

Date: 21 January 2017

## 6.37 Fault Tolerance and Failure Strategies [REU]

### 6.37.1 Description of application vulnerability

Check that the current writeup works how.

AI - to Erhard to rework this vulnerability to focus not on fault tolerance itself, but on vulnerabilities caused by it.

In spite of the best intentions, system components may fail, either from internally poorly written software or external forces such as power outages/variations, radiation or inadmissible user input.

Reasons for failures are plentiful and varied, stemming from both hard- and software. Hence the mechanisms of primary failure can be described only in very general terms:

- omission failures: a service is asked for but never rendered. The client might wait forever or be notified about the failure (termination) of the service.
- commission failures: a service initiates unexpected actions, e. g., communication that is unexpected by the receiver. The service might wait forever, causing omission failures for subsequent calls by clients. The receiver might be hindered to do its legitimate actions in time. At a minimum, resources are consumed that are possibly needed by others.
- timing failures: a service is not rendered before an imposed deadline. System responses will be (too) late, causing corresponding damages to the real world affected by the system.
- Value failures: a service delivers incorrect or tainted results. The client continues computations with these corrupted values, causing a spread of consequential application errors.

Faults are the points in execution where a failure manifests by processing going wrong. If unnoticed or unhandled, they turn into failures at the boundaries of enclosing control units or components. Failures of services are faults to their clients and, if not handled, lead to a failure of the client and consequently to faults and failures in its clients, possibly until the entire system fails.

Detection and handling of faults constitutes the fault tolerance code of the system. The mechanisms of fault tolerance are manifold, corresponding to the nature of the failure and the needs of the application, and range from recovery with subsequent normal continuation of the system ("full fault tolerance") or restricted continuation ("graceful degradation", "fail soft") to termination of the system ("fail stop", "fail safe", "fail-secure"), possibly combined with a subsequent restart.

As such, fault tolerance is itself a potential source of vulnerabilities, particularly when inappropriate or incomplete strategies are implemented. Fault-handling code is difficult to design and program, since it needs to execute in an already damaged environment. Handler code is also difficult to test, since it is executed only when primary failures have occurred. These failures, e.g. radiation damage, may be impossible to recreate with sufficient coverage in a testing environment. Moreover, it is not

Stephen Michell 2017-1-21 10:12 PM  
Formatted: Right

Stephen Michell 2017-1-21 8:49 PM  
**Comment [1]:** Failure strategy is both an application vulnerability and a programming language vulnerability. The programming language pieces, exception handling, concurrency issues, etc. are fairly well covered. The design of fault tolerance strategies is probably for section 7.

Microsoft 2017-1-21 7:51 PM  
Formatted: Font color: Red

Microsoft 2017-1-21 7:51 PM  
Formatted: Font color: Red

Stephen Michell 2017-1-21 8:49 PM  
Deleted: .

Microsoft 2017-1-21 7:51 PM  
Formatted: Font color: Red

Microsoft 2017-1-21 7:51 PM  
Deleted: .

Microsoft 2017-1-21 6:11 PM  
Moved (insertion) [1]

Microsoft 2017-1-21 6:11 PM  
Deleted: .

Microsoft 2017-1-21 6:11 PM  
Deleted: s

Microsoft 2017-1-21 6:11 PM  
Deleted: encounter a

Microsoft 2017-1-21 6:11 PM  
Deleted: ure

Microsoft 2017-1-21 7:03 PM  
Formatted: List Paragraph, Bulleted +  
Level: 1 + Aligned at: 0.63 cm + Indent at:  
1.27 cm

easy to determine the right kind of fault tolerance for a given fault. For security, termination of the malfunctioning system may be the best action; for safety, termination may be more catastrophic than any other fault tolerance mechanism.

Arising vulnerabilities are, for example:

- The fault is not recognized and the system malfunctions or terminates as a consequence
- The fault is recognized but the damage already done is incompletely repaired, with the same consequences as in the first bullet
- A value fault is recognized too late, allowing the incorrect value to be used in the computations of other, thus corrupted, values (which, if not repaired, can cause vulnerabilities such as buffer overflows)
- The fault tolerance processing takes too long to meet timing demands
- Recovery is prevented by the cause of a permanent fault, e.g., a programming error, leading to an infinite series of recovery attempts
- The fault tolerance mechanism causes itself new faults

For vulnerabilities caused by termination issues associated with multiple threads, multiple processors or interrupts also see [Error! Reference source not found](#), [Error! Reference source not found](#), and [Error! Reference source not found](#). Situations that cause an application to terminate unexpectedly or that cause an application to not terminate because of other vulnerabilities are covered in those vulnerabilities. The vulnerability at hand discusses the overall fault treatment strategy applicable to single-threaded or multi-threaded programs.

