Sanitizing Sensitive Data: How to get it Right (or at least Less Wrong…)[1]

Roderick Chapman (orcid.org/0000-0003-2717-760X)

Protean Code Limited, Bath, UK
rod@proteancode.com

Abstract. Coding standards and guidance for secure programming call for sensitive data to be "sanitized" before being de-allocated. This paper considers what this really means in technical terms, why it is actually rather difficult to achieve, and how such a requirement can be realistically implemented and verified, concentrating on the facilities offered by Ada and SPARK. The paper closes with a proposed policy and coding standard that can be applied and adapted to other projects.

Keywords: Security, Sanitization, SPARK, Verification, Volatile, Optimization, Proof.

1 The Problem

Secure systems must be built to resist attack by increasingly sophisticated adversaries. An attacker might be able to observe or provoke a system into "leaking" or revealing secret data, such as cryptographic keys, the plaintext of passwords and so on. A well-documented example is where an intruder manages to read the operating system page file or a core dump of a running (or deliberately terminated) process in order to gain access to unsanitized sensitive data.

Several coding standards and guidance documents exist that call for sensitive data to be *sanitized* when no longer needed, but offer little advice on how this is to be achieved or verified, especially given the complexity of programming languages and hardware. This paper considers this problem in detail and describes the key technical challenges, before going on to consider the facilities offered by Ada and SPARK that can meet these demands, based on experience gained from a recent project.

1.1 Why is Sanitizing Data Hard?

Sanitizing sensitive data might seem simple at first: just overwrite the data with zeros and carry on, right? A less trivial analysis reveals important questions, including:

[1] Original paper appears in proceedings of Ada Europe 2017, LNCS 10300.
DOI: 10.1007/978-3-319-60588-3_3. The final publication is available at link.springer.com

- How do we define "sensitive"? What objects in the program are "sensitive" and how are they identified?
- Imagine that we have two variables A and B which are defined to be "sensitive." We declare and initialize a local variable C with an initial value derived from some function that combines A and B. Is C "sensitive"? Does C need to be sanitized?
- Can constant objects be sensitive? If so, how are they to be sanitized?
- Exactly *when* should sanitization be performed, relating to the scope and lifetime of data objects which is, in turn, intricately entwined with a particular programming language's model of how data should be organized and (de-)allocated?
- Compiler optimization might remove a sanitizing assignment if the assignment is seen to be redundant or "dead" by the optimizer. How can this be prevented?
- How do we verify that sanitization really has been performed correctly, to the satisfaction of ourselves, our customers, and regulators?

1.2 Standards, Guidance and Problems

There are several (possibly far too many) sets of guidance or coding rules for secure systems that call for sensitive data to be sanitized as soon as it is no longer needed, so that (for example) a subsequent buffer over-read will not find any useful data. This section considers an incomplete set of these, and tries to point out problems in meeting their advice.

GCHQ. The UK's national technical authority for secure software, GCHQ, offers a short (but thankfully unclassified) "Guidance Note" on secure coding [1]. It offers some generic advice, but mainly consists of coding rules for C and C++. A need to avoid copying sensitive data is mentioned (to avoid a copy existing even if the original is sanitized), with two paragraphs specifically on sanitization:

> "Sanitise all variables that contain sensitive data (such as cryptovariables and unencrypted data) by overwriting with zeroes once they are no longer needed. This includes all copies of the data: call-by-value functions (as found in C) implicitly copy the value of their parameters, so their parameters should always be sanitised before the function exits. At protective markings higher than Restricted, sanitisation may require multiple overwrites or verification, or both." [1, para 58].

and

> "The sanitisation is needed because errors may result in the disclosure of a block of memory, therefore the risk of that memory containing anything useful needs to be minimised. The size of the data is not a factor: even single bytes need to be sanitised, since in some cases a difference of 8 bits could have a significant impact on the practicality of an attack. On the other hand, the lifetime of data may be a factor: if a variable can be shown to be overwritten shortly afterwards, it may be acceptable not to sanitise it, provided it is sanitised when it is no longer needed. 'Shortly' is not defined more precisely, since it will depend on the situation…" [1, para 59]

This is well-meaning, but offers little in the way of real technical detail of how sanitization is to be achieved or verified. The failure to define "shortly afterwards" is also disappointing.

