Volatile Data Collection

Table of Contents

Volatile Data Collection 2

Why Volatile Data First? 3

Ways to Collect Volatile Data 5

RAM 7

Remote Collection Requirements 10

Local Collection Requirements 12

Imaging RAM 14

Trusted Tools 15

Command Line Tools 17

Scripts 19

GUI-Based Tools 21

Notices 22

Volatile Data Collection

Volatile Data Collection

CERT | Software Engineering Institute | Carnegie Mellon University

3

**003 Okay, we'll go ahead and start out with the volatile data, since we know we want to get at that first, if at all possible, right?

Why Volatile Data First?

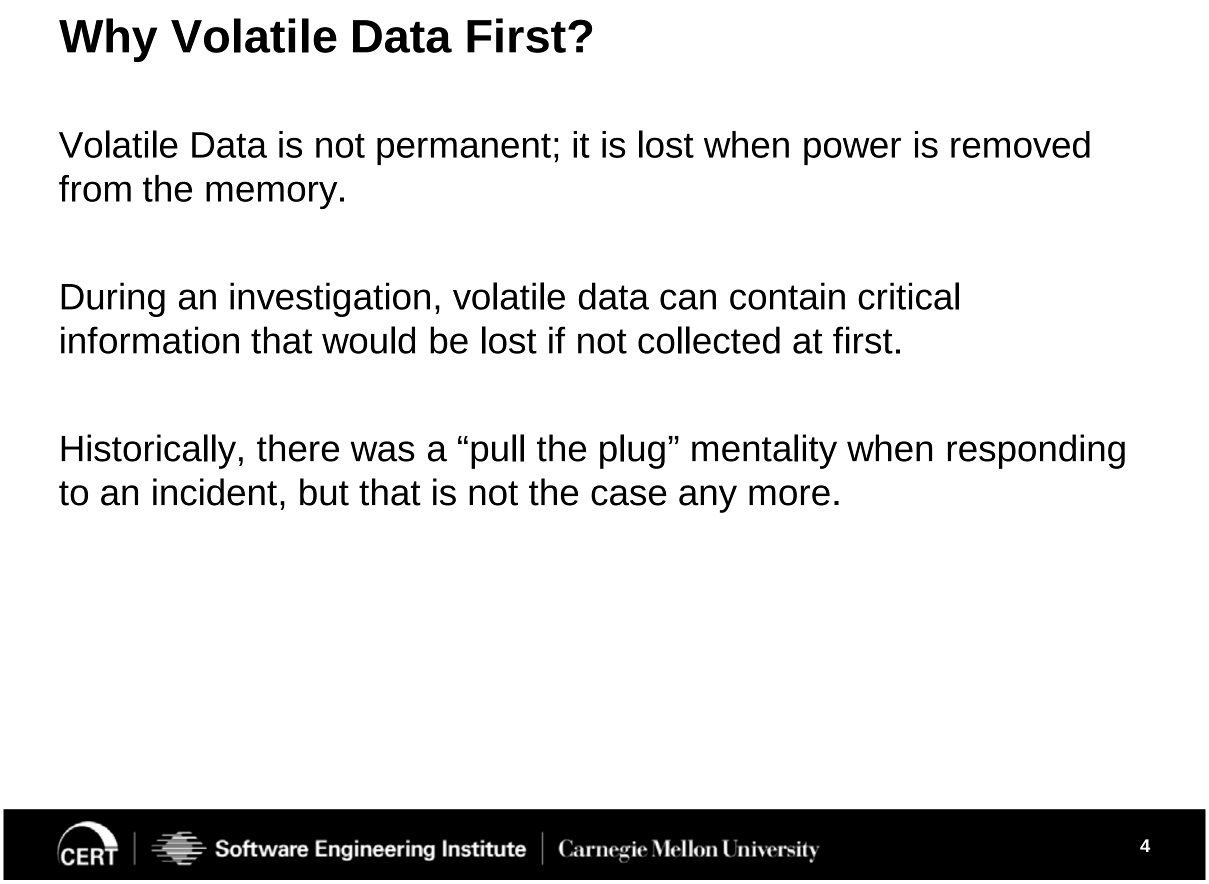

Why Volatile Data First?

Volatile Data is not permanent; it is lost when power is removed from the memory.

During an investigation, volatile data can contain critical information that would be lost if not collected at first.

Historically, there was a “pull the plug” mentality when responding to an incident, but that is not the case any more.

CERT | Software Engineering Institute | Carnegie Mellon University 4

**004 Why do we want it first? Again?

Student: Disappear.

Instructor: Because it'll disappear, right? If you lose power, it's no longer going to be there. There will be no-- gone forever, as far as that's concerned. Right? And there's going to be, during an investigation, there could be some very critical information, so you definitely want to go after this as quickly as possible.

As I mentioned before, historically, there has been a mentality of pull the plug first. And the plug meaning not the CAT-5 cable, but the actual power, right? They're like, "It's malware, it's doing something horrible, so let's just kill the whole thing."

But you cannot pull, if it has-- there's malware out there now, if you've not heard, that resides strictly in memory. It does not have to sit on your drive at all. They can inject it straight into your memory. So you will never find that again, if you just unplug it, right? So that's the sort of thing you're trying to capture.

You're trying to capture the current conditions within your memory. Meaning what processes were you running, what connections were you having? So you know, you want to make sure you take care of that volatile data first, if at all possible.

And if it is in the ranking of importance, if you will, go ahead and go for the volatile data first to make sure that you can get it.

