Design-to-Hull Integration: Analysis and Optimization of the Engineer-to-Order Process Chain

MASTER THESIS

JAVIER CARRERA GUTIÉRREZ

TU Delft

FEADSHIP

ROYAL DUTCH SHIPYARDS

Thesis for the degree of Master of Science (MSc) in Marine Technology in the specialization of Maritime Operations and Management

Design-to-Hull Integration: Analysis and Optimization of the Engineer-to-Order Process Chain

by

Javier Carrera Gutiérrez

Performed at

Scheepswerf Slob

27 November 2026

Report number MT.25/26.011.M

Company supervisor
Responsible Supervisor: J. Kortenoeven

Thesis exam committee
Chair/Responsible Professor: dr. ir. J.F.J. Pruyn
Staff Member: dr. A. Napoleone
Company Member: J. Kortenoeven

Author details
Student Number: 6279759

An electronic version of this thesis is available at http://repository.tudelft.nl/.





DECLARATION OF GENERATIVE AI USE

The research and content of this thesis are my own. I used Claude (Anthropic) to restructure the text, improve its wording and keep the chapters consistent, and ChatGPT (OpenAI) to design some figures, marked in their captions. I take full responsibility for the final content.

ACKNOWLEDGEMENTS

This thesis concludes my MSc in Marine Technology at TU Delft. It was carried out at Scheepswerf Slob, part of the Feadship group.

I would like to thank Dr. ir. Jeroen Pruyn, my supervisor at TU Delft, for guiding me throughout this thesis. From the start he helped me find the right way to approach the problem and to build each part of the work. He also taught me that a good idea is not enough if you cannot tell it well and his feedback on the storyline and the writing has made this a much better thesis.

I am very grateful to Jeroen Kortenoeven, my supervisor at Slob, for opening the doors of the superyacht industry to me, for the many hours he spent with me and for involving me in the whole production chain. He was always there, introduced me to the right people whenever I got stuck and made every contact easy. Beyond the thesis, he gave me the chance to learn about fully custom superyachts, which are my passion.

Many people cannot be named here for reasons of privacy, yet this thesis would not have been possible without them. Engineers, project managers and production staff across the companies of the Feadship group gave me their time, explained how their work really goes and helped me rebuild the production chain, find its problems and improve this work.

Finally, thank you to my family and friends for always being there.

Javier Carrera Gutiérrez
Delft, November 2026

ABSTRACT

Building a custom superyacht means starting the hull before the engineering is finished. Somewhere between design and production, someone has to decide which unfinished parts can go ahead and which risks come with them. That decision is rarely written down. The same happens in any project where design and construction overlap.

This thesis studies that moment at Gate 4, where basic engineering (DOII) is released into detail engineering (DE) at Scheepswerf Slob, the hull yard of the Feadship group. The group's own measurements show that only about one in three structural deliverables reaches detail engineering on time. Interviews and project records show that engineers already cope by working around unfinished areas. What is missing is a shared record of what may go ahead, under which condition and who accepts the remaining risk.

Building on existing maturity assessments (PDRI, FEED MATRS and DECRIS) and on research into working with preliminary information, the thesis designs GAIA, the Gate Assessment of Information Availability. GAIA checks the context of a delivery, rates how mature each discipline is and turns every open item into one of four recorded decisions: GO, CONDITIONAL, ISOLATE or WAIT. It works with the lists and meetings that already exist.

GAIA was applied to three completed deliveries. What it would have decided at the gate was compared with what happened afterwards and other explanations were checked. Several items it would have held back or put under a condition were followed by changes or extra work in the same place. The context check also flagged a risk that no open item showed: a pending vibration result. Other problems had causes outside the design information, such as designs that were hard to build.

GAIA does not predict rework. It offers a way to make expert judgement visible and traceable, without replacing it. How much time and money that saves can only be seen in practice.

CONTENTS

INTRODUCTION 9
1.1 Superyachts as engineer-to-order products 10
1.2 Why design and production overlap 12
1.3 Problem statement: information immaturity at phase transitions 14
1.4 What immature information costs 15
1.5 Research scope and delimitations 17
1.6 Research objective and questions 18
1.7 Research design 19
STATE OF THE ART 21
2.1 How the literature was searched 21
2.2 What has to be managed at the release 24
2.3 PDRI: how well a project is defined 25
2.4 FEED MATRS: maturity and accuracy apart 26
2.5 DECRIS: engineering completeness before construction 27
2.6 LOD: how far a model can be trusted 29
2.7 TRL: how ready a technology is 30
2.8 Deciding with preliminary information 31
2.9 Stage-Gate governance in product development 32
2.10 Comparison of alternatives and research gaps 33
2.11 Synthesis and answer to SQ1 35
SOLUTION DESIGN 36
3.1 Requirements for the solution 36
3.2 Selecting the elements of the solution 36
3.3 Cynefin: where judgement is needed 37
3.4 Towards the framework 37
THE GAIA FRAMEWORK 38
4.1 Requirements and design choices 41
4.2 Level I: the context of the delivery 43
4.3 Level II: where the package is fragile 44
4.4 Level III: a decision for every open item 46
4.5 GAIA at the gate: who does what 49
4.6 Limits of the design 51
4.7 Synthesis and answer to SQ2 52
THE DESIGN-TO-HULL CHAIN AT SLOB 53
5.1 How the diagnosis was built 53
5.2 Industrial context: one-off superyachts 54
5.3 The process on paper (Work-As-Imagined) 55
5.4 What the organisation's own VSM measured 57

5.5 The process in practice (Work-As-Done) 59
5.6 Why intervene at Gate 4 62
5.7 Five findings and answer to SQ3 63
RETROSPECTIVE VALIDATION OF GAIA 64
6.1 Evaluation method and criteria 64
6.2 What is tested and what is not 66
6.3 How each case was assessed 67
6.4 Case A: the reference case 69
6.5 Case B: the second engine room 73
6.6 Case C: the stress case 77
6.7 Across the three cases 81
6.8 Synthesis and answer to SQ4 83
DISCUSSION 84
7.1 What the results mean 84
7.2 What GAIA adds to the release decision 85
7.3 How GAIA fits the organisation 87
7.4 What it costs and what it can give back 87
7.5 Where GAIA stops and what to improve 90
7.6 Next step: a pilot 91
CONCLUSIONS AND RECOMMENDATIONS 93
8.1 Answers to the research questions 93
8.2 Contributions 94
8.3 Recommendations and future work 95
BIBLIOGRAPHY 97
Annexes to Chapter 2 101
Annex 2-A. Literature overview 101
Annexes to Chapter 4 103
Annex 4-A. Traceability of GAIA elements to the literature 103
Annex 4-B. Level I rubric 104
Annex 4-C. The gate today and with GAIA 105
Annexes to Chapter 5 106
Annex 5-A. The seven wastes of the internal VSM exercise 106
Annexes to Chapter 6 107
Annex 6-A. Case A 107
Annex 6-B. Case B 110
Annex 6-C. Case C 113
Annex 6-D. Sources 116

LIST OF ABBREVIATIONS

AEC	Architecture, Engineering and Construction
ATO	Assemble-to-Order
BIM	Building Information Model
CD	Concept Design
CII	Construction Industry Institute
COND	CONDITIONAL (short form in tables)
DD	Design Development
DE	Detail Engineering
DECRIS	Detailed Engineering Completion Rating Index System
DO	Definitief Ontwerp, the integrated multi-disciplinary engineering phase
DOI	Definitief Ontwerp I
DOII	Definitief Ontwerp II, closed by the Final Review at Gate 4
EPC	Engineering, Procurement and Construction
ETO	Engineer-to-Order
EXT	Exterior outfitting
FEED MATRS	Front-End Engineering Design Maturity and Accuracy Total Rating System
FEM	Finite Element Method
FMECA	Failure Mode, Effects and Criticality Analysis
GAIA	Gate Assessment of Information Availability
HVAC	Heating, Ventilation and Air Conditioning
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization (ISOLATE in Table 6.12)
KO	Kickoff of a delivery
LOD	Level of Development
MTS	Make-to-Stock
NASA	National Aeronautics and Space Administration
OE	Operational Excellence programme
P&ID	Piping and Instrumentation Diagram
PDRI	Project Definition Rating Index
PE	Project engineer
PL	Project leader
PME	Project manager engineering
SQ1–SQ4	Sub-questions 1 to 4
TD	Technical Design
TRL	Technology Readiness Level
VFI	Vendor Furnished Information
VSM	Value Stream Mapping
WAD	Work-As-Done
WAI	Work-As-Imagined
WIP	Work-In-Process
WP, WVB	Work Preparation (WVB from the Dutch werkvoorbereiding)
WTB	Mechanical outfitting

INTRODUCTION

Only a third of the structural engineering deliverables arrive on time at Slob. At Scheepswerf Slob, the hull yard of the De Vries Group that builds Feadship superyachts, these deliverables are the engineering packages of each hull section, released from integrated design (DOII) into detail engineering (DE). When the group mapped its own value stream in 2026, it found that 35% of them arrived on their planned date, the steepest drop of the whole chain [1]. The figure measures punctuality, not maturity, but it marks the step where late and changing information starts to cost. This thesis calls that step Gate 4. Past it, a correction no longer stays in a model. It becomes production drawings to redo, work preparation to repeat and, in the worst case, steel to cut again.

This is not carelessness. A Feadship is designed, engineered and built at the same time, because the handover date is as fixed as the yacht is unique. The yacht is split into large blocks, called deliveries, and each one is released into detail engineering on a date set in the planning. The date holds even when the information behind it is late or still changing. The engineers know that some items will still be open at that moment: in their experience, open items always cross. What is not written down is what was accepted with them, by whom and with which risk.

Not every extra hour at the yard comes from immature information. Overruns have three broad sources: information released before it is mature, estimating and execution problems at the yard itself and scope added after the release. This thesis is about the first. The validation treats the other two as rival explanations and keeps them apart from what it tests.

Superyacht builders are not the only ones with this problem. Other industries that build one-off products against a fixed date share it, and some of them have learned to measure what they release. Offshore contractors score how complete the detail engineering is before they cut the first steel. Builders of industrial plants rate how well a project is defined before the money is committed, and better-defined projects have been shown to perform better on cost and schedule. Space agencies rate each new technology on a readiness scale before it goes on a mission, and building designers state how far each element of a model can be trusted. None of these tools was made for a one-off superyacht, and none fits the yard as it is. But they share an idea worth bringing to shipbuilding: measure how mature the information is at the moment it is committed, and let that measure decide what goes ahead and what waits. This thesis brings that idea to Gate 4, in a framework called GAIA. Chapter 2 reviews the tools it learns from.

The rest of this chapter builds that argument step by step. It places the superyacht among the engineer-to-order industries (Section 1.1) and explains why these industries overlap design and production (Section 1.2), where that overlap creates the problem (Section 1.3) and what it costs (Section 1.4). It then sets the scope (Section 1.5), the research questions (Section 1.6) and the design of the research (Section 1.7), where Table 1.1 shows which chapter answers which question.

1.1 Superyachts as engineer-to-order products

Shipbuilding covers a wide range of production strategies, from standardised commercial vessels built in series to one-off ships designed for a single owner [2]. The superyacht sector sits at the one-off end: each project demands an extreme degree of customisation, where every detail matters, which places it within the Engineer-to-Order (ETO) paradigm.

Figure 1.1 places industries on a spectrum defined by where the customer order enters the value chain, the customer order decoupling point [3]. This is a positioning, not a historical sequence: firms move towards standardised production where demand and the product allow it, and stay engineer-to-order where they do not [3]. Shipbuilding itself covers the whole spectrum, from standard designs built in series to fully customised vessels [2]. Custom superyacht builders such as the De Vries Group (Feadship) sit at the far end, where every vessel is a one-of-a-kind project designed and built from scratch.

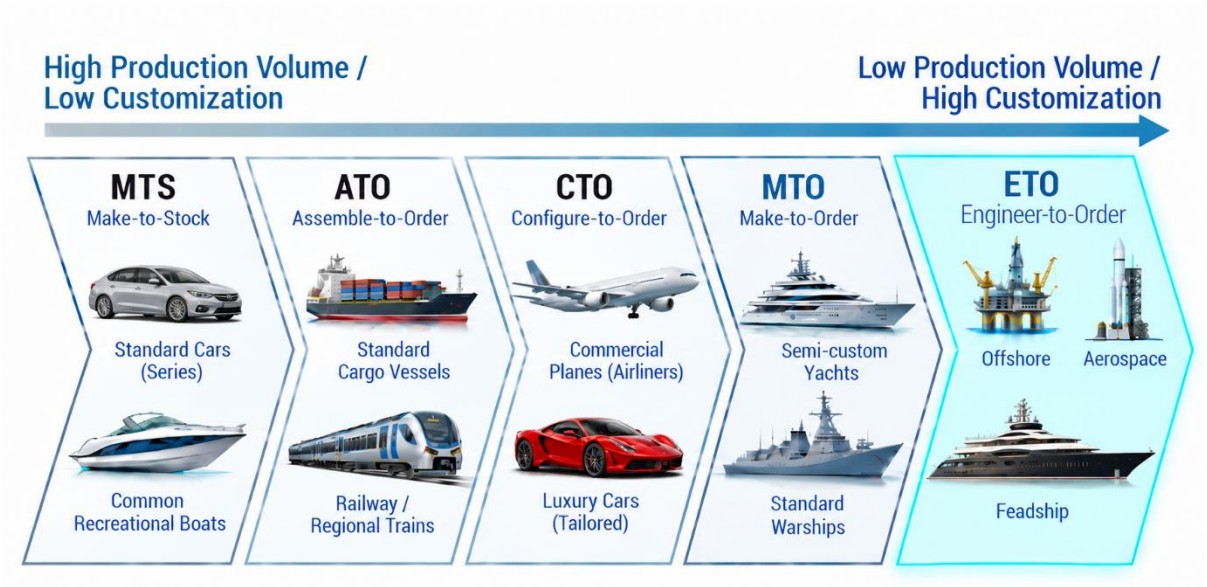


Figure 1.1: Production strategies by customization. Based on [2]–[4]; AI-assisted illustration.

Since each yacht is one-of-a-kind, there is no full-vessel prototype: the first vessel built is the final product. Consequently, management tools developed for MTS or ATO environments are not directly transferable to the full-custom superyacht context [4], [5]. Comparable industrial settings, sharing the combination of high technical complexity and low production volumes, include offshore engineering (Oil & Gas platforms), aerospace and mission-specific military aviation.

This challenge of managing uncertainty in the absence of a validated baseline design is not unique to the superyacht sector. It is a common characteristic that is shared among all Engineer-to-Order (ETO) industries. In the offshore oil and gas sector, Kim et al. [6] documented that premature start of construction based on incomplete engineering is a major driver of cost overruns in EPC megaprojects. In aerospace, the European Space Agency (ESA) developed Concurrent Design Facilities specifically to manage uncertainty in parallel engineering processes [7]. In industrial construction, the Construction Industry Institute (CII) demonstrated that projects with higher levels of front-end planning maturity consistently outperform those

that proceed with incomplete scope definition [8]. Across these industries, releasing incomplete or unstable engineering information can contribute to rework, cost growth and schedule disruption.

Although the end products differ radically among these ETO ecosystems, the day-to-day problems in design, engineering and production are similar. The greatest of them is managing uncertainty without a previously validated design, because engineering has to be developed in parallel with production to meet the client's delivery date [5]. This overlap of design, engineering and production is called concurrent engineering. It causes problems, but it is done for a clear reason: it shortens the lead time, and in these markets a short and fixed delivery time is a competitive requirement (Section 1.2). The price is that information is released to production before it is final, to keep the schedule. As Figure 1.2 shows, this can start a "downward labour productivity spiral" of rework [5].

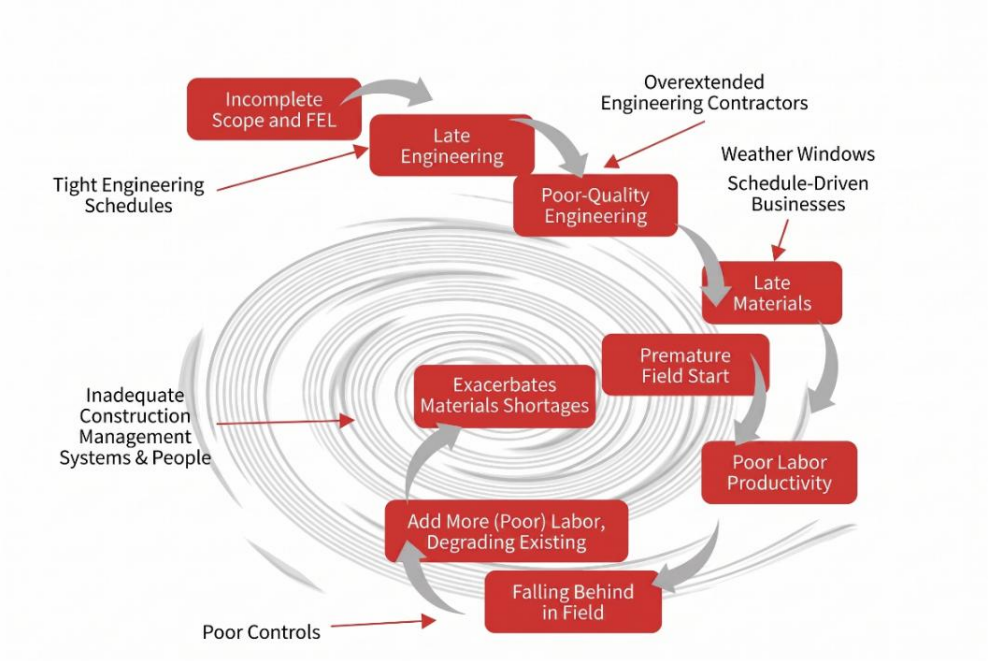


Figure 1.2: The Downward Labour Productivity Spiral. Source: [5]

Rework in ETO environments is not only driven by internal information immaturity. External factors also contribute significantly to downstream variability, including client-requested modifications, evolving regulatory requirements from classification societies and delays in vendor-furnished information. This thesis does not try to control these external drivers. It does require that they are made visible at the release, because an assessment that ignored them would miss some of the largest sources of later rework.

The ETO manufacturing paradigm, whether applied to superyachts, offshore platforms or aerospace systems, creates a fundamental tension between speed and predictability. Concurrent Engineering compresses timelines but introduces information volatility that cascades through the production chain. The next section examines how this tension manifests in the specific constraints of time, cost and scope (the classic 'Iron Triangle' of project management) and why it forces production to begin before all information is fully mature.

1.2 Why design and production overlap

The "Iron Triangle" of project management (Figure 1.3) evaluates project delivery on three interconnected constraints. Barnes [9] framed them as time, cost and quality. They are commonly expressed as scope, cost and time, the form used here. In the full-custom projects studied here, the *scope*, defined as the intersection between the client's ultimate vision and the project's technical feasibility, and the *time*, often dictated by rigid cruising seasons or global events, are fixed by the client. Consequently, *cost* acts as the moving variable for complexity. This financial buffer allows the shipyard to resolve unforeseen technical challenges, accelerate schedules and manage the inherent uncertainty of "one-off" projects without compromising the shipyard's viability or reputation.

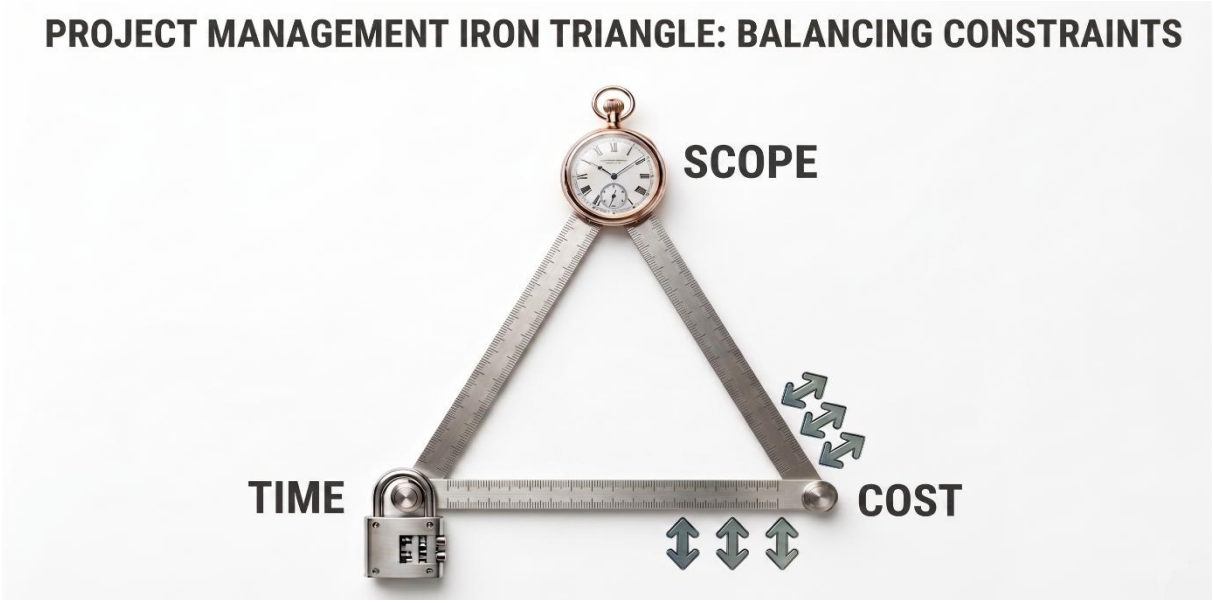


Figure 1.3: The project management iron triangle. Based on [9]; AI-assisted illustration.

In a market directed exclusively at ultra-high-net-worth individual (UHNWI) clients, failing to meet the expected scope or quality is not an option. For the yard, delivering highly integrated vessels within strict timeframes is a primary competitive driver. To achieve this, shipyards like the De Vries Group cannot wait for the entire design phase to be finalized before starting engineering, nor can they wait for engineering to be fully complete before starting hull production. The traditional "plan-then-execute" paradigm is not a competitive option in this environment.

One could ask why the yard does not simply add engineers to close the gap. Additional engineering capacity cannot remove dependencies whose determining information lies upstream, such as supplier data, class approval or owner decisions, which follow their own clock. Adding people does not solve a second problem either, which comes from concurrent work itself. Ford and Sterman [10] found that teams working in parallel under schedule pressure tend to keep known rework to themselves rather than report it, so it is discovered late, when it costs more. Whether the limiting factor at this transition is capacity or information is examined for Slob in Chapters 5 and 6.

To reduce the overall lead time, the industry overlaps the different phases of the production chain. This is the concurrent engineering introduced in Section 1.1, defined by the U.S. Institute for Defense Analyses (IDA) as the integrated, concurrent design of products and their related processes, including manufacture and support, with multidisciplinary teams and shared computer models [11].

Concurrent Engineering is meant to reduce uncertainty and firms facing fast technological and product change use it more [12]. In ETO contexts, however, design continues to evolve throughout engineering and production due to client-driven changes, regulatory feedback, supplier data and innovation maturation. Under these conditions, concurrency can turn from a way to save time into a source of variability. This thesis argues that the transitions between phases are where that variability can still be controlled.

Because upstream work is not finalized, downstream phases have to make decisions based on assumptions. If a structured plan and a rigid quality control system are not in place, the premature release of immature information makes it very likely that some of these assumptions will fail. When this occurs, the resulting rework is not just a loss of material, but a loss of time, triggering the *downward labour productivity spiral*. Therefore, measuring information maturity before it is released to the next phase is a central operational challenge for a full-custom yard such as Slob.

The Iron Triangle constraints in ETO environments make concurrent engineering not a choice but a competitive necessity. However, this overlap between design and production creates a structural weak point: information must be released before it is fully mature and the consequences of a premature release grow as the project moves forward. This raises the question of how the risk of immature information can be identified and controlled at the point where it crosses from engineering into production-oriented Detail Engineering (DE), Work Preparation (WP) and hull construction.

1.3 Problem statement: information immaturity at phase transitions

This section separates two things that are easily mixed up. The problem is the uncertainty that crosses the transition from design into detail engineering. The impact is what that uncertainty costs once it reaches production, and it is treated in Section 1.4. The two are kept apart because the impact also tells which transitions matter most.

In all ETO industries where concurrent engineering is part of the essence, a weak point arises at the interfaces between design phases and production: engineering information is released to downstream phases before it is fully validated, verified and stable. This phenomenon has been documented across different sectors: in offshore construction [6], industrial construction [8], aerospace systems engineering [7], complex shipbuilding [13]–[15]. The root cause is structural: concurrent engineering requires downstream work to begin before upstream decisions are finalized in order to reduce the total lead time. The formal process assumes that what is released, apart from the listed open items, is valid: this is the process as imagined, Work-As-Imagined (WAI). In practice the people downstream know that this is not always true, and they add informal checks and work-arounds to cope: the process as done, Work-As-Done (WAD) [16]. When the released information is less mature than the process assumes, those informal loops have to absorb the difference, and what they miss appears as rework, schedule delays, cost overruns and workflow disruptions.

At the boundary between *Definitief Ontwerp* (DO, the integrated multi-disciplinary engineering phase) and Detail Engineering (DE, the production-oriented engineering phase), called Gate 4 in this thesis, an error is no longer only a correction in a CAD model. It turns into scrapped steel, wasted labour and schedule disruption. Section 1.4 explains why this cost grows with every phase.

Internal evidence, examined in Chapter 5, indicates that the constraints of the Iron Triangle often lead Feadship's Central Engineering to release technical information to DE before the design is fully mature. To maintain the project schedule and protect the shipyard's reputation, engineers must make key assumptions, such as unconfirmed Vendor Furnished Information (VFI), delayed Finite Element Method (FEM) calculations or late-stage client innovations. This creates a discrepancy between the theoretical office planning and the shipyard's reality. From a management perspective, the release of a drawing package represents progress (Work-As-Imagined). However, managers see that a deliverable exists, but not how mature it is. In operational reality, DE and Slob receive data with gaps and assumptions and they cope through their own checks, fenced-off zones and work-arounds (Work-As-Done) [15], [16].

When the actual vendor data, structural calculations or finalized client requirements eventually arrive and DO updates the core design, the premature assumptions end up collapsing. Because DE has already advanced the 3D modelling, and in some cases, Slob has initiated production, this can cause substantial rework. This issue is further intensified by the structure of the chain: DO and DE are different organisations in different locations, with different perspectives. The gap between design intent and production reality (Interface Management) [1] points to the DO-DE frontier as a bottleneck, which Chapter 5 examines in detail. Consequently, evaluating and controlling information maturity at this boundary becomes essential to mitigate rework and optimize the overall Design-to-Hull process.

The question is therefore not whether maturity at this transition matters, but how it can be assessed and turned into a decision at the gate. Existing quality gates in shipbuilding mostly check that documents exist, not how mature their content is. Chapter 2 reviews the methods that already exist for this.

1.4 What immature information costs

In any production system, variability is the main enemy of predictability. Factory Physics [17], [18] establishes that systems with high variability require either excess capacity or excess inventory (Work-In-Process, WIP) to maintain throughput. In manufacturing, WIP refers to partially completed products that occupy space, consume resources and await processing. The greater the amount of WIP in a system, the longer the average lead time and the less predictable the output.

Little's Law [19] uses this simple equation: $L = \lambda W$, where L is the average number of items in the system (WIP), λ is the average throughput rate and W is the average time an item spends in the system (lead time). The relation is direct: if WIP increases without a corresponding increase in throughput, lead time will increase proportionally.

While Factory Physics was developed for physical manufacturing systems, its principles can be directly applicable to information flows in project environments. In ETO industries where engineering and production overlap, information packages released to downstream phases act as WIP in the production chain. When these packages are released prematurely due to unverified assumptions, provisional data or incomplete analyses, they represent defective WIP. Unlike physical WIP, which is visible, information WIP is invisible: a drawing package released to production looks the same whether its underlying data is confirmed or provisional. This invisibility makes defective information WIP particularly dangerous as it is undetectable until the assumptions it contains are proven wrong.

When this defective information WIP enters the production system, it generates rework cycles. Each rework cycle consumes capacity without generating new throughput, consequently, the WIP increases without increasing the output. According to Little's Law, this increases lead time. Moreover, because rework in production-oriented DE, WP and subsequent hull construction is much more expensive than corrections in the engineering phase (the 1-10-100 Rule [20]), defective information WIP carries a compounding cost penalty that grows with each phase transition.

The MacLeamy Curve [21], widely used in Architecture, Engineering and Construction (AEC), illustrates this principle (Figure 1.4): the ability to influence project cost decreases rapidly as the project advances, while the cost of making changes increases correspondingly. In the context of ETO manufacturing, this implies that the optimal strategy is to invest effort in the early phases (Front-End Loading) to ensure that information crossing into production is as mature as possible. The second effort curve in the figure, labelled IPD, refers to Integrated Project Delivery, a project delivery approach in which owner, designer and builder work together from the start so that decisions move earlier [22].

