

## *Chapter 8*

---

# Conduct Accident/Incident Investigations

---

This chapter describes the fifth component of an effective information security/IA program — conducting an accident/incident investigation. The process of how to conduct an accident/incident investigation is explained, as well as the reasons why one should conduct an investigation. The following activities are performed while conducting an accident/incident investigation:

- The cause, extent, and consequences of the failure/compromise are analyzed.
- Recovery mechanisms are initiated.
- The accident/incident is reported.
- Remedial measures are deployed.
- Legal issues are evaluated.

There is extensive interaction between this component and the preceding four components, as the following chapter sections demonstrate.

Before proceeding, it is important to clarify terminology. The terms “accident” and “incident,” “failure” and “compromise” are used to mean different things in diverse publications. Occasionally, these terms are even used interchangeably. This book, taking into account both the technical and legal usage of the terms, defines them as follows:

**accident:** (1) technical — any unplanned or unintended event, sequence, or combination of events that results in death, injury, or illness to personnel or damage to or loss of equipment or property (including data, intellectual property, etc.), or damage to the environment<sup>127,422,425</sup>; (2) legal — any unpleasant or unfortunate occurrence that causes injury, loss, suffering, or death; an event that takes place without one’s foresight or expectation.<sup>214</sup>

**incident:** any unplanned or unintended event, sequence, or combination of events that does not result in death, injury, or illness to personnel or damage to or loss of equipment, property (including data, intellectual property, etc.), or damage to the environment, but has the potential to do so.

In short, an accident results in unexpected loss, physical or cyber. The person or entity incurring the loss may or may not be the responsible party; often a second or third party is involved. An accident can result from accidental or intentional action or inaction. Case law distinguishes between avoidable and unavoidable accidents, a point that is particularly relevant when investigating technology-related accidents/incidents. (*Note:* Some standards refer to accidents as mishaps.)

In contrast, an incident is a near-miss that could have resulted in an accident but did not. Incidents often precede accidents as an early warning of a more serious underlying problem; hence, the need to investigate them as well.

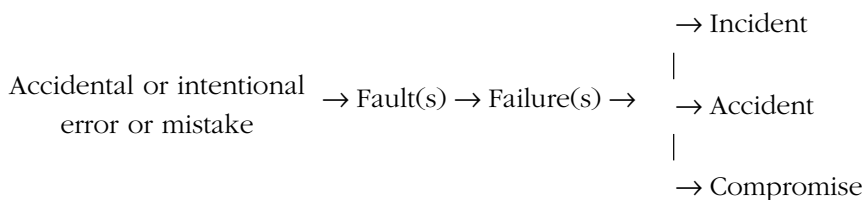
**Failure:** failing to or inability of a system, entity, or component to perform its required function(s), according to specified performance criteria, due to one or more fault conditions.

**Compromise:** an unwarranted and uninvited offensive incursion, infringement, or encroachment of a system, usually by stealth, that defeats safety and security mechanisms to violate and usurp resources and data in a hostile and injurious manner.

A failure implies that a system, entity, or component did not or could not perform its prescribed function(s). Fault tolerance attempts to prevent component and entity failures from becoming system failures. Three categories of failures are commonly recognized: (1) incipient failures are failures that are about to occur; (2) hard failures are failures that result in a complete system shutdown; and (3) soft failures are failures that result in a transition to degraded mode operations or a fail operational status.<sup>44</sup> A failure, particularly of an IA-critical or IA-related function/entity, directly impacts IA integrity.

A compromise represents the digital equivalent of trespassing, infringement, wrongful breaking, entering, and appropriation of data and resources. A system may or may not be rendered inoperable by a compromise; however, sensitive information and/or resources are misappropriated (stolen), usurped, or exposed.

An accident, incident, or compromise is preceded by the failure of one or more safety and security mechanisms (IA design technique/feature, operational procedures, physical security practices, etc.). If one extends the fault tolerance paradigm:



The legal and engineering professions differ on how they define and categorize causes. It is important to be aware of these differences when conducting an accident/incident investigation. The standard engineering definition is<sup>31</sup>:

**Cause:** the action or condition by which a hazardous event (physical or cyber) is initiated — an initiating event. The cause may arise as the result of failure, accidental or intentional human error, design inadequacy, induced or natural operational environment, system configuration, or operational mode(s)/state(s).

The standard legal definition is<sup>214</sup>:

**Cause:** each separate antecedent of an event. Something that precedes and brings about an effect or result. A reason for an accident or condition.

In legal matters, causation is used to assess negligence. [Exhibit 1](#) compares legal and engineering cause categories.

## 8.1 Analyze Cause, Extent, and Consequences of Accident/Incident

Known or suspected accident/incidents should be investigated whenever they occur, whether during the development or operational phases. Accidents/incidents during the development phase are indicative of underlying design deficiencies, misunderstandings about the operation or mission of a system, and incompatible components. Accidents/incidents during the operational phase may result from any of the above causes, as well as deficiencies in the operational environment, operational procedures, physical security practices, and system survivability characteristics.

There are several compelling reasons to investigate accidents/incidents, regardless of their severity, including:

1. To determine in fact what did and did not happen, how it happened, and why it happened or was allowed to happen
2. To ascertain the extent of the consequences and the corresponding need for (immediate) recovery mechanisms, and (long-term) remedial measures
3. To gather the information necessary to file an accurate report of the accident/incident
4. To evaluate legal issues

If an accident/incident is not investigated, useful information is thrown away. The end result is that there are no facts upon which to base immediate recovery mechanisms — recovery efforts will be ineffective or haphazard at best. There are no facts upon which to base long-term remedial measures to ensure that the accident/incident or a similar one does not occur in the future.

## Exhibit 1 Comparison of Legal and Engineering Cause Categories

<i>Legal Category</i>	<i>Engineering Category</i>	<i>Definition</i>
Concurrent cause	No exact engineering equivalent. Concurrent causes might be considered dependent parallel root causes.	Causes acting contemporaneously and together, causing injury that would not have happened in the absence of either. Two distinct causes operating at the same time to produce a given result. <sup>214</sup>
Contributing cause	No exact engineering equivalent. Contributing causes might be considered intermediate causes.	Any factor that contributes to a result, although its causal nexus may not be immediate. <sup>214</sup>
Intervening cause	No exact engineering equivalent. A positive intervening cause may result from an effective threat control measure, defensive layer, or emergency response. A negative intervening cause may result from erroneous human action in response to an accident precursor.	An independent cause that intervenes between the original event and the accident/incident, negates the natural course of events, and produces a different result, positive or negative.
Direct, proximate, or legal cause	Basic, underlying, or root cause	Underlying cause(s), event(s), conditions, or actions that individually or in combination led to the accident/incident; primary precursor event(s) that has (have) the potential for being corrected.
Probable or reasonable cause	No engineering equivalent.	A reasonable ground for belief in certain alleged facts. A set of probabilities grounded in the factual and practical considerations that govern the decisions of reasonable and prudent persons and is more than mere suspicion but less than the quantum of evidence required for conviction. <sup>214</sup>
Remote cause	No engineering equivalent.	A cause which would not according to experience of mankind lead to the event which happened. <sup>214</sup>
No exact legal equivalent	Intermediate cause	An event between the underlying cause and the accident/incident that occurs within the direct chain of events; an epiphenomenon.

Other insights gained from an accident/incident investigation are lost, such as previously unknown latent vulnerabilities/threats. Evidence necessary to pursue legal action is discarded. In summary, the short- and long-term returns on investment from conducting an accident/incident investigation are manifold.

Conducting an accident/incident investigation is a branch of forensic engineering. There is a common misperception that forensic engineering is a mysterious, magical, and occasionally devious endeavor. Not at all. According to Webster's Dictionary, forensic simply means "belonging to, used in, or suitable to courts of justice." Black's Law Dictionary<sup>214</sup> defines forensic engineering as:

*The application of the principles and practice of engineering to the elucidation of questions before courts of law. Practice by legally qualified professional engineers who are experts in their field, by both education and experience, and who have experience in the courts and an understanding of jurisprudence. A forensic engineering engagement may require investigations, studies, evaluations, advice to counsels, reports, advisory opinions, depositions, and/or testimony (expert witness) to assist in the resolution of disputes relating to life or property in cases before courts or other lawful tribunals.*

Accident/incident investigations may be conducted solely for internal purposes, as part of a regulatory process, to share information within or among industrial sectors, or pursuant to legal action. A forensic accident/incident investigation adds the notion that evidence is collected, organized, analyzed, and presented in a manner that is appropriate for a court of law. A regulatory accident/incident investigation is more formal than one conducted solely for internal purposes. Likewise, a forensic accident/incident investigation is more formal than a regulatory one. However, the techniques and methods used in all three are the same.

Petroski<sup>380</sup> explains that conducting a forensic accident/incident investigation is equivalent to performing failure analysis, something with which safety and reliability engineers have extensive experience. Essentially, one determines *what* happened, *how* it happened, *why* it happened or was allowed to happen, and the resultant consequences. The process is similar to performing a digital autopsy; however, in this case, it is necessary to look beyond the "patient" to locate all the contributing factors and sources of accidental and intentional errors, mistakes, faults, and failures that led to the accident/incident. An investigation involves synthesizing scenarios that describe "how could" an accident/incident occur with scenarios that depict "how did" an accident/incident occur through inductive and deductive reasoning. The credibility of an investigation rests on the ability to remain objective, eliminate bias or prejudice, separate fact, opinion, assumptions and theory, and distinguish "symptoms" from the "disease" — all while being thorough and accurate. As Petroski<sup>380</sup> states:

*Although some of the acute interest in accident postmortems no doubt stems from legal and insurance claims, there is considerable engineering experience to be gained in understanding exactly what caused a failure. ...they are necessary and necessarily drawn out because they can involve a painstaking sifting and analysis of clues as subtle as Sherlock Holmes ever had to deal with.*

Today, is it unlikely that during the life of a system an accident/incident will not be experienced. Prior preparation facilitates effective and efficient accident/incident investigations, whether internal, independent, regulatory, or forensic. Poe<sup>385</sup> recommends planning, coordination, and training to address the following issues:

- The speed, accuracy, and completeness of information collection
- Reporting channels and responsibilities, inside and outside the team
- Standardized report forms (draft and final)
- Designated participant lists, per accident scenario, to ensure an interdisciplinary team
- A generic checklist of questions to ask, per accident scenario, to stimulate avenues of investigation
- A fixed chain of custody for evidence, given the fleeting nature of digital evidence
- Procedures for obtaining consent of witnesses prior to conducting critical incident interviews

There are plenty of open sources to examine when collecting evidence for an accident/incident investigation. There is no need to resort to clandestine methods. Analyses of early evidence point to other sources of primary, secondary, and tertiary evidence, in what becomes an iterative process. Digital evidence is supplemented by critical incident interviews. As the evidence accumulates, it is sifted (relevant/irrelevant evidence), validated, and organized (direct/indirect evidence, time sequence, etc.) until a clear picture begins to emerge of what exactly happened. Inductive and deductive reasoning are applied to the evidence to explain how and why the accident/incident occurred. Evidence sources include those that are unique to an accident/incident and those that are common to all accidents/incidents. A sample list of generic evidence sources is given in [Exhibit 2](#). This list is illustrative, not exhaustive. An examination of generic evidence sources often uncovers unique evidence sources; as a result, most investigations begin with generic sources.

Ladkin<sup>321</sup> describes the main tasks of an accident/incident investigation, a process he refers to as “why — because analysis”, as:

1. A general formal definition of causal influence
2. Precise specifications of system and entity behavior
3. Evidence collection and analysis
4. A method of tracing more evidence from the evidence in hand
5. A method of validating the causal reasoning

## Exhibit 2 Generic Accident/Incident Evidence Sources

---

1. Background information
  - System definition
  - System operation characterization
  - System entity control analysis
  - Vulnerability and threat characterizations
  - Transaction paths, critical threat zones
  - IA design techniques/features
  - Threat control effectiveness assessment and summary
  - System design and verification data
  - IA integrity case (assumptions, claims, evidence)
  - Operational procedures, contingency plans, physical security practices
  - System software inventory and configuration (workstation and server); identify what is COTS, custom, authorized, not authorized to be installed
  - Serial numbers of commercial products
  - Authentication parameter file and implementation logic, current and archived
  - Access control rules, definition, and implementation logic, current and archived
2. Previous experiences and observations
  - Prior incident/anomalous activity reports
  - Log of recent preventive, adaptive, and corrective maintenance actions, including system reconfiguration, enhancements, and upgrades
  - History of SPRs, STRs, ECRs, and help desk calls
  - Recent backups
3. Accident/incident characteristics
  - System failure mode, state, and characteristics
  - Chronology of system modes/states and conditions leading up to and including the accident/incident and its aftermath
  - Operational profile, mission, and environment up to and including the accident/incident and its aftermath
  - Names/IDs of all active users and processes
  - System and entity loading characteristics
  - Printouts of anomalous events
  - Message, process, and file header IDs and tags
  - Logical and physical addresses of network nodes, system resources, users, message traffic
  - Message traffic logs, routing tables, e-mail directories, address books
  - Memory, OS, register, buffer, hard disk dumps (server and workstation)
  - System audit trails
  - Keystroke logs
  - Browser screens
  - Printouts of screen freezes (system administrator and end user)
  - Physical security logs, video and audio surveillance tapes
  - Critical incident interviews

---

The causal explanation, evidence, and reasoning must be correct and sufficient; all three must undergo competent validation “so that other people can tell they are correct.”<sup>321</sup> Adequate attention must be paid to the operational environment, operational procedures, and the physical and digital characteristics of the accident. Ladkin and Philley both caution against searching for a single underlying cause of an accident, especially early in an investigation. As Ladkin<sup>321</sup> states:

*In distinction to a common supposition, most complex system accidents are dependent on many factors, not just a single causal chain.*

Philley<sup>381</sup> adds that:

*Isolating one cause as “the” root cause may leave other potential hazards in an uncorrected condition. The investigator should continue to identify, examine, and evaluate all underlying causes.*

Philley notes that identification of a single root cause may cause an investigation to be stopped prematurely and result in over simplistic and incomplete reasoning about the evidence. Similar to Poe,<sup>385</sup> Philley<sup>381</sup> recommends having an interdisciplinary team conduct and validate an investigation.

