

A SCIENTIFIC AND INTELLECTUAL HISTORY

History of Diagnostics, Observability, and Debugging in Engineering Sciences

From control and systems integration to security and resilience

Seven parts and 35 chapters

Including Security Engineering and Resilience Engineering

Concept by Dmitry Vostokov, DumpAnalysis.org

Developed with AI assistance (GPT-6 Astra Max and Fable 5.1 Extra) and checked against the sources listed in this article.

Version 5 | October 2026

Historical synthesis with 83 references,
chronology, glossary, and evidence taxonomy

Abstract

An engineering system can maintain its output while consuming a shrinking control margin. A component can satisfy its specification while its interaction with another component creates a hazardous condition. A display can report an ordinary value while the process it represents has changed. These situations require us to connect physical behavior, control, recorded evidence, and the decisions made from that evidence. This history follows the development of those connections across engineering sciences.

Our organizing thread is the investigation of interacting engineered systems. We follow experimental models and measurement into feedback, state estimation, process diagnosis, systems integration, assurance, and human interpretation. Examples from civil, mechanical, electrical, chemical, transport, environmental, and biomedical engineering show how the boundary of an investigation changes its possible explanations. The companion Technical and Hardware histories provide fuller accounts of physical examination and electronic diagnostic mechanisms; here we examine how those capabilities participate in the behavior of a larger system.

Security Engineering makes deliberate interference with behavior and evidence explicit. Resilience Engineering examines how performance is sustained, restored, and adapted as demands and resources change. Their histories connect diagnosis with the trustworthiness of observations, the authority to act, and the capacity to continue essential work.

Seven parts and 35 chapters connect documented developments with our pattern-oriented approach to diagnostics. Selected accident cases are examined through specific questions about feedback, organizational decisions, and recovery. A worked pumping-station example combines physical diagnosis, software debugging, security analysis, and adaptive performance. A selected chronology, glossary, and evidence taxonomy support comparison across these practices.

How to read this history

This history follows how engineers acquired the means to investigate behavior across the boundaries of components, control systems, software, and organizations. We consider developments that connected observation with models, decisions, intervention, and sustained performance. The parts are thematic and approximately historical. Material examination and instrumented testing provide necessary foundations; control, integration, security, and resilience determine the particular direction of this volume. The selection does not cover every branch of engineering or every diagnostic technique.

Diagnostics concerns the investigation of a condition and the discrimination of possible explanations. In engineering, the condition may be a functional failure, declining efficiency, unexpected deformation, an intrusion, or a loss of capacity to cope with changing demands. We use debugging literally for software and more broadly for locating and correcting defects in an engineered arrangement. The broader use is our analytical interpretation. We do not suggest that an eighteenth-century millwright or a nineteenth-century bridge investigator used the present software vocabulary.

Observability requires particular care. In control theory, it is a property defined relative to a model, available inputs and outputs, and an observation interval. The visibility of a pressure trace or the existence of a dashboard does not establish that property. Elsewhere in the book, we use evidence availability, measurement coverage, diagnostic distinguishability, and operational visibility when these are the more precise expressions. Security additionally requires us to ask whether the observations can be trusted.

Engineering sciences overlap with natural sciences and with practical engineering. A materials experiment may investigate a general mechanism, qualify a manufacturing process, or determine why one component failed. The same instrument can participate in all three activities. Our emphasis is on designed systems and the decisions made about them. Biomedical engineering concerns the behavior of a device here; clinical diagnosis has its own history. Environmental engineering concerns the performance and consequences of interventions as well as the surrounding natural processes.

The surviving literature is uneven. Formal publications give greater visibility to some institutions and countries than to workshop knowledge, maintenance work, and local adaptations. The selected early sources are mainly European; later chapters include Japanese manufacturing and international research and standards. This is a limitation of the selection, not a claim that diagnostic engineering originated in one region or followed a single route. Publication dates establish documented milestones, not unrestricted priority.

Square brackets identify references. We distinguish original publication, later edition, and online reproduction. Cases introduced as hypothetical are constructions for explanation. Our cross-domain comparisons and proposed applications of analysis patterns are identified as interpretations. Standards are discussed as historical and methodological documents, with their editions stated. Their presence does not make this book a compliance manual.

Relationship to the companion histories

Our History of Technical Diagnostics, Observability, and Debugging examines physical condition and integrity through failure physics, measurement, nondestructive examination, machinery diagnostics, and structural monitoring. Our History of Hardware Diagnostics, Observability, and Debugging follows electronic and computing devices through checking circuits, semiconductor test, firmware, diagnostic interfaces, and device or fleet evidence. The versions used here are Technical Version 7 and Hardware Version 14, both dated September 2026. [82, 83]

History	Principal object of investigation	Emphasis in the series
Technical	Physical structures, systems, and components.	Material condition, damage mechanisms, examination methods, machinery, and structural integrity.
Hardware	Electronic and computing devices and their supporting platforms.	Testability, error detection, diagnostic access, firmware records, and device or fleet evidence.
Engineering Sciences	Interacting physical processes, control systems, software, people, and organizations.	Feedback, integration, assurance, trustworthy evidence, security, and sustained performance.

Table 1. The emphasis of the three histories. These are editorial distinctions within overlapping fields.

The same asset can require all three perspectives. A pump may need physical examination; its controller may need hardware diagnosis; the service may require investigation of control logic, operating authority, shared dependencies, and recovery capacity. The engineering question connects these findings to the behavior of the installed system.

Accordingly, the shared foundations are treated briefly, with chapter references to the companion histories where a fuller account is available. We retain primary sources for the developments discussed here. Repeated accident cases serve a stated question about feedback, constraints, decisions, or adaptation. Part V develops diagnosis under adversarial conditions, and Part VI develops the history of adaptive performance. Together with the integration chapters and the worked example, these parts establish the particular scope of this volume.

Neighboring practices and their questions

Practice	Primary question	What the answer does not establish
Inspection and monitoring	What condition or change is visible?	Its cause, significance, or completeness.
Diagnosis	Which explanation is consistent with the evidence, and what distinguishes alternatives?	Future behavior under unexamined conditions.
Prognosis	How may condition and remaining capability develop?	A guaranteed time of failure.
Debugging and corrective engineering	Which change removes a demonstrated defect?	That every failure mechanism has been removed.
Safety and Security Engineering	Which losses must be prevented, and how can protection fail?	That an absence of incidents proves adequate protection.
Resilience Engineering	How is essential performance sustained, restored, and adapted?	Unlimited capacity or immunity to disruption.

Table 2. Related practices. The distinctions organize this history; actual investigations often combine them.

Recurring developments

The chapters follow several changes in diagnostic capability. Inspection and experiment made selected conditions comparable. Instruments converted otherwise inaccessible processes into indications and retained records. Calibration connected local measurements to references and uncertainty. Mathematical models related observed behavior to internal state and possible mechanisms. Distributed acquisition connected records from different places, while software made the acquisition and interpretation procedures programmable.

These changes also altered what could go wrong in an investigation. A recording device could distort the event. A model could exclude the relevant mechanism. A shared data source could make dependent observations look independent. Software could preserve an obsolete value while displaying it as current. Diagnostic progress therefore includes the examination of the observation and analysis arrangements themselves.

The organizing relationship is between what a system does, what can be known about that behavior, and what can be changed in response. Control makes this relationship explicit in feedback. Systems engineering carries it across requirements and interfaces. Security examines how an adversary can manipulate it. Resilience examines whether the resources and coordination needed for action remain available. These questions recur throughout the domain chapters and provide the connection between the historical parts.

Contents

Seven parts and 35 chapters

SECTION	PAGE
Abstract	2
How to read this history	2
PART I. FROM ENGINEERING MODELS TO RECORDED BEHAVIOR AND FEEDBACK	8
1. Vocabulary and historical scope	8
2. Material behavior, scale, and the limits of engineering models	9
3. Recorded cycles and access to internal behavior	9
4. Measurement chains and the comparability of evidence	10
5. Feedback, state estimation, and observability	10
PART II. FROM LOCAL VARIATION AND FAILURE TO SYSTEM DIAGNOSIS	12
6. Process variation and corrective feedback	12
7. Reliability, maintenance, and system dependencies	12
8. Model-based diagnosis and analytical redundancy	13
9. Material failure within an engineered system	14
10. Civil structures, ground conditions, and the observational method	14
PART III. CONTROL, INTERACTION, AND EVIDENCE ACROSS ENGINEERING DOMAINS	16
11. Machinery signatures under changing control and load	16
12. Electrical power systems and synchronized measurements	16
13. Chemical processes, hazards, and alarm interpretation	17
14. Aerospace and transport: interaction, feedback, and recovery	18
15. Environmental systems and biomedical devices	19
PART IV. COMPUTATION, SYSTEMS INTEGRATION, AND HUMAN INTERPRETATION	20
16. Software, embedded control, and distributed evidence	20
17. Numerical simulation, verification, and validation	21
18. Health estimates, digital twins, and the changing system	22
19. Systems engineering, testability, and assurance	22
20. Human factors, displays, and the interpretation of alarms	23

SECTION	PAGE
PART V. SECURITY ENGINEERING: DIAGNOSIS IN THE PRESENCE OF AN ADVERSARY	25
21. Physical protection, computer security, and system design	25
22. Intrusion detection, forensics, and the trustworthiness of traces	26
23. Operational technology and cyber-physical security	27
24. Threat modeling, security testing, and assurance evidence	28
25. Cyber resilience and the preservation of diagnostic capability	29
PART VI. RESILIENCE ENGINEERING: SUSTAINED PERFORMANCE AND ADAPTATION	30
26. From component failure to system accidents	30
27. Organizational conditions, high reliability, and everyday work	31
28. The emergence of Resilience Engineering	32
29. Infrastructure, recovery, and limits of adaptive capacity	33
30. Operational learning, site reliability, and deliberate perturbation	34
PART VII. INTEGRATED INVESTIGATION, CORRECTION, AND CONTINUING DEVELOPMENT	35
31. Accident investigations: feedback, authority, and adaptive capacity	35
32. Connections with pattern-oriented software diagnostics	36
33. A worked example: diagnosing a pumping-station disturbance	37
34. Automated diagnosis and the assurance of investigative tools	39
35. Continuing development of engineering diagnostics	40
Conclusion	42
Appendix A. Selected chronology	43
Appendix B. Glossary	46
Appendix C. Evidence taxonomy	49
References	51

PART I

From engineering models to recorded behavior and feedback

The opening chapters establish the relationships needed for an investigation of a system: between a model and its object, an observation and its acquisition path, and an output and the actions that regulate it. The shared instrument and material histories are foundations for this progression into control and observability.

1. Vocabulary and historical scope

We can begin with a simple engineering situation. A machine no longer performs its intended function. We observe a symptom, collect an artifact, and construct an explanation. However, the machine may still satisfy its local specification while the larger system fails. A pump can rotate at the commanded speed and deliver insufficient water because its inlet conditions have changed. The relevant object of diagnosis depends on the level at which we describe the problem.

This distinction also applies to defects, faults, errors, and failures. Their meanings differ across engineering traditions. In the dependability taxonomy developed by Avizienis and colleagues, a fault can cause an erroneous system state, and an error can lead to a failure at a service boundary. The taxonomy also connects reliability, availability, safety, and security. It helps us avoid treating all undesirable observations as interchangeable. Its definitions belong to a particular conceptual framework, rather than to every historical use of the words. [1]

An engineering science contributes models of structure and behavior. A diagnostic investigation uses those models together with evidence from a particular object or class of objects. The direction of reasoning often changes. Design reasoning asks what an arrangement should do. Diagnostic reasoning asks which arrangements, conditions, and histories could have produced the observed result. Neither direction is secure without assumptions about boundaries, loads, materials, software versions, and operating conditions.

We also need to distinguish the object from its representation. A drawing describes an intended arrangement. A maintenance record describes an intervention. A measured waveform records the response of a measurement chain. A simulation gives the behavior of an implemented model. These artifacts can disagree without any one of them being meaningless. Their disagreement may reveal a modified installation, an omitted action, an instrument defect, or a model outside its useful range.

The engineering boundary in this volume includes the relationships through which a local result affects a service. We ask how components are controlled, how evidence crosses interfaces, who or what can act on it, and which resources support that action. Technical and hardware investigations supply findings at

particular locations within this arrangement. Their combination requires an explanation of the connections as well as the parts.

In our pattern-oriented software diagnostics, the analysis of execution artifacts begins with signs of structure and behavior. We can use the same question in this history: what is actually present in the artifact, and what have we added through interpretation? The comparison is methodological. It does not make a fracture surface a memory dump or make an operator a software component. It helps us preserve the difference between recorded evidence, recognized pattern, and proposed cause. [2]

2. Material behavior, scale, and the limits of engineering models

Engineering judgment connected materials with intended function before engineering sciences acquired their later institutional forms. Vitruvius discussed construction, water, and machines within an organized practical tradition. Galileo's *Two New Sciences*, published in 1638, made strength and scale subjects of mathematical investigation. For our purposes, the important development is the examination of whether a relationship established for one arrangement remains valid in another. A small model and a full-sized structure can have different relationships between weight and resistance. [3, 4]

Smeaton's 1759 experiments on water and wind power compared machine performance under stated conditions. Such work made assumptions about a machine open to measurement. It also made the experimental arrangement part of the explanation: apparatus, load, and operating conditions affect which comparison can be transferred into design. [5]

Let us consider a hypothetical beam test. Unexpected deflection may arise from the material, support movement, loading, or an incorrect boundary condition in the model. The investigation must distinguish the specimen from the apparatus and the model used to interpret it. This is an early form of a recurring systems question: which part of the arrangement produced the observed behavior?

Technical Chapters 2–3 provide the fuller account of proof testing and strength-of-materials reasoning. Here those foundations lead into model adequacy, interface assumptions, and the design of observations that distinguish competing explanations. [82]

3. Recorded cycles and access to internal behavior

The steam-engine indicator related cylinder pressure to motion through a cycle. The Science Museum's account distinguishes Watt's instrument from the recording additions attributed to John Southern. A pencil and moving surface made a relationship within the operating cycle available for later comparison. This was a change in diagnostic capability: engineers could inspect a retained account of internal activity after the activity had ended. [6]

The diagram belongs to the engine together with the indicator and its coupling. Friction, calibration, response, and recording geometry affect the curve. A changed shape can therefore suggest a valve-

setting problem or a change in the observation arrangement. Load, speed, settings, and instrument identity help separate these possibilities.

We use this development as a foundation for investigating recorded behavior. A machine can operate without retaining enough information to explain its operation later. Selecting what to record is consequently a design decision. The later history of control adds a further question: does the record show the process output, the regulating action, or both? Compensation can preserve one while the other changes.