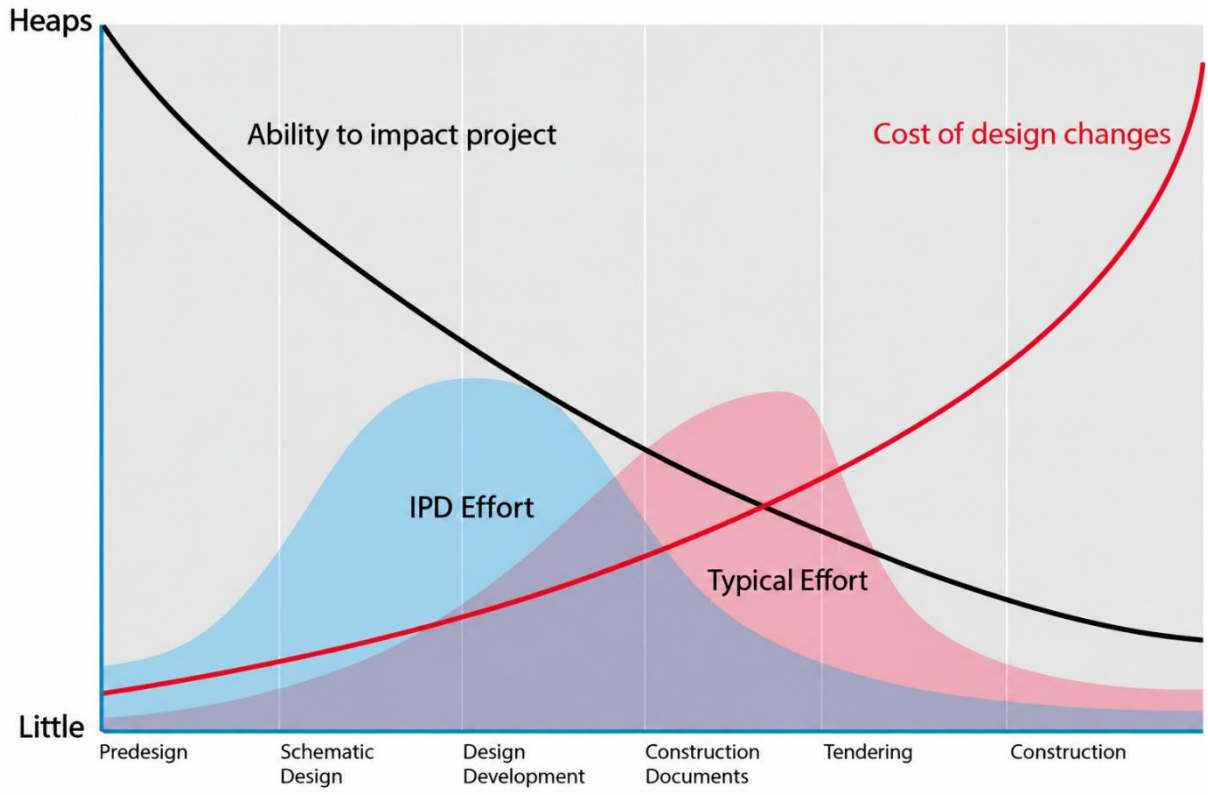


Figure 1.4: The MacLeamy Curve. Adapted from [21].

How large can these costs be? There are no published figures for superyachts, but building construction has measured them. In 161 Australian building projects, rework accounted for about half (52%) of the cost growth [23]. In a separate study, the direct cost of rework alone often reached about 5% of total construction costs [24]. Because no such figures exist for superyachts, this thesis uses the yard's own records (Chapter 5).

Implication for this research: releasing immature information to production can be read as injecting defective WIP into the production chain. Controlling the quality of information at phase transitions, specifically by measuring its maturity before it enters the production chain, should, in theory, reduce lead time and total cost, to the extent that it prevents rework cycles. This provides the theoretical justification for developing a structured maturity assessment at critical gates.

1.5 Research scope and delimitations

To ensure the feasibility of this research, specific boundaries have been established. These delimitations are not randomly chosen as they reflect three converging criteria: the operational point where information immaturity becomes most costly, the availability of empirical evidence and stakeholder access at the De Vries Group and the alignment with the case study methodology adopted in this thesis.

- **In-Scope**: The study focuses on the *Design-to-Hull* process chain up to the point of hull delivery by the Slob yard. This includes structural engineering and the associated "hull outfitting" (e.g., structural foundations, penetrations and primary piping) executed by Slob. The critical assessment boundary is the information handover from Central Engineering (DO) to Detail Engineering (DE), evaluating the accumulated maturity of all upstream deliverables at this specific gate. Gate 4 was selected as the assessment point because it is the moment in the chain where information immaturity escalates from a recoverable design issue into a costly production consequence (the 1-10-100 Rule).
- **Out-of-Scope**: Several elements of the broader process are deliberately excluded from this thesis, each for a specific reason:
  - Early conceptual phases (Concept Design (CD), Technical Design (TD), Design Development (DD) and *Definitief Ontwerp* I (DOI)) generate the foundational data evaluated at Gate 4 but their internal management processes are not the subject of optimization in this work. The principle of maturity-before-transition developed in this thesis is theoretically applicable to these earlier transitions and is discussed as a research direction in Chapter 7.
  - Post-hull delivery phases, such as luxury architectural outfitting, exterior finishing and final commissioning, are excluded as the scope of this thesis is delimited to the completion of all hot work and the subsequent departure of the hull from the Slob shipyard. The applicability of the principles of this framework to these phases is also briefly noted in Chapter 7 as future research.
  - Purely technical IT issues and CAD software interoperability challenges (e.g., automated 2D drawing extraction from 3D models) are excluded because they represent engineering tooling problems rather than process governance problems.

Furthermore, this thesis develops and evaluates the proposed assessment framework through retrospective case analysis. The framework's indicative consistency with historical rework patterns is examined; predictive validation through forward-looking implementation in active projects is identified as future research and is beyond the scope of this work.

Although the root causes of information immaturity may originate in earlier phases such as CD, TD, DD, DOI or DOII, this thesis does not attempt to redesign those phases internally. Gate 4 is selected because it is the last controllable transition before accumulated upstream immaturity enters production-oriented Detail Engineering, Work Preparation and hull construction, where the cost of correction increases significantly. Therefore, the proposed framework assesses the accumulated maturity of upstream work at Gate 4. It does not claim that Gate 4 is the unique origin of the problem.

1.6 Research objective and questions

Within this scope, the thesis pursues one objective and one main research question, divided into four sub-questions.

Research Objective The aim of this research is to reduce rework and its associated costs at the Design-to-Hull transition in custom superyacht construction, thereby improving schedule predictability and overall hull delivery reliability at the De Vries Group. To achieve this aim, this thesis develops an assessment framework, GAIA (Gate Assessment of Information Availability), capable of measuring and structuring the maturity of technical information transferred between Central Engineering (DO) and Detail Engineering (DE), enabling more informed release decisions at Gate 4.

Main Research Question To what extent can rework in hull construction be mitigated by systematically assessing information maturity at the design-to-production interface?

Sub-Research Questions To address the main research question, the following sub-questions have been formulated:

- **SQ1:** What characteristics of engineering information determine its readiness for production release in complex, multi-disciplinary ETO construction environments?
- **SQ2:** How can these characteristics be measured and integrated into a quantitative, verifiable assessment framework at a quality gate?
- **SQ3:** What are the current causes of rework and inefficiency at the design-to-production interface in ETO shipbuilding, and how do planned information flows (Work-As-Imagined) diverge from the information actually delivered (Work-As-Done)?
- **SQ4:** To what extent do the maturity scores generated by the tailored assessment framework show indicative consistency with historical rework and production performance when applied retrospectively to completed projects?

Each sub-question drives one part of the thesis (Table 1.1). SQ1 is answered by the state of the art (Chapter 2). SQ2 is answered by the solution design and the framework (Chapters 3 and 4). SQ3 is answered by the diagnosis of the yard (Chapter 5), which also shows the yard's evidence behind the requirements of the framework. SQ4 is answered by the retrospective validation (Chapter 6). Chapter 7 discusses the results and Chapter 8 answers all the questions, the main one first.

Two clarifications on how these questions are answered. The main question is answered at the level of credibility and indicative consistency: the retrospective cases can show whether the signals of the framework match what happened, not prove that rework would have been avoided, which requires prospective application (Chapter 8). In SQ2, “quantitative” is understood as ordinal, evidence-based scoring, not as a calibrated numerical model.

1.7 Research design

The research follows a Design Science methodology, where the objective is not only to understand a problem but to develop an artifact (a framework, tool or method) that addresses it. Design Science was selected over alternative methodologies (such as action research or a purely descriptive case study) because the research aim is not only to diagnose a problem but to produce a usable instrument that can be evaluated against the problem context [25].

The study is structured in five phases aligned with the thesis assignment:

- **Phase 1 (Analysis):** Mapping the current design-to-hull chain through document review and stakeholder interviews, contrasting the theoretical process design (Work-As-Imagined) with the daily operational reality (Work-As-Done).
- **Phase 2 (Diagnosis):** Analysing the Value Stream Mapping (VSM) exercise carried out by the De Vries Group [1], the first practical step taken with the company to understand the problem and triangulating its bottlenecks with stakeholder interviews and project documentation, while the literature was reviewed in parallel.
- **Phase 3 (Concept Development):** Developing an assessment framework based on identified research gaps, adapting existing methodologies from the literature to the specific context of ETO shipbuilding.
- **Phase 4 (Evaluation):** Evaluating the framework through retrospective application to completed projects, examining the indicative consistency between the framework's output and historical production performance data.
- **Phase 5 (Synthesis):** Formulating conclusions, limitations and recommendations for implementation and future research.

The phases describe the order in which the work was done. Phases 1 and 2 came first, and the diagnosis of the yard shaped the requirements of the framework. Chapter 1 uses their first results to frame the problem. Chapters 2 to 4 develop the framework (Phase 3). Chapters 5 and 6 evaluate it (Phase 4): Chapter 5 sets the full diagnosis of the yard against each requirement, and Chapter 6 applies GAIA to three completed deliveries. Chapters 7 and 8 form the synthesis (Phase 5). Table 1.1 links each chapter to its sub-question and phase.

Chapter	What it does	Answers	Phase
1 Introduction	Sets out the problem, the scope, the questions and the research design	Poses the main question	1 and 2 (framing)
2 State of the art	Reviews existing methods and what they leave open	SQ1	3
3 Solution design	Turns the gaps into requirements and selects the elements	SQ2	3
4 The GAIA framework	Describes the framework, its rules and its roles	SQ2	3
5 The design-to-hull chain at Slob	Diagnoses the chain: the process on paper against the process in practice	SQ3	4
6 Retrospective validation	Applies GAIA to three completed deliveries	SQ4	4
7 Discussion	Interprets the results: value, cost, limits and a pilot	Discusses all	5
8 Conclusions	Answers the questions and gives recommendations	All, main question first	5

Table 1.1: Chapters, sub-questions and research phases.

This chapter has set out the context (engineer-to-order construction under concurrent engineering), the problem (information that crosses into detail engineering before it is stable), its impact (rework that grows with every phase), the scope (the DOII to DE release at Gate 4) and the questions and design of the research. Chapter 2 starts with how the literature was searched and then reviews the methods that exist for assessing and deciding on immature information and what they leave open.

STATE OF THE ART

This chapter answers the first sub-question of the thesis (SQ1):

> *"What characteristics of engineering information determine its readiness for production release in complex, multi-disciplinary ETO construction environments?"*

The chapter answers it through two questions. The first is how the maturity of information can be assessed. The second is what to do with the items that are still open when the gate closes. The chapter starts with how the literature was searched (Section 2.1) and with what has to be managed at the release (Section 2.2). Sections 2.3 to 2.7 review the methods that assess maturity. Sections 2.8 and 2.9 review how to decide on preliminary information and how a gate governs that decision. Section 2.10 compares the methods and names the gaps, and Section 2.11 draws the conclusions and answers SQ1.

2.1 How the literature was searched

The literature review followed an iterative Snowballing strategy [26] rather than a systematic literature review (SLR). The choice was deliberate, as systematic methods are most effective when the research area is stable and the objective is a comprehensive study of a specific field. In this research, the relevant literature was distributed across multiple disciplines (operations management, lean systems, stage-gate governance, BIM, offshore EPC, megaproject management) with no single field providing a complete answer. A snowballing approach was therefore more effective for uncovering related literature through citation chains, allowing the theory to evolve naturally alongside the interview findings.

The initial phase of the research, in consultation with the supervisors, began with an exploration of the operational dynamics of the ETO industry and the concept of Concurrent Engineering. Building upon these two key concepts, the Snowballing method (chain-searching from key references), was adopted to track down relevant literature. Since specific studies on custom superyachts are limited, the research was redirected towards megaproject management methodologies in other ETO and maritime industries.

The databases utilized included Google Scholar, ResearchGate and the TU Delft Education Repository, complemented by fundamental literature on production physics and megaproject management. Several studies, especially extracted from the TU Delft repository, were highly relevant by addressing similar information flow and design-to-build interface issues in comparable Dutch shipyards.

The research process was structured into three conceptual phases that led the research to potential solutions:

1. **ETO Context and Production Dynamics:** the investigation began by analysing the general dynamics of ETO industries and the systemic failures derived from variability. Literature on operations physics (e.g., Factory Physics) and on uncertainty management where design and production overlap was reviewed. Megaproject literature was also reviewed to understand planning bias, uncertainty and recurring project underperformance [27], [28].

2. **Operational Problems and "Work-As-Done":** to align the theory with the specific instability of superyachts, the search focused on information flow and design-to-build interfaces within shipyards. Focusing on literature published after 2015, the digital gap (3D to 2D interoperability), rework cycles caused by the premature release of information, and the discrepancy between management's theoretical process (WAI) and the shop floor reality (WAD) were analysed [13], [15]. Finally, recent literature on engineering waste [14] and documentation flow [29] was used to properly label and categorize the daily frustrations and inefficiencies observed in practice.

3. **Solution Convergence:** in parallel with the theoretical review, the empirical findings from the interviews narrowed the gap between literature and operational reality. The research focused on the diagnosis of information maturity, specifically at the critical interface between engineering and production. Consolidated industrial methodologies such as Advanced Work Packaging (AWP) and the commercial standard Project Definition Rating Index (PDRI) developed by the Construction Industry Institute (CII) [8] were evaluated. Following cross-references on technical maturity, the search converged on the offshore EPC (Engineering, Procurement, and Construction) industry due to its structural complexity similarities. This led directly to the core literature of this thesis: the DECRIS (Detailed Engineering Completion Rating Index System) method [6]. Further analysis of related EPC literature [30], [31] provided the mathematical correlation between engineering maturity and the escalation of construction costs.

To adapt this framework to the ETO shipbuilding reality, literature regarding geometric reliability was integrated, incorporating Level of Development (LOD) standards applied to BIM modelling [32]. Moreover, literature addressing procurement bottlenecks in EPC networks [33] highlighted the critical need to quantify the impact of Vendor Furnished Information (VFI) and the Technology Readiness Level (TRL) of Owner-driven "Specials" [34]. Given the lack of a maturity index for shipbuilding encompassing these variables, the adaptation of the existing framework is the core contribution of this thesis, developed in Chapter 4.

In a later iteration, the design of the item-level logic of the framework led the search to two further strands: the overlapping of activities on preliminary information [35] and the propagation of engineering changes [36]. Both are introduced in Section 2.8. The method used to evaluate the framework on completed projects follows case study research [37] and is described in Section 6.1.1.

Throughout these iterative phases, strict exclusion criteria were applied. Pure Lean Manufacturing methodologies focused on serial naval production (Make-to-Stock or standardized commercial shipyards) were discarded, as their stability principles fail to absorb the extreme volatility of full-custom design.

Much of the search was driven by the interviews. Engineers and managers across the company described how the process works and which ideas they rely on, such as the Iron Triangle of time, cost and scope. Themes emerging from the interviews opened additional search strands. Those searches looked for relevant theory, supporting evidence and alternative explanations, and the sources found became new starting points for snowballing. Table 2.1 summarises where the search started, how each strand was found and what it led to.

Starting point	How it was found	What it led to
The company's VSM exercise	Internal sessions of the De Vries Group	The practical problems already identified by the company, which guided the first searches [1]
Two keywords from the supervisors: concurrent engineering and ETO industry	Google Scholar	Engineering maturity at the start of construction in offshore EPC: DECRIS [6]
DECRIS	Backward snowballing through its references	PDRI [8], then FEED MATRS [38], then LOD [32]
Shipbuilding and superyachts	Google Scholar and TU Delft repository	Theses on information flow in Dutch yards [13], [15], [29], and through them Work-As-Imagined and Work-As-Done [16]
Interviews across the company	Themes from interviews, each opening a new search strand	The Iron Triangle [9], TRL for a project with new technology [34], [39]
Supervisor suggestions	Interviews and a recommended book	Leading and lagging indicators [40], [41], Cynefin [42], megaproject research [27], [28], [43]
The yard's own gate process	Company documents	Stage-Gate [44], [45]
The item-level logic, during the case study	Targeted search	Overlapping on preliminary information [35], change propagation [36], criticality [46], conditional gates [47]
Background from earlier studies	Prior coursework	Little's Law [19], the 1-10-100 rule [20]

Table 2.1: How the literature was found.

A source was kept when it addressed the maturity, release or overlap of engineering information in project-based or engineer-to-order settings, or provided the theory needed to assess it. The search was stopped when new references kept pointing back to sources already included. Literature was abundant on concurrent engineering and ETO, limited to one family of maturity indices in construction and offshore EPC, and almost absent for shipbuilding, which is itself part of the gap described in Section 2.10.

A chronological overview of the key references, with what each contributes to this thesis, is given in Annex 2-A.

Together these references cover the mechanism of the problem, the place to intervene, the structure of the gate decision, the way open items can be decided and the method to evaluate the result.

Selection criteria for the reviewed frameworks. The five frameworks reviewed in this chapter were selected through the iterative literature search described above, applying three criteria. First, the framework must produce a structured, verifiable assessment of engineering or information maturity at a defined milestone, rather than a general process audit or qualitative evaluation. Second, the framework must either have been validated empirically against project performance indicators (cost, schedule or rework) or be an established standard for describing maturity, as is the case for LOD and TRL. Third, the framework must originate from a sector that shares structural similarities with ETO superyacht construction: high project complexity, low production volume and reliance on concurrent engineering. The combination of these criteria resulted in five frameworks that, taken together, cover the relevant dimensions of the maturity assessment problem: front-end scope definition (PDRI), dual-dimension assessment of maturity and accuracy (FEED MATRS), engineering completeness at the construction interface (DECRIS), geometric maturity of 3D models (LOD) and technology risk for innovative components (TRL). Other frameworks were considered during the search (for example, Advanced Work Packaging, ISO 19650 information management standards, Construction Stage-Gate variants) but excluded either because they describe broader management processes rather than maturity assessment instruments or because their empirical validation against project outcomes is not yet available in published literature.

2.2 What has to be managed at the release

Before looking at methods, it helps to name what a method at this transition has to cope with. The problem described in Chapter 1 and the literature point to five kinds of uncertainty that cross the release from design into detail engineering:

- **Incomplete information:** parts of the design are not yet modelled.
- **Unstable information:** parts are modelled but rest on data that can still change, such as preliminary supplier footprints or calculations not yet confirmed.
- **External dependencies:** information owned by suppliers, the classification society or the owner, which follows its own timeline.
- **Open items:** concrete issues still being resolved when the gate closes.
- **Context conditions:** uncertainties that affect a whole package rather than one item, such as a pending vibration analysis, weight growth or new technology.

Four functions follow from this: assessing how defined and how reliable the information is, deciding what may proceed while some of it is still open, judging how far a change would travel, and governing the decision at the gate. Sections 2.3 to 2.9 review the existing methods for each function, and Section 2.10 compares them.

2.3 PDRI: how well a project is defined

The Project Definition Rating Index (PDRI) was developed by the Construction Industry Institute (CII) as a tool to evaluate the completeness of project scope definition during the Front-End Planning (FEP) phase [8]. PDRI assesses 70 elements organised in three sections (Basis of Project Decision, Basis of Design, and Execution Approach) on a scale of 0 to 5, producing an aggregate score between 70 and 1000. A lower score indicates more complete scope definition.

The CII validated PDRI across a large sample of industrial construction projects and demonstrated a statistically significant correlation between PDRI score and project outcomes: projects with better scope definition (lower PDRI) consistently achieved superior cost performance and schedule performance compared to those with incomplete definition. This established PDRI as one of the most widely adopted maturity assessment tools in the construction industry.

However, PDRI presents several limitations when considered for application at a design-to-production interface in ETO shipbuilding. First, PDRI was designed for the Front-End Planning phase, that is, before detailed engineering begins. Its elements evaluate scope definition and conceptual completeness, not the geometric and data maturity of a 3D engineering model ready for production. Second, PDRI was developed for and validated in the industrial construction sector (petrochemical plants, industrial facilities), whose standardised production processes differ substantially from one-of-a-kind superyacht construction. Third, PDRI evaluates projects as a whole, not individual blocks or zones within a project, which limits its granularity for hull block assessment.

Despite these limitations, PDRI sets an important precedent: that a structured, quantitative assessment of maturity at a defined milestone, when validated against project outcomes, can function as a leading indicator of production performance. This principle is directly transferable, even if the specific elements and application point must be adapted.

2.4 FEED MATRS: maturity and accuracy apart

The Front-End Engineering Design Maturity and Accuracy Total Rating System (FEED MATRS) represents an evolution of the CII's maturity assessment approach [38]. While the original PDRI produces a single aggregate score, FEED MATRS introduces a critical innovation: it separates the assessment into two independent dimensions.

The first dimension, maturity, evaluates the completeness and progression of engineering deliverables, basically measuring how much of the required work has been done. The second dimension, accuracy, evaluates the reliability and contextual confidence of the information, essentially measuring how trustworthy the completed work is. A deliverable can be highly mature (complete) but have low accuracy (based on assumptions rather than confirmed data), or vice versa.

FEED MATRS produces two separate scores that are interpreted together in a matrix. The CII validated this approach empirically: projects scoring high on both maturity and accuracy outperformed projects scoring low on both by 24% in cost performance and 12% in change performance [38].

This dual-dimension approach is highly interesting to ETO shipbuilding. In concurrent engineering environments, it is common for a 3D model to appear geometrically complete while the underlying data remains provisional: pipe routing may be fully modelled based on catalogue dimensions rather than confirmed vendor specifications or structural reinforcements may be positioned based on preliminary FEM results that have not yet received Classification Society approval. A single-dimension assessment would rate such a model as "mature," potentially approving its release to production. A dual-dimension assessment would flag the disconnect between geometric completeness and data reliability, revealing the hidden risk.

The FEED MATRS approach provides the conceptual precedent for separating the assessment of engineering information into a geometric dimension (how complete is the physical representation?) and a data dimension (how reliable is the underlying information?). This separation forms one of the key theoretical foundations for adapting existing maturity frameworks to the specific challenges of ETO shipbuilding. Two clarifications are needed. First, accuracy in FEED MATRS also covers factors of the FEED execution context, so reading it as confirmed supplier and design data is this thesis's adaptation to shipbuilding, not a literal transfer. Second, FEED MATRS supports assessing the dimensions separately, but it does not prescribe how to combine them. Letting the weakest dimension set the result would be a good property for this problem: a strong score could then never hide a weak one.

2.5 DECRIS: engineering completeness before construction

The Detailed Engineering Completion Rating Index System (DECRIS) was developed by Kim et al. [6] specifically to address a gap that PDRI does not cover: measuring engineering maturity at the point where construction begins, rather than at the end of Front-End Planning. While PDRI asks "is the scope well-defined before we start engineering?", DECRIS asks "is the engineering sufficiently complete before we start building?"

DECRIS was designed for the offshore EPC (Engineering, Procurement, Construction) industry, where fixed-price contractors must decide the optimal point to initiate steel cutting and fabrication. The system evaluates 69 elements organised in three sections: Basis of Detail Design (12 elements covering project scope, performance requirements and design guidelines), Engineering Deliverables (46 elements covering all engineering outputs by discipline) and Execution Approach (11 elements covering construction methodology, logistics and site preparation).

Each element is assessed on a scale of 1 to 5, with weights derived from a survey of experts. The weighted scores are aggregated into a single index (70–1000) that represents the overall engineering maturity at the assessed milestone.

Kim et al. [6] validated DECRIS retrospectively across 13 completed offshore EPC projects. The results demonstrated a solid correlation between the DECRIS score at the point of construction start and two key performance indicators: Construction Labor hours Increase Rate (CLIR), measuring additional labour hours relative to estimate, and Construction Duration Delay (CDD), measuring schedule delay relative to plan. Projects classified as "High maturity" showed significantly better CLIR and CDD than projects classified as "Low maturity." In a companion study, [30] extended this analysis to five milestones between contract award and the start of construction, using engineering maturity to forecast schedule and cost performance and reinforcing the argument for measuring maturity before moving to production.

DECRIS represents the closest precedent to the framework required for this research, as it measures engineering maturity at the design-to-production interface, which is the main objective of this thesis. However, four specific limitations of DECRIS motivate the development of an adapted framework:

Temporal gap. DECRIS was designed to be applied at the start of fabrication (steel cutting), when detail engineering is already well advanced, measuring whether engineering is sufficiently complete to start construction. In the De Vries Group context, the critical interface is earlier: between *Definitief Ontwerp* (DO) and Detail Engineering (DE), where the decision to release engineering information to the hull yard is made. This temporal shift requires different elements and criteria.

Single-dimension assessment. DECRIS evaluates each element with a single score that does not distinguish between geometric completeness and data reliability. As discussed in Section 2.4, the FEED MATRS precedent shows that these dimensions should be assessed separately. In concurrent engineering, where 3D models may be geometrically complete but based on provisional data, this distinction is critical.

Contextual specificity. DECRIS was developed for offshore EPC megaprojects with large multidisciplinary teams and standardised industrial processes. The superyacht ETO environment differs in scale (single vessel vs industrial plant), customisation (fully bespoke vs partially standardised) and organisational structure (design, hull construction and outfitting split across the companies of one group vs a single EPC contractor). The 69 elements of DECRIS cannot be directly transferred; they require adaptation, consolidation and contextualisation for ETO hull construction.

Governance integration. DECRIS was conceived as a benchmarking and classification tool: it produces a score that categorises projects into maturity groups for comparative analysis. It was not designed as an operational decision-making instrument at a quality gate that decides what happens with each item that remains open, with assigned responsibilities and escalation mechanisms. Integrating the maturity assessment into the organisational governance structure (who evaluates, who decides, what happens with the result) is necessary for the framework to function as a management tool rather than an academic exercise. PDRI and DECRIS do support go or no-go decisions at project level. What they leave open is the local decision on each remaining item, with its condition and owner on record.

2.6 LOD: how far a model can be trusted

The Level of Development (LOD) Specification, maintained by the BIMForum [32], provides a standardized framework for describing the geometric and informational content of Building Information Model (BIM) elements at progressive stages of design and construction. Originally developed for the Architecture, Engineering and Construction (AEC) industry, the LOD framework defines six levels. The five below describe the design, while LOD 500 records the as-built state:

- LOD 100: conceptual representation, approximate geometry, no reliable dimensional data.
- LOD 200: approximate geometry with generalised dimensions, sufficient for coordination but not for fabrication.
- LOD 300: accurate geometry whose size, shape, location and orientation can be measured as designed, though not necessarily free of clashes.
- LOD 350: accurate geometry with connections, supports and interfaces to adjacent elements modelled, sufficient for interdisciplinary coordination.
- LOD 400: fabrication-ready geometry with complete detailing, confirmed dimensions and all necessary data for manufacturing.

The LOD framework establishes a clear and intuitive progression from conceptual placeholder to production-ready model. Its strength lies in its ability to set explicit expectations for what a model element should contain at each project stage, providing a shared reference language between designers, engineers and fabricators.

However, LOD was developed for BIM in building construction, not for 3D product models in shipbuilding. The shipbuilding modelling environment differs in several ways: shipbuilding uses product-model software rather than BIM software, the relationship between model elements and physical assemblies follows naval architecture conventions (blocks, sections, zones) rather than building conventions (floors, rooms, systems), and the integration of vendor-specific equipment data follows procurement processes specific to maritime supply chains.

Despite these differences, the conceptual progression from LOD 100 to LOD 400 provides a solid foundation for defining what "maturity" means in geometric terms. Adapting the LOD framework to ETO shipbuilding requires translating the level definitions to the specific modelling practices and deliverable conventions of the shipyard, while preserving the fundamental principle: that model elements should be assessed not just for their presence, but for their level of geometric definition relative to what production requires.

2.7 TRL: how ready a technology is

The Technology Readiness Level (TRL) framework was originally developed by NASA [34] and subsequently standardised by ISO [39] for space hardware. TRL provides a nine-level scale that assesses the maturity of a technology from basic principles observed (TRL 1) to flight-proven system (TRL 9). The framework has been widely adopted across aerospace, defence and energy sectors as a standard language for communicating technology risk.

In the context of ETO shipbuilding, TRL becomes relevant when owners request the integration of novel or unproven technologies, referred to as "Specials." Examples include alternative fuel systems (methanol, hydrogen), experimental propulsion configurations or novel structural materials. These elements introduce a specific category of risk: even when the engineering information appears complete (high geometric maturity, confirmed supplier data), the absence of operational validation in the specific application context means that the specifications may prove inadequate during production or commissioning.

