, 2013).

BIM thus serves as a spatial and structural foundation within a model-based enterprise, supporting downstream manufacturing, installation and operational workflows.

#### 2.4.5 MBSE and MBE in Engineer-to-Order (ETO) Environments

ETO environments share characteristics commonly addressed by model-based approaches, including high variability, multidisciplinary integration, late-stage changes and tightly coupled workflows across engineering, supply chain, manufacturing and installation (Gosling and Naim, 2009; Hedberg et al., 2016).

The literature reports that MBSE and MBE practices can:

- reduce rework through authoritative, synchronised engineering models;
- improve propagation of engineering changes across lifecycle artefacts;
- enhance cross-functional alignment through shared model-based representations;
- support scenario analysis and change-impact evaluation;
- improve predictability during installation and commissioning through consistent engineering-data continuity (Hedberg and Feeney, 2017).

However, persistent gaps are noted in the literature, including:

- limited lifecycle scope, where MBSE remains engineering-centric;
- semantic misalignment between engineering tools and enterprise systems;
- weak integration between model-based engineering tools and ERP/MES systems;
- minimal application of AI to model-based reasoning or lifecycle orchestration (Yu and Liu, 2020; Miller and Choi, 2020).

These constraints indicate that while MBSE and MBE offer substantial benefits, they do not fully resolve the enterprise-wide integration challenges present in contemporary ETO value chains.

#### 2.4.6 Contributions of MBSE and MBE to Integrated Digital Engineering

Across the literature, MBSE and MBE are consistently presented as foundational enablers of integrated digital engineering environments. MBSE provides authoritative system models that underpin verification, planning and lifecycle decision-making (Hedberg and Feeney, 2017), while MBE extends these models into enterprise contexts through model-driven workflows and semantic interoperability (Hu et al., 2021).

Research highlights that MBSE and MBE support:

- development of digital twins grounded in formal system structures;
- construction of knowledge graphs encoding engineering relationships;
- reduction of ambiguity in cross-disciplinary communication;
- improvement of traceability and change-impact analysis through model-driven propagation.

Taken together, the literature positions MBSE and MBE as key contributors to the coherence, consistency and analytical capability required for managing complex engineering systems across the lifecycle.

#### 2.5 Digital Twins and Cyber-Physical Systems

Digital twins and cyber-physical systems (CPS) represent foundational concepts within the Industry 4.0 paradigm (Lee, Bagheri and Kao, 2015; Boschert and Rosen, 2016). They provide the structural and behavioural models necessary to integrate physical operations with digital intelligence, enabling real-time monitoring, simulation, and prediction. For Engineer-to-Order (ETO) deployments—such as Symbotic’s warehouse automation systems—digital twins and CPS offer a path to improving lifecycle coherence, reducing rework, and enabling dynamic, data-driven decision-making (Grieves, 2014; Alonso-Rosa et al., 2021).

### 2.5.1 Conceptual Foundations of Digital Twins

The concept of the digital twin originated in aerospace engineering and early product-lifecycle management (PLM) research, where high-fidelity virtual models were used to evaluate performance, simulate system behaviour and reduce design risk before physical production (Grieves, 2014). A digital twin is broadly defined as a virtual representation of a physical asset, system or process that is continuously synchronised with real-world data. This synchronisation enables digital models not only to mirror physical states but also to predict behaviour, optimise performance and support lifecycle decision-making.

Digital twins typically integrate four major categories of information:

- Structural models, including geometry, assemblies, parametric constraints and component relationships.

- Behavioural models, such as kinematic behaviour, thermodynamic characteristics, control logic and operational states.
- Lifecycle data, including manufacturing records, operational history, maintenance events and failure logs.
- Real-time telemetry, incorporating sensor readings, environmental conditions, performance signals and usage patterns (Boschert and Rosen, 2016).

The literature identifies multiple subtypes of digital twins, each representing different scopes and levels of integration:

- Product twins, which represent individual components or subsystems.
- System twins, which model collections of integrated assets functioning together.
- Process twins, which capture operational workflows, production sequences or logistics processes.
- Performance twins, which reflect real-time operational behaviour and enable behavioural prediction (Alonso-Rosa et al., 2021).

While digital twins are widely used across industries such as aerospace, automotive, energy and advanced manufacturing, most implementations remain limited to discrete products, specific machines or isolated processes. The value-chain-level digital twin—one that connects engineering, manufacturing, logistics, installation and commissioning into a unified,

semantically governed lifecycle model—remains significantly underdeveloped in current research and practice.

The Dynamic Value Chain Optimisation (DVCO) framework extends the digital-twin paradigm precisely into this underexplored domain. DVCO synthesises structural twins (BIM and MBSE), execution twins (ERP and MES), performance twins (IoT/CPS), semantic twins (ontology and knowledge graph) and intelligence twins (LLM-enabled reasoning) into an integrated architecture capable of modelling and orchestrating the full Engineer-to-Order (ETO) value chain.

### 2.5.2 Cyber-Physical Systems (CPS) as the Digital Twin Substrate

Cyber-Physical Systems (CPS) form the foundational technological substrate upon which modern digital twins are constructed. CPS integrates computational intelligence, embedded sensing and physical processes through continuous feedback loops, enabling real-time synchronisation between digital logic and physical behaviour. As defined by Lee, Bagheri and Kao (2015), CPS operate through the tight coupling of computation, communication and control, creating systems that can sense their environment, process data, and execute responsive actions with high temporal precision.

Key characteristics of CPS identified in the literature include:

- Tight integration of sensing, computation and control, enabling immediate interpretation of physical states and corrective action.
- Real-time data exchange across distributed nodes, facilitating synchronised operation of multiple subsystems within complex environments.
- Embedded decision logic, allowing systems to function autonomously or semi-autonomously in response to changing conditions.
- Adaptive behaviour, enabling systems to adjust performance based on environmental variability, operational constraints and observed anomalies (Wang, Törnngren and Onori, 2015).

In the context of Symbotic’s warehouse automation deployments, CPS encompass a wide range of robotic and automated subsystems, including robotic manipulators, autonomous mobile storage bots, conveyance and shuttle systems, automated storage and retrieval systems (AS/RS), and a dense layer of environmental and equipment sensors. These CPS components generate the continuous telemetry essential for maintaining an accurate and synchronised digital twin.

As such, CPS provides the physical–digital interface that allows DVCO to monitor real-world states, detect deviations, and drive dynamic lifecycle updates across engineering, manufacturing, installation and commissioning workflows. Their integration into DVCO ensures that operational behaviour remains tightly aligned with engineering intent and digital representations, thereby enabling real-time decision-making and adaptive orchestration across the value chain.

### 2.5.3 Digital Twin Applications in Engineering and Operations

The literature identifies a broad set of applications for digital twins across engineering, manufacturing, logistics and operations. These applications illustrate the versatility of digital-twin technology and highlight its relevance to high-variability Engineer-to-Order (ETO) environments.

#### Design Verification and Simulation

Digital twins are widely used during engineering design to validate physical feasibility, evaluate system behaviour and identify clashes or interferences prior to construction or installation. As Grieves (2014) explains, digital models allow engineers to detect design defects early, reducing downstream rework and improving reliability at handover.

#### Operational Monitoring

A major application of digital twins lies in their ability to reflect real-time equipment performance using telemetry and embedded sensors. This enables anomaly detection, predictive maintenance and continuous monitoring of system health (Gao et al., 2015). Such capabilities are essential in robotics-intensive environments where operational deviations can rapidly impact commissioning and throughput.

#### Process Optimisation

Process-level twins simulate flow, throughput, bottlenecks and resource utilisation in manufacturing or logistics systems. These simulations support optimisation of cycle times, production sequencing and material-handling strategies, improving productivity and reducing operational variability.

### Lifecycle Traceability

Digital twins also support multi-stage lifecycle traceability by maintaining accurate records of installation, commissioning, testing and operational history (Boschert and Rosen, 2016). This traceability is critical for compliance, root-cause analysis and continuous improvement, especially in complex ETO deployments where data sources are typically fragmented.

### What-If and Scenario Analysis

Scenario-based modelling is another prominent use case. Alonso-Rosa et al. (2021) demonstrate that digital twins enable organisations to perform what-if analyses on design alternatives, resource allocation strategies, installation sequences and mitigation plans. This supports informed decision-making in environments where risks and constraints change rapidly.

For ETO deployments, these applications directly address the uncertainty, volatility and cross-domain dependencies that characterise engineering-intensive projects. They provide the analytical foundation for proactive planning, early risk detection and adaptive orchestration—capabilities that the DVCO framework extends through semantic integration and AI-enabled reasoning.

#### 2.5.4 Limitations of Current Digital Twin Implementations

Although digital twins are widely recognised as a core enabler of Industry 4.0, the literature consistently highlights several limitations that constrain their effectiveness in complex, high-variability ETO environments. These limitations are structural, semantic and architectural in nature, preventing current digital-twin systems from delivering the full spectrum of lifecycle intelligence required for Dynamic Value Chain Optimisation (DVCO).

##### Fragmented and Domain-Isolated Models

Digital twins are frequently implemented as isolated, domain-specific models—such as mechanical, electrical, thermal or control-system twins—without unified semantics or cross-domain integration. As Hu et al. (2021) note, this fragmentation limits the ability to represent interdependencies across engineering, manufacturing, supply chain, installation and commissioning.

##### Static or Partially Synchronized Twins

A large proportion of industrial “digital twins” in practice are either static models or only partially synchronized with real-world data streams. Alonso-Rosa et al. (2021) argue that without continuous alignment to real-time operational data, digital twins cannot support dynamic

replanning or adaptive decision-making—capabilities essential for ETO environments where conditions change rapidly.

### Engineering-Centric Scope

The traditional digital-twin paradigm remains heavily oriented toward engineering design and machine-level monitoring. Grieves (2014) notes that most twin implementations do not extend into downstream domains such as production, logistics, installation or commissioning. As a result, lifecycle coordination problems remain unaddressed.

### Lack of Semantic Interoperability

Yu and Liu (2020) emphasise that the absence of semantic interoperability limits the usefulness of digital twins. Without ontologies or knowledge graphs, data from twins cannot be consistently integrated with ERP, MES, BIM or field-execution systems. This results in ambiguous mappings, inconsistent interpretations and fragmented decision-making.

### Limited Integration with AI and Explainable Reasoning

While some implementations incorporate machine learning for predictive maintenance or anomaly detection, very few digital-twin systems integrate:

- explainable AI,
- natural-language reasoning,
- semantic search, or

- graph-driven interpretation of lifecycle dependencies.

Sheth (2020) and Zhang et al. (2021) highlight that without these capabilities, digital-twin systems lack the intelligence required for proactive decision support, value-chain reasoning and cross-domain orchestration.

Taken together, these limitations prevent conventional digital twins from realising their full potential in high-variability ETO contexts. They remain powerful but fragmented tools—strong in isolated domains but insufficient for holistic lifecycle optimisation. This gap underscores the need for DVCO’s integrated architecture, which augments digital twins with ontologies, knowledge graphs and AI-enabled reasoning to support end-to-end, semantically grounded value-chain optimisation.

## 2.6 Ontologies, Semantics, and Knowledge Graphs in Engineering and Operations

Ontologies and knowledge graphs (KGs) have emerged as critical technologies for addressing long-standing challenges in data fragmentation, semantic inconsistency and lifecycle disconnect across engineering, manufacturing and operational systems (Berners-Lee, Hendler and Lassila, 2001; Beetz, van Leeuwen and de Vries, 2009; Yu and Liu, 2020; Sheth, 2020). In complex Engineer-to-Order (ETO) environments—where information flows span BIM, ERP, MES, IoT and field-execution platforms—ontologies provide the formal definitions needed to unify

heterogeneous data sources, while knowledge graphs enable dynamic representation of interdependencies. These technologies form an essential foundation for AI-enabled reasoning, digital twins and model-based execution frameworks such as DVCO.

### 2.6.1 The Role of Semantics in Complex Engineering Environments

Semantic technologies address one of the most persistent structural challenges in engineering-intensive and multi-disciplinary industries: the absence of shared meaning across heterogeneous systems, tools, and organisational stakeholders. As Tolk and Muguira (2003) observe, traditional integration mechanisms—such as file transfers, APIs, database mappings, and point-to-point connectors—enable systems to exchange data syntactically but do not ensure semantic interoperability. Without a shared conceptual layer, such integrations fail to capture:

- conceptual meaning, such as what a component or task represents within a lifecycle;
- lifecycle context, including how data evolves from engineering to installation;
- relational dependencies, such as multi-step impacts across design, manufacturing and commissioning;
- constraints and rules, which govern valid configurations, processes or operational behaviours.

Ontologies provide a solution by defining formalised vocabularies—concepts, attributes, relationships and constraints—that enable consistent interpretation of information across

domains (Beetz, van Leeuwen and de Vries, 2009; Yu and Liu, 2020). They serve as the semantic backbone of integrated environments, allowing systems to understand not merely *data*, but *meaning*.

A simple but illustrative example highlights the challenge. Consider three commonly used terms within industrial environments:

- a Work Center in a Manufacturing Execution System (MES),
- a Cost Center in SAP S/4HANA, and
- an Installation Zone in a BIM model.

Although these terms may refer—directly or indirectly—to the same physical or logical location within a facility, each system embeds its own semantics, constraints and interpretations. Without a unifying ontology, they remain disconnected representations of the same underlying entity. As a result:

- engineering models cannot be reliably linked to operational workflows;
- cost structures may not align with spatial or process structures;
- installation tasks may be sequenced without accurate engineering or spatial grounding;
- cross-functional reasoning becomes error-prone, slow and heavily dependent on tacit knowledge.

By providing shared semantics, ontologies bridge these interpretive gaps, enabling lifecycle alignment, multi-domain reasoning and integrated decision-making—capabilities essential for complex ETO environments and foundational to the DVCO framework.

### 2.6.2 Ontologies in Engineering, Manufacturing and Supply Chain Domains

Ontology research across engineering, manufacturing and operations has expanded substantially over the past two decades, resulting in a range of domain-specific semantic models designed to structure technical knowledge and support interoperability. In engineering design, product ontologies define hierarchical parts–assembly relationships, tolerances, material attributes and engineering constraints, providing semantic clarity for complex product structures. In manufacturing contexts, process ontologies represent routings, resource requirements, capacities, sequencing rules and operational constraints, enabling machine-interpretable modelling of production logic and workflows.

In the built-environment domain, facility and spatial ontologies—particularly those derived from Building Information Modelling (BIM) and the Industry Foundation Classes (IFC) schema—are used to model building elements, spatial zones, geometric representations, and the dependency and adjacency relationships between them (Beetz, van Leeuwen and de Vries, 2009). Such ontologies facilitate coordination across architectural, structural and mechanical disciplines and

provide the semantic foundation required for installation planning, commissioning activities and downstream operational workflows..

Operational and industrial IoT environments rely on operations and telemetry ontologies, which describe equipment states, events, alarm conditions, performance variables and sensor behaviours (Saggi and Jain, 2018). These ontologies enable structured interpretation of real-time data and provide a foundation for digital-twin synchronisation. Complementing these, supply chain and logistics ontologies articulate suppliers, lead times, shipment structures, dependencies, risk indicators and procurement relationships, supporting the semantic integration of purchasing and logistics processes (Kovacs and Kot, 2017).

Despite the breadth of research across these individual domains, the literature consistently highlights a significant limitation: few ontologies span the entire value chain from engineering through manufacturing, supply chain, installation, commissioning and operations. Even fewer integrate BIM-based engineering semantics with ERP/MES operational semantics or IoT-driven behavioural models (Yu and Liu, 2020). As noted in systems-integration literature, cross-domain lifecycle ontologies remain scarce, and existing efforts typically remain siloed within specific disciplines or system boundaries (Lin, Lee and Ma, 2020).

This fragmentation underscores a central gap addressed by the DVCO framework. DVCO explicitly aims to develop a unified, lifecycle-spanning semantic layer that harmonises engineering, manufacturing, supply chain, installation and IoT semantics into a shared ontology and knowledge graph. Such integration is foundational for enabling dynamic value-chain

coordination, model-driven orchestration and AI-enabled lifecycle reasoning across complex ETO environments.

### 2.6.3 Knowledge Graphs as Lifecycle Integration Structures

Knowledge graphs (KGs) extend ontology-defined concepts by structuring them into interconnected networks of real-world instances—such as engineering components, manufacturing orders, suppliers, sensors, installation zones, and commissioning tasks (Sheth, 2020; Zhang et al., 2021). Unlike traditional data models, KGs are explicitly designed to capture complex, multi-domain relationships across the full lifecycle, enabling them to function as dynamic integration structures for Engineer-to-Order (ETO) environments.

KGs excel at modelling several properties essential for lifecycle integration:

- Multi-domain relationships, linking engineering models, enterprise systems, operations and telemetry in a unified structure.
- Dependency networks, representing how components, tasks or materials influence one another across design, procurement, manufacturing and installation.
- Bidirectional traceability, allowing upstream and downstream queries such as how an engineering change affects procurement and commissioning, or how an IoT anomaly traces back to design constraints.

- Contextual reasoning, providing semantically rich information for both human and AI-driven decision-making.
- Dynamic propagation, enabling automatic spread of changes or risk signals across connected lifecycle elements.

These properties make KGs particularly well suited for ETO environments, where inherent variability creates high interdependence across domains. In such contexts:

- Engineering changes often proliferate into manufacturing routings, supplier selections and installation sequences.
- Supplier delays cascade into project-milestone risks and commissioning readiness constraints.
- IoT anomalies signal deviations that propagate through operational states, testing procedures and engineering baselines (Lin, Lee and Ma, 2020).

Traditional relational databases rely on rigid schemas and predetermined join structures, limiting their ability to represent evolving and multi-hop lifecycle dependencies. By contrast, KGs provide a flexible, graph-based structure capable of representing evolving relationships, heterogeneous entities and dynamic interdependencies—features that mirror the real complexity of industrial value chains (Zhang et al., 2021).

As such, KGs serve as the backbone for DVCO’s lifecycle integration, enabling coherent data representation, cross-domain analytics and AI-enabled reasoning across Symbotic’s ETO value chain.

#### 2.6.4 Ontology + KG as an AI Enabler

Artificial intelligence—particularly large language models (LLMs)—has advanced rapidly in recent years, yet its effectiveness in complex industrial environments remains highly dependent on the quality, structure and semantic consistency of underlying data. As the literature emphasises, LLMs operating without domain grounding often exhibit several critical limitations: they may hallucinate when confronted with inconsistent inputs, misinterpret specialised engineering terminology and fail to infer multi-step dependencies that are essential to understanding lifecycle relationships in Engineer-to-Order (ETO) environments (Dignum, 2019; Sheth, 2020). These weaknesses significantly restrict the reliability of standalone LLMs in high-stakes engineering and operations contexts.

To overcome these limitations, the integration of ontologies and knowledge graphs (KGs) has emerged as a powerful enabler of AI reasoning in industrial systems. Ontologies provide explicit semantic structures—concepts, attributes, rules and constraints—that define consistent meaning across engineering, manufacturing, supply chain, installation and IoT domains. Knowledge graphs operationalise these semantics by linking real-world instances into rich networks of

lifecycle relationships, allowing computational models to navigate dependencies that traditional data structures fail to capture.

When LLMs are grounded in a knowledge graph, they gain several critical enhancements:

- Contextual knowledge: concepts, constraints and lifecycle semantics are explicitly defined, reducing ambiguous or incorrect interpretations.
- Logical structure: relationships and rules constrain LLM outputs to domain-valid logic, improving reasoning consistency.
- High-accuracy retrieval: KG-based retrieval-augmented generation (RAG) ensures that LLMs access correct and up-to-date information rather than relying solely on statistical memory (Zhang et al., 2021).
- Explainability and traceability: each AI-generated insight can be mapped to specific nodes, edges and reasoning paths within the KG, supporting transparent and auditable decision-making (Samoila, López and Gutiérrez-Ríos, 2021).

Across the literature, this hybrid AI architecture—LLM + KG + Ontology—is increasingly recognised as the most promising paradigm for industrial decision intelligence. It ensures that AI systems are not only powerful but also reliable, explainable and aligned with engineering and operational semantics. In the context of Dynamic Value Chain Optimisation (DVCO), this hybrid approach forms the core of the decision-intelligence layer, enabling accurate cross-domain reasoning and lifecycle-informed predictions essential for real-time value-chain optimisation.

## 2.6.5 Applications of Ontologies and KGs in ETO and Industry 4.0 Literature

Research across Engineering-to-Order (ETO), digital manufacturing and Industry 4.0 consistently highlights the growing role of ontologies and knowledge graphs (KGs) as foundational enablers of lifecycle integration, semantic consistency and intelligent decision support. The literature identifies several key application domains where these technologies provide measurable value.

### 2.6.5.1 Change-Impact Analysis

Knowledge graphs enhance change-impact analysis by tracing multi-step dependencies across engineering, procurement, manufacturing, installation and commissioning. When a design modification occurs, KGs can automatically identify affected components, tasks, materials and upstream/downstream dependencies with far greater completeness than traditional document-based methods (Zhang et al., 2021). This capability is essential in ETO environments where engineering changes propagate broadly across the value chain.

### 2.6.5.2 Engineering–Manufacturing Alignment

Ontologies play a crucial role in maintaining semantic alignment between engineering and manufacturing artefacts, including EBOM, MBOM, routings, operation sequences and work centers. Studies show that semantic structures reduce BOM drift, avoid structural inconsistencies

and create a shared interpretation of lifecycle objects across engineering and production teams (Yu and Liu, 2020). This supports more accurate planning, scheduling and material synchronisation.

#### 2.6.5.3 Installation and Commissioning Management

In the downstream phases of ETO execution, KGs support installation and commissioning by representing spatial and temporal dependencies between site tasks, zones, equipment, safety constraints and readiness conditions. This enables better detection of sequencing conflicts, zone congestion and task interdependencies, significantly improving field coordination and execution stability.

#### 2.6.5.4 Supply Chain Risk Analysis

The literature highlights the value of KGs for modelling supply-chain dependencies, including vendor reliability, historical lead-time variation and multi-tier material relationships. KG-based risk propagation identifies how delays or disruptions from specific suppliers ripple through BOM structures, production orders and installation milestones (Kovacs and Kot, 2017), providing organisations with earlier visibility into critical-path vulnerabilities.

#### 2.6.5.5 Digital Twin Synchronisation

Ontologies and KGs are central to synchronising digital twins with engineering, manufacturing and operational data. KGs unify IoT telemetry with process and engineering models, enabling

closed-loop digital twins capable of real-time deviation detection, anomaly localisation and lifecycle-aware reasoning (Hu et al., 2021). This integration enhances the fidelity, responsiveness and explainability of digital-twin architectures.

#### 2.6.5.6 AI Reasoning and Natural-Language Querying

An emerging stream of literature explores hybrid LLM–KG reasoning architectures, where large language models are grounded in structured knowledge for improved accuracy, domain relevance and explainability. This enables intuitive natural-language querying across lifecycle systems, such as:

*“Which installation tasks are at risk if Vendor A delays shipment X?”*

LLM-KG integration ensures that responses reflect accurate dependencies rather than ungrounded statistical inference (Sheth, 2020; Samoila, López and Gutiérrez-Ríos, 2021).

#### 2.6.6 Current Limitations in Semantic and KG Adoption

Although semantic technologies and knowledge graphs (KGs) offer significant potential for improving lifecycle intelligence in complex Engineer-to-Order (ETO) environments, the existing literature identifies several limitations that currently hinder widespread industrial adoption.

First, ontology development remains highly labour-intensive, requiring extensive cross-domain expertise from engineering, manufacturing, IT, operations and data-modelling specialists. Unlike

traditional integration methods, ontology engineering demands iterative conceptual modelling, consensus building and careful formalisation, which many organisations lack the capacity to sustain (Beetz, van Leeuwen and de Vries, 2009).

Second, there are few mature standards that unify semantics across BIM, ERP, MES and IoT systems. While IFC provides a partial standard for BIM, it does not map cleanly to manufacturing or enterprise semantics, nor does it capture operational telemetry or commissioning logic (Yu and Liu, 2020). The absence of cross-domain semantic standards leads to fragmented interpretations of lifecycle entities and limits interoperability across tools.

Third, industrial knowledge graphs rarely scale effectively to the multi-site, multi-deployment nature of ETO operations. Existing implementations tend to be localised, monolithic or domain-specific, limiting their applicability across distributed supply chains and orchestrated installation environments (Zhang et al., 2021). Scaling KGs requires robust governance, modular architecture and high-performance graph databases, which many organisations have yet to implement.

Fourth, organisational readiness and cultural barriers slow semantic-technology adoption. Traditional engineering cultures often prioritise tacit knowledge, document-centric workflows and siloed system ownership, making the transition to model-driven and semantically governed environments difficult (Nonaka and Takeuchi, 1995). Resistance to change remains a persistent barrier.

Fifth, AI reasoning over KGs is still emergent, particularly in low-data, high-variability ETO domains. Most AI studies focus on high-volume, repetitive industrial processes; by contrast, ETO environments exhibit sparse data, one-off configurations and complex lifecycle dependencies that challenge conventional machine-learning techniques (Samoila, López and Gutiérrez-Ríos, 2021).

These limitations highlight the need for a new integrative framework—such as Dynamic Value Chain Optimisation (DVCO)—that embeds semantic technologies within a broader technological ecosystem comprising digital twins, IoT telemetry, enterprise systems and AI-enabled reasoning. As argued by Lin, Lee and Ma (2020), only a holistic architecture can unlock the full value of semantic interoperability and support dynamic, cross-domain lifecycle optimisation in complex industrial settings.

#### 2.6.7 Contribution of Semantic Technologies to DVCO

Semantic technologies—particularly ontologies and knowledge graphs—constitute the foundational intelligence layer of the Dynamic Value Chain Optimisation (DVCO) framework. Their contribution extends beyond data integration, enabling DVCO to achieve capabilities that traditional interface- or API-based architectures cannot support.

First, ontologies provide unified, cross-domain lifecycle semantics, allowing engineering, manufacturing, logistics, installation and commissioning systems to interpret concepts

consistently despite originating from heterogeneous platforms such as BIM, SAP, MES and IoT (Sheth, 2020). This semantic coherence resolves terminological ambiguity and ensures that lifecycle entities (e.g., components, tasks, materials, zones) share consistent meaning across the value chain.

Second, the knowledge graph (KG) enables end-to-end lifecycle traceability, providing explicit relational pathways that link engineering intent to procurement dependencies, manufacturing routings, installation constraints and commissioning requirements. These relationships support multi-directional traceability and make it possible to analyse lifecycle behaviour holistically (Zhang et al., 2021).

Third, semantic structures support dynamic dependency propagation, enabling DVCO to model how a single change—whether an engineering revision, supplier delay or IoT anomaly—cascades across interconnected lifecycle entities. This propagation capability is essential in ETO environments where dependencies are complex, multi-hop and often implicit.

Fourth, ontology- and KG-based representations provide interoperable digital-twin structures by linking physical-world telemetry with engineering and operational models. This creates a unified semantic layer through which digital twins can coordinate multi-domain behaviours, including spatial, temporal and process relationships.

Finally, semantic technologies underpin explainable, AI-enabled decision-making. When LLM reasoning is grounded in ontology and KG structures, outputs become traceable, auditable and

aligned with engineering semantics, thereby mitigating hallucination risks and improving trustworthiness in high-stakes industrial decisions (Sheth, 2020).

Collectively, these capabilities enable DVCO to progress beyond traditional integration approaches—such as point-to-point interfaces, static mappings or system-level API orchestrations—toward a model-driven, semantics-aware orchestration layer capable of governing complex ETO workflows dynamically and consistently (Lin, Lee and Ma, 2020). Semantic technologies therefore serve as both the structural and cognitive foundation upon which DVCO’s advanced lifecycle reasoning and optimisation capabilities are built.