CERT Coding Standards. The CERT at CMU has produced coding guidance for secure software development, covering C, C++, Java and Perl to date, with several tool vendors claiming compliance. The CERT C Coding Standard [2] provides some advice on sanitization:

- Recommendation 08 (Memory Management), Item 06 is titled "Ensure that sensitive data is not written out to disk" which mostly covers the problem of an operating system "paging out" sensitive data to a disk or an application doing a "core dump" which writes the state of a process to a disk file, potentially revealing the state of sensitive data. These are valid concerns, relevant to any application running on an operating system that supports paging and so on, so not really a "C language issue" per-se, since these problems could affect code written in any language.
- Recommendation 48 (Miscellaneous), Item 06 is titled "Beware of compiler optimizations" and covers the problem of a compiler removing a sanitizing assignment. It goes on to recommend using "optimization safe" C functions such as memset_s(), C's "volatile" qualifier (more of which later…) or operating-system specific functions that are designed to sanitize memory.

Both of these recommendations appear to presume the existence of some sort of operating system (and possibly a "disk"), but what if we're programming an embedded "bare metal" system with no OS at all? How can we sanitize data properly in such an environment?

ISO SC22/WG23 Technical Report 24772. The ISO's SC22 Working Group 23 has produced Technical Report 24772 entitled "Guidance to avoiding vulnerabilities in programming languages through language selection and use." [3] The TR recognizes sanitization as an avoidance mechanism for some vulnerabilities, but does not go into specific details. The language-specific annexes for Ada, C and SPARK offer no additional advice.

Common Weakness Enumeration (CWE). Mitre's CWE [4] includes CWE-14 "Compiler Removal of Code to Clear Buffers" which identifies the risk of removal of sanitizing assignments by optimizing compilers. It advocates the use of volatile objects and suggests "configure your compiler so that it does not remove dead stores."

Cryptography Coding Standard. The Cryptography Coding Standard is "a set of coding rules to prevent the most common weaknesses in software cryptographic implementations" [5]. Their coding rules touch on sanitization in a number of places:

- Coding Rule 5 "Prevent compiler interference with security critical operations" mentions the problem of compilers removing sanitizing assignments, and how even a call to C's standard "memset" function can be optimized away in some cases. It offers the rather vague advice to "Look at the assembly code produced and check that all instructions are there" which hardly seems practical for anything but trivial

code. It also recommends "consider disabling compiler optimizations that can eliminate or weaken security checks" but again this seems impractical – modern compilers have *hundreds* of optimization switches, which makes it almost impossible to "know" which set of them will or won't "interfere" with security. Finally, rule 5 does point out that the 2011 C standard does include a new "memset_s" function, a call to which is explicitly *not* allowed to be optimized.

- Coding Rule 11 "Clean memory of secret data" looks promising, recommending that code should "Clear all variables containing secret data before they go out of scope." It points out the existence of a SecureZeroMemory function in the Win32 API for this purpose. It also offers a portable C function that can be used to overwrite memory that works "for non-buggy compilers" [sic].

1.3 Technical Issues

Having seen that the standards and guidance documents offer well-meaning but imprecise advice, we now turn to a selection of more detailed technical problems.

Unwanted Compiler Optimization. Several of the guidance documents cited above refer to this problem, so it warrants more attention here.

Modern implementations of computer architectures feature a marked difference between the access time of CPU registers, data cache(s), and main memory, sometimes by many orders of magnitude. In short, DRAM access times have not kept pace with CPU clock rates, so the penalty for a "register miss" or a "cache miss" is pronounced. Modern compilers therefore devote significant effort in several, related classes of optimization [6], including:

1. Common sub-expression elimination and partial redundancy elimination. These prevent semantically equivalent expressions from being evaluated more than once.
2. Register allocation and tracking, so that variables and the values of expressions are stored in CPU registers as much as possible.
3. Dead-load and dead-store elimination.

These improve average-case performance, but create some issues for sanitization:

- Guidance calls for the "memory" occupied by a sensitive variable to be overwritten "before the variable goes out of scope", but what does that mean if the variable only ever exists in an internal CPU register and there is no "memory" allocated for it at all?
- A final sanitizing assignment needs to occur just before a variable "goes out of scope", so is (by definition) a "dead store" in the eyes of an optimizer, so might be removed, on the assumption that once a variable has gone out of scope it can't be accessed any more. This creates a conflict in the compiler: we (the programmers) want dead stores to be retained for one or more particular variables, but the compiler is trying its hardest to remove them in the interests of improving performance of the generated code.