ISO 16290 states that it is "applicable primarily to space system hardware, although the definitions could be used in a wider domain in many cases" [39], acknowledging the need for contextual adaptation. In shipbuilding, a TRL assessment would evaluate whether a technology has been demonstrated in a relevant maritime environment (TRL 6-7) or has proven operational history in comparable vessels (TRL 8-9).

The key insight from TRL for maturity assessment is that data reliability depends not only on whether information has been provided by a supplier (a procurement question) but also on whether that information has been validated in the intended operational context (a technology question). This distinction is particularly relevant for elements involving technologies without established track records in the specific application domain.

2.8 Deciding with preliminary information

Sections 2.3 to 2.7 answer the first question of this chapter: how the maturity of information can be assessed. They do not answer the second: what to do with the items that are still open when the gate closes. In concurrent engineering that is the normal situation, so it needs its own basis.

Krishnan, Eppinger and Whitney [35] studied exactly this question: when can downstream work start on upstream information that is still preliminary? Their answer rests on two properties. The first is how quickly the upstream information settles on its final value, which they call evolution. The second is how much downstream work has to be redone if that information changes, which they call sensitivity. Overlapping is reasonable when the information settles early or when the downstream work is not very sensitive to it. It becomes risky when both are unfavourable.

Clarkson, Simons and Eckert [36] add a second point. In complex products a change to one part usually leads to changes in other parts, so the cost of a change depends on how far it travels through connected parts. A small change in a heavily connected element can spread across the design, while a large change in an isolated element may stay local.

Risk engineering offers a simple way to bring these ideas together. Criticality analysis, as used in Failure Mode, Effects and Criticality Analysis (FMECA, standardised in IEC 60812 [46]), ranks each risk by combining the severity of its consequence with how likely it is to occur, so that attention goes first where both are high. And the Stage-Gate literature [44] already treats the gate as the place where such judgements become a documented decision.

Krishnan's model is not the only answer to this question. Set-based design keeps several options open and narrows them only as information firms up, so nothing is committed too early [48]. The Last Planner System, used in construction, releases a piece of work only when everything it needs, including information, is in place [49]. Both work well, but they ask the whole chain to plan in a different way, which a single release meeting cannot change. Krishnan's logic, by contrast, can be applied to each open item at the gate as it stands.

Implication for this research: an open item at the gate should not only be counted, it should be decided. The decision should depend on how many disciplines rely on the item, how far a change would travel if its information moves and how long the item is likely to stay open. These principles come from the literature. The way they are turned into categories, a decision matrix and concrete decisions for the Slob context is a design choice of this thesis. It is developed in Chapter 4 and tested in Chapter 6.

2.9 Stage-Gate governance in product development

The Stage-Gate model, formalized by Cooper [44] and by Cooper and Edgett [45], structures the product development process into a series of stages separated by gates. Each gate serves as a decision point where the project is evaluated against predefined criteria before being allowed to proceed. The gate decisions are typically GO (proceed to next stage), KILL (terminate the project) or HOLD (pause and revisit).

Cooper's original model was designed for new product development in industries with relatively standardized processes. Its application to ETO environments introduces specific challenges. In ETO, projects are unique and the criteria for "readiness" at each gate cannot be standardized across projects in the same way as in Make-to-Stock industries. Furthermore, the schedule pressure inherent in ETO environments creates strong incentives to pass gates prematurely: delaying a gate decision delays the entire project, making it difficult for gate committees to enforce rigorous criteria.

The literature documents a well-known failure mode of Stage-Gate systems in practice: gates become administrative checkpoints rather than quality filters [44]. Gate criteria transform into checklists that verify the existence of deliverables rather than their content quality. Under schedule pressure, decision-makers learn to satisfy the checklist without necessarily ensuring that the underlying work is mature. In a study of 211 businesses, about half held gate meetings that worked as information sessions or project updates rather than decisions [45].

The main consequence is that gates fail to perform their intended function: filtering out immature work before it enters the next phase. In ETO environments where downstream phases involve production-oriented DE, WP and subsequent hull construction with irreversible commitments (steel cutting, welding), the failure of a gate to filter immature information has direct financial consequences.

Cooper himself later loosened the gate. In his third-generation process a gate can give a conditional go: the project moves on under conditions that must be met later [47]. Within the stage-gate literature, this comes closest to releasing work under explicit conditions, which is what happens at Gate 4 in practice.

Implication for this research: effective gate governance in ETO environments requires criteria that assess content maturity, not just deliverable existence. The criteria must be verifiable and resistant to gaming. These requirements define the design principles for any maturity assessment framework at a quality gate.

2.10 Comparison of alternatives and research gaps

Table 2.2 brings the methods of this chapter together per function. For each function it shows which alternatives exist, what they do well and what they leave open at the release studied here.

Function	Existing alternatives	What they do well	What they leave open at the DOII to DE release
Assess definition before a phase decision	PDRI [8], FEED MATRS [38], DECRIS [6]	Element lists, scored with evidence	Weighted sums let strong elements hide weak ones. Made for front-end or EPC phases.
Separate complete from reliable	FEED MATRS [38]	Shows that defined is not confirmed	Rates a whole project, not a discipline.
Describe the maturity of an element	LOD [32], TRL [34], [39]	Shared vocabulary per element	Descriptive: no stability, no decision.
Decide on preliminary information	Overlapping model [35], set-based design [48], Last Planner [49]	When downstream work can start	No ready procedure for a gate.
Judge how far a change travels	Change propagation [36]	Predicts knock-on effects	Needs data not available at a review.
Prioritise open items	FMECA criticality [46]	Ranks by severity and likelihood	Multiplying ordinal scores can mislead [50], [51].
Govern the gate	Stage-Gate [44], [45], conditional gates [47]	A decision with its rationale	Decides on the whole project, not per item.

Table 2.2: Existing methods per function and what they leave open at the DOII to DE release.

No single method covers all functions at this release. The most complete precedents, DECRIS and FEED MATRS, stop at a score, and none of them turns what remains open into a recorded decision per item. It is one of the four gaps below, which together set what the solution has to fill.

The review of existing maturity assessment frameworks reveals a clear progression in the field: from scope definition assessment at the planning stage (PDRI), through dual-dimension evaluation separating maturity from accuracy (FEED MATRS), to engineering completeness assessment at the construction start point (DECRIS), complemented by geometric standardization (LOD) and technology risk evaluation (TRL).

Each method covers part of the problem, but none covers the release from DOII into detail engineering in custom shipbuilding as a whole. Table 2.3 turns what is missing into four requirements.

Gap	Description	Identified From	Requirement for Adapted Framework
Temporal	Existing frameworks measure either too early (PDRI: scope definition) or too late (DECRIS: start of fabrication). No framework measures at the DO-DE interface where the decision to release information to production is made.	PDRI, DECRIS	Temporal positioning at the DO-DE transition point (Gate 4), measuring accumulated maturity of upstream deliverables before DE begins.
Dimensional	DECRIS uses a single-dimension assessment that does not distinguish between geometric completeness and data reliability. FEED MATRS shows that maturity and accuracy should be assessed separately.	DECRIS, FEED MATRS	Separate assessment of completeness and stability, combined with the criticality of what remains open. The weakest dimension sets the result, as a conservative design choice.
Contextual	All reviewed frameworks were developed for industries other than ETO shipbuilding (civil construction, offshore EPC, aerospace). The specific organisational structure, modelling practices and production processes of custom shipbuilding require adaptation.	PDRI, DECRIS, LOD, TRL	Elements, scales and weights adapted to the specific modelling, procurement, and production conventions of ETO hull construction.
Governance	Existing frameworks support phase decisions at project level, but none turns the items that remain open into local decisions with a condition and an owner on record.	DECRIS, PDRI	Integration into organisational governance with verifiable criteria, assigned responsibilities and a graded decision per open item.

Table 2.3: Research gaps identified in the state of the art and the resulting requirements.

These four gaps define the contribution space for this research. The development of a framework that addresses all four gaps simultaneously, measuring at the correct temporal point, with separate dimensions for completeness and stability, adapted to the ETO shipbuilding context, and integrated into operational governance, represents the core technical contribution of this thesis. The design and the retrospective evaluation of this framework are presented in Chapters 4 and 6.

2.11 Synthesis and answer to SQ1

"What characteristics of engineering information determine its readiness for production release in complex, multi-disciplinary ETO construction environments?"

The review leads to three conclusions. Information immaturity at phase transitions is a structural feature of ETO projects under concurrent engineering, not the failure of one organisation [6], [8], [13], [15]. Existing instruments show that maturity can be assessed with evidence before a phase decision [6], [8] and that completeness and reliability are better assessed separately [38]. None of them sits at the DO to DE transition of custom shipbuilding or turns the items that remain open into recorded decisions.

This answers SQ1. Readiness for release is not one property but three, judged per discipline: how complete the information is, how stable its basis is and how critical the remaining open items are. Around them sits the context of the package, such as new technology, a pending analysis or an owner decision, which can move everything at once. Chapter 3 turns these conclusions into requirements and selects the elements of a solution.

SOLUTION DESIGN

This chapter and the next answer the second sub-question of the thesis (SQ2):

> *"How can these characteristics be measured and integrated into a quantitative, verifiable assessment framework at a quality gate?"*

This chapter turns the conclusions of Chapter 2 into requirements (Section 3.1) and selects the elements of the solution (Section 3.2). Section 3.3 explains how Cynefin is used as a design lens, and Section 3.4 leads to the framework that Chapter 4 describes.

3.1 Requirements for the solution

Four requirements follow from the gaps of Section 2.10. The framework must assess at Gate 4. It must assess completeness and stability separately, together with the criticality of what remains open. It must adapt its elements to hull construction. And it must connect the assessment to a decision per open item, with a condition and an owner.

3.2 Selecting the elements of the solution

Section 2.10 identified gaps in four dimensions: no existing method assesses maturity (1) at this point of the chain, (2) separating what is done from what is reliable, (3) in the naval ETO context, and (4) as a decision mechanism at a gate. The solution is built by taking from each method what works and changing what does not, as Table 3.1 shows.

Framework (Ch. 2)	What is used	What is changed
PDRI	Scoring with a rubric and with evidence before a phase decision.	The moment: PDRI is used at the front end of a project. Here it is needed at the close of basic engineering, which is where it hurts.
FEED MATRS	Separating "how much is done" from "how reliable it is".	How the scores combine: FEED MATRS shows the two side by side, here the worst one rules. And a third dimension is added, connecting the profile to the items.
DECRIS	Assessing maturity just before releasing to construction, with viability bands.	Its four shortcomings: it arrives late, measures one thing, comes from offshore platforms, and describes without deciding.
LOD (BIM)	The progression of how developed the geometry is.	LOD says nothing about the data holding that geometry up. Here it is completed by the other dimensions.
TRL	Treating new technology as a risk of its own kind.	Instead of another scale, it enters as a context condition (A7).
Change propagation	A change to one part can spread to connected parts beyond its original scope.	Instead of full network models, three simple levels assessable per item (the impact dimension).

Table 3.1: What the framework takes from each method of the state of the art, and what it changes.

Three choices complete the selection. Because the dimensions are ordinal, multiplying them would give a numerical result with no defensible measurement meaning [50], [52]. The framework therefore uses the minimum as a conservative design rule: a strong dimension cannot compensate for a weak one. Each open item receives a local decision rather than reducing the whole delivery to a go or no-go. This allows unaffected work to continue while uncertainty is contained locally, and Chapter 5 later shows that this principle is consistent with the yard's existing practice of isolating unstable zones. Finally, the decision is built into the existing release review rather than into a parallel process. Cooper's conditional gates provide a precedent for proceeding under explicit conditions [47], while the exact integration into the review is a contextual design choice. Together, these choices favour rules that can be applied with the information available at the gate and without replacing the existing process.

3.3 Cynefin: where judgement is needed

Cynefin, proposed by Snowden and Boone [42], is a sense-making framework that sorts decision contexts by the nature of the relation between cause and effect. In their formulation there are five contexts. In a *simple context* the relation is obvious. In a *complicated* one it exists but has to be found through analysis and expertise. In a *complex* one it can only be understood in retrospect. In a *chaotic* one no relation can be perceived and action has to come first. Disorder is the state of not knowing which of the others applies. Later versions of the framework call the simple context clear, the term used for the lowest level of resolution difficulty in Chapter 4.

The release decision at Gate 4 combines features of the complicated and the complex context. Part of it can be analysed: whether a foundation is modelled, whether a supplier has confirmed a footprint, whether a class comment is still open. Another part cannot be settled in advance: how late owner choices, supplier and class revisions will interact once DE is under way. When nothing separates the two, the analysable part tends to be judged as the uncertain one.

This is the sense in which Cynefin is used in this thesis: as a design lens, not as a classification established for the process studied. The framework does not remove complexity and does not claim to move the decision from one context to another. It separates the part of the decision that can rest on evidence from the part that cannot. The first is scored against explicit criteria. The second is not hidden: it is handled through conditional decisions, isolated zones and an accepted risk that is written down, while expert judgement keeps the final word.

3.4 Towards the framework

The elements of Table 3.1 come together in one framework, the Gate Assessment of Information Availability (GAIA). Chapter 4 describes it in steps and formulas. Chapter 5 introduces the yard where it is applied and confirms the requirements with the yard's own evidence. Chapter 6 applies GAIA retrospectively to three completed deliveries, Chapter 7 discusses whether it is worth its cost and Chapter 8 answers the research questions.

The intended contribution remains the one proposed in the literature review: the codification of maturity criteria at the design-to-production interface of custom shipbuilding, an assessment that keeps completeness and stability apart, and an empirical check of whether its output is consistent with what happened in practice.

THE GAIA FRAMEWORK

This chapter completes the answer to the second sub-question of the thesis (SQ2), begun in Chapter 3:

> *"How can these characteristics be measured and integrated into a quantitative, verifiable assessment framework at a quality gate?"*

The answer this thesis gives is GAIA, the Gate Assessment of Information Availability. Its starting point is a fact of the yard: every delivery crosses Gate 4 with open items and today the risk they carry is accepted without anyone writing down what was accepted, by whom or on what basis. GAIA does not try to stop that crossing. It turns it into a structured decision, built from the elements selected in Chapter 3, by asking three questions before the package is released into detail engineering:

- **Level I.** What around the package could still move it? Seven context preconditions, read on dated evidence.
- **Level II.** Where is the package fragile? A maturity profile per discipline, in which the weakest dimension rules.
- **Level III.** What do we do with each open item? One of four decisions per item: GO, CONDITIONAL, ISOLATE or WAIT.

The answers end in a release record: what proceeds, under which condition, by when, owned by whom and with which risk accepted. In practice this takes five steps, shown on one page in Figure 4.1: freeze the record of the delivery as it stands before the DOII Final Review, check the context, profile the package, decide per item, and record the decisions and follow them until they close. The rest of the chapter explains the rules that make these steps reproducible.

The three levels go from broad to specific and each one needs more detail than the one before. Level I looks at the whole delivery and needs only a dated status for each condition. Level II looks at each discipline: how complete it is against the integration list and how stable its sources are. Level III looks at each open item and needs the most: what depends on it, how far a change would travel and how hard it is to close. Each level answers a question that the other levels cannot, which is why GAIA keeps three levels instead of one score (Section 4.1). The parts are borrowed from existing models: Level I draws on PDRI [8], DECRIS [6] and TRL [34], [39], Level II on FEED MATRS [38], PDRI [8], LOD [32] and criticality analysis [46], and Level III on overlapping [35], change propagation [36], Cynefin [42] and Stage-Gate [44], [47]. Annex 4-A traces every element to its source.

Six terms from the yard's own vocabulary are used throughout the chapter:

- **Delivery (Kickoff, KO):** A large block of the vessel (bow, engine room, superstructure). Each delivery goes through its own DOII and its own release into detail engineering.
- **Functional zone:** A coherent area within a delivery (for example, the engine room).
- **Base / Features / Outfitting / Review:** The four windows of approximately two weeks inside the DOII of each delivery.
- **Integration list:** The internal document that defines, per phase and per discipline, which elements must be integrated in the model.
- **Open item:** Work known to be incomplete or unstable at release, recorded in the close-out list of the delivery.
- **100% Review and Final Review:** The two reviews that close DOII. At the 100% Review the model should be complete, and the Final Review releases the delivery to DE. The DOII Final Review is Gate 4. Detail engineering ends with reviews of the same name, which this thesis always calls DE 100% Review and DE Final Review.

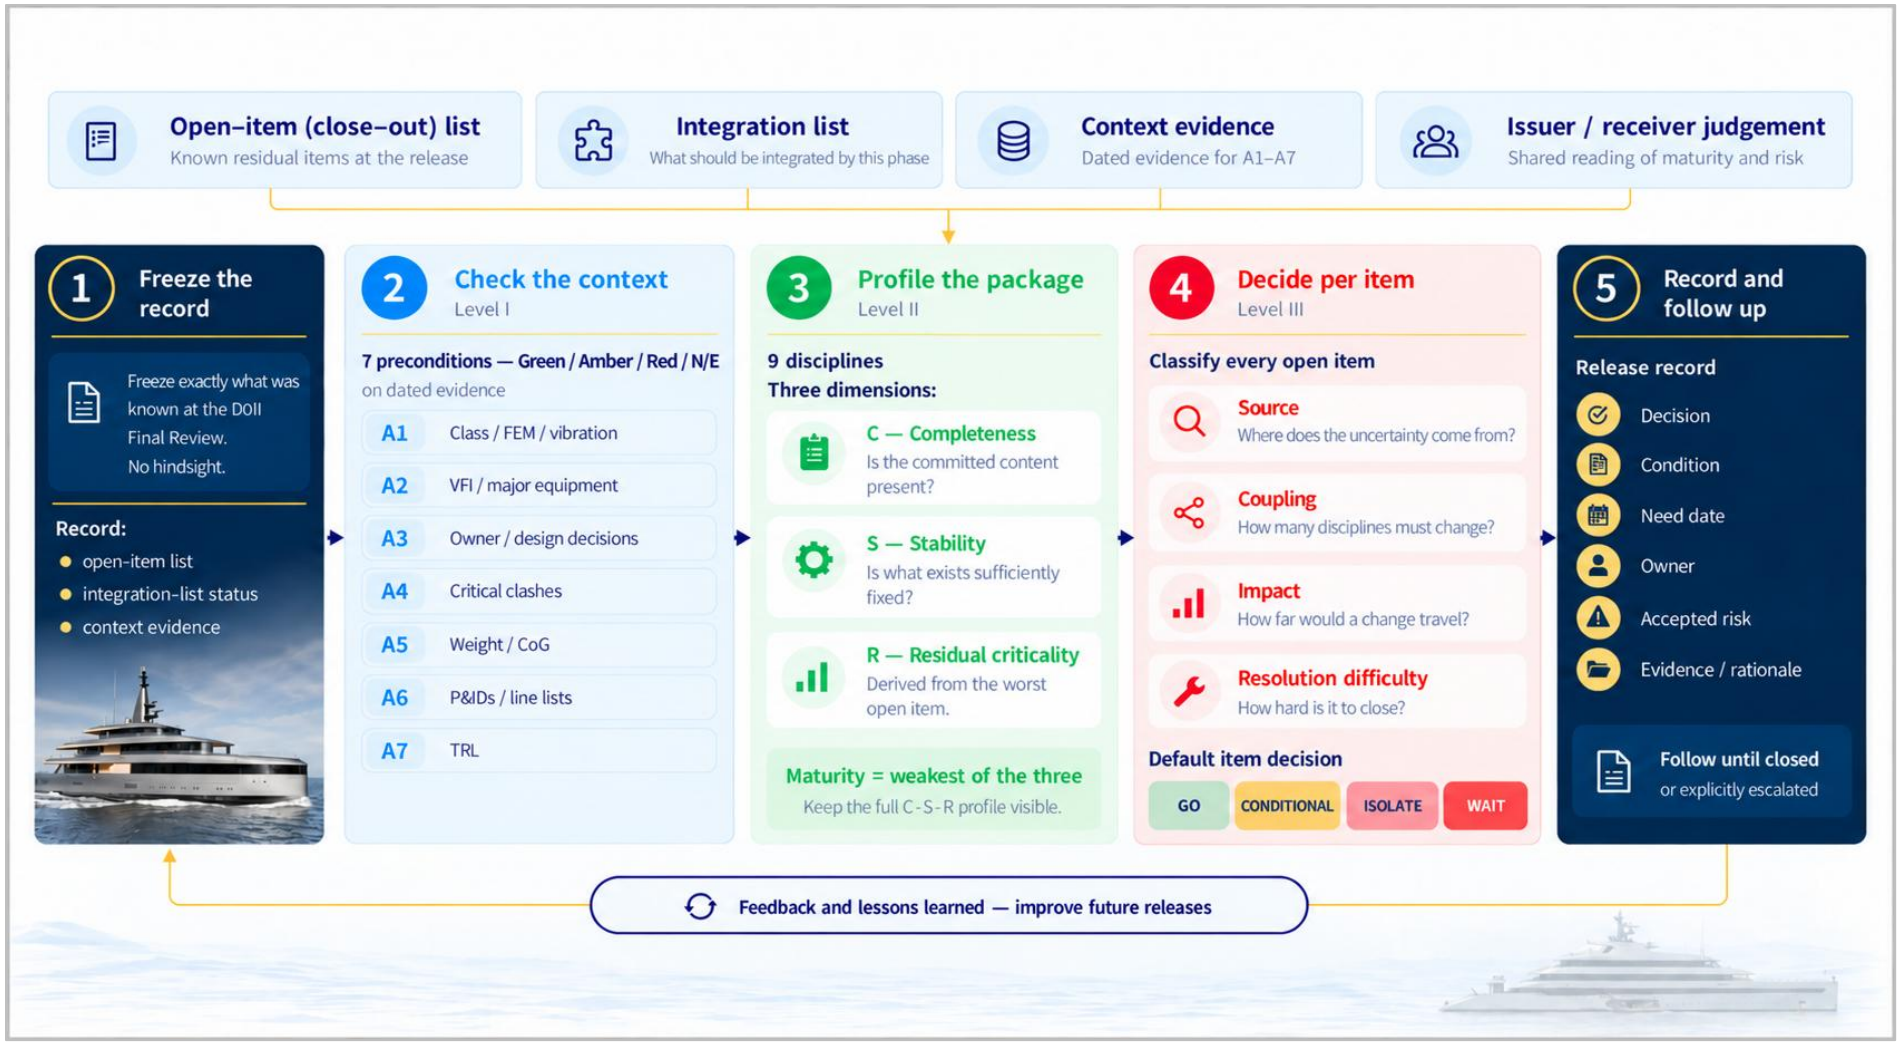


Figure 4.1: The GAIA framework: inputs, five steps and feedback. AI-assisted illustration.

Two things GAIA is not. It is not a predictive model: its scores are ordinal and each one must rest on evidence that can be shown. And it does not ask for every item to be closed before the release. The risk is accepted today anyway. The aim is that it is accepted with open eyes, bounded and with an owner, instead of being absorbed months later by the yard in hours and recovery work.

4.1 Requirements and design choices

The four requirements of Section 3.1 come from the state of the art. During the iterative design they were made operational against the yard's reality, which gives the nine requirements of Table 4.1. Chapter 5 brings the internal evidence behind each of them. The numbers in brackets refer to the four requirements of Section 3.1: (1) assess at Gate 4, (2) assess completeness and stability apart, (3) adapt to hull construction, (4) decide per open item.

Req.	The framework must...	Ch. 3 requirement (X) · Evidence (Ch. 5)
R1	Allow deciding item by item and zone by zone, not a yes/no on the whole delivery	(4) · 5.5.2: the yard has to start anyway, the real decision is what is released, watched or set aside
R2	Accept that some information cannot exist yet (FEM, vibration) and record it as open risk	(1), (2) · 5.6: FEM and vibration need a detailed model that only exists during DE
R3	Distinguish what is modelled from what is confirmed	(2) · 5.4.3, waste 6: information that exists but will change
R4	Look at the delivery context, each discipline and each item at the same time	(2), (4) · 5.5.1, 5.5.4, 5.5.5: problems at context, discipline and item level
R5	Record the real status of what depends on others: suppliers, owner, class, other organisations	(2) · 5.5.5: EXT and WTB depend on decisions the chain does not control
R6	Build on the lists and reviews that already exist, without creating new paperwork	(3) · 5.5.1, 5.5.3: lists, reviews and boxing already work, the assessment is missing
R7	Leave written, for every accepted item: why, under what condition and what risk is being taken	(4) · 5.7, finding 5: risk is accepted without a record of who accepts it or why
R8	Leave a record that can later be checked against hours, change orders, the logbook and failure costs	(4) · 5.2 and 5.6: labour is 80–85% of hull cost and is recorded per section and discipline
R9	Make a shared definition of "ready" explicit across the organisations	(1) · 5.4.3, waste 2: no shared operational definition of ready

Table 4.1: Operational design requirements.

Four design choices follow from these requirements. Each one rejects an obvious alternative.

Three levels, not one score. The diagnosis (Chapter 5) found problems at three heights: risks in the context of the delivery that travel as rumour, disciplines that reach the gate with very different maturity and single items that cross without anyone accepting them. One readiness score, as in DECRIS, would merge the three and answer none of them. GAIA keeps a screening, a profile and a decision.

Built on what already exists. Open-item lists, integration reviews and the boxing of unstable zones already work. A parallel tool would compete with them and lose. GAIA adds fields to the list people already keep and uses the reviews that already take place. It creates no new documents and no new meetings.

The weakest dimension rules. The most dangerous case in the diagnosis is information that exists but will change: it looks complete and it is not stable. A single score would rate it high. GAIA separates completeness from stability, as FEED MATRS separates maturity from accuracy and lets the lowest dimension set the maturity, because releasing something that was not ready costs steel while holding back something that was ready costs days.

A decision per item. The DOII window closes on the calendar whether the information is mature or not, so the real decision was never a yes or a no on the whole delivery. It is what proceeds, what is watched and what is set aside. GAIA puts the decision where it already happens in practice, at the item, and adds what stage-gate theory asks of a good gate: the reasoning behind it.

The three levels are linked. The classified items of Level III feed the discipline profile of Level II, and Level I tells the panel which parts of the delivery need a second look. Each level also works on its own unit. The delivery is what gets released, because the Final Review is held per delivery. The open item is what gets a decision. The functional zone is what gets a maturity profile: when a delivery mixes areas of a different nature, such as an engine room with mooring and anchoring zones, Level II is scored per zone, so that the engine room is not graded against a different risk profile. Items on the boundary of the delivery are assessed with it, because they condition its release. Annex 4-A traces every element of the framework back to the literature.

Throughout, GAIA structures expert judgement instead of replacing it. The people who decide today keep deciding. What changes is that each decision is written down with its reason, its condition and the risk it accepts.

4.2 Level I: the context of the delivery

Some risks live in no item. A pending vibration analysis, or an owner reconsidering a pool on another deck, conditions the reading of every item in the delivery, yet no line on the open-item list captures it. In one of the cases of Chapter 6 the strongest signal at the gate was exactly that: a pending vibration analysis that appeared on no list. Level I makes this context visible through seven preconditions (Table 4.2), structured after the front-end readiness categories of PDRI and the technology screening of TRL.

Precondition	What it checks	Why it is there
A1 Class / FEM / Vibration	Class approval status, strength calculations and vibration topics, for the scope of the delivery. State = the worst of the three.	Three separate structural checks, by approval or calculation, that often arrive after the close and can force changes to steel already released.
A2 VFI / Major equipment	Written supplier confirmation for the equipment that fixes geometry (foundations, clearances).	Equipment without confirmed geometry leaves its foundation hanging. The typical origin of foundation items.
A3 Owner / design decisions	Owner decisions and design loops closed for the scope of the delivery.	Every pending owner choice is a future source of open items.
A4 Critical clashes	No unresolved high-impact clashes in primary structure or major equipment.	An open structural clash means that integration is not done.
A5 Weight / Centre of gravity	Weight estimate within tolerance for the vessel.	A late weight change touches structure and hull form, the most expensive things to change.
A6 P&IDs / line lists	System schematics and routings frozen for the delivery.	When schematics move, penetrations and holes in the structure move with them.
A7 TRL (technology readiness)	No unproven technology in the delivery. If there is, an explicit maturation plan.	What has never been built carries uncertainty that the model does not show.

Table 4.2: The seven Level I preconditions.

Each precondition is rated Green, Amber or Red on dated evidence. The line between Amber and Red is almost always a verifiable fact: there is a committed date or there is not, there is an agreed solution or there is not. Where no dated evidence exists, the state is recorded as not evaluable (N/E) instead of being guessed. The full rubric is given in Annex 4-B.

The colours describe the context. They are not decisions, which belong to the items at Level III. A Red does not stop the delivery either. It flags the scope and items it can affect, which the panel then reviews explicitly at Level III while the rest of the delivery continues.

One rule separates context from content. As long as an uncertainty is not located in specific items, it belongs to Level I. Once it becomes concrete, it moves to Level III as items, and the precondition can return to Green while those items stay open. An owner who wants different engine-room stairs is content: one item, with the owner as its source. An owner reconsidering a pool on another deck, with possible consequences for weight and generator size, is context: A3, and A5 if the weight moves.