Technical Chapter 2 and Hardware Chapter 2 develop the indicator's place in their respective histories. The question carried forward here is how a representation of internal behavior supports comparison, intervention, and reconstruction within a larger engineered arrangement. [82, 83]

4. Measurement chains and the comparability of evidence

A distributed engineering system combines measurements made at different locations and transformed through different instruments and programs. Before using their agreement or disagreement diagnostically, we need to establish what was measured and how the results can be compared. The International Vocabulary of Metrology defines metrological traceability through a documented calibration chain, with uncertainty contributions. The Guide to the Expression of Uncertainty in Measurement supplies a framework for evaluating those contributions. [7, 8]

A complete measurement chain can include excitation, wiring, conversion, filtering, time stamping, transport, and display. Calibration of one instrument does not establish correct operation of every later transformation. For a hypothetical pressure display, the transducer can be accurate while a conversion uses the wrong units or the display retains an obsolete value.

Two agreeing displays may also depend on the same source. Their agreement then establishes less than agreement between independently acquired measurements. The investigator needs the acquisition architecture as well as the displayed numbers. This relationship becomes important in synchronized power measurements, shared control and recording software, and the examination of compromised evidence.

Technical Chapter 4 and Hardware Chapter 4 give the broader histories of measurement, calibration, and uncertainty. Here the emphasis is on their integration: preserving units, source identity, transformation history, and the uncertainty relevant to a decision as evidence moves across system boundaries. [82, 83]

5. Feedback, state estimation, and observability

Feedback changes the relationship between a disturbance and what an observer sees. A regulator may compensate for a developing problem and preserve the output. The condition becomes visible first in

increased control effort, reduced margin, or altered dynamics. This is one reason why a stable output is insufficient evidence that the internal system is unchanged. To diagnose a controlled process, we may need its commands and compensating actions as well as its final result.

Maxwell's 1868 paper *On Governors* examined the dynamics of regulation. He distinguished disturbances that diminish from those that grow, including oscillatory behavior. An intended regulating action could itself contribute to instability. For this history, the paper provides an early mathematical treatment of a diagnostic question: why does a machine with a governor behave in an oscillating or unstable manner, and how do its dynamics constrain a remedy? [9]

Nyquist's 1932 *Regeneration Theory* addressed the stability of feedback arrangements in electrical amplification. The problem arose in communications engineering, but its treatment helped make feedback stability a question that could be investigated through relationships across frequency. A measured oscillation could be related to a loop's dynamics, rather than interpreted only as a local component fault. Mechanical regulation and electrical amplification thus contributed different settings for the mathematical study of feedback. [10]

Twentieth-century state-space control theory made the reconstruction problem more explicit. Kalman's 1960 work on the general theory of control systems developed relationships involving internal state, controllability, and observability. Luenberger's 1964 observer work showed how a dynamical system could be used to estimate the state of another system from available information. These developments supplied a formal language for what can be inferred from outputs under a specified model. [11, 12]

We must distinguish state observability from fault isolation. Two different faults can produce the same measured response even when the nominal state model is observable. Conversely, a particular fault may be detectable without reconstructing every internal state. Parameter identification introduces another question: can different parameter values be distinguished by the selected experiment? A larger collection of sensors does not settle these questions unless their placement and information are relevant.

A hypothetical vessel with one measured liquid level illustrates the issue. A level decrease might indicate reduced inflow, increased outflow, or leakage. The level record alone may not separate them. Additional flow measurements or a controlled operating change can make the alternatives produce different observations. The improvement is not simply more data. It is a change in the relationship between possible explanations and available evidence. That relationship becomes central to model-based diagnosis in Part II.

PART II

From local variation and failure to system diagnosis

A departure can belong to a component, a process, an observation path, or their interaction. We follow methods that made these alternatives explicit and connected them with maintenance and intervention. The central question is what additional relationship must be examined before a local finding explains the behavior of the system.

6. Process variation and corrective feedback

Manufacturing supplies repeated observations of a process and its products. Keeping their sequence can reveal a change that disappears in an overall average. NIST's historical account identifies Shewhart's 1924 memorandum and 1931 book as milestones in statistical quality control. Page's 1954 continuous inspection schemes developed cumulative monitoring of departures over time. These developments connected production history with rules for directing investigative attention. [13, 14]

A monitoring rule and a product specification answer different questions. A stable process can repeatedly produce an unsuitable result, while an acceptable item can occur during a changing process. Diagnosis requires the relationship among the intended product, the process state, and the observations used to judge them.

Toyota's account of *jidoka* connects recognition of an abnormality with stopping work and addressing the problem. Here detection acquires an operational consequence through the production arrangement. This is a distinct tradition with an explicit connection between observation and corrective action. [15]

Consider a hypothetical machining process whose dimensions drift. Records partitioned by tool, fixture, material batch, and measurement method can help discriminate causes. The system question also concerns the response: who receives the signal, who can interrupt production, and how is the effect of correction verified? An accurate detector can have little practical effect when that connection is missing.

Technical Chapter 29 discusses manufacturing and process evidence in greater detail. The emphasis here is the development of a feedback relationship among variation, investigation, authority, and correction. [82]

7. Reliability, maintenance, and system dependencies

Reliability methods connected observations of individual failures with populations, functions, and consequences. Weibull's 1951 paper supplied a distribution later used in lifetime analysis. The NRC's 1981 Fault Tree Handbook organized reasoning from an unwanted event through combinations of contributing conditions. NASA's reliability-centered maintenance guide describes the development of

task selection, including the airline work associated with Nowlan and Heap's 1978 report. These methods examine different relationships between failure behavior and intervention. [16, 17, 18]

The system boundary matters to each relationship. Two available pumps may share a power supply, inlet, controller, or maintenance procedure. Counting the pumps does not establish that the required service has two independent means of support. A maintenance task can improve one component while temporarily removing a protective function elsewhere.

For a hypothetical standby pump, ordinary operating records may contain little evidence of readiness because the pump was never demanded. A test must exercise the needed function under relevant conditions. A successful unloaded start leaves delivery at the required pressure unresolved. The test also needs to preserve the service while one resource is unavailable.

Jardine, Lin, and Banjevic's 2006 review distinguishes condition-data acquisition, diagnosis, and prognosis. Technical Chapter 9 and Hardware Chapter 8 provide fuller accounts of reliability practice. Here we carry forward the dependence of a service on combinations of resources and the distinction between a component's condition and the capability available to the whole system. These questions lead into resilience later in the book. [19, 82, 83]

8. Model-based diagnosis and analytical redundancy

Physical redundancy supplies more than one sensor or component. Analytical redundancy supplies relationships that the observations should satisfy. A measured flow, an estimated flow, and a conservation balance may be compared without installing identical sensors at every location. Their disagreement produces a residual: a discrepancy relative to a model. A residual is therefore constructed evidence. It depends on both the measurements and the assumptions embedded in the prediction.

Willsky's 1976 survey organized methods for detecting abrupt changes in dynamical systems, including failures. Isermann's 1984 survey examined process fault detection through modeling and estimation. These publications mark a mature discussion of how model structure, estimation, and decision rules can support diagnosis. The methods seek sensitivity to relevant faults while limiting sensitivity to noise, disturbances, and normal operating variation. [20, 21]

Another model-based tradition developed through artificial intelligence. Reiter's 1987 theory formulated diagnosis in terms of consistency between a system description and observations. De Kleer and Williams examined the diagnosis of multiple faults and the incremental use of evidence. These approaches make explicit that several candidate diagnoses may remain compatible with what has been observed. They also show why the next measurement can be chosen to discriminate among candidates. [22, 23]

These traditions share a dependence on the model, but their mathematical objects and inference procedures differ. A residual from a dynamical observer is not the same object as a logical conflict set. We should avoid collapsing them into a single method merely because both are described as model-

based. Their historical connection is the deliberate use of expected relationships to interpret departures from expectation.

Consider a hypothetical heat exchanger. An energy imbalance might indicate a sensor error, an unmodeled heat loss, an incorrect flow estimate, or a process fault. If the diagnosis assumes every sensor is correct, the model may force the explanation into the process. If it assumes the process is always nominal, it may force the explanation into a sensor. A useful investigation keeps these alternatives available until evidence separates them. It also records when no unique diagnosis is possible, because that result identifies a limitation of the present diagnostic arrangement.

9. Material failure within an engineered system

A fracture surface, corrosion deposit, or deformation can constrain the history of an engineered system. Griffith's 1921 rupture paper connected flaws with an energy argument. Paris and Erdogan's 1963 work examined crack-propagation laws and the evidence needed to discriminate them. Such models make material findings relevant to loading and service history. They do not make the final damaged surface a complete record of the event. [24, 25]

Nondestructive examination supplies other views of condition. Visual inspection depends on access and the examination arrangement. Firestone's 1946 supersonic-reflectoscope paper described the use of reflected sound pulses and their timing to locate internal boundaries or flaws. The indication must still be related to geometry, material properties, coverage, and acceptance criteria. [26, 27]

The detailed methods are treated in Technical Chapters 6–8 and 12–19, and the electronic-device laboratory in Hardware Chapter 19. Here the question is how a material finding enters a system explanation. A hypothetical cracked pipe may require examination of pressure history, support movement, control settings, and protective action as well as the crack itself. Establishing the physical mechanism helps identify which of those relationships requires correction. [82, 83]

A repair consequently creates several claims to check. The replacement may restore local integrity; altered controls may reduce the loading that caused damage; revised sensing may make recurrence visible. Evidence for one claim does not establish the others. Material diagnosis becomes part of integrated engineering when these claims and their connections remain explicit.

10. Civil structures, ground conditions, and the observational method

Civil and geotechnical systems are difficult to treat as repeatable manufactured specimens. The ground is heterogeneous, boundary conditions evolve, and the completed structure interacts with an environment that cannot be fully controlled. Construction also changes the object while it is being observed. A measurement taken during excavation belongs to a changing system, not merely to a fixed design awaiting acceptance.

Peck's 1969 discussion of the observational method gave a disciplined form to the use of observations during geotechnical work. Predictions, observations, and planned responses to unfavorable developments must be connected. The method is not permission to proceed without adequate preparation and decide what to do only after something goes wrong. Its diagnostic significance lies in designing an investigation and possible interventions together. [28]

Structural health monitoring extends the comparison into service. Farrar and Worden's 2007 introduction describes damage identification for aerospace, civil, and mechanical infrastructure and organizes the problem through statistical pattern recognition. It distinguishes the operational evaluation, data acquisition, feature extraction, and statistical modeling needed to make a monitoring system useful. A change in a measured feature must be interpreted against environmental and operational variability. [29]

A hypothetical bridge can change its dynamic response because of temperature, traffic, support conditions, or damage. A detector trained on one season may identify another season as anomalous. Adding more sensors does not solve the problem if the reference behavior remains inappropriate. We need comparisons that preserve loading and environmental context, together with independent inspection where the consequences justify it.

The same reasoning applies to embankments, tunnels, dams, and foundations. A displacement or seepage trend becomes actionable through its relation to a mechanism, a predicted range, and an available response. The historical contribution of monitoring is therefore wider than collecting a time series. It can make construction and operation responsive to evidence. Its limits also need to be designed: how quickly a significant change can be detected, how uncertain its interpretation remains, and whether intervention can still occur in time.

PART III

Control, interaction, and evidence across engineering domains

The domain examples show how operating conditions, feedback, and dependencies change the meaning of a diagnostic sign. Machinery compensation, synchronized power measurements, process hazards, transport recovery, and variable environmental or biological conditions provide different settings for the same need to investigate the connections around an observed event.

11. Machinery signatures under changing control and load

Vibration, temperature, lubricant, electrical demand, and delivered performance provide different views of machinery behavior. Randall's treatment connects vibration analysis with signal processing and machine mechanisms; NASA's maintenance guide places condition-monitoring methods within the selection of appropriate maintenance tasks. Technical Chapters 11 and 20–21, and Hardware Chapter 24, develop these methods more fully. [30, 18, 82, 83]

Here we examine the relationship between condition signatures and operating control. A vibration change may accompany local damage, altered excitation, or a different speed and load. Comparisons need the operating regime and the actions that established it. Otherwise a controller's adjustment can be mistaken for deterioration, or a comparison at unequal loads can conceal deterioration.

Thermal and fluid processes make compensation especially visible. Deposits in a hypothetical heat exchanger reduce heat transfer while a controller increases cooling or changes flow to preserve the output. A record containing only the regulated temperature can omit the developing loss of margin. Recording both the output and the compensating action provides a different diagnostic view.

The observation also needs a connection to an available response. If an adjustment consumes the remaining control range, an apparently satisfactory output may coexist with limited capacity to handle another disturbance. The engineering investigation therefore relates the condition signature to mechanism, control effort, operating limits, and time for intervention. This connection prepares the later distinction between ordinary compensation and the wider adaptive capacity of an operating system.

12. Electrical power systems and synchronized measurements

Electrical networks make local and system behavior inseparable. A protection device responds to measurements at a particular location, while a disturbance can propagate across a much larger network. An event record may show the operation of a relay without explaining the condition that initiated it. A trip can be a successful protective response, an unnecessary operation, or one step in a wider disturbance. The same recorded action can therefore have different diagnostic meanings.

Phadke, Thorp, and Adamiak's 1983 paper developed measurement techniques for tracking voltage phasors, frequency, and rate of change of frequency. Phadke and Thorp's later book describes the development and applications of synchronized phasor measurements. These capabilities support comparisons across locations with a common time reference. They extend what can be reconstructed about the changing state of a power system. [31, 32]

Synchronization does not eliminate interpretation. The recorded time can refer to acquisition, estimation, transmission, or receipt. Communication delay and missing records can complicate reconstruction. A measurement algorithm also has a response to rapid changes. We must know what each value represents before arranging records into one apparent sequence. An ordered spreadsheet is not sufficient evidence of causal order.

Consider a hypothetical disturbance in which several relays operate. The first recorded trip may not identify the initiating fault if clocks differ or the earlier disturbance was outside the recording coverage. We can compare waveforms, breaker positions, protection settings, network topology, and equipment condition. Some evidence describes the electrical event; other evidence describes the protective system's interpretation of it. These layers should remain distinguishable.

Modern power equipment also contains software, networks, and remote control. A mistaken setting, a firmware defect, a damaged transducer, and unauthorized modification can produce overlapping symptoms. This makes electrical diagnosis one of the places where reliability, security, and resilience meet. The immediate question may be whether to isolate equipment. The later investigation asks why the condition was possible and whether the remaining network can support essential service.

The historical significance of wider measurement is therefore relational. Engineers gained ways to compare distant activity and reconstruct disturbances that could not be understood from one local record. The new visibility also created dependence on common timing, configuration, communications, and interpretation. Those supporting systems became diagnostic objects in their own right.

13. Chemical processes, hazards, and alarm interpretation

Chemical engineering investigates transformations in systems with interacting inventories, flows, temperatures, pressures, and reactions. A disturbance can be delayed, amplified, or temporarily compensated. The absence of an immediate output change does not establish that a process remains within its intended conditions. A diagnostic investigation may need to reconstruct what accumulated inside the system before an alarm appeared.

