I Guess it is Correct - A Survey about Acquiring and Analyzing Main Memory

Lisa Rzepka[a,1], Benedikt Mader[b,1], Zinaida Benenson[b], Harald Baier[a]

[a] *Universität der Bundeswehr München, FI CODE, Germany*
[b] *Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany*

Abstract

In order to detect, mitigate, and respond to stealthy attacks, such as Living off the Land (LOTL) and fileless malware, an in-depth analysis of the affected main memory is unavoidable. While live analysis provides initial insights into malicious activity, acquiring and analyzing a main memory dump offers a more substantial and deeper view of the system. However, acquiring main memory is a complex task that often leads to inconsistencies. Although this is common knowledge, the forensics community currently lacks details about the actual usage of main memory forensics and its pros and cons. To remedy this blind spot, this paper presents the results of a survey conducted with memory forensics practitioners. Based on 31 comprehensively answered questionnaires we gain insights into preferred methods of memory acquisition (i.e., kernel-level software) and related concerns (i.e., system crashes during dumping). With respect to memory analysis, common use cases include obtaining information about processes, network connections as well as indicators of malware on the system. Although inconsistencies in a memory dump are rather common, the survey reveals that the correctness of a memory dump is postulated rather than thoroughly tested. Hence, as a step toward more reliable memory acquisition, we point out key issues during software-based memory dumping.

Keywords: Practitioner Survey, Memory Forensics, Digital Forensics, Memory Acquisition, Memory Analysis, Challenges

1. Introduction

Smart cyberattackers make use of evasion techniques to circumvent detection. Common obfuscation practices are the misuse of ordinary operating system functionality (e.g., Living off the Land, LOTL) or fileless malware (i.e., malicious software, which is only present in main memory rather than on disk). For instance, in their annual report for 2023, the US cybersecurity company ReliaQuest reported that *'most of the critical customer incidents we responded to in 2023 involved fileless malware, accounting for 86.2% of all detections'* [19].

Due to the ephemeral nature of LOTL or fileless malware traces, incident responders and digital forensic practitioners need insights into system activities like processes (both running and recently terminated), network connections (both open and recently closed), and file handles. While live analysis provides fast answers, it relies on the potentially compromised operating system and needs high experience to appropriately prioritize the system analysis in a time-critical situation. Additionally, ephemeral traces of recently terminated processes, network connections or opened files may not be accessible in a live analysis after some minutes. Hence, a common approach is to dump the main memory and subsequently analyze it in the lab in a thorough post-mortem memory analysis.

However, main memory acquisition is a complex task and often leads to inconsistencies. For instance, Case and Richard III [3] introduce the notion of *page smearing*, i.e., inconsistencies due to a non-atomic acquired memory dump. Memory inconsistencies are recently examined and quantified by Ottmann et al. [16], Rzepka et al. [20], and Oliveri and Balzarotti [14]. The overall result is that *'inconsistencies are very common and that their presence can undermine the reliability and validity of memory forensics analysis'* [14].

Although this criticality is common knowledge, the forensics community currently does not have statistics about key aspects of the actual usage of main memory forensics, its benefits and drawbacks. In this paper we address this blind spot and present the results of a survey, which we sent to memory forensics practitioners.

Research Questions. The goal of our study was to investigate how memory forensics is applied in practice and how inconsistencies during analysis are handled. Our research was therefore led by the following research questions:

RQ1 What are common use cases for acquiring main memory in practice?

RQ2 What techniques are used for acquiring main memory in practice?

Email addresses: lisa.rzepka@unibw.de (Lisa Rzepka), benedikt.mader@fau.de (Benedikt Mader), zinaida.benenson@fau.de (Zinaida Benenson), harald.baier@unibw.de (Harald Baier)

[1] These authors contributed equally to this work.

RQ3 How are inconsistencies handled when analyzing a memory dump?

Contributions. Our main contribution is to design, conduct and evaluate a survey on the actual usage of main memory forensics. Although this is a small and difficult-to-reach population, we received 31 completely answered questionnaires with the following key results:

- Digital forensic practitioners weigh the benefits of memory acquisition and live analysis. A key issue of memory dumping is the concern of a possible system crash during acquisition. If a memory dump is acquired, kernel-level software and virtual machine based techniques are the most common ones.
- The most relevant analyzed traces in a memory dump are the processes, network connections, command history, file handles, and indicators of malware at the time of acquisition.
- Although the memory dump often can be analyzed without problems (rating median = 5 on a scale from 1 to 7, where 1 means 'never' and 7 means 'always' in the context of a fully analyzable dump), there is a non-negligible amount of non-analyzable dumps (median = 2 on the same scale).
- The assessment of dump consistency is non-structured and makes use of heuristics, such as inconsistencies in process, timestamp, and network information. Correctness often needs to be assumed to some degree.

Additionally, we present recommendations to acquire memory dumps of high quality.

Outline. The rest of this paper is structured as follows: Section 2 describes the background and related work. Then, in Section 3 the methodology of the survey and the analysis is explained, followed by the presentation of the results in Section 4. We then discuss the survey outcomes in Section 5 and sketch our recommendations for choosing an appropriate memory acquisition method in Section 6. Finally, Section 7 concludes the paper.

2. Background and Related Work

This section first introduces the issue of quality in memory dumps before describing existing surveys about memory forensics.

2.1. Quality in Memory Forensics

Previous work focused on establishing quality criteria for memory dumps [26], in particular, the concept of atomicity. Atomicity means that the resulting memory dump taken from a running machine is expected to be free of inconsistencies caused by concurrency.

Pagani et al. [18] built upon Vömel and Freiling [26] and introduced the notion of time consistency which was later extended by Ottmann et al. [16]. Both demand that a consistent memory dump should be similar to a dump taken from a frozen virtual machine, i.e., acquired during a halted system.

A more practical approach was taken by Ottmann et al. [17], Rzepka et al. [20] and Oliveri and Balzarotti [14] by assessing the amount of inconsistencies present in memory dumps using different methodologies. While Ottmann et al. [17] and Rzepka et al. [20] used a custom program to track inconsistencies in the heap of the said program in combination with the determination of inconsistencies in certain kernel data structures, Oliveri and Balzarotti [14] measured the amount of inconsistencies in kernel data structures using a framework that allows tracking of alternations during a certain time period.

However, none of the previously mentioned works analyzed the possible impact of inconsistent memory dumps during investigations, nor did they investigate the occurrence of inconsistencies in real-world settings. Furthermore, it remains an open question whether memory forensics in general is used in practice and, if so, how frequently.

2.2. Surveys about Memory Forensics

Research in memory forensics has so far focused on technical aspects such as investigating malware, finding artifacts of certain activities in memory, or tool testing [3].