Triggering known fault detection mechanisms can be used to initiate or aggravate Denial-of-Service attacks. Knowledge of a lack of fault detection, particularly of value faults, can be used to initiate system intrusions through mechanisms explained elsewhere in this Tr.

----- drop from here ...

When a fault is detected in a component, there are many ways in which the component can react. The quickest and most noticeable way is to fail hard, also known as fail fast or fail stop. The reaction to a detected fault is then to halt the affected service (or entire system). Alternatively, the reaction to a detected fault could be to fail soft. The system would keep working with the fault present, but the performance of the system would be degraded. Systems used in a high availability environment such as telephone switching centers, e-commerce, or other "always available" applications would likely use such a fail-soft approach, also termed "graceful degradation". Full fault tolerance is achieved when the fault is all but indistinguishable from the normal behavior of the component, e.g. through the use of redundancy. What is actually done in a fail-soft approach can vary depending on whether the system is used for safety-critical or security-critical purposes. For fail-safe systems, such as flight controllers, traffic signals, or medical monitoring systems, there would be no effort to meet normal operational requirements, but rather to limit the damage or danger caused by the fault. A system that fails securely, such as cryptologic systems, would maintain maximum security when a fault is detected, possibly through a denial of service.

Microsoft 2017-1-21 6:35 PM

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.63 cm + Indent at:

Microsoft 2017-1-21 6:12 PM

**Deleted:** floods, or other natural disasters. The reaction to a fault can affect the performance of a system and in particular, the safety and security of the system and its users. - ... [1]

Microsoft 2017-1-21 6:11 PM

**Moved up [1]:** In spite of the best intentions, systems may encounter a failure, either from internally poorly written software or external forces such as power outages/variations, floods, or other natural disasters. The reaction to a fault can affect the performance of a system and in particular, the safety and security of the system and its users. -

Stephen Michell 2017-1-21 10:14 PM

**Deleted:** [Error! Reference source not found](#).

Stephen Michell 2017-1-21 10:14 PM

**Deleted:** 6.61 Concurrency – Directed termination [CGT]

Stephen Michell 2017-1-21 10:14 PM

**Deleted:** 6.63 Concurrency – Premature Termination [CGS]

Stephen Michell 2017-1-21 10:14 PM

**Deleted:** [Error! Reference source not found](#).

Microsoft 2017-1-21 8:02 PM

**Formatted:** Font color: Red

Microsoft 2017-1-21 6:55 PM

**Deleted:** When the software unexpectedly fails to render a requested service or terminates in an unspecified way, safety or security may be compromised. In safety-related systems the results can be catastrophic: for other systems the result can mean failure of the complete system. Failures need not necessarily cause the termination of the failing service; delivering an incorrectly computed result is a failure that, when not discovered, can have even more catastrophic consequences than a termination of the failing service. - ... [2]

Microsoft 2017-1-21 6:50 PM

**Moved (insertion) [3]**

Microsoft 2017-1-21 6:50 PM

**Moved up [3]:** termination issues associated with multiple threads, multiple processors or interrupts also see [Error! Reference source not found](#). 6.61 Concurrency – Directed termination [CGT] and 6.63 Concurrency – Premature Termination [CGS] [Error! Reference source not found](#). Situations that cause an application to terminate unexpectedly or that cause an application to not terminate because of other vulnerabilities are covered in those vulnerabilities. The vulnerability at hand discusses the overall fault treatment strategy applicable to single-threaded or multi-threaded programs.

Microsoft 2017-1-21 6:59 PM

**Deleted:** - ... [3]

Microsoft 2017-1-21 7:00 PM

**Moved down [4]:** Numerous checks on values can and should be made (value range, plausibility within history, reversal checks, checksums, structural checks, etc.) to establish the validity of computed results or input received. Similarly, crucial timing failures should be detected by ... [4]

Whatever the failure or termination process, the termination of an application should not result in damage to system elements that rely upon it. Thus, it should perform "last wishes" to minimize the effects of the failure on enclosing components (e.g., release software locks) and the real world (e.g. close valves).