Derived Values and Copies. In his thorough analysis "Zeroing buffers is insufficient" [7], Percival points out several more pernicious technical issues with a simple "write zeros into memory" approach. Specifically, he points out:

- Sanitizing the *one* memory block where a variable is stored is not good enough. Compilers implicitly make copies of data into registers or implicitly-declared and initialized local variables, so these might also contain a copy of some sensitive information that needs to be sanitized. In the worst case, a compiler might evaluate the value of a sensitive variable into a CPU register *and* spill that register into an implicitly allocated temporary variable on the stack. There is no way to portably sanitize such temporary variables in C or Ada, since those variables do not appear in the source code.
- If a sensitive piece of data is left in a CPU register, you *cannot* assume that that CPU register will be re-used and the data over-written "quickly". Percival points out that some CPU registers (such as the SSE registers on x86) are rarely used, and some registers are specifically designed for cryptographic algorithms such as AES – the problem being that you carefully use a "special" register to hold an AES key (for example), but then that register is not used for anything else in your program, so the key value persists and is never overwritten. Secondly, some CPUs such as x86 can implement register renaming, which further complicates matters.

A related problem is that of derived values. As pointed out in Section 1.1, if two sensitive variables A and B are combined in some way to get a value in variable C, should C be considered to be sensitive and therefore needing sanitization? The answer is "it depends"… on the exact operation used to derive C, the nature of A and B, and so on. It is far from simple to suggest a generic one-size-fits-all policy for such variables.

By-Copy Parameter Passing. If a subprogram parameter is passed by copy, then the value of the actual parameter is copied into the storage associated with the formal parameter (which might be stack memory or a CPU register.) If the actual is sensitive, then so is the formal parameter. In Ada, this is particularly problematic, since "in" mode parameters are constant and so cannot be assigned to at all, and the choice between by-copy and by-reference passing can be unspecified for some types.

CPU Data Caching and Memory Hierarchy. Anyone that has programmed a device-driver on a "bare metal" target will know that the presence of a "write" instruction does *not* guarantee that the data actually reaches the target hardware device at all, or in the order indicated in the source code. Modern CPUs have multiple levels of data caching, which may be in "write back" mode, so an instruction to write a particular word of memory might not actually reach the main memory device until the offending data cache line is flushed or invalidated. Secondly, modern CPUs can execute instructions out-of-order and re-order memory accesses in rather unexpected ways, which can complicate matters further.

Some operating systems offer functions that are specifically designed to securely sanitize memory, such as Win32's SecureZeroMemory function. We presume these functions take care of any required flushing of caches, paged-out data and so on.

On bare-metal targets, we might turn off all data caching or insist on "write through" mode, but this may be Draconian, since disabling all caching for all stack-allocated data would incur a potentially huge performance penalty. Some CPUs might allow special instructions to flush particular cache lines and so-called "memory fence" instructions that instruct the CPU to pause until all queued memory accesses are complete. These techniques are valid (and indeed may be absolutely necessary), but require recourse to obviously non-portable assembly language programming at some level.

The recent 2011 editions of both the C [8] and C++ [9] languages have been extended to define an abstract "memory model" for these languages, plus support in the standard library for atomic types and fence operations, both of which may offer mechanisms that support sanitization more portably.

1.4 An Example – How it Can Go Wrong in Ada

This section closes with a short (and somewhat contrived) example of how sanitization can fail in Ada. In the remainder of the paper, all examples have been compiled with the GPL 2016 Edition of GNAT for 32-bit x86 running on Windows 7 Pro.