Hazard and operability studies examine deviations from an intended process. In his account of ICI's contributions to process safety, Trevor Kletz traces the development of this approach to work in the 1960s. The method gives a structured role to questions about absent, excessive, reversed, or otherwise altered process conditions. It is prospective: possible deviations are examined before they become incidents. The resulting records can also inform later diagnosis. [33]

An alarm is a decision made from selected observations. Its threshold, delay, suppression, and display determine what an operator receives. A large number of alarms can contain less usable information than a smaller set that distinguishes conditions and consequences. Conversely, removing repeated

alarms without examining their cause can erase evidence of a deteriorating process. Alarm reduction and improved diagnosis are not automatically the same achievement.

The CSB investigation of the 2005 BP Texas City refinery disaster connected the physical sequence with instrumentation, procedures, maintenance, and organizational conditions. The report is valuable because it does not reduce the event to one abnormal reading or one operator action. It shows why investigators need evidence from different levels of an engineered operation, and why favorable injury statistics cannot substitute for evidence about major process hazards. [34]

Let us consider a hypothetical reactor whose temperature remains near its target while cooling demand rises. The output looks stable, but the control effort suggests a changing balance. Possible explanations include altered feed, fouling, a measurement problem, or a developing reaction condition. The diagnostic question is which additional evidence separates them. A model that assumes the original feed composition may make the process appear faulty when the boundary conditions have changed.

Chemical-process history thus connects diagnosis to prevention and intervention. A hazard study identifies conditions worth distinguishing; instrumentation makes some of them visible; operators and automated protection respond; investigation examines the resulting sequence. Each connection can fail. A well-written hazard record has limited value if the installed measurements cannot reveal the deviation on which the required response depends.

14. Aerospace and transport: interaction, feedback, and recovery

Transport systems can require protective action before a complete explanation is available. We revisit selected cases to examine that relationship between evidence, authority, and response. Technical Chapters 24 and 34 and Hardware Chapters 23 and 34 provide complementary accounts of physical condition, vehicle diagnostics, and failure reconstruction. [82, 83]

Apollo 13 required operational responses to an unfolding loss of capability, followed by investigation of the oxygen-tank failure and its prehistory. NASA's review-board account connects design, handling, and ground testing with the later event. Our additional question concerns the separation of timescales: what evidence supported continued action during the mission, and what further evidence was needed for the eventual causal explanation? Successful recovery and a supported diagnosis were distinct achievements. [35]

The Challenger investigation examined the solid-rocket motor joint and the launch decision process. Columbia's investigation treated organizational causes alongside physical damage. Here we ask how technical evidence entered decisions: which observations reached those with authority, how uncertainty was represented, and which assumptions limited further investigation. These questions retain the physical mechanisms while examining the arrangements through which knowledge could affect action. [36, 37]

The Qantas Flight 32 investigation connected a manufacturing defect in an oil-feed pipe with an uncontained engine failure, effects across the aircraft, and crew response. The question developed here is how a local origin became a problem of interacting systems and remaining capability. Origin, propagation, and operational recovery require related but separate accounts. [38]

These cases inform later chapters on human interpretation and resilience. The same incident can reveal a defect, a limitation in feedback, and an effective adaptation. A useful investigation preserves all three when supported by evidence, so that corrective work addresses both the initiating condition and the system's ability to recognize and manage its consequences.

15. Environmental systems and biomedical devices

Environmental and biomedical engineering share an important diagnostic difficulty: the engineered object operates within a larger system that is variable and only partly controlled. A treatment plant responds to changing inflow. A medical device interacts with a patient, a clinical workflow, and other equipment. The boundary selected for a model can therefore exclude conditions essential to explaining the observed result.

Water-distribution modeling provides one example. EPANET supports simulation of hydraulic and water-quality behavior in pipe networks. Its 2020 user manual documents a computational instrument for examining flows, pressures, and the movement or transformation of constituents. Such a model can help compare possible explanations for an observed condition, but it requires an appropriate network, demand assumptions, and boundary conditions. A plausible simulation is not independent confirmation of those inputs. [39]

For a hypothetical water-quality anomaly, a changed laboratory result might reflect a process condition, sampling location, contamination during collection, or an analytical problem. A time series from a plant sensor introduces another acquisition path. Agreement is useful only after the two measurements have been related to the same material and time interval. Environmental diagnosis often requires this combination of process records, samples, and spatial context.

Biomedical devices add direct consequences for the person receiving a service. Leveson and Turner's 1993 investigation of the Therac-25 accidents examined software, interface behavior, safety assumptions, and the handling of accident reports. The case demonstrates why a history of prior operation and testing cannot substitute for understanding hazardous interactions. It also shows how reports from users can contain evidence that the manufacturer's existing model does not accommodate. [40]

We should keep device diagnosis distinct from clinical diagnosis. A displayed physiological value may be clinically unusual, incorrectly measured, incorrectly transmitted, or incorrectly associated with a patient. Debugging the display path does not determine the patient's condition. Clinical interpretation does not by itself validate the device. The investigation needs evidence capable of separating these questions.

Across both domains, we examine how process models, instruments, and actual conditions interact. Sensing and computation extend the evidence while adding transformations whose assumptions must remain available to the person making a decision.

PART IV

Computation, systems integration, and human interpretation

Computation can control a process, record its behavior, and construct an explanation of it. We examine how these roles are integrated with requirements, models, assurance, and human decisions. The history now concerns the reliability of the relationships among those activities as well as the accuracy of any one result.

16. Software, embedded control, and distributed evidence

Software changed engineering diagnostics in two directions. It became part of the object under investigation, and it became part of the instrument used to investigate that object. A controller, protection relay, recorder, simulation, and maintenance database can each contain software defects. When the same software participates in control and observation, one defect may affect both the physical behavior and its recorded account.

The dependability framework of Avizienis and colleagues helps separate a fault from the resulting internal error and externally visible failure. In a larger engineered system, these boundaries can be nested. A sensor-processing failure can become an input fault for a controller, while the controller's protective response becomes an observed loss of service. The label attached to an event depends partly on the boundary from which we inspect it. [1]

Distributed systems make sequence reconstruction more difficult. Lamport's 1978 paper formalized a partial ordering of events through their possible causal relationships. Its relevance extends beyond sorting timestamps. Events on different components need not have a known total order, and an arbitrary ordering imposed by a display can suggest a sequence that the evidence does not establish. Logical relationships and physical time answer related but different questions. [41]

In a hypothetical automated plant, the historian may timestamp receipt, the controller may timestamp execution, and an engineering workstation may record a command before transmission. A combined log can therefore appear to show an effect before its cause. Correcting the clock display may help, but the investigation must also recover what each timestamp means. Lost messages, buffering, and retries can change the apparent narrative without changing the underlying physical sequence in the same way.

Memory dumps and software traces add views that ordinary process measurements may lack. They can preserve pending commands, state-machine positions, configuration, or the reason a task stopped progressing. Our trace-analysis pattern reference treats relationships such as Message Context, Time Delta, and Adjoint Thread of Activity as reusable ways of examining such artifacts. In engineering applications, these views complement physical measurements rather than replace them. [42]

The historical result is a larger diagnostic object. We investigate the machine, its control software, and the system that records both. A software fix can restore function while leaving a physical defect unchanged; a mechanical repair can conceal a control defect until another condition activates it. The evidence must support the level at which the claimed correction operates.

17. Numerical simulation, verification, and validation

Simulation made many engineering states available for examination before a corresponding physical system existed. It also made it possible to construct detailed explanations of an incident after the event. These capabilities are valuable, but their diagnostic force depends on the relationship between the implemented model and the actual system. A visually convincing result is not evidence that every relevant mechanism has been represented.

NASA's Standard for Models and Simulations organizes credibility through practices that include verification, validation, uncertainty, and the intended use of results. The 2024 revision provides a contemporary expression of a problem with a much longer history. Verification examines whether the model implementation behaves as intended; validation concerns its adequacy for representing the relevant reality and purpose. Passing one does not establish the other. [43]

A numerical model can be wrong in several different ways. Equations may omit a relevant process, inputs may describe another configuration, boundary conditions may be unsuitable, and the numerical method may introduce an unacceptable error. Software defects can coexist with all of these. Debugging only the code does not make the model appropriate, while adjusting physical parameters can sometimes hide an implementation error.

Consider a hypothetical finite-element reconstruction of a failed support. If the analyst tunes the model until it reproduces the observed deflection, several combinations of stiffness, load, and restraint may fit. Reproduction establishes compatibility under those assumptions. It does not uniquely identify which condition occurred. Independent measurements, construction records, or another response mode may be needed to distinguish the alternatives.

The distinction also applies to engineering experiments. A test used to calibrate a model should not then be presented as wholly independent validation of it. Reusing the same evidence for construction and confirmation can exaggerate confidence. We need to retain the role of each dataset, including the conditions under which the model was expected to apply.

The history of computational engineering therefore adds a new class of artifacts: input files, model versions, solver settings, meshes, scripts, and output transformations. Reconstructing a simulation requires these materials as well as the final plot. The computation itself has a history that can contain the defect. Preserving that history makes a result reviewable and makes later debugging possible when a seemingly physical anomaly proves to belong to the analysis process.

18. Health estimates, digital twins, and the changing system

A maintained digital representation can connect design assumptions with the operating history of an individual asset. Glaessgen and Stargel's 2012 digital-twin proposal combined simulation with vehicle health information, maintenance history, and historical data. Technical Chapter 33 and Hardware Chapter 32 provide fuller accounts of prognosis and digital twins. Here we examine what happens when the represented asset and the decisions made about it change. [44, 82, 83]

A repair, sensor replacement, control update, or temporary operating limit can alter the relationship between model and system. Updating model parameters may improve agreement while concealing an incorrect input or an unmodeled condition. The investigator needs the configuration, measurements, model version, and update history behind the current estimate.

Jardine and colleagues distinguish diagnosis from prognosis. A remaining-life estimate introduces assumptions about future use in addition to present condition. A forecast made before an emergency change in loading may no longer answer the operational question. The current decision must be related to the conditions under which the estimate was produced. [19]

For a hypothetical turbine, the maintained representation should connect measured condition with permissible operation and available alternatives. If a recommendation changes load, later observations describe the system after that intervention. Preserving this sequence supports examination of whether the model helped, whether its assumptions remained applicable, and what evidence is available when the model or its data path becomes unavailable.

19. Systems engineering, testability, and assurance

Large engineering projects must connect requirements, design decisions, implementation, integration, and operation. Diagnostic capabilities can be lost between these activities even when each local component is well designed. A requirement may specify a failure response without specifying how the relevant condition will be detected. An integration test may show that an alarm can be displayed without showing that the installed sensor can produce it under the actual fault.

Shared criteria provided one historical basis for such connections. ASME describes the Boiler and Pressure Vessel Code as conceived in 1911; its first edition was dated 1914 and published in 1915. Technical Chapter 5 develops that institutional history. Here it illustrates how design, fabrication, and inspection records could be related through explicit requirements. [45, 82]

NASA's Systems Engineering Handbook describes technical processes and management activities across the system life cycle. Its treatment of requirements, verification, validation, interfaces, and configuration provides a setting in which diagnostic requirements can be made explicit. A system is easier to investigate when the intended relationships among these artifacts survive implementation and change. [46]

Testability is wider than access to a test connector. We need to know whether a relevant condition can be stimulated, whether its consequences can be observed, and whether competing faults can be distinguished. A built-in test that reports a general failure can support protection while leaving maintenance with little isolation capability. Conversely, detailed service diagnostics may be too slow or intrusive for real-time protection. These are different design purposes.

Functional-safety standards made the relationship between safety functions and their life-cycle treatment explicit. IEC 61508-1:2010 addresses general requirements for electrical, electronic, and programmable electronic safety-related systems. It is cited here as a historical framework, not as a universal prescription for every domain. Diagnostic coverage, proof testing, and management of systematic faults need to be interpreted within the applicable system and standard. [47]

Assurance connects a claim with reasons and evidence. A test report is not an assurance argument merely because it contains many successful results. We must establish which claim the test supports and under which conditions. A test of an earlier software version may remain relevant, become partly relevant, or lose relevance after a change. Configuration history determines the answer.

From our diagnostic viewpoint, systems integration should preserve paths between intended function, possible deviation, observable evidence, and corrective action. This is an application of the book's synthesis. It makes diagnosis a design concern while preserving the different responsibilities of design, operation, and investigation. When an incident occurs, those paths help us determine whether the system departed from its intended behavior, whether the intention was inadequate, or whether the evidence cannot yet distinguish the two.

20. Human factors, displays, and the interpretation of alarms

Engineering evidence becomes operationally useful through interpretation. A display does not simply transfer the physical state of a plant into a person's mind. It selects, transforms, groups, and labels information. The operator must relate that representation to a changing situation, often while acting on it. A technically correct indication can still be insufficient for the decision required.

The Three Mile Island accident provides a documented example. The NRC's account describes a relief valve that remained open and an indication that did not directly reveal its actual position. The discrepancy between the physical condition and the operators' understanding was part of the event. The case shows why a command indication and a state indication should not be treated as the same evidence. [48]

Bainbridge's 1983 paper *Ironies of Automation* examined how automation can leave people with difficult responsibilities while reducing opportunities to maintain the skills and understanding needed for them. The remaining human task may become more demanding rather than disappear. Monitoring an automated system and taking over during an unusual event require a different account of work from simply removing routine manual actions. [49]

Consider a hypothetical control room in which all displays depend on one data concentrator. Several agreeing values may look like independent confirmation, but their agreement can be produced by a shared stale input. An operator who sees no change is not necessarily inattentive. The representation may conceal the freshness problem. A diagnostic design should make the state of the information path visible where it affects the decision.

Human accounts also become evidence after an incident. They can reveal what was visible, which alternatives were considered, and why an action appeared reasonable. Such accounts are not perfect recordings, but they should not be discarded in favor of logs that record a different aspect of the event. A controller log may show the issued command while omitting the workload, competing alarms, or missing information that shaped the decision.

The historical movement toward human factors and cognitive systems approaches expanded the unit of analysis. The question became how people, instruments, procedures, and organizations jointly produced performance. This prepares the ground for Resilience Engineering. It also changes debugging: correcting a label, alarm arrangement, or feedback path may be as consequential as correcting an algorithm, provided the change addresses the demonstrated problem and is examined in the conditions of use.

PART V

Security Engineering: diagnosis in the presence of an adversary

An adversary can interfere with a physical or digital process and with the evidence used to investigate it. Security Engineering makes that possibility explicit in protection models, collection arrangements, testing, and recovery. Its connection to the preceding chapters is the need to establish which observations and actions remain trustworthy.

21. Physical protection, computer security, and system design

Security Engineering adds an intentional actor to the diagnostic situation. Equipment can deteriorate, software can be defective, and people can make mistakes. An adversary may also choose a condition, sequence, or representation that makes protection fail. The security problem therefore includes what an opponent can learn about the system, what access can be obtained, and how a protective arrangement can be bypassed or manipulated.