There exist a few surveys with general digital forensics practitioners [1, 8, 4, 23, 7, 5]. Although most do not specifically address memory forensics, Horsman [7] partly examines the accuracy of digital forensic tools. 79 out of 90 participants reported encountering erroneous results from their tools, and cross-validating tool results with other tools seems to be common practice. Hargreaves et al. [5] asked participants to report their expertise across various digital forensics domains, including memory forensics, as well as how frequently they use digital investigation techniques such as memory analysis. Their findings indicate that memory forensics plays a relatively small role in the field. Of the 122 participants, only 10 rated their expertise as the highest on a scale from 0 to 5, while 17 were in the second-highest category. Additionally, only 9 participants reported using memory forensics almost always (80-100%), and 15 indicated they use it often (60-80%).

To the best of our knowledge, no research currently exists investigating memory forensics practitioners, their practices, or the challenges they face.

3. Survey Methodology

This section first describes the study design choice, the design of the survey, ethics, and the recruiting process. Afterwards, the analysis method is presented.

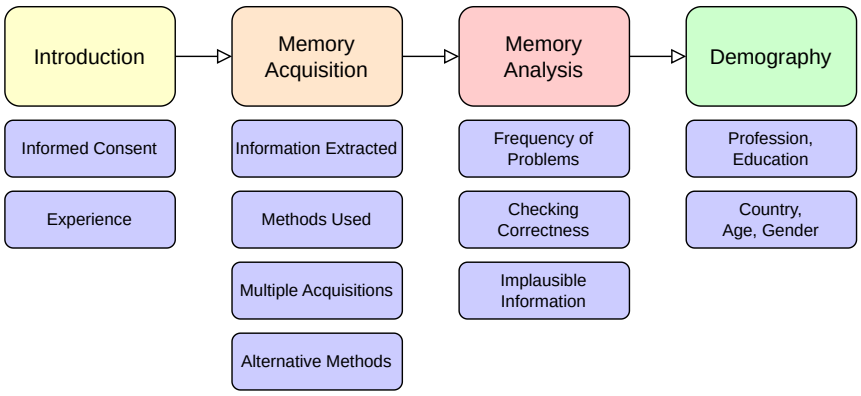

Figure 1: Structure of the survey

3.1. Study Design

As memory forensics practitioners are a difficult-to-reach population, participant recruitment posed a considerable challenge. We therefore decided to conduct an online survey, to keep the barrier to participation as low as possible.

One alternative would have been to conduct interviews with practitioners. A benefit of such an approach would have been that processes and problems could be described in greater detail by the participants. However, it would have required considerably more effort from the participants and would therefore presumably have decreased the willingness to participate markedly.

As there are currently no other studies on memory forensic practitioners, several additional topics remain open for future research. Our study provides a starting point for further work in this area.

3.2. Survey Design

The survey consisted of four parts implemented on a locally hosted instance of LimeSurvey[2]. We chose a survey design with questions that can be analyzed using descriptive statistics or qualitatively. Most questions were answered using seven-point Likert scales ('never' to 'always'; 'no expertise' to 'high expertise') or free-text responses. A few questions had single- or multiple-choice answers. An overview of the survey is shown in Figure 1 with a full version presented in Appendix Appendix A.

After providing informed consent, participants were first asked about their experience with main memory acquisition and analysis. Participants who indicated that they had never created a memory dump before were filtered out at this stage, leaving only those with practical experience in this area.

The second part focused on memory acquisition. The questions included the percentage of time spent on memory forensics and the types of information extracted, based on Osborne [15], as well as the methods used to acquire memory. To focus on the challenges of memory forensics, we asked whether participants conducted memory acquisitions multiple times for one case and what methods other than memory forensics they used and the reasons for doing so.

The third part of the survey investigated memory analysis. It began with the frequency of errors that hinder analysis during the acquisition process. Participants were asked how they check the correctness of information extracted from memory dumps and how they address implausible information. If participants indicated that they had encountered implausible information, they could enter the problems they had encountered and how they addressed them, as well as whether the problems had affected their investigation.

The survey concluded with demographic questions, including their professional roles (in which they acquire and analyze main memory, see Hargreaves et al. [5]), their employment status, year of birth, country of residence, highest educational level, and gender.

To ensure the questions were understandable, we conducted a pre-study with five research colleagues experienced in memory forensics. We then refined the questions based on their feedback.

The survey was designed bilingual in English and German, to improve recruitment and the quality of the free-text responses. For this, three researchers independently translated the English version to German, which were then unified into a single version in a joint discussion of two of these researchers. An additional person, not involved in the research, then translated the German version back into English so that we could compare it to the original and ensure semantic equivalence.

3.3. Ethics

The universities involved in this research do not possess ethical review boards for non-medical studies. However, we ensured that the survey was conducted in an ethical manner based on the following rules. Participation was voluntary and based on informed consent. The survey was anonymous and all demographic questions apart from the

[2] `https://www.limesurvey.org/`

professional role (item 41) were optional. The LimeSurvey instance was hosted locally, ensuring that no data was shared with third parties. Only project members had access to the data.

3.4. Recruiting

Over the span of one month, we distributed the survey to law enforcement agencies, incident response teams, and researchers through personal contacts. We further shared the survey with the digital and memory forensics communities via mailing lists, LinkedIn, and Google Groups. Throughout this process, we employed a snowballing approach, encouraging participants to share the survey with their colleagues and acquaintances.

3.5. Data Cleaning

A total of 48 people completed the questionnaire. Of those, 16 ended the questionnaire automatically after item 13 because they did not have any expertise in memory acquisition.

We manually analyzed the data for plausibility. Special indicators were the frequency of methods used and errors encountered (items 23 and 31) as nonsensical answers were possible here, such as always experiencing both non-analyzable and analyzable memory dumps. We had to remove one participant as the corresponding responses included the value 6, on the scale from 1 (never) to 7 (always), for all statements of item 31 and outlier answers for most other questions, as well as barely plausible demographics. We removed three additional answers only for item 31. Two of these respondents correctly indicated a score of 1 for all options because they had not worked with memory forensics in the last year. The third respondent reported a score of 6 for the last three statements, but otherwise had plausible answers.

3.6. Analysis Methodology

The remaining 31 responses to the survey were analyzed using both quantitative and qualitative methods.

3.6.1. Quantitative Analysis

Given that the Likert scales in this survey consist of seven response points and were designed to measure a single underlying construct, the resulting scores can reasonably be treated as metric data [9]. This, in turn, justified the calculation of mean (M), median, standard deviation (σ), and interquartile range (IQR) for these items. Because most of our data was not normally distributed, we decided to present the median and IQR. For the few single- or multiple-choice questions we provide the resulting frequencies.