Consider a simple procedure GK that takes three seed values A, B, and C, and produces a derived key value K from them. For example:

```
subtype Word32 is Interfaces.Unsigned_32;
procedure GK (A, B, C : in     Word32;
              K       :    out Word32);
```

The body of GK combines A, B, and C using a local, temporary variable T which we have decided is sensitive and needs to be sanitized with a final assignment, thus:

```
procedure GK
  (A, B, C : in     Word32;
   K       :    out Word32)
is
   T : Word32;
begin
   T := A xor B; -- line 15
   T := T xor C;
   K := T;

   -- Now sanitize T
   T := 0; -- line 20
end GK;
```

To see what's going on, we'll compile with both "-g" and "-fverbose-asm" flags. We'll also enable all warnings with "-gnatwa" and "-Wall" as we would on any real project. Compiling GK does yield a warning:

```
p1.adb:20:07: warning: useless assignment to "T", value never
referenced
```

which hints at trouble ahead. Compiling with –O0 (little or no optimization) yields the following assembly language for lines 15 through 20 of GK:

```
movl   8(%ebp), %eax    # a, tmp88  LINE 15
xorl   12(%ebp), %eax   # b, tmp87
movl   %eax, -12(%ebp)  # tmp87, t
movl   16(%ebp), %eax   # c, tmp89  LINE 16
xorl   %eax, -12(%ebp)  # tmp89, t
movl   -12(%ebp), %eax  # t, tmp90  LINE 17
movl   %eax, -16(%ebp)  # tmp90, k
movl   $0, -12(%ebp)    #, t  LINE 20
```

so we can see the final assignment to T on line 20 has indeed been generated as a single "movl" instruction.

Turning on the optimizer at level "-O1" reveals a different story. For the same fragment of code, we get:

```
movl   16(%ebp), %eax   # c, c
xorl   12(%ebp), %eax   # b, D.3010
xorl    8(%ebp), %eax   # a, k
```

and that's all. The local variable T is not allocated at all on the stack – it has completely disappeared, in fact, with the intermediate results left in the CPU register EAX. Our well-intended attempt to sanitize T has been discarded by the compiler, but then again, T has disappeared entirely, so is this sufficient? What about the intermediate value left in EAX – is that overwritten "soon" by the calling subprogram perhaps?

2 Sanitization – Constraints and Goals

In developing the coding standard for a recent project, we had to meet both CESG's guidance for sanitization [1], but also the constraints imposed by the wider demands of the project, including the runtime environment, compilers, features of the target platform and its operating system and so on.

In searching for the most general solution, we tried to respect the following constraints:

1. The approach to sanitization should minimize dependence on predefined library units and the use of language features that require substantial support from the Ada runtime library. In particular, for our project, we required compatibility with GNAT's "Zero Footprint" (ZFP) runtime library.
2. The approach should not depend on any operating system facilities, and so can be deployed on a "bare metal" target system.
3. The approach should be compatible with the SPARK language (either SPARK 2005 [10] or SPARK 2014 [11,12]) and verification tools.

Secondly, what does a "good" approach to sanitization look like? In developing these guidelines for Ada, we tried to respect the following goals:

1. Any proposed approach should be *portable* in that it should not depend on non-standard behavior from the compiler, and should not rely on particular *unspecified* or *implementation-defined* choices made by a compiler.
2. Our approach should permit compiler optimizations to be enabled at all levels, with sufficient confidence that sanitization code would be preserved and implemented correctly.
3. Our approach must prevent (as far as is possible) explicit or implicit copying or assignment of sensitive values. This also affects parameter passing, since a "by-copy" formal parameter involves assignment.
4. Our approach should facilitate (or at least not obstruct) verification with the SPARK toolset.
5. Our approach should meet or exceed the demands of the various regulatory standards, such as [1]. Furthermore, we should be able to explain and justify our approach to those regulators so that we can convince them that it actually works.

3 Sanitization Mechanisms in Ada

Having considered the scope of this problem, this section turns to the language-based mechanisms that are available in Ada. Knowing what mechanisms are available can then lead to a policy that can be adopted for a particular project.

3.1 Volatile

Ada, C and C++ all include a facility to mark an object as "Volatile", meaning that the compiler must respect the exact sequence of reads and writes to such an object that are indicated in the source code. Ada goes further, allowing Volatile *types* as well as objects. The Ada RM [13] offers a clear implementation requirement (Ada 2012 RM, C.6(20)):

"The external effect of a program…is defined to include each read and update of a volatile or atomic object. The implementation shall not generate any memory reads or updates of atomic or volatile objects other than those specified by the program."