Intermediate and root causes are extrapolated from the available evidence. This extrapolation, in turn, points to gaps in the evidence and the need for further investigation. The identification of causes should explore all potential initiating and intermediate events, including:

- Random or time-dependent hardware failures<sup>31</sup>
- Systematic or time-dependent software failures
- Latent vulnerabilities/threats
- Accidental or intentional operator error
- Natural and induced environmental effects<sup>31</sup>
- Deficiencies in operational procedures, contingency plans, and physical security practices
- Accidental or intentional design inadequacies:
  - Ineffective defensive layers and threat control measures
  - Inadequate safety and security margins
  - Unintended operating modes caused by sneak circuits<sup>31</sup>
  - Material inadequacies and incompatibilities<sup>31</sup>
  - Erroneous hardware/software interaction<sup>31</sup>
- Natural or induced transient faults
- Inadvertent operation of IA-critical function
- Accidental or intentional modification of interrupt table<sup>277</sup>
- Accidental or intentional redirection of pointers<sup>277</sup>

A combination of techniques is used to investigate accidents/incidents. [Exhibit 3](#) lists eight proven IA accident/incident investigation techniques. A description of each technique is provided in Annex B, which discusses the purpose, benefits, and limitations of each technique and provides pointers to references for further information.

In addition, several IA analysis and verification techniques can be used during an accident/incident investigation, including BBNs, cause consequence analysis, event tree analysis, HAZOP studies, Petri nets, FMECA, FTA, sneak circuit analysis, and root cause analysis. These techniques are useful for identifying and distinguishing “how did” accident scenarios and “how could” accident scenarios. The IA integrity case is reviewed to determine which assumptions, claims, and evidence were false. [Exhibit 4](#) lists the accident/incident investigation role played by each of these techniques.

**Exhibit 3    IA Accident/Incident Investigation Techniques**

<i>IA Accident/Incident Investigation Techniques</i>	<i>C/R</i>	<i>Type</i>	<i>Life-Cycle Phase in which Technique is Used</i>		
			<i>Concept</i>	<i>Development</i>	<i>Operations</i>
Barrier analysis <sup>a</sup>	C4	SA, SE		x	x
Critical incident interviews	C4	SA, SE		x	x
Damage mode effects analysis <sup>a</sup>	C4	SA, SE		x	x
Event and causal factor charting	R4/C4	SA, SE		x	x
Scenario analysis	C4	SA, SE		x	x
Sequentially timed event plot (STEP) investigation system	R4/C4	SA, SE		x	x
Time/loss analysis (TLA) for emergency response evaluation	C4	SA, SE			x
Warning time analysis	C4	SA, SE			x

<sup>a</sup> These techniques can also be used during verification.

**Legend for Exhibit 3**

<i>Column</i>	<i>Code</i>	<i>Meaning</i>
Type	SA	Technique primarily supports safety engineering
	SE	Technique primarily supports security engineering
	RE	Technique primarily supports reliability engineering
	All	Technique supports a combination of safety, security, and reliability engineering
C/R	Cx	Groups of complementary techniques
	Rx	Groups of redundant techniques; only one of the redundant techniques should be used

Next, the accident/incident investigation techniques are discussed in detail. There is a high degree of interaction and interdependence between the techniques: the output of one technique is used as the input to another technique and the techniques complement or reinforce each other.

***Barrier Analysis***

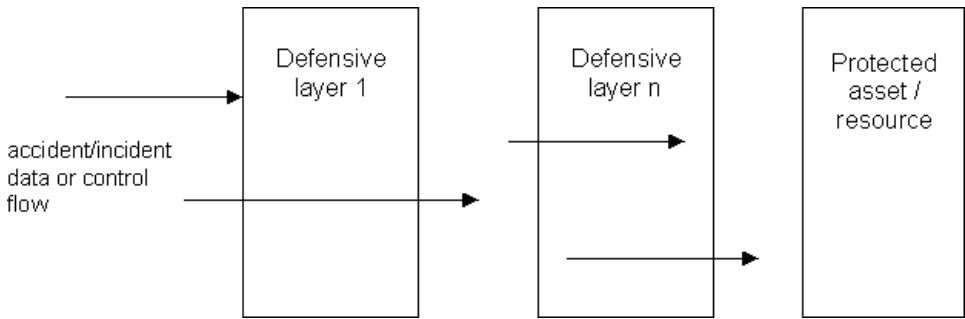
Barrier analysis is used during an investigation to ascertain which defensive layers failed or were missing or inadequate. Barrier analysis helps to determine accident/incident causation by examining each defense in depth layer (or barrier) for accidental or intentional unwanted control, data, or information flow, as illustrated in [Exhibit 5](#). Hazardous control and information flows to/from people and processes are uncovered and how they penetrated or bypassed existing defensive layers is determined. Defensive layers that failed or were missing or inadequate are identified, as well as those that did not fail. As a result, the need for new or modified defensive layers is highlighted. In practice, IA design techniques/features are referred to as “hard barriers,”

## Exhibit 4 Accident/Incident Investigation Role of IA Techniques

<i>Technique</i>	<i>Accident/Incident Investigation Role</i>
<b><i>I. Investigation Techniques</i></b>	
Barrier analysis	Determine which defensive layers failed or were missing or inadequate during an accident/incident.
Critical incident interviews	Collect evidence about an accident/incident and previous related mistakes, anomalies, and near-misses from operational personnel.
Damage mode effects analysis	Postulate which specific threat mechanisms caused an accident/incident from an analysis of the damage modes.
Event and causal factor charting	Graphically reconstruct the events, immediate, intermediate, and root cause(s) of an accident/incident.
Scenario analysis	Develop avenues to investigate from causation theories and hypothetical event chains.
Sequentially timed event plot (STEP) investigation system	Expound a diagram of linked, sequentially timed events and their causal relationships that demonstrates how an accident/incident occurred.
Time/loss analysis (TLA) for emergency response evaluation	Evaluate the: (1) effect of human intervention following an accident/incident, (2) controllability of an accident/incident, and (3) effectiveness of mitigating threat control measures over time.
Warning time analysis	Investigate the delta between the available and actual response times (human and automatic) to an accident/incident and the contributing factors, such as erroneous, unforeseen, or unnecessary delays.
<b><i>II. Analysis Techniques</i></b>	
Bayesian Belief networks (BBNs)	Provide a methodology for reasoning about uncertainty as part of an accident/incident investigation.
Cause consequence analysis	Identify inappropriate, ineffective, and missing threat control measures; verify that all accidental and intentional failure modes had a corresponding threat control measure.
Event tree analysis	Identify inappropriate, ineffective, and missing threat control measures.
HAZOP study	Verify that all accidental and intentional, physical and cyber, hazards associated with the operation of a system had been eliminated or mitigated.
Petri nets	Verify that deadlock, race, and nondeterministic conditions that could cause a system compromise or failure did not exist.
Software, system FMECA	Examine the effect of accidental and intentional, random and systematic failures on system behavior in general and IA integrity in particular.
Software, system FTA	Identify potential root cause(s) of undesired system events (accidental and intentional) to verify the effectiveness of mitigating design features and operational procedures.
Sneak circuit analysis	Verify that all hidden, unintended, and unauthorized hardware and software logical paths or control sequences that could inhibit desired system functions, initiate undesired system events, or cause incorrect timing and sequencing had been removed.

**Exhibit 4    Accident/Incident Investigation Role of IA Techniques (continued)**

<i>Technique</i>	<i>Accident/Incident Investigation Role</i>
<b>III. Verification Techniques</b>	
Control flow analysis	Uncover poor and incorrect program logic structures that could have compromised IA integrity.
Data or information flow analysis	Uncover incorrect and unauthorized data transformations and operations that could have compromised IA integrity.
Review IA integrity case	Determine if the claims made about IA integrity were justified by the supporting arguments and evidence.
Root cause analysis	Identify the underlying cause(s), event(s), conditions, or actions that individually or in combination led to an accident/incident; determine why the defect was not detected earlier.



**Exhibit 5    Barrier Analysis Concept**

while operational procedures and physical security measures are referred to as “soft barriers.” Barrier analysis does not evaluate an entire system, only the defensive layers.

Observations and recommendations are recorded in a barrier analysis report, as shown in [Exhibit 6](#). Existing threat control measures are listed in Part I, along with their intended defensive function. The location of each threat control measure (the TCP/IP or ISO OSI layer and execution point) and its type (anticipate/prevent, detect/characterize, respond/recover) is cited. Next, the accident/incident status is recorded, indicating whether the layer was effective, partially effective, or failed. Part II of the report identifies new defensive layers that are needed. Each recommended new threat control measure is listed, along with the defensive function it will serve, the implementation location, and type. An explanation of the defensive layer this new measure is replacing or reinforcing is provided, with a supporting rationale.

***Critical Incident Interviews***

Critical incident interviews are conducted to collect evidence from operational personnel about an accident/incident and previous related mistakes, anomalies,

Exhibit 6    Barrier Analysis Report

Barrier Analysis Report for:\_\_\_\_\_

as of date:\_\_\_\_\_

I. Existing Defensive Layers

Threat Control Measure	Function	Location <sup>a</sup>	Type <sup>b</sup>	Accident/Incident Status			Remarks
				Effective	Partially Effective	Failed	

II. New Defensive Layers Needed

Threat Control Measure	Function	Location <sup>a</sup>	Type <sup>b</sup>	Defensive Layer Being Replaced or Reinforced	Rationale

<sup>a</sup> TCP/IP or ISO OSI layer and execution point.  
<sup>b</sup> Anticipate/prevent, detect/characterize, or respond/recover.

and near-misses. Key personnel with first-hand experience in developing, using, administering, and maintaining the system that failed or was compromised are interviewed. The interview focuses on experience with or observations about the system immediately before and during the accident/incident and mistakes, anomalies, and near-misses experienced or observed in the past. Operator actions, system modes/states, conditions, functions, malfunctions, etc. are discussed. Printouts, server and workstation OS and memory dumps, audit trails, test results, network and system configuration reports, and such are collected to support verbal accounts. This information is analyzed to expose potential immediate, intermediate, and chronic accident/incident precursors.

People closest to and with the most experience using a system have invaluable insights that other people do not and that may not be readily apparent from technical evidence alone. They also help ensure accurate interpretations of events.

Interviewers need to be careful to separate fact from opinion, subjective from objective data. Interviews must be conducted in an open, positive environment so that witnesses do not feel threatened, intimidated, coerced, or fearful of employment-related retaliation.

## ***Damage Mode Effects Analysis***

Damage mode effects analysis is a deductive technique that provides an early assessment of survivability and the effect of an accident/incident on a system's mission/operation. Damage mode effects analysis is an extension of an FMECA. It examines the damage mode for each IA-critical and IA-related function, entity, component, and subcomponent, specifically<sup>425</sup>:

- The type of damage experienced
- Primary, secondary, and cascading damage effects on this and other functions, entities, and systems
- Variation in the damage mode by operational mode/state, profile, and mission
- The local, next higher level, and end effect(s) of the damage

The damage modes are analyzed to postulate which specific threat mechanisms caused an accident/incident. The survivability assessment provides essential input to recovery efforts and often exposes latent vulnerabilities. The effectiveness of this technique is proportional to the ability to analyze damage modes immediately during or after an accident/incident. If legal action is pursued as the result of an accident/incident, a damage mode effects analysis must be performed.

## ***Event and Causal Factor Charting***

Event and causal factor charts depict a detailed sequence of facts and activities that led to an accident/incident. The right-most block on the chart is the primary event — the accident/incident. The immediate cause is shown in the next block, on the left parallel to the first block. Reasons that permitted or contributed to the immediate causes are listed underneath. This process is continued backward to the underlying root cause(s)/event(s). Unknown events are shown as gaps (?) in the diagram and highlight areas needing further investigation. Causes are categorized as human or system actions. Cascading and lateral events are recorded as well so that all pertinent avenues of investigation are explored.

Event and causal factor charts summarize what is known and unknown about an accident/incident in a format that is easily understood by all stakeholders. The sequential nature of the charts facilitates an unfolding investigation. Arrows connecting cause and event blocks represent potential primary and secondary prevention points; this information can be used to reinforce defensive layers. Event and causal factor charts do not capture the exact timing of events. As a reminder<sup>425</sup>:

*... care must be taken not to limit analysis to merely addressing the symptoms of a problem. The symptoms are sometimes causes in themselves; however, they are often only indications that other factors must be pursued to find the underlying causes.*

[Exhibit 7](#) depicts an event and causal factor chart. The accident/incident in this example is a patient dying after a radiation therapy session. This example was chosen because it (1) illustrates the interaction between safety and security engineering, and (2) highlights the need for security engineering beyond the commercial, defense, and intelligence domains. The same accident/incident scenario is used in [Exhibits 9](#) through [11](#) to demonstrate the similarities and differences between event and causal factor charts and STEP diagrams.