Table 1: Demographics of participants (n=31)

Variable	**Characteristics**	**n**	**%**
Roles (multiple choice)	Law enforcement	14	45.2
	Independent consultant	3	9.7
	Incident response	18	58.1
	Industry	3	9.7
	Research	11	35.5
	Student	5	16.1
	Other (Pentesting)	1	3.2
Main employment	Public sector	15	48.4
	Private sector	14	45.2
	Student	1	3.2
	Prefer not to say	1	3.2
Age	30-39	14	45.2
	40-49	10	32.3
	50-59	2	6.5
	Prefer not to say	5	16.1
Country of residence	Germany	13	41.9
	USA	7	22.6
	Switzerland	3	9.7
	Australia	1	3.2
	France	1	3.2
	UK	1	3.2
	Spain	1	3.2
	Prefer not to say	4	12.9
Language	German	16	51.6
	English	15	48.4
Hightest education	High school degree	2	6.4
	College degree	1	3.2
	Bachelor's degree	3	9.7
	Master's degree	17	54.8
	PhD	6	19.4
	Prefer not to say	2	6.4
Gender	Male	23	74.2
	Female	1	3.2
	Prefer not to say	7	22.6

3.6.2. Qualitative Analysis

The free-text responses were analyzed qualitatively using thematic analysis [2]. This method is used to identify and interpret patterns in data in a structured way. Two team members coded the qualitative responses independently using a LibreOffice[3] document. In this context, coding means assigning one or multiple tags to an answer to describe its semantic content [22]. The codes were then harmonized in a joint discussion. Through an iterative process, the codes were further refined and categorized into overarching themes. The answers were analyzed in their original language. Quotes used in this paper were translated later if necessary. Codes are highlighted

[3] `https://www.libreoffice.org/`

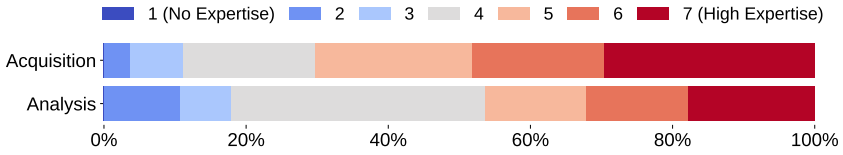

Figure 2: Expertise of participants with memory acquisition and analysis

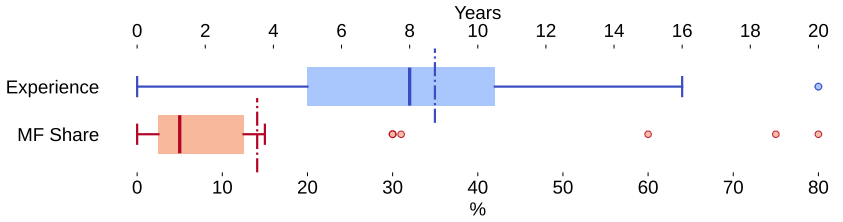

Figure 3: Box-plot and mean of working experience and share of time spent doing memory forensics in the last year

in `monospace` in the following text.

3.7. Demographic Data

The demographic data of the 31 participants is presented in Table 1. Participants primarily worked in incident response, law enforcement, or research, with an even split between the public and private sectors. They came from at least seven countries, with Germany and the USA being the most frequent. An equal number of participants took the questionnaire in German and English. Most of the participants were male, highly educated, and between 30 and 49 years old.

4. Survey Results

This section describes the results of the survey.

4.1. Usage of Memory Forensics

The majority of participants were highly skilled in memory forensics, as can be seen in Figure 2. On a scale from 1 (no expertise) to 7 (high expertise), acquisition was higher (median = 6, IQR = 2.5) compared to analysis (median = 5, IQR = 2.5).

The duration of working with memory forensics is median = 8 years (IQR = 5.5). During the last year, the participants spent median = 5% (IQR = 10.0%) of their working hours on memory forensics. However, there are a few participants with a considerably higher percentage of time doing memory forensics raising the average to $M = 14.1\%$ ($\sigma = 21.5\%$). The distribution of these two variables is shown in Figure 3.

4.2. Memory Acquisition and Extracted Information

The most frequently reported type of information extracted from main memory (Figure 4) was lists of currently running processes and threads (26), followed by network information (26), command history (24), open files (22), and information about malware (22). Additionally, 19 participants search for current users of the system, and 17 for registry information. 16 participants want to extract cryptographic keys. Eight participants indicated types of information other than those listed. These included additional details about `malware`, `additional process information`, `artifact scanning`, and `log files`.

Concerning the acquisition method used, shown in Figure 5, on a scale from 1 (never) to 7 (always) software tools were rated as the most used (median = 6, IQR = 2.0), followed by hypervisor-based acquisition (median = 5, IQR = 4.0). Direct memory access (median = 1, IQR = 1.5) and cold boot methods (median = 1, IQR = 1.0) were both reported to being rarely used, with the majority of participants never using these methods. Only three participants indicated that they occasionally use other methods (median = 1, IQR = 0.0). These other methods included mainly methods which we would classify under direct memory access techniques, such as using `JTAG`.

The most frequently reported software tool for memory acquisition is WinPmem (14), followed by FTK-Imager (8). Seven participants indicated the use of AVML and LiME. Six participants rely on DumpIt for memory acquisition and six participants use products of Magnet. Two participants obtain memory using EDR solutions or Volexity's Surge. Additional mentioned tools include FireEye's Memoryze, Procdump, the minidump functionality of operating systems, LinPmem, KnTDD, and tools by Passware, Arsenal Recon, Palo Alto Networks and Cellebrite.

Regarding direct memory access methods, PCILeech was mentioned by five participants. Other mentioned DMA methods are Doca Appshield, Thunderbolt Cable, products by Belkin, FireWire, and PCIScreamer.