Let's see what happens to our example procedure GK with the declaration of T changed as follows:

```
T : Word32 with Volatile;
```

With that in place, we should be able to turn the optimizer "up to 11" (well…3) and compile with "-O3". Firstly, the warning from the front-end about the useless assignment to T disappears, which is a good sign. The generated code for lines 15 – 20 is:

```
movl   12(%ebp), %eax   # b, b  LINE 15
xorl    8(%ebp), %eax   # a, D.3014
movl   %eax, -12(%ebp)  # D.3014, t
movl   -12(%ebp), %eax  # t, D.3015  LINE 16
xorl    16(%ebp), %eax  # c, D.3014
movl   %eax, -12(%ebp)  # D.3014, t
movl   -12(%ebp), %eax  # t, k  LINE 17
movl   $0, -12(%ebp)    # t  LINE 20
```

so we see that *all* the reads and writes of T have been preserved, including the final sanitizing assignment.

At first glance, this appears to be a perfect match, at least when it comes to preventing the optimization of sanitizing assignments. Unfortunately, it's not that simple for several reasons:

1. Volatile prevents optimization of *all* reads and writes to an object, but we only require that the *final* sanitizing assignment is preserved, so use of Volatile might have a serious but unnecessary impact on the performance of the generated code.
2. SPARK 2014 (release 16.0.2) only permits library level objects to be declared Volatile. Local variables may not be Volatile.
3. Most seriously and worryingly, Regehr and Eide [14] have shown that compilers can mis-compile Volatile and *do* optimize away reads and writes when they shouldn't. Regehr and Eide only tested 13 compilers and their work dates from 2008 so we hope compilers have improved since then. Their tests were based on analysis of C programs, but their concerns are real, especially since their results include those for 9 builds of GCC, which shares its back-end (and optimization code) with GNAT.

So, despite its initial good looks, the use of Volatile is not a panacea for data sanitization. Secondly, it does not address the need to restrict assignment and copying of sensitive data objects at all.

3.2 Controlled Types

Ada's "Controlled Types" offer a tempting approach to supply a "Finalize" procedure that sanitizes an object. At first glance, this seems attractive, but there are several serious problems:

- Controlled types require significant support from the Ada runtime library which conflicts with our requirement for compatibility with the ZFP runtime.
- They are not permitted by SPARK.
- Their semantics and implementation are notoriously difficult to understand [15, 16].

In light of these problems, controlled types were rejected without further investigation.

3.3 Limited Types

Ada's limited types are particularly attractive for holding sensitive data. Firstly, the programmer can have complete control over exactly what set of operations are available to clients. Secondly, and by default, assignment is not defined for limited types, so we can control both copying and creation of derived values. Finally, an explicitly limited record type is defined to be a *by-reference* type (RM 6.2(7)) so we can be sure that all formal parameters of such a type will be passed by reference, not by copy.

3.4 By-Reference Types

Where the use of a limited record type is not appropriate or practical, there are still other means of forcing a type to be a "by-reference" type in Ada, which will, at least, prevent copying by parameter passing where we don't want it. RM C.6 (18) tells us that if any sub-component of a type is Atomic or Volatile, then the type is defined to be a by-reference type. Additionally, RM 6.2(5) specifies that all tagged types are by-reference. Thus we can force by-reference passing for even a simple scalar type by wrapping it in a tagged record or a record which has a single Atomic or Volatile component. For example, instead of declaring a formal "in" mode parameter of type Boolean, we might declare:

```
type Sensitive_Boolean is tagged record
   F : Boolean;
end record;
```

or

```
type Sensitive_Boolean is record
   F : Boolean with Volatile;
end record;
```

to ensure by-reference parameter passing. There are pros and cons to both approaches. The Volatile field has no space overhead and makes the field volatile, but is not compatible with SPARK 2014 at the time of writing. The tagged record is allowed by SPARK, but imposes some space overhead by adding an implicit tag field to the record.

Using GNAT, it is also possible to *verify* the parameter passing mechanism using the "–gnatRm" flag.

3.5 Pragma Inspection_Point

This little-used (and little-understood perhaps) pragma has particular relevance to this problem. Inspection_Point was introduced in Ada 95 as part of the RM's Safety and Security Annex H. It is designed to specify a list of objects that must be *inspectable* at a particular point in a program. A pragmatic interpretation means that the listed objects are supposed to be stored in memory at the inspection point so that their values

can be seen by external means, such as a logic analyser, a JTAG probe, a real-time debugger or similar. From the point of view of optimization, the Ada RM is clear:

'The implementation is not allowed to perform "dead store elimination" on the last assignment to a variable prior to a point where the variable is inspectable. Thus an inspection point has the effect of an implicit read of each of its inspectable objects.' (Ada RM H3.2 (9)).