## **Scenario Analysis**

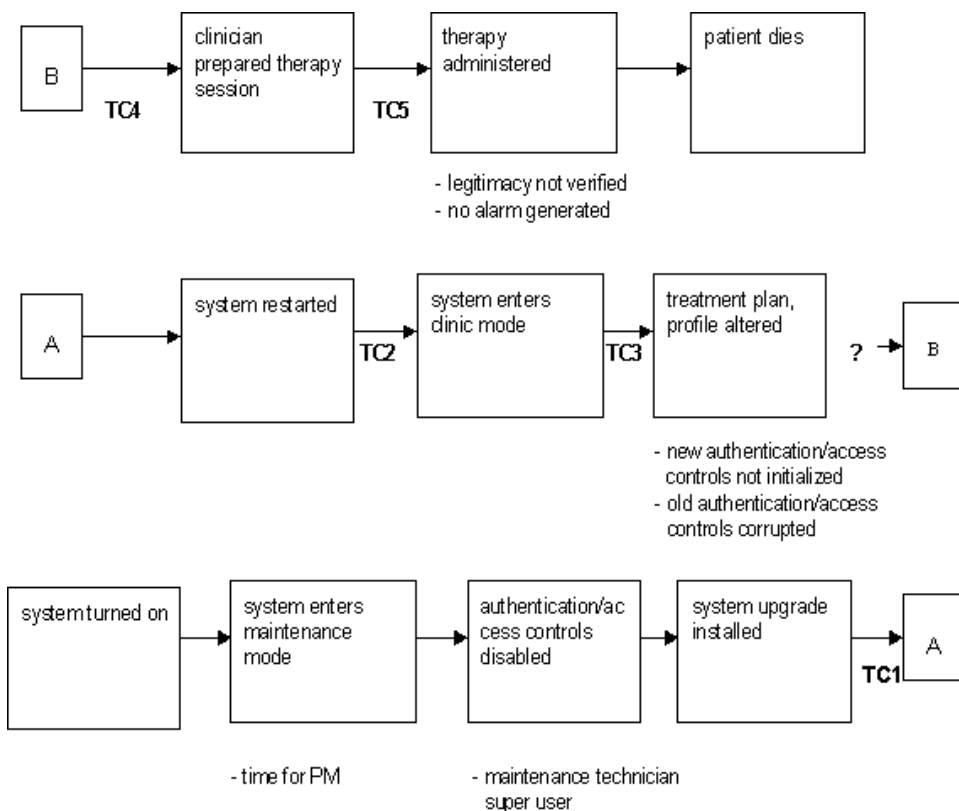
Scenario analysis is conducted to develop avenues to investigate from causation theories and hypothetical event chains. During scenario analysis, all system entities, components, operational profiles, modes/states, environment, and missions and operator actions are examined by an interdisciplinary team. This team, under the guidance of a neutral facilitator, attempts to surmise all possible, credible, and logical scenarios that could have caused or contributed to an accident/incident. The starting point for the team is the fact that the accident/incident occurred. They do not examine evidence; rather, they develop causation theories and hypothetical event chains, based on experience and inductive reasoning, that become avenues to investigate. Scenario analysis, because it is not dependent on extensive evidence, is particularly well suited for investigating novel accidents/incidents for which little or no historical data exists.<sup>425</sup> Successful scenario analysis is dependent on an accurate understanding of the system that failed or was compromised, without letting that knowledge constrain visualization of potential threat scenarios<sup>425</sup>:

*An unfettered mind and an active imagination lead to mastery [of this technique]. ... It can be argued that over-familiarity with the system under analysis restricts the freedom of thought processes necessary to successful application.*

(Note: Do not confuse this technique with formal scenario analysis discussed in Section B.2 and Chapter 6.)

## **Sequentially Timed Event Plot (STEP) Investigation System**

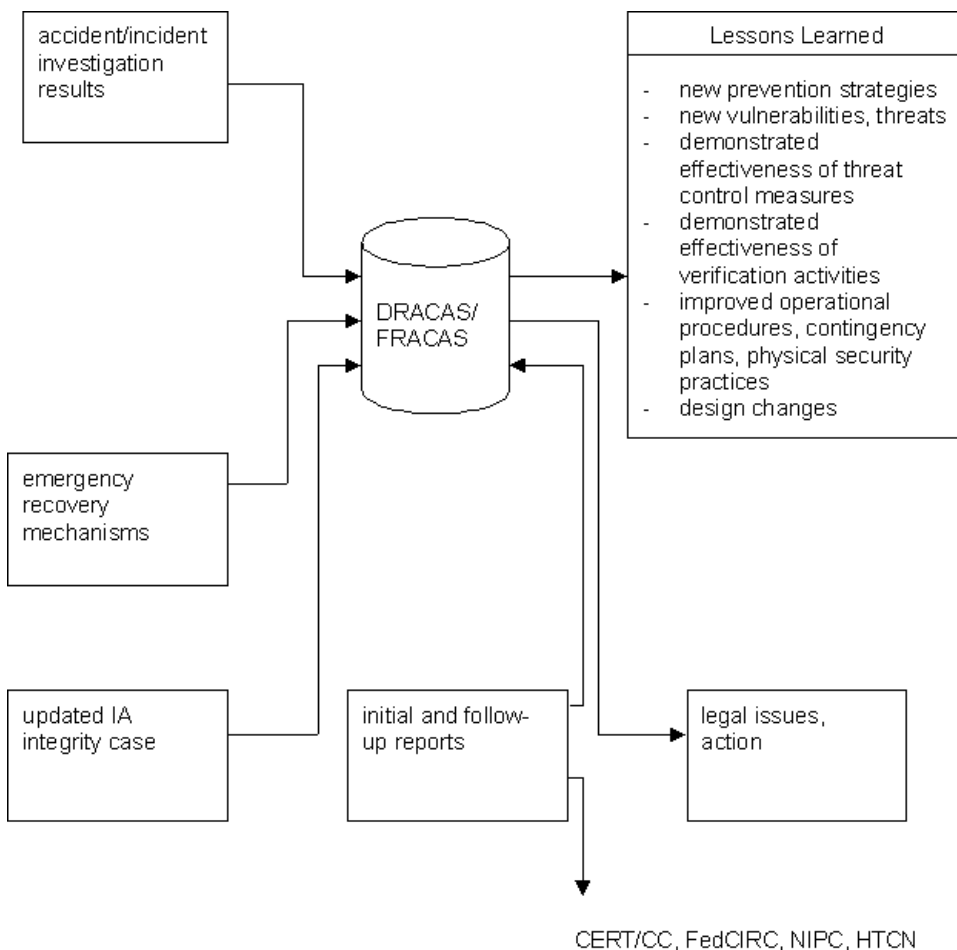
The purpose of the sequentially timed event plot (STEP) investigation system is to expound a diagram of linked, sequentially timed events and their causal relationships, which demonstrates how an accident/incident occurred. The STEP investigation system is an analytical methodology that develops accident process descriptions. A diagram of sequentially timed, multi-linear events depicts acci-



Threat control barriers	
TC1	System is restarted with upgrade. Maintenance technician should have verified that all safety and security features worked correctly, including authentication and access controls.
TC2, TC3	System should have generated an alarm indicating that the new authentication and access controls had not been initialized and that the old authentication and access controls were corrupted. Alarm status should have prohibited access to patient records database and treatment planning system.
TC4	Clinician fails to notice or report that system does not perform authentication.
TC5	Clinician should have verified (manually) legitimacy of treatment profile before initiating therapy. System should have automatically verified legitimacy of treatment profile before initiating therapy. Alarm status should have prohibited therapy session.

## Exhibit 7 Event and Causal Factor Chart

dent/incident causal relationships. Direct, converging, and diverging relationships of immediate, intermediate, and underlying events are illustrated. STEP diagrams visually display the sequence and timing aspects of accident/incident precursors. The event chain necessary to produce the accident/incident outcome is linked together; accident data is transformed into event building blocks.<sup>425</sup> Uncertainties or gaps in the event chain are highlighted for further

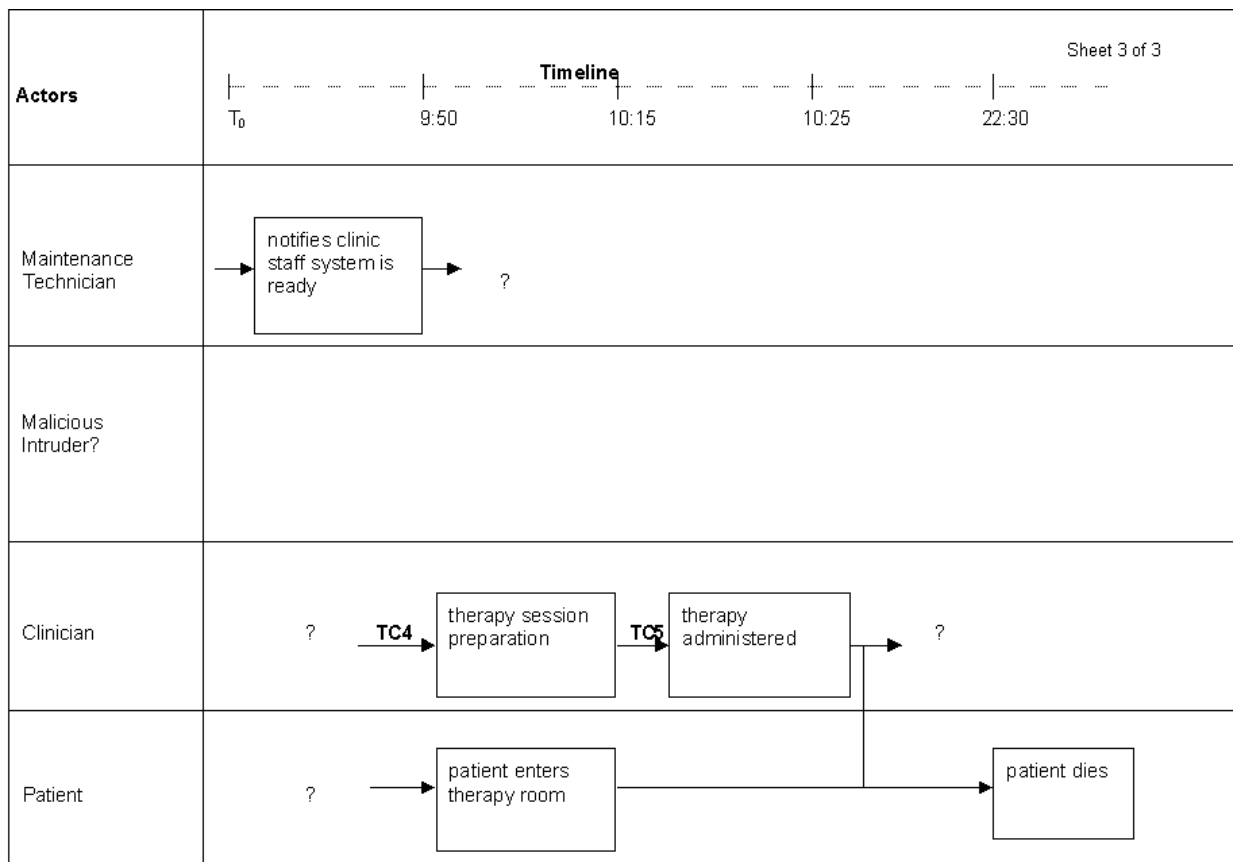


## Exhibit 8 Standard STEP Investigation System Symbols and Notation

investigation. Standard symbols and notation are used to develop a STEP diagram, as shown in [Exhibit 8](#).

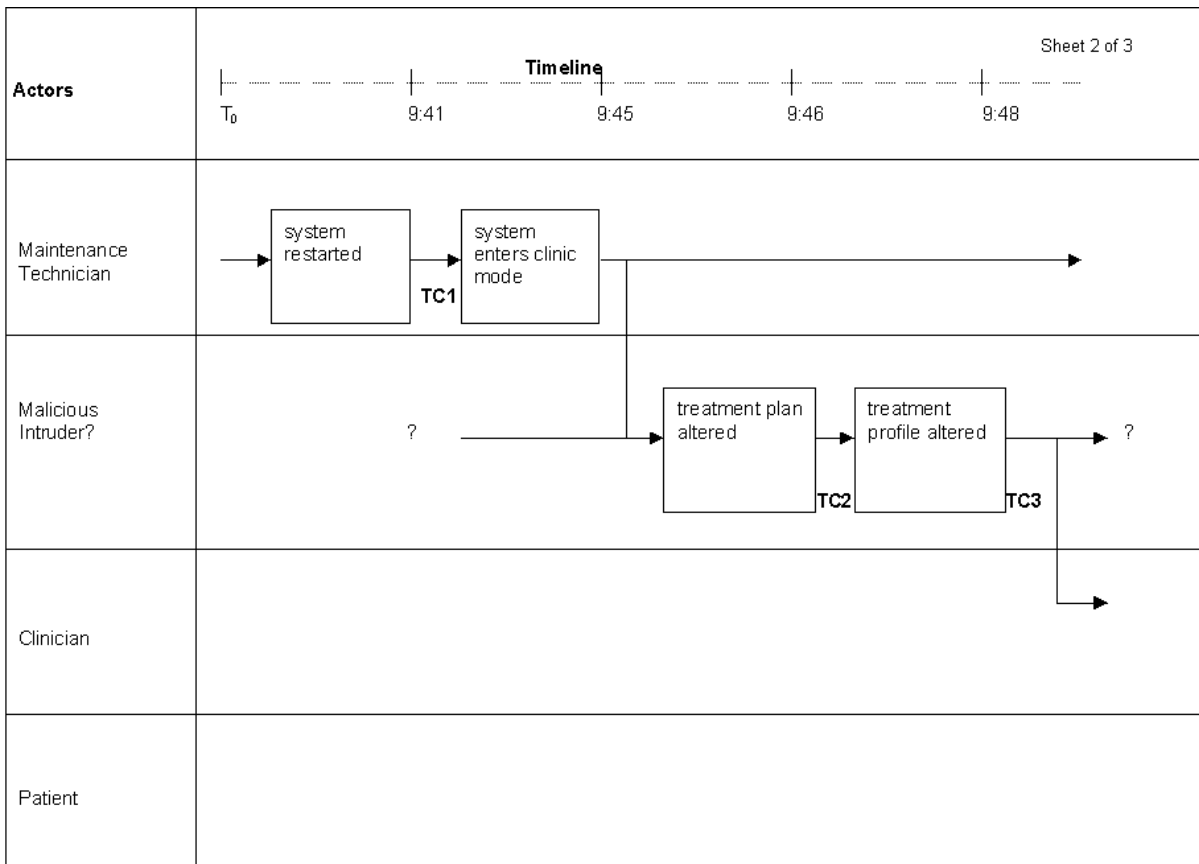
The STEP investigation system supports an in-depth, thorough, and focused analysis of an accident/incident. STEP diagrams are easy to understand; consequently, they can be reviewed and verified by multiple stakeholders. An unlimited number of logical possibilities (accidental/intentional, human/computer action) can be investigated.<sup>425</sup> STEP diagrams expose misunderstandings about how a system “should” versus “does” operate and deficiencies in operational procedures, contingency plans, and physical security practices. A skilled facilitator is needed to keep the analysis proceeding at a level that is meaningful and relevant to the investigation. The analysis should not be at too high or too low a level.