4.3. Reasons Against Using Memory Forensics

When deciding against memory acquisition, participants reported different alternative methods. The most

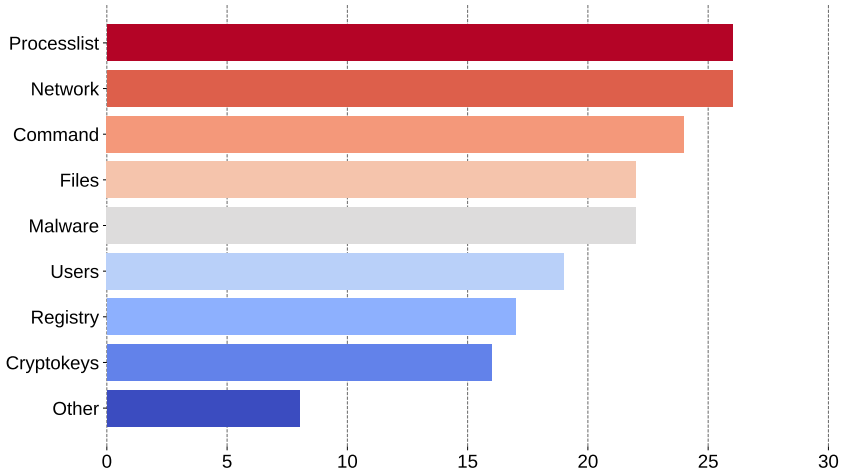

Figure 4: Reported frequency of types of information extracted from memory

common alternative was `disk analysis`, mentioned by eleven participants. However, this is not always possible, as *"some artifacts may only reside in memory" (p30)*. Ten participants use `live analysis` or perform a `selective analysis of specific artifacts` (6), that is, they do not acquire a full memory dump, but obtain specific artifacts (e.g., process memory) or employ *"Triage analysis (KAPE)" (p3)*. Additionally, some participants mentioned `analyzing logs` or utilizing endpoint telemetry from endpoint detection and response (EDR) solutions. As one participant stated: *"Ideally there is an X/EDR solution on the host which has already recorded detailed telemetry/system events enabling adequate investigation of malicious activity." (p22)*. If the purpose of main memory acquisition is to access an encrypted system, participants sometimes attempt to `circumvent the encryption`, for example, to *"find out and test passwords" (p9)*.

In addition to indicating the alternative methods used, the survey also asked participants to state their reasons for relying on such methods. A major challenge of memory forensics is its `reliability`, which includes both errors during acquisition or tool-based analysis and errors arising from false interpretations of data, i.e., both the reliability of the tools themselves and the reliability of the analysis results. This makes memory forensics somewhat `cumbersome`, as it requires substantial time and effort. If the data from main memory is `not necessary`, `other techniques` may be preferable.

> *"When analysis tools such as Volatility do not work appropriately." (p30)*
>
> *"Memory analysis does not often provide new results (in addition to other methods), is tedious and prone to false interpretations." (p3)*

Many participants also highlighted the `risk of a system crash` as a deterrent to using memory forensics. The process of memory acquisition may crash the system, resulting in data loss or reduced availability. One participant argues as follows:

> *"Memory acquisition poses a significant risk of crashing a system, and this is often an unacceptable risk for physical ("bare metal") systems in production. [...] Where memory cannot be collected safely or is unlikely to yield results, alternative collection methods are preferable versus getting no information." (p22)*

In other instances, memory forensics is simply `not possible`. Practitioners may have `no access` to the system, either physically or in terms of privileges, have `no tools available`, the system may have `too much memory`, or `ransomware` may occupy all CPU resources. In law enforcement contexts, investigators may also lack a warrant to seize main memory.

> *"No admin/sudo privileges" (p21, translated)*
>
> *"In the case of ongoing encryption (ransomware), it is virtually impossible to perform a dump because almost all computing time is used by the ransomware and the active time until the ransomware effectively self-destructs, including keys, is sometimes only a few minutes. Ideally, you would have a paused VM—but that is wishful thinking on the part of first responders/forensic experts :-)" (p14, translated)*

Often, memory acquisition may also be fruitless because the `data is already lost`, for example, if the program has already finished execution, too much time has passed since the incident, or `after a system shutdown`, as this resets the main memory.

> *"Very often memory acquisition isn't feasible because systems were rebooted or it's been too*

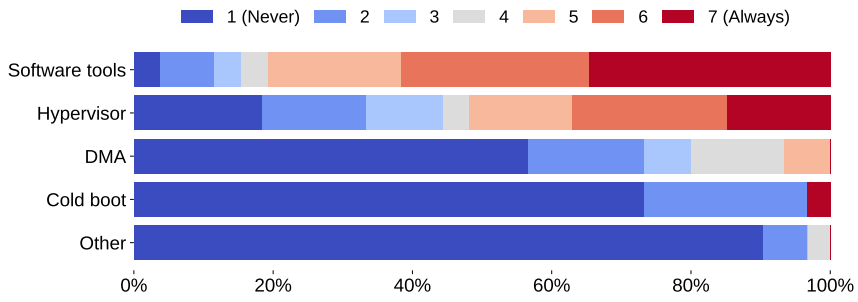

Figure 5: Acquisition methods used by participants

long since the incident to even consider memory acquisition" (p20)

"Not helpful if the admin has already 'cleaned' the system and restarted it several times." (p10, translated)

4.4. Problems during the Memory Forensics Process

Problems that occur during memory acquisition and analysis are generally infrequent, but most participants reported encountering them at least occasionally. An overview of these results is given in Figure 6. For acquisition, both the problem of not being able to start the process (median = 2, IQR = 2.3) and the process not completing successfully (median = 2, IQR = 2.0) were observed. While roughly one third of participants never experienced these issues, several reported that they occur regularly. A similar pattern emerged in the analysis of memory dumps. Although it was uncommon to have a dump that was completely unanalyzable (median = 2, IQR = 1.0) or only partly analyzable (median = 2, IQR = 2.0), only a single participant reported always obtaining fully analyzable dumps (median = 5, IQR = 3.0). Moreover, nearly one quarter of participants gave ratings above 4 for problems with non- or partly analyzable dumps.

4.5. Encountering Implausible Information

Participants mentioned three main strategies to verify the correctness of a memory dump: `check plausibility`, `comparison with other methods`, and `dual tool verification`. Checking plausibility refers to identifying known information within the memory dump as stated by a participant:

"Check for system processes that I'm sure are running on the system and check their information (start time and up time of the process, user associated etc)" (p5)

This strategy is related to comparing the dump with other sources, such as disk analysis. Many participants also reported performing `dual tool verification`. However, as most participants never or rarely acquire memory multiple times for the same case (median = 2.0, IQR = 2.0, see Figure 7), this often does not address the possibility of errors during acquisition.