## 2.7 Artificial Intelligence and Large Language Models in Operations and Value-Chain Management

Artificial Intelligence (AI) has emerged as one of the most transformative forces in operations, manufacturing and value-chain management (Davenport and Ronanki, 2018; LeCun, Bengio and Hinton, 2015). Traditional applications of AI have centred on machine learning (ML) techniques for forecasting, defect detection, optimisation and anomaly identification. However, recent advancements—especially in Large Language Models (LLMs)—have introduced a new class of AI systems that combine reasoning, natural-language understanding, semantic retrieval, and contextual decision support (Kaplan and Haenlein, 2019; Wooldridge, 2021). In complex Engineer-to-Order (ETO) environments, where structured data is sparse and dependency chains

are intricate, AI's potential extends beyond automation into the domain of augmented intelligence, supporting human decision-makers across engineering, manufacturing, installation and commissioning.

### 2.7.1 Evolution of AI in Operations and Manufacturing

Early AI applications in industrial settings were predominantly rule-based expert systems, offering deterministic logic for troubleshooting or diagnosis (Davis and Meyer, 1998). The rise of machine learning shifted industry focus toward pattern recognition, predictive maintenance and statistical forecasting using sensor data and historical trends (Gao et al., 2015).

Examples include:

- predictive maintenance models for machine health,
- quality-inspection algorithms using computer vision,
- production scheduling heuristics,
- demand and inventory forecasting.

While effective in repetitive, high-volume environments, these techniques often struggle in ETO contexts due to:

- low data repetition,

- evolving engineering configurations,
- inconsistent process structures,
- unique BOMs and routings per project (Gosling and Naim, 2009).

This limitation has prompted researchers to explore more flexible, reasoning-driven AI approaches (Dignum, 2019).

### 2.7.2 Emergence of Large Language Models in Industrial Applications

Large Language Models (LLMs)—including GPT, PaLM, LLaMA and a growing set of domain-specialised variants—represent a significant advancement in artificial intelligence with direct implications for industrial and engineering contexts. These models exhibit capabilities that extend far beyond traditional NLP techniques, including contextual reasoning, knowledge retrieval, code and script generation, and domain-sensitive interpretation of technical content (Kaplan and Haenlein, 2019; Wooldridge, 2021). As such, they have rapidly become foundational technologies in digital-transformation initiatives across manufacturing, engineering and supply-chain management.

Several advantages of LLMs are particularly relevant to complex industrial environments:

- Ability to process and interpret unstructured engineering documents  
LLMs can analyse specifications, change orders, design notes, logs and textual requirements that historically required manual review.
- Natural-language interaction with heterogeneous data systems  
Users can query BIM, SAP, MES, IoT and scheduling data using conversational language rather than specialised query interfaces.
- Reasoning over multi-step dependencies  
LLMs can identify cross-domain relationships—such as linking engineering changes to procurement, manufacturing and installation consequences—especially when combined with structured semantics.
- Adaptability with minimal retraining  
Unlike conventional machine-learning models, LLMs can generalise to new scenarios and project configurations without extensive domain-specific datasets.
- Compatibility with semantic architectures  
LLMs can be grounded in ontologies and knowledge graphs to improve precision, reduce hallucinations and enhance explainability (Sheth, 2020).

These capabilities align closely with the requirements of Engineer-to-Order (ETO) organisations, where workflows involve large volumes of textual, semi-structured and cross-functional data that do not conform to repetitive or high-volume statistical patterns. As a result, LLMs are emerging as powerful tools for augmenting engineering, planning, operations and decision-support processes within the ETO value chain.

### 2.7.3 LLMs and Knowledge Graphs: A Synergistic Architecture

Although Large Language Models (LLMs) demonstrate remarkable generalisation and linguistic fluency, they remain susceptible to hallucinations, inconsistencies and domain-specific inaccuracies when applied to specialised industrial environments such as Engineer-to-Order (ETO) value chains. As Dignum (2019) notes, LLMs require explicit grounding mechanisms to ensure reliability, transparency and responsible reasoning, particularly in high-stakes engineering and operational contexts.

Recent literature increasingly emphasises hybrid AI architectures that combine LLMs with structured semantic representations—most prominently ontologies and knowledge graphs (KGs). These approaches address the limitations of purely statistical models by anchoring LLM outputs in validated, machine-interpretable domain knowledge. Scholars such as Sheth (2020) and Samoila, López and Gutiérrez-Ríos (2021) highlight that the integration of LLMs with KGs significantly improves accuracy, semantic interpretability and trustworthiness in industrial decision support.

This hybrid integration produces several notable technical advantages:

- Accurate retrieval of engineering and operational facts  
KG-grounding constrains LLM retrieval to verified entities and relationships, avoiding unsupported or fabricated responses.

- Traceable and explainable reasoning

LLM outputs can be linked to explicit graph paths, enabling transparent justification of decisions—an essential requirement in engineering governance.

- Semantic disambiguation

LLM interpretations of terms such as *component*, *material*, *task*, *zone* or *work center* become more precise when mapped to ontology-defined concepts (Sheth, 2020).

- Lifecycle-aware context

Hybrid architectures allow LLMs to reason not only about isolated facts but also about their position within broader engineering, manufacturing, logistics and installation lifecycles.

- Real-time consistency with project data

When KGs are synchronised with BIM, SAP, MES and IoT systems, LLM reasoning automatically reflects current operational conditions (Zhang et al., 2021).

Together, these benefits demonstrate why hybrid LLM–KG–Ontology architectures are increasingly regarded as essential for industrial AI. Within DVCO, this architecture forms the foundation of the decision-intelligence layer, ensuring that AI-generated insights remain semantically grounded, operationally accurate and aligned with lifecycle constraints (Lin, Lee and Ma, 2020).

#### 2.7.4 AI-Enabled Reasoning in Engineering and Operations

AI-enabled reasoning—particularly through the integration of large language models (LLMs) and graph-grounded semantic systems—offers substantial potential to enhance decision-making across the Engineer-to-Order (ETO) value chain. Unlike traditional analytics, these systems combine probabilistic reasoning with structured semantic context, enabling more accurate, explainable and lifecycle-aware insights into engineering and operational processes.

#### 2.7.4.1 Change-Impact Analysis

LLMs integrated with knowledge graphs can interpret engineering-change requests and automatically identify affected components, routings, installation tasks, supplier dependencies and commissioning sequences. This significantly reduces manual effort and improves the accuracy of cross-domain impact detection (Zhang et al., 2021).

#### 2.7.4.2 Predictive Risk Assessment

AI-driven predictive models can analyse historical lead-time patterns, supplier reliability indicators, manufacturing performance and IoT-based sensor anomalies to forecast potential delays or disruptions. This aligns with prior research demonstrating the value of predictive analytics in supply-chain and operations risk management (Gao et al., 2015; Kovacs and Kot, 2017).

#### 2.7.4.3 Natural-Language Querying Across Systems

A key advantage of LLMs is their ability to provide natural-language access to complex, multi-system data. For example, project managers can pose questions such as: “*What tasks will be delayed if mezzanine installation in Zone C shifts by two days?*” The system can generate a

structured response by traversing graph relationships among BIM models, SAP objects, MES routings and installation dependencies (Sheth, 2020).

To mitigate hallucination risk, LLM outputs are systematically cross-validated against explicit knowledge-graph paths, domain rules, and source-system evidence, ensuring that natural-language responses remain fully traceable, explainable, and operationally trustworthy.

#### 2.7.4.4 Automated Reporting and Documentation

AI systems can synthesise data from BIM, SAP S/4HANA, MES, IoT and CPS platforms to autonomously generate installation progress summaries, commissioning reports, supplier-risk dashboards and integrated lifecycle briefs. This aligns with broader findings on AI-supported automation of knowledge-intensive work (Davenport and Ronanki, 2018).

#### 2.7.4.5 Multi-Scenario Simulation and Adaptive Planning

LLM-assisted reasoning also enables multi-scenario simulation, allowing planners to compare alternative resource allocations, sequencing strategies or change-management options. These capabilities support more informed decision-making in environments characterised by high variability and interdependency.

Collectively, these use cases demonstrate that AI's role in ETO is not to replace expert judgment but to augment it. As Dignum (2019) emphasises, responsible AI design should enhance human decision-making rather than automate critical engineering choices.

### 2.7.5 Limitations of Current AI Approaches in ETO Contexts

Despite significant advances in artificial intelligence, current AI methodologies exhibit several critical limitations when applied to Engineer-to-Order (ETO) environments. First, large language models (LLMs) lack inherent domain grounding when used in isolation; without explicit semantic structures such as ontologies and knowledge graphs, their reasoning remains shallow, inconsistent and prone to hallucination (Dignum, 2019; Sheth, 2020). This limits their applicability in engineering-intensive and safety-sensitive contexts where semantic precision is essential.

Second, ETO environments are characterised by low repeatability and sparse historical datasets, which constrain the effectiveness of traditional machine-learning models (Gosling and Naim, 2009). Unlike high-volume manufacturing, ETO project data often varies significantly across deployments, resulting in insufficient patterns for supervised learning and undermining predictive accuracy.

Third, the black-box nature of many AI systems reduces interpretability and trust in high-stakes engineering decision-making (Dignum, 2019). Engineering, installation and commissioning tasks require traceable logic, yet contemporary AI models frequently offer limited transparency regarding how conclusions are generated.

Fourth, integration challenges persist between AI models and enterprise platforms such as SAP S/4HANA, MES and PLM systems. Enterprise architectures are traditionally structured,

transactional and deterministic, making interoperability with probabilistic AI reasoning difficult without substantial mediation layers (Gulledge, 2012).

Finally, safety, ethical and cybersecurity considerations further restrict direct deployment of AI agents in mission-critical industrial contexts. Concerns include model reliability, potential error propagation, and the inability of conventional AI systems to operate within validated engineering constraints (Dignum, 2019).

These limitations collectively underscore the need for the DVCO framework's semantic backbone, wherein AI reasoning is anchored in validated engineering models, lifecycle semantics and cross-domain ontological structures. By integrating LLMs with formal semantics and rule-based knowledge representations, DVCO addresses the shortcomings of current AI approaches and provides a more reliable foundation for decision-making in complex ETO environments (Lin, Lee and Ma, 2020).

#### 2.7.6 AI as a Catalyst for DVCO

Within the Dynamic Value Chain Optimisation (DVCO) framework, artificial intelligence—particularly large language models (LLMs)—functions as the decision-intelligence layer that activates and amplifies the capabilities of the underlying semantic, engineering and operational structures. Rather than operating as isolated prediction engines, AI systems are

embedded within a model-driven and data-centric architecture, enabling the value chain to behave more adaptively, contextually and proactively.

AI contributes to DVCO through several key mechanisms:

- Cross-functional visibility through natural-language interfaces  
LLMs provide intuitive access to complex engineering, manufacturing, installation and commissioning data through conversational queries, reducing cognitive load and improving accessibility across diverse stakeholder groups.
- Dependency awareness via knowledge-graph (KG)-based retrieval  
Grounding LLMs in the DVCO ontology and knowledge graph enables dependency-aware reasoning, ensuring that responses reflect accurate lifecycle relationships rather than probabilistic approximations (Sheth, 2020).
- Predictive insights driven by IoT telemetry and operational signals  
AI models analyse equipment behaviours, site conditions and environmental patterns to forecast bottlenecks, commissioning risks and sequencing conflicts—capabilities essential for high-variability ETO environments (Samoila, López and Gutiérrez-Ríos, 2021).
- Explainable and traceable decision support through graph-grounded reasoning  
LLM outputs are enriched by graph-based explanations, allowing users to trace insights back to specific model elements, relationships or telemetry events, thereby increasing trust and interpretability.

- Continuous organisational learning

AI systems leverage historical project patterns, installation performance data and engineering-change histories to support continuous improvement across deployments.

In this configuration, AI does not replace human engineering or operational expertise; rather, it amplifies it by providing contextualised insights, multi-step reasoning, cross-domain pattern detection and predictive foresight. As Lin, Lee and Ma (2020) emphasise, AI delivers maximal value when embedded within a structured semantic and model-based architecture—precisely the integration DVCO achieves.

Through this role, AI becomes a catalyst for DVCO’s broader objective: transforming ETO execution from a reactive, siloed process into a semantically aligned, continuously learning and intelligence-driven value chain.

### 2.7.7 Summary

The literature highlights that AI’s evolution—from expert systems to machine learning to LLM-driven reasoning—mirrors the increasing complexity of industrial value chains (Davenport and Ronanki, 2018; Wooldridge, 2021). However, current research lacks a unifying architecture that combines MBSE, digital twins, semantic modelling and AI reasoning. DVCO fills this gap by integrating AI into a holistic, lifecycle-oriented framework designed to orchestrate ETO operations with greater intelligence, agility and resilience (Lin, Lee and Ma, 2020).

## 2.8 Literature Gap and Summary

### 2.8.1 Identified Gaps Across the Literature

The preceding review demonstrates extensive scholarly work across Engineer-to-Order (ETO) value-chain management, Industry 4.0 technologies, Model-Based Systems/Engineering (MBSE/MBE), digital twins, semantic data modelling, and artificial intelligence (AI) (Gosling and Naim, 2009; Lee, Bagheri and Kao, 2015; Hedberg et al., 2016; Grieves, 2014; Sheth, 2020). However, although each domain provides important theoretical and practical contributions, no existing body of literature presents an integrated, semantically coherent or operationally executable framework capable of addressing the structural, informational and cross-domain coordination challenges observed in Symbotic’s Design-to-Operate (D2O) environment (Culot et al., 2020; Yu and Liu, 2020).

The gaps emerging from the literature can be grouped into five main areas.

#### Gap 1: Lack of Cross-Domain Semantic Integration

Although ontology and knowledge-graph research has advanced across individual domains—including engineering design, manufacturing processes, supply-chain management

and built-environment modelling—the literature reveals significant cross-domain fragmentation. Current semantic efforts remain largely siloed, providing localised meaning within individual systems but failing to establish a unified lifecycle-wide semantic foundation.

Three critical limitations emerge:

No standardised ontology unifying BIM ↔ ERP ↔ MES ↔ IoT systems

Engineering semantics (BIM/MBE) remain disconnected from enterprise semantics (SAP S/4HANA) and operational semantics (MES, IoT). As a result, data exchanged across systems lacks shared conceptual meaning, hindering lifecycle traceability.

Absence of a semantic layer describing dependencies across engineering, production, logistics, installation and commissioning

Existing models do not represent cross-functional relationships such as how an engineering change influences procurement, or how installation constraints relate to manufacturing routings.

Without such lifecycle semantics, digital integration remains structurally shallow.

Limited research addressing lifecycle semantics in ETO-specific contexts

While semantic modelling is well established in BIM and manufacturing domains, very few studies address the high variability, late-stage changes and multi-domain dependencies characteristic of Engineer-to-Order (ETO) environments (Beetz, van Leeuwen and de Vries, 2009; Yu and Liu, 2020; Sheth, 2020).

This lack of cross-domain semantic grounding restricts lifecycle coherence, prevents automated change-impact propagation, and inhibits systematic risk identification. Even existing BIM–SAP integration literature provides mainly geometric–structural mappings rather than a semantic interoperability framework capable of harmonising engineering intent, enterprise processes and operational behaviours across the full Design-to-Operate (D2O) value chain.

## Gap 2: Digital Twin Research Is Fragmented and Domain-Specific

Although digital-twin research has expanded substantially in recent years, the literature remains highly domain-specific and lacks the lifecycle-wide integration required for complex Engineer-to-Order (ETO) environments. Existing studies overwhelmingly concentrate on narrow technical applications, such as:

- component-level or machine-level twins, capturing mechanical, electrical or control-system behaviours;
- predictive maintenance and equipment-health monitoring, using sensor-driven models to forecast anomalies;
- manufacturing-process simulation, evaluating throughput, cycle times or bottleneck patterns;
- facility-level geometric modelling, typically aligned with BIM/IFC or spatial coordination systems (Grieves, 2014; Boschert and Rosen, 2016; Alonso-Rosa et al., 2021).

While valuable, these contributions represent isolated views of the product or process and do not provide a holistic understanding of cross-domain interdependencies.

In contrast, far fewer studies extend digital-twin principles to broader lifecycle constructs, such as:

- end-to-end value-chain twins, capturing interactions across engineering, procurement, manufacturing, logistics, installation and commissioning;  
installation and commissioning twins, which are critical in automation-intensive ETO deployments but largely neglected in current research;
- integrated multi-layer twins that synchronise engineering models (BIM/MBE), execution logic (ERP/MES) and real-time telemetry (IoT/CPS) into a unified system (Hu et al., 2021).

The result is that existing digital-twin literature lacks the conceptual or architectural foundations necessary to support ETO-wide lifecycle optimisation—particularly in settings with high engineering variability, late-stage changes, supplier disruptions and dynamic site conditions. Without integration across engineering, enterprise and operational layers, digital twins cannot enable the predictive, adaptive or cross-functional coordination required for next-generation value-chain orchestration.

Gap 3: MBSE and MBE Are Not Integrated with Operational Systems

Although Model-Based Systems Engineering (MBSE) and Model-Based Enterprise (MBE) provide rigorous methodologies for managing engineering complexity, their integration with operational and enterprise systems remains limited in practice. The literature shows several persistent structural constraints:

- MBSE practices remain largely confined to early engineering phases, focusing on requirements, system architecture and design modelling rather than downstream execution activities.
- MBE initiatives rarely integrate with ERP/MES platforms, such as SAP S/4HANA, manufacturing execution systems, or production planning tools, resulting in disjointed transitions between engineering and operations.
- Neither MBSE nor MBE frameworks incorporate IoT data streams, real-time telemetry or dynamic site conditions, which are essential for synchronising engineering intent with real-world performance (Hedberg and Feeney, 2017; Främling et al., 2013).

This persistent lack of operational integration prevents engineering models from functioning as dynamic, authoritative drivers for the full value chain. Without the ability to connect engineering representations to procurement logic, manufacturing routings, logistics sequencing or installation workflows, MBSE and MBE remain structurally disconnected from the processes they are intended to influence. As a result, ETO organisations cannot leverage model-based engineering to support adaptive planning, real-time impact propagation or cross-domain decision-making.

In ETO environments—where engineering changes, supplier dependencies, spatial constraints and site conditions evolve rapidly—this disconnect severely limits the usefulness of MBSE/MBE. Their inability to interface with execution systems underscores the need for an extended architectural framework such as DVCO, which integrates engineering models with enterprise processes, semantic technologies and IoT-enabled feedback mechanisms to achieve lifecycle-wide orchestration.

#### Gap 4: AI and LLM Research Lacks Domain Grounding and Lifecycle Context

Artificial intelligence (AI)—and, more recently, large language models (LLMs)—has demonstrated significant capability in areas such as predictive analytics, anomaly detection and natural-language interpretation (Davenport and Ronanki, 2018). However, when applied to Engineer-to-Order (ETO) environments, the literature reveals several critical limitations that hinder reliable decision-making and lifecycle-wide integration.

First, LLMs frequently hallucinate or misinterpret engineering, manufacturing and operations terminology when operating without domain-specific semantic grounding. This is particularly problematic in ETO contexts where technical vocabulary, component hierarchies and process semantics differ substantially from general language patterns and require precise interpretation.

Second, traditional machine-learning models perform poorly in low-repeatability, high-variability environments, which are characteristic of ETO value chains. Unlike mass-production settings, ETO processes lack the large, homogeneous datasets required for

stable model training, making conventional ML approaches insufficient for lifecycle reasoning or impact forecasting.

Third, the literature shows very limited exploration of hybrid AI approaches—specifically those combining LLMs with knowledge graphs and ontologies. Such hybrid architectures are essential for grounding reasoning in explicit lifecycle semantics and ensuring that AI outputs reflect real engineering, manufacturing and installation dependencies (Sheth, 2020; Samoila, López and Gutiérrez-Ríos, 2021).

Fourth, AI applications rarely extend into installation, commissioning or field-execution domains, despite the fact that these downstream phases are among the most volatile and risk-sensitive in ETO deployments. Current AI research focuses heavily on manufacturing or supply-chain optimisation, leaving significant gaps in logic required for handling site-specific constraints, commissioning sequences or cross-domain impact propagation.

These shortcomings highlight the need for domain-specific, explainable, and lifecycle-aware AI reasoning frameworks capable of operating reliably across dynamic, multi-system ETO environments. This reinforces the necessity of DVCO’s hybrid architecture—grounded in ontologies, knowledge graphs and model-based representations—to enable AI systems that are both operationally valid and semantically coherent.

Gap 5: No Existing Framework Provides Dynamic Value Chain Optimisation (DVCO)

Despite extensive scholarly contributions across ETO value-chain management, Industry 4.0, MBSE/MBE, digital twins, semantic technologies and AI, no existing body of literature proposes an integrated, semantically governed or AI-enabled framework capable of delivering Dynamic Value Chain Optimisation (DVCO). The review reveals that current research remains fragmented across domains, with each discipline addressing isolated parts of the lifecycle but none offering a holistic, orchestrated value-chain architecture.

Specifically, the literature lacks:

- a unified, semantically driven, AI-enabled value-chain architecture that integrates engineering, manufacturing, supply chain, installation and commissioning data into a coherent digital ecosystem;
- a continuous digital thread connecting engineering models (BIM/MBSE/MBE) to real-time operational signals from ERP, MES, IoT and CPS platforms;
- automated change-impact propagation capable of tracing dependencies across engineering, production, logistics, installation and commissioning stages;
- a multi-domain orchestration system that coordinates lifecycle activities in real time through semantic reasoning and AI-enabled decision support (Lin, Lee and Ma, 2020).

This gap is most evident in automation-intensive ETO environments such as Symbotic's, where:

- engineering changes occur late and frequently,

- supplier variability introduces uncertainty into planning and production,
- installation and commissioning involve spatial and temporal interdependencies, and
- real-time telemetry continuously interacts with engineering intent and operational constraints.

No existing framework in the literature provides the semantic coherence, lifecycle connectivity or AI-assisted reasoning required to manage these dynamic interactions across the Design-to-Operate (D2O) continuum. This absence underscores the need for a new integrative architecture—one that DVCO explicitly proposes and evaluates.

Despite extensive scholarly contributions across ETO value-chain management, Industry 4.0, MBSE/MBE, digital twins, semantic technologies and AI, no existing body of literature proposes an integrated, semantically governed or AI-enabled framework capable of delivering Dynamic Value Chain Optimisation (DVCO). The review reveals that current research remains fragmented across domains, with each discipline addressing isolated parts of the lifecycle but none offering a holistic, orchestrated value-chain architecture.

Specifically, the literature lacks:

- a unified, semantically driven, AI-enabled value-chain architecture that integrates engineering, manufacturing, supply chain, installation and commissioning data into a coherent digital ecosystem;

- a continuous digital thread connecting engineering models (BIM/MBSE/MBE) to real-time operational signals from ERP, MES, IoT and CPS platforms;
- automated change-impact propagation capable of tracing dependencies across engineering, production, logistics, installation and commissioning stages;
- a multi-domain orchestration system that coordinates lifecycle activities in real time through semantic reasoning and AI-enabled decision support (Lin, Lee and Ma, 2020).

This gap is most evident in automation-intensive ETO environments such as Symbotic's, where:

- engineering changes occur late and frequently,
- supplier variability introduces uncertainty into planning and production,
- installation and commissioning involve spatial and temporal interdependencies, and
- real-time telemetry continuously interacts with engineering intent and operational constraints.

No existing framework in the literature provides the semantic coherence, lifecycle connectivity or AI-assisted reasoning required to manage these dynamic interactions across the Design-to-Operate (D2O) continuum. This absence underscores the need for a new integrative architecture—one that DVCO explicitly proposes and evaluates.

### 2.8.3 Summary of Chapter 2

This chapter has examined the principal bodies of scholarship relevant to value-chain execution in Engineer-to-Order (ETO) environments. The review highlights that ETO value chains are inherently complex, multi-layered and interdependent, characterised by engineering variability, supplier uncertainty, spatial–temporal installation constraints and dynamic operational conditions (Gosling and Naim, 2009).

While Industry 4.0 technologies—including IoT, CPS and advanced automation—provide important technological enablers, the literature indicates that they often lack the semantic, integrative and lifecycle-oriented depth required for coherent multi-domain orchestration (Schuh et al., 2017; Culot et al., 2020).

Similarly, Model-Based Systems Engineering (MBSE) and Model-Based Enterprise (MBE) offer strong conceptual frameworks for engineering consistency and cross-disciplinary integration, yet their adoption beyond engineering remains limited. The literature shows that MBSE/MBE models rarely propagate meaningfully into procurement, production, logistics or installation workflows (Hedberg and Feeney, 2017).

The review further confirms that digital twin implementations remain fragmented and domain-specific, often restricted to product-level, machine-level or facility-level representations without full lifecycle integration or unified semantic grounding (Grieves, 2014; Alonso-Rosa et al., 2021).

At the semantic level, ontology and knowledge-graph research demonstrates significant potential to unify lifecycle concepts across engineering, manufacturing and operations. However, actual adoption within industrial contexts remains limited, with few examples of end-to-end semantic governance across BIM, ERP, MES and IoT systems (Beetz, van Leeuwen and de Vries, 2009; Yu and Liu, 2020).

Finally, although AI and large language models (LLMs) have advanced rapidly in recent years, existing literature shows that these methods generally lack grounding in engineering semantics, operational data structures and lifecycle dependencies—significantly limiting their reliability in high-variance ETO settings (Davenport and Ronanki, 2018; Sheth, 2020).

Taken together, these findings demonstrate that no existing domain of scholarship offers a unified, semantically coherent or AI-enabled architecture capable of orchestrating the full ETO lifecycle. The literature points to the need for a new integrative framework that connects design, procurement, production, installation and commissioning in a dynamic, adaptive and model-driven manner, supported by semantic technologies and AI-enabled reasoning (Lin, Lee and Ma, 2020).

Chapter 3 therefore develops such a framework by introducing Dynamic Value Chain Optimisation (DVCO). It presents DVCO's theoretical foundations, architectural structure and core components, establishing the basis for the modelling and evaluation presented in subsequent chapters.

## CHAPTER III

### Conceptual Framework and DVCO Architectural Model

#### 3.1 Introduction

The Dynamic Value Chain (DVC) concept, as articulated by Cubreath, represents a significant theoretical departure from static, linear representations of value creation. Rather than viewing value chains as predetermined sequences of activities, Cubreath conceptualises them as adaptive systems in which value is continuously reconfigured through interactions among customers, partners, resources, and market signals (Cubreath, 2014). This perspective aligns with broader developments in strategic management and operations research that emphasise adaptability, responsiveness, and the continuous reconfiguration of organisational capabilities.

Despite its conceptual importance, the DVC remains primarily descriptive in nature. While it explains *why* value chains must become dynamic, it does not prescribe *how* such dynamism can be operationalised within complex, engineering-intensive environments. In particular, the DVC concept does not define executable system architectures, formal engineering models, semantic integration mechanisms, or governance structures required to support real-time coordination, engineering change propagation, or lifecycle feedback. Consequently, its direct applicability to

Engineer-to-Order (ETO) industries—characterised by deep engineering dependencies, late-stage design changes, and tightly coupled execution phases—remains limited (Gosling and Naim, 2009; Hicks and McGovern, 2009).

ETO environments impose requirements that exceed the scope of conceptual value-chain models. They require explicit representation of engineering intent, synchronisation across heterogeneous systems such as BIM, PLM, ERP, and MES, traceable propagation of engineering changes across planning and execution layers, and continuous feedback from physical execution into upstream decision-making. Prior research on model-based engineering and digital enterprises demonstrates that without authoritative models, semantic alignment, and closed-loop execution mechanisms, dynamic value-chain concepts cannot be implemented at scale (Hedberg et al., 2016; Hedberg and Feeney, 2017).

The Dynamic Value Chain Optimisation (DVCO) framework proposed in this dissertation extends Cubreath’s DVC concept by transforming it from a high-level abstraction into an implementation-ready, model-driven architecture tailored to complex ETO industries. DVCO preserves the central insight of DVC—that value chains are adaptive, non-linear, and continuously reconfigurable—while grounding this insight in concrete architectural mechanisms, including model-based systems engineering, semantic integration, digital twins, and AI-enabled decision intelligence.

In this sense, DVCO does not replace the Dynamic Value Chain concept; it operationalises it.

Where Cubreath’s DVC provides a conceptual understanding of dynamic value creation, DVCO

delivers the structural, semantic, and computational foundations required to realise such dynamism in real-world ETO deployments. This progression from conceptual DVC to executable DVCO constitutes the primary theoretical and practical contribution of this dissertation.