Ways to Collect Volatile Data

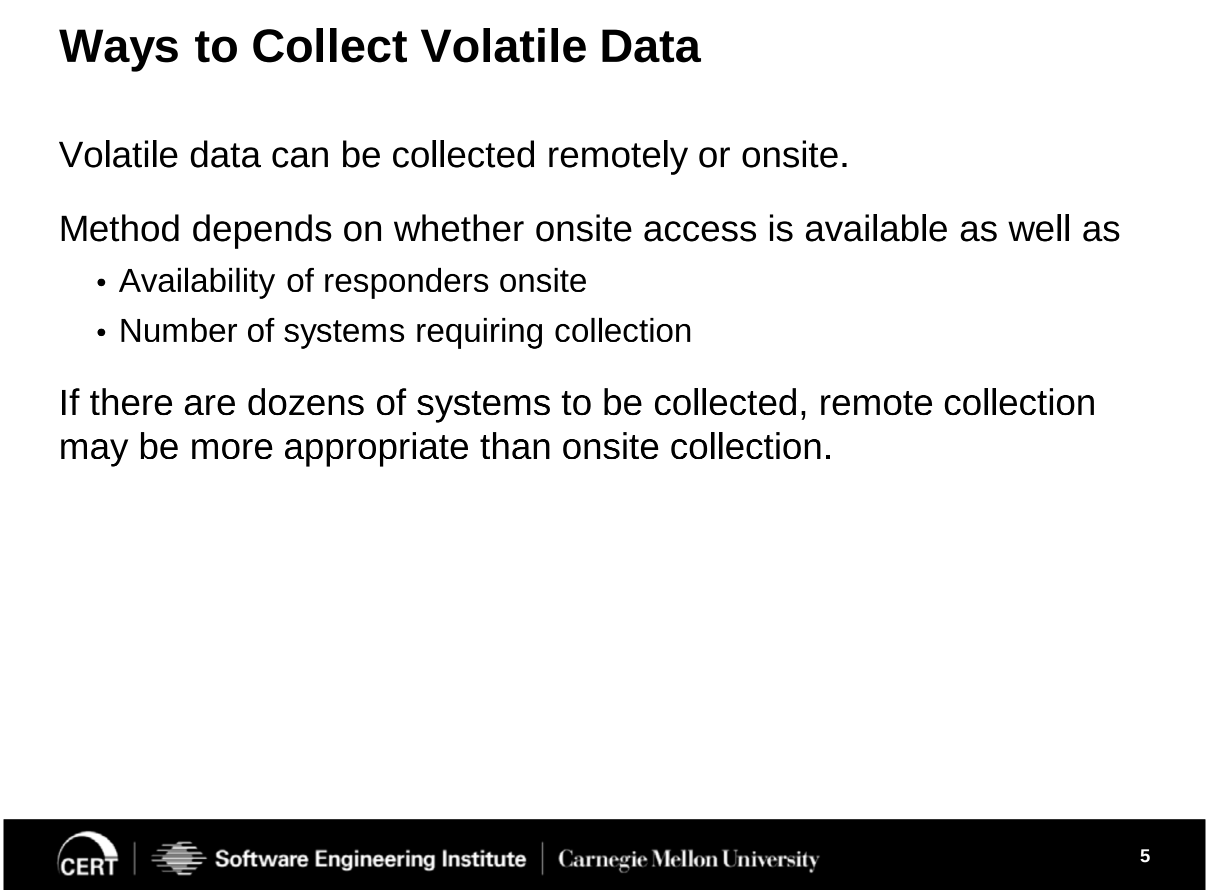

Ways to Collect Volatile Data

Volatile data can be collected remotely or onsite.

Method depends on whether onsite access is available as well as

- Availability of responders onsite
- Number of systems requiring collection

If there are dozens of systems to be collected, remote collection may be more appropriate than onsite collection.

CERT
Software Engineering Institute
Carnegie Mellon University
5

**005 You can collect it either remotely, or onsite, of course.

But remotely, you may have the challenges of trying to pull back eight or sixteen gigabytes of RAM across, what kind of network connection might you have? You may have to ask somebody to punch a hole in their firewall when you're that far away.

So sometimes the answer to that is to have somebody onsite there. You have them-- send them some sort of package that has instructions on how to do it, or you hope somebody there onsite is capable of pulling that memory for you.

And depending on the number of responders, and the number of systems required, but that, like I said, if you really need that volatile memory, and considering the size of it, right, it's no longer four megs of RAM, or even four gigs of RAM. You're talking eight, sixteen and more so on servers have thirty-two, sixty-four, and more gigabytes of RAM. So you're talking about a significant amount of pull, especially across the network.

So might be a better idea to have somebody there onsite do it.

RAM

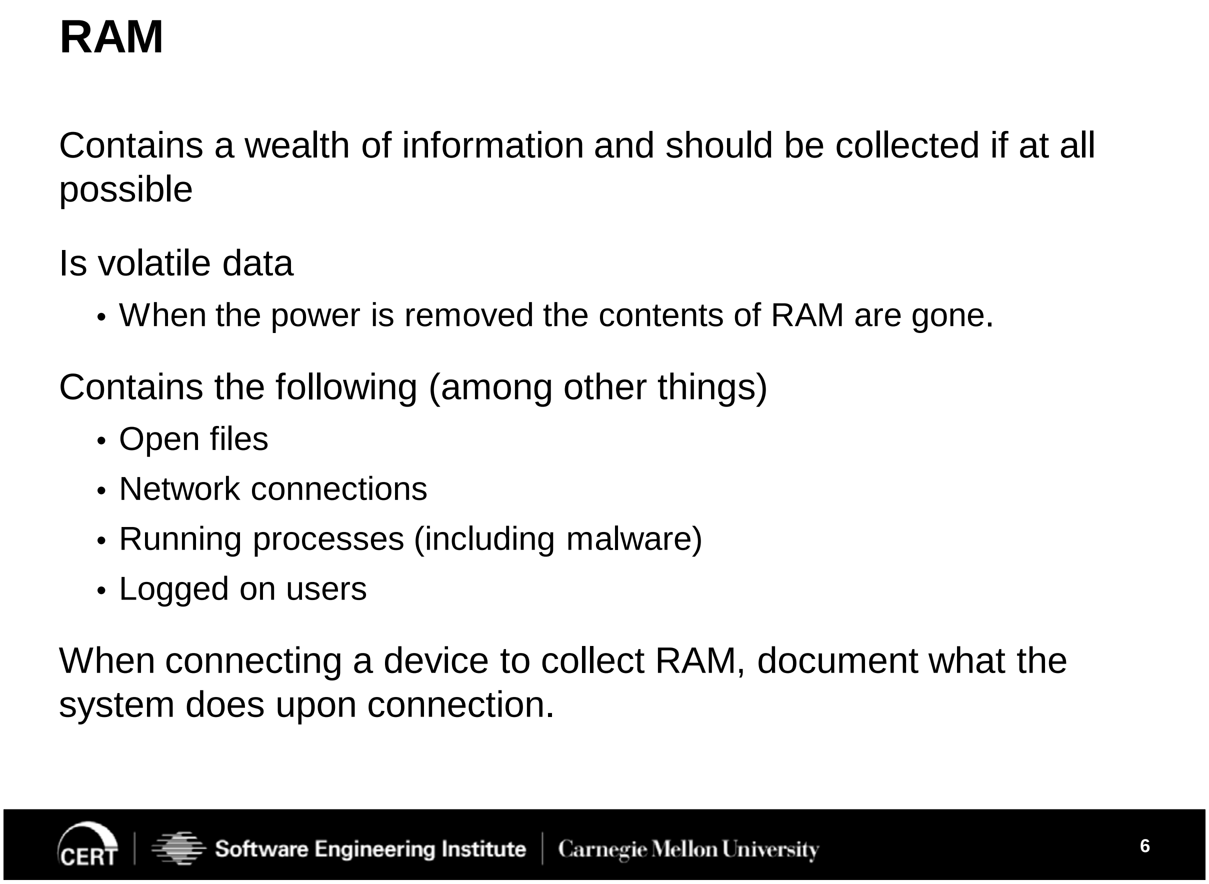

RAM

Contains a wealth of information and should be collected if at all possible

Is volatile data

- When the power is removed the contents of RAM are gone.