Additional strategies mentioned included the `manual analysis` of a dump or simply verifying whether it `is parsable`. Nevertheless, the most frequent emerging theme was `assume correctness`. Participants either reported that verifying correctness is `not really possible` or relied on strategies that serve only as weak indicators of correctness, such as `hashing`, or those mentioned above (e.g., `is parsable` or `dual tool verification` limited to the analysis stage):

"Verifying correctness is difficult in practice and can usually only be determined on a sample basis." (p16, translated)

"If the memory can be parsed, it's probably okay." (p4, translated)

15 participants indicated that they had encountered implausible information while checking the correctness of a memory dump in the past. The participants reported several problems, which can be grouped into the following themes: In their text responses, participants described various types of `inconsistencies`. These include issues with `processes`, such as faulty names, associated data, *"unreachable processes and missing/broken relations between parent, child and sibling processes" (p5)*, or a *"too short process list" (p10, translated)*. Further inconsistencies concerned `timestamps`, for example, *"impossible dates and time" (p15)*, and the `network`, such as *"invalid IP addresses in [the] connection record" (p25)*. These inconsistencies were sometimes linked to data being `incomplete`. Problems also became apparent when comparing the results from analyses from other tools or dumps (`different analysis results`), for instance *"different processes for the analysis with different analysis tools" (p16, translated)*. Participants also reported cases of `faulty analysis`, their tools either produced *"incorrectly parsed/corrupted information" (p22)* or failed entirely. A particular category of issues was associated with `special systems`. Examples included problems with the profile layout for *"specific custom kernel builds" (p13)* of Linux or complications caused by *"a dual boot system retaining RAM through reboot instead of power off event" (p15)*.

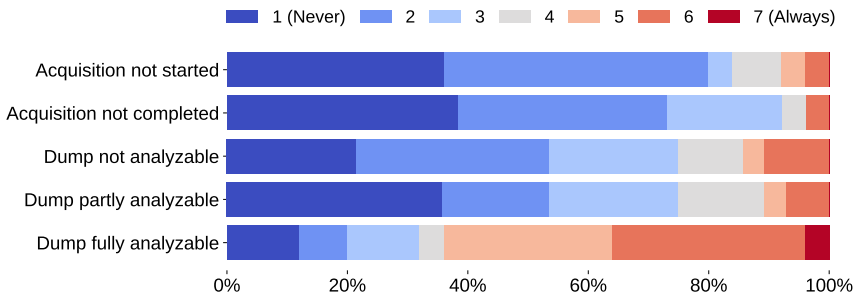

Figure 6: Experiences of participants during memory acquisition and analysis

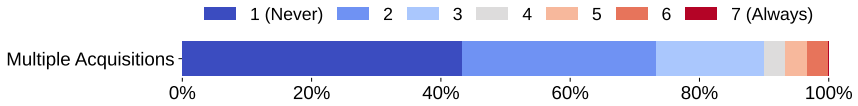

Figure 7: Frequency of performing multiple acquisitions for the same case

Participants in these situations often had no other choice than to `repeat the acquisition` or `disregard the results`:

> *"Take the snapshot again" (p5)*
>
> *"Move on because client never resolved the issue lol" (p20)*
>
> *"Ignore memory capture given it was useless" (p11)*

Participants also attempted to `use other tools` to check whether the problem persisted or to `compare` the results with other sources, e.g., *"cross validate finding e.g. with DNS logs" (p11)*. In specific cases, it was possible to perform some `manual analysis`, such as *"manually write [a] parser/hex editor template (010)" (p22)*, manually `fix the errors`, or to `contact the tool maker`.

Encountering such problems with implausible information can have negative consequences for practitioners' investigations. While there are fortunate cases in which there is `no effect`, participants more often reported at least `wasting time` or the analysis requiring `more effort` than otherwise needed:

> *"Much time was wasted." (p2)*
>
> *"Yes, in general they are more difficult to resolve and I rely on multiple snapshots to try to overcome them" (p5)*

In a few instances, participants mentioned `unspecific effects`, such as *"probably missing data but couldn't get around it" (p26)*. However, in the worst cases, the memory dump `could not be used`:

> *"Sometimes it means not relying on certain information or disregarding it altogether" (p22)*
>
> *"Direct use in argumentation questionable" (p14, translated)*

5. Discussion

In this section we first discuss the results of the survey regarding main memory acquisition and analysis. Then, we focus on the challenge of verifying the correctness of a memory dump. Finally, we outline the limitations of the survey.

5.1. The Need for Reliable Tools

The study indicates that participants with experience in memory forensics spent about 14% of their working hours acquiring or analyzing main memory in the last year, which is relatively high given the niche nature of memory forensics.

A common problem with main memory acquisition is reported to be a higher risk of the system crashing, which would sometimes mean to loose all data which is not desirable at all for some scenarios, for example in production. While it is well known that accessing certain memory regions with hardware-based acquisition techniques can lead to system crashes [27, 24, 10], we are not aware of any publications that investigate system crashes during memory acquisition with software-based techniques.

With respect to common problems regarding memory acquisition, participants reported that the process of acquiring main memory is cumbersome, i.e., time consuming, or even impossible due to access restrictions (physical or privilege-based). Moreover, problems can also arise during analysis. For example, participants indicated that the resulting memory dumps can contain process, timestamp and network inconsistencies or are incomplete. Special systems, e.g., custom Linux-based operating systems, hold additional challenges to forensic analysts. Lastly, the tools used to do the analysis can report wrong, implausible or no results.

As a result, when confronted with in some way broken memory dumps, practitioners are often forced to repeat

the acquisition, do a manual analysis, or disregard the results entirely. However, a repetition of the acquisition process is not always possible in some cases, so they refer to alternative methods to obtain the desired artifact, including live or disk analysis and a selective extraction of specific artifacts.

Since some information that justifies memory acquisition and analysis is only resident in main memory, there is a need for reliable acquisition and analysis tools which are easy to use, do not corrupt the system or expose either the system or the resulting memory dump to previously mentioned risks. Alternatively, a robust and straightforward way to assess the correctness of the extracted information from a memory dump is required.

5.2. Assuming Correctness

Although practitioners are generally aware of inconsistencies in memory dumps and attempt to verify their correctness, their options are often limited. In practice, this usually means either manually assessing whether the data appears plausible or comparing results with those obtained through alternative methods. Rzepka et al. [21] suggest more systematic strategies, such as dual-tool verification or testing tools against a known ground truth [6]. Some participants reported using dual-tool verification, but its effectiveness is debatable when applied only during the analysis phase: if inconsistencies are already introduced during acquisition, subsequent analysis cannot fully resolve them. Applying dual-tool verification at the acquisition stage presents its own difficulties, since tools cannot run simultaneously. As a result, they are likely to capture different system states, making direct comparison difficult. Only in the case of virtual machines can acquisitions of the exact same state be carried out reliably, yet these memory dumps tend to exhibit fewer inconsistency issues to begin with, as the system is halted.

That the currently employed strategies may not be particularly effective is further suggested by our findings: more than half of the participants reported never encountering implausible information in their dumps. Considering prior work highlighting the prevalence of inconsistencies [16, 20, 14], it seems improbable that their reports accurately reflect reality.