[Exhibits 9](#) through [11](#) are a STEP diagram of the radiation therapy session accident/incident introduced in [Exhibit 7](#). Note that a STEP diagram captures more detail than an event and causal factor chart. Events are associated with actors and timestamps. Relationships between actors and events are recorded.



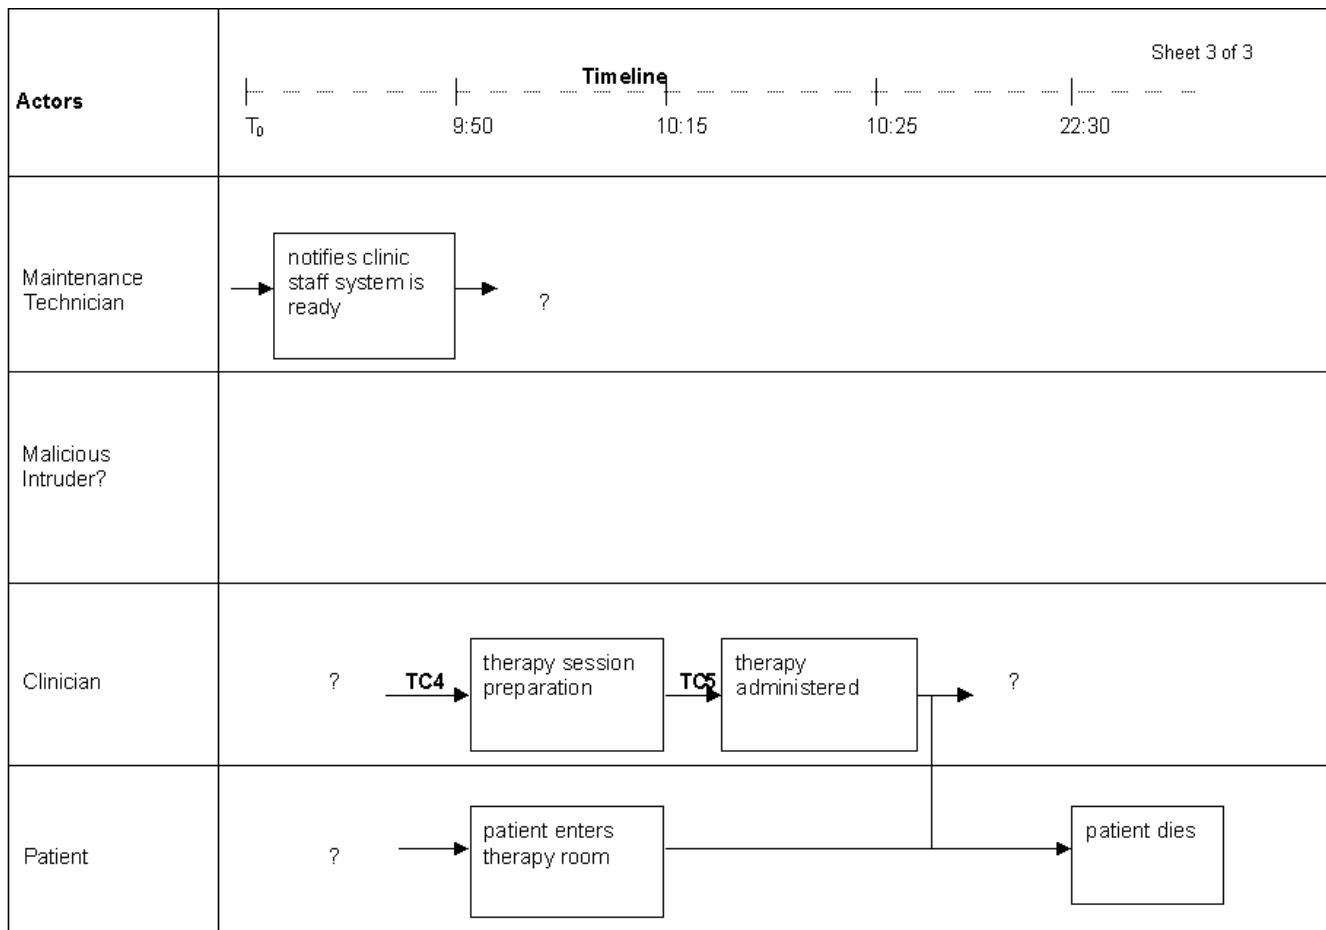
TC: threat control barriers.

## Exhibit 9 STEP Investigation Diagram



TC: threat control barriers.

## Exhibit 10 STEP Investigation Diagram (continued)



TC: threat control barriers.

## Exhibit 11 STEP Investigation Diagram (continued)

Legend for Exhibits 9 through 11

Threat Control Barriers

TC1	System is restarted with upgrade. Maintenance technician should have verified that all safety and security features worked correctly, including authentication and access controls.
TC2, TC3	System should have generated an alarm indicating that the new authentication and access controls had not been initialized and that the old authentication and access controls were corrupted. Alarm status should have prohibited access to patient records database and treatment planning system.
TC4	Clinician fails to notice or report that system does not perform authentication.
TC5	Clinician should have verified (manually) legitimacy of treatment profile before initiating therapy. System should have automatically verified legitimacy of treatment profile before initiating therapy. Alarm status should have prohibited therapy session.

Barrier analysis is performed to determine why defensive layers failed and the accident/incident was allowed to progress. Scenario analysis and critical incident interviews are conducted to resolve gaps in the STEP diagram. Pieces of the STEP diagram are filled in as an investigation unfolds. It is often useful to develop an event and causal factor chart as draft input to a STEP diagram.

An event and causal factor chart illustrates what happened during an accident/incident. A STEP diagram explains what happened and how. An accident/incident investigation must also answer the question of why did it or why was it allowed to happen. For example, looking at Exhibits 9 through 11, the following questions arise:

1. Who ordered and approved the system upgrade? Who knew about it?
2. Who altered the treatment plan and profile? How did they know about the system upgrade?
3. Was the maintenance technician collaborating with the malicious intruder or just negligent?
4. Was the clinician collaborating with the malicious intruder or just negligent?
5. Is negligence always accidental or can it be intentional?
6. What do physical security logs show?
7. Why was the system audit trail not archived before the upgrade and restarted afterward?
8. Is the malicious intruder a person or a process?
9. Who should the patient's family sue?
10. How many other patients were affected?

**Time/Loss Analysis (TLA) for Emergency Response Evaluation**

TLA “defines and organizes data needed to assess the objectives, progress, and outcome of an emergency response” to an accident/incident.<sup>425</sup> TLA serves

several purposes. It evaluates the (1) effect of human intervention following an accident/incident, (2) controllability of an accident/incident, and (3) effectiveness of mitigating threat control measures over time. The results of TLA are recorded in TLA graphs.

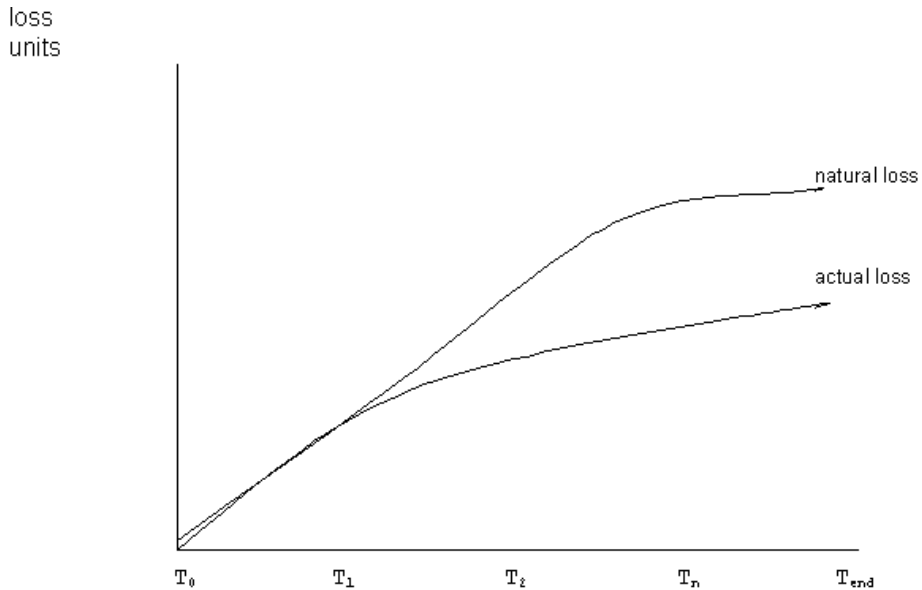
TLA graphs measure and compare actual versus natural loss following an accident/incident. Intervention data is recorded at vertical points on the  $x$ -axis time line. Loss units (number of fatalities or injuries, property damage, financial loss, loss of productivity, environmental damage, etc.) are recorded on the  $y$ -axis.  $T_0$  is when the accident/incident commences.  $T_{\text{end}}$  correlates to the time of the last directly related loss. The natural loss curve is estimated over time given no human intervention. The actual loss curve plots the sequential effect of each intervening action  $T_n$ . The slope between  $T_0$  and  $T_1$  is the same for both curves and represents the effectiveness of automatic mitigating (detect/characterize, respond/recover) threat control measures over time. The delta between the actual and natural loss curves from  $T_1$  on is a function of the controllability of the accident/incident and the value of human intervention. The general shape of the curves is more important than precise data points.<sup>425</sup> TLA graphs can also be used to analyze alternative hypothetical intervention strategies<sup>425</sup> and contingency plans. Criteria for measuring loss units must be standardized and objective. TLA must be performed, or at least begun, promptly after an accident/incident because the evidence tends to dissipate.

[Exhibit 12](#) present TLA graphs for four different accident/incident emergency response scenarios. The first graph ([Exhibit 12a](#)) illustrates the TLA for a single system in which human intervention was effective and lowered the total loss experienced. [Exhibit 12b](#) illustrates the TLA for a single system in which human intervention was ineffective and actually increased the total loss experienced. There are several possible explanations for this, including:

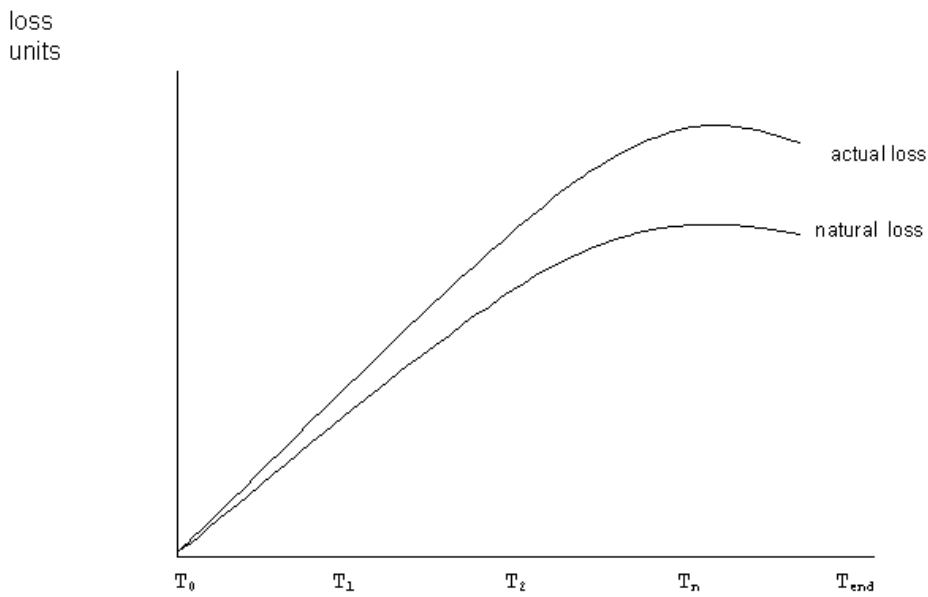
1. The situation was misdiagnosed and the wrong corrective action was applied.
2. The operational procedures or contingency plans were in error in describing how to respond to this situation.
3. A critical step was omitted or performed in the wrong sequence during the emergency response.
4. The emergency response was applied after the interval during which it could have been effective.
5. The operational procedures or contingency plans were deficient; they did not cover this situation and operational personnel guessed how to respond.
6. The operational procedures and contingency plans were correct, but operational personnel had no training or familiarity with them.

The reason for counterproductive human emergency response will be significant if legal action is taken subsequent to an accident/incident.