Physical protection provides a concrete example. Garcia's 2001 treatment organizes protection around facility characterization, threat definition, target identification, detection, delay, response, and evaluation. A barrier has meaning in relation to detection and the time available for response. A sensor has meaning in relation to assessment and action. Security is a property of their arrangement with people and procedures, rather than the sum of independently impressive devices. [50]

Computer security developed its own architectural principles. Saltzer and Schroeder's 1975 paper examined mechanisms for protecting information and articulated principles including least privilege, complete mediation, and fail-safe defaults. Such principles constrain how a system should behave when access is requested or a mechanism fails. They also supply diagnostic questions: which check was expected, which authority was used, and where did the actual behavior depart from the protection model? [51]

Ross Anderson's Security Engineering connects technical mechanisms with their economic, organizational, and human setting. Its third edition, published in 2020, provides a broad view of building dependable systems under adversarial conditions. For our history, this is evidence of a mature discipline whose scope extends beyond cryptographic algorithms and isolated vulnerabilities. The surrounding incentives, interfaces, and operational arrangements are also engineering concerns. [52]

Consider a hypothetical access-controlled equipment room. An electronic log records that a credential was accepted. It does not by itself identify the person who entered, establish whether the door closed, or prove that no one followed another person through it. A camera, door-position sensor, and maintenance record provide different evidence. Their agreement must be examined rather than assumed, particularly if they share an administrator or power source.

The same distinction appears in software. A recorded authentication success establishes something about an authentication transaction. It does not establish that every subsequent action was authorized or that the authenticated endpoint was uncompromised. Security diagnostics must reconstruct the relevant chain of authority and behavior. The historical contribution of Security Engineering is to make the adversary and the protection assumptions explicit within the system model, including the assumptions on which the investigator depends.

22. Intrusion detection, forensics, and the trustworthiness of traces

An operational log becomes security evidence when it is used to investigate possible misuse, intrusion, or unauthorized change. This introduces questions that ordinary monitoring can leave implicit. Who could create the record? Who could alter or suppress it? Does the recorded identity refer to a person, a process, a credential, or a network location? An apparent sequence can be accurate as a record of messages while remaining misleading as an account of agency.

Denning's 1987 intrusion-detection model described the use of audit records and profiles to identify abnormal activity. It is an important milestone in treating recorded system behavior as a source of automated security detection. The underlying distinction remains necessary: an unusual pattern can justify investigation without proving an attack. Legitimate change, faults, and incomplete observation can also produce departures from a learned or specified profile. [53]

Digital forensics adds preservation and reconstruction. NIST's 2006 guide addresses data sources including files, operating systems, network traffic, and applications within incident response. The sources differ in volatility, scope, and meaning. A memory artifact may reveal state that was never written to a log. A network record may show a transaction that a compromised endpoint no longer reports. Neither source is automatically complete. [54]

MITRE ATT&CK organizes observed adversary behavior into a shared knowledge base. Its Design and Philosophy paper explains the motivation and structure of the approach. A behavioral vocabulary can improve communication between detection engineering and investigation. A mapped technique is nevertheless a description supported to some degree by evidence; it is not proof of an actor's identity, intent, or complete activity. Coverage of a catalog is also different from successful detection under actual conditions. [55]

Let us consider a hypothetical gap in an audit sequence. The gap may indicate deletion, transport loss, rotation, filtering, a stopped service, or a period without the recorded activity. Its meaning depends on the recording mechanism. An independent sequence number or remote copy can help distinguish explanations. The absence of a message becomes evidence only after we establish why the message should have existed and survived.

This history connects security diagnosis with the analysis of evidence production. The collection path must itself be observable enough to support confidence in the collected artifact. Hashes and protected storage can help establish that a retained artifact has not changed after a particular point. They do not by themselves establish that the original event description was true or complete. That distinction is essential when the system producing the evidence may already be compromised.

23. Operational technology and cyber-physical security

Industrial control brought computer security into direct contact with physical processes. A digital action can change a valve, motor, setpoint, or protective function. The consequences depend on the plant and its operating state. An action that is harmless in a test environment can be consequential during a particular process phase. Security analysis therefore needs engineering knowledge of the controlled process as well as knowledge of the network.

The Stuxnet disclosures in 2010 made this connection visible to a wider audience. The archived ICS-CERT advisory documented malware targeting Siemens control software and associated industrial environments. The historically important point for this book is the deliberate interaction between malicious software and control-system operation. We do not infer geopolitical attribution or a complete physical effect from the advisory alone. [56]

NIST's 2023 Guide to Operational Technology Security explicitly addresses performance, reliability, and safety requirements alongside protection. Its scope includes industrial control, building automation, transportation, physical access control, and environmental monitoring. This breadth matters because operational technology is defined by interaction with the physical environment, not by one industrial sector or communication protocol. [57]

A process historian may appear to provide the plant's authoritative account. However, its records can be affected by compromised endpoints, configuration changes, filtering, time manipulation, or communication failures. If the same source controls a process and reports its state, its report is not an independent observation. A physical measurement outside that path may be needed to test the claimed state.

Consider a hypothetical motor whose controller reports the expected speed while an independent instrument detects a discrepancy. The possibilities include a reporting fault, an incorrect sensor mapping, a control defect, or deliberate manipulation. The first diagnostic task is to preserve those alternatives and determine the independence of the measurements. Treating every discrepancy as an attack can be as misleading as treating every discrepancy as accidental.

Security intervention must also respect the physical situation. A restart, active test, or communication isolation changes the engineered system. The appropriate response depends on the process state and the available protective arrangements. This is not a reason to avoid investigation. It is a reason to integrate the investigation with engineering control of its consequences.

The historical development of cyber-physical security therefore enlarges both the threat model and the evidence model. We need to know what an adversary might change, what the physical process would then do, and which observations would remain trustworthy. The resulting diagnosis can cross electrical, mechanical, software, and organizational boundaries without allowing any one of them to stand for the whole event.

24. Threat modeling, security testing, and assurance evidence

Security diagnosis is not confined to incidents that have already happened. Threat modeling examines how a system might be attacked, and testing examines selected claims about its behavior. These activities can identify defects in architecture, implementation, configuration, and operation before they produce an observed loss. Their historical development parallels safety analysis in some respects, but deliberate strategy changes the assumptions about how conditions arise.

Schneier's 1999 presentation of attack trees provides one way to organize alternative routes to an adversary's objective. It makes a goal and its possible means explicit. A tree can help compare defenses, but its structure is a model built by the analyst. Omitted routes remain omitted, and a numerical comparison can give false precision if its probabilities or costs lack an adequate basis. [58]

Security testing supplies evidence about particular arrangements under particular conditions. A penetration test can demonstrate a reachable weakness. Its failure to find one does not prove that none exists. Fuzzing can reveal unexpected behavior under generated inputs, while static analysis can identify selected properties or suspicious constructions in code. The techniques have different coverage, and their findings require interpretation in the operational context.

NIST's 2022 Secure Software Development Framework connects vulnerability reduction with practices throughout development. It also includes responding to vulnerabilities and addressing their underlying causes. This moves corrective work beyond a single patch: repeated findings can become evidence of a problem in requirements, design, implementation practices, or the development environment. The framework supplies a process vocabulary, rather than a guarantee that following it eliminates every vulnerability. [59]

NIST's Engineering Trustworthy Secure Systems, revised in 2022, places security within systems engineering across the life cycle. It treats protection as a concern for the whole system and its contested environment. This supports the inclusion of security requirements in architecture and trade studies, together with verification and validation of the resulting behavior. [60]

For a hypothetical remote maintenance interface, we might test authentication, authorization, session termination, audit generation, and the restrictions on physical commands. Passing an authentication test leaves the other claims unresolved. If an engineer can issue a hazardous command through an authorized session, the problem may be a missing safety constraint rather than an authentication defect. A useful assurance argument connects each test to the specific claim it supports.

The historical lesson is the development of a more explicit relationship between security intentions and evidence. A control can exist without being effective in its installed context. A test can succeed while exercising an unrepresentative configuration. Diagnostic records should therefore retain the tested version, assumptions, stimulus, observations, and unresolved questions. This makes later investigation capable of explaining what an earlier assurance claim did and did not cover.

25. Cyber resilience and the preservation of diagnostic capability

Prevention cannot be treated as an assumption that removes the need for continued operation and recovery under compromise. Cyber resilience asks how a system can sustain essential objectives when some resources are unavailable, misleading, or controlled by an adversary. This introduces a different diagnostic question from the discovery of a vulnerability: which capabilities remain trustworthy enough to support a decision now?

NIST's 2021 *Developing Cyber-Resilient Systems* organizes cyber resiliency around anticipating, withstanding, recovering from, and adapting to adverse conditions involving cyber resources. It connects systems security engineering with resilience engineering. The document supports selecting and adapting techniques to mission, operational, and threat conditions; it does not define resilience as installing a single product or achieving an absolute state. [61]

The 2025 revision of NIST's incident-response guidance places detection, response, and recovery within broader cybersecurity risk management. This recognizes that an incident-handling capability depends on preparations and decisions made before the incident, together with learning afterward. A report that closes an incident administratively may leave important questions about recurrence, evidence, and recovery unresolved. [62]

We can apply these ideas to diagnostic capability itself. A system may retain essential service while losing the measurements needed to understand its condition. Conversely, it may retain excellent logs while losing the ability to act. Resilience requires a relationship between knowledge and action. Independent evidence paths, protected configuration records, and tested restoration procedures can support that relationship, but their usefulness depends on the assumed disruption.

Consider a hypothetical compromise of a plant-management server. Isolation may protect the controllers, but it can also remove trend displays and maintenance information. Operators may continue using local indications and an established degraded mode. The event should be examined at several levels: the initial compromise, the containment action, the loss of visibility, and the ability to sustain the required physical function. Success at one level does not settle the others.

Recovery also requires diagnosis. A backup may restore old vulnerable software, a corrupted configuration, or an attacker's persistence. Restoring bytes is different from re-establishing a trustworthy operating state. The recovery claim should identify what was restored, what was independently checked, and which uncertainty remains. A new baseline should preserve enough history to explain the transition.

Security Engineering and Resilience Engineering therefore meet through the continuity of warranted action. We need evidence that supports decisions despite disruption, along with ways to revise those decisions as new evidence appears. The connection is especially close in systems where the adversary can interfere with the very instruments used to detect the attack. Protection of evidence is then part of protection of the system's capacity to respond.

PART VI

Resilience Engineering: sustained performance and adaptation

A supported explanation does not by itself supply the capacity to act. Resilience Engineering examines the resources, coordination, and adjustments through which performance continues under changing demands. We follow its development through system safety, organizational research, everyday work, and the limits of adaptation.

26. From component failure to system accidents

Reliability analysis can identify how components fail, but a harmful system outcome does not always require a broken component. Correctly functioning elements can interact under inadequate constraints. An operation can satisfy local goals while placing the larger system in a hazardous condition. The development of system-safety approaches made these relationships a more explicit subject of investigation.

Perrow's Normal Accidents, first published in 1984, examined the implications of interactive complexity and tight coupling for high-risk technologies. The argument directs attention to combinations that are difficult to anticipate and situations that leave little time or flexibility for recovery. It is an argument about particular system characteristics, not a claim that every accident is unavoidable or that detailed technical investigation has no value. [63]

Leveson's 2004 presentation of STAMP proposed an accident model centered on inadequate enforcement of safety constraints. The approach examines control structures, process models, feedback, and interactions across levels. Her later Engineering a Safer World develops the systems perspective across analysis, design, and operation. These works broaden the diagnostic question from which component failed to how the control of hazardous behavior became inadequate. [64, 65]

Fault trees, failure analyses, and systems approaches need not be treated as mutually exclusive historical stages. A damaged component may be central to one event, while an unsafe interaction between functioning components may be central to another. The appropriate model depends on the questions and evidence. What matters is recognizing when a selected model excludes the mechanism that the investigation needs to consider.

Let us consider a hypothetical automated transfer. Each subsystem accepts a valid command, but their combined timing permits an unsafe material state. Testing each subsystem independently can miss the interaction. The investigation must reconstruct the coordination, the conditions assumed by each controller, and the feedback available to it. Correcting one component's code is effective only if it changes the relevant system relationship.

This development prepared one of the foundations for Resilience Engineering. Performance depends on coordinated activity under constraints, not only on the absence of local failures. Diagnostic work consequently needs to examine how the system remained within acceptable conditions before an event and how those conditions changed. The failure sequence is part of a longer history of successful and unsuccessful control.

27. Organizational conditions, high reliability, and everyday work

The engineer who encounters a failed component often encounters an organizational history as well. Inspection intervals were selected, maintenance was deferred or performed, information was reported or lost, and someone decided how to interpret an earlier anomaly. These activities shape the conditions in which physical and software mechanisms operate. They become part of the evidence when they help explain the event, rather than merely providing background commentary.

Reason's distinction between person and system approaches to error emphasizes the conditions that make errors consequential and the role of defenses. His account of latent conditions directs attention beyond the last visible action. This does not remove individual actions from the reconstruction. It changes the question from identifying a person to understanding how the arrangement allowed an action or condition to produce harm. [66]

Rasmussen's 1997 model of risk management examined interacting levels of a dynamic society and the pressures that shape behavior. A system can migrate toward a boundary as local decisions respond to cost, workload, and performance. No single decision need announce that the safety margin is being consumed. The diagnostic problem is then to reconstruct a changing pattern of constraints and adaptations, not merely an isolated deviation. [67]

High-reliability research studied organizations that maintained demanding operations despite substantial hazards. Rochlin, La Porte, and Roberts's 1987 account of aircraft-carrier operations examined organizational arrangements and the capacity to reorganize work. Its value here is the attention to how reliable performance is produced in practice. An absence of catastrophe becomes a subject for investigation rather than the end of the inquiry. [68]

For a hypothetical maintenance team, an informal cross-check may compensate for an ambiguous procedure. Routine records can show successful completion while omitting the extra coordination that made completion possible. Removing a person, changing a shift, or automating part of the task can remove that compensation. The resulting failure may appear sudden even though the formal procedure was inadequate throughout.

We therefore need evidence of work as performed, including handovers, interruptions, adaptations, and checks. Such evidence should be interpreted with participants' knowledge of the situation, rather than reduced to counts of deviations from a document. Some deviations introduce danger; others keep work within acceptable conditions. Determining which occurred is a diagnostic question. It cannot be answered from the word deviation alone.