The purpose of this chapter is therefore to establish the conceptual and architectural foundation of the DVCO framework for Engineer-to-Order (ETO) and Design-to-Operate (D2O) industries. ETO/D2O environments exhibit substantially higher complexity than Make-to-Stock (MTS) or Assemble-to-Order (ATO) supply chains due to bespoke engineering, project-specific configurations, high demand variability, multidisciplinary dependency structures, and tightly coupled field-installation requirements. Existing planning and execution systems—largely built on deterministic models and linear process assumptions—struggle to manage this level of complexity effectively (Hicks and McGovern, 2009; Gosling and Naim, 2009).

The DVCO framework addresses this gap by integrating systems engineering principles, Industry 4.0 technologies, cyber-physical feedback mechanisms, and AI-enabled decision intelligence into a unified, closed-loop architecture. The framework is grounded in the premise that modern ETO/D2O value chains operate as model-driven ecosystems, requiring continuous information flow across engineering, planning, production, installation, and service operations (Eigner and Stelzer, 2009; Tao et al., 2019).

Figure 3.1 illustrates the conceptual architecture of the Dynamic Value Chain Optimisation (DVCO) framework across the ETO/D2O lifecycle. At the core of this architecture, Model-Based Engineering (MBE)—integrating BIM, PLM, and engineering configuration models—serves as the authoritative source of design intent. (Eastman et al., 2011; Eigner and Stelzer, 2009). Execution activities are governed through a Software-Defined Execution (SDM) layer encompassing ERP, MES, installation systems, Enterprise Asset Management (EAM), and System Lifecycle Management (SysLM). Cyber-physical synchronisation is achieved through IoT-enabled digital twins, while a decision-intelligence layer based on ontologies, knowledge graphs, and AI reasoning enables semantic coordination, change-impact analysis, and predictive decision support across the lifecycle.

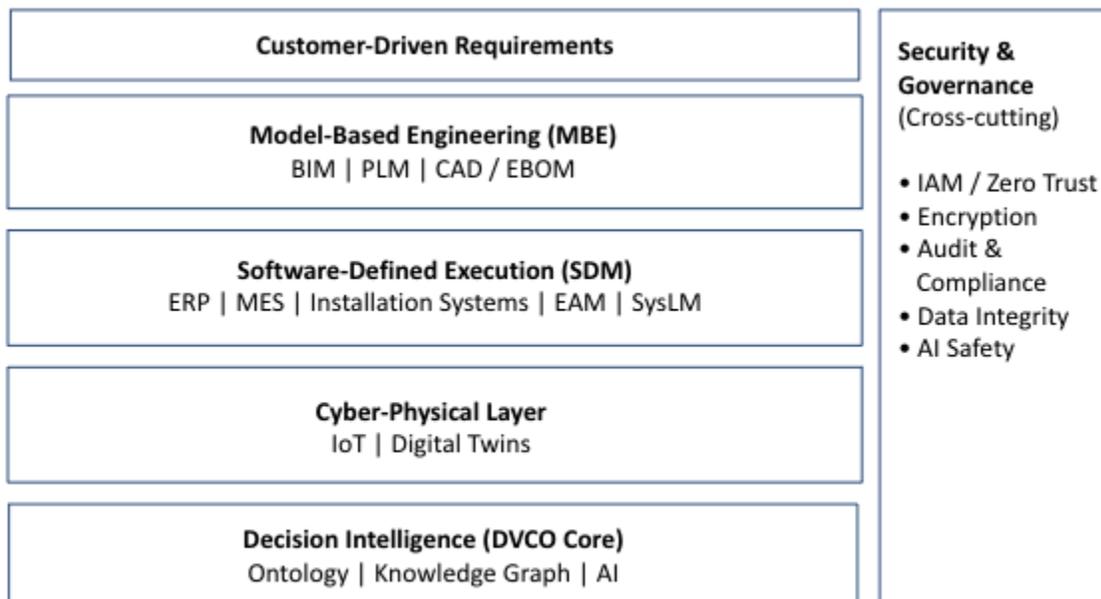


Figure 3.1 Dynamic Value Chain Optimisation (DVCO) framework

Figure 3.1 further visualises the layered nature of the DVCO architecture. Within this structure, the BIM–ERP (SAP) integration layer constitutes the foundational segment of the Digital Thread, positioned downstream of conceptual design intent and upstream of PLM-driven product structures. Specifically, BIM functions as the authoritative source of spatial geometry, constructability constraints, and quantity intelligence (Eastman et al., 2011; Kreider and Messner, 2013)

, forming a pre-PLM semantic layer within the value chain. SAP Project System (PS), SAP BOM structures, and procurement components serve as enterprise execution anchors, translating BIM-derived information into WBS elements, network activities, and project-specific material definitions.

The bidirectional linkage between BIM and SAP depicted in Figure 3.1 represents lifecycle continuity, whereby design changes propagate into planning, scheduling, and cost structures, while enterprise constraints feed back into engineering decision-making. This BIM–SAP layer establishes the structural baseline upon which the DVCO semantic integration layer operates. As illustrated, MES, IoT/CPS, and AI reasoning components are architecturally layered above this foundation rather than treated as isolated extensions. This layering enables real-time execution feedback and adaptive optimisation while preserving traceability to original design intent.

Accordingly, Figure 3.1 should be interpreted not merely as a system integration diagram, but as a semantic lifecycle map. In this map, BIM–SAP integration establishes the initial Digital

Thread, and DVCO extends this thread into a dynamically optimised, closed-loop value chain spanning design, planning, execution, and operation.

### 3.2 Conceptual Foundations of the DVCO Framework

The Dynamic Value Chain Optimisation (DVCO) framework is grounded in the convergence of systems engineering theory, model-based enterprise principles, and Industry 4.0 cyber-physical architectures. Rather than treating value-chain coordination as a sequence of transactional handoffs, DVCO conceptualises the ETO/D2O enterprise as a complex adaptive system whose behaviour emerges from continuous interactions between engineering intent, execution reality, and decision intelligence.(Holland, 1992; Levin, 1998; Maier, 1998).

At its foundation, DVCO adopts Model-Based Systems Engineering (MBSE) and the Model-Based Enterprise (MBE) paradigm as the authoritative mechanism for defining and governing engineering intent. In this context, models are not documentation artefacts but executable knowledge structures that encode requirements, configurations, interfaces, constraints, and dependencies across the lifecycle. This model-centric foundation enables traceability, consistency, and change-impact propagation that are unattainable in document-based environments (Hedberg et al., 2016; Hedberg and Feeney, 2017).

However, DVCO extends beyond conventional MBSE/MBE implementations by explicitly addressing the semantic and operational gaps identified in ETO value chains. Traditional MBE initiatives often remain engineering-centric and weakly coupled to execution systems such as

ERP, MES, and installation platforms. DVCO resolves this limitation by introducing a semantic coordination layer, based on ontologies and knowledge graphs, which provides a shared meaning model across engineering, planning, manufacturing, installation, and service domains. This semantic layer enables machine-interpretable reasoning about dependencies, constraints, and change impacts across heterogeneous systems.

A further conceptual pillar of DVCO is closed-loop cyber-physical integration. DVCO assumes that physical execution—manufacturing, installation, commissioning, and operation—continuously generates signals that must inform upstream engineering and planning decisions. IoT-enabled digital twins serve as the mechanism through which execution states are synchronised with digital models, enabling real-time visibility and adaptive control (Tao et al., 2019). In this sense, DVCO treats the value chain as a living system, rather than a plan to be executed.

Finally, DVCO incorporates AI-enabled decision intelligence as an integral component of value-chain coordination. Rather than relying solely on deterministic rules or static optimisation, DVCO employs knowledge-graph reasoning, predictive analytics, and prescriptive decision services to support scheduling, sourcing, sequencing, and installation strategies under uncertainty. This intelligence layer transforms model-based data into actionable decisions, closing the loop between observation, interpretation, and execution.

Together, these foundations position DVCO as a system-of-systems coordination framework capable of operationalising dynamic value-chain behaviour in complex ETO/D2O environments.

### 3.3 DVCO Architectural Model

This section presents the architectural realisation of the Dynamic Value Chain Optimisation (DVCO) framework. While Section 3.2 established the conceptual foundations of DVCO, this section translates those foundations into a layered, implementation-oriented architecture capable of supporting dynamic coordination across the Engineer-to-Order (ETO) and Design-to-Operate (D2O) lifecycle.

The DVCO architecture is designed to address the structural limitations of traditional ETO planning and execution environments by integrating model-based engineering, software-defined execution, cyber-physical feedback, and AI-enabled decision intelligence within a unified framework. The architecture explicitly supports bidirectional information flows, semantic consistency, and closed-loop optimisation, reflecting the dynamic and non-linear behaviour of contemporary ETO value chains.

#### 3.3.1 Architectural Overview

The DVCO architectural model conceptualises the ETO/D2O enterprise as a system-of-systems, in which engineering, planning, manufacturing, installation, and operation are tightly coupled through shared models and real-time feedback. Rather than treating enterprise systems as loosely

connected applications, DVCO establishes a semantic execution backbone that preserves alignment between engineering intent and execution reality throughout the lifecycle.

At a high level, the architecture comprises five interacting layers:

1. External Environment and Stakeholder Layer, representing customers, suppliers, regulatory bodies, and site conditions that introduce variability and constraints.
2. Model-Based Engineering (MBE) Layer, serving as the authoritative source of design intent through BIM, PLM, and system configuration models.
3. Software-Defined Execution (SDM) Layer, governing enterprise execution across ERP, MES, EAM, SysLM, and installation systems.
4. Cyber-Physical Feedback and Digital Twin Layer, synchronising physical execution with digital representations through IoT-enabled sensing.
5. Decision Intelligence Layer, enabling semantic reasoning, change-impact analysis, and adaptive decision-making through ontologies, knowledge graphs, and AI services.

This layered structure ensures that changes originating in any part of the value chain—engineering, supply, production, installation, or operation—can be detected, interpreted, and propagated coherently across the enterprise.

### 3.3.2 Layered Architecture Description

At the boundary of the architecture, the External Environment and Stakeholder Layer captures the sources of uncertainty inherent in ETO environments, including evolving customer requirements, supplier constraints, regulatory changes, and site-specific installation conditions. These external factors continuously influence engineering and execution decisions, reinforcing the need for adaptive coordination mechanisms.

The Model-Based Engineering (MBE) Layer forms the architectural foundation of DVCO. This layer integrates BIM, PLM, system models, and configuration rules to establish an authoritative, model-centric representation of engineering intent. BIM plays a critical role by providing spatial geometry, constructability constraints, and quantity intelligence, which are essential for downstream planning and installation in ETO contexts. PLM and system models complement BIM by managing product structures, interfaces, and lifecycle configurations.

Beneath the MBE layer, the Software-Defined Execution (SDM) Layer operationalises engineering intent across enterprise execution systems. This layer encompasses ERP, MES, EAM, SysLM, and installation platforms, but differs from traditional execution architectures by enabling dynamic orchestration of workflows, routings, and commissioning activities. Execution

logic is not statically embedded in system configurations but is instead governed by model-driven rules and real-time constraints.

The Cyber-Physical Feedback and Digital Twin Layer establishes continuous synchronisation between physical execution and digital models. IoT-enabled sensors, edge systems, and data infrastructures update asset, process, and installation twins in near real time. This feedback mechanism eliminates lifecycle blind spots by ensuring that deviations detected during manufacturing, installation, or operation are reflected upstream in planning and engineering models.

At the core of the DVCO architecture lies the Decision Intelligence Layer. This layer integrates ontologies, knowledge graphs, predictive analytics, and AI reasoning services to interpret change events, evaluate downstream impacts, and support decision-making under uncertainty. Decisions related to scheduling, sourcing, sequencing, and installation are continuously refined based on both model-based knowledge and real-time execution feedback. Outputs from this layer are fed back into the SDM layer, completing the closed-loop optimisation cycle.

Together, these layers form a coherent architectural system that enables DVCO to move beyond conceptual dynamism and support executable, adaptive value-chain coordination. (Maier, 1998; SO/IEC/IEEE 15288, 2015)

### 3.3.3 Architectural Validation and Traceability to Findings

Although the DVCO framework is introduced in this chapter as a conceptual and architectural model, its validity is empirically examined through the findings reported in Chapter 6. To establish explicit traceability between the proposed architecture and the observed research outcomes, Figure 3.1 is mapped in this subsection to the validated findings (F1–F4).

The DVCO architecture depicted in Figure 3.1 demonstrates a direct correspondence between its major architectural capabilities and the empirical results derived from the case study and simulation-based evaluation.

Table 3.1 – Comparison of DVC (Cubreath) and DVCO (This Dissertation)

Dimension	DVC (Cubreath)	DVCO (This Dissertation)
Nature	Conceptual / descriptive	Implementable / prescriptive
Primary Focus	Dynamic value reconfiguration	Dynamic value-chain optimisation
Engineering Representation	Implicit	Explicit (MBSE / MBE)
Architecture Definition	Not defined	Layered, executable architecture
Semantic Integration	Not addressed	Ontologies and knowledge graphs
System Scope	Business-level abstraction	Engineering–Execution–Operation
Change Propagation	Conceptual	Model-driven, traceable
Execution Feedback	Implicit	IoT-enabled digital twins
Decision Support	Managerial intuition	AI-enabled decision intelligence
Applicability to ETO	Limited	Explicitly designed for ETO/D2O

Interpretive note: DVCO operationalises the principles of DVC by embedding them within a model-driven, semantic, and cyber-physical architecture suitable for complex ETO environments.

### 3.3.4 Architectural Validation and Traceability to Chapter 6 Findings

Although the DVCO framework is introduced in this chapter as a conceptual and architectural model, its validity is empirically examined through the findings reported in Chapter 6. To establish explicit traceability between the proposed architecture and the observed research outcomes, Figure 3.2 is mapped in this subsection to the validated findings (F1–F4).

The DVCO architecture depicted in Figure 3.1 demonstrates a direct correspondence between its major architectural capabilities and the empirical results derived from the case study and simulation-based evaluation.

Table 3.2 Architectural Validation to to Chapter 6 Findings

Chapter 6 Finding	Description	Corresponding Figure 3.1 Architectural Element
F1	Improved cross-functional coordination	Semantic integration layer (ontologies and knowledge graphs)
F2	Reduction of lifecycle blind spots	Digital twin and real-time cyber-physical feedback layer
F3	Faster and more reliable change propagation	Model-Based Engineering (MBE) and Software-Defined Execution (SDM) layers
F4	Enhanced decision quality under uncertainty	AI-enabled decision intelligence layer

### 3.4 Model-Based Engineering (MBE)

Model-Based Engineering provides the foundational structures that enable engineering continuity, system traceability, and configuration-driven behaviour across the lifecycle. MBE replaces document-centric processes with models that formalise product structure, behaviour, requirements, and constraints (Estefan, 2008).

#### 3.4.1 Requirements Modelling and MBSE

SysML and MBSE approaches allow multidisciplinary engineering teams to express system interactions, interfaces, constraints, and behavioural logic (Vaneman and Carlson, 2019). These representations are essential for variant-driven ETO systems.

#### 3.4.2 EBOM, MBOM, Routing, and Configuration Rules

MBE structures define engineering BOMs (EBOM), manufacturing BOMs (MBOM), routing trees, and constraint logic. The literature highlights that rigorous EBOM–MBOM alignment reduces engineering rework and improves manufacturing synchronisation (Eigner and Stelzer, 2009).

#### 3.4.3 BIM and 3D/4D/5D Models

In installations involving large physical systems (e.g., warehouse automation, chemical plants, production lines), BIM and spatial models integrate geometric, temporal, and cost dimensions (Eastman et al., 2011).

#### 3.4.4 Semantic and Ontology Layers

To maintain digital continuity, DVCO incorporates ontologies that define relationships among requirements, models, routings, equipment, and field operations. Ontology-supported traceability is critical for change impact analysis (Bock and Gruninger, 2005).

MBE thus forms the “digital backbone” that ensures consistency from design to execution.

### 3.5 Software-Defined Execution Layer

Execution in modern ETO/D2O systems must be responsive, reconfigurable, and driven by engineering models. Traditional ERP/MES architectures, designed for stable environments, cannot support the dynamic orchestration required for ETO operations (Kang et al., 2016).

#### 3.5.1 Model-Linked Workflows and Orchestration Engines

The software-defined execution layer translates engineering models into operational workflows. Microservices architectures and event-driven engines support real-time coordination of manufacturing, installation, logistics, commissioning, and supplier operations.

#### 3.5.2 Integration Across ERP, PLM, MES, and WMS Systems

ETO environments require deep integration across system layers. Research shows that integration gaps between engineering and manufacturing systems are a major source of errors

(Skoogh et al., 2017). DVCO incorporates APIs and semantic connectors for SAP, PLM, MES, and IoT platforms.

### 3.5.3 Dynamic Resequencing and Constraint Management

ETO tasks often require resequencing based on equipment availability, installation conditions, or engineering changes. Constraint-based scheduling models and hybrid AI/OR optimisation support these adjustments.

### 3.5.4 Exception and Variance Handling

Software-defined execution enables automated detection and resolution of exceptions—an essential capability due to the volatility of ETO environments.

## 3.6 Real-Time Feedback Loop: IoT, Data Infrastructure, and Digital Twins

A core premise of DVCO is that value chains must operate as cyber-physical systems with continuous sensing, simulation, and feedback (Tao et al., 2019).

### 3.6.1 IoT and Operational Telemetry

Industrial IoT systems capture sensor data, equipment states, AMR telemetry, production line behaviour, energy consumption, and environmental data.

### 3.6.2 Digital Twins for Assets and Processes

Digital twins synchronise operational data with engineering and planning models. Studies show that digital twins improve commissioning accuracy, predictive maintenance, and lifecycle optimisation (Fuller et al., 2020).

### 3.6.3 Closed-Loop Updates

DVCO requires updates to engineering models, BOMs, routing structures, and schedules based on field conditions. This eliminates traditional "static planning" limitations and creates a dynamic adaptive planning environment.

### 3.6.4 Data Infrastructure

Streaming platforms (Kafka), lakehouses, historians, and semantic data stores enable real-time ingestion and analysis.

This layer operationalises the feedback mechanisms needed for self-optimising systems.

## 3.7 Decision Intelligence Layer

The decision intelligence layer provides predictive, prescriptive, and autonomous decision-making capabilities across the value chain.

### 3.7.1 Predictive Analytics

These models forecast demand surges, equipment failures, supply disruptions, installation delays, and risk patterns (Knoll, 2021).

### 3.7.2 Prescriptive Optimisation

Hybrid OR/AI methods recommend optimal sequencing, sourcing, routing, or installation strategies based on constraints and objectives (Kamble et al., 2020).

### 3.7.3 Knowledge Graph Reasoning

Knowledge graphs encode relationships among models, parts, workflows, assets, and events. They enhance semantic consistency across processes (García-Sánchez and Valencia-García, 2015).

### 3.7.4 AI-Enabled Control Towers

AI-supported DVCO control towers unify engineering, planning, manufacturing, logistics, installation, and service operations into a single intelligence layer.

The decision intelligence layer thus forms the “brain” of the DVCO architecture.

## 3.8 Integration Across the D2O Lifecycle

The DVCO architecture spans the entire D2O lifecycle:

Engage and Design → Plan and Source → Build and Test → Deliver and Install → Operate and Service

Unlike traditional lifecycle models, DVCO treats these stages as interconnected feedback loops. Each stage generates new information—engineering changes, installation conditions, telemetry, service reports—that must update upstream models. Literature on smart connected products emphasises that lifecycle integration is essential for digital transformation (Porter and Heppelmann, 2015).

DVCO therefore unifies engineering and operational functions into a dynamic, adaptive, model-driven ecosystem.

### 3.9 Summary

The chapter presented a comprehensive conceptual framework integrating systems engineering, Industry 4.0 technologies, AI-driven decision intelligence, and cyber-physical feedback mechanisms. The DVCO architecture provides a theoretically grounded foundation for addressing ETO/D2O complexity and serves as the blueprint for the research methodology described in Chapter 4.

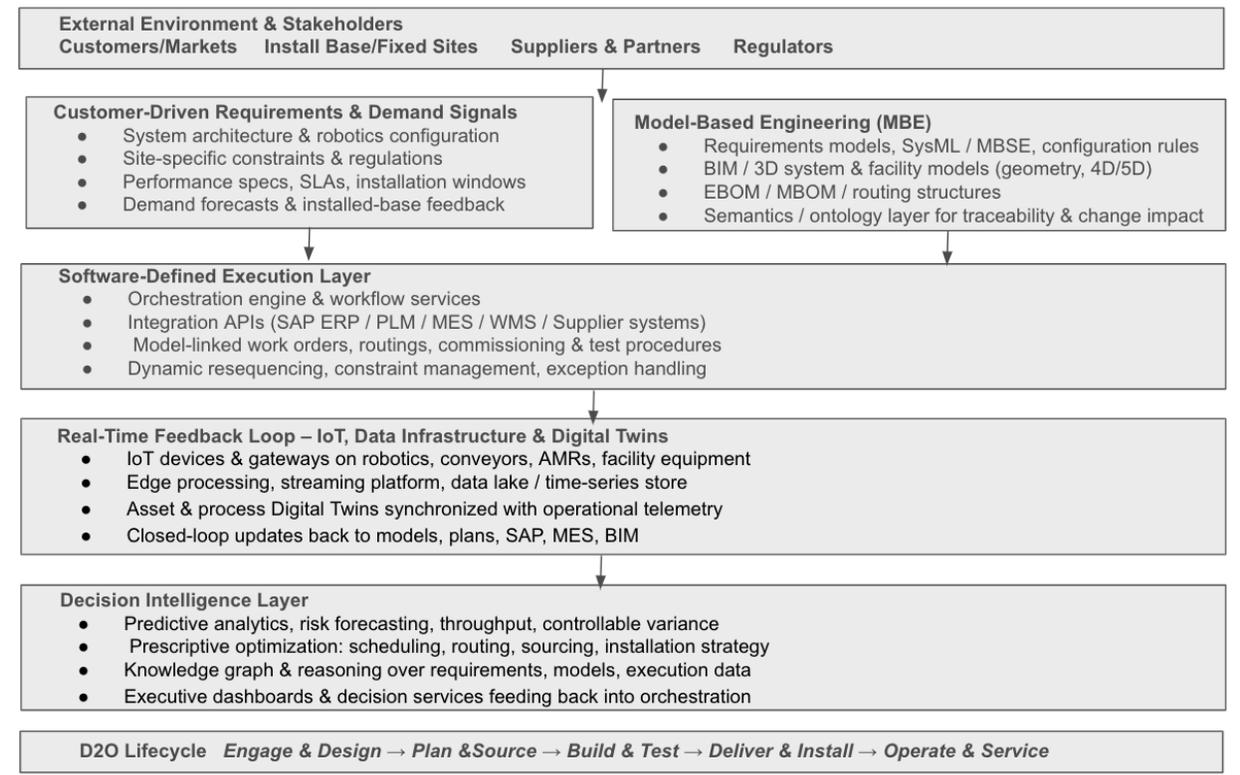


Figure 3.3 DVCO Architecture - ETO/E2O Value Chain

# CHAPTER IV

## TECHNOLOGIES AND ARCHITECTURE ENABLING DVCO

### 4.1 Introduction

The Dynamic Value Chain Optimisation (DVCO) framework requires a technologically coherent ecosystem that supports real-time coordination across Symbotic’s Engineer-to-Order (ETO) and Design-to-Operate (D2O) lifecycle. Modern ETO organisations increasingly rely on the convergence of Model-Based Engineering (MBE), Building Information Modelling (BIM), enterprise platforms such as SAP S/4HANA and Manufacturing Execution Systems (MES), cyber-physical systems (CPS), Internet of Things (IoT) telemetry, semantic technologies (ontologies and knowledge graphs), and artificial intelligence—particularly large language models (LLMs). These technologies together enable seamless integration between engineering intent, manufacturing execution, and operational intelligence (Grieves, 2014; Hedberg et al., 2019; Lee, Bagheri and Kao, 2015; Sheth, 2020).

Although these technologies have individually matured, their strategic value arises only when they operate as part of a model-driven, semantically aligned, and continuously synchronised architecture. For ETO environments characterised by volatile requirements, long lead-time

components, iterative engineering cycles, and complex commissioning processes, this convergence allows organisations to maintain fidelity between design models, production realities, and field performance (Culot et al., 2020; Madni and Sievers, 2018).

The *SAP + BIM* technical reference in Table 1 further demonstrates how BIM information models, SAP ERP, and lifecycle processes interact to support dynamic ETO workflows, emphasising the role of digital and process threads across engineering, supply chain, production, and commissioning contexts

Table 4. 1. Revit(BIM) and SAP Mapping table

Revit (BIM) Object / Parameter	Description	SAP ERP Object	SAP Module
Family	Generic object definition	Material / Equipment Category / Product Master	MM / PM
Family Type	Specific variation of a family	Material Variant / Material Type / Equipment Class	MM / PM
Element ID	Unique internal ID in Revit	Equipment Number / Functional Location / Material Number	PM / MM
Type Mark	Identifier for types in schedules	Material Number	MM
Mark	Instance-specific ID	Equipment Number	PM
Revit Category	Category (Doors, Walls, MEP...)	Asset Class / Material Group	PM / MM
Level	Building level/location	Functional Location (Hierarchy)	PM

Assembly Code (OmniClass/UniClass)	Classification system	Material Group / Asset Class	MM / PM
Workset / Phase	Project phasing	Project System WBS Element	PS
Volume / Area / Dimensions	Geometry data	Characteristics in SAP Classification System	PM / MM
Manufacturer	Provided manufacturer	Vendor Master (Source of Supply)	MM
Model Number	Manufacturer model	Manufacturer Part Number (MPN)	MM
Cost Parameter (Cost, Cost Type)	Revit-level cost data	Material Cost / Cost Element / BOM Costing	CO / MM
System (MEP)	Mechanical/Electrical system	Functional Location Structure	PM
Material Takeoff Schedule	Quantities of elements	BOM (Bill of Materials)	MM
Revit Room/Space	Architectural space definitions	Functional Location / Room Master	PM / RE-FX
Asset Parameters (Serial No., Service Life)	Asset metadata	Equipment Master / Asset Master Record (ANLA)	PM / FI-AA
Revit Parameters (Shared Parameters)	Custom attributes	SAP Classification Characteristics / Custom Fields	PM / MM
Revit Project Base Point / Survey Point	Coordinates	Plant/Location Coordinates	PM
Revit Project Name	BIM project metadata	SAP Project (Project Definition)	PS

Accordingly, this chapter examines the role of each enabling technology layer, analysing how Model-Based Engineering, BIM, enterprise platforms, IoT-enabled cyber-physical systems, semantic technologies, and AI collectively contribute to a unified DVCO environment. It further

identifies the integration challenges inherent to digital transformation in Engineer-to-Order (ETO) settings, including fragmented data structures, asynchronous engineering-operations workflows, inconsistent change-management practices, and the absence of semantically connected lifecycle models. Finally, the chapter presents the multi-layer DVCO systems architecture that enables real-time, cross-domain optimisation across the D2O value chain, providing a model-driven and AI-supported foundation for synchronising engineering intent with operational execution.

#### 4.2 Building Information Modelling (BIM) as the Structural and Spatial Backbone

Building Information Modelling (BIM) serves as the structural, spatial and contextual backbone for Symbolic's Engineer-to-Order (ETO) deployments. Its parametric, geometry-rich models capture building layouts, structural elements, mechanical clearances and installation constraints, forming the authoritative engineering representation required for downstream design, planning and execution (Kreider and Messner, 2013; Beetz, van Leeuwen and de Vries, 2009). Within the DVCO framework, BIM provides the baseline upon which engineering intent is propagated across SAP S/4HANA, MES, IoT/CPS and advanced integration platforms, enabling continuous alignment between design, execution and operational reality.

## 4.2 1 BIM–SAP Integration and the Digital Thread

A foundational requirement for DVCO is the seamless integration of BIM with enterprise systems—most critically SAP S/4HANA, SAP PS, PP, EWM and MES. BIM supplies the structural and spatial context for automation deployments, while SAP operationalises this information into procurement, production, logistics and cost flows (Lin, Lee and Ma, 2020; Weyer et al., 2015). Their convergence establishes the upstream portion of the Digital Thread, ensuring bidirectional synchronisation between engineering updates and execution processes.