[Exhibit 12c](#) illustrates the TLA across multiple parallel systems in which human intervention was effective and lowered the actual loss experienced by all three systems. This example depicts the results of accident/incident propagation across



a. single system, intervention effective

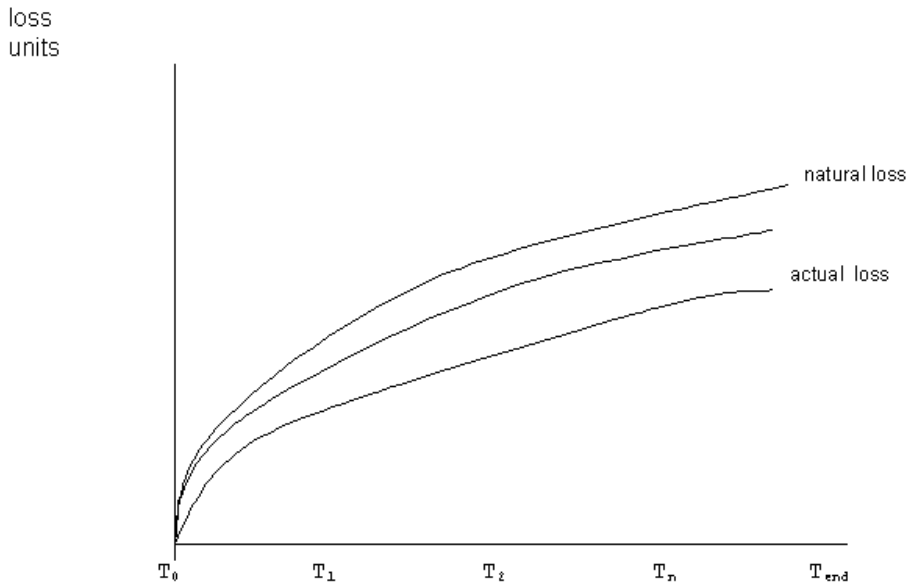


b. single system, intervention counterproductive

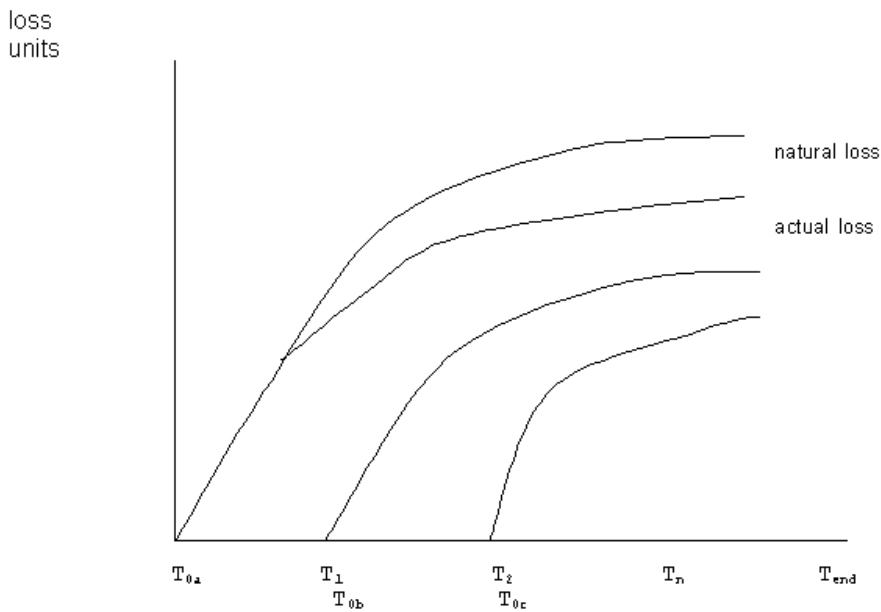
## Exhibit 12 TLA Graphs

parallel systems, usually within one organization. The time taken to respond to the second, third, and  $n^{\text{th}}$  system is longer, due to notification/response logistics, but the actual loss is still less than the natural loss.

[Exhibit 12d](#) illustrates the TLA across multiple cascading systems in which human intervention was effective and lowered the actual loss experienced by all three systems. This example depicts the results of accident/incident



c. multiple systems, intervention effective



c. multiple cascading systems, intervention effective

## Exhibit 12 TLA Graphs (continued)

propagation across cascading systems belonging to multiple organizations, analogous to accidents/incidents that propagate via the Internet. If accident/incident data is shared and reported quickly (see chapter section “Report Accident/Incident”), the time taken to respond to the second, third, and  $n^{\text{th}}$  system is less, lowering the actual loss.

Another interesting avenue to investigate is the  $T_0$  to  $T_1$  curve. If TLA and barrier analysis are combined, points on the curve can be identified to indicate the interval during which each defensive layer was effective before it failed. In summary, TLA can be used to evaluate a variety of different scenarios.

## **Warning Time Analysis**

Warning time analysis investigates the delta between the available and actual response times (human and automatic) to an accident/incident and the contributing factors, such as erroneous, unforeseen, or unnecessary delays. Warning time analysis examines various intervals along the time line from when an accident/incident occurred and recovery mechanisms were initiated. Specific intervals scrutinized include<sup>31</sup>:

- **Propagation time:** time from occurrence of initiating event to time when accident/incident occurred
- **Detection time:** time from occurrence of initiating event to earliest indication or alarm
- **Response time<sub>A</sub>:** time for automatic corrective action
- **Response time<sub>H</sub>:** time for human-initiated corrective action

Warning time analysis evaluates the effectiveness of anomaly detection, the time available for a response, and the adequacy of emergency operational procedures and contingency plans, especially when system reconfiguration or redundant switchover was needed. A comparison between the available and actual response times is made.

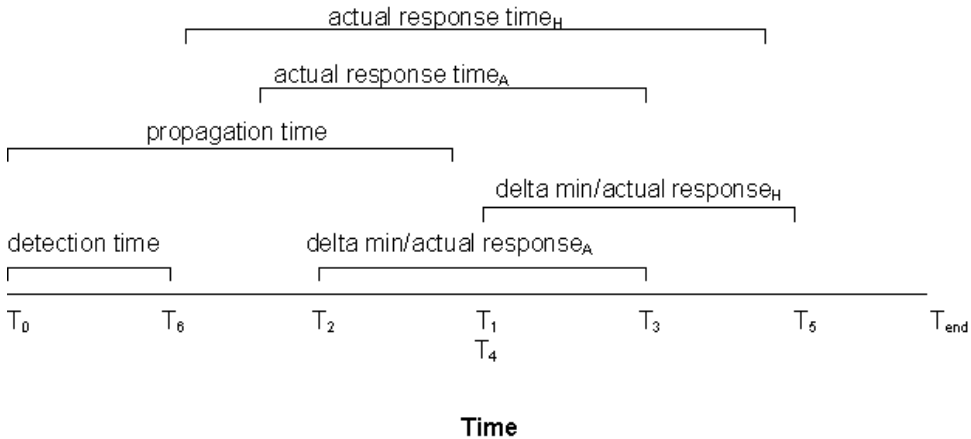
[Exhibit 13](#) presents a warning time analysis report. Several observations can be made from a report such as this. In this example, both the actual automatic response and actual human response exceeded the minimum time available. The responsiveness of automatic system and human actions in preventing an accident/incident can be measured as a function of the delta between the minimum and actual response times. Deltas between minimum and actual response times, both automatic system and human, need to be explained as part of an accident/incident investigation; that is, (1) what caused the delays, (2) were the delays accidental or intentional, (3) were the delays avoidable or unavoidable, and (4) what can be done to eliminate or reduce delays in the future.

[Exhibit 14](#) summarizes the interaction between the different accident/incident investigation techniques. The figure is not meant to show a chronology in which to use the techniques, but rather the information flow among the techniques. Some techniques are used to develop “how did” accident scenarios; others are used to develop “how could” accident scenarios. Some techniques are usually only performed once, while others are repeated as evidence accumulates and unknowns or uncertainties are resolved.

Barrier analysis reports have multiple uses beyond their original scope. When combined with TLA graphs, they measure the effective interval of each

# Warning Time Analysis Report

System: \_\_\_\_\_ as of date: \_\_\_\_\_



## Data Points

- $T_0$ : initiating event
- $T_1$ : accident/incident
- $T_2$ : minimum response<sub>A</sub>
- $T_3$ : actual response<sub>A</sub>
- $T_4$ : minimum response<sub>H</sub>
- $T_5$ : actual response<sub>H</sub>
- $T_6$ : earliest alarm indication
- $T_{end}$ : accident/incident surceases

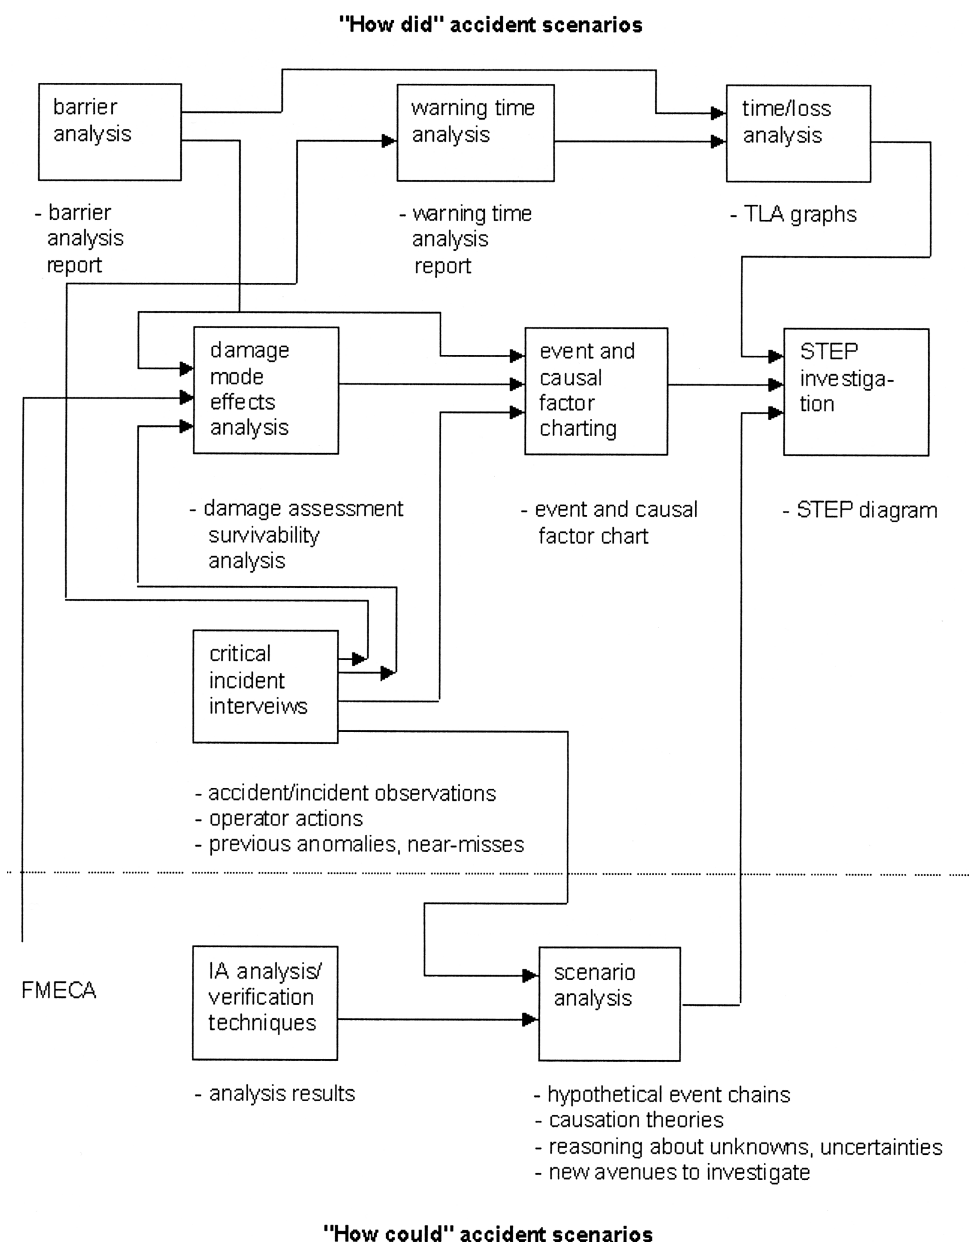
## Data Ranges

- $T_0 - T_4$ : propagation time
- $T_0 - T_6$ : detection time
- $T_0 - T_3$ : actual response time<sub>A</sub>
- $T_0 - T_5$ : actual response time<sub>H</sub>
- $T_2 - T_3$ : delta minimum/actual response time<sub>A</sub>
- $T_4 - T_5$ : delta minimum/actual response time<sub>H</sub>

## Exhibit 13 Warning Time Analysis Report

defensive layer prior to failure. The failure of defensive layers is captured on event and causal factor charts and STEP diagrams. Failed defensive layers are analyzed when performing damage mode effects analysis.

Information collected during critical incident interviews provides useful input to damage mode effects analysis, event and causal factor charts, and scenarios analysis. As mentioned earlier, event and causal factor charts can be developed as draft input to a STEP diagram. TLA graphs capture timing specifics about intervention actions that are recorded on STEP diagrams. Scenario analysis supports reasoning about unknowns, uncertainties, and gaps in STEP diagrams; it also identifies new avenues to investigate.



**Exhibit 14 Interaction Between Accident/Incident Investigation Techniques**

## 8.2 Initiate Short-Term Recovery Mechanisms

After an accident/incident occurs, the cause, extent, and consequences are investigated. The initial report of the accident/incident triggers immediate short-term recovery mechanisms. This process is discussed below. Complete follow-up accident/incident reports stimulate long-term remedial measures. That process is discussed in the chapter section “Deploy Remedial Measures.” Initiating recovery mechanisms and reporting an accident (chapter section

## Exhibit 15 Accident/Incident Recovery Steps

---

1. Review preliminary investigation results about cause(s), extent, and consequences of the accident/incident.
  2. Determine what can and cannot be recovered in the short term:
    - a. Systems
    - b. Communication equipment
    - c. Hardware
    - d. System software
    - e. Application software, services
    - f. Data
  3. Ascertain when each system, entity, and component can and should be restored:
    - a. Technical considerations
    - b. Operational priorities
    - c. Safety and security priorities
  4. Decide how each system, entity, and component can and should be restored:
    - a. Level of service to be restored
    - b. Actions, commands necessary to effect recovery
    - c. Verifying effectiveness of recovery efforts
  5. Notify customers, end users, system administrators, maintenance staff, etc.
    - a. Problem experienced
    - b. Emergency precautions
    - c. Estimated recovery time
- 

“Report Accident/Incident”) are parallel activities. [Exhibit 15](#) summarizes the five steps involved in initiating recovery mechanisms.