Practitioners, however, have limited options in this regard. On the one hand, they may employ suboptimal strategies in an effort to increase the likelihood of obtaining a correct dump. On the other hand, they ultimately need to assume, at least to some degree, that both acquisition and analysis have worked as intended, effectively operating under a presumption of reliability. Such an assumption can be problematic in judicial contexts as it can lead to false accusations and convictions [13, 12, 25]. So, this reliance on presumed correctness diminishes the evidential weight of memory forensics in court proceedings.

Instead, memory forensics appears to offer greater value in incident response, where the emphasis is not on producing evidence of high forensic quality but on quickly reconstructing events to mitigate security threats. This trend is already visible in our demographics compared with those reported by Hargreaves et al. [5] for general digital forensics. Whereas the majority of their respondents were affiliated with law enforcement (68.0% vs. 45.2%), our participant pool included a substantially larger proportion from incident response (58.1% vs. 12.3%).

5.3. Limitations

The number of participants (n=31) is relatively small compared to other surveys. However, given that memory forensics constitutes a highly specialized and relatively niche field, with practitioners employed in domains that are difficult to access, such as law enforcement, this sample size can nonetheless be considered a success. For comparison, in the survey conducted by Hargreaves et al. [5], only 10 participants (8.2%) indicated expertise in memory forensics.

6. Toward more reliable memory dumps

Based on the survey results of Section 4 and discussion in Section 5, we now come up with aspects to acquire memory dumps of high quality. We consider three aspects.

Decision Tree. First, we need an accepted and validated decision tree on how to choose the appropriate acquisition technique. Based on the well-known acquisition diagram by Ligh et al. [11], we propose an adapted decision tree on preferable acquisition techniques in Figure 8.

In the case of a virtualized system, the hypervisor provides the option to acquire an atomic RAM snapshot. Hence, in Figure 8, we recommend making use of this option in the case of a virtualized system. This suggestion aligns with the survey results. If the system is shut down, however, a RAM snapshot will be useless (as some participants stated, see Section 4.5), so only persistent RAM artifacts, such as hibernation files, are used.

If the system provides an EDR (Endpoint Detection and Response), we recommend using it for RAM acquisition. This situation mostly occurs in an institutional environment, so we recommend selecting an EDR that can acquire high-quality RAM snapshots. Our survey reveals that participants use this approach as well (see Section 4.2).

Next, we suggest using kernel-level software acquisition if the required privilege is available. The survey showed that this acquisition method is the most widely used (see Figure 5). However, we are aware that privileges are sometimes missing (`no access`, see Section 4.5).

Finally, if software-based acquisition is unavailable, we recommend carefully considering a hardware-based acquisition approach using Direct Memory Access (DMA) and cold boot, provided that the Memory Overwrite Request (MOR) bit is not set. Alternative extraction methods, such as live analysis, should also be considered.

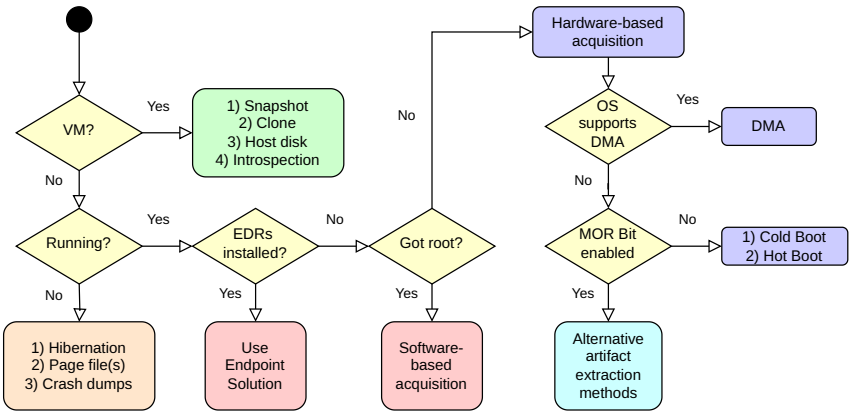

Figure 8: Adapted memory acquisition diagram from Ligh et al. [11]

Software Choice. Figure 5 reveals that software-based acquisition is the most common method. However, Section 4.5 indicated that in practice an unstructured evaluation of dump quality is common. Instead, practitioners use heuristics such as inconsistencies in the process, timestamp, and network information.

Our primary recommendation is to evaluate the quality of the acquisition software *a priori*, i.e., independently of a forensic examination. One key metric of an acquisition tool is its acquisition time (see [20, 21]). Hence, the preferred acquisition tool should be as fast as possible. For an administered system (e.g., an institutional environment), we recommend enforcing forensic readiness by installing EDR with a high-priority RAM acquisition facility. If possible, the software selection should be based on in vitro tests with respect to well-known ground truth, including cross-validation with results from disk or live analysis.

RAM vs. live analysis. Third, the survey indicates that a key issue is not acquiring main memory due to presumed system instability during acquisition. While we are not aware of any studies or experiences demonstrating the actual (in-) stability of the system during RAM acquisition, recent work by [21] does not reveal system instability during software-based acquisition. Nevertheless, we recommend testing a priori under which circumstances the system crashes during acquisition.

7. Conclusion

The paper presented the results of a conducted survey with practitioners about memory acquisition and analysis, in particular common use cases for memory forensics, memory acquisition techniques and the handling of inconsistencies present in memory dumps.

As a result, memory forensics is still used to extract information otherwise unobtainable, e.g., encryption keys or network information. However, the participants highlighted challenges of memory forensics, including the risk of system crash and the reliability of evidence from RAM. Even though it seems to be apparent that inconsistencies caused by commonly used software tools in memory dumps can lead to disregarding analysis results due to reliability problems, the correctness of a acquired memory dump is often assumed rather than thoroughly tested. Moreover, participants indicated that the mentioned problems with memory analysis often result in more effort and wasted time.

The results suggest to research a robust and reliable way to assess the correctness of memory dumps, as well as to develop reliable acquisition and analysis tools which take the problem of inconsistencies into account. Until then, we recommended to choose a suitable acquisition method based on an acquisition diagram. In case of software-based acquisition, which is the most common method used by practitioners, the software should be tested independently of the forensic examination to investigate the likelihood of errors during the acquisition like system crashes.

Acknowledgments

We would like to thank Maximilian Eichhorn and Gaston Pugliese for helping translate the survey questions. Special thanks go to all participants, testers, and to those who distributed our survey. Without them, our study would not have been possible. This work was partially supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Research Training Group 2475 “Cybercrime and Forensic Computing” (grant number 393541319/GRK2475/2-2024).

CRediT Authorship Contribution Statement

This section follows the guide from `https://credit.niso.org/`.