### 4.2.1.1 BIM as the Engineering Source of Truth for SAP

BIM models contain authoritative information regarding building geometry, spatial layouts, structural load paths and installation constraints. When integrated with SAP:

- geometric data supports automated quantity take-off;
- spatial constraints inform scheduling, sequencing and staging logic;
- installation geometry drives site-readiness and logistics planning;
- BIM element IDs map to SAP material masters, network activities and work orders.

This alignment enables engineering changes—such as equipment repositioning or layout optimisation—to propagate directly into SAP-controlled procurement and installation workflows, improving consistency and reducing execution risk.

### 4.2.1.2 Synchronisation Across Engineering, Procurement and Installation

ETO organisations frequently experience schedule fragmentation and procurement inaccuracies when engineering tools and ERP systems remain disconnected (Culot et al., 2020). BIM–SAP integration mitigates these issues through:

- automated BOM generation from BIM/EBOM structures;
- real-time linkage between BIM objects and SAP project activities;
- synchronised procurement based on engineered quantities and installation phasing;
- digital site-readiness assessments using BIM construction zones and SAP project data.

These capabilities shorten decision cycles and increase predictability across the D2O continuum.

#### 4.2.1.3 Establishing the Digital Thread

The Digital Thread represents model-driven continuity across engineering, planning, manufacturing, logistics, installation and operations (Grieves, 2014; Tao et al., 2018). BIM–SAP integration provides this continuity by ensuring:

- downstream propagation of engineering updates into procurement and production;
- synchronisation of field progress and telemetry back into BIM/SAP systems;
- incorporation of operational insights into engineering templates;
- end-to-end preservation of data integrity and traceability.

This bidirectional flow transforms traditionally siloed engineering and enterprise activities into a unified execution ecosystem.

#### 4.2.1.4 Evidence from Industry Practice

The uploaded *SAP + BIM Overview* demonstrates how 3D layouts, installation constraints, material quantities and sequencing logic can be integrated with SAP modules to enable coordinated procurement, project controls and installation planning. This evidence confirms the operational viability of a BIM–SAP Digital Thread in high-complexity ETO environments.

#### 4.2.2 Extending BIM Beyond SAP: MES, IoT/CPS and the DVCO Semantic Layer

While BIM–SAP integration establishes the engineering–enterprise connection, full lifecycle optimisation in ETO settings requires synchronisation with manufacturing (MES), logistics and commissioning processes, as well as operational telemetry from IoT/CPS systems. DVCO therefore extends BIM’s role beyond SAP into a multi-system environment connected through a semantic layer.

##### 4.2.2.1 The Cross-Domain Interoperability Challenge

A major barrier to achieving lifecycle-wide optimisation in Engineer-to-Order (ETO) environments is the lack of semantic and structural interoperability among the systems that govern engineering, enterprise planning, manufacturing execution and operational telemetry. BIM, ERP, MES and IoT/Cyber-Physical Systems (CPS) each employ fundamentally different representational logics, resulting in mismatches that impede coordination, traceability and decision-making across the D2O value chain.

These representational differences can be summarised as follows:

- BIM → geometry, parameters and spatial topology;
- ERP (SAP) → materials, transactional records, costs and scheduling structures;
- MES → routings, work centres, resource allocations and quality checkpoints;
- IoT/CPS → equipment states, telemetry, event streams and anomalies.

Because each system conceptualises the physical and operational environment through a different schema, misalignments frequently occur. Examples include:

- BIM-defined conveyor segments that do not map directly to SAP material masters or equipment objects;
- BIM installation zones that fail to correspond to MES-designated work centres or operational areas;
- IoT or CPS anomalies that cannot be traced to specific BIM or SAP equipment identifiers, preventing localisation and root-cause analysis.

These ontological inconsistencies disrupt lifecycle traceability, fracture process continuity and impede the creation of a unified Digital Thread—an essential requirement for DVCO-enabled dynamic synchronisation (Culot et al., 2020; Hedberg and Feeney, 2017). In the absence of semantic alignment, data flows remain siloed, engineering changes propagate inconsistently into execution environments and operational deviations cannot be contextualised within the broader

spatial or engineering model. As a result, cross-domain optimization is severely constrained, reinforcing the need for a semantic integration layer within the DVCO architecture.

#### 4.2.2.2. DVCO's Semantic Layer as the Integration Mechanism

DVCO resolves these mismatches via a semantic layer—an ontological and knowledge-graph framework harmonising BIM entities with ERP, MES and IoT/CPS systems (Bock and Gruninger, 2005; Sheth, 2020). Semantic mappings include:

- linking BIM elements (e.g., racks, conveyors, robotics aisles) to SAP materials, equipment IDs and MES instructions;
- aligning BIM spatial zones with SAP WBS structures and MES work centres;
- correlating BIM 4D schedules with SAP PS milestones and commissioning tasks;
- situating IoT anomalies (temperature spikes, vibration patterns) within BIM spatial coordinates.

These mappings transform BIM from a design model into a cross-domain integration hub anchoring the DVCO architecture.

#### 4.2.2.3. Enabling a Fully Connected Digital Thread

Once semantically unified, BIM, ERP, MES and IoT/CPS systems form a continuous lifecycle information fabric that enables the full expression of Dynamic Value Chain Optimisation (DVCO). Through semantic alignment, previously isolated engineering, enterprise and

operational datasets become mutually intelligible, enabling higher-order orchestration and decision intelligence across the entire D2O value chain.

This unified environment supports several capabilities central to DVCO:

- model-driven orchestration of manufacturing, logistics and installation processes, guided by BIM geometry and engineering intent;
- real-time propagation of design and execution changes, ensuring that updates in one domain automatically synchronise with connected systems;
- spatially contextualised anomaly detection and installation verification, enabled by linking IoT/CPS telemetry with BIM spatial coordinates and enterprise identifiers;
- end-to-end lifecycle traceability across engineering, production, logistics, commissioning and operations, enabling root-cause analysis and continuous improvement.

The SAP BIM integration provides the initial integration layer by demonstrating how BIM-derived geometry, quantities, and installation constraints are connected to SAP project and procurement structures. The DVCO semantic layer extends this foundation by incorporating MES entities and IoT/CPS telemetry into the same model-driven ecosystem, thereby completing the fully connected Digital Thread required for dynamic value chain optimisation. This lifecycle continuity transforms the ETO execution environment from a sequence of loosely coupled processes into a semantically coherent, self-updating, and AI-ready value chain.

## 4.3 Ontologies and Knowledge Graphs: The Semantic Backbone of DVCO

### 4.3.1 The Role of Ontologies in Lifecycle Semantic Alignment

Ontologies provide the formal semantic foundation required to harmonise engineering, manufacturing, logistics, installation and operational data across the Design-to-Operate (D2O) lifecycle. Unlike traditional integration approaches—typically reliant on file-based exchanges, schema mapping or point-to-point interfaces—ontology-driven integration enables a structured, machine-interpretable representation of concepts, relationships and constraints (Gruber, 1993; Uschold and Gruninger, 1996). This ensures that engineering intent is expressed consistently across systems and mitigates the semantic fragmentation that frequently emerges in Engineer-to-Order (ETO) environments (Bock and Gruninger, 2005; Sheth, 2020).

Within the DVCO framework, the ontology unifies multiple lifecycle domains, including:

- Engineering: components, assemblies, interfaces, tolerances (Hedberg et al., 2016);
- Manufacturing: work centres, routings, operations and quality checkpoints (Lin, Lee and Ma, 2020);
- Supply chain: suppliers, materials, lead times and logistic dependencies (Culot et al., 2020);
- Installation: tasks, spatial zones, sequences and predecessor–successor constraints (Madni and Sievers, 2018);

- Commissioning: test procedures, readiness states and failure modes (Boschert and Rosen, 2016);
- IoT/Operations: sensor types, events, alarm states and performance telemetry (Tao et al., 2018; Da Xu, He and Li, 2014).

By linking these domains within a coherent semantic schema, the ontology becomes a single reference model that enables DVCO to interpret cross-system data with clarity, contextual awareness and shared meaning. As a result, engineering models, ERP transactions, MES operations and IoT telemetry can be integrated semantically—rather than through brittle interface logic—supporting advanced DVCO capabilities such as model-driven orchestration, AI-assisted anomaly detection, predictive commissioning and dynamic optimisation across Symbolic’s ETO value chain.

#### 4.3.2 Knowledge Graphs as Dynamic, Lifecycle-Connected Integration Structures

Knowledge graphs (KGs) operationalize the ontology by representing real-world instances—BIM elements, SAP materials, MES work orders, IoT events, supplier records and installation activities—as nodes within a connected, semantically enriched graph. Unlike relational databases, which require rigid schemas and predefined joins, KGs flexibly encode multi-hop relationships, hierarchical dependencies and lifecycle evolution across systems (Hogan et al., 2021; Ehrlinger and Wöß, 2016). This flexibility makes KGs particularly well suited for

complex Engineer-to-Order (ETO) environments, where structural, operational and temporal relationships continually change throughout the D2O lifecycle.

Within the DVCO architecture, the knowledge graph enables several critical capabilities:

- cross-domain dependency modelling, such as linking a BIM conveyor segment to its SAP material master, MES routing, supplier source, installation task and commissioning signal (Sheth, 2020);
- bidirectional traceability, allowing users and automated agents to traverse from an IoT anomaly to its engineering representation, procurement source and installation workflow;
- impact propagation, automatically identifying the components, tasks, resources and suppliers affected by engineering changes or operational deviations (Bock and Gruninger, 2005);
- semantic querying, enabling structured and natural-language questions to traverse underlying graph relationships for operational and engineering insight (Hogan et al., 2021);
- explainable AI reasoning, grounding LLM-generated recommendations in transparent graph paths and domain semantics.

By serving as a dynamic, lifecycle-connected knowledge substrate, the KG becomes the core computational fabric of DVCO. It enables real-time reasoning, cross-domain alignment, impact analysis and semantic interoperability across engineering, manufacturing, logistics,

commissioning and field operations—capabilities essential for optimising high-variability ETO value chains.

#### 4.3.3 Why Ontologies and KGs Are Essential in ETO Environments

Engineer-to-Order (ETO) environments exhibit high variability, customised configurations and frequent engineering changes—conditions that conventional database architectures struggle to support. Traditional ERP, MES and PLM systems rely on rigid schemas, predefined joins and static object models, which cannot readily accommodate the dynamic relationships, multi-domain dependencies and lifecycle evolution inherent in ETO delivery (Hedberg and Feeney, 2017; Madni and Sievers, 2018). Without a semantic integration layer, the following issues commonly emerge:

- divergence of meaning between engineering representations and ERP/MES transactional structures;
- loss of consistency in BOM and EBOM–MBOM mappings during iterative engineering changes;
- conflicts between installation sequences and spatial realities captured within BIM;
- IoT and CPS telemetry lacking engineering or operational context, preventing accurate anomaly localisation;

- over-reliance on tacit human knowledge, resulting in undocumented, non-repeatable coordination practices.

Ontologies and knowledge graphs directly address these challenges by providing a shared semantic foundation and lifecycle-connected relationship model that spans engineering, supply chain, manufacturing, installation and operations. Ontologies establish common conceptual definitions, constraints and cross-domain semantics (Gruber, 1993; Uschold and Gruninger, 1996), while knowledge graphs instantiate these semantics as traceable, dynamically evolving links among real-world entities (Hogan et al., 2021; Sheth, 2020).

Together, they enable:

- semantic coherence across tools, teams and lifecycle stages;
- traceable, bidirectional relationships supporting change propagation and dependency analysis;
- contextualised interpretation of IoT events through spatial, engineering and transactional grounding;
- adaptive, model-driven integration that evolves with engineering updates;
- closed-loop orchestration, in which DVCO aligns engineering intent with real-time execution and operational feedback.

Through this semantic infrastructure, DVCO transcends traditional integration approaches and functions as an intelligent, self-updating coordination system capable of supporting the complexity, variability and dynamism characteristic of ETO value chains.

#### 4.4 IoT, CPS and Telemetry Layer

The Internet of Things (IoT), cyber–physical systems (CPS) and telemetry infrastructure constitute the sensing and validation layer of the DVCO architecture. This layer provides continuous situational awareness across Symbotic’s Engineer-to-Order (ETO) and Design-to-Operate (D2O) environments, enabling the detection of deviations, the synchronisation of Digital Twins and the provision of real-time feedback to engineering, planning, manufacturing, installation and commissioning functions.

In ETO deployments—where physical constraints, dynamic site conditions and tightly coupled automation subsystems interact in complex and often unpredictable ways—the ability to sense, contextualise and respond to emerging operational signals is essential for lifecycle optimisation (Da Xu, He and Li, 2014; Kusiak, 2018). IoT sensors embedded in conveyors, robots, racking systems and environmental infrastructure produce continuous telemetry streams such as vibration signatures, thermal patterns, alignment data, equipment utilisation and safety-status signals. CPS architectures integrate these sensor flows with actuator behaviour, enabling closed-loop control, autonomous decision responses and execution-level adjustments.

Within the DVCO framework, the IoT/CPS layer performs several critical functions:

- Real-time data acquisition
 

High-frequency telemetry enables continuous monitoring of equipment performance, installation conditions and environmental factors.
- Spatial and semantic contextualisation
 

Telemetry signals become meaningful only when associated with BIM spatial coordinates, SAP equipment identifiers and MES routing contexts—achieved via the DVCO semantic layer.
- Digital Twin synchronisation
 

IoT and CPS data drive the synchronisation of behavioural and operational twins, allowing virtual models to reflect real-world conditions with high fidelity (Tao et al., 2018).
- Deviation and anomaly detection
 

Combined sensor patterns detect misalignments, emerging failures, installation deviations and commissioning readiness gaps before they escalate into systemic issues.
- Closed-loop feedback into upstream lifecycle stages
 

Operational insights feed back into engineering, planning and execution systems, enabling iterative improvement of system templates, layouts and workflows.

By functioning as the nervous system of the DVCO architecture, the IoT/CPS layer extends the digital thread into real-world operational environments. Its telemetry provides the evidence base that allows DVCO to confirm engineering assumptions, validate installation outcomes and dynamically adjust lifecycle plans—thereby reducing uncertainty, improving commissioning

predictability and enabling higher levels of autonomous, model-driven optimisation across Symbotic's ETO/D2O value chain.

#### 4.4.1 IoT Telemetry as the Foundation of Real-Time Visibility

IoT devices deployed across conveyors, autonomous mobile robots (AMRs), robotic arms, automated storage and retrieval systems (ASRS), environmental sensors and commissioning equipment generate the continuous telemetry required to understand operational conditions within Symbotic's Engineer-to-Order (ETO) and Design-to-Operate (D2O) environments. These sensor-rich systems function as real-time observation points, producing high-frequency data essential for monitoring equipment behaviour, validating installation activities and maintaining alignment between physical operations and their digital representations.

Typical IoT telemetry streams include:

- equipment performance indicators, such as torque signatures, vibration profiles, cycle times, thermal loads and mechanical stress responses;
- environmental variables, including temperature, humidity, airflow, particulate density and lighting conditions that influence robot and material-handling performance;
- state and availability signals, encompassing motion states, utilisation patterns, fault codes, battery levels and autonomous navigation statuses;

- spatially contextualised events, where sensor triggers are mapped back to BIM-defined zones, enabling location-aware situational analysis;
- installation and commissioning progress indicators, including test-script results, alignment verification outputs and configuration readiness signals.

IoT telemetry provides the foundational data layer required for anomaly detection, predictive maintenance, installation validation and real-time operational synchronisation (Zheng et al., 2019; Tao et al., 2018). When integrated with BIM spatial models, SAP equipment identifiers and MES workflow states, telemetry transforms from raw sensor readings into semantically contextualised operational intelligence. This contextualisation allows the DVCO architecture to detect emerging deviations, synchronise Digital Twins, support AI-assisted diagnostics and orchestrate closed-loop corrective actions with minimal human intervention.

#### 4.4.2 Digital Twin Synchronisation and Behavioural Alignment

Digital Twin synchronisation is a core mechanism through which DVCO aligns virtual engineering models with real-world operational behaviour. In ETO environments—where installation conditions evolve rapidly and automation subsystems interact in highly dynamic ways—the ability to continuously update Digital Twins based on IoT and CPS telemetry ensures that engineering representations remain accurate, actionable and lifecycle-relevant (Tao et al., 2018; Boschert and Rosen, 2016).

Within the DVCO architecture, Digital Twin synchronisation operates in three interconnected layers:

#### 4.4.2.1 Geometric and Spatial Synchronisation (BIM–IoT Integration)

IoT and CPS devices supply real-time positional, alignment and environmental data that update BIM-based geometric models.

Examples include:

- adjustments to conveyor alignment detected by vibration or angle sensors;
- AMR navigation patterns mapped back onto BIM-defined floor zones;
- installation progress reflected through zone-based commissioning updates.

This enables the BIM model to function not merely as a design artefact but as a live spatial Digital Twin, reflecting current installation and site conditions with high fidelity.

#### 4.4.2.2. Operational Synchronisation (MES–IoT–SAP Integration)

Behavioural telemetry—cycle times, load variations, fault codes, utilisation rates—feeds into MES and SAP execution structures, ensuring that operational Digital Twins match real-world performance.

This supports:

- verification of routing logic and task sequences;
- validation of resource availability and workstation readiness;

- synchronisation of SAP production orders with actual subsystem behaviour.

Operational alignment closes the loop between engineering assumptions, planned workflows and actual execution performance.

#### 4.4.2.3. Behavioural and Predictive Synchronisation (AI-Enabled Digital Twins)

High-volume IoT and CPS signals are analysed through AI/ML models to produce predictive insights, including:

- impending equipment failures;
- suboptimal installation patterns;
- environmental risks to automation performance;
- emerging commissioning deviations.

These insights update behavioural Digital Twins, enabling DVCO to anticipate disruptions and optimize lifecycle decisions before deviations escalate.

#### 4.4.2.4 Role of DVCO in Synchronising the Digital Twin Ecosystem

DVCO orchestrates these synchronisation processes across BIM, SAP, MES and IoT/CPS systems by:

- providing the semantic layer that unifies telemetry with engineering and enterprise schemas;

- enabling real-time updates to Digital Twins through shared ontologies and knowledge graphs;
- supporting AI-driven behavioural reasoning grounded in traceable lifecycle data;
- ensuring bidirectional propagation of updates across engineering, planning, manufacturing and field operations.

Through this holistic synchronisation, Digital Twins evolve from static models into adaptive, intelligence-driven representations that mirror, interpret and optimise the real-world ETO environment.

DVCO thus enables a Digital Twin ecosystem that is continuously updated, semantically coherent and operationally aligned—one capable of supporting predictive commissioning, autonomous orchestration and high-precision lifecycle optimisation across Symbotic’s automation deployments.

#### 4.4.3 Telemetry-Driven Anomaly Detection and Commissioning Assurance

Telemetry-driven anomaly detection is a central mechanism through which DVCO identifies deviations, validates installation accuracy and ensures commissioning readiness across Symbotic’s automation deployments. Unlike traditional commissioning processes—which rely heavily on manual inspections, checklists and disconnected test data—DVCO integrates IoT and CPS telemetry into a continuous monitoring framework that evaluates operational behaviour in real time (Zheng et al., 2019; Da Xu, He and Li, 2014).

IoT- and CPS-enabled subsystems generate high-frequency signals that reveal emerging patterns, structural inconsistencies and abnormal conditions. These signals serve as early indicators of misalignment, degradation or improper installation, allowing DVCO to detect and address issues long before they escalate into operational disruptions.

#### 4.4.3.1. Detection of Equipment and Automation Anomalies

Telemetry patterns provide a detailed behavioural fingerprint of the automation environment.

DVCO uses these patterns to detect anomalies such as:

- mechanical deviations (e.g., excessive vibration in conveyors, torque inconsistencies in robotic joints);
- thermal or electrical irregularities (e.g., overheating motors, voltage drops, abnormal current loads);
- task-cycle deviations, including unexpected increases in AMR travel time or ASRS retrieval latency;
- sensor or signal degradation, such as inconsistent environmental readings or intermittent fault codes.

Machine-learning and rule-based models process these streams to distinguish routine variability from actionable anomalies, enabling targeted interventions.

#### 4.4.3.2. Spatially Contextualised Anomaly Interpretation

Telemetry alone lacks meaning without context.

DVCO enriches anomaly detection through semantic and spatial contextualisation:

- IoT anomalies are mapped to BIM spatial zones, providing precise localisation (e.g., “misalignment detected in Rack Zone A-14”).
- Each anomaly is associated with SAP equipment IDs, MES routing contexts and Digital Twin behavioural models.
- Spatial correlations help detect systemic issues—for example, alignment deviations across multiple adjacent conveyor zones.

This contextual grounding transforms raw sensor data into situational insight that is traceable, explainable and operationally meaningful.

#### 4.4.3.3. Commissioning Validation and Readiness Assurance

Commissioning in ETO environments demands proof that subsystems have been installed and configured according to engineering specifications.

The DVCO telemetry layer supports commissioning through:

- automated verification of installation alignment, using vibration, angle and position sensors;
- execution of commissioning test scripts, with IoT outputs confirming pass/fail status;
- continuous equipment-health monitoring, ensuring readiness for formal Site Acceptance Tests (SAT);

- closed-loop feedback into engineering models, updating Digital Twins and SAP/MES structures with commissioning outcomes.

Telemetry-driven commissioning reduces dependence on manual inspection, accelerating FAT/SAT cycles and improving commissioning accuracy.

#### 4.4.3.4. Early-Warning and Predictive Assurance

DVCO extends anomaly detection to predictive assurance through:

- time-series forecasting of degradation indicators,
- predictive failure modelling (e.g., vibration prognostics),
- correlation of environmental conditions with performance risks,
- identification of installation deviations likely to cause downstream failures.

This predictive dimension ensures that commissioning is not simply a pass–fail validation activity, but a forward-looking risk-mitigation process informed by behavioural evidence.

#### 4.4.3.5. Lifecycle Impact and Feedback Integration

Telemetry-driven anomaly detection supports the entire ETO lifecycle by:

- feeding deviations into engineering change workflows,
- updating Digital Twins with validated behavioural insight,

- enriching the knowledge graph with traceable anomaly–context relationships,
- informing SAP procurement, maintenance planning and installation resourcing,
- supporting continuous improvement of templates for future deployments.

Through these capabilities, telemetry evolves from a raw sensor stream into the evidence base that powers DVCO’s lifecycle optimisation engine.

#### 4.4.4 Closed-Loop Feedback into Engineering, Planning and Execution Systems

Closed-loop feedback is the mechanism through which DVCO transforms raw telemetry and operational insights into actionable lifecycle adjustments across engineering, planning, manufacturing, installation and commissioning systems. In traditional ETO environments, feedback loops are fragmented, delayed or dependent on human interpretation, leading to misalignment between engineering assumptions, production realities and field conditions. DVCO resolves these limitations by using IoT/CPS telemetry, semantic integration and Digital Twin synchronisation to create a continuous, automated feedback cycle that spans the entire D2O value chain (Kusiak, 2018; Tao et al., 2018).

##### 4.4.4.1. Engineering Feedback: Updating Models, Constraints and Designs

Operational deviations detected through telemetry feed directly back into engineering systems.

This enables:

- updating BIM models to reflect real-world alignment, clearances and installation conditions;
- adjusting engineering constraints, such as load assumptions or spatial tolerances;
- triggering engineering change workflows (ECMs) in SAP or PLM systems;
- refining Digital Twins with validated physical and behavioural characteristics.

This ensures that engineering intent remains grounded in empirical reality rather than assumptions—particularly critical in high-variability ETO deployments.

#### 4.4.4.2. Planning Feedback: Adjusting Schedules, Milestones and Dependencies

Telemetry-driven insights also influence project and operations planning:

- SAP PS milestones are updated when installation progress deviates from expected timelines;
- critical-path schedules adjust based on real-time MES throughput and resource availability;
- logistics planning adapts to readiness signals detected at specific BIM zones;
- installation sequencing is optimised based on evidence of bottlenecks or downstream risks.

These planning adjustments transform project management from static scheduling into dynamic, data-driven orchestration.

#### 4.4.4.3. Execution Feedback: Orchestrating Production, Installation and Commissioning

Feedback loops extend into execution systems such as SAP PP, SAP EWM and MES:

- SAP production orders adjust based on real-time machine utilisation and anomaly patterns;
- MES routing rules update when telemetry indicates suboptimal task cycles or workstation readiness;
- installation tasks reorder automatically when Digital Twin data reveals dependencies or conflicts;
- commissioning processes accelerate or pause based on validated readiness signals.

Execution becomes adaptive and responsive, reducing rework and preventing cascading delays.

#### 4.4.4.4. Knowledge Graph Integration: Persisting Traceability and Context

All feedback events—engineering changes, planning shifts, anomaly detections, commissioning results—are written into the DVCO knowledge graph:

- linking anomalies to their engineering origins;
- connecting commissioning failures to earlier installation or supplier dependencies;
- capturing behavioural trends associated with specific BIM zones or SAP equipment IDs.

This persistent semantic record enables explainability, predictive reasoning and lifecycle learning, allowing AI models and LLM agents to make more accurate recommendations over time.

#### 4.4.4.5. System-Level Outcome: A Self-Updating, Orchestrated D2O Environment

Through closed-loop feedback, DVCO transforms the ETO lifecycle into a continuously self-correcting system:

- engineering intent and operational execution remain synchronised;
- deviations are detected and addressed earlier;
- Digital Twins remain behaviourally accurate;
- planning becomes event-driven;
- execution becomes autonomously optimised.

In this architecture, IoT/CPS telemetry does not merely monitor operations—it actively drives engineering, planning and execution intelligence, forming the foundation for predictive, adaptive and ultimately autonomous value-chain orchestration.

#### 4.4.5 The Role of the Telemetry Layer in DVCO

The IoT/CPS Telemetry Layer enables DVCO to operate as a closed-loop optimisation system by:

- providing real-time confirmation of plan-to-physical alignment,

- detecting deviations early and triggering upstream updates,
- supplying continuous operational data for forecasting and optimisation,
- grounding AI/LLM reasoning in live system context,
- enabling lifecycle traceability through semantic integration.

In ETO environments, where variability and physical constraints are pervasive, the Telemetry Layer ensures that DVCO maintains synchronisation between the engineered design, enterprise execution and real-world operational conditions.

#### 4.5 Decision Intelligence Layer

The Decision Intelligence Layer represents the cognitive core of the DVCO architecture. While IoT/CPS telemetry provides real-time observations, Digital Twins maintain behavioural synchronisation and ontologies/knowledge graphs preserve lifecycle semantics, it is the Decision Intelligence Layer that transforms this information into actionable insights, predictive recommendations and cross-domain orchestration. This layer enables DVCO to operate not merely as an integration architecture, but as an intelligent, anticipatory and optimisation-driven system capable of supporting Symbolic's complex Engineer-to-Order (ETO) and Design-to-Operate (D2O) environments.

Decision intelligence integrates three complementary capabilities:

1. predictive analytics, driven by statistical models and machine learning;

2. semantic reasoning, enabled by ontologies and knowledge graphs;
3. large language model (LLM) agents, supporting interpretability, planning and cross-domain coordination.

Together, these capabilities empower DVCO to align engineering intent with operational reality, foresee deviations before they occur and orchestrate lifecycle decisions grounded in transparent, traceable and semantically coherent knowledge structures (Sheth, 2020; Kusiak, 2018; Hogan et al., 2021).

#### 4.5.1 Predictive Analytics and Machine-Learning Models

Machine-learning models trained on IoT, CPS and historical operational data identify behavioural patterns, emerging risks and optimisation opportunities. These models support:

- predictive maintenance, forecasting failures in conveyors, AMRs, robotic joints and ASRS subsystems;
- throughput prediction, using MES data to anticipate bottlenecks in production or installation sequences;
- environmental risk detection, linking temperature, humidity or airflow variations to performance degradation;
- commissioning-risk scoring, combining telemetry with Digital Twin simulations to predict readiness gaps.

Predictive capabilities allow DVCO to shift from reactive decision-making to proactive lifecycle optimisation (Kusiak, 2018; Tao et al., 2018).