28. The emergence of Resilience Engineering

Resilience has several histories. Holling's 1973 ecological paper distinguished resilience from stability and examined persistence under disturbance. The ecological object and its questions differ from those of an engineered organization. We can recognize a conceptual connection without claiming that all later engineering uses simply transferred one ecological definition. The meaning of resilience changes with what is expected to persist and what changes are permitted. [69]

The 2006 collection *Resilience Engineering: Concepts and Precepts* gave a prominent shared expression to a developing safety perspective associated with Hollnagel, Woods, Leveson, and other contributors. The focus extended from preventing failures toward understanding how systems adjust to conditions and sustain performance. Its publication is a clear milestone, but the work drew on earlier studies of human performance, safety, and complex systems. [70]

The 2011 guidebook developed the practical discussion. Hollnagel's formulation identifies abilities to respond, monitor, learn, and anticipate. These are related capabilities, not four independent boxes whose completion establishes resilience. Monitoring must concern conditions that matter for response; learning must influence future capability; anticipation must be connected to possible action. An organization can be strong in one ability and constrained by another. [71]

Hollnagel's Safety-I and Safety-II distinction, developed in the 2014 book, directs attention to everyday successful performance as well as to unwanted outcomes. The point is not to stop investigating failure. It is to investigate the adjustments that usually allow work to succeed and the conditions under which similar adjustments can produce a different result. Failure evidence alone gives an incomplete account of the system's performance. [72]

FRAM, presented in Hollnagel's 2012 book, provides a way to model functions and their couplings in complex sociotechnical work. It examines how performance variability can combine. It should not be read as a physical resonance equation applied unchanged to organizations. The method concerns relationships among functions and the circumstances of their performance. The resulting model still requires empirical grounding in the work being examined. [73]

Consider a hypothetical emergency repair that succeeds because a technician recognizes an unfamiliar signal and contacts a colleague outside the usual escalation route. The result reveals an adaptive capability. It does not demonstrate that the arrangement can repeat indefinitely or survive the absence of those people. A resilience investigation asks what resources, relationships, and knowledge made the adjustment possible, and what demands it placed on them.

For diagnostic history, the expansion is substantial. We now need artifacts of successful adaptation, capacity limits, and changing demands, in addition to artifacts of failure. A quiet incident log may coexist with increasing effort required to maintain service. The absence of recorded failure can therefore conceal a deterioration in the conditions that make success possible.

29. Infrastructure, recovery, and limits of adaptive capacity

Resilience is sometimes used for returning to a previous state, sometimes for resisting disturbance, and sometimes for the ability to change. These meanings support different measurements. Woods's 2015 paper distinguishes rebound, robustness, graceful extensibility, and sustained adaptability. The distinctions help prevent an engineering claim about one capability from being presented as evidence of all of them. [74]

His 2018 theory of graceful extensibility examines how systems extend their capacity to adapt when events challenge their boundaries. It emphasizes finite resources, saturation, and networks of adaptive units. This differs from assuming that enough effort can overcome every disturbance. A system that depends on continuing extraordinary effort may be consuming the capacity on which its apparent resilience rests. [75]

Infrastructure makes these relationships visible. Water, power, communications, transport, and buildings support one another and the activities of communities. A restored component does not necessarily restore the service people need. A pumping station may be physically intact but lack power, access, staff, or usable communications. Its functional condition must be evaluated through those dependencies.

NIST's 2015 Community Resilience Planning Guide connects buildings and infrastructure with social and economic functions and recovery goals. It provides a planning approach in which desired performance and recovery periods help set priorities. This makes the definition of success more explicit than a statement that assets should be resilient. Different functions can have different tolerances for interruption. [76]

Let us consider a hypothetical power loss affecting a water system. Backup generation may maintain pumping for a limited period. Recovery then depends on fuel, access, staffing, and the condition of the wider grid. A diagnosis that records only generator availability misses the duration and dependency questions. The relevant evidence includes the operating history of the backup arrangement and the resources needed to sustain it.

Recovery can also change the system. Temporary configurations, bypasses, and altered service priorities may remain after the immediate event. Their continued use needs an explicit account because they can become the starting point of another incident. Returning a dashboard to green does not establish that the original protection, capacity, and diagnostic coverage have been restored.

The historical movement toward resilience planning therefore widens the time horizon of diagnostics. We examine preparation, response, restoration, and adaptation as related periods. We also distinguish resilience of a component from resilience of a service and from the consequences for the people depending on that service. These levels can support one another, but their outcomes should not be assumed to be identical.

30. Operational learning, site reliability, and deliberate perturbation

Digital-service operations developed practices that connect engineering changes with continuing evidence from production. Google’s 2016 Site Reliability Engineering book distinguishes symptoms from causes and black-box monitoring from internal instrumentation. Its monitoring discussion also emphasizes the need for comprehensible alerting and evidence that supports investigation. These are operational formulations of questions that appear elsewhere in this history. [77]

The book’s treatment of postmortems describes learning from significant incidents through accounts that examine contributing conditions and lead to follow-up work. A blameless approach concerns the quality of explanation and learning. It does not mean that actions have no consequences or that every interpretation is equally supported. The investigation still needs a clear account of what happened and evidence for the proposed changes. [78]

Netflix’s 2011 description of the Simian Army documented deliberate introduction of selected failures and checks for abnormal conditions in its cloud environment. Such practices made recovery assumptions testable in an operating setting. They should be understood within their particular technological and organizational context. A cloud experiment is not a general license to inject disturbances into any safety-critical physical process. [79]

Deliberate perturbation has a diagnostic role because it can make hidden dependencies visible. If a service is expected to tolerate the loss of one component, a controlled test can examine that claim. The test needs a defined hypothesis, relevant observations, limits, and a means of ending the intervention. Merely producing a disruption provides little evidence unless the result is related to a prior expectation.

We should also distinguish recovery practice from the full scope of Resilience Engineering. Restarting a process or shifting traffic may test an engineered fallback. Examining how people coordinate under an unfamiliar combination of demands asks another question. The two activities can inform one another without being interchangeable. Automation can preserve one service while increasing the workload or reducing the options available elsewhere.

For a hypothetical incident review, a useful record includes the trigger, developing conditions, detection, response, recovery, and unresolved uncertainty. It also records what helped: a redundant measurement, a person’s recognition, a useful runbook, or an improvised coordination. Learning becomes concrete when these observations change a design, test, procedure, staffing arrangement, or information path and the effect of that change is examined. Closing the report is an administrative event; improving the system is an engineering result.

PART VII

Integrated investigation, correction, and continuing development

The final part brings feedback, authority, trustworthy evidence, and adaptive capacity into the same investigation. Selected cases are read for additional system questions, and a hypothetical pumping station connects physical and software diagnosis with security and resilience. The resulting explanation must support examination of the proposed corrective work.

31. Accident investigations: feedback, authority, and adaptive capacity

The companion histories discuss several cases revisited here. We ask how the observations available at a particular stage could influence action. Physical evidence, recordings, decision records, and accounts of work allow us to investigate feedback, authority, and response. Technical Chapter 34 and Hardware Chapter 34 provide complementary accounts of methods and mechanisms. [82, 83]

Apollo 13, Three Mile Island, Challenger, Therac-25, Columbia, and Qantas Flight 32 provide different settings for this comparison. The investigations document their particular mechanisms and circumstances. The questions in the final column below are our comparative interpretation of those sources; they do not imply that the reports used one shared framework. [35, 48, 36, 40, 37, 38]

Case and date	Evidence and interaction	Additional system question
Apollo 13, 1970	Flight observations, equipment history, and operational recovery.	What could support action during the event, before the full causal history was known?
Three Mile Island, 1979	Displayed control information and actual valve condition.	Which feedback distinguished a command from the achieved physical state?
Challenger, 1986	Seal behavior, earlier observations, and launch decisions.	How did technical evidence and uncertainty reach those with decision authority?
Therac-25, 1980s	Software sequences, machine behavior, and incident reports.	How did assumptions about safeguards affect the interpretation and investigation of reports?
Columbia, 2003	Damage, imagery, analysis, and organizational evidence.	How did the organization handle uncertainty and opportunities to obtain further evidence?
Qantas Flight 32, 2010	Manufacturing origin, effects across systems, and crew response.	Which remaining capabilities and adaptations supported continued control?

Table 3. Selected cases examined through feedback, authority, and adaptive capacity. Dates identify the events or incident period; the questions are this volume's comparative interpretation.

Three Mile Island shows why the meaning of an indication matters to action. A display associated with a valve command did not directly establish the valve's physical position. Challenger and Columbia add a

different relationship: evidence must enter an organizational process capable of examining uncertainty and changing a decision. Neither a correct local measurement nor a technically sound analysis guarantees that this connection exists. The investigator must reconstruct the particular information and authority available. [48, 36, 37]

Recovery introduces another source of evidence. Apollo 13 and Qantas Flight 32 can be examined for what supported continued action as well as for what failed. A successful response does not establish that the same resources or options will exist next time. The resilience question concerns the conditions that made the response possible and the demands it placed on the system. [35, 38]

This reading also limits hindsight. A later report can combine material, experiments, and records unavailable to participants during the event. We should distinguish what was known, what could have been investigated then, and what was established afterward. Corrective work can then address a particular missing observation, constraint, decision path, or recovery capability. Repeating a case becomes useful when it supports such an additional, explicit question.

32. Connections with pattern-oriented software diagnostics

Our pattern-oriented approach begins with recurring structures and behaviors in diagnostic artifacts. A pattern gives us a way to recognize a configuration, describe it, relate it to other observations, and continue an investigation. Recognition is valuable because it organizes attention. It does not turn resemblance into a demonstrated cause. The same visible form can have different mechanisms, and one mechanism can leave different forms under different recording conditions. [2, 42]

Theoretical Software Diagnostics and our trace and log analysis reference provide vocabulary for reasoning about these relationships. Visibility Limit directs attention to what the artifact cannot expose. No Activity requires an account of expected activity and its recording. Message Context connects a message to information needed for interpretation. Adjoint Thread of Activity provides a way to organize related activity that may not coincide with one directly recorded execution thread. We use these ideas here as methodological connections, without claiming that an engineering predecessor employed the same pattern vocabulary. [2, 42]

Let us consider a hypothetical industrial trace containing an accepted start command and no later confirmation. A No Activity interpretation remains incomplete until we know whether confirmation was expected, whether its channel was active, and whether the collection preserved that period. A physical inspection may show that the machine did start. We have then diagnosed a limitation in the trace or its interpretation, rather than demonstrated a failure to execute the command. The distinction mirrors the separation of command and state examined in earlier chapters.

In Bridging the Great Divide, we connect patterns of software construction with patterns of execution. This supplies another useful direction for engineering diagnostics. Architecture describes intended relationships; execution artifacts expose particular realized relationships. We can ask which design element should account for an observed interaction, and which artifact would expose a suspected

defect in the design. The connection is an organizing proposal, not a claim that one pattern catalog replaces an engineering model. [80]

For a hypothetical pump-control system, the intended architecture includes a controller, actuator, independent protection, and a historian. An execution investigation may reveal that the supposedly independent records all originate in the controller. The design description and the diagnostic artifacts then expose a mismatch in assumed independence. Corrective work may require a change in acquisition architecture rather than an additional message from the existing source.

This comparison also helps us connect defects with their manifestations. A defect in a configuration can produce different symptoms as operating conditions change. Conversely, several defects can produce the same alarm. A useful pattern description should preserve the relation among context, observed form, possible mechanisms, and the next discriminating observation. Giving a new name to a symptom without preserving these relationships adds vocabulary but little diagnostic capability.

Security and resilience add further questions to this approach. Could an adversary manufacture the observed pattern or suppress its expected counterpart? Does the trace expose only nominal operation while omitting the adaptations that sustained service? These questions extend the investigation around the artifact. They do not establish that software patterns have already been validated for every physical or organizational setting. Such applications require their own cases, limits, and evaluation.

33. A worked example: diagnosing a pumping-station disturbance

Let us construct a hypothetical pumping station serving a storage reservoir. The station contains a variable-speed pump, inlet and outlet measurements, a programmable controller, an operator display, and a historian. A separate instrument can measure outlet flow locally. An alternate station can sustain part of the service. The example is invented to connect the book's methods; its values and sequence are not a report of an actual installation.

We use the station to connect the three histories' perspectives. Physical examination investigates delivery and condition. Controller and acquisition analysis investigates the production of the displayed evidence. The engineering investigation relates both to service demand, control, authorized change, and the resources available for continued operation.

At the start of the investigation, the display reports ordinary flow, but the reservoir level is falling faster than expected for the recorded demand. The pump's electrical demand has changed. Operators also report an intermittent vibration. None of these observations establishes a unique cause. A leak, reduced supply, an incorrect level measurement, a pump condition, an unexpected demand, and a recording problem remain possible. Some can coexist.

Establishing the evidence

The investigation begins by recording the configuration and identifying the source of each displayed value. The display and historian receive flow from the same controller tag, so their agreement is one observation represented twice. The independent local instrument reports reduced flow. A check of the reservoir level measurement supports the observed level trend within the relevant uncertainty. The discrepancy now concerns the flow report and its relation to the physical process, rather than an unsupported assumption that the reservoir measurement is wrong.

Relative period	Preserved observation	Interpretation still requiring examination
Before the disturbance	Historian values appear steady; a configuration change is recorded.	The change may be relevant, incidental, or incompletely described.
Detection	Reservoir level falls; the display shows ordinary flow.	Actual flow, demand, level measurement, or reporting may be wrong.
Independent comparison	Local flow measurement supports reduced delivery.	Shared display and historian values no longer count as independent confirmation.
Controlled examination	A retained controller value fails to refresh after an acquisition fault.	Stale data explain the display discrepancy, but not yet the reduced physical delivery.
Mechanical inspection	The inlet path contains an obstruction.	Its significance depends on the observed pressure and performance changes.
Verification and recovery	Corrected acquisition and cleared inlet restore expected relationships.	Wider operating conditions and the origin of the changes remain separate questions.

Table 4. Hypothetical evidence sequence. Relative periods avoid implying clock precision that the example has not established.

Separating explanations

The software examination finds that an acquisition fault can leave the previous flow value in the controller while the display continues to label it as current. Reproducing that behavior under controlled conditions establishes a defect in the information path. It does not establish the physical cause of reduced delivery. This is an important stopping point: a successful software diagnosis has explained one symptom, not the entire incident.

Mechanical inspection then finds an obstructed inlet. The associated pressure measurements and a comparison after clearing the obstruction support its contribution to reduced flow. The system therefore contained at least two consequential conditions: restricted delivery and a reporting defect that concealed part of its development. If the investigation had ended after correcting the displayed value, the service problem would have remained. If it had ended after clearing the inlet, the diagnostic defect would have remained.

Hypothesis	Evidence that would help discriminate it	What the initial display cannot establish
Restricted inlet or another hydraulic condition	Local pressure and flow, physical examination, response to a controlled change.	The actual hydraulic condition from a stale value.
Wrong measurement or conversion	Reference comparison, acquisition status, units and scaling, signal path.	Whether a plausible number represents a fresh measurement.
Unauthorized modification	Preserved configuration, authenticated change records, access evidence, independent records.	Whether a change was accidental or deliberate.
Unexpected service demand	Independently reconciled demand and storage history.	Whether lower storage results entirely from the station.

Table 5. Hypotheses are evaluated through discriminating evidence. The rows are not mutually exclusive.