Lisa Rzepka: Conceptualization, Methodology, Investigation, Data Curation, Validation, Writing - Original Draft, Writing - Review & Editing, Visualization

Benedikt Mader: Conceptualization, Methodology, Investigation, Data Curation, Validation, Writing - Original Draft, Writing - Review & Editing, Visualization
Zinaida Benenson: Conceptualization, Methodology, Writing - Review & Editing, Supervision
Harald Baier: Conceptualization, Methodology, Writing - Review & Editing, Supervision

References

[1] Al Fahdi, M., Clarke, N., Furnell, S., 2013. Challenges to digital forensics: A survey of researchers & practitioners attitudes and opinions, in: 2013 Information Security for South Africa, pp. 1–8. URL: https://ieeexplore.ieee.org/document/6641058, doi:10.1109/ISSA.2013.6641058. iSSN: 2330-9881.

[2] Braun, V., Clarke, V., 2021. Can I use TA? Should I use TA? Should I not use TA? Comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches. Counselling and Psychotherapy Research 21, 37–47. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/capr.12360, doi:10.1002/capr.12360.

[3] Case, A., Richard III, G.G., 2017. Memory forensics: The path forward. Digit. Investig. 20, 23–33. URL: https://doi.org/10.1016/j.diin.2016.12.004, doi:10.1016/J.DIIN.2016.12.004.

[4] Franqueira, V.N.L., Bryce, J., Al Mutawa, N., Marrington, A., 2018. Investigation of Indecent Images of Children cases: Challenges and suggestions collected from the trenches. Digital Investigation 24, 95–105. URL: https://www.sciencedirect.com/science/article/pii/S1742287617302669, doi:10.1016/j.diin.2017.11.002.

[5] Hargreaves, C., Breitinger, F., Dowthwaite, L., Webb, H., Scanlon, M., 2024a. DFPulse: The 2024 digital forensic practitioner survey. Forensic Science International: Digital Investigation 51, 301844. URL: https://www.sciencedirect.com/science/article/pii/S2666281724001719, doi:10.1016/j.fsidi.2024.301844.

[6] Hargreaves, C., Nelson, A., Casey, E., 2024b. An abstract model for digital forensic analysis tools - a foundation for systematic error mitigation analysis. Forensic Science International: Digital Investigation 48, 301679. URL: https://www.sciencedirect.com/science/article/pii/S2666281723001981, doi:https://doi.org/10.1016/j.fsidi.2023.301679. DFRWS EU 2024 - Selected Papers from the 11th Annual Digital Forensics Research Conference Europe.

[7] Horsman, G., 2019. Tool testing and reliability issues in the field of digital forensics. Digital Investigation 28, 163–175.

[8] James, J.I., Gladyshev, P., 2016. A survey of mutual legal assistance involving digital evidence. Digital Investigation 18, 23–32. URL: https://www.sciencedirect.com/science/article/pii/S174228761630072X, doi:10.1016/j.diin.2016.06.004.

[9] Joshi, A., Kale, S., Chandel, S., Pal, D., 2015. Likert Scale: Explored and Explained. British Journal of Applied Science & Technology 7, 396–403. URL: https://journalcjast.com/index.php/CJAST/article/view/381, doi:10.9734/BJAST/2015/14975.

[10] Libster, E., Kornblum, J.D., 2008. A proposal for an integrated memory acquisition mechanism. SIGOPS Oper. Syst. Rev. 42, 14–20. URL: https://doi.org/10.1145/1368506.1368510, doi:10.1145/1368506.1368510.

[11] Ligh, M.H., Case, A., Levy, J., Walters, A., 2014. The art of memory forensics: detecting malware and threats in windows, linux, and Mac memory. John Wiley & Sons.

[12] Marshall, P., 2020. The harm that judges do – misunderstanding computer evidence: Mr Castleton's story: 'an affront to the public conscience'. Digital Evidence and Electronic Signature Law Review , 25–48URL: https://journals.sas.ac.uk/deeslr/article/view/5172, doi:10.14296/deeslr.v17i0.5172.

[13] Mason, S., 2014. Electronic evidence: A proposal to reform the presumption of reliability and hearsay. Computer Law & Security Review 30, 80–84. URL: https://www.sciencedirect.com/science/article/pii/S0267364913002057, doi:10.1016/j.clsr.2013.12.005.

[14] Oliveri, A., Balzarotti, D., 2025. A comprehensive quantification of inconsistencies in memory dumps, in: ACM (Ed.), RAID 2025, 28th International Symposium on Research in Attacks, Intrusions and Defenses, 19-22 october 2025, Gold Coast, Australia, Gold Coast.

[15] Osborne, G., 2013. Memory Forensics: Review of Acquisition and Analysis Techniques. Technical Report DSTO-GD-0770. Australian Department of Defence, Defence Science and Technology Organisation.

[16] Ottmann, J., Breitinger, F., Freiling, F., 2022. Defining Atomicity (and Integrity) for Snapshots of Storage in Forensic Computing, in: Proceedings of the Digital Forensics Research Conference Europe (DFRWS EU) 2022. URL: https://dfrws.org/presentation/defining-atomicity-and-integrity-for-snapshots-of-storage-in-forensic-computing/.

[17] Ottmann, J., Breitinger, F., Freiling, F.C., 2024. An experimental assessment of inconsistencies in memory forensics. ACM Trans. Priv. Secur. 27, 2:1–2:29. URL: https://doi.org/10.1145/3628600, doi:10.1145/3628600.

[18] Pagani, F., Fedorov, O., Balzarotti, D., 2019. Introducing the temporal dimension to memory forensics. ACM Trans. Priv. Secur. 22, 9:1–9:21. URL: https://doi.org/10.1145/3310355, doi:10.1145/3310355.

[19] ReliaQuest, T.R.T., 2024. Living off the land and fileless malware. URL: https://reliaquest.com/blog/living-off-the-land-fileless-malware/.

[20] Rzepka, L., Ottmann, J., Freiling, F., Baier, H., 2024. Causal inconsistencies are normal in windows memory dumps (too). Digital Threats URL: https://doi.org/10.1145/3680293, doi:10.1145/3680293. just Accepted.

[21] Rzepka, L., Ottmann, J., Stoykova, R., Freiling, F., Baier, H., 2025. A scenario-based quality assessment of memory acquisition tools and its investigative implications. Forensic Science International: Digital Investigation 52, 301868. URL: https://www.sciencedirect.com/science/article/pii/S2666281725000071, doi:https://doi.org/10.1016/j.fsidi.2025.301868. DFRWS EU 2025 - Selected Papers from the 12th Annual Digital Forensics Research Conference Europe.

[22] Saldaña, J., 2025. The Coding Manual for Qualitative Researchers. 5 ed., SAGE. URL: https://uk.sagepub.com/en-gb/eur/the-coding-manual-for-qualitative-researchers/book287917.