Three candidates were left out on purpose. Engineering capacity and planning pressure describe resources and calendar, not the state of the information, and capacity returns in Chapter 6 as a rival explanation. Hours against budget were excluded because they are a lagging indicator: useful to validate afterwards, useless to decide at the gate.

4.3 Level II: where the package is fragile

Level II gives the maturity profile of the zone: one score per discipline, so that the panel sees which disciplines cross the gate solid and which do not. Today nobody has that view.

The disciplines come from the DOII integration list, the document that already says what each discipline must integrate in each phase. Scoring against it means scoring against a commitment the organisation has already made, not against a standard invented for this thesis. Its four families are refined where the risk differs. Structure is split into primary steel, which conditions class and hull form, and secondary steel. Foundations and major equipment form a discipline of their own, because that is where structure meets supplier data. HVAC and process piping are split from mechanical outfitting, because they create the penetrations that keep appearing on the item lists. Electrical covers cable trays and their penetrations, and class approvals follow the delivery's own class drawings. The result is nine disciplines: primary steel, secondary steel, foundations and major equipment, exterior outfitting, interior interfaces, electrical, HVAC, process piping and class approvals. The list is the same for every zone, only the relevance of each discipline changes. Manufacturability was considered as a tenth discipline and left out, because it could not be anchored to objective evidence.

Each discipline is scored on three dimensions:

- **Completeness.** Is everything the integration list commits for this phase modelled and integrated? It is measured against the list, not against the items, because a missing element that nobody noticed never becomes an item. A missing structuring element caps the score at 2.
- **Stability.** Does what exists rest on confirmed sources? A source closes when the supplier confirms in writing, the owner decision is registered, the class approval is received or another organisation freezes its routing. A source that has already missed commitments on the project cannot score above 2, whatever date it now promises.
- **Residual criticality.** How serious is what remains open? It is not scored separately. Once the open items have been classified at Level III, it follows from the worst of them, combining its impact and resolution difficulty (Table 4.3). The mapping is inspired by the criticality matrices of risk engineering (FMECA, IEC 60812 [46]) but is a design rule of GAIA, not an FMECA calculation.

Worst open item of the discipline	Residual criticality
No open items	5
Low impact, with low or medium difficulty	4
Medium impact with low or medium difficulty, or low impact with high difficulty	3
High impact with low or medium difficulty, or medium impact with high difficulty	2
High impact with high difficulty, or extreme severity	1

Table 4.3: Residual criticality, derived from the worst open item of the discipline.

Completeness and stability are scored from 1 to 5 against the anchors of Table 4.4, and every score carries a short written reason that points at something that can be shown: the state of the model or specific items on the list. If two assessors disagree, they should be able to point at the exact piece of information they read differently.

Level	Completeness	Stability
1	Not started.	Key sources without a date or warning of changes.
2	Preliminary: concept only, with a structuring element of the scope missing.	Several key confirmations pending, some without a committed date.
3	Partially developed: recognisable model, with open items named and located.	Confirmations pending, all with a committed date that falls before the receiver's need date.
4	Substantially complete: only minor elements missing.	Almost everything confirmed; remaining items minor and bounded.
5	Complete: everything committed is in the model.	All sources confirmed in writing. Frozen.

Table 4.4: Anchors for completeness and stability.

The split answers the most dangerous case of the diagnosis. A foundation can be perfectly modelled on unconfirmed supplier data: completeness 4, stability 2. One score would rate it high and hide the risk. Three scores expose it. The maturity of the discipline is then the lowest of the three:

Discipline maturity = MIN(completeness, stability, residual criticality)

Three reasons support the minimum. Real maturity is limited by the weakest dimension: a complete model on unstable data is only as mature as its data. The cost of error is asymmetric: overrating maturity ends in rebuilt steel, underrating it costs days. And a weighted average would need weights that three cases cannot justify. A product, as in the risk priority number of FMEA, was rejected because the scores are ordinal: a 4 is better than a 2 but not twice as good [50], and products of ordinal scores are a known source of error in FMEA and in readiness

metrics [51], [52]. For the same reason the nine maturities are not averaged into a single index. The output of Level II is the profile itself, always reported with its three dimensions.

A low maturity is an alert, not a stop. A 2 in foundations means that a critical element is open, and what happens to it is decided at Level III, where the item behind the score can often be isolated while the rest of the delivery proceeds. Because the minimum is not the final decision, it can afford to be strict. It does not prevent overrating when evidence is missing, and a low score without later disruption is not automatically a false alarm: the risk may simply have been managed. Chapter 6 reads both situations with that in mind.

4.4 Level III: a decision for every open item

Level III is where GAIA decides. It takes the open-item list the delivery already has, with its usual fields, and adds a few columns to each line: the source of the uncertainty, the affected disciplines, the impact, the resolution difficulty, the decision, the condition with its need date, the accepted risk and, when the panel departs from the default, a justification. Only items still open at the gate receive a decision. Items closed during the review window show the review doing its job, and items raised after the gate are downstream evidence, not inputs.

Each item is classified on three scales (Table 4.5). Coordination coupling counts how many disciplines must change for the item to move and is High whenever the item needs an approval that crosses primary structure, class or major equipment. Impact describes how far a change would travel if the information moves after release, from the item alone to a chain reaction across zones, following the change propagation literature of Section 2.8. Resolution difficulty describes what is still needed to close the item, from one missing datum to a problem without an agreed solution, using the Cynefin distinction of Section 3.3. The scales are independent on purpose: coupling looks at the present, impact at the future and difficulty at the way out.

Dimension	Low	Medium	High
Coordination coupling	One discipline must change.	Two disciplines must change.	Three or more disciplines, or an approval that crosses primary structure, class or major equipment.
Impact	Affects the item itself and, at most, one other. Local, bounded rework.	Affects two or three items in a defined zone. Rework containable by isolating the zone.	Risk of a chain reaction across zones or disciplines (for example: foundation → stairs → escape route → class). The rework cannot be bounded in advance.
Resolution difficulty	Only a datum or a confirmation is missing, with a clear origin. When it arrives, the item closes. (A clear problem.)	The solution is known, but it takes work or an agreement between parties to execute it. (A complicated problem.)	There is no agreed solution. It has to be designed or negotiated before the item can close. (A complex problem.)

Table 4.5: Classification scales for open items.

The source of the item (supplier, owner, class, FEM or vibration, design office, coordination, review finding, purchasing or yard) plays no part in the decision. It records where the uncertainty comes from, so that over time it shows which origins cause the most trouble across projects.

Each item ends in one of four decisions (Table 4.6). Four, because the real decision is never binary, and because the yard's practice of boxing unstable zones had already shown that isolating is a legitimate decision. It only needed a name, a boundary and a record.

Decision	What it means
GO	Released. No material residual restriction identified. A "watch" flag can optionally be set if the item is worth a look in the weekly meetings, with no formal condition.
CONDITIONAL	Released only if something concrete is fulfilled by a date. What, who checks it and by when are recorded. If the date passes without evidence, the item escalates to ISOLATE or WAIT.
ISOLATE (BOX)	The delivery is released except a bounded zone, which is left untouched (boxing). The boundary, the information that unblocks it and its expected date are recorded. Only available when the zone can truly be bounded: a definable boundary, the rest can advance without mortgaging what is inside, and the unblocking information is identified.
WAIT	That part is not released. It requires sitting down and deciding before continuing. What is missing, and where it has to come from, are recorded.

Table 4.6: The four decisions per item.

In practice engineers use "wait" loosely for both WAIT and ISOLATE, and Chapter 6 keeps each decision as the engineer recorded it. CONDITIONAL and ISOLATE carry a condition with a need date: the date by which the information must arrive so that the affected work is not executed on assumptions. A condition closes only with released information, not with information issued for approval.

One field runs across all decisions: the accepted risk. For every item that proceeds, the consequence of the information changing after release is written down, with its author. This is the field that does not exist today, and the one that later allows each decision to be compared with what actually happened.

So that two assessors start from the same place, the classification gives a default decision (Figure 4.2). Coupling and impact set the base, and impact weighs more, because a change that cascades later is more dangerous than several disciplines watching it now. A High resolution difficulty moves the decision one step towards the restrictive end: an easy fix does not soften a decision, but not knowing how to solve an item makes its risk worse. An ISOLATE whose zone cannot be bounded becomes WAIT, a question the experts answer, not the matrix. Extreme severity, such as a change of hull form or unresolved FEM or class on primary structure, is WAIT by default, whatever the matrix says.

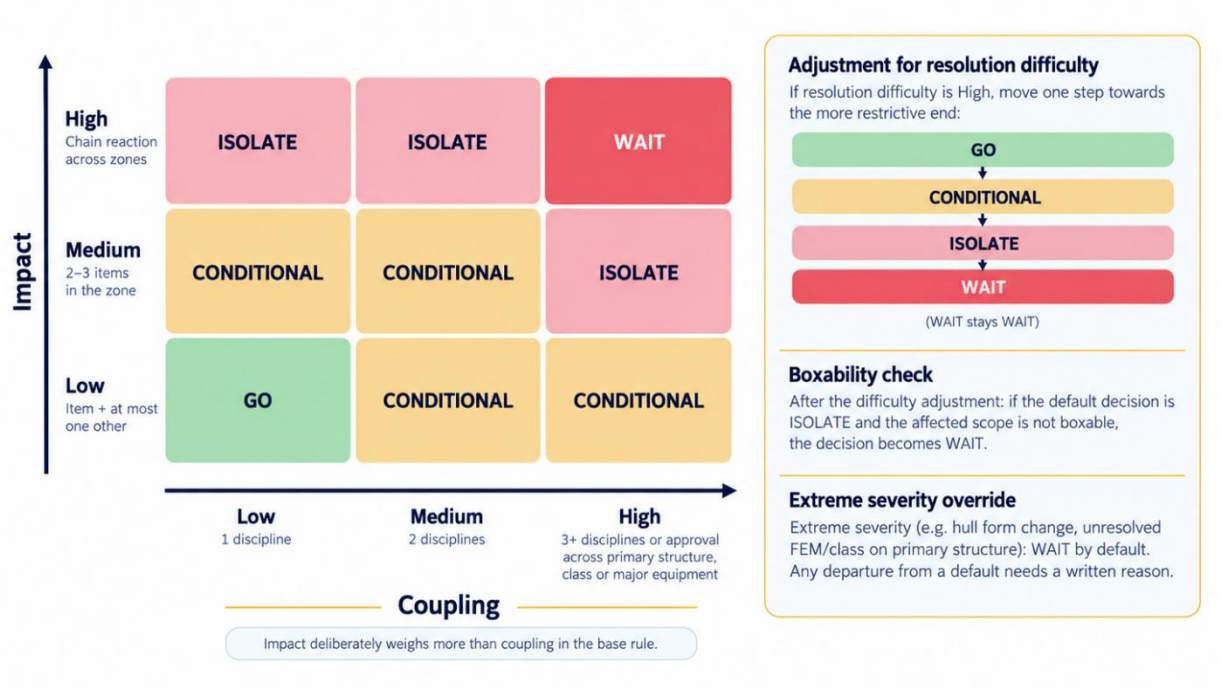


Figure 4.2: Default decision per open item. AI-assisted illustration.

An example from the reference case (Case A) shows the logic at work. The foundation of a major piece of equipment reaches the gate without written confirmation of its footprint. Coupling is Medium. If the preliminary geometry is expected to hold, impact is Low and the default is CONDITIONAL, with the written confirmation and its need date as the condition. If a change of footprint would propagate to adjacent structure and piping, impact is High and the default is ISOLATE. The yard confirms in the panel that the zone can be bounded, so the foundation is boxed and the rest of the delivery proceeds.

The default is a starting point, not a verdict. The panel can depart from it, provided the reason is written down.

4.4.1 Formal definition of the decision rules

The rules above can be summarised formally. The purpose is not to build a mathematical model but to make the framework reproducible: two assessors using the same evidence should reach the same default. For each discipline d, maturity is set by its weakest dimension:

$$M_d = \min(C_d, S_d, R_d) \quad (1)$$

where C_d is completeness, S_d stability and R_d residual criticality, each on the 1 to 5 scale. A strong score in one dimension cannot compensate for a weak one. Residual criticality is derived from the most critical open item of the discipline:

$$R_d = \min \{ r(I_i, D_i) : i \in O_d \} \quad (2)$$

where O_d is the set of residual open items of discipline d, I_i and D_i are the impact and resolution difficulty of item i, and r is the mapping of Table 4.3. A discipline without open items has $R_d = 5$.

The default decision for each open item follows in three steps, with the decisions ordered GO < CONDITIONAL < ISOLATE < WAIT:

Step 1, base decision. Coupling and impact set the starting decision through the matrix of Figure 4.2.

Step 2, difficulty adjustment. If resolution difficulty is High, the decision moves one step towards the restrictive end (GO → CONDITIONAL → ISOLATE → WAIT). WAIT stays WAIT.

Step 3, boxability check. If the result is ISOLATE but the affected scope cannot be bounded, the default becomes WAIT.

Extreme severity is WAIT regardless of these steps. Level I does not enter this logic: a Red precondition requires the affected scope and items to be reviewed explicitly, after which the same rules apply.

These are default rules, not automatic verdicts. The panel may deviate from the default when project-specific evidence justifies it, provided the reason is recorded. GAIA structures expert judgement, it does not replace it.

4.5 GAIA at the gate: who does what

GAIA invents no roles. The three-tier structure of every project already has them: the project leader (PL), the project manager engineering (PME) with a project engineer (PE) per discipline, and the organisations that receive the work. Figure 4.3 shows what each of them does before, at and after the DOII Final Review, and what GAIA adds to work they already do.

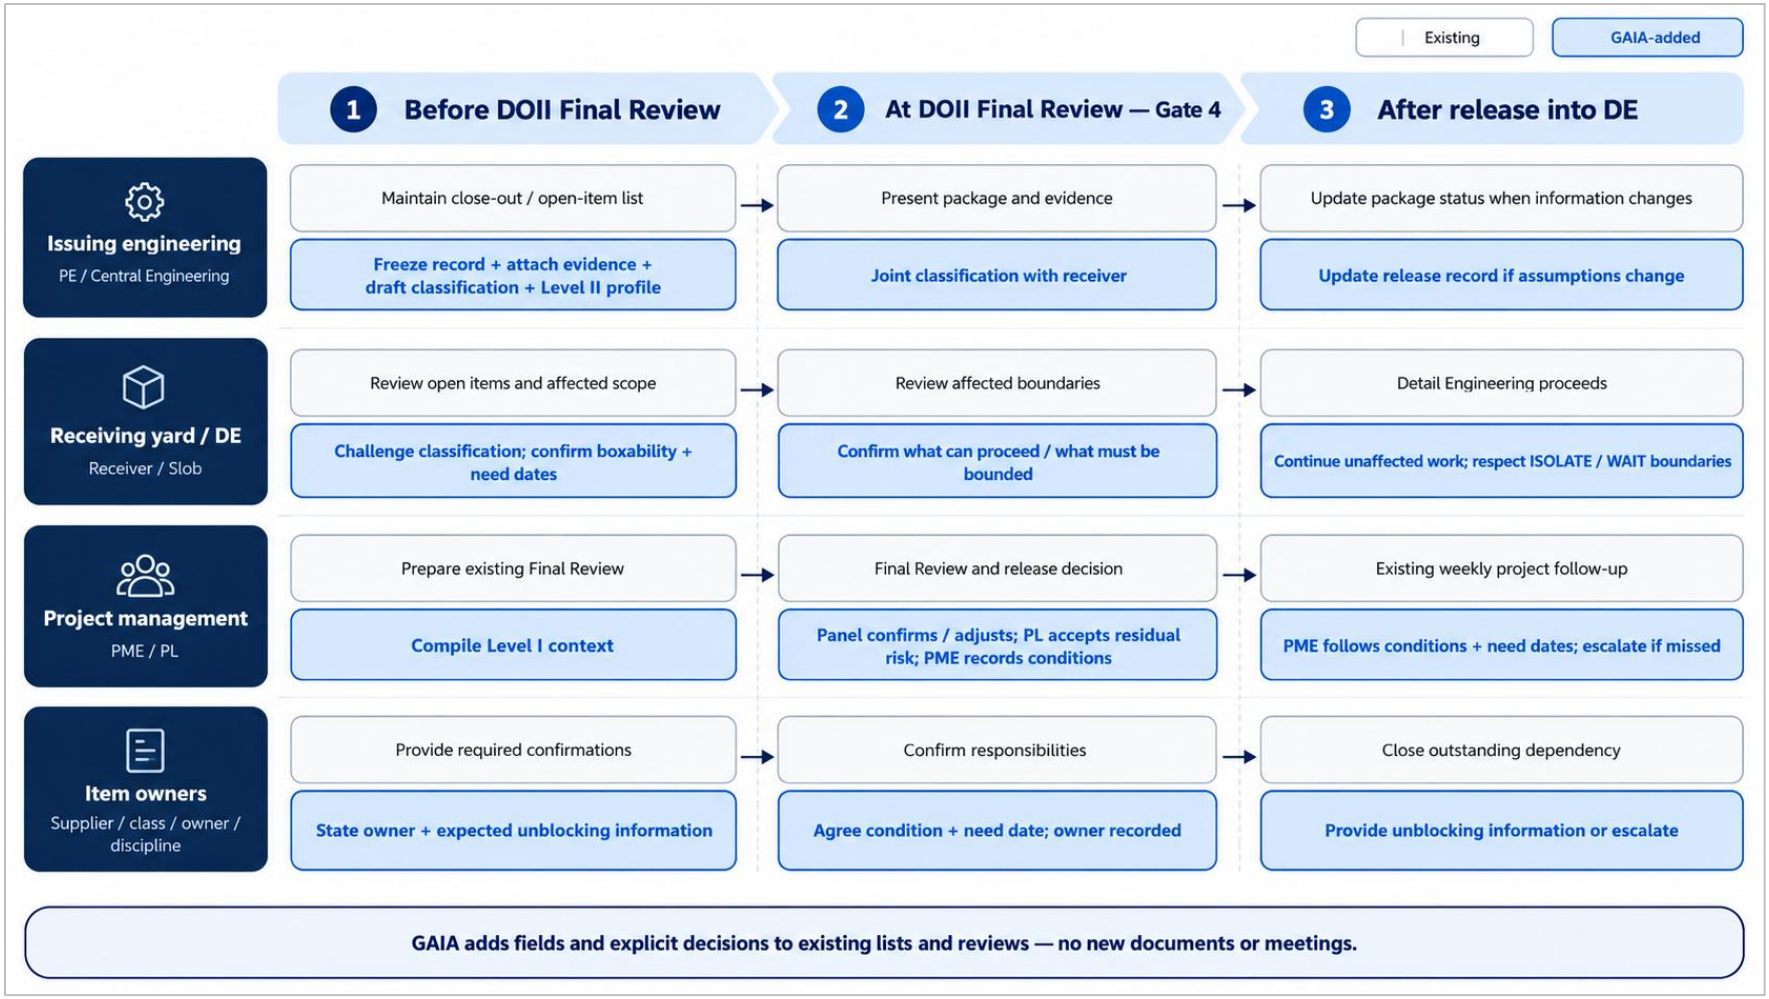


Figure 4.3: Gate 4 with GAIA by role and moment: existing activities (white) and GAIA additions (blue).

In short, each PE classifies the items of the discipline together with the counterpart on the receiving side, which builds a challenge into every classification. The Final Review panel, chaired by the PME, confirms or adjusts the classification and the decisions and the PME follows conditions and need dates in the weekly meetings. A Red precondition gets one named owner at its natural scope. The PL accepts the residual risk of the delivery in writing, because the PL is the role that sees all deliveries and organisations. Release authority does not move. What changes is the material the decision is made on: a classified list with its evidence instead of a flat list.

Two situations that are settled informally today get an explicit rule. If issuer and receiver disagree on a maturity or a risk, the disagreement is written into the item and the PL decides. If an accepted risk later materialises, the record shows who accepted it, under which condition and on what evidence. The point is not blame. It is that the cost of a decision can finally be traced to the decision itself.

Everything happens at moments that already exist: a context snapshot when the delivery starts, the discipline profile before the 100% Review, Levels I and III at the Final Review, and follow-up in the weekly meetings. The essential core is small: a context check and a decision per open item, recorded with its condition and owner. Annex 4-C compares the current process and the process with GAIA step by step, and Section 7.4 estimates the effort. The allocation of decision rights is a working version, discussed with project leadership and the yard's management team (Section 7.3).

4.6 Limits of the design

The limits of the design are stated before the validation, not after it. Four matter most:

- **Not calibrated.** The scales and the default logic are operationally defined but not empirically calibrated. Three cases cannot calibrate them.
- **Judgement remains.** Anchors, mandatory evidence, derived dimensions and default rules are designed to reduce the variability between assessors. They cannot remove it, and agreement between independent assessors has not been tested.
- **Some information matures late by nature.** A full vibration analysis is triggered by the contractual comfort criteria and, according to a structural engineer, can take about fifteen weeks from start to finish, so its results usually fall in DOII or later. GAIA cannot force them. It records them as open risk, which is the most a gate can do.
- **Follow-up is untested.** Conditions, need dates and escalation cannot be tested on closed projects. Chapter 6 tests the classification, not the follow-up, and for disciplines that are finished after the hull leaves the yard, the yard's records show less of what happened to them.

GAIA's structure comes from the literature and the diagnosis. Case A was its first retrospective application: it took more sessions and several versions of its case file, which helped to settle how the files were filled in, but GAIA was not changed to fit its results. The same levels, dimensions and decision rules were applied in the same way to the three cases (Cases A, B and C). Chapter 6 reads Case A as the reference case, a project with a good outcome.

4.7 Synthesis and answer to SQ2

"How can these characteristics be measured and integrated into a quantitative, verifiable assessment framework at a quality gate?"

This answers SQ2. GAIA measures the characteristics found in Chapter 2 and turns them into a decision. Seven preconditions read the context of the delivery. Nine disciplines, scored on completeness, stability and residual criticality, show where the package is fragile. Every open item ends in one of four decisions, with its condition, owner and accepted risk on record. The scores are ordinal and every one rests on evidence that can be shown. It does not ask for every item to be closed at Gate 4. The risk was being accepted anyway. With GAIA, it would be visible, bounded and owned before the yard absorbs it.

What no design can show about itself is whether its classification works: whether it separates the items that later hurt from those that did not. Chapter 5 first shows the yard's own evidence behind the requirements of GAIA. Chapter 6 then tests the classification on three cases: the reference case with a good outcome (Case A), a second engine room with the most complete records (Case B), and a stress case, chosen for its documented overrun, to see where the framework stops: a large deck structure in the superstructure (Case C). The test follows indicative consistency (SQ4): whether what GAIA flags at the gate matches what the evidence shows afterwards.

THE DESIGN-TO-HULL CHAIN AT SLOB

This chapter presents the empirical diagnosis of the design-to-hull chain and answers the third sub-question of this thesis (SQ3):

> *"What are the current causes of rework and inefficiency at the design-to-production interface in ETO shipbuilding, and how do planned information flows (Work-As-Imagined) diverge from the information actually delivered (Work-As-Done)?"*

This diagnosis precedes the retrospective analysis of the three deliveries in Chapter 6 and does not depend on its results.

The chapter moves from setting to findings. Section 5.1 sets out the evidence. Sections 5.2 and 5.3 describe the industrial context and the process on paper. Section 5.4 presents what the organisation's own VSM measured, and Section 5.5 how the release is handled in practice. Section 5.6 explains why Gate 4 is the point to intervene, and Section 5.7 condenses the diagnosis into five findings. The findings do not derive the framework, which was designed in Chapters 3 and 4. They test whether its requirements hold in the yard.

5.1 How the diagnosis was built

The diagnosis compares the process as it is designed on paper (Work-As-Imagined, WAI) with the process as it is carried out (Work-As-Done, WAD) [16]. The WAI side was reconstructed from company documents: the stage-gate definitions of the Operational Excellence (OE) programme and its generic gate plan, the working instructions of DOII, the integration list, the delivery planning and the close-out lists of the projects used in Chapter 6. The WAD side comes from interviews and working conversations held over the months of this research with professionals across the chain (Table 5.1), with open questions on how the release works in practice, what arrives late or changes, how open items are handled and where the effort ends up.

Code	Function in the chain
I1	Detail engineering — hull
I2	Detail engineering — superstructure
I3	Detail engineering — WTB
I4	Design-phase engineering (DD–DO)
I5	Chain coordination between organisations
I6	Project management
I7	Engineering management
I8	Work preparation and production — WVB

Table 5.1: Interviewee codes and roles.

A third source sits between the two: the Value Stream Mapping (VSM) exercise carried out by the organisation itself, before and alongside this research [1]. A VSM maps the flow of materials and information from order to delivery and separates value-adding work from waste [53], [54]. It shows where waste accumulates, not how to remove it. In this thesis it is used as diagnostic evidence produced independently of the author, not as an analysis by the author.

A finding was kept only when more than one role supported it independently, or when an interview observation matched a VSM measurement or a document. Quotes illustrate the findings, they do not carry them on their own. To protect the people involved, interviewees are identified by role codes, and the organisations of the group are referred to by their function: the design office, central engineering, the hull yard, the mechanical outfitting organisation and the finishing yards. Figures come from internal records and are reported as recorded, with their known limitations.

5.2 Industrial context: one-off superyachts

Feadship designs and builds fully custom superyachts. There is no series effect and little standardisation between projects, which places the yard in the Engineer-to-Order setting described in Chapter 1: no client would accept the lead time of a sequential process, so detail engineering starts before the design is complete and hull construction before engineering is finished. Concurrency is not a choice here. It is the condition under which the chain operates.

The chain runs across five organisations. The design office develops the naval architecture and styling. Central engineering turns the design into an integrated 3D model through successive definition phases. The hull yard does the detail engineering of the steel structure, prepares the work and builds the hull, including structural outfitting. In parallel, a specialised organisation does the detail engineering of the mechanical outfitting. After hull delivery, the finishing yards complete systems, interior and exterior. Responsibility for the product therefore moves across separate companies and budgets.

Two features shape how the evidence is read. Labour is around 80–85% of the production cost of the hull, so information problems become visible as hours, which is why Chapter 6 uses budgeted against actual hours as its main evidence. And steel, once cut and welded, is in practice irreversible: a late correction no longer means editing a model but rebuilding a structure, the late-correction pattern described in Chapter 1.

Quality has no single measure linked to the release decision. It shows up as rework, which leaves traces in three places recorded by different systems: hours spent correcting at the yard, work passed on to the finishing yards and inspection records accepted after repair. Chapter 6 keeps them apart on purpose.

5.3 The process on paper (Work-As-Imagined)

The design process is a sequence of overlapping phases that progressively define the vessel (Figure 5.1). Concept Design (CD) and Technical Design (TD) establish the vessel. Design Development (DD) sets the general arrangement, the specifications and a first structural model. Definitief Ontwerp I (DOI) prepares the main structural plan for class approval, with its strength checks. Definitief Ontwerp II (DOII) integrates all disciplines into one coordinated 3D model. Detail Engineering (DE) turns it into buildable information, Work Preparation (WVB) into work packages, and production builds the hull. Each transition is a gate of the corporate OE programme, which defines the deliverables required before the next phase starts. This thesis studies the release from DOII to DE, called Gate 4 throughout.

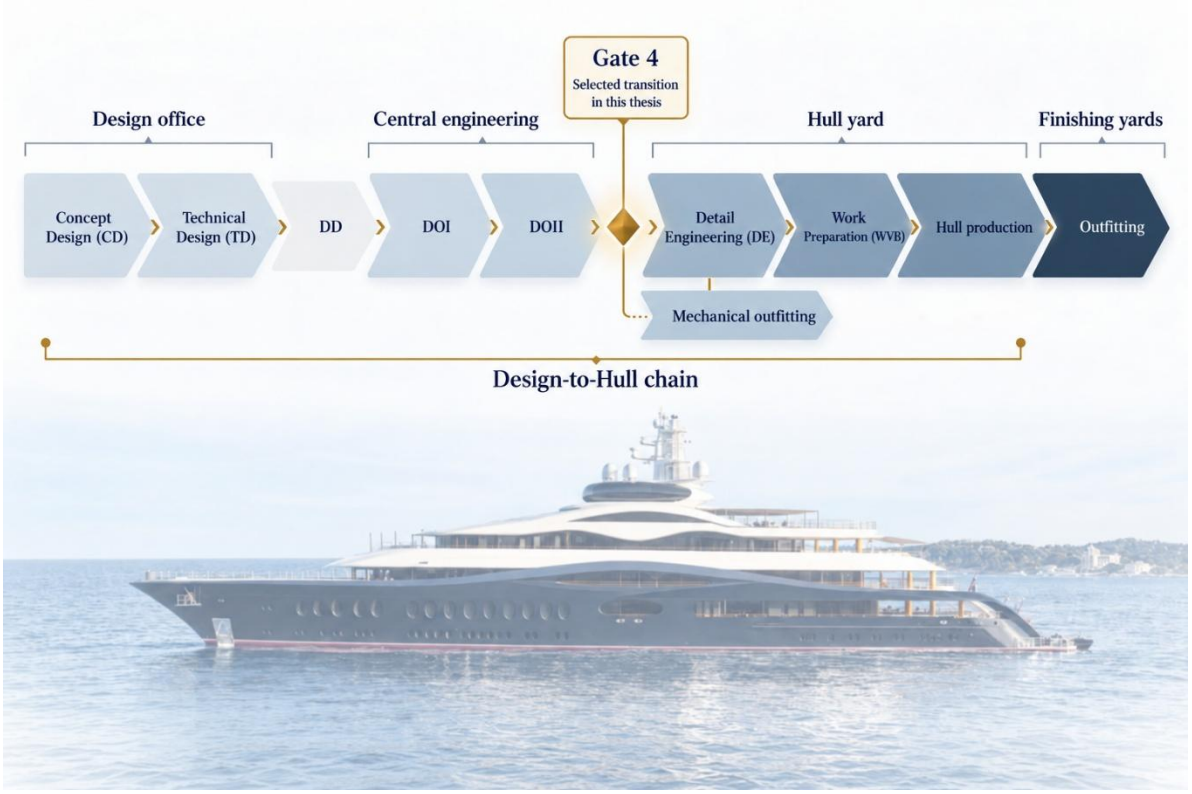


Figure 5.1: The design-to-hull chain and Gate 4. AI-assisted illustration.