The reaction to a detected fault in a system can depend on the criticality of the portion in which the fault originates. When a program consists of several tasks, each task may be critical, or not. If a task is critical, it may or may not be restartable by the rest of the program as a fault handling measure. A task that detects a fault within itself but must leave the fault handling to a higher authority, should be able to halt leaving its resources available for use by the rest of the program, halt clearing away its resources, or halt the entire program. The latency of task termination and whether tasks can ignore termination signals should be clearly specified.

----- ... to here ----->

### 6.37.2 Cross reference

JSF AV Rule: 24

MISRA C 2012: 4.1

MISRA C++ 2008: 0-3-2, 15-5-2, 15-5-3, and 18-0-3

CERT C guidelines: ERR04-C, ERR06-C and ENV32-C

Ada Quality and Style Guide: 5.8 and 7.5

### 6.37.3 Mechanism of failure

Reasons for failures are plentiful and varied, stemming from both hard- and software. Hence the mechanisms of failure from fault tolerance or the lack thereof can be described only in very general terms:

- [Fault tolerance code, in particular fault checking code, may interfere with the timeliness of the components to meet their deadlines](#)
- [An inappropriate fault tolerance mechanism or strategy may lead to failures in fault detection and other secondary failures](#)
- [Considerable latency and processor use can arise from finalization and garbage collection caused by the termination of a task. Thus, termination must be designed carefully to avoid causing timing failures of other tasks. The termination of tasks can be maliciously used to prevent on-time performance of other active tasks.](#)
- [Having inconsistent approaches to detecting and handling a fault or a lack of overall design for the fault tolerance code can potentially be a vulnerability, as faults might escape the necessary attention.](#)
- [If faults are not detected in time and repaired completely, the following failures arise:](#)
  - [omission failures: a service is asked for but never rendered. The client might wait forever or be notified too late about the failure \(termination\) of the service.](#)
  - [commission failures: a service initiates unexpected actions, e.g., communication that is unexpected by the receiver. The service might wait forever, causing omission failures for subsequent calls by clients, or the actions might interfere with the regular](#)

ploedere 2016-8-15 7:01 PM  
**Formatted:** Font:Not Italic

Microsoft 2017-1-21 8:03 PM  
**Formatted:** Font color: Red

Microsoft 2017-1-21 7:01 PM  
**Deleted:** -

ploedere 2016-8-15 6:11 PM  
**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.63 cm + Indent at: 1.27 cm

Microsoft 2017-1-21 7:36 PM  
**Moved (insertion) [5]**

Microsoft 2017-1-21 7:36 PM  
**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.63 cm + Indent at:

Microsoft 2017-1-21 7:36 PM  
**Deleted:** -

ploedere 2016-8-15 6:11 PM  
**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.63 cm + Indent at:

Microsoft 2017-1-21 7:33 PM  
**Formatted:** List Paragraph, Bulleted + Level: 2 + Aligned at: 1.9 cm + Indent at: 2.54 cm

[processing going on in the meantime. At a minimum, it consumes resources possibly needed by others to meet deadlines.](#)

- [timing failures: a service is not rendered before an imposed deadline. System responses will be \(too\) late, causing corresponding damages to the real world affected by the system.](#)
- [Value failures: a service delivers incorrect or tainted results. If not the client continues computations with these corrupted values, causing a spread of consequential application errors and implementation vulnerabilities caused by corrupted values as discussed elsewhere in this TR.](#)

### 6.37.4 Applicable language characteristics

This vulnerability description is intended to be applicable to all languages.

### 6.37.5 Avoiding the vulnerability or mitigating its effects

Software developers can avoid the vulnerability or mitigate its ill effects in the following ways:

- Decide on a strategy for fault handling. Consistency in fault handling should be the same with respect to critically similar parts.
- Use a multi-tiered approach of fault prevention, fault detection and fault reaction.
- Unambiguously describe the failure modes of each possibly failing task as fail-stop, fail-safe, fail-secure, or fail-soft as explained in 6.37.1.
- [Check early for any faults, particularly value faults](#)
- [Numerous checks on values can and should be made \(value range, plausibility within history, reversal checks, checksums, structural checks, etc.\) to establish the validity of computed results or input received.](#)
- [Timing failures should be detected by "Watch-dog timers" and similar mechanisms that can be used to stop rogue tasks.](#)
- [Always validate incoming data. Validate computed results at strategic points to discover value failures. See also pre- and postconditions in << reference to BLP, Liskov>>.](#)
- [Use environment- or language-provided means to stop tasks that substantially exceed deadlines.](#)
- [Always prepare for the possibility that a service does not return with a requested result in due time.](#)
- [Keep fault handling simple. If in doubt, decide for a lesser level of fault tolerance.](#)
- [In the case of continued execution, make sure that any corrupted variables of the program state have been corrected to an actual and correct or at least safe value.](#)
- System-defined components that assist in uniformity of fault handling should be used when available. For one example, designing a "runtime constraint handler" (as described in Annex K of 9899:2012 [4]) permits the application to intercept various erroneous situations and

Microsoft 2017-1-21 7:18 PM  
**Deleted:** T

Microsoft 2017-1-21 7:18 PM  
**Deleted:** .