This seems ideal for our needs – if a final, sanitizing assignment to a sensitive object is immediately followed by a pragma Inspection_Point for that object, then that final assignment should not be optimized away. This provides much finer control than pragma Volatile. For the curious, GNAT actually implements pragma Inspection_Point by generating a dummy volatile read to each of the objects specified in the pragma. See the file gcc-interface/trans.c in the GNAT sources for details [17] and search the file for "Inspection_Point".

Returning to our simple example, we revert to declaring T as a normal (non-volatile) local variable, but now follow the final assignment with an Inspection_Point, thus:

```
-- Now sanitize T
T := 0; -- line 20
pragma Inspection_Point (T);
```

The generated code at –O1 is:

```
# 21 "p3.adb" 1
# inspection point: t is in $0  #
# 0 "" 2
movl  12(%ebp), %eax  # b, b
xorl  16(%ebp), %eax  # c, D.3010
.loc 1 16 0
xorl  8(%ebp), %eax   # a, k
```

which is interesting. Again, the variable T has been entirely eliminated, but commentary has been added that "t is in $0" since T does not have an accessible address in memory at all.

3.6 No_Inline and Sanitizing Operations

Having identified the problems with Volatile objects, Regehr and Eide go on to recommend that all reads and writes of a volatile variable should be performed by a subprogram call that can never be inlined, since inlined code has the potential to be optimized away during the compilation of any calling units. They demonstrate how this works well for C, and the equivalent mechanism exists for Ada with the GNAT-defined pragma No_Inline.

Combining this idea with the use of a limited private type for sensitive data yields the following pattern for a sensitive abstract data type:

```
package Sensitive is
   type T is limited private; -- so no assignment

   procedure Sanitize (X : out T);
   pragma No_Inline (Sanitize);
private
   type T is limited record -- so by-reference
       F : … -- and so on…
   end record;
end Sensitive;
```

The body of Sensitive.Sanitize might depend on the target platform and operating system, so we recommend implementing it as a separate subunit of package Sensitive to allow for alternative implementations to be chosen at build-time. Let's imagine that the field F of type T is of type Word32. In that case, a suitable implementation for a bare-metal/ZFP target might be:

```
separate (Sensitive)
procedure Sanitize (X : out T) is
begin
   X.F := 0;
   pragma Inspection_Point (X);
end Sanitize;
```

At –O3, the generated code for the assignment statement and the pragma is:

```
movl   8(%ebp), %eax # x, x
movl   $0, (%eax)  #, x_2(D)->f
# 6 "sensitive-sanitize.adb" 1
# inspection point: x address is in %eax # x
# 0 "" 2
```

We can also check the parameter passing mechanism using –gnatRm which yields:

```
procedure sanitize declared at sensitive.ads:5:14
  convention : Ada
  x : passed by reference
```

4 Verification and SPARK

The SPARK toolset offers two major forms of static verification—information-flow analysis and proof of user-defined contracts. This section briefly considers the interplay between sanitization and these forms of verification.

4.1 Information Flow Analysis

As expected, both the SPARK 2005 and SPARK 2014 tools will reliably report that a final sanitizing assignment to a local variable is *ineffective*, meaning that the assignment has no influence on the final value of any exported variable of the subprogram under analysis. This is perfectly correct and reasonable. At first, such errors being

reported might seem an annoyance, we can turn this to our advantage using pragma Warnings to document the expectation and need for the sanitization. For our earlier example, we would add:

```
pragma Warnings (Off, "unused assignment",
                 Reason => "Sanitization");
T := 0;
pragma Inspection_Point (T);
```

4.2 Proof

At first glance, it might be possible to prove that sanitization of variables has been performed, but closer inspection reveals two main issues:

- The final value of a *local* variable cannot be asserted in the post-condition of the subprogram that declares it, owing to its very local-ness. It would be possible to assert the value of a sanitized library-level variable.
- Writing an assertion regarding the *value* of a sensitive variable means that we need to decide on a (constant) value that should be used. The naive approach of "zero all bits" might not be appropriate, since "all zeros" might not be a valid value. SPARK and Ada have no "memset" or similar, so we need to be able to write an assignment statement which is legal and itself free from runtime errors.