Contains the following (among other things)

- Open files
- Network connections
- Running processes (including malware)
- Logged on users

When connecting a device to collect RAM, document what the system does upon connection.

**006 Okay. So this is the RAM, the random access memory, which is interesting, because much of memory now is pretty random. It's not like tape when it had to be sequential and you couldn't get access to it. But there's a ton of information that you'll be pulling if you do get the RAM, if you're able to do so.

In addition to the things I mentioned already, there are the open files. And also logged on users who's actually on a which account, right? It says "users" but actually it's "what account," because sometimes it's a service account that's doing something, if it's been set up to do a

CRON job or some sort of previously scheduled task.

You find out, "Hey, there's a CRON job, or a scheduled task that none of our administrators set up. But it doesn't look like a regular user, it just looks like a service account that did it." So that's the kind of stuff you're also looking for when you look at this. Excuse me.

Student: Full-stop memory on a BM, a snapshot would be sufficient?

Instructor: That's my understanding as well, but since volatility can pull memory right off of-- if you've never heard, there's a tool called Volatility. And I'm sure other things do it, too, but Volatility definitely can pull memory off of a virtual machine. So if you're able to do that, it'd probably be recommended to go ahead and capture that as well.

And that way if, for some reason something goes wrong with your snapshot, because after all a snapshot is just a copy onto your drive, something that happens to you drive or that part of it. Then having the actual stuff pulled out of the Volatility might be a good idea.

And then this here, last statement on the bottom, "When connecting a device to collect to RAM, document what the system does upon connection." This has to deal with possible anti-forensic measures that are a part of, say, malware, that pays attention to maybe a USB connection, whatever way you're connecting to it. Since that's a common way to gather the information is either a USB key, or a USB drive with something.

Or it may even pay attention to putting a disc into the ROM drive. And all of a sudden you see the hard drive light going like crazy. Well, it might be erasing itself. Or it might go as far as being destructive, and it might start doing other things.

There's been cases where the anti-forensic software, or malware, I guess actually went and destroyed the hard drive. Because it was like, "Oh, yeah? Well, you're not going to get anything on me, because I'm going to wipe just about everything off I can," type of thing. And that's truly where you want to yank that cord in a hurry, if at all possible to maintain the integrity of whatever's left on the drive.

Remote Collection Requirements

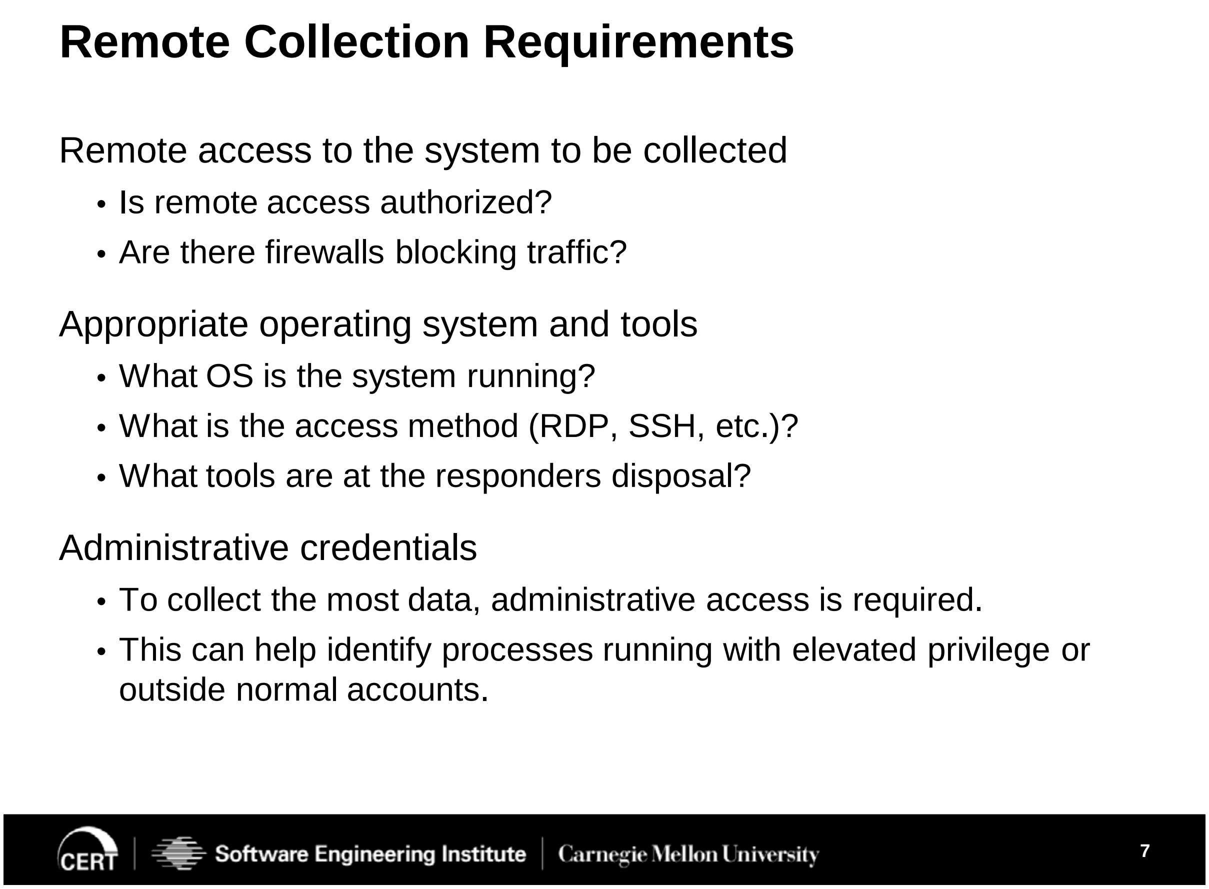

Remote Collection Requirements

Remote access to the system to be collected

- Is remote access authorized?
- Are there firewalls blocking traffic?