Microsoft 2017-1-21 6:28 PM  
**Deleted:** <#>Faults are the points in execution, where a failure manifests by processing going wrong. If unnoticed or unhandled, they turn into failures at the boundaries of enclosing control units or components. Failures of services are faults to their clients and, if not handled, lead to a failure of the client and consequently to faults and failures in its clients, possibly until the entire system fails. Detection and handling of faults constitutes the fault tolerance code of the system. As such, it is itself a potential source of failures. Fault-handling code is particularly difficult to design and program, since it ne... [5]

Microsoft 2017-1-21 6:28 PM  
**Moved up [2]:** Faults are the points in ex... [6]

Microsoft 2017-1-21 6:28 PM  
**Moved (insertion) [2]**

Microsoft 2017-1-21 7:36 PM  
**Moved up [5]:** Considerable latency an... [7]

ploedere 2016-8-15 6:51 PM  
**Formatted:** Font:Not Italic

Microsoft 2017-1-21 7:35 PM  
**Formatted** ... [8]

ploedere 2016-8-15 6:51 PM  
**Formatted:** Font:Not Italic

Microsoft 2017-1-21 7:35 PM  
**Formatted:** Font:Not Italic

ploedere 2016-8-15 6:38 PM  
**Deleted:** The reaction to a fault in a system... [9]

ploedere 2016-8-15 6:55 PM  
**Deleted:** .

ploedere 2016-8-15 6:56 PM  
**Deleted:** 3

Microsoft 2017-1-21 7:38 PM  
**Formatted:** Normal

Microsoft 2017-1-21 7:00 PM  
**Moved (insertion) [4]**

Microsoft 2017-1-21 7:39 PM  
**Deleted:** Similarly, crucial t

Microsoft 2017-1-21 7:38 PM  
**Deleted:** . ... [10]

Microsoft 2017-1-21 7:38 PM  
**Formatted**

Microsoft 2017-1-21 7:40 PM  
**Deleted:**

Microsoft 2017-1-21 7:41 PM  
**Deleted:**

perform one consistent response, such as flushing a previous transaction and re-starting at the next one. << is this example appropriate? >>

- [Prior to any abnormal termination of a component, perform “last wishes” to minimize the effects of the failure on enclosing components \(e.g., release software locks held locally\) and the real world \(e.g. close valves opened by the component\).](#)
- [Specify a fault-handling policy whereby a task, in the absence of simple full fault tolerance or graceful degradation, may](#)
  - Halt, and keep its resources available for other tasks (perhaps permitting restarting of the faulting task).
  - Halt, and release its resources (perhaps to allow other tasks to use the resources so freed, or to allow a recreation of the task).
  - Halt, and signal the rest of the program to likewise halt.

<<< I consider this last advice a bit too specific to one particular model of execution. In fact, I disagreed with the original, since it excluded full fault tolerance altogether. simplify to “kill everything or do the right thing about resources” ? i.e. incorporate in “last wishes” above? >>>

### 6.37.6 Implications for standardization

In future standardization activities, the following items should be considered:

- Languages should consider providing a means to perform fault handling. Terminology and the means should be coordinated with other languages.

Microsoft 2017-1-21 8:03 PM

**Formatted:** Font color: Red

Microsoft 2017-1-21 8:03 PM

**Formatted:** Font color: Red

Microsoft 2017-1-21 7:43 PM

**Formatted:** Font:

Microsoft 2017-1-21 7:43 PM

**Formatted:** List Paragraph

Microsoft 2017-1-21 7:47 PM

**Deleted:** When there are multiple tasks,

Microsoft 2017-1-21 7:48 PM

**Deleted:** should be specified

Microsoft 2017-1-21 8:04 PM

**Formatted:** Font color: Red

ploedere 2016-8-15 7:40 PM

**Formatted:** Indent: Left: 1.9 cm, No bullets or numbering

Microsoft 2017-1-21 8:04 PM

**Formatted:** Font color: Red