5 A Policy for Sanitization

In light of the difficulties described above, and the facilities offered by Ada and SPARK, and our experience on one project, we would offer the following policy for sanitization of sensitive data for future work.

5.1 Identification and Naming of Sensitive Variables

A project must document a clear policy for what exactly is and isn't considered to be a "sensitive" object. This is clearly project- and application-specific. In cryptographic applications, for example, sensitive data might include cryptographic keys, single-use random "nonce" values, and initialization vectors for encryption algorithms.

The definition of "sensitive" may also have to consider the visibility and lifetime of the objects—local variables and library level states might have to be treated very differently, for example.

Having chosen a policy for deciding which states are sensitive, we propose a naming convention as follows:

- The names of *types* used for sensitive data should be prefixed with "Sensitive_".
- The names of *variables* that are sensitive should have the suffix "_SAN" meaning that such variables should be sanitized.
- The name of a formal subprogram parameter that *might* be associated with a sensitive actual parameter shall also have the suffix "_SAN".

- Sensitive *constants* are not permitted.

5.2 Types and Patterns for Sensitive Data

- By-reference types should be used for all sensitive data.
- Preferably, and if possible, a limited type should be used for sensitive data to forbid assignment. In this case:
  - A "Sanitize" procedure should be supplied, as shown in section 3.6, which has a No_Inline pragma applied to it.
  - The body of such a "Sanitize" procedure should be a separate subunit to allow for multiple implementations for different platforms and operating systems.
  - The body of "Sanitize" shall include a pragma Inspection_Point immediately following the final assignment to the formal parameter. Note that the presence of the pragma is sufficient to suppress the "useless assignment" warning illustrated in section 1.4. This is useful for verification, since presence of this warning is a strong indication that the programmer has forgotten to add the pragma.
- If a limited type is not possible, then a Sanitize procedure shall still be supplied for any sensitive type, implemented as above. In this case, code review checklists must include a check that assignment is not used for objects of such types.
- For SPARK code, a pragma Warnings shall always precede a final sanitizing assignment (or the call to a Sanitize procedure) to document the need for the sanitization and to suppress the information-flow warning.

5.3 Compiler Switches and Analysis

- All code should be compiled with "-gnatwa" to ensure that the "useless assignment" warning is generated. This should be expected for sanitizing assignment, but suppressed with pragma Inspection_Point.
- The "-gnatRm" switch should be used to verify that the compiler has chosen by-reference parameter passing mechanism for all sensitive formal parameters. This is easy if the naming convention above has also been followed.
- Analysis of the generated assembly language should be performed using the "-g" and "-fverbose-asm" flags to verify that the inspection points are present and correct.
- Additional analysis of the generated code might be required to verify that cache manipulation instructions and memory fences are as required.

6 Related and Further Work

Several authors have called for compilers to help automate sanitization via some sort of special compilation switch ("-fsanitize_local_data" perhaps?). This could go further than source-based techniques since the compiler could arrange to sanitize *all* local states, derived variables, temporaries, and CPU registers for example. How a

compiler designer would convince others of the correctness of such an approach remains unknown.

A compiler switch seems a rather blunt instrument though. Sanitizing *all* local data might produce an unacceptable performance overhead, so we return to the idea of how objects in a program can be marked as sensitive and therefore requiring sanitization. We might imagine a new "Sensitive" aspect in Ada 2012 that can be applied to types and/or objects, rather like Volatile.

A standardized Ada binding to the C11 "stdatomic" library might be a useful exercise to supply portable access to memory fence operations.

Another compiler-related issue is that of link-time optimization (LTO). This style of optimization has appeared recently in compilers like GCC [18] and LLVM [19]. Studies are needed to verify that sanitization code is preserved in the presence of LTO.

There has been significant interest in the verification of compilers, particularly owing to the CompCert effort [20, 21]. The proof of CompCert covers the correct compilation of Volatile objects, which could carry over to the correctness of sanitizing assignments and inspection points.

The problem of sensitive derived variables could be addressed through more advanced information flow analysis. If a tool like GNATProve, for example, knew that variables A and B were sensitive, then could it automatically infer that C (derived from A and B) were also sensitive? This can also be seen as a variant of the taint analysis embodied in languages like Ruby and Perl.