In this thesis "steel" or "hull structure" means plates, stiffeners, primary and secondary members and the foundations of equipment. Pipes, ducts and cable trays are not steel, but they cross it: penetrations, piping routes, HVAC and cable trays are assessed as separate disciplines, because much of the impact of a late change reaches the hull through them.

DOII work is not organised by vessel but by Deliveries. A Delivery is a large block of the vessel: the bow, the engine room, the superstructure (Figure 5.2). Each one goes through its own DOII and its own entry into DE. Within each Delivery there are several sections, which are the physical units with which the hull is later built, and within a section, some outfitting disciplines, such as inbouw or piping, are in turn organised into blocks. The real production cost is measured precisely at section level, a detail that will matter in Chapter 6.

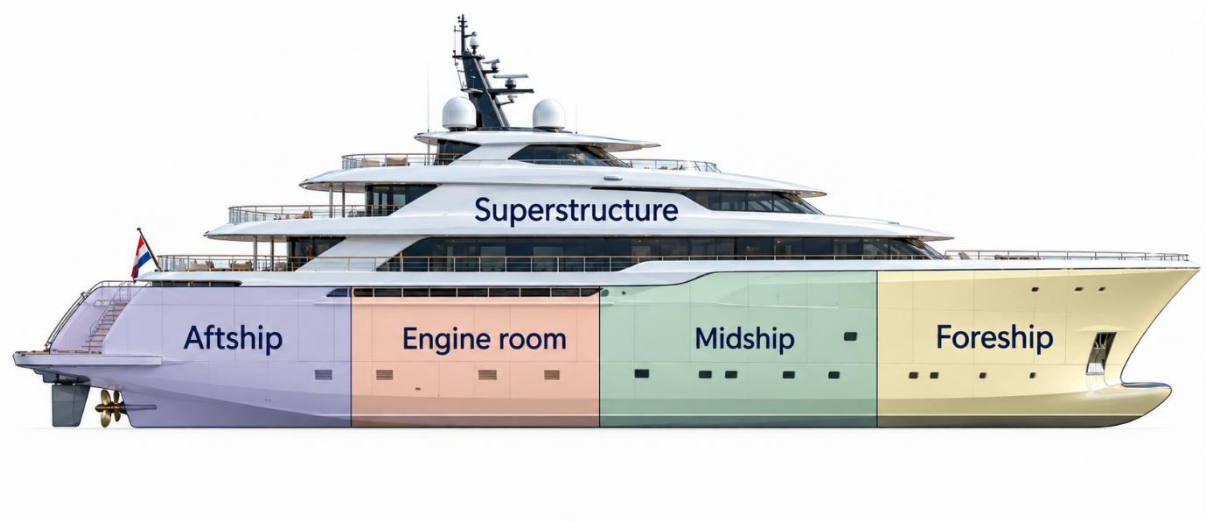


Figure 5.2: The deliveries of a yacht. AI-assisted illustration.

Each delivery takes around eight weeks of DOII, split into four windows of roughly two weeks: Base, Features, Outfitting and Review, each with its own integration list. The order goes from what is most structural and hardest to change towards secondary detail. DOII closes with two reviews: a 100% Review, at which the model should be complete, and a Final Review, which releases the delivery to DE. That Final Review is Gate 4. Where the DOII close, the delivery kickoff and the planning milestone fall on different dates, Chapter 6 states which cutoff is used for each case. Elements not integrated by then pass to DE as pending work, the open items, recorded in the close-out list. On paper the gate therefore expects three things: that the integrated model is complete, that the exceptions are listed and that the receiver accepts the release. It does not ask for a structured assessment of what accepting those exceptions implies.

There is a second feature of the formal process worth highlighting: although on paper it looks like a cut, the transition is actually progressive. DE does not receive a closed package, nor does it start from scratch. It keeps working on the same 3D model that came out of DOII and, over its own eight weeks, takes it from model to construction detail and fabrication drawings. That is why the gate looks more like a decision window over something still moving than like the handover of a finished package.

This leads to the central idea of the chapter. The formal process is aware: it knows that things will remain open at release. But there is something it does not do. It offers no structured way of measuring how much immaturity a specific release actually carries, nor of estimating what that immaturity may end up costing downstream. This is the gap GAIA addresses.

5.4 What the organisation's own VSM measured

The organisation itself carried out a VSM exercise involving the design office, central engineering and the yard: a current-state map, timeliness and budget measurements, and a bottleneck analysis. This section presents the three findings most relevant to the diagnosis.

5.4.1 On-time delivery drops most at Gate 4

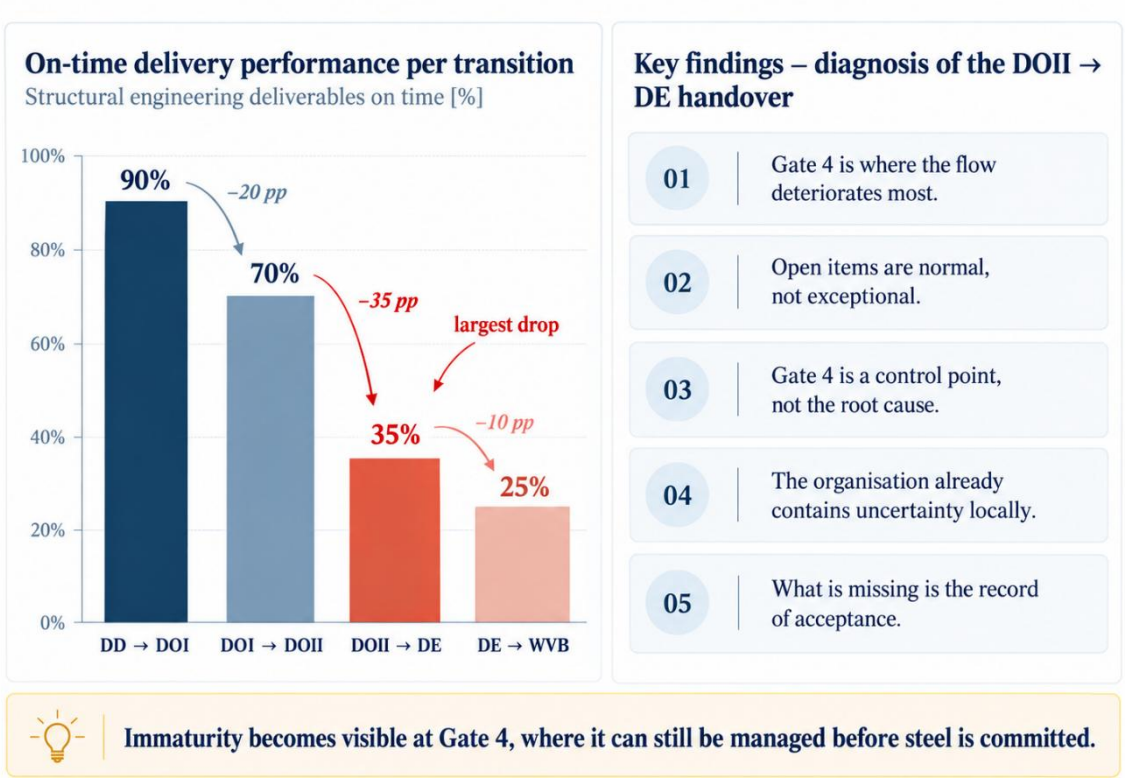


Figure 5.3: On-time structural engineering deliverables per transition (internal VSM) and the five findings. AI-assisted illustration.

Timeliness degrades as a cascade (Figure 5.3): each phase inherits the delay of the previous one and adds its own. At the release towards DE only about one structural engineering deliverable in three arrives on time, the steepest drop of the chain. The figure measures delivery against the MS Project planning, not the maturity of the content, but it is the best available indicator of the transitions. It also explains why the gate cannot simply demand that everything be finished: if only finished work could pass, releases would be held back routinely and the planning could not be kept. What is needed is not more strictness but a way to manage what is still immature without loading its cost onto the receiver.

5.4.2 Every phase uses more hours than budgeted

Phase	Hours above budget
Structural engineering (Constructie)	+20–30%
Detail engineering (DE)	+10%
Work preparation (WVB)	+5–10%
Hull production (Productie)	+5%

Table 5.2: Hours above budget per phase (internal VSM, budgets approved in DD).

Every phase uses more hours than budgeted (Table 5.2), against budgets fixed in DD, before anyone knows in what state the information will leave DOII. The largest relative overrun is upstream, in structural engineering, which matches the VSM finding that DOII works as a rework phase. Downstream the percentages are smaller, but there a change means redoing drawings, work packages or steel rather than editing a model. The figures come from different budgets and do not measure the cost of fixing the same problem, so they are context, not proof of the 1-10-100 pattern of Chapter 1.

5.4.3 The seven wastes

The most interesting result of the VSM exercise is its waste analysis. The internal report summarising the sessions concludes that the dominant waste in the current value stream is not technical, but a consequence of process maturity, activity sequencing, coordination between the different organisations and chain management. Moreover, problems tend to become visible late (DE, WVB or production) where their impact on cost, schedule and predictability is greatest. That conclusion, supported by the organisation itself, could serve as the problem statement of this thesis.

The sessions identified seven recurring wastes. Table 5.3 groups them by their relevance to the release under study, and Annex 5-A gives each with its documented manifestation.

Relevance to the release	Wastes	What they describe
Direct	1 Rework and late changes. 2 Waiting for information and approvals. 6 Misaligned and unstable assumptions.	Immature input crossing into DE, no shared definition of “ready”, and information that exists but will change.
Supporting	5 Insufficient chain coordination.	A shared release assessment can support the conversation between organisations, but does not fix coordination by itself.
Outside the direct scope	3 Duplication of work and conversions between systems. 4 Load imbalance and batching. 7 Local versus system efficiency.	Tooling, batch sizing and incentives, which a release assessment does not solve.

Table 5.3: The seven VSM wastes by relevance to the release.

Wastes 1, 2 and 6 describe, in the organisation’s own words, immature information crossing into detail engineering without a shared definition of ready. Waste 6 also separates incompleteness from instability, information that exists but will change, which is exactly the distinction between completeness and stability in GAIA. GAIA therefore addresses a defined subset of the VSM diagnosis, not the whole value-stream problem. And the diagnosis is not external criticism: it matches what the organisation has already found.

5.4.4 The organisational causes behind the wastes

The seven wastes describe symptoms. Read together with the interviews, they point at causes that are organisational rather than technical and the internal diagnosis of the chain names five of them. Different objectives and priorities between the organisations (waste 7: each department is efficient on its own while the project flow suffers). Chain governance exists at project level (gate reviews to the board, weekly contact between project leadership and engineering management, active push-back on late owner changes) but there is no consistent, cross-organisational way of assessing what remains pending at a delivery release and of recording under which conditions it is accepted (waste 5: no structural project consultation between the organisations). Local optimisation instead of global optimisation (wastes 4 and 7: batch releases that suit the issuer and overload the receiver). Different interpretations of what "ready" means (waste 2: no shared operational definition). And the way decisions and risks are accepted at Gate 4 (wastes 1 and 6: immature and unstable input crossing the gate without an explicit acceptance).

This redefines the problem. Information maturity is not the root cause. It is the measurable symptom in which these organisational causes appear at Gate 4. That is why Gate 4 works as a thermometer of the chain: it is where the accumulated state of the upstream process becomes visible and starts to cost. GAIA is designed on that reading: it measures the symptom, and it structures the decision, the ownership and the accepted risk that the diagnosis identifies as the underlying problem.

5.5 The process in practice (Work-As-Done)

The interviews confirm the conclusions of the documents and show how the release is experienced and managed day to day. Five patterns come out from the conversations, repeated across independent witnesses. The workforce already applies its own versions of the mechanisms this thesis proposes to structure.

5.5.1 Releases almost always carry open items

> *"In my almost 20 years experience I've never started the DE with zero open DO2 items."*
>
> — *Superstructure detail engineer (I2)*

Project leadership confirms that the aim is zero open items at the Final Review. On another project, still under way at the time of the interview, the project leader said it seemed on track so far, helped by pushing back on late owner requests. The interviews suggest this remains the exception rather than the norm.

Interviewees describe open items at the DOII–DE release as a recurring pattern. The open-item list issued at close is normal practice. Its format, level of detail and follow-up vary between projects. The information a structured release assessment would need already exists, but in inconsistent formats, which limits its use for decision-making and, in addition, complicates its retrospective reconstruction (a limitation, Chapter 6).

(Pattern supported independently by the hull and the superstructure engineers, I1 and I2, and consistent with the formal process itself: the close-out list exists by design, see Section 5.3).

5.5.2 The handover is gradual and work continues

> *"We have the same model, but we just keep adding more and more details."*
>
> *— Hull detail engineer (I1)*

Professionals describe the transition exactly as Section 5.3: a decision window over a shared model, not a handover. Except for important reasons, the window closes under schedule pressure regardless of maturity: as a superstructure engineer put it, "we can actually never choose to not start with DE as the planning dictates we have to finish at a certain date" (I2). The realistic decision at the gate is therefore never a binary go/no-go on the whole delivery, but a portfolio of smaller decisions on which parts to release, which to watch and which to isolate. Any framework that treats the gate as a binary switch would fail against this reality, which is why GAIA decides per item (Chapter 4).

Continuing with pending items is not a failure of discipline: engineers describe it as the rational response to planning and capacity constraints and the practical question is not whether to continue but which work can continue, under what condition and before which activity the pending information must arrive.

(Supported by I1, I2 and a project manager, I6. Compatible with the low on-time figure of Section 5.4.1, without showing a cause.)

5.5.3 Engineers box what is not ready

> *"You can always make some kind of box of the area that will change, and the rest you can go further."*
>
> *— Hull detail engineer (I1)*

Both structure and superstructure engineers independently described the same practice, here called boxing: when a zone of the model is known to be unstable, they isolate it (a "box" is marked around it) and keep detailing everything outside. Engineers keep their own records of, per section and per review, which information came in, when and whether it was processed (I2). Boxing is effective at avoiding work in immature zones and avoiding non value-added effort. This supports the isolation idea of the framework.

(Described independently by I1 and I2, from the hull and the superstructure side respectively).

5.5.4 Risk arrives as rumour, not as information

> *"Usually we do hear some rumours about this or that, like vibration simulations, strength calculations, possible design changes and/or likely problems with Plan Approval. But often we don't know the impact for construction."*
>
> *— Superstructure detail engineer (I2)*

Warning signals do reach DE: pending analyses, likely design changes, unapproved drawings. But they arrive informally, without the full impact being known. Imprecise information cannot be properly assessed. Interviewees added that this depends a lot on who is upstream: some engineers are very careful and know what the person downstream needs and when. Others miss it or do not give it as much importance (I2). Basically today, whether information arrives well depends mostly on which person you get, when it should be guaranteed by the process itself.

The decision to accept a DOII model with pending items does exist in practice and so does the acceptance of consequences. What does not exist is a record of that decision consistent across projects and across the two project managers who successively own a delivery.

(I2. Consistent with VSM wastes 2, 5 and 6: waiting for information, insufficient chain coordination and unstable assumptions.)

5.5.5 Where immaturity hurts most

Asked where immaturity hurts most, the interviewees pointed to several specific situations. Most interviewees work in the structural engineering of the hull or the superstructure, so their answers centre on steel and its interfaces. The project managers added a view of the yacht as a whole.

In mechanical outfitting (WTB), DE starts at the same time as the hull yard's but finishes later. This means that late pipe reroutes or bracket changes can arrive when the structural model is already closed (I1, I7).

In exterior outfitting (EXT), maturity depends on supplier and owner decisions that the chain does not control (I1, I2).

In addition, the finishing yards sometimes ask to add work during hull construction (actions easier and cheaper to carry out at the yard than later) and those additions enter as change orders that add to the release baseline (I5).

For the framework, the lesson is that maturity is best looked at discipline by discipline and that external inputs such as supplier data, owner decisions and other organisations' schedules deserve a clear status at the gate.

(I1, I2, I5 and I7. Consistent with waste 1 (rework and late changes) and with the scope demarcation between the yard and the mechanical outfitting organisation).

5.6 Why intervene at Gate 4

Chapter 1 placed GAIA at the DOII → DE release. The diagnosis supports that choice for four reasons (Figure 5.4). Empirically, it is where on-time delivery of structural engineering deliverables drops most sharply, to about 35%. Economically, it is the last cheap decision point: before the gate a correction means editing a model, after it drawings, work packages and eventually steel. As practitioners put it, what determines the cost of a pending item is not how difficult it is but the stage in which it is processed. Institutionally, the organisation itself points here: its VSM names the lack of an operational definition of ready and of a firm quality control per phase as main causes, and its own improvement proposal asks for the class review status to be part of the DO–DE transition [1]. Practically, it is where the evidence is: the hull yard's scope is contractually delimited and its hours, changes and failure costs are recorded per section and discipline, which makes the retrospective test of Chapter 6 possible.

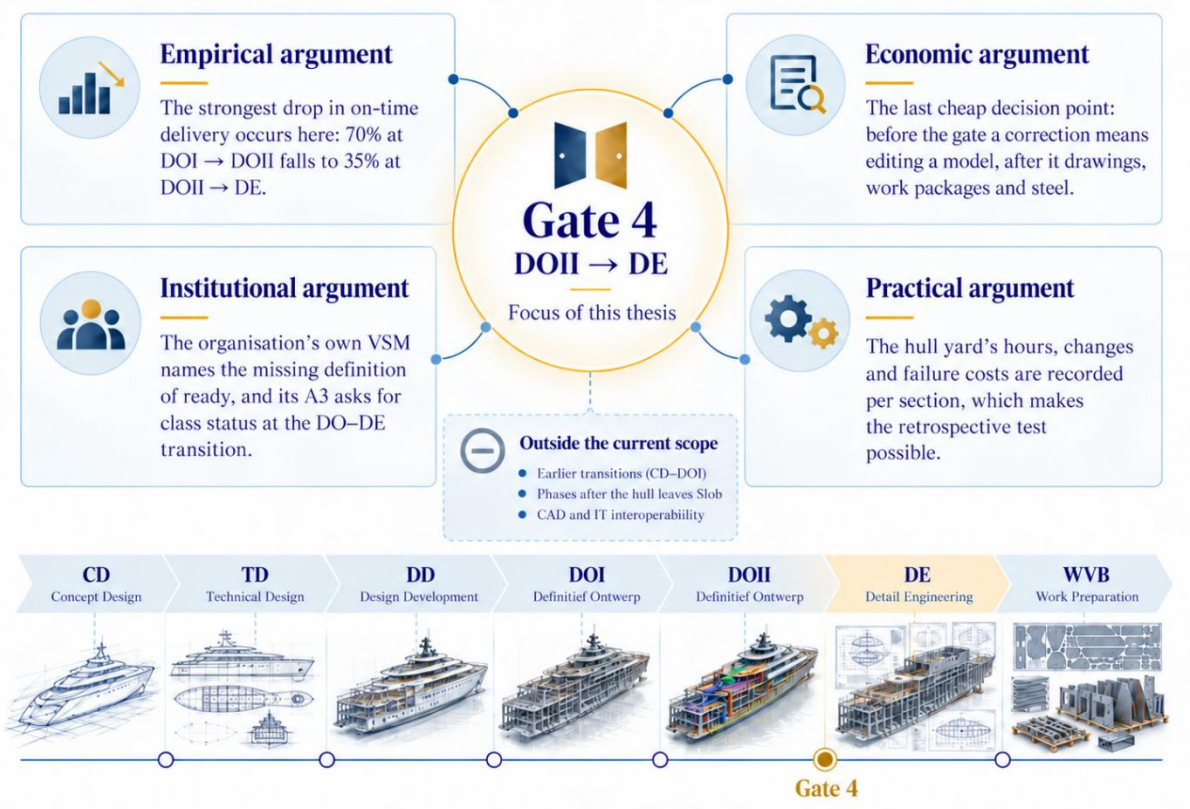


Figure 5.4: Why Gate 4: four arguments and what stays out of scope. AI-assisted illustration.

The gate cannot force every piece of information to exist. Some results, such as finite element or vibration calculations, need a model detailed enough that it only exists during DE or later. For these, the gate cannot demand the answer, but it can ask for the status: which analysis is pending, when the result is expected, what the specialists already expect it to show and which scope is exposed. Level I rates that status, and an expectation alone is not yet a confirmation (Chapter 4). This is why GAIA assesses maturity instead of demanding completeness.

5.7 Five findings and answer to SQ3

"What are the current causes of rework and inefficiency at the design-to-production interface in ETO shipbuilding, and how do planned information flows (Work-As-Imagined) diverge from the information actually delivered (Work-As-Done)?"

The diagnosis comes down to five findings, which together answer SQ3.

1. Gate 4 is where the flow deteriorates most. Only about one third of structural engineering deliverables reaches DE on time, clearly below the earlier transitions. The figures come from the organisation's own VSM and do not establish causality.

2. Open items are always there. Work runs in parallel, the DOII window closes on a date and the yard has to start, so every delivery crosses the gate with items still open. On paper the release looks like a yes or no on the whole delivery. In practice the decision is what to release, what to watch and what to hold back.

3. Gate 4 is a control point, not the root cause. Supplier data, owner decisions, class, coordination and capacity originate upstream or outside the chain. Gate 4 is where their accumulated effect becomes visible and where the remaining risk can still be managed before steel is committed.

4. The organisation already contains uncertainty locally. Close-out lists exist, reviews are held and engineers box unstable zones. The practice works, but it depends on individuals and its format varies between projects.

5. What is missing is the record of acceptance. Items cross the gate without a structured assessment of what accepting them implies, and without a record of the condition, owner and risk accepted. Nobody can later check whether an accepted item had consequences. This is the gap GAIA is designed to close.

In short, the answer to SQ3 is this. Rework and inefficiency at the design-to-production interface come from information that crosses the gate before it is stable (supplier data, owner decisions, class and coordination), made worse when engineering capacity falls short, and their effect builds up until Gate 4. On paper the release is a cut, with a complete model and a list of exceptions. In practice it is a progressive handover over a shared model, every delivery carries open items, and what accepting them implies is judged by individuals and not recorded.

These findings do not derive GAIA after the fact. They corroborate, with the yard's own evidence, the requirements developed in Chapters 3 and 4 (Table 4.1). Chapter 6 now tests whether the signals GAIA produces at this point are consistent with what happened downstream.

RETROSPECTIVE VALIDATION OF GAIA

Chapter 4 designed GAIA: a way of structuring the release of a delivery from DOII into Detail Engineering, at three levels. Level I checks the context of the delivery, Level II the maturity of the package per discipline, and Level III turns each open item into one of four decisions: GO, CONDITIONAL, ISOLATE or WAIT. This chapter puts that design to the test on three real projects. It answers the fourth sub-question of the thesis (SQ4):

> *"To what extent do the maturity scores generated by the tailored assessment framework show indicative consistency with historical rework and production performance when applied retrospectively to completed projects?"*

In practical terms, this asks whether the signals GAIA produces at Gate 4 are consistent with what the yard recorded afterwards.

Each case tests something different. Case A, a good project and the first application, is the reference: it shows whether GAIA stays quiet when it should. Case B, a second engine room with the best records, tests whether the procedure repeats. Case C, a deck structure in the superstructure with a large overrun, tests where the framework stops. Section 6.1 sets out the evaluation method and its criteria, Sections 6.2 and 6.3 the test, Sections 6.4 to 6.6 apply it to each case, Section 6.7 compares the cases and Section 6.8 answers the question.

6.1 Evaluation method and criteria

Before the cases, this section sets out how GAIA is evaluated (Section 6.1.1) and against which criteria (Section 6.1.2).

6.1.1 Retrospective case study

The framework is evaluated through a retrospective multiple-case study [37]. Three completed deliveries were selected to differ in outcome and context: a reference project with a good outcome (Case A, engine room), a recent project in the same zone with more complete records (Case B, engine room) and a stress case in a different zone and organisation (Case C, superstructure deck). In this version the three projects are anonymised as Case A, B and C. For each case, the state at Gate 4 is reconstructed from dated documents and reviewed with engineers involved in the project, the framework is applied, and the result is compared with the downstream record of hours, planning, quality and change orders.

The comparison follows pattern matching with rival explanations [37]. The question is whether the pattern expected from the framework matches the pattern observed, not whether a statistical correlation exists. Every link between an open item and a later event receives a verdict, and the best alternative explanation is checked before the link is accepted. The detailed procedure is given in Section 6.3.

The case material combines interviews with engineers, project managers and production staff across the design office, central engineering and the yard; the yard's project documents (open-item lists, close-out notes, planning, change-order registers and quality records); and its registers of hours. Annex 6-D lists the sources per case.

6.1.2 Leading versus lagging indicators

Kaplan and Norton [40], [41] introduced the Balanced Scorecard framework as a response to the limitations of purely financial performance measurement. Their main argument is that financial metrics (revenue, profit, cost) are lagging indicators: they report what has already happened but arrive too late for management to intervene. Effective performance management requires a balanced combination of lagging indicators (confirming past results) and leading indicators (predicting future results and enabling proactive action).

A leading indicator must satisfy three conditions to be useful. First, it must precede the outcome temporally: it must be measurable before the result it predicts materializes. Second, it must correlate with the outcome: changes in the leading indicator must be associated with changes in the lagging indicator. Third, it must be actionable: management must be able to intervene based on the leading indicator to influence the outcome.

In project environments, typical lagging indicators include total project cost, schedule deviation, rework hours and quality defects at delivery. These are essential for accountability and learning but cannot prevent the problems they measure. Leading indicators, by contrast, might include design completeness at key milestones, resource allocation versus plan or risk register status. The challenge is that few leading indicators in project environments have been empirically validated against outcomes with sufficient rigour to justify their use as decision-making tools.

The Construction Industry Institute (CII) has made significant contributions in this area. The PDRI [8] demonstrated that front-end planning completeness measured before construction correlates with project cost and schedule performance. The FEED MATRS methodology [38] extended this by producing separate maturity and accuracy scores and showing that projects scoring high on both dimensions outperform those scoring low. In the offshore EPC sector, Kim et al. [6] demonstrated that engineering completeness measured at the point of construction start correlates with labour inefficiency and schedule deviation.

These precedents establish that maturity indices measured at phase transitions can function as leading indicators of production performance, provided the correlation is empirically demonstrated under the specific conditions of the application context. The transition from a lagging measurement ("how much rework did we have?") to a forward-looking assessment ("given the current maturity level, how does it compare with patterns observed in completed projects?") is the fundamental step that converts a descriptive index into a management tool.

Implication for this research: a maturity index measured before production has the potential to function as a leading indicator of production performance if three conditions are met: temporal precedence (measured before production starts), demonstrated correlation (validated against lagging KPIs) and actionability (connected to a decision mechanism at the gate). These conditions define the validation requirements for the assessment framework developed in this thesis. The proposed framework is designed to meet the first and third conditions. The second can only be examined in this thesis as indicative consistency across three cases. A statistical correlation needs a larger sample and is left to future research.

6.2 What is tested and what is not

The test follows the pattern-matching logic of Yin. An expected pattern, the classification at the gate, is compared with an observed pattern, the downstream record. Each link receives a verdict that grades its strength, from consistent to partial or inconclusive (Table 6.2), and is tested against a rival explanation. The aim is **indicative consistency**: whether cost and rework appeared where GAIA pointed and not where it did not. The test does not claim statistical prediction, causality or savings. If effort later landed where GAIA saw a weakness, that shows consistency, not that GAIA would have prevented it.

The delay of the vessel is not the variable being tested. In all three cases the hull left the yard on or close to its date. What moved was where the effort was absorbed. The downstream record is therefore read in four layers, always separately and never added together: hours (budget against actual, per section and element), local planning (sections and blocks), quality and recovery effort (straightening and rework after inspection, the hours Slob books to correct a section after it has been inspected) and the register of change orders and failure costs.

The version of GAIA fixed in Chapter 4 is applied in the same way to the three cases. This is not a blind test. The cases were selected for their different outcomes (Section 6.1.1), so the researcher knew their general result. The engineers who helped to rebuild each gate knew how Detail Engineering had gone and some production events, but not the hours, quality records and change costs, which were collected separately from the yard's systems. Some events are also several years old. Hindsight cannot be removed, but three rules limit it. Every status at the gate rests on dated evidence. Engineers were asked how they would classify an item given what was on record, not what they remember happened. And, once fixed, the framework was not adjusted to fit any case. Where the record does not allow a status to be set, it is marked N/E (not evaluable). N/E does not mean nothing was known. It means the information that exists is partial, local or undated and not enough to set a status for the whole package.