Security and resilience interpretations

The recorded configuration change now requires examination. Its existence does not prove malicious action. The investigation compares the installed configuration with a retained version, checks the authorized work record, and examines the available access evidence. Suppose the change can be traced to an authorized maintenance task, while the evidence does not support an intrusion. The conclusion concerns the evidence available for this case. It should not be rewritten as proof that the system cannot be attacked.

The response also depended on operators noticing the reservoir trend, obtaining an independent measurement, and using the alternate station within its capacity. These activities preserved part of the service while the causes remained uncertain. A resilience analysis asks how that capacity was made available, whether the staff and resources could sustain it, and which dependencies could have defeated it. The alternate station's existence alone would not answer those questions.

Corrective work follows the demonstrated relationships. The inlet condition needs its maintenance explanation. The acquisition path needs a freshness indication and behavior appropriate to loss of measurement. Configuration work needs adequate testing of failure conditions. Recovery planning needs an account of alternate capacity and duration. Verification then examines whether these changes produce the intended behavior. The integrated diagnosis consists of several supported explanations joined by the incident, not one convenient root-cause label.

34. Automated diagnosis and the assurance of investigative tools

Statistical monitoring, model-based diagnosis, signal processing, and machine learning contribute different ways of examining engineering evidence. Technical Chapters 32–33 and Hardware Chapters 31–32 provide their complementary histories and applications. The question developed here concerns the diagnostic tool as part of an operating system: what evidence enters it, what claim comes out, and what action follows? [20, 21, 22, 19, 82, 83]

An engineering label can originate in a confirmed inspection, a maintenance decision, or an unverified alarm. A model trained on those records inherits their differences. If its recommendations alter maintenance or operating conditions, they also alter the later evidence. Evaluation needs to preserve that intervention history when assessing the tool's contribution.

AI assistance can retrieve records, compare configurations, and suggest discriminating tests. It can also merge incompatible accounts or introduce unsupported details. NIST's 2023 AI Risk Management Framework places trustworthiness within the tool's life cycle and context of use. For diagnostics, this makes validation, traceable supporting evidence, and continuing observation part of integrating the tool. [81]

Consider a hypothetical assistant that recommends changing a controller setting. The engineering claim requires evidence about the physical process and the installed configuration. Permission to propose the change, authority to perform it, operating limits, and verification of the result are further requirements. The recommendation itself becomes an artifact to preserve.

Security asks whether the inputs, retrieved sources, or proposed actions can be manipulated. Resilience asks how investigation and essential operation continue when the tool is unavailable or misleading. These questions connect automated diagnosis to the same feedback, authority, and capacity relationships examined throughout this history.

35. Continuing development of engineering diagnostics

A fracture examination, a control-theoretic observer, a software trace, and an operator's account can be required in one investigation. The companion Technical and Hardware histories develop many of the methods that produce these artifacts. This volume has followed the relationships through which their findings support understanding and intervention across an engineered system. Feedback, integration, authority, security, and adaptive capacity make those relationships explicit. [82, 83]

Several recurring developments can nevertheless be identified. Measurements became retained records. Records became comparable through standards, calibration, and contextual information. Models made internal relationships explicit. Networks connected observations from different locations. Software made acquisition, interpretation, and response programmable. Safety, security, and resilience perspectives expanded the investigation beyond the condition of an individual component. These developments overlap and continue at different rates in different domains.

We can describe the resulting engineering obligation as maintaining a path from a claim to its supporting artifacts. If we claim that a structure is undamaged, which examinations and sensitivities support that claim? If we claim that a control action occurred, do we have a command record or evidence of the resulting state? If we claim that a service is resilient, which disturbances, resources, dependencies, and time periods have been considered? The questions become more precise when the object and limits of the claim are explicit.

The same path should extend toward correction. A diagnostic result is incomplete for some purposes until it supports a useful decision, but the decision may be to acquire better evidence rather than to change the equipment immediately. Other situations require protective action before a unique diagnosis is available. The urgency of action and the certainty of explanation are separate dimensions. A historical account that records only the final repair can conceal this important part of engineering practice.

There are also limits to historical inference. A famous accident can preserve unusually detailed evidence while ordinary successful operation remains poorly documented. New terminology can make an older practice look newly invented. A formal standard can appear more influential in the surviving literature than informal coordination was in the actual work. We have therefore used documented milestones and selected cases without assigning every development a single origin or assuming a uniform direction of progress.

Our synthesis is methodological. Engineering diagnostics advances when the relationships among system, artifact, interpretation, intervention, and outcome become more explicit and more open to examination. Security Engineering adds the possibility that those relationships are deliberately manipulated. Resilience Engineering adds the need to understand how performance is sustained as demands, resources, and conditions change. Together, they extend the history from finding failed parts to investigating the continued possibility of trustworthy and effective engineering action.

Conclusion

We have followed engineering diagnosis from models and recorded behavior into feedback, integration, security, and resilience. A material finding, an instrument reading, a software state, and an operator's account can each be correct while describing different parts of an event. The engineering explanation must establish how those parts interacted and which relationships made the observed outcome possible.

This provides the connection with the companion Technical and Hardware histories. Their methods help establish physical condition and the behavior of electronic and computing devices. The system investigation connects such findings with the process being controlled, the evidence available for decisions, the authority to intervene, and the capacity to sustain a service. Each perspective can identify a defect that remains unresolved by correction at another level.

Security and resilience extend the same investigation in different directions. Evidence and actions can be manipulated deliberately. Successful operation can also depend on coordination and resources that ordinary records omit. We therefore need accounts of how evidence was produced, why it could be trusted, what options were available, and what made effective action possible.

The continuing task is to preserve these relationships through design, operation, and change. A useful correction has supporting evidence, an explicit purpose, and an examined effect on the wider system. Records of unresolved alternatives, useful adaptations, and changed dependencies allow later investigators to revisit the explanation. They also make it possible to learn from continued performance as well as from failure.

Appendix A. Selected chronology

Dates identify publications, documented instruments, events, or specified editions. They are not claims of unrestricted priority. Event dates and later report dates are distinguished where relevant. The final column states the diagnostic significance emphasized in this history.

Date	Development or documented milestone	Diagnostic significance
First century BCE	Vitruvius, On Architecture; cited through Morgan's 1914 translation. [3]	Organized discussion of materials, construction, water, and machines.
1638	Galileo, Two New Sciences. [4]	Mathematical examination of strength and scaling.
1759	Smeaton's experimental study of water and wind power. [5]	Quantitative comparison of machine performance.
Late eighteenth century	Watt's indicator and Southern's recording additions, as distinguished in the museum record. [6]	Internal machine activity becomes a retained diagram.
1868	Maxwell, On Governors. [9]	Regulation and instability become objects of mathematical analysis.
1914–1915	First ASME boiler-code edition: dated 1914, published 1915. [45]	Shared criteria connect design, fabrication, and inspection.
1921	Griffith's rupture paper. [24]	Flaws and energy enter a theory of fracture.
1924; 1931	Shewhart's control-chart memorandum and book. [13]	Recorded variation supports investigation of process change.
1932	Nyquist, Regeneration Theory. [10]	Feedback stability is examined in electrical amplification.
1946	Firestone's reflectoscope paper. [27]	Sound reflections reveal internal material boundaries and flaws.
1951	Weibull's statistical-distribution paper. [16]	A flexible distribution supports analysis of failure-related variability.
1954	Page, Continuous Inspection Schemes. [14]	Accumulated evidence supports detection of process shifts.
1960	Kalman's general theory of control systems. [11]	State, controllability, and observability acquire a common framework.
1963	Paris and Erdogan examine crack-propagation laws. [25]	Crack growth is connected to testable relationships and data requirements.
1960s	Development of HAZOP at ICI, discussed retrospectively by Kletz. [33]	Structured examination of process deviations.
1964	Luenberger's observer paper. [12]	State estimation is formulated through a dynamical observer.
1969	Peck's Rankine Lecture on the observational method. [28]	Predictions, field observations, and planned responses are connected.
1970	Apollo 13 accident and review-board report. [35]	Flight evidence is connected to prior equipment history.
1973	Holling's ecological resilience paper. [69]	A distinct account of persistence under disturbance enters the later conceptual background.
1975	Saltzer and Schroeder on protection of computer information. [51]	Security design principles organize protection mechanisms.

Date	Development or documented milestone	Diagnostic significance
1976	Willsky's survey of failure detection in dynamic systems. [20]	Analytical relationships support fault detection.
1978	Nowlan and Heap's reliability-centered maintenance work, discussed in NASA's guide. [18]	Maintenance selection is related to functions and failure consequences.
1978	Lamport on time, clocks, and event ordering. [41]	Distributed reconstruction distinguishes logical order from a single clock sequence.
1979	Three Mile Island accident. [48]	Physical state and displayed information become central to investigation.
1981	NRC Fault Tree Handbook. [17]	A documented methodology organizes combinations leading to a top event.
1983	Phadke, Thorp, and Adamiak on phasor and frequency measurement. [31]	Electrical-system measurements support wider dynamic reconstruction.
1983	Bainbridge, Ironies of Automation. [49]	Remaining human responsibilities become part of automation analysis.
1984	Isermann's process-fault review; Perrow's Normal Accidents. [21, 63]	Model-based detection and system-interaction analysis develop distinct perspectives.
1986	Challenger accident and commission report. [36]	Technical mechanisms and decision conditions are investigated together.
1987	Reiter; de Kleer and Williams on model-based diagnosis. [22, 23]	Logical models support reasoning about inconsistent observations and multiple faults.
1987	Denning's intrusion-detection model. [53]	Security monitoring becomes an explicit detection architecture.
1987	Rochlin, La Porte, and Roberts on aircraft-carrier reliability. [68]	Successful organizational performance becomes an object of investigation.
1993	Leveson and Turner's Therac-25 investigation. [40]	Software, equipment, and operational evidence are analyzed together.
1997	Rasmussen's dynamic risk-management framework. [67]	Interacting social and organizational levels shape risk.
1999	Schneier's attack-tree presentation. [58]	Alternative routes to an adversarial objective become explicit.
2003	Columbia accident and investigation report. [37]	Physical and organizational evidence inform a wider causal account.
2004	Avizienis and colleagues' dependability taxonomy; Leveson's systems accident model. [1, 64]	Service boundaries and control relationships sharpen different diagnostic questions.
2005; 2007	BP Texas City disaster and CSB final report. [34]	Process evidence is connected to maintenance and organizational conditions.
2006	Resilience Engineering: Concepts and Precepts. [70]	Adaptive performance receives a prominent shared treatment.
2006–2007	Jardine and colleagues on condition-based maintenance; Farrar and Worden on structural health monitoring. [19, 29]	Condition information is connected to decision models and damage identification.
2008; 2012	Cited JCGM editions of the uncertainty guide and metrology vocabulary. [8, 7]	Common measurement concepts support interpretation and comparison.

Date	Development or documented milestone	Diagnostic significance
2010	Stuxnet advisory; Qantas Flight 32 engine failure. [56, 38]	Deliberate control-system interference and physical failure propagation require different evidence.
2011	Resilience Engineering in Practice; Netflix’s Simian Army account. [71, 79]	Adaptive abilities and deliberate operational perturbation enter different practical settings.
2012	Hollnagel’s FRAM; Leveson’s Engineering a Safer World; NASA digital-twin paper. [73, 65, 44]	Functional coupling, system constraints, and maintained models support different forms of analysis.
2014	Hollnagel’s Safety-I and Safety-II. [72]	Everyday successful adjustments become an explicit subject of safety inquiry.
2015	Woods’s resilience distinctions; NIST community-resilience guide. [74, 76]	Different adaptive capabilities and recovery goals are separated.
2016–2017	Google’s Site Reliability Engineering (2016); NASA’s Systems Engineering Handbook, Rev 2 (published 2017). [77, 78, 46]	Operational evidence and life-cycle engineering remain connected concerns.
2018	Woods on graceful extensibility. [75]	Adaptive capacity is examined at its limits.
2021–2023	Cited NIST cyber-resilience, secure-systems, secure-development, and OT publications. [61, 60, 59, 57]	Security is integrated with engineering life cycles and contested operation.
2023	NIST AI Risk Management Framework 1.0. [81]	Diagnostic AI can itself be examined as a contextual engineering system.
2023–2024	Cited editions of our trace-pattern reference and theoretical diagnostics collection. [42, 2]	Analysis patterns connect artifacts, interpretation, and software construction and execution.
2024	NASA-STD-7009B. [43]	Model and simulation credibility are treated through explicit practices.
2025	NIST SP 800-61 Revision 3. [62]	Incident response is related to wider cybersecurity risk management.

Table 6. Selected chronology. The thematic chapters explain the significance and limitations of these milestones.

Appendix B. Glossary

The definitions below summarize usage in this history. Some terms have narrower meanings in particular standards or mathematical frameworks. Where that matters, the chapter context and cited source take precedence over a general definition.

Adaptive capacity. The resources, relationships, knowledge, and available adjustments through which a system can respond to changing demands. Its limits can change during an event.

Alarm. An indication intended to prompt attention or action when selected conditions are detected. An alarm embodies acquisition, decision, and presentation choices.

Anomaly. A departure from a stated or implicit expectation. Its presence does not determine whether the object, measurement, reference, or interpretation is defective.

Artifact. A preserved object or representation available for analysis, such as a specimen, waveform, log, configuration, simulation result, or incident account.

Assurance. A reasoned basis for confidence in a claim, supported by relevant evidence and an account of the conditions and limitations of that evidence.

Attack tree. A model organizing alternative means of achieving an adversarial objective. Its diagnostic usefulness depends on the adequacy of the modeled routes.

Calibration. Establishment of a relationship between reference values and measurement indications under specified conditions. Calibration is distinct from adjustment and does not alone verify every operating condition.

Causal explanation. An account of mechanisms or relationships producing an outcome, supported by evidence that goes beyond similarity or temporal succession alone.

Common-cause failure. Multiple failures arising from a shared cause. Apparent redundancy can be ineffective when channels share the relevant vulnerability.

Condition monitoring. Collection and interpretation of evidence about equipment or system condition over time, often to support maintenance and operating decisions.

Configuration. The identified arrangement of hardware, software, settings, connections, and relevant documents at a particular stage of a system's history.

Controllability. In a specified control model, a property concerning whether admissible inputs can move the system between relevant states. It is distinct from observing those states.

Cyber resilience. Capacity to anticipate, withstand, recover from, and adapt to adverse conditions involving cyber resources, within the scope of the cited engineering framework.

Debugging. Locating and correcting defects in software or, in this book's broader analytical usage, in an engineered arrangement and its investigation procedures.

Defect. A shortcoming relative to an intended requirement, property, or use. It should not be treated as synonymous with every anomaly.

Detection. Establishing an indication that a specified condition may be present. Detection alone does not necessarily identify the condition's cause or location.

Diagnosis. Investigation of a condition through artifacts, models, and comparisons, with the aim of identifying or narrowing explanations and their consequences.

Diagnostic distinguishability. The ability of available evidence, under stated assumptions, to separate candidate conditions or explanations. Different faults can remain indistinguishable with existing measurements.