[23] Sanchez, L., Grajeda, C., Baggili, I., Hall, C., 2019. A Practitioner Survey Exploring the Value of Forensic Tools, AI, Filtering, & Safer Presentation for Investigating Child Sexual Abuse Material (CSAM). Digital Investigation 29, S124–S142. URL: https://www.sciencedirect.com/science/article/pii/S1742287619301549, doi:10.1016/j.diin.2019.04.005.

[24] Stüttgen, J., Cohen, M., 2013. Anti-forensic resilient memory acquisition. Digital Investigation 10, S105–S115. URL: https://www.sciencedirect.com/science/article/pii/S1742287613000583, doi:https://doi.org/10.1016/j.diin.2013.06.012. the Proceedings of the Thirteenth Annual DFRWS Conference.

[25] Van Buskirk, E., Liu, V.T., 2006. Digital Evidence: Challenging the Presumption of Reliability. Journal of Digital Forensic Practice 1, 19–26. URL: https://doi.org/10.1080/15567280500541421, doi:10.1080/15567280500541421. publisher: Taylor & Francis _eprint: https://doi.org/10.1080/15567280500541421.

[26] Vömel, S., Freiling, F.C., 2012. Correctness, atomicity, and integrity: Defining criteria for forensically-sound memory acquisition. Digit. Investig. 9, 125–137. URL: https://doi.org/10.1016/j.diin.2012.04.005, doi:10.1016/J.DIIN.2012.04.005.

[27] Vömel, S., Freiling, F., 2011. A survey of main memory acquisition and analysis techniques for the windows operating system. Digital Investigation 8, 3–22. doi:10.1016/j.diin.2011.06.002.

Appendix A. Survey

This section contains the complete English version of the survey. All questions marked with an asterisk (*) are mandatory.

1. Introduction
   - 1.1 Thank you for taking the time to participate in this survey! Please be sure to answer all the questions, as only completed surveys can be analyzed. Your answers will be anonymized and treated confidentially.
     The survey can be completed in about 10-15 minutes.
     The survey seeks to better understand the use and challenges of acquiring and analyzing main memory (RAM) in practice. This refers to creating a copy of the main memory of a running system for analysis purposes. The resulting copy is usually called a "memory dump".
     [Contact information for questions]
   - 1.2 I hereby consent to participating in this survey. The goal of this survey is to gain insights into the usage and common challenges of memory acquisition and analysis. I consent that my survey data will be collected, processed and used by [name of university] for research purposes, and will not be given to any third parties. Under Article 6 (1) (a) GDPR, I have the right to know exactly what data is stored and to have this data deleted. I may abort the survey at any time and delete my answers without any disadvantages. All survey data are anonymized for processing.*
     • I consent.
     • I do not contest *(end of questionnaire)*
   - 1.3 Have you ever created a copy of the main memory (RAM) for analysis purposes?*
     • Yes • No *(end of questionnaire)*
   - 1.4 For how many years has your work included memory acquisition and analysis?* *(numeric input)*
   - 1.5 How would you rank your expertise in acquiring main memory?* *(scale: 1 no expertise – 7 high expertise)*
   - 1.6 How would you rank your expertise in analyzing main memory?* *(scale: 1 no expertise – 7 high expertise)*
2. Memory Acquisition
   - 2.1. What percentage of your working hours did you spend acquiring or analyzing main memory last year? *(slider 0-100%)*
   - 2.2. What types of information do you want to extract when acquiring main memory?* *(multiple choice, [15])*
     • Lists of currently executing processes and threads
     • Cryptographic keys
     • System registry information
     • Established network connections and related network data such as Internet Protocol (IP) addresses and port information
     • Command history
     • Open files
     • Current users
     • Lists of rootkits or other malware
     • Other:
     • I don't know
   - 2.3. How often do you use the following methods when acquiring main memory?* *(matrix with scale: 1 never – 7 always)*
     • Software tools (kernel- or user-level) running on the target system
     • Direct memory access (DMA)
     • Cold boot
     • Hypervisor-based acquisition in virtual machines
     • Other:
   - 2.4. Which other methods do you use for acquiring main memory? *(free text, if Other >1)*
   - 2.5. Which software tools do you use for acquiring main memory? *(free text, if Software tools >1)*
   - 2.6. Which direct memory access techniques do you use for acquiring main memory? *(free text, if DMA >1)*
   - 2.7. How often do you perform memory acquisition multiple times for the same case?* *(scale: 1 never – 7 always)*
   - 2.8. If you decide against acquiring main memory in a specific case, what alternative methods do you use? *(free text)*
   - 2.9. What are your reasons for using alternative methods? *(free text)*
3. Memory Analysis
   - 3.1. How often did you experience the following situations in the last year?* *(matrix with scale: 1 never – 7 always)*
     • The memory acquisition could not be started
     • The memory acquisition could not be completed successfully
     • The memory dump could not be analyzed
     • Only parts of the memory dump could be analyzed
     • The memory dump could be analyzed completely
   - 3.2. How do you check the correctness of information extracted from a memory dump? *(free text)*
   - 3.3. While checking the correctness of a memory dump, have you ever encountered implausible information?*
     • Yes • No
   - 3.4. What problems with implausible information did you encounter? *(free text, if encountered problems: yes)*
   - 3.5. How did you address the encountered problems? *(free text, if encountered problems: yes)*
   - 3.6. Did the encountered problems affect your investigation in any way? *(free text, if encountered problems: yes)*
4. Demography
   - 4.1. In which of these roles did you acquire and analyze main memory?* *(multiple choice, [5])*
     • Law enforcement
     • Independent consultant
     • Incident response
     • Industry
     • Research
     • Student
     • Other:
   - 4.2. What is your current main employment status?
     • Employed in private sector
     • Employed in public sector (e.g., civil servant / government official)
     • Self-employed / freelancer
     • In training / apprenticeship / student
     • Other:
   - 4.3. What is your year of birth? *(list of years)*
   - 4.4. What is your current country of residence? *(list of countries)*
   - 4.5. Please select your highest educational level.
     • Less than high school (no university/college entrance certificate)
     • Entrance certificate for university/college (high school diploma, GED, GCE, etc.)
     • Some college, associate degree or equivalent
     • Bachelor's degree or equivalent
     • Master's degree or equivalent
     • Professional degree (M.D., J.D., etc), Ph.D. (doctoral degree) or equivalent
     • Other:
   - 4.6. Please select the gender you most identify with.
     • Female
     • Male
     • Non-binary / diverse
     • Other:
5. End Page
   - 5.1. Thank you for your participation!
     Please help us to distribute this survey by forwarding it to other potential participants who have experience with main memory acquisition and analysis by using this link: [link]