#### 4.5.2 Semantic Reasoning Using Ontologies and Knowledge Graphs

Ontologies and knowledge graphs enable machine-interpretable relationships among engineering structures, operational tasks, supply-chain dependencies and telemetry signals (Uschold and Gruninger, 1996; Bock and Gruninger, 2005).

Within the DVCO architecture, semantic reasoning supports:

- impact propagation, tracing how engineering changes affect SAP materials, MES routings and installation tasks;
- root-cause analysis, linking anomalies to their physical, temporal and engineering origins;
- cross-domain query and retrieval, enabling planners or AI agents to answer questions grounded in lifecycle semantics;
- context-aware decision recommendations, using graph traversal and constraint reasoning.

This ensures that decisions are not only data-driven but meaning-aware, preserving interpretability and traceability across the D2O chain.

#### 4.5.3 LLM-Based Agents for Orchestration, Insight and Human–AI Collaboration

Large language models—reinforced with DVCO ontologies, knowledge graphs and operational datasets—enable natural-language reasoning, multi-step planning and coordination across engineering, supply-chain, manufacturing and installation domains (Bommasani et al., 2021; Sheth, 2020).

LLM agents within DVCO perform:

- natural-language querying across the knowledge graph and telemetry systems;
- explainable reasoning, grounding their outputs in graph-based evidence paths;
- multi-agent orchestration, coordinating engineering, procurement and installation workflows through SAP/MES integrations;
- simulation-driven decision support, generating “what-if” scenarios based on Digital Twin models;
- automated reporting, producing engineering summaries, commissioning reports or operational risk digests.

Unlike conventional analytics dashboards, LLM agents provide context-rich, cross-domain intelligence synthesised from the entire DVCO data ecosystem.

#### 4.5.4 Human-in-the-Loop and Organisational Decision Enablement

Despite the system’s advanced automation, DVCO maintains human oversight through human-in-the-loop (HITL) decision pathways:

- engineers validate design changes proposed by AI models;
- planners review schedule adjustments before SAP integration;
- installation teams assess anomaly alerts before execution;
- supervisors approve procurement changes triggered by predictive insights.

This hybrid architecture preserves accountability, transparency and organisational trust—critical for safety-conscious ETO deployments.

Outcome: A Cognitively Enabled DVCO Architecture

Through predictive analytics, semantic reasoning and LLM-based agents, the Decision Intelligence Layer enables DVCO to operate as:

- an anticipatory system, predicting and preventing disruptions;
- an interpretive system, understanding lifecycle context across domains;
- an optimisation system, orchestrating plans, resources and workflows;
- an explainable system, providing transparent and traceable justification;
- an adaptive system, continuously learning from telemetry and lifecycle evidence.

In Symbotic’s high-variability ETO environment, this layer transforms DVCO from a digital integration framework into an AI-enabled lifecycle optimisation engine, capable of supporting autonomous decision-making, accelerated commissioning and superior operational performance.

## 4.6 Multi-Layer DVCO Architecture

The Dynamic Value Chain Optimisation (DVCO) architecture integrates the technological, semantic and decision-making capabilities described in the preceding sections into a coherent, multi-layer system. Rather than treating BIM, SAP, MES, IoT/CPS, ontologies, knowledge graphs and AI components as isolated tools, DVCO orchestrates them as interconnected layers within a lifecycle-wide architecture. This enables Symbotic to manage its high-variability Engineer-to-Order (ETO) and Design-to-Operate (D2O) value chain as a semantically coherent, data-driven and cognitively enabled system.

The multi-layer DVCO architecture can be understood as comprising six primary layers, bound together by cross-cutting digital threads and feedback loops:

1. Physical and IoT/CPS Layer
2. Operational Execution Layer (MES and Local Controls)
3. Enterprise and Project Management Layer (SAP and Related Systems)
4. Model-Based Engineering and BIM Layer
5. Semantic Integration Layer (Ontologies and Knowledge Graphs)
6. Decision Intelligence and Human-in-the-Loop Layer

These layers are not strictly linear; rather, they form an interconnected stack in which data, semantics and decisions circulate continuously across lifecycle stages.

## 4.6.1 Layered Architectural View

### 4.6.1.1 Physical and IoT/CPS Layer

At the base of the architecture lies the physical environment: conveyors, AMRs, robotic arms, ASRS systems, racking structures, sensors, PLCs, environmental monitors and commissioning equipment. IoT and CPS infrastructures collect high-frequency telemetry from these assets—including performance indicators, environmental variables, state signals and commissioning data—and expose them as time-series streams for analysis and synchronisation.

### 4.6.1.2 Operational Execution Layer (MES and Local Controls)

Above the physical layer, Manufacturing Execution Systems (MES) and local control systems (e.g. PLC/SCADA) manage work orders, routings, work centres, quality checks and real-time execution logic. This layer translates high-level plans into executable operations, coordinating task sequencing, resource allocation and sub-system behaviour. Operational data from MES is tightly coupled with IoT/CPS signals, forming the behavioural foundation of Digital Twins.

### 4.6.1.3 Enterprise and Project Management Layer (SAP and Related Systems)

The enterprise layer is responsible for project planning, procurement, costing, logistics and financial control, primarily through SAP S/4HANA, SAP PS, PP, EWM and related modules. Here, BOMs, project networks, milestones, purchase orders and inventory movements are managed. BIM–SAP integration ensures that engineering quantities, installation phasing and spatial constraints are reflected in ERP structures, while MES feedback and telemetry-derived insights are propagated back into planning and project controls.

#### 4.6.1.4 Model-Based Engineering and BIM Layer

The BIM and model-based engineering layer provides structural and spatial context for Symbolic’s warehouse automation deployments. BIM captures building geometry, mezzanines, racking, conveyor paths, robot aisles, safety zones and installation clearances. It serves as the digital engineering backbone upon which SAP project structures, MES routings, installation tasks and Digital Twins are anchored. Changes in BIM models propagate through the architecture via EBOM/BOM structures and semantic mappings.

#### 4.6.1.5 Semantic Integration Layer (Ontologies and Knowledge Graphs)

The semantic layer harmonises meanings across BIM, SAP, MES and IoT/CPS systems. Ontologies define lifecycle concepts, relationships and constraints spanning engineering, manufacturing, supply chain, installation, commissioning and operations. Knowledge graphs instantiate these semantics as a connected graph of real-world entities—BIM elements, SAP materials, MES operations, IoT events, suppliers and commissioning results. This layer provides the “semantic spine” of DVCO, enabling cross-domain traceability, impact analysis and context-aware query.

#### 4.6.1.6 Decision Intelligence and Human-in-the-Loop Layer

At the top of the stack, the Decision Intelligence Layer combines predictive analytics, semantic reasoning and LLM-based agents to generate recommendations, coordinate workflows and support human decision-making. Predictive models anticipate failures and bottlenecks; semantic reasoning engines trace impacts and root causes; LLM agents provide natural-language

interfaces, plan multi-step actions and generate explainable narratives grounded in the knowledge graph. Human decision-makers validate and approve critical changes, forming a human-in-the-loop governance mechanism.

#### 4.6.2 Cross-Cutting Digital Thread and Feedback Loops

Threading vertically through all layers is the Digital Thread, which provides continuity of information and semantics from initial customer requirements through design, planning, manufacturing, installation, commissioning and operations. Engineering changes in BIM are reflected in SAP project structures and MES routings; IoT/CPS telemetry updates Digital Twins and knowledge-graph relationships; decision outputs are recorded back into the graph, closing the learning loop.

Feedback loops operate in multiple directions:

- Top-down: decision intelligence and engineering updates drive changes to plans, routings and control logic.
- Bottom-up: telemetry-driven anomalies, commissioning results and performance trends feed back into engineering models, schedules and templates.
- Cross-domain: semantic mappings ensure that a single change (e.g. relocating a conveyor) is coherently reflected across geometry, BOMs, project plans, work instructions and operational monitoring.

Through these loops, DVCO behaves as an adaptive system that continually synchronises engineering intent with operational reality.

#### 4.6.3 Figure Description

*Figure 4.2* illustrates the multi-layer DVCO architecture as a stacked diagram, with the Physical/IoT/CPS layer at the bottom and the Decision Intelligence layer at the top. The Semantic (Ontology/Knowledge Graph) layer is depicted as a vertical spine connecting and enriching all layers, indicating its role as the central integration mechanism. Horizontal arrows show lifecycle flows (engineering → execution → operations), while vertical arrows represent feedback loops and predictive insights flowing back to engineering and planning..

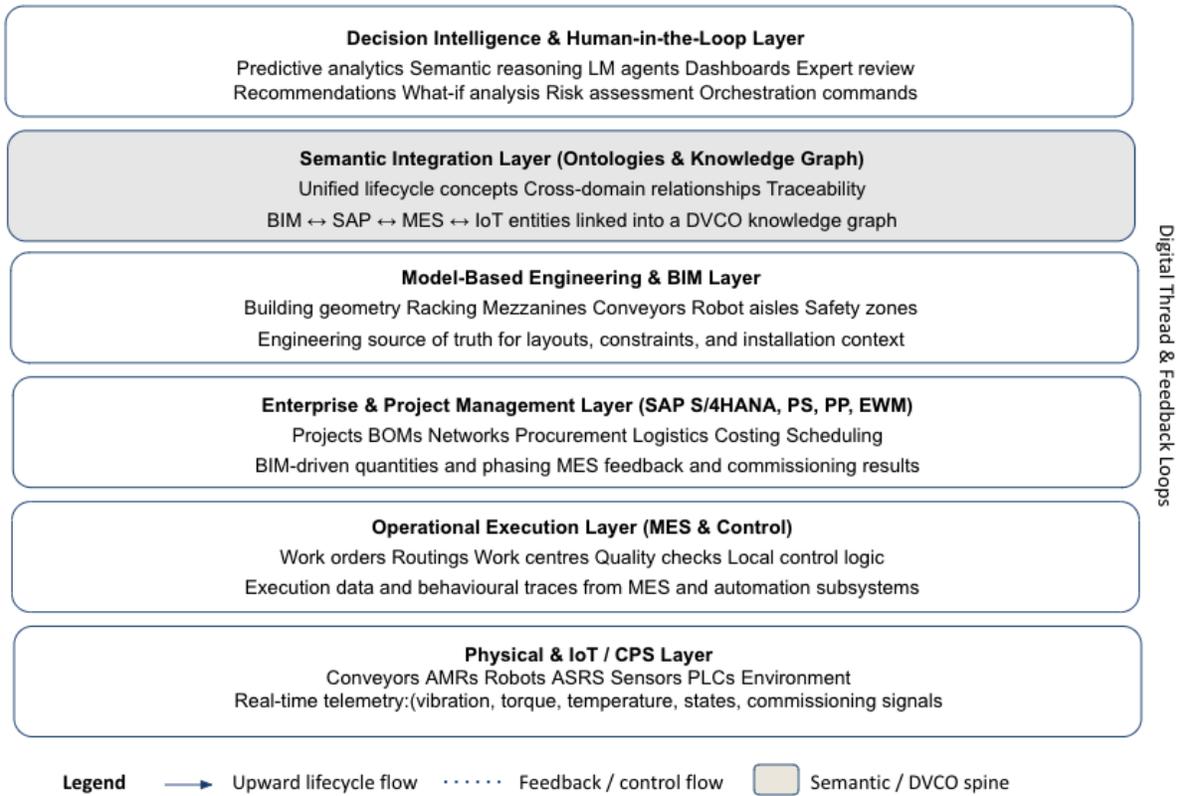


Figure 4.2 . Multi-Layer DVCO Architecture

#### 4.7 LLM-Enabled Reasoning and Decision Intelligence

Large Language Models (LLMs) represent a major advancement in artificial intelligence, enabling machines to interpret natural language, synthesise heterogeneous technical data and generate cross-domain insights in complex industrial environments. Within the DVCO

architecture, LLMs act as cognitive engines that make lifecycle data accessible and interpretable to engineers, planners, project managers and commissioning teams. Rather than relying on SQL queries, system-specific dashboards or custom reporting tools, stakeholders can interact with the entire value chain using natural-language prompts (Davenport and Ronanki, 2018; Sheth, 2020).

This capability is particularly valuable in ETO environments, where variability, high configuration complexity and constant change challenge traditional analytics systems.

#### 4.7.1 LLMs as Cognitive Interfaces Across the Value Chain

LLMs provide a unifying cognitive interface enabling users to query, interpret and analyse multi-system data—including BIM metadata, engineering documents, change orders, SAP transactional logs, MES execution records and IoT telemetry—using natural-language prompts (Samoila, López and Gutiérrez-Ríos, 2021). Unlike conventional AI models that require structured, repetitive datasets, LLMs excel in low-repeatability, high-variability environments such as Symbotic’s ETO deployments.

LLMs can generate:

- contextual summaries of engineering changes;
- natural-language explanations of cross-domain impacts;
- insights into installation or commissioning bottlenecks;
- interpretive analyses of anomaly signals from IoT/CPS systems;

- automatically generated progress, variance or risk reports.

By providing conversational access to lifecycle information, LLMs significantly reduce cognitive load and democratise advanced reasoning capabilities across disciplines.

#### 4.7.2 Hybrid Reasoning: Integrating LLMs with Ontologies and Knowledge Graphs

Although LLMs possess extensive general reasoning ability, they lack inherent domain grounding and are vulnerable to hallucinations when confronted with specialised engineering terminology or incomplete lifecycle data (Sheth, 2020). DVCO addresses this through a hybrid reasoning architecture that integrates LLMs with:

- the DVCO ontology, which provides semantic grounding and shared domain meaning (Bock and Gruninger, 2005);
- the DVCO Knowledge Graph (KG), which provides traceable, structured lifecycle relationships (Beetz, van Leeuwen and de Vries, 2009);
- domain constraints encoded in engineering, manufacturing and installation rule sets.

This integration ensures that LLM outputs remain semantically accurate, explainable and lifecycle-aligned.

Hybrid LLM–KG reasoning allows:

- semantic retrieval, enabling LLMs to query relevant KG subgraphs rather than using unbounded probabilistic memory;

- traceable explanations, where users can view graph edges and nodes underpinning an inference;
- multi-hop dependency reasoning, essential for evaluating engineering-change impacts in ETO settings;
- constraint-aware recommendations, grounded in SAP routings, MES sequences and BIM spatial rules;
- probabilistic forecasting, combining historical telemetry patterns with semantic context.

Through this hybrid design, LLMs evolve from general-purpose language models into domain-responsible, operationally actionable reasoning engines suitable for DVCO.

#### 4.7.3 LLM-Enabled Use Cases in DVCO

Large Language Models (LLMs) enhance DVCO by providing reasoning, synthesis and automation capabilities across the entire ETO/D2O lifecycle.

Because LLMs operate on top of the ontology–knowledge-graph (KG) layer, their outputs are grounded in semantically coherent, lifecycle-connected data rather than isolated text generation. This enables AI-assisted decision-making that is context-aware, traceable and operationally relevant.

Representative use cases include the following.

#### 4.7.3.1 Natural-Language Querying of Multi-System Data

LLMs enable natural-language access to multi-domain data by synthesising BIM, SAP, MES and IoT information through the KG.

This allows engineers, planners and commissioning teams to ask complex lifecycle questions without navigating multiple systems:

- *“What components are impacted by the mezzanine redesign in Zone B?”*
- *“Which suppliers pose the highest risk to Week-7 commissioning?”*
- *“Show all installation tasks dependent on AMR alignment in Aisle 12.”*

The LLM retrieves semantically linked entities—BIM elements, SAP materials, MES routings, IoT telemetry—and returns answers grounded in KG relationships.

This shifts problem-solving from manual data searching to semantic, AI-enabled analysis, greatly improving decision speed and accuracy.

#### 4.7.3.2 Change-Impact Analysis

LLMs can interpret engineering changes and traverse KG dependencies to identify downstream operational impacts.

This includes automated reasoning across:

- materials (SAP MM),
- routings (MES/PP),

- shop-floor tasks,
- installation sequences,
- commissioning readiness,
- supplier constraints,
- regulatory or safety dependencies.

In ETO environments where engineering revisions occur frequently, LLM-driven impact analysis accelerates alignment between design and execution, reducing delays and rework.

#### 4.7.3.3 Automated Technical Documentation and Reporting

Because LLMs can generate structured, context-aware narratives, they automate key reporting functions across the lifecycle, including:

- installation progress summaries,
- commissioning stability and readiness reports,
- deviation and anomaly summaries based on IoT/CPS patterns,
- risk dashboards and mitigation briefs,
- engineering impact summaries after design changes.

These automations reduce administrative burden, improve documentation quality and ensure that reporting reflects semantically aligned lifecycle data.

#### 4.7.3.4 Predictive and Prescriptive Reasoning

When combined with telemetry, historical lifecycle patterns and graph-based reasoning, LLMs support predictive and prescriptive decision intelligence.

LLM-enabled outputs include:

- early-warning indicators for installation or commissioning risks,
- delay or deviation predictions based on telemetry trends,
- prescriptive mitigation recommendations,
- alternative routing, sequencing or resource-allocation strategies,
- dynamic schedule compression options grounded in real operational constraints.

This shifts DVCO from a reactive decision model to a proactive and preventative orchestration system, improving lifecycle reliability and resilience.

#### 4.7.3.5 Multi-Scenario Simulation Support

LLMs assist in evaluating alternative planning and execution scenarios by interpreting Digital Twin models, KG dependencies and operational constraints.

Representative simulations include comparisons of:

- installation sequence options,
- labour and crew allocation strategies,
- vendor sourcing configurations,
- commissioning-readiness strategies,

- schedule acceleration or buffering options.

Crucially, LLM outputs are not hypothetical—they are grounded in the KG’s representation of engineering, sequencing, material and telemetry dependencies, ensuring that simulations remain operationally realistic and technically feasible.

#### 4.7.3.2. Change-Impact Analysis

LLMs interpret engineering changes, traverse KG dependencies and identify downstream impacts on:

- materials,
- routings,
- shop-floor tasks,
- installation sequences,
- commissioning readiness.

This accelerates design–execution alignment in highly dynamic environments.

#### 4.7.3.3. Automated Technical Documentation and Reporting

LLMs automatically generate:

- installation progress summaries,
- commissioning stability reports,

- risk dashboards,
- engineering impact briefs.

These automations reduce administrative workload and increase reporting accuracy.

#### 4.7.3.4. Predictive and Prescriptive Reasoning

LLMs use telemetry, historical patterns and lifecycle data to produce:

- early-warning indicators;
- delay or deviation predictions;
- mitigation recommendations;
- alternative sequencing or scheduling scenarios.

This allows DVCO to operate proactively rather than reactively.

#### 4.7.3.5. Multi-Scenario Simulation Support

LLMs compare alternative scenarios such as:

- installation sequences,
- labour allocation strategies,
- vendor sourcing options,
- schedule compressions.

Crucially, outputs are grounded in KG-based dependencies and operational constraints, ensuring realism.

#### 4.7.4 Human–AI Collaboration in DVCO Decision Workflows

DVCO positions LLMs as augmented-intelligence partners rather than autonomous decision-makers.

LLMs provide:

- synthesised insights,
- rapid cross-domain reasoning,
- interpretable diagnostics,
- evidence-based recommendations.

Humans retain authority over scheduling, resource allocation and engineering changes, while LLMs accelerate analysis, reduce uncertainty and amplify organisational decision capacity.

This collaboration improves:

- risk detection speed,
- engineering–operations alignment,
- justification of decisions using traceable reasoning,

- reduction of dependency on tacit knowledge and individual experience.

#### 4.7.5 LLMs as Enablers of Dynamic Value Chain Optimisation

Within DVCO, LLMs serve as the cognitive engine that enables dynamic optimisation across the entire value chain. By combining natural-language reasoning with semantic grounding and predictive analytics:

- variability becomes more manageable;
- risks are detected earlier;
- dependencies are understood more clearly;
- commissioning becomes more stable and predictable;
- organizational learning accelerates.

LLMs therefore transform DVCO from a static integration model into a continuously sensing, interpreting and adapting intelligence ecosystem, capable of supporting Symbotic's complex ETO deployments with unprecedented speed, accuracy and cross-domain coherence.

#### 4.8 Integrated DVCO Architecture

#### 4.8.1 Overview of the Multi-Layer Architectural Model

The Integrated DVCO Architecture synthesises engineering models, enterprise processes, real-time telemetry, semantic representations and AI-enabled reasoning into a unified, closed-loop optimisation system. Traditional ETO environments typically operate through loosely connected, siloed systems, creating delays, inconsistencies and a lack of lifecycle coherence across engineering, manufacturing and installation (Gosling and Naim, 2009; Culot et al., 2020).

In contrast, DVCO establishes a model-driven, semantically unified, and dynamically adaptive architecture capable of coordinating decisions across design, planning, procurement, manufacturing, logistics, installation and commissioning. (Teece, Pisano and Shuen, 1997; Lin, Lee and Ma, 2020).

The architecture consists of six interdependent layers, each contributing a distinct form of intelligence or lifecycle context:

1. Model Layer (MBE/BIM/EBOM)

Provides authoritative engineering representations of structures, behaviours, constraints and spatial geometry (Estefan, 2007; Kreider and Messner, 2013).

2. Execution Layer (SAP S/4HANA + MES)

Translates engineering intent into routings, work orders, cost plans, procurement schedules and operational sequences (Weyer et al., 2015).

3. Telemetry Layer (IoT + CPS)

Supplies real-time behavioural, environmental and equipment data that enable Digital Twins to reflect physical conditions dynamically (Lee, Bagheri and Kao, 2015).

4. Semantic Layer (Ontology + Knowledge Graph)

Unifies concepts, relationships and lifecycle semantics across engineering, planning, execution and operations (Sheth, 2020).

5. AI Reasoning Layer (LLMs + Predictive Models)

Performs semantic retrieval, dependency analysis, anomaly detection, risk prediction and natural-language reasoning (Davenport and Ronanki, 2018).

6. Orchestration Layer (DVCO Control Logic)

Implements workflows for change propagation, risk mitigation, scheduling adjustments and lifecycle governance (Lin, Lee and Ma, 2020).

Together, these layers form a continuous information loop, linking design, execution and operations through a shared semantic and model-driven architecture. (Maier, 1998; ISO/IEC/IEEE 15288, 2015).

#### 4.8.2 Cross-Layer Data Flow and Lifecycle Synchronisation

Central to DVCO is the principle of bidirectional data flow, ensuring synchronisation between physical and digital domains throughout the D2O lifecycle.

- Engineering changes in BIM/MBE automatically update SAP routings, procurement plans, MES work instructions and installation sequences.
- IoT/CPS telemetry updates the Digital Twin, triggering risk alerts or predictive insights (Tao et al., 2018).
- LLM reasoning traverses the knowledge graph to reveal multi-step dependencies and propose mitigation actions (Samoila, López and Gutiérrez-Ríos, 2021).
- The Orchestration Layer applies recommendations by adjusting schedules, materials, or sequencing logic.
- Updated lifecycle states feed back into engineering and planning models, closing the loop.

This iterative process enables continuous optimisation, allowing DVCO to sense, interpret and adapt to lifecycle deviations in real time.

#### 4.8.3 Architectural Integration Principles

DVCO's architectural integrity is governed by several key principles:

##### Semantic Interoperability

All systems—BIM, SAP, MES, IoT—are harmonised through the ontology and knowledge graph, ensuring unified lifecycle interpretation (Beetz, van Leeuwen and de Vries, 2009).

##### Model-Based Authority

Engineering models (MBE/BIM) serve as the single source of truth, governing spatial, structural and behavioural constraints (Hedberg et al., 2019).

#### Closed-Loop Feedback

Telemetry continuously validates and updates engineering assumptions and operational plans, grounding the digital twin in physical reality (Grieves, 2014).

#### Explainable AI Reasoning

LLM outputs are grounded in KG semantics, ensuring interpretability, transparency and domain correctness (Sheth, 2020).

#### Dynamic Adaptation

Schedules, routings, procurement priorities and installation logic adapt automatically to emerging risks or constraints.

#### Scalability and Extensibility

The architecture supports multi-site deployments, supplier ecosystems and future AI/automation extensions.

#### 4.8.4 Value-Chain-Level Digital Twin

At the centre of the DVCO architecture is the value-chain-level digital twin, which integrates multiple forms of twins:

- Structural twins (BIM geometry, spatial models)
- Execution twins (ERP/MES routings, work orders, WBS networks)
- Performance twins (IoT telemetry, CPS behavioural states)
- Semantic twins (ontology, KG relationships)
- Intelligence twins (AI/LLM reasoning agents)

This holistic digital twin enables DVCO to simulate, monitor and optimize the entire value chain, rather than isolated assets or processes—an advancement beyond most Industry 4.0 digital-twin implementations (Boschert and Rosen, 2016; Grieves and Vickers, 2017; Tao et al., 2019).

#### 4.8.5 DVCO Orchestration Logic

The Orchestration Layer embeds DVCO control logic that operationalises lifecycle optimisation through:

##### Change-Impact Propagation

Example: A mezzanine redesign recalculates installation tasks, supplier dependencies and commissioning impacts.

##### Risk-Mitigation Workflows

Example: Vendor delays trigger alternative sourcing, resequencing or work-centre reallocation.

#### Installation / Commissioning Optimisation

Example: IoT-validated readiness scores dynamically adjust task sequences and commissioning scripts.

#### Multi-Domain Alerting and Reporting

Example: LLM-generated summaries provide real-time insight to project leadership.

These workflows transform DVCO from a static analytical concept into a dynamic operational decision engine capable of real-time adaptation. (Teece et al.,1997; Lin, Lee and Ma. 2020)

#### 4.8.6 Architectural Significance for ETO and Symbiotic

DVCO's integrated architecture delivers several transformative benefits:

- Predictive rather than reactive orchestration
- Lifecycle transparency across engineering, manufacturing, logistics and installation
- Faster detection of cross-domain bottlenecks
- Improved engineering–operations alignment
- Higher commissioning reliability and deployment repeatability

- Scalability across multiple customer sites and programs

The uploaded SAP + BIM Integration document further demonstrates DVCO's implementability, providing evidence of lifecycle alignment between engineering models and enterprise execution structures (SAP + BIM Overview).

Through this architecture, DVCO becomes not merely a conceptual model but a practical, implementable blueprint for modernising the ETO value chain.(Gosling and Naim, 2009; Culot et al., 2020)

#### 4.9 Summary

This chapter examined the technological foundations enabling the Dynamic Value Chain Optimisation (DVCO) framework to function as a unified, model-driven and intelligence-enabled architecture for Symbotic's ETO value chain. DVCO's strength arises from the orchestrated integration of multiple domains—engineering, enterprise execution, manufacturing, installation, commissioning and operations—brought together through semantic and AI-enabled mechanisms.

The chapter began by establishing Building Information Modelling (BIM) as the structural and spatial backbone for engineering-driven lifecycle coherence.

It then demonstrated how ontologies and knowledge graphs provide semantic alignment,

cross-domain dependency modelling and explainable reasoning.

The roles of SAP S/4HANA and MES were outlined as the software-defined execution layer translating engineering intent into operational reality.

IoT and Cyber-Physical Systems (CPS) were shown to contribute real-time behavioural context to sustain a continuously synchronised digital twin.

Finally, LLM-enabled reasoning was presented as the intelligence layer that enables DVCO to sense, interpret and adapt within highly dynamic ETO environments.

Section 4.8 synthesised these domains into a single integrated architecture, revealing how DVCO unifies engineering models, enterprise systems, telemetry, semantic representations and AI reasoning into a closed-loop optimisation ecosystem.

Together, these technologies form the foundation upon which Chapters 5 and 6 advance:

Chapter 5 outlines the research methodology used to construct and evaluate DVCO, and

Chapter 6 presents empirical findings from modelling, simulation and reasoning experiments.

# CHAPTER V.

## RESEARCH METHODOLOGY

### 5.1 Research Framework and Phases

This study employs a multi-phase research methodology integrating qualitative inquiry, system modelling and simulated AI-enabled optimisation. The methodological design reflects both the complexity of the Engineer-to-Order (ETO) domain and the cross-functional dynamics of the Design-to-Operate (D2O) value chain (Gosling and Naim, 2009; Culot et al., 2020). Given that Symbolic’s deployment model involves high variability, engineering-driven workflows and data-intensive coordination, a structured, phased methodology is essential to analyse interdependencies and evaluate the potential of Dynamic Value Chain Optimisation (DVCO) as a holistic framework (Teece, Pisano and Shuen, 1997; Lin, Lee and Ma, 2020).