Three things are kept apart throughout: the decision GAIA would have produced at the gate, which is reconstructed, what the project actually did at the time, where it is documented and what happened afterwards. Two further limits apply. Level II takes its residual criticality from the Level III items, so agreement between the two levels is internal consistency, not a second confirmation. And a released item without a later trace may reflect incomplete records rather than the absence of a problem.

6.3 How each case was assessed

The procedure (Table 6.1) is the same for the three cases and would be the same for any other zone of the vessel. Only two things change between zones: the scope reference used to measure completeness (hull vs superstructure) and the sources available downstream.

Gate 4 is the DOII Final Review, which releases the delivery into Detail Engineering (Chapter 5). It is dated by the kickoff of the delivery in the project planning, confirmed with the project engineer, by the same rule in the three cases. The DE reviews that follow weeks later belong to Detail Engineering, not to the gate. Only what was on record at the gate is used to classify, and anything logged later belongs to the downstream record. The unit of analysis is **the Gate 4 handover that contains the functional zone**. Open items are sorted into three groups: DOII carry-over (scope that crossed the gate unresolved), normal Detail Engineering work and items created after the reviews. Only the carry-over is assessed. The other two groups are used as controls.

Step	What is done	Output
1. Reconstruct the gate	Close-out list or engineering log, kickoff and review minutes, scope reference, planning. Each open item is assigned to a group.	Dated gate file
2. Apply GAIA	Level I from dated evidence or N/E. Level II: completeness against the scope reference, stability from the state of the sources, residual criticality from the worst open item. Level III: decision rule, reviewed with the project engineer	Classification at the gate
3. Downstream evidence	Hours, planning, quality and change register, kept separate, nothing attributed yet.	Four layers
4. Pattern matching	Per item: same physical element, timing and mechanism; verdict, which grades the strength of the link. Per discipline: Level II profile against the observed behaviour.	Verdicts
5. Rival explanations	For each link, the best explanation other than immature information, settled one by one.	Rivals

Table 6.1: The five steps applied to each case.

Verdict	Meaning
Consistent	The risk flagged at the gate left a trace in the same element.
Consistent (context)	A Level I or Level II signal matches the downstream behaviour of the package as a whole.
Partial	Trace in the same zone or shared with another item. The link cannot be separated.
Negative control	Item released (GO) with no trace afterwards. Never counted as strong: absence of evidence is not proof.
Inconclusive	No trace either way or item no longer identifiable.
Possible false negative	Item released (GO) that later left a direct trace.
Execution error	Item released (GO) whose later trace was booked as a production error. Counted neither as a match nor as a miss.
Not rated at the gate	The effect is visible downstream but its status at the gate could not be set (N/E, Level I only).

Table 6.2: Verdicts used to compare the gate with the record.

The change register is read through the yard's own categories: changes to the hull scope paid by the client, added scope, failure costs absorbed by the yard and administrative corrections. These are reading categories, not damage figures. They are never added to each other or to hours and the register total is never presented as the impact of GAIA. A change is linked to an item only when element and mechanism match.

How to read the cases. Each case uses the same three tables. The first shows GAIA at the gate: how many of the seven preconditions were green, amber, red or not evaluable (Context), the maturity of each discipline from 1 to 5, the lowest of its three dimensions, where 5 means complete and frozen (Package), and how the open items split over the four decisions (Items). The second answers five questions about what happened afterwards. The third follows each signal to the downstream record and gives a verdict from Table 6.2. Its column GAIA shows the decision given at the gate for an open item, a colour for a Level I precondition or a score for a Level II discipline. Items are named by what they are. The number in parentheses points to the annex, which gives the item tables in full.

6.4 Case A: the reference case

6.4.1 The case and the gate

Case A was a good project: a cooperative owner, a stable team and a hull delivered twelve days after plan. That is exactly why it comes first. A framework that raises alarms on a healthy project would not be used. The question here is whether GAIA stays quiet when it should and whether the little it does flag is what later left a trace. As the first case, it took more sessions to rebuild, but it was assessed with the same rules as the other two.

The gate of the delivery containing the engine room (Figure 5.2) was held on its planned date. The close-out list, made five days before the gate, is a dated snapshot: 23 open items in the delivery, 11 of them in the engine room. One item (128, exhaust foundations) was already marked done and is kept as a closed control. That leaves 10 open items at the release.

6.4.2 What GAIA shows at the gate

At the gate GAIA sees a healthy package with two weak spots, both in information that was drawn but not yet confirmed (Table 6.3).

Level	Result	What it means
I. Context	3 green · 4 N/E · no red	No red signal in the reconstructed evidence. Supplier and vibration issues existed, but each was already an open item.
II. Package	Primary steel 2 · Secondary steel 3 · Foundations and equipment 2 · Piping 3 · HVAC 3 · Interior 3	Mostly at 3. Two disciplines at 2 because what was drawn still rested on unconfirmed data: supplier footprints for the foundations, an open pillar detail for primary steel.
III. Items	10 items: 2 GO · 4 CONDITIONAL · 4 ISOLATE · 0 WAIT	Ten open items become three different treatments. None requires stopping the engine room.

Table 6.3: Case A at the gate. Details in Annex 6-A, Tables 6-A.1 to 6-A.3.

6.4.3 What happened afterwards

The engine room overran by 5.6%. The structure stayed on budget and the extra hours went elsewhere (Table 6.4).

Question	What the record shows
Did the package overrun?	Slightly: +5.6% over the package budget.
Where did the extra effort go?	The structure family stayed on budget (+0.8%). The piping and installation family rose by 20.6%. At element level the largest deviations are hard piping (+46%) and foundations and equipment installation (+31% and +24%): the discipline GAIA marked as the limiter.
Did the schedule move?	The gate was on time and the hull only twelve days late. Inside the zone there were delays: supplier information on foundations reached the yard 12 to 14 weeks late, straightening took 42 days instead of 14, and rework after inspection ran for 49 and 74 days with no baseline in the plan.
Was quality affected?	Final quality was accepted. The signal is the effort to get there: rework after inspection reached more than twice its budget.
What changed in the register?	36 change orders and failure costs in the zone. The two largest were client-paid scope changes (switchboard structure, piping scope reassigned eleven months after the gate). Neither had an open item at the gate.

Table 6.4: Case A downstream.

6.4.4 Does it match?

Three of the ten items left a clear trace in the same element, three a partial one, and both released items stayed clean (Table 6.5).

Signal at the gate	GAIA	What happened	Verdict
Generator foundations: final supplier footprint missing (78)	CONDITIONAL	Foundation revisions in the same family. Foundations +31%; supplier information 12-14 weeks late	Consistent
Pillar in walkway: solution under discussion (99)	CONDITIONAL	Later connection and vibration-damping work on the same pillar	Consistent
Holes in girders: piping check pending (112)	CONDITIONAL	Penetrations added and installed equipment moved to fit the piping	Consistent
Foundations and equipment	Level II: 2	Foundations +31%, equipment installation +24%, structure on budget	Consistent (context)
Manholes at the generator foundations, ventilation penetrations, lightening holes in frames (77, 127, 131)	ISOLATE / COND.	Work in the same zone, shared with other items or booked as execution failure	Partial
Lift foundations, door at frame 28 (80, 89)	GO	No trace	Negative control
Lift service space, sliding doors at frame 32 (100, 125)	ISOLATE	No specific trace. The sliding-door item can no longer be identified	Inconclusive

Table 6.5: Case A: what GAIA flagged and what happened. Full table in Annex 6-A, Table 6-A.4.

Rival explanations (Annex 6-A, Table 6-A.5). The foundation link is kept at family level, because the foundation hours include other foundations. Normal detail development and shared changes explain part of the pillar, ventilation and girder work, and the frame rework was an execution failure, not a hit. The switchboard and the later piping scope lie outside anything an open item could show.

6.4.5 What Case A tells us

1. **GAIA stays calm on a healthy project.** The context shows no red and most disciplines sit at 3. That matters for use: a tool that flags everything is soon ignored.
2. **It points to the right place, though not to every place.** Foundations and equipment, one of the two disciplines GAIA scored 2, overran clearly: foundations +31% and equipment installation +24%, against +5.6% for the whole package. The other discipline at 2, primary steel, belongs to the structure family, which stayed on budget (+0.8%): its open pillar detail left a local trace and the frame rework was booked as an execution failure. The largest single deviation, hard piping (+46%), sat in a discipline scored 3. Its isometrics arrived twelve weeks late and part of its scope was reassigned eleven months after the gate, as a client-paid change. Neither was an open item at the gate.
3. **It sorts instead of counting.** Ten open items became three kinds of decision. The original list already carried an informal risk mark and the two items with the most direct later trace, the generator foundations and the girder holes, did not have it. The same information, read with explicit criteria, would have put them in front of the release meeting. This does not show that the framework outperforms an experienced engineer in general. It shows what explicit criteria add to the record that already exists.
4. **It shows what an item list cannot see.** The largest groups of changes, the switchboard structure, part of it budgeted but not yet approved, and a piping scope reassigned eleven months after the gate, were client-paid scope changes, not rework from immature information. No open item could show them. A scope still undecided at the release is a question for the context level (A3). A scope decided later lies outside GAIA.

In a good project, the value of GAIA is not alarm but focus: it tells the release meeting where to spend its attention.

6.5 Case B: the second engine room

6.5.1 The case and the gate

Case B is a recent project in the same zone as Case A (Figure 5.2), with better records and a normal outcome. It shows whether the procedure repeats without adjustment and what a second engine room adds.

The gate was held on its planned date, confirmed by the project engineer and documented in his DOII close-out note. The open-item board of this project is cumulative: it keeps growing during Detail Engineering. The population at the gate was therefore rebuilt with the project engineer. It holds 22 open items carried over from DOII. Fourteen further items are collisions between structure and piping found during Detail Engineering. The mechanical engineering organisation confirmed they are normal DE work so they are excluded and used as a control.

6.5.2 What GAIA shows at the gate

Here GAIA raises one red flag, and it is not on the item list (Table 6.6).

Level	Result	What it means
I. Context	1 red · 6 N/E	The close-out note states that the impact of the vibration analysis is still unknown. That is an open structural uncertainty with no location and it appears in none of the items on the board.
II. Package	Primary steel 2 · Secondary steel 4 · Foundations and equipment 2 · Piping 3 · HVAC 2 · Interior 2	Almost everything was drawn, but on bases still moving: scantlings waiting for the vibration result, the generator footprint, the lift, an earlier owner change to the air-conditioning room. Piping at 3 is its planned level.
III. Items	22 items: 12 GO · 6 CONDITIONAL · 3 ISOLATE · 1 WAIT	Classified by the project engineer. The WAIT is the generator footprint: the tank and foundation zone was held and the rest released.

Table 6.6: Case B at the gate. Details in Annex 6-B, Tables 6-B.1 to 6-B.3.

6.5.3 What happened afterwards

The package overran by 5.5%, all of it in the structure (Table 6.7).

Question	What the record shows
Did the package overrun?	Slightly: +5.5% over the package budget.
Where did the extra effort go?	All into the structure: metal and welding +12.4%. Mechanical installation came in 27% under budget and the built foundations were under budget as well.
Did the schedule move?	The gate and both DE reviews were on their date. Afterwards, the section-building activity of the lower section closed 50 days late against the revised baseline, mostly at block assembly according to production, and straightening finished 29 days late. The hull was placed and launched on time.
Was quality affected?	The engine-room welds were accepted. Shell fairness needed local correction with filler, left to the finishing yard. Rework after inspection rose far above its small budget, with no baseline in the plan.
What changed in the register?	The largest change linked to vibration and class, extra structure around the pillars (1.4% of the package budget, 92% in production), had no open item, but the red context signal at the gate pointed to it.

Table 6.7: Case B downstream.

6.5.4 Does it match?

The three strongest signals at the gate left a trace, one of them from the context. Of the twelve released items, one later needed work, a small change, all before production, and two shared a failure cost booked as a production error (Table 6.8).

Signal at the gate	GAIA	What happened	Verdict
Vibration impact unknown	Level I: Red	Extra structure around the pillars, 1.4% of the budget, mostly in production. Structural overrun of the section	Consistent (context)
Generator footprint (229)	WAIT	At the time, the engineer held the tank-top and foundation zone. The related change was solved in engineering and preparation, with no production hours on that change. Foundations under budget. This is compatible with management before steel, but does not show that the hold caused it	Consistent
Guest-lift pit (173)	ISOLATE	Change adding structure for the lift, same element	Consistent
Void adjustment (301)	ISOLATE	Shares a change with the void-access question (343)	Partial
Recirculation-duct structure, casing above the upper section, casing construction (341, 339, 181)	ISOLATE / COND.	No change isolates the effect of the item	Inconclusive
Are the void tanks accessible? (343)	GO	An access was created a month later, with a small change	Possible false negative
Sea-inlet details (304, 322)	GO	Failure cost on the plates around the sea inlet, booked as production error	Execution error
Other released items	GO	No trace	Negative control

Table 6.8: Case B: what GAIA flagged and what happened. Full table in Annex 6-B, Table 6-B.4.

The void-access question (343) is recorded as a possible false negative. It was a small acceptance question, asked and closed more than a year before this reconstruction and the engineers recalled it with less certainty than the other items. It points to a gap in the rule rather

than to a failure of judgement: an open question about acceptance does not look like a geometry risk, so the rule had no place to catch it.

Rival explanations (Annex 6-B, Table 6-B.5). Part of what happened later has causes outside GAIA. The largest change of the package, foundation strips, was added after the gate. The largest failure cost came from execution on the upper section. The switchboard stairs were decided after the DE reviews. None of these was open at the release. Hours were also not always booked against the change that caused them, so they show a direction, not an exact measure.

6.5.5 What Case B tells us

1. **On time is not the same as mature.** The gate and the reviews met their dates and 22 carry-over items were resolved during Detail Engineering, followed by local tails in production.
2. **The item list is necessary, but not sufficient.** The one hard signal of the case, vibration, was written in the close-out note and in none of the open items. An item-only tool would have missed it. This is the empirical case for keeping context as a separate level. It is also not negligence: a full vibration response analysis takes around fifteen weeks, so its result normally arrives during DOII or later.
3. **A held zone was resolved before production.** The zone held at the gate was solved in engineering and preparation and the foundations finished under budget. This is compatible with managing the risk before steel, not proof that holding it was the cause.
4. **Many open items do not mean many risks.** Fourteen of the items on the board were normal development work and the piping and installation family finished 8% under budget. Counting open items would have painted the wrong picture.

6.6 Case C: the stress case

6.6.1 The case and the gate

Case C changes zone, organisation and scope reference: a large deck structure in the superstructure (Figure 5.2). It is the case with the largest downstream disturbance and the one that tests the framework hardest, because much of the extra effort was not a matter of information maturity.

The functional zone is two sections, confirmed by the detail engineer as one unit. The gate is the kickoff of the delivery in the project planning, by the same rule as in the other two cases. Entries logged after it, up to the end of DOII a week later, are read as downstream and using that later date would not change the conclusions. There is no close-out list for this delivery. Its place is taken by the engineering log of the section, where every entry is dated. Eight entries were still open at the gate. Grouped by element and decision, they form four open items.

6.6.2 What GAIA shows at the gate

On paper the deck structure looked drawn. GAIA shows that it was not stable (Table 6.9).

Level	Result	What it means
I. Context	1 amber (weight) · 1 green · 5 N/E	Weight was an active constraint: nine scantling reductions in two weeks, one change adding 900 to 1,000 kg "to investigate" and a strength check still pending. There is partial information on class and suppliers, but no general status for the package.
II. Package	Primary steel 2 · Secondary steel 3 · Foundations and equipment 1 · Piping 3 · Exterior 2	The package looked drawn: primary steel scores 4 on completeness. But stability falls to 2 in three disciplines: scantlings waiting for calculation, a crane with no supplier data and an exterior definition still changing. Foundations drops to 1 because its worst item, the crane support, combines high impact with high difficulty.
III. Items	4 items: 2 CONDITIONAL · 1 ISOLATE · 1 WAIT	The WAIT is the crane support structure and the ISOLATE the storage hatches. The items were classified by the superstructure project engineer and an engineer from work preparation. The detail engineer confirmed the mechanism and severity of these two.

Table 6.9: Case C at the gate. Details in Annex 6-C, Tables 6-C.1 to 6-C.3.

6.6.3 What happened afterwards

This is the case with the heaviest consequence: the package needed 58.6% more hours than budgeted (Table 6.10).

Question	What the record shows
Did the package overrun?	Heavily: +58.6% over the package budget.
Where did the extra effort go?	Into the steel itself. Metal and welding +70%, and section building alone +79%. Engineering and work preparation stayed on or under budget.
Did the schedule move?	Detail Engineering and its reviews took place on the dates shown in the planning, but with no baseline, compliance with the original plan cannot be assessed. Section building took about nine months and, with straightening, ran past the launch, which was held. Finishing continued at the finishing yard.
Was quality affected?	No inspection records specific to the zone were available to this study. The corrective effort is visible: straightening far above its budget and rework after inspection, unbudgeted, 2% of the package budget. The project evaluation notes deformation during building and flat bars where production wanted profiles.
What changed in the register?	17 change orders above 0.25% of the package budget, about 9% in total: 8.2% specific to the zone and 0.6% shared with other sections. By description: class-related 34%, exterior releases 30%, buildability and execution 22%, interfaces 14%.

Table 6.10: Case C downstream.

6.6.4 Does it match?

The two items GAIA would have held back or fenced off both came back, one in Detail Engineering and one in production (Table 6.11).

Signal at the gate	GAIA	What happened	Verdict
Storage hatches modelled on two conflicting source models (R2)	ISOLATE	Both hatches redone at the last minute with the section already in production. Change of 0.6% of the package budget for brackets and stiffening	Consistent
Crane support on a web proposal not accepted (R4)	WAIT	Beams moved within days of the release, web dimensions revised about seven weeks later, flat bars added. Absorbed in DE with extra hours, according to the detail engineer	Consistent
Storage mass and lightening holes, both awaiting checks (R1, R3)	CONDITIONAL	Later work in the same zones: drain holes added about nine months after the gate, closure of lightening holes requested about five months after it. Identity with the original items not confirmed	Partial
Weight under study	Level I: Amber	Stiffeners reversed after the strength calculation. Flat-bar change during Detail Engineering	Consistent (context)
Exterior definition still moving	Level II: 2	New form revision five days after the gate, deck gutter released two months later, railing pots moved with the section already in production	Consistent (context)
Class approval	Level I: N/E	Class comments generated work from about four months after the gate. Review needed six weeks against two planned	Not rated at the gate

Table 6.11: Case C: what GAIA flagged and what happened. No item was released at the gate, so there is no negative control. Full table in Annex 6-C, Table 6-C.4.

Rival explanations (Annex 6-C, Table 6-C.5). Four alternatives explain much of the overrun, and the study does not divide the hours between them. Engineering capacity was limited, and part of the steel was detailed with flat bars where production wanted profiles. The planning left no slack. The class review needed six weeks against two planned and later class comments reopened some approved details. And heat input and working methods in thin plate contributed to deformation. Two events had no signal at the release at all: removing hard spots (1% of the budget) and a piping clash (0.5%).

6.6.5 What Case C tells us

1. **The two most serious items at the release left the strongest trace.** The hatches, which GAIA would have isolated and the crane support, which it would have held are the two items with the clearest downstream consequences.
2. **A package that looks drawn can still be unstable.** Primary steel scored 4 on completeness and 2 on stability. Without a stability dimension, the deck structure would have looked ready.
3. **The release was a date, not a freeze.** In the week after the kickoff, eight more entries reached the log, four of them still open at the end of DOII. The formal release and the moment the information stopped moving were different events and neither recorded what was accepted.
4. **Much of the overrun has explanations beyond information.** Limited engineering capacity, a design hard to build in thin plate, working methods and a compressed planning all play a part, and the study does not divide the hours between them. Information maturity and readiness to build are different things.
5. **That limit is useful.** A tool that is open about what it cannot see is easier to trust when it does raise a flag.

6.7 Across the three cases

The comparison is about mechanisms and patterns, not effects. There is no control group between vessels, and no claim that a decision on one project would have changed another. Table 6.12 puts the three cases side by side.

	Case A	Case B	Case C
Zone and role	Engine room · reference	Engine room · repetition	Superstructure deck · stress case
Level I	3 green · 4 N/E	1 red (vibration) · 6 N/E	1 amber (weight) · 1 green · 5 N/E
Level II at 1 or 2	Primary steel, foundations	Primary steel, foundations, HVAC, interior	Foundations 1; primary steel and exterior 2
Level III	10: 2 GO · 4 COND · 4 ISO · 0 WAIT	22: 12 GO · 6 COND · 3 ISO · 1 WAIT	4: 0 GO · 2 COND · 1 ISO · 1 WAIT
Package hours	+5.6%	+5.5%	+58.6%
Clearest matches	Generator foundations, walkway pillar, girder holes	Vibration (context), generator footprint, guest-lift pit	Storage hatches, crane support
Largest events with no item	Client-paid scope changes: switchboard, and piping scope decided eleven months later (outside GAIA)	Vibration (seen by Level I) · foundation strips added after the gate (outside GAIA)	Weight and class review window (Level I) · buildability in thin plate, engineering capacity (outside GAIA)
Main lesson	Focus without false alarm	The item list is not enough	Where the framework stops

Table 6.12: The three cases side by side.

1. **Maturity is not the same as looking finished.** Primary steel and foundations score 2 on stability in all three projects. Most things were drawn, but on bases that were still

moving: supplier data, scantlings waiting for a calculation, a vibration result. Punctuality hides the same gap. The hulls left on or near their dates, while the hours booked to correct sections after inspection rose well above their budget, or had no budget at all in Case C, with no baseline in the planning.

2. **GAIA sorts instead of sounding a general alarm.** In every case the items spread over at least three of the four decisions. Twenty-six of the 36 items would have moved on, freely or under a condition. WAIT appears twice, each time on one element and never in the healthy project.

3. **The strongest signals meet the clearest traces.** In every case the items flagged at the gate include those with the clearest later trace: the generator foundations, pillar and girder holes in Case A, the vibration warning, generator footprint and lift pit in Case B, the hatches and crane support in Case C. GAIA points to where to look, not to how much it will cost. Impact says how far a change would travel, not how many hours it would take: the generator foundations of Case A had a low impact and their trace stayed within the foundations, yet it was the clearest trace of the case.

4. **The largest disturbances came from outside the item list and they were of two kinds.** Some were inside GAIA, at the context level: the vibration warning in Case B, and the weight constraint and the class review window in Case C. Others were outside GAIA, because they were not a question of information at the release: client-paid scope changes in Case A, foundation strips added after the gate in Case B and, in Case C, a design hard to build in thin plate, limited engineering capacity and a planning without slack. The rule is simple. GAIA sees what is open or undecided at the release, in the item list or in the context. What is added or decided later, and how well a design can be built, lies outside it. The first kind of event is why GAIA has a context level. The second is where it ends, and why every link faces a rival explanation.

The context level was the hardest to rebuild, and that is a finding too. For most preconditions there was information, in logs, notes and the engineers' memory, but not a dated status for the whole package. GAIA's rule is strict: without dated evidence it marks N/E rather than guess with hindsight. That keeps the test honest, and it points to the gap Chapter 5 ended with: nobody recorded what was known and accepted at the moment of release. GAIA is designed to leave that record.

The same patterns show where the framework can be sharpened: an open acceptance question should not be released as GO. A red context signal should come with an exposed zone, an owner and a date, the class review window should be weighed against the date the receiving yard needs the information and a buildability check should be added where a difficulty is already known at the gate. These refinements, and the future research needed to test them, are discussed in Chapter 7.

6.8 Synthesis and answer to SQ4

"To what extent do the maturity scores generated by the tailored assessment framework show indicative consistency with historical rework and production performance when applied retrospectively to completed projects?"

SQ4 asked whether the signals GAIA gives at the gate are consistent with the rework and production performance recorded afterwards. The answer is: to a partial and indicative extent. In all three cases the items GAIA would have held back, isolated or conditioned include those with the clearest later traces and the weakest disciplines include those where extra effort appeared, though not one to one. Of the 14 items it would have released, eleven showed no documented downstream trace in the records reviewed, two shared a failure cost booked as a production error and one, a small access question, needed a small change before production. Level I caught a risk that the item list missed. The evidence therefore supports GAIA as a way to structure the release decision and to record the risk accepted. It does not show that GAIA predicts project performance or how much rework it would avoid. That needs a live release. Measured against the criteria of Section 6.1.2, GAIA meets two of the three conditions for a leading indicator by design: it is measured before production and it is tied to a decision at the gate. The third, a demonstrated correlation, is shown here only as indicative consistency.

The evidence has limits. Three cases, two in the same zone. Case C released no item at the gate, so it has no negative control. The gates were rebuilt from records that are in part several years old, by a researcher and engineers who knew the general outcome, though not the hours, quality records and change costs collected afterwards. Many context preconditions could only be marked N/E. Hours carry booking inaccuracies and overlapping rows, one planning was re-baselined and another has no baseline and the change register mixes settlement types. The cases are not the worst projects of the yard. These limits restrict how far the result can be generalised. They do not undo the matches observed.

The answer to "is it consistent?" is therefore a qualified yes. That is enough to move to the question that matters for practice, what it would take for the yard and the group to use GAIA, which Chapter 7 addresses.

DISCUSSION

Chapter 5 showed that information often crosses from DOII into Detail Engineering before it is stable and that what is accepted at that moment is rarely written down. Chapter 4 designed a way to make that decision explicit and Chapter 6 tested it on three finished deliveries. This chapter asks what that means in practice: what GAIA changes at a real release, what it costs and what is still missing before the yard can rely on it.

In one sentence: **GAIA helps to decide on the uncertainty that is already known at the release, the three cases support it in part and the next step is to try it on a live release, one that is still being decided.** Less rework is a reasonable expectation, but it has not been shown.

7.1 What the results mean

Chapter 6 gives three results. In all three cases, some items that GAIA would have held back, isolated or released under a condition left a trace later in the same element: changes, extra hours or recovery work. Many items it would have released left no trace in the records that were checked. And a large part of what went wrong downstream had other causes, most of all in the third case. Each case adds its own piece (Table 7.1).

Case	What it adds	Why it matters
Case A · engine room · reference case	A good project and the first application. It does not raise a general alarm: ten open items become three kinds of decision. Foundations and equipment, scored 2, overran by a quarter to a third, while the structure family stayed on budget.	A framework that flags everything on a healthy project would not be used twice.
Case B · engine room · recent project	The item list is needed but not enough. The strongest signal, the unknown vibration result, was written in the close-out note and in none of the open items. Many open items were normal development work.	Counting open items gives the wrong picture. Context needs its own level.
Case C · superstructure deck · stress case	GAIA works in another zone, another organisation and another scope list, and the two most serious items at the release left the clearest traces. But much of the overrun had other causes: a design that was hard to build, limited engineering capacity and a tight planning.	It shows where GAIA stops: mature information is not the same as being ready to build.

Table 7.1: What each case adds.

What the cases do not show must be just as clear. They show that signals and later events match in part. They do not show that GAIA predicts problems, and they show no saving. Whether GAIA helps while a decision is being taken can only be tested on a live release.

The test has one weak point that no retrospective study can remove: hindsight. The researcher and the engineers knew the general outcome of each case, though not the hours, quality records and change costs, which were collected separately. The rules of Section 6.2 limit this risk. Where an engineer's decision differed from the default rule, both are recorded in the annexes. In every such case the engineer was stricter and none of these deviations turned a released item into a held one or the reverse, so the counts of released and flagged items do not depend on them. Under the default rule, four items recorded as ISOLATE would have been CONDITIONAL and the WAIT of Case B would have been an ISOLATE. Only a live release, before anyone knows the outcome, removes hindsight completely.

7.2 What GAIA adds to the release decision

7.2.1 From drawn to confirmed

In the three packages, much of what crossed the gate was drawn while its basis could still move: supplier footprints, scantlings waiting for a calculation, a vibration result. Stability scored 2 for primary steel and foundations in all three cases. This matches the idea behind FEED MATRS that being defined is not the same as being confirmed. It also fits the work of Krishnan et al. [35] on starting work with preliminary information. The cases show information that kept changing more often than information that was missing: the model looked complete, but its bases kept moving. In Case C the two clearest traces came from the only two items that touched three disciplines. That is what research on change propagation predicts [36]: the more parts a change reaches, the higher its risk. Taking the lowest score as the maturity is a cautious rule of GAIA. The cases agree with it, but three cases cannot prove that it works better than an average.

7.2.2 From counting open items to deciding on them

Detail Engineering never starts with zero open items. The real question is not how many remain, but what to do with each one. The Stage-Gate literature warns that under time pressure a gate checks whether deliverables exist, not what state they are in. In a yacht yard there is also no real choice between going on and stopping. GAIA replaces one decision on the whole delivery with one decision per item: release it, release it with a condition, isolate the zone or hold it. In Case A the list already carried an informal risk mark, but the two items with the clearest later trace did not have it. The information was the same. The criteria were not. Table 7.2 shows how one of those items would look as a GAIA record. It is an example, not a record that existed at the time.