The preliminary accident/incident investigation results drive the emergency recovery response. In many cases, there is a one-to-one if this cause/consequences, then take this recovery action. Pre-accident modes/states and conditions, dependencies and interactions among systems and entities, and automatic system and human-initiated corrective action already taken (whether or not it was correct or effective) are taken into account before initiating recovery mechanisms.

The second step is to determine what systems and entities can and cannot be recovered in the short term. In particular, which internal and external systems, entities, hardware, communications equipment, system software, applications software/services, and data can or cannot be restored. The extent of the damage, whether or not the entities are under your control, and other logistical matters will factor into this determination.

The third step is to ascertain when each system and entity can and should be restored. This step can be complex. Technical requirements and priorities are interleaved with operational requirements and priorities. Several factors are meshed together:

1. When is it technically feasible to restore each entity and component?
2. What is the optimum sequence in which entities and components should be restored, from a technical perspective, taking into account dependencies and logistical considerations?
3. What are the operational priorities for restoring various systems and services?
4. What are the safety and security priorities for restoring various threat control measures?

In short, an entire system and all of its entities cannot be restored all at once after an accident/incident; instead, recovery is effected gradually.

Prior to initiating recovery mechanisms, one must verify that the accident/incident has been contained and mitigated, that it is in effect “over” and not likely to recur or take new unanticipated twists and turns. Premature recovery efforts may interfere with automatic system and human-initiated action taken to contain and mitigate an accident/incident. In the worst-case scenario, this can result in a prolonged cycle of an accident/incident, failed automatic and human responses, failed recovery efforts, and accident/incident propagation and recurrence.

The fourth step is to decide how each system, entity, and component can and should be restored. This step involves several decisions, such as what level of service should be restored, what actions and commands are necessary to effect recovery, and how to verify the effectiveness of recovery efforts.

The level of service to be restored in the short term will depend on the extent and consequences of the accident/incident and operational priorities. Some sessions and processes may need to be terminated immediately to reduce system load and the possibility of further failure or compromise. Critical transactions may be “frozen” or checkpointed to prevent loss of critical data. A limited number of IA-critical services may be restored immediately while further recovery is pending. The number and location of users who can access the limited services may also be restricted initially.

Operational procedures and contingency plans should spell out the step-by-step actions and commands necessary to effect recovery. They should also identify who to contact for further help. Transaction paths and critical threat zones should be consulted to locate attack points and the corresponding recovery points. Recovery actions and commands can involve any of the following:

- Activating cold spare or hot standby redundant hardware
- Reconfiguring a system or network
- Restarting, reloading, reinitializing a system or data from local or offsite archives
- Switching operations to a remote location
- Switching to an alternate service provider
- Restoring and restarting access control rules, authentication parameters and processing, security audit trail/alarm, and other threat control measures

Before the restored services can be turned over to users, the effectiveness of the recovery efforts must be verified. This involves verifying the robustness of the restored services and the robustness of the restored threat control measures. In particular, the ability of the restored threat control measures to withstand a new or repeat accident/incident needs to be verified; a survivability assessment is paramount.

The first four steps are more or less sequential. The fifth step, notification, is an ongoing activity that occurs in parallel with the other four steps. A critical

component of recovery is notifying customers, end users, system administrators, maintenance staff, etc. that an accident/incident has occurred. Initially, they should be informed about the type of problem experienced; the description should be concise and not overly technical. Concurrently or immediately afterward, they should be advised of any emergency precautions to take. The third piece of information to convey is the estimated time required for a limited recovery and for a full recovery. It is important to keep all stakeholders informed during recovery efforts; without accurate and timely information, they might accidentally sabotage or further complicate recovery efforts.

In summary, the effectiveness and timeliness of recovery efforts are entirely dependent on prior planning, preparation, coordination, and training. All the information needed to perform steps 2 through 5 should be recorded in the operational procedures and contingency plans. The thoroughness of the operational procedures and contingency plans and the familiarity of operational personnel with them are key ingredients for a quick and successful recovery. The other option, of course, is panicked guessing.

### 8.3 Report Accident/Incident

Reporting an accident/incident is an essential part of investigating, responding to, and recovering from it. An initial report should be filed as soon as an accident/incident is known or suspected. The initial report describes the characteristics of the accident/incident, the time it was first detected, suspected cause and source (if this information is available at the time), the consequences to date, and estimated recovery time and actions needed. Later, as more is known about the accident/incident, follow-up reports are filed. Facts and objective observations replace earlier theories and suppositions.

Organizations need to have clearly defined accident/incident reporting channels and responsibilities, inside and outside the organization. Employees, end users, and customers should be encouraged to report known or suspected accidents/incidents; this is not the responsibility of a system administrator alone. People who report an accident/incident may or may not participate in the subsequent investigation; most likely, they will participate in critical incident interviews.

There are several reasons to report an accident/incident, inside and outside an organization, and benefits to be derived from doing so. As Bond<sup>218</sup> observes:

*...an accident is the invasion of the unaware by the unknown. To reduce accidents, we must make the person aware of the hazards unknown by him but known too well by others.... A wise man learns from his own experience, but a wiser man learns from the experiences of others.*

First, an accident/incident must be reported before the situation can be corrected. If the accident/incident is reported in a timely manner, the damage/loss experienced by this and other systems can be minimized, as shown in

the TLA graphs (Exhibit 12). Second, reporting the results of an accident/incident investigation and what was learned from it reduces the likelihood of recurrence, within and among organizations.<sup>218</sup> Third, customers and employees will have more confidence in an organization that reports accidents/incidents; they gain the impression that the organization is being open and is on top of the situation. This is another example of perception management. Fourth, an organization may have a legal responsibility to report accidents/incidents to stockholders, customers, the public, or a regulatory agency, depending on the nature of the organization and the legal jurisdiction in which it resides.<sup>248</sup> Fifth, an accident/incident must have been reported if subsequent legal action is to be taken. Finally, as Rathmell<sup>391</sup> points out:

*...national information assurance can only be achieved if threat assessments and early warnings are distributed widely across industry and to the public.*

There are a variety of potential sources to notify of known or suspected accidents/incidents outside an immediate organization. The computer emergency response team coordination center (CERT/CC) was established November 1988 at Carnegie Mellon University Software Engineering Institute (SEI). CERT/CC provides a 24-hour central point of contact and clearinghouse for identifying vulnerabilities and responses. CERT/CC maintains a knowledge base of computer network and system vulnerabilities and best practices.<sup>350</sup> The center also studies Internet security vulnerabilities, provides incident response services, publishes security alerts, and researches security and survivability issues in wide-area computing. During the first ten years in operation, CERT/CC responded to 14,000 incidents, published 180 advisories, replied to 200,000 e-mail messages, and answered 15,000 hotline calls.<sup>379</sup> The time to get acquainted with CERT/CC is before one experiences an accident/incident. They can be contacted at [www.cert.org](http://www.cert.org).<sup>460</sup>

CERT/CC resources are available to the public. U.S. government agencies have the option of subscribing to the federal computer incident response center (FedCIRC), managed by the General Services Administration (GSA). FedCIRC was established October 1998 and provides 24-hour hotline, e-mail, and help desk support, as well as security alerts and advisories. FedCIRC can be contacted at: [www.fedcirc.gov](http://www.fedcirc.gov).<sup>469</sup>

The U.S. National Infrastructure Protection Center (NIPC) is managed by the FBI.<sup>326</sup> Established in 1998, the NIPC's mission is to "detect, deter, warn, and respond to attacks on the nation's critical infrastructures," both physical and cyber, and "to serve as the government's information clearinghouse for both security and responses to attacks by individuals and foreign governments."<sup>196</sup> Among other things, the NIPC is developing a multi-source database of known threats and actual intrusions, successful and unsuccessful responses, prevention strategies, and computer crime trends. Indispensable to NIPC's success is the cooperation and exchange of information between public and private organizations. Accordingly, the formation of the public-private Partnership for Critical Infrastructure Security was announced January 2000.<sup>355</sup>

Initial members of the Partnership included: RSA Security, Cisco Systems, Network Associates, Microsoft, and CERT/CC.<sup>355</sup> Information about NIPC and the Partnership can be found at [www.fbi.gov/nipc/index.htm](http://www.fbi.gov/nipc/index.htm).<sup>468</sup>

Readers may also want to contact local and national law enforcement agencies. In the United States, the damage threshold for computer crime to be prosecuted is very low — only a few thousand dollars. Lost productivity most certainly should be included in the damage estimate. A good organization to contact in this regard is the High Tech Crime Network at [www.htcn.org](http://www.htcn.org).<sup>472</sup>

The news media should also be contacted if one needs to alert customers quickly, for example, to not use ATMs, to not access online banking systems, or that credit card information has been compromised.

Software vendors are also beginning to post information about security features, vulnerabilities, alerts, and patches; see, for example, Microsoft's security site at [www.microsoft.com/security](http://www.microsoft.com/security).<sup>480</sup>

A major challenge facing the NIPC and CERT/CC is the need for standardized reporting of accidents/incidents. Without standardized reporting elements, notation, and criteria, the data accumulated and shared has little value. Vulnerabilities, threats, responses, consequences, etc. need to be reported and categorized uniformly. In this way, keyword searches, queries about how to respond to an emergency situation or prevention strategies for a particular threat, and statistical reports showing the prevalence of certain types of malicious activity can be generated quickly with a reasonable degree of accuracy and shared across multiple organizations.<sup>226</sup> A main premise of the Partnership for Critical Infrastructure Security and NIPC is that the privacy of individuals and corporations will be protected when collecting, reporting, storing, and aggregating this information; without demonstrated privacy protections, it is unlikely that many organizations will participate.

This is not a new challenge. The U.S. Department of Energy (DoE), Federal Aviation Administration (FAA), Food and Drug Administration (FDA), Occupational Safety and Health Administration (OSHA), and National Safety Transportation Board (NSTB) have had accident/incident reporting systems for years. The degree of success these agencies have had in collecting and reporting standardized data varies. One of the more successful efforts has been the DoE Computerized Accident/Incident Reporting System (CAIRS); the need for standardized data elements and codes was recognized early. Between 1981 and 1993, 50,000 accident/incident investigation reports were entered into CAIRS.<sup>226</sup> What is unique about the NIPC and CERT/CC challenge is that the data is voluntarily reported; unlike the other systems, there is no legal mandate to report the information.

It was previously mentioned that safety and reliability engineers have considerable forensic engineering experience, although they do not refer to it as such. They also have considerable experience in developing and using standardized failure reporting and analysis systems, usually referred to as a data reporting and corrective action system (DRACAS) or failure reporting and corrective action system (FRACAS). A DRACAS or FRACAS is a good model to follow when reporting or collecting accident/incident information.

All in-service anomalies are recorded in a DRACAS or FRACAS. This information is used to determine the safety and security significance of the

anomaly, whether the anomaly is a symptom of a more serious underlying problem, and the immediate precautions to be taken.<sup>130</sup> Anomaly reporting often results in the identification of a new vulnerability/hazard or reclassification of known vulnerabilities/hazards.<sup>130,131</sup> In addition, anomaly reporting supports change impact analysis as part of initiating recovery mechanisms. It is important that accident/incident reports be submitted promptly<sup>131</sup>:

*All reported faults/failures need to be considered for possible action as quickly as possible. If action is not taken quickly it is likely that the affected items will be returned for repair and valuable data which could have been obtained by a detailed investigation will be lost.*

A DRACAS or FRACAS accurately and consistently categorizes accidents/incidents according to their cause, significance, and frequency.<sup>131</sup> The cause establishes the exact root and intermediate causes of the accident/incident. Care must be taken not to draw premature or unsubstantiated conclusions. As an analogy, the fact that a light is on or a car is parked in front of a house **only** means that the light is on and the car is parked in front of a house; it does not mean that someone is home. The significance of an accident/incident rates its effect on the ability to perform essential mission capabilities. A quantitative measurement of the frequency and duration of the accident/incident is captured for trend analysis purposes. The remaining useful life of the system is also assessed after an accident/incident.

Exhibits 16 and 17 present a standardized accident/incident report template. Confidential information, such as individual name, organization name, address, and contact information has been removed. The report is in two parts. The first part (Exhibit 16) contains the accident/incident description. Information about the anomaly, its severity, and the conditions under which it was experienced is captured through a combination of category codes and free text fields. This information is necessary to classify the accident/incident so that (1) the information can be accurately shared with other organizations via alerts, and (2) the appropriate emergency response can be determined. As much of this information is provided as possible in the initial report, without unduly delaying it. Additional detail and corrections are provided in follow-up reports. An important section (11) is the identification of other systems/entities inside and outside the organization that may also be impacted by the accident/incident.

The second part of the report (Exhibit 17) is the accident/incident assessment. This part of the report records the causes and consequences of the accident/incident, including estimated damages. Two important sections are the recommendations for short-term recovery mechanisms (11) and the lessons learned to be applied to long-term remedial measures (12).

## 8.4 Deploy Long-Term Remedial Measures

An accident/incident investigation is undertaken to discover what exactly happened, how it happened, and why it happened or was allowed to happen.

## Exhibit 16 Accident/Incident Report: Part I-Description

### I. Accident/Incident Description

<i>Report Field</i>	<i>Initial Report</i>	<i>Follow-up Report</i>
1. Report reference number	x	x
2. Anomaly classification (see Chapter 6, <a href="#">Exhibit 12</a> )	x	x
3. Description of failure/compromise	x	x
4. Severity:	x	x
a. Catastrophic      c. Critical		
b. Marginal          d. Insignificant		
5. Date/time first detected or experienced	x	x
6. Frequency experienced	x	x
7. Duration	x	x
8. Mission significance:	x	x
a. Failure of IA-critical functions/entities (cite)		
b. Failure of IA-related functions/entities (cite)		
c. Failure of MWFs (cite)		
d. Failure of MNWFs (cite)		
e. No option but to fail safe/secure		
f. No option but to fail operational		
g. Total loss of system		
h. Loss of critical/sensitive data		
i. Number of personnel affected		
9. Primary systems/entities affected:	x	x
a. Entity/system ID and origin		
b. System/entity type		
c. Number of systems/entities affected		
10. Time in operation prior to accident/incident	x	x
11. Other systems/entities inside and outside organization that may be impacted	x	x
12. System configuration, version numbers, etc.	x	x
13. Network configuration, version numbers, etc.	x	x
14. Assumptions	x	x

The results of an accident/incident investigation drive emergency short-term recovery mechanisms; form the basis of reports submitted to CERT/CC, FedCIRC, NIPC, HTCN, and other internal and external organizations; and may precipitate legal action. Equally important, accident/incident investigation results stimulate long-term remedial measures.