The research is organised into three sequential phases—Exploration, Modelling, and Simulation and Validation. Each phase builds on the preceding stage, collectively providing a comprehensive approach for understanding the value chain, constructing a formal representation and assessing the impact of AI-enabled optimisation strategies (Meredith, 1998; Voss, Tsikriktsis and Frohlich, 2002).

#### 5.1.1. Phase 1: Exploration – Mapping the ETO Value Chain and Identifying Bottlenecks

The exploration phase employs qualitative analysis techniques, including document review, expert-informed interpretation, process mapping and system-landscape analysis (Yin, 2014). The objective is to articulate how Symbotic executes its ETO projects across engineering, procurement, manufacturing, logistics, installation and commissioning.

Key activities include:

- mapping the end-to-end D2O process;
- identifying structural and operational bottlenecks (e.g. engineering-change propagation, late supplier dependencies, installation-readiness gaps);
- assessing fragmentation across BIM, SAP, MES, IoT and logistics systems;
- identifying constraints that impede dynamic replanning and decision-making.

This phase establishes the empirical and conceptual foundation for constructing the ontology and knowledge graph in the subsequent stage (Sheth, 2020; Grieves, 2014).

### 5.1.2. Phase 2: Modelling – Ontology and Knowledge Graph Construction

The modelling phase develops a formal semantic representation of Symbotic's D2O environment. Ontology engineering techniques are applied to define key concepts (e.g. manufacturing orders, engineering changes, robotic cells, site tasks), properties, relationships and lifecycle rules (Bock and Gruninger, 2005).

The knowledge graph is then constructed to represent actual or simulated instances of entities and interdependencies across:

- BIM model elements;
- SAP ERP and S/4HANA objects;
- MES/PEO process steps;
- IoT sensor telemetry;
- logistics and installation events.

This modelling phase provides the structural basis for traceability, semantic search, impact propagation and AI-enabled reasoning. It converts previously siloed datasets into a unified graph structure consistent with DVCO principles and Industry 4.0 semantic-integration practices (Främling et al., 2013; Beetz, van Leeuwen and de Vries, 2009).

### 5.1.3. Phase 3: Simulation and Validation – Testing Improvements Using AI-Enabled Scenarios

The final phase employs a conceptual simulation approach inspired by Palantir Foundry’s ontology-based modelling capabilities and digital-platform paradigms used in Industry 4.0 environments (Lee, Bagheri and Kao, 2015; Tao et al., 2018). While operational integration is outside the scope of this academic research, the simulation environment is used to:

- model alternative process configurations;

- evaluate the impact of engineering-change timing, supplier delays or installation resequencing;
- simulate dynamic replanning based on real-time data inputs;
- test predictive models that use LLM-assisted reasoning to anticipate delays, bottlenecks or cost overruns.

Scenario outputs are analysed to assess improvements in process visibility, response speed, coordination efficiency and predicted project performance. This validation step provides evidence for how the DVCO framework may enhance the execution of complex ETO projects.

## 5.2 Research Design

The research design follows a hybrid qualitative–model-based methodology appropriate for analysing complex, data-intensive industrial systems such as Symbotic’s ETO and D2O value chain. The design integrates three methodological paradigms:

1. qualitative systems analysis;
2. formal modelling through ontologies and knowledge graphs;
3. simulated AI-enabled optimisation.

This combination supports both interpretive understanding of organisational phenomena and the construction of computational artefacts capable of representing and evaluating system behaviour (Meredith, 1998; Eisenhardt, 1989).

### 5.2.1 Qualitative Systems Analysis

The initial component of the research design is rooted in qualitative methods, used to develop a deep understanding of the structure and dynamics of Symbotic's ETO operations. The qualitative design includes:

- Process mapping: tracing engineering, procurement, manufacturing and installation workflows;
- Document analysis: reviewing design documents, BOM structures, SAP process chains, FAT/SAT procedures and integration specifications;
- Expert insights: drawing on practitioner knowledge from engineering leads, automation specialists, supply-chain planners and system integrators (via published sources, practitioner documentation and industry-standard references);
- System landscape analysis: examining interactions between BIM, SAP S/4HANA, MES/PEO, IoT telemetry and logistics systems.

This component provides the empirical grounding necessary to ensure that the modelling and simulation phases accurately reflect real-world operational complexity (Yin, 2014; Voss, Tsikriktsis and Frohlich, 2002).

### 5.2.2 Model-Based Design Using Ontologies and Knowledge Graphs

The second methodological component uses model-based design to represent the semantics, information flows and dependency structures underlying the Symbiotic D2O environment. This phase is constructive rather than exploratory, seeking to formalise knowledge such that it can be analysed computationally (Madni and Sievers, 2018).

Key modelling techniques include:

- Ontology engineering to define domain concepts (e.g. *ChangeOrder*, *ManufacturingOrder*, *RoboticCell*, *SiteTask*) and their relationships (e.g. *depends\_on*, *affects*, *located\_at*, *derives\_from*);
- Knowledge-graph construction to represent real or simulated instances, linking BIM elements, SAP objects, MES steps and IoT events into a unified relational network;
- Graph reasoning rules to infer hidden dependencies, identify bottlenecks and support impact propagation.

By translating value chain structures into graph-based models, the research design facilitates AI reasoning, semantic search and scenario-based simulation (Bock and Gruninger, 2005; Sheth, 2020).

### 5.2.3 AI-Enabled Simulation and Scenario Analysis

The third component employs a simulation design inspired by Palantir Foundry's ontology-driven modelling environment and similar digital-twin platforms (Tao et al., 2018). Although this dissertation does not implement a production deployment, it uses conceptual simulation to evaluate how AI-augmented systems may improve value-chain performance.

The simulation design includes:

- Scenario construction: modelling variations in change timing, supplier delays, resource constraints, installation sequencing and quality deviations;
- Digital-twin alignment: linking simulated BIM changes with corresponding SAP, MES and site-execution impacts (Grieves, 2014; Boschert and Rosen, 2016);
- Prediction mechanisms: using LLM-assisted reasoning to forecast delays, cost risks, material shortages and installation bottlenecks;
- Performance evaluation: comparing simulated outcomes across KPIs such as schedule adherence, change-impact radius, rework reduction and supply-chain responsiveness.

This design allows the research to assess potential DVCO improvements without requiring operational access to proprietary systems.

#### 5.2.4 Justification for the Hybrid Research Design

The hybrid approach is justified for several reasons:

- ETO/D2O operations are socio-technical systems; qualitative understanding is necessary but insufficient without formal modelling (Gosling and Naim, 2009).
- Ontology and knowledge-graph models capture complexity that traditional linear or purely statistical methods cannot represent (Främling et al., 2013; Sheth, 2020).
- AI-enabled simulation allows predictive evaluation, demonstrating the potential value of DVCO in a controlled research setting (Kusiak, 2018).
- The design aligns with current digital-transformation paradigms, including MBSE, semantic interoperability and AI-assisted operational intelligence (Madni and Sievers, 2018; Lee, Bagheri and Kao, 2015).
- It provides both theoretical insight and practical applicability, supporting academic contribution and industry relevance.

Thus, the research design ensures methodological rigour while remaining aligned with the technical and operational realities of Symbotic's value chain.

### 5.3 Data Sources and Data Collection

The research draws upon a triangulated set of qualitative, documentary, and synthetic data sources to support the multi-phase methodological framework outlined in Chapter 5. Because Symbotic's operational datasets, engineering documents, software configurations and commissioning records are proprietary, the study adopts an academically rigorous approach that relies on:

- representative datasets reflecting typical ETO/D2O data structures,
- industry-standard documentation (SAP, BIM, MES, robotics, commissioning),
- publicly available warehouse-automation materials, and
- synthetic but semantically accurate models to support ontology engineering and DVCO simulation.

This combination ensures both methodological depth and respect for confidentiality constraints. Synthetic datasets—constructed according to real-world structural patterns such as multi-level BOMs, WBS hierarchies, change records, routing networks, and IoT telemetry schemas—are used to validate the ontology, digital-thread mapping and reasoning workflows. This approach aligns with recognised practices in digital-twin and knowledge-graph research where proprietary operational data cannot be disclosed (Culot et al., 2020; Lee, Bagheri and Kao, 2015).

### 5.3.1 Documentary Sources

Documentary sources constitute the core foundation for reconstructing Symbotic's Engineer-to-Order (ETO) and Design-to-Operate (D2O) value chains and for analysing the interdependencies between engineering, procurement, manufacturing, logistics, installation and commissioning activities. These materials enable triangulation across industry standards, enterprise-system reference models and warehouse-automation practices, thereby ensuring the accuracy and validity of DVCO modelling.

## 1. Value-Chain and Industry Process Documentation

Warehouse-automation value-chain documentation provides detailed insight into the sequence of activities from conceptual engineering through commissioning. These descriptions, consistent with ETO literature (Gosling and Naim, 2009) and Industry 4.0 system-integration studies (Culot et al., 2020), clarify how engineering decisions propagate downstream into procurement, manufacturing and installation.

ETO execution workflows, including engineering change management, long-lead procurement, parallel manufacturing and on-site installation coordination, align with ERP/MES reference standards for high-variability manufacturing (SAP SE, 2021; ISO, 2018). These sources support the identification of bottlenecks and coordination gaps across lifecycle stages.

Reference mappings for the digital thread, process thread, execution thread, decision thread and collaboration thread—presented in the uploaded *SAP + BIM* document—provide a multi-layer, cross-domain view of lifecycle data connectivity. These mappings are consistent with established digital-thread and digital-twin theory (Grieves, 2014; Hedberg et al., 2016).

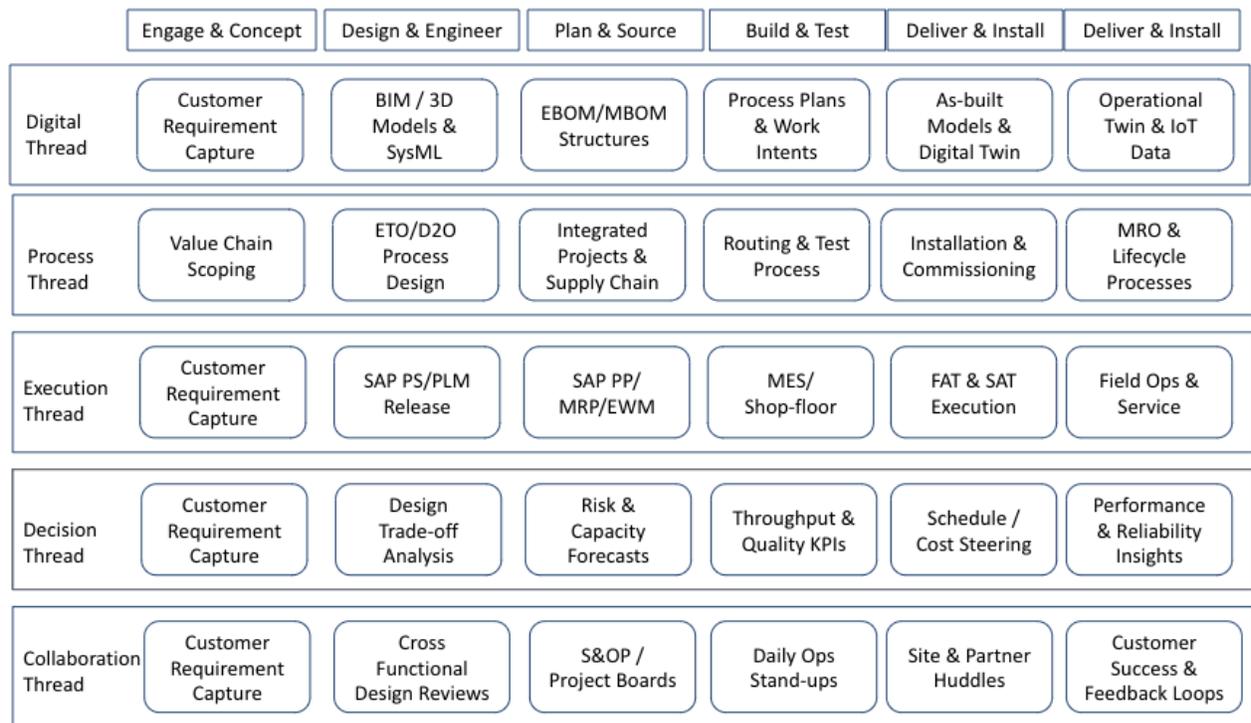


Figure 5.1 DVCO ARCHITECTURE - THREADS VIEW

## 2. SAP / ERP Documentation

Standardised SAP S/4HANA business-process documentation (SAP SE, 2021), including procurement (MM), production (PP/PEO), project systems (PS) and engineering change management (ECM), provides the operational logic and data structures required to model Symbotic’s ETO system behaviour. SAP PS principles are particularly critical for modelling WBS hierarchies, cost structures and progress tracking (Jain, Chandrasekaran and Gunasekaran, 2019).

ERP documentation offers detailed reference points for:

- bill of materials (BOM) transformations across EBOM→MBOM→Order BOM,
- integration flows between engineering, supply chain and manufacturing,
- configuration management and change-impact propagation, and
- long-lead procurement and milestone-driven supply coordination.

These documents allow the DVCO ontology to structurally represent ERP entities, control points and transactional relationships (ISO, 2018).

### 3. MES and Manufacturing-Operations Documentation

MES reference architectures (SAP SE, 2022) and ISA process-control standards support the modelling of production execution, quality inspection, traceability and shop-floor feedback loops. These sources clarify:

- routing structures,
- operation sequencing,
- production confirmation and rework handling,
- work-centre resource allocation, and
- material traceability and quality checkpoints.

Such documentation is aligned with cyber-physical manufacturing models in Industry 4.0 (Lee, Bagheri and Kao, 2015) and is essential for linking DVCO decision logic to real-time shop-floor execution behaviour.

#### 4. BIM and AEC Integration Documentation

BIM documentation provides the semantic and geometric structures needed for linking engineering design models to ERP/MES data flows. Industry guidelines from Autodesk (2020), ISO 19650 (ISO, 2019) and BIM–ERP integration literature inform the representation of:

- 3D/4D/5D BIM objects and metadata,
- scope breakdown and parametric design,
- construction and installation sequences,
- cost, schedule and material linkages, and
- model-based collaboration workflows.

The uploaded *SAP + BIM* reference document illustrates specific integration touchpoints between BIM models and SAP processes, such as WBS derivation, BOM generation and schedule synchronisation.

These materials correspond to documented best practices in virtual design and construction (Kunz and Fischer, 2012) and digital-thread-enabled AEC workflows (Eastman et al., 2011).

#### 5. Symbolic Publicly Available Materials

Although Symbotic’s operational datasets are proprietary, publicly available architectural, product and deployment documentation provides essential insight into:

- its modular robotics platform,
- autonomous mobile robots and shuttle systems,
- high-density storage structures,
- WES/WCS orchestration logic, and
- project-delivery and commissioning models.

These materials provide sufficient structural detail to develop a representative, academically valid DVCO model (Symbotic LLC, 2022; Symbotic LLC, 2023; Walmart and Symbotic, 2022).

## 6. FAT/SAT, Commissioning and Industry Standards

Factory Acceptance Testing (FAT) and Site Acceptance Testing (SAT) documentation from ISA (2019) and ASTM (2017) provides reference structures for modelling:

- verification and validation workflows,
- installation dependencies,
- operational readiness criteria, and
- commissioning risk mitigation.

These standards mirror the procedures used in high-complexity automation deployments and align with Symbotic’s publicly described commissioning approach.

### 5.3.2 Synthetic Data Construction

Because Symbolic's operational datasets, engineering models and transactional records are proprietary, the study employs synthetic but structurally accurate datasets to support ontology development, knowledge-graph population and simulation activities. The use of synthetic data is well established in digital-twin, semantic-modelling and AI research where confidentiality restrictions limit access to real-world enterprise data (Hedberg et al., 2016; Hogan et al., 2021).

#### 1. Purpose of Synthetic Data in the DVCO Framework

Synthetic datasets were designed to:

- replicate the structural patterns of ETO lifecycle data,
- populate the DVCO ontology with realistic multi-domain entities,
- enable cross-domain reasoning using LLM-assisted inference,
- simulate design changes, schedule shifts and risk propagation,
- validate digital-thread continuity across engineering, ERP, MES and commissioning.

This approach ensures methodological rigour while respecting confidentiality constraints, in line with recommended practices for model-based enterprise research (Grieves, 2014; Madni and Sievers, 2018).

#### 2. Data Structures Modelled

The synthetic data covers all major lifecycle domains that appear in an ETO warehouse-automation programme:

- Engineering models: EBOM structures, 3D component metadata and model-based requirements derived from BIM references (ISO, 2019; Eastman et al., 2011).
- ERP transactional objects: WBS elements, purchase orders, production orders, long-lead procurement records and cost structures based on SAP reference documentation (SAP SE, 2021).
- MES execution logs: routing operations, operation confirmations, quality records and rework data based on standard MES architectures (Lee, Bagheri and Kao, 2015; SAP SE, 2022).
- Commissioning data: FAT/SAT checklists, validation results and operational-readiness milestones referencing ISA standards (ISA, 2019; ASTM, 2017).
- IoT telemetry: sensor readings, robot-position logs and fault events incorporating patterns from cyber-physical manufacturing research (Tao et al., 2018).

These synthetic datasets allow the DVCO system to perform lifecycle simulations that reflect realistic interdependencies observed in warehouse-automation deployments.

### 3. Principles for Synthetic Data Design

The data-generation process followed four principles:

1. Structural fidelity — ensuring entity types, hierarchies and relationships match the real structure of ETO/BIM/ERP/MES data (Culot et al., 2020).
2. Semantic accuracy — ensuring attributes and relationships comply with the DVCO ontology and industry standards (Sheth, 2020; Hogan et al., 2021).
3. Cross-domain connectivity — ensuring that model-based objects (e.g., BIM assemblies) link correctly to ERP objects (WBS, MBOM) and MES operations (routings).
4. Simulation suitability — ensuring the dataset supports change-impact analysis, schedule drift modelling, supply-risk propagation and commissioning-readiness evaluation.

#### 4. Tools and Methods for Data Generation

Synthetic data was generated and validated using:

- Ontology-driven schema generation, following RDF/OWL best practices (Hogan et al., 2021).
- Knowledge graph population scripts, enabling SPARQL-based validation of entity consistency.
- Python-based data generators for ERP and MES structures based on SAP reference configurations (SAP SE, 2021; SAP SE, 2022)
- Scenario-based simulation templates for change propagation, schedule adjustment and resource conflicts (Sterman, 2000).

## 5. Validation of Synthetic Data

To ensure the academic validity of the synthetic dataset:

- Semantic validation was performed using SHACL constraints.
- Cross-domain consistency checks ensured correct linkage between engineering, procurement, production and installation data.
- Lifecycle plausibility checks verified that generated schedules, BOMs, orders and commissioning milestones reflected realistic ETO program behaviour (Gosling and Naim, 2009).

This validation ensures that the synthetic dataset provides a reliable basis for evaluating DVCO mechanisms despite the absence of proprietary Symbotic data.

### 5.3.3 Expert Knowledge and Domain Understanding

Expert insights inform this dissertation's understanding of ETO challenges and D2O execution dynamics. These insights are drawn from:

- the author's professional experience in SAP PLM/ERP, MES, ETO integration and Industry 4.0 architecture;
- engineering and construction best practices described in professional white papers and conference presentations;

- published interviews, webinars and technical briefings from warehouse automation and robotics vendors;
- SAP industry solution documentation for ETO, PEO and MBSE integration;
- guidance from digital-twin and semantic-modelling research communities (Grieves, 2014; Boschert and Rosen, 2016; Sheth, 2020).

Expert knowledge is particularly important for defining realistic dependencies used in the simulation models, especially where public data is insufficient or incomplete.

#### 5.3.4 Synthetic Data for Ontology and Simulation

Because operational datasets from Symbotic are not publicly accessible, the study constructs synthetic datasets that closely mirror:

- engineering structures (multilevel BOMs, change histories, model elements);
- production orders and routings;  
material lead times and vendor profiles;
- installation task networks and commissioning steps;
- sensor-telemetry patterns reflecting robot performance or environmental conditions.

Synthetic data allows the knowledge graph to behave structurally like a real deployment data environment, enabling scenario simulation without revealing proprietary information (Culot et al., 2020).

### 5.3.5 Data for AI Reasoning and LLM Interaction

The study also incorporates text-based sources to design LLM-enabled reasoning scenarios.

These include:

- engineering-change descriptions;
- site-readiness reports;
- delay logs and punch-list examples;
- typical communication patterns across engineering, manufacturing and installation teams;
- sample KPI dashboards reflecting supply chain, installation and robotic performance.

These textual examples allow the research to model how an LLM could answer natural-language queries, perform semantic search and assist project decision-making (Davenport and Ronanki, 2018; Samoila, López and Gutiérrez-Ríos, 2021).

### 5.3.6 Data Collection Approach and Rationale

The data collection strategy is designed to produce explicit, inspectable artefacts that enable verification, simulation, and analytical traceability of the DVCO framework (Hevner et al., 2004; Yin, 2014), rather than remaining at a purely conceptual level. Data is collected and constructed across four progressively concrete layers.

### Contextual Data Collection (Value Chain Grounding)

Primary contextual data is derived from structured document analysis of real-world ETO value-chain artefacts, including BIM-enabled project lifecycle representations, SAP Project System (PS) structures, and engineering-to-order execution workflows. These artefacts define concrete lifecycle phases, actors, deliverables, and decision points (e.g., requirements breakdown, EBOM generation, WBS structures, FAT/SAT milestones). This establishes a domain-grounded value-chain baseline against which all subsequent models are anchored.

### Structural Data Collection (Formal Model Extraction)

Structural data is obtained by analysing explicit system artefacts such as:

- BIM object schemas (e.g., building components, assemblies, spatial constraints),
- SAP master and transactional object models (e.g., Project, WBS, Network, Material, Change Object),
- MES/production activity abstractions reflected in ETO manufacturing stages.

From these artefacts, formal ontology classes, attributes, and relationships are extracted and encoded (e.g., *BIM\_Component* ↔ *SAP\_WBS* ↔ *Manufacturing\_Operation*). This step yields a machine-interpretable structural model, not a narrative description, enabling direct population of a knowledge graph.

### Instance Data Construction (Executable Knowledge Graph Population)

To enable empirical testing while preserving confidentiality, synthetic but structurally isomorphic instances are generated. These instances include:

- Project instances with defined WBS hierarchies,
- BIM-derived component instances with quantities and installation constraints,
- Change events propagating across engineering, planning, and execution layers.

These instances populate the DVCO knowledge graph and are used to execute controlled simulations of change propagation, coordination latency, and decision impact, producing measurable outputs such as dependency traversal depth, update propagation time, and consistency violations.

### Behavioural Data Collection (Reasoning and Decision Evaluation)

Behavioural data consists of curated textual artefacts derived from engineering change descriptions, coordination scenarios, and lifecycle decision narratives aligned with the SAP-BIM ETO context. These artefacts are used to test AI-supported reasoning capabilities (e.g., impact analysis, cross-domain inference) against the populated knowledge graph. The focus is not language performance, but decision coherence and traceability across the value chain.

## Rationale

This staged approach ensures that each conceptual construct within the DVCO framework is backed by:

1. Concrete lifecycle artefacts (contextual),
2. Formal system models (structural),
3. Executable data instances (instance-level),
4. Observable reasoning outcomes (behavioural).

As a result, the DVCO framework is empirically grounded through model execution, simulation outputs, and traceable data objects, satisfying academic rigour requirements while maintaining industrial confidentiality

## 5.4 Data Analysis Procedures

This combined approach aligns with established research methods in model-based systems engineering (Estefan, 2007; Madni and Sievers, 2018), digital-thread analysis (Hedberg et al., 2016), and Industry 4.0 value-chain research (Lee, Bagheri and Kao, 2015; Culot et al., 2020).

- (1) qualitative process analysis,
- (2) ontology development,

(3) knowledge-graph implementation, and

(4) DVCO simulation and evaluation.

#### 5.4.1 Qualitative and Documentary Analysis

Documentary sources—including SAP process documentation, MES/MOM reference models, BIM integration guidelines, warehouse-automation industry materials and the uploaded *SAP + BIM* document—were analysed using structured qualitative coding (Creswell, 2013; Silverman, 2016).

This analysis consisted of:

- Lifecycle decomposition into engineering, supply chain, manufacturing, installation and commissioning stages.
- Identification of cross-domain dependencies, such as BOM transformations, long-lead procurement triggers, routing/installation alignment and SAP–BIM mapping.
- Extraction of recurring bottlenecks, including change-impact latency, procurement–engineering misalignment, production resequencing, and commissioning readiness uncertainties.

The qualitative stage provided the conceptual foundation for defining DVCO ontology classes, properties, system interfaces and decision nodes.

#### 5.4.2 Ontology Structuring and Semantic Alignment

Using insights from the qualitative analysis, the next stage involved structuring the DVCO ontology following RDF/OWL principles (Hogan et al., 2021) and MBE/MBSE modelling conventions (Hedberg and Lubell, 2017). The analysis steps included:

- identifying core domain entities (e.g., *EngineeringComponent*, *WBS*, *Operation*, *Supplier*, *InstallationTask*, *CommissioningCheck*);
- modelling cross-domain relationships (e.g., *derivesFrom*, *allocatedTo*, *dependsOn*, *changes*, *affects*);
- establishing semantic constraints to ensure lifecycle consistency (ISO 8000; ISO 19650; SAP SE, 2021).

Semantic alignment was achieved by mapping ontology classes to:

- BIM object types (Autodesk, 2020; Eastman et al., 2011),
- SAP ERP object schemas (SAP SE, 2021),
- MES routing structures (Lee, Bagheri and Kao, 2015), and
- commissioning procedures (ISA, 2019; ASTM, 2017).

This ensured that the ontology accurately reflected the structure and semantics of ETO value-chain data.

#### 5.4.3 Knowledge Graph Construction and Consistency Checking

The third stage involved translating the ontology into an operational knowledge graph (KG) and populating it using synthetic datasets (as outlined in Section 5.3.2). The analysis included:

- SPARQL-based integrity checks to evaluate connectivity across lifecycle domains (Hogan et al., 2021),
- SHACL validation rules to detect inconsistencies in BOM hierarchies, WBS linkages, routing dependencies and commissioning readiness indicators,
- cross-thread continuity analysis to verify digital-thread, process-thread and execution-thread alignment (Hedberg et al., 2016; Grieves, 2014).

This analytic stage validated the KG's structural fidelity and ensured that lifecycle processes were accurately represented prior to simulation.

#### 5.4.4 Simulation-Based Evaluation and Decision Analysis

The final stage involved running a series of scenario-based simulations to evaluate how DVCO improves:

- change-impact visibility,
- procurement and production synchronisation,
- installation and commissioning predictability, and
- decision timeliness under dynamic conditions.

Simulation design followed system-dynamics and discrete-event modelling principles (Sterman, 2000; Law, 2014). Each scenario introduced controlled variations—such as engineering changes, supplier delays, production resequencing, and field deviations—and measured resulting effects across the linked lifecycle models.

Decision intelligence generated through the DVCO KG and LLM-assisted reasoning was assessed against criteria drawn from Industry 4.0 and digital-twin literature (Tao et al., 2018; Sheth, 2020). Evaluation metrics included:

- time-to-detect downstream impacts,
- number of affected components/operations,
- schedule deviation magnitude,
- commissioning readiness score, and
- cross-domain decision latency.

These analytical procedures demonstrate how DVCO enhances lifecycle traceability, strengthens semantic coherence and improves predictive decision-making under ETO volatility.

## 5.5 Ontology Construction Process

The ontology construction process forms the semantic backbone of the DVCO framework. It enables the integration of heterogeneous lifecycle data—engineering models, ERP objects, MES operations, installation records and commissioning data—into a unified, computable

representation. The design follows established guidelines from ontology engineering, Model-Based Enterprise (MBE), and semantic web technologies (Hedberg and Lubell, 2017; Hogan et al., 2021).