	As listed at the gate	Proposed GAIA release record
Item	Generator foundations: final supplier footprint needed before model update. Hours estimated. No risk mark.	Same description.
Source and coupling	—	Supplier. Medium coupling: structure and mechanical assembly.
Impact and difficulty	—	Low impact: a new footprint changes the foundation, its brackets and its holes, not the zone around it. Low difficulty: only the supplier's confirmation is missing.
Decision	None: the item goes to DE like any other.	CONDITIONAL: medium coupling with low impact (Figure 4.2).
What can proceed	—	Section structure, tank top and piping routes outside the foundation footprint. The foundation top plates, their brackets and the equipment holes wait for the confirmation.
Condition	—	Written footprint confirmation by the date the yard needs it; escalation if the date is missed.
Accepted risk and owner	—	Proceeding around a preliminary footprint until that date, to be accepted by project management with the yard as receiver.

Table 7.2: One Case A item, as listed at the gate and as a GAIA record would show it. Illustrative.

7.2.3 From a conversation to a record

The release decision already exists. Engineers ask whether a model can be accepted with an item still open, and they accept the consequences knowingly. What is missing is a record of it. Rebuilding the release state of three packages took weeks and several engineers, and some preconditions could still only be marked N/E. A frozen release record would keep together what was open, what was decided, under which condition and who accepted the remaining risk. The effort of this study suggests that such a record would make later reviews and learning easier. Only real use can prove it.

7.3 How GAIA fits the organisation

GAIA does not land in an empty space. Talks with project leadership, the yard's management team and the project engineers show that governance exists and works.

1. **Reviews and authority are in place.** Project leadership aims at zero open items at the Final Review, pushes back on owner changes before a section is released and reports to the board at each gate review, where the board decides whether the project moves on:

   "You always need to move forward. Otherwise there is a problem. So you need to make sure that you are in a good state."

2. **What is missing is consistency across organisations.** The management team described gate boxes ticked on informal information passed between companies and concluded that missing supplier information "should be written down" at the kickoff:

   "*We will continue, but listen, this is a risk we shouldn't be taking alone.*"

3. **Some statuses are unknown at the release.** Class comments are recorded in the classification society's software and at the moment a drawing goes to production it is not always clear whether they are closed. The management team proposed a class-status report at each kickoff.

4. **The yard already thinks in this structure, after the fact.** Its project evaluations are written as improvement point, cause, action, owner and date. GAIA asks for the same structure before the consequences.

Three simple rules follow and none of them gives anyone new authority. The sending and receiving parties classify the items together and project management keeps the decision to accept the risk. A disagreement is recorded with both views and settled by that same authority. If a condition or an expected date is missed, the item is escalated: it moves to ISOLATE or WAIT, or its risk is stated and accepted again. A GO item carries no condition and is not followed.

7.4 What it costs and what it can give back

How much extra time does GAIA take, and what can it give back? The cost can be estimated now. The benefit can be argued and compared in the same unit, but it has not been measured. The comparison covers the assessment only. GAIA finds and records the risk. Whether a fix, such as an extra week of engineering, is worth its price remains a decision for the project.

7.4.1 What it costs

GAIA adds four tasks to a review that already exists (Figure 4.3, Annex 4-C): freezing the record, classifying the carry-over items, recording the decisions and following up the conditions. Engineers already know their open items when they reach the review. What GAIA adds is a shared classification and a written record. How long it takes depends less on the method than on the items: how many there are and how hard they are to discuss. Many are routine DE work and are settled in a minute; a few need a real discussion. Its cost was estimated from the classification sessions of this study, which were not timed. Each session covered one functional zone. Classifying took about an hour for the ten items of Case A, including a first

explanation of the method, and about forty minutes for Case B, with some items checked later with colleagues. The four items of Case C also took about an hour, but part of that time went into finding them in the engineering log, which a live release would not need. Once an item was understood, most decisions took a few minutes. Follow-up could not be observed in finished projects. None of this includes rebuilding a release from years ago, which took weeks and is not part of normal use.

Task	Who	Estimated time per zone	Person-hours per zone
Prepare the frozen record	Project engineer	30–60 min: freeze the open-item list, draft Level I and II	0.5–1
Classification session	Issuer, receiver, project manager	30–60 min, within the Final Review	1.5–3
Record decisions	Project engineer	About 30 min	0.5
Follow-up during DE	Item owners	Not observed in the retrospective cases	To be measured in a pilot
Total per zone			2.5–4.5, follow-up excluded

Table 7.3: Effort GAIA adds per functional zone, estimated from the classification sessions of this study.

Similar assessments are short: a PDRI evaluation reportedly takes about two hours or less [55]. What they try to avoid is much larger. Rework accounted for about half of the cost growth in 161 building projects [23], and its direct cost alone often reaches about 5% of construction cost [24]. Late changes hurt labour productivity more than early ones [56], and in one study of 226 projects productivity never reached plan once change passed about a fifth of the labour hours [57]. Projects with a mature and accurate front-end design did about a quarter better on cost growth than weak ones [38]. Outside engineering, a one-page surgical checklist that only structures a conversation that already happens was followed by a fall in deaths from 1.5% to 0.8% and complications from 11% to 7% [58]. None of this proves that GAIA pays off. It does show that structuring a decision early is usually cheap compared with what it can avoid. Holding work back without need also costs something, in idle capacity or replanning, and that cost is not measured here.

7.4.2 What it can give back, and to whom

Beneficiary	Mechanism	How sure we are
Slob, operational	Knowing before steel is cut which work can go ahead and which needs a condition can reduce late changes to built structure.	Seen in some cases, not measured
Slob, commercial	A clearer basis to discuss scope and responsibility: what was open and accepted at the gate, what came later and what was an execution error. On its own, it does not decide who pays for a later change.	Depends on how the record is used
The group	Disagreements between sender and receiver come up earlier, coordination improves and yachts can learn from each other. Moving a cost from one yard to another is not a saving for the group.	Needs several projects
Everyone	Decisions taken on recorded facts that can be checked afterwards.	The problem is shown, the solution is not yet tested

Table 7.4: Benefits by beneficiary and by strength of evidence.

Two consequences do not show up in hours. A correction in built structure leaves deformation that the finishing yard has to fill and fair, as on the shell of Case B. And recovery work squeezed into a short period puts many people in one space, which the yard's management sees as a safety risk. Learning is the slowest benefit. A structural engineer described how a late request is absorbed with extra effort and then comes back on the next yacht. A project engineer warned that every yacht is a prototype, so lessons do not carry over by themselves. A record makes the comparison possible. Whether people will use it is not yet shown.

7.4.3 Is it worth it?

GAIA adds a few person-hours per zone to a meeting that already exists (Table 7.3). A single correction in built steel easily costs more: the two Case C hatches, redone with the section already in production, led to a change of 0.6% of the package budget for brackets and stiffening, partly class-driven, and the extra Case B pillar structure needed 1.3% of its budget in production. These figures show scale, not savings. The question is therefore not whether GAIA is cheap, but whether it is used well. It would do more harm than good if it held work back without need, produced records nobody reads or added a list or a meeting instead of using the existing ones. As the group's lean lead put it, a tool that is too complex is simply not used.

The larger benefit grows with use. In the first yachts, GAIA gives the engineers, who know their open items day by day, a shared and defensible basis to say when to continue, when to isolate and when to ask for a firm date. After a few yachts, the frozen records show which sources keep arriving late and which open items keep turning into rework, with real numbers instead of impressions. That gives the yard and the group a basis to move the fix upstream, to where the information is created and a correction is cheapest (Chapter 1). This is an expectation to test, not a result.

7.5 Where GAIA stops and what to improve

The main limit comes from Case C. Information maturity asks whether the design information is complete, stable and understood well enough to release. Being ready to build asks more: whether the yard has the skills, the capacity and a design that can be built well. GAIA looks only at the first. DECRIS also covers buildability and the execution approach. GAIA chose a narrower scope so that it fits in one review meeting and Case C shows the price of that choice. The cases point to a few targeted improvements (Table 7.5). They are proposals for a next version, not tested results, and none was applied to the cases they came from.

Finding	Proposed refinement	What remains to be tested
Case B: an open acceptance question was released as GO and later needed work	An open acceptance or verification question cannot be released as GO	Whether it catches such items without inflating conditions
Case B: a red context signal was coloured but not acted on	Every red or amber context signal gets an exposed zone, an owner and a need date	Whether owners take responsibility for context signals
Case C: class comments arrive after DE starts and their closure is unknown	A1 assessed on the class review window against the yard's need date, with a class-status report at each kickoff	Who compiles the report and how often
Case C: thin plate, curved shapes and flat bars explain much of the overrun	Known building difficulties are checked with production before release, with shared criteria for what counts as hard to build and a named person who approves the solution	Which criteria work, and whether it repeats an existing review
Structural engineer: a vibration response analysis takes about fifteen weeks	A red context signal can reorder the release within a delivery, starting with zones far from the source	Whether planning can absorb a different order
All three cases: the release state was hard to reconstruct. In Case C information kept arriving a week after the kickoff	A frozen release record at every delivery kickoff	Effort to prepare and maintain it

Table 7.5: Targeted refinements for a next version of GAIA.

GAIA was built for one release, from DOII to DE, but its logic is not tied to that gate. The earlier transitions, such as DOI to DOII, also pass on information that is not yet stable, and the same questions apply there: what is still open, how far a change would travel and who accepts the risk. The same holds after the hull leaves Slob, when outfitting, exterior finishing and commissioning build on what came before. Applying GAIA there would need other disciplines and criteria and it is left for future research.

7.6 Next step: a pilot

This thesis tested GAIA on past releases, as set out in its scope (Section 1.5). Using it on a live release is the natural next step, because it can measure clarity and added effort at once, and consistency with later events over the following months.

A first live check was attempted on a superstructure delivery of a fourth project, in Detail Engineering during this study. It could not be done: the open-item log of that delivery was not complete enough to support it, as a superstructure detail engineer explained. The cases point the same way. In the hull, both engine rooms had a dated record of their open items at the gate. In the superstructure, Case C had none, so its gate had to be rebuilt from the engineering log, and the log of the fourth project was also incomplete. Where keeping the record depends on each engineer, it is easily left behind, which is the first of the risks below.

The proposal is deliberately small: **one delivery, one functional zone, one kickoff, the existing meeting and the existing tools. Table 7.6 lists what it measures.** The items are classified by the issuing and receiving engineers together, the record is frozen, and the release proceeds as the project team decides.

What is measured	When	What it tells
Added time for preparation, classification and follow-up	At the kickoff and after	The real cost, confirming or correcting Table 7.3
Disagreements between issuer and receiver made visible	At the kickoff	Decision value, even if no rework follows
Decisions that change, or keep their outcome with a better justification and a condition	At the kickoff	Whether the structure adds clarity to experienced judgement
Accepted risks with an owner; conditions met on time	Weeks after	Whether the record is used or ignored
Downstream trace of the classified items	Months later	The first prospective evidence of consistency

Table 7.6: Pilot measures.

A pilot does not need to change decisions to be useful. It is also useful if it keeps a decision but makes its justification explicit, its condition traceable and its follow-up reliable. It is not useful if it takes clearly more time than the estimate of Table 7.3 and adds neither clarity nor follow-up. In that case GAIA should be simplified.

Four things could still make GAIA fail in practice, and the pilot should watch them. The first is a record filled in only to tick a box. The yard's management team already described gate boxes ticked on informal information (Section 7.3), and a record that nobody keeps up to date protects no one, as the attempt on the fourth project showed. The second is openness. Under schedule pressure, people tend to keep known problems to themselves [10], and a risk that slows the release is the hardest one to declare. The third is follow-up. Conditions and need dates only work if someone checks them during Detail Engineering, and that part could not be observed in this study (Table 7.3). The fourth is effort: GAIA has to stay light and must not hold work back without need (Section 7.4.3). None of these is a reason not to try GAIA. Each is a reason to watch it closely in the pilot.

CONCLUSIONS AND RECOMMENDATIONS

8.1 Answers to the research questions

Main research question. *To what extent can rework in hull construction be mitigated by systematically assessing information maturity at the design-to-production interface?*

Partly, in principle. A check at the release cannot remove all rework, but it has the potential to reduce the part that comes from uncertainty already known at that moment. It would do so by making that uncertainty visible, deciding what can go ahead and holding back the rest before steel is cut. The three cases are consistent with this. They do not show how much rework would be saved and part of the rework had other causes: designs that are hard to build, limited experience with such designs and lack of capacity. Only a live release can show the real effect.

SQ1. *What characteristics of engineering information determine its readiness for production release in complex, multi-disciplinary ETO construction environments?*

Three characteristics, judged per discipline: how complete the information is, how stable it is and how critical the open items are. Above them sits the context of the delivery: class, vibration, weight or supplier data can affect the whole package without ever appearing as an open item. In the cases, stability was the characteristic that most often held readiness back (Chapters 2 to 4, and the cases of Chapter 6).

SQ2. *How can these characteristics be measured and integrated into a quantitative, verifiable assessment framework at a quality gate?*

Through GAIA. Seven context preconditions are checked on dated evidence. Nine disciplines are scored on the three characteristics and the lowest score sets the maturity. Every open item is classified by source, coupling, impact and difficulty into one of four decisions, recorded with its condition, owner and accepted risk. The scores are ordinal rather than strictly quantitative, and they can be checked because every status rests on evidence that can be shown (Chapters 3 and 4).

SQ3. *What are the current causes of rework and inefficiency at the design-to-production interface in ETO shipbuilding, and how do planned information flows (Work-As-Imagined) diverge from the information actually delivered (Work-As-Done)?*

Rework does not come only from late information. It comes from information that crosses from DOII to DE without a shared idea of what “ready” means, from outside parties (suppliers, class, the owner) whose timing the planning cannot control, and from release decisions taken in conversation, without a consistent record. The group’s own measurement shows the gap: only 35% of structural engineering deliverables reached Detail Engineering on their planned date. What is planned as a clean handover is in practice a gradual transition (Chapter 5).

SQ4. *To what extent do the maturity scores generated by the tailored assessment framework show indicative consistency with historical rework and production performance when applied retrospectively to completed projects?*

To a partial and indicative extent. In the three cases, the items GAIA would have held back, isolated or conditioned included those with the clearest later traces and the weakest disciplines included those where extra effort appeared. Of the 14 items it would have released, eleven showed no documented downstream trace in the records reviewed. Level I caught a risk that the item list missed. There were also unclear links, one small possible miss and large disturbances with other causes, above all in Case C. The cases do not prove that GAIA can reliably tell high-risk from low-risk items, or predict what will happen (Chapter 6).

8.2 Contributions

	What the thesis contributes
Developed	A release framework for engineer-to-order shipbuilding under concurrent engineering. It links the context, each discipline and each open item to an explicit decision and it uses the lists and meetings that already exist. It adapts ideas from maturity indices, stage-gate practice, work on preliminary information and change propagation to a setting where stopping is not an option.
Observed	In three finished deliveries, in two zones and two organisations, several items GAIA would have held back or conditioned left a later trace, and the weakest disciplines included those where extra effort appeared. Of the 14 items GAIA would have released, eleven showed no documented downstream trace in the records reviewed. There were also unclear links, one possible false negative and much rework with other causes. Rebuilding past release decisions proved hard.
Still to be proven	That GAIA improves a decision while it is being taken, what it costs in practice and whether it reduces rework. This needs a live application and more projects.

Table 8.1: What was developed, what was observed and what remains to be proven.

8.3 Recommendations and future work

The advice to Slob and Feadship is to start small: **test GAIA on one live delivery kickoff before deciding anything larger** and measure the time it adds, the decisions it clarifies and the disagreements it brings to light (Section 7.6). Three cheap practices can start right away: freeze the list of open items at each kickoff, give every accepted risk an owner and a need date and have project management report at the kickoff the status of the signals that can still change the structure, such as class comments, strength and vibration results, supplier data and owner decisions.

Future work has two parts. The first is to make GAIA more reliable. A live application should show whether engineers can use it inside the existing Final Review, how much time it adds and whether the agreed conditions are followed. More applications should check whether two engineers score the same package alike and sharpen the scores of Level II and the rubric of Level I. Weights between disciplines should wait until there is data to support them.

The second part is to strengthen GAIA where the cases showed the most room and to let it grow with use.

A stronger Level I, connected upstream. In the three cases, the only red signal came from the context: the vibration warning of Case B. Yet most Level I preconditions ended as N/E, because the context is assessed at the release while it is created earlier, in DOI and DOII. Level I would give much more value if project management fed it as each signal arises, with a date, a status and the zones it may still change, so that the release reads a status instead of rebuilding one. Faced with a vibration impact “still unknown”, a member of the yard’s management team was clear: “That’s not a status. I cannot work with that.” Two signals deserve priority.

Vibration. A full response analysis takes about fifteen weeks and often reports in DOII or after the release. At each kickoff, Level I could record how far the analysis has come (global strength model, frequency model, response results), the frequency ranges the structure must stay out of and the zones a result could still change, such as pillars and long spans. Zones far from the source could then be released first.

Class comments. Plan-approval comments are kept in the classification society’s software and their status at the release is not always clear. Level I could list, for each drawing package of the delivery, whether it is submitted, approved or still has open comments and weigh the expected review time against the date the yard needs the drawings. In Case C the review took six weeks against two planned.

Buildability. A design can be complete and stable and still be hard to build. Adding it to the minimum of Level II would mix two questions: whether the information is mature and whether the design can be built well. It fits better as a separate flag per zone, next to the Level II profile, set with production on a few yes or no questions drawn from Case C: thin plate, flat bars instead of profiles, curved shapes, a team with little experience in the material and no earlier reference. Any yes would call for a short review with production before the release, with a named person who approves the solution.

A tool that learns. Every frozen record keeps what was open, what was decided and what happened next. After a few yachts, the records can show which sources keep arriving late and which items keep coming back. The generator footprint, for instance, was open at the release of both engine rooms. This is how a check at one gate becomes knowledge for the next design. It only works if GAIA is fed at every release and polished with each use: scores calibrated, rules sharpened and anything that adds no clarity removed.

Each of these steps should be tested on its own before it becomes part of GAIA. The aim is not a heavier gate, but a release that knows more for little extra effort.

The thesis started by asking whether information was mature enough. It ends with a more useful question for a yard that will always work on overlapping phases. Not whether everything is complete, but whether what remains open is visible, bounded, owned and accepted on purpose. GAIA is a way to answer that question at the moment it matters. In short, a framework has been built, retrospective matches have been observed and its value in daily use is still to be tested.

Release what is ready. Bound what is not. Record the decision.

BIBLIOGRAPHY

[1] De Vries Group, "Value stream mapping: Centrale Engineering and yacht hull construction," internal session reports, unpublished, 2026.

[2] M. Semini, D. E. G. Haartveit, E. Alfnes, E. Arica, P. O. Brett, and J. O. Strandhagen, "Strategies for customized shipbuilding with different customer order decoupling points," *Proc. Inst. Mech. Eng. M, J. Eng. Marit. Environ.*, vol. 228, no. 4, pp. 362–372, 2014, doi: 10.1177/1475090213493770.

[3] J. Wikner and M. Rudberg, "Integrating production and engineering perspectives on the customer order decoupling point," *Int. J. Oper. Prod. Manage.*, vol. 25, no. 7, pp. 623–641, 2005, doi: 10.1108/01443570510605072.

[4] J. Gosling and M. M. Naim, "Engineer-to-order supply chain management: A literature review and research agenda," *Int. J. Prod. Econ.*, vol. 122, no. 2, pp. 741–754, 2009, doi: 10.1016/j.ijpe.2009.07.002.

[5] E. W. Merrow, *Industrial Megaprojects: Concepts, Strategies, and Practices for Success*, 2nd ed. Hoboken, NJ, USA: Wiley, 2024.

[6] M.-H. Kim, E.-B. Lee, and H.-S. Choi, "Detail engineering completion rating index system (DECRIS) for optimal initiation of construction works to improve contractors' schedule-cost performance for offshore oil and gas EPC projects," *Sustainability*, vol. 10, no. 7, Art. no. 2469, 2018, doi: 10.3390/su10072469.

[7] M. Bandecchi, B. Melton, and F. Ongaro, "Concurrent engineering applied to space mission assessment and design," *ESA Bull.*, no. 99, pp. 34–40, 1999.

[8] Construction Industry Institute, *PDRI: Project Definition Rating Index – Industrial Projects, Version 5.0*, Implementation Resource 113-2. Austin, TX, USA: CII, 2019.

[9] M. Barnes, "Construction project management," *Int. J. Project Manage.*, vol. 6, no. 2, pp. 69–79, 1988, doi: 10.1016/0263-7863(88)90028-2.

[10] D. N. Ford and J. D. Sterman, "The liar's club: Concealing rework in concurrent development," *Concurrent Eng.: Res. Appl.*, vol. 11, no. 3, pp. 211–219, 2003, doi: 10.1177/106329303038028.

[11] R. I. Winner, J. P. Pennell, H. E. Bertrand, and M. M. G. Slusarczuk, "The role of concurrent engineering in weapons system acquisition," Inst. Defense Analyses, Alexandria, VA, USA, IDA Rep. R-338, 1988.

[12] X. Koufteros, M. Vonderembse, and W. Doll, "Concurrent engineering and its consequences," *J. Oper. Manage.*, vol. 19, no. 1, pp. 97–115, 2001, doi: 10.1016/S0272-6963(00)00048-6.

[13] E. T. Douma, "Identifying and improving the internal and external information exchange for shipbuilding processes," M.S. thesis, Delft Univ. Technol., Delft, The Netherlands, 2020.

[14] F. Schulze and P. Dallasega, "Lean and Industry 4.0 mitigating common losses in engineer-to-order theory and practice: An exploratory study," *Flex. Serv. Manuf. J.*, vol. 36, no. 3, pp. 780–820, 2024, doi: 10.1007/s10696-023-09503-z.

[15] B. Visser, “Bridging operations and management on the shipyard production floor: The application of FRAM within a shipbuilding environment,” M.S. thesis, Delft Univ. Technol., Delft, The Netherlands, 2025.

[16] E. Hollnagel, “Why is work-as-imagined different from work-as-done?,” in *Resilient Health Care, Vol. 2: The Resilience of Everyday Clinical Work*, R. L. Wears, E. Hollnagel, and J. Braithwaite, Eds. Farnham, U.K.: Ashgate, 2015, pp. 249–264.

[17] W. J. Hopp and M. L. Spearman, *Factory Physics*, 3rd ed. New York, NY, USA: McGraw-Hill/Irwin, 2008.

[18] E. S. Pound, J. H. Bell, and M. L. Spearman, *Factory Physics for Managers: How Leaders Improve Performance in a Post-Lean Six Sigma World*. New York, NY, USA: McGraw-Hill Education, 2014.

[19] J. D. C. Little, “A proof for the queuing formula: L = λW,” *Oper. Res.*, vol. 9, no. 3, pp. 383–387, 1961, doi: 10.1287/opre.9.3.383.

[20] G. Labovitz, Y. S. Chang, and V. Rosansky, *Making Quality Work: A Leadership Guide for the Results-Driven Manager*. New York, NY, USA: HarperBusiness, 1993.

[21] Construction Users Roundtable (CURT), “Collaboration, integrated information and the project life cycle in building design, construction and operation,” CURT, Cincinnati, OH, USA, White Paper WP-1202, 2004.

[22] American Institute of Architects and AIA California Council, *Integrated Project Delivery: A Guide*, version 1. Washington, DC, USA: AIA, 2007.

[23] P. E. D. Love, “Influence of project type and procurement method on rework costs in building construction projects,” *J. Constr. Eng. Manage.*, vol. 128, no. 1, pp. 18–29, 2002, doi: 10.1061/(ASCE)0733-9364(2002)128:1(18).

[24] B.-G. Hwang, S. R. Thomas, C. T. Haas, and C. H. Caldas, “Measuring the impact of rework on construction cost performance,” *J. Constr. Eng. Manage.*, vol. 135, no. 3, pp. 187–198, 2009, doi: 10.1061/(ASCE)0733-9364(2009)135:3(187).

[25] A. R. Hevner, S. T. March, J. Park, and S. Ram, “Design science in information systems research,” *MIS Quart.*, vol. 28, no. 1, pp. 75–105, 2004, doi: 10.2307/25148625.

[26] C. Wohlin, “Guidelines for snowballing in systematic literature studies and a replication in software engineering,” in *Proc. 18th Int. Conf. Eval. Assessment Softw. Eng. (EASE)*, London, U.K., 2014, Art. no. 38, doi: 10.1145/2601248.2601268.

[27] B. Flyvbjerg, “What you should know about megaprojects and why: An overview,” *Project Manage. J.*, vol. 45, no. 2, pp. 6–19, 2014, doi: 10.1002/pmj.21409.

[28] B. Flyvbjerg, “Top ten behavioral biases in project management: An overview,” *Project Manage. J.*, vol. 52, no. 6, pp. 531–546, 2021, doi: 10.1177/87569728211049046.

[29] A. Reumer, “Documentation flow analysis to improve project execution,” M.S. thesis, Delft Univ. Technol., Delft, The Netherlands, 2020.

[30] M.-H. Kim, E.-B. Lee, and H.-S. Choi, “A forecast and mitigation model of construction performance by assessing detailed engineering maturity at key milestones for offshore EPC mega-projects,” *Sustainability*, vol. 11, no. 5, Art. no. 1256, 2019, doi: 10.3390/su11051256.

[31] M.-H. Kim and E.-B. Lee, “A forecast model for the level of engineering maturity impact on contractor’s procurement and construction costs for offshore EPC megaprojects,” *Energies*, vol. 12, no. 12, Art. no. 2295, 2019, doi: 10.3390/en12122295.

[32] BIMForum, “Level of Development (LOD) Specification, Part I,” 2025. [Online]. Available: https://bimforum.org/resource/lod-level-of-development-lod-specification/ (accessed Oct. 19, 2026).

[33] S.-H. Jo, E.-B. Lee, and K.-Y. Pyo, “Integrating a procurement management process into critical chain project management (CCPM): A case-study on oil and gas projects, the piping process,” *Sustainability*, vol. 10, no. 6, Art. no. 1817, 2018, doi: 10.3390/su10061817.

[34] J. C. Mankins, “Technology readiness levels: A white paper,” NASA Office of Space Access and Technology, Washington, DC, USA, 1995.

[35] V. Krishnan, S. D. Eppinger, and D. E. Whitney, “A model-based framework to overlap product development activities,” *Manage. Sci.*, vol. 43, no. 4, pp. 437–451, 1997, doi: 10.1287/mnsc.43.4.437.

[36] P. J. Clarkson, C. Simons, and C. Eckert, “Predicting change propagation in complex design,” *J. Mech. Des.*, vol. 126, no. 5, pp. 788–797, 2004, doi: 10.1115/1.1765117.

[37] R. K. Yin, *Case Study Research and Applications: Design and Methods*, 6th ed. Thousand Oaks, CA, USA: SAGE, 2018.

[38] Construction Industry Institute, “Assessing the maturity and accuracy of front end engineering design,” CII, Austin, TX, USA, Res. Summary RS331-1, 2018.

[39] *Space Systems — Definition of the Technology Readiness Levels (TRLs) and Their Criteria of Assessment*, ISO 16290:2013, International Organization for Standardization, 2013.

[40] R. S. Kaplan and D. P. Norton, *The Balanced Scorecard: Translating Strategy into Action*. Boston, MA, USA: Harvard Business School Press, 1996.

[41] R. S. Kaplan and D. P. Norton, “The balanced scorecard—Measures that drive performance,” *Harvard Bus. Rev.*, vol. 70, no. 1, pp. 71–79, 1992.

[42] D. J. Snowden and M. E. Boone, “A leader’s framework for decision making,” *Harvard Bus. Rev.*, vol. 85, no. 11, pp. 68–76, 2007.

[43] B. Flyvbjerg and D. Gardner, *How Big Things Get Done*. New York, NY, USA: Currency, 2023.

[44] R. G. Cooper, “Perspective: The Stage-Gate idea-to-launch process—Update, what’s new, and NexGen systems,” *J. Prod. Innov. Manage.*, vol. 25, no. 3, pp. 213–232, 2008, doi: 10.1111/j.1540-5885.2008.00296.x.

[45] R. G. Cooper and S. J. Edgett, “Best practices in the idea-to-launch process and its governance,” *Res.-Technol. Manage.*, vol. 55, no. 2, pp. 43–54, 2012, doi: 10.5437/08956308X5502022.

[46] *Failure Modes and Effects Analysis (FMEA and FMECA)*, IEC 60812:2018, International Electrotechnical Commission, 2018.

[47] R. G. Cooper, “Third-generation new product processes,” *J. Prod. Innov. Manage.*, vol. 11, no. 1, pp. 3–14, 1994, doi: 10.1016/0737-6782(94)90115-5.

[48] D. K. Sobek II, A. C. Ward, and J. K. Liker, “Toyota’s principles of set-based concurrent engineering,” *Sloan Manage. Rev.*, vol. 40, no. 2, pp. 67–83, 1999.