Appropriate operating system and tools

- What OS is the system running?
- What is the access method (RDP, SSH, etc.)?
- What tools are at the responders disposal?

Administrative credentials

- To collect the most data, administrative access is required.
- This can help identify processes running with elevated privilege or outside normal accounts.

CERT | Software Engineering Institute | Carnegie Mellon University 7

**007 So remote collection requirements, right? You have to get authorization. That seems like a really common thing. But sometimes people try to do things across the network that they forget to ask permission for, and you can get in trouble for that sort of thing.

Especially when the stakeholder is, you know, perhaps around the world and it's an international thing. You don't want to be in trouble for steeling something internationally, I'm sure. Right?

And then it's the same thing when I mentioned before, you may have to ask somebody to punch a hole through a firewall, and/or if they have a proxy, or some other device that will prevent you from actually getting access internally to that system.

Same thing holds true for remote, right? You gotta have the approximate operating system, even though you're however many hundreds of mi-- thousands of miles away, you may forget. "Gee, I'm trying to collect this thing on a Windows machine, and I forget that it's a Linux or an Apple machine." So you want to make sure you have that stuff worked out well in advance for remote collection.

Your access method sometimes can limit or give you more abilities, right? Are you connecting through a remote desktop? You have a GUI, but you can only do so much through that, right? You can't access some of the things through a remote desktop that you can through SSH, with your secure shell and other methods of connecting. And then it all depends on what tools that responder's got access to, it can also help or hamper the investigation collection piece.

Administrative credentials. This is where when you do something remotely, a lot of times systems are set up in such a way that you have to have that administrator or root account in order to access the full extent of what's available. So you'll want to work that ahead of time so you have the appropriate credentials

at the right level that you need,
because you don't want to sit there
and do something to try to play with
the-- escalate your own privileges.

So make sure you get the
right ones to begin with.

Local Collection Requirements

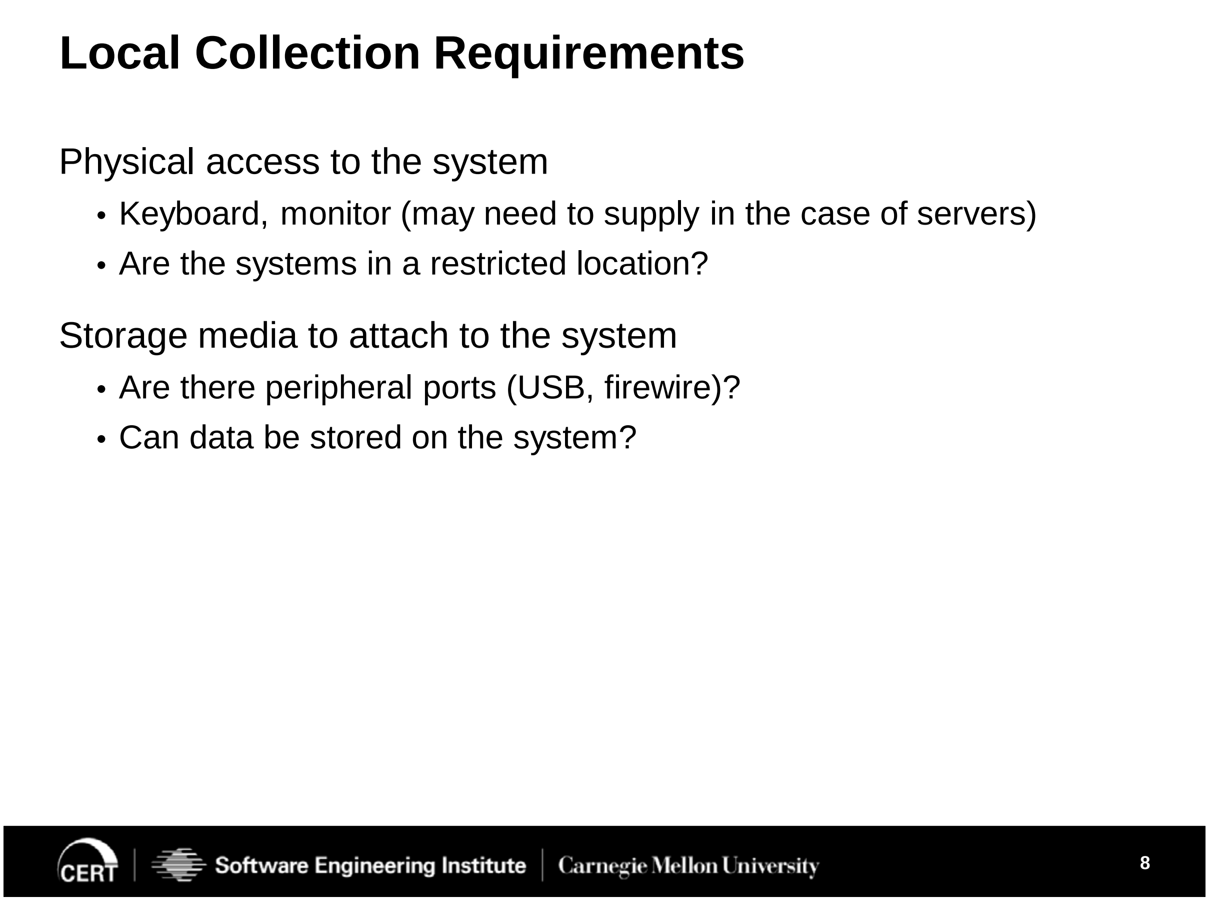

Local Collection Requirements

Physical access to the system

- Keyboard, monitor (may need to supply in the case of servers)
- Are the systems in a restricted location?

Storage media to attach to the system

- Are there peripheral ports (USB, firewire)?
- Can data be stored on the system?

**008 And then the local access collections,
right? Sometimes we forget, if they
use terminals, and you need to
access a server, you may need to
bring a keyboard and a monitor to be
able to look at what you're dealing
with. Right?

And then they're-- if these systems are in restricted areas-- for the Federal government, is it a classified area? Is it a sensitive compartmented information facility, or a SCIF, that you're going to? Because you have to pass your clearances. You have to get permission to go there. There's a lot of times, there's a Visitor Request that you have to send forward, in addition to your clearance information.

And there's different industrial based companies, right? The companies that support the Department of Defense, like, excuse me, the Lockheed's and the Raytheon's and those companies like that, that have the same sort of requirements, even though they're not technically government, they work with government and they have SCIF facilities at their buildings.

So when you're dealing with that sort of thing, be sure you understand that, for showing up on a local site, that you have all your visitor stuff. If you need badges, get those badges done early. All that stuff.