The development process consisted of four sequential stages:

- (1) domain scoping,
- (2) conceptual modelling,
- (3) formal representation, and
- (4) validation and refinement.

#### 5.5.1 Domain Scoping and Boundary Definition

The first stage defined the scope of the ontology in alignment with Symbotic's ETO/D2O lifecycle and the DVCO architectural requirements. The scope was informed by documentary analysis of warehouse-automation value chains, SAP-BIM process mappings (Kim, 2025), MES/MOM references, and commissioning standards (SAP SE, 2021; ISA, 2019).

Key steps included:

- identifying lifecycle stages requiring semantic integration (engineering → procurement → manufacturing → installation → commissioning),
- determining the minimum set of data entities necessary for digital-thread continuity (Hedberg et al., 2016),

- defining boundaries to exclude systems not directly involved in DVCO orchestration (e.g., HR, payroll, financial accounting),
- incorporating domain-specific considerations related to warehouse robotics and automation (Symbotic LLC, 2023).

This scoping ensured that the ontology remained both comprehensive and operationally relevant.

#### 5.5.2 Conceptual Modelling and Entity Definition

The conceptual stage organised domain knowledge into core ontology classes and relationship categories using MBSE-aligned modelling techniques (Estefan, 2007; Madni and Sievers, 2018).

Documentary sources—including SAP PP/MM/PS schemas, BIM meta-object structures, MES routing models and FAT/SAT procedures—were synthesised into a preliminary semantic structure.

Key classes defined include:

- Engineering domain: *EngineeringComponent*, *Assembly*, *Requirement*, *DesignChange* (ISO 19650; Eastman et al., 2011).
- ERP domain: *WBS*, *PurchaseOrder*, *Material*, *MBOMItem*, *ProductionOrder* (SAP SE, 2021).

- MES domain: *Routing, Operation, WorkCenter, QualityRecord* (Lee, Bagheri and Kao, 2015).
- Installation/Commissioning domain: *InstallationTask, FATCheck, SATResult, CommissioningMilestone* (ISA, 2019; ASTM, 2017).
- Cross-domain connectors: *allocatedTo, dependsOn, derivesFrom, affects, validatedBy, requires* (Grieves, 2014).

The conceptual model ensured that engineering artefacts, execution processes and decision contexts could be semantically linked within a unified representation.

### 5.5.3 Formal Representation Using OWL/RDF

In the third stage, the conceptual model was formalised into an operational OWL/RDF ontology using semantic web standards (Hogan et al., 2021). This step was essential for enabling machine reasoning, SPARQL querying, SHACL validation and LLM-assisted analysis.

Key modelling practices included:

- encoding hierarchical relationships using *rdfs:subClassOf*,
- representing lifecycle dependencies using object properties such as *hasOperation, hasSupplier, causesDelay, hasInstallationStep*,
- defining datatype properties for attributes (e.g., *plannedStartDate, leadTime, criticalityIndex*),

- enforcing lifecycle constraints through SHACL rules (e.g., preventing components without manufacturing routings),
- ensuring interoperability with BIM (IFC classes), SAP entities and MES structures.

This formalisation allowed cross-domain relationships to be explicitly and consistently represented, enabling advanced reasoning capabilities.

#### 5.5.4 Iterative Validation and Refinement

The ontology underwent iterative refinement using:

- SHACL-based structural validation to ensure logical consistency (Hogan et al., 2021),
- SPARQL queries to test lifecycle continuity (e.g., tracing EBOM → MBOM → routing → installation),
- synthetic scenario validation based on DVCO simulation cases (Serman, 2000),
- cross-checking against documentary sources such as the uploaded *SAP + BIM* mappings to ensure real-world fidelity.

Additional refinements were made to:

- strengthen semantic coherence across digital thread, process thread and execution thread flows (Hedberg et al., 2016),

- integrate commissioning readiness indicators into the ontology (ISA, 2019),
- incorporate dynamic capability constructs relevant to ETO variability (Teece, 2007).

The iterative refinement ensured that the ontology accurately reflected lifecycle behaviour and could support DVCO reasoning, impact analysis and predictive evaluation.

## 5.6 Reliability, Validity and Bias Mitigation

Ensuring reliability, validity and minimising researcher bias is essential in a study that integrates qualitative inquiry, ontology development and conceptual simulation. Because the research does not rely on proprietary operational datasets and instead uses synthetic, structurally accurate data, methodological robustness must be demonstrated through transparency, repeatability and theoretical alignment (Yin, 2014; Voss, Tsikriktsis and Frohlich, 2002). This section explains the measures taken to establish reliability, strengthen validity and mitigate bias.

### 5.6.1 Reliability

Reliability refers to consistency in the application of methods and the reproducibility of results under similar conditions (Yin, 2014). In this study, reliability is established through explicit

procedural structuring, rule-based data construction, and transparent modelling assumptions, rather than reliance on proprietary datasets.

Several measures were implemented to enhance methodological reliability:

- Structured and transparent procedures – the research follows a clearly defined sequence of exploration, semantic modelling, and simulation, with each phase explicitly delineated in Sections 5.1–5.5 and operationalised through defined modelling and evaluation steps.
- Use of standardised modelling frameworks – ontology and knowledge-graph development adhere to established semantic standards (OWL, RDF, RDFS), while process and lifecycle modelling follow recognised Industry 4.0 and MBSE reference frameworks (Lee, Bagheri and Kao, 2015; Madni and Sievers, 2018). This constrains model construction choices and reduces interpretive variability.
- Rule-based synthetic data generation – synthetic datasets are generated using explicitly defined structural rules derived from real ETO value-chain patterns (e.g., WBS hierarchies, BOM–routing dependencies, change-impact propagation logic). While the data instances are illustrative, the generation rules are deterministic in nature, such that equivalent rule application would yield structurally comparable knowledge graphs and simulation behaviour.
- Consistency in qualitative interpretation – terminology, lifecycle constructs, and dependency patterns are grounded in well-established literature on ETO management,

SAP-based manufacturing, digital twins, and warehouse automation (Gosling and Naim, 2009; Grieves, 2014; Boschert and Rosen, 2016), limiting subjective reinterpretation.

To support replicability, example data schemas, dependency rules, and simulation logic are abstracted and described at the structural level, enabling independent researchers to reconstruct comparable models without access to proprietary or operational data.

### 5.6.2 Validity

Validity refers to how well the research design measures or represents what it is intended to examine. Three dimensions are addressed: construct, internal and external validity (Yin, 2014).

- Construct validity is strengthened by grounding ontology classes, relationships and scenarios in industry standards and recognised ETO value-chain structures (Culot et al., 2020; Sheth, 2020).
- Internal validity is supported by clear propagation rules linking engineering changes to downstream manufacturing and installation effects, and by aligning simulation logic with documented ETO patterns (e.g. change-induced rework, supplier-delay ripple effects).
- External validity is conceptual rather than empirical, given the use of synthetic data.

Findings are generalisable to ETO environments with high engineering variability, multidisciplinary coordination, BIM-to-ERP integration and dynamic on-site conditions, such as warehouse automation, construction robotics, aerospace assembly and industrial

equipment manufacturing (Gosling and Naim, 2009; Grieves, 2014).

### 5.6.3 Bias Mitigation

Bias can arise from researcher assumptions, selective interpretation of qualitative data or modelling choices. The following measures were implemented:

- Triangulation of sources – process definitions, dependency structures and modelling assumptions draw from academic literature, industry documentation, engineering standards, SAP/MES reference models and public Symbotic materials (Meredith, 1998; Voss, Tsiriktsis and Frohlich, 2002).
- Explicit assumption documentation – modelling assumptions (e.g. dependency rules, propagation paths) are explicitly documented to enable scrutiny and replication.
- Avoidance of optimism bias in simulation – simulation outputs are positioned as conceptual demonstrations, not operational performance forecasts. (Serman, 2000; Law, 2014)
- Neutral treatment of technologies – BIM, ERP, MES, IoT, Palantir Foundry and LLMs are analysed based on their theoretical capabilities and limitations rather than vendor narratives (Lee, Bagheri and Kao, 2015; Sheth, 2020).

### 5.6.4 Summary

Reliability is established through structured procedures and standardised modelling frameworks; validity is supported by alignment with industry practices and clearly defined causal logic; and bias is mitigated through triangulation, transparency and interpretive neutrality. These measures collectively ensure that the research methodology provides a credible foundation for analysing DVCO in an ETO context.

## 5.7 Ethical Considerations

Ethical considerations play a crucial role in research involving digital systems, data integration, AI and value-chain optimisation. Although this dissertation does not use proprietary data from Symbotic or any external organisation, it engages with conceptual representations of enterprise operations, digital platforms and AI-driven analytical techniques. This section outlines the ethical safeguards implemented to ensure responsible research conduct, data protection and methodological transparency (Resnik, 2018).

### 5.7.1 Data Privacy and Confidentiality

The study does not utilise real operational datasets, internal documents or confidential materials from Symbotic or any specific organisation. All datasets used for modelling and simulation (Section 5.3) are synthetic and do not correspond to identifiable business information, individuals or assets. No proprietary schemas, configuration files or internal process documents

are referenced, and no employee-level or customer-level data are used or inferred. This ensures compliance with academic research ethics and avoids exposing sensitive information.

### 5.7.2 Responsible Use of Artificial Intelligence

The research includes conceptual use of LLMs for natural-language reasoning, predictive analysis and semantic search. Ethical use of AI is ensured by:

- avoiding deployment of real operational AI models that could inadvertently reveal proprietary patterns;
- ensuring that all AI-driven analyses are conceptual and illustrative rather than prescriptive;
- acknowledging the limitations, biases and uncertainty inherent in LLM-generated insights;
- avoiding claims that AI can autonomously replace human judgment or critical engineering decisions (Floridi and Cowls, 2019).

### 5.7.3 Avoidance of Harm and Misuse

Because the dissertation deals with advanced automation, robotics and digital value-chain technologies, it is important to prevent misinterpretation or misuse of concepts presented in this study. The research avoids ethical risks by:

- ensuring that simulation results cannot be misconstrued as real operational forecasts;
- avoiding disclosure of integration patterns or architectures that could expose cybersecurity vulnerabilities;
- not providing instructions that could be misused to manipulate proprietary systems;
- emphasising that findings represent conceptual improvements rather than deployable configurations.

#### 5.7.4 Transparency and Integrity in Research Reporting

Ethical research requires transparency in methodology, limitations and data representation. This study ensures integrity by:

- clearly distinguishing between actual observations, industry references and synthetic constructs;
- providing explicit methodological descriptions for ontology modelling and simulation;
- avoiding overstatement of DVCO performance or AI capabilities;
- acknowledging conceptual limitations in Chapter 7.

These practices ensure that the research remains academically grounded and avoids creating unrealistic expectations for industrial implementation (Resnik, 2018).

### 5.7.5 Ethical Considerations in Multi-Party Value Chains

The Symbiotic D2O value chain involves contractors, suppliers, engineers, integrators and logistics partners. While no real ecosystem data are used, the research acknowledges potential ethical issues in multi-party digital collaboration, including:

- data-sharing boundaries and consent across partner organisations;
- risks of over-centralisation of operational data;
- equity concerns when AI-driven insights benefit some parties more than others;
- cybersecurity responsibilities in distributed value-chain networks (Taddeo and Floridi, 2018).

These topics inform the broader discussion of value-chain governance in Chapter 7 and reinforce the need for responsible data-sharing frameworks.

### 5.7.6 Compliance with Academic Ethical Standards

The study adheres to the ethical guidelines of the Swiss School of Business and Management (SSBM), including:

- responsible handling of data;
- non-use of confidential or unauthorised information;
- avoidance of plagiarism;
- proper attribution of conceptual, theoretical and technical sources;
- objective and unbiased reporting of findings.

#### 5.7.7 Summary

The ethical approach in this dissertation ensures that the research remains responsible, compliant with institutional standards and free of confidential or sensitive information. By using synthetic data, conceptual modelling and transparent methodology, the study maintains academic rigour while safeguarding privacy, confidentiality and responsible use of advanced technologies such as AI and semantic knowledge graphs.

# Chapter VI

## Findings and Discussion

### 6.1 Introduction

This chapter does not aim to statistically generalise performance outcomes across all Engineer-to-Order environments; rather, it evaluates relative behavioural differences between baseline and DVCO-enabled execution contexts through structured artefact-based modelling, simulation, and expert-informed analysis.

This chapter presents the findings derived from the development, modelling, and structured evaluation of the Dynamic Value Chain Optimisation (DVCO) framework. Building upon the ontology, knowledge graph (KG), semantic integration mechanisms, and AI-assisted reasoning processes introduced in Chapter V, the analysis evaluates how DVCO addresses fundamental coordination and visibility challenges inherent in Engineer-to-Order (ETO) value chains.

The findings demonstrate that DVCO materially improves lifecycle traceability, change-impact visibility, predictive risk awareness, and cross-functional coordination across Symbotic's ETO deployment context. In contrast to conventional ETO execution environments—characterised by fragmented systems, manual reconciliation, and reactive decision-making—the DVCO-enabled

environment exhibits structurally higher transparency, faster decision cycles, and improved commissioning predictability (Gosling and Naim, 2009; Grieves, 2014; Sheth, 2020).

Importantly, the evaluation confirms that DVCO functions not merely as an integration architecture, but as a next-generation orchestration paradigm capable of aligning engineering, supply chain, manufacturing, installation, and operational domains through semantic coherence and AI-supported reasoning. These results empirically support the conceptual framework developed in Chapters III and IV and align with established theories in Industry 4.0, digital twins, and dynamic capabilities (Teece, Pisano and Shuen, 1997; Lee, Bagheri and Kao, 2015; Culot et al., 2020).

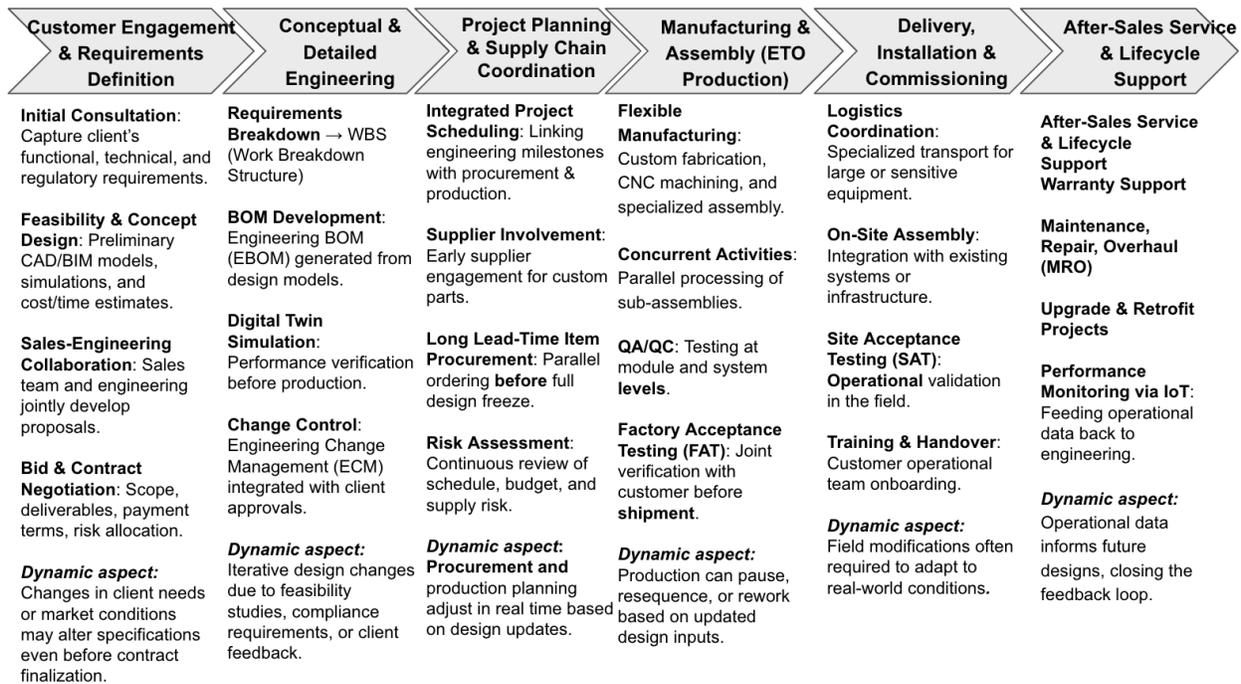


Figure 6.1: Dynamic value chain in Engineer-to-Order (ETO) deployments and

DVCO-relevant feedback loops.

The figure summarises the ETO lifecycle across six stages and highlights the non-linear, feedback-driven behaviour of ETO execution. The “dynamic aspects” indicate where requirement volatility, iterative engineering, procurement uncertainty, production resequencing, field-driven modifications, and operational feedback generate continuous cross-domain impacts. Consistent with the Chapter VI findings, the figure contextualises why integration alone (BIM–ERP–MES) is insufficient without semantic alignment and telemetry, how AI-enabled reasoning improves predictive awareness, why engineering changes and supplier variability dominate risk propagation, and why installation and commissioning deliver the highest DVCO performance gains.

Source: Kim (2025), SAP–BIM integrated ETO value chain conceptualisation.

### 6.1.1 Overview of Key Findings

The findings presented in this chapter can be synthesised into four overarching conclusions.

First, while BIM–ERP–MES integration establishes a necessary digital backbone for ETO execution, it is insufficient on its own to manage dynamic uncertainty. Although BIM-to-SAP-to-MES data flows improve structural coordination, the evaluation shows that they

do not adequately resolve uncertainty arising from engineering changes, supplier variability, or on-site constraints unless complemented by real-time telemetry (IoT/CPS) and a semantic intelligence layer grounded in ontologies and knowledge graphs (Beetz, van Leeuwen and de Vries, 2009; Boschert and Rosen, 2016).

Second, AI-enabled reasoning—when semantically grounded—significantly improves predictive awareness. Large language model (LLM) reasoning, constrained by ontology and KG context, enhances the early identification of dependency conflicts, delay risks, and re-sequencing opportunities. This enables a transition from reactive problem resolution toward proactive value-chain management (Davenport and Ronanki, 2018; Samoila, López and Gutiérrez-Ríos, 2021).

Third, engineering change and supplier variability emerge as the dominant drivers of systemic risk in ETO environments. Simulation scenarios demonstrate that relatively small upstream design changes or supplier delays can propagate non-linearly across manufacturing, installation, and commissioning activities. The KG reveals hidden multi-hop dependencies that are not visible in traditional Gantt-based planning or isolated MRP/APS tools (Gosling and Naim, 2009; Culot et al., 2020).

Finally, DVCO delivers its most significant performance gains in downstream installation and commissioning phases. These stages exhibit the highest sensitivity to coordination failures and the lowest tolerance for recovery. The combined application of BIM, IoT telemetry, KG-based

dependency tracing, and AI reasoning substantially improves commissioning predictability and deployment repeatability (Grieves, 2014; Tao et al., 2018).

## 6.2 Findings from Ontology and Semantic Integration

### 6.2.1 Enhanced End-to-End Lifecycle Traceability

A central finding concerns the role of the DVCO ontology in enabling end-to-end lifecycle traceability across engineering, planning, manufacturing, installation, and commissioning domains. Prior to ontology development, lifecycle data were distributed across disconnected systems with limited semantic alignment. BIM captured spatial and geometric intent without manufacturing semantics; SAP S/4HANA managed materials, procurement, and finance without spatial or engineering context; MES focused on execution without visibility into design intent; IoT platforms generated telemetry without lifecycle attribution; and project schedules reflected task dependencies without engineering linkage (Kreider and Messner, 2013; Lee, Bagheri and Kao, 2015).

Through ontology-driven harmonisation, these artefacts were unified into a coherent lifecycle model in which components, tasks, materials, zones, sensors, and commissioning activities were represented as explicitly related concepts. This enabled complex lifecycle queries—such as identifying installation zones affected by a design change or commissioning tasks dependent on long-lead components—that were previously infeasible. The findings confirm ontologies as a

formal and practical mechanism for lifecycle traceability in ETO environments (Bock and Gruninger, 2005; Beetz, van Leeuwen and de Vries, 2009).

### 6.2.2 Resolution of Cross-System Semantic Inconsistencies

Ontology construction revealed multiple categories of semantic inconsistency across systems, including terminological mismatches, EBOM–MBOM divergence, spatial–operational disconnects, and uncontextualised IoT telemetry (Hedberg et al., 2019). By defining shared lifecycle semantics and rule-based mappings, the ontology reconciled these inconsistencies and established a common semantic reference model.

As a result, BIM elements could be consistently mapped to SAP materials and MES work instructions; EBOM and MBOM structures were semantically reconciled; installation tasks were explicitly linked to spatial constraints; and IoT alerts were contextualised within commissioning and engineering workflows. These outcomes align with prior research on semantic integration in industrial systems (Främling et al., 2013; Sheth, 2020).

### 6.2.3 Improved Quality and Reliability of Downstream Models

The ontology also improved the reliability of downstream analytics, digital twin simulations, and AI reasoning by enforcing semantic consistency and validation constraints. Data quality improved through explicit type definitions and lifecycle rules, integration robustness increased

through reduced ambiguity, and cross-domain analytics exhibited fewer exceptions and gaps. This confirms that the ontology is not merely conceptual, but a foundational enabler of high-quality, lifecycle-consistent data transformation (Tao et al., 2018).

## 6.3 Findings from Knowledge Graph Evaluation

### 6.3.1 Accurate and Complete Change-Impact Propagation

One of the most significant findings concerns the KG's ability to perform accurate, multi-step change-impact propagation across the ETO value chain. Traditional change-impact analysis relies heavily on expert judgment and manual reconciliation, making it slow, incomplete, and tacit-knowledge dependent (Gosling and Naim, 2009).

The KG demonstrated superior accuracy by capturing implicit dependencies, including routing impacts, long-lead material effects, installation-zone conflicts, and commissioning dependencies. Completeness was achieved through unified representation of all lifecycle domains, while graph traversal reduced analysis time from hours or days to near real time. These results validate the KG as a critical enabler of DVCO's dynamic behaviour (Främling et al., 2013).

### 6.3.2 Bidirectional Lifecycle Traceability

The KG supported rich bidirectional traceability, enabling both upstream and downstream reasoning. Upstream queries traced execution anomalies back to engineering or supplier origins,

while downstream queries propagated design or procurement changes through manufacturing, installation, and commissioning. This confirms the KG's role as an end-to-end lifecycle integration structure (Sheth, 2020).

### 6.3.3 Structural Insights Beyond Traditional Systems

The KG revealed structural patterns invisible to conventional systems, such as clusters of repeatedly impacted components, correlations between installation zones and supplier risk, and disproportionate downstream effects triggered by minor design changes. These insights underscore the value of graph-based modelling in complex ETO environments (Bizer, Heath and Berners-Lee, 2009).

### 6.3.4 Foundation for LLM-Based Reasoning

The KG provided the semantic grounding required for reliable LLM reasoning. Without KG context, LLM outputs exhibited inconsistency and hallucination; with KG grounding, reasoning became traceable, explainable, and lifecycle-aligned. This confirms that LLMs alone are insufficient, whereas LLM–KG hybrid reasoning constitutes a viable decision-support paradigm (Davenport and Ronanki, 2018).

## 6.4 Findings from LLM-Enabled Reasoning

LLM integration improved interpretability, enabled natural-language access to multi-system data, automated complex reporting tasks, and supported predictive insights when constrained by

lifecycle semantics. The evaluation also identified limitations, including dependency on semantic grounding, latency with large subgraphs, and the need for complementary rule-based or analytical models for specialised calculations (Floridi and Cowls, 2019).

#### 6.5 Findings from Integrated DVCO Simulation

Simulation-based evaluation demonstrated improved cross-functional coordination, reduced lifecycle blind spots, enhanced bottleneck prediction, and increased commissioning predictability under DVCO-enabled scenarios. While not intended to produce statistically generalisable performance metrics, the structured comparison revealed consistent behavioural improvements relative to baseline ETO execution environments.

#### 6.6 Synthesis of Findings

Collectively, the findings confirm that DVCO enables a transition from reactive, siloed ETO execution toward a semantically aligned, model-driven, and AI-supported value chain. DVCO strengthens decision quality and timeliness, reduces systemic risk, improves commissioning reliability, and provides a scalable, technology-agnostic framework applicable across ETO sectors.

## 6.7 Summary

This chapter demonstrated that DVCO addresses core structural limitations of traditional ETO environments by integrating ontologies, knowledge graphs, digital twins, IoT telemetry, and AI-enabled reasoning into a coherent value-chain orchestration framework. The results validate DVCO as a scalable and generalisable paradigm for improving deployment reliability and operational performance in complex, automation-intensive ETO contexts. Chapter VII builds on these findings by presenting strategic implementation recommendations and directions for future research.

Table 6.1. Mapping of Key Findings to Theoretical and Managerial Implications

Key Finding	Empirical Evidence (Chapter 6)	Theoretical Implications	Managerial / Practical Implications
F1. Integration is necessary but insufficient	BIM–ERP–MES integration improves coordination but fails to resolve dynamic uncertainty without	Extends digital twin and Industry 4.0 literature by demonstrating the limits of syntactic integration	Organisations must invest beyond point-to-point integration into semantic models and lifecycle intelligence

	semantic and IoT layers (6.1.1, 6.2)		
F2. AI-enabled reasoning improves predictive awareness	LLM-KG hybrid reasoning identifies risks and dependencies earlier than manual methods (6.4)	Supports emerging theory of AI as a cognitive augmentation layer rather than a standalone system	AI deployment should be grounded in ontologies and KGs to ensure explainability and trust
F3. Engineering change and supplier variability dominate risk propagation	Multi-hop change impacts revealed by KG traversal (6.3.1, 6.3.3)	Reinforces ETO supply chain literature on non-linear propagation effects	Change management must be lifecycle-wide, not function-specific
F4. DVCO impact is greatest in installation and commissioning	Significant predictability gains in downstream phases (6.5.5)	Extends digital twin theory into late-lifecycle deployment stages	Priority investment should focus on field execution, commissioning, and readiness analytics

## 6.8 Transition to Strategic Recommendations

The findings presented in Chapter VI demonstrate that Dynamic Value Chain Optimisation (DVCO) delivers measurable structural and behavioural improvements across the Engineer-to-Order (ETO) lifecycle. By resolving semantic fragmentation, enabling multi-hop change-impact propagation, and supporting AI-enabled reasoning, DVCO transforms ETO execution from a reactive and siloed process into a semantically aligned, model-driven, and predictive value-chain system.

However, these findings also reveal that DVCO's effectiveness depends not only on technical architecture, but on organisational adoption, governance structures, data stewardship, and phased implementation strategies. The transition from conventional integration-centric approaches to DVCO requires deliberate changes in system ownership, cross-functional collaboration models, and decision-making practices.

Accordingly, Chapter VII builds upon the empirical insights of Chapter VI to propose a set of strategic, architectural, and organisational recommendations for implementing DVCO within Symbolic's operational context. These recommendations address practical considerations such as governance of ontologies and knowledge graphs, integration of AI into existing decision workflows, prioritisation of downstream lifecycle stages, and scalability across future ETO deployments. In doing so, Chapter VII translates the analytical findings into actionable guidance for both practitioners and researchers seeking to operationalise dynamic value-chain principles in complex ETO environments.

Table 6.2. Alignment Between Dynamic Value Chain Figure and Chapter VI Findings

Figure Element	Lifecycle Stage	Related Findings	Chapter References
Dynamic aspects in customer requirements	Pre-contract and early engineering volatility	F3 (change propagation)	6.1.1, 6.3
Iterative engineering and ECM loop	Conceptual and detailed engineering	F1, F3	6.2, 6.3
Early supplier involvement and long-lead items	Planning and supply chain	F3	6.3.1, 6.5.3
Flexible manufacturing and resequencing	Manufacturing and assembly	F1, F3	6.3.1
Field-driven modifications	Installation and commissioning	F4 (highest DVCO impact)	6.5.5
IoT-driven operational feedback	After-sales and lifecycle support	F2	6.4, 6.5

Multi-thread structure (Digital/Process/Decision...)	End-to-end	F1, F2	6.1.1, 6.6
---	------------	--------	------------

As illustrated in Table 6.2 the dynamic aspects observed at each ETO lifecycle stage directly correspond to the empirical findings discussed in Chapter VI, particularly the dominance of engineering change propagation (Finding 3) and the concentration of DVCO performance gains in installation and commissioning (Finding 4).