Accident/incident investigation reports are analyzed to learn from the what, how, and why of an accident/incident. The most obvious reason is to determine what remedial measures are necessary to prevent the same or similar accidents from recurring. As Petroski<sup>380</sup> succinctly states:

*No matter how tragic a failure might be, it is obviously more tragic if it could have been anticipated and prevented.*

## Exhibit 17 Accident/Incident Report: Part II-Assessment

---

### II. Accident/Incident Assessment

<i>Report Field</i>	<i>Initial Report</i>	<i>Follow-up Report</i>
1. Conditions that produced accident/incident	?	x
2. Critical event sequence		x
3. Related near-misses		x
4. Consequences		
a. Likely	x	
b. Actual		x
5. Corrective action taken		
a. Automatic system	x	x
b. Human initiated	x	x
6. Investigation techniques used		x
a. Barrier analysis		
b. Critical incident interviews		
c. Damage mode effects analysis		
d. Event and causal factor charting		
e. Scenario analysis		
f. STEP investigation system		
g. TLA for emergency response		
h. Warning time analysis		
i. Other		
7. Investigation results: include reports, diagrams, graphs, etc.		x
8. Intermediate and root causes		x
a. Design error		
b. Implementation error		
c. Operational procedures error		
d. Contingency plan error		
e. Physical security practices error		
f. Accidental human action		
g. Malicious intentional human action		
h. Inadvertent operation		
i. Failure or unavailability of key infrastructure system		
j. Other		
9. Remaining useful life of system		x
10. Estimated loss/damages		x
11. Recommendations for short-term recovery	?	x
12. Observations/lessons learned for long-term remedial measures		x

---

Not preventing avoidable accidents may have legal consequences (see chapter section “Evaluate Legal Issues”).

Second, during the course of an accident/incident investigation, latent vulnerabilities and threats are often uncovered. An analysis of the investigation results presents an opportunity to employ remedial measures to eliminate or mitigate the latent vulnerabilities and threats before they cause damage or loss.

Third, because an accident/incident investigation is not limited to technical issues, needed improvements in operational procedures, contingency plans, physical security measures, training, logistical matters, etc. are exposed as well.<sup>31,320,381</sup> Again, an analysis of accident/incident investigation results presents an opportunity to employ remedial measures to eliminate these deficiencies before they cause damage or loss.

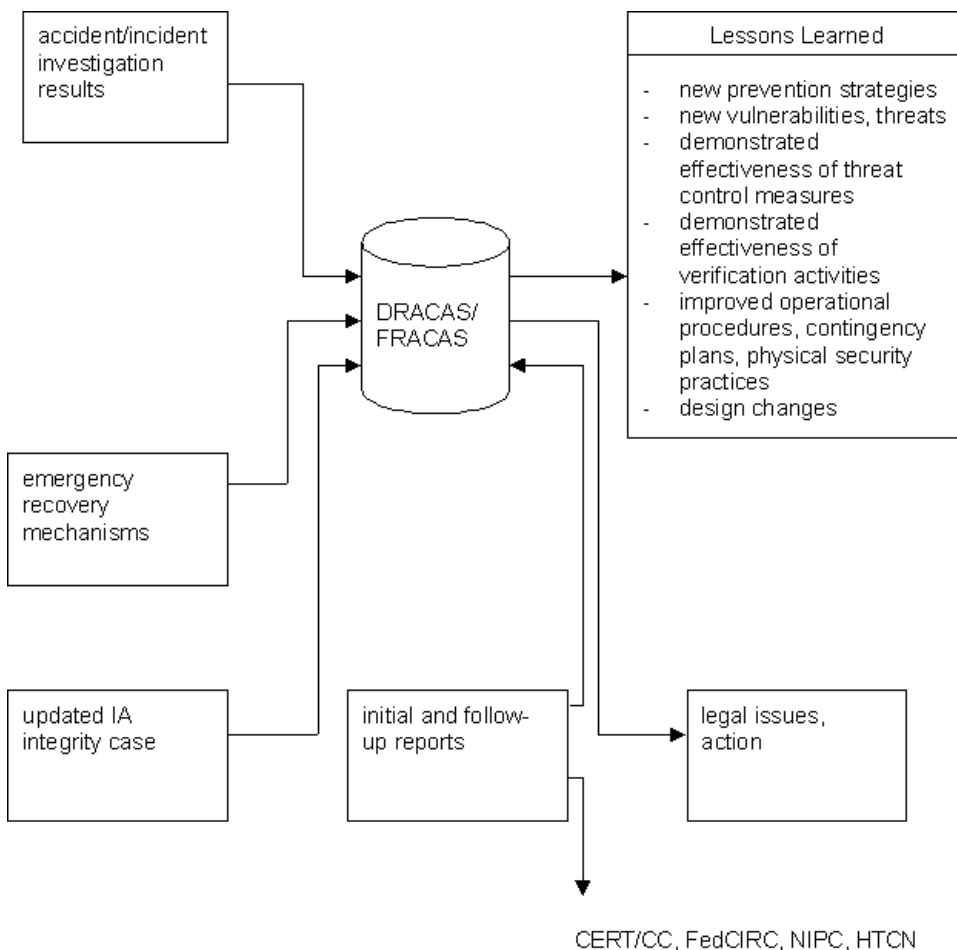
In general, catastrophic accidents are investigated more thoroughly than less-serious accidents/incidents because of their severity and less frequent occurrence.<sup>320</sup> It is debatable whether this is a wise approach. Accident/incident investigations should be conducted for the purpose of organizational learning and feedback, not just for the purpose of collecting actuarial or other legal data.<sup>320</sup> As mentioned, incidents are often precursors to or early warnings of impending accidents. There is much to be learned about the need for remedial measures from analyzing them. A thorough incident investigation, if acted upon, may forestall an accident. Conversely, failing to thoroughly investigate an incident that subsequently leads to an accident may have legal consequences. Reporting and analysis of known and suspected near-misses should be encouraged. Many organizations mistakenly discourage the reporting and analysis of incidents, leading to “surprise” when a serious accident occurs. Under-reporting and under-analysis of incidents can have solemn technical, financial, and legal consequences.

The analysis of accident/incident results reinforces the need for standardized reporting of accident/incident data. To be meaningful, uniform data elements, categories, and evaluation criteria must be used. [Exhibit 18](#) illustrates the information flow between accident/incident investigations, reports, and remedial measures.

The DRACAS/FRACAS functions as the central accident/incident information storage and retrieval system. Information is collected about the accident/incident, how and why it occurred, and what corrective action was needed to respond to and recover from it so that individuals and organizations can learn from their mistakes and those of others. Accident/incident information is collected, reported, and analyzed throughout the life of a system. All malfunctions, accidents, anomalies, near-misses, deviations, and waivers should be analyzed.<sup>31</sup> Lessons learned are derived from analyzing accident/incident reports, which can then be deployed as remedial measures. Lessons learned can take many forms, including the<sup>31,320</sup>:

- Discovery of new prevention strategies, tools, and techniques
- Identification of new vulnerabilities and threats
- Demonstrated effectiveness of threat control measures
- Demonstrated effectiveness of verification activities
- Improved operational procedures, contingency plans, and physical security practices
- The need for design changes

In summary, the lessons learned from an accident/incident make clear what worked, what did not work, and what needs to be changed. To ensure



## Exhibit 18 Information Flow Between Accident/Incident Investigations, Reports, and Remedial Measures

that the correct lesson is learned from an accident/incident, it is essential that an investigation be thorough, accurate, unbiased, and that fact, opinion, theory, and assumptions are clearly delineated. The IA integrity case is updated to reflect the accident/incident investigation results, particularly the lessons learned.

## 8.5 Evaluate Legal Issues

This chapter section is not part of a legal textbook, nor is it offering legal advice. Rather, the purpose of this chapter section is to make the reader aware of the legal issues involved in information security/IA and the need to seek appropriate legal counsel.

Several legal terms are used outside their precise legal meaning in everyday speech. Hence, it is important to clarify the legal definition and usage of these

terms before discussing the related issues. The following definitions are from Black's *Law Dictionary*<sup>214</sup>:

**Defect:** deficiency; imperfection; insufficiency; the absence of something necessary for completeness or perfection; a deficiency in something essential to the proper use for the purpose for which a thing is to be used; a manufacturing flaw, a design defect, or inadequate warnings. A design defect exists whenever the design itself poses unreasonable dangers to consumers.

**Damage:** loss, injury, or deterioration, caused by the negligence, design, or accident of one person to another, in respect of the latter's person or property; the harm, detriment, or loss sustained by reason of injury.

**Injury:** any wrong or damage done to another, either his person, rights, reputation, or property; the invasion of any legally protected interest of another.

**Negligence:** failure to use such care as a reasonably prudent and careful person would use under similar circumstances; the doing of some act which a person of ordinary prudence would not have done under similar circumstances or failure to do what a person of ordinary prudence would have done under similar circumstances; conduct which falls below the norm for the protection of others against unreasonable risk of harm. It is characterized by inadvertence, thoughtlessness, inattention, recklessness, ...

**Liability:** condition of being or potentially subject to an obligation; condition of being responsible for a possible or actual loss, penalty, evil, expense, or burden; condition that creates a duty to perform an act immediately or in the future; including almost every character of hazard or responsibility, absolute, contingent, or likely.

**Assumption of risk:** a plaintiff may not recover for an injury to which he assents, that is, that a person may not recover for an injury received when he voluntarily exposes himself to a known and appreciated danger. The requirements for the defense ... are that: (1) the plaintiff has knowledge of facts constituting a dangerous condition, (2) he knows that the condition is dangerous, (3) he appreciates the nature or extent of the danger, and (4) he voluntarily exposes himself to the danger. Secondary assumption of risk occurs when an individual voluntarily encounters known, appreciated risk without an intended manifestation by that individual that he consents to relieve another of his duty.

At the beginning of this chapter, an accident was defined to involve death, injury, loss or damage. Consequently, an accident may involve multiple legal issues, regardless of whether the accident is the result of accidental or malicious intentional action. Defects include design defects, such as inadequate or ineffective threat control measures, and inadequate warnings to customers and end users. Injuries and loss may be to a person, his reputation, or property, physical or cyber. For example, if sensitive personnel, financial, medical, or other information is compromised, a person's reputation could be damaged; identity theft and employment discrimination are two of many possible scenarios. Negligence incorporates errors of commission and errors of omission that did not prevent the accident from occurring. Legal liability determines who was responsible for preventing the accident and thus paying damages to the injured

party. Assumption of risk is equivalent to the concept of informed consent prior to authorizing risky medical procedures. Depending on what risk was knowingly assumed, damages paid by the liable party or parties may be reduced. In summary:

Negligence → Defect → Injury → Damage → Liability to pay  
Damages for injury

In the realm of information security/IA, legal issues and responsibilities can arise from many different perspectives, including system owner, end user, system administrator, victim, customer, stockholder, and test/certification lab. Legal issues, responsibilities, and liability can be distributed among individuals and organizations. Falla<sup>259</sup> points out that:

*Liability can fall on the manufacturer, supplier, distributor, or certifier of products. ... Suppliers of components can also be liable. In cases where the component is used in products which are exposed to the general public, the extent of such liability can be enormous.*

Little case law exists in the field to date, it is still evolving. As Wood<sup>442</sup> observes:

*...the risks of cyberspace go far beyond the military and raise complex and unprecedented ethical and legal issues that current policies, organizations, laws, and procedures cannot readily answer.*

Given this situation, most sources recommend a proactive legal risk management approach.<sup>230,248,259</sup> The first step is to define appropriate boundaries of authority and responsibility for technical and legal decisions and oversight. Liability issues should be reviewed throughout the development and operation of a system and at predefined milestones. The connection between technical and legal risk should be examined regularly.<sup>230</sup> As Falla<sup>259</sup> notes:

*...attaching a particular legal step or procedure to each event in the life cycle ... helps legal precautions to be taken at appropriate times and often also preempts the escalation of legal problems.*

Given that technology itself is continuously evolving, it is sometimes possible to use the legal defense that current industry best practices were followed. To do so, (1) best practices, like those described in this book, actually have to have been followed, and (2) all stakeholders have to be aware of and live up to these legal responsibilities. Burnett<sup>230</sup> summarizes these responsibilities as follows:

**Designer:** the system designer/developer is responsible for ensuring that the system will fail safe/secure or fail operational, as appropriate, in all situations so that no damage or loss is incurred.

**Technical experts:** technical experts, whether employees or consultants, are responsible for maintaining complete, in-depth, and current competence in their field, such that this competence is above average, but not necessarily at the genius level.

**Component suppliers:** component suppliers, such as COTS vendors, are responsible for accurately representing component capabilities, limitations, claims, labelling, and instructions for use.

**Testing and certification labs:** testing and certification labs are responsible for accurately explaining what was and was not tested or evaluated, providing accurate test results, an accurate description of test coverage, and defensible reliability, safety, and security claims. Test and certification labs are responsible for employing competent people to perform the tests/evaluation and verifying that facts, opinions, and assumptions are separated.