Because again, you're trying to get there, and you're trying to deal with an incident, you know, in an expedient way as possible. So all that stuff, right?

Storage media to attach, right? You know, make sure you got the right cables. And make sure that you have enough room on the data.

Imaging RAM

Imaging RAM

Challenging without the proper tools

- Tools like FTK Imager can be run from a USB drive and make it easy to collect RAM.

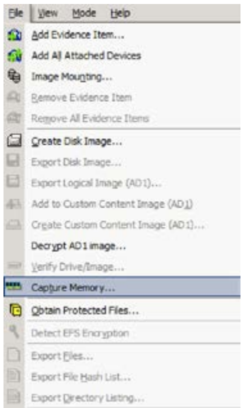

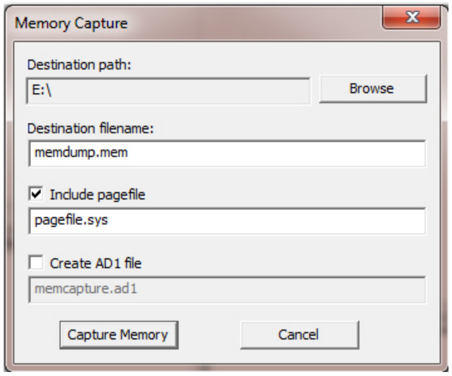
Memory Capture

Destination path: E:\

Destination filename: memdump.mem

☑ Include pagefile

pagefile.sys

☐ Create AD1 file

memcapture.ad1

**009 So imaging RAM, you can do this with FTK Imager. This just happens to be an example.

I don't know if you can see this very well, but this is the verified drive image, and that's where it does the hashing. It doesn't happen to have that option as some do it here. But there's that capture memory. And then you go ahead and you put the information on naming it, and where you want it to go. Pretty straightforward.

Any questions?

Trusted Tools

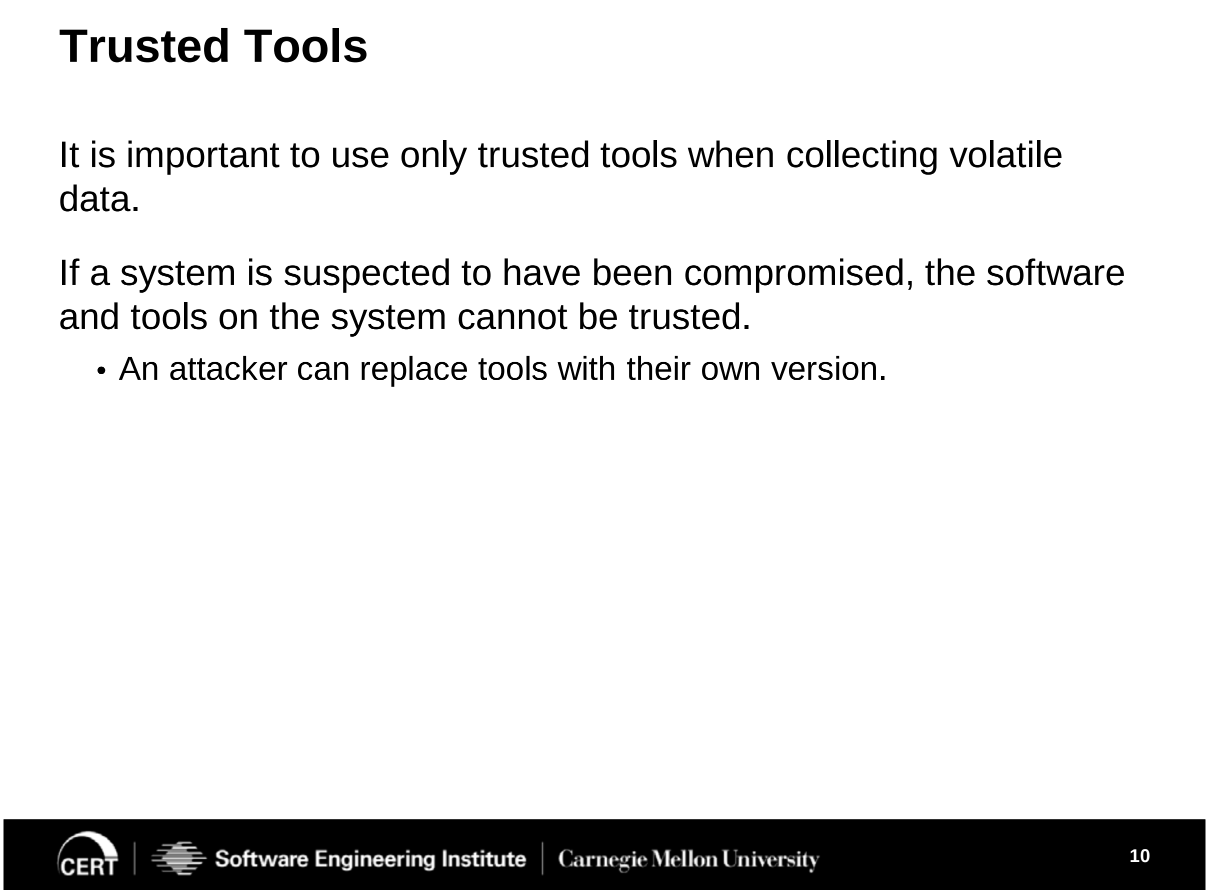

Trusted Tools

It is important to use only trusted tools when collecting volatile data.

If a system is suspected to have been compromised, the software and tools on the system cannot be trusted.

- An attacker can replace tools with their own version.

CERT | Software Engineering Institute | Carnegie Mellon University

10

**010 So trusted tools, this is kind of a side note, but sometimes when a system has been compromised, and you want to go into it, and you're not doing the full disc image, necessarily, and you're just gathering certain pieces of information. You might be tempted-- maybe not you, but the forensic team person-- might be tempted to use some tools that are native or already on the machine. This is not recommended.

There's sometimes where a threat actor will replace a particular file that has a payload or some shell code in it that just calls back. You're able to do this with some of the tools that are available out there today, where you just take a file, and you tell it to put this shell code into the file, and when it executes, everything executes about that same file the same way. Like Calculator or something like that, Notepad, will run as normal, but it will also inject this shell code into memory and run and make a connection back if it's able to do so.

So we're talking about tainted tools, right, versus trusted tools. So this doesn't happen all the time. And it doesn't, you know, doesn't apply to every situation. But if you're dealing with a host that has been compromised and you're only grabbing certain pieces, and you're like, "Oh, it's not--," or maybe it's not even compromised as much as maybe something happened on it. They staged something. They put something there. Or it's a key piece of evidence you need to go pull.