7 Conclusions

Sanitization of sensitive data remains a thorny issue: standards call for it to be done, but offer little advice on how it should be achieved in practice or verified. This paper has illustrated some of the problems and shown how they can be addressed in Ada and SPARK and developed into a policy, coding standard, and verification strategy for a particular project.

Acknowledgements. The author would like to thank Robert Seacord, Florian Schanda, Bill Ellis and the conference reviewers for their comments on earlier drafts of this paper.

References

1. CESG. Coding Requirements and Guidance (IA Developers' Note 6), CESG, Issue 1.1, October 2015. www.ncsc.gov.uk/guidance/coding-requirements-and-guidance-ia-developers-note-6
2. US CERT. SEI CERT C Coding Standard. www.securecoding.cert.org/confluence/display/c/SEI+CERT+C+Coding+Standard

3. ISO/SC22/WG23. Information Technology — Programming Languages — Guidance to avoiding vulnerabilities in programming languages through language selection and use. TR 24772:2013. http://www.open-std.org/JTC1/SC22/WG23/
4. Mitre Corp. Common Weakness Enumeration (CWE). http://cwe.mitre.org/
5. Cryptography Coding Standard Project cryptocoding.net/index.php/Cryptography_Coding_Standard
6. Aho, A.V., Lam, M. S., Sethi, R., Ullman, J. D. Compilers: Principles, Techniques and Tools, Second Edition. Pearson, 2013. ISBN: 978-1292024349.
7. Percival, C.: Zeroing Buffers is Insufficient. www.daemonology.net/blog/2014-09-06-zeroing-buffers-is-insufficient.html
8. Programming Langauges – C. ISO/IEC 9899:2011. http://www.open-std.org/jtc1/sc22/wg14/www/standards.html
9. Programming Langauges – C++. ISO/IEC 14822:2011 http://www.open-std.org/JTC1/SC22/WG21/docs/standards.html
10. Barnes, J with Altran Praxis. SPARK: The Proven Approach to High-Integrity Software. 2012. ISBN: 978-0-9572905-0-1.
11. McCormick, J. W., Chapin, P. C. Building High-Integrity Applications with SPARK. Cambridge University Press 2015. ISBN: 978-1-107-04073-1.
12. SPARK 2014 Community Site. www.spark-2014.org
13. Consolidated Ada 2012 Language Reference Manual. ISO/IEC 8652:2012/Cor 1:2016. Available from www.ada-auth.org/standards/ada12_w_tc1.html
14. Regehr, J., Eide, E.: Volatiles are Miscompiled and What to Do About It. In: Proceedings of the Eighth ACM and IEEE International Conference on Embedded Software (EMSOFT), Atlanta, Georgia, October 2008. DOI: 10.1145/1450058.1450093 www.cs.utah.edu/~regehr/papers/emsoft08-preprint.pdf
15. Comar, C., Dismukes, G., Gasperoni, F. The GNAT implementation of controlled types. In: Proceedings of Tri-Ada 1994. Baltimore, MD, USA. ACM Press. DOI: 10.1145/376503.376724
16. Kirtchev, H. A New Robust and Efficient Implementation of Controlled Types in the GNAT Compiler. In: Proceedings of High-Integrity Language Technology 2012. ACM SIGAda Letters 32(3) pp. 43-50. DOI: 10.1145/2402676.2402693.
17. GNAT sources at gcc.gnu.org. http://gcc.gnu.org/viewcvs/gcc/trunk/gcc/ada/gcc-interface/trans.c
18. GCC Online Documentation. Chapter 24 – Link Time Optimization. https://gcc.gnu.org/onlinedocs/gccint/LTO.html
19. LLVM Compiler Infrastructure. Link Time Optimization: Design and Implementation. http://llvm.org/docs/LinkTimeOptimization.html
20. Leroy, X.: Formal verification of a realistic compiler, Communications of the ACM, 52 (7), July 2009. DOI: 10.1145/1538788.1538814]
21. Kang, J., Kim, Y., Hur, C-K., Dreyer, D., Vafeiadis, V.: Lightweight verification of separate compilation. In Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL) 2016. pp. 178-190. ACM Press. DOI: 10.1145/2837614.2837642.