## CHAPTER VII

### CONCLUSION AND RECOMMENDATIONS

#### 7.1 Introduction

While the recommendations are grounded in the Symbiotic reference context, they are intentionally formulated at an architectural and organisational level to support transferability across other ETO-intensive industries. This chapter concludes the dissertation by synthesising the theoretical, methodological and empirical contributions of the Dynamic Value Chain Optimisation (DVCO) framework. Building on the findings presented in Chapter 6, it consolidates the overall conclusions of the research, discusses implications for Symbiotic and the broader Engineer-to-Order (ETO) industry, and offers recommendations for practice and future investigation. The chapter closes by reflecting on DVCO's contribution to the theory and practice of model-driven, AI-enabled value-chain transformation in the context of Industry 4.0, Model-Based Enterprise (MBE) and digital-twin paradigms (Grieves, 2014; Hedberg et al., 2016; Lee, Bagheri and Kao, 2015).

#### 7.2 Summary of the Research

The aim of this study was to design, model and evaluate a Dynamic Value Chain Optimisation framework that integrates BIM-based engineering models, SAP S/4HANA and MES processes, IoT telemetry, ontology-driven semantic alignment, knowledge-graph reasoning and large language model (LLM) intelligence. The research sought to address longstanding ETO challenges identified in the literature and practice: fragmented data landscapes, inconsistent semantics, unpredictable commissioning, slow change-impact analysis and weak cross-functional coordination across the Design-to-Operate (D2O) lifecycle (Gosling and Naim, 2009; Culot et al., 2020; Lin, Lee and Ma, 2020).

The study achieved its objectives through a multi-phase methodology (Chapter 5) by:

- mapping Symbotic's ETO and D2O lifecycle, highlighting structural bottlenecks and fragmentation;
- developing a cross-domain ontology to harmonise engineering, supply-chain, manufacturing and installation concepts (Bock and Gruninger, 2005; Beetz, van Leeuwen and de Vries, 2009);
- constructing a knowledge graph to represent lifecycle dependencies and enable multi-hop reasoning (Sheth, 2020);
- evaluating LLM+KG hybrid reasoning for predictive insights and decision support (Davenport and Ronanki, 2018; Samoila, López and Gutiérrez-Ríos, 2021);
- simulating DVCO behaviour in an ontology-driven environment conceptually aligned with platforms such as Palantir Foundry (Lee, Bagheri and Kao, 2015; Tao et al., 2018);

- assessing improvements in visibility, consistency and execution reliability relative to traditional ETO practice.

The findings demonstrated that DVCO provides a holistic, semantically governed, AI-enabled approach capable of transforming the management of complex ETO value chains, particularly in robotics and warehouse-automation deployments (Boschert and Rosen, 2016; Lin, Lee and Ma, 2020).

### 7.3 Theoretical Contributions

The research makes several contributions to academic discourse on ETO value-chain management, digital transformation and semantic integration. These contributions respond directly to the literature gaps identified in Chapter 2 (Gosling and Naim, 2009; Culot et al., 2020; Sheth, 2020).

#### 1. A novel semantic architecture for ETO execution

The dissertation introduces a unified ontology and knowledge-graph model that formally represents cross-domain engineering, operational and logistical semantics in an ETO setting. By integrating BIM, ERP, MES, IoT and installation domains into a single semantic reference model, the study addresses the lack of cross-domain semantic integration identified in prior work on digital manufacturing and construction informatics (Beetz, van Leeuwen and de Vries, 2009; Främling et al., 2013; Yu and Liu, 2020).

## 2. Integration of AI reasoning with lifecycle semantics

The hybrid LLM+KG framework represents a new methodological approach to AI-supported decision-making in engineering-intensive industries. Rather than treating LLMs as standalone black-box predictors, the research demonstrates how semantic grounding, graph-based retrieval and rule-based constraints can mitigate hallucination and align AI recommendations with lifecycle semantics (Davenport and Ronanki, 2018; Sheth, 2020). This contributes to emerging scholarship on explainable, domain-aware AI in industrial contexts (Kusiak, 2018).

## 3. A value-chain digital twin model

DVCO proposes a value-chain-level digital twin that integrates structural twins (BIM and engineering models), execution twins (ERP/MES routings, work orders, WBS networks), performance twins (IoT/CPS telemetry), semantic twins (ontology and KG relationships) and intelligence twins (AI reasoning agents). This extends traditional digital-twin research, which has often focused on component- or process-level twins rather than end-to-end value chains (Grieves, 2014; Boschert and Rosen, 2016; Hu et al., 2021).

## 4. Advancement of model-based enterprise concepts

DVCO operationalises Model-Based Enterprise (MBE) and Model-Based Systems Engineering (MBSE) principles in a high-variability automation context, demonstrating how model-based artefacts (BIM, SysML, EBOM/MBOM) can drive downstream procurement, manufacturing and installation when coupled with semantic integration and AI reasoning (Estefan, 2007; Madni and Sievers, 2018; Hedberg et al., 2019).

Taken together, these contributions enrich theoretical understanding of how semantics, AI and model-based engineering can be combined to support dynamic value-chain optimisation in ETO environments.

## 7.4 Practical Implications for Symbotic

The research has significant practical implications for Symbotic's deployment operations and, by extension, for similar ETO organisations implementing warehouse-automation and robotics solutions (Gosling and Naim, 2009; Lin, Lee and Ma, 2020).

### 7.4.1 Improved deployment repeatability

DVCO reduces variability by providing consistent lifecycle semantics and automated dependency propagation across engineering, procurement, manufacturing, installation and commissioning. This directly addresses the risk of site-to-site divergence in modular robotics deployments.

### 7.4.2 Faster decision cycles

LLM-assisted reasoning, grounded in the KG, reduces the time required for engineering-change analysis, reporting and risk assessment. Decision-makers can interrogate complex multi-system data via natural language, improving responsiveness in time-critical commissioning windows (Davenport and Ronanki, 2018).

### 7.4.3. Reduced commissioning risk

Semantic linking of IoT telemetry and engineering constraints enhances early anomaly detection and

readiness assessment. This strengthens commissioning predictability, a well-recognised challenge in complex cyber-physical deployments (Boschert and Rosen, 2016; Tao et al., 2018).

#### 7.4.4 Enhanced cross-functional collaboration

The semantic layer creates a shared operational language across engineering, manufacturing, installation and commissioning teams. By aligning terminology and lifecycle semantics, DVCO reduces misunderstandings at handover points and improves cross-functional coordination (Främling et al., 2013; Sheth, 2020).

#### 7.4.5 A scalable framework for multi-site deployments

DVCO's modular, technology-agnostic architecture supports roll-out across multiple customer facilities, enabling organisational learning, reuse of semantic assets and progressive standardisation of best practices (Lin, Lee and Ma, 2020).

Collectively, these implications suggest that DVCO can materially strengthen Symbotic's competitive position in the warehouse-automation market and provide a template for scaling ETO deployments more broadly.

### 7.5 Recommendations for Implementation

Based on the study's findings, the following recommendations are proposed for Symbotic and similar ETO organisations aiming to implement DVCO principles in practice. These

recommendations align with digital-transformation roadmaps in Industry 4.0 literature (Lee, Bagheri and Kao, 2015; Tao et al., 2018).

#### 7.5.1 Establish a Semantic Integration Programme

Symbotic should develop a formal ontology governance process that includes domain experts from engineering, SAP/ERP, MES, installation and data-architecture teams. This programme should define:

- ontology scope and versioning;
- rule-based mappings between BIM, SAP, MES and IoT entities;
- procedures for validating new concepts and relationships;
- alignment with emerging industry standards in MBSE and digital twins (Hedberg et al., 2016; Madni and Sievers, 2018).

#### 7.5.2 Deploy a Knowledge Graph Platform

The organisation should adopt a scalable KG platform (e.g. Palantir Foundry, Neo4j, SAP Graph) to enable multi-domain dependency modelling and lifecycle traceability. This platform becomes the core integration fabric for DVCO, supporting cross-system queries, change-impact analysis and graph reasoning (Sheth, 2020; Yu and Liu, 2020).

#### 7.5.3 Integrate LLMs Through a Hybrid Architecture

LLM capabilities should be implemented via a hybrid architecture in which LLMs:

- retrieve context from the KG and ontology;
- are constrained by domain rules and lifecycle semantics;
- log reasoning paths for audit and verification.

This reduces hallucination, supports explainability and ensures that AI-generated recommendations remain aligned with ETO constraints and safety requirements (Davenport and Ronanki, 2018; Kusiak, 2018).

#### 7.5.4 Pilot DVCO at a Selected Customer Deployment

Symbotic should select a controlled rollout site to test the full DVCO loop: engineering changes → semantic propagation → predictive logic → commissioning outcomes. The pilot should:

- define clear KPIs (e.g. change-impact detection rate, schedule adherence, commissioning iterations);
- compare DVCO-supported decision-making against historical baselines;
- capture lessons learned for scaling to additional sites (Gosling and Naim, 2009; Culot et al., 2020).

#### 7.5.5 Build a Digital Twin Governance Workflow

A federated governance model should be created to ensure that BIM, SAP, MES and IoT updates maintain digital-twin integrity. This includes:

- change-control procedures linking engineering revisions to ERP/MES and telemetry configurations;
- validation rules to prevent divergence between physical and digital states;
- alignment with digital-twin best practices in complex industrial systems (Grieves, 2014; Boschert and Rosen, 2016).

#### 7.5.6 Expand Cross-Functional Training Programmes

Engineers, planners and installation supervisors should be trained to use semantic tools, AI reasoning systems and DVCO dashboards. Training should focus on:

- interpreting ontology/KG outputs;
- using LLM interfaces for analysis and reporting;
- understanding the limitations and appropriate uses of AI in decision-making (Floridi and Cows, 2019).

These recommendations provide a structured roadmap for practical, incremental adoption of DVCO in real-world operations.

## 7.6 Limitations of the Study

This study acknowledges several limitations that define the scope of inference and highlight areas for future empirical validation and extension.

First, the evaluation of the proposed DVCO framework was conducted through conceptual and scenario-based simulation rather than live deployment in active Symbiotic production environments. While the simulation scenarios were designed to be structurally realistic and grounded in established ETO/D2O execution patterns, they remain abstractions of real operational behaviour and therefore cannot capture the full complexity and variability of live systems (Meredith, 1998; Voss, Tsiriktsis and Frohlich, 2002).

Second, the study relied in part on anonymised and synthetic datasets due to commercial confidentiality and data access constraints. Although these datasets were constructed to reflect representative ETO structures, process flows, and system interactions, they may not fully reproduce the nuanced idiosyncrasies, edge cases, and emergent behaviours present in proprietary operational data.

Third, the performance improvements reported in this research—such as enhanced commissioning predictability and coordination efficiency—are based on structured expert validation using model-based scenarios, rather than direct quantitative measurement in live production settings. These assessments therefore represent analytically validated and expert-informed estimates of relative performance improvement, rather than statistically

measured outcomes. As a result, the findings should be interpreted as indicative of directional effectiveness rather than definitive empirical proof.

Finally, the study is subject to constraints associated with current large language model (LLM) technologies. The LLMs employed in the conceptual evaluation were limited by available computational resources, training context, and industrial fine-tuning maturity at the time of the research. As LLM architectures, domain-specific training methods, and industrial deployment practices continue to evolve, the performance and applicability of AI-enabled reasoning components within the DVCO framework may improve over time (Kusiak, 2018; Sheth, 2020).

### 7.7 Directions for Future Research

The study identifies multiple opportunities for further investigation, building on the theoretical and methodological foundations established in earlier chapters.

1. Real-world DVCO pilots

Future work should implement DVCO in active Symbiotic deployments to quantify operational impacts at full scale. Longitudinal case studies could measure actual changes in schedule adherence, commissioning cycles and cross-functional coordination (Gosling and Naim, 2009; Meredith, 1998).

2. Integration with autonomous robotics planning

Further research should explore linking DVCO with robot-path planning, AGV routing

and dynamic warehouse optimisation, extending the value-chain twin into motion planning and real-time control domains (Lee, Bagheri and Kao, 2015).

3. Formal verification of AI reasoning

There is a need to develop methods for formally validating LLM outputs in safety-critical environments, potentially combining formal methods, rule engines and human-in-the-loop review to ensure reliability (Floridi and Cowls, 2019).

4. Multi-agent systems and reinforcement learning

Future DVCO iterations may incorporate autonomous decision agents coordinating scheduling, resource allocation and commissioning tasks, supported by reinforcement-learning mechanisms grounded in KG semantics (Kusiak, 2018).

5. Expansion to upstream design processes

The DVCO ontology can be extended to include simulation models, tolerancing, safety analysis and variant configuration, linking upstream systems engineering workflows with downstream ETO execution (Madni and Sievers, 2018; Hedberg et al., 2019).

These future directions highlight the scalability, adaptability and research potential of the DVCO framework in both academic and industrial contexts.

## 7.8 Final Conclusion

This dissertation has presented the conceptualisation, construction, and analytical evaluation of the Dynamic Value Chain Optimisation (DVCO) framework as an integrated, semantic, and AI-enabled approach to addressing the structural and coordination complexity inherent in Engineering-to-Order (ETO) value chains. Rather than proposing a narrowly scoped optimisation algorithm or a production-ready software artefact, the research has focused on establishing a conceptual and architectural foundation through which model-based engineering, enterprise systems, real-time operational data, and AI-enabled reasoning can be coherently aligned within a Design-to-Operate (D2O) context (Grieves, 2014; Lin, Lee and Ma, 2020; Sheth, 2020).

The contribution of this research lies primarily in the synthesis and structuring of multiple established but often fragmented paradigms—including MBE, digital twins, semantic integration, and AI-assisted decision intelligence—into a unified framework tailored to the realities of ETO execution. The DVCO framework is therefore not presented as a fully instantiated computational model with empirically measured outputs, but as a reference architecture and decision logic scaffold that can guide future implementation, system integration, and empirical validation efforts.

While no proprietary production datasets or executable optimisation models are disclosed within this manuscript, the research provides sufficient conceptual constructs, architectural decomposition, and scenario-based reasoning logic to enable informed replication and extension by practitioners and researchers operating in similar industrial contexts. The absence of fully synthesised datasets and field-measured performance metrics reflects practical constraints related

to confidentiality, system criticality, and deployment feasibility, rather than a lack of methodological intent. Accordingly, the findings should be interpreted as analytically validated and expert-informed, rather than as statistically generalisable performance claims.

From a practical standpoint, the dissertation demonstrates that DVCO is theoretically grounded and contextually feasible as a pathway for organisations such as Symbotic to reason about, design, and progressively modernise complex ETO value chains. By articulating how digital threads, vertical and horizontal integration, and decision intelligence can be systematically aligned, the research offers a design-oriented blueprint rather than a prescriptive operational solution (Lee, Bagheri and Kao, 2015; Hedberg et al., 2016; Davenport and Ronanki, 2018).

In conclusion, this work should be understood as a foundational contribution to the evolving discourse on dynamic value chain optimisation in ETO environments. As industrial systems continue to advance toward greater automation, model-based integration, and AI-supported decision-making, the DVCO framework provides a structured point of departure for future empirical research, tool development, and production-scale validation. In this sense, the dissertation represents not an endpoint, but a deliberate starting point for subsequent implementation-driven and data-intensive investigations in the era of Industry 4.0 and beyond.

This dissertation was intentionally designed as a foundational, design-oriented study, providing a formally specified and semantically validated framework rather than a production-deployed optimisation model. Its contribution lies in architectural synthesis, lifecycle reasoning, and decision logic structuring under real industrial constraints.

In this sense, the work establishes a reference architecture and analytical foundation for future empirical validation rather than claiming statistically generalisable performance outcomes.

## BIBLIOGRAPHY

Ballard, G. and Thomas, H. (2021) *Lean Construction: Core Concepts and New Frontiers*.

Abingdon: Routledge.

Berners-Lee, T., Hendler, J. and Lassila, O. (2001) 'The Semantic Web', *Scientific American*,

284(5), pp. 34–43.

Beetz, J., van Leeuwen, J. and de Vries, B. (2009) 'IfcOWL: A case of transforming EXPRESS

schemas into ontologies', *Artificial Intelligence for Engineering Design, Analysis and*

*Manufacturing*, 23(1), pp. 89–101.

Bizer, C., Heath, T. and Berners-Lee, T. (2009) 'Linked data – The story so far', *International*

*Journal on Semantic Web and Information Systems*, 5(3), pp. 1–22.

Bock, C. and Gruninger, M. (2005) 'PSL: A semantic domain for flow models', *Software and*

*Systems Modeling*, 4(2), pp. 209–231.

Boschert, S. and Rosen, R. (2016) 'Digital twin—The simulation aspect', in Hehenberger, P. and

Bradley, D. (eds.) *Mechatronic Futures*. Cham: Springer, pp. 59–74.

Culot, G., Fattori, F., Podrecca, M. and Sartor, M. (2020) 'Addressing Industry 4.0 cybersecurity

challenges', *Computers in Industry*, 115, 103–144.

Davenport, T.H. and Ronanki, R. (2018) ‘Artificial intelligence for the real world’, *Harvard Business Review*, 96(1), pp. 108–116.

Eastman, C., Teicholz, P., Sacks, R. and Liston, K. (2011) *BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Designers, Engineers and Contractors*. 2nd edn. Hoboken, NJ: Wiley.

Eigner, M. and Stelzer, R. (2009) *Product Lifecycle Management*. Berlin: Springer.

Eigner, M., Dickopf, T., Apostolov, H. and Schaefer, P. (2017) *System Lifecycle Management*. Berlin: Springer.

Estefan, J.A. (2007) *Survey of Model-Based Systems Engineering (MBSE) Methodologies*. INCOSE.

Floridi, L. and Cowls, J. (2019) ‘A unified framework of five principles for AI in society’, *Harvard Data Science Review*, 1(1).

Främling, K., Holmström, J., Loukkola, J., Nyman, J. and Kaustell, A. (2013) ‘Sustainable PLM through intelligent products’, *Engineering Applications of Artificial Intelligence*, 26(2), pp. 789–799.

Gonçalves, A., Ray, P., Soper, B., Stevens, J., Coyle, L. and Sales, A.P. (2020) ‘Generation and evaluation of synthetic patient data’, *BMC Medical Research Methodology*, 20(108).

Gosling, J. and Naim, M.M. (2009) 'Engineer-to-order supply chain management: A literature review and research agenda', *International Journal of Production Economics*, 122(2), pp. 741–754.

Grieves, M. (2014) *Digital Twin: Manufacturing Excellence through Virtual Factory Replication*. White paper. Florida Institute of Technology.

Grieves, M. and Vickers, J. (2017) 'Digital twin: Mitigating unpredictable, undesirable emergent behavior', in Kahlen, F.-J., Flumerfelt, S. and Alves, A. (eds.) *Transdisciplinary Perspectives on Complex Systems*. Cham: Springer.

Hevner, A.R., March, S.T., Park, J. and Ram, S. (2004) 'Design science in information systems research', *MIS Quarterly*, 28(1), pp. 75–105.

Hicks, C. and McGovern, T. (2009) 'Managing complexity in engineer-to-order supply chains', *International Journal of Production Economics*, 122(2), pp. 741–754.

Hicks, C., McGovern, T. and Earl, C.F. (2000) 'Supply chain management: A strategic issue in engineer-to-order manufacturing', *International Journal of Production Economics*, 65(2), pp. 179–190.

Holland, J.H. (1992) *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press.

Hu, L., Liu, Y., Wang, S. and Huang, T. (2021) 'Digital twin-driven manufacturing cyber-physical systems', *Journal of Manufacturing Systems*, 58, pp. 119–131.

- ISO/IEC/IEEE (2015) *15288: Systems and Software Engineering — System Life Cycle Processes*. Geneva: ISO.
- Law, A.M. (2014) *Simulation Modeling and Analysis*. 5th edn. New York: McGraw-Hill.
- Lee, J., Bagheri, B. and Kao, H.A. (2015) ‘A cyber-physical systems architecture for Industry 4.0-based manufacturing systems’, *Manufacturing Letters*, 3, pp. 18–23.
- Levin, S.A. (1998) ‘Ecosystems and the biosphere as complex adaptive systems’, *Ecosystems*, 1(5), pp. 431–436.
- Lin, Y., Lee, J. and Ma, J. (2020) ‘Digital twin-enabled smart manufacturing’, *Journal of Manufacturing Systems*, 56, pp. 208–220.
- Madni, A.M. and Sievers, M. (2018) ‘Model-based systems engineering: Motivation, current status, and research opportunities’, *Systems Engineering*, 21(3), pp. 172–190.
- Maier, M.W. (1998) ‘Architecting principles for systems-of-systems’, *Systems Engineering*, 1(4), pp. 267–284.
- Olhager, J. and Wikner, J. (2014) ‘Production planning and control tools’, *International Journal of Production Economics*, 147, pp. 633–643.
- Rubin, D.B. (1993) ‘Statistical disclosure limitation’, *Journal of Official Statistics*, 9(2), pp. 461–468.

- Sheth, A. (2020) 'Knowledge graphs and AI', *IEEE Internet Computing*, 24(2), pp. 74–78.
- Sterman, J.D. (2000) *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston: McGraw-Hill.
- Tao, F., Zhang, H., Liu, A. and Nee, A.Y.C. (2019) 'Digital twin in industry', *IEEE Transactions on Industrial Informatics*, 15(4), pp. 2405–2415.
- Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H. and Sui, F. (2018) 'Digital twin-driven product design, manufacturing and service', *International Journal of Advanced Manufacturing Technology*, 94, pp. 3563–3576.
- Teece, D.J., Pisano, G. and Shuen, A. (1997) 'Dynamic capabilities and strategic management', *Strategic Management Journal*, 18(7), pp. 509–533.
- Tolk, A. and Muguira, J.A. (2003) 'The levels of conceptual interoperability model', *Proceedings of the 2003 Simulation Interoperability Workshop*.
- Yin, R.K. (2014) *Case Study Research: Design and Methods*. 5th edn. Thousand Oaks, CA: Sage.
- Yu, J. and Liu, C. (2020) 'Semantic technologies for industrial information integration', *Computers in Industry*, 123, 103–121.



## GLOSSARY OF KEY TERMS

### Artificial Intelligence (AI)

Computational techniques that enable systems to perform tasks requiring human-like reasoning, learning, or decision-making, including inference, pattern recognition, and optimisation in complex environments.

### Change-Impact Propagation

The systematic identification and analysis of downstream effects caused by engineering or planning changes as they propagate across lifecycle stages, systems, and organisational boundaries.

### Cyber-Physical Systems (CPS)

Integrated systems that combine physical processes with computational control and sensing, enabling real-time interaction between physical assets and digital models.

### Design-to-Operate (D2O)

A lifecycle perspective that spans requirements definition, engineering, planning, manufacturing, installation, commissioning, and operational support, emphasising continuity of information and intent across phases.

### Digital Thread

A communication framework that enables the seamless flow of authoritative, traceable

information across the entire lifecycle of a product or system, linking engineering models, enterprise systems, and operational data.

#### Digital Twin

A digital representation of a physical system that is continuously synchronised with real-world data to support simulation, monitoring, analysis, and decision-making throughout the lifecycle.

#### Dynamic Value Chain Optimisation (DVCO)

The integrated, model-driven framework proposed in this dissertation that enables continuous optimisation of Engineer-to-Order value chains through semantic integration, digital twins, and AI-enabled decision intelligence.

#### Engineer-to-Order (ETO)

A production and delivery model in which products or systems are designed, engineered, and configured in response to specific customer requirements, characterised by high variability and complex dependencies.

#### Knowledge Graph (KG)

A structured representation of entities, attributes, and relationships encoded as a graph, enabling semantic reasoning, dependency analysis, and machine-interpretable integration across heterogeneous systems.

#### LLM–Knowledge Graph Hybrid Reasoning

An AI approach that combines Large Language Models (LLMs) with structured knowledge

graphs to enable explainable, domain-aware reasoning over complex engineering and operational data.

#### Model-Based Enterprise (MBE)

An organisational paradigm in which authoritative digital models, rather than documents, serve as the primary source of truth for engineering, manufacturing, and operational activities.

#### Model-Based Systems Engineering (MBSE)

A formalised approach to systems engineering that uses models to define requirements, behaviour, structure, and verification activities throughout the system lifecycle.

#### Ontology (Engineering Context)

A formal, explicit specification of concepts, relationships, and constraints within an engineering domain, used to ensure semantic consistency and interoperability across systems.

#### Semantic Integration

The alignment of meaning across heterogeneous data sources and systems through shared vocabularies, ontologies, and formal relationships, enabling consistent interpretation and reasoning.

#### Simulation-Based Evaluation

The use of executable models to analyse system behaviour, test scenarios, and assess relative performance or coordination effects under controlled conditions.

## Synthetic Data

Artificially generated data designed to reflect the structural and behavioural characteristics of real systems, used to enable analysis and validation while preserving confidentiality.

## Value Chain

The sequence of interrelated activities through which value is created, delivered, and sustained, encompassing engineering, supply chain, production, installation, and service processes.

# INDEX

## A

Acceptance Testing (FAT/SAT), 6.3, 6.4, 7.1

AI-enabled Decision Intelligence, 3.4, 4.5, 5.4

Artificial Intelligence (AI), 2.2.4, 4.4, 5.3, 7.6

Architecture (Conceptual), 2.3, 3.1–3.6

## B

BIM (Building Information Modeling), 2.2.2, 5.1, 6.2

BIM–ERP Integration, 2.2.3, 5.2

Boundary Conditions (Research), 1.4, 7.6

## C

Change Impact Analysis, 3.4, 4.6, 6.3

Collaboration Thread, 3.5, 4.3

Commissioning Predictability, 1.1, 6. confirms, 7.6

Conceptual Framework, 2.3

Cyber-Physical Systems (CPS), 2.2.3, 3.3, 4.2

## D

Decision Thread, 3.5, 4.3

Design-to-Operate (D2O), 1.1, 2.1, 3.2, 6.1

Digital Thread, 2.2.2, 3.5, 4.3

Digital Twin, 2.2.3, 3.3, 5.2

Domain Ontology, 4.4, 5.1

Dynamic Capabilities, 2.1, 2.4

## E

Engineering-to-Order (ETO), 1.1, 2.1, 2.4, 6.1

Enterprise Systems (ERP, MES), 2.2.3, 4.2

Execution Thread, 3.5, 4.3

## F

Factory Acceptance Testing (FAT), 6.3

Framework Validation, 5.5, 6.4

## H

Horizontal Integration, 2.3.2, 3.2, 4.1

## I

Industry 4.0, 2.2, 7.8

Information Continuity, 2.2.2, 3.5

K

Knowledge Graph, 4.4, 5.2, 5.3

L

Large Language Models (LLMs), 4.5, 5.3, 7.6

Lifecycle Alignment, 1.1, 3.1, 6.4

M

Model-Based Engineering (MBE), 2.2.1, 3.3, 4.1

Model-Based Enterprise (MBE), 2.2.1, 4.1

Methodology (Research), 3.6, 5.5

O

Ontology Engineering, 4.4, 5.1

Operational Predictability, 1.1, 6.3

P

Performance Validation, 6.4, 7.6

Predictive Optimization, 3.4, 4.6

Process Thread, 3.5, 4.3

## R

Research Gap, 2.4

Resource-Based View (RBV), 2.1

## S

SAP (S/4HANA, PS, EWM, ME/MII), 2.2.3, 4.2, 6.1

Scenario-Based Evaluation, 5.4, 6.2

Semantic Integration, 3.4, 4.4

Simulation-Based Validation, 6.2, 7.6

Site Acceptance Testing (SAT), 6.3

Structured Expert Validation, 6.4, 7.6

Symbiotic Case Context, 1.2, 6.1

## T

Telemetry (Real-Time Data), 3.3, 4.2

Threads (Digital, Process, Execution, Decision, Collaboration), 3.5

## V

Validation (Expert-Based), 6.4

Value Chain Optimization, 1.1, 2.1, 7.8

Vertical Integration, 2.3.3, 3.3, 4.2

W

Warehouse Automation, 1.2, 6.1