Digital twin. A representation maintained in relation to an identified physical system and used for specified engineering purposes. The name does not establish model fidelity or valid prediction.

Error. In the dependability terminology used here, a system-state condition that may lead to a service failure. Other fields use the term differently.

Evidence. An artifact or established observation considered in relation to a claim. Its strength depends on provenance, relevance, interpretation, and the alternatives it can distinguish.

Failure. Loss of, or departure from, required service or function at a stated boundary. A local action can succeed while a larger system fails.

Fault. In the dependability framework, a cause of an error. In other engineering usage, the term may identify a physical or functional abnormal condition.

Fault tree. A model of combinations of events leading to a defined top event. Its completeness and assumptions determine the scope of its conclusions.

Feedback. Use of information about a system's behavior to influence its subsequent behavior. Feedback can compensate for disturbances or contribute to instability.

Forensic preservation. Retention of artifacts and their acquisition history so that later examination can address integrity, provenance, and the effects of handling.

Functional safety. Safety associated with the correct functioning of specified safety-related arrangements. Its detailed treatment depends on the system and applicable framework.

Graceful extensibility. Capacity to extend adaptive performance when demands approach or exceed existing boundaries, as developed in Woods's account of adaptive systems.

Hazard. A condition with the potential to contribute to harm or loss in a specified setting. A hazard is not identical to an observed accident.

HAZOP. Hazard and operability study: structured examination of deviations from design intent and their possible causes and consequences in a defined process.

Historian. A system that retains process or operational data over time. Its record depends on the sources, sampling, transformations, timing, and preservation arrangements.

Independence of evidence. Separation of relevant origins and failure mechanisms between observations. Different displays of one source are not independent confirmation of its accuracy.

Indication. A response or representation produced by an examination or measurement process. Its relationship to an actual feature or condition still requires interpretation.

Latent condition. A condition that can remain inconspicuous until it combines with circumstances that make its consequences significant. It need not originate near the eventual event.

Measurement uncertainty. A parameter characterizing the dispersion of values attributed to a measurand from the available information. It is not simply an admission of careless measurement.

Metrological traceability. A measurement result's relationship to a reference through a documented calibration chain, with each link contributing to uncertainty.

Model residual. A difference or consistency measure relating observed behavior to behavior expected under a model. A residual can reflect a process fault, measurement problem, or model inadequacy.

Nondestructive examination. Investigation of relevant material or component features without making the object unusable for its intended purpose through the examination itself.

Observability. In control theory, the ability to determine the internal state from known inputs and measured outputs over a specified finite interval, under a stated system model. Broader operational usage should identify what evidence becomes available.

Observer. A model-based dynamical arrangement used to estimate state from available information. Its estimates depend on the model, inputs, observations, and assumptions.

Operational technology. Programmable systems or devices that monitor or affect physical processes and environments. Their security has direct relationships with performance, reliability, and safety.

Pattern. A recurring and describable arrangement of structure or behavior that helps organize analysis. Recognizing a pattern does not establish a unique mechanism.

Prognosis. An estimate of how condition or capability may develop under stated future assumptions. It should preserve uncertainty about those assumptions and the estimated development.

Provenance. Information about an artifact's origin, acquisition, transformations, custody, and identity. Provenance supports interpretation without guaranteeing that the original account was complete or true.

Recovery. Restoration of a specified function or capability after disruption. The recovered state may differ from the earlier state and may retain temporary arrangements.

Reliability. Ability to provide required service over a stated period and conditions, or a corresponding measure. The boundary, conditions, and duration matter.

Resilience Engineering. An engineering and safety perspective concerned with how sociotechnical systems adjust to conditions and sustain performance, including the resources and limits of adaptation.

Robustness. Ability to maintain specified performance against a defined range of disturbances or variation. Performance beyond that range requires separate evidence.

Root cause. A selected causal explanation regarded as sufficiently fundamental for a particular investigation. The label should not conceal multiple mechanisms or depend on an unstated stopping rule.

Safety. A concern with preventing unacceptable harm or loss. Its engineering treatment depends on hazards, interactions, constraints, and the conditions of operation.

Security Engineering. Design and maintenance of protection against adversarial actions, including examination of threats, vulnerabilities, controls, and the trustworthiness of supporting evidence.

Structural health monitoring. Acquisition and interpretation of information intended to identify changes relevant to structural condition, with operational and environmental variability considered.

Symptom. An observed manifestation that motivates or supports investigation. A symptom can have several causes and can change with the observation method.

Testability. The extent to which relevant conditions can be stimulated, observed, and distinguished through tests. Access alone does not establish adequate diagnostic discrimination.

Threat model. A representation of protected interests, adversarial capabilities, possible actions, and relevant system conditions used to organize security analysis.

Time synchronization. Establishment of sufficiently related time references for a stated purpose. Synchronization does not itself identify which stage of an observation a timestamp describes.

Validation. Evaluation of suitability for an intended use, including whether the selected requirements or model adequately address that use. Terminology varies by domain.

Verification. Evaluation of conformity to specified requirements or correct implementation of a model or design. Conformity alone does not establish fitness for every intended use.

Visibility limit. A boundary on what an artifact or acquisition arrangement can expose. The absence of recorded evidence beyond that boundary cannot settle the corresponding question.

Appendix C. Evidence taxonomy

This taxonomy classifies artifacts by the relationships they preserve, rather than by the profession that uses them. A single investigation can require several classes. The same artifact can also play more than one role: a relay record describes a physical disturbance through a measurement path and records a protection system's response.

Artifact class	What it can preserve	Principal interpretive limit
Physical remains and specimens	Material, geometry, damage, deposits, and traces of loading or exposure.	The event and subsequent handling may change or destroy earlier features.
Inspection images and indications	Features accessible through a particular examination method and geometry.	Resolution, sensitivity, coverage, and method-specific artifacts constrain inference.
Time series and waveforms	Sampled evolution of measured quantities.	Sampling, filtering, saturation, timing, and operating context affect the apparent behavior.
Event and protection records	Selected transitions, decisions, trips, and alarms.	Recorded action need not reveal the initiating condition or achieved physical state.
Calibration and reference records	Relationships between indications and reference values under specified conditions.	Laboratory relationships may not cover all installed or transient behavior.
Software traces and memory artifacts	Recorded activity or retained state of a computational arrangement.	Collection scope, loss, transformations, and execution timing limit reconstruction.
Configuration and change records	Versions, settings, connections, and documented interventions.	The recorded configuration may differ from the configuration actually installed.
Design and requirement documents	Intended functions, assumptions, interfaces, and constraints.	Intended structure is not direct evidence of realized behavior.
Models and simulations	Consequences of selected assumptions and implemented relationships.	Agreement with one result may not identify the correct mechanism or intended-use validity.
Test and experiment records	Responses to specified interventions and comparisons.	Coverage and environmental differences limit transfer to untested conditions.
Maintenance and manufacturing history	Material origins, work performed, acceptance evidence, and replacements.	Administrative completion does not guarantee technical adequacy or complete recording.
Security and forensic records	Access, actions, recovered data, preservation history, and selected adversary behavior.	The producer or collection path may be compromised; integrity after collection is only one property.
Human accounts and work observations	Perceived conditions, reasoning, coordination, and adaptations.	Memory, perspective, timing, and recording purpose shape the account.
Incident and investigation reports	Integrated reconstructions, findings, uncertainties, and proposed changes.	Summaries can hide conflicts or inherit weaknesses from underlying artifacts.
Recovery and exercise records	Restoration sequence, resource use, dependencies, and response to selected disturbances.	One successful recovery does not establish capability under all combinations or durations.
Baselines and learned representations	Selected regularities used for comparison, classification, or prediction.	Selection, labels, regime change, and omitted cases can distort the reference.

Table 7. Evidence classes and their limits. The taxonomy is our synthesis of the practices discussed in the chapters.

From artifact to supported claim

We can describe an artifact through its object, acquisition, preservation, transformations, and use. The object is what the artifact concerns. Acquisition describes how the signal or material was obtained. Preservation describes what survived. Transformations include filtering, normalization, reconstruction, translation, and summarization. Use identifies the claim for which the artifact is being considered. These dimensions help separate a weak artifact from an unsuitable interpretation of an otherwise useful artifact.

For example, an hourly average can be accurate for an energy-accounting purpose and unsuitable for reconstructing a short disturbance. A signed configuration file can have intact provenance while containing an erroneous setting. A faithful interview transcript can preserve a participant's account without establishing that every perceived condition matched the physical state. Evaluating evidence means matching its properties to the particular question.

Combining evidence without counting the same source twice

Agreement matters when the observations have relevant independence. Two displays fed by one sensor can help identify a display defect if they disagree, but their agreement does not independently validate the sensor. Two analyses of one retained waveform can expose a processing mistake, yet they still share the waveform's acquisition limits. A physical inspection and a log may constrain different parts of an explanation even when neither is decisive alone.

Contradiction can be more useful than agreement. If a model predicts a pressure change that no reliable measurement shows, the discrepancy can direct attention to assumptions, topology, or the proposed mechanism. The contradiction should remain visible until it is explained. Removing it through averaging, selective quotation, or premature classification destroys information that could distinguish competing accounts.

Preserving the history of investigation

The investigation itself changes the system and produces new artifacts. A probe can alter electrical loading. A software diagnostic mode can alter timing. A shutdown can remove transient conditions. A repair can destroy the original configuration. Recording these interventions allows later readers to distinguish evidence from the original event from evidence created while testing an explanation.

We should also preserve why a hypothesis was rejected or left unresolved. Otherwise, a later investigator may repeat the same inconclusive test or mistake an abandoned conjecture for a finding. A useful diagnostic record connects each important claim to its supporting artifacts, material assumptions, and remaining alternatives. That connection is the common thread through the histories of engineering inspection, observability, debugging, security, and resilience.

References

The dates below identify the cited works or editions. Later online copies do not change the original publication dates. URLs are included explicitly for printed copies. Institutional pages without a publication year are marked n.d. Online sources consulted on 6 October 2026. Standards are cited by edition as historical sources; a citation does not imply that an edition is the latest or that it governs every application.

- [1] Avizienis, Algirdas, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing* 1(1) (2004): 11–33.
<https://doi.org/10.1109/TDSC.2004.2>
- [2] Vostokov, Dmitry, and Software Diagnostics Institute. *Theoretical Software Diagnostics: Collected Articles*. Fourth edition. OpenTask, 2024.
<https://www.dumpanalysis.org/theoretical-software-diagnostics-book>
- [3] Vitruvius. *The Ten Books on Architecture*. First century BCE; translated by Morris Hicky Morgan, Harvard University Press, 1914. Books I–II, VIII, and X.
<https://www.gutenberg.org/ebooks/20239>
- [4] Galilei, Galileo. *Discorsi e dimostrazioni matematiche intorno a due nuove scienze*. Leiden: Elzevir, 1638. Historical copy and description, University of Oklahoma Libraries.
<https://galileo.ou.edu/exhibits/discourse-two-new-sciences.html>
- [5] Smeaton, John. An experimental enquiry concerning the natural powers of water and wind to turn mills, and other machines, depending on a circular motion. *Philosophical Transactions* 51 (1759): 100–174.
<https://doi.org/10.1098/rstl.1759.0019>
- [6] Science Museum Group. Watt’s steam engine indicator. Collection record, object co51439. n.d. The record distinguishes the instrument from Southern’s recording additions.
<https://collection.sciencemuseumgroup.org.uk/objects/co51439>
- [7] Joint Committee for Guides in Metrology. *International vocabulary of metrology—Basic and general concepts and associated terms (VIM)*. Third edition, 2008 version with minor corrections, JCGM 200:2012.
https://www.bipm.org/documents/20126/2071204/JCGM_200_2012.pdf
- [8] Joint Committee for Guides in Metrology. *Evaluation of measurement data—Guide to the expression of uncertainty in measurement*. JCGM 100:2008.
<https://doi.org/10.59161/JCGM100-2008E>
- [9] Maxwell, James Clerk. On governors. *Proceedings of the Royal Society of London* 16 (1868): 270–283. Original paper also hosted on a University of Cambridge teaching page.
<https://doi.org/10.1098/rspl.1867.0055>
https://www.damtp.cam.ac.uk/user/gold/pdfs/teaching/old_literature/Maxwell1868.pdf
- [10] Nyquist, Harry. Regeneration theory. *Bell System Technical Journal* 11, no. 1 (1932): 126–147.
<https://doi.org/10.1002/j.1538-7305.1932.tb02344.x>
- [11] Kalman, Rudolf E. On the general theory of control systems. *IFAC Proceedings Volumes* 1(1) (1960): 491–502.
[https://doi.org/10.1016/S1474-6670\(17\)70094-8](https://doi.org/10.1016/S1474-6670(17)70094-8)
- [12] Luenberger, David G. Observing the state of a linear system. *IEEE Transactions on Military Electronics* MIL-8(2) (1964): 74–80.
<https://doi.org/10.1109/TME.1964.4323124>
- [13] NIST/SEMATECH. *e-Handbook of Statistical Methods*, section 6.1.1, How did Statistical Quality Control Begin? Institutional account of Shewhart’s 1924 memorandum and 1931 book. n.d.
<https://www.itl.nist.gov/div898/handbook/pmc/section1/pmc11.htm>