You may not want to use the software that's already installed on that. If you can get your own clean verified version of that software, you would be better off, just in case, so you're not propagating any malware, you're not causing any more problems than absolutely necessary.

Right? So they'll replace it with their own version. Sometimes that thing connects back, sometimes it does all kinds of other things.

So I think connecting out
would be relatively benign compared

to trying to delete the entire hard drive or something like that.

Command Line Tools

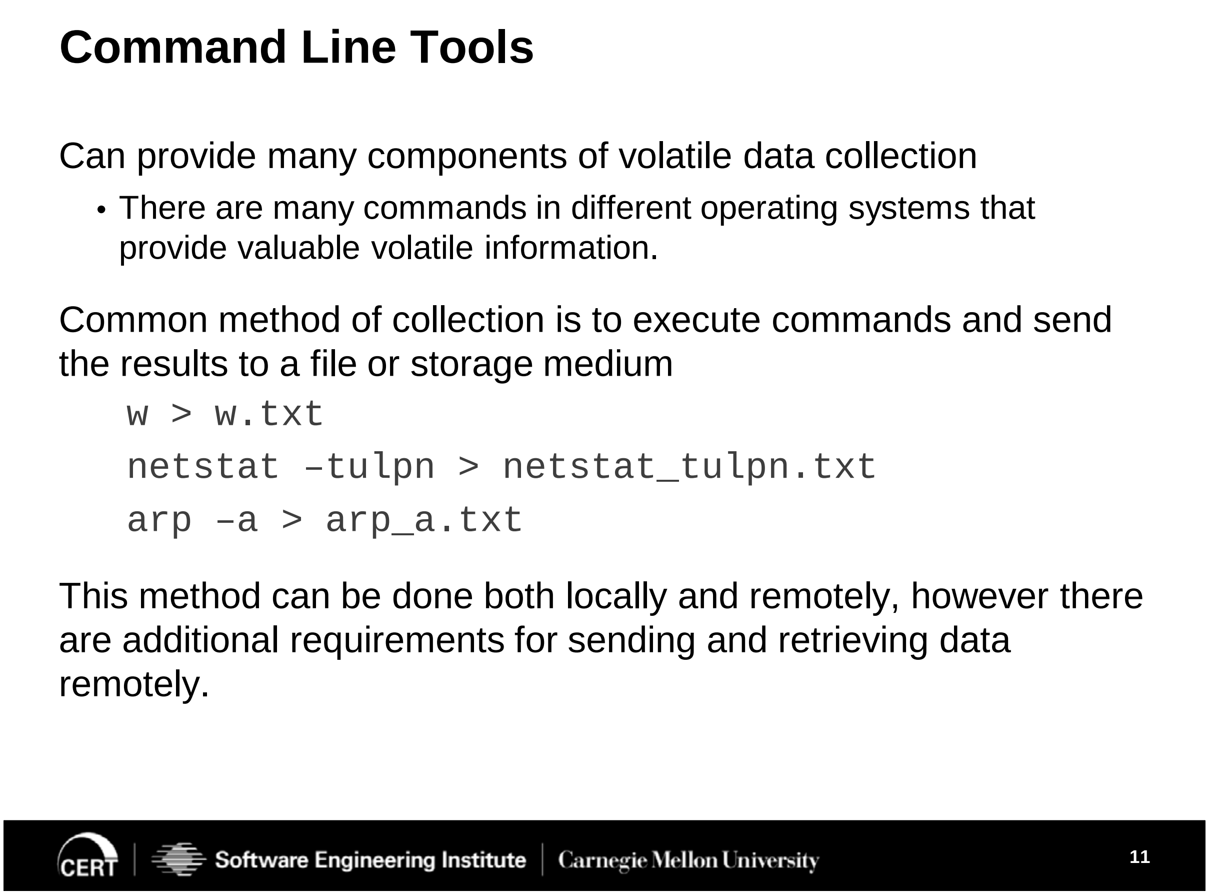

Command Line Tools

Can provide many components of volatile data collection

- There are many commands in different operating systems that provide valuable volatile information.

Common method of collection is to execute commands and send the results to a file or storage medium

```
w > w.txt
netstat –tulpn > netstat_tulpn.txt
arp –a > arp_a.txt
```

This method can be done both locally and remotely, however there are additional requirements for sending and retrieving data remotely.

**011 The command line tools, these are the, you know, the less thought of as forensic tools. But taking the W, that'll list out who's logged in, and you use the redirection, that's the less than or greater than, whichever way you want to put that. And you put that into a text file, that's a perfectly legitimate way to gather things using the command line.

You have your netstat command, and you use the -tulpn flags on it, and you can drop it into a text file that can be used later on to see what it was that was going on. If you want to check out the ARP cache, you just do an ARP command with a -a, and it'll pull all of it, and you can dump it into a text file.

Now this can be done both locally and remotely, so that's really nice. You still need access to it, so there might be firewall issues, or however it is that you need to get at it. You may need to get specific administrator-- some of these are attached to administrator capabilities. The sudo command, or I'm sorry, sudoer ability, or root ability to do this. So you may have to ask for elevated privileges, and those accounts either made or given to you. But they're perfectly legitimate, and they're just as good as anything out there, as far as pulling data straight off your machine, right?

And Windows has the same basic stuff, netstat works there. You can use the set command to pull environment variables. There's a lot of things that you can do natively on the operating system. A lot of those are designed for the admin, right? People that need that information, just to operate it.

But if it's the piece of the information that you're looking for as part of your investigation, you can pull it and you can dump it right into a text file, or another type of file as you need.