Product liability stems from the concept that products must be fit for all purposes for which goods of that kind are commonly used. This concept, often referred to as “fitness for use” or “fitness for purpose,” implies that products are free from defects, safe, reliable, and secure. If a technical failure results in loss or damage, liability for damages may be incurred. Liability can be limited by two types of remedies: (1) replacement/repair and (2) limiting damages to a particular amount. Liability cannot be restricted when negligence is a factor and death or serious injury (physical or cyber) results.

Liability may be civil or criminal, depending on the nature of the accident, and falls under the due-care or strict liability criteria. Without going into too much legalese, the plaintiff must establish that<sup>230,259</sup>:

- The manufacturer and/or supplier owed the plaintiff a duty of care.
- There has been a breach of this duty that caused the damage/loss; for example, failing to adequately verify the safety or security of a system or component.
- The kind of damage sustained was reasonably foreseeable as a consequence of that breach.
- Damage/loss has, in fact, occurred.

The concept of strict liability in tort eliminates the due-care criteria; the plaintiff only has to demonstrate the last two items.<sup>420</sup>

A product may be considered defective if it exhibits safety, reliability, or security behavior that is less than that to which people are generally entitled to expect.<sup>230</sup> Damage may result in personal injury, material damage, economic loss, environmental corruption, administrative chaos, and privacy violations.

A related legal issue to consider is warranties. The U.S. Uniform Commercial Code (UCC S2-315) equates fitness for purpose to implied warranties. In the past, the concept of express and implied warranties has been applied to commercial products, such as COTS software. In the future, as information security/IA case law evolves, it is not inconceivable that the concept of implied warranties could be applied to online banking systems and other IT services.

Another legal issue concerns the assumption of risk, primary and secondary. If the plaintiff knew about the risk, understood the potential consequences, appreciated the nature and extent of the risk, and voluntarily accepted this risk, liability for damages may be limited.<sup>214,420</sup> If, however, these four criteria are not met, the risk is considered to be unassumed. Accident/incident investigations clarify risk assumed and unassumed by all stakeholders.

System owners may be liable for the theft or fraudulent use of information or resources, whether these crimes are committed by insiders or outsiders. It is even conceivable that stockholders could sue a corporation for negligence if such a crime occurred.<sup>248</sup> A proactive legal risk management approach is the best defense in this situation. This approach should include what is referred to as a good-faith effort to prevent such crimes, for example<sup>248</sup>:

*...written policies and procedures against crime, security awareness programs, disciplinary standards, monitoring and auditing systems that represent applicable industry practice, reporting detected crimes to law enforcement agencies, and cooperating with investigations.*

As an analogy, if there are locks on the doors and windows of one's house and an electronic burglar alarm, one is more likely to collect on an insurance policy after a burglary than if one leaves the doors and windows open.

This discussion has been presented from the perspective of current U.S. law to enlighten the reader on the legal issues involved in information security/IA and the need to seek legal counsel. The laws pertaining to computer crime, civil and criminal liability, and mandatory privacy protections vary from country to country. Often, computer crime laws are not enacted until after a major offense has been committed. Campen<sup>236</sup> observes that:

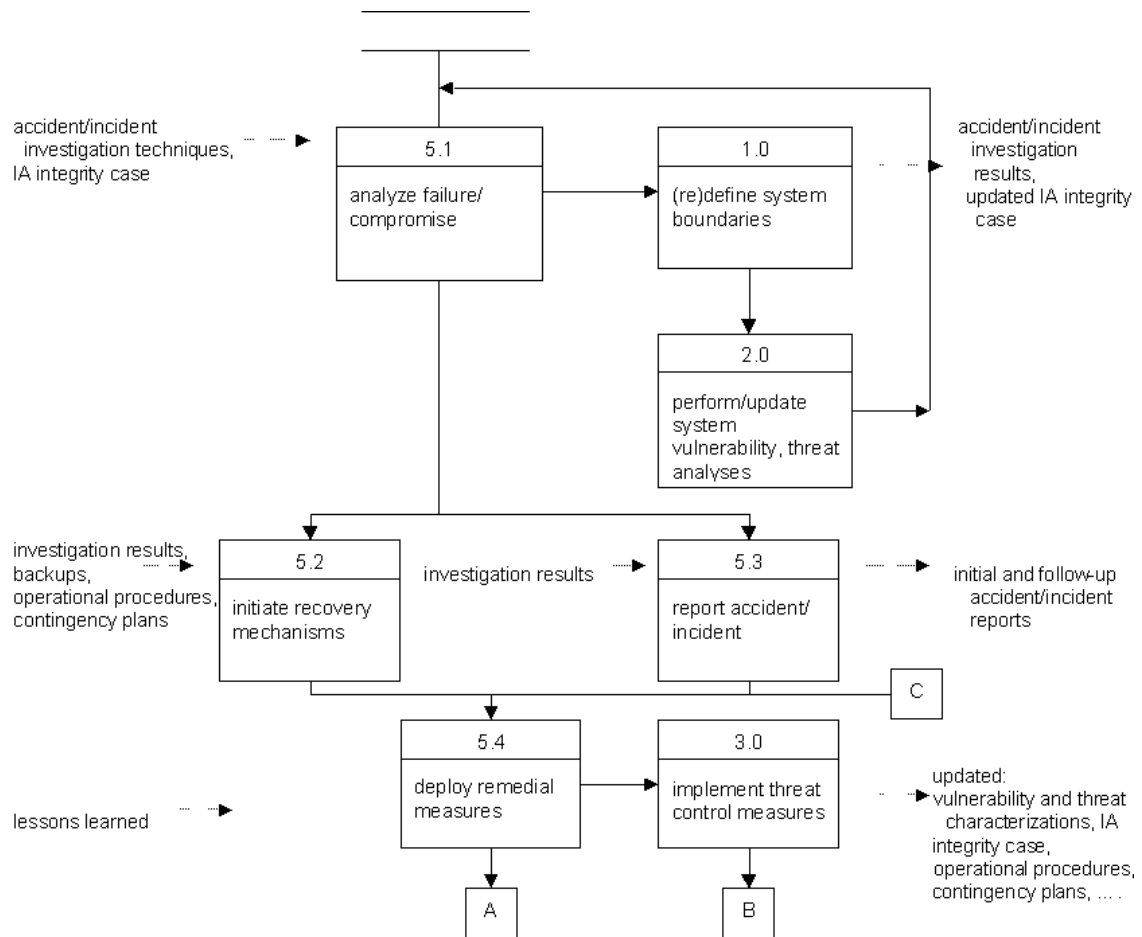
*Information security is an international issue, involving diverse cultures that do not hold a common view of personal privacy in cyberspace or how it should be secured.*

## 8.6 Summary

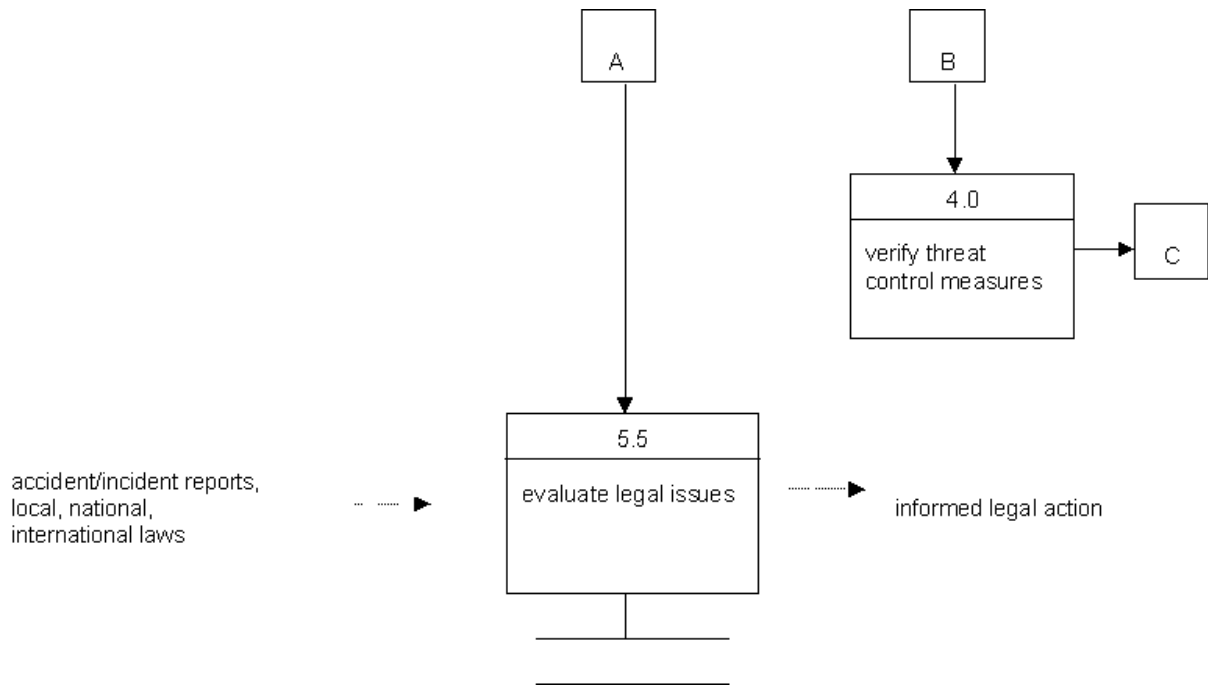
The fifth component of an effective information security/IA program is conducting an accident/incident investigation. Five activities are performed while conducting an accident/incident investigation, as shown in [Exhibits 19](#) and [20](#).

- The cause, extent, and consequences of the failure/compromise are analyzed.
- Recovery mechanisms are initiated.
- The accident/incident is reported.
- Remedial measures are deployed.
- Legal issues are evaluated.

Accident/incident investigations are conducted for legal and engineering reasons. As a result, it is important to understand and distinguish between the



**Exhibit 19 Summary of Activities Involved in Conducting Accident/Incident Investigations**



---

**Exhibit 20 Summary of Activities Involved in Conducting Accident/Incident Investigations (continued)**

legal and engineering usage of terms, such as accident, incident, failure, compromise, and cause.

Accident/incident investigations are conducted to determine, in fact, what did and did not happen, how it happened, and why it happened or was allowed to happen. There are plenty of open sources to examine when collecting evidence for an accident/incident investigation. Analyses of early evidence point to other sources of evidence in what becomes an iterative process. Inductive and deductive reasoning is applied to the evidence to explain how and why an accident/incident occurred. A combination of techniques is used to investigate an accident/incident by developing “how could” and “how did” accident scenarios. The same techniques are used to conduct an internal, independent, regulatory, or forensic investigation. The credibility of an accident/incident investigation rests on the ability to remain objective; eliminate bias or prejudice; separate fact, opinion, assumptions, and theory; and distinguish “symptoms” from the “disease” — all while being thorough and accurate.

Initial accident/incident reports trigger short-term recovery mechanisms. Preliminary investigation results about the cause(s), extent, and consequences of the accident/incident are reviewed. A determination is made about what can and cannot be recovered in the short term. When and how each system, entity, and component can and should be restored are ascertained. Customers, end users, system administrators, maintenance staff, etc. are notified about the accident/incident and the status of recovery efforts. The effectiveness and timeliness of recovery efforts are entirely dependent on prior planning, coordination, and training as well as complete operational procedures and contingency plans.

Reporting an accident/incident is an essential part of investigating, responding to, and recovering from it. There are several reasons to report an accident/incident, inside and outside an organization, and benefits to be derived from doing so:

1. An accident/incident must be reported before the situation can be corrected.
2. Reporting the results of an accident/incident and what was learned from it reduces the likelihood of recurrence, within and among organizations.
3. Customers and employees will have more confidence in an organization that reports accidents/incidents.
4. An organization may have a legal responsibility to report an accident/incident.
5. An accident/incident must have been reported if subsequent legal action is to be taken.

Standardized initial and follow-up reports should be filed to internal and external organizations such as CERT/CC, FedCIRC, NIPC, and HTCEN.

Follow-up accident/incident reports stimulate long-term remedial measures. The results of accident/incident investigations are analyzed to derive lessons learned, such as<sup>31,320</sup>:

- Identification of new vulnerabilities and threats
- Demonstrated effectiveness of threat control measures
- Demonstrated effectiveness of verification activities
- Improved operational procedures, contingency plans, and physical security practices
- The need for design changes

There are several legal issues involved in information security/IA. Legal issues and responsibilities may arise from many different perspectives, including system owner, end user, system administrator, victim, customer, stockholder, test/certification lab, etc. Engineers need to be aware of these issues, seek appropriate legal counsel, and pursue a proactive legal risk management strategy.

## 8.7 Discussion Problems

1. Explain the relationship, if any, between a compromise and: (a) an accident, (b) a vulnerability, (c) a threat control measure, (d) a failure, and (e) intentional malicious action.
2. When are accidents/incidents investigated? Why?
3. Why would an organization want or not want to conduct: (a) an internal accident/incident investigation, (b) an independent investigation, and (c) a forensic investigation?
4. What is determined during an accident/incident investigation?
5. How are sources of evidence located? Give some examples.
6. How are IA analysis and verification techniques used during an investigation?
7. Describe the similarities and differences between event and causal factor charting and the STEP investigation system.
8. Describe the similarities and differences between warning time analysis and time/loss analysis.
9. Why should accident/incident reports be filed?
10. When are accident/incident reports filed? Who are they submitted to?
11. Create an accident/incident report for the most recent serious anomaly or near-miss experienced by your organization. Include a warning time analysis report and a TLA graph.
12. Why is it important to delineate fact, opinion, theory, and assumptions during an accident/incident investigation?
13. What is the first priority during accident/incident recovery?
14. What is the key to a quick and successful recovery effort?
15. Which accidents/incidents are avoidable?

16. Is the return on investment from investigating incidents worthwhile?
17. Who has the legal responsibility and liability for ensuring that the following are safe, secure, and reliable: (a) a system, (b) a component, (c) an external entity, (d) COTS products, and (e) a service?
18. Explain the connection, if any, between fitness for purpose and assumption of risk.