- [14] Page, E. S. Continuous inspection schemes. *Biometrika* 41(1–2) (1954): 100–115.
<https://doi.org/10.1093/biomet/41.1-2.100>
- [15] Toyota Motor Corporation. Toyota Production System. Corporate account of jidoka and just-in-time production. n.d.
<https://global.toyota/en/company/vision-and-philosophy/production-system/>
- [16] Weibull, Waloddi. A statistical distribution function of wide applicability. *Journal of Applied Mechanics* 18(3) (1951): 293–297. Original article hosted on a McGill University course page.
<https://jhanley.biostat.mcgill.ca/bios601/CHchapters040506/Weibull-ASME-Paper-1951.pdf>
- [17] Vesely, W. E., F. F. Goldberg, N. H. Roberts, and D. F. Haasl. Fault Tree Handbook. NUREG-0492. U.S. Nuclear Regulatory Commission, January 1981.
<https://www.nrc.gov/regulations-legislation/nureg-series-publications/publications-prepared-by-nrc-staff/sr0492>
- [18] NASA. Reliability-Centered Maintenance Guide for Facilities and Collateral Equipment. September 2008. Includes the development of reliability-centered maintenance and discussion of Nowlan and Heap’s 1978 work.
<https://www.nasa.gov/wp-content/uploads/2023/06/nasa-rcmguid.pdf>
- [19] Jardine, Andrew K. S., Daming Lin, and Dragan Banjevic. A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing* 20(7) (2006): 1483–1510.
<https://doi.org/10.1016/j.ymssp.2005.09.012>
- [20] Willsky, Alan S. A survey of design methods for failure detection in dynamic systems. *Automatica* 12(6) (1976): 601–611.
https://willsky.lids.mit.edu/publ_pdfs/15_pub_AUTO.pdf
- [21] Isermann, Rolf. Process fault detection based on modeling and estimation methods—A survey. *Automatica* 20(4) (1984): 387–404.
[https://doi.org/10.1016/0005-1098\(84\)90098-0](https://doi.org/10.1016/0005-1098(84)90098-0)
- [22] Reiter, Raymond. A theory of diagnosis from first principles. *Artificial Intelligence* 32(1) (1987): 57–95.
[https://doi.org/10.1016/0004-3702\(87\)90062-2](https://doi.org/10.1016/0004-3702(87)90062-2)
- [23] de Kleer, Johan, and Brian C. Williams. Diagnosing multiple faults. *Artificial Intelligence* 32(1) (1987): 97–130.
[https://doi.org/10.1016/0004-3702\(87\)90063-4](https://doi.org/10.1016/0004-3702(87)90063-4)
- [24] Griffith, A. A. The phenomena of rupture and flow in solids. *Philosophical Transactions of the Royal Society of London A* 221 (1921): 163–198.
<https://doi.org/10.1098/rsta.1921.0006>
- [25] Paris, P., and F. Erdogan. A critical analysis of crack propagation laws. *Journal of Basic Engineering* 85(4) (1963): 528–533.
<https://doi.org/10.1115/1.3656900>
- [26] American Society for Nondestructive Testing. NDT Method Overview: Visual Testing. ASNT Pulse, 2023.
<https://www.asnt.org/standards-publications/blog/ndt-method-overview-visual-testing>
- [27] Firestone, Floyd A. The supersonic reflectoscope, an instrument for inspecting the interior of solid parts by means of sound waves. *Journal of the Acoustical Society of America* 17, no. 3 (1946): 287–299.
<https://doi.org/10.1121/1.1916330>
- [28] Peck, Ralph B. Advantages and limitations of the observational method in applied soil mechanics. *Géotechnique* 19(2) (1969): 171–187.
<https://doi.org/10.1680/geot.1969.19.2.171>
- [29] Farrar, Charles R., and Keith Worden. An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A* 365(1851) (2007): 303–315.
<https://doi.org/10.1098/rsta.2006.1928>
- [30] Randall, Robert Bond. *Vibration-based Condition Monitoring: Industrial, Aerospace and Automotive Applications*. Wiley, 2011.
<https://doi.org/10.1002/9780470977668>
- [31] Phadke, A. G., J. S. Thorp, and M. G. Adamiak. A new measurement technique for tracking voltage phasors, local system frequency, and rate of change of frequency. *IEEE Transactions on Power Apparatus and Systems* PAS-102(5) (1983): 1025–1038.
<https://doi.org/10.1109/TPAS.1983.318043>

- [32] Phadke, Arun G., and James S. Thorp. Synchronized Phasor Measurements and Their Applications. Springer, 2008.
<https://doi.org/10.1007/978-0-387-76537-2>
- [33] Kletz, Trevor. ICI's contribution to process safety. IChemE Symposium Series 155, Hazards XXI, 2009.
<https://www.icheme.org/media/9503/xxi-paper-005.pdf>
- [34] U.S. Chemical Safety and Hazard Investigation Board. Investigation Report: Refinery Explosion and Fire, BP, Texas City, Texas, March 23, 2005. Report 2005-04-I-TX, March 2007.
<https://www.csb.gov/bp-america-texas-city-refinery-explosion/>
- [35] NASA. 50 Years Ago: Apollo 13 Review Board Report. 2020. Account of the investigation and report released on 15 June 1970.
<https://www.nasa.gov/history/50-years-ago-apollo-13-review-board-report/>
- [36] Presidential Commission on the Space Shuttle Challenger Accident. Report to the President. 6 June 1986. NASA archive, especially Volume I, chapters 4–6.
<https://www.nasa.gov/history/rogersrep/>
- [37] Columbia Accident Investigation Board. Report, Volume I. August 2003.
<https://ntrs.nasa.gov/citations/20030066167>
- [38] Australian Transport Safety Bureau. In-flight uncontained engine failure, Airbus A380-842, VH-OQA, overhead Batam Island, Indonesia, on 4 November 2010. Investigation AO-2010-089, final report, 2013.
<https://www.atsb.gov.au/investigations/ao-2010-089>
- [39] Rossman, Lewis A., et al. EPANET 2.2 User Manual. EPA/600/R-20/133. U.S. Environmental Protection Agency, 2020.
https://cfpub.epa.gov/si/si_public_record_Report.cfm?LAB=CESER&dirEntryID=348882
- [40] Leveson, Nancy G., and Clark S. Turner. An investigation of the Therac-25 accidents. Computer 26(7) (1993): 18–41.
<https://doi.org/10.1109/MC.1993.274940>
- [41] Lamport, Leslie. Time, clocks, and the ordering of events in a distributed system. Communications of the ACM 21(7) (1978): 558–565.
<https://www.microsoft.com/en-us/research/publication/time-clocks-ordering-events-distributed-system/>
- [42] Vostokov, Dmitry, and Software Diagnostics Institute. Trace, Log, Text, Narrative, Data: An Analysis Pattern Reference for Information Mining, Diagnostics, Anomaly Detection. Fifth edition. OpenTask, 2023.
<https://www.dumpanalysis.org/trace-log-analysis-pattern-reference>
- [43] NASA. Standard for Models and Simulations. NASA-STD-7009B, 5 March 2024.
<https://standards.nasa.gov/standard/NASA/NASA-STD-7009>
- [44] Glaessgen, Edward H., and David S. Stargel. The digital twin paradigm for future NASA and U.S. Air Force vehicles. AIAA Paper 2012-1818, 2012.
<https://ntrs.nasa.gov/api/citations/20120008178/downloads/20120008178.pdf>
- [45] Canonico, Domenic. The history of ASME's Boiler and Pressure Vessel Code. ASME, 1 December 2010; page updated 31 October 2024.
<https://www.asme.org/topics-resources/content/the-history-of-asmes-boiler-and-pressure>
- [46] NASA. NASA Systems Engineering Handbook. NASA/SP-2016-6105 Rev 2, published 17 February 2017.
<https://ntrs.nasa.gov/citations/20170001761>
- [47] International Electrotechnical Commission. Functional safety of electrical/electronic/programmable electronic safety-related systems—Part 1: General requirements. IEC 61508-1:2010, edition 2.
<https://webstore.iec.ch/en/publication/5515>
- [48] U.S. Nuclear Regulatory Commission. Backgrounder on the Three Mile Island Accident. Institutional historical account. n.d.
<https://www.nrc.gov/regulations-legislation/fact-sheets-brochures/backgrounder-on-the-three-mile-island-accident>
- [49] Bainbridge, Lisanne. Ironies of automation. Automatica 19(6) (1983): 775–779. Developed from a paper presented in 1982.
[https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- [50] Garcia, Mary Lynn. The Design and Evaluation of Physical Protection Systems. First edition. Butterworth-Heinemann, 2001.
<https://shop.elsevier.com/books/the-design-and-evaluation-of-physical-protection-systems/garcia/978-0-08-050391-2>

- [51] Saltzer, Jerome H., and Michael D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE* 63(9) (1975): 1278–1308.
<https://web.mit.edu/Saltzer/www/publications/protection/>
- [52] Anderson, Ross. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Third edition. Wiley, 2020. Author's archive.
<https://www.cl.cam.ac.uk/archive/rja14/book.html>
- [53] Denning, Dorothy E. An intrusion-detection model. *IEEE Transactions on Software Engineering* SE-13(2) (1987): 222–232.
<https://doi.org/10.1109/TSE.1987.232894>
- [54] Kent, Karen, Suzanne Chevalier, Tim Grance, and Hung Dang. *Guide to Integrating Forensic Techniques into Incident Response*. NIST SP 800-86, 2006.
<https://doi.org/10.6028/NIST.SP.800-86>
- [55] Strom, Blake, Andy Applebaum, Douglas Miller, Kathryn Nickels, Adam Pennington, and Cody Thomas. MITRE ATT&CK: Design and Philosophy. MITRE, publication page dated 31 March 2020.
<https://www.mitre.org/news-insights/publication/mitre-attck-design-and-philosophy>
- [56] Industrial Control Systems Cyber Emergency Response Team. Primary Stuxnet Advisory. ICSA-10-272-01, advisory originating in 2010; archived by CISA.
<https://www.cisa.gov/uscert/ics/advisories/ICSA-10-272-01>
- [57] Stouffer, Keith, et al. *Guide to Operational Technology (OT) Security*. NIST SP 800-82 Rev. 3, September 2023.
<https://doi.org/10.6028/NIST.SP.800-82r3>
- [58] Schneier, Bruce. Attack trees. *Dr. Dobbs's Journal*, December 1999. Author's archive.
https://www.schneier.com/academic/archives/1999/12/attack_trees.html
- [59] Souppaya, Murugiah, Karen Scarfone, and Donna Dodson. *Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities*. NIST SP 800-218, 2022.
<https://doi.org/10.6028/NIST.SP.800-218>
- [60] Ross, Ron, Mark Winstead, and Michael McEvilly. *Engineering Trustworthy Secure Systems*. NIST SP 800-160 Vol. 1 Rev. 1, November 2022.
<https://doi.org/10.6028/NIST.SP.800-160v1r1>
- [61] Ross, Ron, Victoria Pillitteri, Richard Graubart, Deborah Bodeau, and Rosalie McQuaid. *Developing Cyber-Resilient Systems: A Systems Security Engineering Approach*. NIST SP 800-160 Vol. 2 Rev. 1, December 2021.
<https://doi.org/10.6028/NIST.SP.800-160v2r1>
- [62] Nelson, Alexander, Sanjay Rekhi, Murugiah Souppaya, and Karen Scarfone. *Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile*. NIST SP 800-61 Rev. 3, April 2025.
<https://doi.org/10.6028/NIST.SP.800-61r3>
- [63] Perrow, Charles. *Normal Accidents: Living with High-Risk Technologies*. Basic Books, 1984; expanded reissue, Princeton University Press, 1999.
<https://press.princeton.edu/books/paperback/9780691004129/normal-accidents>
- [64] Leveson, Nancy G. A new accident model for engineering safer systems. *Safety Science* 42(4) (2004): 237–270.
[https://doi.org/10.1016/S0925-7535\(03\)00047-X](https://doi.org/10.1016/S0925-7535(03)00047-X)
- [65] Leveson, Nancy G. *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, 2012.
<https://mitpress.mit.edu/9780262016629/engineering-a-safer-world/>
- [66] Reason, James. Human error: models and management. *BMJ* 320(7237) (2000): 768–770.
<https://doi.org/10.1136/bmj.320.7237.768>
- [67] Rasmussen, Jens. Risk management in a dynamic society: a modelling problem. *Safety Science* 27(2–3) (1997): 183–213.
[https://doi.org/10.1016/S0925-7535\(97\)00052-0](https://doi.org/10.1016/S0925-7535(97)00052-0)
- [68] Rochlin, Gene I., Todd R. La Porte, and Karlene H. Roberts. The self-designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review* 40(4) (1987), article 7.
<https://digital-commons.usnwc.edu/nwc-review/vol40/iss4/7/>

- [69] Holling, C. S. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics* 4 (1973): 1–23.
<https://doi.org/10.1146/annurev.es.04.110173.000245>
- [70] Hollnagel, Erik, David D. Woods, and Nancy Leveson, eds. *Resilience Engineering: Concepts and Precepts*. Ashgate, 2006.
<https://www.routledge.com/Resilience-Engineering-Concepts-and-Precepts/Hollnagel-Woods-Leveson/p/book/9781315605685>
- [71] Hollnagel, Erik, Jean Parihès, David D. Woods, and John Wreathall, eds. *Resilience Engineering in Practice: A Guidebook*. Ashgate, 2011. See Hollnagel’s prologue, *The Scope of Resilience Engineering*.
<https://www.routledge.com/Resilience-Engineering-in-Practice-A-Guidebook/Paries-Wreathall-Hollnagel/p/book/9781315605692>
- [72] Hollnagel, Erik. *Safety-I and Safety-II: The Past and Future of Safety Management*. Ashgate, 2014. Author’s description.
<https://erikhollnagel.com/ideas/safety-i-and-safety-ii-2014>
- [73] Hollnagel, Erik. *FRAM: The Functional Resonance Analysis Method—Modelling Complex Socio-technical Systems*. Ashgate, 2012.
<https://www.routledge.com/FRAM-The-Functional-Resonance-Analysis-Method-Modelling-complex-socio-technical-systems/Hollnagel/p/book/9781315255071>
- [74] Woods, David D. Four concepts for resilience and the implications for the future of resilience engineering. *Reliability Engineering & System Safety* 141 (2015): 5–9.
<https://doi.org/10.1016/j.res.2015.03.018>
- [75] Woods, David D. The theory of graceful extensibility: Basic rules that govern adaptive systems. *Environment Systems and Decisions* 38(4) (2018): 433–457.
<https://doi.org/10.1007/s10669-018-9708-3>
- [76] NIST. *Community Resilience Planning Guide for Buildings and Infrastructure Systems, Volume I*. NIST SP 1190v1, 2015; cited PDF dated May 2016.
<https://doi.org/10.6028/NIST.SP.1190v1>
- [77] Ewaschuk, Rob. Monitoring distributed systems. In *Site Reliability Engineering: How Google Runs Production Systems*, edited by Betsy Beyer, Chris Jones, Jennifer Petoff, and Niall Richard Murphy. O’Reilly, 2016, chapter 6.
<https://sre.google/sre-book/monitoring-distributed-systems/>
- [78] Lunney, John, and Sue Lueder. Postmortem culture: Learning from failure. In *Site Reliability Engineering: How Google Runs Production Systems*. O’Reilly, 2016, chapter 15.
<https://sre.google/sre-book/postmortem-culture/>
- [79] Izrailevsky, Yury, and Ariel Tseitlin. The Netflix Simian Army. *Netflix Technology Blog*, 19 July 2011.
<https://netflixtechblog.com/the-netflix-simian-army-16e57fbab116>
- [80] Vostokov, Dmitry. *Bridging the Great Divide*. Software Diagnostics Institute, n.d.; also collected in *Theoretical Software Diagnostics*, fourth edition.
<https://www.dumpanalysis.org/the-great-divide>
- [81] NIST. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1, January 2023.
<https://doi.org/10.6028/NIST.AI.100-1>
- [82] Vostokov, Dmitry, concept and direction. *History of Technical Diagnostics, Observability, and Debugging*. Version 7, September 2026. DumpAnalysis.org. Companion history; chapter references in this volume refer to this version.
<https://www.dumpanalysis.org/history-technical-diagnostics-observability-debugging>
- [83] Vostokov, Dmitry, concept and direction. *History of Hardware Diagnostics, Observability, and Debugging*. Version 14, September 2026. DumpAnalysis.org. Companion history; chapter references in this volume refer to this version.
<https://www.dumpanalysis.org/history-hardware-diagnostics-observability-debugging>