Scripts

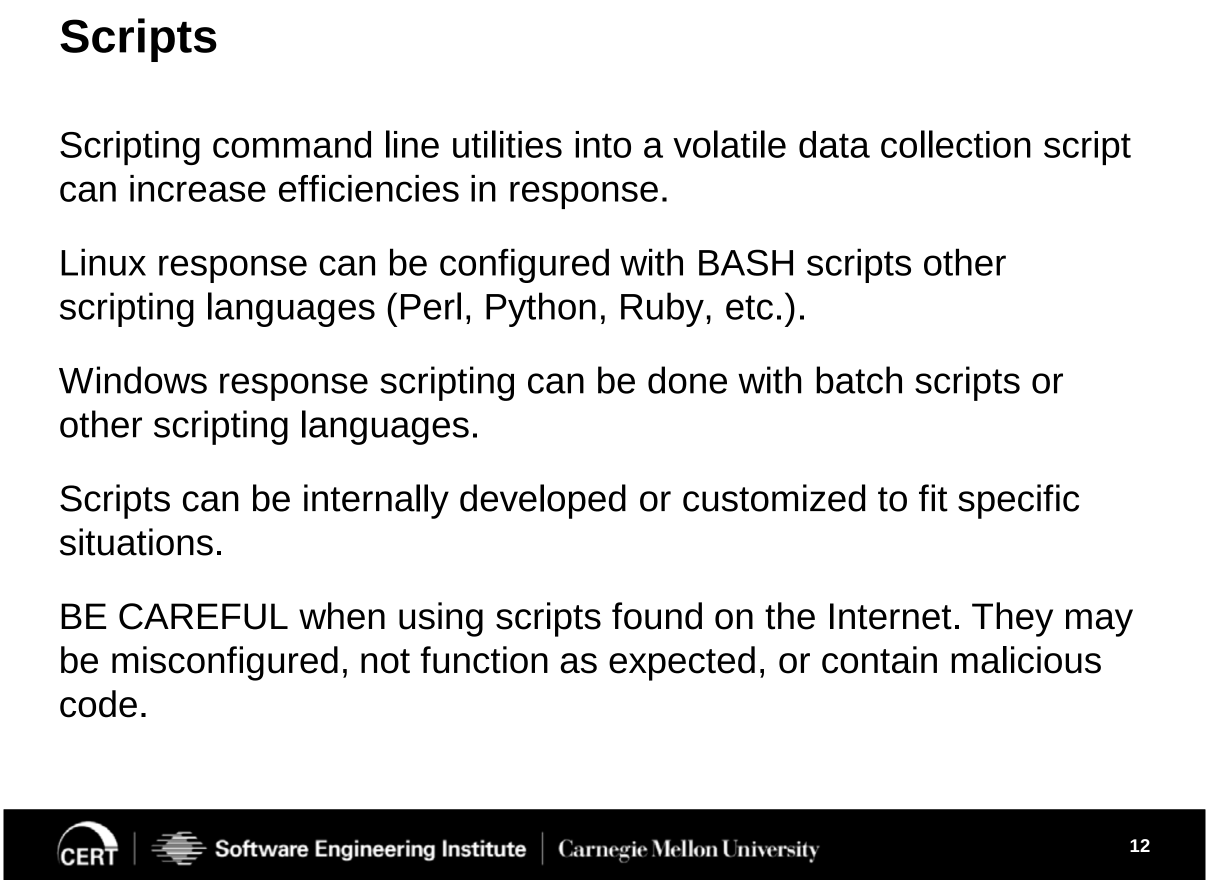

Scripts

Scripting command line utilities into a volatile data collection script can increase efficiencies in response.

Linux response can be configured with BASH scripts other scripting languages (Perl, Python, Ruby, etc.).

Windows response scripting can be done with batch scripts or other scripting languages.

Scripts can be internally developed or customized to fit specific situations.

BE CAREFUL when using scripts found on the Internet. They may be misconfigured, not function as expected, or contain malicious code.

CERT | Software Engineering Institute | Carnegie Mellon University 12

**012 Scripts are quite handy for those of you that have used them before, you know this. In windows, you have the batch files, batch scripts. Essentially, those are, a list of those commands, are some of those that I mentioned that you could put in one file and you can run it and they'll do them pretty much sequentially.

In the Linux world, they use BASH scripts. That's the Born-Again Shell Scripts. And there are scripting languages that are made, not explicitly for this, but are very good at doing these sort of things. The Pearls, the Python, the Rubies of the world, are-- you can use these scripts and you can automate multiple steps on multiple machines, and it just makes your life that much easier.

Again, you got to be careful right here, it mentions it. If you pull any scripts off the internet, make sure you check it. It's really easy to go, "Oh, some admin probably has done this before," and you'll do a search, a Google search. And somewhere in that file they may have put in something malicious.

So you definitely want to be careful. You don't want to introduce anything extra when you're trying to do an investigation under these conditions.