[49] H. G. Ballard, “The Last Planner System of production control,” Ph.D. dissertation, Univ. Birmingham, Birmingham, U.K., 2000.

[50] S. S. Stevens, “On the theory of scales of measurement,” *Science*, vol. 103, no. 2684, pp. 677–680, 1946, doi: 10.1126/science.103.2684.677.

[51] J. B. Bowles, “An assessment of RPN prioritization in a failure modes effects and criticality analysis,” in *Proc. Annu. Reliab. Maintainability Symp.*, Tampa, FL, USA, 2003, pp. 380–386, doi: 10.1109/RAMS.2003.1182019.

[52] E. Kujawski, “Analysis and critique of the system readiness level,” *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 43, no. 4, pp. 979–987, 2013, doi: 10.1109/TSMCA.2012.2209868.

[53] M. Rother and J. Shook, *Learning to See: Value Stream Mapping to Add Value and Eliminate Muda*. Brookline, MA, USA: Lean Enterprise Institute, 1999.

[54] T. Ohno, *Toyota Production System: Beyond Large-Scale Production*. Cambridge, MA, USA: Productivity Press, 1988.

[55] G. E. Gibson Jr., P. R. Dumont, and J. R. Fish, “Scope management using the project definition rating index (PDRI),” in *Proc. PMI 27th Annu. Seminar/Symp.*, Boston, MA, USA, 1996.

[56] W. Ibbs, “Impact of change’s timing on labor productivity,” *J. Constr. Eng. Manage.*, vol. 131, no. 11, pp. 1219–1223, 2005, doi: 10.1061/(ASCE)0733-9364(2005)131:11(1219).

[57] W. Ibbs, “Construction change: Likelihood, severity, and impact on productivity,” *J. Legal Aff. Dispute Resolut. Eng. Constr.*, vol. 4, no. 3, pp. 67–73, 2012, doi: 10.1061/(ASCE)LA.1943-4170.0000089.

[58] A. B. Haynes et al., “A surgical safety checklist to reduce morbidity and mortality in a global population,” *N. Engl. J. Med.*, vol. 360, no. 5, pp. 491–499, 2009, doi: 10.1056/NEJMsa0810119.

Annexes to Chapter 2

Annex 2-A. Literature overview

This annex lists the key references of the literature review (Chapter 2) by year.

Year	Title / Authors	Focus / Contribution	Applicability to Thesis
1961	**Proof of the Queuing Formula L = λW (Little, J.D.C.)**	**Fundamental relationship between WIP, throughput and lead time.**	**Theoretical basis for understanding how immature information (defective WIP) increases lead time.**
1988	Toyota Production System (Ohno, T.)	Classification of seven wastes in manufacturing systems.	Framework for categorizing waste in ETO information flows (rework, waiting, overprocessing).
1992	The Balanced Scorecard (Kaplan, R.S. & Norton, D.P.)	Distinction between leading and lagging indicators for performance management.	Justifies the need for a predictive (leading) maturity index rather than reactive (lagging) rework measurement.
1995	**Technology Readiness Levels (Mankins, J.C. / NASA)**	**Definition and standardized scale of technology maturity (TRL 1-9).**	**Provides the academic baseline for evaluating the risk of innovative "Specials" in ETO projects.**
1997	A Model-Based Framework to Overlap Product Development Activities (Krishnan, V., Eppinger, S.D. & Whitney, D.E.)	When downstream work can start on preliminary upstream information: evolution and sensitivity.	Basis for deciding what can continue on open items at the gate (Level III).
1999	Learning to See (Rother, M. & Shook, J.)	VSM methodology for mapping material and information flows.	Diagnostic tool for identifying waste at phase transitions in the design-to-production chain.
2004	**MacLeamy Curve (CURT)**	**Illustrates decreasing ability to influence cost as project advances.**	**Justifies Front-End Loading and the 1-10-100 Rule for measuring maturity before production.**
2004	Predicting Change Propagation in Complex Design (Clarkson, P.J., Simons, C. & Eckert, C.)	A change to one part can spread to connected parts.	Basis for the impact dimension of open items (Level III).
2007	A Leader's Framework for Decision Making (Snowden, D.J. & Boone, M.E.)	Cynefin framework: Complex vs Complicated decision domains.	Used as a design lens: makes the evidenced part of the gate decision more analysable.
2008	**Factory Physics (Hopp, W.J. & Spearman, M.L.)**	**Dynamics of variability, bottlenecks and WIP control in production chain.**	**Theoretical basis for treating immature information as defective WIP in the production pipeline.**
2008	**The Stage-Gate Idea-to-Launch Process: Update and NexGen Systems (Cooper, R.G.)**	**Update of the Stage-Gate model and its common failure modes.**	**Theoretical basis for gate governance. Identifies failure mode where gates become administrative checkpoints.**
2014	Factory Physics for Managers (Pound, E.S. et al.)	Accessible translation of Factory Physics principles for management application.	Supports the argument that controlling information quality at gates reduces variability and lead time.

Year	Title / Authors	Focus / Contribution	Applicability to Thesis
2018	Detail Engineering Completion Rating Index System, DECRIS (Kim, M.-H. et al. — Sustainability)	DECRIS, applied to 13 offshore EPC projects at the start of fabrication.	Core model. Mathematical basis for measuring engineering maturity and its correlation with construction labour inefficiency and schedule deviation.
2018	**Integrating a Procurement Management Process (Jo, S.-H. et al.)**	**Impact of procurement delays on critical chain management in EPC projects.**	**Justifies the inclusion of VFI timeliness as a factor in maturity assessment.**
2018	FEED MATRS (Construction Industry Institute / CII)	Dual-dimension assessment separating maturity and accuracy scores.	Key precedent for separating geometric completeness from data reliability (completeness vs stability in this thesis).
2018	IEC 60812 Failure Modes and Effects Analysis (International Electrotechnical Commission)	Criticality as the combination of severity and likelihood.	Basis for residual criticality and the ranking of open items.
2018	Case Study Research and Applications (Yin, R.K.)	Case study design, pattern matching and rival explanations.	Method for the retrospective evaluation of the framework (Chapter 6).
2019	**Project Definition Rating Index — PDRI (Construction Industry Institute)**	**Maturity measurement tool for Front-End Planning in construction.**	**Commercial baseline for maturity assessment. Demonstrates correlation with cost/schedule performance but limited applicability to late-stage engineering interfaces.**
2019	**A Forecast Model for the Level of Engineering Maturity (Kim, M.-H. et al. — Energies)**	**Correlation between engineering maturity and construction cost escalation.**	**Connects information immaturity directly to financial impact in production-heavy environments.**
2020	Information Exchange for Shipbuilding Processes (Douma, E.T.)	Analysis of the impact of prematurely releasing information in shipbuilding.	Justifies the risk of pushing immature 3D models to production in ETO shipbuilding.
2020	**Documentation Flow Analysis (Reumer, A.A.)**	**The value of information and batching of documents in shipbuilding.**	**Supports the analysis of VFI delays and information flow bottlenecks.**
2023	How Big Things Get Done (Flyvbjerg, B. & Gardner, D.)	Dynamics of megaproject failure: optimism bias and planning fallacy.	Justification for investing effort in Front-End planning ("think slow, act fast").
2024	**Lean and Industry 4.0 Mitigating Losses (Schulze, F. & Dallasega, P.)**	**Scientific classification of engineering waste and losses.**	**Structure for categorizing failures in the design-to-production phase.**
2025	Bridging Operations and Management (Visser, B.J.)	Discrepancies between management perception (WAI) and operational reality (WAD).	Academic foundation for the WAI vs WAD internal analysis methodology.
2025	**Level of Development Specification (BIMForum)**	**Standardization of 3D model reliability by development phases (LOD 100-400).**	**Provides the geometric maturity scale adapted for the completeness dimension of the assessment framework.**

Table 2-A.1: Key references of the literature review, by year.

Annexes to Chapter 4

Annex 4-A. Traceability of GAIA elements to the literature

This annex traces every element of GAIA (Chapter 4) to the literature it comes from.

GAIA element	Anchor in the literature review
Level I (context screening, upstream inheritance)	PDRI and DECRIS assess the basis of the project before the deliverables. TRL for new technology.
Level II dimensions (completeness, stability)	FEED MATRS (maturity vs accuracy). PDRI (rubric with evidence). LOD (geometric progression).
Level II as ordinal, indicative scale	Maturity scales: ordinal by nature, valuable in comparison, calibrated only with sample size.
Residual criticality (impact and difficulty, Table 4.3)	Criticality matrices in risk engineering (FMECA, IEC 60812 [46]). (Section 2.8)
Coordination coupling	Overlapping on preliminary information and change propagation (Section 2.8): how much other work depends on the item. Counting the disciplines is a design choice of GAIA.
Impact dimension	Engineering change propagation (Section 2.8).
Resolution difficulty	Cynefin: clear / complicated / complex, operationalised at item level.
Four decisions per item	Overlapping on preliminary information (Section 2.8), Stage-Gate decisions, and the boxing practice observed in Chapter 5.
Gate output as decision + rationale	Stage-gate theory: what a well-designed gate produces.
Validation by indicative consistency	Case study methodology (Yin, Section 6.1.1), SQ4.

Table 4-A.1: Traceability of every GAIA element to the literature review.

Annex 4-B. Level I rubric

The rubric below assigns the state of each Level I precondition. A1 and A5 are stated as the working version and should be validated with the people who own those variables upstream. Where no dated evidence exists, the state is N/E.

Precondition	Green	Amber	Red
A1 Class / FEM / Vibration	Class approval with no open comments affecting the delivery, strength calculations up to date for the scope, and no open vibration topics. All three at once.	Comments or analyses pending, with a committed resolution date and no identified structural impact.	Comments or analyses pending without a date, or with possible impact on primary structure.
A2 VFI / Major equipment	Written supplier confirmation received for all geometry-fixing equipment in the delivery.	Some confirmation pending, with a date committed by the supplier.	Confirmation pending without a date, or the supplier has warned of possible changes.
A3 Owner / design decisions	No open owner or design decisions affecting the delivery.	Open decisions with a decision date set, and their effect bounded to identified zones.	Open decisions affecting structure, equipment or several zones or without a date.
A4 Critical clashes	No open clashes in primary structure or major equipment.	Open clashes with a solution proposed and agreed between the affected disciplines.	Open clashes without an agreed solution.
A5 Weight / CoG	Estimate within tolerance.	Deviations under study, with their effect bounded to identified zones.	A deviation that requires changes to structure or to equipment.
A6 P&IDs / line lists	Schematics and routings frozen for the delivery.	Minor revisions ongoing, with no effect on penetrations already defined.	Schematics under revision with effect on penetrations or structural passages.
A7 TRL	No unproven technology in the delivery.	New technology present, with a dated maturation and testing plan.	New technology without a maturation plan, or with pending tests that can change geometry.

Table 4-B.1: Level I rubric.

Annex 4-C. The gate today and with GAIA

The meeting, the people and the open-item list stay the same. GAIA adds a short preparation, a structured classification, a written record and a focused follow-up.

Step	Today	With GAIA	Who	Added work
Before the DOII Final Review	Open-item list kept by engineering	Same list, plus the frozen record: context evidence and integration-list status	Level I context by project management	Preparing the record
At the Final Review (Gate 4)	Items discussed, delivery released	Each open item classified, then GO, CONDITIONAL, ISOLATE or WAIT	Issuer and receiver together	Classification and resolving disagreements
Decision	Implicit, in the conversation	Recorded: condition, need date, owner, accepted risk	Project management	Recording
During DE	Open items followed informally	Only CONDITIONAL, ISOLATE and WAIT items followed until closed or escalated	Owner of each item	Follow-up
After the project	Evaluation written afterwards	Evaluation compared with the release record	Project team	None beyond today

Table 4-C.1: The gate step by step, today and with GAIA.

Annexes to Chapter 5

Annex 5-A. The seven wastes of the internal VSM exercise

This annex describes the seven wastes of the internal VSM exercise, summarised in Section 5.4.3.

#	Waste	Manifestation	Relevance for this thesis
1	Rework and late changes	Structural calculation results, class comments and supplier information arrive late. Engineering continues while production has already started.	Directly relevant: immature input crossing the gate becomes rework. This is what the assessment aims to make visible.
2	Waiting for information and approvals	Downstream phases wait for final engineering, class approvals and supplier data. There is no shared operational definition of "ready".	Core of this thesis: the gate lacks a definition of ready.
3	Duplication of work and conversions between systems	No single 3D source of reference; model conversion between phases consumes significant hours.	Outside the scope of this thesis: a tooling and systems issue that a release assessment does not solve.
4	Load imbalance and batching	Simultaneous releases of large batches of sections. Overload and multitasking in the receiving phases.	Only indirectly related: batching concentrates immaturity at the gate, but the framework does not address batch sizing itself.
5	Insufficient chain coordination	Lack of structural project consultation between the organisations. Each organisation optimises its own part.	Partially relevant: a shared release assessment can support the conversation but it does not by itself fix chain coordination.
6	Misaligned and unstable assumptions	Arrangements, specifications and equipment assumptions incomplete or subject to change. Pressure to continue.	Directly relevant: information that exists but will change. This is the reason for the separation of completeness and stability in Chapter 4.
7	Local versus system efficiency	Departments are individually busy and efficient, while project flow and delivery reliability suffer.	Outside the direct scope: a management and incentive issue beyond a gate assessment.

Table 5-A.1: The seven wastes of the internal VSM: manifestations and relevance.

Annexes to Chapter 6

Each case has the same five tables: Level I, Level II, Level III, pattern matching and rival explanations. Rows in bold are the ones the chapter relies on. In this version the cases are anonymised and the detailed records of hours, planning, quality and change orders are not included.

Annex 6-A. Case A

This annex gives the detailed tables of Case A (Section 6.4).

What GAIA shows at the gate

Precondition	Status	Evidence at the gate
A1 Class / FEM / vibration	N/E	No package-level status recovered. The pillar issue (99) is local and sits in Level III.
A2 Supplier data (VFI)	N/E	No supplier register. Supplier dependencies located as items 78, 80, 100.
A3 Owner / design decisions	N/E	No dated decision log. Later testimony describes a cooperative owner.
A4 Critical clashes	**Green**	**Known clashes listed as located items (77, 112, 127); none unbounded.**
A5 Weight / CoG	**Green**	**Weight-reduction work on record; no open weight uncertainty in the engine room.**
A6 P&IDs / routing	N/E	No frozen status recovered; 112 and 127 are local.
A7 Technology	**Green**	**No new technology in the engine room.**

Table 6-A.1: Case A. Level I: the seven preconditions at the gate.

Discipline	Completeness	Stability	Residual criticality	Maturity	Basis
Primary steel	**4**	**2**	**3**	**2**	**Structure developed; pillar and frame details open (99, 131)**
Secondary steel	3	3	3	3	Manhole and pillar interfaces open (77, 99)
Foundations / major equipment	**3**	**2**	**2**	**2**	**Geometry on unconfirmed supplier footprints (78, 80, 100)**
Piping / penetrations	3	3	4	3	Routing exists; piping check pending (112)
HVAC	3	3	4	3	Final penetration check pending (127)
Interior interfaces	4	4	3	3	Door interface of item 125
Exterior	N/A	N/A	N/A	N/A	No exterior scope in the zone
ELE / ICT · Class	N/E	N/E	N/E	N/E	No reliable gate record

Table 6-A.2: Case A. Level II. Maturity is the lowest of the three dimensions.

Nr	Open item	Source	Disciplines	C / I / D	Decision
77	Engine-room manholes colliding with generator foundations, cross-over and stair	Supplier	CNS-WTB	M / M / M	ISOLATE
78	**Generator foundations: final supplier footprint needed**	**Supplier**	**CNS-WTB**	**M / L / L**	**CONDITIONAL**
80	Lift foundations: stiffer proposal, yard can take over	Mechanical eng.	CNS	L / L / L	GO
89	Door at frame 28: rotation defined	Design change	CNS	L / L / L	GO
99	**Pillar in walkway: solution under discussion**	**FEM / vibration**	**CNS-WTB**	**M / M / M**	**CONDITIONAL**
100	Lift service space against fixed frame 30	Mechanical eng.	CNS-WTB	M / M / H	ISOLATE
112	**Holes in girders: piping check still needed**	**Mechanical eng.**	**CNS-WTB**	**M / L / M**	**CONDITIONAL**
125	Sliding doors at frame 32: structural reduction to coordinate	Mechanical eng.	CNS-WTB-INT	H / M / L	ISOLATE
127	Ventilation penetrations: final check pending	Mechanical eng.	CNS-WTB	M / L / L	CONDITIONAL
131	Lightening holes in frames: confirmation pending	Mechanical eng.	CNS-WTB	M / M / M	ISOLATE
128	Exhaust top foundations (marked done)	Mechanical eng.	CNS	L / L / L	Closed control

Table 6-A.3: Case A. Level III. C / I / D = coupling, impact, resolution difficulty. Decisions reviewed item by item with the project engineer.

Note. Deviations from the default rule: for items 77 and 131 the default is CONDITIONAL (medium coupling, medium impact, medium difficulty). The project engineer chose the more restrictive ISOLATE. Both the default and the recorded decision are kept in the case file.

Does it match?

Nr	Decision	What happened	Verdict
78	**CONDITIONAL**	**Foundation revisions in the same family; foundations +31%; supplier information 12-14 weeks late**	**Consistent**
99	**CONDITIONAL**	**Connection and vibration damping on the same pillar**	**Consistent**
112	**CONDITIONAL**	**Penetrations and relocation of installed equipment**	**Consistent**
77	ISOLATE	Work in the generator zone; no record names the element	Partial
127	CONDITIONAL	Shared with 112; penetration hours neutral	Partial
131	ISOLATE	Frame rework booked as execution failure	Partial
80, 89	GO	No trace	Negative control
100	ISOLATE	No specific trace	Inconclusive
125	ISOLATE	Element no longer identifiable	Inconclusive
128	Closed	Exhaust elbows 52% under budget; its large change order turned out to include other work	Closed control

Table 6-A.4: Case A. Pattern matching per item.

Rival explanation	Applies to	How it is handled
Foundation hours include other foundations	78	Link kept at family level, not claimed item by item
Normal detail development	99, 127	Accepted as partial explanation
Several changes in the same zone	112	Shared attribution declared
Execution failure	131	Booked as yard failure cost. Not counted as a hit
Client-paid scope change	Switchboard, piping	Outside GAIA's evidence. No item existed

Table 6-A.5: Case A. Rival explanations.

Annex 6-B. Case B

This annex gives the detailed tables of Case B (Section 6.5).

What GAIA shows at the gate

Precondition	Status	Evidence at the gate
A1 Class / FEM / vibration	**Red**	**Close-out note: "impact of the vibration analysis still unknown". Not located in any item**
A2 Supplier data (VFI)	N/E	No package-level status. The generator dependency is located in item 229
A3 Owner / design decisions	N/E	No dated decision log; the lift dependency is located in item 173
A4 Critical clashes	N/E	No package-level clash status
A5 Weight / CoG	N/E	No status at the gate
A6 P&IDs / routing	N/E	Piping pre-routed on the DOII model, as expected at this stage
A7 Technology	N/E	No record either way

Table 6-B.1: Case B. Level I: the seven preconditions at the gate.

Discipline	Completeness	Stability	Residual criticality	Maturity	Basis
Primary steel	**4**	**2**	**3**	**2**	**Scantlings not confirmed while vibration was open**
Secondary steel	4	4	4	4	Local items only; negative control of the level
Foundations / major equipment	**2**	**2**	**2**	**2**	**Generator footprint not confirmed (229)**
Piping / penetrations	4	3	5	3	Planned development level at DOII
HVAC	3	2	3	2	Owner change to the air-conditioning room not yet absorbed (341)
Interior interfaces	3	2	2	2	Guest-lift arrangement not fixed (173)
Exterior	N/A	N/A	N/A	N/A	No exterior scope in the zone
ELE / ICT · Class	N/E	N/E	N/E	N/E	No gate record; vibration counted once, in A1

Table 6-B.2: Case B. Level II: maturity per discipline.

Nr	Open item	Source	C / I / D	Decision
229	**Adjust construction for the generator footprint**	**Supplier**	**M / H / M**	**WAIT**
173	**Guest-lift pit construction**	**Design change**	**M / H / M**	**ISOLATE**
301	Void adjustment in the engine room	Design change	M / M / M	ISOLATE
341	Structure for the recirculation duct after the air-conditioning room change	Owner decision	M / M / M	ISOLATE
181	Construction in the casing open again	Coordination	M / M / M	CONDITIONAL
339	Is the casing above the upper section correct?	Design verification	M / M / L	CONDITIONAL
179	Move stiffeners behind cabinets	Coordination	M / L / L	CONDITIONAL
183	Frame 33 port side: flip profile	Coordination	M / L / L	CONDITIONAL
235	Remove surplus manhole	Coordination	M / L / L	CONDITIONAL
260	Adjust manholes	Coordination	M / L / L	CONDITIONAL
343	**Are the void tanks sufficiently accessible?**	**Acceptance question**	**L / L / L**	**GO**
182	Sludge-tank manhole to be corrected in the model	Model transfer	L / L / L	GO
184	Frame 33 starboard: flip profile	Design change	L / L / L	GO
236	Plate at pillars and stairs may be removed	Review finding	L / L / L	GO
270	Enlarge lightening hole in transverse bulkhead	Review finding	L / L / L	GO
288	Structural option (finishing-yard designer)	Design change	L / L / L	GO
289	Frames through top plate	Review finding	L / L / L	GO
302	Manholes frames 33-36 unnecessary	Review finding	L / L / L	GO
303	Bracket support frames 47/48 missing	Review finding	L / L / L	GO
304	Sea inlet: smaller opening in bulkhead	Review finding	L / L / L	GO
322	Sea-inlet boxes: vent holes in shell	Review finding	L / L / L	GO
358	Brackets added after the 100% review	Class	L / L / L	GO

Table 6-B.3: Case B. Level III: 22 carry-over items classified by the project engineer (12 GO, 6 CONDITIONAL, 3 ISOLATE, 1 WAIT). Excluded and declared: 14 items of normal DE work.

Note. Deviations from the default rule: 229 has ISOLATE as default (medium coupling, high impact, medium difficulty); the project engineer recorded WAIT on the tank-top and foundation scope ("footprint from supplier not final at gate; box tank top and foundation zone, release the rest"). 301 and 341 have CONDITIONAL as default; the engineer chose ISOLATE because the void zone and the duct route could be fenced while the tank arrangement and the owner change settled. Item 358 was recorded on the day of the DE Final Review.

Does it match?

Signal	Decision	What happened	Verdict
A1 vibration	**Red**	**Extra pillar structure, mostly in production; structural overrun with recovery tail**	**Consistent (context)**
229	**WAIT**	**Related change solved in engineering and preparation; foundations under budget**	**Consistent**
173	**ISOLATE**	**Change on the same element**	**Consistent**
301	ISOLATE	Shares a change with 343	Partial
341	ISOLATE	Originating owner change before the gate; no record isolates the item	Inconclusive
339, 181	CONDITIONAL	Coordination in the same family; no one-to-one change	Inconclusive
343	GO	A change creates the access that was in question	Possible false negative
304, 322	GO	Failure cost booked as production error	Execution error
179, 183, 235, 260	CONDITIONAL	Resolved routinely; no material trace	Inconclusive
Other GO items	GO	No trace	Negative control
14 normal-DE items	Excluded	Piping family finished 8% under budget	Control

Table 6-B.4: Case B. Pattern matching per item and signal.

Rival explanation	Applies to	How it is handled
Scope added after the gate	Foundation strips, the largest change of the package	No item existed; outside GAIA
Execution on the upper section	Largest failure cost of the package	Outside GAIA
Later owner and class decision	Switchboard stairs	Created after the reviews. Not a gate item
Booking accuracy	Hours and change orders in general	Hours were not always booked against the change that caused them. Hours are read as direction, not as exact values

Table 6-B.5: Case B. Rival explanations.

Annex 6-C. Case C

This annex gives the detailed tables of Case C (Section 6.6).

What GAIA shows at the gate

Precondition	Status	Evidence at the gate
A1 Class / FEM / vibration	N/E	One welding list approved and one local check pending; no plan-approval status for the package
A2 Supplier data (VFI)	N/E	Crane web proposal on record; no supplier register for crane, davit or hatches
A3 Owner / design decisions	N/E	Exterior form revised three times; no pending decision with a date
A4 Critical clashes	N/E	One source-model conflict at the hatches; no clash report
A5 Weight / CoG	**Amber**	**Nine scantling reductions in two weeks; one change adding 900 to 1,000 kg "to investigate"; strength to be recalculated before go-ahead**
A6 P&IDs / routing	N/E	No routing status before the gate
A7 Technology	**Green**	**Conventional steel construction; thin-plate difficulty is buildability, not technology**

Table 6-C.1: Case C. Level I: the seven preconditions at the gate.

Discipline	Completeness	Stability	Residual criticality	Maturity	Basis
Primary steel	**4**	**2**	**4**	**2**	**Modelled; scantlings changed for weight, strength recalculation pending (R3)**
Secondary steel	3	3	3	3	Storage and hatch details on pending checks (R1, R2)
Foundations / major equipment	**3**	**2**	1	1	Crane structure on a web proposal, no supplier basis (R4: high impact, high difficulty)
Piping / penetrations	3	3	3	3	Piping below the hatches attached to R2
Exterior	**3**	**2**	**3**	**2**	**Exterior definition still being revised (R2)**
HVAC · ELE / ICT · Class	N/E	N/E	N/E	N/E	No gate record
Interior interfaces	N/A	N/A	N/A	N/A	Exterior functional zone

Table 6-C.2: Case C. Level II. Completeness against the DO-2 scope guideline for the superstructure.

Nr	Open item	Source	C / I / D	Decision	Condition
R1	Storage: closure, drainage and mass not agreed; plate adds 900-1,000 kg	Design verification	L / M / M	CONDITIONAL	Mass and strength effect accepted by the design office before work preparation
R2	**Storage hatches on two conflicting source models**	**Coordination**	**H / M / L**	**ISOLATE**	**Openings and brackets not committed until the hatch model is released**
R3	Lightening holes in primary structure awaiting check	Design verification	H / L / L	CONDITIONAL	Holes not cut until the check is received
R4	**Crane support structure on a web proposal not accepted**	**Design change**	**H / H / H**	**WAIT**	**Web solution and crane interface confirmed before the beams are committed**

Table 6-C.3: Case C. Level III: 4 open items at the gate (2 CONDITIONAL, 1 ISOLATE, 1 WAIT), grouped from the 8 log entries still open that day. Classified by the superstructure project engineer and an engineer from work preparation.

Note. Entries logged in the days after the gate are read as downstream evidence: an opening for a duct penetration, a deck gutter, a plate for the hatch locking pins and the davit foundation.

Does it match?

Signal	Decision	What happened	Verdict
R2	**ISOLATE**	**Hatch openings corrected during DE; both hatches redone in production; change of 0.6% of the package budget**	**Consistent**
R4	**WAIT**	**Beams moved within two weeks of the gate; repositioned and web dimensions revised about seven weeks after it; flat bars added; absorbed in DE**	**Consistent**
R1	CONDITIONAL	Drain holes added about nine months after the gate; no change order	Partial
R3	CONDITIONAL	Hole closure requested about five months after the gate; identity not confirmed	Partial
A5 weight	**Amber**	**Stiffeners reversed after the strength calculation; flat-bar change during DE; later stiffening required by class**	**Consistent (context)**
Exterior (Level II)	2	Form revisions five days and two months after the gate; deck gutter released two months after; railing pots moved in production	Consistent (context)
A1 class	N/E	Class comments from about four months after the gate; six weeks needed against two planned	Not rated at the gate

Table 6-C.4: Case C. Pattern matching per item and signal. No item was released at the gate, so there is no negative control.

Rival explanation	What it explains	How it is handled
Limited engineering capacity. Flat bars instead of profiles	Section building, straightening, thin-strip corrections	Outside GAIA: capacity and buildability
Planning without slack	Amplifies every other cause	Outside GAIA
Class review window and later class comments on approved details	Class-related changes, some in production	The window falls under Level I; later comments on approved details are outside GAIA
Heat input and working methods in thin plate	Deformation and recovery effort	Outside GAIA. Not quantifiable per section
Events with no signal at the release	Removing hard spots (1% of budget), a piping clash (0.5%)	Declared as blind spots

Table 6-C.5: Case C. Rival explanations. No percentages are assigned between them.

Annex 6-D. Sources

This annex lists the documents and people consulted for each case (Section 6.1.1).

Case	Documents	People consulted
Case A	DOII close-out list; delivery plan; steel and project planning; review minutes; change register; hours; radiography, class inspection and straightening records	Project engineer (item-by-item review); production project manager; quality department
Case B	DOII close-out note; open-item board; hours; project planning; change register; class inspection; shell as built; section and project evaluations	Project engineer (population and classification); mechanical engineering organisation; planner; production project manager
Case C	Engineering log of the section; DO-2 scope guideline for the superstructure; project planning; hours; change register; project evaluation; commercial change list	Project engineer of the superstructure and an engineer from work preparation (classification); detail engineer of the superstructure; exterior structure engineering; quality department (general mechanisms)

Table 6-D.1: Sources per case.