GUI-Based Tools

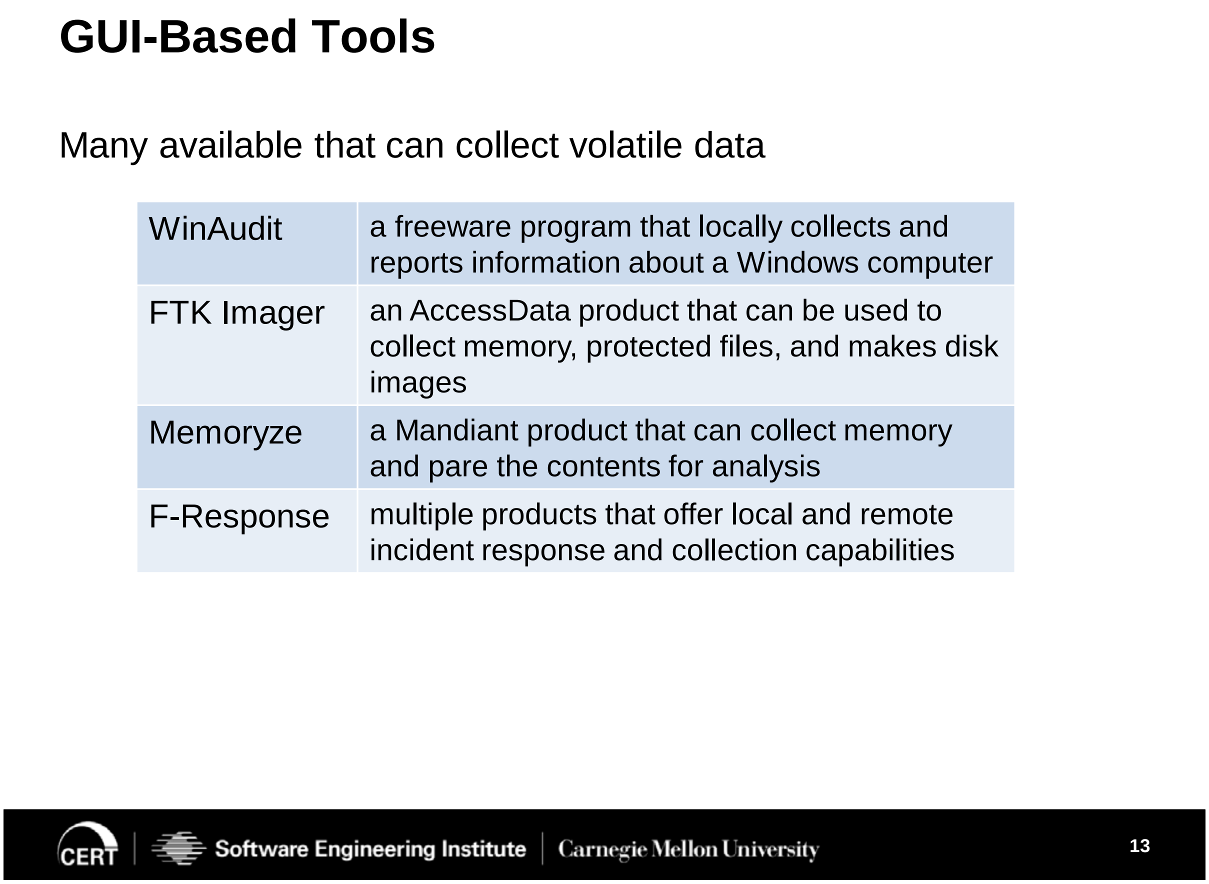

GUI-Based Tools

Many available that can collect volatile data

WinAudit	a freeware program that locally collects and reports information about a Windows computer
FTK Imager	an AccessData product that can be used to collect memory, protected files, and makes disk images
Memoryze	a Mandiant product that can collect memory and pare the contents for analysis
F-Response	multiple products that offer local and remote incident response and collection capabilities

**013 Okay, and then here are some of the GUI based tools. Some I'd mentioned before like the FTK imager, right? Access data creates that. WinAudit is a free program. And that's one of those where it collects, reports information about a Windows computer, in particular, so it specializes in that operating system.

Mandiant has a product for collecting memory, and it's able to pare down some of that information for analysis.

EP Response is actually one of those tools that offers, in particular, I think it's better known for, even though it allows you to do things locally, it's known for the remote capability of incident response.

I don't exactly the technology, kind of like if it's using tunneling or what it is, but they are known to be able to do remote collection very well.

Notices

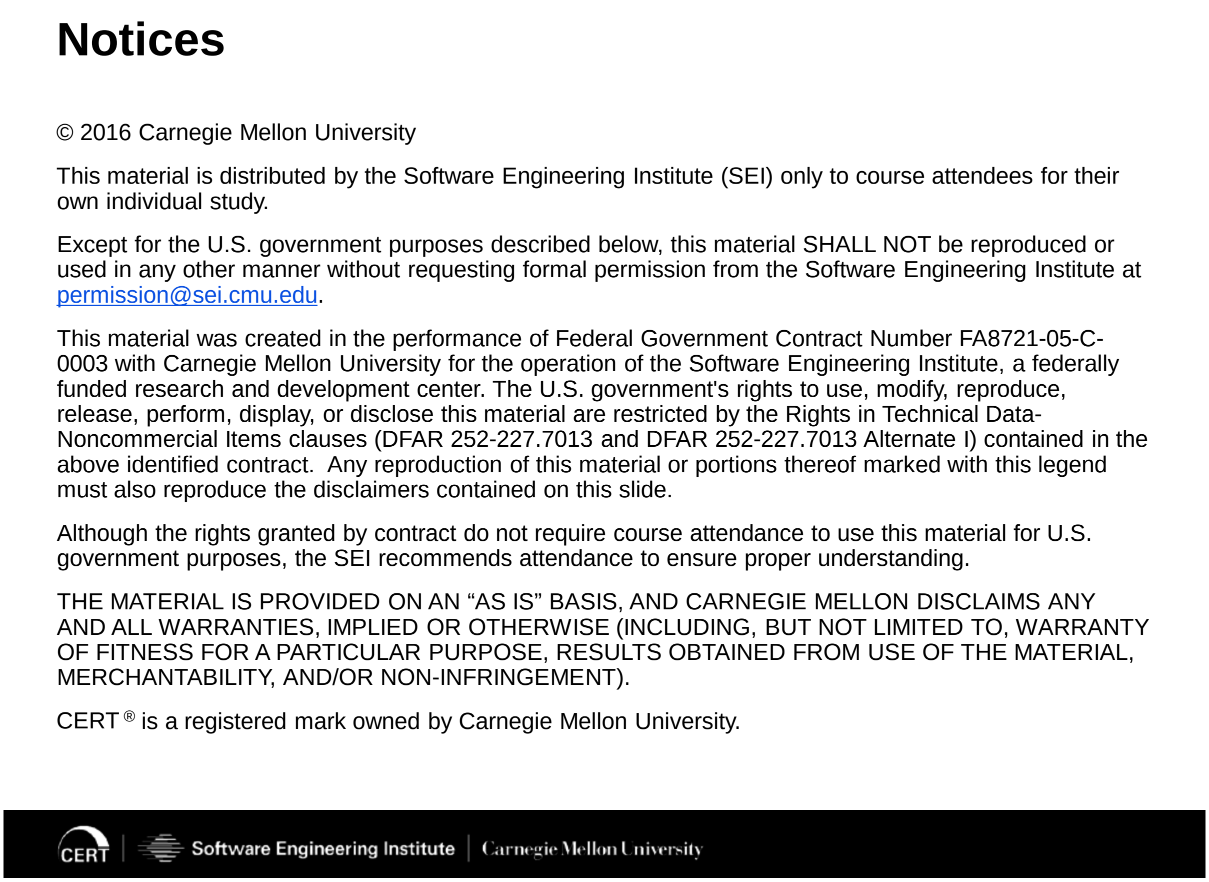

Notices

This material is distributed by the Software Engineering Institute (SEI) only to course attendees for their own individual study.

Except for the U.S. government purposes described below, this material SHALL NOT be reproduced or used in any other manner without requesting formal permission from the Software Engineering Institute at permission@sei.cmu.edu.

This material was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The U.S. government's rights to use, modify, reproduce, release, perform, display, or disclose this material are restricted by the Rights in Technical Data-Noncommercial Items clauses (DFAR 252-227.7013 and DFAR 252-227.7013 Alternate I) contained in the above identified contract. Any reproduction of this material or portions thereof marked with this legend must also reproduce the disclaimers contained on this slide.

Although the rights granted by contract do not require course attendance to use this material for U.S. government purposes, the SEI recommends attendance to ensure proper understanding.

THE MATERIAL IS PROVIDED ON AN “AS IS” BASIS, AND CARNEGIE MELLON DISCLAIMS ANY AND ALL WARRANTIES, IMPLIED OR OTHERWISE (INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE, RESULTS OBTAINED FROM USE OF THE MATERIAL, MERCHANTABILITY, AND/OR NON-INFRINGEMENT).

CERT ® is a registered mark owned by Carnegie Mellon University.